

Fast Translation Invariant Classification of HRR Range Profiles in a Zero Phase Representation

Joris Portegies Zwart^{1 2}

René van der Heiden³
Frans Groen^{1 2}

Sjoerd Gelsema²

May 23, 2004

¹University of Amsterdam, IAS Group, Kruislaan 403, 1098 SJ, Amsterdam, The Netherlands

²TNO Physics and Electronics Laboratory, P.O. Box 96864, 2509 JG, The Hague, The Netherlands

³NATO C3 Agency, P.O. Box 174, 2501 CD, The Hague, The Netherlands

Abstract

In this paper we describe a new method for aligning High Resolution Radar (HRR) range profiles for classification purposes. We describe the effects of Translational Range Migration on the phase of the Fourier transformed profiles, and show how the problem of finding an absolute alignment of HRR profiles can be described as a problem in phase estimation.

Using this description, we propose a new alignment method, the Time-Smoothed Zero Phase Representation, which we compare to existing alignment methods in terms of classification performance. Classification is performed using a training set of simulated profiles to classify both simulated and measured test profiles.

The method we propose has two main advantages compared to those described in the known literature. Since translation invariance is achieved early on in the classification process, it becomes possible to use more advanced feature extraction methods. Furthermore, the computational cost of classification is considerably lower compared to traditional alignment methods based on cross-correlation.

1 Introduction

An HRR range profile can be viewed as a one dimensional ‘image’ of an aircraft, where the parts of the aircraft that reflect the radar radiation, that is, the *scatterers*, are projected onto the line of sight (see figure 1). A description of the main properties of HRR profiles are given in section 2.

It has been shown (see [1] for an overview) that HRR profiles are promising candidate signatures for automatic classification of aircraft, since they provide potentially discriminative information on the geometry of aircraft.

The use of HRR profiles for classification purposes is challenging due to extreme within-class variability. Within each class, i.e. for each aircraft type, there exist three main sources of variation in HRR profiles: speckle, Rotational Range Migration (RRM) and Translational Range Migration (TRM). In section 2 we provide a description of these effects.

This work deals with the effects of TRM, which cause the location of scatterers within a profile to depend critically on the distance between radar and aircraft. Since this distance is not known with sufficient accuracy to correctly align profiles, i.e. within a fraction of the range resolution, some form of translational invariance has to be achieved during the classification process.

The remainder of this paper is built up as follows. In section 3, we give an overview of approaches to translation invariant classification in the recent radar literature. It is important to note that this problem is often ignored in literature. Also, many papers report on experiments using simulated profiles only—in this case the exact distance between radar and aircraft is obviously known, and so TRM is not a source of error.

In section 4 we show how the problem of determining an absolute alignment of HRR profiles can be described as a problem in phase estimation. Based on this description, we introduce a new alignment method in section 4.2.

In section 5 we describe a number of alignment and classification performance experiments. Experimental results are presented and discussed in section 6, followed by our conclusions in section 7.

2 HRR Range Profile Variability

In this section we discuss the main sources of HRR range profile variability: speckle, Rotational Range Migration and Translational Range Migration.

2.1 Measurement of HRR Range Profiles

The range profiles were produced by emitting a bandwidth B using K pulses with linearly increasing frequencies, called a stepped frequency waveform [2]. The coherent responses (K complex numbers) can be windowed and/or zero-padded, increasing the dimensionality from K to $d \equiv zK$, where $z > 1$ is the oversampling factor. From a subsequent Fourier transformation the phases are discarded—only the magnitudes are considered. Finally, d -dimensional range profiles are produced by squaring the magnitudes.

The aircraft aspect angle can be expressed as a coordinate pair (α, θ) where α is the aspect azimuth and θ is the aspect elevation (see figure 2). We define the aspect elevation θ as the angle between the radar line of sight and the plane through the wing-tips and nose of the aircraft. The elevation is positive if the aircraft is viewed from underneath. We define the aspect azimuth α as the angle between:

- The direction of the nose of the aircraft and
- The direction of the radar line of sight projected on the plane through nose and wingtips.

Thus, the aspect azimuth is zero if the aircraft is viewed from nose-on and π if viewed from tail-on. Finally, the aspect azimuth is chosen positive if the aircraft is viewed from the starboard side and negative if viewed from the port side. We will assume, however, that the aircraft is symmetric, so

that a range profile measured at aspect azimuth $-\alpha$ can be regarded to be identical to a profile measured at α .

A typical measurement consists of multiple profiles of an aircraft in flight as the full waveform is repeatedly transmitted while the radar is tracking the aircraft. Such a sequence of profiles is called a *leg*. Since the aircraft will be moving during the measurement of a leg, both the aspect angle and the distance between radar and aircraft will vary. This obviously leads to variability in the measured profiles.

2.2 Speckle and Rotational Range Migration

If an aircraft rotates over a large azimuth angle, such that the outermost scatterers move from one resolution cell to the other, the range profiles collected during this rotation suffer from Rotational Range Migration (RRM) [3].

Another effect, speckle, causes range profile variability for much smaller changes in aspect angle. It is caused if in a single resolution cell at least two distinct scatterers are present—then, only a slight rotation of the aircraft in aspect azimuth or elevation is enough to change the differential path length to the radar over half the wavelength. This causes the sum of the two scatter contributions to turn from constructive to destructive interference within tiny changes of aspect angle; generally between one and two orders of magnitude smaller than the aspect angle changes associated with RRM.

The effect of speckle is that the amplitudes of the range profile elements may vary rapidly if a sequence of consecutively measured range profiles is considered—the change in aspect angle is due mainly to small aircraft yaw motions during the recording time. The peak *positions*, however, do not alter.

2.3 Translational Range Migration

In addition to speckle and RRM, a third source of variation in HRR profiles is Translational Range Migration (TRM).

TRM occurs when a change in distance between the radar and the aircraft causes scatterers to move from one resolution cell to the next. Since all scatterers are translated by the same amount, the relative distance between two scatterers does not change. Therefore, the shape of the profile does not change due to TRM, and so *TRM causes a translation of the original profile*. In case of a stepped frequency waveform, this shift is cyclical.

To deal with the problem of TRM one must *align* the profiles. Alignment can be either *relative* or *absolute*. Relative alignment is achieved when a profile is aligned relative to one or more other profiles in a data set. Profiles aligned relatively to each other are not necessarily aligned to profiles in other data sets.

Absolute alignment means that profiles are shifted to optimise some *external* (i.e., not depending on other profiles in the same set) optimisation criterion.

3 Literature on HRR Range Profile Alignment

A general classification method consists of several discrete stages: *data acquisition* (measuring radar returns, computing profiles within a simulated environment), *pre-processing* (scaling, noise removal), *feature extraction* (including dimension reduction) and finally the actual *classifier*, which can be either similarity- or model-based.

Translation invariance can in principle be achieved in any of these stages. In this section, we categorise possible approaches to translation invariance proposed in recent radar literature.

3.1 Registration

During the pre-processing phase of a classification method, data is prepared for later feature extraction. Pre-processing typically includes transformations such as scaling and noise removal.

Translation invariance in this phase is achieved by *registering* the profiles: each profile is translated such that some optimisation criterion is maximised.

In the literature both relative and absolute alignment procedures are found in this stage. Recall that only absolute alignment leads to true translation invariance - if a relative alignment technique is used, true translational invariance must be achieved in one of the later stages.

Relative alignment is commonly achieved by aligning two profiles such that their correlation is maximised. The third profiles is the aligned with respect to the second, the fourth with respect to the third, etc. One drawback of these and similar approaches, extensively discussed in [4], is that if one profile is misaligned, this error propagates and disturbs the alignment of subsequent profiles in the data set. In [4] a method is proposed to deal with these ‘rogue’ profiles.

A possible absolute alignment method is given in [5]. This method, which the authors call *auto-aligning*, registers profiles by translating them such that the entropy of an energy vector E , whose elements are the inner products between the translated profile and K Gaussian windows with increasing standard deviation σ_k , is minimal.

3.2 Extracting Translation Invariant Features

The second opportunity for obtaining translation invariant classification is during the feature extraction phase. A common approach [6] is to use the magnitude of the Fourier transformed profiles as feature vectors, discarding all phase information. In [5] translation invariant features are obtained by first calculating the bi-spectrum (defined as the 2-D Fourier transform of the third-order autocorrelation function) for each profile, which is then integrated and inverse Fourier transformed to obtain features (called *accumulated bi-spectral features*).

Using translation invariant features can significantly decrease the computational cost of classification. The main disadvantage of these and similar approaches is that extracting translation invariant features involves discarding potentially discriminating information, which can result in a decrease of classification performance. Exactly which features should be considered, and whether this approach in general provides acceptable performance, is still an open question.

3.3 Translation Invariant Classifiers

Obtaining translation invariant classification in the actual classifier is probably the most common approach towards translation invariant classification taken in the literature. In [1], several similarity-based classifiers using the Sliding Euclidean Distance (SED) are discussed. The SED is a metric defined as the minimum Euclidean distance between two profiles $\mathbf{x}_1, \mathbf{x}_2$ over all possible cyclical translations:

$$\text{SED}(\mathbf{x}_1, \mathbf{x}_2) = \underset{i}{\operatorname{argmin}} \|(S^i \mathbf{x}_1) - \mathbf{x}_2\|^2, \quad (1)$$

for $i = 1, \dots, d$. The operator S^i shifts its argument vector i elements cyclically to the left, and d is the number of bins in the profile (and so $S^d \mathbf{x} \equiv \mathbf{x}$). Since

$$\|(S^i \mathbf{x}_1) - \mathbf{x}_2\|^2 = \|\mathbf{x}_1\|^2 + \|\mathbf{x}_2\|^2 - 2(S^i \mathbf{x}_1) \cdot \mathbf{x}_2, \quad (2)$$

we see that the SED is equivalent to the normal Euclidean distance after first aligning \mathbf{x}_1 and \mathbf{x}_2 such that their cross-correlation

$$\lambda(i) = (S^i \mathbf{x}_1) \cdot \mathbf{x}_2 \quad (3)$$

is maximal. This procedure can be implemented efficiently using the Fast Fourier Transform. Let \mathbf{X} denote the FFT $\mathcal{F}(\mathbf{x})$ of \mathbf{x} . The *discrete correlation theorem* [7] states that

$$\lambda = \mathcal{F}^{-1}[\mathbf{X}_1 \mathbf{X}_2^*], \quad (4)$$

where \mathbf{X}_2^* is the complex conjugate of \mathbf{X}_2 .

The SED as presented here is only approximately translation invariant, since the possible shifts are quantised – only shifts of an integer number of range bins are possible. However, if the number

of range bins is high (which can be enforced using zero-padding when constructing the profile), the approximation error can be made arbitrarily small.

In [8, 9] similar approaches using cross-correlation are described. In [10] translation invariance is achieved in a model-based classifier by modelling the statistical variations of profile centroids with respect to a set of reference profiles. During actual classification, the likelihood that an unknown test profile \mathbf{x} belongs to a given class C is integrated over all possible centroid positions, weighted by the centroid distribution.

The disadvantage of obtaining translation invariance in the actual classification stage is that meaningful feature extraction becomes very difficult. Many often-used extraction methods (like Principal Components Analysis) depend heavily on statistical properties of the data. Incorrect alignment will affect these properties, decreasing the accuracy of the feature extraction, which in turn decreases classification performance.

4 Zero Phase Representation

The effects of Translational Range Migration can be best described in terms of the phases of the Fourier transformed profiles. One of the well-known symmetries of the Fourier transform is that for any function $f(x)$ with Fourier transform $\mathcal{F}(f)$, the Fourier transform of $f(x - s)$ is given by

$$\mathcal{F}(f(x - s)) = \mathcal{F}(f(x))e^{i\omega s}, \quad (5)$$

i.e. a phase shift of $F(f(x))$. In the discrete case, where we are dealing with sampled functions f_n ($n = 1, \dots, d$) of f , a similar property holds for cyclically shifted versions of f_n :

$$\mathcal{F}(f_{(n+k) \bmod d}) = \mathcal{F}(f_n)e^{2\pi ink/d}. \quad (6)$$

So for a discrete shift k , the phase ϕ of the first AC component will be shifted by $2\pi k/d$, and higher order phases are adjusted accordingly.

Suppose we have a leg consisting of N profiles \mathbf{x}_i , $i = 0, \dots, N - 1$, taken from an aircraft in flight at discrete times t_i . If we calculate for each profile the phase ϕ of the first AC component of its FFT, we end up with a sequence $\phi(t_i) \equiv \phi_i$.

Variation in ϕ_i is caused by both rotational and translational effects. Suppose the aircraft remains at a fixed distance R w.r.t. the radar while rotating. This means ϕ_i is solely a function of the aspect angles α and θ :

$$\phi_i = \phi(\alpha(t_i), \theta(t_i)) \equiv \phi_i^a. \quad (7)$$

Since the aircraft remains at a fixed position, all variation in ϕ is the result of speckle and RRM.

In a more realistic setting, the distance $r(t_i) \equiv r_i$ between radar and aircraft also varies with time. This is a linear effect, so we can write

$$\begin{aligned} \phi_i &= \phi_i^a + \frac{2\pi}{L}(r_i - R) \\ &= \phi_i^a + \gamma(r_i - R), \end{aligned} \quad (8)$$

where $\gamma = 2\pi/L$ and L is the total length of the profile in meters.

The last term in (8) is the source of TRM. If the r_i were known exactly, reversing the effects of TRM would be trivial: we could just adjust the phase sequence ϕ_i by subtracting $\gamma(r_i - R)$ at each time step and adjust higher phases accordingly. Since r_i is not known, we have to find another way of eliminating the last term in (8).

4.1 Pure Zero Phase Representation

A simple method for finding a translation invariant representation (which has been previously used for alignment of panoramic images [11]), is setting each ϕ_i to zero and adjusting the higher order phases accordingly. This eliminates the range term in (8), but at the same time introduces an alignment error due to discarding the ϕ_i^a .

As an example, in figure 3(a) a stack of simulated range profiles of a Fokker 100 aircraft viewed over a 180 degree turn is shown. These profiles are perfectly aligned as the aircraft rotated around its centre without changing its distance to the radar. For comparison the same stack in the Pure Zero Phase Representation is shown in figure 3(b).

In general, setting the first phase exactly to zero will not result in an integer shift of the original sampled signal. In some cases it can therefore be beneficial to perform a phase shift such that the resulting phase of the first AC component is as close to zero as possible, while ensuring that the resulting shift in the time domain is an exact integer. If an integer shift of the original signal is not required, it may be necessary to again use zero-padding and windowing in the Fourier domain to minimise the effects of frequency leaking and the resulting occurrence of side-lobes [12].

4.2 Time-Smoothed Zero Phase Representation

We propose a more sophisticated method for dealing with the range terms in (8): the Time-Smoothed Zero Phase Representation (SZPR). It is based on the assumption that the relative location of the majority of prominent scatterers within the profile will remain stable over small changes in aspect angle. Therefore, two successive profiles in a leg can be accurately aligned *relative to each other* by maximum correlation (as in (3)). Relative alignment through maximal correlation is a popular approach in the literature for obtaining translation invariant classification within similarity-based classifiers.

Our aim is to combine the advantages of relative alignment through maximal correlation with the advantages of an absolute alignment using the pure ZPR. Consider again equation (8). The phases ϕ_i and ϕ_{i-1} measured at times t_i and t_{i-1} are related through

$$\begin{aligned}\phi_i - \phi_{i-1} &= \phi_i^a - \phi_{i-1}^a + \gamma(r_i - r_{i-1}) \\ &= \phi_i^a - \phi_{i-1}^a + \gamma\Delta r_i,\end{aligned}\tag{9}$$

for $i = 1, \dots, N-1$. Now, for each profile \mathbf{x}_i , $i > 0$ we can determine the optimal shift with respect to the previous profile \mathbf{x}_{i-1} using maximal correlation. This shift corresponds to a phase shift $\delta\phi_i^c$. *The assumption that this phase shift correctly aligns profile \mathbf{x}_i with respect to \mathbf{x}_{i-1} is equivalent to the assumption $\delta\phi_i^c = \gamma\Delta r_i$.* Since γ is known, this means that for each measurement (except the first) we have an estimate of the relative translation Δr_i of the aircraft between measurements.

Using this estimate we can correct the original phases ϕ_i for the aircraft translation between measurements to find a corrected phase sequence ϕ_i^c :

$$\begin{aligned}\phi_i^c &= \phi_i^a + \gamma(r_i - R - \sum_{j=1}^i \Delta r_j) \\ &= \phi_i^a + \gamma(r_0 - R)\end{aligned}\tag{10}$$

for $i = 1, \dots, N-1$. In other words, all profiles are now registered at $r = r_0$, the distance between radar and aircraft at the moment the first measurement was made.

To obtain absolute alignment, all profiles should be registered at the reference distance R . This can only be achieved if r_0 is known. Since r_0 is not known, we have to find a way to estimate it.

To estimate r_0 , we take a heuristic approach. At each time-step, $\gamma(r_0 - R)$ is estimated *locally* by a moving average over ϕ_i^c . This estimate is then subtracted from ϕ_i^c to obtain a final phase estimate ϕ_i^f :

$$\phi_i^f = \phi_i^c - \frac{1}{w+1} \sum_{k=-w}^w \phi_{i+k}^c.\tag{11}$$

By varying the window size w of the moving average filter, we can interpolate our alignment procedure between the absolute alignment of the pure ZPR ($w = 1$) and relative alignment through pure correlation ($w = N$).

In theory the optimal window size w depends on the differences in aspect angle between successive profiles, which in turn depends on the pulse repetition frequency (PRF) and the target

rotation rate. For small differences in successive aspect angles (high PRF and/or low target rotation rate), alignment by maximum correlation works well, and so we can use a large value of w . When the PRF is small, or the target rotates quickly, the quality of alignment by maximum correlation degrades, and we should use a smaller value of w .

However, in practical situations, i.e. when looking at targets of opportunity, the target rotation rate can not be measured accurately enough to allow for dependable estimates of w . In this paper we therefore try to find a globally acceptable value for w by experimentation.

Figures 3(c) and 3(d) show an example of the effect of the SZPR on a stack of profiles. Since filtering with a windows size w makes aligning the first $w - 1$ profiles impossible, these are left blank in the plots.

4.3 Implementation and Computational Cost

In this section we present our implementations of the SED and the SZPR, and compare the computational cost of classification when using either one. The example task is the classification of a leg consisting of M profiles \mathbf{x}_m , given a training set containing N pre-aligned (using the SZPR) training profiles \mathbf{x}_n .

To classify the leg using the SED, we have to perform the following computations for each of the MN possible pairs of test profiles \mathbf{x}_m and training profiles \mathbf{x}_n :

1. Calculate the FFT's \mathbf{X}_m and \mathbf{X}_n of \mathbf{x}_m and \mathbf{x}_n .
2. Calculate the correlation function $\lambda = \mathcal{F}^{-1}(X_n X_m^*)$ and find the maximum of λ . This determines the optimal alignment from maximum correlation.
3. Collect the phase ϕ_m of the first AC component of the \mathbf{X}_m 's, and adjust it such that it reflects the optimal alignment.
4. Use the adjusted phase to align \mathbf{X}_m with \mathbf{X}_n .
5. Calculate the Euclidean distance between the two profiles (which can be done directly in the Fourier domain.)

To classify the leg using the SZPR and a Nearest Neighbour classifier, we first have to perform the following computations to align the leg in the SZPR:

1. Calculate the FFT's \mathbf{X}_m and \mathbf{X}_n of \mathbf{x}_m and \mathbf{x}_n .
2. Calculate the $(M - 1)$ correlation function $\lambda = \mathcal{F}^{-1}(X_m X_{m-1}^*)$ and find the maximum of λ . This determines the optimal alignment from maximum correlation.
3. Collect the phases ϕ_m of the first AC component of the \mathbf{X}_m 's, and adjust them such that the leg is now aligned using maximum correlation.
4. Filter this phase sequence using a moving average filter with window size w .
5. Use the the filtered phase sequence to align \mathbf{X}_m in the SZPR.
6. Calculate the Euclidean distance between each of the MN possible pairs of test and train profiles (which can be done directly in the Fourier domain).

Table 1 summarises the computational costs, showing the number of times different computations have to be performed. From table 1 it is clear that aligning using the SZPR is significantly faster than using the SED. The tasks directly associated to aligning the profiles are roughly a factor N faster when using the SZPR than when using the SED. The filtering of the phase sequence adds some computational cost to the SZPR, but this is negligible since the filtering can be implemented efficiently using the Fourier transform, similar to the correlation function. [7]

5 Experiments

5.1 Available Data

For the classification experiments, two collections of range profiles are used in this paper. The first consists of *measured* range profiles, the second profile collection is *simulated* using the RCS-prediction code RAPPORT, developed at the TNO Physics and Electronics Laboratory [13]. As the bandwidth was chosen to be 450 MHz, the resolution of the range profiles was (after applying the Hamming window) about half a meter. It consisted of 324 steps in frequency, giving an unambiguous range of over 100 m. The coherent responses were routinely zero-padded to a dimensionality of 512 before applying the Fourier transforms.

Both collections contain profiles from five civil aircraft, seen at approximately broadside aspect angles: the Boeing 737-500 and the 747-400, the Fokker 100, the Airbus A310 and the McDonnell Douglas of the 80-88 series.

For the simulated data, the models of these five aircraft that are used in this paper have been purchased at Viewpoint Datalabs Intl., Orem, Utah, USA.¹ Figure 4 shows a rendered representation of the five models at approximately broadside aspect angles. For each of these models the predictions are carried out using RAPPORT.

The number of flat facets ranges from 5,238 for the Airbus to 17,935 for the Boeing 747. This demonstrates that the number of facets for the Airbus is relatively low. From the figure 4 the rougher surface of the Airbus can actually be observed. Clearly, the predictions for the Airbus will therefore be less accurate than for the other aircraft types.

For each aircraft profiles are predicted at 505 grid-points in aspect azimuth and elevation: one profile every 2.5 degrees in aspect elevation and one every 0.2 degrees in aspect azimuth. See figure 5.

The measured data comprise of 6 legs containing a total of 836 profiles. They are collected with the FELSTAR S-band radar at TNO-FEL located in The Hague, The Netherlands. During the measurements information from a secondary radar was available, which gave the ground truth on the aircraft type. The measured range profiles used in this paper are free of influences of radial velocity and acceleration² and the error on the estimated aspect angle was estimated to be about 5 degrees in both azimuth and elevation. The estimated aspect angles are plotted in figure 5 as well.

5.2 Construction of simulated legs

In the case of measured profiles, the concept of a leg is clear: profiles are collected from each aircraft a number of times in rapid succession. Each set of measurements is therefore a leg: a sequence of profiles ordered in time.

For the simulated profiles, the profiles are only labelled by aspect angle. Since our alignment procedure relies on ordered sequences of profiles, we have to artificially construct legs from the set of simulated data. For each aircraft in our database, we collect the profiles predicted at equal aspect elevation. These are then ordered according to aspect azimuth. Therefore, one leg of simulated data consists of 105 profiles at a constant aspect elevation, ordered by aspect azimuth which ranges from 80 to 100 degrees with a step size of 0.2.

5.3 Pre-processing

Each profile \mathbf{x} was transformed using a parametric power transform know as the Box-Cox transformation with parameter η , given by

$$\mathbf{x}^{(\eta)} = \begin{cases} (\mathbf{x}^\eta - 1)/\eta & \text{if } 0 < \eta \leq 1, \\ \log \mathbf{x} & \text{if } \eta = 0 \end{cases} \quad (12)$$

¹See <http://www.viewpoint.com/>.

²This is achieved through the use of a Velocity Tolerant Waveform, described in [1].

This transformation to *normality* has been shown to improve classification performance [14]. Following the reference we use an optimal value of $\eta = 0.2$.

After Box-Cox transforming the profiles, the baseline level of each profile was estimated by searching for that consecutive part of 30 elements in the profile which contained the lowest average energy. This estimate was then subtracted from the profile. This has been shown to be necessary when both simulated profiles and measured profiles are involved [1]. Finally, all profiles were normalised to unit energy.

5.4 Range Profile Classification

Since we are only interested in the relative performance of different alignment methods, the choice of classifier is arbitrary. The 1-nearest neighbour algorithm is straightforward to implement, and appropriate for all considered alignment methods. All classification experiments reported in this paper are therefore based on this algorithm, using the Euclidean metric as a distance metric.

To examine the effects of the different alignment methods on classifier performance (expressed as the percentage of correctly classified profiles), two sets of experiments were performed.

For the *first* set of experiments we split the set of simulated profiles into a training and a test set. The training set was produced by taking the odd numbered profiles from each leg. The test set consisted of all remaining, even-numbered profiles. In reality we obviously do not have such a highly correlated training and test set and the associated overestimate of classifier performance. However, our purpose in this paper is to investigate the relative merits of the various alignment techniques, and therefore absolute classifier performance is not relevant.

The effect of different alignment procedures on the the classification performance was measured using the classification performance on correctly aligned sets as a comparison baseline.

For the *second* set of experiments, the training set consisted of all available simulated profiles. The test set consisted of the legs of *measured* profiles. Again, we measured the effect of different alignment procedures on classification performance.

6 Results

6.1 Principal Component Analysis

To show the potential benefits of an absolute alignment, we performed a Principal Components Analysis [15] (PCA) on a data set including all profiles, both synthetic and measured, aligned using the SZPR with window size 8. A PCA dimension reduction projects the data onto a low-dimensional linear subspace such that the amount of variance in the projected data is maximal. Figure 6 shows the results in terms of the amount of variance of the projected data as a function of the dimensionality of the projection space.

From figure 6 we see that 90% of the total variance is contained in a linear subspace spanned by the first 40 PCA basis vectors. Such a statistical dimension reduction is clearly not possible using unaligned data.

6.2 Classification Performance

The classification (using a 1-nearest neighbour classifier) of the simulated test data using simulated training data was optimal, i.e. all 1751 test profiles were classified correctly. A perfect classification was also obtained by using the SED instead of the normal Euclidean distance.

Using the pure Zero Phase Representation for absolute alignment results in the confusion matrix shown in table 2. In this case, 13 profiles were classified incorrectly. Classification performance in the SZPR was also measured for a window size ranging from 1 to 20. Optimal performance was obtained at a window size of 8. In this case, only one profile was classified incorrectly.

These experiments show that in theory the SZPR is a viable method for absolutely aligning HRR profiles. In practice, however, one wants to classify measured profiles from a training set of simulated profiles. To mimic this situation (and simultaneously research the effect of noisy

data on alignment), we added noise to our simulated test set. Gaussian noise with variance σ^2 and zero mean was superimposed on the (I, Q) -data of the simulated test set. Then, overall classification performance was measured for correctly aligned profiles, profiles aligned using the SED, and profiles in the ZPR and SZPR as a function of σ . For the SZPR we used a window size $w = 8$. As a reference, we additionally provide the results for correctly aligned profiles. All results are shown in figure 7.

This process was repeated 5 times, each time drawing independent samples from the Gaussian noise distribution. The classification results were averaged over these 5 experiments. The error bar in 7 shows the standard deviation over these 5 measurements.

As could be expected, using correctly aligned profiles is optimal in terms of classification performance. Classification using the SED outperforms the SZPR, and is more robust with respect to noise. Recall, however, that classification in the ZPR/SZPR is roughly 500 times faster compared to the SED. Also, observe that for low signal-to-noise ratios, the lines converge to an error of 0.8 since the classification is effectively random over five classes.

Finally, we performed classification experiments using the measured profiles for testing. The training set consisted of *all* simulated profiles. Again, the SZPR was tested for window sizes ranging from 1 (which is equivalent to the pure ZPR) to 20. For this set, again a window size $w = 8$ produced the best classification in the SZPR. In table 3, confusion matrices for the SED, the ZPR, and the SZPR with $w = 8$ are shown. As in the case where we used simulated profiles for both training and testing, classification using the SED outperforms both the ZPR and the SZPR, though the differences are smaller. The reason for the overall higher classification error is the imperfect CAD modelling of the aircraft, and the imperfections in the RCS prediction code RAPPORT [16].

7 Conclusion

In this paper we have described a new method for finding an absolute alignment of HRR profiles. This method, the Time-Smoothed Zero Phase Representation, is a mixture of absolute alignment in the pure Zero Phase Representation and relative alignment using maximal correlation.

Absolute alignment is possible at the cost of a decrease in classifier performance. However, there are important advantages in using an absolute alignment. Because the effects of Translational Range Migration are minimised early on in the classification process, new possibilities for feature extraction arise. Furthermore, the computational cost of translation invariant classification in the SZPR is dramatically lower compared to the SED or similar methods, which is a major advantage considering the huge databases involved in aircraft recognition with radar range profiles.

Therefore, a combination of the SZPR and an advanced feature extraction step is expected to be a viable alternative for traditional approaches to translation invariant classification of HRR profiles.

References

- [1] HEIDEN, R. v. *Aircraft Recognition with Radar Range Profiles*. PhD thesis, University of Amsterdam, 1998. Available from <ftp://ftp.wins.uva.nl/pub/computer-systems/aut-sys/reports/ThesisRvdH.pdf>.
- [2] WEHNER, D. *High-Resolution Radar*, second ed. Artech House, 1994. ISBN 0-89006-727-9.
- [3] HUDSON, S., AND PSALTIS, D. Correlation filters for aircraft identification from radar range profiles. *IEEE Transactions on Aerospace and Electronics Systems* 29, 3 (July 1993), 741–748.
- [4] KOSIR, P., DEWALL, R., AND MITCHELL, R. A multiple measurement approach for feature alignment. In *Proceedings of the IEEE 1995 National Aerospace and Electronics Conference NAECON 1995* (1995), pp. 94–101. ISBN: 0-7803-2667-9.
- [5] LIAO, X., AND BAO, Z. Two new categories of shift-invariant features of high-resolution radar range profiles. In *Proceedings of ICSP '98* (1998), pp. 1485–1488.
- [6] ZYWECK, A., AND BOGNER, R. Radar target classification of commercial aircraft. *IEEE Transactions on Aerospace and Electronics Systems* 32, 2 (april 1996), 598–606.
- [7] PRESS, W. H., TEUKOLSKY, S. A., VETTERLING, W. T., AND FLANNERY, B. P. *Numerical Recipes in C++*, second ed. Cambridge University Press, 2002.
- [8] BALLARD, J. P., AND LEONARD, T. P. A maximum likelihood range-profile classifier for maritime targets. In *RTO SCI Symposium on "Non-Cooperative Air Target Identification Using Radar"* (1998).
- [9] NAKANO, Y., HARA, Y., AND SAITO, J. Radar target recognition system using 3-d mathematical model. In *SPIE Conference on Automatic Target Recognition VIII* (1998), pp. 83–91.
- [10] WEBB, A. Gamma mixture models for target recognition. In *NATO RTO-SCI Symposium on Non-Cooperative Air Target Identification Using Radar* (Mannheim, 1998). ISBN: 92-837-0000-7.
- [11] PAJDLA, T., AND HLAVÁČ, V. Zero phase representation of panoramic images for image based localization. In *8th International Conference on Computer Analysis of Images and Patterns* (1999), F. Solina and A. Leonardis, Eds., Springer Berlin, pp. 550–557. ISBN 3-540-66366-5.
- [12] HARRIS, F. J. On the use of windows for harmonic analysis with the Discrete Fourier Transform. *Proc. of the IEEE* 66, 1 (1978), 51–83.
- [13] BRAND, M. Radar signature analysis and prediction by physical optics and ray-tracing. The RAPPORT code for RCS prediction. Tech. Rep. FEL-95-A097, TNO-FEL, 1995.
- [14] HEIDEN, R. v., AND GROEN, F. The Box-Cox metric for nearest neighbour classification improvement. *Pattern Recognition* 30, 2 (1997).
- [15] JOLLIFFE, I. *Principal Component Analysis*. Springer, New York, 1986.
- [16] VAN DER HEIDEN, R., GROEN, F., AND VAN EWIJK, L. Aircraft recognition with radar range profiles using a synthetic database. In *RTO SET Symposium on "High Resolution Radar Techniques"*, (1999).

Task	SED	SZPR
FFT	MN	MN
Calculating λ	MN	$M - 1$
Phase adjustments	MN	$M - 1$
Filtering phase sequence	-	1
Shifting profiles	MN	M
Euclidean distance	MN	MN

Table 1: Computational cost of classification using the SED and the SZPR.

Class	B73S	B74F	EA31	FK10	MD80	Error
B73S	347	0	0	1	2	0.86 %
B74F	0	347	0	1	1	0.86 %
EA31	0	0	349	1	0	0.29 %
FK10	1	0	0	346	3	1.14 %
MD80	1	0	0	2	347	0.86 %

Table 2: Confusion matrix using a 1-NN classifier on profiles in the Zero Phase Representation. Overall classification error is 1.08 %.

(a) Confusion matrix for classification using the SED. Overall classification error is 30.4 %.

Class	B73S	B74F	EA31	FK10	MD80	Error
B73S	280	0	0	12	1	4.4 %
B74F	0	173	0	4	0	2.3 %
EA31	52	0	75	6	27	53.1 %
FK10	2	0	0	51	9	17.7 %
MD80	13	0	0	128	3	97.9 %

(b) Confusion matrix for classification in the ZPR. Overall classification error is 38.5 %.

Class	B73S	B74F	EA31	FK10	MD80	Error
B73S	240	0	0	49	4	18.1 %
B74F	0	170	0	5	2	3.9 %
EA31	57	0	70	4	29	56.2 %
FK10	4	1	0	33	24	56.8 %
MD80	23	0	0	120	1	99.3 %

(c) Confusion matrix for classification in the SZPR. The case reported is for a window size of 8. Overall classification error is 35.5 %.

Class	B73S	B74F	EA31	FK10	MD80	Error
B73S	232	0	0	56	5	20.8 %
B74F	0	170	0	1	6	3.9 %
EA31	43	2	76	3	36	52.5 %
FK10	1	0	0	54	7	12.9 %
MD80	17	0	0	120	7	95.1 %

Table 3: Confusion matrices for the various alignment methods. The training set consists of synthetic profiles, while the test set consists of measured profiles.

List of Captions to Illustrations

1. A range profile of an aircraft viewed from the left hand side. Responses from the aircraft scatterers (circles) are projected onto the line of sight, resulting in a radar range profile (bottom). (Geometrical data by Viewpoint Datalabs International.)
2. Definition of aspect elevation, θ , and aspect azimuth α . In this particular orientation both α and θ are positive.
3. Outputs of the several alignment procedures for one leg of simulated profiles. The profiles were obtained at zero elevation angle, with azimuth ranging from 0 to π . On the right, the phases of the resulting profiles are plotted.
 - (a) Correct alignment.
 - (b) Zero Phase Representation.
 - (c) Time-Smoothed ZPR, $w = 20$.
 - (d) Time-Smoothed ZPR, $w = 200$.
4. The five models available for RCS-predictions. Above each model the name of the aircraft type is shown. Between brackets, the code-name as used by secondary radar systems is given.
5. The aspect angles of the measured and predicted range profiles. The solid lines represent the measurements. Each curve is labelled by the call sign of the aircraft: Boeing 737 (B73S, 293 profiles), Boeing 747 (B74F, 177 profiles), Airbus A310 (EA31, 160 profiles), Fokker 100 (FK10, 62 profiles) and McDonnell-Douglas (MD80, 144 profiles). Each of the dots represent the aspect angle of five predicted range profiles, one for each aircraft type.
6. Variance of a data set consisting of all profiles in our data set as a function of the dimensionality of the PCA projection space.
7. Classification error for correctly aligned profiles (solid), using the SED (dotted), the ZPR (dashed) and profiles in the SZPR using a window size $w = 8$ (dash-dotted) as a function of the SNR.

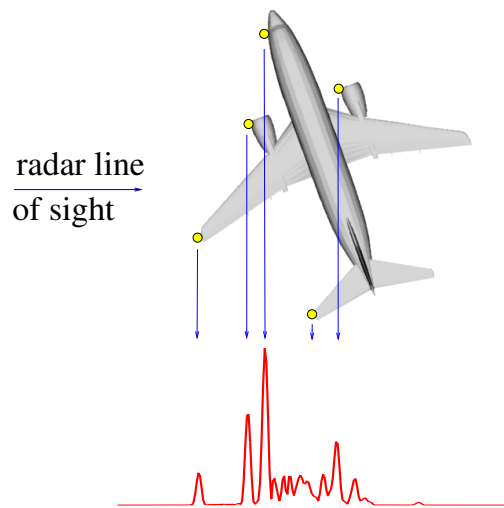


Figure 1: A range profile of an aircraft viewed from the left hand side. Responses from the aircraft scatterers (circles) are projected onto the line of sight, resulting in a radar range profile (bottom). (Geometrical data by Viewpoint Datalabs International.)

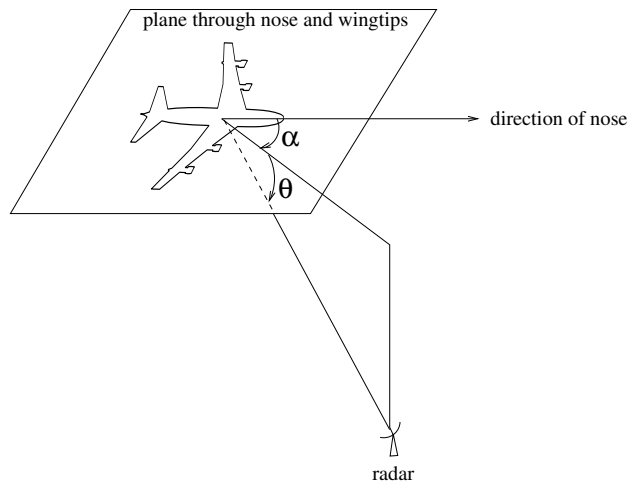


Figure 2: Definition of aspect elevation, θ , and aspect azimuth α . In this particular orientation both α and θ are positive.

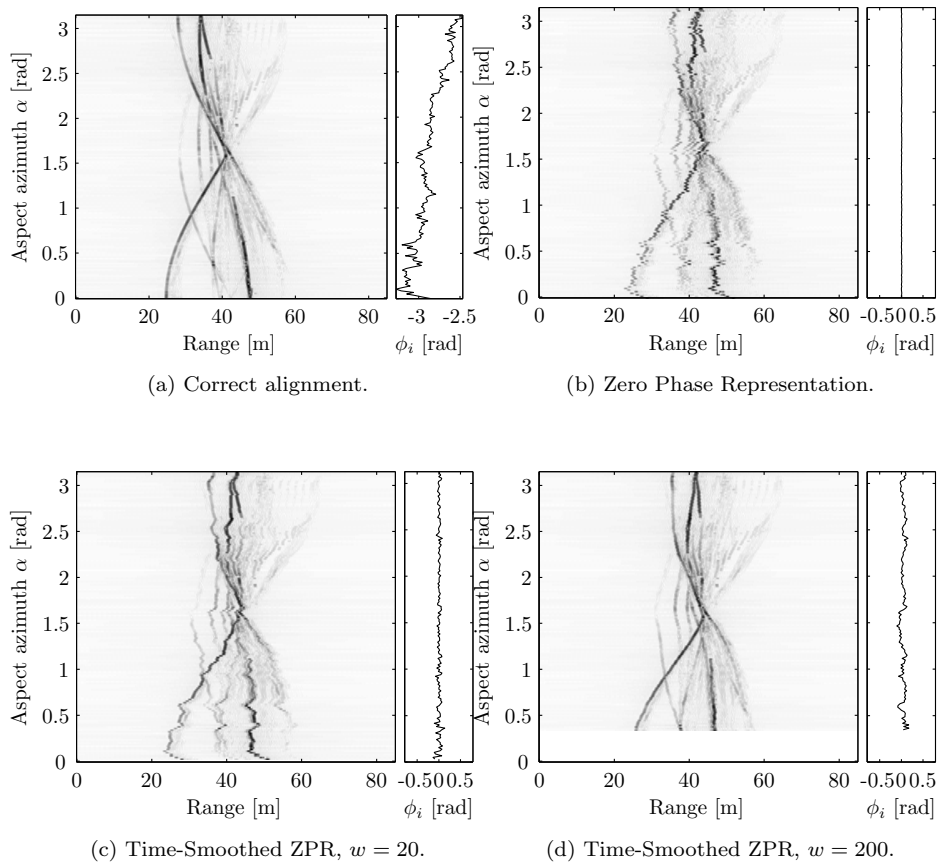


Figure 3: Outputs of the several alignment procedures for one leg of simulated profiles. The profiles were obtained at zero elevation angle, with azimuth ranging from 0 to π . On the right, the phases of the resulting profiles are plotted.




Figure 4: The five models available for RCS-predictions. Above each model the name of the aircraft type is shown. Between brackets, the code-name as used by secondary radar systems is given.




Figure 5: The aspect angles of the measured and predicted range profiles. The solid lines represent the measurements. Each curve is labelled by the call sign of the aircraft: Boeing 737 (B73S, 293 profiles), Boeing 747 (B74F, 177 profiles), Airbus A310 (EA31, 160 profiles), Fokker 100 (FK10, 62 profiles) and McDonnell-Douglas (MD80, 144 profiles). Each of the dots represent the aspect angle of five predicted range profiles, one for each aircraft type.




Figure 6: Variance of a data set consisting of all profiles in our data set as a function of the dimensionality of the PCA projection space.




Figure 7: Classification error for correctly aligned profiles (solid), using the SED (dotted), the ZPR (dashed) and profiles in the SZPR using a window size $w = 8$ (dash-dotted) as a function of the SNR.