

Ground Target Classification Using Combined Radar and IR with Simulated Data

Andris Lauberts, Mikael Karlsson, Fredrik Näsström, Rahman Aljasmi

Swedish Defence Research Agency

P.O. Box 1165, SE-581 11 Linköping, Sweden

{andris, mikkar, frenas, rahman}@foi.se

Abstract - Classification of ground targets is studied using a combination of simulated IR and radar data. The data are presumed to be collected by a missile, approaching a target area containing one of six possible targets. IR data are generated from textured models of the scene and targets, using a 3D rendering tool. High-resolution radar range profiles are calculated using physical optics based software, applied to high definition CAD models. Modeled radar clutter is added coherently. It is assumed that the targets have been detected. They are classified using an artificial neural network on IR and radar data individually, as well as in combination. Two different data fusion methods are studied, feature fusion, using a common feature vector, and decision fusion using past experience. The results show that feature fusion outperforms the best sensor, given knowledge of uncertainty in data. Decision fusion performs overall better than feature fusion as far as a priori data can be trusted.

Keywords: data fusion, IR, radar, multisensor, classification

1 Introduction

Classification of ground vehicles poses a difficult problem in view of a variable background and diverse obscuring effects. A number of approaches have been made in this area, see e.g. [7], [8] and [9]. In many cases, data from a single sensor lacks essential information to identify the unknown object. The aim of this work is to analyze the performance of fusing simulated IR and radar data to identify stationary ground-based targets. An airborne platform is presumed to carry a system of an IR sensor (8-9 μm) and a high range resolution (HRR) radar (16 GHz). The platform approaches one of six different targets from a distance of 8000 m, at a constant altitude of 300 m. A simple target background is assumed, consisting of short grass, to which a reasonably accurate statistical modeling has been applied to describe the radar back-scatter, verified with some experimental checkpoints (section 2.3). Future work will include more complicated scenarios, backed with further measurements.

Optimal target recognition calls for all information from available sensors. One way to reduce the expected large amount of data is to extract features. These features are later compared with a library of features with the aim to classify the observed targets. Ideally, the selected features should maximize the similarity of objects within the same class in a multi-dimensional feature space, and at the same time maximize the dissimilarity between objects of different classes. For optimal performance, the chosen

features should be statistically independent, geometrically invariant and offer short computation times under robust conditions.

For IR, the distance independent features, a) min/max target diameter, b) normalised standard deviation of the intensity, c) 2nd order invariant moment, d) area of target divided by its periphery squared, and e) mean edge strength, have proven to perform especially well for classification purpose.

For radar two different cases have been tested. The first one uses the individual high-resolution range profile with different polarisations as features. The second one uses a more elaborate singular value decomposition (SVD) based method for feature refinement. Comparative approaches may be found in e.g. [8] and [9].

Data fusion can be seen as a high level classification when data from different sources is combined to expand the available information space. Some examples of typical classifiers are: Bayes method, Artificial neural networks (ANN), Fuzzy logic (soft decision rules), and different kinds of matched filters. Sometimes we also need a method to determine density distributions of probabilities, such as the Parzen method (weighed distance measure) or k -nearest neighbour. In this study, all classification is done using a multi layer ANN.

In this study, we have applied two different data fusion methods. With *decision* fusion, the targets are classified by IR and radar individually. Their results are then combined using confusion matrices containing the statistical experience from earlier experiments. The *feature* fusion uses a common feature vector built of features extracted from individual sensor data associated with the same target. The merits of each method are weighed in terms of information age and data uncertainty.

2 Data generation

This section describes the scenario used, as well as details of the generation of the synthetic data which this study is based on.

2.1 Scenario

The scenario features a multisensor platform flying in a straight trajectory towards one of six different stationary targets in a grass field. The part of the trajectory that is of interest in this study begins at a slant range distance of

8000 m and ends at 600 m. The height above the target is constant at 300 m, see Fig. 1.

The sensors collect data on their flight towards the target, with constant azimuth, and increasing elevation angle from 2.15° at 8000 m distance to 30° at 600 m. The final part of the flight (600 - 0 m) has been omitted in this study.

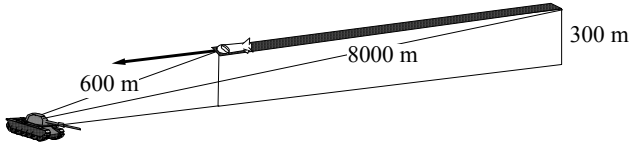


Fig. 1. The classification is studied from a distance of 8000 m to 600 m.

Two cases have been studied in this work. Case 1 is designed to study how the classification depends on target distance with the range, elevation and azimuth choices given in section 2.2. It uses a single IR image and radar range profile in each point to make the classification. This reduces the number of data that needs to be generated per trajectory. Therefore it has been possible to evaluate this classification method in target aspects covering the full circle in azimuth.

Case 2 addresses mainly the strong variations in the radar return (*glint*) with changing aspect angle and uses a much finer angular grid for the radar data, as specified in section 2.3. The SVD method used needs data in very fine elevation steps, which takes a lot of resources to generate. The number of azimuths evaluated in this case has therefore been reduced. It is though possible to compare the methods in the common aspects. For IR, the data generated in case 1 are also used for case 2.

We assume that the target is detected and is stationary. Conforming with the aim in this phase of our work to put the main effort on algorithmic aspects of target recognition with the dual sensor, we assume some further simplifying conditions, *viz.* that the target is located in a field with short grass, with the barrel pointing in the forward direction. No rain or clouds are expected in the scenario. Complications with respect to these points will be introduced in future progress of this work.



Fig. 2. The image shows the six CAD models with textures for visual wavelengths that is used in the simulations. From the upper left: BMP 1, BTR 80, M109A6, T72, T80, S44 launch vehicle.

The six targets used in the simulation are shown in Fig. 2. To classify these targets we need information about their characteristic IR signatures. Each CAD model has an

IR texture to simulate the IR signature. The radar signatures were calculated using a high definition geometry model of each target. No texture information is handled by the software that calculates the radar cross section, i.e. only the high-definition geometry is used.

2.2 IR data

The IR textured CAD models were placed in a virtual environment. The terrain model used in the IR simulations was created from high spatial resolution data collected using a helicopter carrying a laser measuring system. The system was also equipped with a high-resolution digital camera and an IR camera. The surface model covers an area of 1500 x 1500 metres and has a geometric resolution of 0.25 metres per pixel. Geometrically corrected image data from the IR camera, with resolution of 0.2 metres per texel, has been used to texture the model.

The IR images have been generated using a tool named SceneServer [1], which is an OpenGL based scene rendering tool that is controllable via a socket interface. Fig. 3 shows an example of a part of the terrain with a tank positioned in the grass field in front of a forest line. The tree line is not included in the immediate background of the target. The different target models were placed in the virtual environment and controlled from a MATLAB client. The interface between MATLAB and the 3D environment enables the user to manipulate the models and background to be rendered in detail.

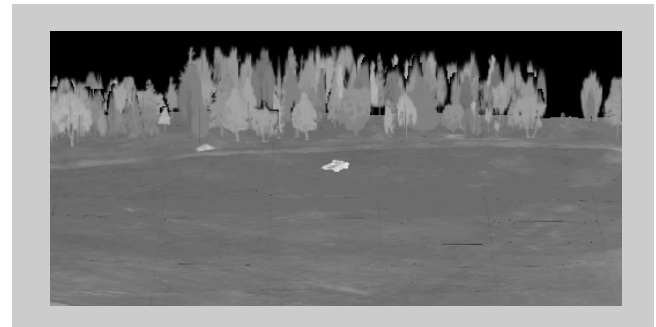


Fig. 3. The image shows an IR scene with a tank in front of the forest line.

The so generated IR database consists of 512 x 512 pixel grey-scale images with field-of-view $5 \times 5^\circ$ at 8-9 μm waveband. The training data set has 13 target ranges (or elevation angles) differing by a constant step size in $\log_{10}(\text{range})$. The azimuth angles are varied from 0° through 355° stepped by 5° . The evaluation data set has elevation angles differing by 1° and azimuth angles differing by 2° , from the corresponding nearest training data set.

2.3 Radar data

Relevant radar parameters in the simulations are listed in Table 1.

The frequency is chosen to give a compromise between angular resolution and all-weather capability; however, no specific study of the weather dependence has been made

in this work. The assumed bandwidth gives a high range resolution, 10 cm. The radar is fully polarimetric with coherent processing of the polarization channels.

A preprocessing in the radar channel is assumed to produce high range resolution (HRR) profiles, which form the primary input to the classification unit. For the simulations, the raw data input to the preprocessing is obtained from a radar cross section calculation program. This program, based on physical optics, delivers frequency-stepped, far-field, coherent back-scatter data for a specific target, aspect angle, and polarizations HH, VV, HV, and VH.

Table 1: The parameters of the simulated radar.

| Feature | Implementation |
|-------------------|------------------|
| Frequency | 16 GHz (Ku-band) |
| Bandwidth | 1.5 GHz |
| Signal processing | Coherent |
| Range resolution | 0.1 m |
| Antenna diameter | 0.3 m |
| Antenna lobe | 3.6° |
| Polarisation | HH,VV,HV,VH |

The frequency is stepped in increments of 12.5 MHz, giving calculated returns at $n = 121$ specific frequencies between 15.25 and 16.75 GHz, for each polarization. An IFFT with a Hamming shaped weight function is applied to the data to produce a complex HRR profile (with amplitude and phase).

A clutter background is added coherently to each range gate. The clutter model assumes gamma distributed amplitudes and uniformly distributed phases. The coefficients in the gamma distribution for different elevations have been estimated from extrapolation of values taken from the literature [2] [3] [4]. A specific measurement campaign has been carried out to check clutter values [5]. The final range profile is obtained by taking the magnitude of the complex values.

In the simulations the radar is assumed to collect range profiles on its flight towards the target, with a specific, constant azimuth, and with the elevation increasing from 2.15° at 8000 m distance to 30° at 600 m. For case 1, range profiles are generated for azimuth and elevation values given in section 2.2, and are used individually for the classification. In case 2 the step is 0.025° in elevation between the collection points. This gives a total of 1114 profiles in the mentioned distance interval. The profiles are accumulated column-wise into a matrix with n rows. No implicit dynamics is involved in the data collection; the production of a profile is taken to be instantaneous. Another considerable simplification in our study is that the collected profiles are assumed to have the same relative alignment in the range dimension, with respect to a fixed reference point in the target. One consequence of this is that specific peaks in the profiles do not move between range gates from profile to profile.

3 IR signal processing

3.1 Features

Table 2 shows the nine IR features used in this study. The distance independent features, 1, 3, 4, 5 and 7, have proven to perform especially well for classification purpose. Features 8 and 9 are included only when the angular measures can be transformed to absolute quantities using distance estimates provided by the radar.

Table 2: IR features

| No. | Description |
|-----|--|
| 1 | Ratio min/max target diameter (smallest and largest sides of a circumscribed rectangle with minimum area). |
| 2 | Mean intensity of the segmented target. |
| 3 | Standard deviation of the intensity of the target normalised by its mean intensity. |
| 4 | 2nd order invariant moment computed on the binary image of the thresholded target. |
| 5 | Area of target divided by its periphery squared. |
| 6 | A measure of the signal to clutter (SNR). |
| 7 | Mean edge strength using the Canny filter. |
| 8 | Maximum diameter of target (converted to distance in meters using radar). |
| 9 | Target area (meters squared, cf. no. 8). |

3.2 Linear discriminant analysis

Feature reduction aims at choosing those features that save class separation with minimum loss. The class separability is mainly independent of the coordinate system. However, the choice of features depends on both the class distributions and the classifier used. The idea of the Linear Discriminant Analysis (LDA) is to find a projection of features that maximizes the separation between different classes and simultaneously minimises their internal variations [6]. In the case of only two classes the LDA method is usually named Fisher Discriminant. We wish to find a feature transform keeping the class-separability high in a lowest possible order of dimension. Let us project the r -dimensional original feature vector \mathbf{x} onto an s -dimensional \mathbf{y} vector via the transform $\mathbf{y} = \mathbf{A}^T \mathbf{x}$, where $s < r$ and \mathbf{A} is a $r \times s$ matrix with linearly independent column vectors. The variance within classes is defined by the covariance matrix \mathbf{S}_w ,

$$\mathbf{S}_w = \sum_{i=1}^L P_i E \{ (\mathbf{x} - \mathbf{m}_i)(\mathbf{x} - \mathbf{m}_i)^T | \omega_i \}. \quad (1)$$

The variance between classes is defined by the covariance matrix \mathbf{S}_b ,

$$\mathbf{S}_b = \sum_{i=1}^L P_i (\mathbf{m}_i - \mathbf{m}_0)(\mathbf{m}_i - \mathbf{m}_0)^T. \quad (2)$$

Here, L is the number of classes, P_i is *a priori* probability of class ω_i , \mathbf{m}_i is the mean value of \mathbf{x} of class ω_i and \mathbf{m}_0 is the mean value of \mathbf{x} over all classes. A qualitative measure of class separability is given by combining \mathbf{S}_w and \mathbf{S}_b into $J = \text{trace}(\mathbf{S}_w^{-1} \mathbf{S}_b)$ to be maximized for best performance. It turns out that the optimal transform \mathbf{A} is given by a matrix with columns being eigenvectors corresponding to the s largest eigenvalues of the matrix $\mathbf{S}_w^{-1} \mathbf{S}_b$.

4 Radar signal processing

In case 1, individual range profiles are used as target features for the classification. However, in case 2 the training of the classification unit (the neural network) is not made with individual range profiles, but with profile “templates” obtained from a singular value decomposition (SVD) of the $n \times m$ matrix \mathbf{Y} consisting of m accumulated profiles, each with n samples in range, as described in [7]. According to this procedure the template is taken as the eigenvector of the covariance matrix $\mathbf{R} = \mathbf{Y}\mathbf{Y}^T$, corresponding to the largest singular value, giving a data compression of m times. This template is a feature vector, having all the physical characteristics of a range profile, representing the target over the aspect angle spanned during the collection of the m profiles. This choice of eigenvectors of the covariance matrix \mathbf{R} as features for the classification was made for its simplicity of implementation, and can also be shown to be optimal in the minimum mean squared error sense [7].

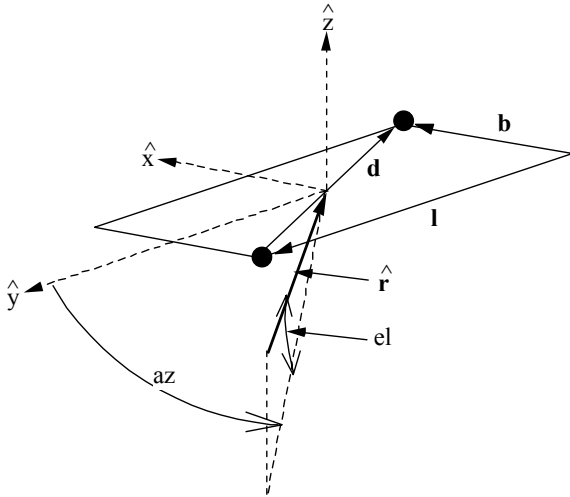


Figure 4. Simple 2D model of the top surface of a target for an estimation of the radar glint rate, experienced by a radar viewing the target in the direction $\hat{\mathbf{r}}$, while approaching it with constant azimuth az , but changing elevation el . The rectangular top of the target is of length l and width b . The interference of the radar returns from two diagonally opposed scattering centers, separated by \mathbf{d} , is taken to produce the maximum glint rate, as quantified in the text.

For the continuation of this work we plan to include model-based methods for robust feature generation, which can be expected to perform better. Such methods have been applied in previous work on HRR target recognition, reported in the literature. For example, in [8] an autoregressive modeling approach is used to extract amplitudes and locations of peaks in the radar return, corresponding to size and range location of target scatterers, followed by application of statistical methods to capture the variability of data. The latter problem has been addressed with another approach using a hidden Markov model [9].

The number of profiles, m , entering the SVD procedure is chosen to vary in inverse proportion to $\sin(el)$, where el is a mean elevation of the accumulation interval. This choice can be shown to give approximately the same expected “glint”-like variations of the received signal from a target, for all elevation intervals for which the SVD procedure is applied, assuming a simple 2D model shown in Figure 4. The target is taken to have a rectangular top surface, of length l and width b , with scattering centers as sources for the radar return. If the radar sensor views the target along the direction $\hat{\mathbf{r}}$, and approaches it in a vertical plane, *i.e.* with changing elevation (el) but constant azimuth ($az \approx 45^\circ$), the greatest glint rate originates from the changing interference between the radar returns from two diagonally opposed scatterers, separated by the vector \mathbf{d} . The path difference between the returns is

$$2\mathbf{d} \cdot \hat{\mathbf{r}} = 2[-b \cos(el) \sin(az) - l \cos(el) \cos(az)], \quad (3)$$

$$\equiv -2C \cos(el)$$

where C is constant, assuming a constant az . The change in path difference for an increase $\Delta el (= 0.025^\circ)$ in elevation corresponds to a phase change of

$\Delta\phi = 2C \sin(el) \Delta el (2\pi / \lambda)$, where λ is the wavelength. The phase change over the angle interval of the collection of m profiles is approximately $m\Delta\phi = 2Cm \sin(el) \Delta el (2\pi / \lambda)$. We have taken as a criterion for choosing the number of profiles to go into the SVD procedure that this phase change should be equal for all elevation intervals for which the procedure is applied, giving that m should vary as $1/\sin(el)$. Hence, if one starts the template generation during a flight using the SVD procedure applied to m_0 profiles, collected about elevation el_0 , then the number of profiles to go into the procedure about elevation el is

$$m = m_0 \sin(el_0) / \sin(el). \quad (4)$$

Fig. 5 shows the result of the classification with the radar channel alone, using a neural network trained with templates from every second of the azimuth angles ($45.0, 45.1, 45.5, 47.0, 53.0^\circ$) each containing 13 SVD-processed m -profile groups over the elevation interval of 1114 single profiles. The test data consists of *single* range profiles,

taken from the remaining four azimuths of the generated catalogue. (We have not yet implemented compression of the test data in the same fashion as the training set.) In the multi-polarization case, a combined feature vector is formed as the concatenation of the individual profiles HH, HV, VV, and VH.

Fig. 5 represents single-trial results using radar only. Independent repetitions will give improved accumulated performance. In order to attain a total probability of correct classification of, e.g., 0.99, six single classifications will suffice, given that they are independent. The independency will have to be given a separate consideration in future work.

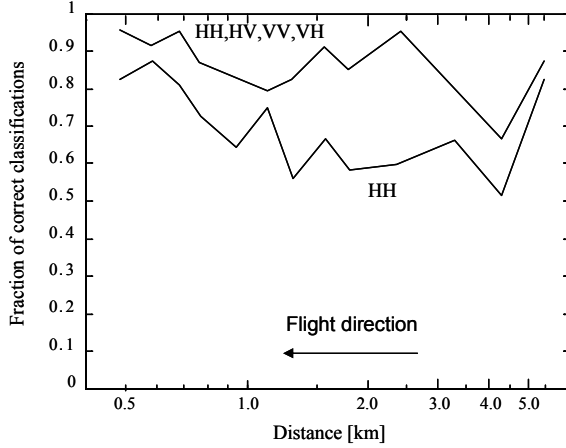


Fig. 5. Neural network classification results, showing the dependence on distance of single-profile classification using training feature vectors obtained from the SVD procedure described in the text. The figure also illustrates the gain of using multi-polarization data compared to one single channel (HH).

5 Classification and data fusion

All classification has been carried out using a single artificial network (ANN) classifier. The *Resilient Backpropagation* (RPROP) training method has been used and the ANN is used directly as a fusion engine on the sensor data features. The RPROP method is characterized by using only the sign of partial derivatives to determine the direction of the weight update. The structure of the ANN forms three layers. The input layer is fed with feature vectors through six channels, one for each class. These vectors are combined by weights before entering and after leaving a single hidden layer with (no. of classes \times no. of features) neurons. The result compiles into an output layer with six slots, one for each class. The ANN is then trained against known template classes, adjusting its internal weights to reach an optimal solution. Test data are finally fed to the net, now with all weights fixed, and the most likely classes are estimated from the output layer.

In this study, we have applied two different methods for data fusion. With *decision* fusion, the targets are classified by IR and radar individually. Their results are then combined using the confusion matrices that contain the statistical experience from earlier experiments.

Feature fusion uses a common feature vector built of the features extracted from the individual sensor data, associated with the target. There remains the question of independent observations. Truly, both sensors see the same target. This common link to sensor data advocates for feature fusion. On the other hand, the extracted features complement each other since they reflect different physics of the target. This fact would support the use of decision fusion. In practice there will be a mix of factors, a motive to study both feature and decision fusion.

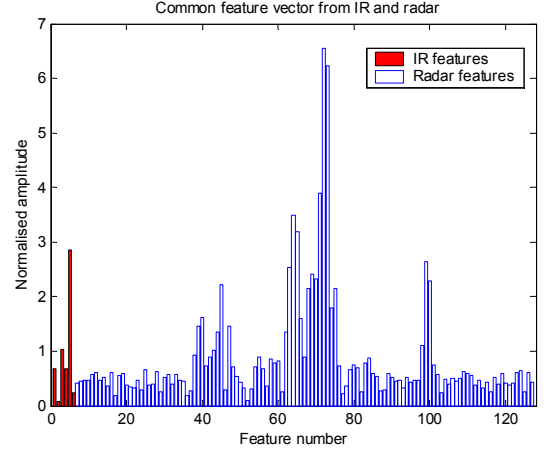


Figure 6. Common feature vector composed of IR (first six elements) and radar (last 121 elements) normalised features.

Figure 6 depicts a thought example of feature fusion where the normalised feature vectors from IR and radar are concatenated into a common vector entering a single classification algorithm.

Decision fusion, here named *type A*, combines the individual decisions from IR and radar with corresponding rows in a confusion matrix (cf. Fig. 9) resulting from earlier experiments, given a particular aspect angle. Here, columns refer to true types and rows refer to proposed types after classification, with matrix elements showing the relative numbers of hits. The element values for rows corresponding to estimated classes (e.g. no. 4 for IR and no. 5 for radar) are given in Table 3.

Table 3: Decision fusion, type A, using estimated classes from individual sensors combined with historical data.

| Class, true \rightarrow | 1 | 2 | 3 | 4 | 5 | 6 |
|---------------------------|---|---|----|-----|----|---|
| Estimated \downarrow | | | | | | |
| IR (row): 4 | 1 | 1 | 33 | 238 | 9 | 8 |
| Ra (row): 5 | 6 | 4 | 2 | 3 | 5 | 0 |
| Max prod at 4 | 6 | 4 | 66 | 714 | 45 | 0 |

Note that not necessarily the maximum values appear in elements having the same column number as the estimated class number. For instance, the class estimate from radar, no. 5, differs from the position of maximum element value, no. 1. Now, we look for the maximum of the products of corresponding elements in these rows. The

product of elements, taken column-wise, has its maximum at element no. 4, which we consider as the number of preferred class after fusion. This is a natural result since IR decides on class 4 with much higher confidence than radar decides on class 5 (class 1 nearly as probable).

Given that the ANN outputs can be translated to approximate probabilities, decision fusion using Bayes rule (see [6], pp. 51-52), here named *type B*, is a possible alternative. Assuming further, at least to some degree, uncorrelated observations we may update the sequential estimates using Bayes rule in recursive form. Once again we resort to historical data, summarised in confusion matrices with each element supplemented with the average ANN output value. Two categories of matrix elements are considered. Diagonal elements stand for true classifications, off-diagonal elements represent false classifications. We calculate the number densities of each category of values as a function of ANN output. The ratio of $\text{density}(\text{true})/(\text{density}(\text{true}+\text{false}))$ is finally taken as an approximate probability for estimating true class as function of ANN output value. This procedure is repeated separately for matrices based on IR, radar and fusion data. The resulting ANN output vs probability curves are not always monotonically rising, the reason being that sometimes the ANN output is considerably less than the maximal value for a correct classification (and, normally, even lesser for the false cases). However, given enough training data, the ANN classifier may indeed approximate the *a posteriori* class probabilities, see [10, pp. 184-187].

In contrast to feature fusion, decision fusion may be based on different classification methods, e.g. a neural network for IR and correlation for radar, whereas feature fusion normally uses only one classifier. On the other hand, feature fusion offers classification with no *a priori* knowledge of classification confidence in a particular scenario (it should, however, be noted that training results do not reveal the *true* sensor performance).

6 Results

Results from the two cases are shown in Fig. 7, where comparison is made between results from individual classifications feature fusion and decision fusion of type A. The average fraction of correctly classified targets is given as a function of range from 489 m to 5460 m in log scale. The IR targets have been degraded by up to 25 percent pepper-and-salt noise. Training with the neural network involves data using every second target orientation aspect in azimuth, and in a wide range interval, limited by ranges estimated from the elevation angle combined with platform altitude. Radar targets have been degraded by clutter (see section 2.3) plus extra noise for robust training using the neural network. Here, training is carried out using a restricted set of orientation aspects in azimuth, 45, 45.1, 45.5, 47 and 53 degrees, and the more precise ranges from the radar. The benefit of data fusion depends on the choice of method.

Case 1 uses a sparse sample of IR and radar data without SVD analysis. As can be seen from the top picture in Fig. 7, feature fusion outperforms the best sensor (IR) for ranges above 1600 m (log value 3.2). At shorter ranges feature fusion only outrivals the worst performing sensor

(radar), an expected result since fusion does not work with really bad data lacking a measure of their quality. Another factor is the unbalanced number of IR and radar features, with an increased risk for the ANN method to be trapped at local optima. Considering this fact it is perhaps not surprising that decision fusion of type A, much surpassing feature fusion, performs at least as well as the best sensor. As said before, however, feature fusion uses only information acquired in real time. The *a priori* information used for decision fusion may be out of date. The true performance should lie somewhere between these two extremes.

Case 2 is focused on SVD analysis of a rich sample of radar data used for training the neural network, see Fig. 7, bottom picture. The training data set for the IR is the same as in case 1. Test data for both IR and radar are still sampled sparsely with equal $\log_{10}(\text{range})$ step size. In this case, the radar performance matches IR well, even surpassing IR at shorter ranges. Data fusion is carried out under the thought assumption that the IR and radar sensors observe the target from different platforms showing widely different azimuth angles.

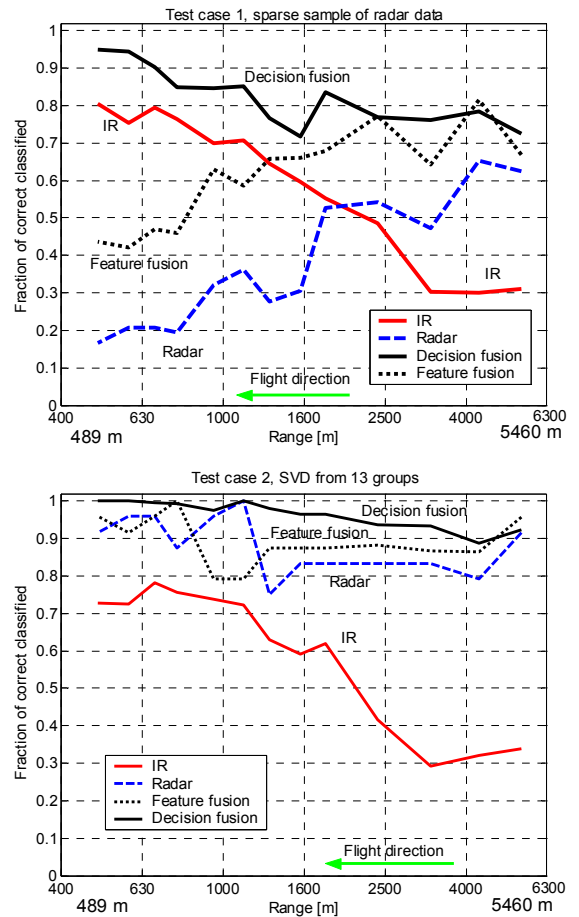


Fig. 7. Comparison of classification from individual sensors, feature fusion and decision fusion of type A. The average fraction of correctly classified targets is shown as function of range from 489 m to 5460 m. Top: Case 1, Bottom: Case 2 (SVD analysis).

Fig. 7, bottom picture, shows such a case with IR observing a target at azimuth angles 2, 32, 62 and 92 degrees while, simultaneously, the radar observes the target at azimuth 45.05, 45.2, 46 and 49 degrees. Note the much better performance of radar when using SVD. Here, feature fusion outperforms the best sensor at ranges exceeding 1600 m, while at shorter ranges, with both IR and radar displaying high performance, fusion matches the best sensor well. Again, decision fusion of type A outperforms feature fusion on all ranges. Additional experiments with articulated targets (gun turrets turned off the forward position) show, as expected, degraded classification performance depending on the degree of articulation.

As can be seen from Fig. 7, feature fusion, in its present form, does not fulfil all expectations. Already pointed out for case 1, one sensitive factor may be the way of balancing the IR and radar features. We are currently working on a method to reduce the number of radar features based on underlying physics. However, given that the ANN outputs can be converted to approximate probabilities, decision fusion is also possible using Bayes rule (see section 5, decision fusion type B).

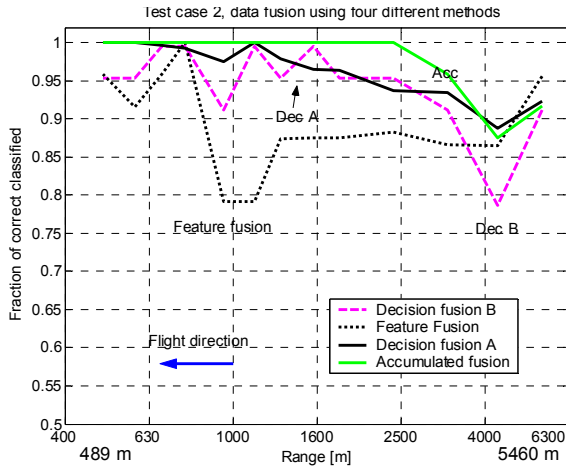


Figure 8. Comparison of four different data fusion methods. Case 2, mean over all targets and chosen orientation angles. In addition to feature fusion and decision fusion of type A, already shown in Fig. 7, decision fusion of type B (see text) and accumulated fusion are based on Bayes method using approximate probabilities for correct class. Note that vertical scale begins at 0.5.

Figure 8 shows the result of applying Bayes rule to do decision fusion. Here, decision fusion of type B refers to local fusion of individual IR and radar classifications using estimated probabilities for correct class. Taken in sequence, class estimates from decision fusion B are updated (accumulated fusion) using Bayes rule in recursive form. We note that, except for the first two starting steps beyond 4000 m range, decision fusion of both types (A & B) result in 95 percent correct classification on the average. Of these two, decision fusion of type A performs most consistently towards

correct classification approaching a target. Feature fusion lies much behind with an occasional dip about range 1000 m. Superior performance is achieved by accumulated fusion which attains 100 percent correct classification at ranges closer than 2500 meters. As a safeguard, the Bayes recursion algorithm has been run with *a priori* probabilities within [0.1, 0.9]. The question of partly correlated observations is postponed to further investigations.

Averaging the results over all aspects and ranges the classification performance for each class is conveniently shown in so-called confusion diagrams, see Fig. 9, case 1. The top two pictures show the performance for the individual sensors, IR and radar. The bottom two pictures show the results for feature and decision fusion type A in turn. A perfect classification should give results on the diagonal only. Generally, feature fusion improves classification compared with classification using individual sensor data. The major off-diagonal elements in columns 1-2 and rows 4-5 depend on too sparse training data for the radar, something that was remedied by the SVD approach in case 2. The bottom right diagram shows the outstanding performance of decision fusion type A, assuming that the *a priori* data can be trusted.

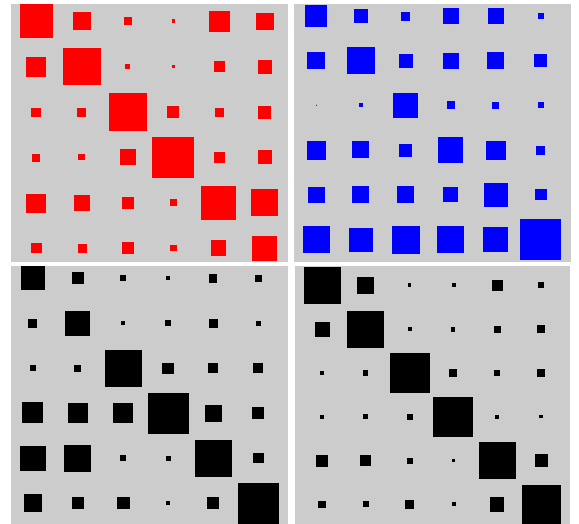


Fig. 9. Confusion diagrams for IR and radar (top), feature fusion and decision fusion of type A (bottom), case 1. Average results over all aspect angles. Columns: correct target types 1-6. Rows: proposed target type after classification. The number of associated correct-proposed pairs is proportional to the area of squares.

7 Conclusions

Ground targets have been classified by a neural network using simulated IR and radar data with variable success, from almost undetermined to about 95 percent correct or more, depending on target slant range, aspect angle and sensor combination. A small number of orientation invariant features extracted from IR data, together with a much larger set of HRR features from radar, participate in classification of (assumed) already detected objects. The classification is much improved by scaling distance dependent IR features by ranges from radar. A big

improvement in classification is also achieved by sampling high-resolution radar range profiles with equal steps in sine of elevation angle, followed by a singular value decomposition. Generally, decision fusion surpasses feature fusion, which in turn outperforms individual IR or radar sensors with proper respect to data uncertainty. Finally, assuming independent observations, sequential update of class estimates eventually attain an ultimate performance of 100 percent correct classification.

References

- [1] Fredrik Bennet and Stefan Fenelius, "SceneServer – a 3D software assisting developers of computer vision algorithms", Technical Report FOI-R--0831-SE, Swedish Defence Research Agency (FOI), Sensor Technology, Linköping, Sweden, 2003.
- [2] Nicolas C. Currie, Robert D. Hayes, and Robert N. Trebits. *Millimeter-Wave Radar Clutter*. Artech House, Norwood, MA, 1992.
- [3] L.M. Novak and C.M. Netishen. Polarimetric Synthetic Aperture Radar Imaging. *Int. J. Imaging Systems and Technology*, 4:306-318, 1992.
- [4] Chris Oliver and Shaun Quegan. *Understanding Synthetic Aperture Radar Images*. Artech House, Norwood, MA, 1998.
- [5] Rahman Aljasmi, Leif Carlsson, Staffan Gadd, Jan Gustavsson, Magnus Gustafsson, Nils-Uno Jonsson, Mikael Karlsson, Nils Karlsson, Andris Lauberts, Fredrik Näsström, Ain Sume, Mathias Wilow, Anders Örbom. Multisensor Target Seeker – Classification of ground targets using datafusion – 3. User Report FOI-R--1119--SE, Swedish Defence Research Agency (FOI), Sensor Technology, Linköping, Sweden, 2003 [In Swedish].
- [6] Keinosuke Fukunaga, *Introduction to statistical pattern recognition*. 2nd Ed., San Diego, CA, Academic Press, 1990.
- [7] Arnab K. Shaw and Vijay Bhatnagar. Automatic target recognition using eigen-templates. In Edmund G. Zelnio, editor, Proc. SPIE Conf. Algorithms for Synthetic Aperture Radar Imagery V, volume 3370, pages 448-459, SPIE, Bellingham, Washington, 1998.
- [8] Richard A. Mitchell, John J. Westerkamp. Robust Statistical Feature Based Aircraft Identification. *IEEE Transactions on Aerospace and Electronic Systems*, 35:1077-1094, July 1999.
- [9] X. Liao, P. Runkle, and L. Carin. Identification of Ground Targets From Sequential High-Range-Resolution Radar Signatures. *IEEE Transactions on Aerospace and Electronic Systems*, 38:1230-1242, October 2002.
- [10] Simon Haykin. *Neural Networks - A Comprehensive Foundation*. 2nd Ed., Prentice Hall, New Jersey, 1999.