Aspects of NCTR for Near-Future Radar

Robert Miller, David Shephard, Mark Newman
BAE SYSTEMS Advanced Technology Centre
West Hanningfield Road
Great Baddow, Chelmsford
Essex CM2 8HN
United Kingdom

robert.j.miller@baesystems.com

SUMMARY

This paper considers a number of aspects related to the achievement of non-cooperative target recognition capabilities in current and near-future radar systems. The scope of the paper is restricted to consideration of the use of high range-resolution profiles. Three particular aspects are discussed. Firstly, the problem of achieving a high range-resolution capability on radars which typically have only narrow instantaneous bandwidths is considered; an approach is described in which the usual shortcomings associated with step-frequency waveforms are avoided. Secondly, the consequences of having less than ideal performance from the radar system are considered. The loss in classification performance which occurs when returns are degraded in terms of resolution and signal-to-noise ratio are described. The results given apply to civil aircraft and compare performance from a feature-matching and a profile-correlation algorithm. The third aspect considered relates to the nature of the classifier itself. There are numerous choices to be made; we discuss what data should be used for classification, sources of reference data for classifiers and different types of classification algorithm. A focus is placed on the representation of reference data as a scattering centre model of each aircraft of interest; such a model attempts to give an abstract representation of key features in a form which may incorporate both radar and non-radar data, and which is not particular to any one radar system.

1.0 INTRODUCTION

The desirability of the inclusion of some degree of non-co-operative target recognition (NCTR) capability in current and future radar systems is widely recognised. Discussions of NCTR techniques frequently assume the use of a purpose-built radar; however, in view of the long life-times of contemporary radar systems, this is not realistic. It is highly desirable to utilise if at all possible the untapped potential of current radar systems to perform NCTR. The paper starts from this view-point, and discusses a number of issues associated with introducing NCTR capability to contemporary or near-future radar systems.

Two different approaches for using radar to provide NCTR capability are commonly discussed, i.e. (i) analysis of frequency modulation of returns (jet engine modulation, helicopter rotor modulation) and (ii) comparison of high-resolution range profiles (HRRP) with reference data on signatures of targets of interest. This paper considers only the latter approach. The scope of the paper is also limited to a consideration of air targets.

The use of existing or modestly upgraded radar systems often implies that there are limitations on the waveforms that can be used for NCTR, and these will impact on NCTR performance. The first part of this paper discusses a novel technique for generating HRRP waveforms from a sequence of narrow-band pulses. This technique avoids the shortcomings of a conventional step-frequency approach, which may lead to aliasing, high range sidelobes and wrap-round of long targets. The new technique combines pulses

Paper presented at the RTO SET Symposium on "Target Identification and Recognition Using RF Systems", held in Oslo, Norway, 11-13 October 2004, and published in RTO-MP-SET-080.

using motion compensation which is accurate to within a fraction of a wavelength to achieve a high bandwidth synthetic waveform which is free of phase discontinuities.

The second part of the paper considers results obtained from degrading high-resolution profiles to show how classifier performance changes with variation in resolution and signal-to-noise ratio. Two different kinds of classifier are considered, the first based on feature matching, and the second based on direct comparison of profiles with reference data.

The third part of the paper considers in greater depth the performance of different types of classifier, and the nature of the reference data used to perform the classification and its impact on performance. This consideration also overlaps with fundamental concerns regarding the source of reference data on many targets of interest – while it may sometimes be possible to obtain detailed radar measurements of friendly aircraft, different sources of data will need to be used for other aircraft. Particular consideration is given to the use of a reference model which describes the principal scattering centres on each aircraft of interest. The validity of such a model is considered, as is the inclusion of both radar and non-radar data in such a model.

2.0 GENERATION OF HRRP WAVEFORMS

The ability of a radar system to identify a target from its radar range profile is directly dependent on the achieved range resolution, which is in turn dependent on the transmitted bandwidth of the signal. Early work indicated that a bandwidth of about 400 MHz is required to identify air targets. Ideally this would be achieved using a wideband instantaneous waveform. Most radars currently in service have a very limited instantaneous bandwidth. This makes it impossible without a major redesign to achieve the required bandwidth using an instantaneous waveform. The overall RF bandwidth of the system is generally wide enough that an alternative is to synthesise the required bandwidth by frequency stepping.

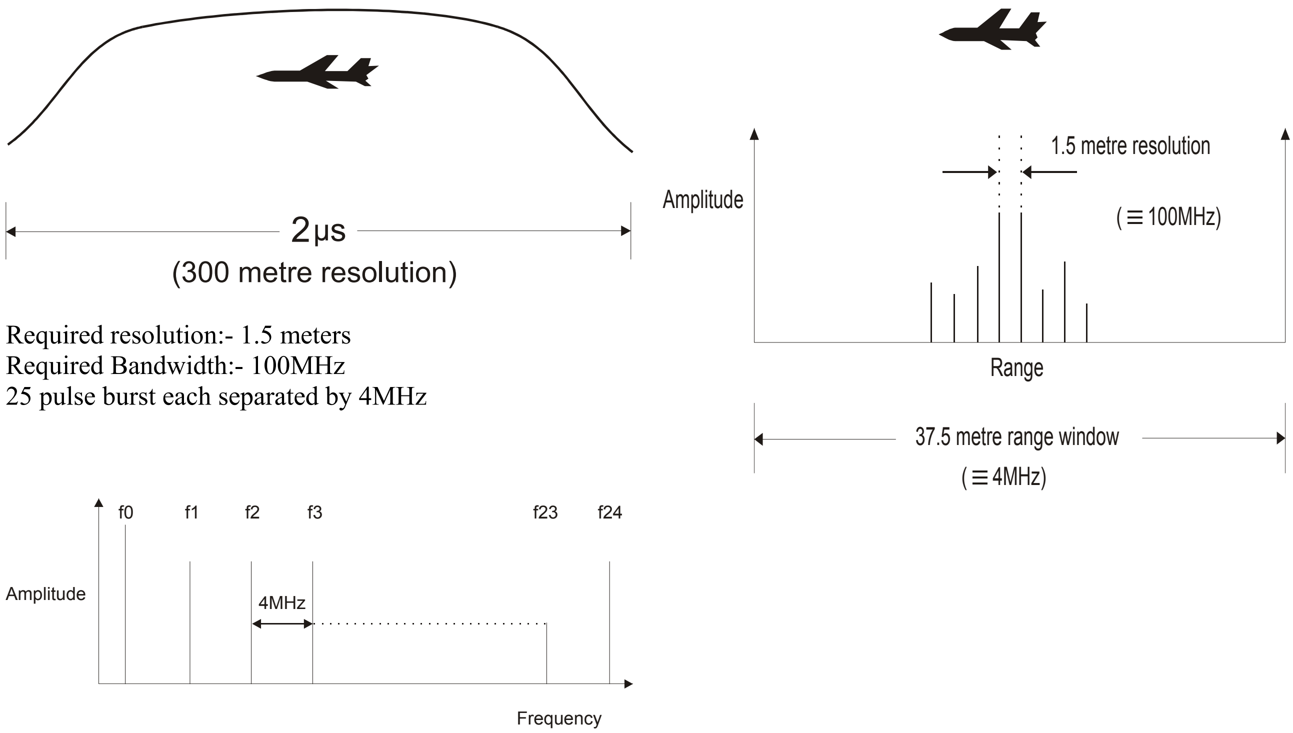

Figure 1: Traditional Step Frequency Technique

This frequency-stepping technique has been shown to work successfully and is described in detail in standard reference books such as that by Wehner [1]. This type of technique is shown schematically in Figure 1. In this example, a 25-pulse burst with a 4 MHz separation between pulses is used to generate a synthetic bandwidth of 100 MHz. It is important to note that the size of the resulting time-domain window is proportional to the reciprocal of the frequency step; in the example given the 4 MHz frequency step results in a synthesised range window 37.5 metres long. Aircraft longer than this window will be incorrectly profiled due to wraparound effects. This will restrict the length of targets that can be profiled using the traditional step frequency technique. To address this issue, an alternative technique known as Hybrid Stepped Frequency Range Profiling has been developed by BAE SYSTEMS ATC.

The hybrid stepped-frequency method involves the transmission of a series of narrow-band FM chirps, transmitted at stepped carrier frequencies. The returned signals are combined to form a wide bandwidth result spanning a continuous range of frequency – there are no gaps or phase discontinuities. The FM chirps can be compressed against a reference either individually before summation, or all together after summation. This technique has been successfully demonstrated at our radar test site in Great Baddow, UK.

The technique has been used to produce range profiles of both stationary calibration targets and aircraft in flight. Figure 2 shows a range profile of a stationary test target. The range swath is in excess of one kilometre and shows that the hybrid step-frequency technique is not limited by the 1/frequency-step range ambiguity of the traditional method.

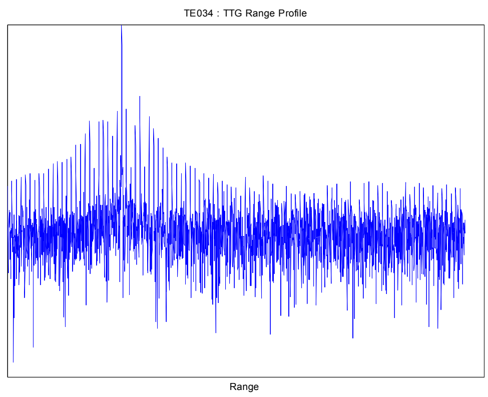

Figure 2: Range profile of test target

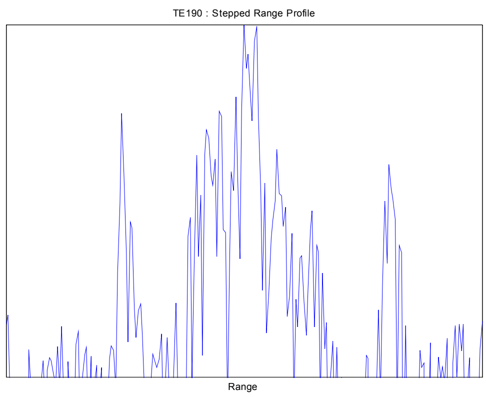

Figure 3: Profile of airborne target

Figure 3 shows a range profile of a target of opportunity identified as a Boeing 747 measured at a range of 21 km and having a radial speed of 351 knots. The 747 has a length of over 70 m; it is evident that no fold-over is occurring in the profile.

The detailed sidelobe structure of the motion-compensated waveform obtained from airborne targets has been compared with that obtained from calibration targets; they are found to be virtually identical. This demonstrates that the motion-compensation techniques employed are sufficiently accurate and robust to generate high-quality profiles. The range profiles obtained using the hybrid technique are generally found to be of very high fidelity and compare favourably with data measured using an instantaneous wideband waveform.

The processing can be taken a stage further by applying the technique to each Doppler channel in a burst. The resulting range-Doppler map is shown in Figure 4(a). The distribution of the Jet Engine Modulation, either side of the skin return, is within the anticipated range window and does not show any range-Doppler coupling. Figure 4(b) shows a second result measured from a Boeing 737. This target was at 13.5 km and a radial speed of 272 knots. In this example, each of two engines give rise to two distinct Doppler sidebands.

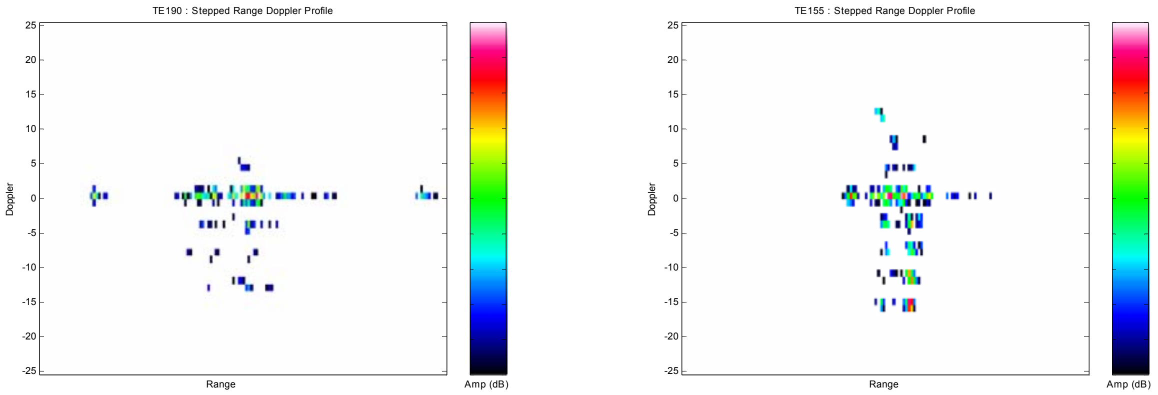

(a) range-Doppler map for Boeing 747 **(b) range-Doppler map for Boeing 737**

Figure 4: Range-Doppler maps

This technique for producing high range resolution profiles of air targets is robust. It can be applied to targets of opportunity with sufficient information available from the measured data to enable motion compensation to be undertaken to high accuracy.

3.0 VARIATION OF RESOLUTION AND SNR

This section considers the effects of variation of range-resolution and signal-to-noise ratio on classifier performance. Two forms of classifier are considered.

3.1 Data

High-resolution data was collected consisting of a sequence of 50 1 µS pulses, each modulated with a 270 MHz linear FM chirp. 1000 samples of each pulse were taken at a rate of 400 MHz. Pulses are compressed using a reference profile from a point target and the results are motion-compensated to produce a stack of aligned range profiles. Doppler processing is then applied to produce a range-Doppler map. This separates the skin echo from engine returns. A sample set of 58 aircraft datasets, consisting of six different types, was selected from a database of trial data. Table 1 lists the aircraft selected. The identity of these aircraft has been confirmed by National Air Traffic Services (NATS).

Aircraft Type	Class	No. of Datasets
Boeing 747	large civil	11
Lockheed Tristar	large civil	8
Boeing 707	large civil	8
Boeing 737	medium civil	10
Boeing 757	medium civil	11
McDonald Douglas MD80	medium civil	10
Total		*58*

Table 1: Aircraft types in trial data

3.2 Classification Algorithms

Two types of algorithm have been considered, a *feature-matching* algorithm and a *profile-correlation* algorithm.

The feature-matching algorithm attempts to match features of the observed aircraft to a list of features stored in a database. These features include length, distance of engines from the nose and number of engines. A range-Doppler map (see Figure 4) is used to distinguish the skin return from engine returns. The algorithm is designed to recognise aircraft over aspect angles between 0° and ±40° from the nose since the JEM required to distinguish the engines is typically visible only over this range. The reference database used contains 71 aircraft types and variants. These aircraft have been separated into 4 generic classes: large civil, medium civil, small civil and large military. The information about each aircraft in the database has been gathered from technical data available in the public domain. The aircraft type is determined according to the greatest number of satisfactory matches; the class is determined from the type classification. An aircraft may be unclassified if no satisfactory matches occur.

The profile-correlation algorithm uses only the range profile of the skin echo from the data, i.e. the contents of the zero-Doppler bin in the range-Doppler map. Profiles are correlated with profiles in a reference database which consists of 612 sampled range profiles. The identity of most aircraft has been confirmed by NATS. The profiles in this database have been separated into the same four generic types as before: large civil, medium civil, small civil and large military. The trial dataset used is a subset of the reference profiles, but the algorithm excludes correlation of a profile with itself. A k nearest-neighbour algorithm is used with classification being based on the majority type or class from the 10 best matches.

In normal use, any NCTR algorithm would be used in conjunction with tracking data to estimate the aspect angle of the target aircraft; such data was not available for the profiles considered here, so classifier performance is generally not as good as it could be.

3.3 Reduction of Resolution and SNR

To assess the effect of variation of range resolution and SNR on the performance of the classifiers, the trial data was degraded in one of two ways, either (i) by reducing the sampling rate to approximately emulate the effect of reducing range resolution or (ii) by decreasing the signal-to-noise ratio by injecting additional noise into the original data.

3.3.1 Range Resolution

To estimate the effects of degradation in range resolution, the data in each of the trial datasets was methodically reduced by re-sampling. The effect of an increase in sampling interval (SI) from 0.375 m to 1.875 m on the range-Doppler map and the range profile is illustrated in Figure 5.

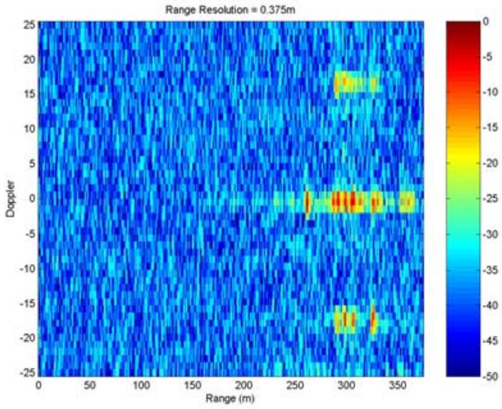

(a) range-Doppler map SI = 0.375 m

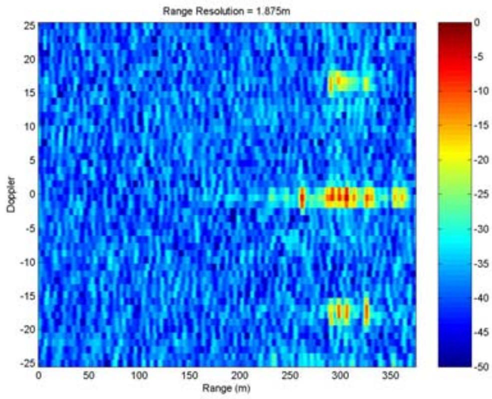

(b) range-Doppler map SI = 1.875 m

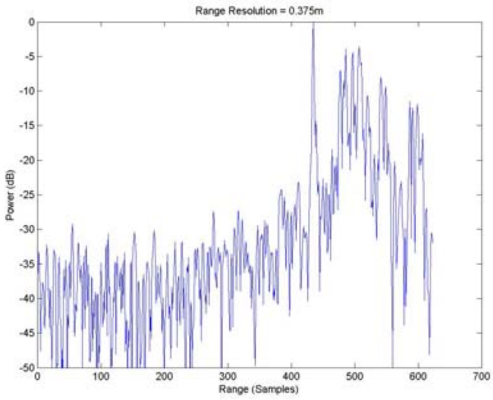

(c) range-profile SI = 0.375 m

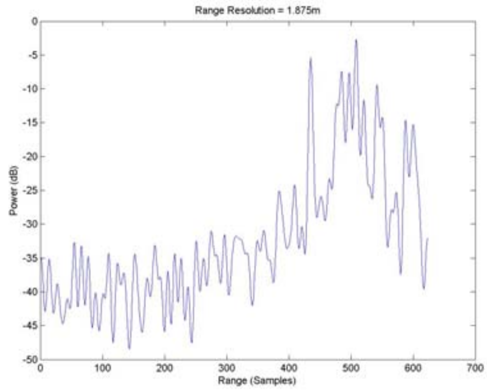

(d) range profile SI = 1.875 m

Figure 5: Effects of reduced range resolution on range-Doppler map and range profile

3.3.1.1 Feature-Matching Algorithm

Figure 6(a) shows that the feature-matching algorithm initially classifies around 50% of the trial datasets correctly to type, and that this performance falls away approximately linearly as resolution is decreased. Figure 6(b) shows that the algorithm is able to correctly identify the generic class of all classified datasets. The number of unclassified datasets rises as the sampling interval is increased. Datasets are unclassified when the algorithm is unable to match the aircraft length and the position of its engines to any reference set in the database.

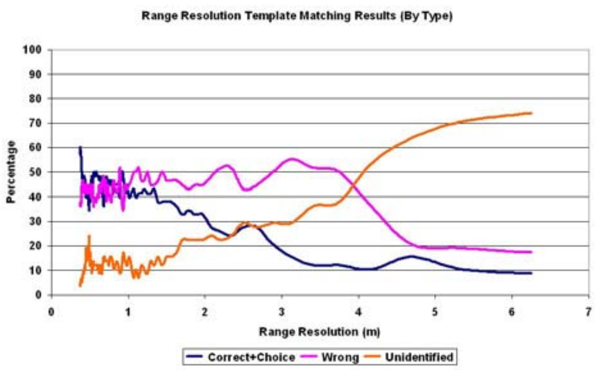

(a) classification to type

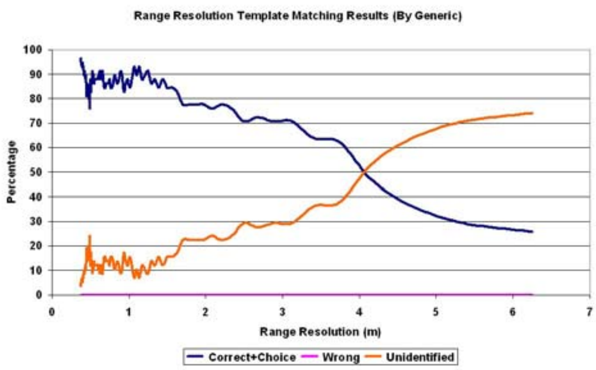

(b) classification to class

Figure 6: Effects of reduction in range resolution for feature-matching classifier

More detailed analysis of the results indicates that smaller aircraft are mis-identified, but that they are always identified with an aircraft of the same class. This is why the class results are much better than the type-specific results.

3.3.1.2 Profile-Correlation Algorithm

Figure 7(a) shows that the profile-correlation algorithm displays an approximately linear drop in performance as the sampling interval is increased. This algorithm is capable of correctly identifying over 80% of aircraft measured with a range resolution of 2.5 meters or less. Figure 7(b) shows that performance for identification of target class is similar to that for identification of target type. Datasets are never unclassified under the profile-correlation algorithm, but performance to class may be worse than performance to type. For example, suppose that 4 of the 10 matches are of a Boeing 747, 3 matches are a Boeing 767 and the remaining 3 are a Boeing 757. The type chosen would be a Boeing 747 (correct) but the generic class would be a medium civil aircraft (incorrect).

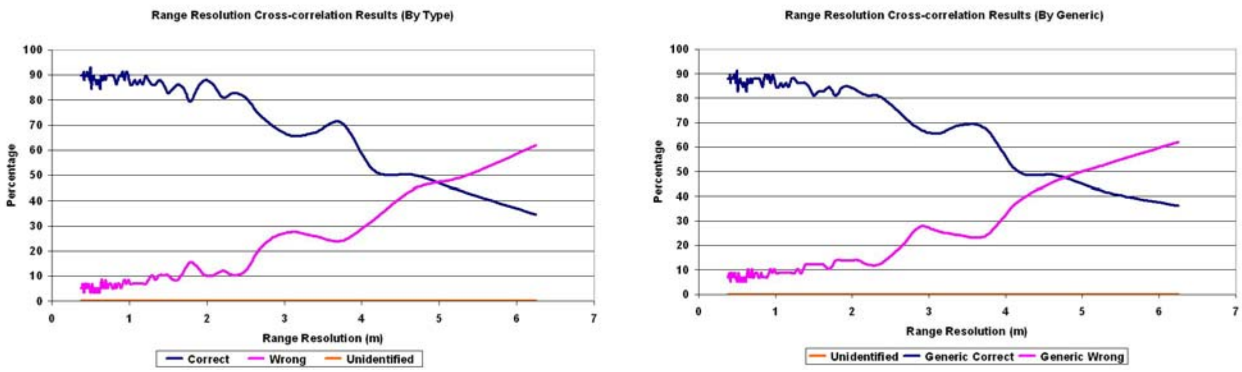

(a) classification to type **(b) classification to class**

Figure 7: Effects of reduction in range resolution for profile correlation classifier

Again, smaller aircraft tend to be more frequently mis-identified. The large Boeing 747 is correctly identified for 10 of the 11 Boeing 747 datasets, even with a sampling interval of 6.5 meters.

3.3.2 Signal to Noise Ratio

To simulate a reduced signal to noise ratio, the original datasets were combined with sampled noise data taken from the radar system with the transmitter turned off. Repeatable results were produced by using a single noise dataset injected at varying intensities.

3.3.2.1 Feature-Matching Algorithm

Figure 8(a) shows that there is a linear drop in the number of correct identifications as the SNR decreases. Figure 8(b) shows that the algorithm’s ability to identify the generic class of the aircraft is not significantly affected until the signal to noise ratio is decreased by around 15 dB. Up to this point, over 90% of the target aircraft are correctly identified to class.

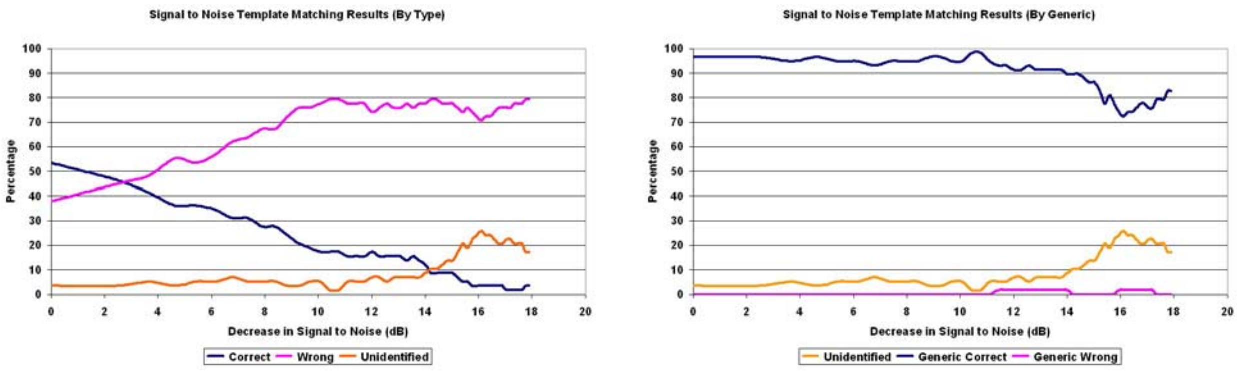

(a) classification to type (b) classification to class

Figure 8: Effects of reduction in SNR for feature-matching classifier

As the signal to noise ratio was reduced the smaller aircraft begin to be mis-identified. Since the larger aircraft like the Boeing 747 produce a larger return, and therefore have a higher signal to noise ratio to start with, the degradation of the signal is less than that of the smaller aircraft. The signal to noise change has a more significant effect on the signature of smaller aircraft than on larger aircraft.

3.3.2.2 Profile-Correlation Algorithm

Figure 9(a) and Figure 9(b) show that initially there is very little drop in the performance of the algorithm as the signal to noise ratio is decreased. However, once the signal to noise level has been decreased by around 8 dB, the performance begins to drop linearly. As previously, the performance of this classifier when identifying targets to class is slightly worse than when identifying the type of the aircraft.

As with the feature-matching algorithm, smaller targets are mis-identified first.

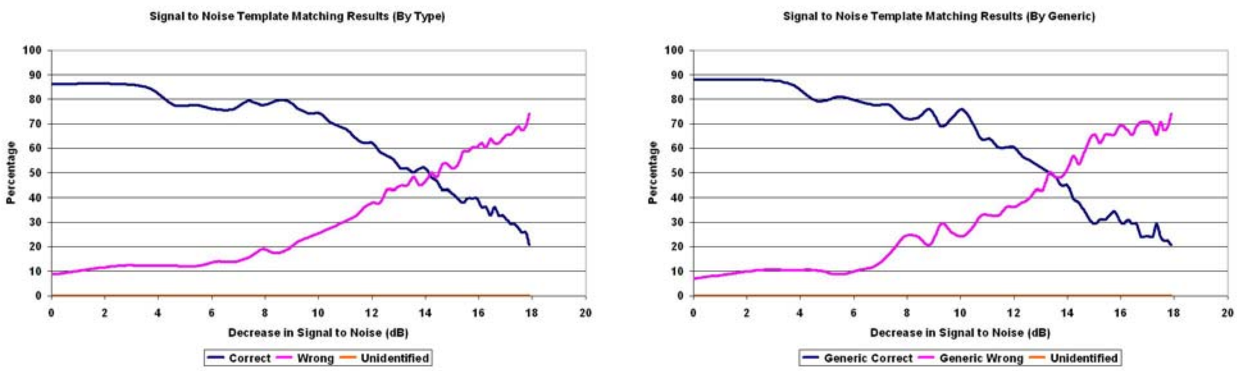

(a) classification to type (b) classification to class

Figure 9: Effects of reduction in SNR for profile-correlation classifier

3.4 Conclusions

Even with the relatively large aircraft considered here, it is apparent that decrease in range resolution leads to quite a rapid fall-off in classifier performance for both algorithms. It therefore seems desirable to use sub-metre range resolution if at all possible. Profile-correlation seems to be more robust to decrease in SNR for classification to type, whereas template-matching seems to be more robust for classification to class. It would be of interest to determine which features of the algorithms are responsible for these different performance characteristics. Finally, the performance of the profile-correlation algorithm is generally substantially better than the feature-matching algorithm, which indicates that it may be possible to identify a richer feature set to improve the performance of the latter algorithm.

4.0 CLASSIFICATION ISSUES

In the present context, a classifier is a mechanism which makes an association between radar measurements of an aircraft and a particular class or type of aircraft. In order to make this association, it makes reference to some form of prior knowledge about each type of aircraft which may be of interest. Statistical classification algorithms work on the basis that the measurements, possibly after some processing, may be regarded as a point in a high-dimensional 'decision space'. The prior knowledge serves to divide up the decision space into non-overlapping regions, each of which is associated with a particular target type. The algorithm notes which of these regions the measurement point falls into, and makes its decision accordingly. Provision may also be made for measurement points to be unclassified – the measurements may, for example, be too noisy.

There are many different types of classification algorithm - nearest neighbour, neural nets, tree pruning and so on. These vary in the complexity of the decision surfaces they may represent. However, they all depend critically on the quality of the information they are supplied with to form the decision surfaces. The information that is available and its quality is the focus of our emphasis here; the classification algorithm used is in all cases a simple, classical k-nearest neighbour algorithm [2].

4.1 Target Measurements

Classification is generally performed on the basis of several measured profiles of the target aircraft, not just one. In the extreme case, hundreds of profiles may be used in order to form an ISAR image of the aircraft. Rihaczek and Hershkowitz [3] use this approach in order to filter out off-fuselage returns; this has a distinct advantage, in that returns from the wings of an aircraft can often be confusing since many different stores configurations may be employed with a single type of aircraft. The use of ISAR has, however, at least two distinct disadvantages. Firstly, in order to obtain the large number of profiles required, the dwell time must be long, and this is often operationally unattractive. Secondly, ISAR relies on the use of cross-track motion relative to the radar, and it is therefore not at all clear that ISAR may be used for the important case of aircraft flying more or less directly towards the radar. For these reasons, ISAR has not been considered further in the present work.

Radar returns from aircraft vary rapidly with aspect angle, principally due to multiple scatterers occurring in the same range gate and interfering with each other. This variability is unhelpful to classifiers and it is therefore desirable to reduce it where possible. Some degree of reduction in variability may be achieved by averaging over a small number of profiles; this technique has therefore been adopted in the current work. Separation of engine returns from fuselage returns may also be achieved by applying Doppler processing to a small number of profiles.

4.2 Prior Knowledge

Prior knowledge of the backscatter characteristics of an aircraft may be obtained from many sources. The best classification results have been obtained using detailed radar measurements of target aircraft of interest. Such measurements must be made over all aspect angles from which the aircraft is likely to be observed, and are therefore time-consuming and expensive to obtain. Exemplars of likely hostile aircraft may also be difficult to obtain. Consequently, alternative sources of prior knowledge have been sought.

Good results have also been obtained using scale measurements of detailed aircraft models. Such models are again expensive, and the quality of classification achieved depends on the level of detail which goes into the models. It is not sufficient to represent backscatter from the just the skin of the aircraft, since returns may be obtained from components hidden under the skin by radar-transparent materials; perhaps the most obvious case is radar equipment covered by a radome in the front of many fighter aircraft.

A third source of prior knowledge is the use of detailed computer models of aircraft used in conjunction with computational electromagnetics (CEM) codes to infer backscatter characteristics. Errors may occur from two principal sources. Firstly, the model may be insufficiently accurate or lack critical components such as antennas or small air inlets. Secondly, CEM codes used are often based on approximations such as physical optics and GTD (geometric theory of diffraction). Rigorous 'full-wave' codes such as method of moments and FDTD (finite difference time domain) also exist, but they are computationally expensive, and it is not yet feasible to apply such codes to, say, fighter-sized aircraft illuminated with X-band frequencies. However, there is continual progress in this area [4].

Prior knowledge from any of the sources noted above can be used to provide templates to which measured profiles are matched, or may be used as training data to set up decision surfaces. Rather than use reference profiles directly in this manner, we have chosen to identify prominent returns in each profile and to incorporate these into a model from which reference profiles may be regenerated. The notion is that, for each aircraft of interest, we generate a model of the prominent scattering centres on the aircraft. Each scattering centre is described in terms of its position and how the amplitude and phase of its return varies with aspect angle – note that this is quite different from describing the returns in terms of those from a number of isotropic point scatterers. For the scattering centre model to be of value in classification, there should be only a small number of scattering centres for each aircraft (say less than 10), and the returns from each scattering centre should persist over an appreciable range of aspect angle (say greater than 10°).

An advantage of this approach is that other forms of prior knowledge than those mentioned above may be easily incorporated into the model. In particular, the approximate location of at least some scattering centres may be inferred from material such as engineering drawings, CAD models and photographs of aircraft.

The scattering centre model is only a model – it is at best an approximate representation of radar backscatter from an aircraft. This model may be directly validated by forming a map of reflectivity over the whole aircraft constructed from profiles taken over a range of accurately known aspect angles using tomographic principles. Scattering centres should be evident as 'hot-spots' on tomograms generated in this manner.

Figure 10 shows a set of range profiles obtained at different aspect angles (a) and the tomogram derived from them (b). The range profiles are aligned in such a way that the aircraft appears to rotate about some fixed point. Prominent scattering centres in the range profiles will describe sinusoidal arcs – several of these are evident in Figure 10(a). The tomographic reconstruction maps these arcs into single points in the tomogram.

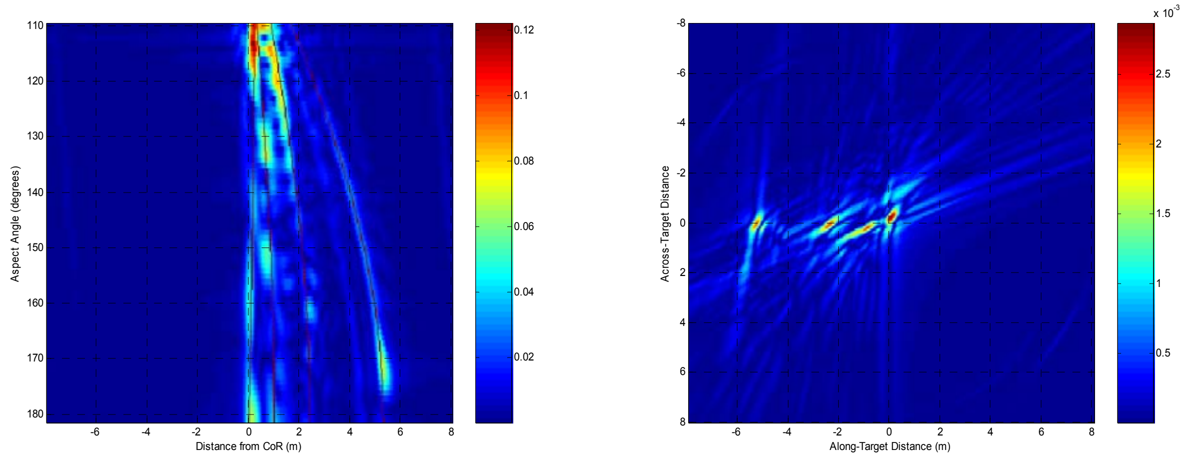

(a) aligned range profiles **(b) tomogram constructed from profiles**

Figure 10: Range profiles and tomogram

The range profiles shown in Figure 10 were obtained by applying a CEM code to a computer model of an aircraft. The fact that the tomogram exhibits only a small number of discrete points supports the validity of representing backscatter from an aircraft as a scattering centre model. A similar analysis may be applied to actual radar measurements of an aircraft to determine to what extent a scattering centre model is appropriate.

4.3 Classification

The nearest neighbour method of classification relies on a comparison of measured data, or features derived from it, with reference data. For aircraft targets, comparison is complicated by the following two factors.

Firstly, aircraft signatures depend on aspect angle, but the aspect angle at which an aircraft is measured is seldom known at all accurately. The course of the aircraft may be estimated from tracking data, but this does not accurately indicate its aspect, since aircraft commonly do not point in the direction they are heading due to cross-winds – they 'crab' with respect to their course. The crab angle may be several degrees. For this reason, it is necessary to compare measurements with reference data applicable to a range of aspect angles – we use ±5° about the estimated course.

Secondly, the range at which features occur in range profile measurements is potentially a powerful discriminant amongst aircraft, so that, before reference data can be used, it must be aligned with the measured profiles. Such alignment is a significant part of the overall classification algorithm.

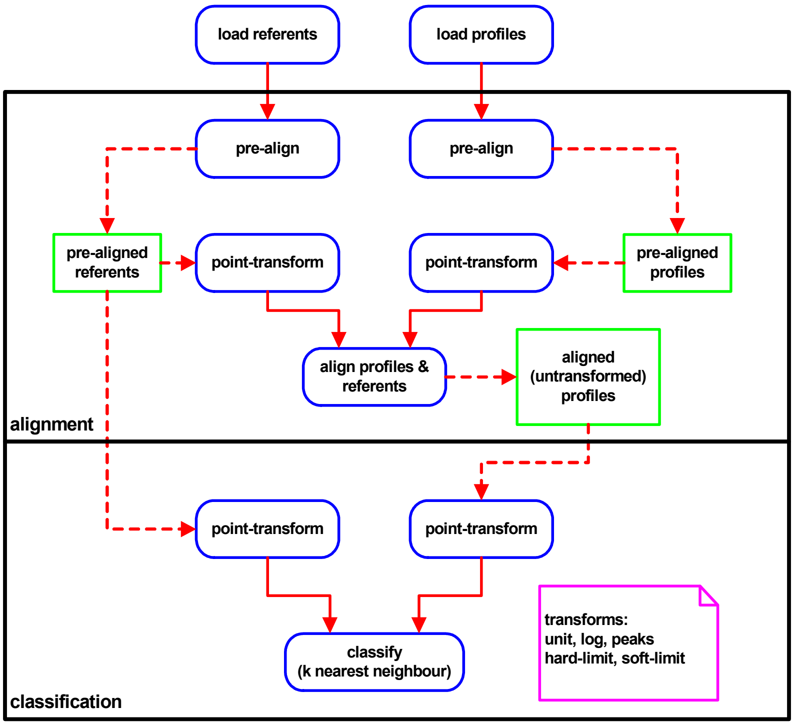

Figure 11: Classification scheme overview

An overview of the classification scheme used in the current work is given in Figure 11. First, relevant (measured) profiles and referents are loaded. These are then roughly aligned by centring them within the range swath. Next, the pre-aligned profiles and referents are subject to a *point transformation*; the use of these point transformations is a key feature of the scheme described here. Point transformations may be as simple as taking the logarithm of each point on the referent or profile, or may be more complicated such as detecting the positions of peaks then taking a hard limit, i.e. applying a threshold to reduce the value to one or zero. The point transformations allow particular features in the profiles to be emphasised or de-emphasised. A soft threshold, i.e. reducing all values below a specified threshold to zero, is useful in screening out noise in measurements. Alignment is achieved by one of a variety of methods, such as minimising the correlation between profile and referent with respect to shift in range.

Following alignment, the referents and profiles are subjected to another point transform, which is generally different to that used for alignment. An identity transform leaves data unchanged, giving a comparison between the full profile and referent; the referent acts as a template, so the algorithm amounts to template-matching. Comparisons of this form tend to give undue emphasis to noise. A combination of soft-limiting and peak detection allows just the positions and amplitudes of peaks to be compared. With this point transform, features are essentially extracted from the profile so that the classification scheme is feature-based rather than template-based. A slightly different transform may be used to hard-limit the peak amplitudes; this variation of the algorithm thus compares just the positions of high returns in the profile. This is valuable, given that use of non-radar data to form referents may give reasonably accurate indications of peak positions, but will seldom give accurate information on the amplitudes of these returns. Using the scheme above, it is possible to directly gauge the effect of discarding or ignoring parts of the original data.

5.0 CONCLUSION

A number of aspects of NCTR for near-future radar have been discussed. It has been shown that it is feasible to provide such radars with high-resolution capability, and the need for such capability, even for classifying large aircraft, has been demonstrated. The use of a codified form of reference data in the form of a scattering centre model has been discussed – a model of this kind allows both radar data and non-radar data to be utilised as reference material. A classification scheme has also been described which allows a range of options to be explored in exploiting reference data of this kind. What remains is to establish the performance of such algorithms both in classifying the type and the generic class of aircraft of interest.

6.0 ACKNOWLEDGEMENTS

The work reported here on high-resolution waveforms was sponsored by BAE SYSTEMS Avionics; the work on the effect of decreased range resolution and signal-to-noise ratio was sponsored by AMS (UK). Work on classification issues has been performed under the joint industry/MoD sponsored Electromagnetic Remote Sensing Defence Technology Centre (EMRS DTC). The authors wish to thank all of these sponsors for their encouragement in this work and for their permission to publish it.

7.0 REFERENCES

[1] D. R. Wehner, “High Resolution Radar”, Artech House, 1987.

[2] P. A. Devijver, J. Kittler, “Pattern Recognition - A Statistical Approach”, Prentice/Hall International, 1982.

[3] A. W. Rihaczek, S. J. Hershkowitz, “Theory and Practice of Radar Target Identification”, Artech House, 2000.

[4] J. A. Lord, J. D. S. Langley, “Finite-difference time-domain (FDTD) analysis of antennas on a supercomputer”, GEC-Marconi RF Technology Conference Digest, pp.115-119, April 1998.