

Correlation Filters for Aircraft Identification From Radar Range Profiles

SCOTT HUDSON

Washington State University

DEMETRI PSALTIS

California Institute of Technology

The potential for identifying aircraft using one or more radar range profiles, in conjunction with a correlator, is investigated. Two types of filters are described which maximize the expected value of certain correlation peaks. The effectiveness of one type of filter was investigated in identification experiments using an extensive data set of real radar range profiles of 24 different aircraft. The results suggest that reliable identification is possible provided aircraft aspect information is used and identifications are based on multiple profiles.

Manuscript received April 8, 1991; revised April 14, 1992.

IEEE Log No. T-AES/29/3/07978.

This work was supported, in part, by the Litton Corporation and the Rome Air Development Center.

Authors' addresses: S. Hudson, Electrical Engineering and Computer Science, Washington State University, Pullman, WA 99164-2752; D. Psaltis, Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125.

0018-9251/93/\$3.00 © 1993 IEEE

I. INTRODUCTION

Over the last half century radar has proven to be a very powerful tool for the detection, ranging, and tracking of aircraft. However, these capabilities leave a key question unanswered: What has been detected? This is often a crucial piece of information as it determines the appropriate response. Ideally radar systems should have the ability to provide a reliable identification of every aircraft they detect and/or track.

The most complete approach to aircraft identification would be to solve the inverse scattering problem [1] and fully characterize the electromagnetic properties of the aircraft. However, this is a formidable undertaking, and we usually are forced to settle for a partial characterization. Fortunately, we typically know *a priori* that the aircraft cannot be an arbitrary object, but that it must represent one of a finite number of possibilities, e.g., an existing model of aircraft. In such cases a partial characterization may suffice to specify the the aircraft identity from among the possibilities, in much the same way that a person can be identified by his or her fingerprints even though fingerprints do not fully characterize a person.

There is a variety of ways to partially characterize an aircraft using radar. A considerable amount of research has centered on identifying aspect-independent scattering resonances and using them as features for identification purposes [2]. In practice it is found that such techniques are very sensitive to noise [3]. Two-dimensional radar images of aircraft can be formed using the techniques of inverse synthetic aperture radar [4]. Drawbacks are that precise motion compensation is required to form well-focused images, the signal processing requirements can be significant, and a considerable dwell time may be required [5-7]. A very simple and rapid way to characterize an aircraft is through the use of radar range profiles, essentially one-dimensional radar images.

We investigate the use of radar range profiles as aircraft "fingerprints" for identification purposes. We concentrate on a correlator-based system. While a correlator is only one of many ways to approach the identification problem, it is attractive because it can take advantage of high-speed/low-power implementations such as those afforded by optics [8]. However, for a correlator-based system to achieve its maximum effectiveness care must be exercised in designing the correlation filters, particularly in the limit of short radar wavelengths where range profiles can fluctuate greatly with small aspect changes. Accordingly, we devote Section III to the development of techniques for producing correlation filters. This is followed in Section IV with identification experiments using real radar range profiles. We begin with a brief description of radar range profiles.




Fig. 1. Relation between radar range profile and aircraft.




Fig. 2. Small aspect change $\Delta\theta$ causes relative range change of $W\Delta\theta$ for points separated by W .

II. RADAR RANGE PROFILES

If an aircraft is illuminated by a short radar pulse, and the backscattered power is recorded as a function of time, the result is a radar range profile. A radar range profile can be thought of as the one-dimensional projection of the spatial distribution of radar reflectivity of the aircraft onto the radar line of sight. This is illustrated in Fig. 1. At a given time the pulse illuminates only a single strip, or *range bin*, of the aircraft, consisting of all points at a certain range, or distance, from the radar. The backscattered field received at a given time consists of the sum of scatter from only the points of the aircraft in a corresponding range bin. The resulting radar range profile thus resolves the points of the aircraft in range to within the range resolution Δr of the radar.

If the *aspect* of the aircraft (the angle its nose makes with the radar line of sight) changes then the relative ranges of points may change. The result may be *range migration*, i.e., that two points which were once in the same range bin are no longer. Consider the case shown in Fig. 2. If the wingtips are originally at the same range, and are a distance W apart, a small aspect change of $\Delta\theta$ will cause a relative range change of approximately $W\Delta\theta$ between them. If this is larger than the range resolution of the radar, i.e., if $W\Delta\theta > \Delta r$, then the wingtips will no longer fall into the same range bin. This places a limit

$$\Delta\theta < \frac{\Delta r}{W} \quad (1)$$

on the range of aspects for which the radar profile can be considered the same projection of the aircraft. For example, if $\Delta r = 0.5$ m and $W = 10$ m, we find $\Delta\theta < 3^\circ$.

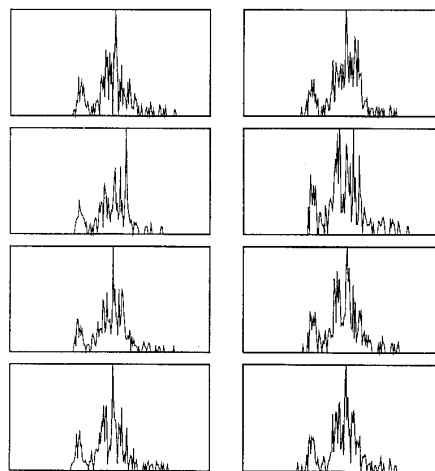


Fig. 3. Example of speckle fluctuations in profiles of flying aircraft collected 2.5 ms apart.

Even if the aspect does not change enough to cause range migration, changes in the relative range of different scatterers of a fraction of the radar wavelength can cause the radar profile to fluctuate due to interference effects. If $W\Delta\theta \approx \lambda/4$, where λ is the radar wavelength, then the scatter from the wingtips will change from constructively interfering to destructively interfering, or *vice versa*, and the range profile will fluctuate significantly. This is the phenomenon of *speckle* [9] that plagues all coherent imaging systems. To avoid speckle fluctuations the aspect would have to stay constant to within

$$\Delta\theta \ll \frac{\lambda}{4W}. \quad (2)$$

When using short wavelengths, we do not expect the aspect of a flying aircraft to stay constant to within the limit given in (2) for any significant period, so aircraft profiles can be expected to display speckle fluctuations as time passes. (Note that additional fluctuations may arise if the aircraft is not a rigid body, e.g., it has exposed moving parts such as propellers. Also, all real observations will contain fluctuations due to system noise.) An example of speckle fluctuations, at S-band, in the profiles of a flying aircraft is shown in Fig. 3. Even though the eight profiles were collected only 2.5 ms apart, there are considerable variations from profile to profile. These variations are more than an order of magnitude larger than the system noise.

The fact that very small aspect changes can drastically alter the profile of an aircraft implies that a huge number of profiles would be required to characterize an aircraft from all possible aspects. This, and the infeasibility of estimating a flying aircraft's aspect to the precision required by (2), make a "look-up table" approach to identification impractical in the short-wavelength limit. That is, in contrast with the long-wavelength case [10, 11], it is not practical

to record a profile of an unknown aircraft, use an estimate of the aspect of an aircraft as an index to a library of aircraft profiles, pull out the profile for each aircraft that has the given aspect, and look for a match.

III. CORRELATION FILTERS

Speckle fluctuations complicate the use of short-wavelength range profiles for aircraft identification. However, for a given aircraft over a limited range of aspects, i.e., (1), the distribution of scattering centers in range will be similar. Hence, the corresponding profiles should display common features, although the amplitudes of the features will fluctuate from profile to profile. At very different aspects the profiles will generally display a different set of features due to the different distribution of scattering centers. Accordingly, it is useful to consider a given aircraft over a given range of aspects as a distinct "target." For example, aircraft "A" over aspects $0 \leq \theta \leq \Delta\theta$ could be considered target "A1," the same aircraft over aspects $\Delta\theta \leq \theta \leq 2\Delta\theta$ could be considered target "A2," etc. Aircraft "B" over aspects $0 \leq \theta \leq \Delta\theta$ could be considered target "B1," and so on. Because the profiles of a given target share common features, it might be possible to construct a filter for each target which correlates better, on average, with the profiles of that target than with the profiles of other targets. A profile from an unknown target could then be correlated against a library of such filters for different known targets. The filter which produced the greatest correlation would provide an identification for the unknown target.

The first filter design approach we consider is to concentrate solely on maximizing the correlation of a filter with the profiles of a given target. We refer to these as "type I" filters. In this case the goal is to make the filter "look like" the profiles of that target as much as possible. We then know that any radar profiles that correlate strongly with this filter look very much like those of the target in question.

A second approach is to form a filter which simultaneously maximizes its correlation with the profiles of a given target and minimizes its correlation with the profiles of all other targets. The result is a filter which emphasizes the differences between the profiles of the given target and those of other targets, and we refer to these as "type II" filters. This approach is potentially more powerful, but it is also much more computationally intensive. Although we develop expressions for both types of filters, we implemented only type I filters in the identification experiments described below.

The situation for type I filters is that for a given target we want to compute the filter $\mathbf{f} = [f(0), f(1), \dots, f(K)]$ which maximizes

$$\Phi_I = E[\text{peak}(\mathbf{f} \star \mathbf{p})] \quad (3)$$

the expected value of the correlation peak between it and $\mathbf{p} = [p(0), p(1), \dots, p(K)]$, where \mathbf{p} is any randomly chosen profile of the target, and the expected value is over all possible profiles. The correlation between the real-valued vectors \mathbf{f} and \mathbf{p} is given by

$$(\mathbf{f} \star \mathbf{p})(s) = \sum_{j=0}^{K-1} f(j)p(j-s) \quad (4)$$

for $0 \leq s \leq K-1$. K is the number of range bins, and we assume the correlation is circular, i.e., $p[j \pm K] = p[j]$.

The goal is to create a library of filters, $\mathbf{f}^{(1)}, \mathbf{f}^{(2)}, \dots, \mathbf{f}^{(M)}$, with M the number of targets to be identified. If a profile from an unknown target is correlated against this library producing numbers

$$\alpha^{(k)} = \text{peak}(\mathbf{f}^{(k)} \star \mathbf{p}) \quad (5)$$

(correlation peaks) the hope is that

$$\alpha^{(j)} > \alpha^{(k)} \quad \text{for all } k \neq j \quad (6)$$

where the j th filter is the "correct" one, i.e., the one which corresponds to the true identity of the unknown target.

In practice we have N profiles recorded for a given target, and our task is to use these to compute the filter which maximizes the correlation peaks with future, randomly drawn profiles from the same target. Towards this end we form the estimate

$$\Phi_I = \frac{1}{N} \sum_{n=1}^N \text{peak}(\mathbf{f} \star \mathbf{p}_n) \quad (7)$$

where $\{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N\}$ is the set of recorded radar profiles for this particular target. Now, assume we already know what \mathbf{f} is, and the peak value of $(\mathbf{f} \star \mathbf{p}_n)(s)$ occurs at $s = s_n$. Then we could form \mathbf{p}'_n , a shifted version of \mathbf{p}_n ,

$$p'_n(k) = p_n(k - s_n) \quad (8)$$

so that

$$\begin{aligned} \text{peak}(\mathbf{f} \star \mathbf{p}_n) &= (\mathbf{f} \star \mathbf{p}_n)(s_n) \\ &= \sum_j f(j)p_n(j - s_n) \\ &= \sum_j f(j)p'_n(j) \\ &= \mathbf{f} \cdot \mathbf{p}'_n \end{aligned} \quad (9)$$

(Since we assume circular correlations, these shifts should also be circular; any data shifted off one end must be shifted onto the other. If desired, the profiles can always be zero padded to avoid such edge effects.) In other words, we could "align" \mathbf{p}_n with the filter so

that the correlation peak occurs at zero shift. Then

$$\begin{aligned}\Phi_I &= \frac{1}{N} \sum_{n=1}^N \mathbf{f} \cdot \mathbf{p}'_n \\ &= \mathbf{f} \cdot \left(\frac{1}{N} \sum_{n=1}^N \mathbf{p}'_n \right).\end{aligned}\quad (10)$$

This dot product is the quantity that the filter \mathbf{f} must maximize.

The magnitude of \mathbf{f} is unimportant for identification purposes since changing it merely scales its correlation with all vectors, so there is no loss in generality if we constrain \mathbf{f} to be a unit vector. The unit vector \mathbf{f} which maximizes (10) is simply

$$\mathbf{f} = \frac{\sum_{n=1}^N \mathbf{p}'_n}{\left\| \sum_{n=1}^N \mathbf{p}'_n \right\|}\quad (11)$$

and so

$$\begin{aligned}\Phi_I &= \frac{\sum_{n=1}^N \mathbf{p}'_n}{\left\| \sum_{n=1}^N \mathbf{p}'_n \right\|} \cdot \left(\frac{1}{N} \sum_{n=1}^N \mathbf{p}'_n \right) \\ &= \frac{1}{N} \left\| \sum_{n=1}^N \mathbf{p}'_n \right\|.\end{aligned}\quad (12)$$

While we don't know \mathbf{f} *a priori*, the above analysis makes it clear that whatever it is, Φ_I is given by (12), and the filter that achieves this is given by (11). Therefore, to maximize $E[\text{peak}(\mathbf{f} \star \mathbf{p})]$ we need only find the shifts s_n which maximize (12) and then evaluate (11) to compute the desired filter.

The practical difficulty is how to find the set of shifts $\{s_1, s_2, \dots, s_N\}$ that take the set $\{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N\}$ into the set $\{\mathbf{p}'_1, \mathbf{p}'_2, \dots, \mathbf{p}'_N\}$ that gives the filter via (11). A global search over all possible shifts is out of the question. For example, suppose $N = 100$, and we need to consider 10 shifted versions of each profile, then there are 10^{100} different configurations $\{s_1, s_2, \dots, s_N\}$ possible.

In the experiments described here, we have used an iterative procedure for maximizing (12) in which a global search is performed over only the shifts of one profile at a time. To see how this can be carried out, we write

$$\begin{aligned}\left\| \sum_{n=1}^N \mathbf{p}'_k \right\|^2 &= \left\| \mathbf{p}'_j + \sum_{n \neq j} \mathbf{p}'_n \right\|^2 \\ &= \|\mathbf{p}'_j\|^2 + \left\| \sum_{n \neq j} \mathbf{p}'_n \right\|^2 + 2\mathbf{p}'_j \cdot \sum_{n \neq j} \mathbf{p}'_n\end{aligned}\quad (13)$$

where \mathbf{p}'_j is any one of the profiles. The only term that will change when \mathbf{p}'_j is shifted is $2\mathbf{p}'_j \cdot \sum_{k \neq j} \mathbf{p}'_k$, and increasing this will increase the quantity we are trying to maximize. Therefore, with respect to the shift s_j , we are doing the best we can if we choose it to

make $\mathbf{p}'_j \cdot \sum_{k \neq j} \mathbf{p}'_k$ as big as possible. A global search over s_j can be done efficiently using the fast Fourier transform. Having done this for one profile we can go on to do it for the others and then start the process over. Each shift increases $\left\| \sum_{k=1}^N \mathbf{p}'_k \right\|$, and eventually we will reach a maximum.

Note that the result is not guaranteed to be a global maximum. The procedure is similar to iterative line maximization techniques used to find extrema of functions of several variables [12], and it may potentially suffer the same problems of local maxima. However, in the applications described below, we investigated this problem by starting with different initial profile configurations, and we found that the algorithm converged to the same result.

This same reasoning can be applied to type II filters. Let \mathbf{p}_n , $1 \leq n \leq N_p$, denote the profiles of a specific target, and let \mathbf{q}_n , $1 \leq n \leq N_q$, denote the profiles of all other targets. We seek the filter \mathbf{f} which maximizes

$$\Phi_{II} = E[\text{peak}(\mathbf{f} \star \mathbf{p}) - \text{peak}(\mathbf{f} \star \mathbf{q})]\quad (14)$$

where \mathbf{p} and \mathbf{q} are randomly drawn, and the expected value is over all possible profiles. For a given data set we estimate this using

$$\Phi_{II} = \frac{1}{N_p} \sum_{n=1}^{N_p} \text{peak}(\mathbf{f} \star \mathbf{p}_n) - \frac{1}{N_q} \sum_{n=1}^{N_q} \text{peak}(\mathbf{f} \star \mathbf{q}_n).\quad (15)$$

If we knew what \mathbf{f} was we could shift the vectors \mathbf{p}_n and \mathbf{q}_n into \mathbf{p}'_n and \mathbf{q}'_n such that

$$\begin{aligned}\Phi_{II} &= \frac{1}{N_p} \sum_{n=1}^{N_p} \mathbf{f} \cdot \mathbf{p}'_n - \frac{1}{N_q} \sum_{n=1}^{N_q} \mathbf{f} \cdot \mathbf{q}'_n \\ &= \mathbf{f} \cdot \left(\frac{1}{N_p} \sum_{n=1}^{N_p} \mathbf{p}'_n - \frac{1}{N_q} \sum_{n=1}^{N_q} \mathbf{q}'_n \right).\end{aligned}\quad (16)$$

This expression is analogous to (10), and by the same reasoning followed there, the desired unit-magnitude filter is

$$\mathbf{f} = \frac{\frac{1}{N_p} \sum_{n=1}^{N_p} \mathbf{p}'_n - \frac{1}{N_q} \sum_{n=1}^{N_q} \mathbf{q}'_n}{\left\| \frac{1}{N_p} \sum_{n=1}^{N_p} \mathbf{p}'_n - \frac{1}{N_q} \sum_{n=1}^{N_q} \mathbf{q}'_n \right\|}\quad (17)$$

where \mathbf{p}'_n and \mathbf{q}'_n are the shifted versions of \mathbf{p}_n and \mathbf{q}_n that maximize

$$\Phi_{II} = \left\| \frac{1}{N_p} \sum_{n=1}^{N_p} \mathbf{p}'_n - \frac{1}{N_q} \sum_{n=1}^{N_q} \mathbf{q}'_n \right\|.\quad (18)$$

Note that although Φ_I and Φ_{II} have analogous forms, maximizing Φ_{II} is much more computationally intensive as N_q is typically much larger than N_p . For example, in the data set described below we




Fig. 4. Structure of single encounter of RADC data set.

typically have $N_p \approx 50$ while $N_q \approx 6000$. To keep the computations reasonable we limited ourselves to type I filters for our experiments.

IV. CORRELATION EXPERIMENTS

The data set used in our experiments consists of 11,968 S-band radar profiles provided by the Air Force's Rome Air Development Center (RADC). The profiles have 200 range bins with a range resolution of approximately 0.5 m. They are grouped into sets of about 100 profiles each called an *encounter* (Fig. 4). Each encounter arose from a several-second series of observations of a single aircraft in flight. There are 119 encounters in all, collected on various dates, and representing 24 distinct aircraft types at various aspects. Within each encounter the profiles are grouped into *frames* of eight profiles each. The profiles within a frame were collected 2.5 ms apart. An average of 12 frames were collected during each encounter. The frames are separated by roughly 2 s of dead time. The aircraft aspects were estimated to within roughly 20° .

We had little technical information available to us concerning the absolute calibration of the radar from date to date. Without precise calibration data the absolute scale of a profile is meaningless. Therefore, we removed all scale considerations by normalizing each profile to have unit magnitude. That is, each profile $p(k)$, $1 \leq k \leq 200$, where k indexes the range bins, was scaled so that

$$\sum_{k=1}^{200} p^2(k) = 1. \quad (19)$$

As a result we are concerned only with the profile shapes, and we have discarded information about the absolute radar cross sections of the aircraft.

Ideally there would be two separate encounters for each target, and one would go into a training set to form filters, and the other into a test set for identification experiments. Unfortunately, in many cases only a single encounter was available for a given




Fig. 5. Examples of profiles assigned to training and test sets. Profiles on left went into training set, profiles on right went into test set.

target. What we did was to divide the 11,968 profiles into a training set of 6,256 profiles and a test set of 5712 profiles, by assigning the 8 profiles of each odd numbered frame to the training set and the even numbered frames to the test set.

Given its limitations, we believe this was the best way to use this data set to approximate how an operational system might function. We would want to collect training profiles over every range of aspects, in the sense of (1), for which we would later seek to identify an aircraft, but we would not expect to sample exactly the same aspects, in the sense of (2), as we would view in the future. The relatively long time between frames (~ 2 s) allowed the aircraft to move enough to give us distinct aspects, in the sense of (2), in the training and test sets (Fig. 5). At the same time, hopefully the total aspect change during the encounter did not seriously violate (1). If it did then the resulting filter would be essentially be trying to simultaneously represent two different projections of the aircraft and would probably fail to well represent either. Unfortunately the aspects estimates are too coarse for us to be sure of this.

With the profiles from each encounter suitably divided into training and test sets, the (roughly 50) training profiles of each encounter were used with the algorithm of Section III to produce a type I filter for that encounter. The result was a library of 119 filters, one filter for each encounter. Thirteen of these filters, each representing a different aircraft at nominally nose-on aspect, are shown in Fig. 6. The solid curves are the filters while the dashed curves show the rms deviation about the filter of the profiles of that encounter and so give an indication of the amplitude of the speckle fluctuations.

The filter library was used in a series of identification experiments. Each experiment consisted

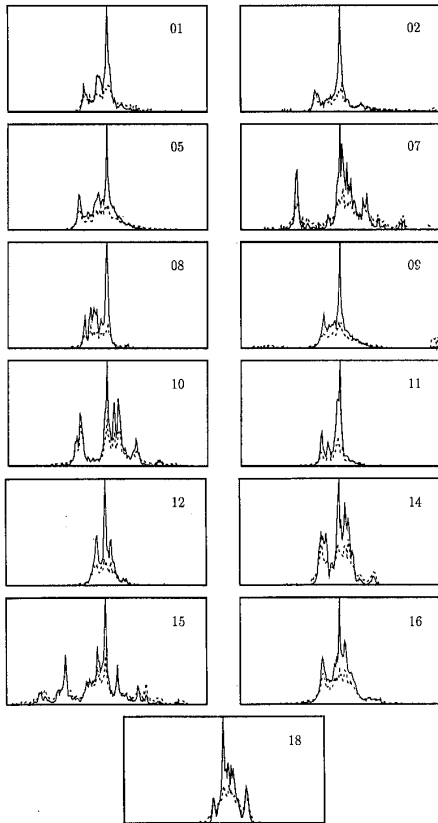


Fig. 6. Filters for thirteen aircraft at nose-on aspects. Solid curves are filters, dashed curves show rms deviation of training profiles about the filters.

of correlating profiles from the test set against each of the 119 filters, noting the correlation peaks, and recording the largest. These results were used to identify the aircraft that produced the profile(s). The identifications were compared with the known identities, and statistics were kept on the rate of correct identifications.

The first experiment was to feed each of the 5712 profiles of the test set through the correlator and identify each according to the largest correlation peak, without regard to aspect. The result was that 3262 of the 5712 profiles were correctly identified for an overall correct-identification rate of 57%. Not all targets were correctly identified at the same rate. A few were correctly identified 100% of the time while others were correctly identified less than 10% of the time. This is not surprising since, as Fig. 6 indicates, many targets in this data set have similar profiles and hence speckle fluctuations are bound to cause confusion. At the same time, a few of the filters are very distinctive, indicating that even with speckle fluctuations the profiles of the corresponding target will still be clearly identifiable. While a correct-identification rate of 57% is certainly




Fig. 7. Similarities between correlation filters of different aircraft at different aspects. Left filter is for aircraft at nose-on aspect, right filter is for different aircraft at broadside aspects.

an improvement over the $\frac{1}{24} = 4\%$ rate that would be achieved by random guessing (there are 24 aircraft in the data set), it is not high enough to suggest that this approach would be of much practical use.

Interestingly, many of the identification errors involve an aircraft at nose-on (0° – 20°) being mistaken for a different aircraft at broadside (80° – 100°), and *vice versa*. Fig. 7 shows two filters, the one at left from an aircraft at nose-on, the one at right from a different aircraft at broadside. Considering the similarity of the filters, it is easy to see how speckle fluctuations could cause the profiles from these aircraft, at these aspects, to be mistaken for one another.

The existence of similar profiles from different aircraft at different aspects suggests that incorporating aspect information into the identification process would improve things. Therefore, in our next experiment we took the estimated aspect of the aircraft into account; profiles of the test set were correlated against only those filters which corresponded to the same estimated 20° aspect range. The results, broken down by aspects, are given in Table I under the column labeled "profile." Not surprisingly, the identification rates are almost all better than the overall 57% achieved without using aspect estimates. However, several of the aspect ranges contain only a few different aircraft, so the corresponding correct identification rates probably do not have much statistical significance. Most of the profiles are at nose-on or tail-on. For this reason, for brevity, and because of the special importance of nose-on, i.e., approaching, aircraft, we limit the discussion below to nose-on aspects.

Excluding data from different estimated aspects, the test profiles of the 33 encounters at nose-on aspect, representing 13 distinct aircraft types, were correctly identified at an average rate of 79%. Given the relatively minor differences in the filters of different aircraft, our only hope for significantly improving the correct-identification rate is to somehow reduce the effects of speckle.

To combat speckle, we need to perform some sort of averaging operation, and this will require the use of multiple profiles. One possibility is to average several profiles using an alignment technique like that of Section III and then put the result through the correlator. More straightforward is a compound identification approach where we put the series of profiles through the correlator, identify each, and

TABLE I
Summary of Results of Aircraft Identification Experiments

aspect	correct identification rates			
	random	profile	frame	encounter
0° - 20°	$\frac{1}{13} = 8\%$	$\frac{1256}{1584} = 79\%$	$\frac{167}{198} = 84\%$	$\frac{33}{33} = 100\%$
20° - 40°	$\frac{1}{6} = 17\%$	$\frac{354}{424} = 83\%$	$\frac{50}{53} = 94\%$	$\frac{9}{9} = 100\%$
40° - 60°	$\frac{1}{2} = 50\%$	$\frac{88}{88} = 100\%$	$\frac{11}{11} = 100\%$	$\frac{3}{3} = 100\%$
60° - 80°	$\frac{1}{2} = 50\%$	$\frac{96}{96} = 100\%$	$\frac{12}{12} = 100\%$	$\frac{3}{3} = 100\%$
80° - 100°	$\frac{1}{3} = 33\%$	$\frac{183}{216} = 85\%$	$\frac{22}{27} = 81\%$	$\frac{4}{4} = 100\%$
100° - 120°	$\frac{1}{4} = 25\%$	$\frac{109}{248} = 44\%$	$\frac{15}{31} = 48\%$	$\frac{4}{5} = 80\%$
120° - 140°	$\frac{1}{8} = 13\%$	$\frac{423}{704} = 60\%$	$\frac{55}{88} = 63\%$	$\frac{12}{15} = 80\%$
140° - 160°	$\frac{1}{4} = 25\%$	$\frac{212}{280} = 76\%$	$\frac{28}{35} = 80\%$	$\frac{5}{5} = 100\%$
160° - 180°	$\frac{1}{17} = 6\%$	$\frac{1468}{2072} = 71\%$	$\frac{204}{259} = 79\%$	$\frac{40}{42} = 95\%$
all	$\frac{1}{24} = 4\%$	$\frac{3262}{5712} = 57\%$	$\frac{465}{714} = 65\%$	$\frac{102}{119} = 86\%$

Note: Left column shows aspect ranges. "All" refers to experiments in which no aspect information was used. Columns under "correct identification rates" show various rates for data at given aspect. "Random" is rate that would be obtained by guessing and is simply one over the number of different aircraft at that aspect represented in the data. Rightmost three columns give percentage rates and show number correctly identified over number tested. "Profile" is rate at which correlator correctly identified individual profiles; "frame" is rate for compound identifications based on frames; "encounter" is compound identifications based on encounters.

average the identifications to arrive at a single "best" identification. That is, we identify the aircraft by observing which filter most often produced the largest correlation peak. The identification of a single profile can be considered as a "vote" for one of the filters, and the compound identification then corresponds to choosing the filter which obtains the most "votes."

Because of its structure (Fig. 4), the natural level at which to perform compound identification with this data set is by frame, or by encounter. Accordingly, we next performed an identification experiment in which each of the eight profiles of a test-set frame were identified by the correlator followed by a compound identification for the frame as a whole. The result for the nose-on data was a correct-identification rate of 84%. This is a modest improvement over the 79% rate for single profiles, suggesting that the advantage of performing identification using a short burst of profiles, as opposed to using a single profile, is limited. The results for all aspects are given in Table I under the column "frame."

We then investigated the effectiveness of compound identification at the encounter level. Each profile of the test set of an encounter (approximately 50 collected over a time period of about 10 s) was sent through the correlator and identified, and a compound identification was made based on all of them. At this level the correct-identification rate for the nose-on data was 100%, i.e., the test data from all thirty-three

encounters at this aspect could be correctly identified. The results for all aspects are given in Table I under the column "encounter." Although one needs to be careful about making general statements based on results obtained with this one particular data set, these results do suggest that reliable aircraft identification is possible using range profiles if compound identification techniques are employed over a sufficient number of profiles.

Finally, we briefly investigated what can happen if the training and test sets do not cover the same range of aspects. Instead of assigning alternate frames of data to the training and test sets (Fig. 4), we assigned the first half of the frames to the training set and the last half to the test set. Assuming the aircraft aspects were changing monotonically, this led to different ranges of aspects being represented in the training and test sets. We did this for nose-on aspects, formed new filters, and then ran a correlation with compound identification at the encounter level. The result for the nose-on data was 32/33 encounters correctly identified, as opposed to 33/33 in the "even/odd" experiment.

Potentially even greater aspect discrepancies would result if entirely different encounters were used in the training and test sets, possibly as much as a 20° difference. Unfortunately, not all aircraft had more than one encounter at each aspect. Of the thirteen aircraft types represented in Fig. 6, there were nine for which there were at least two nose-on encounters, the aircraft numbered 01, 05, 07, 08, 09, 10, 12, 14, 18. We ran an experiment in which one complete encounter was used as the training set, and the other complete encounter was used as the test set. Compound identifications based on the test encounters were able to correctly identify 7/9 aircraft. Not surprisingly, for the two aircraft incorrectly identified, the profiles of the training and test sets had significant and persistent qualitative differences, indicating that in these cases the aspects were very different between the training and test encounters.

V. CONCLUSION


We have described two types of correlation filters for identifying aircraft using radar range profiles, and have offered an algorithm for computing them. The results of experiments with an extensive data set suggest that reliable aircraft identification is possible provided aspect information is used, and provided identifications are based on sufficient number of profiles so as to decrease errors due to speckle fluctuations.

ACKNOWLEDGMENTS

The authors wish to thank the staff of Litton Data Systems Division, particularly Ken Snapp, Bill Ross, John Lucas, and Lou Kelly, for their help.


REFERENCES

- [1] Lax, P. D., and Phillips, R. S. (1989) *Scattering Theory*. New York: Academic Press, 1989.
- [2] Fok, F. Y. S., Moffat, D. L., and Wang, N. (1987) K-pulse estimation from the impulse response of a target. *IEEE Transactions on Antennas and Propagation*, AP-35 (1987), 926-933.
- [3] Dudley, D. G. (1988) Progress in identification of electromagnetic systems. *IEEE Antennas and Propagation Society Newsletter* (Aug. 1988), 5-11.
- [4] Wehner, D. R. (1987) *High Resolution Radar*. Norwood, MA: Artech House, 1987.
- [5] Steinberg, B. D. (1988) Microwave imaging of aircraft. *Proceedings of the IEEE*, 76 (1988), 1578-1592.
- [6] Werness, S., Carrara, W., Joyce, L., and Franczak, D. (1990) Moving target imaging algorithm for SAR data. *IEEE Transactions on Aerospace and Electronic Systems*, 26 (1990), 57-67.
- [7] Hudson, R. S., and Psaltis, D. (1988) Optical network that learns to perform motion compensation in radar imaging. *Proceedings of SPIE*, 881 (1988), 154-161.
- [8] *Proceedings of the IEEE*, Special issue on acousto-optic signal processing (Jan. 1981).
- [9] Goodman, J. W. (1985) *Statistical Optics*. New York: Wiley, 1985.
- [10] Chen, J. S., and Walton, E. K. (1986) Comparison of two target classification techniques. *IEEE Transactions on Aerospace and Electronic Systems*, AES-22 (1986), 15-21.
- [11] Chamberlain, N. F., Walton, E. K., and Garber, F. D. (1991) Radar target identification of aircraft using polarization-diverse features. *IEEE Transactions on Aerospace and Electronic Systems*, 27 (1991), 58-67.
- [12] Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T. (1988) *Numerical Recipes in C*. New York: Cambridge, 1988.



Scott Hudson received the B.S. degree in engineering and applied science, and the M.S. and Ph.D. degrees in electrical engineering in 1985, 1986, and 1991, all from California Institute of Technology, Pasadena.

Since 1990 he has been an Assistant Professor in the School of Electrical Engineering and Computer Science, Washington State University. His research interests are radar astronomy, radar imaging and target identification, and optical signal processing.



Demetri Psaltis received the B.Sc. in electrical engineering and economics in 1974 and the M.Sc. and Ph.D. degrees in electrical engineering in 1975 and 1977, respectively, all from Carnegie-Mellon University, Pittsburgh, PA.

After the completion of the Ph.D., he remained at Carnegie-Mellon, as a Research Associate and later as a Visiting Assistant Professor, for a period of three years. In 1980, he joined the faculty at the California Institute of Technology, Pasadena, where he is now Professor of Electrical Engineering and consultant to industry. His research interests are in the areas of optical information processing, holography, radar imaging, pattern recognition, neural networks, optical memories, and optical devices.

Dr. Psaltis has authored or co-authored over 170 publications. He is a fellow of the Optical Society of America and received the the International Commission of Optics Prize in 1989.