AIAA JOURNAL
Vol. 30, No. 3, March 1992

Identification of Helicopter Noise Using a Neural Network

Randolph H. Cabell,* Chris R. Fuller,† and Walter F. O'Brien†
Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24061

An artificial neural network (ANN) has been trained to distinguish between the noise of two helicopters. The performance of the ANN is compared with that of a conventional recognition system. The conventional system uses the ratio of the main-rotor blade passage frequency (bpf) to the tail-rotor bpf. The ANN was trained to use similar main/tail-rotor information, in addition to information describing the distribution of spectral peaks of the main rotor. It is shown that this additional information allows the ANN to distinguish between the helicopters when tail-rotor noise is removed from the spectrum. The performance of the two methods is given as a function of signal-to-noise strength, and propagation distance, using a model of atmospheric sound propagation. The conventional method outperforms the ANN when main- and tail-rotor noise are present, but the conventional method cannot identify helicopters when tail-rotor noise is removed. At 20-dB signal-to-noise ratio (SNR), when tail-rotor noise is not present in the spectrum, the ANN correctly identifies the helicopters 100% of the time, compared to 50% for the conventional method. The performance of the ANN drops as signal strength decreases. At 8-dB SNR, the ANN is correct 77% of the time, while at 0 dB it is correct 58% of the time. Similar results are obtained for the performance when the signal is propagated through the model of the atmosphere.

Nomenclature

E = weighted sum of input values
H = harmogram value
P = sigmoid function
S_x = background noise
s = output of neural network unit
u = iteration number
w = interconnection weight
Y = noise spectrum
α = smoothing parameter for weight updates
δ = error gradient
Δw = weight update
ε = learning rate
μ = mean value
ω = frequency
σ = standard deviation

Subscript

i = unit number in neural network

Superscripts

n = layer in neural network, where n = 3 is output layer
$*$ = target output value

Introduction

ARTIFICIAL neural networks (ANNs) are becoming more commonplace due to their processing power and promise of processing speed. Loosely based on the structure of parts of the human brain, ANNs are characterized by massively parallel collections of simple computational elements. An individual computational element sums input values, then produces a nonlinear output. Large assemblies of these simple elements can solve problems requiring massive constraint satisfaction.[1] ANNs have the additional advantage of learning the optimal connection weights between processing elements. This learning process eliminates the tedious programming that often accompanies complex problems.[1]

Several interesting applications of ANNs involve processing of acoustic-type signals. Examples can be found in speech recognition,[1-3] where it has been shown that ANN-type structures can learn relevant acoustic-phonetic attributes of speech.[1] Other acoustic applications include the recognition of whale calls,[4] the modeling of dolphin echolocation,[5] and the analysis of sonar signals.[6] Numerous examples can be found where ANN architectures are used to classify patterns. A wonderful example of the learning power of an ANN is demonstrated in the article by Sejnowski and Rosenberg.[7] They describe an ANN that is taught to convert text to speech.

These applications all require the evaluation of many hypotheses at a high rate of speed. The text-to-speech converter must learn a variety of phonetic rules such as hard and soft consonants, then make use of these rules instantaneously when pronouncing text. A speech recognizer must also evaluate many possible responses in a very short time. The superior performance of ANNs in these areas has been demonstrated.[1,7]

Conventional methods of pattern recognition have been applied to acoustic identification problems. Examples can be found in helicopter recognition,[8] and the more general case of distinguishing between classes of vehicles.[9,10]

Although ANNs have an inherent appeal due to their "intelligent" nature, they may not be suited for every problem. Conventional algorithms can outperform ANN methods in some areas. The best applications for ANNs require very fast processing, or problems that involve incomplete or missing information. A conventional method of helicopter identification uses the ratio of the main- and tail-rotor frequencies of the helicopter, comparing the calculated value with a table of known values.[8] This method will work as long as tail-rotor noise is measured at the receiver. Tail-rotor noise will not be measured if the helicopter does not have a tail rotor, or if the higher frequency tail-rotor noise is attenuated by the atmosphere. These situations have driven the search for a new method of helicopter identification.

An ANN recognition system requires features, or descriptive characteristics, in order to distinguish between pattern

Presented as Paper 90-3973 at the AIAA 13th Aeroacoustics Conference, Tallahassee, FL, Oct. 22–24, 1990; received March 9, 1991; revision received July 17, 1991; accepted for publication July 17, 1991.

*Research Associate; currently On-Site Contractor, NASA Langley Research Center, Hampton, VA 23665.
†Professor. Member AIAA.

classes. However, an ANN can still classify patterns if a feature is approximate or missing.[11] In order to do this an ANN needs information from several sources. When information from one source becomes ambiguous, other sources of information are available for classifying the patterns. In helicopter recognition, as long as some useful feature type exists that is independent of tail-rotor noise, then the ANN will still be able to identify helicopters if tail-rotor noise is removed.

The purpose of this work is to demonstrate the ability of an ANN to identify helicopters regardless of the presence of tail-rotor noise. The ANN is taught to identify helicopters using two types of characteristics, or features. One type is associated with the ratio of the main-rotor to tail-rotor bpf. The other type describes the distribution of peaks in the main-rotor spectrum, which is independent of the tail rotor. It is shown that the ability of the ANN to identify helicopters compares with that of the conventional method when main- and tail-rotor noise are present. The addition of the spectral shape feature improves the performance of the ANN, compared to that of the conventional recognition method, when tail-rotor noise is absent. Identification performance of the methods is given at various signal-to-noise ratios and as a function of simulated propagation distance.

In the first section of the paper we describe the ANN and conventional methods. The next section contains a description of the recognition experiments, including the simulated helicopter noise and model of the atmosphere. The results are given in the third section, and discussed in the fourth section. The conclusions are given in the fifth section.

Identification Methods

Helicopter Identification

An artificial neural network is a collection of simple nonlinear processing elements operating in parallel. The arrangement of the elements is loosely drawn from the structure of biological neural networks.[12] Each element computes a weighted sum of its inputs, which is then passed through a nonlinearity. The nonlinearity is often a sigmoidal function, as shown in Fig. 1, chosen for its convenient mathematical properties.[13] The external inputs to the ANN are called features. Features are quantities that describe the patterns, which in this case are helicopter acoustic signatures. The feature values are assembled into a feature vector, so that all patterns are then described by their corresponding feature vector.[14]

The ANN designed for helicopter identification has three layers of processing elements, shown in Fig. 2. It has been shown that a three-layer neural network can form arbitrarily complex decision functions.[15] Processing elements in the first layer, called first-layer hidden units, form weighted sums of input feature values according to Eq. (1)[7]:

$$E_i^{(n)} = \sum_j w_{ij} s_j^{(n-1)} \tag{1}$$

The nonlinear output of each unit is computed as

$$s_i = \frac{1}{1 + e^{-E_i}} \tag{2}$$

where s_i and E_i refer to the same unit in the network.

The input feature values are numerical characteristics describing helicopter noise and are discussed in the second section of this paper. The outputs of the first-layer hidden units are the inputs to the second layer of hidden units, whose outputs in turn are input to the output units. The decision rule on the output layer is to choose the class corresponding to the unit with the highest output value.[12] Each unit on a layer is connected to every unit on the surrounding layers, and connected to no units on the same layer. All outputs are thus only fed forward; hence the name feed-forward, fully connected ANN.

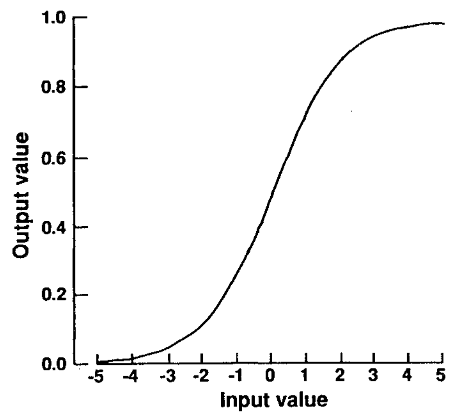

Fig. 1 Sigmoid activation function.

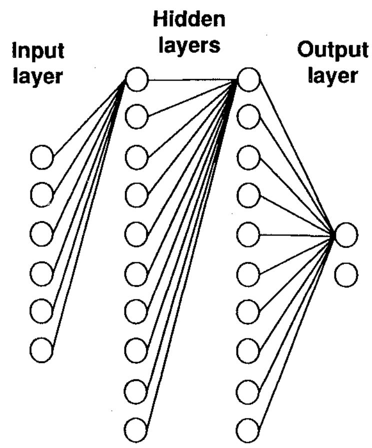

Fig. 2 ANN for helicopter identification.

The ANN used for helicopter identification has six features at the input layer, 10 units in each of the hidden layers, and two units at the output layer. The number of units in the input and output layers is set according to the number of input features and number of classes, respectively. The number of units in the hidden layers is determined by a trial-and-error method. The number of units must be large enough to be able to solve the given problem, but should not be so large that the interconnecting weights cannot be determined from the training data.[12] For this application, the hidden layers contain 10 units each.

The 10 units per hidden layer in the chosen network could be excessive for this application. The size of the network was chosen after comparing the performance of different-sized networks. The chosen network had a faster convergence time than smaller networks. Although convergence time can be an important factor in some applications, it is actually of secondary importance in this application. A more useful selection strategy would be to minimize the number of hidden layer units while maximizing recognition performance on a test data set. This strategy would most likely result in a network with fewer hidden units producing the same recognition performance as the larger network.

The experience of the authors is that the trade-off between network size and performance on test data is most important for data sets that are not normally distributed, or data sets that are undersampled during training. Since the current work involves training and testing data generated under controlled conditions, the recognition performance is less sensitive to the size of the network.

Training Algorithm

A training algorithm is used to adjust the interconnecting weights. This training algorithm requires an external "teacher" to provide feedback on the performance of the network. The teacher supplies correct outputs for all training patterns.

The back-propagation training algorithm is used in the ANN discussed in this paper. Back-propagation is an error-correcting algorithm that applies a gradient descent technique to the multiple layers of the ANN.[7,13] Errors are computed on the output layer, then propagated backwards through the network, changing weights up to the input layer. A complete derivation of the algorithm can be found in Ref. 13. A FORTRAN version of the algorithm was written by the authors based on a description in Ref. 7.

The goal of training is to arrive at a set of interconnection weights such that each training pattern is classified correctly, or a minimum number are incorrectly classified. This can be achieved by minimizing the mean square error between the output of the network and the desired output,

$$\text{Error} = \sum_{i=1}^{J}(s_i^* - s_i^{(3)})^2 \qquad (3)$$

The error is summed over J output units. The desired output for each unit, s_i^*, is supplied by the teacher.

The interconnection weights are adjusted after each pattern presentation. An error gradient specifies the amount and direction of change for each weight based on the error between the network and the desired output. An error gradient is computed for each output unit and hidden unit in the network. First, the gradient is computed at the output units

$$\delta_i^{(3)} = (s_i^* - s_i^{(3)})P'(E_i^{(3)}) \qquad (4)$$

where P' denotes the derivative, or slope, of the sigmoid at the value E_i, $\mathrm{d}P/\mathrm{d}E_i$. This error gradient is propagated backwards through the network to each hidden unit

$$\delta_i^{(n)} = \sum_j \delta_j^{(n+1)} w_{ji}^{(n)} P'(E_i^{(n)}) \qquad (5)$$

The error gradient for a given unit i is proportional to the error at each unit j in the next highest layer multiplied by the weight between units i and j. This quantity is then multiplied by the first derivative of the sigmoid function. The derivative is greatest when the sigmoid value is 0.5, and least when the value is near +1 or 0. This tends to drive the output of each unit to one of the extremes, either +1 or 0.

The preceding equations specify an incremental correction for each weight at each pattern presentation. The corrections are noisy approximations to the actual error gradient, so a smoothing technique is used to accelerate learning.[7] A running average of the correction increment is computed,

$$\Delta w_{ij}^{(n)}(u) = \alpha \Delta w_{ij}^{(n)}(u-1) + (1-\alpha)\delta_i^{(n+1)} s_j^{(n)} \qquad (6)$$

The quantity α is chosen to be 0.9. The iteration counter is incremented after each pattern presentation. The weights are then adjusted according to

$$w_{ij}^{(n)}(u+1) = w_{ij}^{(n)}(u) + \varepsilon \Delta w_{ij}^{(n)} \qquad (7)$$

where ε is 0.1. The learning rate affects the speed of convergence and the stability of convergence.

The patterns used to train the ANN are different examples of noise from the two helicopters. The patterns are presented to the network in alternating order, first helicopter A, then helicopter B, and so on, to improve learning.[1,7] The weights are adjusted after each training pattern, and the mean square error of the output saved for that pattern. Training is stopped when the mean square error, summed over the entire training set, stabilizes to an acceptable value. In some applications the error does not stabilize, which indicates the ANN cannot correctly classify the training patterns.

Conventional Recognition System

The ANN is proposed as an alternative identification system. Hence we wish to compare it to a conventional identification system. The conventional method of helicopter identification relies on the ratio of the blade passage frequencies of the main and tail rotors. There are production systems to carry out this identification; however, the algorithms are proprietary. A custom main/tail ratio identification system is designed for this experiment, and will subsequently be referred to as the HARMO system. The HARMO system performs quite well, as will be evident in the results, but it should be noted that it is an experimental system.

The HARMO system was written to use frequency domain information to distinguish between helicopters. The first step in the identification process is to find the fundamental of the main rotor, as described below. Then, the frequencies of the tail rotors of all known helicopters are calculated, using stored main-/tail-rotor ratios. The energy at each tail-rotor frequency is calculated from the spectrum. The tail-rotor frequency with the greatest energy is used to identify the corresponding helicopter type. This process is described in more detail in the following paragraphs.

A harmonic energy detector is used to locate the fundamental frequency of the main rotor. Although the blade-passage frequency may be known for each helicopter, the flight speed of the helicopter will change the measured frequency due to Doppler shifting. The harmonic energy detector is designed for detecting periodic signals whose period is unknown.[16] The harmogram is calculated as part of the detection process. Calculation of the harmogram requires the spectrum

$$Y(\omega_k) = N^{-1/2} \sum_{j=0}^{N-1} y(j\tau)\exp(-i2\pi kj/N) \qquad (8)$$

calculated from a time series sampled at times $t_j = j\tau(j = 0, 1, \ldots, N-1)$. In addition, an estimate of the background noise, $S_x(\omega_k)$, at frequencies ω_k is required. The harmogram, summed over M harmonics, is defined as

$$H(\omega_j) = 2 \sum_{m=1}^{M} \hat{S}_x^{-1}(\omega_{mj})|Y(\omega_{mj})|^2 \qquad (9)$$

where $j = 1, \ldots, [N/2M]$. This is a sum of the spectrum values normalized by background noise values for M harmonics of frequency ω_j. Estimates of the background noise, $\hat{S}_x$, are computed by smoothing the spectrum after removing significant peaks. Reference 16 outlines a complete procedure used for detecting a periodicity. For this work, only the frequency with the maximum harmogram value, summed over eight harmonics, is kept. The frequency search range of the harmonic energy detector is restricted such that any detected periodicity is assumed to belong to a main rotor.

Given this main-rotor frequency, the tail-rotor frequency is calculated for each helicopter from the corresponding main-/tail-rotor ratio. Equation (9) is then used to calculate a harmogram value for each tail-rotor frequency, summing over four harmonics. The highest tail-rotor harmogram value identifies the corresponding helicopter. In this work, since there are only two helicopters, only two tail-rotor harmogram values are calculated.

Recognition Experiments

The purpose of this work is to compare the ability of an ANN with that of a conventional method for distinguishing between the noise of two helicopters. The ANN requires training to learn the characteristics of the helicopters. The conventional method is preprogrammed to distinguish between the helicopters. The two methods are written in FORTRAN code, then tested with spectra for identification. This section describes the features used by the ANN to identify the helicopters and the data used for training and testing the identification methods.

The ANN uses a set of features to distinguish between helicopter types. Features are characteristic attributes calculated to discriminate between the classes of patterns.[14] The largest part of an identification problem is often the search for a set of invariant, or useful, features. Once found, almost any decision algorithm will work. If a useful set cannot be found, then the patterns will not be correctly classified, no matter how complicated the decision algorithm.

The search for features began with the ratio of the fundamental frequencies of the main and tail rotors. These frequencies are easily detected using the harmogram algorithm since each rotor produces a series of harmonically related tones. Two feature values are calculated to describe the harmonic energy at the tail-rotor frequencies of each helicopter. The calculation of these values is described in the previous section of this paper. These two features are referred to as tail-rotor harmogram features. Another set of features is calculated to be used in identifying helicopters regardless of the presence of tail-rotor noise. The additional set of features should contain information common to all helicopters, but the information should be independent of tail-rotor noise.

The additional set of features is used to describe the shape of the main-rotor noise spectrum. First-order and third-order polynomial curves are fitted, in a least squares sense, to the amplitudes of the first eight harmonics of main-rotor noise. The coefficients of the resulting best-fit line and cubic are used as features. That is, the slope of the best-fit line is one feature, and the coefficients of the x, x^2, and x^3 terms from the best-fit cubic are three more features. Note the slope of the best-fit line and the coefficient of the x term from the best-fit cubic are usually different, depending on the shape of the cubic.

Simulated helicopter noise, described in the next section, was created such that the shape of the main-rotor spectrum could be used for identification. Having such control over the training and testing data simplified the feature selection process, since, in effect, a relevant feature was created. The feature selection process for measured helicopter noise would require more careful consideration of the physics of the noise generation mechanisms on a helicopter. The purpose of this work is only to demonstrate the feasibility of a helicopter identification system using some feature that is independent of tail-rotor noise in parallel with a main-rotor/tail-rotor ratio feature. The relevance or significance of the spectral shape features is not established in this paper.

The feature values are scaled using Eq. (10), to avoid numerical problems during neural network training[17]:

$$Z = \frac{X - \mu}{\sigma} \tag{10}$$

The mean value and standard deviation of each feature are calculated for all patterns in the set of training data. These values are stored in order to scale features of patterns in the set of testing data. The scaling creates a bimodal distribution of feature values centered about zero. The tail-rotor harmogram values are scaled by dividing each by the main-rotor harmogram value.

Simulated Helicopter Noise

The performance of the ANN was evaluated using simulated noise of two helicopters. Two prototype spectra were designed, then random noise added to the background and to the amplitudes of the tones. Examples were generated for training the ANN, then a different set was generated for testing the ANN and HARMO methods.

The fundamental frequency of the main rotor was 12 Hz for both prototypes, but the spectrum for each was shaped differently. These prototypes are shown in Figs. 3a and 3b for helicopters A and B, respectively. The distributions of the main-rotor peaks were arbitrarily specified to have the shapes shown in the figure. Each spectrum contained 18 harmonics of the main rotor, and five harmonics of the tail rotor. The distribution of tail-rotor peaks was identical for the two helicopters. The ratio of the main-rotor bpf to tail-rotor bpf was 0.23 for helicopter A, and 0.17 for helicopter B.

The simulated noise was generated in the time domain using Eq. (11):

$$x(t) = \sum_{i=1}^{N} a_i \sin(\omega_i t) + a_n n(t) \tag{11}$$

where $n(t)$ is a Gaussian random process with $\mu = 0$ and $\sigma = 1$. The N amplitudes a_i were specified for each helicopter and the amplitudes were randomly perturbed using uniformly distributed perturbations between $\pm 10\%$. This perturbation, in addition to the $n(t)$ term, produced randomly varying spectra, based on the prototypes. Two such spectra are shown in Figs. 4a and 4b for helicopters A and B, respectively. The identification methods were written to read in power spectra, so the time series were transformed to the frequency domain, then classified by the ANN and HARMO methods.

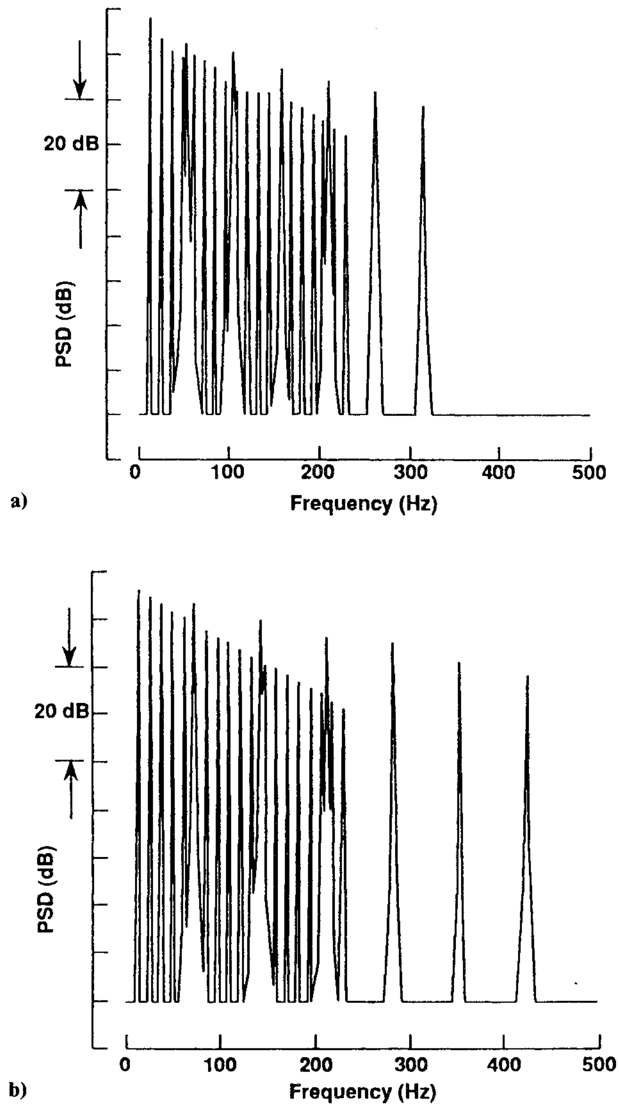

Fig. 3 Prototype for: a) helicopter A; b) helicopter B.

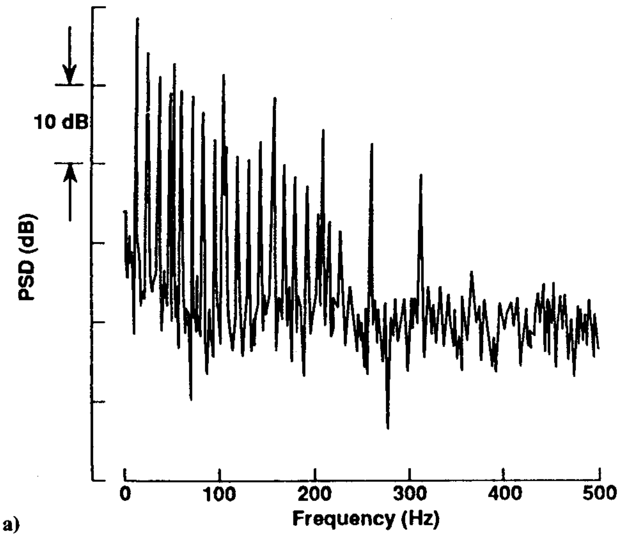

a)

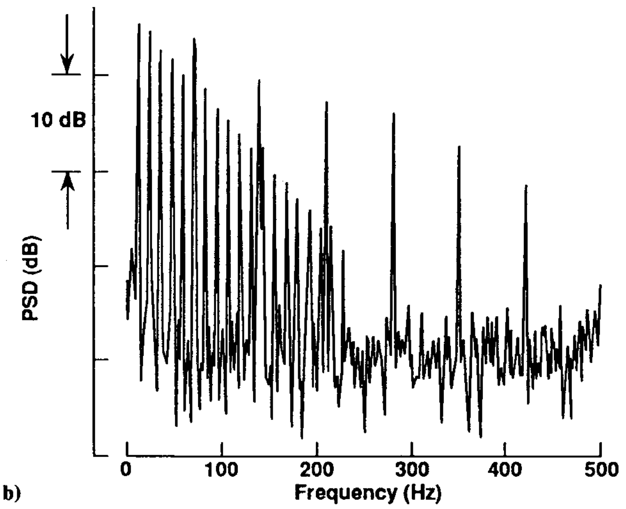

b)

Fig. 4 Prototype plus noise for: a) helicopter A; b) helicopter B.

Training data were made up of 48 spectra, consisting of 24 examples of each helicopter. The training patterns were generated at three signal-to-noise ratios (SNRs): 8, 10, and 20 dB. Eight patterns were generated at each SNR level. The neural network was trained for 100 cycles through the set of 48 training patterns.

Propagation effects were modeled to evaluate the effect of shape changes on the performance of the ANN. For sound propagation through the atmosphere, higher frequencies are attenuated greater than lower frequencies. This can change the shape of the spectrum. At two very different distances, the spectra of the same helicopter could have very different shapes. If the ANN were trained to recognize the spectral shape of helicopter noise recorded at close range only, the shape at long ranges could be too different for the ANN to distinguish between the helicopters.

The model of the atmosphere contained losses due to spherical spreading, atmospheric absorption, and ground impedance. The atmospheric absorption was calculated using ANSI standard S1.26-1978[18] at a temperature of 20°C, and relative humidity 75%. A medium-hard ground, as found in the desert, was used for calculating ground effect.[19] This effect causes a doubling of received power at very low frequencies, then quickly drops off to an asymptotic value as frequency increases. The helicopter was assumed to be flying at an altitude of 30 m.

Results

The identification methods were tested under conditions approximating those encountered with measured helicopter noise. The helicopter signals were embedded in various amounts of background noise in order to test the performance as a function of signal strength. In a separate test, the signals were propagated from various distances through a model of the atmosphere. This gave performance as a function of signal strength and spectral shape, due to the shape-changing nature of atmospheric propagation.

Two types of helicopter noise were used for testing the performance of the identification methods. The first type contained main- and tail-rotor noise. The second type contained main-rotor noise only. The second type of data was generated by eliminating the tail-rotor terms from Eq. (11). All of the training examples for the ANN contained main- and tail-rotor noise.

The following plots show the percent of test sources correctly identified by each method. Data were generated at eight signal-to-noise ratios. Due to the impulsive nature of helicopter noise, the SNR values given in the following results were calculated by taking the average of the SNR of the first eight tones in the spectrum. In a separate test, data were propagated from seven distances through a model of the atmosphere. Two hundred spectra were generated for each data point shown in the figures. The methods classified each test spectrum as either helicopter A or helicopter B. Random guessing would give correct results 50% of the time.

Main-Rotor and Tail-Rotor

The performance of the methods when tested with signals containing main- and tail-rotor noise is shown in Figs. 5 and 6. Both methods can distinguish between the helicopters when the signal is strong, relative to the background noise. The performance of the ANN drops off more quickly than that of the HARMO method as the signal strength diminishes. In Fig. 5, at 10-dB SNR, the ANN drops to 84% correct, while that of the HARMO method is correct 99% of the time. As the SNR becomes smaller, the performance of the two methods drops off rapidly, however at 0-dB SNR the performance of both is still better than random guessing.

Similar results can be seen in Fig. 6. Both methods distinguish between the helicopters 100% of the time at distances under 3 km. As the propagation distance increases, performance drops, with that of the ANN dropping most rapidly. At 4.5 km the ANN is correct 84% of the time, and the

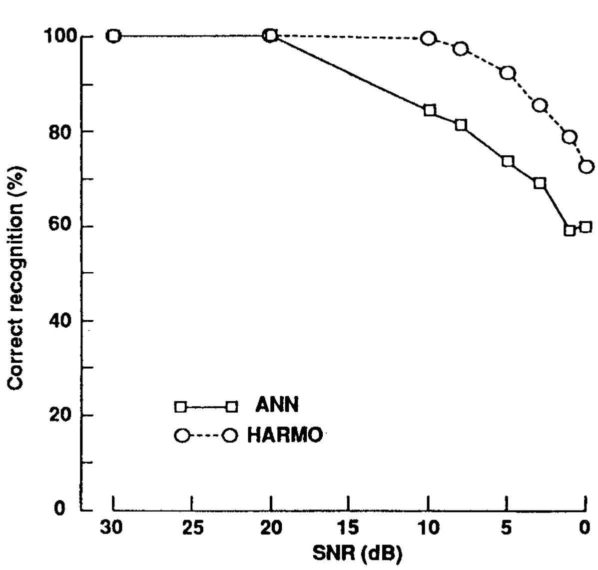

Fig. 5 Recognition performance vs signal strength, main- and tail-rotor noise in spectrum.

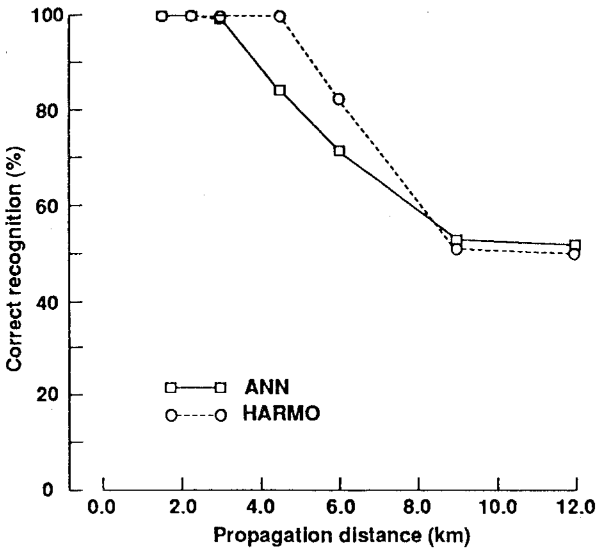

Fig. 6 Recognition performance vs propagation distance, main- and tail-rotor noise in spectrum.

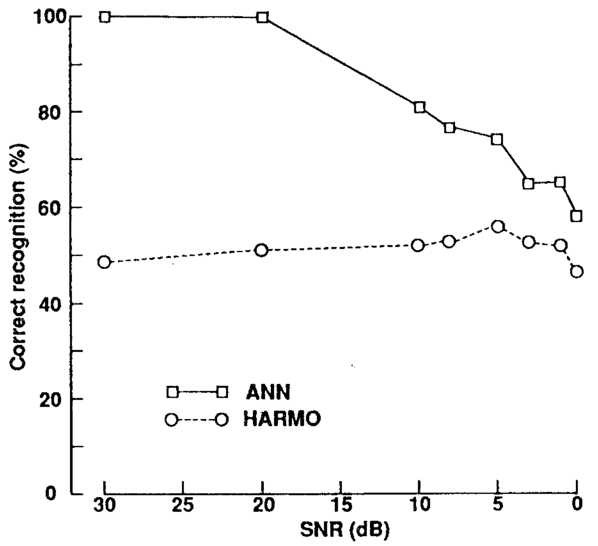

Fig. 7 Recognition performance vs signal strength, main-rotor noise only in spectrum.

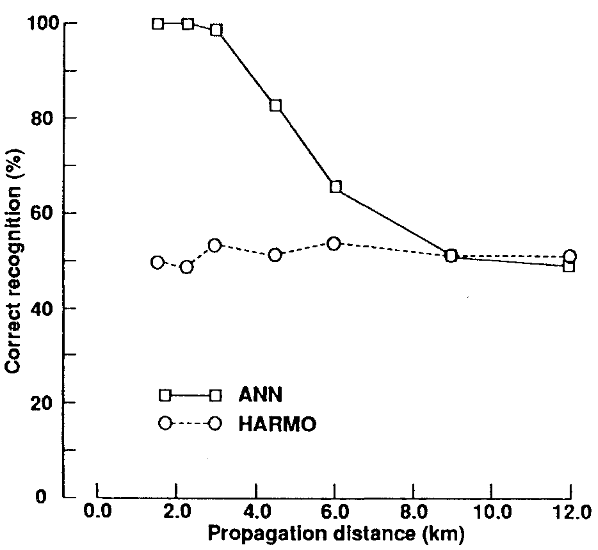

Fig. 8 Recognition performance vs propagation distance, main-rotor noise only in spectrum.

HARMO method is correct 100% of the time. At distances equal to or greater than 9 km, neither method can distinguish between the helicopters greater than 50% of the time.

Main-Rotor Noise Only

The results shown in Figs. 7 and 8 demonstrate the advantage of using the ANN for helicopter identification. For these results, the methods were tested with helicopter signals containing main-rotor noise only. The HARMO method is correct 50% of the time, on average, for all SNR and propagation distances. This is equivalent to random guessing. The ANN, on the other hand, correctly identifies the helicopters 100% of the time at SNR of 20 dB and above, and distances up to 3 km. The ability of the ANN to correctly identify the helicopters drops in a similar fashion to that seen in Figs. 5 and 6. At 4.5 km, the ANN was correct 84% of the time on data containing main- and tail-rotor noise and 83% of the time on data without tail-rotor noise. Similarly, at 8-dB SNR, the ANN was correct 81% of the time on data containing main- and tail-rotor noise, and 77% on data without tail-rotor noise.

Discussion

The results indicate the ANN and HARMO methods can distinguish between two helicopters when main- and tail-rotor noise are present. The HARMO method is more robust with respect to changing signal strength, shown in Fig. 5, and propagation distance, shown in Fig. 6. The performance of the ANN was affected by the changing shape of the spectrum as the signal became more buried in the background noise. As the signal strength drops, relative to the background noise, the higher harmonics blend in with the noise. As the higher harmonics disappear, the shape of the spectrum changes, until it is very different from its original shape. The feature values change as the shape changes, affecting the performance of the ANN. The ANN could be made less sensitive to changes in spectral shape if the training set included examples of the spectral shapes that are to be encountered. In this work, that would mean including examples of helicopter noise propagated from a variety of distances. The performance of the ANN with respect to signal strength would improve with an expanded training set, assuming the network converges to a solution.

The advantage of using the ANN for helicopter identification is that it can identify helicopters even if tail-rotor noise is not measured. The results shown in this paper illustrate the sensitivity of the ANN to changes in spectral shape. However, it is that sensitivity that allows the ANN to distinguish between helicopters when tail-rotor noise is not present. The ANN was trained with two types of features, both of which were useful for distinguishing between the helicopters given during training. It was shown that when the signal was strong, the ANN used these features to identify helicopters accurately. However, when one feature type became useless, i.e., the information it was providing became ambiguous, the performance of the ANN dropped only slightly.

Figures 6 and 8 illustrate the advantage of using the ANN. When the tail-rotor noise was removed from the spectrum, the HARMO algorithm became useless. The energy at tail-rotor frequencies was that of the background noise itself. In contrast, the ANN was trained with spectral shape information in addition to the tail-rotor information used by the HARMO method. This additional information allowed the ANN to distinguish correctly between the helicopters when tail-rotor information was removed, with only a slight drop in performance compared to the full-spectrum results.

Conclusions

An artificial neural network was used to identify simulated helicopter noise regardless of the presence of tail-rotor noise. Features, or descriptive characteristics, were calculated to describe the main-rotor and tail-rotor noise of the helicopters.

The performance of the ANN dropped only slightly when tail-rotor noise was removed from the helicopter noise spectrum. In comparison, the performance of the conventional recognition system, using only the tail-rotor features, dropped significantly when tail-rotor noise was removed from the spectrum. The ANN could identify helicopter noise without tail-rotor noise due to the addition of the main-rotor feature. The incorporation of additional features, such as those describing main-rotor noise, is simplified with the ANN since it learns how to recognize the helicopters using a given set of features. The performance of the ANN is a function of the diversity of signals used during training. The conventional method performed better than the ANN on helicopter signals propagated through a model of the atmosphere. The performance of the ANN on these propagated signals could be improved if the training set for the ANN included helicopter signals propagated from a greater variety of distances.

Acknowledgment

The authors are grateful to the Applied Acoustics Branch of NASA Langley Research Center for its support for this research under Contract NAS1-18471.

References

[1]Waibel, A., Hanazawa, T., Hinton, G., Shikano, K., and Lang, K. J., "Phoneme Recognition Using Time-Delay Neural Networks," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. ASSP-37, No. 3, 1989, pp. 328–339.

[2]Bengio, Y., Cardin, R., De Mori, R., and Merlo, E., "Programmable Execution of Multi-Layered Networks for Automatic Speech Recognition," *Communications of the ACM*, Vol. 32, No. 2, 1989, pp. 195–199.

[3]Bourlard, H., and Wellekens, C., "Links Between Markov Models and Multilayer Perceptrons," *Advances in Neural Information Processing Systems I*, Morgan Kaufmann, San Mateo, 1989, pp. 502–510.

[4]Taber, W. R., Deich, R. O., Simpson, P. K., and Fagg, A. H., "The Recognition of Orca Calls with a Neural Network," *Proceedings of the International Workshop on Fuzzy System Applications*, Iizuka City, Japan, Aug. 21–24, 1988.

[5]Roitblat, H. L., Moore, P. W. B., Nachtigall, P. E., Penner, R. H., and Au, W. W. L., "Dolphin Echolocation: Identification of Returning Echoes Using a Counterpropagation Network," *Proceedings of the International Joint Conference on Neural Networks*, Vol. 1, Inst. of Electrical and Electronics Engineers, New York, 1989, pp. 295–300.

[6]Khotanzad, A., Lu, J. H., and Srinath, M. D., "Target Detection Using a Neural Network Based Passive Sonar System," *Proceedings of the International Joint Conference on Neural Networks*, Vol. 1, Inst. of Electrical and Electronics Engineers, New York, 1989, pp. 335–339.

[7]Sejnowski, T. J., and Rosenberg, C. R., "Parallel Networks that Learn to Pronounce English Text," *Complex Systems*, Vol. 1, 1987, pp. 145–168.

[8]Hemdal, J. F., "An Investigation of Acoustic and Seismic Sources, Final Report," Willow Run Lab., CR 86, Ann Arbor, MI, Feb. 1973.

[9]Cabell, R., and Fuller, C., "A Smart Pattern Recognition System for the Automatic Identification of Aerospace Acoustic Sources," AIAA Paper 89-1114, 1989.

[10]Moukas, P., Simson, J., and Norton-Wayne, L., "Automatic Identification of Noise Pollution Sources," *IEEE Transactions on Systems, Man and Cybernetics*, Vol. SMC-12, No. 5, 1982, pp. 622–634.

[11]Obermeier, K. K., and Barron, J. J., "Time to Get Fired Up," *BYTE Magazine*, Aug. 1989, pp. 217–224.

[12]Lippmann, R. P., "An Introduction to Computing with Neural Nets," *IEEE ASSP Magazine*, April 1987, pp. 4–22.

[13]Rumelhart, D. E., and McClelland, J. L., *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Vol. 1: Foundations*, MIT Press, Cambridge, MA, 1986.

[14]Tou, J. T., and Gonzalez, R. C., *Pattern Recognition Principles*, Addison-Wesley, Reading, MA, 1974.

[15]Hornik, K., Stinchcombe, M., and White, W., "Multilayer Feedforward Networks are Universal Approximators," *Neural Networks*, Vol. 2, 1989, pp. 359–366.

[16]Hinich, M. J., "Detecting a Hidden Periodic Signal When Its Period Is Unknown," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. ASSP-30, No. 5, 1982, pp. 747–750.

[17]Lapedes, A., and Farber, R., "How Neural Nets Work," *Evolution, Learning, and Cognition*, edited by Y. C. Lee, World Scientific, Teaneck, 1988, pp. 331–346.

[18]"American National Standard Method for the Calculation of the Absorption of Sound by the Atmosphere," ANSI S1.26-1978 (ASA 23-1978), June 23, 1978.

[19]Chien, C. F., and Soroka, W. W., "Sound Propagation Along an Impedance Plane," *Journal of Sound and Vibration*, Vol. 43, No. 1, Nov. 8, 1975, pp. 9–20.