Distance based range profile classification techniques for aircraft recognition by radar - a comparison on real radar data

R. van der Heiden
TNO Physics and Electronics Laboratory
PO Box 96864
2509 JG The Hague
The Netherlands

F.C.A. Groen
University of Amsterdam
Faculty of Mathematics and Computer Science
Kruislaan 403, Watergraafsmeer
1098 SJ Amsterdam
The Netherlands

ABSTRACT

Aircraft identification is essential in any air-defence scenario. Without a robust classification capability no effective threat evaluation can be performed. A prominent aircraft recognition technique is based on the exploitation of a one-dimensional image of a target, a *range profile*. In this paper, we employ four different classification techniques, all based on shift invariant distances, and a method to compare them. Two of the techniques are based on Radial Basis Functions for which a novel technique to optimize the number of free parameters is presented. The application is on real radar data, where a true separation between train- and test profiles is accomplished. The classification results are encouraging. As an example, a qualitative statement is given about the best of the four classifiers to be used in combinations of two scenarios and four applications.

1. INTRODUCTION

An important aircraft identification technique, Identification Friend Foe, relies on the cooperation of the target. If, in war- or crisis time, the aircraft fails to cooperate for whatever reason the only *save* conclusion for the interrogator is that the aircraft is hostile.

This incomplete decision process caused serious cases of fratricide. In April 1994, two Blackhawk (friendly) helicopters were shot down in the no-fly zone of Iraq. This incident underlined again the importance of an additional identification capability such as NCTR (Non-Cooperative Target Recognition).

Currently we are investigating the NCTR potential of High Range Resolution (HRR) range profiles. Measurement of these signatures is relatively easy and the requirements for motion compensation are moderate or the compensation may even be omitted. Additionally, range profile classification is applicable at almost all aircraft orientations.

In the literature, several approaches to classify range profiles are reported. Selection of a feature vector from the spectral components of a range profile is reported by Garber *et al* [1], DeWitt [2] and Kouba [3]. Classification of this vector is carried out by a nearest neighbour rule, a *Hidden Markov Model* and a recurrent neural network, respectively. The latter two have the interesting ability to process *sequences* of range profiles. Baras and Wolk [4] showed the feasibility of range profile classification on multiple resolution levels using wavelets.

In real measurements, the absolute positions of the scatter returns in a range profile are undefined. It requires that the classification method should be shift-invariant. A promising solution to this problem is the use of correlation filters [5]. In this paper we investigate the potential of a shift invariant profile-to-profile distance. Once it is defined, all classification techniques that are based on these pair distances are available. Earlier results on such a distance metric using a nearest neighbour classifier (section 3.4) are reported by Novak [6].

Four different classification methods are devised and tested. Two of them are based on the Nearest Neighbour rule, the other two are implementations of a Radial Basis Functions network. In this study a thorough test on real radar data from inflight aircraft is carried out. An important property of the used data set is that the train- and test profiles were measured independently.

Furthermore, we present a method to compare the classifiers. Clearly, an important comparison criterion is the error on an independent test set. With an eye on future applications for a range profile classifier we believe that it is important to include the classification speed in the comparison as well. Two less important parameters are the time needed for training a classifier and the size of the classifier.

Finally, we select for all combinations of two scenarios and four applications the most appropriate classifier. Although these choices are rather tentative given the moderate amount of data that is considered, it clearly demonstrates the employed method.

The organisation of this paper is as follows: The next section reviews the physics of range profiles, section 3 describes the used distance metric and the four classification techniques. Then, section 4 establishes the approach to compare classification techniques. Section 5 shows the results on real radar data and, finally, section 6 gives the conclusions.

Paper presented at the AGARD MSP 3rd Symposium on "Tactical Aerospace C³I in Coming Years", held in Lisbon, Portugal from 15th May to 18th May 1995, and published in CP-557.

2. RANGE PROFILES

Figure 1 shows the contour of an aircraft and its range profile. The profile can be viewed as a projection of the aircraft scatterers onto the line of sight. It thus shows the radar cross section as a function of range.

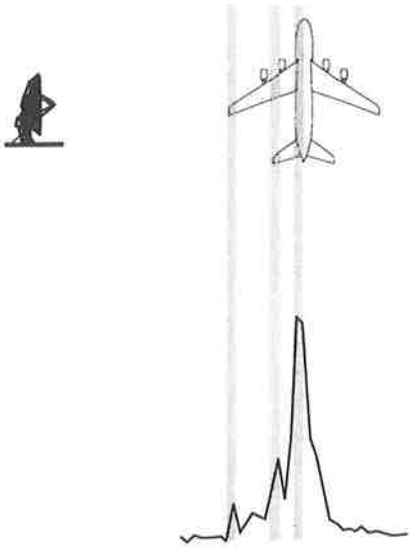

Fig. 1: The aircraft is illuminated from the left side. Each strip represents a range cell. The contributions of the scatterers in each strip are summed to constitute a single range profile element.

For the generation of range profiles, we need a radar that is able to emit a high bandwidth-waveform. This can be done either with a single short pulse, or with a burst of pulses at linearly increasing carrier frequencies [7].

Due to coherent summation of aircraft scatterers (speckle), the exact shape of the range profiles depends strongly on aspect angle. However, the overall profile shape does not change significantly (see for example figure 3) as long as the aircraft scatterers do not move outside one range resolution cell.

For the available data, the maximum change in aspect angle to avoid this *rotational range migration* is approximately 1.5 degrees. It is very difficult to determine the aircraft aspect angle with this accuracy. Consequently, a simple look-up table approach - a measured profile and its aspect angle is compared to the profiles in a data base with the same aspect angle - is not applicable [5].

Another approach is therefore to consider aspect angle bins that are several times larger than the *error* in aspect angle. The procedure is to construct a classifier for each bin. Then, for an unknown profile, retrieve its aspect angle, select the appropriate bin and assign a class with the corresponding classifier. Evidently, an extra mismatch probability will occur, because a profile from class 1 may look like a profile from class 2, as seen at a different aspect within the bin.

For the data set we consider in this paper the errors on the aspect angles are witnin five degrees. All profiles have aspect angles with absolute values ranging from 0 to 20 degrees and are placed in a single bin.

3. RANGE PROFILE CLASSIFICATION

3.1 Definitions

Two sets of independently measured range profiles within a single bin are available, the *input set* and the *test set*. A subset of the input set, the *train set*, is used for training a classifier. The profiles in the input set, train set and test set are randomly ordered and are named $\mathbf{r}_i$, $i = 1, \ldots, N_{\text{input}}$, $\mathbf{p}_i$, $i = 1, \ldots, N_{\text{train}}$, and $\mathbf{q}_i$, $i = 1, \ldots, N_{\text{test}}$ respectively.

A *classifier* is fully determined by

1. the *classification technique* and
2. the *train set*.

The classes of all profiles are known. This enables us to train and test a classifier. Clearly, in an operational situation there is no test set available.

3.2 Sliding Euclidean Distance

All our classification methods are based on profile-to-profile distances. The absolute positions of the reflections in the profile depend strongly on the distance to the target. As we cannot estimate this distance accurately enough to place the reflections on an objective position, we must use a shift invariant distance.

Suppose we have two range profiles $\mathbf{x}_1$ and $\mathbf{x}_2$, length α, elements $x(1), \ldots, x(\alpha)$. Then we define the distance D as the minimum Euclidean distance over all shifts:

$$D(\mathbf{x}_1, \mathbf{x}_2) \equiv \min_{j=0,\ldots,\alpha-1} \sqrt{\sum_{i=1}^{\alpha} \left[x_1(i+j) - x_2(i)\right]^2}$$

The shifts are cyclical, that is $x_1(\alpha + j) \equiv x_1(j)$.

3.3 Compression and normalisation of profiles

In profile classification using the Sliding Euclidean distance it is advantageous to lift the weak scatterers in the range profile relative to the strong scatterers so that they can play a role in the profile matching as well. Several choices can be made for such a *compression*, e.g. a log-scale or a power function with a power less than one.

Current investigations concern the search for the optimum compression function. Preliminary results show that a power function with a power ¼ works satisfactorily.

After the compression, the profiles need to be normalized, as the magnitudes of the reflections depend strongly on the absolute sensitivity of the radar and the distance at which the aircraft was measured. As neither the sensitivity nor the exact distance of the aircraft is known, we normalise the compressed profiles such that the sum of squares of the profile elements equals one.

3.4 Nearest Neighbour

The nearest neighbour rule decides that the class of a profile from the test set is the class of the nearest profile in the train set. Here 'nearest' is with respect to the chosen distance metric D.

A simple extension to this technique is to search for the k ($k \geq 2$) nearest neighbours, giving k class declarations. The class that occurred most frequently is assigned to the profile

from the test set. Experiments showed that this extension did not give significant differences in the classification results. Therefore we will only consider a 1-nearest neighbour in this paper.

3.5 Condensed Nearest Neighbour

A drawback of the nearest neighbour technique is the large computational effort necessary for the classification. For each profile for which classification is desired, we have to compute all distances to the profiles in the train set again. This is even more a problem in our application, because the chosen distance measure D is computationally expensive.

The technique we will apply here to reduce the computational burden is based on the idea that a profile that is far from the decision boundary has, on average, far less influence on the outcome of the nearest neighbour classifier than a profile that is near the decision boundary. Therefore we might as well skip this profile and save the computation time.

It is possible that a profile does not contribute to the decision boundary at all, as it is completely surrounded by other profiles from the same class. Skipping it does not alter the outcome of a nearest neighbour rule. However, in our application this situation seldom occurs as a profile is of very high dimension and thus almost always defines a part of the decision boundary. This means that in virtually all cases the classification accuracy is reduced if a profile is removed.

In this paper, we use the *condensing* algorithm [8]. To arrive at the condensed subset of the train set, two complementary subsets of this set, named A and B, are defined. Place the first profile from the train set, $\mathbf{p}_1$, in A, the remaining profiles, $\mathbf{p}_2$, . . . , $\mathbf{p}_{N_{train}}$, in B. The method proceeds as follows:

1. Use the nearest neighbour rule to classify the first profile in B with the profile(s) in A. If it is classified correctly with the nearest neighbour rule, leave it in B, otherwise, place it in A. Repeat this operation for all profiles that are left in B.
2. If in step 1 not a single profile has been transferred from B to A, or if B is empty then terminate. Else return to step 1.

After termination, A contains the condensed subset. For classification, the nearest neighbour rule is applied using the condensed subset instead of the full train set.

3.6 Radial Basis Functions

Radial Basis Functions (RBF) provide a way to construct a function that maps vectors from a high dimensional space onto a lower dimensional space [9]. As the only inputs for this method are distances between profiles we can use the sliding Euclidean distance D to make the method suitable for range profile classification. The advantage of the used RBF implementations compared to a nearest neighbour technique is the large reduction of classifier size and classification effort.

From the train set, L profiles are selected to serve as *centres* $\mathbf{c}_l$, $l = 1, \ldots, L$. (The next two subsections 3.7 and 3.8 describe the used selection methods.) The pair distances between the centres and a train profile $\mathbf{p}_i$ enter the RBF network and form the elements of a distance vector $\mathbf{d}_i$ with elements:

$$d_{il} = D(\mathbf{p}_i, \mathbf{c}_l) \tag{1}$$

Then, a non-linear transform using a Gaussian function

$$\phi(r) = e^{-\frac{r^2}{2}} \tag{2}$$

is applied to each of the elements in the distance vector, giving

$$b_{il} = \phi(d_{il}). \tag{3}$$

Multiplication of the vector $\mathbf{b}_i = (b_{i1}, \ldots, b_{iL})^T$ by a weight matrix W and addition of a bias vector $\mathbf{w}_0$ gives, for each train profile, the output $\mathbf{o}_i$.

$$\mathbf{o}_i = \mathbf{w}_0 + W\mathbf{b}_i \tag{4}$$

In the training phase, the weights $\mathbf{w}_0$ and W are chosen such that the outputs are as close as possible to unit vectors in a γ-dimensional space, where γ is the number of classes.

Train profiles from class 1 are mapped as close as possible onto the output $\mathbf{e}_1 \equiv (1, 0, \ldots, 0)^T$, train profiles from class 2 onto $\mathbf{e}_2 = (0, 1, \ldots, 0)^T$, et cetera.

Hence the training of the Radial Basis Functions network boils down to finding the least squares solution for $\mathbf{w}_0$ and W using equation 4 for all train profiles. This is an attractive property of the Radial Basis Functions approach: although it is able to construct any complex non-linear decision boundary, the weights can be found by linear methods [9].

For classification we simply compute the output for a *test* profile $\mathbf{q}_i$ using the distances D to the centres and equations 3 and 4. If the closest unit vector to the output is $\mathbf{e}_j$ then class j is assigned to the test profile.

In the next sections we will address the problems of choosing the centres and selecting the number of centres.

3.7 Radial Basis Functions with Random Centre Selection

A good first choice for the centres is to select them randomly from the train profiles. One must be careful, however, about the number of centres to choose. Each extra centre adds an extra degree of freedom to fit the train profiles. If we take too few centres, the approximation will be too coarse. If we take too many centres (but less than the number of train profiles), also the noise on the profiles will be fitted ('overfitting'). In both cases, the generalisation capabilities of the classifier will be worse than with an intermediate number of centres.

To find the optimum number of centres we devised the following algorithm:

1. Select, randomly, half of the profiles from the train set, and use them as *evaluation* set. Use the other half as *design* set.
2. Choose, randomly, a profile from the design set (one that has not been chosen earlier) and copy it to the *centre* set.
3. Find the weights using the centre set and the design set from equation 4.
4. Compute the outputs $\mathbf{o}_i$ for each of the evaluation profiles. Compute the sum of the errors $\| \mathbf{o}_i - \mathbf{e}_j \|$ where $\mathbf{e}_j$ are the desired outputs corresponding to the class of the evaluation profiles. This is called the *evaluation error*.
5. Repeat steps 2, 3 and 4 until all design profiles (apart from one) are used as centre.

Similarly to step 4 it is instructive to compute the *design error* as well.
Although this error decreases monotonically as the number of centres increases, the evaluation error will reach a minimum for a certain number of centres. See figure 2. The best choice for the number of centres is therefore at the minimum value of the evaluation error.

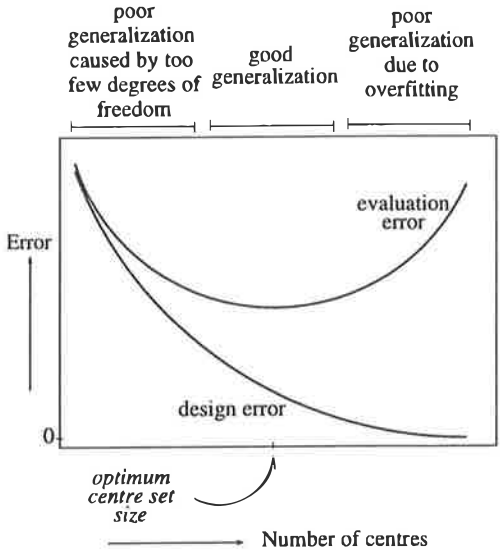

Fig. 2: As the number of centres increases the network is able to represent the design profiles better, i.e. the design error tends to zero. However, the true classification capability is revealed by the error on an independent evaluation set.

3.8 Radial Basis Functions with Gram-Schmidt Centre Selection

A procedure to select the best centres from the design set is to use a Gram-Schmidt orthonormalization technique. Here we will confine ourselves to a qualitative description, for details we refer to [10].
As in the Random Centre Selection, the first step is to employ the first part of the train set for designing the classifier and the other part for evaluation. Then we search for *that* profile in the design set that gives the best least-squares solution if used as a centre. At each next step, we add another profile to the centre set that gives the best improvement of the least-squares solution.
Instead of the computation of the least-squares solutions, Chen *et al* [10] devised an efficient Gram-Schmidt orthonormalization procedure to select the best centres.
As in the random centre selection, we compute the classification error on the independent evaluation set and choose that number of centres where the evaluation error has a minimum.

4. COMPARISON OF CLASSIFIERS

Often, classification techniques are compared in terms of their errors on a test set only. For most practical applications three more properties define the usefulness of a classifier. The following list gives the four measurable classifier properties that are of interest.

- a_1 Classification error [% false on independent test set]
- a_2 Computational effort needed for one classification [# floating point operations]
- a_3 Computational effort needed for training the classifier [# floating point operations]
- a_4 Memory required to store the trained classifier [# of bytes].

Ideally, each of these quantities equals zero. In practice, for each classification technique there will be a trade-off between these four properties which can be found by varying the size of the train set.
For example, let us consider a classifier that uses the nearest neighbour technique ($a_3 \equiv 0$) and a small sized train set. Then the classification error, a_1, will decrease if the train set increases. This also implies, however, that more distances have to be computed and it thus results in a larger a_2 and a_4.
To choose the right classifier we would like to have weight functions, ω_i (monotonously increasing) so that the quantity

$$\sum_{i=1}^{4} \omega_i(a_i) \qquad (5)$$

is minimized with respect to $a_1, \ldots, a_4$. Unfortunately, we do not have these functions available, but we can make a few simplifying but realistic assumptions to tackle the problem.
The first one is that the most important parameters in a military context are a_1 and a_2. The time needed for training (a_3) is of much lesser importance, because it can be done off-line. The size of the trained classifier is generally also less significant. Besides that a_4 is (almost) linearly related to a_2 for the classification techniques we consider. Therefore we do not have to minimize a_4 by itself. For the remainder of this paper, we will therefore focus on a_1 and a_2 only.
We do not make a choice for ω_1 and ω_2 either, but construct a large number of classifiers to demonstrate the trade-off between a_1 and a_2. For example, if the user wishes a certain a_1 he may find in a single curve the classifier that has the smallest a_2.
At this point, we also want to stress that not only the application (e.g. surveillance or aircraft radar) is decisive for the classifier choice, but also the scenario (crisis or wartime). As an illustration table 1 shows roughly the importance of correct classification and fast classification as a function of the application and the scenario.

This table shows in qualitative terms that in times of crisis it is more important to have a reliable answer then to have a quick answer. In wartime it is of greatest important to have a fast answer.

5. RESULTS

5.1 Available data

We have an input set available of 357 profiles of four different aircraft from an S-Band radar. The number of elements of the profiles is 128. These profiles were gained in six different aircraft flights.

	scenario	crisis		wartime	
	classifier property	correct class.	fast class.	correct class.	fast class.
application	SHORAD	+	0	0	+
	HIMAD	+	−	0	0
	Fighter aircraft	+	0	0	+
	Surveillance	+	−	0	−

Table 1: Relative importance of classification properties for application and scenario in terms of minus signs (less importance), zeroes (moderate importance) and plus signs (high importance). Here SHORAD means SHOrt Range Air Defense and HIMAD HIgh to Medium Air Defense (e.g. HAWK, PATRIOT).

In five other measurements, 339 profiles were obtained from the same four aircraft. These profiles made up an independent test set.

For each profile, an approximate aspect angle is available. The absolute aspect angles (we assume symmetry around aspect angle 0) are in the range of 0 to 20 degrees from head-on. The errors on the angles are believed to be within 5 degrees, the elevation is approximately zero.

As stated in subsection 3.3 the profiles were compressed with a power of 1/4 and normalised. Figure 3 shows some examples of compressed and normalised profiles from the four different classes and from the input- and test set. Each of the three profiles in one class and one set was measured during the same flight. This means that, although the aspect angle is inaccurate, the change in aspect angle (as indicated above each profile) is more accurately defined.

5.2 Classification experiment

In this section we investigate the properties of the classification techniques of chapter 3. To this end, we construct a large number of classifiers using the four techniques and varying train sets to monitor the trade-offs between the classification speed and the classification error.

Carry out the following steps for $N_{\text{train}} = 8, 24, 40, \ldots, 152$:

1. Choose, randomly, $N_{\text{train}}/4$ profiles per class from the input set and use them as train set.
2. Construct the classifiers
   - **NN** (Nearest Neighbour)
   - **CNN** (Condensed Nearest Neighbour)
   - **RR** (Radial Basis Functions with Random Centre Selection)
   - **RGS** (Radial Basis Functions with Gram-Schmidt orthonormalization)

   for this train set.
3. Classify all profiles in the test set using these classifiers. Compute the percentage of false classifications. This gives a_1. Also keep track of the number of flops used for classifying a single profile (a_2).
4. Repeat steps 1-3 thirty times and average a_1 and a_2.

The results are shown in figures 4 and 5.

Figure 4 shows that a fairly good classification rate can be achieved with only a small number of profiles per class in the train set. For example the nearest neighbour technique needs only 40 train profiles (on average one profile per class per two degrees) to achieve a classification error less than 11%. It suggests that a rather crude coverage of aspect angle suffices for reasonable classification, although one must be aware that the performance is favoured by the small number of classes. For all sizes of the train set, the nearest neighbour technique has the best classification rate. This technique apparently makes the best use of the available data. For small sizes of the train set, both Radial Basis Function techniques have poor classification rates. This is because half of the profiles has to be used for evaluation. If the train set increases, this effect becomes less important.

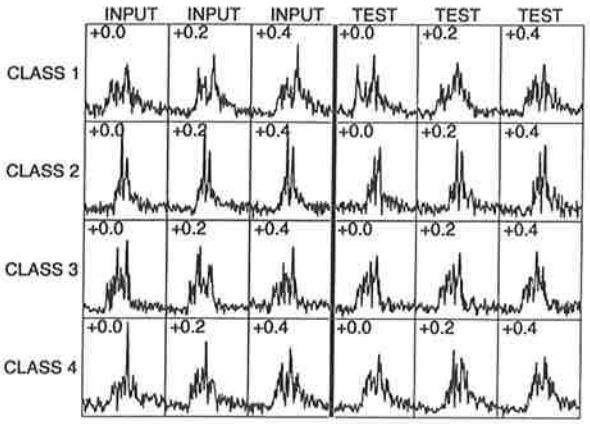

Fig. 3: Examples of compressed and normalised profiles of four different aircraft, near head-on. The input set is shown on the left hand side and the independent test set on the right hand side. In the upperleft corner, the aspect angle difference (in degrees) relative to the leftmost of the three profiles is shown. As can be seen, the small scale variations (speckle) are unpredictable whereas the overall appearance, in most cases, is similar.

The condensed nearest neighbour has an approximately 6% higher classification error rate than the normal nearest neighbour for all sizes of the train set. As stated in section 3.5, the condensing procedure deletes profiles that somehow contribute to the decision boundaries.

The Radial Basis Functions using a Gram-Schmidt centre selection has a somewhat better classification rate than the RBF using a random centre selection. For larger input sets, the difference tends to vanish.

The classification effort (figure 5) is closely related to the number of profiles that is present in the classifier.

In the nearest neighbour case, all profiles are used in the classifier - the CNN classifiers use the condensed profiles only. The classification effort is exactly linearly related to the

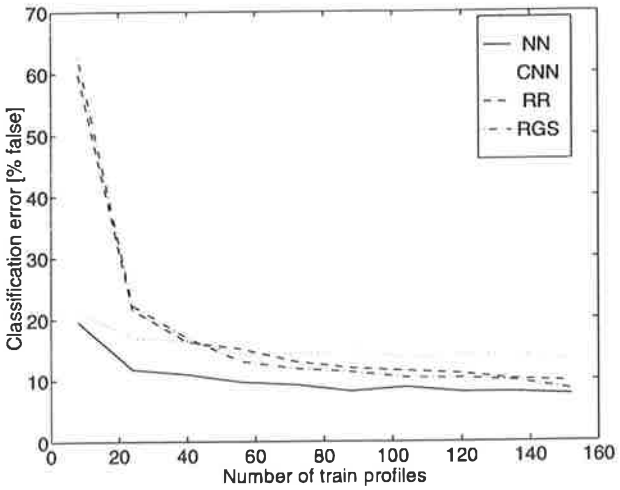

Fig. 4: Average percentage wrong (a_1) as a function of train set size.

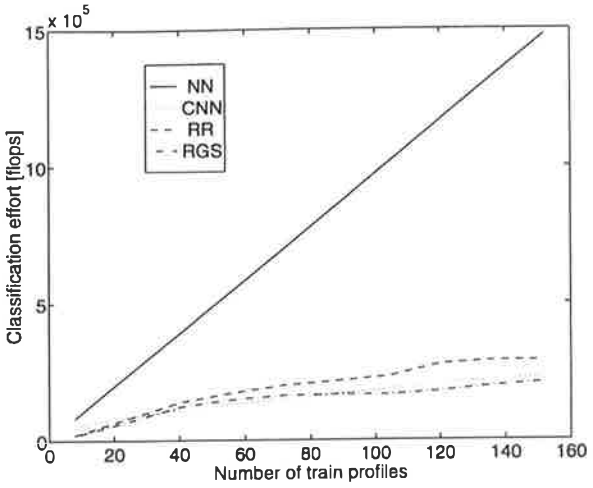

Fig. 5: Average classification effort for resulting classifier (a_2) as a function of train set size.

number of profiles in the train set (NN) or the number of condensed profiles (CNN).
The left-over profiles in an RBF classifier are the centres. The major part of the computations arises from the profile-to-profile distance evaluations - a small number of extra computations is necessary for the non-linear transform (equation 2) and the matrix multiplication (equation 4).
The two plots show that in the CNN-, RR- and the RGS classifiers only a very small number of profiles is left over, compared to the nearest neighbour. It means that redundant or nearly redundant profiles are removed at the cost of an increased classification error.
As the important parameters are a_1 and a_2, figure 6 shows the trade-offs between classification rate and classification speed for the four classification techniques.
From this figure one can decide which classifier is most appropriate for a particular classification purpose. The simple approach is to choose a desired classification error on the vertical axis, move horizontally until the first curve in the plot is reached. This classifier should be used as it is the most rapid one. For example, if one desires a classification between approximately 9% and 14% an RGS classifier is the best choice. The required size of the train set can be found in figures 4 or 5.
Conversely, figure 6 can be utilised to find the best classifier given a desired classification speed. E.g. if one is willing to carry out 10,000 flops for one classification, a CNN classifier is the best choice, because it has the minimum error of approximately 17%.
If one desires the minimum classification error possible, a nearest neighbour is appropriate, but it will take a long time to answer.
Returning to the table 1 we may insert, using figure 6, the most appropriate classification techniques; see table 2. We want to stress that filling in this table is merely a demonstration of the method of classifier selection - for a decisive answer on which techniques to use, larger scaled experiments have to be carried out.

6 CONCLUSIONS

In this paper we described a successful classification test on range profiles. The profiles used for training and those for testing were acquired in strictly separated aircraft flights and covered a wide aspect angle range of 20°. Still, the best classification results were within 10% error. Although these results are favoured by the small number of classes, they are very encouraging for the applicability of this technique for NCTR.
All classification techniques we considered were based on profile-to-profile distances. The best rates were found using a simple nearest-neighbour rule. In this paper we demonstrated, however, that for most cases this is not the best technique if one includes the classification speed into the comparison as well. It shows that the Condensed Nearest Neighbour and the Radial Basis Functions with Gram-Schmidt orthonormalization have a more favourable trade-off between classification rate and -speed.

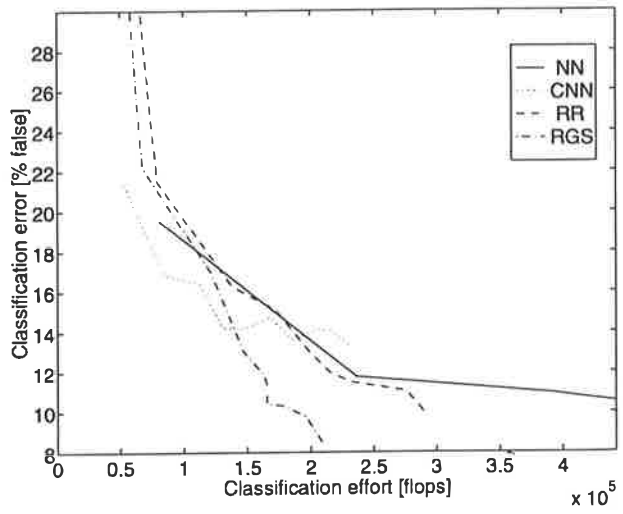

Fig. 6: Average classification effort (a_2) vs Average classification rate (a_1) (zoomed).

	scenario	crisis	wartime
application	SHORAD	RGS medium train set	CNN small train set
	HIMAD	RGS large train set	CNN small train set
	Fighter aircraft	RGS medium train set	CNN small train set
	Surveillance	RGS large train set	CNN medium train set

Table 2: Best classification technique given scenario and application.

7. ACKNOWLEDGEMENT

This research is sponsored by the Ministry of Defence of the Netherlands.

REFERENCES

[1] F.D. Garber, N.F. Chamberlain, and Ö. Snorrason. Time-domain and frequency domain feature selection for reliable radar target identification. In *IEEE National Radar Conference*, pages 79-84, 1988.

[2] M.R. DeWitt. High range resolution radar target identification using the Prony model and hidden Markov models. Master's thesis. AFIT/GE/92-d-15, Air Force Inst. of Tech., Wright-Patterson AFB, OH. School of Engineering, 1992.

[3] E.T. Kouba. Recurrent neural networks for radar target identification. Master's thesis. AFIT/GSO/ENG/92d-02, Air Force Inst. of Tech., Wright-Patterson AFB, OH. School of Engineering, 1992.

[4] J.S. Baras and S.I. Wolk. Wavelet based progressive classification of high range resolution radar returns. In *SPIE Vol. 2242 Wavelet Applications*, pages 967-977, 1994.

[5] S. Hudson and D. Psaltis. Correlation filters for aircraft identification from radar range profiles. *IEEE Trans. on Aerospace and Electronic systems*, 29(3):741-748, July 1993.

[6] L.M. Novak. A comparison of 1-D and 2-D algorithms for radar target classification. In *IEEE International Conference on Systems Engineering*, pages 6-12, 1991.

[7] D.R. Wehner. *High Resolution Radar*. Artech House, 1987.

[8] P.A. Devijver and J. Kittler. *Pattern recognition, a statistical approach*. Prentice Hall International, 1982.

[9] D.S. Broomhead and D. Lowe. Radial basis functions, multi-variable functional interpolation and adaptive networks. Memorandum 4148, RSRE, 1988.

[10] S. Chen, C.F.N. Cowan, and P.M. Grant. Orthogonal least squares learning algorithm for radial basis function networks. *IEEE Trans. on Neural Networks*, 2(2):302-309, March 1991.

AGARD-CP-557

AGARD

ADVISORY GROUP FOR AEROSPACE RESEARCH & DEVELOPMENT
7 RUE ANCELLE, 92200 NEUILLY-SUR-SEINE, FRANCE

AGARD CONFERENCE PROCEEDINGS NO. 557

Tactical Aerospace C³I in Coming Years

(Commandes, pilotage, communications, renseignements
tactiques aérospatiaux dans les prochaines années)

Papers presented at the Mission Systems Panel 3rd Symposium held in Lisbon, Portugal from 15th May to 18th May 1995.

North Atlantic Treaty Organization
Organisation du Traité de l'Atlantique Nord