/ 9940/9748

N94-24221

ACOUSTIC TARGET DETECTION AND CLASSIFICATION USING NEURAL NETWORKS[1]

S14-71

201686

P_10

James A. Robertson
Mark Conlon
IIT Research Institute
4140 Linden Ave., Suite 201
Dayton, OH 45432

SUMMARY

This research demonstrates a neural network approach to the classification of acoustic emissions of ground vehicles and helicopters. Data collected during the Joint Acoustic Propagation Experiment conducted in July of 1991 at White Sands Missile Range, New Mexico was used to train a classifier to distinguish between the spectrums of a UH-1, M60, M1 and M114. An output node was also included that would recognize background (i.e. no target) data. Analysis revealed specific hidden nodes responding to the features input into the classifier. Initial results using the neural network were encouraging with high correct identification rates accompanied by high levels of confidence.

INTRODUCTION

The strong and definable acoustic emissions from ground vehicles, helicopters and aircraft make systems employing acoustic sensing attractive. Sources such as engines, tracks, rotor systems and propulsion systems generate emissions which acoustic sensors can use to determine target line of bearing, range and identification. These sensors can provide passive detection at relatively large distances without the line-of-sight restrictions radar systems impose.

The fidelity of an acoustic target classifier becomes crucial in applications such as identification friend or foe (IFF), border monitoring and smart mines. It is vital that the identification is correct with a high level of confidence. Traditional approaches to designing a classifier consist of extracting a number of candidate features from a training set from which a final feature set is selected for the logic design. The performance of the classifier depends upon how closely the test or recall database resembles the training database. If the classifier does poorly, the database could be extended to include more data; however, this could lead to a situation where individual classes might not be separable. In general, traditional classifiers will

[1] Work performed on contract at IIT Research Institute, DLA900-86-C-0022

do well over test databases which use training databases which encompass the range of target conditions anticipated.

Unfortunately, databases are rarely this comprehensive. The signature variations due to the environment, terrain, vehicle maintenance, and other dynamic conditions are difficult to predict and impossible to fully characterize.

An exceptional classifier should be flexible, robust and be able to cope with varying levels of noise and still correctly identify most target samples. It should be able to deal with a complex system which may not be fully understood. Most importantly, it must be able to generalize from a limited amount of training data and maintain good performance on data which may contain only some similarities to the training set.

The difficulty of the problem suggests that a neural network (NN) may provide a viable solution.

NEURAL NETWORK OVERVIEW

A NN is a system which mimics the computational ability of biological systems. They consist of large numbers of interconnected neurons (nodes). These neurons take data from sensors or other neurons, perform simple operations on the data and pass it on to another neuron.

One of the most popular networks for applications is backpropagation(BP). BP is a multi-layer feed forward NN. The meso-structure of a typical three layer feed forward NN is shown in Figure 1. These layers are referred to as the input, hidden and output layers. The interconnection between the ith input node and the hth hidden node is referred to as W_{ih}, whereas the interconnection between the hth hidden node and the jth output node is referred to as W_{hj}. A set of features is applied to the input node, then the NN processes this data calculating the activation levels of the hidden and output nodes. The output of a neural network used for classification may be referred to as the class activation level. The number of input features determines the number of input nodes. The number of output nodes is determined by the number of target classes. The number of hidden nodes, and if necessary, hidden layers, is generally application specific.

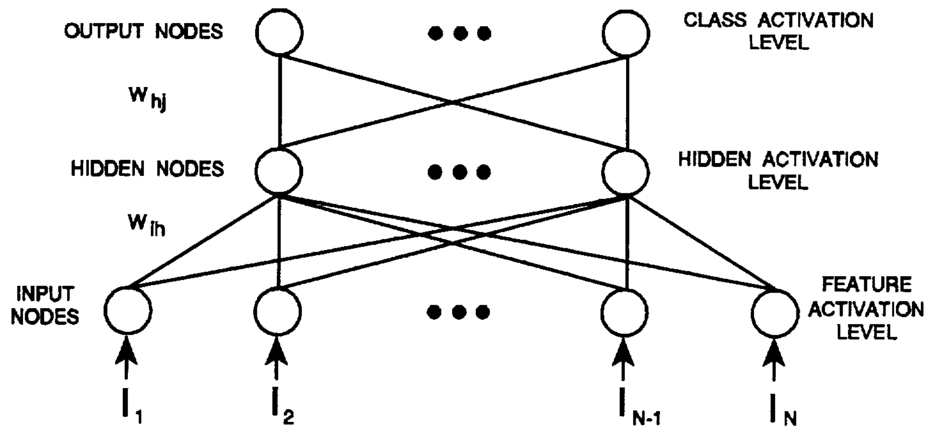

Figure 1. Meso-structure of a multi-layer feed forward NN

Figure 2 illustrates the processing that occurs at the neuron level in a NN. After summing the input values multiplied by the interconnections plus the jth nodes threshold associated with it, a transfer function is used to scale the neuron's responses to incoming signals. Many types of transfer functions exist including threshold-logic, hard-limit, continuous-function and radial basis. Two of the more common continuous transfer functions, the sigmoid and modified sigmoid, are shown. A sample calculation of the jth activation level is also shown.

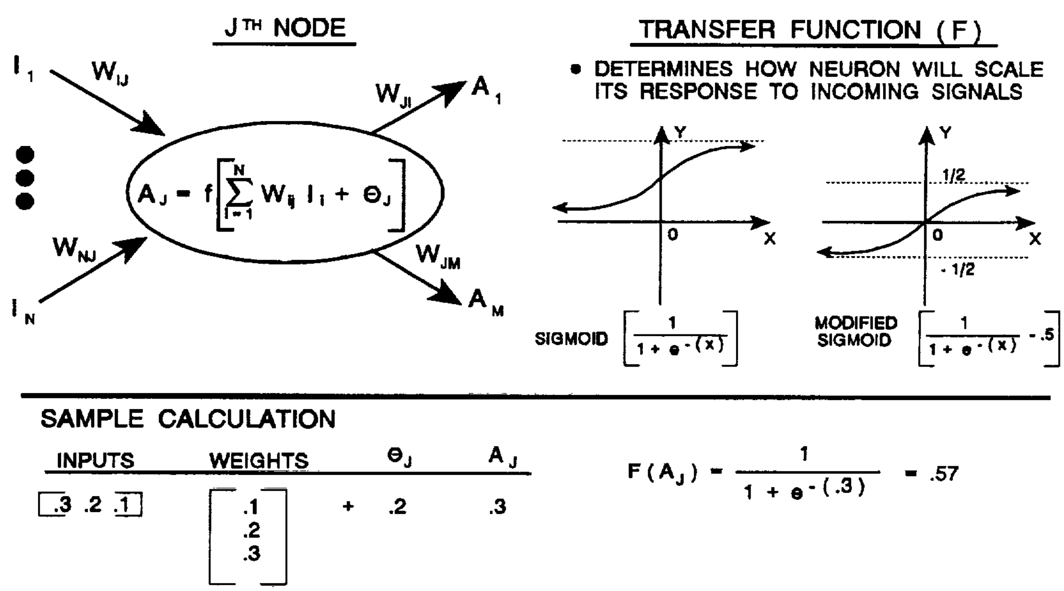

Figure 2. Micro-structure of a NN

BP uses a generalized delta rule for learning. This rule allows the error to affect all layers of the interconnection weights. The method of learning is supervised where actual training data is used. Initially, the weights and threshold are randomized to small values, usually between 0 and 1 or -0.5 and +0.5. Adjustments to the weights and thresholds in all layers are made according to the difference between the desired output activation level and the actual activation level as shown in equation 1.

$$\mathbf{ERROR} = C_j \,(1 - C_j)\, (C_j^k - C_j) \quad (1)$$

where:

c_j **= actual output of jth node**

c_j^k **= desired output of jth node**

The advantages of NNs have been reported by many researchers (ref. 1). The most attractive reason for using a NN, particularly for target classification, is its ability to generalize. A NN has the ability to generalize and find similar features to that of the training database. For a classifier to be successful in an unknown and poorly characterized environment, it must have the ability to generalize. Another advantage is a NNs' ability to store and distinguish many patterns. This is alluring as both the number of classes and the variability within the class increase.

Researchers have also noted some limitations and disadvantages using a NN. BP in particular suffers from lengthy training sessions. There are ways to reduce the training time by adding momentum (ref. 2) , scaling inputs, thresholds and weights, and adapting the learning rates after each iteration(ref. 3). However, even after optimizing for speed, training sessions may still be lengthy. A NN is specific for a certain application. After it has been trained to identify n classes, adding a new class, n + 1, requires retraining the NN. Another disadvantage is the difficulty associated with selecting the number of hidden nodes in a NN. General formulas to determine the number of hidden nodes (e.g. Lippman ref. 4, Hecht-Nielson ref. 5) may help for an initial guess; however, the particular application appears to be the driving force. Selecting too many hidden nodes may cause the NN to memorize the input patterns as opposed to generalize. Selecting too few hidden nodes may yield an unstable NN incapable of forming complex decision regions.

EXPERIMENT

Acoustic data was collected at White Sands Missile Range, Dirt Site in July, 1991 during the Joint Acoustic Propagation Experiment(JAPE) by personnel from MIT Lincoln Laboratory. All data was lowpass filtered at 670 Hz and sampled at 2kHz at MIT Lincoln Laboratory. Single channel data was selected from six different trials notated by trial numbers as shown in Table 1.

Table 1. JAPE data set

JAPE TRIAL NUMBER	TARGET	DESCRIPTION
015507	UH1	100 knots, 150 m Alt.
080507	M1	20 mph
092508	M60	20 mph
115509	M114	15-20 mph
095amb	none	background
084508	M60	idle, 750 rpm

The selected data was segmented into 1 second samples, and Hanning windowed. The power spectrum was then estimated for each sample. The amplitude values from 1 Hz to 150 Hz were used as input into the NN. The NN meso-structure consisted of 150 input nodes, 80 hidden nodes and 5 output nodes. The output nodes represent each target class: UH1, M1, M60, M114 and background (no target). In order to increase convergence all inputs, weights, thresholds and outputs were normalized between -0.5 and +0.5. The error term was adjusted to properly apply the modified sigmoid transfer function as shown in equation 2.

$$\mathbf{d_j = (c_j + 1/2) (1/2 - c_j) (c_j^k - c_j)} \quad (2)$$

where:

$\mathbf{c_j}$ **= actual output of jth neuron**

$\mathbf{c_j^k}$ **= desired output of jth neuron**

A momentum term was used to decrease oscillations and decrease training time. All training continued until the rms error over the entire training set was less than 1%. Approximately 20-30% of the data set was not used in training sessions but saved for effectiveness testing.

RESULTS

A closer look at the hidden activation levels may provide insight into the operation of the NN. The hidden activation level is the actual output of a hidden node. Ideally the knowledge stored in the hidden layer is abstracted from the information contained in the input pattern. A wide variety of features can be represented in the hidden layer. This layer often shows which hidden nodes become activated in response to a particular input pattern.

Figure 3 shows sample input, hidden activation level and class activation level when the target was a UH-1. The NN was able to map the differing inputs into a relatively invariant set of hidden activation levels and class activation levels. Comparing the hidden activation levels for the UH-1 target to the hidden activation levels from other target class samples revealed that certain nodes were responding to the input patterns. The 22nd and 32nd hidden nodes appeared to be most useful for distinguishing the UH-1 from the M60. The 20th hidden node appeared to be most useful for distinguishing between the UH-1 and the APC, whereas the 40th node was the most useful for distinguishing between the UH-1 and the M1.

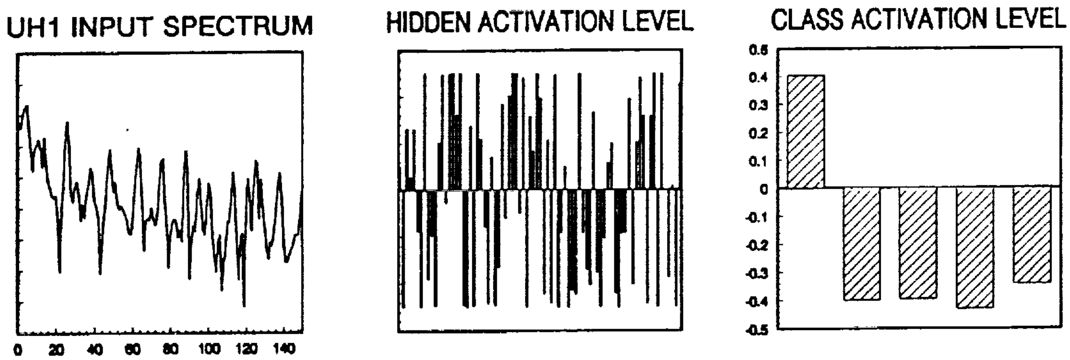

Figure 3 UH1 NN input, hidden and class activation levels

Figures 4 through 7 show samples of the M1, M60, APC and background NN results. A similar hidden node analysis was done to yield the distinguishing nodes as listed in Table 2.

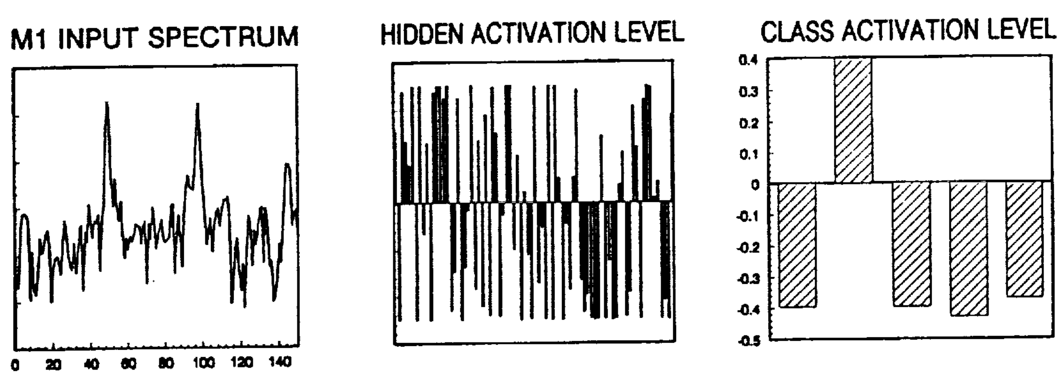

Figure 4 M1 NN input, hidden and class activation levels

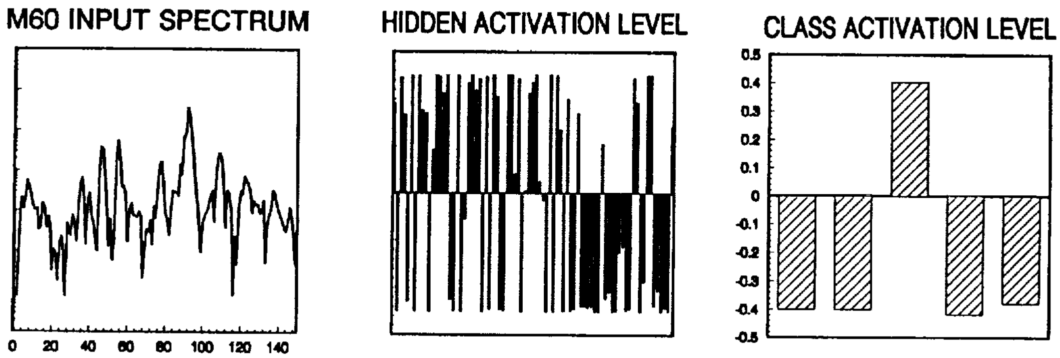

Figure 5 M60 NN input, hidden and class activation levels

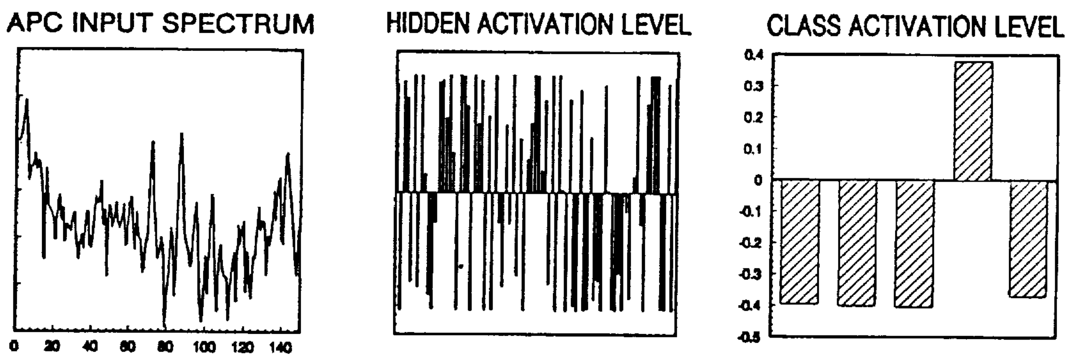

Figure 6 M114 NN input, hidden and class activation levels

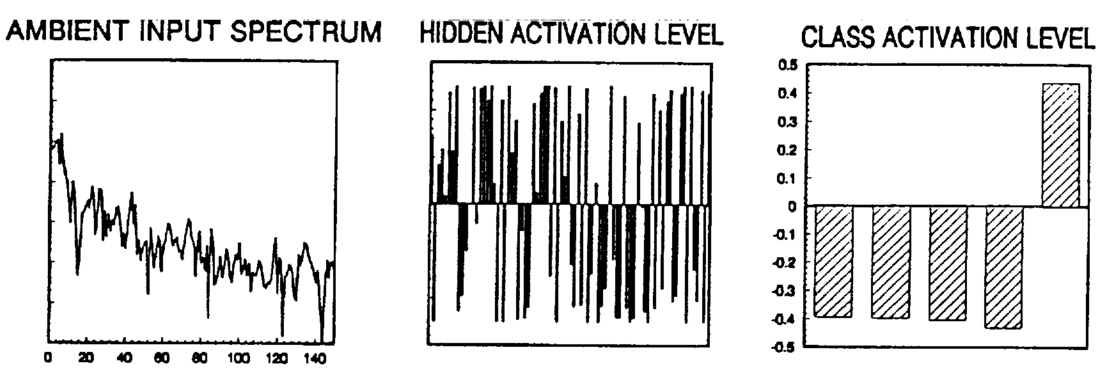

Figure 7 Ambient NN input, hidden and class activation levels

Table 2. Hidden nodes for class separation

Target	UH-1	M1	M60	APC	Ambient
UH-1	-	19,40,74	18,19,22,32, 52,74,75	18,19,20, 52,74	-
M1	19,40,74	-	40	20,40	17,29,40,64, 74,78
M60	18,19,22,32, 52,74,75	40	-	20,22,75,78	18,22,29,32, 52,56,64,68, 74,75,78
APC	18,19,20, 52,74	20,40	20,22,75,78	-	18,20,29,52, 74
Ambient	-	17,29,40,64, 74,78	18,22,29,32, 52,56,64,68,. 7,4,75,78	18,20,29,52, 74	-

Analysis of the hidden nodes also revealed that some nodes did not assist in the classification of any of the targets. Hidden nodes, numbers 2, 16, 34, 46, 73 and 79, yielded the same hidden activation level for all inputs. This suggests that the NN could have learned the same amount of information with less hidden nodes.

The test set was used to determine overall correct classification. Results showed greater than 98% classification for all classes. A system user may want to know how confident an identification is at a particular time. Confidence levels were calculated for each class by using the difference of the highest activation level and the second highest class activation level divided by the maximum activation level difference. Values should range between 0.0 and 1.0. Ideally confidence levels should be high for correct identifications and low for incorrect identifications. The confidence levels of the NN shown in Figure 8 adhere to these guidelines. Notice that for each of the classes, if the NN identification was correct the confidence level was 0.9 or above. However, when the NN identification was incorrect the confidence level was 0.6 or below.

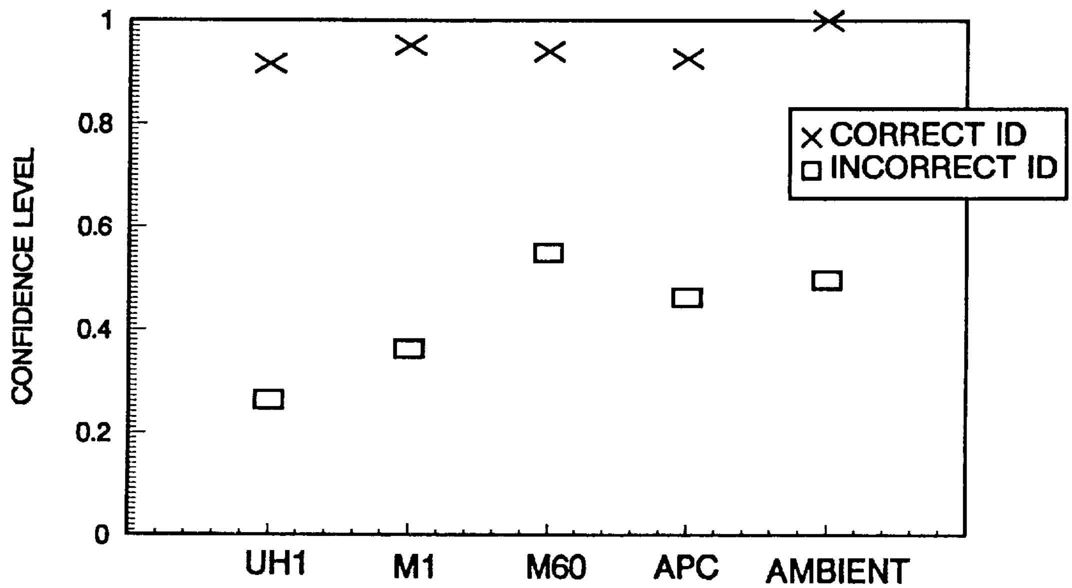

Figure 8 Confidence levels for a trained NN: correct ID vs incorrect ID

CONCLUSIONS

A NN has been used to successfully identify the acoustic emissions of ground vehicles and helicopters. Initial analysis indicates that a high level of confidence can be associated with the identification using a NN classifier. The hidden node analysis demonstrated that the hidden layer is distinguishing between classes using the target specific input features. The analysis also indicated that a smaller number of hidden nodes would suffice for this particular example. The use of ambient or background data as an output class could prove quite useful in determining when no target is present.

A NN trained using a fairly large database could improve the classification performance of existing acoustic sensors. The generalization capability characteristic of a NN will enhance the performance of acoustic sensors.

REFERENCES

1. Maren, A., Harston, C., Pap, R., Handbook of Neural Computing Applications, Academic Press, Inc., San Diego, 1990.

2. Rumelhart, D., McClelland, J., Parallel Distributed Processing, MIT Press, Cambridge, Massachusetts, 1988.

3. Fahlman, S.E., " An empirical study of learning speed in backpropagation networks", Carnegie Mellon Technical Report, CMU-CS-88-162, June 1988.

a-3

4. Lippman, R.P., " An introduction to computing with neural nets", IEEE ASSP magazine, April 1987.

5. Hecht-Nielsen, R., "Kolmogorov's mapping neural network existence theorem", Proc. First Int'l Joint Conf. Neural Networks, San Diego, CA, June 1987.

6. Marple, S., Digital Spectral Analysis, Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 19