Multi-level Data Fusion for the Detection of Targets using multi-spectral Image Sequences

D. Borghys, P. Verlinde, C. Perneel and M. Acheroy

Royal Military Academy, Signal & Image Centre,
Renaissancelaan 30, 1000 Brussels, Belgium

ABSTRACT

This paper presents an approach to the long range automatic detection of vehicles, using multi-sensor image sequences. The method was tested on a database of multi-spectral image sequences, acquired under diverse operational conditions. The approach consists of two parts.

The first part uses a semi-supervised approach, based on texture parameters, for detecting stationary targets. For each type of sensor one learning image is chosen. Texture parameters are calculated at each pixel of the learning images and combined using logistic regression into a value that represents the conditional probability that the pixel belongs to a target given the texture parameters. The actual detection algorithm applies the same combination to the texture features calculated on the remainder of the database (test images). When the results of this feature level fusion are stored as an image, the local maxima correspond to likely target positions. These feature-level-fused images are calculated for each sensor. In a sensor fusion step the results obtained per sensor are then combined again. Region growing around the local maxima is then used to detect the targets.

The second part of the algorithm searches for moving targets. In order to detect moving vehicles, any motion of the sensor needs to be detected first. If sensor motion is detected, it is estimated using a Markov Random Field model. Available prior knowledge about the sensor motion is used to simplify the motion estimation. The estimate is used to warp past images onto the current one in a temporal fusion approach and moving targets are detected by thresholding the difference between the original and warped images.

Decision level fusion combines the results from both parts of the algorithm.

Keywords: Sensor Fusion, Decision Fusion, Feature Fusion, Temporal Fusion, Target Detection, Texture, Motion Estimation

1. INTRODUCTION

Long range automatic detection of vehicles is of great military importance to modern armed forces. The most critical factor of any system for automatic detection is its ability to find an acceptable compromise between the probability of detection and the number of false targets. A lot of work has already been carried out on the detection of single vehicles and target formations.[1–3] However, detection and tracking of small, low contrast vehicles in a highly cluttered environment using a single sensor, still remains a very difficult task.

This paper describes an approach to tackle this problem using data fusion at different levels. The approach was implemented and then tested on a set of six image sequences obtained from different sensors under diverse operational circumstances. The target area in the images varies from small, typically less than 20 pixels-on-target, up to 2000 pixels-on-target. The vehicles can be either moving or stationary. In most of the sequences, the sensor was mounted on a stationary platform and could only perform a tilt and pan operation. Most of the images are highly cluttered. This clutter is caused by sensor noise, natural background texture and the presence of human artifacts in the scene (e.g. buildings).

The approach presented in this paper consists of two independent parts. The next section presents the image database, then an overview of the method is presented. The subsequent sections describe the two parts of the algorithm. The last two sections show results and conclusions.

2. IMAGE DATABASE

For the development and testing of the algorithm, a database of 6 multi-spectral image sequences, numbered MS01 to MS06 *, was compiled. Table 1 presents some properties of the sequences. In figure 1 some typical images are shown.

Sequence	Targets	Target Motion	Sensors
MS01	Helicopter Tow truck	Across Field of View (FOV)	Infrared (LW) Visual (B/W)
MS02	Truck	Toward sensors	Infrared (LW) Visual (B/W)
MS03	2 Armoured Vehicles	Stationary	Infrared (LW) Infrared (SW) Visual (B/W)
MS04	2 Armoured Vehicles	Stationary	Infrared (LW) Infrared (SW) Visual (B/W)
MS05	Helicopter	Across FOV	Infrared (LW) Visual (Red, Green, Blue)
MS06	Jeep Armoured Vehicle	Across FOV	Infrared (LW) Visual (Red, Green, Blue)

Table 1. Overview of the image sequence database

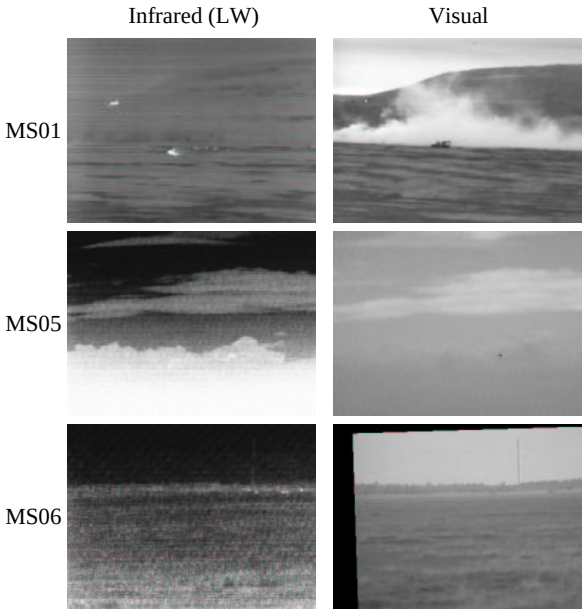

Figure 1. Image Sequence Database

*MS01-MS02: Courtesy of Defense Research Establishment Valcartier, Quebec, Canada
MS03-MS04: Courtesy of Naval Air Warfare Center, China Lake, US
MS05-MS06: Courtesy of ASIAT-DTT, Peutie, Belgium

3. OVERVIEW OF THE APPROACH

Figure 2 shows an overview of the method. The proposed algorithm consists of two independent parts. The first part searches for targets in single images while the second part uses multiple subsequent images in order to specifically find moving targets.

For the first part of the algorithm we have chosen an approach based on texture feature extraction. We are not interested in explicitly modelling or measuring texture but only in augmenting the difference between targets and background. Nevertheless, these texture features are interesting because they are independent measurements of the local spatial distribution of grey values within an image and it is likely that some of these parameters will highlight the difference between targets and background. The texture parameters are even more appealing because it can easily be seen that features that are classically used for target detection such as intensity and gradient are just special cases of these texture parameters. Feature level fusion is used to combine the texture features from each image into a new image in which the grey value at each pixel is proportional to the probability that the pixel belongs to the target. These images from the different sensors are fused in a sensor fusion step.

The second part of the algorithm detects moving targets in sub-sequences from each sensor separately.

Each part of the algorithm behaves as an expert indicating the possible presence of vehicles in the scene. Decision fusion is used to combine the outcomes from all experts.

Figure 2 shows the global overview of the method.

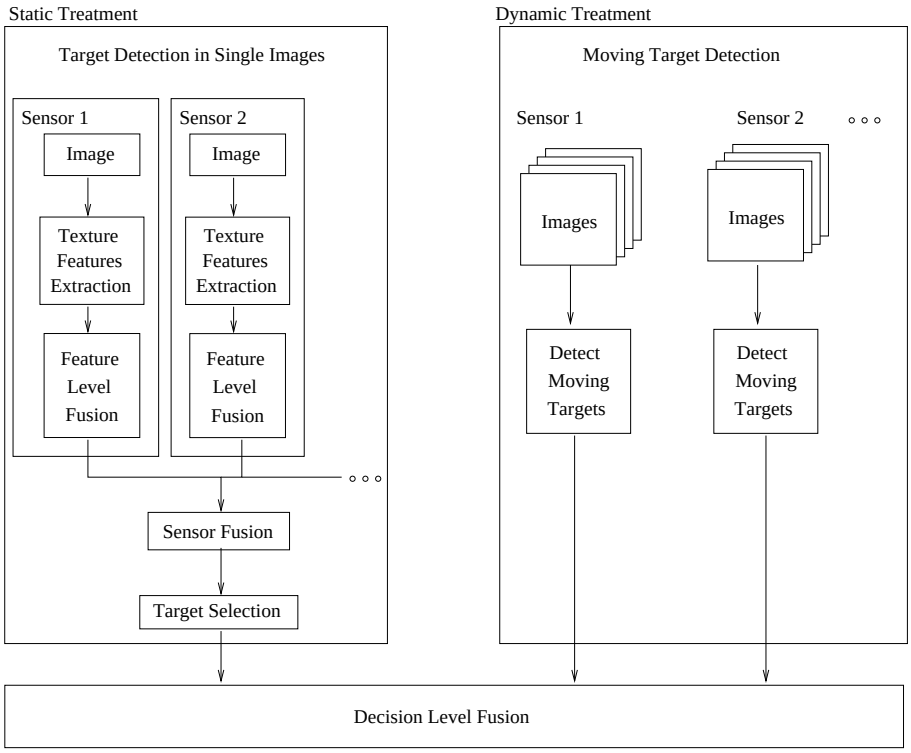

Figure 2. Overview of the approach

4. TARGET DETECTION IN SINGLE IMAGES (TDSI MODULE)

4.1. Introduction

For the detection of targets in single images, a semi-supervised approach based on texture features was chosen. For each sensor type, one image was selected to constitute the learning database. On these images the true targets were delimited. Then several texture parameters were calculated at each pixel of these learning images and logistic

regression[8] was used to find a combination of the texture parameters that is proportional to the probability of finding a target at the corresponding image location.

The actual detection algorithm then applies the same combination to the texture features calculated on the remainder of the image database (test images). When this function is applied to the texture features calculated at each pixel of a test image, a new image, called feature-level-fused image, can be formed in which the maxima correspond to likely target positions. These feature-level-fused images, obtained from all the different sensors, are then fused again in a subsequent sensor fusion step.

To find the possible target positions, first the local maxima are determined in this sensor-fused image and then available prior knowledge about possible target size and aspect ratio is used to reject false targets.

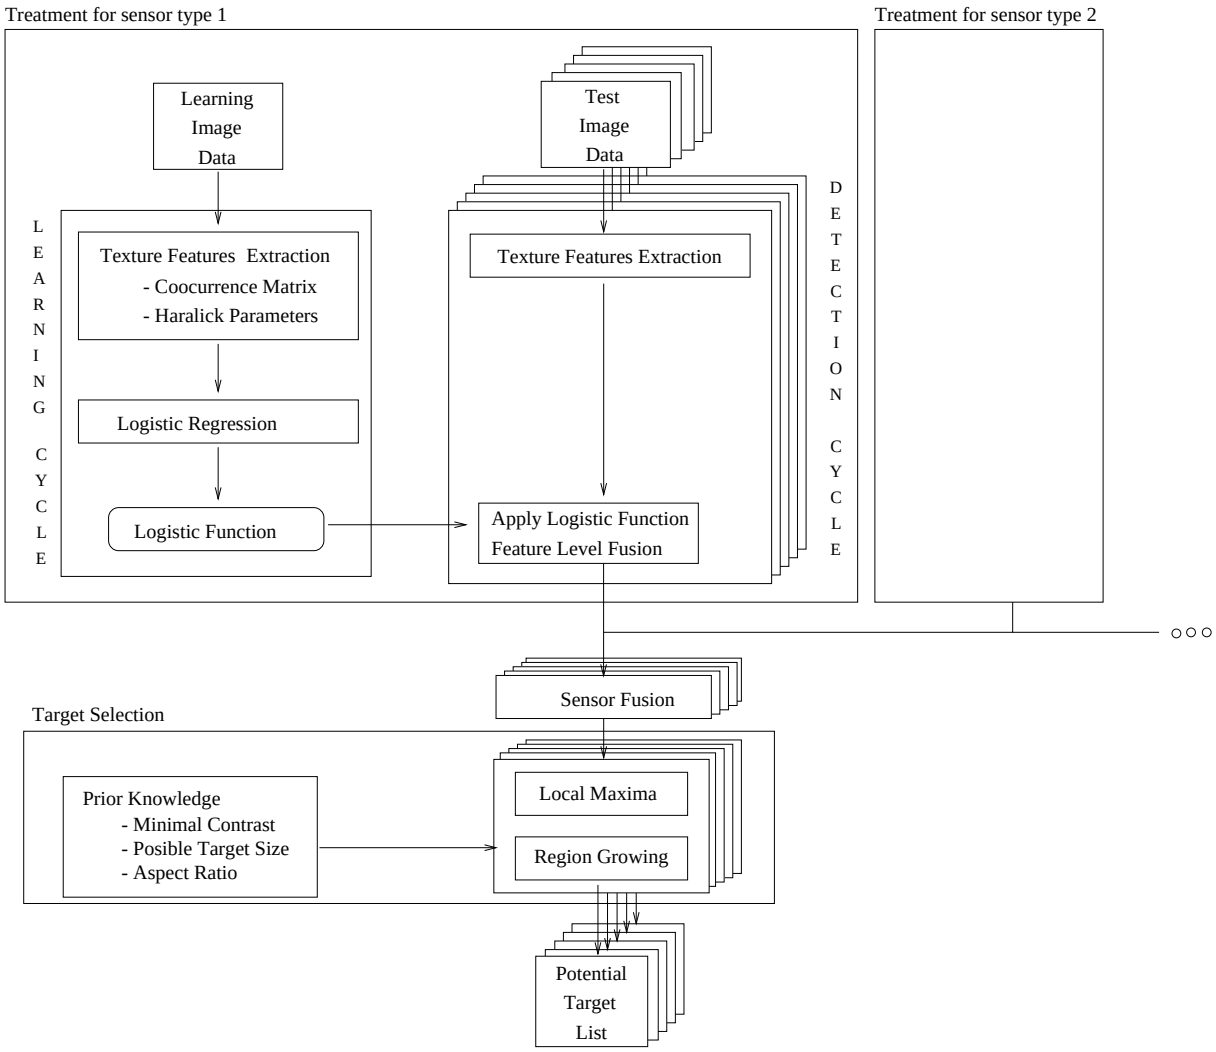

Figure 3. TDSI module

4.2. Texture parameters

The calculation of the texture features is based on the co-occurrence matrix. The co-occurrence matrix is defined as a function of a given direction and distance, or alternatively, as a function of a displacement (dx,dy) along the x and y direction in the image. For a given displacement (dx,dy), the (i, j) element of the co-occurrence matrix is the number of times the grey value G at the current position (x,y) is i when the value at the distant position (x+dx,y+dy) is j.

$$C^{dx,dy}(i,j) = \#(G(x,y) = i \mid G(x+dx, y+dy) = j) \tag{1}$$

The co-occurrence matrix can be calculated on the whole image. However, by calculating it in a small window scanning the image, a co-occurrence matrix can be associated with each image position. The centre of the window is denoted (x_c, y_c) and the corresponding co-occurrence matrix is $C^{dx,dy}_{x_c,y_c}(i,j)$ In figure 4 an example of a co-occurrence matrix is shown. The matrix corresponds to the small window of the image on the left and was calculated for a displacement of $dx = 1, dy = 2$.

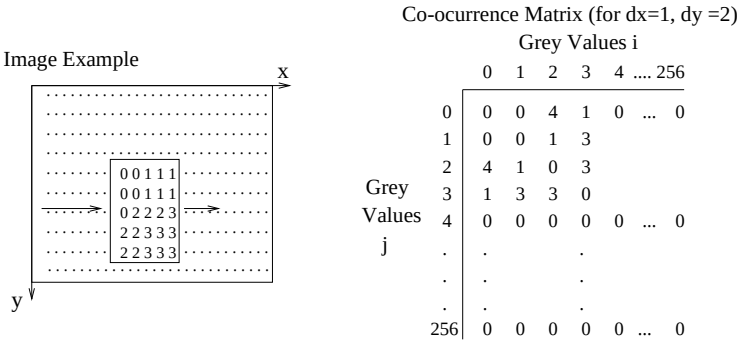

Figure 4. Co-occurrence matrix

The textural features that were used, were introduced by Haralick[4–6] and are widely used in texture analysis. Based on the local co-occurrence matrix, the used parameters are defined as follows:

$$\begin{aligned}
F_1(x_c,y_c) = &\quad Energy &&= \textstyle\sum_i \sum_j C^{dx,dy}_{x_c,y_c}(i,j)^2 \\
F_2(x_c,y_c) = &\quad Contrast &&= \textstyle\sum_i \sum_j \left[(i-j)^2 C^{11}_{x_c,y_c}(dx,dy)\right] \\
F_3(x_c,y_c) = &\quad Max.Prob. &&= max\ \left[C^{dx,dy}_{x_c,y_c}(i,j)\right] \\
F_4(x_c,y_c) = &\quad Entropy &&= C^{dx,dy}_{x_c,y_c}(i,j) log\left[C^{dx,dy}_{x_c,y_c}(i,j)\right] \\
F_5(x_c,y_c) = &\quad Homogeneity &&= \textstyle\sum_i \sum_j \frac{max\left[C^{dx,dy}_{x_c,y_c}(i,j)\right]}{[1+(i-j)^2]} \\
F_6(x_c,y_c) = &\quad Variance &&= \left[\textstyle\sum_i (i-E_i)^2 \sum_j C^{dx,dy}_{x_c,y_c}(i,j)\right]\left[\textstyle\sum_j (j-E_j)^2 \sum_i C^{dx,dy}_{x_c,y_c}(i,j)\right] \\
&&&\quad with\ E_i = \textstyle\sum_i i \sum_j C^{dx,dy}_{x_c,y_c}(i,j) \\
&&&\quad and\ E_j = \textstyle\sum_j j \sum_i C^{dx,dy}_{x_c,y_c}(i,j)
\end{aligned} \tag{2}$$

We are not interested in modelling or measuring texture but only in detecting a difference between target and background pixels. The "texture parameters" are only used as features of which we hope that some will highlight the difference between target and background. Because we do not intend to measure the texture within the target, the parameters are useful even for small targets and we can chose an arbitrary displacement $(dx = 1, dy = 1)$ for all calculations of the co-occurrence matrix.

The window used to calculate the local co-occurrence matrix had a size of 5×5. The results for each texture feature can be converted into an image. Figure 5 shows the texture images corresponding to the first image set (IR and VIS) of sequence MS01, which served as the learning image set.

As can be noticed in figure 5, the vehicles are clearly visible in some of the texture images. Hence the idea to combine the texture features to get an optimal discrimination between background and targets. Discriminant

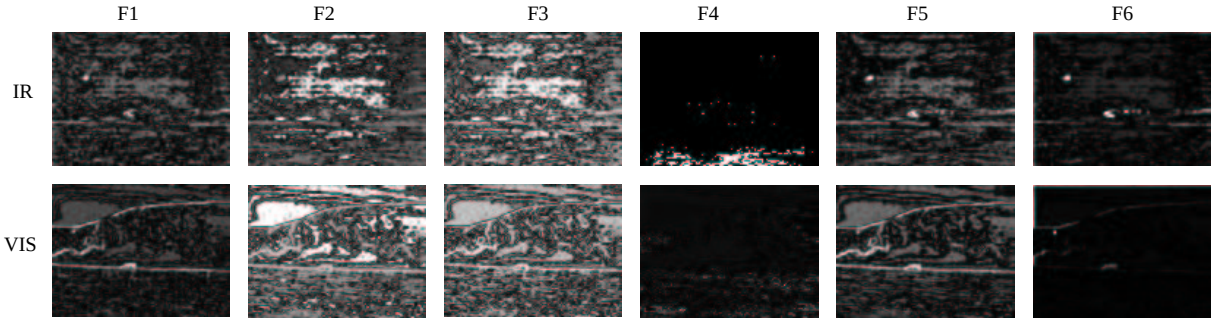

Figure 5. Texture Images

analysis can be used to obtain such a combination,[7] but it is based on assumptions that are not always verified (e.g. normal distribution with same co-variance matrix for all classes). When only two classes are involved, as is the case here (targets/background), logistic regression offers a more appropriate approach.[8]

In the learning phase, at each pixel of the learning image(s), the texture features are calculated and stored in a table. Then the human operator interactively indicates the bounding rectangles surrounding the targets in the learning image(s) and a column is automatically added to the table assigning each measurement in the table to either class 0 (background), when it corresponds to an image pixel that falls outside the bounding rectangles, or class 1 (targets) when it is inside one of the rectangles.

Logistic regression is then used to find a combination of the form :

$$p_{xy}(target \mid \mathbf{F}) = \frac{e^{b_0+\sum_i b_i F_i(x,y)}}{1+e^{b_0+\sum_i b_i F_i(x,y)}} \tag{3}$$

in which $p_{xy}(target \mid \mathbf{F})$ is the conditional probability that a pixel (x, y) belongs to the class 1 (target class) given the vector of texture parameters $\mathbf{F}$ at the given pixel. The logistic regression was carried out using Wald's forward method. In this method, at each step, the most discriminant feature is added and the significance of adding it to the model is verified. This means that not all feature will necessarily be included into the model. Table 2 shows the results for the visual and infrared learning images.

Coefficient	Feature	IR type sensors	VIS type sensors
B_0	Constant Term	-15.478	-17.927
B_1	Energy	2.356	0
B_2	Contrast	0	0
B_3	Max. Prob.	0	4.733
B_4	Entropy	4.664	5.604
B_5	Homogeneity	0	2.604
B_6	Variance	0.0107	0

Table 2. Coefficients from the logistic regression

4.3. Feature level fusion

If the learning images are representative for the images of a given sensor type, the most discriminating features for each sensor will have the highest weights b_i. Therefore, when using the same weights (table 2) to combine the feature images of the remainder of the database into new images using equation 3, targets will appear as local maxima. This is the feature level fusion. Results of applying the feature-level fusion to the images of figure 1 are shown in figure 6.

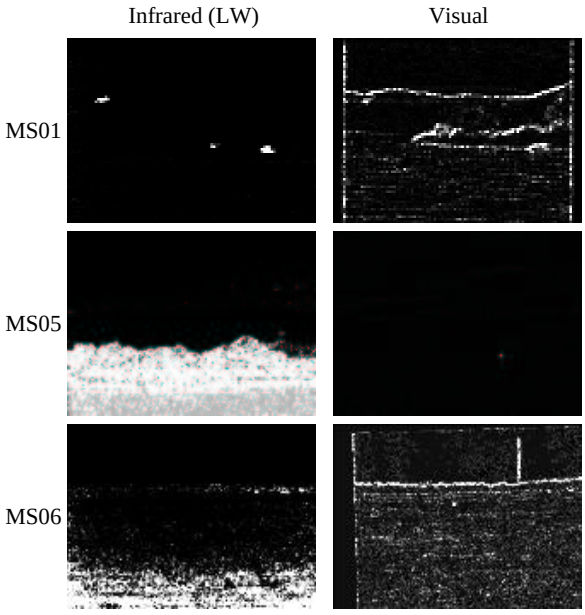

Figure 6. Results of feature fusion

4.4. Sensor Fusion

The sensor fusion step combines the images obtained by the feature level fusion step. In the feature-level-fused images, for each sensor, targets appear as local maxima. Therefore it is possible to fuse these images by a simple multiplication. In the new images the targets will still appear as local maxima. Figure 7 shows the result of applying this simple sensor fusion to the feature-level-fused images of figure 6.

Figure 7. Results of sensor fusion

4.5. Region Growing around Local Maxima

In the sensor-fused image the local maxima will correspond to likely target positions. To detect the targets it is thus necessary to find these local maxima. A region growing procedure around the maxima is then used to incorporate available prior knowledge about target size and aspect ratio.

4.5.1. Local Maxima

The detection of local maxima is based on a succession of morphological operations.[9,10] The basic operator is a dilation with a 2×2 structuring element. Figure 8 shows the different steps of the method.

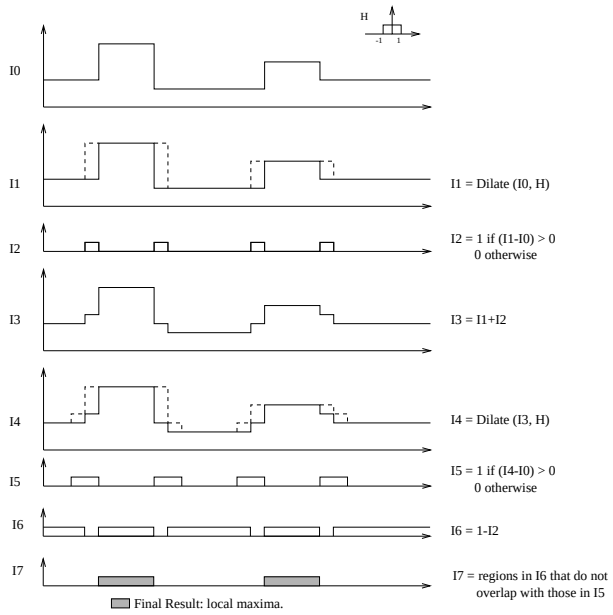

Figure 8. Detection of Local Maxima

4.5.2. Region Growing

To incorporate any available prior knowledge about the possible range of target size or aspect ratio, a region growing procedure is used. The initial regions for the region growing are the local maxima in the image. Surrounding pixels are added to these regions as long as their grey level differs less than a given threshold from the value at the local maxima. If the region becomes too large it is discarded. If the region growing of a given region stops before it reaches the upper size-limit, the other constraints are checked. If a constraint is not satisfied, the region is discarded.

5. MOVING TARGET DETECTION (MTD MODULE)

The second part of the algorithm focusses on the detection of moving targets. In order to detect moving targets, any sensor motion needs to be detected and its effects compensated first. Then, in a temporal fusion step, preceding images can be warped onto the current one. Moving objects will appear as a difference between the original image and the warped ones.

5.1. Detection of sensor motion

The detection of sensor motion is again based on co-occurrence matrices. This time the co-occurrence matrix is calculated between an image and the preceding one (temporal co-occurrence matrix).

$$C_{x_c,y_c}^{dx,dy,dt}(i,j) = \#\left(G(x,y;\ t) = i \mid G(x+dx, y+dy;\ t+dt) = j\right)$$

If no sensor motion occurred between the two images, ideally, for $dx = dy = 0$ (i.e. no spatial displacement), all non-zero elements of the temporal co-occurrence matrix should lie on the diagonal. However, due to noise, there will be a small spread along the diagonal. If one calculates the spatial co-occurrence matrix for a small displacement, the spread along the diagonal is due to noise and to the fact that the image is not homogeneous. Therefore, when comparing this spatial co-occurrence matrix with the temporal co-occurrence matrix, the spread along the diagonal is expected to be the largest in the former one if no motion occurred between the two images that were used to calculate the temporal co-occurrence matrix. When motion is present, the spread along the diagonal quickly becomes larger. The measurement we used to detect sensor motion is based on the percentage of off-diagonal points in both co-occurrence matrices:

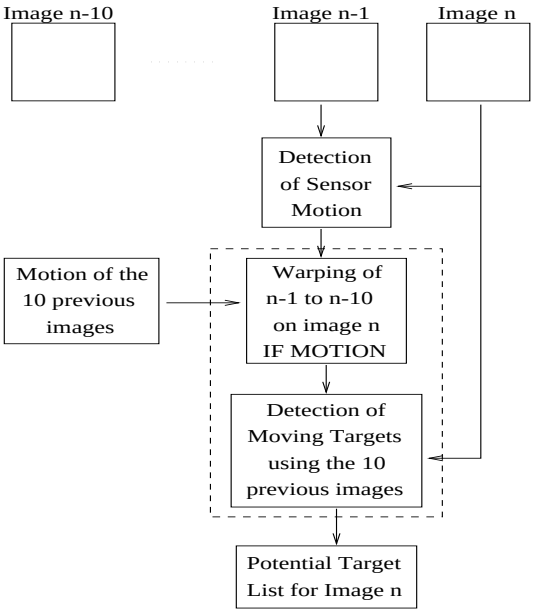

Figure 9. MTD module

$$MC = \frac{\sum_j \sum_{j \neq i} C(i,j)}{\sum_j \sum_i C(i,j)}$$

This is calculated for both the temporal (MC_{temp}) and for the spatial co-occurrence matrix (MC_{spat}). Sensor motion is said to be present if $\frac{MC_{temp} - MC_{spat}}{MC_{spat}} \geq 0.005$. In figure 10 the spatial and temporal co-occurrence matrix are shown. The upper images show the matrices for a part of a sequence where no sensor motion was present. The lower images show an example of both matrices calculated in a part of the same sequence where the sensor was moving.

5.2. Motion Estimation

If sensor motion is detected, we need to estimate it and compensate its effects on the images. The estimation of sensor motion relies on prior knowledge. Three cases are distinguished.

- If it is known that the sensor was fixed and can only do a pan or tilt, the corresponding motion in the image will consist of a uniform translation. In this case we search for the translation by optimising the correlation for a few horizontal and vertical lines.
- If the sensor was mounted in an aircraft, the terrain can be approximated by a plane and a model of the perspective projection of a rigid plane moving in three dimensions is used.[11] This model has only 8 parameters, so finding 4 displacement vectors between two images is sufficient to determine these parameters. In fact we use a threshold that is progressively lowered until at least 15 corresponding regions are found and then a least square method is used to find the parameters.[3]
- In any other case we calculate the optical flow. The calculation of optical flow is based on a Markov Random Field model[12,13] which implements a conservation and a smoothness constraint but at the same time handles motion discontinuities and occlusions.

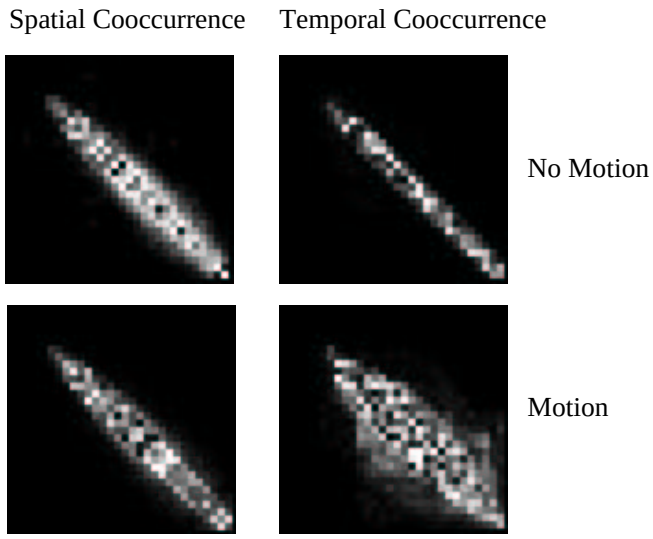

Figure 10. Detection of sensor motion

The three methods given above are listed in order of increasing computational complexity. The method based on correlation would typically need a few seconds (on an HP900/715) to estimate the motion between two images whereas the optical flow method would take about half an hour. Furthermore, because in the latter case, no model can be used to describe the motion field, at least one of the sensors should present enough contrast everywhere in the background in order to obtain an accurate flow field. When only infrared sensors are used, this is not always the case. Therefore it is important to use prior knowledge about the type of sensor motion whenever it is available in order to simplify the used model.

5.3. Detection of moving targets

Once the sensor motion is estimated, preceding images are warped onto the current one. Then the original image is subtracted from the warped ones. If a moving object is present in the scene, we should find a large value at its position. The resulting images after subtraction are therefore thresholded and objects with acceptable size and aspect ratio are selected using a region growing procedure. Tracking is used to get the target list. Figure 11 shows the result of subtracting the original image from the warped ones.

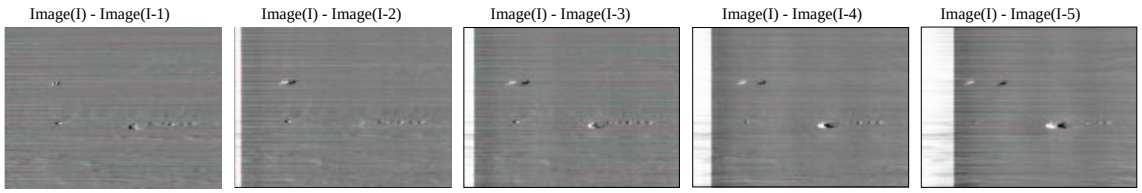

Figure 11. Detection of moving targets

6. DECISION FUSION

The two parts of the algorithm each behave as experts indicating the possible position of targets in the scene. The final decision is reached by fusing the results of these experts.

For the detection of targets in single images, the decision was made after fusing the feature images from all sensors. For the detection of moving targets however, each single sensor acts as an expert. Therefore, the single decision from

the first part of the algorithm should have a higher weight than the different decisions from the second part of the algorithm. Furthermore, because each expert only provides a binary decision - i.e. either a target is present or it is not - the decision fusion is implemented as a weighted "k out of N" voting-rule.[14,15]

7. RESULTS AND DISCUSSION

Table 3 shows the results for the two parts of the algorithm. For the first part of the algorithm the decision is reached after the sensor fusion step. For the second part of the method, the decision is presented for each individual sensor. In the table Pd is the probability of detection and Nft is the average number of false targets per image.

Sequence	Sensor	Results of TDSI After sensor fusion		Results of MTD For each sensor	
		Pd (in %)	Nft	Pd (in %)	Nft
1	LW	95	2.6	92	1
	VIS			1	0
2	LW	63	4	20	0
	VIS			36	0
3	LW	95	0	4	0
	SW			0	0
	VIS			4	0
4	LW	34	1	0	0
	SW			0	0
	VIS			0	0
5	LW	95	0.2	1	3
	Red			42	0
	Green			41	0
	Blue			46	0
6	LW	50	5.8	1.8	0
	Red			19	0.38
	Green			28	0.43
	Blue			29	0.52

Table 3. Results of the different "experts"

The motion detection only gives useful results in the first and fifth sequence. This is basically due to the sequences themselves. In sequences 2 the target is approaching the sensor straight on and therefore it is only detected as a moving target at the end of the sequence where its apparent size increases. In sequence 3 and 4 the targets are stationary. The few detections that are made by the MTD algorithm in sequence 3 are due to a moving antenna on top of the vehicles.

In table 4 the results after the decision fusion are shown. In sequence 3 the fusion improves results. The results for sequences 5 and 6 are better using a single sensor than those found after the fusion. This is due to the error on the image registration. For the calculation of Pd and Nft after fusion, all results are mapped in the coordinate space of the infrared sensor using image registration methods. Table 3 shows that, in the second part of the algorithm, the visual sensor in sequence 5 and 6 is almost solely responsible for all detections. Therefore, if the registration from the visual image to the infrared image is not accurate enough, a position declared as being a target by the visual sensor might just fall within the true target region while, when it is mapped into the coordinate space of the infrared image, it falls just outside and hence is counted as a false target.

8. CONCLUSIONS

In this paper an approach is presented for the detection of targets using multi-sensor image sequences. The developed algorithm consists of two parts. The first part detects targets in single images and is based on texture measurements. For this part of the algorithm a semi-supervised approach is followed. For each sensor a learning image is chosen

Sequence	Results for k=2	
	Pd (in %)	Nft
1	95	2.6
2	63	4
3	93	0
4	34	1
5	96	2.8
6	50	8.6

Table 4. Results after decision fusion with "k out of N" voting-rule

and targets are indicated manually. Then logistic regression is used to find the combination of the different texture parameters that optimises the discrimination between target and background. The weights for these measurements are specific for each of the sensor types (Infrared/Visual). The same weights are then used for the actual target detection in other images of the same sensor type, even for images from different sequences. The second part of the algorithm specifically focusses on the detection of moving targets. If a target is moving in a direction that does not coincide with the viewing direction of the sensors, this part of the algorithm gives good results. In any case, the number of false targets produced by this part of the algorithm is very low. Each part of the algorithm behaves as an expert indicating the possible presence of a target. The final decision of the algorithm is reached by fusing the results of the experts for the different sensors. Because the information to be fused is binary (either a target is detected or it is not), the decision fusion is based on a weighted "k out of N" voting-rule.

REFERENCES

1. J. Knecht, L. Sévigny, C. Birkemark, R. Gabler, B. Hoeltzener-Douarin, , J. Haddon, W. Beck, R. Cusello, E. Oestevold, and L. Garn, *Detection of Target Formations in IR Image Sequences*, NATO/AC243/P3/RSG9, NATO Headquarters, Brussels, 1989.
2. T. Peli, L. Vincent, and V. Tom, "Morphology-based algorithms for target detection/segmentation," in *Proceedings on Architecture, Hardware, and Forward-Looking Infrared Issues in Automatic Target Recognition - SPIE - USA (Orlando)*, Apr. 1993.
3. D. Borghys, C. Perneel, and M. Acheroy, "Long range automatic detection of small targets in sequences of noisy thermal infrared images," in *Proceedings on Signal and Data Processing of Small Targets - SPIE - USA (Orlando)*, 4-8 April 1994.
4. R. Haralick, "Statistical and structural approaches to texture," *Proc. IEEE* **67**(5), pp. 786–804, 1979.
5. P. Verlinde, "Numerical evaluation of the efficiency of a camouflage system in the thermal infrared (in dutch)," Master's thesis, Katholieke Universiteit Leuven (KUL/ESAT), Leuven, 1989.
6. P. Verlinde and M. Proesmans, "Global approach towards the evaluation of thermal infrared countermeasures," in *Proceedings of Chracterization, Propagation and Simulation of Source and Backgrounds - SPIE - USA (Orlando)*, vol. 1486, Apr. 1991.
7. D. Borghys, P. Verlinde, C. Perneel, and M. Acheroy, "Long range target detection in a cluttered environment using multi-sensor image sequences," in *Proceedings on Signal Processing, Sensor Fusion and Target Recognition IV - SPIE - USA (Orlando)*, 20-25 April 1997.
8. D. Hosmer and S. Lemeshow, *Applied Logistic Regression*, John Wiley and Sons, 1989.
9. C. Perneel, M. de Mathelin, and M. Acheroy, "Detection of important directions on thermal infrared images with application to target recognition," in *Proceedings of Forward Looking Infrared Image Processing - SPIE - USA (Orlando)*, 12-16 April 1993.
10. P. DeBeir, *Détection d'objectifs a longue distance dans une séquence d'images infrarouges*, Royal Military Academy, 30 Avenue de la Renaissance, B-1040 Brussels, 1992.
11. R. Y. Tsai and T. S. Huang, "Estimating three-dimensional motion parameters of a rigid planar patch," *IEEE Trans. on Acoustics, Speech and Signal Processing* **29**(6), pp. 1147–1152, 1981.
12. R. Morris, *Image Sequence Restoration using Gibbs Distributions*. PhD thesis, Dept. of Egineering and Trinity College, Cambridge University, 1995.

13. R. Morris and W. Fitzgerald, "Discontinuous motion and occlusion estimation - theory and application," in *Proceedings of the Iternational Conference for Young Computer Scientists, Beijing*, July 1995.
14. E. Waltz and J. Llinas, *Multisensor Data Fusion*, Artech House, Boston, 1990.
15. R. Antony, *Principles of Data Fusion Automation*, Artech House, Boston, 1995.

Dirk Borghys received his MSc degree in Physics in 1986 at the Vrije Universiteit Brussel, where he also obtained a postgraduate degree in Computer Science in 1988. Between 1988 and 1993, he was researcher at the Technical Department of the Army involved in infrared imaging and the development of objective evaluation methods for countermeasures in the infrared and visual band. Since 1993 he continues his research activities in the electrical engineering department at the Belgian Royal Military Academy. His main research interests are sensor fusion, motion analysis, target detection, Markov Random Fields and restoration of image sequences.

Patrick Verlinde graduated in 1983 as a telecommunications engineer from the Polytechnical Division of the Royal Military Academy in Brussels (Belgium). Until 1985 he worked in the maintenance company of the Belgian Army Light Aviation. In 1987 he joined the Technical Department of the Army, where he was responsible for the development of objective evaluation methods for multi-spectral countermeasures. In 1989 he obtained the degree of mechatronics engineer from the Catholic University of Leuven (Belgium) and in the same year he became an Engineer of the Military Material. From 1991 until 1995 he worked within the framework of the Belgian Gendarmerie, where he was the technical project manager for the realisation of a nationwide radio trunking network for all Belgian public safety services. He is since 1996 an assistant professor in telecommunications at the Royal Military Academy. His current research interests include data fusion, statistical pattern recognition and neural networks.

Christiaan Perneel was born in 1963. In 1986, he received the master's degree of engineer in telecommunications at the Royal Military Academy of Brussels, Belgium, and he received the Ph.D. degree in 1994 at the Vrije Universiteit Brussel. Currently, he is Lecturer at the Department of Applied Mathematics of the Royal Military Academy. His research interests include image processing, pattern recognition and sensor fusion. He is the head of the pattern recognition cell of the Signal and Image Centre of the Royal Military Academy and he is the Belgian governing board member of IAPR (International Association for Pattern Recognition)

Marc Acheroy was born in 1948. In 1971 he received his master's degree in engineering at the Royal Military Academy. In 1981 he obtained the master's degree in control and automation from the Université Libre de Bruxelles where he also obtained his PhD in applied sciences in 1983. Currently he is head of the electrical engineering department and of the Signal & Image Centre at the Royal Military Academy. He is professor in signal processing and automatic control. His main scientific interests include thermal imaging, pattern recognition, image restoration and compression.