Image Cover Sheet

CLASSIFICATION	SYSTEM NUMBER 103862
UNCLASSIFIED	

TITLE

THE HOLOGRAPHIC NEURAL NETWORK PERFORMANCE COMPARISON WITH OTHER NEURAL NETWORKS

System Number:

Patron Number:

Requester:

Notes:

DSIS Use only:

Deliver to: JR

National Defence   Défense nationale

THE HOLOGRAPHIC NEURAL NETWORK: PERFORMANCE COMPARISON WITH OTHER NEURAL NETWORKS (U)

by

Robert Klepko

DEFENCE RESEARCH ESTABLISHMENT OTTAWA
TECHNICAL NOTE 91-18

Canada

October 1991
Ottawa

National Defence / Défense nationale

THE HOLOGRAPHIC NEURAL NETWORK PERFORMANCE COMPARISON WITH OTHER NEURAL NETWORKS (U)

by

Robert Klepko
Airborne Radar Section
Radar Division

DEFENCE RESEARCH ESTABLISHMENT OTTAWA
TECHNICAL NOTE 91-18

PCN
021LA

October 1991
Ottawa

ABSTRACT

Artificial Neural Networks (ANNs) use weighted interconnections of computational elements in an attempt to mimic the ability of the human brain to learn pattern associations. The type and architecture of an ANN is defined by the following items: the coding and decoding schemes used for the inputs and outputs, respectively; the interconnection pattern between the elements; the computation performed by each element; and the method used to change the weighting of values passed on the connections (i.e. the learning mechanism). These networks can be trained, or can train themselves, to perform tasks such as image recognition, speech understanding, robot control, and data filtering, just to name a few of the more popular areas of application. The ANN of interest in this report is the Holographic Neural Network (HNN). This network is specifically designed for pattern association storage, and its designers claim that it has an exceptional generalization capability and storage capacity. The HNN mathematically realizes the storage and retrieval mechanisms of holograms; hence the origin of its name. The operation of this network is briefly described within this report.

In addition to describing the HNN, this report presents pattern recognition performance results of the HNN and compares them with the results achieved by other more common ANNs: Adaline, Hamming, Bidirectional Associative Memory, Recirculation, and Back Propagation. The patterns used for testing represent artificial high resolution radar images of ships. These images appear as a two-dimensional topology of peaks with various amplitudes.

The performance comparisons show that the HNN does not perform as well as the other ANNs when using the same test data. However, modification of the data, to make it appear more Gaussian distributed, does improve the performance of the HNN. It is shown that the HNN performs best if the input data is completely Gaussian distributed.

RÉSUMÉ

Les Réseaux Neuronaux Artificiels (NRAs) utilisent des interconnections qui relient des éléments de calculs afin d'imiter l'habileté du cerveau humain à apprendre à reconnaître et à associer les formes. Le type et l'architecture d'un RNA est défini par les caractéristiques suivantes: le codage et le décodage utilisés respectivement pour les entrées et les sorties; la structure des interconnections entre les éléments; les calculs faits par chaque élément; et la méthode utilisée pour changer le poids des valeurs rencontrées par les connections (c'est-à-dire le mécanisme d'apprentissage). Ces réseaux peuvent être entraînés ou peuvent s'entraîner eux-mêmes pour accomplir des tâches telle que la reconnaissance d'images, la compréhension de la parole, le contrôle de robot et le filtrage de données. Ces applications ne sont que quelques'unes des plus populaires. Le Réseau Neuronal Holographique (RNH) est discuté dans ce rapport. Ce réseau est spécifiquement conçu pour le rangement d'association de formes, et ses concepteurs prétendent qu'il a des capacités exceptionelles pour la généralisation et le rangement. Le RNH réalise mathématiquement les mécanismes holographiques de rangement et d'extraction, d'où l'origine de son nom. Le mode d'opération d'un tel réseau est brièvement décrit dans ce rapport.

En plus de décrire le RNH, ce rapport présente des résultats de la performance du RNH pour la reconnaissance de formes et les compare avec ceux associés à d'autres RNAs plus communs: Adaline, Hamming, Association Bidirectionelle de Mémoire, Recirculation et Propagation Arrière. Les formes utilisées pour les tests représentent des images radars de bateaux artificielles à haute résolution. Ces images sont constituées de pics d'amplitude variable dans un espace à deux dimensions.

La comparaison des performances, déterminées à l'aide du même set de données dans tous les cas, montre que le RNH n'est pas aussi performant que les autres RNAs. Cependant, une modification des données pour obtenir une distribution plus rapprochée d'une distribution Gaussienne améliore la performance du RNH. Il est démontré que le RNH donne les meilleurs résultats quand les données d'entrée ont une distribution complètement Gaussienne.

EXECUTIVE SUMMARY

The Airborne Radar Section at the Defence Research Establishment Ottawa is interested in determining and evaluating techniques which may be used to recognize high resolution radar images of ships. An Artificial Neural Network is one class of techniques which has drawn much attention in the past few years. These networks attempt to mimic the operation of the human brain to perform tasks such as speech understanding, robot control and pattern recognition. These networks do not make any assumptions about the distribution statistics of the input data that is processed. Some networks must be trained, while others can learn on their own. These networks can learn data which is uncorrupted, but recall the same data after it has been buried in noise.

The neural network of most interest within this report is called the Holographic Neural Network (HNN). This network promises a very large data storage capacity and excellent generalization capability. Both of these attributes can be achieved with only a few learning trials, which is unlike most neural networks that require on the order of thousands of learning trials.

This report will study the pattern recognition capability of the HNN for the application area mentioned above. Its performance will be compared with five other commonly used neural networks. It is shown that the best performance achieved by the HNN occurs when the training data has a Gaussian distribution. This appears to be the only drawback of this neural network.

TABLE OF CONTENTS

Contents **Page**

ABSTRACT . . . iii

RÉSUMÉ . . . v

EXECUTIVE SUMMARY . . . vii

1.0 INTRODUCTION . . . 1

2.0 THE HOLOGRAPHIC NEURAL NETWORK AND ITS PERFORMANCE . . . 3

2.1 Description of the Holographic Neural Network . . . 3

2.2 Performance Measurements of the Holographic Neural Network . . . 6

3.0 PERFORMANCE COMPARISON OF THE HOLOGRAPHIC NEURAL NETWORK WITH OTHER NETWORKS . . . 12

3.1 Application Area Used for Performance Testing . . . 12

3.2 Holographic Neural Network Performance Results Compared With Five Other Networks . . . 15

4.0 MODIFICATION OF APPLICATION DATA TO IMPROVE PERFORMANCE . . . 20

5.0 CONCLUDING REMARKS . . . 23

ACKNOWLEDGEMENTS . . . 24

REFERENCES . . . 24

LIST OF FIGURES

Figure		Page
Figure 1	Sigmoidal function - Transforming real-valued inputs to the 0 to 2π phase domain [7].	4
Figure 2	Schematic of the Encoding and Decoding scheme of the HNN [6].	5
Figure 3	Mapping of Gaussian (Normal) Distributed input through the Sigmoid Function [7].	6
Figure 4	A plot of the mean analog response error, as a function of both the number of storage patterns and the number of learning trials. The length of the stimulus vector is fixed at 1024 elements [8].	7
Figure 5	A plot of the mean analog response error, as a function of the stored pattern density, for different lengths of the stimulus vector. Only one learning trial is used to derive these results [8].	8
Figure 6	A plot of the mean analog response error, as a function of both the percentage of distortion of the input stimulus pattern and the number of learning trials. The length of the stimulus vector is fixed at 1024 elements, and only 256 pattern associations are stored [8].	10
Figure 7	A plot of the mean analog response error, as a function of the percentage of distortion of the input stimulus pattern, for different numbers of encoded pattern associations. The length of the stimulus vector is fixed at 1024 elements [8].	10
Figure 8	A plot of the mean analog response error, as a function of both the number of storage patterns and the number of learning trials. The length of the stimulus vector is fixed at 128 elements.	11

LIST OF FIGURES (CONTINUED)

Figure		Page
Figure 9	A plot of the mean analog response error, as a function of the stored pattern density for one length of the stimulus vector equal to 128 elements. Only one learning trial is used to derive these results.	11
Figure 10	An example image pattern used to test the ANN [9].	12
Figure 11	(A) Three-dimensional plot of the image pattern in Figure 10, and (B) its corresponding plot with 12 modified or added peaks [9].	14

LIST OF TABLES

Table		Page
Table 1	Performance Summary of five ANNs after training with uncorrupted patterns [9]. NOT TESTED implies that no test was performed because of anticipated poor results.	15
Table 2	Performance summary for the Adaline and BackProgation networks using corrupted training data [9].	17
Table 3	Performance results for the HNN using uncorrupted patterns for training, and 300 randomly selected patterns for testing. NOT TESTED implies that no test was performed because of anticipated poor results.	18
Table 4	Performance results for the HNN using corrupted patterns for training, and 300 randomly selected patterns for testing.	19
Table 5	Performance results for the HNN using more randomized input patterns for training, and 300 randomly selected patterns for testing. NOT TESTED implies that no test was performed because of anticipated poor results.	21
Table 6	Performance results for the HNN using input patterns with Gaussian distributed noise added to each pattern used for training and testing. Testing was done with 300 randomly selected patterns.	22

1.0 INTRODUCTION

Artificial Neural Networks (ANNs) attempt to model our present understanding of the human brain. Our vast interconnection of neurons is modelled as a dense interconnection of nonlinear computational elements or nodes. The links between the nodes are weighted, and the ANN is able to learn by varying the values of these weights. Typical things that these networks can learn to perform are image, written text and speech recognition, and data filtering. The ANNs can either learn on their own or through supervision. Self-learning implies the association of an input stimulus to an output response, in which the input data is clustered by the ANN to derive an output. Supervised learning implies that the ANN is presented with both an input stimulus and its associated or desired output response.

An ANN contains the following components: the coding and decoding mechanisms for its input and output data, respectively; the interconnection topology of the nodes; the nonlinear computation performed in each node; and the method used to change the weights on the links between the nodes. The coded format of the input and output data can either be binary, bipolar or (scaled) analog. A typical interconnection topology has each node connected to every other node in the network. The computation performed at each node is identical, and consists of passing the summation of all the weighted input values through a nonlinear threshold function, and then feeding the result to the input of all the other nodes. The method for changing the weights defines the learning scheme for the network. The type and architecture of the ANN depends upon the learning scheme, the interconnection pattern, the computation performed by each node, and the data coding format. A description of various types of ANNs, including a good introduction to the topic, can be found in [1]. For the interested reader, reference [2] describes more recent (i.e. experimental) ANNs.

The ANN of interest in this report is the Holographic Neural Network (HNN) because of the exceptional performance it has demonstrated [8]. This network mathematically realizes the storage and retrieval mechanisms of holograms; hence the origin of its name. Its operation is briefly described in the next section. The creators of this network claim that it has an outstanding generalization capability and storage capacity. These claims are presented and substantiated in the next section. The network's superb performance is achieved using Gaussian distributed data.

The Airborne Radar Section at DREO is interested in studying the performance of various neural networks in recognizing high resolution radar images of ships. These images can be grossly described as a two-dimensional topology of peaks with various amplitudes. To simplify the testing process, artificial as opposed to real or simulated ship imagery is used. Unlike real or simulated images, which require extensive data processing, an artificial image is simply created with the knowledge that ship images are composed of a distribution of peaks. Each artificial ship image consists of a fixed number of peaks arranged in a predetermined manner. Each peak is assigned one of three predefined amplitude values.

The distributions of the peaks - one distribution per ship image - are further described in Section 3.

This distribution of amplitude values is non-Gaussian. Consequently, the performance of the HNN is severely degraded, as is shown in Section 3. The HNN's recognition capability for this application is compared with that obtained by other more common ANNs: Adaline [3], Hamming [1], Bidirectional Associative Memory [4], Recirculation [5], and Back Propagation [1]. A description of these networks is beyond the scope of this report, but can be found in the references given.

Since the HNN had difficulty recognizing the artificial radar images it was trained with, but performed very well with Gaussian distributed data, the artificial data was modified to make its distribution appear slightly more Gaussian. The results presented in Section 4 indicate that the HNN's performance improved after this modification was made. Section 4 gives a description of the modified distribution of the test data.

2.0 THE HOLOGRAPHIC NEURAL NETWORK AND ITS PERFORMANCE

2.1 Description of the Holographic Neural Network

The Holographic Neural Network (HNN) (a complete description can be found in references [6] and [7]) is a heteroassociative ANN which uses supervised learning of analog stimulus/response vector pairs. This means that the network can accept real-valued inputs arranged as an N element vector, called the stimulus, and output a real-valued M element vector, called the response. Values for N and M range from 1 to some maximum, limited by the amount of memory available to implement the network. Since N does not have to equal M, the network is considered heteroassociative. Learning is supervised because the input stimulus and its corresponding output response must be presented to the network at the same time.

There are two main advantages that this network provides over other networks. They are the ability to learn or retain large amounts of data, which can be recalled with little error, and the small number of trials required to learn that data. The first advantage is achieved by using a sophisticated transformation scheme for the input data, prior to its storage, and a complementary scheme for the output data. Each real-valued input element of the stimulus vector is coded as phase information by using the sigmoidal transformation

$$\theta_k = \frac{2\pi}{1 + e^{\frac{\mu - s_k}{\sigma}}} \quad , \qquad k = 1, 2, 3, \ldots , N \qquad (1)$$

where s_k is the input value for element k of the stimulus vector, μ is the mean value over all N input elements, and σ is the standard deviation of the N input element values. The phase of input element k, θ_k, has values between 0 and 2π, as shown in Figure 1 [7]. A level of confidence in, or weighting of, the value s_k is heuristically determined, where 0 corresponds to no confidence and 1 corresponds to full confidence in the value. Usually this confidence level is set to 1. The confidence level represents the magnitude, and θ_k represents the phase, of the complex number that is actually stored within the HNN. Since the network's stored values are complex, its output values are also complex. However, since the input was originally real-valued, the output must be transformed to a real value. This is simply done by computing the inverse tangent of the imaginary part divided by the real part of the complex number. The magnitude of the output complex number provides a level of confidence in, or weighting of, the real-valued output.

With the sigmoidal transformation scheme, it is possible to overlay or store a large number of input values onto a single input element, and yet be able to retrieve their corresponding or associated output response values with little error. The number of input values that can be stored on an element is dependent upon the statistical distribution of the input data.

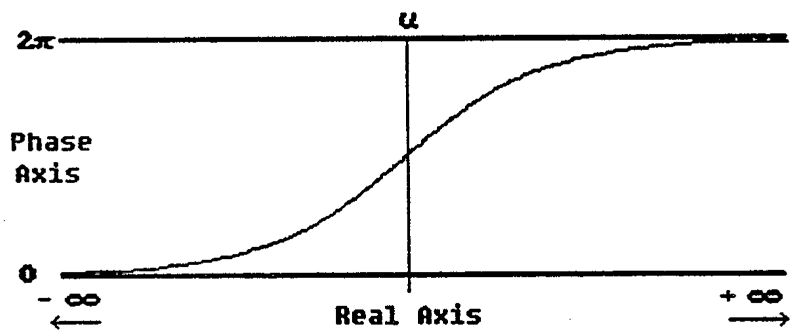

Figure 1 Sigmoidal function - Transforming real-valued inputs to the 0 to 2π phase domain [7].

One major advantage of the HNN over other ANNs is its ability to learn very quickly. This is due to its encoding (learning) scheme, which creates a correlation matrix, X, of complex numbers, by executing a complex inner product over the stimulus and response vector elements. This is similar to the scheme used in the Bidirectional Associative Memory [4]. Each element in the correlation matrix, $x_{k,j}$, is computed by multiplying the complex conjugate of the transformed stimulus vector element k, with the transformed response vector element j. In vector notation,

$$X = S^{*T} \cdot R \qquad (2)$$

where S^{*T} is the transpose of the complex conjugate of the stimulus vector, and R is the response vector. To decode, or retrieve the response vector stored in the correlation matrix X, the following operation is performed

$$R = 1/c\ (S \cdot X) . \qquad (3)$$

where c is a constant equal to the sum of the confidence measures over all k stimulus vector elements. The encoding and decoding operation is depicted in Figure 2 [6]. By modifying the encoding scheme so that the changes made to the correlation matrix after each learned stimulus/response pattern are minimized, superior network performance is achieved [6]. The modified or enhanced encoding method is given by (in vector notation)

$$X_{DIFFERENCE} = S^{*T} \cdot (R - 1/c\ (\ S \cdot X\ )) \tag{4a}$$

whereby

$$X_{NEW} = X + X_{DIFFERENCE} \tag{4b}$$

With this learning scheme, no more than three iterations of the learning cycle (i.e. three encodings) using the same stimulus/response patterns are required to achieve the same level of performance as other ANNs, which typically require thousands of learning cycles. The maximum levels of performance for the HNN are achieved with Gaussian distributed data, as is discussed in the following subsection.

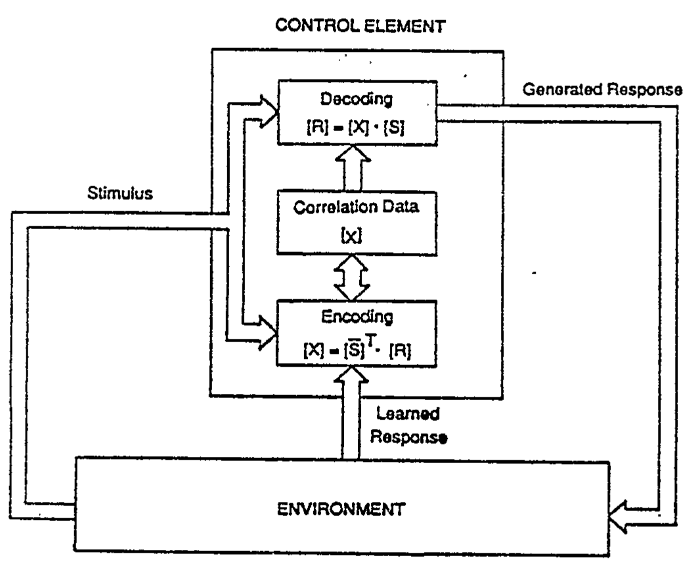

Figure 2 Schematic of the Encoding and Decoding scheme of the HNN [6].

2.2 Performance Measurements of the Holographic Neural Network

Two basic sets of performance measurements were derived by the company that developed the HNN. The first is the storage capacity, and the second is the generalization capability, of the HNN. The enhanced encoding scheme, described in the previous subsection, was used. It should be noted that during the testing procedures, the length of the stimulus vector is varied for different tests, but the length of the response vector is fixed at one element. The values of the input and output data range between 0 and 255, and are randomly selected from the same Gaussian distribution. In other words, the stimulus and response vectors are associated using numbers from the same distribution. The first N values are associated with the (N+1)th value, the next N values are associated with the (2N+2)th value, and so on. This testing demonstrates that by associating input vectors, taken from a Gaussian distribution, with output numbers from the same distribution, the network is still able to distinguish between the individual numbers. It should also be noted that by using the sigmoidal transformation function on the input data, the Gaussian distributed data is mapped into a more uniform distribution, which is stored in the HNN. This mapping is depicted in Figure 3 [7].

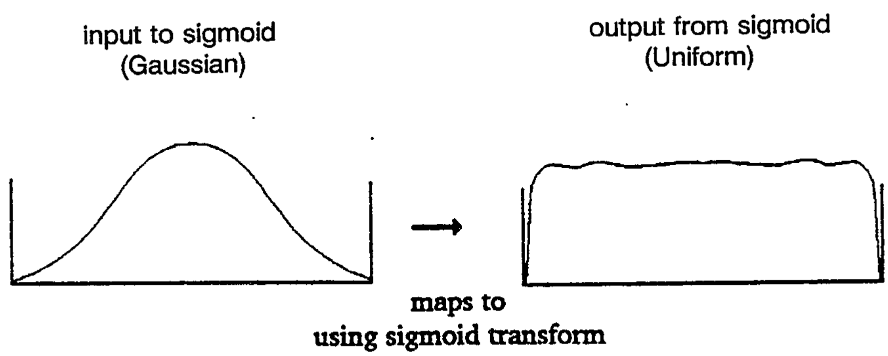

Figure 3 Mapping of Gaussian (Normal) Distributed input through the Sigmoid Transform [7].

Figure 4 [8] is a plot of the storage density of the HNN using a 1024 element stimulus vector. The plot shows how the mean (over all test patterns) analog response error in the output response vector, which is a percentage of the full scale value of 255, varies as a function of both the number of storage pattern associations and the number of learning trials. The error for each stored association is measured by dividing the absolute value of the expected response minus the output response, by the maximum valued output response.

The errors for all associations are then averaged to obtain the mean response error.

With a conventional storage device, 512 patterns of a vector with 1024 elements would require 512 Kilobytes of storage (one byte per element). With the HNN process, all 512 patterns are overlaid onto 1024 complex numbers, where each number is represented by a magnitude (two bytes of storage) and a phase (two bytes of storage). Thus, only 4 Kilobytes are required to store these patterns. The storage requirements are reduced by a factor of 128. However, with conventional storage, there is no error when the response vector is recalled, but the HNN has about a 2.2 percent error (after three learning trials).

Figure 5 [8] is another plot representing storage density. It illustrates the variation of the mean analog response error, as a function of the number of stored pattern associations (which is a percentage of the length of the stimulus vector), for different lengths of the stimulus vector. Note that as the length of the stimulus vector increases, the performance improves slightly.

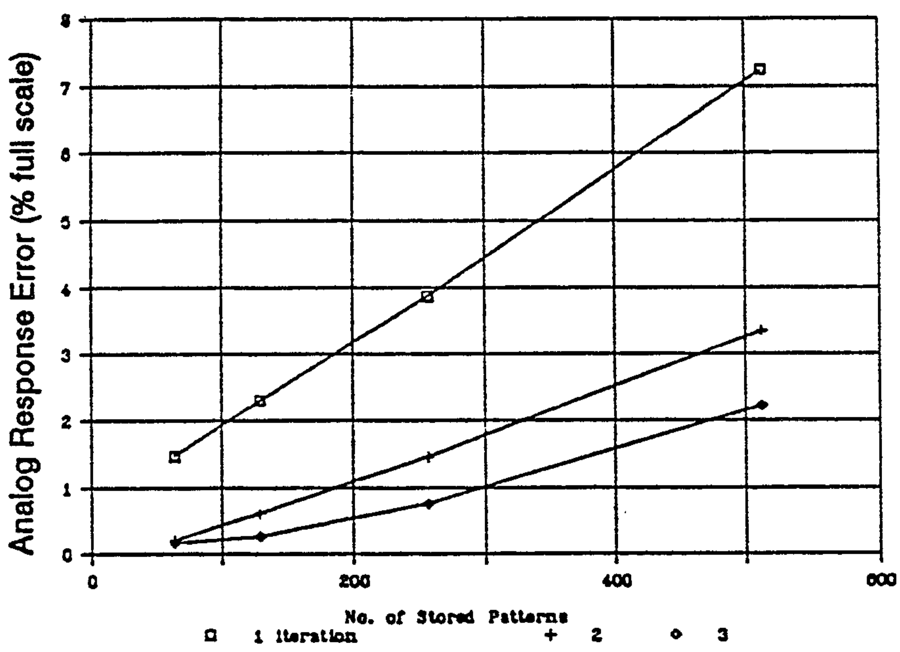

Figure 4 A plot of the mean analog response error, as a function of both the number of storage patterns and the number of learning trials. The length of the stimulus vector is fixed at 1024 elements [8].

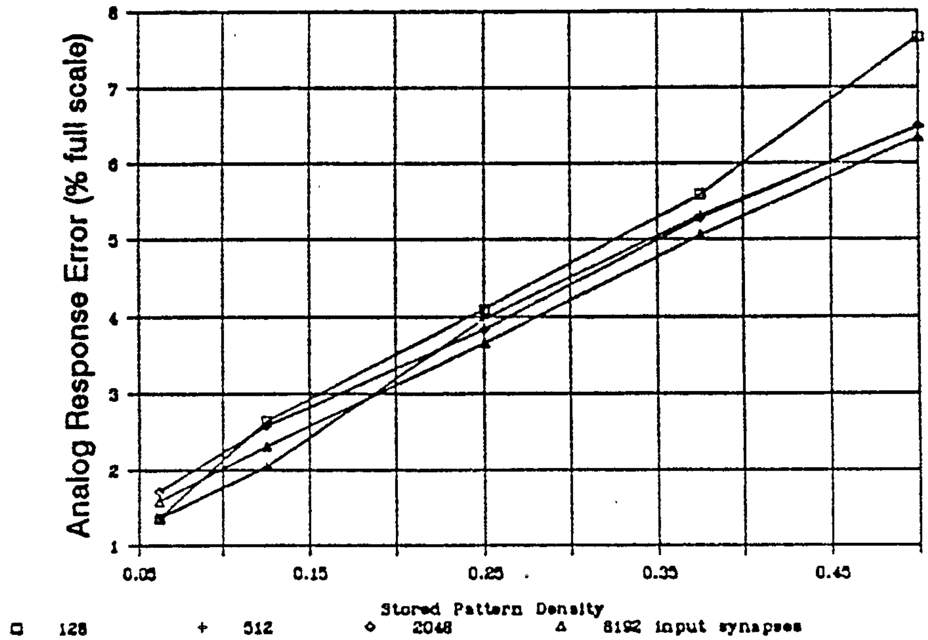

Figure 5 A plot of the mean analog response error, as a function of the stored pattern density, for different lengths of the stimulus vector. Only one learning trial is used to derive these results [8].

Figure 6 [8] illustrates the generalization capability of the HNN. Generalization is a performance measure of the amount of error in the output response, for various amounts of distortion of the input stimulus pattern. The plot in Figure 6 shows the mean analog response error, as a function of both the number of learning trials and the distortion of the stimulus pattern. The distortion is measured as a percentage of the number of vector elements which have been altered from their original values. The length of the stimulus vector is fixed at 1024 elements and the number of stored pattern associations is 256. Note that even with a distortion of 50 percent, the error is below 10 percent.

As a final illustration of the performance of the HNN, Figure 7 [8] also depicts its generalization capability. The plot in Figure 7 shows how the mean analog response error varies, as a function of the distortion of the stimulus pattern, for different numbers of encoded pattern associations. The length of the stimulus vector is fixed at 1024 elements.

The results in both Figures 6 and 7 indicate that the HNN has a large immunity to very high distortion levels, even with high pattern storage densities. High levels of immunity to distortion indicate that the learned pattern associations form generalizations, in the sense that each pattern association maps a wide region within the stimulus state space to the desired response.

Prior to comparing the performance of the HNN with other ANNs in our application area, software to implement the HNN had to be written. It was coded in FORTRAN-77 and implemented on an IBM compatible 80286 microcomputer. Tests with Gaussian distributed data were conducted to verify that the HNN was coded correctly. This allowed its measure of performance, for storage capacity, to be directly compared with the results plotted in Figures 4 and 5. Figure 8 illustrates the mean analog response error, as a function of the number of stored pattern associations, for different numbers of learning trials. The results in this figure can be compared with the results shown in Figure 4, except that the "Number of Stored Patterns" axis in Figure 8 should be multiplied by 8 to allow a direct comparison with Figure 4. The reason for this multiplication is that the results in Figure 4 apply to a 1024 element stimulus vector, while the results in Figure 8 were obtained with a 128 element stimulus vector. Comparing the results in these two figures indicate that they are similar.

Figure 9 displays the mean analog response error versus the stored pattern density. The stimulus vector length is 128 elements. The results in this figure are similar to the results plotted in Figure 5.

Since the results in Figures 8 and 9 are similar to those in Figures 4 and 5, respectively, we can assume that the HNN was coded correctly. Although the mean error results are very similar, they are not identical. There are two simple reasons for this difference. First, the length of the stimulus vector used for testing by the company was 1024 elements, while only 128 elements (due to memory limitations) were used in our testing. To a lesser degree, the second reason is that the Gaussian distributed data used to generate the results in Figures 4 and 5 were not exactly the same as the data used to obtain the results shown in Figures 8 and 9.

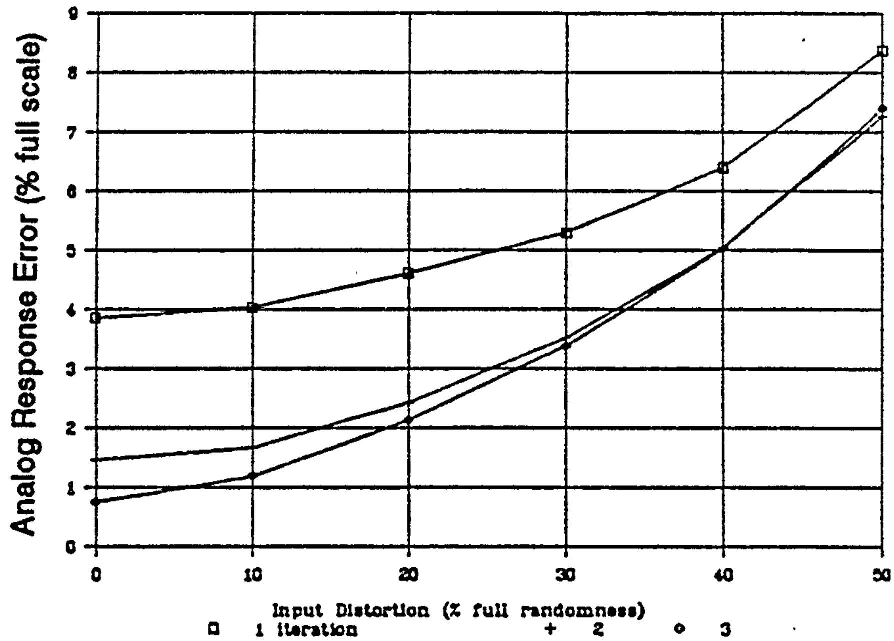

Figure 6 A plot of the mean analog response error, as a function of both the percentage of distortion of the input stimulus pattern and the number of learning trials. The length of the stimulus vector is fixed at 1024 elements, and only 256 pattern associations are stored [8].

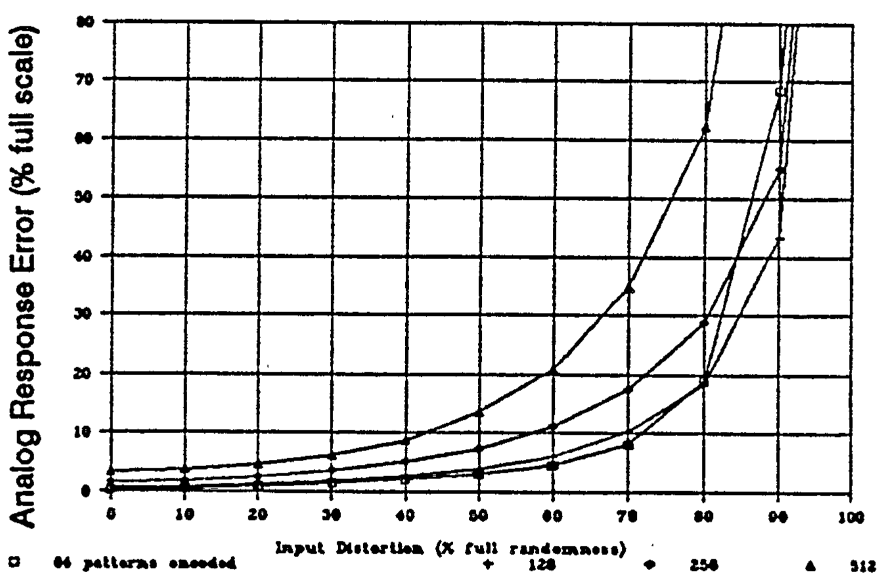

Figure 7 A plot of the mean analog response error, as a function of the percentage of distortion of the input stimulus pattern, for different numbers of encoded pattern associations. The length of the stimulus vector is fixed at 1024 elements [8].

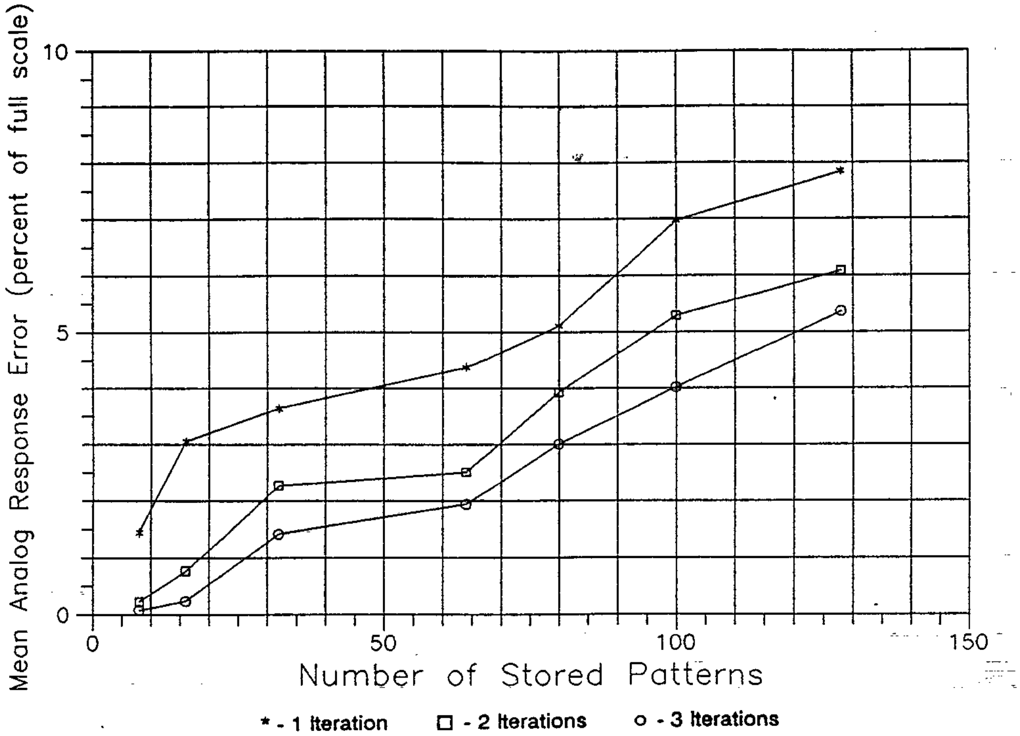

Figure 8 A plot of the mean analog response error, as a function of both the number of storage patterns and the number of learning trials. The length of the stimulus vector is fixed at 128 elements.

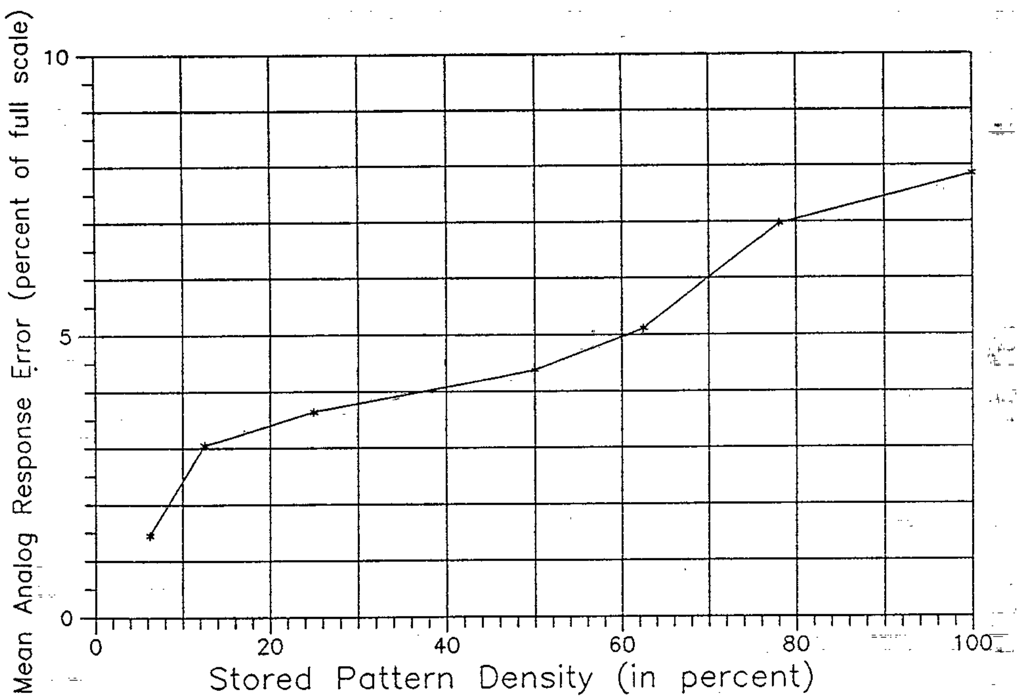

Figure 9 A plot of the mean analog response error, as a function of the stored pattern density for one length of the stimulus vector equal to 128 elements. Only one learning trial is used to derive these results.

3.0 PERFORMANCE COMPARISON OF THE HOLOGRAPHIC NEURAL NETWORK WITH OTHER NETWORKS

One goal of the Airborne Radar Section at the Defence Research Establishment Ottawa is to develop and evaluate algorithms that can be used to classify high resolution radar images of ships. An evaluation of the pattern recognition capability of five common ANNs was performed, and their performance is reported in [9]. In order to evaluate the performance of the HNN in this application area, we shall test the HNN with the same data as was used to evaluate the other ANNs in [9]. Thus, the pattern recognition capability of the HNN can be directly compared with the other ANNs for this particular application. The performance results and comparisons are presented in Section 3.2. A more detailed description of the application area is given in the following subsection. It should be pointed out that due to the limited funding for performing the evaluation described in [9], real or simulated radar images of ships were not used since they would require far more time for testing than was available. Instead, artificial images were used since they could be easily generated and provided for a well-controlled data test set.

3.1 Application Area Used for Performance Testing

The data set, used to evaluate the ability of the ANNs to recognize patterns, can be roughly described as a two-dimensional topology of peaks with various amplitudes. These patterns crudely represent high resolution radar images of real ships. A total of ten distinct patterns were created and used for testing. Each pattern consists of 6 rows of 27 columns each. These patterns are presented in [9], but one is shown in Figure 10 for convenience.

	1				5					10					15					20					25		27
1							M	M	M										L	L	L						
2							M	M	M										L	L	L						
3	H	H	H																						H	H	H
4	H	H	H																						H	H	H
5																											
6																											

Figure 10 An example image pattern used to test the ANN [9].

The pattern within this figure is composed of a certain number of peaks, with each peak being assigned one of three predefined amplitude values. A peak occupies six pixels or elements in the two-dimensional grid pattern. The three peak amplitude values are designated low, medium and high, and correspond to values of 5, 10 and 20, respectively. The arrangement of these peaks and the values assigned to each, define the artificial radar image of a ship. The pattern in Figure 10 has four peaks: two are high, one is medium, and the fourth is low.

Not only are the image patterns used for testing, but corrupted versions of these patterns are used as well. A predetermined number of image pixels was randomly selected and assigned random amplitude values. If a pixel's amplitude was zero, then a value was added, and if it was non-zero, then this amplitude was modified. These altered image pixels are called "added or modified peaks" in the remainder of this report. Figure 11(a) [9] shows a (crude) three-dimensional plot of the image pattern of Figure 10, and Figure 11(b) [9] shows a three-dimensional pattern of the same image with 12 added or modified peaks.

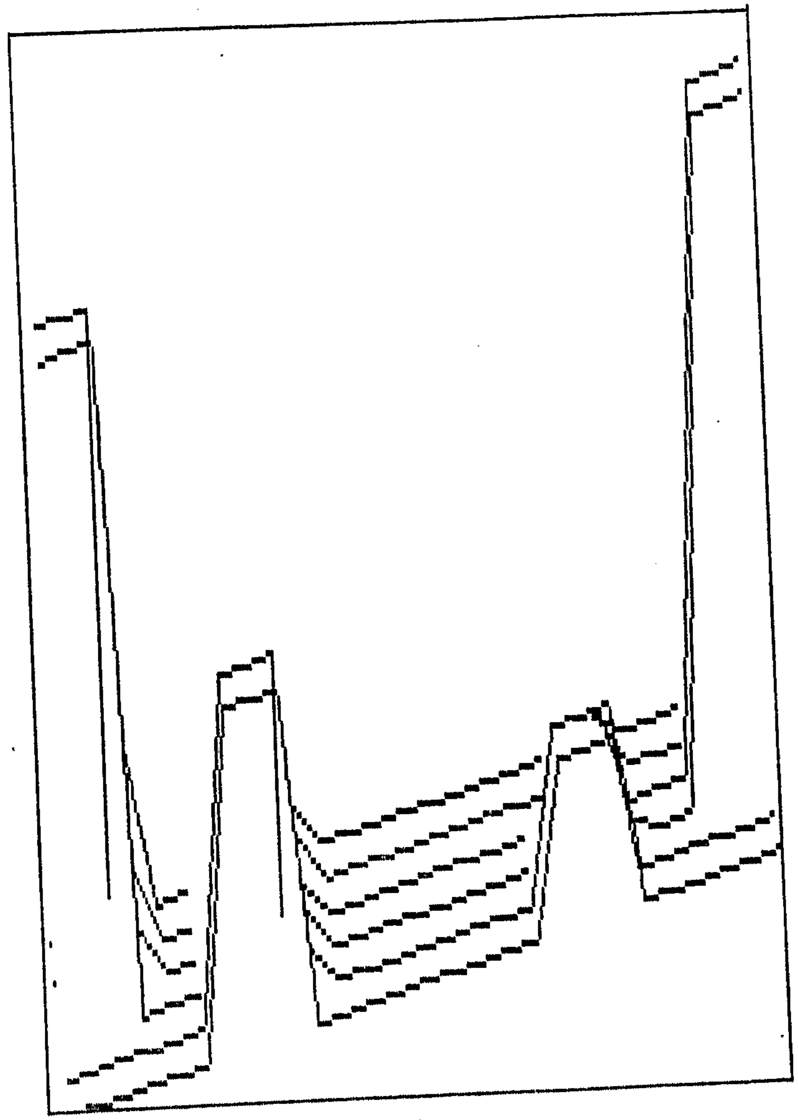

(A)

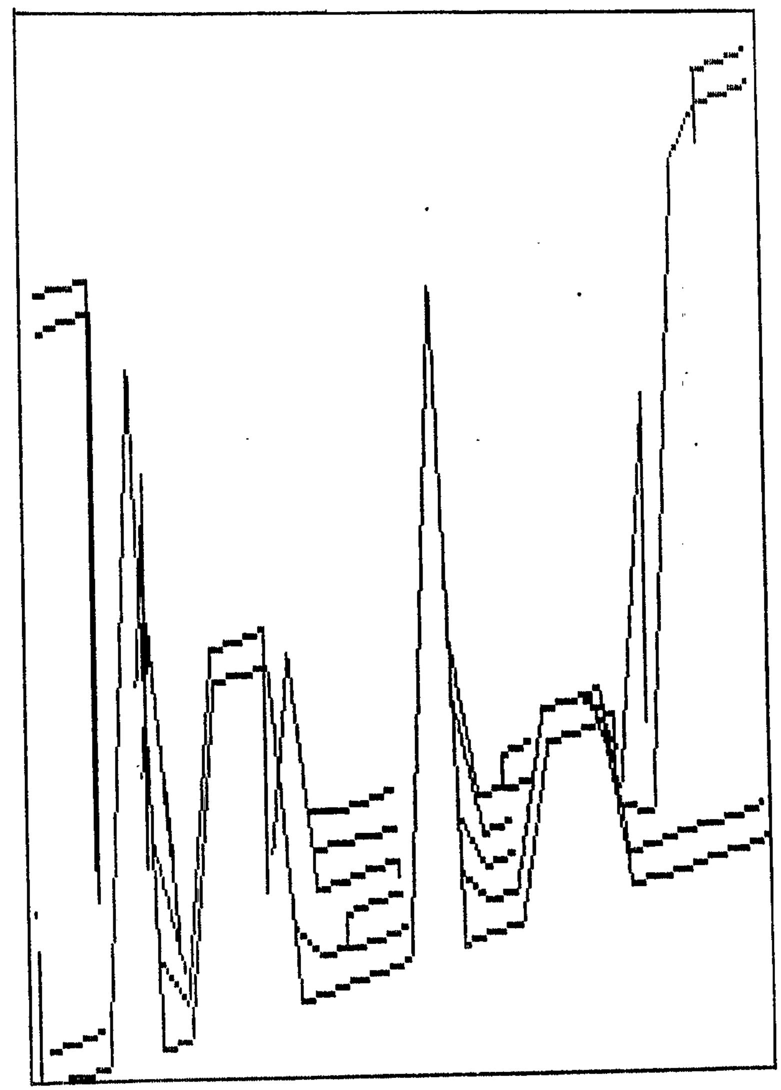

(B)

Figure 11 (A) Three-dimensional plot of the image pattern in Figure 10, and (B) its corresponding plot with 12 modified or added peaks [9].

3.2 Holographic Neural Network Performance Results Compared With Five Other Networks

A performance summary of the five ANNs evaluated in reference [9], using ten distinct image patterns similar to the one given in Figure 10, can be found in Table 1 [9]. The commercial neural networks software development package called Professional II™, created by Neuralware Inc. [10], was used to evaluate the five ANNs. This package provides these precoded ANNs, and a graphical interface to display input and output data.

NETWORK	NUMBER OF ADDED OR MODIFIED PEAKS (Test Data)	PERFORMANCE	CONFIDENCE INTERVAL in Percent (Confidence level is set at 99.0%)
Adaline	0	100.0%	97.84 - 100.0
	3	100.0%	97.84 - 100.0
	6	95.0%	90.68 - 97.38
	9	82.3%	75.95 - 87.26
	12	71.0%	63.86 - 77.23
Hamming	0	100.0%	97.84 - 100.0
	3	100.0%	97.84 - 100.0
	6	100.0%	97.84 - 100.0
	9	100.0%	97.84 - 100.0
	12	100.0%	97.84 - 100.0
Bidirectional Associative Memory	0	30.0%	23.68 - 37.19
	3	46.3%	39.05 - 53.71
	6	NOT TESTED	N/A
	9	NOT TESTED	N/A
	12	0.0%	0.00 - 2.16
Recirculation	0	100.0%	97.84 - 100.0
	3	98.3%	95.09 - 99.42
	6	95.3%	91.06 - 97.58
	9	87.0%	81.19 - 91.21
	12	81.3%	74.85 - 86.40
BackPropagation	0	100.0%	97.84 - 100.0
	3	89.7%	84.29 - 93.39
	6	79.3%	72.67 - 84.66
	9	73.0%	65.96 - 79.05
	12	62.3%	54.90 - 69.17

Table 1 Performance Summary of five ANNs after training with uncorrupted patterns [9]. NOT TESTED implies that no test was performed because of anticipated poor results.

During the learning process, the patterns presented to each network were randomly selected from the training set of ten images. The Adaline network completed its training after 4500 patterns (i.e. each of the distinct patterns was used 450 times) had been presented. The Hamming and Bidirectional Associative Memory networks only required one learning trial for each of the ten images. The Recirculation network was trained with each pattern 750 times, and training of the BackPropagation network was completed after each pattern had been presented 1950 times.

All test results were arrived at by presenting 300 patterns to the trained networks for classification. These were randomly selected from the 10 distinct images, with each being chosen 30 times using a different corrupted variation each time. With only 300 patterns available for testing, the measured performance results may not be statistically accurate. In other words, if many more patterns were used for testing, then more accurate or statistically consistent results would have been obtained. Since a very large number of patterns was not used for testing, a confidence region is defined, which provides an interval within which the true network performance would lie with a given level of confidence. The bounds of the confidence interval are provided in the following tables of performance results. These limits have been computed under the assumption that the probability distribution of the number of errors is Gaussian [11]. Note that if the confidence level is reduced, then the confidence interval becomes smaller.

In addition to training with uncorrupted data, training was done with patterns that had a number of added or modified peaks. Table 2 shows the performance results, for two of the networks, using corrupted patterns for both training and testing [9]. As with the results in Table 1, only 300 patterns were used to test the networks.

The performance results for the Adaline network were achieved after using 3300 corrupted (i.e. 330 times for each pattern) patterns for training. The BackPropagation network required 19200 corrupted patterns (i.e. 1920 times for each pattern) in the learning process to achieve its performance results. Note that, for this test data, the Adaline network performed better when trained with corrupted data than with uncorrupted data. However, this was not the case for the BackPropagation network.

In order to test the HNN with the same image patterns, its architecture was created to accept input stimulus vectors with 162 elements (i.e. 27 columns by 6 rows each) having scaled (from the original pixel values of 0, 5, 10 and 20) analog inputs ranging in value from 0 to 255. The output response vector consisted of 10 elements with analog values ranging between 0 and 128. Each output element corresponded to one of the patterns. Hence, all the elements, except one, should be zero at any time. In reality, all the output elements had some non-zero value, but the element with the highest value was considered to be the pattern which was associated to the input stimulus vector.

NETWORK	NUMBER OF ADDED OR MODIFIED PEAKS (Training Data)	NUMBER OF ADDED OR MODIFIED PEAKS (Test Data)	PERFORMANCE	CONFIDENCE INTERVAL in Percent (Confidence level is set at 99.0%)
Adaline	12	0 3 6 9 12	100.0% 100.0% 98.7% 95.0% 90.7%	97.84 - 100.0 97.84 - 100.0 95.68 - 99.62 90.68 - 97.38 85.46 - 94.18
Adaline	24	0 3 6 9 12	100.0% 100.0% 99.7% 97.7% 96.3%	97.84 - 100.0 97.84 - 100.0 97.28 - 99.97 94.23 - 99.10 92.35 - 98.25
Back-Propagation	24	0 3 6 9 12	60.0% 60.0% 59.3% 59.7% 59.7%	52.58 - 66.99 52.58 - 66.99 51.87 - 66.33 52.27 - 66.71 52.27 - 66.71

Table 2 Performance summary for the Adaline and BackPropagation networks using corrupted training data [9].

Using the same training and testing procedure for the HNN as was used for the other ANNs, the performance results of Tables 3 and 4 were achieved. These can be compared to Tables 1 and 2, respectively. It is apparent that the HNN performed very poorly in this application. It was mentioned in reference [7] that because of the data coding format used, perhaps only 3 or 4 learning trials (i.e. present each pattern association 3 or 4 times to the network) would be needed to achieve the maximum pattern recognition performance. For testing the HNN, 300 training patterns were used, and this corresponds to using each of the 10 patterns 30 times. Hence, the suggested maximum number of training trials was exceeded, but this did not result in good performance. It should be pointed out that evaluations with only 3 training trials were made, but resulted in a maximum performance of about 30.0% when tested with uncorrupted data.

From the results in Table 3, it appears as though the performance of the HNN becomes worse with more training. For many ANNs, their performance tends to get better with more training, but levels out at some maximum regardless of the amount of training done. However, there are some ANNs whose performance tends to reach some maximum level, and then decrease with additional training.

TOTAL NUMBER OF TRAINING PATTERNS	NUMBER OF ADDED OR MODIFIED PEAKS (Test Data)	PERFORMANCE	CONFIDENCE INTERVAL in Percent (Confidence level is set at 99.0%)
300	0 3 6 9 12	100.0% 37.67% 18.67% 15.0% 10.7%	97.84 - 100.0 30.81 - 45.07 13.58 - 25.12 10.45 - 21.06 6.93 - 16.17
600	0 3 6 9 12	100.0% 35.3% 16.3% 14.67% 12.67%	97.84 - 100.0 25.58 - 42.65 11.55 - 22.51 10.18 - 20.69 8.52 - 18.43
900	0 3 6 9 12	100.0% 34.67% 14.67% 13.67% 10.0%	97.84 - 100.0 28.00 - 42.01 10.18 - 20.69 9.34 - 19.57 6.37 - 15.36
1200	0 3 6 9 12	100.0% 37.0% 16.0% 15.67% 11.67%	97.84 - 100.0 30.18 - 44.39 11.29 - 22.18 11.02 - 21.81 7.71 - 17.29
1500	0 3 6 9 12	100.0% NOT TESTED NOT TESTED NOT TESTED 10.67%	97.84 - 100.0 N/A N/A N/A 6.90 - 16.14

Table 3 Performance results for the HNN using uncorrupted patterns for training, and 300 randomly selected patterns for testing. NOT TESTED implies that no test was performed because of anticipated poor results.

From the results in Table 4, it appears that if the HNN is trained with corrupted data, its performance becomes worse. This is opposite of the behaviour of the Adaline network, as indicated in Table 2.

The performance results for the HNN, as discussed in Section 2, showed that this network has the potential to achieve high recognition capability using very few training trials.

The major difference between the tests performed in Sections 2 and 3 is that Gaussian distributed data was used in Section 2, while the patterns used in this section did not have this distribution. To improve the performance of the HNN, the distribution of its input data should be made to appear more Gaussian. Attempts at improving the performance of the HNN, based on the requirement for Gaussian distributed input data, is presented in the following section.

TOTAL NUMBER OF TRAINING PATTERNS	NUMBER OF ADDED OR MODIFIED PEAKS (Training Data)	NUMBER OF ADDED OR MODIFIED PEAKS (Test Data)	PERFORMANCE	CONFIDENCE INTERVAL in Percent (Confidence level is set at 99.0%)
300	3	0	60.0%	52.58 - 66.99
		3	22.67%	17.08 - 29.45
		6	12.3%	8.22 - 18.01
		9	11.0%	7.17 - 16.52
		12	10.0%	6.37 - 15.36
1200	3	0	90.0%	84.64 - 93.63
		3	29.3%	23.04 - 36.45
		6	13.3%	9.04 - 19.15
		9	12.0%	7.97 - 17.67
		12	9.3%	5.82 - 14.54
300	12	0	10.0%	6.37 - 15.36
		3	11.3%	7.41 - 16.87
		6	11.0%	7.17 - 16.52
		9	10.0%	6.37 - 15.36
		12	10.0%	6.37 - 15.36

Table 4 Performance results for the HNN using corrupted patterns for training, and 300 randomly selected patterns for testing.

4.0 MODIFICATION OF APPLICATION DATA TO IMPROVE PERFORMANCE

Two methods were used to try to improve the pattern recognition performance of the HNN tabulated in Section 3. First, the input patterns were transformed to make them appear more random and have less zero-valued pixels. To a human, zero-valued pixels are not weighted heavily in the recognition process. This implies that zero-valued pixels should not be part of the input patterns since they are not considered of importance to the recognition process. Zero-valued pixels have no weight in the learning process of a typical ANN. However, in the HNN, these pixels have the same weighting as nonzero-valued pixels and are thereby important to the recognition process of the HNN. The second method attempted to improve the performance of the HNN by manipulating the input patterns to make their distributions more Gaussian.

To remove the zero-valued pixels, and make the input patterns appear more random, the pixel values along the six rows were summed and the values along the 27 columns were summed to form a 33 element input stimulus vector. This summation procedure was done to each of the ten input patterns, for both training and testing. The HNN was evaluated with these transformed input data patterns, and the results are presented in Table 5.

Comparing the results in this table with those in Tables 3 and 4, we find that the transformation of the input patterns did not achieve the desired result of improved recognition performance. In fact, the performance got worse.

To make the input pattern distributions appear more Gaussian, Gaussian distributed noise was added to each pattern used in training and testing. The following range of noise levels were tried (recall that pixel values within a pattern were 0, 5, 10 or 20): 0 to 1; 0 to 0.5; 0 to 0.2; 0 to 0.1; and 0 to 0.05. Preliminary tests showed that the range 0 to 0.2 produced the best performance, and so further testing was done with this range of noise levels.

Table 6 lists the recognition performance results achieved by the HNN using corrupted and uncorrupted input patterns, with Gaussian distributed noise added to each pattern. Comparing these results with Tables 3 and 4 shows that the results are almost identical, except when 300 uncorrupted patterns were used for training. In this case, about a 5% improvement in the performance of the HNN was achieved. Although this is a small improvement, it is significant enough to indicate that if the input pattern distributions were more Gaussian, then better performance could be achieved. However, due to the small number of test patterns used, this improvement may only be the result of the variance in the performance measures, as indicated by the confidence intervals provided in each table.

TOTAL NUMBER OF TRAINING PATTERNS	NUMBER OF ADDED OR MODIFIED PEAKS (Training Data)	NUMBER OF ADDED OR MODIFIED PEAKS (Test Data)	PERFORMANCE	CONFIDENCE INTERVAL in Percent (Confidence level is set at 99.0%)
300	0	0 3 6 9 12	70.0% 23.3% 13.0% 11.0% 11.67%	62.81 - 76.32 17.63 - 30.12 8.79 - 18.81 7.17 - 16.52 7.71 - 17.29
1200	0	0 3 6 9 12	70.0% 19.67% NOT TESTED NOT TESTED NOT TESTED	62.81 - 76.32 14.44 - 26.21 N/A N/A N/A
300	3	0 3 6 9 12	20.0% 18.3% 9.0% 10.0% 8.67%	14.73 - 26.57 13.26 - 24.71 5.59 - 14.19 6.37 - 15.36 5.33 - 13.80
1200	3	0 3 6 9 12	30.0% 20.3% NOT TESTED NOT TESTED NOT TESTED	23.68 - 37.19 15.0 - 26.9 N/A N/A N/A

Table 5 Performance results for the HNN using more randomized input patterns for training, and 300 randomly selected patterns for testing. NOT TESTED implies that no test was performed because of anticipated poor results.

Comparing the results in Table 6 with Tables 1 and 2 shows that, for this application, the recognition abilities of the HNN need a lot more improvement to achieve the maximum performance of the other five ANNs. Based on the performance results given in Section 2, it is anticipated that the HNN could indeed achieve better performance results than the other ANNs in this application. This is provided an appropriate transform can be found to make the input pattern distributions appear more Gaussian.

TOTAL NUMBER OF TRAINING PATTERNS	NUMBER OF ADDED OR MODIFIED PEAKS (Training Data)	NUMBER OF ADDED OR MODIFIED PEAKS (Test Data)	PERFORMANCE	CONFIDENCE INTERVAL in Percent (Confidence level is set at 99.0%)
300	0	0	100.0%	97.84 - 100.0
		3	43.3%	36.16 - 50.73
		6	23.3%	17.63 - 30.12
		9	18.67%	13.58 - 25.12
		12	16.0%	11.29 - 22.18
1200	0	0	100.0%	97.84 - 100.0
		3	34.67%	28.0 - 42.01
		6	15.0%	10.45 - 21.06
		9	13.0%	8.79 - 18.81
		12	10.0%	6.37 - 15.36
300	.3	0	60.0%	52.58 - 66.99
		3	22.67%	17.08 - 29.45
		6	12.3%	8.22 - 18.01
		9	10.67%	6.90 - 16.14
		12	10.0%	6.37 - 15.36
1200	3	0	90.0%	84.64 - 93.63
		3	29.3%	23.04 - 36.45
		6	13.3%	9.04 - 19.15
		9	12.0%	7.97 - 17.67
		12	10.0%	6.37 - 15.36

Table 6 Performance results for the HNN using input patterns with Gaussian distributed noise added to each pattern used for training and testing. Testing was done with 300 randomly selected patterns.

5.0 CONCLUDING REMARKS

The Airborne Radar Section at DREO is interested in evaluating various techniques for recognizing high resolution radar images of ships. This report has presented performance results on the pattern recognition capability of a novel Artificial Neural Network (ANN) called the Holographic Neural Network (HNN). It has compared the HNN's performance with five more commonly used ANNs.

A brief description of the architecture and operation of the HNN was given in Section 2. The network's large storage capacity and generalization capability were demonstrated with the use of Gaussian distributed data. Based upon these results, it was anticipated that this network would be a prime candidate in achieving the goal of recognizing radar images of ships. However, when the HNN was tested with an artificial application, which resembled the pattern recognition area mentioned above, the network was found to perform poorly. When the HNN's performance was compared with the performance of the five other networks described in Section 3, it was shown to be far worse. However, based upon the excellent results reported in Section 2, it was decided to modify the image patterns used to train and test the HNN. An improvement in the recognition performance was achieved after adding Gaussian distributed noise to the input image patterns.

Further improvement in the HNN's recognition performance requires that an appropriate transform be devised which can make the input pattern distributions appear more Gaussian. One drawback to this network is that it performs best when using Gaussian distributed input data. Present time constraints prevented any further investigation into devising such a transform, but this does remain a topic for future study. Further study is justified because, as shown in Section 2, the HNN is capable of a large storage capacity requiring only a few learning trials, and the network has excellent generalization capabilities.

ACKNOWLEDGEMENTS

This work was supported by the Research and Development Branch of the Department of National Defence. Thanks also go to G. Vrckovnik for his constructive comments.

REFERENCES

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

[2] R.P. Lippmann, "Pattern Classification Using Neural Networks," IEEE Communications Magazine, November 1989, pp. 47 - 63.

[3] B. Widrow, R.G. Winter and R.A. Baxter, "Layered Neural Nets for Pattern Recognition," IEEE Trans. on ASSP, Vol. 36, No. 7, July 1988, pp. 1109 - 1118.

[4] B. Kosko, "Bidirectional Associative Memories," IEEE Trans. on Systems, Man and Cybernetics, Vol. 18, No. 1, January/February 1988, pp. 49 - 60.

[5] G.E. Hinton and J.L. McClelland, "Learning Representations by Recirculation," Proc. of IEEE Conference on Neural Information Processing Systems, November 1988.

[6] J.G. Sutherland, "A Holographic Model of Memory, Learning and Expression," International Journal of Neural Systems, Vol. 1, No. 3 (1990), pp. 259 - 267.

[7] J.G. Sutherland, "HNeT - Applications Development System for Holographic Neural Technology," Proc. of the Inter. Joint Conf. on Neural Networks, Seattle Washington, 1991.

[8] AND FILE Update, "Error Testing Procedures in Holographic Neural Technology," AND Corporation Internal Memorandum, 1990.

[9] W.E. Thorp Associates Ltd., "Evaluation of Neural Network Performance in Target Classification Applications," Report for DND Contract # W7714-9-9151, December 1989.

[10] NeuralWorks Professional II™, Neural Network Development Software by NeuralWare Inc., 103 Buckskin Court, Sewickley, PA 15143, USA.

[11] P. Meyer, "Introductory Probability and Statistical Applications - Second Edition," Addison-Wesley Publishing Company, Massachusetts, 1970.

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (the name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Establishment sponsoring a contractor's report, or tasking agency, are entered in section 8.)
Defence Research Establishment Ottawa
3701 Carling Ave
Ottawa, Ont K1A 0Z4

2. SECURITY CLASSIFICATION (overall security classification of the document, including special warning terms if applicable)
UNCLASSIFIED

3. TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C or U) in parentheses after the title.)
The Holographic Neural Network - Performance Comparison with Other Neural Networks(U)

4. AUTHORS (Last name, first name, middle initial)
R.D. Klepko

5. DATE OF PUBLICATION (month and year of publication of document)
October 1991

6a. NO. OF PAGES (total containing information. Include Annexes, Appendices, etc.)
32

6b. NO. OF REFS (total cited in document)
11

7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)
DREO Technical Note

8. SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include the address.)
Defence Research Establishment Ottawa
3701 Carling Ave
Ottawa, Ont K1A 0Z4

9a. PROJECT OR GRANT NO. (if appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant)
Project # 021LA

9b. CONTRACT NO. (if appropriate, the applicable number under which the document was written)
N/A

10a. ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique to this document.)
DREO TECHNICAL NOTE 91-18

10b. OTHER DOCUMENT NOS. (Any other numbers which may be assigned this document either by the originator or by the sponsor)
N/A

11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

(X) Unlimited distribution
( ) Distribution limited to defence departments and defence contractors; further distribution only as approved
( ) Distribution limited to defence departments and Canadian defence contractors; further distribution only as approved
( ) Distribution limited to government departments and agencies; further distribution only as approved
( ) Distribution limited to defence departments; further distribution only as approved
( ) Other (please specify):

12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availabilty (11). However, where further distribution (beyond the audience specified in 11) is possible, a wider announcement audience may be selected.)
FULL UNLIMITED ANNOUNCEMENT

13. ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), or (U). It is not necessary to include here abstracts in both offical languages unless the text is bilingual).

Artificial Neural Networks (ANNs) use weighted interconnections of computational elements in an attempt to mimic the ability of the human brain to learn pattern associations. The type and architecture of an ANN is defined by the following items: the coding and decoding schemes used for the inputs and outputs, respectively; the interconnection pattern between the elements; the computation performed by each element; and the method used to change the weighting of values passed on the connections (i.e. the learning mechanism). These networks can be trained, or can train themselves, to perform tasks such as image recognition, speech understanding, robot control, and data filtering, just to name a few of the more popular areas of application. The ANN of interest in this report is the Holographic Neural Network (HNN). This network is specifically designed for pattern association storage, and its designers claim that it has an exceptional generalization capability and storage capacity. The HNN mathematically realizes the storage and retrieval mechanisms of holograms; hence the origin of its name. The operation of this network is briefly described within this report.

In addition to describing the HNN, this report presents pattern recognition performance results of the HNN and compares them with the results achieved by other more common ANNs: Adaline, Hamming, Bidirectional Associative Memory, Recirculation, and Back Propagation. The patterns used for testing represent artificial high resolution radar images of ships. These images appear as a two-dimensional topology of peaks with various amplitudes.

The performance comparisons show that the HNN does not perform as well as the other ANNs when using the same test data. However, modification of the data, to make it appear more Gaussian distributed, does improve the performance of the HNN. It is shown that the HNN performs best if the input data is completely Gaussian distributed.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Neural Networks
Pattern Recognition
Holographic Neural Network
Performance
Adaline
Hamming
Bam
Recirculation
Back Propagation
Synthetic Aperture Radar
Ship Images

NO. OF COPIES NOMBRE DE COPIES	COPY NO. COPIE N°	INFORMATION SCIENTIST'S INITIALS INITIALES DE L'AGENT D'INFORMATION SCIENTIFIQUE
1	1	JL

AQUISITION ROUTE FOURNI PAR ▶	DREO
DATE ▶	17 JANUARY 1992
DSIS ACCESSION NO. NUMÉRO DSIS ▶	92-00969

DND 1158 (6-87)

103862

National Defence / Défense nationale

~~PLEASE RETURN THIS DOCUMENT TO THE FOLLOWING ADDRESS:~~ ~~PRIÈRE DE~~ [redacted] ~~CE DOCUMENT À L'ADRESSE SUIVANTE:~~

~~DIRECTOR~~
~~SCIENTIFIC INFORMATION SERVICES~~
~~NATIONAL DEFENCE~~
~~HEADQUARTERS~~
~~OTTAWA, ONT. - CANADA K1A 0K2~~

~~DIRECTEUR~~
~~SERVICES D'INFORMATION SCIENTIFIQUES~~
~~QUARTIER GÉNÉRAL~~
~~DE LA DÉFENSE NATIONALE~~
~~OTTAWA, ONT. - CANADA K1A 0K2~~

CRAD / DRDIM 3
DEPARTMENT OF NATIONAL DEFENCE
OTTAWA ON K1A 0K2
CANADA