Int. Conf. Requirements and Technologies for the Detection, Removal and Neutralization of Landmines and UXO
15-18 September 2003
VUB, Brussels, Belgium

Feature level fusion of polarimetric infrared and GPR data for landmine detection

F. Cremer, W. de Jong, K. Schutte

TNO Physics and Electronics Laboratory
P.O. Box 96864, 2509 JG, The Hague
The Netherlands
E-mail: FCremer@etro.vub.ac.be

A.G. Yarovoy, V. Kovalenko
Faculty of Information Technology and Systems
Delft University of Technology
Mekelweg 4, 2628 CD Delft
The Netherlands
E-mail: A.Yarovoy@its.tudelft.nl

R.F. Bloemenkamp
Department of Applied Earth Sciences
Delft University of Technology

Abstract

Feature-level sensor fusion is the process where specific information (i.e. features) from objects detected by different sensors are combined and classified. This paper focuses on the feature-level fusion procedure for a sensor combination consisting of a polarimetric infrared (IR) imaging sensor and a GPR: a video impulse radar (VIR).

The single sensor detection methods and the feature-level sensor-fusion methods are evaluated. The detection results of both single sensors and the sensor-fusion methods are presented in receiver operator characteristics (ROC) curves. They show that on the training set feature-level sensor-fusion always outperforms the best single sensor. Furthermore, on the independent evaluation set there are ROC points of the feature-level sensor-fusion methods that are better than the best sensor.

1 Introduction

Sensor fusion is the process in which information from different sensors is used for a unified declaration of objects as detected by these sensors. Commonly a distinction is made between three different levels of sensor fusion: data-level fusion, feature-level fusion and decision-level fusion [1].

This paper describes a feature-level fusion procedure for a polarimetric IR system and a GPR system and shows some preliminary results.

The measurements with the polarimetric IR system and the VIR have been performed at the TNO-FEL test facility for humanitarian demining [2] in July 2002.

In Sec. 2 the polarimetric IR measurements and classification results are discussed. Sec. 3 describes the extraction of features in the video impulse radar followed by the single sensor classification results. The object associations and sensor-fusion methods are described in 4. The results and a discussion of the performance are given in 5. Finally the conclusions are given and discussed in 6.

2 Polarimetric infrared measurements and features

Polarimetric IR is measured using a rotating filter setup, consisting of a Radiance HS mid-wave IR camera, a wire-grid linear polarisation filter and a custom built rotational setup [3].

Special pre-processing algorithms, which take into account the motion of the platform and bending of the ray by the filter, result into 3 independent Stokes images [4]. One of the 3 Stokes images, the intensity image I, is input for the region selection algorithm. This region selection algorithm selects regions using a tophat filter. For the selected regions the features are measured in the feature extraction process. The 10 features of each object that are measured and that are available for feature-level sensor fusion are: mean values of I, Q and U, contrasts between the object and the background for the values of I, Q, U, the area of the object, the length of the major axis, the fraction of the area and the convex area of the object, and the fraction of the minor axis and major axis.

By means of exhaustive search over all possible feature combinations ($2^{10} - 1$ combinations), the best feature selection for the training set is made. This optimisation is performed for every point in the ROC.

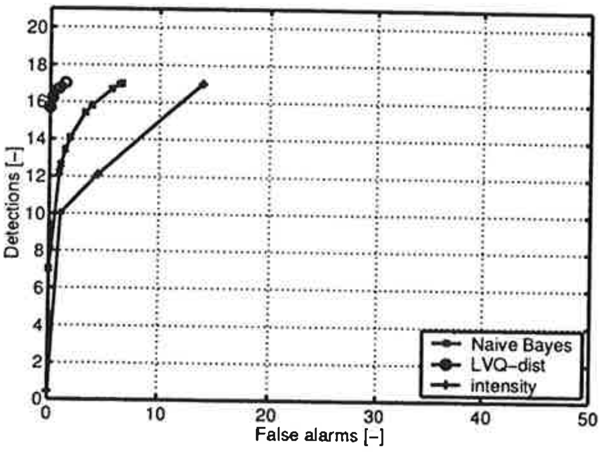

(a) Training set

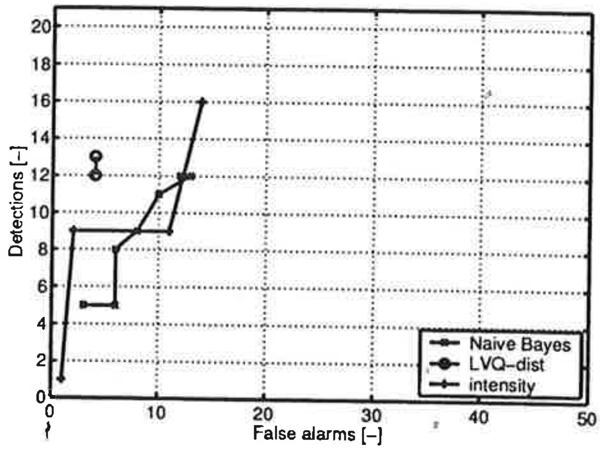

(b) Evaluation set

Figure 1: The detection and classi£cation results of the polarimetric infrared system on the training set (a) and on an independent evaluation set (b) using leave-one-out as evaluation method. Boths sets contain 21 landmines and 20 m^2.

Two different classi£ers have been trained on this data set: Naive Bayes and LVQ-dist, see Sec. 4.3. These classi£ers along with a threshold on the intensity are trained and evaluated using the leave-one-out evaluation method, see Sec. 4.4. The classi£cation results are shown in an ROC curve in Fig. 1. To re□ect the size of the training and evaluation set, absolute numbers are used in the ROC curves that are shown in this paper. On the training set a large improvement is shown for the two classi£ers (Naive Bayes and LVQ-dist) compared to using only a threshold on the intensity.

3 Video Impulse Radar measurements and features

GPR data has been acquired with a video impulse radar developed in IRCTR especially for landmine detection [5]. Only part of the data set, which corresponds to a quasi-monostatic co-polar scattering geometry (single transmit - single receive polarization), has been used for fusion purposes. This data has been pre-processed according to a diffraction stack algorithm [6]. As a result a focused three dimensional image (two horizontal coordinates and estimated depth) of the subsurface has been obtained. Unlike IR images, GPR images contain both amplitude and phase information of the received electromagnetic £eld.

Regions of interest are selected based on thresholding the energy projected to the 2-D surface. For these regions a number of features that may be suitable for classi£cation can be extracted from GPR data (e.g., 3D object position, object shape, dielectric permittivity of the object). Extraction of some of these features requires considerable processing.

For this paper, only the most basic and easy to measure features are used namely: position, average depth, spatial variation of depth, average projected energy, spatial variation of projected energy, average highest positive signal, spatial variation of highest positive signal, area of the object, major axis of the object, fraction of the area and the convex area, and fraction of the minor and major axis.

The same classi£ers and evaluation methods as for the polarimetric IR system have been used to determine the single sensor performance of the VIR see Sec. 2. The classi£cation results are shown in Fig. 2. For the training set the best features for each ROC point are determined using an exhaustive search ($2^{10} - 1$ combinations).

4 Sensor-fusion process

The feature-level sensor fusion process starts with the regions of interest with their features as measured by the individual sensors (see Sec. 2 and Sec. 3) and consists of three steps. The £rst step is object association. In this step the features from the objects from the different sensors are combined to form an associated object. This step is further discussed in Sec. 4.1. Which features are used is discussed in Sec. 4.2. The second step is to classify the features from these associated objects. This is performed by the feature classi£cation algorithms as discussed in Sec. 4.3. The third and £nal step is the performance evaluation. In this step it is discussed how the false alarm rate and the detection rate are determined, see Sec. 4.4.

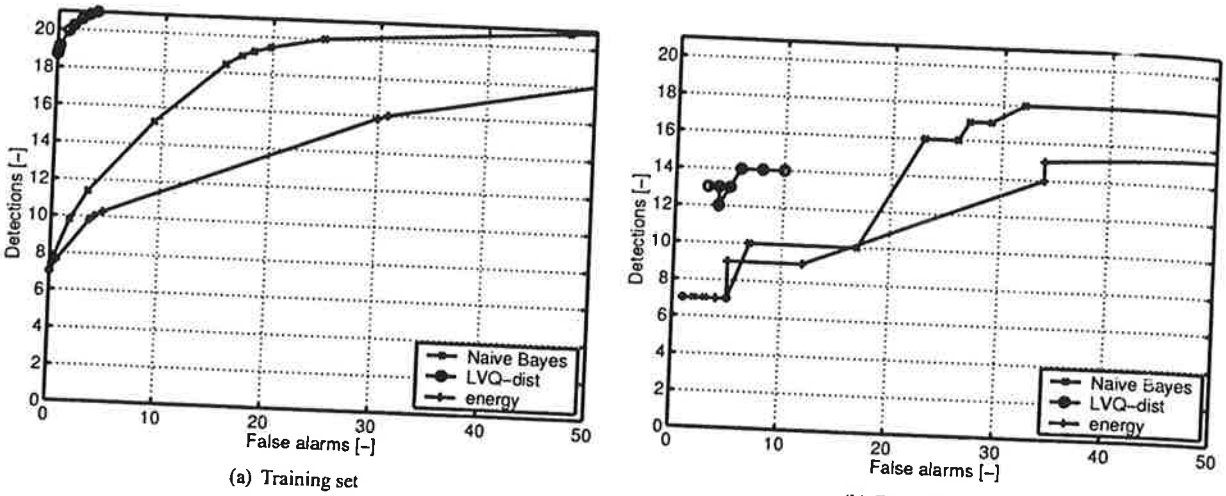

(a) Training set

(b) Evaluation set

Figure 2: The detection and classi£cation results of the VIR on the training set (a) and on an independent evaluation set (b) using leave-one-out as evaluation method. Boths sets contain 21 landmines and 20 m^2.

4.1 Object association

Object association is the most important part of feature-level sensor-fusion. The other remaining steps are general steps for feature-based classi£cation and detection and are not speci£c for feature-level fusion.

A simple object association algorithm is used in our fusion process. For each object from one sensor, the object from the other sensor that is closest in distance is found. If this distance is within its maximum bound (smaller than 5.5 cm) then the two objects are associated.

With this form of object association there is always a maximum of two associated sensor objects. This one-to-one object association avoids ambiguity that might arise when a one-to-many object association is allowed. The disadvantage is that there is a chance that not the right objects are associated.

By including more information the object association may be improved. One source of information may be the object depth [7], but this requires a different object association approach. For other features it is not easily determined how and if they have some correlation between sensors. Therefore, it was chosen to keep the object assocation simple and base it solely on the Euclidian distance.

4.2 Features for fusion

In principle feature-level fusion can handle all features from both sensors. Furthermore two additional features can be de£ned: the number of observations for each sensor. For the polarimetric infrared processing there are more observations of the same object, since one object can be detected in more than one image [4]. However when there is a VIR object that has not been associated to a polarimetric IR object then the number of IR observations is zero. For a polarimetric infrared object the number of VIR objects is only one or zero.

By taking all possible features from both sensors a feature vector of length 22 is reached. An exhaustive search over all feature combinations would require $2^{22} - 1$ (4 million) evaluations and is not practical feasible (this will take over 6 months of calculation).

We have selected only 5 of the 10 features of each sensor in order to keep the calculation time practical. These 5 features are the ones that are used most extensively for the single sensor classi£ers. For polarimetric infrared these features are: the level above tophat threshold, the area, the average Q value, the average U value and the contrast in I. For the VIR these features are: the variation of depth, the variation in energy, the average energy, the depth and the variation of highest positive value.

4.3 Classi£ers

Two classi£ers have been used for feature classi£cation, see also [4]. The £rst classi£er is Naive Bayes also known as Bayes plug-in classi£er [8]. This classi£er works under the assumption that the features are independent. For both classes (landmines and background) it estimates the mean value and the standard deviation of each feature and consequently returns an equivalent likelihood ratio. Due to the assumption of independence of features, it is justi£able to have accurate estimates of the mean and standard deviation for each class. However, some of the features are not likely to be independent; therefore the classi£er is called

Naive Bayes and may produce suboptimal results.

The second classi£er is the LVQ-dist classi£er. It is an extension based on the learning vector quantisation (LVQ) algorithm [9] and uses the distance to the closest vector in each class to determine a measure for classi£cation. This classi£er can model the probability distribution more accurately, but may have not enough data to correctly estimate the parameters.

4.4 Performance evaluation

A leave-one-out method [10] has been used for the performance evaluation. In a leave-one-out evaluation method the classi£er is trained on all but one sample and tested on the remaining sample. This process is repeated until all samples have been part of the evaluation set. The training set results are averaged over all possible training sets and the evaluation set results are summed up over all evaluation sets.

The sample size we used is not a single object, but one landmine and its surrounding area [4]. Furthermore leave-one-out by itself does not have a way to generate ROC curves. A modi£ed approach that uses a range of cost functions solves this problem [4].

The training set is small with only 21 landmines and 20 m^2 area. The number of potential false alarms is high (1485 for the polarimetric IR and 2616 for the VIR). However, these false alarms originate from test lanes that have a low amount of clutter. In a real mine£eld there are many more sources of clutter. In this respect the sensor-fusion performance should be related to the single sensor performance, as presented in Fig. 1 and Fig. 2, and not seen as measures of detection. To re□ect this all curves show the absolute number of detected landmines and false alarms and not percentages (implying detection measures).

5 Results

The objects generated by the polarimetric infrared sensor and the VIR are associated using the algorithm described in Sec. 4.1. Additional measurements of features are performed for objects that are present at one speci£c location for one sensor but not for the other. Using this form of association a complete joint data set is generated.

This joint data set has been classi£ed by the same two feature classi£cation methods that have also been for the single sensor results in Sec. 2 and 3. The feature-level fusion results are presented in Fig. 3.

The LVQ-dist classi£er performs best on the training set. It reaches the ideal ROC point where it detects all landmines with zero false alarms. This point is obviously better than the two single sensor results with the LVQ-dist classi£er. Since there is only one point in the training set (all cost functions naturally select this point), there is only one point in the evaluation set. This point in the evaluation set detects the most landmines of all LVQ-dist evaluation set results. However, since the number of false alarms is also higher, it depends on the cost function whether or not this ROC point is indeed better than the single sensor LVQ-dist results.

The Naive Bayes classi£er performs on the training set also better than the same classi£er on either single sensor. On the evaluation set it performs better than the polarimetric infrared sensor and for some part better than the VIR. A very remarkable point on the evaluation set is the point where all 21 landmines are detected with 19 false alarms (The VIR started with 2612 and the polarimetric IR with 1485 false alarms at the point where all mines were detected). Compared to the training set there is only a small increase in the number of false alarms, but it still detects all landmines.

6 Conclusions

A basic feature-level sensor-fusion implementation has been applied to polarimetric IR data and quasi-monostatic co-polarized GPR data. Basic steps in data pre-processing and feature extraction for each sensor have been brie□y addressed. Implementation of the fusion algorithm and its object association has been described. The feature-level sensor-fusion methods have been compared to the single sensor results. Based on the results, we conclude that sensor fusion is useful for the training set, since for both classi£ers the fusion results are always better than the single sensor results with the same classi£er. Furthermore a couple of sensor-fusion ROC points in the evaluation set are better than the single sensor evaluation set results.

The best classi£er is the LVQ-dist classi£er, but it seems to have a larger difference between the training and classi£cation set results.

In short we conclude that we have described and implemented a procedure for performing feature-level sensor fusion that proves to be working on a limited data set.

References

[1] E. Waltz and J. Llinas, *Multisensor datafusion*. Artech House, 1990.

[2] W. de Jong, H. A. Lensen, and Y. H. L. Janssen, "Sophisticated test facility to detect land mines," in *Proc. SPIE Vol. 3710, Detection and Remediation Technologies for Mines and Minelike Targets IV*, A. C.

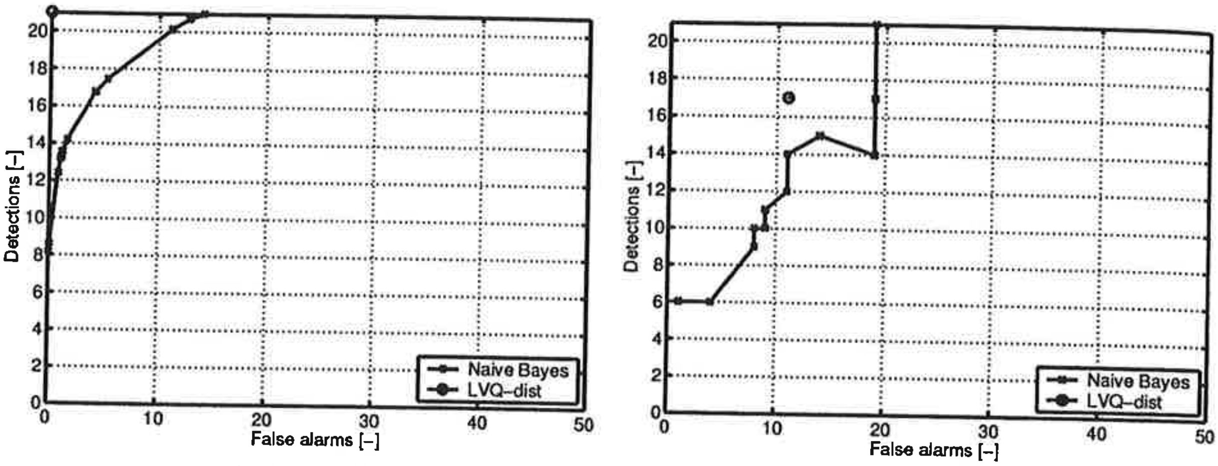

(a) Training set

(b) Evaluation set

Figure 3: The sensor-fusion results of the VIR and the polarimetric infrared system on the training set (a) and on an independent evaluation set (b) using leave-one-out as evaluation method. Boths sets contain 21 landmines and 20 m^2.

Dubey and J. F. Harvey, Eds., Orlando (FL), USA, Apr. 1999, pp. 1409–1418.

[3] F. Cremer, W. de Jong, and K. Schutte, "Infrared polarisation measurements and modelling applied to surface laid anti-personnel landmines," *Optical Engineering*, vol. 41, no. 5, pp. 1021–1032, May 2002.

[4] F. Cremer, J. G. M. Schavemaker, W. de Jong, and K. Schutte, "Comparison of vehicle-mounted forward looking polarimetric infrared and downward-looking infrared sensors for landmine detection," in *Proc. SPIE Vol. 5089, Detection and Remediation Technologies for Mines and Minelike Targets VIII*, R. S. Harmon, J. T. Broach, and J. H. Holloway, Eds., Orlando (FL), USA, Apr. 2003.

[5] A. G. Yarovoy, L. P. Ligthart, A. Schukin, and I. Kaploun, "Polarimetric video impulse radar for landmine detection," *Subsurface Sensing Technologies and Applications*, vol. 3, no. 4, pp. 271–293, Oct. 2002.

[6] J. Groenenboom and A. G. Yarovoy, "Data processing and imaging in GPR system dedicated for landmine detection," *Subsurface Sensing Technologies and Applications*, vol. 3, no. 4, pp. 387–402, Oct. 2002.

[7] J. G. M. Schavemaker, E. den Breejen, F. Cremer, K. Schutte, and K. W. Benoist, "Depth fusion for anti-personnel landmine detection," in *Proc. SPIE Vol. 4394, Detection and Remediation Technologies for Mines and Minelike Targets VI*, A. C. Dubey, J. F. Harvey, J. T. Broach, and V. George, Eds., Orlando (FL), USA, Apr. 2001, pp. 1071–1081.

[8] A. K. Jain, R. P. W. Duin, and J. Mao, "Statistical pattern recognition: a review," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 22, no. 1, pp. 4–37, Jan. 2000.

[9] J. Laaksonen and E. Oja, "Classi£cation with learning k-nearest neighbors," in *Proceedings of ICNN'96*, Washington DC, USA, June 1996, pp. 1480–1483.

[10] S. M. Weiss and C. A. Kulikowski, *Computer systems that learn*. Morgan Kaufmann Publishers, 1991.

H. Sahli
A. M. Bottoms
J. Cornelis (Eds.)

EUDEM2-SCOT - 2003

International Conference on Requirements and Technologies for the Detection, Removal and Neutralization of Landmines and UXO

15-18 September 2003 - Vrije Universiteit Brussel, Brussels, Belgium, Proceedings

Volume - 2

SCOT
Society for Counter Ordnance Technology