# Multi-Spectral Target Detection Fusion

G S Beard

*QinetiQ – University College London*

**Abstract:** Scenarios for modern air to surface missile seekers require a high probability of target detection with a minimal false alarm rate to avoid collateral damage. Active and passive countermeasures are increasingly expected to appear on enemy targets, designed to counteract a seeker in a narrow or wide band of frequencies. A technique to improve on the performance of a single sensor is the use of two separate sensors such as radar and infrared. The detection and false-alarm performance of each individual sensor on data of military targets is examined, and compared to that achieved through the use of detection level fusion.

## 1 Introduction

The expected model of battlefield engagements has moved from the attack of massed armour on the battlefield to one comprising small-scale engagements of small groups of targets in urban areas. The emphasis is focused on the attack of smaller groups of targets while minimising collateral damage. Defensive aid suites in the form of active and passive countermeasures such as radar netting, infrared (IR) sheeting and decoy targets are increasingly expected to appear. These can be designed to counteract a seeker in one waveband or over a range of frequencies. The operational use of the seeker sensor for this work is envisaged as being a Terminally Guided Sub-Munition (TGSM) at the end of a larger chain of targeting sensors.

An important question for future seekers is whether a single-band sensor can deliver the desired level of performance against countermeasured threats, or if sensors operating in multiple wavebands are required. Taking advantage of the different physical nature of microwave and optical sensor systems, the fusion of the data can result in an enhancement in system performance.

In this paper the probability of target detection and rate of false alarm are calculated for both radar and IR sensors. Detection level fusion is then applied to the sensor imagery and the effects assessed.

Section 2 describes the data available and how the imagery was formed. Section 3 covers the radar detection processing and performance. Section 4 describes the IR processing and performance. Section 5 introduces the fusion options and the benefits these bring. Conclusions are drawn in section 6.

## 2. Data gathering and Image formation

Data for this work was gathered using a seeker sensor in both the radar and IR domains. The seeker consists of a real-beam radar and an uncooled imaging infrared sensor. A range of military vehicles, some with camouflage deployed, were imaged in a rural environment by the seeker mounted on a helicopter. The weather was clear and sunny resulting in a high solar loading of the scene with consequent high contrast in the IR images. Most of the vehicles were deliberately heated by running their engines prior to the start of measurement to further aid the IR sensor.

The radar sensor is typical of the type that could be utilised in low-cost TGSM roles. The down-range resolution of the radar is limited by bandwidth, and is far better than the cross-range resolution which is limited by the size of the antenna. Hence the sensor can resolve fine detail down-range but only coarse structure in cross-range. The radar data consists of a series of range-profiles for each swath (side-to-side scan) recorded with information on where the antenna was pointing. From this data, a ground map image can be formed for each swath as shown in Figure 1.

The IR sensor is co-located with the RF sensor; as the radar antenna scans across the field of view, the IR sensor looks in the same direction. Each IR swath comprises the data gathered by the middle column of a 64 by 64 pixel detector array, and so is comprised of a sequence of vertical strips 64 pixels high. By placing these strips side-by-side as the sensor scans, a 2D image of the scene is formed.

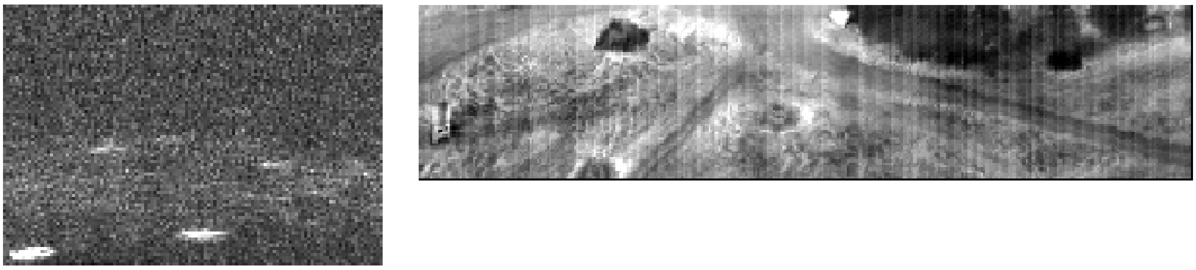

**Figure** 1 Radar (left) and IR (right) imagery of the same scene

## 3. Radar Processing

The radar target detection processing operates in a number of stages before a target is declared.

An adaptive threshold is used to detect pixels with intensities brighter than their local background. Every pixel in the image is processed sequentially. The pixel under test is surrounded by a square of pixels which are ignored – the guard area. This prevents returns from the target affecting an outer-ring of pixels – the statistics ring. The size of the guard is set such that the distance from the test pixel to the statistics ring exceeds the expected target dimensions. In order to be declared a possible target, the pixel value needs to exceed a threshold calculated from the statistics ring. The approach used was to calculate an ordered statistic threshold. This procedure orders the pixels in the statistics ring by brightness then sets the threshold such that k% of pixels are below this threshold. This procedure does not assume a particular clutter model so has robust performance across boundaries, but is computationally expensive.

The output from the thresholding is an image with scattered pixels that have exceeded the ordered statistic threshold. From this, little information is available about regions of interest and some clustering of pixels is needed. A technique known as snake delineation [2] is used. This process finds the region of pixels that bound a potential target after it has been initially detected. The snake consists of point nodes which are joined together to form a polygon that is initially seeded around a detection. The algorithm moves each node by a small amount and tests to find whether the new snake is a better delineation. If the delineation is better the node is left where it is, otherwise the node is moved back. The program keeps trying to move all the nodes one by one in random directions. When no improvement is found the delineation is complete and an outline is formed as shown in Figure 2.

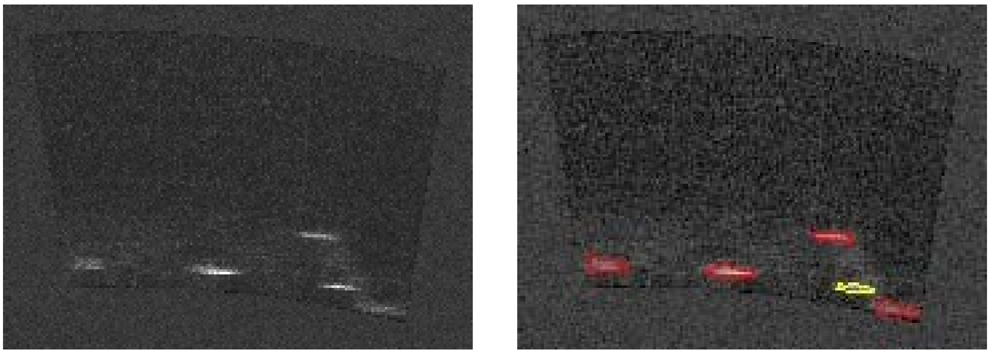

**Figure 2** Radar image (left) and with snake segmentation (right)

The snake segmentation can be seen to produce an outline that matches a human visual interpretation of the object's extent in this example scene. Each segmented object has a number of features calculated. These are spatial features such as length and width as well as statistical features such as target to clutter ratio and fractal dimension. A number of data sets were examined to provide 'training data'. These allowed selection of certain features, and for constraints to be applied on the range of values within which a target should fall. These features were fairly successful at removing false alarms caused by clutter, and artefacts due to saturation in the radar imagery. The yellow outline of the second target from the right in Figure 2 indicates an object that has been classified as clutter by the feature discriminator.

| | |
|---|---|
| Overall PD | 0.95 |
| Open PD | 1.00 |
| Treeline PD | 0.67 |
| FA/km$^2$ | 14.55 |

**Table 1 Pd and Pfa Even channel**

The detection results shown in Table 1 are split into three categories, detection of all targets, detection of those targets in the open, and detection of targets in the treeline. The rates of false alarm shown are normalised for a square kilometre, an area approximately 25 times larger than the sample search area.

## 4. IR Processing

The IR detection algorithm is designed to identify objects that are unusually bright or dark relative to the background and whose sizes and shapes are consistent with those expected from military targets. Statistical models are used to characterise the background and unusual objects are then identified as statistical outliers. The process follows similar stages to the radar technique, although using slightly different methods at each. Pixels with unusually strong edges, i.e. a steep gradient with respect to the whole image are found, clusters of pixels are formed and features are extracted from these clusters.

Figure 3 shows an IR swath with the detections enclosed by an ellipse. Any pixels within the ellipse are considered to be target returns for the purposes of calculating statistical measures.

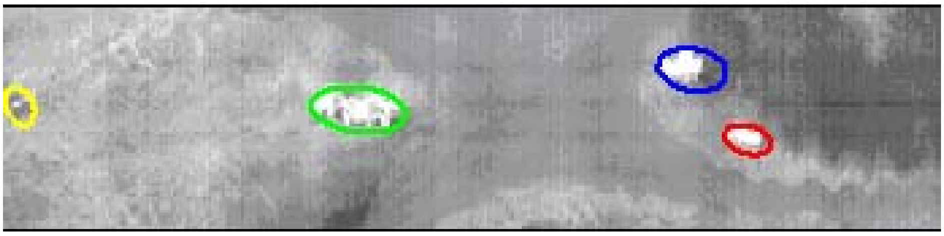

**Figure 3** IR image with detections highlighted

| | |
|---|---|
| Overall PD | 0.97 |
| Open PD | 0.96 |
| Treeline PD | 1.00 |
| FA/km$^2$ | 99.86 |

**Table 2** Pd and Pfa IR sensor

IR detections and false alarms were counted in the same manner as for the radar sensor. The IR had fewer problems detecting targets in the treeline but was more susceptible to the effects of camouflage. Targets with IR sheeting were detected via patches of heat escaping from the edges, and should such a patch be selected as the target location the seeker could well miss the vehicle. Should the target have been stationary without the engines running, IR detection of the targets would be very difficult. The IR also suffered from a significantly higher false alarm rate than the radar sensor, as shown in Table 2.

While the IR sensor has a high probability of detection for the targets, it also has a large number of false alarms present. These are caused by man-made objects in the scene, but also by trees and ‘bright’ patches of ground.

## 5. Detection level fusion

Fusion using the two sensors aims at the integration of disparate and complementary data to enhance the information present in the images. This should lead to more accurate data and increased utility.

Fusion is a possibility at three levels: pixel, feature and detection. Pixel level fusion involves the data from each sensor being mapped onto an image with the same spatial dimensions such that each pixel in the new image represents a combination of the two sensors at that point. For sensors that have dissimilar geometries such as RF and IR this combination is a difficult problem. In feature level fusion, features are calculated from regions of interest and are combined, the aim being that the combined features produce better performance in a detection algorithm than the features from each sensor independently. This scheme would be practical using the data gathered, however for this work detection level fusion was employed since this is the simplest fusion scheme to implement. Detection algorithms are performed separately on data from each sensor to produce regions of interest and the locations of the regions are then combined. This gives an ability to reinforce valid target decisions and reject false alarms.

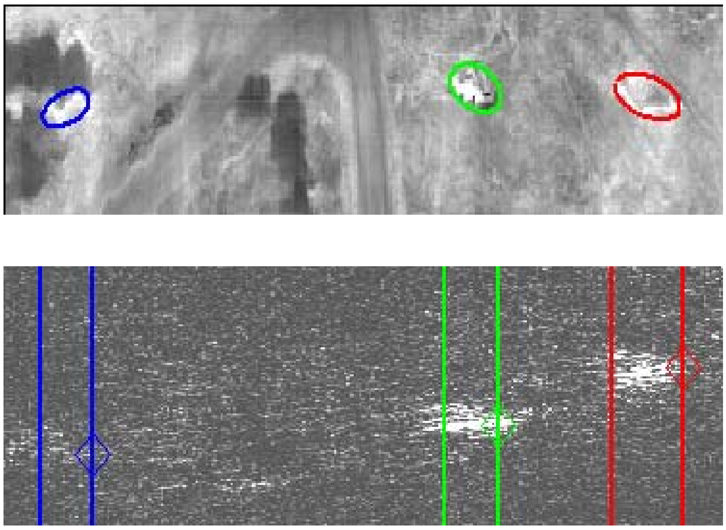

**Figure 4 IR detections with azimuth mapped to RF image**

Due to the co-location of the sensors and the antenna azimuth information available it is a simple process to correctly register the direction of detections from one sensor to the other. Figure 4 shows an IR swath with detected areas, and the azimuth extent of the detections mapped onto the RF image. While the azimuth maps across accurately, the different geometries of the sensors - azimuth and elevation angle for the IR and azimuth and range for the RF - makes the mapping in the vertical direction less trivial. For the results shown here the fusion has effectively been performed manually, however for future work, the automatic registration of the detections should be possible.

| | |
|---|---|
| Overall PD | 0.94 |
| Open PD | 0.96 |
| Treeline PD | 0.80 |
| FA/km$^2$ | 12.02 |

**Table 3** Pd and Pfa – AND Fusion

Each sensor's detection algorithms are currently configured for single-sensor detection performance which may not be the optimal setting for the best performance when their results are combined. Table 3 shows the effects of applying 'AND' detection fusion to the outputs of both sensors. Corresponding IR results were available for 4 of the 5 radar scenes and hence the fused treeline PD is constrained by a radar PD of 0.8 instead of 0.67.

None of the tree-based false alarms from the IR sensor are passed through since the tree areas do not give strong returns in the radar domain. The only source of false alarms that passed through the fusion process were man-made metallic objects in the scene.

## 6. Conclusions

Applying detection level fusion to the sensor system resulted in improved performance over that from each single sensor based on the reduction in false alarm rate. There was a slight drop in detection probability, however the reduction in the false alarm rate compensated for this.

In order for fusion to produce the best results it would be ideal for each individual sensor's detection scheme to be optimised for a fusion role. This could involve raising the detection probability slightly for each sensor, relying on clutter objects not producing a false alarm in both sensors.

The system would also offer flexibility for different scenarios. Scenes where targets are expected to have IR camouflage or in misty/rainy weather conditions could place emphasis on the radar sensor. Scenes where targets are expected to be operational and producing heat would favour the IR sensor.


## Acknowledgment

Sharon Watson at QinetiQ Farnborough for the IR detections.


## References


[1] G.T. Maskall and S.K. Watson, "DMS Multi-Spectral Signatures and Data Fusion Options". QinetiQ/S&E/APC/CR032156. 24 Nov 2003. Unpublished QinetiQ report.

[2] D. Pedlar, "Target Delineation and Classification using a Region-Based Active Countour and a support Machine classifier on SAR Imagery". EUSAR Conference, Proceedings Volume 1, Ulm, May 2004.