MULTI-PERSPECTIVE IMAGING AND IMAGE INTERPRETATION

Chris J. Baker, H. D. Griffiths and Michele Vespe
Department of Electronic and Electrical Engineering
University College London, London, UK

Abstract. High resolution range profiling and imaging have been the principal methods by which more and more detailed target information can be collected by radar systems. The level of detail that can be established may then be used to attempt classification. However, this has typically been achieved using monostatic radar viewing targets from a single perspective. In this chapter methods for achieving very high resolutions will be reviewed. The techniques will include wide instantaneous bandwidths, stepped frequency and aperture synthesis. Examples showing the angular dependency of high range resolution profiles and two-dimensional imagery of real, full scale targets are presented. This data is examined as a basis for target classification and highlights how features observed relate to the structures that compose the target. A number of classification techniques will be introduced including statistical, feature vector and neural based approaches. These will be combined into a new method of classification that exploits multiple perspectives. Results will be presented, again based upon analysis of real target signatures and are used to examine the selection of perspectives to improve the overall classification performance.

Key words: ATR, radar imaging, multi-perspective classification

1. Introduction

Target classification by radar offers the possibility of remotely identifying objects at ranges well in excess of those of any other sensor. Indeed radar has been employed in operational systems for many years. However, the systems use human interpretation of radar data and performance is generally unreliable and slow. Nevertheless, the benefits of fast and reliable classification are enormous and have the potential for opening huge areas of new application. Research in recent years has been intense but still, automated or semi-automated classification able to work acceptably well in all conditions seems a long way off. The prime approach to developing classification algorithms has been to use higher and higher spatial resolutions, either one dimensional range profiles (Hu and Zhu, 1997) or two-dimensional imagery (Novak et al., 1997). High resolution increases the level of detail in the data to be classified and this has generally been seen as providing more and better information. However, the performance of classifiers, whilst very good against a limited set of free space measurements is much less satisfactory when applied to operationally realistic conditions. In this chapter we will review

methods for obtaining high resolution, use these to generate high resolution target signatures and subsequently illustrate some of their important aspects that require careful understanding if classification is to be successful. We then examine some typical classifiers before considering in more detail an approach to classification that uses a multiplicity of perspectives as the data input. This also enables much information about the nature of target signatures and their basis for classification to be evaluated.

Firstly though, the concepts of resolution and classification are discussed as these terms are often used with imprecise or varying meanings. Most often resolution is defined as the ability of radar (or any sensor) to distinguish between two closely spaced scatterers. A measure of this ability is captured in the radar system point spread function or impulse response function with resolving power being determined by the 3-dB points. This is a reasonable definition but care needs to be taken as there is an implicit assumption that the target to be considered has point like properties. This is well known not to be the case, but nevertheless it has proved a useful descriptor of radar performance. As will be seen later, if resolution is improved more scatterers on a target can be distinguished from one another and there is undoubtedly an improved amount of detail in the resulting signature. Care also has to be taken with synthetic aperture imaging in two dimensions such as SAR and ISAR. These imaging techniques again have an implicit assumption that targets have point like properties and are continuously illuminated during formation of the image. Once again most targets are not points and quite often there is occlusion of one target by another. For example high placed scatterers near the front of a target often place the scatterers further back in shadow and hence they are not imaged. Thus the resolution in a typical SAR image is not necessarily constant, a fact often overlooked by developers of classification algorithms. However, once again, these assumptions have still resulted in robust SAR and ISAR image generation and as larger apertures are synthesised there is an undeniable increase in detail in the resulting imagery. We will return to these themes as we explore high resolution techniques and approaches to classification of the resulting target signatures.

2. High down range resolution techniques

Despite the reservations discussed above we nevertheless begin by modelling the reflectivity from a target as the coherent sum of a series of spatially separated scatterers (Keller, 1962). Its simplicity offers valuable insight into some of the scattering processes observed. Thus the frequency domain reflectivity function $\zeta_\theta(f)$ of a complex target illuminated at a given aspect angle θ by an incident field of frequency $\lambda \ll D$, where D is the physical dimension of the target, is

given by:

$$\zeta_\theta(f) = \sum_{i=1}^{N} \zeta_\theta^i(f) \tag{1}$$

where

$$\zeta_\theta^i(f) = A_\theta^i(f)\ \exp[-j\,\vartheta_\theta^i(f)]. \tag{2}$$

Thus we see that the amplitude and phase of the *i*-th scatterer depend on both frequency and aspect angle. An inverse FFT will yield the complex reflectivity function or range profile, where the magnitude in each of the IFFT bins represents the magnitude of the reflections from the scatterers in a range resolution cell.

The resolution in the range dimension is related directly to the pulse width (T_p) of the waveform. Two targets of equal RCS are said to be recognized as being resolved in range when they are separated from each other by a distance:

$$d = \frac{c\,T_p}{2}. \tag{3}$$

This equation tells us the range resolution of a pulsed radar system when the pulse is not modulated. Thus a very short pulse is needed if high range resolution is required. Indeed to resolve all the scatterers unambiguously the pulse length has to be short enough such that only one scatterer appears in each range cell. Normally as high a range resolution as possible is used and it is accepted that not all scatterers will be resolved. The resulting ambiguity is one source of ensuing difficulty in the next stage of classification.

In long range radar, a long pulse is needed to ensure sufficient energy to detect small targets. However, this long length waveform has poor range resolution. The use of a short duration pulse in a long range radar system implies that a very high peak power is required. There is a limitation on just how high the peak power of the short pulse can be. Ultimately the voltage required will ‘arc’ or breakdown, risking damage to the circuitry and poor efficiency of transmission.

One method which is commonplace today that overcomes this limitation is pulse compression (Knott et al., 1985).

Pulse compression is a technique that consists of applying a modulation to a long pulse or waveform such that the bandwidth B of the modulation is greater than that of the un-modulated pulse (i.e. $1/T_p$). On reception the long pulse is processed by a matched filter to obtain the equivalent resolution of a short pulse (of width $1/B$). In the time domain the received signal is correlated with a time reversed replica of the transmitted signal (delayed to the chosen range). This compresses the waveform into a single shorter duration that is equivalent to a pulse of length given by $1/B$. In other words it is determined by the modulation bandwidth and not the pulse duration. Thus the limitation of poor range resolution using long pulses is overcome and long range as well as high resolution can both be achieved

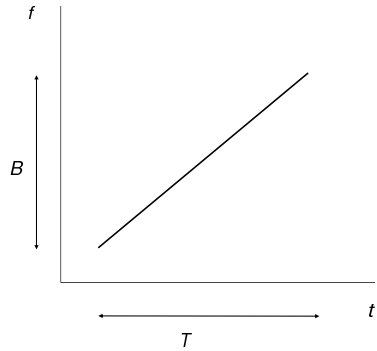

Figure 1. Changing the frequency as a function of time across the duration of a pulse. B is the modulation bandwidth and T the pulse duration.

together. Here pulse compression can be achieved by frequency, phase and amplitude modulation. The most common form of modulation used is to change the frequency from the start to the end of the pulse such that the required bandwidth B is swept out as shown in figure 1. This leads to a resolution given by:

$$d = \frac{c}{2B}. \tag{4}$$

To achieve even higher range resolutions a frequency modulated stepped-frequency compressed waveform may be employed. This reduces the instantaneous modulation bandwidth requirement while increasing the overall bandwidth. In other words the necessary wider bandwidth waveform is synthesised using a number of pulses. However, note that this has the disadvantage of collecting the individual waveforms over a finite period of time making the required coherency vulnerable to target motion. High Range Resolution (HRR) profiles are subsequently produced by processing a wideband reconstruction of a targets reflectivity spectrum in the Frequency Domain (Wilkinson et al., 1998). By placing side by side N narrower chirp waveforms or pulses of bandwidth B and using an inter-pulse frequency step increment also equal to B, it is possible to synthesise a total bandwidth of $B_t = (N + 1)B/2$. This is illustrated in figure 2.

A range profile may then be defined as a time sequence of the vector sum of signals reflected back by different scatterers within a range cell. By matched filtering this stepped frequency waveform a range resolution $d = c/(2NB)$ is achieved. The received signal is the transmitted pulse modulated by the corresponding subspectrum representing the target reflectivity function. By adding the compressed individual portions of reflectivity function, which result from time convolution between each received pulse with the complex conjugate of the corresponding transmitted pulse, the entire spectrum is eventually obtained commensurate with the extended bandwidth. The HRR profile may then be synthesised from an in-

verse FFT applied to each row of the time history of the target's frequency domain signature matrix.

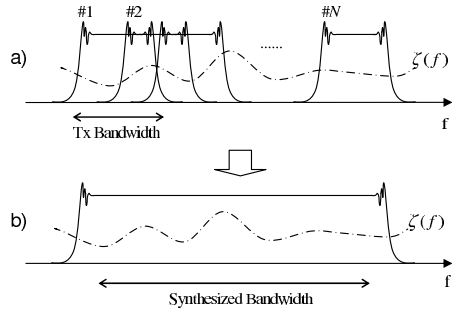

Figure 2. Spectrum reconstruction of the target reflectivity function using a series of chirp pulses (a) and the corresponding synthesised bandwidth (b).

3. High cross range resolution

In real aperture radar the cross range resolution is determined by the antenna beamwidth. This is often thought of as being so large that it effectively produces no resolution of any value and is why HRR profiles are often referred to as one-dimensional target signatures. In reality they are two dimensional images where one of the dimensions of resolution is very poor.

The real beamwidth (or cross range resolving power) is defined as the width at the half power or 3-dB points from the peak of the main lobe. The beamwidths may be the same in both the elevation (vertical) and the azimuth (horizontal) dimensions, although this is by no means mandatory. The beamwidth in radians is a function of antenna size and transmission wavelength. For a circular aperture (i.e. a circular profile of a parabolic dish) the beamwidth in radians at the half power or 3-dB points is given approximately by:

$$B_{az} = \frac{\lambda}{D}. \tag{5}$$

This means that the cross-range extent of the beam at a range R is given by:

$$R_{az} = R\frac{\lambda}{D}. \tag{6}$$

Thus for a range of only 10 km and a wavelength of 3 cm, a 0.5 m diameter antenna will have a cross range resolution of 600 m. This contrasts with a frequency modulated pulse bandwidth of 500 MHz leading to a range resolution of 30 cm. It is typical to have equal down and cross range resolutions.

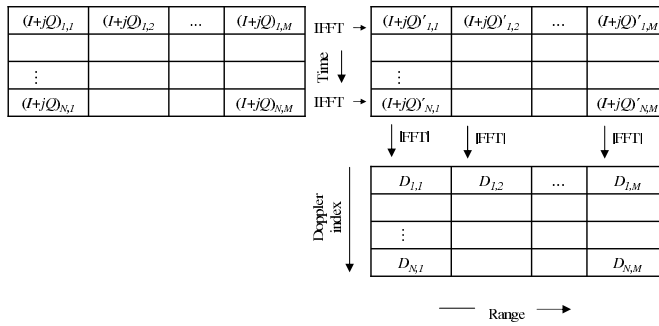

Figure 3. Matrix decomposition for HRR profiles and ISAR imagery.

Two techniques by which much higher cross range resolution may be obtained are SAR and ISAR. These are essentially completely equivalent and rely on viewing a target over a finite angular ambit to create a synthetic aperture with a length in excess of that of the real radar aperture and hence able to provide a higher resolution. The greater the angular ambit traversed, the greater the length of the synthesised aperture and the higher the cross range resolution. Here we will confine our discussions to ISAR imaging only and the interested reader is referred to excellent textbooks that cover SAR imaging (Curlander and McDonough, 1991; Carrara et al., 1995).

To form an ISAR image the HRR profiles assembled using the step frequency technique reviewed in the previous section are used. To obtain the magnitude of the ISAR image's pixel in l-th range cell and j-th Doppler cell ($D_{l,j}$) an FFT is applied to each column of the matrix representing the time history of target's range profile. This is demonstrated in figure 3.

For ISAR, the resolution achieved in cross-range depends upon the minimum resolvable frequency Δf_D between two adjacent scatterers (Wehner, 1995). Doppler resolution is also related to the available coherent time of integration T which is equal to the time required to collect the N chirp returns. Therefore, consecutive reflectivity samples from the same range cells are taken every $N\Delta T$ seconds:

$$\Delta f_D = \frac{1}{N\Delta T} \approx \frac{1}{T}. \tag{7}$$

As a consequence, the cross-range resolution Δr_c can be written as:

$$\Delta r_c = \frac{c\,\Delta f_D}{2\,\omega_0\,f_c} = \frac{\lambda}{2\,\omega_0\,T} \tag{8}$$

where $\lambda = c/f_c$ is the illuminating wavelength and ω_0 is the angular velocity of the target rotational motion.

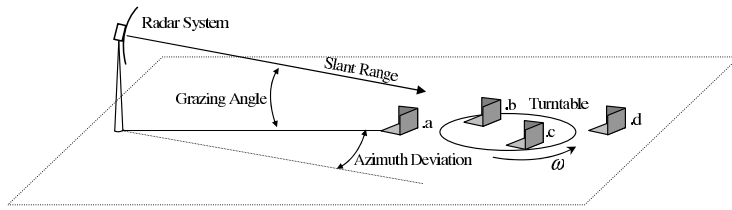

Figure 4. ISAR geometry: two stationary corner reflectors are in front of and behind the turntable, while two rotating ones are placed on the turntable.

In ISAR image processing the motion of the target is usually unknown. The target motion can be seen as the superposition of rotational and translational motions with respect to the radar system. If the former contributes to the ability to resolve in cross-range, in order to obtain a focused image of the target it is necessary to compensate for phase errors due to the translational motion occurring during data collection. This is usually referred to as motion compensation. After this correction an ISAR image may be obtained by processing data collected over an arc of circular aperture whose dimensions depend on the rotational speed of the target ω_0 and the integration time T. There may still be further residual motions that require correction using autofocus techniques. We now introduce experiments by which HRR profiles and ISAR images are generated.

4. High resolution target signatures

In this section we exploit the high resolution techniques introduced above to examine the form of the resulting signatures. To do this we use measurements made of calibration and vehicle targets that have been mounted on a slowly rotating platform. Figure 4 shows the experimental set up. Here the radar system views a turntable at a shallow grazing angle of a few degrees. The figure shows corner reflector calibration targets in place. Two are located on the turntable and two are located in front of and behind the turntable. Additional experiments have also been performed with the corners removed and a vehicle target situated on the turntable instead. The profiles are generated from an X-band radar having an instantaneous bandwidth of 500 MHz. Eight frequency steps spaced by 250 MHz are used to synthesise a total bandwidth of 2.25 GHz. The turntable data is first enhanced by removing any stationary clutter (Showman et al., 1998). Estimation and subtraction have been performed in the frequency domain.

Figure 5 shows the resulting range profiles and their variation as the turntable rotates through 360 degrees. The two stationary trihedrals show a constant response at near and far range as expected. For the two rotating trihedral targets,

when the line-of-sight is on the trihedral bisector, a peak of reflection occurs. This is consistent with the expected theoretical response. As the trihedral targets rotate, the backscattered field decreases progressively until a point is reached where there is a peak of specular reflection. This is a reflection from one of the sides making up the trihedral which is orthogonal to the illuminating radar system (i.e. it faces the radar beam and looks like a flat plate reflector). At increasing rotation angles the RCS of the target drops since the orientation of the trihedral is such that it tends to reflect incident radiation away from the radar.

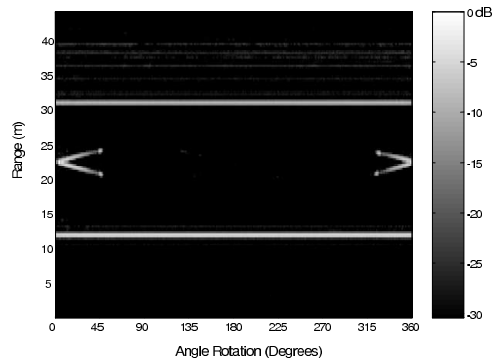

Figure 5. History of HRR range profiles (30 cm of range resolution) from four corner reflectors, two rotating and two stationary. The zero degrees angle is in correspondence of the two corners having the same range relative to the system and facing directly the radar beam.

This angular dependency of the RCS of a well known reflector such as a trihedral begins to illustrate how the backscattering properties of real targets may vary with the orientation of observation. For example if a target has part of its structure that mimics a trihedral it may show this feature over a similarly limited angular range. Thus in a Multi-Perspective (M-P) environment, different angular samples of a target's signature should improve the likelihood of observing a corner or corner like reflector and recognising it as such. Particular shapes such as flat plates and corners can be common on many manmade structures and are often quite dominant features that may prove useful for classification.

In figure 6, the complete angular ambit of range profiles spanning 360 degrees from a Land Rover vehicle rotating on the turntable is shown.

This highlights a number of different scattering behaviours: the strong peaks from specular reflections (0°, 90°, 180°) appear over a very limited angle range and obscure nearby point-like backscattering. Corner-like returns can be observed at far range ($\sim$ 6 m) for two range angular spans ($\sim [10° - 60°]$ and $[130° - 180°]$). These returns correspond to the trihedral like structures formed at the rear of the Land Rover. This is a vehicle without the rear soft top and has a metallic bench seat that makes a corner where it joins the rear bulkhead of the drivers cabin. At $\sim$ 8 m range there is a double bounce return corresponding to one of the corners.

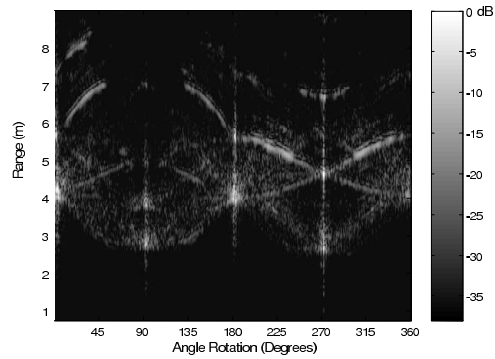

Figure 6. History of HRR range profiles (range resolution less than 15 cm) from a series of X-Band stepped frequency chirps illuminating a ground vehicle as it rotates over 360 degrees. At zero degrees, the target is broadside oriented, while at 90 degrees has its end-view towards the radar.

This type of effect increases the information that can be processed which would be otherwise impossible to reconstruct by a traditional single perspective approach. It also illustrates the complexity and subtlety of radar target scattering. Note also that here we are dealing with targets on turntables where there is no clutter or multipath and the signal to noise ratio is high. In more realistic scenarios the scattering from targets will in fact be even more complicated.

The turntable data can be processed using the ISAR technique to yield two-dimensional imagery.

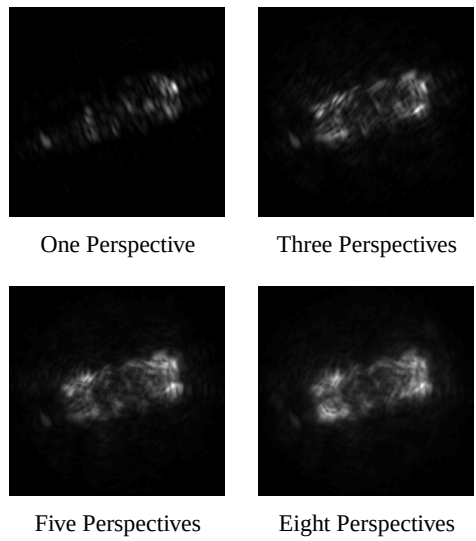

Figure 7. Multi-Look image reconstruction: (a) is generated using a single perspective, (b) three perspectives, (c) five perspectives and (d) eight perspectives.

Figure 7 shows a series of ISAR images of a Ford Cougar car in which the number of perspectives used to form the imagery is slowly increased. It is clear that when only a single perspective is used there is considerable self-shadowing by the target and the shape cannot be determined with good confidence. When all eight perspectives are employed then the effects of self-shadowing are largely eliminated and much more complete information is generated. Note, however, that the concept of resolving scatterers as seemed to be observed in the range profiles of the Land Rover is much less obvious. As with most modern vehicles the Cougar is designed to have low wind resistance and has little in the way of distinct scattering centres. This begins to place in question both the image formation process and classification techniques if they are using a point target scatterer model assumption.

Figure 8 shows a multi-look ISAR image of the Land Rover target used to generate the HRR profiles seen earlier. There are two 'point-like' scatterers that are consistent with the area where the rear metal bench seat meets the bulk head behind the drivers cabin and forms a corner like structure. There are some 'pointish-like' scatterers at the rear of the vehicle although these are rather less distinct. Otherwise the scattering appears quite 'un-point-like' and resembles in many ways the form of image for the Cougar. Nevertheless, the shape of the Land Rover is easily discernable and is clearly quite different to that of the Cougar. In this way we start to appreciate the complexity of electro-magnetic backscatter and the ability to reliably and confidently classify real targets.

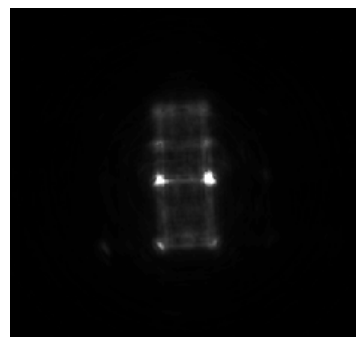

Figure 8. Multi-look image of the Land Rover target.

5. Target Classification

Having examined the detailed structure and composition of target signatures we now consider operating on them to make classification decisions. Robust and reliable target classification has the potential to radically increase the importance and value of radar systems. Interest in radar target classification has intensified due

to the increase of information available from advanced radar systems. As seen in the previous section this is mainly due to the backscattered signature having high resolution in range and cross-range. Although the description of target scattering is detailed enough to attempt radar target classification, the reliability of this crucial task is influenced by a number of unpredictable factors and is generally poor. For example, one and two-dimensional signatures can both appear significantly different even though the measurements are collected with the same radar and a very similar target geometry (measurement noise, rotational and translational range migration, speckle, etc). In addition, global or local shadowing effects also noticeably challenge the attempts to classify targets reliably, as will multipath.

As many sources of target signature variability depend on target orientation as viewed by the radar system, it may be possible to aid target recognition by using more than one perspective and hence average out the sources causing misclassification. Furthermore, the possible employment of bistatic modes of operation between the nodes of the network might additionally provide a valuable counter to stealth technology (Baker and Hume, 2003). In a monostatic netted radar scenario, the classification performance increases with the number of perspectives. Although the improvements are significant, those benefits are dependent on the location of the nodes with respect to the position of the target and its orientation.

6. Classification techniques

There are two main aspects to target classification. The first is to isolate the target returns from the clutter echoes (e.g. by filtering) and to extract the features that can help to distinguish the class of the target. The second aspect is related to the method used for performing the decision as to which class or target type the feature data belongs. When target classification is achieved through automatic computation it is usually referred to as Automatic Target Recognition (ATR).

In ATR the classification task requires complex techniques and there are a number of approaches that can be used. For example in a model based technique a model of the target is made by Computer Aided Design (CAD) and Electro-Magnetic simulations. This enables many simulated versions to be compared with the target signature to be classified. This is a computationally intensive technique. Alternatively, in a template matching based technique, many real versions of the target signatures (at a large number of geometries) are stored in a database and subsequently compared with the target detected in order to assign it to a class. Consequently, a very large database is needed. Further, if the target is altered in some way (e.g. a tank may carry some additional equipment) then the templates may no longer represent the modified signature and the classification can fail. Finally, pattern based techniques exploit features extracted from the input signature. These might include peak amplitudes and their locations in a HRR profile or ISAR image. These then are used to make a multi-dimensional feature vector which

can be compared with the stored feature vectors from previous measurements of known targets in order to perform classification. This technique is less costly in terms of computation and it is consistent with a netted radar framework as a number of perspectives of the same object are available.

Classification typically requires a high probability of declaration P_{dec} which is the probability that a detected target will be classified either correctly or incorrectly as a member of the template training set. A target is known when its feature-vector belongs to the training data set. A second performance requirement is to have a high probability of correct classification P_{cc}. These two parameters are often related each other: if P_{dec} is low, the ATR system may declare just those cases of very high certainty of correct classification. As a result, the system is able to achieve a very high P_{cc} but may not classify all possible targets. Finally, a low probability of false alarm P_{fa} is required, i.e. a low probability that an unknown target is incorrectly classified as a different object in the ATR database. This final task is particularly difficult since it is impossible to include a specific class of unknown targets in the template. This is because there is always a possibility that the classifier would need more information to make a correct decision.

The decision, which is usually made automatically, may be performed with two different data input types. The first is one-dimensional target classification by HRR profiles. The HRR profile can be thought of as representing the projection of the apparent target scattering centers onto the range axis. Hence the HRR profile is a one-dimensional feature vector.

Three classifiers are considered here, a more exhaustive treatment may be found in (Tait, 2006). These are (i) a Naïve Bayesian classifier, (ii) a Nearest Neighbour classifier and (iii) a classifier using Neural Networks.

6.1. NAÏVE BAYESIAN MULTI-PERSPECTIVE CLASSIFIER

The Naïve Bayesian classifier is a pattern recognition statistical method (Looney, 1998). The decision-making is reduced to pure calculations of feature probabilities. The training set is stored in different databases of templates and therefore the number of classes is known as well as sets of their representative vectors (i.e. the learning strategy is supervised). We consider a set of n_c classes $\{C_i : i = 1, ..., n_c\}$ and a single HRR profile X formed by a sequence of n elements $x_1, x_2, ..., x_n$ as the feature vector. The classifier decides that X belongs to the class C_i showing the highest posterior probability $P(C_i \mid X)$. This is done using Bayes' theorem to calculate the unknown posterior probability of classes conditioned on the unknown feature vector to be classified:

$$P(C_i \mid X) = \frac{P(X \mid C_i)P(C_i)}{P(X)}. \tag{9}$$

The Naïve Bayesian classifier is based on the assumption of class independence of each attribute of the feature vector. Furthermore, since the values $x_1, x_2, ..., x_n$ are

continuous, they are assumed Gaussian distributed. Their statistical parameters are deduced from the training set and used to calculate the conditional probability $P(x_i \mid C_i)$ for the attribute x_i and, eventually, the likelihood ratio test:

$$\text{If} \quad \frac{P(X \mid C_i)}{P(X \mid C_j)} > \frac{P(C_j)}{P(C_i)} \implies X \in \text{class } i. \tag{10}$$

These concepts are integrated for the M-P Naïve Bayesian classifier. Here we consider a network of N radars and the sequence of 1-D signatures $\{X_j : j = 1, ..., N\}$ is the information collected by the system. The posterior probability $P(C_i \mid X_1, ..., X_N)$ of the sequence of range profiles conditioned to C_i is the probability that the sequence belongs to that class and by applying Bayes' theorem, can be expressed as follows:

$$P(C_i \mid X_1, ..., X_N) = \prod_{j=1}^{N} \frac{P(X_j \mid C_i)P(C_i)}{P(X_j)}. \tag{11}$$

Assuming constant the probability $P(X_1, ..., X_N)$ with a fixed number of perspectives, the final decision is made for the class C_i that maximises $P(C_i \mid X_1, ..., X_N)$. This procedure enables the distinction of a single-perspective stage where all the conditional probabilities $P(X_j \mid C_i)$ are computed separately for each perspective.

6.2. *K*-NN MULTI-PERSPECTIVE CLASSIFIER

The *K*-Nearest Neighbours (*K*-NN) is a non-parametric approach to classification (Duda et al., 2001). It consists of measuring and selecting the minimum K distances from the feature vector to be classified and comparing with the templates of the different classes. Consider an input vector X and a population made up of a set of classes $\{C_i : i = 1, ..., n_c\}$. Then the distances from $d_{i,j}$ to $T_{i,j}$, the t-th template vector of the i-th class, can be computed and stored:

$$d_{i,j} = d(X, T_{i,j}) = \|X - T_{i,j}\|. \tag{12}$$

Subsequently, the K minimum scalars $d_{i,j}$ are selected from each class forming a K-dimensional vector D labelled in ascending order. The final decision is made on the basis of the largest number of votes over the K dimensional vector obtained.

There are three stages of the M-P *K*-NN classifier which are implemented as follows:

1. The Mono-Perspective stage: after the collection of the sequence of feature vectors $\{X_j : j = 1, ..., N\}$ where N is the number of sensors in the network, the same number of single-perspective classifiers is implemented. The j-th classifier computes a vector D_j consisting of the K minimum distances from the templates.

2. M-P processing: the whole of the vector D_j is processed and the minimum K distances are selected giving a weight for the decision.
3. Classification: the input sequence of feature vectors is associated with the class with the greatest number of weights.

Different K values have been tested for this problem: the best trade-off between complexity and classification performance suggests a value $K = 5$ minimum distances.

6.3. FANN MULTI-PERSPECTIVE CLASSIFIER

Given a feature vector X, Artificial Neural Networks (ANN) learn how to execute the classification task by means of examples (Christodoulou and Georgiopoulos, 2001). They are able to analyse and associate an output corresponding to a particular class of objects. Feed-forward ANN (FANN) supervised with a back-propagation strategy can be implemented. During the *learning phase* the training samples are used to set internal parameters of the network, that is, after giving the templates as inputs to the classifier, the weights are modified on the basis of the distance between the desired and actual outputs of the network. Considering an input vector X and a population made up of a set of n_c classes, the *execution mode* consists of the calculation of the output vector $Y = (y_1, y_2, ..., y_{n_c})$. Since the unipolar sigmoid (or logistic) function was chosen as activation function, the elements of the output vectors range from zero to one. The ultimate decision is made for the i-th class, where i is the index of the maximum value of Y. If we assume N perspectives, the sequence of feature vectors $\{X_j : j = 1, ..., N\}$ is the input for the first stage. Each single-perspective network accepts a vector and the partial outputs $\{Y_j : j = 1, ..., N\}$ are calculated. Subsequently, the output $\overline{Y}$ of the M-P stage is the mean value of the partial outcomes, and the classification decision is finally made on the basis of the maximum index of $\overline{Y}$.

7. Multi-perspective classification

We concentrate on the particular situation in which a single non-cooperative target has been previously detected and tracked by the radar system. Pre-processing raw data is necessary in order to increase the quality of the radar signatures: the target region is isolated and made more prominent thanks to the noise level subtraction from the rest of the range profile (Zyweck and Bogner, 1996).

Principal discriminating factors for classification purposes are Range Resolution, Side-Lobe Level (SLL) and Noise Level. Higher resolution means better point scatterers separation but the question of compromise regarding how much resolution is needed for good cost-recognition is difficult to resolve. Generally, a high SLL means clearer range profiles but this also implies deterioration in

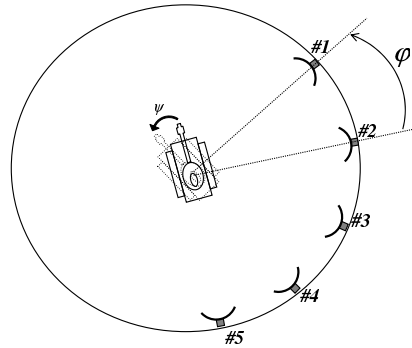

Figure 9. Multi-perspective environment deduced from ISAR geometry.

resolution. Eventually, a low noise level means high quality range profiles for classification.

Real ISAR turntable data have been used to produce HRR range profiles and images. Three vehicles classified as *A, B* and *C* constitute the sub-population problem. Each class is described by a set of range profiles covering a 360 degrees rotation of the turntable. Single chirp returns are compressed giving 30 cm range resolution. The grazing angle of the radar is 8 degrees and $2''$ of turntable rotation is the angular interval between two consecutive range profiles. Therefore, approximately 10000 range profiles are extracted from each data file over the complete rotation of 360 degrees. The training set of representative vectors for each class is made by 36 range profiles, taken approximately every 10 degrees for rotation of the target. The testing set of each class consists of the remaining range profiles excluding the templates. The features extracted after Principal Component Analysis (PCA) are the input attributes to the classifier.

The three algorithms have been implemented and tested in both single and multi-perspective environments. In this way any bias introduced by a single algorithm should be removed. Figure 9 represents a possible approximation of the multi-perspective scenario: each node of the network is assumed as having a fixed position as the target rotates by an angle ψ. The perspective angle φ is the angular displacement between the line-of-sights of two consecutive radars. From each of the radar positions either a series of range profiles can be generated as inputs to a one-dimensional classifier or they can be processed into an ISAR image that can be input to a two-dimensional classifier. It is therefore possible to perform classification using multiple perspectives.

In figure 10, the classification performance of the three previously described classifiers is plotted versus the number of perspectives used by the network. Because of the nature of the data and the small available number of targets, the

classifiers present a high level of performance when using only a single aspect angle. Improved performance is achieved increasing the number of radars in the network but the greatest improvement in performance can be appreciated with a small number of radars.

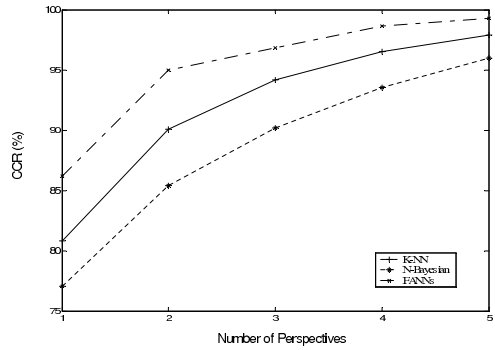

Figure 10. Multi-perspective classification rates using different numbers of perspectives.

Since the number of perspectives, and therefore the number of radars, is strictly related to complexity, costs and time burden of the network, in terms of classification purposes, it may be a reasonable trade-off to implement networks involving a small number of nodes. However, this analysis is against a small number of target classes and these conclusions require further verification.

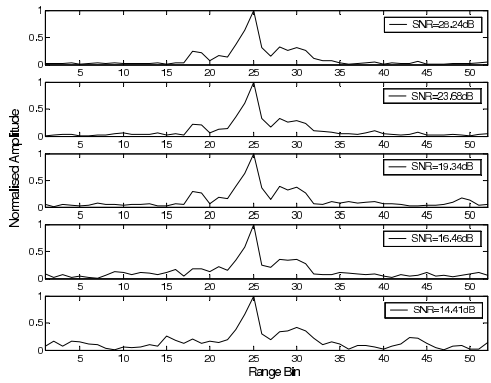

Figure 11. Different SNR levels of the same range profile.

The extent to which SNR affects classification and whether multi-perspective scenarios are effective at different SNR levels are now examined. The FANN classifier has been used for this particular task. The range profiles I and Q components are corrupted with additive white Gaussian noise. The original data, after noise removal, has a 28.24 dB SNR. Subsequently, the classifier is tested with range profiles presenting a progressively lower SNR.

In figure 11, five different SNR levels of the same range profile from class *A* are represented. From the particular orientation, the length of the object is 5.5 m (spanning almost 18 range bins). As the SNR decreases, some of the useful features become less distinct, making the range profile more difficult to be classified.

In figure 12, the performance of the FANN classifier only is plotted versus the number of perspectives used and SNR levels, showing how the enhancement in classification varies with different noise levels. The graph illustrates an increase in classification performance with numbers of perspectives in each case, particularly valuable for lowest SNR levels. However, below an SNR of 17 dB the performance quickly degrades indicating that classifiers will be upset by relatively small amounts of noise.

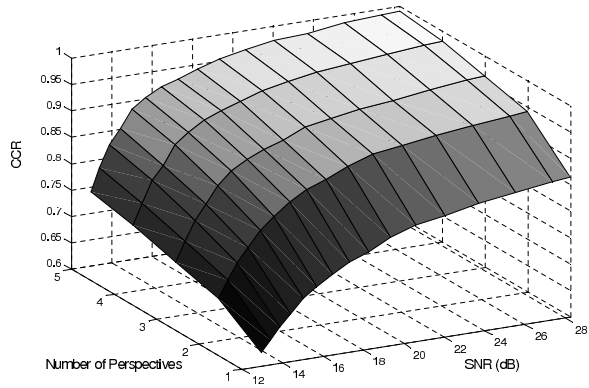

Figure 12. Correct classification rates for different SNR levels and netted radars.

8. Choice of perspectives

We now examine the geometrical relationships of nodes and target location and orientation using full scale measurements of representative targets. We start by re-visiting the neural network classification process in a little more detail and examine the classification performance as a function of the angle between the two perspectives taken. In a Two-Perspective (2-P) scenario, the parameter that distinguishes the perspective node locations is their relative angular displacement $\Delta\phi_{1,2} = \phi_2 - \phi_1$. Hence, after fixing $\Delta\phi_{1,2}$, the 2-Perspective classifier is tested with all possible pairs of HRR profiles displaced by that angle covering all the possible orientations of the target. Having a test set consisting of N profiles, the same number of pairs can be formed to test the 2-Perspective classifier.

The training set of representative vectors for each class is made up by 36 range profiles, taken approximately every 10° degrees of target rotation. The testing set

of each class consists of the remaining range profiles neglecting the templates. The M-P classifier can be seen as the combination of N single-perspective classifiers (i.e. N nodes in the network) whereas the eventual decision is made by processing the outputs of each FANN. For this first stage of investigation, the angle $\Delta\phi_{1,2}$ is not processed as information by the M-P classifier.

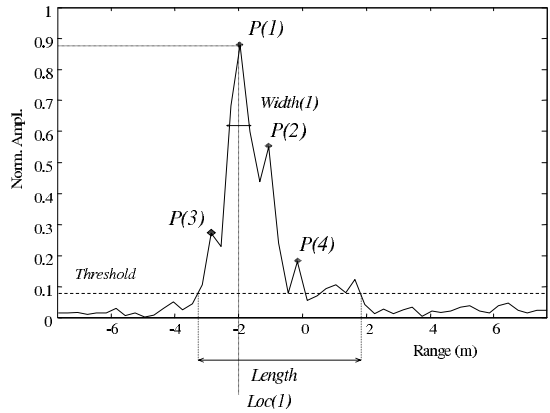

Figure 13. Feature extraction from HRR profile.

Features from radar signatures have been extracted in order to reduce the intrinsic redundancy of data and simplify the model. This methodology, to a certain extent, decreases the overall correct classification rates but, on the other hand, makes the representation model in the feature space less complex yielding a classification process consistent with those features effectively characterising the object. A typical HRR profile is shown in figure 13, where a threshold is applied to the HRR profile. The threshold is determined by measuring the mean intensity value neglecting the maximum peak, after normalising the profile. This guarantees adaptive target area isolation and less sensitiveness to the main scatterer reflection. Therefore, the radar length of the target for that particular orientation is measured as the distance between the first and last threshold crossings. This is the first component of the feature vector f. The second component is a measure of the average backscattering of the target whilst the successive M triples contain the information of the M peaks extracted in terms of amplitude, location and width. If the number P of peaks above the threshold is less than M, the last $M-P$ triples are set to zero (Huaitie et al., 1997). Different numbers of peaks have been extracted until the classification process revealed a certain degree of robustness. For $M = 4$, the feature vector has a dimension of 14 elements while 52 range bin values make up the raw echo profile.

Next, PCA has been applied to the profiles in order to better separate the classes in the feature space. By choosing the largest K eigenvectors from the original 14 of the covariance matrix, the dimension of the new feature vector f'

is additionally reduced. In table I the resulting confusion matrix for $K = 10$ is shown for the four-class population problem for a single perspective classifier.

TABLE I. Single-perspective FANN confusion matrix on features after PCA using $K = 10$ (correct classification rate = 74.4%).

Output → Input ↓	Class *A*	Class *B*	Class *C*	Class *D*
Class *A*	75.85	11.19	9.57	3.39
Class *B*	4.75	80.91	13.26	1.08
Class *C*	2.81	18.98	69.94	8.27
Class *D*	12.91	0.76	15.26	71.07

Radar classification is highly dependent on the object orientation. Since radar targets are usually man-made objects, they often present a number of 3-D symmetries. In this work, a low grazing angle is used to simulate a ground-based scenario, where the radar system, the target and their relative motion vectors lie on the same plane. This geometry allows us to consider the 2-D problem, whereby the perspectives represented by 1-D signatures collected by the measurement system are on the same plane. For example, for most of the possible 2-D ground-vehicle orientations, with 180° between the two perspectives, the corresponding profiles might expect to be quite highly correlated. This is due to the 180° symmetry typically exhibited by vehicles and hence little extra information is added. If this is the case it will cause a reduction in target characterisation and eventually, of M-P Correct Classification Rate (CRR) improvement. However, details such as rear-view mirrors, the antenna position and any non-symmetrical or moving parts will change the two signatures, producing CRR benefits when compared with single-perspective classifiers.

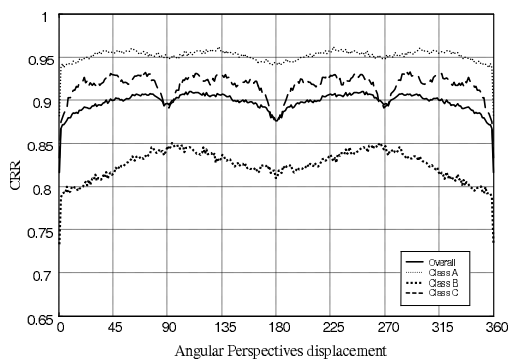

Figure 14. 2-P correct classification rates versus the angular displacement $\Delta\phi_{1,2}$ for the three-class problem.

We now consider the relationship between nodes location and target orientation, investigated for the cases of both two and three-perspective classifiers. In a 2-P scenario the angular perspective displacement between radars is the discriminant factor for the combined CRR: as the two range profiles de-correlate (i.e. $\Delta\phi_{1,2}$ increases) the information content of the pair-increases.

Figure 14 shows the classification performance as a function of angular separation of the two perspectives. The equivalent monostatic CRR is shown at the 0° and 360° positions. The single target classification rates show how the global accuracy depends on the peculiar geometric features of the target. The drop at $\Delta\phi_{1,2} = 180°$, as previously hypothesised, is due to the multiple axes of symmetries of the targets and it is visible for all the classes. For targets A and B there is also a drop at 45° indicating a possible further degree of symmetry.

The relationship between M-P classification and range profiles information content can be deduced from figure 15, where the cross-correlation between profiles collected from different perspectives is represented for the class *C* target.

The regions of high cross-correlation influence the M-P classifier: when 90° of separation occurs (dotted line in figure 15), the profiles between two perspectives, the input profiles taken in the range [120-150] degrees are highly correlated with the ones belonging to the orientations [210-240] degrees. M-P classification maxima and minima are mainly related to geometric symmetries.

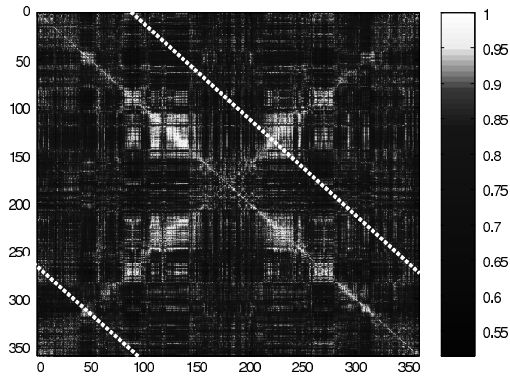

Figure 15. Non-coherent cross-correlation between profiles belonging to class *C*.

We now add to the dataset another sub-population class *D* and a new M-P FANN classifier. All the internal parameters of the Neural Network were changed by the new learning phase. As a result, the decision boundaries between the classes are modified. For this reason, as can be observed in figure 16, class *B* is more likely misclassified than class *C*, whose CRR remain almost unaltered. The four-class CRR shows the overall performance deterioration in terms of classification. On the other hand, the internal symmetries that influence the classifier remain unchanged by adding elements to the population. For example, it is thought that

the target belonging to class *C* has a number of multiple-bounce phenomena from corner-like scatterers. Their persistence is less than 15°, causing the number of relative maxima within less than 90° of separation.

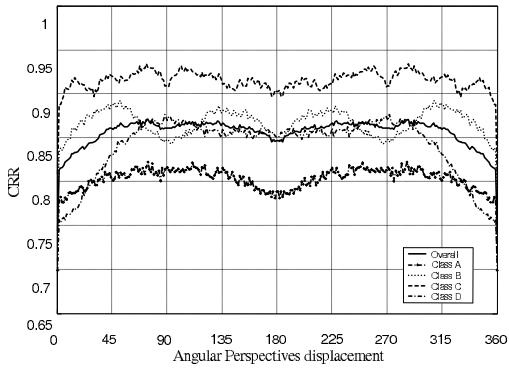

Figure 16. 2-P correct classification rates versus the angular displacement $\Delta\phi_{1,2}$ for the four-class problem.

Every target shows different trends: 90° angular perspective displacement can mean either an improvement or reduction in information content from a backscattering information point of view. This is due to the detailed make up of the target and the way in which differential symmetries can be exhibited. This can make the overall M-P classification rates less sensitive to the angular perspective displacement of the nodes.

The effects of symmetries on CRR performance can also be seen in the neighborhood of $\Delta\phi_{1,2} = 90°$ and $\Delta\phi_{1,2} = 270°$. This is verified when the relationship between two perspective displacements $\Delta\phi_{1,2}$ and $\Delta\phi'_{1,2}$ is:

$$\Delta\phi_{1,2} = \pi - \Delta\phi'_{1,2}. \tag{13}$$

When the perspective condition expressed in the above equation is verified, the 2-P performance is similar because of the intrinsic geometrical symmetries of classes *A*, *B*, and *C*. This is verified in both the three and four-class problems shown in figures 14 and 16, illustrating the M-P independence of the cross-affinity between different targets.

The de-correlation rate of the HRR profiles also seems to vary depending on the particular target. In general, it appears proportional to the number of wave trapping features and their persistency. If a single main scatterer has a high persistency over a wide range of orientations, the cross-correlation of the profiles over the same range of target orientations is expected to be high and hence we conclude there is less extra information to improve classification.

The more distinguishing features that exist, the greater the separation benefits achieved for single-aspect classification. This concept is amplified for the M-P

environment since those crucial features appearing for few orientations affect a greater number of inputs now represented by the perspective combination. For example, the length of the target is an important classification feature. For the M-P environment those targets presenting lengths not common to the other objects have even greater benefits. The mean CRR of the 2-P classifier is increased from 74.4% to 85.7% (+11.3%) with respect to the single-perspective case, while for the single class *D* the probability of correct classification increases by 13.1%. However, this result is highly averaged but is indicative of the overall performance improvement.

In figure 17 the classification rates for three perspectives are shown. Here the first perspective is fixed whilst the other two slide around the target covering all 360°. Thus these are shown as a function of the angular displacement $\Delta\phi_{1,2}$ between the second and the first node, and $\Delta\phi_{1,3}$ between the third and the first node.

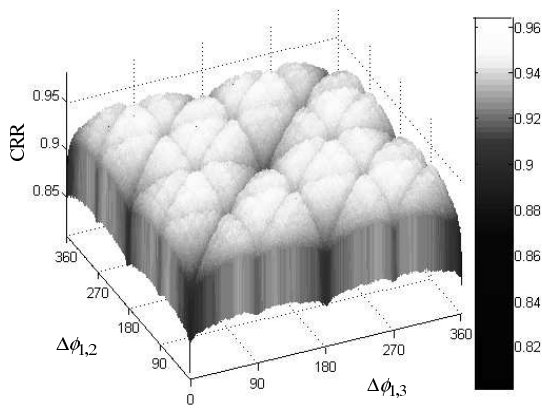

Figure 17. 3-P correct classification rates versus the angular displacements $\Delta\phi_{1,2}$ and $\Delta\phi_{1,3}$ for the three-class problem.

The origin represents the mono-perspective case and again indicates the overall improvement offered by the multi-perspective approach. The bisector line corresponds to two radar systems at the same location, while the lines parallel to the bisector symbolise two perspectives displaced by 90° and 180°. The inherent symmetry in the radar signature of the vehicle gives rise to the relatively regular structure shown in figure 17.

In a 2-Perspective scenario, if the two nodes view the target from the same perspective (i.e. the two radars have the same LOS) then this gives the mono-perspective classification performance. This is not the case for a 3-Perspective network. This is a consequence of weighting twice the perspective of two nodes and once the third node's perspective. As a result, the 3-Perspective performance

when $\Delta\phi_{1,2} = 0$ and $\Delta\phi_{1,3} \neq 0$ is worse than the 2-Perspective scenario that simply neglects one of the coinciding perspectives.

In figure 18, the CRR of the four-class population problem is represented with respect to the angular displacements between the three nodes of the network. As has been observed for the 2-Perspective case, the classification performance is less sensitive to the aspect angle when considering a greater number of classes problem. This may suggest that in a real environment with a large number of classes, the M-P classification could be equally effective for any target position provided that each node LOS is spaced enough from the others in order to collect uncorrelated signatures. The 3-P classifier probability of correct classification is increased from 74.4% to 90.0% (+15.6%) with respect to the single-perspective case.

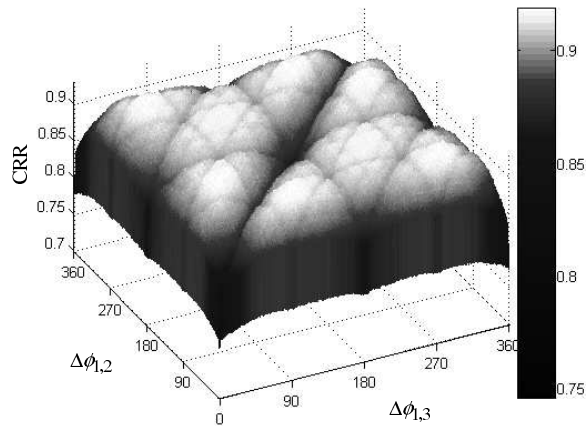

Figure 18. 3-P correct classification rates versus the angular displacements $\Delta\phi_{1,2}$ and $\Delta\phi_{1,3}$ for the four-class problem.

9. Summary and Conclusions

After implementing a multi-perspective classifier for ATR, the results in terms of classification rates have been examined using features extracted from HRR profile signatures. The benefits of the M-P classifier implementation have been analysed, showing a non linear but very clear CCR improvement with the number of perspectives. Furthermore, the correct classification gain from employing multi-perspectives is achievable for different SNRs and any radar node location. However, there is a small variation in the correct classification rate for a relatively constrained set of node locations. This is due to inherent geometrical symmetries of man-made objects. The M-P classification performance has been described for

two and three perspectives and applied to a three- and a four-class problem. The multiple perspectives affect the single class probability of correct classification differently depending mainly on the number, nature and persistency of scattering centers appearing in the profiles. As a consequence, the node location dependence of the global classification rate decreases when a greater number of classes is involved, making the classifier equally reliable for a wide range of angular displacements. The M-P classification improvements are reduced when the nodes are closely separated since the perspectives exhibit a significant degree of correlation. Nevertheless, the overall probability of correct classification is well above the mono-perspective case (+11.3% and 15.6% using two and three perspectives respectively). In addition the complexity and variability of reflectivity from real targets has been highlighted. Multiple perspective classification doesn't necessarily offer a trouble free route to acceptable classification and requires further testing under more realistic conditions. It also helps to indicate what information in the radar signature is important for classification. However, much further research remains before routine and reliable classification by radar becomes the norm.

Acknowledgements

The work reported herein was funded by the Electro-Magnetic Remote Sensing (EMRS) Defence Technology Centre, established by the UK Ministry of Defence and run by a consortium of Selex, Thales Defence, Roke Manor Research and Filtronic. The authors would also like to thank Thales Sensors for providing ADAS files to investigate HRR range profiles and ISAR images from real targets.

References

Baker, C. J. and Hume, A. L. (2003) Netted Radar Sensing, *IEEE Aerospace and Electronic Systems Magazine* **18**, 3–6.

Carrara, W. G., Goodman, R. S., and Majewski, R. M. (1995) *Spotlight Synthetic Aperture Radar: Signal Processing Algorithms*, Boston, Artech House.

Christodoulou, C. and Georgiopoulos, M. (2001) *Applications of Neural Networks in Electromagnetics*, Artech House.

Curlander, J. C. and McDonough, R. N. (1991) *Synthetic Aperture Radar: Systems and Processing*, New York, Wiley.

Duda, R. O., Hart, P. E., and Stork, D. G. (2001) *Pattern Classification*, John Wiley and Sons.

Hu, R. and Zhu, Z. (1997) Researches on radar target classification based on high resolution range profiles, In *Proceedings of the IEEE Aerospace and Electronics Conference*, Vol. 2, pp. 951–955.

Huaitie, X., Zhaowen, Z., Zhenpin, C., Songhua, H., and Biao, G. (1997) Aircraft target recognition using adaptive time-delay neural network, In *Proceedings of the Aerospace and Electronics Conference*, Vol. 2, pp. 764–768.

Keller, J. B. (1962) Geometrical theory of diffraction, *Journal of the Optical Society of America* pp. 116–130.

Knott, E. F., Shaeffer, J. F., and Tuley, M. (1985) *Radar Cross Section*, Artech House.

Looney, C. G. (1998) *Pattern Recognition Using Neural Networks*, Oxford University Press.

Novak, L. M., Halversen, S. D., Owirka, G., and Hiett, M. (1997) Effects of polarization and resolution on SAR ATR, *IEEE Transactions on Aerospace and Electronic Systems* **33**, 102–116.

Showman, G. A., Richards, M. A., and Sangston, K. J. (1998) Comparison of two algorithms for correcting zero-Doppler clutter in turntable ISAR imagery, In *Proceedings of the Conference on Signals, Systems & Computers*, Vol. 1, pp. 411–415.

Tait, P. (2006) *Introduction to Radar Target Recognition*, IEE, Peter Perigrinus Publishing.

Wehner, D. R. (1995) *High Resolution Radar*, Artech House.

Wilkinson, A. J., Lord, R. T., and Inggs, M. R. (1998) Stepped-frequency processing by reconstruction of target reflectivity spectrum, In *Communications and Signal Processing, '98*, pp. 101–104.

Zyweck, A. and Bogner, R. (1996) Radar target classification of commercial aircraft, *IEEE Transactions on Aerospace and Electronic Systems* **32**, 598–606.