A Doppler-Based Target Classifier Using Linear Discriminants and Principal Components

A. G. Stove
Thales Airborne Systems
Manor Royal
Crawley RH10 9PZ
United Kingdom

andy.stove@uk.thalesgroup.com

SUMMARY

This paper describes the design of the automatic target classifier which has been introduced into the AMSTAR Battlefield Surveillance Radar. It discusses the requirements which have driven the design of the classifier, the data which is used to make the classification, the choice of Linear Discriminant Analysis as one of the classification techniques used and the use of Principal Components Analysis to simplify the training of the discriminator. It also discusses the addition of other classes by the use of other data about the targets. It includes a discussion of the testing of the classifier and the performance achieved.

1.0 INTRODUCTION

The classifier described in this paper has been designed and implemented for the AMSTAR (Advanced Man Portable Surveillance and Tracking Radar) radar. This is a man-portable battlefield surveillance radar, derived from the earlier MSTAR radar [1]. Fig. 1 shows a picture of the radar. The principal features of the radar are:

- small size and light weight
- low power consumption enabling operation from long periods from rechargeable batteries and
- very good Doppler performance.

The small size, light weight and low power consumption allow the radar to be man-portable. The Doppler processing allows it to perform its battlefield surveillance task, detecting slow-moving personnel and other targets, but rejecting clutter. The original MSTAR radar had been in service with the UK and other armed forces since 1989. It was designed primarily as a artillery spotting aid, for locating targets of interest and observing fall of shot. The advanced MSTAR (AMSTAR) variant originated as a Mid-Life Improvement (MLI) of the original design which was undertaken during the year 2000. In common with most, if not all, radars of its general class, MSTAR is a coherent radar and has from the outset possessed a so called 'audio' output which can be used by the operator to classify targets. This audio mode uses a sample-and-hold to 'stretch' the train of pulses returned from a target so that they form a continuous signal. The frequency spectrum of the resulting signal is the Doppler spectrum which would be obtained from CW illumination of the same target. The radar operates in the upper J-band and for this carrier frequency the Doppler frequencies lie within the audio band, being of the order of a kilohertz, and so can be presented as an audio tone to the radar operator, via headphones. After listening to the target for a few seconds a trained operator can classify it with a high degree of accuracy.

Paper presented at the RTO SET Symposium on "Target Identification and Recognition Using RF Systems",
held in Oslo, Norway, 11-13 October 2004, and published in RTO-MP-SET-080.

The Doppler audio frequently sounds surprisingly like the actual sound of the object, perhaps because the same underlying mechanisms modulate both the sound which the target makes and its Doppler spectrum.

The audio-based classifier has three deficiencies, however. The first deficiency is that it must stop scanning and stare at the target which is to be classified, interrupting the execution of its other tasks. The second deficiency is that classification imposes a high workload on operator. He must use the cursor to set the radar to the desired bearing to 'stare' at the target of interest, and must frequently search in range to find the range cell which contains the target of interest. He must then listen intently to the audio to decide the target class. The third deficiency with such a classifier is that the operator must be trained to discriminate the different sounds made by different targets, which require a considerable investment in time and resources.

An automatic classification aid which operates whilst the radar continues to perform its surveillance is thus a significant aid in reducing the operators' workload. The classifier is thus required to operate 'on the fly' during the normal surveillance dwell of the radar so the operator does not have to interrupt normal surveillance operations to obtain a classification - classifications must be generated automatically for all plots and made available on demand.

For the surveillance application, the classifier has to be sufficiently accurate that after observing a target for a few scans the operator can be reasonably confident of its class. It is important to realise that this application does not need the same level of reliability as some other non-cooperative target classification applications (see, for example, reference 2), since the operator can make use of other information, such as target speed and, if desired, the audio mode, to supplement the information provided by the automatic classifier. It should be emphasised again that the purpose of the automatic classifier is to relieve the operators of the workload required to listen to the Doppler signals of each target one at a time. In an environment with few targets, the operators still generally prefer the confidence they obtain by listening to the Doppler from a target, but when there are many targets present, that is not always possible and then the automatic classifier can be used as an aid to help to decide which targets should be examined individually.

2.0 CLASSIFICATION DURING THE SCAN

Since the radar must classify the target as it scans past during its normal surveillance operations, the sample of the Doppler which it can use is limited in length. The dwell during which the classification must be performed is typically less than 100ms. This time is much shorter than that used for audio classification, where the operator can in principle listen to the Doppler for as long as he likes and will typically listen for several seconds. This difference between the 'audio' case and the 'scanning' case means that the techniques which are appropriate for audio recognition are not necessarily appropriate for the automatic classifier. Indeed, the fact that audio-based classification is possible does not in itself prove that classification with a scanning radar is possible.

Jahangir et al. [3] describe a Doppler classifier which uses speech-recognition principles, but they concentrate more on longer audio samples and show that the performance degrades with shorter samples. They point out that a classifier, such as ours, which treats the spectrum as stationary and performs a single transform over all of the data is likely to be sub-optimal for longer data sets, where the spectrum generally changes with time. In our case, however, the shortness of the samples means that the spectrum can reasonably be considered to be stationary. To continue the analogy with speech: conventional speech recognition looks at the chain of successive phonemes, whereas this 'scanning mode' classifier is trying to perform a task which is closer to trying to recognize a single phoneme.

3.0 CLASSIFIER ARCHITECTURE

For the MSTAR MLI, the set of possible target classes between which discrimination was required was:

- Personnel
- Wheeled Vehicle
- Tracked Vehicle

Internally the classifier also generated the classes of

- Unknown
- Reject.

The later versions also includes the classes of

- Helicopter
- Small ship
- Large ship

The original MSTAR MLI variant, which distinguished between Personnel, Wheeled Vehicles and Tracked Vehicles, used a linear discriminant[4] for the classification, with separate templates for each of a number of velocity bands, and used principal component analysis[5] to simplify the training process. This was preceded by pre-processing to organize the data into a form which was suitable for the linear discrimination process. The set of velocity bands ensured that the class of 'Personnel' could not be returned from a target which was moving faster than a running man.

3.1 Rejected and Unknown

It is important to be able to recognize the cases where the classification breaks down. A class of 'Unknown' is returned if the Doppler sideband information is likely to have been corrupted by noise sidebands returned from a much larger clutter return in the same range/azimuth cell. This process is also refined to allow the target itself to be rejected if it may be purely a detection of the transmitter noise sidebands. The stability of the radar is such that this almost never happens, but it is a theoretical possibility in the presence of extremely large clutter returns at relatively short range, and is an easy condition for which to check within the classifier, since all of the preliminary work the has already been done in order to check the validity of the Doppler sideband information.

3.2 Helicopters

As shown in figure 12 of reference 6, the signature of the helicopter blade 'flash' is very distinctive. The approach which has been taken is to look for significant energy in the part of the spectrum which is moving in the opposite sense from the body of the object. Since the blade flash is a fleeting phenomenon, it is not appropriate to process it using the relatively long sampling windows used by the linear discriminant system. The actual detection process is therefore performed in the time domain, where the energy from the blade flash is concentrated, rather than in the frequency domain, where it is diffuse, so the detection process can then be more efficient. The ability to detect the blades, together with the basic radar's ability to detect a crawling man, i.e. a target which is moving extremely slowly, means that the radar now has an effective capability to detect as well as recognise hovering helicopters.

The process of detecting the blade-flashes was found to be significantly more reliable than attempting to use the rotor hub motion as a classification 'feature' in a linear discriminant system.

3.3 Ships

Ships are in effect defined as targets which are not helicopters but which are over the sea. Preliminary studies were made to look at the effectiveness of a Doppler-based classifier for ships. Although under some circumstances the audio obtained from the Doppler from ships can be quite distinctive, it was found to be difficult reliably to incorporate ships into the scheme of the linear discriminant, because of the variability of their signatures.

It may be possible in some cases to determine whether the clutter background is land or sea from its statistical variations and Doppler spectrum, but it was reckoned that, again, it would be very difficult to do this reliably, particularly with a radar like MSTAR which was designed to reject the clutter, so in the end this decision was made using maps and this determines whether a target is assumed to be a land target or a ship. The details of this process are discussed in section 8.

3.4 Architecture

When the classifier was extended to include the additional classes, the decision was taken to move away from the almost pure 'linear discriminant-based' approach originally used. The positive reason for doing this was to take account of additional information, such as radar cross section, which although not a reliable discriminant in itself, does contain useful information. The negative reason was to avoid the possibility of the classification being degraded as more classes are added. For example, if there are N classes then N-1 binary comparisons might be made to determine the target class. If each classification has a probability of p of being correct, then the overall probability of correct classification is only

$$p_0 = p^{N-1}$$

For example, if p = 90% and N =3, typical of the original classifier, then p_0 = 81%. Extending the number of classes to 5 would therefore be expected to reduce the probability of correct classification to only about 66%.

Another desire is to be able, where the resultant update rate is acceptable to the user, to correlate the classification obtained on the same object on different scans, again to improve the performance.

The overall algorithm of the classifier is thus:

if *target is below the clutter sidebands* **then** *Class* **:=** '*Reject*'
    else **if** *Doppler sidebands are masked by clutter sidebands* **then** *Class* **:=** '*Unknown*'
        else **if** *target is small and slow-moving* **then** *Class* **:=** '*Personnel*'
            else **if** *'blade flash' detected* **then** Class **:=** '*Helicopter*'
                else *perform linear discrimination to distinguish between*
                    '*Wheeled*,' '*Tracked*' *and* '*Personnel*'**;**
if *the target is over the sea* **and** *Class* **in** **[**'*Personnel*,' '*Wheeled*,' '*Tracked*'**]** **then**
    if *radar cross section > threshold* **then** *Class* := '*Small Ship*'
    else *Class* **:=** '*Large Ship*'**;**
*compare class with other classifications of the same object***;**
if *there is a clear winner* **then** *return weighted majority class*
else *return 'Unknown.'*

4.0 PRE-PROCESSING THE SIGNAL SPECTRA

The sequence of samples obtained during the radar's dwell is first Fourier transformed to form a spectrum, representing the Doppler shifts of the body of the target and of any parts of the target which are moving independently of the main body at that moment. Figs. 2 to 4 show typical spectra of a wheeled vehicle, a tracked vehicle and a man walking. The clean spectrum of the wheeled vehicle can be contrasted with the much more complex, but asymmetrical, spectrum of the tracked vehicle, and this can again be distinguished from the more symmetrical spectrum of the walking man.

This spectrum is also used to assess whether the target should be 'Rejected' or classified as 'Unknown' on account of the level of clutter present.

4.1. Pre-Processing for the Linear Discriminant Algorithm

Before the spectrum is presented to the linear discriminant classifier it is edited in the frequency domain to remove Doppler cells which it is calculated can only contain noise. Typically 100 Doppler bins are kept to ensure that all significant features of the signal are retained. The spectrum is also regularized by moving its peak, which is assumed to represent the Doppler shift of the bulk of the target, to the centre of the spectrum.

To avoid too great a degree of reliance on the normalization of the spectra with respect to different target velocities, six separate classifiers are used, covering differing velocity bands. The body velocity is estimated from the peak of the spectrum and the signal is presented to the appropriate classifier.

The spectrum is also normalized in amplitude. This normalization presents some theoretical problems which can be solved only by experimentation to provide an acceptable ad-hoc solution: if the spectrum is not normalized then changes in signal level can affect the signature, confusing the classification process. If it is normalized, however, potentially useful information about signal levels is lost. Normalization also assumes implicitly that the amplitudes of all the Doppler components change together, but this has in practice been found to be the more successful approach.

The normalized spectrum is then treated as a vector of features which is passed to the classifier.

5.0 ROBUSTNESS TO SIGNAL TO NOISE RATIO

The Doppler-based classification approach is relatively robust to the signal to noise ratio of the targets. Provided that the signal to noise ratio is sufficiently high for the Doppler sidebands to be distinguished from the noise, the classification performance is substantially independent of that ratio.

Although the power level in any one Doppler sideband is relatively low compared with the body return, the total power in all the sidebands is quite high. The 'classic' radar range equation also means that reduction in classification range compared with the detection range is proportional only to the fourth root of the ratio of 'body' to 'sideband' power, so good classification performance is maintained over most of the radar's detection range against a given target.

The detection process, and likewise the classification process, of the AMSTAR radar is substantially immune to the effects of clutter because:

- low power radars such as this have relatively low dynamic ranges, making good clutter suppression easier,
- the all-solid state design makes the radar very stable

- the fine Doppler resolution makes it easy to distinguish even slow-moving targets (and their Doppler sidebands) from the Doppler sidebands of the nominally-stationary clutter.

6.0 CLASSIFICATION APPROACH

Since the radar must have low power consumption any algorithms used in the real-time classification process must be simple, although complex, time-consuming processes can be used, if required, for off line, training.

6.1. Fisher Linear Discriminant

A statistical classification technique, Fisher's Linear Discriminant Analysis (LDA), was chosen [4]. This provides an optimum discrimination solution for a linearly separable problem. A possible alternative approach would be to use a neural network [6], of which a multi-layer perceptron, trained by back-propagation would perhaps have been the most appropriate form. The principal theoretical advantage of the neural network is that it can cope with problems in which the two classes are not linearly separable. The potential difficulties of training the network mean, however, that it is generally best to try a linear technique first and stick to that if it can give adequate results.

It is sometimes very hard to understand the inner workings of a neural network, although the behaviour of the LDA can also sometimes be hard to understand in a multi-dimensional problem. Understanding how the algorithm is working is of considerable practical benefit when it is to be used in a radar system. It can help to show whether the performance which has been achieved is as good as it should be, whether the pre-processing and the choice of classification parameters have been appropriate and, on a more mundane level, even whether the algorithm has been coded properly.

6.1.1. Choice of Decision Point

As mentioned above, the basic LDA treats the normalized Doppler spectrum as a feature vector, treating each frequency bin as a separate dimension. As is well-known, the algorithm determines the hyperplane which best separates the points in one class from another.

The position of this hyperplane is determined from the statistical distributions of the positions of the returns from the two classes. One 'tunable' parameter is the criterion for the position of this line. Alternative definitions allow it to be, for example, at the point where the probability distributions overlap, so that the most probable answer is given for any particular point in the hyperspace. If the distributions have different covariances, however, this is not necessarily the same as the position which gives equal probabilities of error for the two classes. This point is illustrated in Fig. 5. In practice, a hybrid approach was used in the AMSTAR classifier, which was chosen to give the best balance of performance overall.

6.1.2. Run-Time Complexity

The actual output of the training stage is the vector representing a normal to the discriminant line, or, more generally, a normal to the discriminant hyperplane. The test of a particular spectrum is performed by calculating the scalar product of the feature vector and this reference normal, which projects the point in the hyperspace onto that normal. The choice of class is made by determining to which side of the discrimination hyperplane the point falls. The run-time complexity is thus one multiplication and one addition per element of the spectrum plus one comparison, which is all very simple compared, for example, with the Fourier transformation required to generate the spectrum in the first place.

6.2 Principal Component Analysis

The amount of training data required increases with the number of elements in the vector. As mentioned above, typically 100 elements are used in the vector, but it is clear that the number of independent 'features' in the data is far fewer than that. If the dimensionality of the test vector can be reduced to reflect the number of degrees of freedom in the underlying data then the amount of training data can be reduced. A well-known way of doing this is to use the Principal Components Analysis (PCA) technique [5], also known as the Karhunen-Loeve transform..

As is well known, the principal directions found by the PCA are the eigenvectors of the autocorrelation matrix of the data, and the significance of each dimension is indicated by the relative value of the associated eigenvalue, so the simplification of the problem is achieved by retaining only a relatively small number of the most significant components.

In general the PCA complements rather than replaces the LDA, because the PCA is not looking for directions which optimize the between-class discrimination, but rather for directions which explain as much as possible of the variance within the union of both classes.

It is not commonly recognized that it is theoretically possible that the most significant directions chosen by the PCA could *only* explain components of the covariance which are common to both classes, in which case if the other 'less significant' dimensions are discarded, it would no longer be possible to discriminate between the classes. Figure 8 illustrates how this can occur by taking a hypothetical example where the data contains two dimensions. The ellipses represent the distributions of the two classes in these dimensions. It can be seen that most of the variance in this case is in dimension 1, whereas all the discrimination information is in dimension 2. A 'degenerate' PCA which, took only the first principal component, would therefore in this case discard all the discrimination information. Although this is clearly an extreme, simplified, case it will be appreciated that similar effects can in theory happen in more complex, practical, cases.

In practice, however the principal dimensions do usually contain most of the classification information. This has been specifically checked for this classifier by showing that if the principal components are *excluded* from the classification vector, very poor discrimination is achieved between the classes.

Removal of only a few dimensions, with the lowest eigenvalues, on the other hand, can sometimes significantly help the discrimination process by eliminating dimensions which explain virtually none of the variance, but which serve only to compromise the stability of the algorithm. It is sometimes difficult to relate some of these phenomena directly back to the original data set, but dimensions which explain nothing of the variance are probably due to artefacts in the pre-processing.

Although the directions chosen by the PCA contain most of the discrimination information, the principal direction is in general not aligned exactly with the normal to the original discrimination plane. In order to obtain good discrimination the PCA is thus followed by an LDA process as described previously. The PCA is thus used during the training process as a 'pre-processor', to reduce the dimensionality of the problem. About ten principal components are typically retained, leading to a ten-fold reduction in the amount of training which must be obtained.

It is important to emphasize that the PCA does not affect the complexity of the real-time classification, since the dimensionality of the reference vector can be increased again after training to match the size of the data vector, so that no additional real-time processing of the data is required.

7.0 TESTING AND PERFORMANCE

7.1 Tests during Development

During the training phase, 10% of the available data, chosen at random, was set aside for testing, i.e. 90% was used to train the discriminator but the other 10% was only used to test it. The individual types of the vehicles used to gather the training data, their registration numbers and the identities of the personnel who acted as the 'targets' were all recorded to allow checks to be made that the classifier was not accidentally discriminating between the identities of individual targets when it should generalize for broad classes.

Tests were made of the correlation of false alarms between adjacent data samples to assess what time delay was needed before samples of training data could be considered to be independent. It was concluded that, very approximately, samples of personnel which were separated by 50ms could be considered independent whereas the delay had to be 100ms for wheeled vehicles and 150ms for tracked. Since each sample was about 50ms along, all the 'personnel' data could be used to provide independent training samples, but only alternate samples of the wheeled vehicle data and one in every three of the tracked vehicle data. This 'thinning out' of the data meant that we avoided the possibility that some of the 'test' data could have been highly correlated with particular adjacent examples of the training data.

During the training phase, tests were carried out to assess the sensitivity of the classifier to the aspect of the targets, by training a discriminator using only data gathered at one nominal aspect angle and testing it with data gathered at another. As expected, it was found that the Doppler signature is relatively insensitive to changes in aspect angle. The signatures remained correlated over about 30 degrees, although large changes in aspect can causes significant changes in the signatures, probably due to obscuration of parts of targets.

The sensitivity to variations in range was tested by adding noise to the test data to simulate a reduction in the signal to noise ratio. The sensitivity to changes in velocity was tested by varying the number of velocity bands and checking that this did not significantly change the performance.

7.2. Acceptance Tests

Proposing a test plan for the classifier can present a significant theoretical problem if one does do not know *a priori* what factors, such as target range or speed or direction, might affect its performance. In that case field tests would have to cover all possible target conditions and would be prohibitively expensive. The problem of testing the classifier can, however, be made manageable by making use of the tests which have been performed during the development to assess the sensitivity to aspect angle, range, velocity etc. This allows a suitable range of tests to be devised, but there may still be a relatively large number of combinations to be tested.

For AMSTAR most of the testing was performed using the 'test set' extracted from the training data, as described above, which had been used for the tests performed during the training. This contained enough data to obtain statistically significant results in different scenarios to ensure that the classifier's performance was satisfactory in all combinations of scenarios. Since the whole training and testing system operated in software, the process of obtaining the results in these laboratory based tests was also many times faster than reading each result individually from the radar display and recording it by hand. Although this procedure is efficient, it suffers from the fact that it is still dependent on the correctness of the sensitivity analysis carried out during the training phase and can only test the classifier against data gathered in the same scenarios as the training data. A relatively small number of completely independent 'field' tests were therefore performed to give confidence that the results of the 'laboratory' testing were reliable. Whilst these could confirm that the classifier worked in the scenarios tested, it could not give statistically significant information on a large number of the scenarios individually.

The overall classification performance obtained for the three-class case used in the MLI variant has already been reported before [8]. The comparable results with four classes plus 'Unknown' are:

Table 1: Overall Classification Accuracy for the AMSTAR Classifier

True Class		Reported As			
	Personnel	Wheeled	Tracked	Helicopter	Unknown
Personnel	83%	14%	3%	1%	0%
Wheeled	0%	82%	10%	5%	4%
Tracked	0%	15%	77%	8%	0%
Helicopter	3%	6%	4%	82%	6%

The totals are not exactly 100% due to rounding errors. It will be appreciated that these results are for a single observation of the target and the results seen by the operator will be better, once successive looks have been compared as discussed below.

It will be noted that, compared with the three-class case, the performance on tracked vehicles has dropped slightly below 80% for this 'single look' case reported above, although the change in the numbers is probably not significant. It is noteworthy that these results are still generally better than those reported in [3] for short dwells, although better results were obtained with the hidden Markov models when longer dwells were available.

The corresponding field trials of this 'four class' classifier showed correct classifications of between 73% and 80% on a single 'look', which, is in equivalent to a value of the order of 90% after the scan-to-scan comparisons, and is thus comparable with the results reported in reference 8. The field trials also showed that, as predicted, correct classification was achieved at ranges of the same order as the detection range.

8.0 POST PROCESSING

8.1. Scan to Scan Processing

The classifier follows the radar's plot extractor. The output of the classifier is used to tag the plot when it is sent to the radar's Control and Display Assembly. The operator accesses the classification by putting the cursor near the plot. This placement does not have to be as precise as is needed for audio classification. If the nominated plot is part of a trail, the display processing can be arranged to look at the classification of adjacent plots to reject occasional false classifications or to report 'unknown' if the classifications are inconsistent. This process ensures that the accuracy after a few scans is at least as 80% for any combination of target type and velocity band, i.e. it gives good results not only in 'average' scenarios but in the worst case. In other variants, an automatic tracker can be used to associate the plots together, and the individual classifications within the track can be compared to eliminate the occasional erroneous classifications.

If the classification is inconsistent, this is used in some variants as an additional 'trap' which also allows the class to be reported as 'Unknown.' User feedback suggests that operators are much happier with a relatively high proportion of returns being classed as 'Unknown' than they are with firm classifications which are in fact erroneous.

8.2. Sea Target Classifications

It is also at this stage that the classification of sea targets is incorporated. The map in the MMI is used to ascertain whether the target is on the land or on the sea. If it on the sea and not a helicopter or an 'unknown,' it is classified as a ship. Ships are distinguished between 'large' and 'small' on the basis of their radar cross sections. The ratio of radar cross section between 'small' and 'large' targets is 50:1 (17dB) and by placing the threshold 5dB above the limit for a 'small' target there is a 94% chance that a target with a Swerling 2 fluctuation and an RCS at the limit of the values to be classified will be classified correctly.

9.0 CONCLUSION

An automatic classifier has been successfully incorporated into the AMSTAR Battlefield Surveillance Radar. It uses Linear Discriminant Analysis to exploit the Doppler signatures of the targets to provide classification between the classes of Personnel, Wheeled vehicle and Tracked vehicle on land, Small and Large Ships on the sea and Helicopters in both domains. This provides classification aids which can significantly reduce the operator's workload without compromising the radar's surveillance mode performance. A performance equivalent to better than 90% average and better than 80% worst-case after scan-to-scan comparisons, equal to that reported in reference 8 for the three-class case has now been achieved for the four-class case. The technique of Principal Components Analysis has also been used to minimize the amount of training data required.

10.0 ACKNOWLEDGEMENTS

The Mid Life Improvement of the MSTAR radar, which led to the evolution to AMSTAR, was funded by the UK Ministry of Defence. The initial design of the algorithms was performed by Mr Rick Rickman and Mr Stephen Sykes of Racal Research Laboratories (now Thales Research and Technology). The incorporation of the Helicopter and Ship classes was designed and executed by Mr Graeme Evans and Mr Alastair McIver of Thales Airborne Systems. Much of the exploration of the behaviour of the PCA with this data set was carried out by Mr. Simon Gibbs.

11.0 REFERENCES

[1] C. J. Baker and B. D. Trimmer, "Short-Range Surveillance Radar Systems," *Electronic Commun. and Engineering J.*, August 2000, pp.181-192

[2] R. van der Heiden, L. J. van Ewijk and F. C. A. Groen, "A Comparison of Radar Range Profiles Between In-Flight Measurement and RCS Prediction," *Proc. NATO RTO Meeting on "Non-Cooperative Air Target Identification Using Radar 22-24 April 1998,* RTO Meeting Proceedings 6, published November 1998.

[3] M. Jahangir, M., K. M. Pointing and J. W. O'Loghlen, "A Robust Doppler Classification Technique Based on Hidden Markov Models," *Proc. Radar 2002*, IEE Conf. Publ. 490, October 2002, pp.162-6.

[4] P. A. Devijver and J. Kittler, "Pattern Recognition: A Statistical Approach," London, Prentice-Hall, 1982.

[5] D. F. Morrison, "Multivariate Statistical Methods," New York, McGraw-Hill, 1978.

[6] P. M. Grant, "Artificial Neural Network and Conventional Approaches to Filtering and Pattern Recognition," *Electronic Commun. and Engineering J.*, September/October1989, pp.225-232

[7] M. Moruzzis and N. Colin, "Automatic Recognition of Air targets for Future SHORAD Radars," *Proc. NATO RTO Meeting on "Non-Cooperative Air Target Identification Using Radar 22-24 April 1998,"* RTO Meeting Proceedings 6, published November 1998. pp.30-1 to 30-10.

[8] A. G. Stove and S. R. Sykes, "A Doppler-Based Automatic Target Classifier For A Battlefield Surveillance Radar," *Proc. Radar 2002*, IEE Conf. Publ. 490, October 2002, pp.419-423.

Figure 1: The AMSTAR Radar

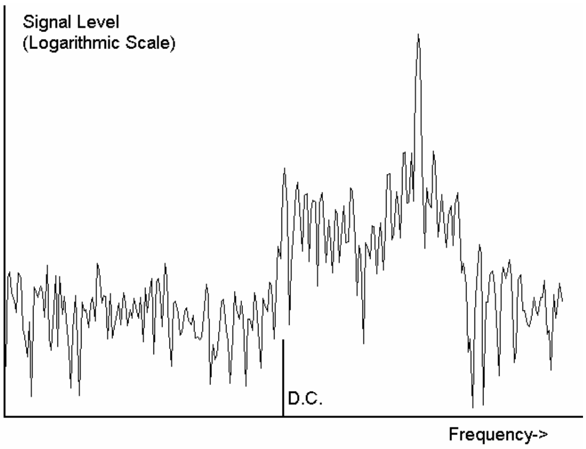

Figure 2: Spectrum of a Tracked Vehicle

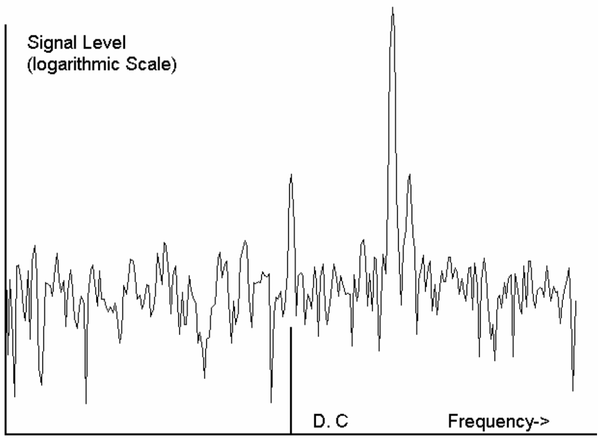

Figure 3: Spectrum of a Wheeled Vehicle

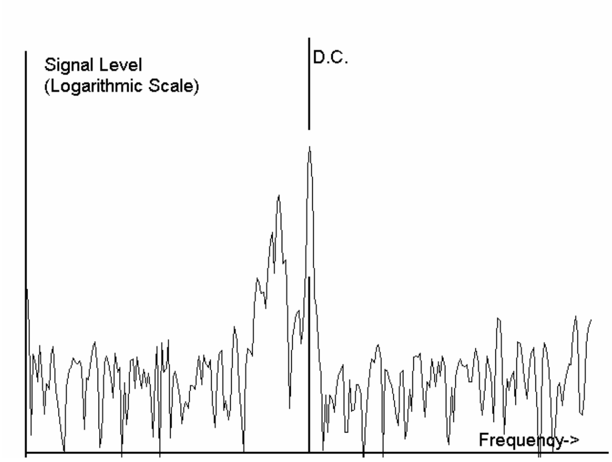

Figure 4: Spectrum of a Walking Man

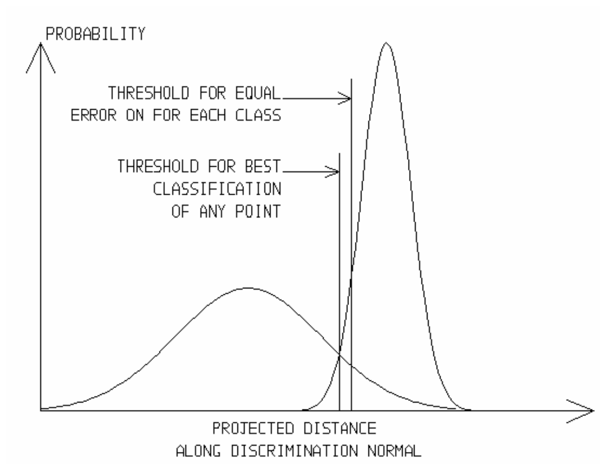

Figure 5: Different Discrimination Choices for the Linear Discriminant Classifier

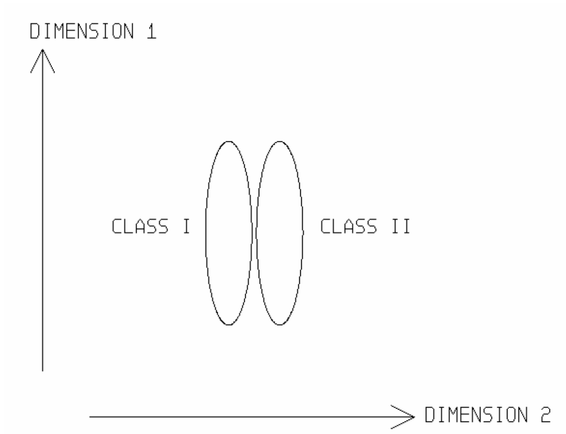

Figure 6: Hypothetical a Case when PCA would not work with LDA