Progress in Radar Recognition of Aircraft Without Using Radar-Derived Databases

R. J. Miller, G. C. Sarno, D. J. Shephard
BAE SYSTEMS Advanced Technology Centre
West Hanningfield Road, Great Baddow, Chelmsford, Essex CM2 8HN

Abstract

Most techniques for recognising aircraft using high range resolution radar profiles rely on prior collection of extensive bodies of radar data for targets of interest. This paper considers progress made on an alternative approach which considers what can be achieved without the use of such radar-derived databases. Motivated by earlier work by Rihaczek and Hershkowitz, the current work considers the use of scattering centre models of candidate aircraft to provide features used in classification of measured aircraft profiles. Two particular aspects of this approach are discussed, namely the validation of the scattering centre model itself, and the impact of the approach on classification algorithms.

Keywords: Radar, Non-Co-operative Target Recognition

Introduction

The desirability of identifying aircraft in a battlespace is widely recognised; identification of aircraft serves both to increase situation awareness and to prevent inappropriate engagements. In the event that aircraft of interest do not have a co-operative IFF system, identities need to be determined by some other means – a process referred to as non-co-operative target recognition or NCTR.

Radar has the potential to provide long-range, all-weather, day-night recognition capabilities. There are two main ways in which radar may be used in this context. Firstly, Doppler shifts induced on the radar waveform by moving engine components may be utilised (JEM, or jet-engine modulation) – this is not considered further here. Secondly, high range resolution radar waveforms may be used to obtain profiles of the radar cross section of an aircraft as a function of distance along its length. Use of such high resolution range profiles (HRRPs) is what is considered here.

Background

The use of HRRPs in aircraft recognition has been contemplated for some time, see, e.g., Wehner (1). The most commonly adopted approach is to make detailed radar measurements of relevant aircraft, then to use these measurements as a referent in the recognition process. Many schemes have been proposed for doing this. A relatively simple scheme is to use the measurements to form reference templates which are subsequently compared directly with profiles observed in action in order to achieve the classification. This amounts to a form of nearest-neighbour classifier.

In order to obtain a workable scheme, there are various obstacles which need to be overcome. Two such obstacles are as follows. Firstly, range profiles change rapidly with the aspect at which the aircraft is viewed, so it is necessary to accumulate data at all relevant aspect angles. Secondly, military aircraft may carry many different types of stores, generally carried under the wings; different stores configurations result in differences in the range profile, which must also be accounted for. Both of these difficulties add to the burden of measurements which must be obtained for reference purposes.

Furthermore, it is often the case that it is not possible to make detailed measurements of hostile aircraft of interest because exemplars of such aircraft are not available. Accordingly, a different approach has been considered.

Approach

The current work started by making an appraisal of more recent work of Rihaczek and Hershkowitz (2) (R&H) in this area. The work of these authors is somewhat controversial, but nonetheless seems to have had substantial backing from the U.S. DoD. A detailed report appraising this work has been prepared under the EMRS DTC.

For present purposes, the key aspects of this work are as follows.

Firstly, aircraft recognition is performed on the basis of a feature-based classifier rather than on template matching. The features used are primarily the positions of strong regions of backscatter on the aircraft. Secondly, these features are derived, not from dedicated radar measurements of the aircraft, but from material such as photographs, plans and schematics of the aircraft. It is these aspects of their approach which we have pursued.

R&H also rely heavily on the use of ISAR imagery of aircraft which they use, for example, to separate out backscatter from the wings and backscatter from the fuselage of the aircraft. Their classifiers are based on scattering from the fuselage, and thus avoid many of the difficulties presented by variable stores configurations carried on the wings. However, gathering ISAR imagery often seems to be infeasible from an operational perspective – for example, it is not possible to gather ISAR imagery from an aircraft heading straight towards the radar – so we have not chosen to follow this aspect of their approach. Rather, we have elected to see what can be achieved using a small number of range profiles in conjunction with other aspects of R&H's approach.

In order to perform classification, observed range profiles must still be compared with some form of referent. In R&H's approach, the locations of strong regions of backscatter which persist over appreciable ranges of aspect angle are identified. This may be expressed in terms of a model, which we term a 'scattering centre model'. This model is discussed in more detail in a later section.

There are numerous issues which arise in association with the approach espoused by R&H. We consider two of them in more detail below.

Scattering Centre Model

In our interpretation of R&H's technique, the referents used to determine classification take the form of models describing the principal scattering centres of each type of aircraft.

A scattering centre is described by its location on the aircraft and by how the amplitude of its return varies with aspect angle. Thus, the model of a single scattering centre is similar to the polar diagram of an antenna, which shows its response as a function of angle. The description of a scattering centre may also take into account that it may be occluded by parts of the aircraft structure over some ranges of aspect angle. The phase of the response of the scattering centre may also vary with aspect angle, so that the possibility of its phase centre wandering with aspect may be also taken into account.

The scattering centre model consists of the totality of scattering centres on the aircraft. Nominal range profiles of the aircraft may be determined from the scattering centre model by convolving the returns obtained from a particular direction with a function representing the pulse-shape of the radar. Returns are added in a complex fashion, so interference of returns from scattering centres lying in the same range gate occurs.

Does the scattering centre model provide a reasonable description of backscatter from an aircraft? The model needs to be validated before it can have any credibility as a basis for performing NCTR.

Validation of Scattering Centre Model

In order to demonstrate the validity or otherwise of the scattering centre model, the following approach has been adopted.

If we have a set of measured profiles of an aircraft taken over a closely-spaced range of aspect angles, it is possible to construct a map of the origins of the backscatter using tomographic reconstruction techniques. If this map reveals a small number of well-defined centres of scattering, this will lend weight to the scattering centre model. To effect this reconstruction technique, it is essential that the aspect angles appertaining to each range profile are accurately known.

At the time of writing, we are still in the process of obtaining appropriate data to validate the scattering centre model. However, we are able to illustrate the process of validation as follows.

Detailed range profiles of an aircraft may be simulated using computational electromagnetics codes. The steps in the process are: construct a mesh structure representing the object of interest; use the code to compute the backscattered field as a function of frequency; transform from the frequency to the time domain to obtain a range profile. There are many similarities between range profiles computed in this way and measured range profiles, though there are also significant differences. This is not the issue here; we are simply using the CEM code to generate a set of reasonably realistic profiles based on rigorous scattering theory.

A full-wave finite-difference time-domain code was used to calculate profiles from a crude mesh model of a small aircraft; profiles were calculated at zero degrees elevation and over azimuth angles of 0 to 180 degrees in 1 degree increments. A subset of these profiles is shown as a range-azimuth map in Figure 1. Arcs in this map cover partial sinusoids and are indicative of returns from well-defined scattering centres.

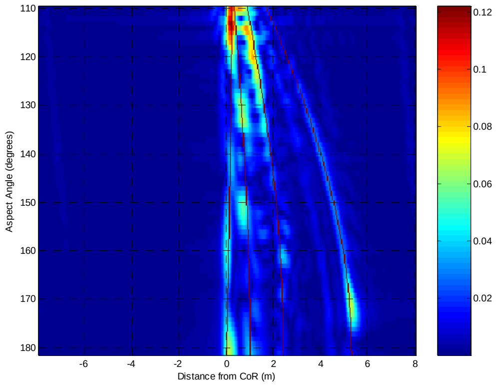

Figure 1: Profiles from small aircraft

The data shown were processed using a tomographic reconstruction algorithm (3). The essence of such algorithms starts from the recognition that each point on the target generates a sinusoidal arc with a distinct amplitude and phase; summing the returns over each arc gives the total return from each point. In this way, it is possible to construct a reflectivity map indicating the return from any point on the target. A tomogram is generated by calculating the returns over a grid of points. The result is shown in Figure 2. Four well-defined peaks are evident, lending credence to the scattering centre model.

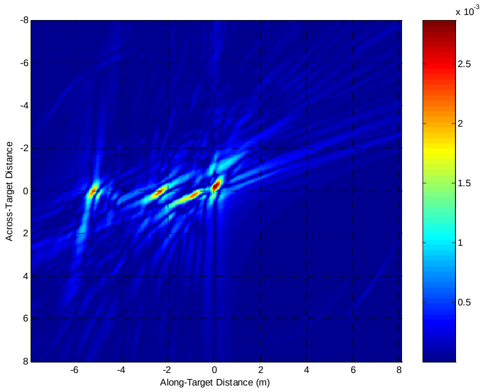

Figure 2: Tomogram from small aircraft

The co-ordinate system for the tomogram has its origin at the centre of rotation of the target. In this case, the centre of rotation is the tail of the aircraft. Hence, locations ahead of the tail have negative along-target distance. The points of high reflectivity shown in Figure 2 correspond to the nose, wing, tail-planes and the rear of the aircraft. The scattering centre model posits that, not only do returns arise from a small number of well-defined scattering centres, but also these returns are high over particular ranges of azimuthal aspect angles. To investigate this, it is possible to trace through those sinusoidal arcs corresponding to the positions of the scattering centres, and to show the reflectivity as a function of aspect angle for each point.

The sinusoidal arcs for each of the prominent scattering centres shown in Figure 2 are shown as dark lines in Figure 1. Reflectivity as a function of aspect angle for each point is shown in Figure 3.

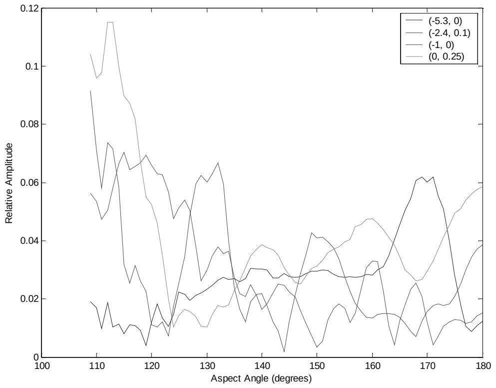

Figure 3: Angular response of scattering centres for small aircraft

Figure 3 demonstrates that the pattern of returns from a single scattering centre is generally rather more complicated than a single peak. The impact of this kind of behaviour on classification algorithms will need to be given close consideration. The next step will be to apply the validation techniques outlined above to radar measurements rather than to profiles obtained from CEM codes.

Classification Algorithms

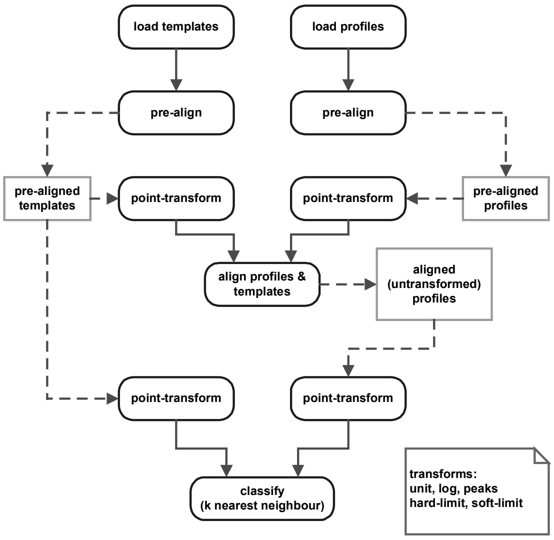

Figure 4: Alignment and classification

A set of comparatively simple classification algorithms has been implemented, preceded by alignment algorithms. The aim of what is currently implemented has not so much been to determine the best classification that can be achieved, but rather to gauge the effects of using limited knowledge about aircraft signatures to inform the classification process.

The classification process starts by generating a set of templates from scattering centre models for various aircraft. These templates take the form of idealised, noise-free range profiles; they may subsequently be processed to extract features such as peaks and their locations. Templates are generated over a range of aspect angles for each aircraft being considered.

The templates are subsequently compared with measured profiles to yield a classification of the profiles. In what is considered here, an attempt is made to classify each profile individually; this allows some estimate to be made of the probability of achieving correct classification from the small sets of data currently available to us. In future work, it is envisaged that classification will be performed using small sets of profiles rather than single profiles.

Before any comparison between template and profile can be made, they must be aligned. For each profile, a number of duplicates of the profile is generated to correspond with the number of templates. Each duplicate profile is then shifted to give a best match with its corresponding template over all possible shifts. Alignment is achieved using one of a number of methods, examples being correlation and minimising the Euclidean distance between profiles. Prior to alignment, both templates and profiles may be subjected to a point transformation. One of a variety of transformations may be used; a simple transformation would be a logarithmic one. However, more complex transformations may be employed, such as a peak detection followed by a threshold selection, which allows only major peaks in templates and profiles to be considered.

Once alignment of templates and profiles has been achieved, a classifier is used to determine the class of the profile from the aligned templates. One of a variety of mis-match measures is used to determine the mis-match between each template and the profile. These mis-match measures are used by a k-nearest neighbours classifier to determine the class. Essentially, the k-nearest neighbour algorithm considers all templates within a defined neighbourhood of the profile, then classifies according to which type of aircraft the majority of templates in the neighbourhood derive from.

As for alignment, prior to classification, both templates and profiles are subjected to a point transformation. Point transformations may be different to those adopted for alignment. Point transformations prior to classification are of particular importance, since different transformations allow different kinds of comparisons between profile and template to be considered. Examples are given below.

First of all, consider the use of a unit transformation, so that both profile and template are unchanged. This implies that the comparison is between the full profile and template, so that the classifier compares the amplitude of the template with that of the profile for each point along the profile. No attempt is made to extract features from the profile.

Secondly, consider the use of a peak detection algorithm; by this, we mean an algorithm which detects local maxima then forms a vector whose elements are non-zero only at the positions of the local maxima; the value of each of these non-zero elements is that of the amplitude of the corresponding peak. If peak detection is followed by a soft threshold to eliminate peaks below a specified amplitude, the derived vector describes the positions and amplitudes of the major peaks in the profile or template; hence, particular features have been extracted, and comparison and classification is now performed on the basis of these features alone.

Finally, consider the use of peak detection followed by a hard threshold, so that the derived vector consists of ones and zeroes, with the ones at locations in the vector corresponding to the positions of major peaks. This form extracts features (peaks), but retains only their locations; thus, classification is performed only on the basis of positions of major peaks – it does not take into account amplitude information. This latter form is particularly important for current work, since, while it is feasible that positions of scattering centres may be located from, e.g. photographs of aircraft, good estimates of the RCS of these scattering centres will be much harder to obtain.

Some preliminary experiments using data with 30 cm range resolution have indicated that the losses in classification accuracy due to using only the positions of peaks is small as compared to using the full profile. However, only very limited data has been considered, so these results cannot in any way be thought of as being definitive.

Conclusion

Current work has shown that it is at least feasible to consider classification of aircraft using data derived from non-radar sources.

The key approach adopted here is the use of a scattering centre model from which prominent scattering features on each candidate aircraft are derived. These features are then compared with features derived from measured profiles in order to classify the radar target.

Some progress has been made in validating the scattering centre model, though to date it has only been possible to use data derived from computational electromagnetic calculations rather than controlled radar measurements. The results given here support the notion of a model consisting of a small number of scattering centres, but indicate that the variation of backscatter with aspect angle from these scattering centres may be quite complicated.

Preliminary work on classification of aircraft using limited knowledge of the likely backscatter from aircraft, namely the positions of peaks but not their amplitudes, indicates that the approach is viable, but further work needs to be done to confirm this.

Future Work

Future work envisaged is as follows:

- validate the scattering centre model using controlled radar measurements obtained from a real aircraft,
- verify that a respectable degree of classification performance can be achieved by exercising classification algorithms based on non-radar data on a greater variety of measured aircraft range profiles,
- consider robustness of approach to giving very reliable discrimination between broad classes of aircraft, e.g. military/civil,
- incorporate the classification scheme into an on-site radar facility to gauge performance in real-world conditions.

References

1. Wehner, D. R. 1987. High Resolution Radar, Artech House
2. Rihaczek, A. W. and Hershkowitz, S. J. 2000. Theory and Practice of Radar Target Identification, Artech House
3. The MathWorks, Inc., Image Processing Toolbox.