

# Depth fusion for anti-personnel landmine detection

John G.M. Schavemaker<sup>a</sup>, Eric den Breejen<sup>a</sup>, Frank Cremer<sup>a,b,c</sup>, Klamer Schutte<sup>a</sup> and Koen W. Benoist<sup>a</sup>

<sup>a</sup>TNO Physics and Electronics Laboratory, P.O. Box 96864, 2509 JG, The Hague, The Netherlands

<sup>b</sup>Pattern Recognition Group, Delft University of Technology, Delft, The Netherlands

<sup>c</sup>Section of Applied Geophysics, Delft University of Technology, Delft, The Netherlands

## ABSTRACT

In this paper we introduce the concept of depth fusion for anti-personnel landmine detection. Depth fusion is an extension of common sensor-fusion techniques for landmine detection. The difference lies within the fact that fusion of sensor data is performed in different physical depth layers. In order to do so, it requires a sensor that provides depth information for object detections. Our ground-penetrating radar (GPR) fulfills this requirement. Depth fusion is then taken as the combination of the output of sensor fusion of all layers. The underlying idea is that sensor fusion for the surface layer has a different weighing of the sensors when compared with the sensor fusion in the deep layers because of apparent sensor characteristics. For example, a thermal infrared (TIR) sensor hardly adds information to the sensor fusion in the deep layers. Furthermore, GPR has difficulties suppressing clutter in the surface layer. As such, the surface fusion should emphasize on the TIR sensor, whereas sensor fusion in the deep layers should have a higher weighing of the GPR. This *a priori* information can be made explicit by choosing for a depth-fusion approach. Experimental results from measurements at the TNO-FEL test facility are presented that validate our depth-fusion concepts.

**Keywords:** Sensor fusion, depth fusion, GPR, MD, TIR, landmine detection.

## 1. INTRODUCTION

The detection of landmines by any (technical) means is an important research issue. Current research focuses on the use of multiple sensors to landmine detection, as opposed to approaches which use a single sensor. The use of one sensor is generally believed to be insufficient for landmine detection meeting the requirements of humanitarian demining. A single sensor has a detection rate which is too low or a false-alarm rate which is too high. The aim of sensor fusion is to achieve a higher probability of detection with a lower probability of false alarms.

TNO-FEL has reported work on sensor fusion for landmine detection.<sup>2,1,7,4-6,14</sup> Our continuing goal is to improve the performance of fusion concepts for humanitarian demining. The ESPRIT project LOTUS, in which we are currently involved, strives towards a vehicle-mounted sensor-fusion system for the detection of anti-personnel (AP) landmines. The sensor suite of this system consists of a ground penetrating radar, a metal detector and a thermal infrared camera.

The employed sensor-fusion methods apply decision-level sensor fusion. However, some of the sensors are capable of providing additional features from their data which can not be easily incorporated into a decision-level scheme. For example, GPR can provide object detections with depth information and TIR can give shape features as well as estimates of texture of background clutter and objects. To improve sensor fusion these features must be included in some way. We focus on the extension of decision-level sensor fusion with depth information in this paper.

## 2. SENSOR-FUSION METHODS

Sensor-fusion methods for landmine detection differ in the fusion level they apply. The fusion level defines the abstraction level at which the information of the different sensors is combined. Commonly, a distinction is made between data-, feature- and decision-level fusion<sup>16</sup> (see also Figure 1(a)):

**Data-level** fusion is the fusion of sensor data like camera images. Landmine detection systems that apply data-level fusion usually have multiple sensors of the same kind that differ in wavelength range<sup>3</sup> or polarization.<sup>9</sup> The fusion methods are in most cases based on a physical model of the sensor and convert different sensor data into one image for visual display or further analysis.

---

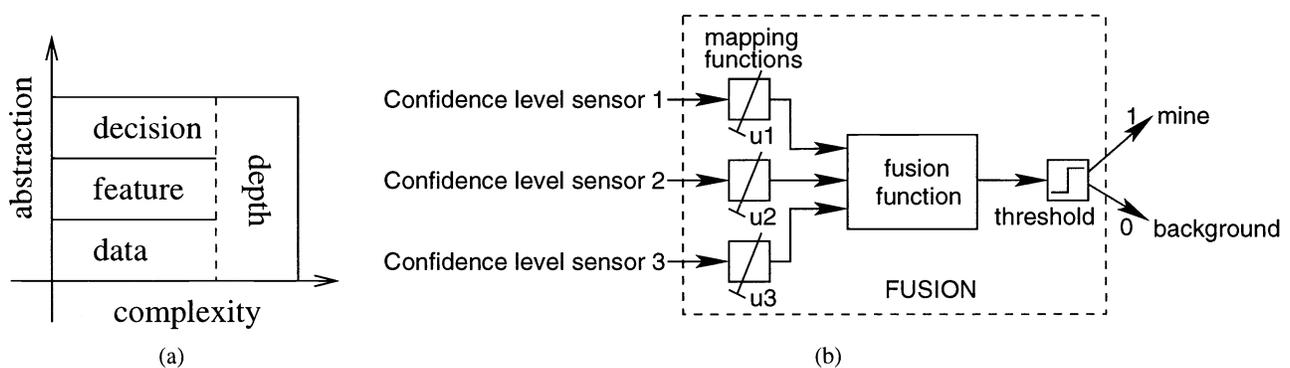
Further author information: (Send correspondence to Eric den Breejen)

Eric den Breejen: E-mail: Breejen@fel.tno.nl Phone: +31 70 374 0426 Fax: +31 70 374 0654

**Feature-level** fusion is the fusion of features extracted from the different sensor data that represent object properties like size and material type. Feature-level methods fuse information by means of selecting or concatenating features obtained from the different sensors. Multiple features are converted into one feature or a detection (by means of classification) at the feature level resulting in a (probability of) detection.

**Decision-level** fusion fuses detections (or probabilities of detections) of different sensors. When features are not available, one can choose for a decision-level fusion which takes landmine detections from the different sensors as input. Decision-level methods differ in the fusion method applied as well as the kind of input and output of the method. Popular fusion methods are Bayesian approaches, applications of Dempster-Shafer theory, fuzzy logic, rules and voting techniques.

The choice of a suitable fusion level depends on the available sensor types: when the sensors are alike, one can opt for fusion at data level. Feature-level fusion is the proper level when features obtained from different sensors can be combined in such a way that the combination provides sufficient information for landmine detection. When the sensors are very different, decision-level fusion is more suitable and also is computationally more efficient. Furthermore, 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 they can use this expertise to optimize the detection performance. For these reasons and the availability of this expert knowledge, we have chosen for decision-level fusion.



**Figure 1.** (a) A taxonomy of sensor-fusion methods. (b) The general layout of our decision-level sensor-fusion methods for three sensors.

Notwithstanding our choice, we still are in the process of researching extensions or alternatives to improve our sensor-fusion system. One of the foreseen improvements is to incorporate more features in the decision making process. Currently, we regard the depth of object detections as an important feature to investigate, for the reason that the performance of most sensors depends on the depth of the objects to be detected. Furthermore, our decision-level fusion is inherently a scheme that works on two-dimensional data, whereas, with the current rate of development, some sensors (GPR and MD) are capable of providing three-dimensional sensor data. To facilitate three-dimensional sensor fusion we introduce the new concept of depth fusion. Before we do so, we shortly describe decision-level fusion methods in the next sections as they form the basis of depth fusion.

### 3. DECISION-LEVEL SENSOR-FUSION METHODS

Figure 1(b) shows the general layout of a decision-level sensor-fusion method. The input of a sensor-fusion method is a confidence level for each location and sensor (with the assumption that all sensor confidence levels are co-registered on a two-dimensional grid), and the output after fusion is one confidence level per location. The output can be thresholded to obtain one for a landmine detection and zero for no detection (background).

A confidence level of a sensor or sensor fusion at a certain location expresses a confidence or belief in a landmine detection on that position given that sensor or sensor-fusion method. Confidence levels are numbers in the range from zero to one and are not necessarily a probability of detection in a statistical sense. They 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 landmine, but these do not have to scale linearly and may even differ per sensor.

A decision-level sensor-fusion method adjusts the weight of each sensor by means of a mapping function. This mapping function requires one parameter ( $u_i$ ) per sensor and can remove differences in definitions of the confidence levels among

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

**Table 1.** Decision-level sensor-fusion methods and their corresponding mapping functions.

sensors. The mapped confidence levels are combined with a decision-level fusion function to acquire a single confidence value per grid cell. Common decision-level sensor-fusion methods and their corresponding mapping functions are listed in Table 1.

The next sections contain short descriptions of three decision-level sensor-fusion methods that are used in our experiments. For more detail and descriptions of other decision-level sensor-fusion methods we refer to earlier work.<sup>4</sup> The fusion methods we describe and use in this paper are: best sensor, naive Bayes and Dempster-Shafer.

### 3.1. Best sensor

The best sensor is the most simple form of a decision-level sensor-fusion method: it just selects the best sensor from its input. The selection is based on the detection and false-alarm rate of each sensor. The best sensor is chosen as the sensor with the lowest number of false alarms given a certain detection rate. Both the detection rate as well as the false-alarm rate depend on the threshold applied. Hence, for different thresholds, a corresponding number of so-called working points (pairs of detection and false-alarm rates) can be obtained for each sensor. When we choose the best working point for each threshold from the available sensors, we obtain the best sensor.

### 3.2. Naive Bayes

The naive Bayes' approach assumes that the confidence levels scale with the likelihood ratio per sensor. The likelihood ratio is the quotient of the conditional probability densities for both classes (mines and background). Based on this likelihood ratio the optimal Bayes decision can be made.<sup>10</sup> If these conditional probabilities are sensor independent, then the joint likelihood ratio is the product of these individual likelihood ratios. Since the confidence levels most likely do not linear scale with the likelihood ratios, a mapping factor is included for each sensor. The naive Bayes fusion function  $f(c_1, c_2, c_3)$  and its mapping are given by:

$$f(c_1, c_2, c_3) = \left( (1 - u_1)c_1 + \frac{1}{2}u_1 \right) \cdot \left( (1 - u_2)c_2 + \frac{1}{2}u_2 \right) \cdot \left( (1 - u_3)c_3 + \frac{1}{2}u_3 \right), \quad (1)$$

with  $c_1, c_2, c_3$  the confidence levels and  $u_1, u_2, u_3$  the mapping parameters.

### 3.3. Dempster-Shafer theory

For application of the Dempster-Shafer theory<sup>8</sup> to sensor fusion, three inputs per sensor are needed: the probability mass assigned to a landmine  $m(M)$ , the probability mass assigned to background  $m(\bar{M})$ , and the unassigned probability mass  $m(M \cup \bar{M})$ . The sum of these masses always equals one, so there are only two independent masses ( $m(M)$  and  $m(\bar{M})$ ). The mass  $m(M)$  represents a belief in a landmine detection, the mass  $m(\bar{M})$  represents the opposite, and the mass  $m(M \cup \bar{M})$  reflects the uncertainty of the sensor. Each sensor produces only one confidence level at each sample location, which must be mapped onto the three required probability masses. In our case, the unassigned probability mass is kept constant for each of the sensors, and the probability mass assigned to a landmine is chosen linear with the confidence level. Given the confidence level of sensor  $i$ , the following mapping is used to acquire the probability masses:

$$m_i(M) = (1 - u_i)c_i \quad (2)$$

$$m_i(\bar{M}) = (1 - u_i)(1 - c_i) \quad (3)$$

$$m_i(M \cup \bar{M}) = u_i \quad (4)$$

where  $u_i \in [0, 1]$  is the unassigned probability mass or the uncertainty of sensor  $i$ . This mapping assures that higher confidence levels lead to higher probability masses assigned to a landmine.

The probability mass assigned to a landmine ( $m(M)$ ) is zero at confidence level zero, and the probability mass assigned to background ( $m(\bar{M})$ ) is zero for confidence value one. The uncertainty  $m(M \cup \bar{M})$  is optimized for each sensor using a training set of examples.

Given two sensors with probability masses  $m_1(M)$ ,  $m_1(\bar{M})$ , and  $m_1(M \cup \bar{M})$ , respectively  $m_2(M)$ ,  $m_2(\bar{M})$ , and  $m_2(M \cup \bar{M})$ , the resulting fused probability masses  $m_{1,2}$  are calculated using the following equations:

$$m_{1,2}(M) = \frac{m_1(M)m_2(M) + m_1(M)m_2(M \cup \bar{M}) + m_1(M \cup \bar{M})m_2(M)}{k} \quad (5)$$

$$m_{1,2}(\bar{M}) = \frac{m_1(\bar{M})m_2(\bar{M}) + m_1(\bar{M})m_2(M \cup \bar{M}) + m_1(M \cup \bar{M})m_2(\bar{M})}{k} \quad (6)$$

$$m_{1,2}(M \cup \bar{M}) = \frac{m_1(M \cup \bar{M})m_2(M \cup \bar{M})}{k} \quad (7)$$

$$k = 1 - m_1(M)m_2(\bar{M}) - m_1(\bar{M})m_2(M)$$

To combine three sensors, the procedure is to first combine two sensors, followed by the combination of the resulting probability masses  $m_{1,2}$  and the probability masses of the third sensor  $m_3$ .

There are three resulting probability masses after fusion of three sensors, denoted by  $m_{1,2,3}(M)$ ,  $m_{1,2,3}(\bar{M})$  and  $m_{1,2,3}(M \cup \bar{M})$ . To use Dempster-Shafer theory as a classifier, a single value is needed. The probability mass assigned to a landmine, also called support, can be used for this purpose. Another option is to use the probability mass assigned to a landmine increased with the uncertainty, also called the plausibility. We use the average of the support and plausibility:

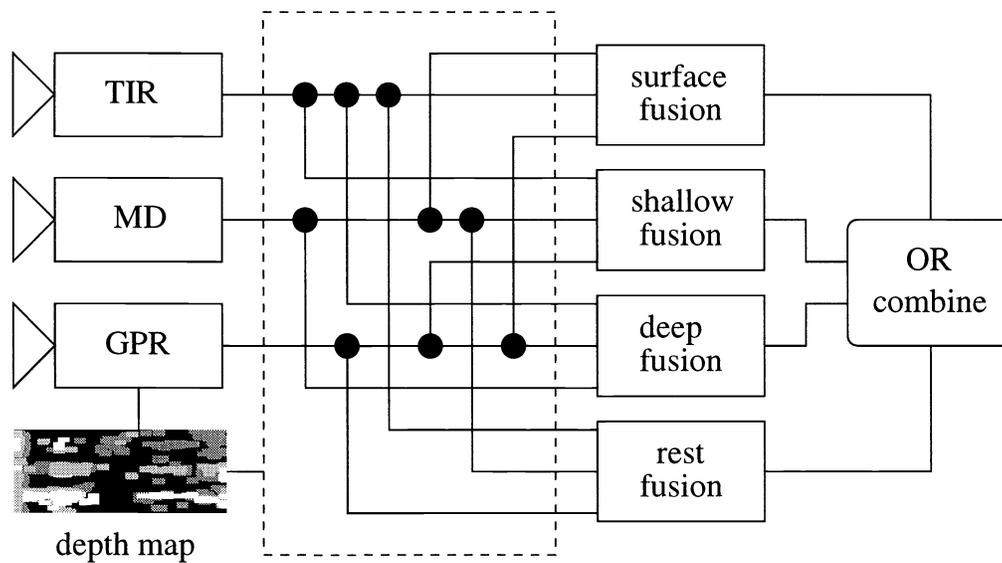
$$m_{res}(M) = m_{1,2,3}(M) + \frac{1}{2}m_{1,2,3}(M \cup \bar{M}) \quad (8)$$

#### 4. DEPTH FUSION

Depth fusion takes into account that performance of a sensor depends on the depth of the objects to be detected. Such depth-related performance is apparent for GPR, MD as well as TIR. For example, GPR performance deteriorates for the surface due to reflections at the air-to-ground interface. These reflections cause clutter that cannot be suppressed easily. TIR, on the other hand, performs at its best when objects are laid on the surface but has difficulties detecting objects that are buried. Those objects can only be detected with a high number of false alarms. The performance of MD is also affected by depth, though in a less discrete way than GPR or TIR. Our MD processing is not capable of discriminating between metal content and object depth; both features have a relationship with the observed signal strength. As such, it is not possible to distinguish one object from another that has less metal but that is buried less deep.

Depth fusion attempts to overcome these depth depending sensor performances by dividing the sensor fusion task in a number of (physical) depth layers. For each layer decision-level sensor fusion is applied to combine the sensor readings for that specific layer. All layers are combined together by means of a disjunction of all sensor-fusion detections, see Figure 2. The sensor fusion at a specific layer is performed only for the grid cells for which the GPR found object evidence pointing to that specific layer. The GPR therefore is the root sensor to generate a so-called depth map of the area under investigation. The depth map indicates for each grid cell at which layer the sensor fusion should take place. As such, the map gives a tessellation of the grid into the different layers. To cover all grid cells an additional rest layer is introduced to include the cells for which the GPR found no evidence for any layer.

All decision-level fusion methods incorporate some optimization parameters. The optimization parameters are determined for each layer with an exhaustive search. For naive Bayes and Dempster-Shafer the mapping parameters (one per sensor) are varied in 11 steps and the threshold is set to the minimum value to detect all landmines. This gives a total number of 1331 possible parameter settings for each layer. Each parameter setting corresponds with a detection and false-alarm rate. The optimal parameter settings are a subset of all possible parameter settings that have the lowest false-alarm rate for certain detection rates (and correspond with working points on the ROC curve, see next section). To optimize sensor fusion for a specific layer the depth map is used in conjunction with the ground truth. The optimal parameter settings of sensor fusion for a specific layer are then found by calculating the detection and false-alarm rates for the grid cells of that layer.



**Figure 2.** Schematic overview of depth fusion.

Depth fusion is not optimized by means of optimization of its parameters. Optimization of depth fusion is performed analogous to building the best sensor instead. Depth fusion combines the (binary) detections of sensor fusion for all layers by means of a logical OR. To obtain different optimal settings of depth fusion that have the lowest false-alarm rate for certain detection rates, optimal settings for the different layers are combined. As such, the optimization of depth fusion is an exhaustive search among the optimal settings of sensor fusion for all layers.

## 5. PERFORMANCE EVALUATION

The performance of sensors and sensor-fusion techniques is evaluated using the Receiver Operator Characteristics (ROC) curve.<sup>13</sup> In a ROC curve the detection rate is plotted against the false-alarm rate for adjustable (optimization) parameters. The detection rate is defined as the fraction of detected landmines. The corresponding number of false alarms per unit area is calculated using the SCOOP<sup>1</sup> performance-evaluation method. Each working point on the ROC corresponds with a set of values for the optimization parameters.

The evaluation method called Split Clusters On Oversized Patches (SCOOP) yields a less biased probability of false alarm by taking into account the spatial arrangement of false alarms. The name SCOOP refers to the scoop-size area that needs to be checked by the landmine-clearance personnel. The method takes into account that it takes landmine-clearance personnel more time to reject a false alarm occupying a large area (multiple 'scoops') than a false alarm occupying a small area. As such, this performance-evaluation method finds a balance between the surface that false alarms occupy and the number of false alarms. This kind of performance measure is requested by the demining experts as an indication of the reduction in workload of the human operator and landmine-clearance personnel.

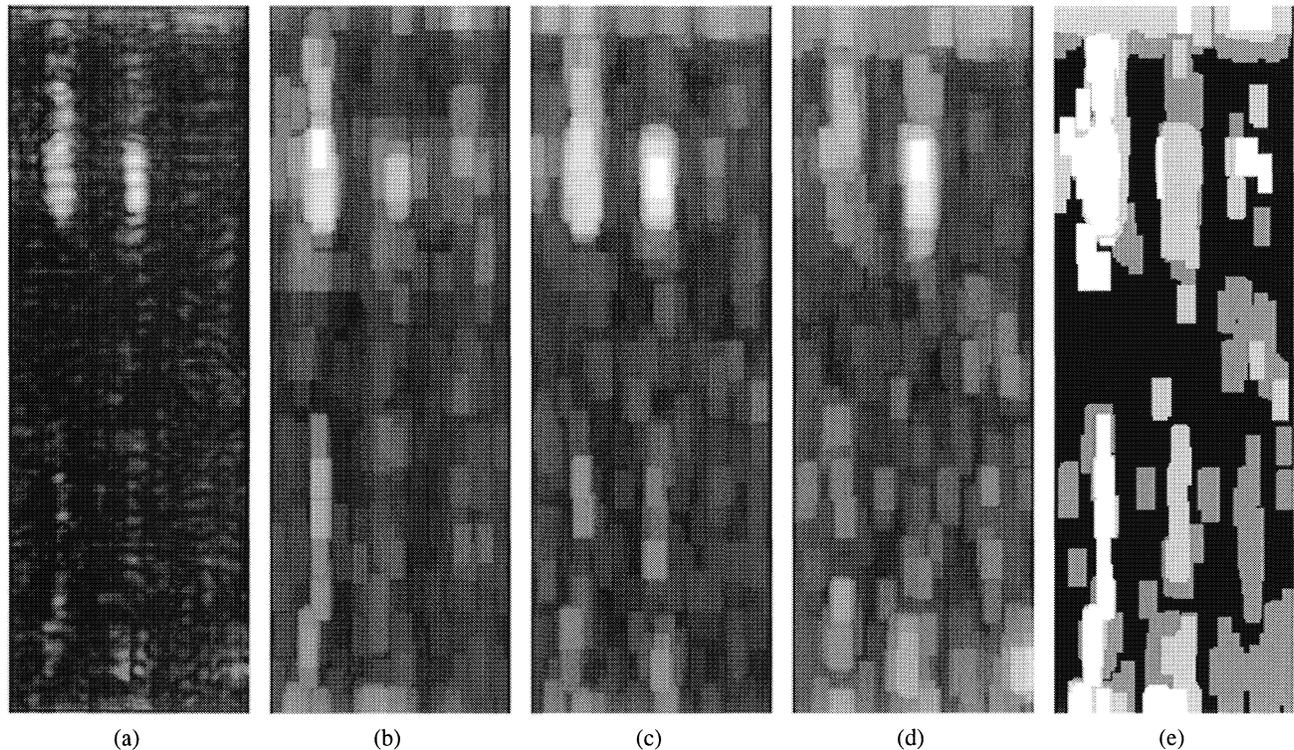
## 6. EXPERIMENTS

In this section we show results of experiments with sensor-fusion methods on sensor data recorded on the sand test lane at the TNO test facility<sup>11</sup> with the 1998 layout. We show results of single sensors, common decision-level sensor-fusion methods and depth fusion. All results are evaluated with the SCOOP algorithm as described in the previous section. First, we describe sensors and their processing.

### 6.1. GPR depth processing

The GPR used is a Geocenter 301B. It is an array with 2D beam-steering capabilities. For the test lane three sensor sweeps are combined to a single lane-wide 3D data set. This data set is resampled to have 5 cm gridcells in the  $x$  and  $y$  directions, and to have 4 times more gridcells in the  $z$  direction. For every  $(x, y)$  position the height of the surface layer is estimated. This height map is smoothed by subsequently applying a median filter in horizontal and vertical direction. The complete 3D data set is for

every  $(x, y)$  position translated by this smoothed surface layer estimate. This will compensate for timing artefacts between the GPR channels and the slope in the test lane, so after this correction the  $z$  coordinate of the 3D data set corresponds better to true depth.



**Figure 3.** (a) The GPR confidence level map. (b)-(d) The local maxima images for the standard deviation of the top, middle and bottom layers. (e) Result of the depth processing. The white areas are assigned to the top layer, the middle gray to the medium-depth layer, and the dark gray area to the bottom layer. The black area is not assigned to any depth layer, indicating unknown depth.

From this data set the local mean in the scan direction is subtracted. This should reveal local structure in the GPR return. Using the local mean estimate in the cross scan direction did not work due to non uniform readout from the different GPR channels. Also, the Geocenter 301B GPR used already has processing in place for the cross-scan direction. Subsequently, four depth layers are defined: top (0-31), middle (32-63), bottom (100-131), and full (0-159). For these layers the standard deviation in the vertical direction is calculated. The standard deviation of the full layer is used as confidence map, see Figure 3(a). In a  $3 \times 9$  (cross-scan  $\times$  scan) region the local maximum is calculated of the three standard deviation images of the top, middle and bottom layers resulting in the images  $t_m$ ,  $m_m$ , and  $b_m$  for the top, middle and bottom layer. These images are depicted in Figure 3. For the top layer is used:

$$t_b = T[C_l[t_m, 65]/m_m, 0.4]$$

For the middle layer is used:

$$m_b = T[m_m, 0.94] \wedge t_b$$

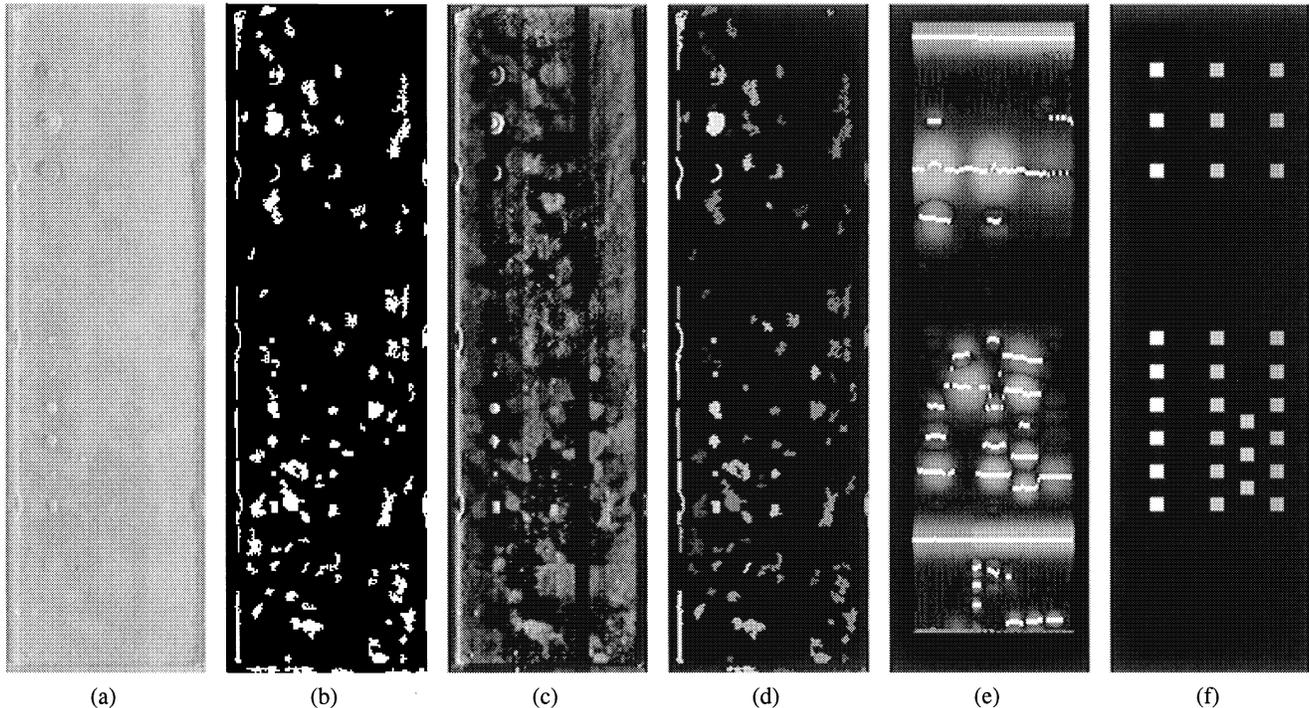
For the bottom layer is used:

$$b_b = T[C_l[b_m, 65]/m_m, 0.63] \wedge (t_b \vee m_b)$$

where the operator  $T[I, t]$  thresholds image  $I$  at level  $t$  and operator  $C_l[I, l]$  maps all pixel values in  $I$  less than  $l$  to 0. The result is depicted in figure 3(e).

## 6.2. MD processing

The Foerster Minex 2FD is a dual-frequency, continuous-wave, metal detector. When the detector moves over a metal object, a zero crossing occurs in its signal. The larger the signal, the more metal is present (or closer it is to the detector). The MD processing detects zero crossings and converts them to confidence levels in the confidence map of the sensor. With this detector it is difficult to detect small metal objects which are close to a large metal object, because the weak signals of the small objects are superimposed on the strong signal of a large object. In such areas where a strong signal of a large metal object is present, the detector is to some degree blind to other small metal objects. To compensate for this, areas with a strong signal are set to a non-zero confidence level.



**Figure 4.** Infrared processing steps: (a) resampling (b) blob search (c) local contrast enhancement (d) maximum. (e) MD confidence map. (f) Ground truth (depth legend as for result of depth processing by GPR). Note that the TIR processing steps, MD confidence map, and the ground truth are given for the entire test lane; evaluation of sensors and sensor fusion is, however, done for a subset of the lane for which there is valid GPR data.

## 6.3. TIR processing

In this section we describe the processing<sup>15</sup> for the infrared data. The method relies on the principle that the landmines have a higher (apparent) temperature than their surrounding background. This is reflected in a higher intensity in the infrared image. In Figure 4(a) the infrared image data is shown for the test lane. One can clearly distinguish some of the mines. Note that depending on the demining scenario (type of soil, time of day) the landmines can also have a lower apparent temperature. In that case, one has to make use of negative contrast.

As there is a correlation between apparent temperature and landmines an obvious choice for the conversion of sensor data to confidence levels is to use the raw infrared data as confidence levels. However, global thresholding does not take into account the changing surroundings of landmines. As such, our proposed method applies local contrast enhancement to make it invariant for the local background intensity. Furthermore, global thresholding would not reduce the number and size of false alarms. Our proposed method performs false-alarm reduction by selecting blobs on morphological and size attributes.

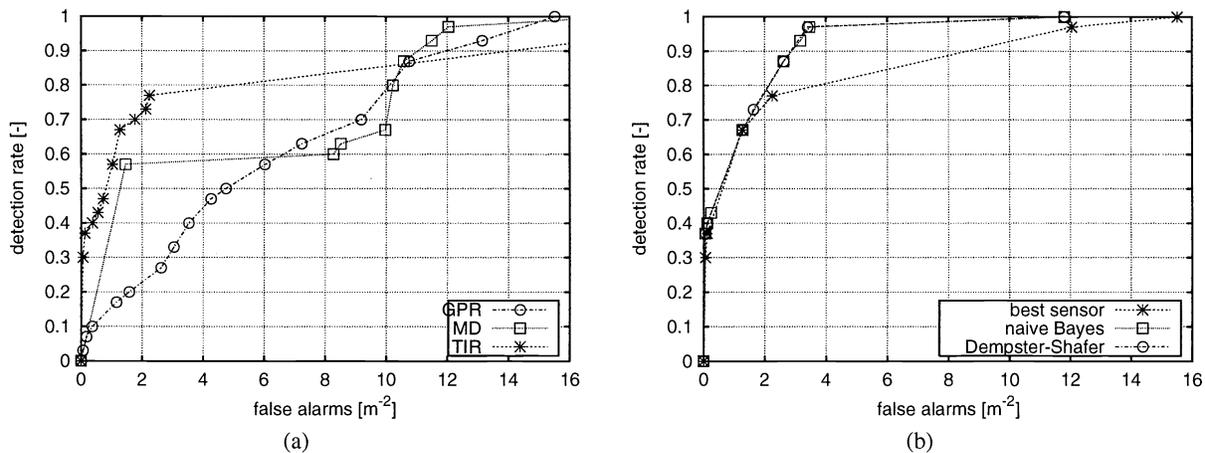
Because the spatial resolution of the infrared camera does not match the resolution imposed by the  $2.5 \times 2.5$  cm grid cells for sensor fusion, the infrared data must be resampled. Our method is a sequence of resampling, blob search and contrast enhancement. It consists of the following processing steps:

1. **resampling** : reduce the size of the infrared image to a defined 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  $l$  (minimum) and  $u \times u$  (maximum). Optimal  $l$  and  $u$  on the data are found to be 9 and 11 grid cells 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 4(a) to 4(d). The ROC for landmine detection using an infrared camera can then be calculated using the above-mentioned SCOOP algorithm, the result is shown in Figure 5(a). The ROC expresses landmine detection as a function of false alarms for different global thresholds on the grid with confidence levels, i.e. infrared image data.

#### 6.4. Single-sensor and sensor-fusion results

Figure 5(a) shows the ROC curves for landmine detection with single sensors (without sensor fusion) on the test lane. For lower number of false alarms, the TIR sensor is superior to both the GPR and MD. However, the TIR sensor is not able to detect all landmines as both GPR and MD can for higher number of false alarms. Furthermore, the MD performs better than the GPR for this test lane because it has a good working point with 57% detections and  $1.5m^{-2}$  false alarms. This point probably corresponds with detection of all landmines that contain metal.



**Figure 5.** (a) Single-sensor results for GPR, MD and TIR. (b) Sensor-fusion results for best sensor, naive Bayes and Dempster-Shafer.

Figure 5(b) presents the ROC curves of three decision-level sensor-fusion methods: best sensor, naive Bayes, and Dempster-Shafer. All fusion methods combine the GPR, MD, and TIR sensors on the test lane. From the results we may conclude that naive Bayes and Dempster-Shafer perform better than the best sensor on this test lane. Furthermore, the differences between naive Bayes and Dempster-Shafer are negligible. This probably stems from the fact that Bayes and Dempster-Shafer have discriminant functions that belong to the same class of functions.<sup>4</sup>

#### 6.5. Depth-fusion results

Figures 6 and 7 present depth-fusion results for the test lane. All depth-fusion methods combine the GPR, MD, and TIR sensors on the test lane using the depth map as calculated by the GPR. Figures 6(a), 6(b) and 7(a) show the differences between a common decision-level sensor-fusion method and its corresponding depth-fusion method. The results are shown for best sensor, naive Bayes, and Dempster-Shafer. Figure 7(b) gives a summary of these results. It is clear that for best sensor, naive

Bayes, and Dempster-Shafer the depth-fusion results are better than the results of sensor fusion with one single layer. When we compare all depth-fusion methods with all one-layer fusion methods, we see that the best sensor with depth fusion performs worse than Bayes and Dempster-Shafer with one-layer fusion.

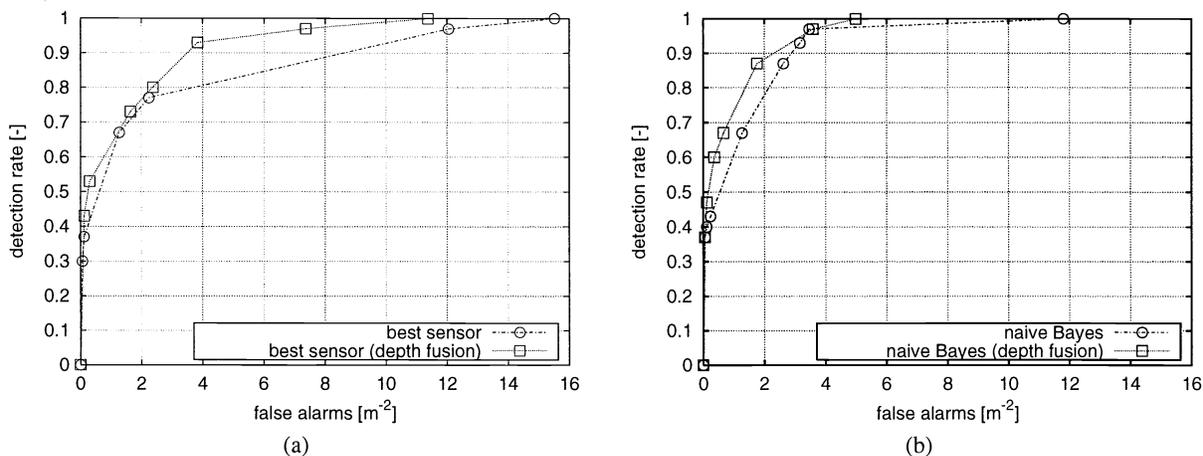


Figure 6. (a) Depth-fusion results for best sensor. (b) Depth-fusion results for naive Bayes.

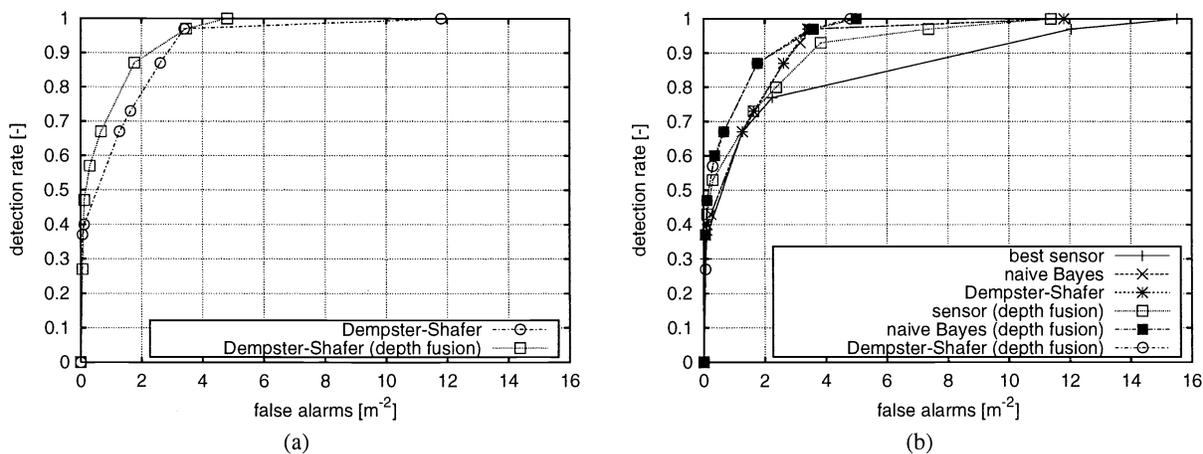


Figure 7. (c) Depth-fusion results for Dempster-Shafer. (d) Depth-fusion results.

## 7. CONCLUSIONS AND DISCUSSION

The proposed new concept of depth fusion is a valuable extension of common decision-level sensor-fusion methods. Our experiments validate the notion that a priori knowledge on sensor characteristics can be made explicit to improve fusion results. Furthermore, the experiments show that depth fusion performs better than normal (decision-level) fusion for all employed decision-level sensor-fusion methods. The ranking among depth-fusion methods is: (1) naive Bayes/Dempster-Shafer, (2) best sensor.

Notwithstanding these results, note that the results were obtained by optimizing and evaluating the sensor processing and fusion on the same data set. This biases the results optimistically. On the other hand, similar data sets<sup>4</sup> show that the differences are small for naive Bayes and Dempster-Shafer when the results are compared with results obtained with a leave-one-out<sup>17</sup> evaluation method. However, in this case we have four times (layers) as much optimization parameters. That could bias the results more optimistically than what we saw in earlier work.<sup>4</sup>

That said, there is still room for alterations in our depth-fusion concept to improve depth-fusion results. Improvements we foresee to explore in future work are the following:

- application of a GPR with a higher bandwidth and a better timing;
- determination of the depth map by means of all sensors;
- optimal choice of depth layers (number and size);
- specialized sensor processing for each depth layer.

Regarding humanitarian demining, the sensor-fusion results still should be improved for an operational multi-sensor system for AP landmine detection. Therefore, future work should focus on improvement of sensors, their processing and on sensor fusion. For example, improvement of sensors can be achieved by using additional imaging techniques. One can think of polarization of infrared and visual imagery<sup>12</sup> or the utilization of a multi-spectral imager.<sup>18</sup>

## ACKNOWLEDGMENTS

This research is partly funded by the European Commission as ESPRIT project LOTUS, number 29812.

## REFERENCES

1. 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.
2. 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.
3. G. A. Clark, J. E. Hernandez, N. K. DelGrande, R. J. Sherwood, S.-Y. Lu, P. C. Schaich, and P. Durbi. Computer vision for locating buried objects. In *Proceedings of the Twenty-Fifth Annual Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove (CA), USA, Nov. 1991.
4. F. Cremer, K. Schutte, J. G. M. Schavemaker, and E. den Breejen. A comparison of decision-level sensor-fusion methods for anti-personnel landmine detection. *Information fusion*. accepted for publication.
5. 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.
6. F. Cremer, J. G. M. Schavemaker, E. den Breejen, and K. Schutte. Detection of anti-personnel land-mines 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.
7. 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, J. F. Harvey, J. T. Broach, and R. E. Dugan, editors, *Proc. SPIE Vol. 4038, Detection and Remediation Technologies for Mines and Minelike Targets V*, pages 792–803, Orlando (FL), USA, Apr. 2000.
8. A. P. Dempster. Upper and lower probabilities induced by a multi-valued mapping. *Annals of Mathematical Statistics*, 38:325–339, 1967.
9. S. L. Earp, E. S. Hughes, T. J. Elkins, and R. Vickers. Ultra-wideband ground-penetrating radar for the detection of buried metallic mines. *IEEE AES Systems Magazine*, pages 30–34, Sept. 1996.
10. K. Fukunaga. *Introduction to statistical pattern recognition*. Academic press, Inc., Boston, USA, 2 edition, 1990.
11. W. de Jong, H. A. Lensen, and Y. H. L. Janssen. Sophisticated test facility to detect 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 1409–1419, Orlando (FL), USA, Apr. 1999.
12. W. de Jong, F. Cremer, K. Schutte, and J. Storm. Usage of polarisation features of landmines for improved automatic detection. In A. C. Dubey, J. F. Harvey, J. T. Broach, and R. E. Dugan, editors, *Proc. SPIE Vol. 4038, Detection and Remediation Technologies for Mines and Minelike Targets V*, pages 241–252, Orlando (FL), USA, Apr. 2000.
13. N. A. Macmillan and C. D. Creelman. *Detection theory: a user's guide*. Cambridge university press, Cambridge (UK), 1991.

14. N. Milisavljevic, S. P. van den Broek, I. Bloch, P. B. W. Schwering, H. A. Lensen, and M. Acheroy. Comparison of belief functions and voting method for fusion of mine detection sensors. In A. C. Dubey, J. F. Harvey, J. T. Broach, and V. George, editors, *Proc. SPIE Vol. 4394, Detection and Remediation Technologies for Mines and Minelike Targets VI*, Orlando (FL), USA, Apr. 2001.
15. J. G. M. Schavemaker, F. Cremer, K. Schutte, and E. den Breejen. Infrared processing and sensor fusion for anti-personnel land-mine detection. In *Proceedings of IEEE Student Branch Eindhoven: Symposium Imaging*, pages 61–71, Eindhoven, the Netherlands, May 2000.
16. E. Waltz and J. Llinas. *Multisensor datafusion*. Artech House, 1990.
17. S. M. Weiss and C. A. Kulikowski. *Computer systems that learn*. Morgan Kaufmann Publishers, 1991.
18. P. J. Withagen, E. den Breejen, E. M. Franken, A. N. de Jong, and H. Winkel. Band selection from a hyperspectral data-cube for a real-time multispectral 3ccd camera. In S. S. Shen and M. R. Descour, editors, *Proc. SPIE Vol. 4381, Algorithms for Multispectral, Hyperspectral, and Ultraspectral Imagery VII*, Orlando (FL), USA, Apr. 2001.