Advances in Radar NCTR Using Non-Radar Referents

R. J. Miller, C. B. Wyllie, D. J. Shephard
BAE SYSTEMS Advanced Technology Centre
Great Baddow, Chelmsford, CM2 8HN

Abstract

The use of scattering centre models to recognise aircraft in conjunction with high-resolution radar profiles is examined. Using a set of data obtained from three types of civil aircraft, the performance of a conventional classifier and classifiers using scattering centre models derived from both radar data and non-radar data is compared.

Keywords: radar, aircraft recognition

Introduction

This paper considers an aspect of the problem of recognising aircraft using high-resolution radar profiles. Recognising aircraft in this way has been considered for many years [1]. Common practice is to assemble large databases of radar measurements of aircraft of interest and to use these as reference material to classify aircraft seen in theatre. Problems with this approach arise for several reasons. First of all, even an aircraft of one particular type may have many different configurations, due, for example, to the carriage of a variety of under-wing stores. Secondly, detailed radar measurements of aircraft constitute both commercially and militarily sensitive information, and are unlikely to be available for non-allied aircraft.

So, another way must be found.

In order to address the problem of recognising aircraft without using large databases of previously acquired radar measurements, the notion of modelling the returns from an aircraft in terms of a small number of scattering centres has been adopted. Previous work [2], [3] has examined high-quality radar data obtained from an aircraft on a turntable to see whether such a model is valid; it has been found that much of the backscatter can indeed be represented in terms of a scattering centre model or SCM. The objective of the current paper is to explore the use of such scattering centre models in classifying aircraft.

A scattering centre model is illustrated below. Each scattering centre has a well defined position and a polar diagram, which may extend in both azimuth and elevation, although only azimuth is shown below. In general, the polar diagram gives the amplitude and phase of the return from the scattering centre as a function of aspect angle, although, in practice, accurate phase information is difficult to obtain.

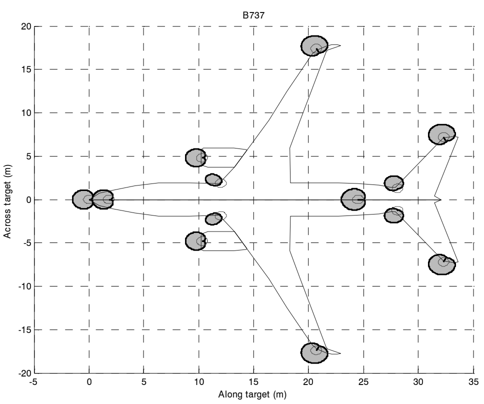

Figure 1: Scattering centre model

Scattering centre models may be built up from either radar or non-radar data, or a mixture of the two; however, the precision of an SCM constructed from non-radar data is likely to be less than that obtained using radar data. The principal difficulty is that of accurately predicting the amplitude of the backscatter from any given scattering centre. So, how important is the loss of amplitude information to aircraft classification? This question is addressed in much of the remainder of this paper by way of an example using profiles from civil aircraft. The paper also illustrates the versatility and effectiveness of the use of scattering centre models.

Data

In order to illustrate the approach adopted, we will use some radar measurements collected over a number of years from a radar on the roof-top of one of the buildings at the BAE SYSTEMS site at Great Baddow. The radar operates at a frequency of 3.5 GHz (8.6 cm wavelength), and has a nominal range resolution of 0.4 m.

The dataset considered consists of a number of profiles for each of three types of commercial aircraft – details are given in the following Table.

Aircraft	No. profiles
Boeing 737-700	104
Boeing 747-100	89
Boeing 757-200	83

Table 1: Dataset content

Aircraft were identified using SSR codes. The aspect angle for all profiles is similar, since the aircraft are constrained by designated routes. The dataset is illustrated in Figure 2 below. Each row of the figure shows a single profile, with its amplitude indicated by the colour scale. The profiles are grouped so that all profiles from the same type of aircraft are shown together.

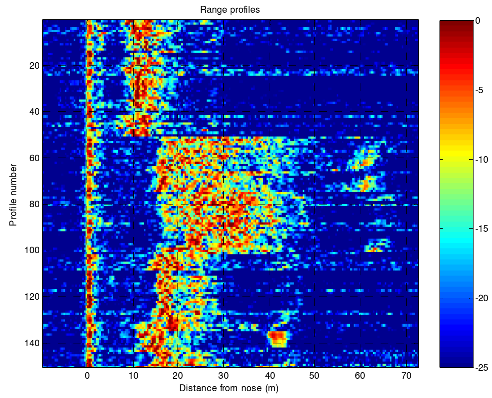

Figure 2: Range profiles

Conventional Classification

It has been found that a reasonable classification of the profiles can be obtained by choosing a single profile from each type of aircraft in the dataset; each profile is used as a template for all other profiles and a classification is formed using a nearest-neighbour classifier. In more detail, the classification process is as follows.

For each profile, before classification can be attempted, the template must be aligned to the profile. Alignment is achieved by lining up the first prominent return in each profile with that on each template; 'prominence' is defined as exceeding some specified threshold. In the particular dataset being considered, a prominent return is consistently obtained from the nose of each aircraft – this significantly assists the alignment process. Good alignment is important in achieving good classification performance, and, in other cases, more general procedures such as applying a shift which maximises the cross-correlation between profiles and templates may need to be considered.

Once all templates have been aligned with a given profile, classification may be attempted. Prior to classification, a point transformation is applied to both template and profile; such point transformations may include transformations of amplitudes to a logarithmic scale, clipping, hard-limiting and re-normalisation. Transformations are used to improve the performance of the classification process compared to what would be obtained if it were applied to untreated data. A nearest-neighbour algorithm is used, which means that some measure of distance (or, more generally, a measure of mis-match) is determined for each profile/template pair, and a class is assigned to the profile according to which template lies closest.

The general process is summarised in the following Figure.

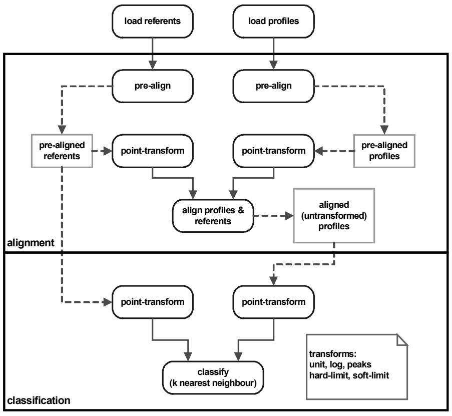

Figure 3: Classification process

The results of the classification process may be summarised in a 'confusion matrix', an example of which is given below. The confusion matrix shows the percentage of correct classifications of each aircraft type along the diagonal, and the percentage with which one aircraft type has been confused with a different type in off-diagonal entries. The particular confusion matrix given in Table 2 gives the performance of the classification process described above.

	B737	B741	B752
B737	89.1	0.0	15.0
B741	0.0	97.7	8.8
B752	10.9	2.3	76.3
Overall: 88.0 % correct			

Table 2: Confusion matrix

Thus, in Table 2, B737 is not confused at all with B741, but 10.9 % of the time is mis-classified as B752; B752 is mis-classified as B737 15 % of the time.

The profile used as a template for the Boeing 7300 is shown as the blue line in the following Figure (the red line will be discussed later). The profile is normalised to have a maximum of one, i.e. 0 dB.

.

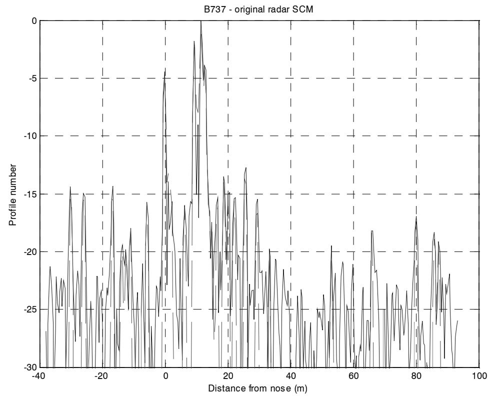

Figure 4: Original template

The horizontal scale on the figure shows the estimated distance from the nose of the aircraft for each point on the profile. There are a number of significant peaks in front of the nose (due to range sidelobes) which indicates that this is perhaps not a particularly good choice; nonetheless, reasonable classification results have been obtained using it.

Radar-Derived SCMs

In this section, we will consider the use of several different SCMs derived directly from the radar data. The aim is to indicate how performance changes as the model is successively simplified; in particular, the effect of discarding amplitude information will be considered.

An SCM may be derived from a range profile over a narrow range of aspect angles by noting the positions and amplitudes of significant peaks in the profile. This technique may be extended to cover a large range of aspect angles by using the type of tomographic analysis described in [2].

Once the SCM has been constructed, a range profile may be re-constructed from it taking into account the range resolution and sampling rate of the radar used to measure profiles.

The result of this process is illustrated in Figure 4 by the red line. Here, all peaks greater than –20 dB have been included in the SCM. The results obtained by approximating the range profile in this way are shown below.

	B737	B741	B752
B737	92.1	0.0	15.0
B741	0.0	96.5	1.3
B752	7.9	3.5	83.8
Overall: 91.0 % correct			

Table 3: Initial SCM results

There is a slight, but probably insignificant, increase in performance.

Next, we consider an SCM with a much smaller number of peaks chosen to coincide with knowledge of the geometry of each aircraft – we know that there should be no peaks ahead of the nose, and the length of each aircraft is known. A B737 template derived from the resulting SCM is shown below.

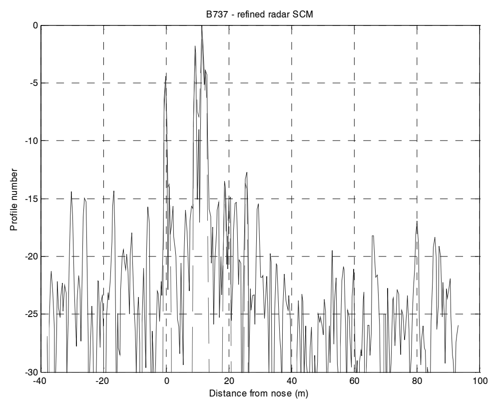

Figure 5: Refined SCM template

The results of using this refined template are shown in Table 4.

	B737	B741	B752
B737	90.1	0.0	13.8
B741	0.0	96.5	1.3
B752	9.9	3.5	85.0
Overall: 90.6 % correct			

Table 4: Refined SCM results

There is little appreciable change in performance.

Finally, we consider the effect of discarding amplitudes from the SCMs. This is achieved by replacing the amplitude of each peak by 0 dB. We refer to this type of SCM as a *reduced* SCM; a B737 template derived from this type of SCM is shown below.

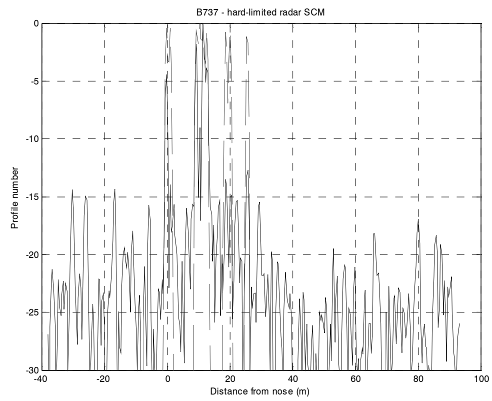

Figure 6: Reduced SCM template

A reduced SCM simply indicates the expected positions of peaks in the profile.

To perform classification using templates derived from reduced SCMs, both the templates and profiles are reduced to binary vectors by applying an appropriate threshold then hard-limiting. The binary vectors indicate where significant returns are found or are expected.

The results of using the reduced SCM are given in Table 5.

	B737	B741	B752
B737	97.0	0.0	17.5
B741	0.0	82.6	0.0
B752	3.0	17.4	82.5
Overall: 88.0 % correct			

Table 5: Reduced SCM results

This demonstrates that, with this particular data, the classification achieved using position-only information is similar to that achieved using full profile information. If this result can be replicated using a wider range of data, it is of significant importance in considering how radar range profiles may be classified.

Non Radar-Derived SCMs

A major goal of the current work is to consider how well aircraft can be classified without using radar measurements as reference data in the classifier. The use of SCMs constructed without using radar data are now considered.

Initially, an SCM was constructed by noting the likely positions of significant scattering centres as illustrated in Figure 1 given in the introduction. SCMs were constructed for each aircraft type; the amplitudes of all scattering centres were set to 0 dB. Classification was performed using the same approach as that used for the reduced SCMs discussed above. The initial results, given below, are unimpressive.

	B737	B741	B752
B737	70.3	20.9	23.8
B741	0.0	7.0	31.3
B752	29.7	72.1	45.0
Overall: 42.3 % correct			

Table 6: Non-radar SCM results using full models

It is particularly noticeable that the B741, which is much longer than the other two aircraft, is almost never classified correctly.

It was noted that returns from the tails of the aircraft (see Figure 2) do not appear consistently, and are thus a rather weak feature of the range profiles. This may reflect the geometry of the radar measurements, in which the radar looks up towards approaching aircraft, so that much of the tail may be obscured by the body of the aircraft. It was therefore possible that tail returns were being over-emphasised in the model. To test this, tail returns were removed from each of the SCMs – this is very easy to do, and demonstrates one aspect of the versatility of using SCMs. Significantly improved results were obtained as follows.

	B737	B741	B752
B737	66.0	3.5	0.0
B741	1.0	83.7	45.0
B752	33.7	12.8	55.0
Overall: 68.2 % correct			

Table 7: Non-radar SCM results with tail removed from models

These results indicate that even rather crude models of radar backscatter can give a useful degree of classification performance. The earlier results given in Table 4 also indicate that there is significant room for improvement in our development of non-radar SCMs.

Future Work

We have demonstrated that scattering centre models may be used to classify measured aircraft profiles, and have given results for SCMs derived using both radar and non-radar data.

As is always the case with classification studies, the results are expected to be highly dependent on the data from which they are obtained; it is therefore a primary concern to apply these techniques to a broader class of radar data, including different types of aircraft to those considered here, and in particular, military aircraft.

The use of non-radar SCMs poses a spectrum of new challenges for classifiers. The lack of amplitude information changes both the classification and alignment problems, and it is planned to apply a variety of techniques to address the novel issues which arise.

Finally, the construction of adequate scattering centre models from non-radar data is a significant challenge in itself, to which a variety of techniques may be applied. These techniques may be refined by comparison of profiles derived from the models with measured profiles to give an improved understanding of which physical features correspond to significant scattering centres.

References

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

[2] R. Miller, C. Sarno, D. Shephard, "Progress in Radar Recognition of Aircraft Without Using Radar-Derived Databases", Electromagnetic Remote Sensing DTC Conference Proceedings, May 2004.

[3] R. Miller, D. Shephard, M. Newman, "Aspects of NCTR for Near-Future Radar", NATO RTO Proceedings, Symposium on Target Identification and Recognition, Oslo, Norway, October 2004.

Acknowledgement

The work reported in this paper was funded by the Electro-Magnetic Remote Sensing (EMRS) Defence Technology Centre, established by the UK Ministry of Defence and run by a consortium of SELEX Sensors and Airborne Systems, Thales Defence, Roke Manor Research and Filtronic.