Knowledge Based Diversity Processing

Christopher John Baker[1] and Hugh Duncan Griffiths[2]

Abstract In the past, radar sensing has tended to consist of relatively monolithic, single entity systems that present their output (often in the form of detections on a PPI display) as reports to an operator. The function of the operator is to interpret these reports and subsequently either provide information via a command chain for a decision on how to react, or to make such a decision themselves. In this way the human operator has been the source of intelligence in the sensing system, often aided and abetted by training and experience that has allowed a remarkably wide set of tasks to be performed with a high level of ability. However, the advent of electronic scanning coupled with advances in digital signal processing leads to a class of radar known as 'Multi-Function' and these are now challenging traditional methods by placing demands on the radar itself to make well informed, reliable decisions as to how a mission should be conducted. This is leading to the concept of intelligent or cognitive sensing. As a simple example an electronically scanned radar system is able to re-point its beam in timescales that are much faster than human reaction times. Where the beam should next be pointed therefore has to be a decision made by the radar itself. To understand and exploit its environment as fully as possible the system has the option of varying its parameters in a way that is tailored to the information it is seeking. This we term 'diversity processing'. A logical strategy is to do this in the light of prior experience, its own perception of the world and an appreciation of the task to be carried out. This we term 'knowledge based processing'. In this chapter we explore the early development of the concept of 'knowledge based diversity', drawing upon examples from both synthetic and natural echo locating systems to indicate how, eventually, true intelligence might be incorporated into future sensors.

Keywords: Multi-Function Radar, knowledge based processing, intelligent processing, diversity, resource management

---

[1]University College London
[2]Defence College of Management and Technology, Shrivenham, Cranfield University

J. Byrnes (ed.), *Unexploded Ordnance Detection and Mitigation,*

1 Introduction

We begin by examining what is meant by terms such as 'intelligence' and 'knowledge' to set the scene for the challenges faced by knowledge based diversity processing systems. Artificial intelligence has been a topic of intense research for many years and so called artificial intelligent methods of processing radar data, such as neural networks, have been exploited for some time. However, these have been implemented in a relatively simplistic way not really representative of true intelligence. In fact intelligence might be better used to describe a property of the mind (or in our case the radar signal processor) that encompasses many related abilities, such as the capacities to reason, to plan, to solve problems, to think abstractly, to comprehend ideas, to use language, and to learn. There are several ways to define intelligence. In some cases, intelligence may include traits such as creativity, personality, character, knowledge, or wisdom. These are characteristics that go substantially beyond any existing radar sensor system, as they imply a dynamic interaction with the world in which they exist whilst attempting to achieve a chosen goal or objective. Similarly, knowledge is defined in the Oxford English Dictionary variously as (i) *expertise, and skills acquired by a person through experience or education; the theoretical or practical understanding of a subject*, (ii) *what is known in a particular field or in total; facts and information or* (iii) *awareness or familiarity gained by experience of a fact or situation*. Philosophical debates in general start with Plato's formulation of knowledge as 'justified true belief'. There is, however, no single agreed definition of knowledge presently, nor any prospect of one, and there remain numerous competing theories. Knowledge acquisition involves complex cognitive processes: perception, learning, communication, association and reasoning. The term 'knowledge' is also used to mean the confident understanding of a subject with the ability to use it for a specific purpose if appropriate.

Although imprecise, these descriptions convey a sense of known information that can be exploited to create a desired effect in the context of a perceived understanding of the local and globally significant operating environments. Together they will form the basis of future truly intelligent sensing systems and are providing the impetus for much current research as demonstrated by the US 'KASSPER' [1] and 'Sensors as Robots' [2] programmes. In this chapter we explore the underlying concepts that provide the keys to greater and greater levels of cognition in sensor systems. In particular we draw upon the lessons that are being learnt from nature and especially from echo-locating mammals which imply that a much more holistic approach to systems operation may be necessary. We begin by reviewing the nature of information that can be inferred from radar echoes, and then go on to explore how this is enhanced by exploiting the full parametric diversity available to maximize the quality of sensed information upon which radar systems can make robust and reliable decisions, and subsequently how this might be further enhanced by exploiting prior knowledge.

2 Acquiring Information from Radar

It is intuitive that the information acquired from a radar sensor is dependent on the parameters used to design the system. The more we are able to adaptively adjust these parameters, the more it is possible to maximize the quality and relevance of the information acquired and hence to maximize the likelihood of ‘mission’ success. Changes to the radar of this type we term diversity and the range of diversity possible is embraced by the following:

– ***Bandwidth and or frequency*****:** this might be wide bandwidths for high spatial resolution or narrow bandwidth, long duration signals for high Doppler resolution.
– ***Orientation*****:** this provides for spatial diversity. A simple example is SAR or ISAR which use angular diversity to reduce speckle in multi-looked SAR or ISAR imagery.
– ***Waveform design*****:** where parameters such as pulse length, modulation and PRF can be adjusted dynamically.
– ***Signal strength*****:** it is clear that a low echo strength of received signals will be corrupted by the effects of noise, and hence the signal-to-noise ratio must be sufficiently high to avoid this.
– ***Time*****:** the time evolution of behavior of an object can give important clues as to its nature and intent.
– ***Phase centre*****:** the use of multiple and adjustable phase centres can provide important location information.
– ***Polarization*****:** man-made targets are often made up of polarization sensitive structures that allow them to be more easily recognized against a background comprising natural scatterers.
– ***Knowledge diversity*****:** this can take almost any form of priors, memory mapping and their associations.

In the following sections we examine aspects of these forms of diversity, showing how some are already used in radar sensors but how performance could be improved if the concepts of true intelligence could be appropriately adopted.

3 Information in Current Radar Systems

We begin this section by examining how resolution in the radial-range and cross-range dimensions provides useful improvements in information content by creating localized zones of scattering. Frequency or bandwidth span is conventionally used to provide spatial resolution to isolate targets from one another. Wide-bandwidths can result in improvements in down range resolution which further increase information content by isolating scattering into separate discrete areas. This is an example of using coherence between the transmitted and received signals. Synthetic aperture techniques such as SAR and ISAR, e.g. [3], combine angle and bandwidth data coherently to provide high resolution in two dimensions. We also note that frequency diversity employing a multiplicity of illuminating frequencies of non-overlapping bandwidths can be used non-coherently to reduce the effects of speckle as in multi-looked SAR or MIMO radar [4]. In this case the multi-looking is attempting to provide an improved estimate of the underlying radar cross-section and hence improved information. Here we will examine in more detail how coherent use of frequency diversity provides better spatial information.

Radar is a relatively simple sensor able to provide information about the position of an object in 3-D space as a function of time. In essence this is facilitated by the measurement of radial range from the radar to a target and the rate of radial change of position of the target, both as a function of azimuth and elevation angles as determined by the azimuth and elevation beamwidths. In this way targets can be detected and tracked over time. By transmitting a short or modulated pulse a radar is able to resolve between multiple targets with a resolution given by:

$$\triangle r = \frac{c}{2B} \tag{1}$$

where

$\triangle r =$ range resolution (m)
$c =$ velocity of light (ms^{-1})
$B =$ bandwidth of radar signal (Hz)

Thus wideband signals are necessary to achieve high spatial resolution. If $\triangle r$ is a fraction of the target dimension as presented radially to the radar then it is possible to begin to measure important target characteristics such as length. Indeed, it is potentially possible to use difference in echo strengths from different parts of the target to uniquely discriminate the target from other possible candidates. Such an echo is termed a High Range Resolution Profile (HRRP) and Figure 1 shows schematically the type of response that might be observed.

However, there are limits on how wide a bandwidth it is realistically possible to transmit and receive within a single pulse. This limit can be overcome by

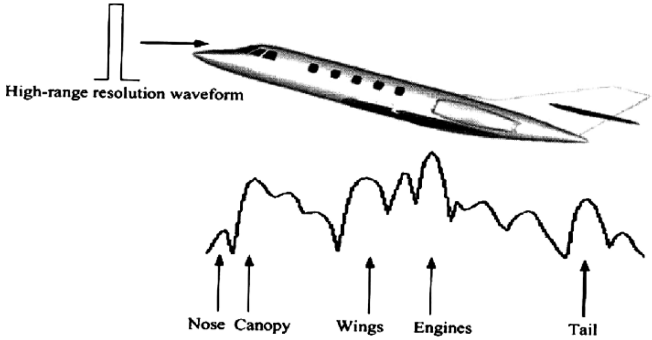

Fig. 1 The HRRP generated by a wide band pulse for an aircraft target.

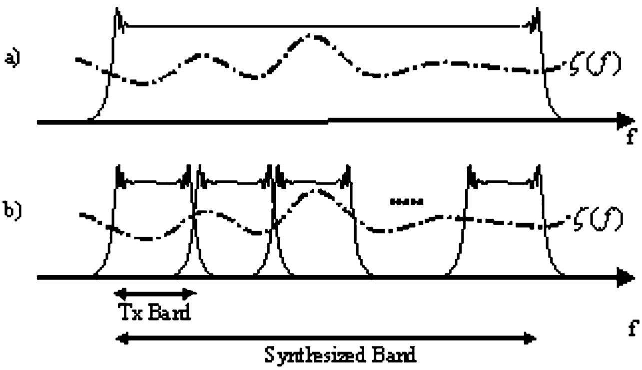

Fig. 2 Spectrum reconstruction of the target reflectivity function by one single chirp pulse **(a)** and by stepped-frequency coherent addition of sub-spectra **(b)**.

transmitting and receiving a series of pulses where the centre frequency of each is stepped incrementally such that the total bandwidth spanned is much greater that of any single pulse. Each echo is digitised and recorded so that the full bandwidth span can be re-constructed and significantly higher resolution achieved. Figure 2 shows this schematically for a series of pulses.

In this example the separation of pulses in the frequency domain is equal to their bandwidth. More typically they will be overlapped by as much as 50%. This reduces the achieved bandwidth span and degrades resolution but avoids potentially awkward bandwidth reconstruction at the band edges. Figure 3 shows an example of range profiles of a Land Rover vehicle plotted as a series of intensity modulations covering a total azimuth extent of 360. Discrete areas of high echo strength are clearly visible as is their angular span. The scattering characteristics appear quite varied, nevertheless this represents the base information for classification. However, this method only provides high resolution in the radial direction and hence any scatters lying at the same range will be measured as a single echo response. This means that small changes in the viewing geometry will cause large and rapid fluctuations in the range profile as scatters will interfere with one another constructively and destructively. One method to militate against this is to provide resolution in the cross range dimension using synthetic aperture techniques.

The real beamwidth of any radar antenna is determined approximately by the ratio of the wavelength to the physical size of the antenna. In simple form this is expressed as:

$$\triangle\Phi = \frac{\lambda}{D} \tag{2}$$

where
$\triangle\Phi$ = angular beamwidth (rads)
λ = wavelength (m)
D = diameter of antenna

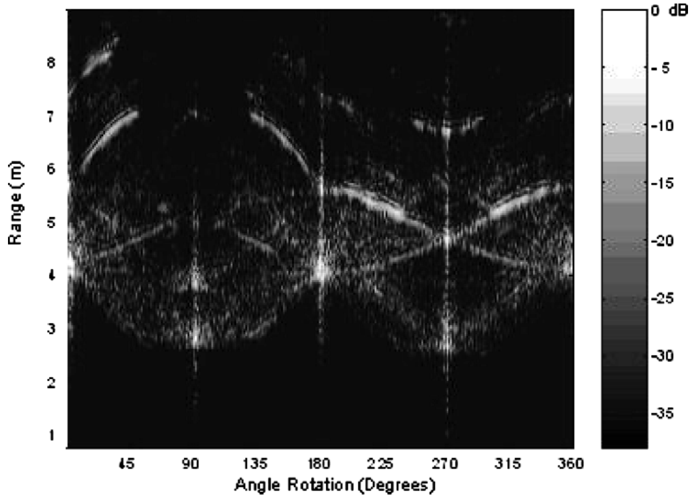

Fig. 3 History of HRR range profiles (8 cm of range resolution) from a series of X-Band stepped frequency chirps illuminating a ground vehicle as it rotates over 360 degrees. At zero degrees, the target (a Land Rover) is broadside oriented, while at 90 degrees has its end-view towards the radar.

If Equation (2) is multiplied by the range we get the width of the beam at that range. Thus at a range of 10 km the beam size of a 1 m dish antenna operating at a wavelength of 3 cm is 300 m which is much larger than, for example, vehicles and most aircraft. To overcome this limitation large apertures can be synthesised (i.e. effectively increasing D in Equation (2)) using the Synthetic Aperture Radar (SAR) or Inverse Synthetic Aperture Radar (ISAR) techniques. These techniques create large apertures (albeit on signal reception only) by collecting data over as a function of viewing angle, thus effectively mapping out a much larger aperture than that of the physical antenna. Figure 4 shows an example of ISAR imagery generated from the HRRPs displayed in Figure 3.

The image in Figure 4 has the expected rectangular outline typical of a Land Rover vehicle and the bonnet, cab and truck areas are clearly discernable. The scatterers observable as a function of angle in Figure 3 are now 'focussed' into discrete zones showing that in fact they emanate from the same physical part of the truck. In forming a 2-D image in this way we have now exploited both range resolution (frequency diversity) and angle resolution (angular diversity) and hence this is an example of exploiting diversity (i.e. combining range and angle) to improve the information obtained by the radar to make better decisions. Indeed, it seems much less of a leap of faith to make the assumption that the data shown in Figure 4 is in a form to support human based classification. However, automating this process has proved to be an extremely challenging problem in all but a few restrictive cases. The reasons for this are chiefly to do with the fact that often 'resolution cells' still contain multiple scatterers and hence they may again, scintillate rapidly with small changes in viewing angle. Additionally, there is often a multipath component which adds to

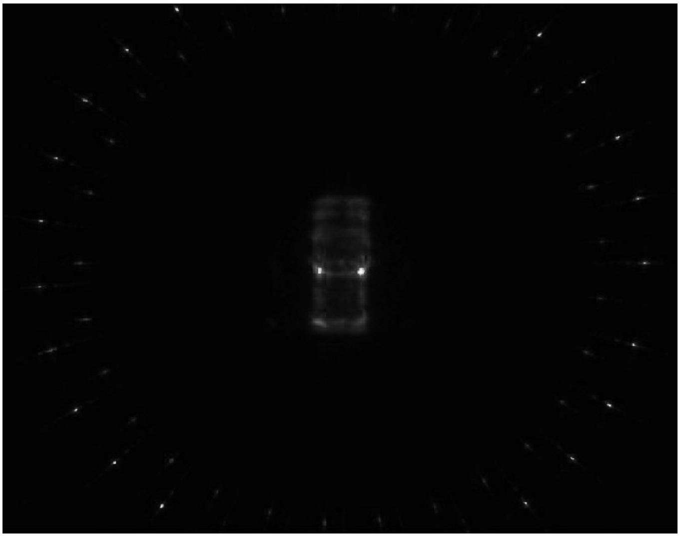

Fig. 4 Multi look ISAR image of the Land Rover target.

this scintillation. It is also possible for some resolution cells to contain part target and part clutter which further exacerbates these effects. In addition not all objects are comprised of discrete scatterers alone and other forms of re-radiation can take place. Overall the form of backscatter from extended targets is extremely complex and subject to great variability hence making consistent and reliable interpretation by automatic means a substantial challenge. However, tasks such as navigation, collision avoidance and object classification are executed with seeming ease by echo locating mammals. In the next section we examine this in more detail with a view to learning lessons that can be valuably employed in synthetic sensors.

Firstly, we examine one more component of diversity as provided by polarization. It is well known that polarimetric information may be useful for classification since it completes the information which can be obtained from the target returns. Radar targets have different responses to different polarization signals. By emitting a mixture of polarizations and using the receiving antenna with a specific polarization, several different signals can be collected and used for recognition. For this purpose it is necessary to illuminate the target with two signals having different polarization vectors, for instance when linear polarization is used the polarizations are vertical and horizontal. The polarization properties of the target can be completely represented by its scattering matrix [5]. The scattering matrix is defined as:

$$\begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix} \tag{3}$$

where S represents the state of polarisation. The components of the matrix are called scattering components and are derived from the following relation:

Table 1 Scattering matrix for some common geometrical shapes.

Shape	Linear polarization	Circular polarization	Aspect
Sphere 	$S = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$	$S = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$	Every aspect
Flat Plate 	$S = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$	$S = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$	Radar signal perpendicular to the plate
Dihedral comer reflector 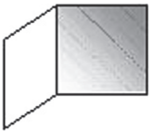	$S = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$	$S = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$	Radar signal perpendicular to the comer's axis
Trihedral comer reflector 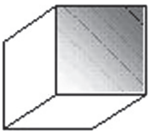	$S = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$	$S = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$	Radar signal perpendicular to the comer's axis

$$\begin{cases} E_H^R = S_{11}E_H^T + S_{12}E_V^T \\ E_V^R = S_{21}E_H^T + S_{22}E_V^T \end{cases} \tag{4}$$

where E is the electric field, the subscript means horizontal polarization, H, or vertical polarization, V, and the superscript means reflected, R, or transmitted, T. S_{11} and S_{22} are the co-polar coefficients and S_{12} and S_{21} are the cross-polar coefficients. Simple geometrical shapes present well-know polarization matrices; some examples are show in Table 1.

In HRRPs the scattering matrix obtained in each range bin is the result of the combination of the matrices of all the scattering points in the same range bin. Therefore, the shape extracted from this matrix is not necessarily reliable for recognition. In ISAR images, instead, if the resolution is high it is possible to begin to isolate the scattering matrix of the main scattering point and consequently to estimate the geometrical shapes of these points. The estimation of the shapes of scattering points can be subsequently used as features for target recognition [6].

The polarimetric information content for both HRRPs and Images are examined here qualitatively. Figure 5 shows the linear polarimetric range profiles as a function of rotation angle again for the Land Rover target. The information is largely the

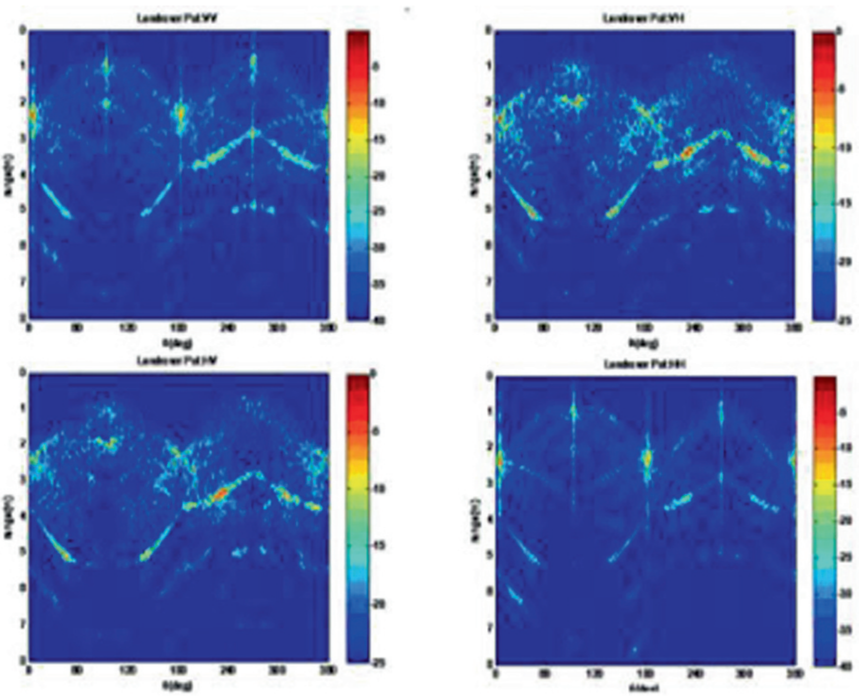

Fig. 5 Polarimetric range profiles of the Land Rover target.

same for features such as length and width but there are discernable differences in detail with some scatters appearing in one polarisation and not another as well as differences in echo strength and their angular persistence. Figure 6 shows the same base data but this time processed into 2-D imagery. Again the main characteristics are similar and there are small differences in detail but, perhaps, even less obvious. The overall advantage of polarimetric data in terms of improved classification performance is still to be proven.

4 Diversity in Biologically Sensing System

Reliable and robust navigation, collision avoidance and object classification are carried out with great success by echo locating mammals such as bats that are able to detect, select and attack prey even in a dense clutter and hostile countermeasures environment. Although the frequencies and waveform parameters used by synthetic sensors and by echo locating mammals are not the same there remain close parallels that suggest lessons can be usefully learnt. Here we examine methods that are exploited by mammals such as bats and dolphins and to see how they may be applied, particularly concentrating on how diversity is utilized.

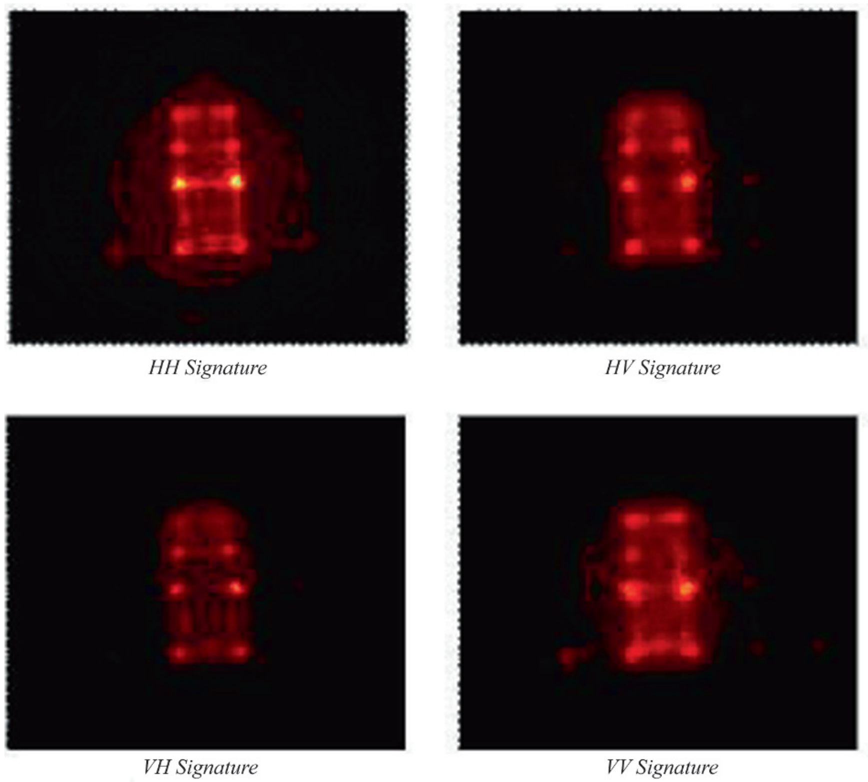

Fig. 6 Polarimetric images of the Land Rover target.

Bats have evolved echo location as a means of detecting, selecting and attacking prey over a period of more than 50 million years into a highly sophisticated capability on which they depend for their survival. It therefore seems self evident that there is potentially a great deal that can be learnt from understanding how they use this capability and applying such an understanding to synthetic sensing systems such as radar and sonar. Bats are able to modify their PRF, transmitted power, waveform type, and bandwidth. They also use multiple perspectives as part of their hunt strategies. These are all examples of exploiting diversity and could be implemented in modern radar systems if we are able to understand how to do so appropriately.

Bats use two very different types of waveform for classification. The first has the form of an un-modulated pulse that contains a number of harmonics. This is well suited to the classification of moving targets exhibiting a micro-Doppler component and there is evidence from biological studies that supports this hypothesis. The second waveform consists of a near hyperbolic chirp, also with two or three harmonics typically present. As the bats approaches the target the pulse length decreases and the degree of hyperbolic curvature increases. In this way a greater bandwidth is generated leading to finer and finer resolution as the bat gets closer and closer to the target. Again there is evidence from biological studies that this type of waveform is used when engaging static targets such as flowers. There also seems to be a strong relationship between the orientation of the bat, the position of the target, the precise nature of the transmitted waveform and the clutter and reverberation environment.

Here, we examine both forms of real bats calls recorded during the classification phase when attempting to identify potential prey supported by full scale measurements of real targets. A data base of real bat calls has been used to extract digitised versions that seem best suited to classification tasks. These include calls of horseshoe bats (Rhinolophidae) that use long constant frequency components that encode acoustic glints in echoes from fluttering targets and the hyperbolic modulation of an Eptesicus Nilssoni used when extracting nectar from a flower head.

Figure 7 shows the spectrogram of the calls of a Pearson's horseshoe bat Rhinolophus pearsonii when attempting target classification. The waveform is characterized by a relatively large constant frequency component with a short wider bandwidth modulation at either end. There are a series of harmonic replicas spanning quite a considerable overall bandwidth. The method of classification being used is based on exploiting the micro-Doppler signature which is best achieved with an extended constant frequency component. The short wider band part of the sign is probably used for ranging. The response to a simple model for the beating wing of an insect is shown in Figure 8. The Doppler signal generated is a function of the harmonic frequency and shows different degrees of sensitivity to the different constant frequency transmissions. Thus the bat is able to select the best signal for classification as well as using multiple copies for further confirmation.

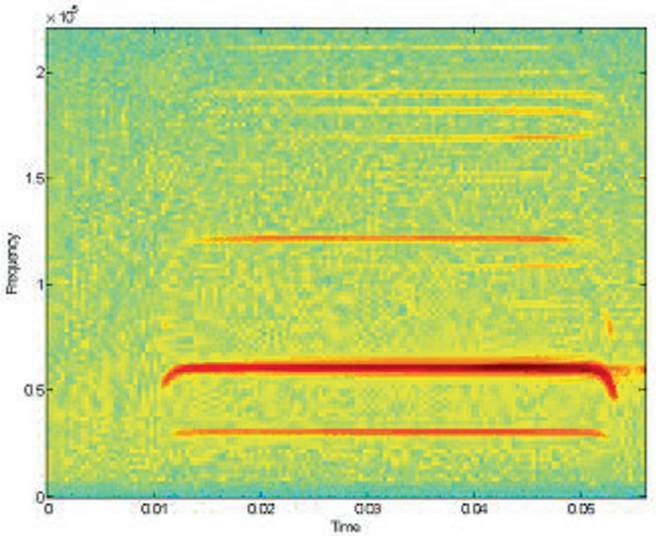

Fig. 7 Spectrogram of the call made by a bat when attempting target classification.

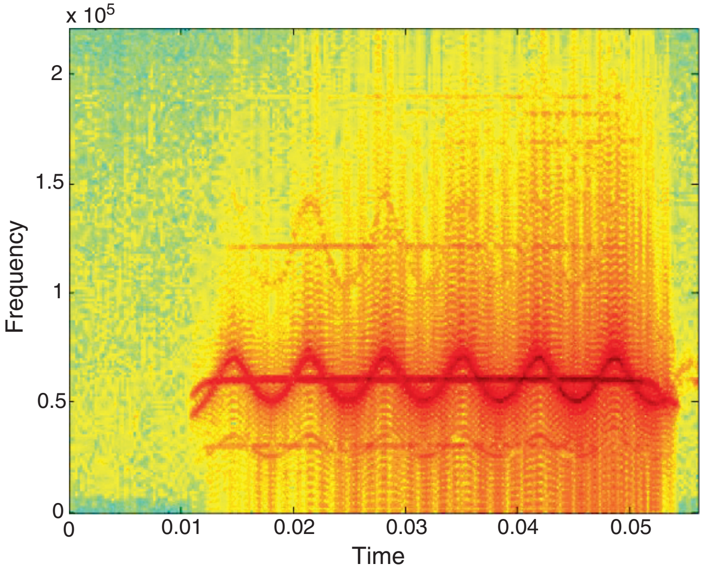

Fig. 8 Change in the received signal when modified by a simulation of a wing beat at an orientation broadside to the target.

An additional possibility is that the bat also uses this harmonic type of waveform to generate high resolution in range by effectively mimicking a frequency stepping system (albeit in a single pulse). It does, potentially, have very attractive properties that make it worthy of further investigation for a variety of applications. For example, Passive Bistatic Radar often uses illuminators of opportunity that emit simultaneously on several frequencies that collectively span a much wider bandwidth than that of any single frequency. By utilizing the total bandwidth transmitted a relatively high resolution mode could be developed. However, the signals will be under-sampled and there could be potentially significant sidelobes. This is another example of the exploitation of diversity.

A further factor that is generally very important in the attack of prey is the trajectory that is used. For example, when approaching a target located in a background of leaf clutter the bats uses a low line of attack, presumably, to minimise backscatter (i.e. clutter) from the leaves. Figure 9 shows the resulting interaction between the illuminating waveform and the simple wing beat model when the relative angle between the two is altered to 45 degrees. Here a reduction in the amplitude of the modulation can be observed. The bat is able to deliberately alter its orientation and transmission frequency with respect to the prey such that it maximises its sensitivity to the micro-Doppler signature and hence has the best information set for classification. Also, this provides important information as to the orientation of the prey and that micro-Doppler aspect angle dependent information may be gleaned that is used in the classification process. Clearly, there is a degree of speculation in these hypotheses and further research is required to develop a more complete understanding.

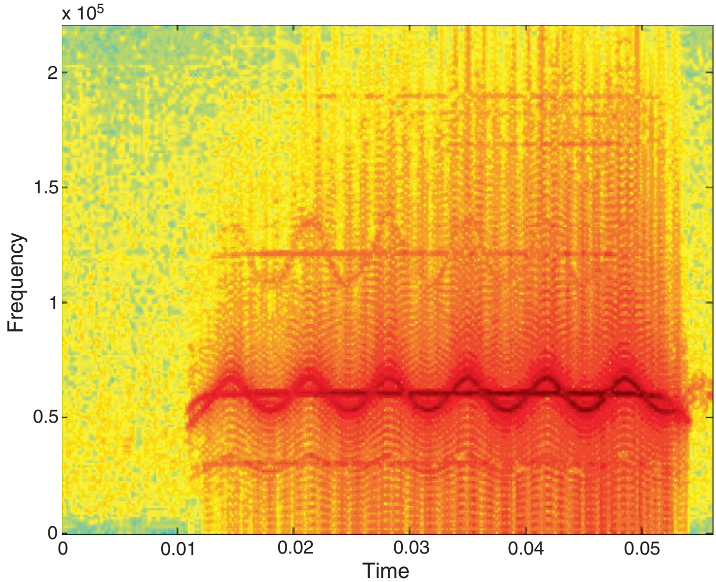

Fig. 9 Change in the received signal when modified by a simulation of a wing beat at an orientation angle of 45 degrees to that of Figure 6.

Figure 10 shows a sequence of waveforms that an Eptesicus Nilssoni bat used when attacking a stationary target. The bat was constantly changing its orientation such that it views the target over an angle range of approximately 270 degrees and is gradually getting closer and closer to the target. In other words the bat is benefitting by using a combination of angle and waveform diversity.

For each pulse the waveform approximates well to a hyperbolic function. As successive pulse are emitted and received and the bat closes in on the target the pulse length is reduced to conserve energy and avoid ambiguity. The degree of hyperbolic curvature increases to improve resolution and tolerance to any differential Doppler. Therefore, just prior to attack the bat operates with the highest resolution (of the order of 1 mm) and has gained multi perspectives for classification. Indeed, the use of the two receivers (ears) combined with very high resolution is suggestive of the selection of a particular part of the target. Figure 11 shows one of the final pulses and its ambiguity diagram in more detail. The high degree of Doppler curvature and subsequent high range resolution and Doppler tolerance may be observed. The waveform and its dynamic adjustment as a function of the previously received echoes are quite typical of nectar feeding bats. Thus we see that the waveform is being altered in a number of ways as the viewing geometry is also changed. This suggests diversity in a wide sense may be a key ingredient along the path to employment of truly intelligent synthetic sensing systems. For synthetic sensor systems the equivalent dynamical adjustment of sensor parameters is well within the scope of current technology. The more demanding question is ‘what are the appropriate adjustments that need to be made to maximise the success of a mission? This requires considerable further research.

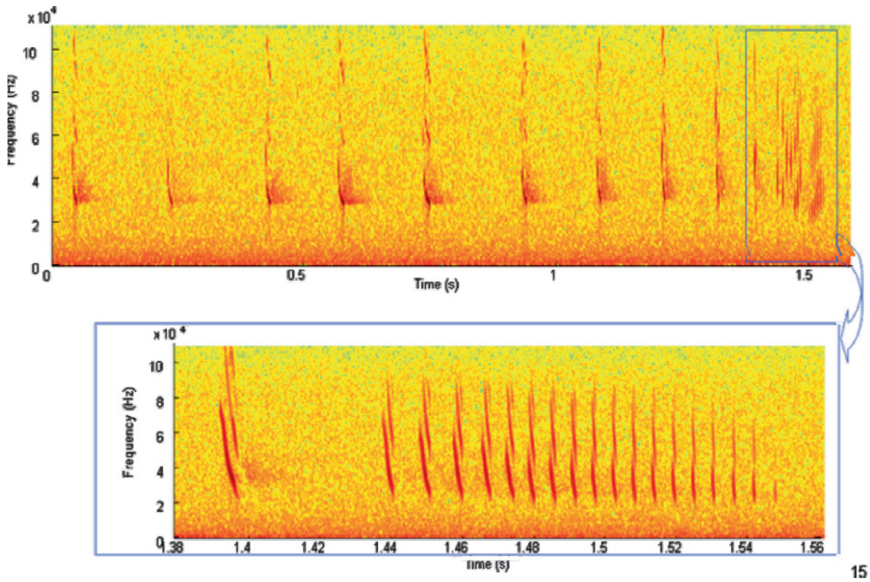

Fig. 10 Sequence of waveforms that an Eptesicus Nilssoni bat used when attacking a stationary target.

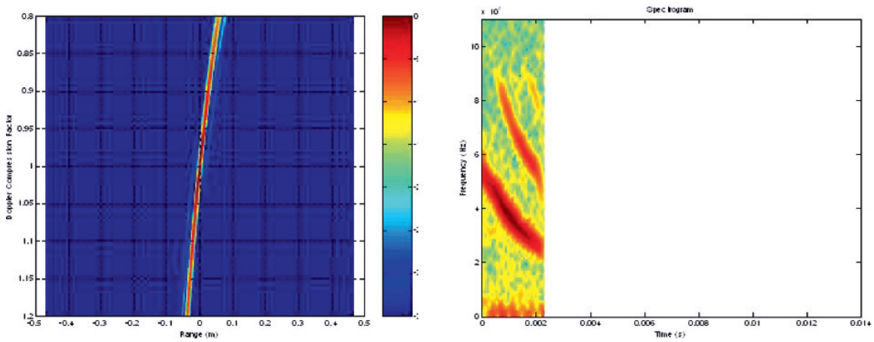

Fig. 11 Example ambiguity diagram and waveform for an Eptesicus Nilssoni used in the attack phase.

We now examine how the bat is able to recognize nectar providing plants against a complex background as classification is a key component in intelligent sensing. To investigate static target classification in more detail, high resolution range profile data was acquired from experiments carried out at Bristol University [7]. The experiments generate a wideband waveform that provides a resolution of approximately 1mm. This is used to illuminate a number of flower head targets mounted on a narrow pin and recording the received echo as a function of orientation angle in the ground plane. Figure 12 shows this experimental layout in schematic form.

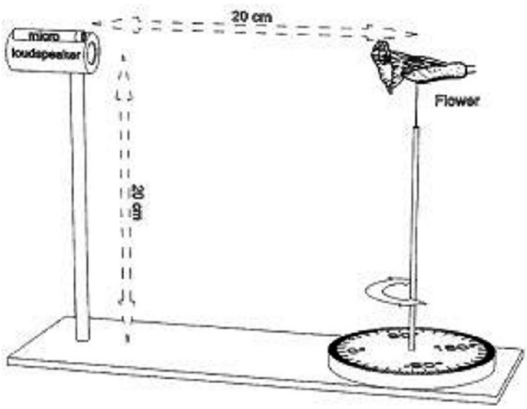

Fig. 12 Experimental geometry used to gather high range resolution profiles of flower heads as a function of orientation angle (Courtesy of G Jones and M Holderoid, University of Bristol.)

This particular configuration has been chosen as bats and plants have co-evolved to enable classification to take place reliably. In other words it is in the interests of the bat to be able to recognise the flower as it is a source of food and for the flower to be recognised by the bat as the bat is a means of pollination and thus the ingredients for successful classification are in-built.

One strategy used by plants is to expose their flowers e.g. by suspending them on modified long leafless braches into the sub-canopy. This helps bats to approach the flower and the uncommon presentation stands out against other vegetation. Indeed it might be that the local scattering environment offers many clues (prior knowledge) that there are likely to be flowers present and is this an important classification aid. However, many flowers are also presented closer to the plant or grow on stems and branches (cauliflory). In such cases several echo acoustic cues might render the flower unique such as:

– Echo strength can be higher than in leaves of the same size, because the bell-shaped corolla of many chiropterophilous flowers collects and focuses sound back towards the bat.
– Floral echoes last longer than echoes of leaves and branches, because sound is reflected within the corolla.
– Floral echo fields are often omni-directional, which means that most sound energy is always projected back into the direction of sound incidence. This contrasts with other plant structures, such as leaves, which produce echoes of maximum amplitude only when ensonified perpendicularly.
– Because flowers are complex targets and consist of many different reflectors at different distances, interference generates specific peaks and notches in the echo spectra, giving them a ‘coloured’ spectral appearance. This not usually occurs in convex structures such as branches or leaves.

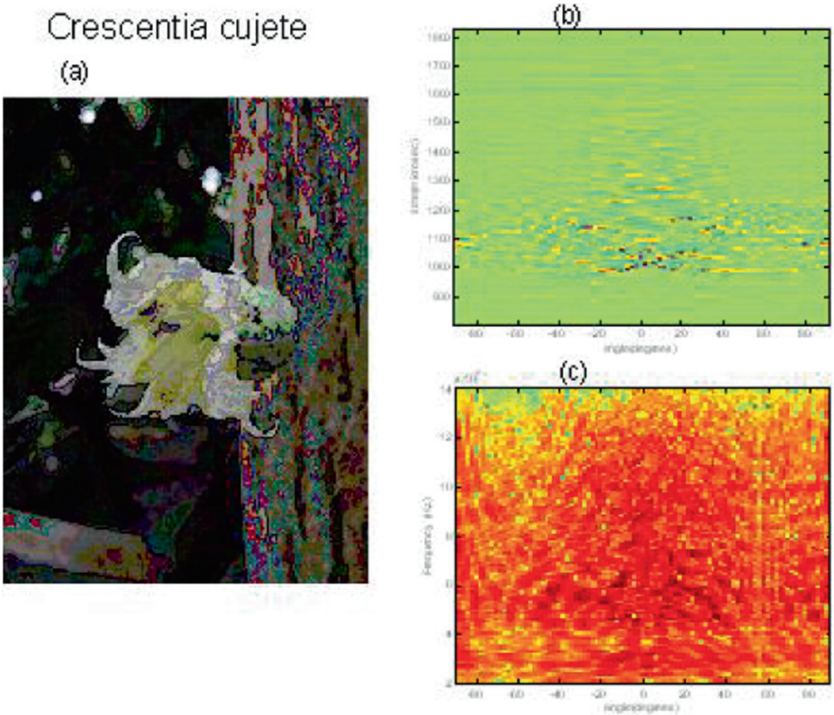

Fig. 13 Example flower head (Crescentia cujete) and its corresponding high range resolution profile together with a spectral representation (both as a function of orientation angle.

Figure 13 shows an example flower head (Crescentia cujete) and its corresponding high range resolution profile together with a spectral representation (both as a function of orientation angle).

The range profiles show clear structure that is present at all angles and exhibits both similarities and differences to the high resolution radar profiles of the Land Rover target. One difference that seems to be apparent from a visual inspection of the profiles is that more features appear to persist over larger angular ambits. In the radar profiles there appears much more variability and indeed, it is the angular dependence that is responsible for improved classification performance as additional perspectives are included. Figure 14 shows the same form of plot as 13 but for a Viresa gladiflora. Again, similar behaviour may be observed although the detail of the form and structure is significantly different.

In order to examine classification using such data in more detail a multi perspective classifier was applied. Three flower targets are used to make the case quite demanding as in reality it is unlikely that the bat has to differentiate between two flower heads in order to feed. Figure 15 shows the results of applying a Forward feed neural network multi-perspective classifier. The plot also includes increasing levels of noise.

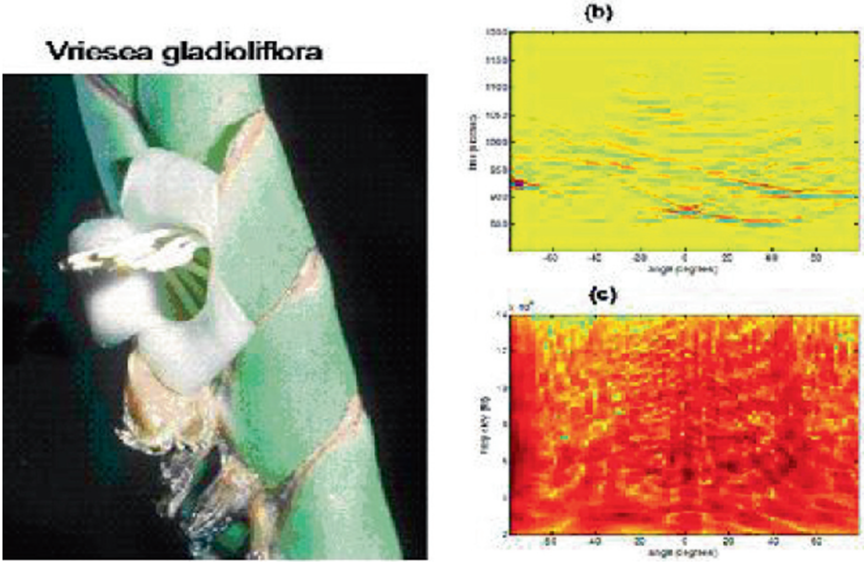

Fig. 14 Example flower head (Vires gladiflora) and its corresponding high range resolution profile together with a spectral representation (both as a function of orientation angle.

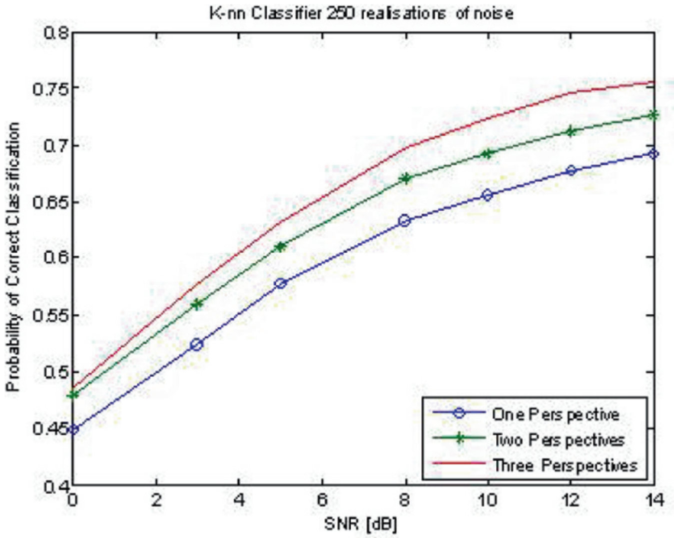

Fig. 15 Multi perspective classification performance of the three flowers versus signal to noise ratio.

There are two main conclusions that may be drawn. The first is that there is a significant increase in classification performance in going from one to two to three perspectives i.e. exploiting angle diversity. Secondly, as the signal to noise ratio increases the classification performance, as might be expected, also increases. Indeed without noise the classification performance, even with a single perspective is close to 100%. As noise is added eventually there is a more rapid drop off in performance, indicative of the loss of key information, probably embedded in scatterers of smaller echo value? This therefore highlights the importance of having a sufficiently high signal to noise ratio when echo locating.

In this section it was seen how the bats use a combination of waveform design and orientation to carry out their mission. In other words they are exhibiting all the requirements of a truly intelligent system (and with a processor the size of a 'pea') and are utilizing diversity in doing so. The bat gathers information about its local environment and then adjusts the parameters of the next call as well as changing its position to ensure that it extracts the right or best information to continue and complete its task as successfully as possible. Indeed, this simple examination of how diversity is exploited by bats highlights that this natural sensing system is continually adapting its entire diversity range, seemingly to ever improve the information the bat needs to fulfill its mission. In the next section we examine how angle diversity via multiple viewing angles can be used to improve classification performance in much the same way it might be used by bats.

5 Exploiting Diversity in Synthetic Sensors

In this section we look at the role of trajectory in the recognition process as invoked in a synthetic radar sensor [8]. The same form of radar data as displayed in Figure 3 is utilized as an input to a classifier. As in the previous section we compare classification performance as a function of the number of perspectives used. In this example three classifiers are employed to ensure that the results are not biased by the classifier itself. The three classifiers are (i) A Bayesian classifier, (ii) a neural network and (iii) a K nearest neighbor classifier. Range profiles at approximately every eighteen degrees are used to train the classifiers and then removed from the data set to be classified. The results are displayed in Figure 16.

Figure 16 shows that the classification performance increases with increasing number of perspectives no matter which type of classifier is used consistent with that seen for the case of the bats data. This is also consistent with our own experiences. Generally, if we wish to positively identify something that we expect to recognize we will move our viewing position either to get a 'better look' or to re-enforce an expectation of the objects identity. We might further conjecture that there will be preferential look directions to exploit determined by the first look direction. For example, objects that exhibit symmetry would not yield large additional amounts of information if the viewing angle were altered by 180 degrees. This is further illustrated in Figure 17 which shows the multiperspective classification performance for a Land Rover target as a function of viewing angle for a three perspective classifier.

Figure 17 shows that at certain angle combinations the classification performance falls. The performance at position 0,0 is the monostatic case and is the worst of all.

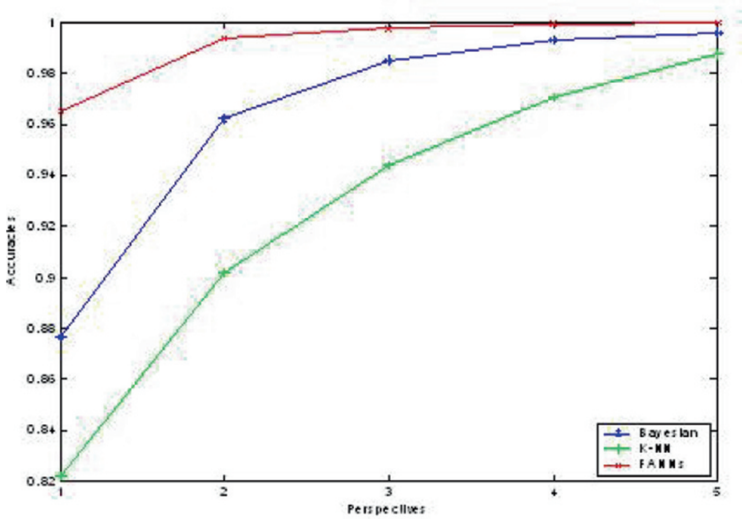

Fig. 16 Multi perspective classification performance of four vehicle targets versus number of perspectives using high resolution range profiles as input data.

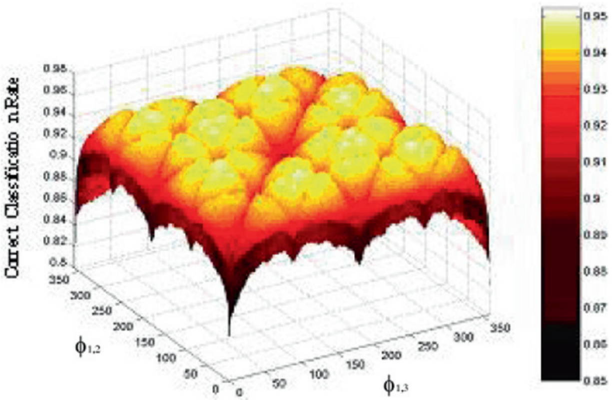

Fig. 17 3-P correct classification rates versus the angular displacements and for the three-class problem.

The diagonal represent the two perspective case and again performance is inferior to the peak. The other troughs tend to occur at separation angles such as 180 degrees, where the symmetry of the target means that there is little additional information to be gleaned that can benefit the classifier. If the orientation of the target can be determined then this prior information can be used to arrange a second look to be one that compliments the first adding maximal new information thus optimising overall classification performance. This would be an example of a simple but genuinely intelligent radar sensor i.e. one that uses previous information to improve subsequent information and hence maximizes its chances of carrying out the desired task successfully. In the next section we examine other prior information methods to improve mission performance.

6 Prior Knowledge

As we have begun to explore in the previous sections, knowledge based radar systems might be thought of as the precursor to future systems that will employ artificial intelligence [9]. Indeed, it is logical that information already known from other sources about the target area or task to be carried out can be used to direct the operation of the radar or to interpret the radar's findings more usefully. As we saw in the section above the possibility also exists for data gathered in real time (or 'on the fly') to be used for generation of relevant information or for better cuing of further information. Another example could include an MTI radar that is also able to operate in a SAR mode which can be used to provide coarse image information as an aid to more accurate detection of moving targets. Historical data of the same type (i.e. MTI data) for a given scene should also theoretically provide a useful 'memory' if a way can be found to exploit it when performing new MTI detection in real time.

The electronically scanned radar system is another example of a sensor that brings the need for intelligent radar operation into sharp focus. Instantaneous, adaptive beam pointing enables combinations of functions such as tracking, surveillance, and weapons guidance (previously performed by single dedicated radar systems) to be implemented simultaneously. The decision as to where and when to re-point the radar beam can and sometimes has to be taken in the interpulse period, far faster than the rate at which a human operator could intervene. Thus the radar itself must make these decision based upon a combination of what it has already observed, prior knowledge and the mission objectives which is embracing the concept of intelligence. One area where this might occur is that of sector prioritization. Military understanding is used as the prior knowledge to allocate the sectors and determine high level metrics for assigning priority. We consider such an example that also utilizes fuzzy logic to give both a more humanistic decision making logic and to preserve where possible radar resources for further allocation of tasks.

The attribution of priority to regions and targets of interest may be done in a variety of ways [10]. For example the decision tree structure shown in Figure 18 could be used. The information required to take the decision is supplied by radar data operating in tracking and high resolution modes. The different variables provide differing types of information used to set priority. These are threat, hostility, weapons system capabilities, track quality and the target position. The selection of these variables has been by expert judgment based upon operational experience. They may require further refinement with further experience. The logical relationships between these variables then determine the setting of priorities. Fuzzy logic can be used to provide a softer way of making decision by allowing variables to take

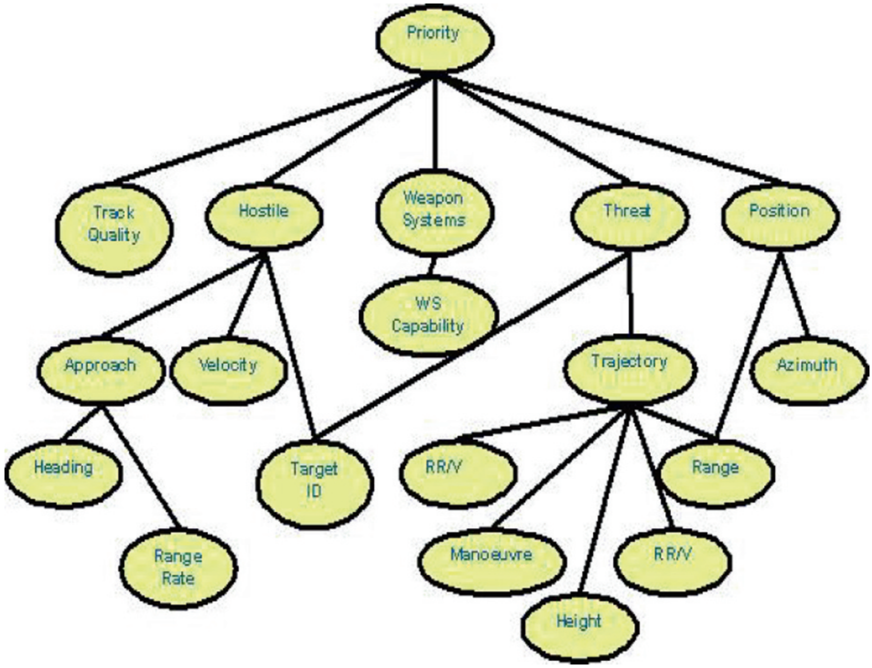

Fig. 18 Decision tree for sectors for target priority assessment.

values in between an either 'on' or 'off' state. The nature of the inferential rules linking the fuzzy variables can be written again using expert judgment and then tuned using simulation and real experience. In fact the actual number of rules used in the inference system may be less than the number initially set. This is because some combinations of rules are unlikely to be found in real systems. The reduced number of rules does not reduce system performance as associations used to determine the truth of an assertion is largely determined by the dominating term.

The evaluation of the fuzzy rules must follow the sequence proposed in the decision tree. Thus the system inputs are fuzzified and successively used to assess other fuzzy variables in the cascade to the point where the final priority is evaluated. Graphic representations are invaluable in helping to assess how the fuzzy rules are operating. These may be generated by fixing all the variables except the two being assessed. Figure 19 shows an example of this where it is assumed that three variables (track quality, position and weapons capabilities) are maintained at a fuzzy value of 0.5 and both the threat and the hostility are varied over their entire ranges. This configuration might represent a situation in which the target is located at a medium range and has medium importance with respect to the weapons system of the radar platform.

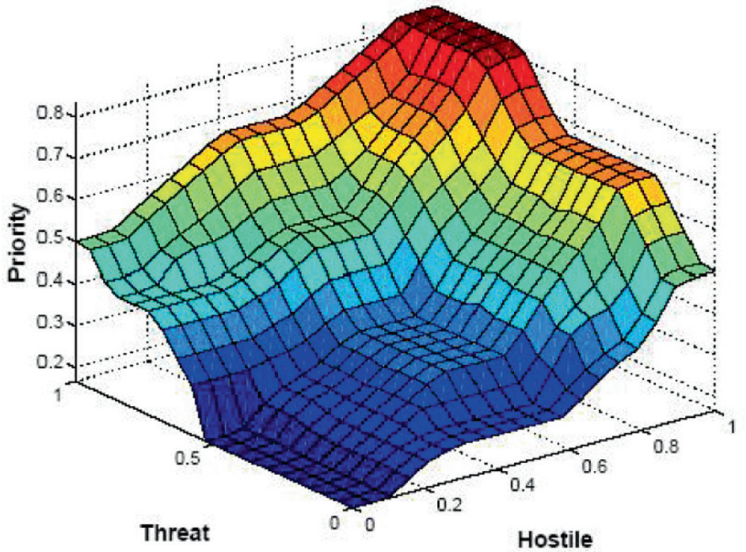

Fig. 19 Graphic representation of the fuzzy rules with the position, track quality and weapons systems fixed.

It is observed that, as might be expected, the priority increases as a consequence of increases in the degree of threat and hostility of a target. Conversely, low degrees of threat and hostility place the priority at a low level. Two other areas may also be identified on this surface. The first is related to the degree of hostility varying between 0.5 and 1 (medium to very high) and the degree of threat varying between 0 and 0.5 (very low to medium). The resulting priority increases as a result of rises in the degree of threat or hostility. However, the sensitivity to rise in the degree of hostility is greater than that of the threat. This behavior is explained by examining situation in which targets with medium and high probability of being the enemy are moving away from the radar platform. The hostility of the target is determined by its probability of being an enemy but the threat is determined more by its trajectory and position. Thus the situation is dominated by the identity of the target. The second area corresponds to degrees of threat varying between 0.5 and 1 and low levels of hostility. The resulting priority increases are a consequence of rises in the degree of threat or hostility. However, the behavior is different to the previous area. Thus the sensitivity to increases in the degree of threat is greater than the sensitivity to increases in the degree of hostility. This is explained by considering situations where, having low probabilities of being the enemy, targets move on threatening trajectories towards the radar platform. In this case, the way the target is approaching the radar platform has a greater effect in determining the final priority than its identity. Of course the manner in which these relationships are formulated is itself a variable and is one in which the expert judgment plays a key role. Thus there is a learning process during which these rules and relationships will be refined in the light of experience.

Having defined and tuned the fuzzy if-then rules the method for prioritising the relative importance of tracked targets can be validated against test trajectories. In the example presented here the scenario consists of targets with different identities and velocities. The analysis shows that by knowing the identity of the targets their priorities may vary. This provides valuable information to be accounted for when deciding how to allocate radar resources in overload situations. Two cases are presented for targets moving towards the radar platform on constant-velocity straight line trajectories. These have been chosen as they represent situations of a high degree of threat where targets may be moving towards the radar platform in order to start an attack. In addition, they represent the behaviour of the method when a variable such as approach is fixed. This helps simplify the analysis and the evaluation of the reasons for the results of the prioritisation. The system can also be examined in more complex scenarios where all variables involved in the prioritisation are changing over the simulation.

The left hand side of Figure 20 shows the first test trajectory where a target moves towards the radar platform on a straight line, having a constant velocity of 300 m/s. The red dot indicates the origin of the trajectory. Three targets are assumed in the analysis. They have the same dynamics and flight height; however, their probabilities of being enemy are different as follows: 1 (enemy), 0.5 (unknown) and 0.1 (friendly), corresponding to the red, blue and green curves respectively. The evolution of the resulting priorities is seen in right hand-side figure shows that, in general, all priorities increase as the targets move towards the radar platform; and the greater the probability of being enemy, the greater the resulting priority. Figure 18 also suggests that priorities of targets which have unknown identity present a similar behaviour to friendly targets in the early stages of the trajectory. This may be explained by the fact that during that period, the range of the targets is longer than the tactical range of the platform weapon systems. This happens until around 80 s. From that instant, as the target is moving close to the boundaries of this weapon systems

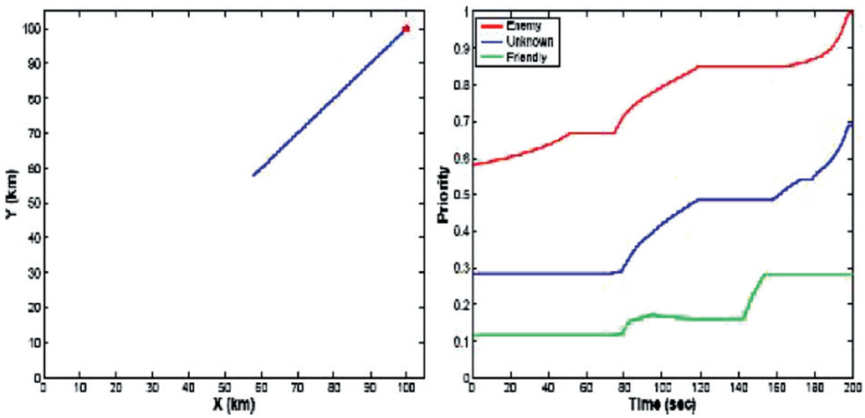

Fig. 20 Resulting priorities for three targets with different probabilities of being enemy, moving on the same trajectory.

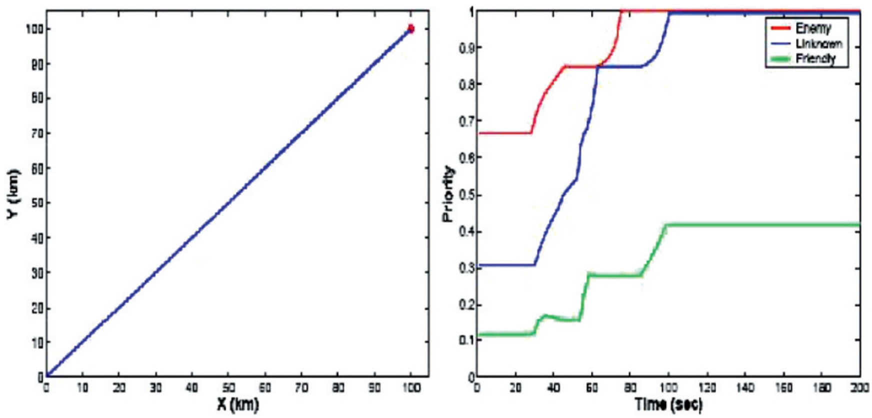

Fig. 21 Resulting priorities for three targets with different probabilities of being enemy, moving on the same trajectory. Target velocity: 800 m/s.

tactical range, the degree of threat of the unknown target is likely to increase. Thus, its priority evolution has the similar behaviour to the priority evolution of the enemy target. The closer the unknown target is, the higher and the closer to the enemy target its priority will be. At short ranges, if the identity of the target is still unknown, the target is assumed to be enemy, and its resulting priority is assessed as that.

Figure 21 presents the results of a simulation where targets are assumed to move on a straight line trajectory but this time with a velocity of 800 m/s. The same probabilities of being enemy as in the previous case are considered. Due to the high velocity and short ranges, the evolution of the priorities is now rather different. During the first few seconds of simulation, both unknown and enemy targets have slightly higher priorities than in the first example. This may be explained by their high velocities.

All target priorities remain fixed until about 30 s, when the target position is getting close to the weapon systems operational range. Before 30 s, all targets have the maximum priority possible for the set of characteristics of their dynamics, identity and the capabilities of the weapon systems. Thereafter, the priorities are increased in order to allow the radar platform to face the threat. It is observed that, from around 30 s to 60 s of simulation time, the priority of the unknown target presents a high rate of increase. The analysis indicates that more importance is progressively given to this target which is gradually assumed to be like an enemy target, because its velocity is very high, the target is approaching the radar fast, and its identity is unknown over this period. From around 60 s to 85 s, the unknown target has the highest priority possible for the combination of input variables which determine its importance. From 85 s, its priority increases again, reaching its highest at around 100 s, when the target position is within operational range of the platform weapon systems and as a consequence both enemy and unknown targets have the same priority. Such an unknown approaching target is considered to be of highest importance because of its potential degree of danger, represented by its velocity, the way it is approaching the radar platform. Like the unknown target, the priorities of both enemy and friendly targets increase from around 30 s, as they are getting close to the weapon system operational range. These priorities continue to increase reaching their maximum values not later than 100 s of simulation, when the position in within the operational range with a degree of membership of 100%.

The results of the situations examined here suggest that the fuzzy logic approach is an intelligent and valid means for evaluating the priority of targets. By imitating the human decision-making process, and by combining dynamic characteristics about radar tracking and military aspects, such as the ability of the weapon systems of the radar platform to face potential threats, the fuzzy approach may represent an effective and intelligent support for decisions regarding radar resource management. It also demonstrates the range of behaviours that such an approach can cope with and that it does so in a manner consistent with our definition of an intelligent system.

Prior knowledge can also be introduced in a variety of other ways, but has only been addressed by a few researchers [1, 2, 10]. For example, whilst there has been much work carried out on DPCA and STAP for GMTI there has been much less examining the role of prior information. This may well be a reflection of the fact that the processing needed for useful simulation and implementation of such Knowledge Based (KB) systems has only become available in recent years. However, it seems intuitive that there is a lot to gain from using external information to improve the performance of the MTI process. Some initial studies have examined the use of Geographic Information Systems (GIS), which have information relating to aspects such as land use, building disposition, terrain data etc and hence may be rewarding sources of information.

A major question to be answered is ‘how to integrate these disparate data types with the MTI process and hence what is the likely maximum improvement in MTI accuracy that can be expected’? As yet there appears to be no emergent convincing answers and methodologies. It is complicated by the fact that there are different qualities of terrain data, GIS data etc and so one would expect there to be a relationship between the quality of data and the benefit it can bring to the MTI process. For example, there are data of higher resolution than others but there has been no work to confirm that improvements in resolution of the data used by KB techniques are proportional to the performance improvement in MTI radar. Indeed, the key to good KB data may lie in another property such as the existence and relative location of roads and buildings. Similarly, should there be a lower limit on the quality of KB data needed to ‘add value’ to the MTI processor and is there a level of MTI fidelity that is required for KB to work? The issue of data fusion is clearly important in KB research. Additionally, the amount of data available will continue to increase and the ability of human analysis of it all will become an unrealistic prospect, if this is not already the case. Yet refraining from any analysis of this data will almost certainly mean missing out on useful information that could be the difference between detecting a moving target of interest on the ground or not.

There are two strands for information fusion in this context: fusing data of the same type and fusing data of different types. A KB MTI system can be thought of as a fusion of different types of data. Fusion of the same type of data is of equal relevance: obviously the cumulative information from two maps at different scales will be higher than the information from just one map, so how should the two maps be fed into an MTI process? Should the maps be fused beforehand, or should the fusion be one of a fusion of different types within the MTI system as described above? The fusion of similar-type data may be thought of as a separate area. For example, how should two maps or SAR images be fused for performance improvement as part of a KB MTI system? There is significant work on the fusion of SAR images from different radars, but not on using the result to provide knowledge for GMTI. Using SAR data itself for MTI processing has been attempted with promising results, and this type of approach would involve the fusion of two different types of MTI data, in which there has been little exploration. Furthermore maps and terrain data can be considered historical compared to real-time MTI data. Could therefore historical MTI data be used to provide extra knowledge? The literature on KB techniques for MTI radar appears to show only that there is good potential for increased MTI performance. The main focus of KB techniques in MTI radar has been to intelligently select training data for STAP. There are many more areas that are not yet investigated. Finally, it should be noted that, in published literature both on STAP/DPCA and KB techniques, is that the application to real radar systems is lacking. Real data provides the ultimate test for new systems because it brings with it real world errors that are often difficult to simulate or even unexpected.

A summary of information sources that could be obtained before live GMTI gathering but used for knowledge based processing with live data is shown below. Multiple sources of the same time, but gathered at different times, could be used with data fusion methods outlined above. Some sources are evidently more readily available than others.

- GIS data from digital maps, encompassing terrain and land elevation information, ground cover and transport routes
- Airborne, look-down optical imagery from reconnaissance aircraft or satellites
- GMTI data from historical runs over the area of interest
- Known clutter information of the scene (clutter maps for a radar with given parameters)

In a STAP system GIS data could be used to select more accurate training data for generation of the interference covariance matrix. Optical imagery or high resolution SAR imagery from the scene of interest could also be used to either select suitable training data or to identify targets of interest. Image processing techniques would be required in this case. Clearly this is an area of development with much further research necessary before KB diversity systems become common place.

7 Conclusions and Summary

In this chapter we have examined and demonstrated the value of exploiting sensor and platform diversity together with prior knowledge in improving system performance. Analogy with echo locating bats provides encouragement that such techniques will lead to future developments that embody genuine intelligence, potentially offering vastly superior capability over the systems of today. However, there remains much research to be done before this potential can be realized and there are still many questions remaining.

References

1. J.R. Guerci and E.J. Baranoski, 'Knowledge aided adaptive radar at DARPA', *IEEE-AES Magazine*, **23**(1), 41–50, Jan 2006.
2. M.C. Wicks, M. Rangaswamy, R. Adve and T.B. Hale, 'Space-Time adaptive processing', *IEEE-AES Magazine*, **23**(1), 51–65, Jan 2006.
3. W.G. Carrara, R.S. Goodman and R.M. Majewski, *Spotlight Synthetic Aperture Radar: Signal Processing Algorithms*, Artech House, Boston, MA, 1995.
4. P.F. Sammartino, C.J. Baker and H.D. Griffiths, 'Adaptive MIMO radar system in clutter', *IEEE Radar Conference 2007*, Waltham, MA, 276–281, 17–19 April 2007.
5. F.T. Ulaby and C. Elachi, *Radar Polarimetry for Geosciences Applications*, Artech House, Norwood, MA, 1990.
6. P. Tait, *An Introduction to Radar Target Recognition*, IET Publications, London, UK, 2005.
7. M.W. Holderied and O.V. Helversen, 'Binaural echo disparity as a potential indicator of object orientation and cue for object recognition in echolocating nectar-feeding bats', *Journal of Experimental Biology*, **209**, 3457–3468, 2006.
8. M. Vespe, C.J. Baker and H.D. Griffiths, 'Radar target classification using multiple perspectives', *IET Radar, Sonar and Navigation*, **1**(4), 300–307, 2007.
9. S. Haykin, 'Cognitive Radar', *IEEE-AES Magazine*, **23**(1), 30–40, Jan 2006.
10. S. Miranda, C.J. Baker, K. Woodbridge and H.D. Griffiths, 'Knowledge based resource management for Multifunction Radar', *IEEE-AES Magazine*, **23**(1), 66–76, Jan 2006.