NEURODYNAMICAL SYSTEMS FOR COGNITION AND TARGET IDENTIFICATION

FINAL TECHNICAL REPORT

N.H. Farhat

OCTOBER, 1994

U.S. ARMY RESEARCH OFFICE

DAAL03-91-G-0209

UNIVERSITY OF PENNSYLVANIA

APPROVED FOR PUBLIC RELEASE;
DISTRIBUTION UNLIMITED.

NEURODYNAMICAL SYSTEMS FOR COGNITION AND TARGET IDENTIFICATION

FINAL TECHNICAL REPORT

N.H. Farhat

OCTOBER, 1994

U.S. ARMY RESEARCH OFFICE

DAAL03-91-G-0209

UNIVERSITY OF PENNSYLVANIA

APPROVED FOR PUBLIC RELEASE;
DISTRIBUTION UNLIMITED.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE
October, 1994

3. REPORT TYPE AND DATES COVERED
Final - 7/1/91 to 6/30/94

4. TITLE AND SUBTITLE
NEURODYNAMICAL SYSTEMS FOR COGNITION AND TARGET IDENTIFICATION

5. FUNDING NUMBERS
DAAL03-91-G-0209

6. AUTHOR(S)
N.H. Farhat

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Pennsylvania
Department of Electrical Engineering
200 South 33rd Street
Philadelphia, PA 19104-6390

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
U. S. Army Research Office
P. O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSORING/MONITORING AGENCY REPORT NUMBER
ARO 28125.9-MA

11. SUPPLEMENTARY NOTES
The view, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation.

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

Our study of cognitive automated target recognition based on the neural paradigm for information processing reveals that inclusion of bifurcation and synchronicity (phase-locking) in network dynamics can markedly improve the performance of ATR systems. This gave impetus to our study of how synchronicity could arise in cortical networks when it is known the brain has no central clock. Raising this question has led us, through analysis of models of biological neurons employing the tools of non-linear dynamics, to the development of the bifurcating neuron concept and model. This spiking neuron model combines functional complexity comparable to that of biological neurons with structural simplicity and low power consumption when implemented electronically or optoelectronically. These attributes make the bifurcating neuron ideally suited for use as building block of a new generation of spiking neural networks that employ phase-locking, bifurcation and chaos, on the single processing element level, to emulate higher-level cortical functions such as feature-binding and cognition that are essential for advanced ATR systems, and other operations like separation of object from background, inferencing and rudimentary reasoning.

14. SUBJECT TERMS
Advanced ATR (Automated Target Recognition), Complexity, Dynamical Networks, Cognition, Bifurcation, Synchronicity Chaos, the bifurcating neuron model and bifurcating neural nets, Higher-level processing.

15. NUMBER OF PAGES
80

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	**18. SECURITY CLASSIFICATION OF THIS PAGE**	**19. SECURITY CLASSIFICATION OF ABSTRACT**	**20. LIMITATION OF ABSTRACT**
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

1. FOREWARD

A justifiable criticism of artificial neural net models frequently voiced by biologists and neuroscientists is that they are minimal in nature as evident in the extreme functional simplicity of the neuron models employed in comparison to the biological neuron.

Neural networkers are quick to respond to this criticism by pointing out that despite such simplification, neural networks consisting of simple processing elements (neurons) exhibit rich collective emergent properties and that significant progress in machine learning, associative storage and recall, and solution of optimization problems have taken place in the past decade leading to significant growth in basic knowledge about self-organizing systems and collective computing and to realistic applications.

Despite this progress, neural networks continue to be plagued by several widely acknowledged limitations. These include (a) general inefficiency of learning algorithms, (b) inability to handle spatio-temporal information in a natural way, and (c) general inability to provide higher-level functionality such as feature binding, cognition, distortion invariance*, separation of object from ground (background), inferencing, reasoning and other functions known to be carried out by the cortex almost effortlessly. It is reasonable to assume that the functional complexity of the cortex is a consequence of both the functional complexity of cortical neurons and the intricate interaction patterns between different neuronal pools in the cortex.

Motivated by these observations, and by our findings in the study of cognitive networks for automated target recognition, we have carried out a study aimed at producing biology-oriented neuronal models that duplicate as much as possible the functional complexity of the living neuron while being realized in a structurally simple and power efficient embodiment. The availability of such functionally complex but structurally simple neurons of low power consumption can lead to computing structures (neural networks) in which one can model and study the dynamics of cortical networks and some of the higher-level processing functions they exhibit. Introducing higher-level functionality in neural networks will significantly enhance the power of neurocomputing, leading to a host of new applications, and emphasizing the viability of the neural paradigm for information processing.

The results of the above study was the development of the bifurcating neuron concept and model. Our work to date shows that the bifurcating neuron combines functional complexity approaching that of the biological neuron with structural simplicity and low power consumption because of its spiking nature. All of these are attractive attributes for simulation or hardware implementation of a new generation of neural networks possessing greater functional complexity and computing power than present day networks and specially suited for study and development of higher-level functionality. To date our work shows that under periodic activation the bifurcating neuron is capable of firing in several modalities and can bifurcate (rapidly switch) between these modalities depending on the nature of its input. As such, it appears capable of encoding its spatio-temporal input, the aggregate of all spike trains incident at any time on synaptic sites of its dendritic-tree, (which we call incident spike wavefront), in a complicated manner. This functional complexity stems from the ability of certain incident spike wavefronts to produce periodic episodes in the neuron's activation potential. The focus on periodic activation stems from the fact that in a population of synchronized (phase-locked) bifurcating neurons, the activation potentials formed by dendritic-tree processing are periodic. Depending on the nature of the incident spike wavefront, i.e., whether it is incoherent, partially coherent or coherent**, the bifurcating neuron can behave

* Invariance of object or signal recognition in the presence of changes in object size, orientation, position, and signal-to-noise ratio.

** A coherent incident spike wavefront is one in which all the spike trains incident on the neuron are correlated.

as a sigmoidal neuron or as a periodically driven oscillator neuron capable of producing a host of regular phase-locked firing patterns or chaotic firing. There is mounting evidence, from physiological observations and numerical simulation, that phase-locked (synchronized) firing states of cortical networks underlie cognitive functions and that chaos might be playing a useful role in their dynamics. We expect the functional complexity of the bifurcating neuron to manifest itself in the complexity of operations and computing power of bifurcating neural networks which are specially suited for use in the modeling and study of cortical functions.

2. TABLE OF CONTENTS

3. NEURODYNAMICAL SYSTEMS FOR COGNITION AND TARGET IDENTIFICATION

The research effort described in this final report was concerned with the study and development of algorithms and systems for automated target recognition based on the neural paradigm for information processing that are specifically intended to operate in *complex uncontrolled* environment like that frequently encountered in automated target recognition (ATR), robotics, and autonomous systems in general. An automated recognition system for identifying handwritten zip code numerals for the postal service constitutes an example of an automated recognition system operation in a *complex controlled* environment. Complex, because of the wide variation in handwriting between individuals; controlled, because the system is strictly designed to recognize handwritten zip code numerals, and once operational no one is going to recognize anything else other than handwritten numerals. Another example of a *complex controlled* environment occurs in automated recognition by industrial robots of manufactured parts. Clearly there are many situations where an automated identification system is required to operate in the more challenging complex but *uncontrolled* environment where it can encounter objects or patterns other than those it was intended for. In such instances the recognition task becomes considerably more difficult.

We discuss next the reasons for this difficulty, then go on to describe earlier work we have carried out to overcome these difficulties by adopting an approach based on nonlinear dynamical systems and certain general attributes of higher-level cortical information processing. Finally we discuss how this earlier work has led us to develop the concept of bifurcating neuron as a building block for a new generation of neural networks suitable for the study of higher-level functions such as feature-binding and cognition. We also give a summary of the most important results obtained from a detailed investigation of the bifurcating neuron concept.

A. Background and Statement of the Problem Studied: The first step in any automated object recognition system is feature-extraction which is the production of invariant object features from sensory data. The invariance is with respect, distance, orientation, displacement and signal-to-noise ratio (SNR) which includes illumination level and variability. The invariant features are needed to make the recognition system robust. The literature and methodology of feature-extraction is quite varied and extensive and it is not the intent to discuss it here. For our purposes here, it suffices to note that once a suitable working feature extraction method is selected the next and crucial step in the automated recognition process is feature-binding or linking where the identity of the object is inferred from its invariant feature vector. The most straight forward means for feature-binding is a look-up table where the feature vector of an unknown object is compared against a library of feature vectors belonging to objects known to occur in the system's working environment. A best fit criterion is used then to identify the object. A second, and more sophisticated approach to feature-binding is to use a multilayer feed-forward neural network, usually trained by an error-back-propagation, (e.b.p.) algorithm, to map the feature-vectors of a set of objects into associated identifying labels. When this training is carried out properly, the resulting network has generalization ability and certain level of robustness with SNR.

It is well known to practitioners in the field that both the look-up table approach and the feed-forward e.b.p. neural net classifier approach can not be used in systems intended to work in an uncontrolled environment. There are two reasons for this. One is that the number of objects that can occur in an uncontrolled environment is not limited but can be very large indeed and the system must be able then to distinguish between all the possible objects or at least between the set of objects it is designed for and novel objects. This usually makes the learning task very complicated and lengthy if not impractical because learning in neural networks is NP-complete which means that learning time and network complexity grow exponentially with the size of the learning task, i.e. with the number and complexity of the objects the system must learn. This constitutes a major issue in neurocomputing (artificial neural networks (ANNs)) and machine learning in general and is summarized by the simple question: *how can effective learning be achieved in a network or machine intended to operate in a complex uncontrolled environment.*

Our studies indicate that the problem of learning in complex uncontrolled environments may be traced to the fact that most ANNs and learning algorithms today have no cognitive ability. By *cognition*, is meant here *ability of the network to distinguish on its own between familiar objects, i.e., objects belonging to its training set and novel objects not belonging to the training set.* In many operating environments of practical interest, the occurrence of novel objects is unavoidable. The danger then is that without cognition an ANN can end up misclassifying a novel object as one belonging to its training set and this is obviously not acceptable and can be even catastrophic in certain situations. To overcome this problem, the training of ANNs or machine is often modified to either: (a) Include training on negative examples, i.e., on the class of novel objects that could occur in the ANNs environment. This approach is unproductive because it increases the size of the network and training time becomes unacceptably long. (b) Incorporation of novelty detectors that would detect and measure attributes of the objects that could help in deciding whether an object is novel or not. This approach is unattractive because novelty detectors often add complexity and cost to the system.

To make progress in this difficult problem we have adopted a nonlinear dynamical system approach to feature-binding and cognition which leads to ways of circumventing the issue of NP-completeness of learning. The approach draws on attributes of cortical information processing. The cortex is that part of the brain where higher-level functions, such as feature binding, cognition, reasoning and all the other interesting complex information processing functions we humans do, are believed to reside. Cortical neurons and populations are nonlinear and highly interconnected. Therefore one can view the cortex as high-dimensional nonlinear dynamical systems. Nonlinear dynamical systems exhibit three types of phase-space attractors: Point, periodic, and chaotic. Most attractor type neural net models being dealt with today employ point attractors to provide associative memory, optimization, and learning functions, but lack cognition. An inavoidable question then is: what role could periodic, and chaotic attractors play and could they be used to achieve higher-level neural functions such as feature-binding and cognition, and how could they be incorporated in the design of ANNs to enhance their performance by enabling them to compute with such attractors. The pioneering work of Freeman and co-workers (see for example: C. Skarda and W. Freeman, Behavioral and Brain Sciences, 10, 161-165, Cambridge Univ. Press, 1987) suggests that bifurcation in networks that compute with diverse attractors could be a mechanism for cognition. We have applied this hypothesis successfully to the design of a composite cognitive neural network for automated target identification [1] (see also more detailed account given in Appendix I) which provided distortion invariant identification of microwave test objects from a single echo or signature, thus solving the long-standing problem of 3-D object recognition independent of range, orientation, signal-to-noise ratio, and location within the field of view for this particular sensing and recognition modality. This network computes with diverse attractors and is capable not only of differentiating between and identifying familiar objects successfully, but also employs bifurcation* between a periodic attractor and a point attractor as the mechanism for feature-binding and cognition; and differentiating between familiar and novel objects. An important aspect of the system is the use of segmentation of the signature vector (echo or response vector of the target for a given aspect) during the training and interrogation phases in order to avoid ambiguities and enhance the probability of recognizing novel objects as such without sacrificing performance in recognizing learned (familiar) objects.

* Bifurcation means sudden change in behavior or computing modality depending on change in a parameter of the system (here whether the signature presented to the network belongs to a familiar or novel object.

The bifurcation/cognition capability in the system we just described was furnished by a periodic attractor network which required synchronous updating of the neurons for proper operation and delicate setting of learned weights to make novel objects trigger the bifurcation from periodic to point attractor. Although there is no problem in providing synchronous update in a real neural system, the question of how would synchronous update occur in cortical networks is a valid one to raise in this connection because the original approach in our designing cognitive nets, was brain-inspired. Since it is generally agreed that the brain does not contain a central clock and there is no evidence that the α rhythm serves such function, one can ask next: how could synchronicity and coherence spontaneously emerge in cortical networks especially when noise in biological (cortical) neurons is known to cause them to respond inconsistently to the same repeated stimulus? Raising this question has led us to develop the concept of *bifurcating neuron* as a model of the excitable biological membrane which is capable of providing synchronicity through phase-locking and of exhibiting functional complexity paralleling that of the living neuron but in an extremely simple and power-efficient structure which is important for hardware implementation of cortical neuron models and networks. Bifurcations between attractors in such networks could provide a more natural and reliable mechanism for feature-binding and cognition then the aforementioned periodic attractor network and may have other useful applications.

B. Summary of the Most Important Results: Adopting the nonlinear dynamical systems view of the cortex and applying it to the ATR problem and to neurocomputing in general has so far led to the following accomplishments in our work:

1. We introduced a new function to neural networks to be added to the repertoire of functions they already possess: association, optimization, and learning with generalization. We add now cognition and this enhances the power of neurocomputing because: **(a)** Without cognition a neural-based identification system, intended to operate in a complex uncontrolled environment is useless because a novel object can trigger erroneously the response belonging to one of the familiar (learned) objects. **(b)** With cognitive ability the system can be made to respond more appropriately; for example, ignore its response in instances of novel objects or alter its mode of operation by reverting to a learning mode where it can proceed to learn the novel object when it occurs, and add it to its repertoire. **(c)** With cognition one can consider now designing banks of relatively small neural networks (neural modules) which can be trained to recognize only a subset of the objects in the environment to the exclusion of all others. This leads to neural modules of manageable size, each designed to recognize a small set of objects with the entire assembly of modules being able collectively to recognize a larger set of objects. This is perhaps the most significant implication for introducing cognition in neural systems. The training time of the smaller neural modules is considerably shorter than learning the entire problem with one large network which for many practical-sized problems is not feasible with present-day learning algorithms. Cognition circumvents therefore the scaling problem associated with learning large tasks, which as stated earlier is NP-complete. **(d)** Cognition provides a system with a rudimentary level of awareness of its environment and this is a step in the direction of imparting other higher-level functions to neural networks.

2. Development of the concept of bifurcating neuron [2]-[4] that combines functional complexity paralleling that of the living neuron with structural simplicity that facilitates hardware implementations, opens the way to constructing a new generation of neural networks that could exploit synchronicity and coherence in performing higher-level functions and which can employ all three types of attractors to achieve such functions. This would introduce essentially a new paradigm in neurocomputing where complexity, bifurcation, and chaos on the single neuron level become important aspects of neurocomputing.

Specific accomplishments in our bifurcating neuron and bifurcating neural networks research are:

- Developed a bifurcating neuron theory that is descriptive, predictive, and quantitative.
- Developed analytical and numerical simulation tools for characterizing the way a bifurcating neuron encodes periodic components or episodes appearing in its activation potential. Periodic activation arises when a network of bifurcating (spiking) neurons enters phase-locked firing. The characterization is mostly in terms of a bifurcation diagram. (see discussion below).
- Obtained increasing evidence that erratic (or chaotic) firing of the bifurcating neuron, which occurs under specific periodic activation conditions, can be a source of adaptive noise for annealing bifurcating nets, i.e., can aid network entrainment (help it arrive at a phase-locked firing state) which is analogous to annealing of sigmoidal nets into states of local or global energy minima in order to arrive at optimal solutions.

The bifurcating neuron effort was also motivated by the observation that the functional complexity of present-day dynamical (recursive or attractor-type) neural networks stems primarily from the collective behavior of neurons that are functionally simple nonspiking processing elements (e.g. sigmoidal or binary (McColloch-Pitts) neurons). In contrast, biological neurons in the cortex, where feature-binding, cognition, inferencing and other higher-level processing are believed to take place, are functionally very complex processing and encoding elements. It is reasonable to believe that the functional complexity of such neurons, traceable to the rich and complex dynamics of the driven excitable biological membrane responsible for their spiking behavior, would underlie the functional complexity and collective computating power of cortical networks. Development of artificial model neurons that emulate the functional complexity of the cortical neuron is therefore desirable because it yields the *ultimate processing element* for use as building-block in a new generation of neural networks that compute with diverse attractors and seek to achieve vastly enhanced processing and learning power. The spiking nature of neurons in such networks would enable preserving the relative timing of action potentials and the introduction of notions of coherence, synchronicity and phase-locking. The emergence of synchronicity and coherence means that neurons in such networks can find themselves being subjected to correlated incident spike patterns which give rise via linear and nonlinear dendritic-tree processing (filtering and smoothing operations) to periodic activation potentials that drive the excitable "membrane" dynamics of the neuron which is the origin of complexity alluded to earlier.

Thus motivated by these observations and also by the results of our preceding work on *cognitive automated target recognition* (ATR) [1] we have carried out a systematic study aimed at producing biology-oriented neuronal models that preserve as much as possible of the signal-processing-related functional complexity of real cortical neurons but can be realized in structurally simple and power efficient embodiment. The result was the bifurcating neuron model. The investigation involved analyzing the dynamics of the periodically driven integrate-and-fire (I&F) model neuron, a mono-ionic simplification of the well known Hodgkin-Huxley model for action potential generation in the excitable biological membrane and revealed that the firing behavior can be described by an iterative map of the phase interval [0-2π] onto itself which we call a *phase-transition map* (PTM) [2]-[4]. Like other maps of the interval onto itself, the PTM can be studied employing the tools of *nonlinear dynamics*. This provides a novel way for viewing and characterizing the micro-neurodynamics (neurodynamics on the single neuron level) in recursive networks in terms of a *bifurcation diagram* which shows the extreme functional complexity of the periodically driven I&F neuron model that is achieved despite its relatively simple structure. In the absence of periodic activation the I&F neuron reverts to the usual sigmoidal response (sigmoidal dependence of firing frequency on activation potential). Because the complex behavior of such model neuron can be described best by a bifurcation diagram we have named it the *bifurcating neuron.*

To achieve these results we developed unique analytical, simulation, and experimental tools for characterizing the performance of several embodiments of the periodically driven I&F neuron. As a result we were successful in developing a bifurcating neuron circuit whose bifurcation diagram (see Figure 1) exhibited, *full-blown* chaotic firing in addition to several phase-locked periodic firing modalities that include period-m phase-locked firing, aperiodic firing, and intermittency, i.e. a complex range of firing modalities, depending on parameters of the driving signal. In this diagram θ_n is the relative-phase of the n-th spike fired by the neuron measured relative to the immediately preceding peak (or zero crossing) of the periodic driving signal. The parameters f_s and a are respectively the frequency and amplitude of the cosinusoidal driving signal. The bifurcating neuron circuit employed utilized time-delayed modulation of the restoring current source (circuit diagram omitted for lack of space). The chaotic firing ability of this neuron was verified by computing the Lyapunov exponent of the orbits θ_n, n=1,2 .. observed for certain values of the (f_s, a) parameters. For complete description of the behavior of the bifurcating neuron one needs obviously a set of such bifurcation diagrams, one for every possible value of the amplitude a. Contrasting this with the simple transfer function of firing frequency vs. activation potential for sigmoidal neurons gives immediately an idea of the complexity and richness of behavior one can expect to observe in bifurcating neural networks. Learning to harness such richness and complexity to achieve feature-binding, cognition, and other higher-level functions is the goal of our research.

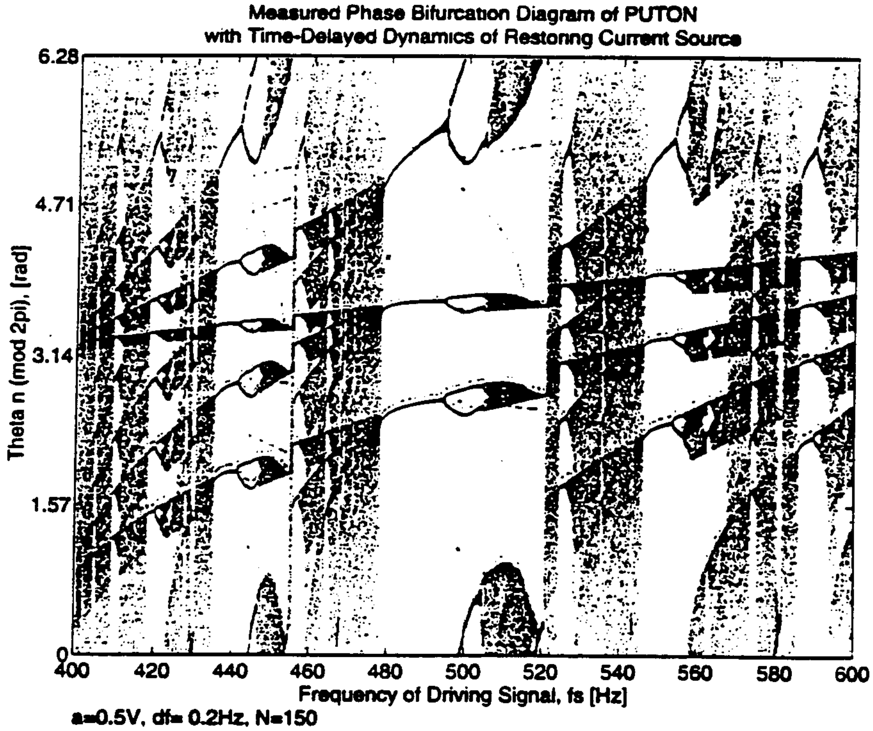

Fig. 1. Measured bifurcation diagram for a Programmable Unijunction Transistor Oscillator Neuron (PUTON) embodiment of the bifurcating neuron employing time delayed dynamics of restoring current source. The fractal (self-similar) and complex structure of the diagram, which includes phase-locked ordered firing and chaotic firing regimes, promise to make the bifurcating neuron the processing element of choice in *dynamical neurocomputers* that compute with diverse attractors and employ synchronicity, bifurcation and chaos in their operation in order to achieve significant improvement in capabilities and performance over present-day neural networks especially for feature-binding and cognition.

Continuation of the research reported here is being focused on further simplification of the chaotic bifurcating neuron circuit. Our goal is to develop the simplest bifurcating neuron circuit that can serve as paradigm for complexity and chaos on the single neuron level in dynamical artificial spiking neural networks. A more general and long-term goal of our research program is to demonstrate that such functional complexity on the single processing element level is the instrument by which significant enhancement of the capabilities of present-day neural networks can be achieved in order to make them more suitable for use in solving practical problems besides feature-binding and cognition, like continuous speech processing, complex control, and in many other diverse applications such as modeling and simulation of populations of coupled biological oscillators for better understanding of biological clocks and cardiac dynamics and arrythmia. We believe we are at the dawn, if not in the midst, of a new era in computing, that of *dynamical computing* in networks of structurally simple but functionally complex processing elements.

C. List of All Publications and Technical Reports:

1. N.H. Farhat and B. Bai, "Optimal Spectral Windows for Microwave Diversity Imaging," IEEE Trans. on Ant. and Prop., vol. 39, pp. 985-993, July, 1991.

2. N.H. Farhat and H. Babri, "Cognitive Networks for Automated Target Recognition and Autonomous Systems Applications," Also presented by invitation at PIERS'93 Proc. Progress in Electromagnetic Research Symposium, Pasadena, CA (1993) and at the Second Government Neural Network Applications Workshop, Huntsville, AL, Sept. 1993.

3. N.H. Farhat and M. Eldefrawy, "The Bifurcating Neuron: Characterization and Dynamics," in Photonics for Computers, Neural Networks and Memories, SPIE Proceedings, vol. 1773, SPIE, Bellingham, Wash. (1992) pp. 23-35.

4. B. Bai and N.H. Farhat, "Learning Networks for Extrapolation and Radar Target Identification," Neural Networks, vol. 5, pp. 507-529, 1992.

5. Z. Zhao and N.H. Farhat, "Tomographic Microwave Diversity Image Reconstruction Employing Unitary Compression," IEEE Trans. on MTT, vol. 40, pp. 315-322, Feb. (1992).

6. N.H. Farhat, S-Y Lin, and M. Eldefrawy, "Complexity and Chaotic Dynamics in a Spiking Neuron Embodiment," in Adaptive Computing: Mathematics, Electronics, and Optics, J. Caulfield (Ed.), SPIE, (1994), pp. 77-88 (**Invited**). Also presented in the Proc. 1993 Int. Symp. on Nonlinear Theory and Its Applications (NOLTA'93), (**Invited**).

7. N.H. Farhat, M. Eldefrawy and S-Y Lin, "Theory of a Bifurcating Model Neuron: A Nonlinear Dynamic Systems Approach," Transactions of the 11th Army Conference on Applied Mathematics and Computing, Sponsored by the Mathematics and Computer Sciences Division, ARO Report 94-1 (1994), pp. 145-197. Also to be published in Origins: Brain and Self-Organization, K. Pribram (Ed.), L. Erlbaum Assoc. Publishers, Hillsdale, N.J. (1994), (**Invited**).

D. List of All Participating Scientific Personnel and Advanced Degrees Earned:

H. Babri - Ph.D. Degree Awarded (1992)

M. Eldefrawy - Ph.D. dissertation completion date, Dec. 1994

S-Y Lin - Ph.D. dissertation completion date, Oct. 1994

N.H. Farhat - Principal Investigator

4. REPORT OF INVENTIONS

No applications for inventions were made during the period of this final report. However, shortly before the start of this grant the following related invention was assigned to the University of Pennsylvania.

"Super-Resolution and Signal Recovery Using Models of Neural Networks"
U.S. Patent No. 4,995,088
Asignee: University of Pennsylvania
Issued: February 19, 1991

5. BIBLIOGRAPHY

1. N. Farhat and H. Babri, "Cognitive Networks for Automated Target Recognition and Autonomous System Applications," Proc. Int. Joint Conf. on Neural Networks, Vol. 1, p. 535, 1992, also in Proc. of the 1993 Progress in Electromagnetic Res. Symp., Pasadena, CA, (1993). **(Invited)**.

2. N.H. Farhat, S-Y Lin and M. Eldefrawy, "Complexity and Chaotic Dynamics in a Spiking Neuron Embodiment," SPIE, Vol. CR55, Adaptive Computing: Mathematics, Electronics, and Optics, To be published by SPIE, Bellingham, April (1994), **(Invited)**.

3. N. Farhat, M. Eldefrawy and S-Y Lin, "Theory of Bifurcating Neuron: A Nonlinear Dynamic System Approach," Proc. of the Second Appalachian Conference on Behavioral Neurodynamics, to be published by L. Erlbaum Assoc. Publishers, Hillsdale, NJ (1994) **(Invited)**, (in press). This paper was also presented at the Eleventh Army Conference on Applied Mathematics and Computing, Carnegie Mellon University, June 1993.

4. N. Farhat, "Complexity and Chaotic Dynamics in a Photonic Neuron Embodiment," Proc. Int. Symp. on Nonlinear Theory and Applications (NOLTA'93), Vol. 1. Published by the Research Society of Nonlinear Theory and Its Applications (NOLTA) and The Institute of Electronics, Information and Communications Engineers (IEICE), Japan 1993, pp. 91-97. **(Invited)**.

APPENDIX I

Cognitive Networks for Automated Target Recognition

Haroon A. Babri* Nabil H. Farhat

University of Pennsylvania
Electrical Engineering Department
200 South 33rd Street
Philadelphia, PA 19104

*Now at the School of Electronic and Electrical Engineering, Nanyang Technological University, Nanyang Avenue, Singapore 2263

Abstract

We define a cognitive neural network as one capable of not only differentiating between familiar objects (those it has been trained on) but to also differentiate on its own between familiar and novel objects (the set of all other objects). We maintain that imparting such cognitive ability to neural networks has far reaching implications on the ability to design practical networks. We illustrate our thesis by an example of designing a composite hierarachial network for cognitive automated target identification. The main thesis is: By imparting cognition to a network we control the set of objects within its awareness domain. The awareness domain is defined as the set of all objects the network is supposed to identify correctly. We show that by combining cognition with segmentation and bifurcation in a dynamical network that computes with diverse attractors we are able to circumvent the scaling problem associated with learning practical problems.

1 Introduction

A longstanding problem in pattern recognition that has resisted satisfactory solution for a long time is that of recognizing three-dimensional objects irrespective of orientation (aspect), distance (range), location within the field-of-view (f.o.v.) and signal-to-noise ratio (SNR). This problem has come to be known as *distortion invariant recognition* and belongs to the class of *inverse problems* (see for example [1], [2]).

In this paper we present a solution to this problem in the context of Automated Target Recognition (ATR) of 3-D radar scattering objects. There are two approaches to distortion invariant recognition of 3-D microwave scattering objects. One is based on forming images to be identified by human observers. The second approach involves machine recognition from individual *signature vectors* of the target. A discussion of these approaches and the reasons why the signature vector approach is preferable, specially when the neural paradigm is applied, are given in [3]. There it is also argued that the neural paradigm has potential for obviating the imaging approach altogether because it circumvents the practical and cost limitations of the latter.

The paper is organized as follows. In Section 2, we discuss what cognition means and why it is important in the context of automatic target recognition (ATR) and other applications. Section 3 briefly describes the ATR concept and the associated terminology and framework. This will enable one to the understand examples from ATR used to illustrate certain issues in learning, which we discuss in section 4. In section 5 we examine the potential of garden variety neural networks as applied to the ATR problem and see that for a satisfactory solution of the problem one has to appropriately address the issues of generalization, cognition, and robustness. This is the focus of Section 6, in which we discuss how this can be achieved by computing with diverse attractors and using multisensory information. In section 7, we describe how a practical ATR system can be designed. A design example is given in section 8. Section 9 gives the conclusions and discusses the contributions of our work. In the appendix, we briefly describe the periodic attractor network.

2 The importance of cognition

In everyday usage, *cognition* is usually defined as the act or process of knowing, perceiving or becomimg aware of something. In the context of our work, cognition means the ability of the system or algorithm to tell on its own when the viewed object is familiar or novel. More specifically, in the ATR context, cognition is the ability of a machine being able on its own, without the use of auxiliary novelty detectors, filters or other gear, to tell that data presented to it belongs to a familiar or novel object. In the context of cognitive neural networks, a familiar object is one that belongs to the training set. A novel object is one that does not belong to the training set.

Cognition is important due to reasons we enumerate below.

- It introduces a new function to neural networks to be added to the repertoire of functions they possess now, i.e. association, optimization, learning and generalization. Adding cognition enhances the power of neural information processing.

- Imparting cognition to a neural network is important in pattern recognition when the network is required to operate in a *complex uncontrolled environment*; it enhances the capabilities of *autonomous systems*. The ATR environment is an example of a complex uncontrolled environment whereas recognition of handwritten zip code numbers (in some postal setting) is an example of complex controlled environment.

- It helps mis-identifying a novel object. Without cognition, a neural-based identification system operating in a complex uncontrolled environment is useless because a novel object can trigger the response identifying one of the familiar (learned) objects.

- With cognitive ability a neural net system can be made to respond appropriately, for example by giving an indication to disregard the network's response in instances of novel objects. In certain situations, this can be used as a cue to alter the network's mode of operation by reverting to a learning modality (when unsupervised learning is involved) where it can proceed to learn the novel object.

- Cognition, combined with smart sensing, segmentation, and bifurcation in dynamical neural networks that compute with diverse attractors, solves as shown below the longstanding problem of distortion invariant recognition of 3-D objects in the context of ATR and enables circumventing scaling problems related to learning when designing practical autonomous ATR systems. With cognition one can consider now designing banks of relatively small neural networks (neural modules) which can each be trained to recognize a subset of the set of all objects, i.e. a finite set of objects to the exclusion of all others. This results in neural modules of manageable size, each designed to recognize a small set of objects, with the entire assembly of modules being able to recognize a large set of objects.

- Cognition imparts to a system a rudimentary level of awareness of its environment. The set of all objects that can induce a response in the cognitive system is divided into two sets: the *targeted* or *crucial* set which the system is specifically designed to recognize, and the *untargetted* or *non-crucial* set (consisting of all the other objects that can possibly occur in the system's environment) and to which the system is not intended to respond.

- Cognition and the ensuing level of awareness resulting from it is a step in the direction of imparting higher-level function to neural networks.

The philosophy of the approach followed in our method to achieve cognition is to apply the power of nonlinear dynamical systems to the problem while being guided by broad general features of biological signal processing known to us today. One such general feature is that feature-extraction in early stages of our sensory system with the exception perhaps of the olfactory system, is carried out by predominantly feedforward networks, while feature binding and cognition are carried out by cortical networks involving heavy feedback and nonlinearity which makes them essentially nonlinear dynamical information processors. The second general feature is the possible occurence of segmentation of data in the various sensory mappings formed by the early stages of our sensory system. The third general feature is that our brains use and fuse multisensory information to overcome ambiguities (and possibly also for unsupervised learning). Our method of solution entails evidence in support of the hypothesis we have made earlier [4], that *in order to make a neural net cognitive it must be nonlinear, dynamical, and capable of computing with diverse attractors and be able to bifurcate between them depending on whether the input presented to the network is familiar or novel.* Introducing cognition in neural networks increases their signal processing power and obviates the need to use novelty detection or novelty filters which usually entail auxiliary equipment that adds to system cost. Achieving cognition turns out to be intimately related to the ability to exert control over the *phase-space trajectory* and hence over the behavior of the network. We call this *Phase-Space Engineering*: the art of synthesizing prescribed trajectories in the phase-space of a network through control of network parameters. Achieving distortion invariant recognition turns out to be intimately related to data acquisition and representation issues.

This inability of a network to distinguish independantly between familiar and novel objects may be termed as its lack of cognition and is one of the major outstanding issues in pattern recognition. The second major issue is how to achieve distortion invariant recognition which is often referred to as displacement, rotation, scale, and SNR (signal-to-noise ratio) independant recognition. Both issues are of crucial importance in remote sensing and autonomous systems that are meant to operate in a complex uncontrolled environment. Both issues also have consistently resisted attempts at their solution for a long time. The third issue basically defines the spectrum of problems to which neural networks can be applied with great advantage. It also affords a criterion for evaluating the capabilities of a given neural network model when applied to a subset of these problems.

3 The ATR Concept

The Automated Target Recognition (ATR) problem is one of longstanding interest and aims at recognizing radar targets irrespective of aspect or orientation and range from the radar, and in the presence of noise and clutter.[1] Historically, there have been two approaches to this problem. One consists in attempting to recognize a target from its image. To obtain a good enough image the hardware requirement is that of synthesizing a large enough aperture, either physically or in time. On the analytical side one needs to establish an explicit relationship between scattered field on the one hand and target shape and characteristics as well target illumination on the other hand. Researchers involved with inverse problems know that this is a tough problem. It is usually simplified by making some scalar approximation which may be essential to the formulation

[1] In the case of ATR of aerospace objects of interest here, clutter is minimal since objects are observed against empty sky or space and the only clutter can arise due to antenna side-lobes that see the ground.

of an algorithm (direct and indirect). This usually means sacrifising polarization information. However, the main reason for pursuing non-imaging methods is the technical complications and economic considerations of pursuing the imaging option.

In the second approach, the electromagnetic response of a target is processed with a view to extracting a set of parameters that defines a target uniquely and therefore set it apart from other targets. The success of this method therefore depends not only on the data processing method employed but also on the suitability of the signature parameters chosen. If a particular signature does not change sufficiently for different targets, or if the signature for the same target changes drastically due to some form of distortion (possibly noise) in the data, the demands on the processing method would be too burdensome. The ease or difficulty of choosing appropriate signatures in a given application is also fundamentally related to the complexity of the process (in our case, electromagnetic scattering) that generates data from which the signatures are to be extracted, a theme that we will expand upon in the next paragraph. Also, most of the methods proposed todate in the signature based strategy, have relied on the digital computer for processing data. Therefore the question of choosing a small number of optimum or near optimum parameters to identify a possibly large number of targets has been at the core of this problem.

Recapitulating, the information about targets is conveyed through the complex scattering phenomenon which relates the material and geometric properties of the target and those of the electromagnetic waves illuminating the target to the measured vector fields scattered by the target. Due to the complexity of this relationship, the different techniques developed and applied to the problem to date are often based upon scalar approximations and hence neglect polarization. That polarization information significantly improves classification and detection is borne by evidence from many problem areas, as documented in the NATO Report [5] and other papers (see for example [6]). However, a successful algorithmic approach that incorporates polarization and other information in a comprehensive fashion is improbable due to the complexity of the scattering process. The obvious advantage of applying the neural paradigm is that it takes the alternate route of extracting complex relationships between the target and its echoes from examples that are made available to it. The reader is referred to [3] for a discussion of ATR based upon models of neural networks.

Our approach is based on two sets of concepts. First is *smart sensing* which enables forming signature vectors of the targets so as to facilitate distortion invariant operation as well as to be amenable to training suitable neural networks with robust, learning and generalization ability. In radar the target is usually tracked so it is always located on or very close to the line of sight of the tracking and data acquisition radar which measures the target signature. Thus because of tracking we need not be concerned with the question of location within the field of view as far as distortion invariant recognition is concerned. This takes care of the lateral displacement of the object. Slight displacement of the target from line of sight would change the aspect of the target proportionately. In our approach, independance of target aspect is achieved through learning and generalization by appropriately designed neural networks. Interrogating the target with impulsive plane wave illumination and measuring the far field provides echoes (impulse responses of the target) whose shape is independant of range and this provides range independance. The SNR of the echoes would change with range and this should be handled by robustness of the neural net design.

Second is the utilization of *segmentation*, *bifurcation* and *computing with diverse attractors*, and *multisensory information* to achieve and enhance cognition. The next subsection describes how target representations invariable under target displacement can be acquired under controlled conditions.

Finally one needs to link that data acquisition and learning in a controlled laboratory environment to recognition of actual radar targets in real environment. The philosophy is that libraries of signature vectors are produced for scale models of actual targets of interest. These are used to train suitable neural networks with due attention given to the principle of electromagnetic similitude as applicable to perfectly conducting bodies (see Section 3.2). This principle states that electromagnetic scattering experiments carried out on scale model and the target itself would be equivalent if the frequencies would scaled in the same proportion.

3.1 Target representation

The term RADAR (Radio Detection and Ranging) has come to refer to active electromagnetic remote sensing methods primarily used for detecting a class of natural or manmade objects that respond to electromagnetic waves by scattering them in a manner that depends on the characteristics of the objects as well as the interrogating waves. Targets may be single objects such as aero-space objects or ships or distributed objects such as terrain, vegetation, ocean waves, clouds or rain. Discrete objects are also characterized by their shape and by such intrinsic parameters as their conductivity, permeability and permittivity functions which determine the intimate interaction of the waves with the object. The story of this interaction is told by the scattered wave through changes in its four basic parameters : amplitude, frequency, phase and polarization. Variations of some or all of these basic parameters may be used to construct signatures which help in distinguishing different objects. One example of a signature of the target is the first N prominent resonances of an object illuminated by an impulse [7]. The problem is how to extract the resonances (or poles) from the available data, possibly corrupted by noise. Whereas different methods have been proposed to take into account the possiblity of multiple poles and prior indeterminacy of the actual number of poles that can represent the scattering data [8], the effect of noise on extracting resonances is very serious due to the nature of the scattering phenomenon as explained in [9].

Another example is a library of range profiles of an object collected over some solid aspect angle. The range profile of a target aspect is simply defined as the real part of the inverse Fourier Transform of the band-limited frequency response of the target at that aspect. The reason for choosing the real part is explained in [10]. The method used in generating range-profile information in an anechoic chamber environment using scale models of actual targets is described in some detail in [3]. One strives to produce a library of range-profiles for scale models for targets of interest for a wide range of aspect angles. The number of aspect angles depends on the angular sampling criterion and the solid angle of encounter of the target (the solid angle formed by all possible aspects of the target that can be encountered in a realistic situation). Such libraries of range profiles furnish the data used to train a neural network to recognize the target from a single "look" or signature.

A range profile does not contain the depolarization information about the target. Because of the complexity of the scattering phenomenon, it is not an easy task to "extract" comprehensive signatures that would uniquely belong to an object. Since two independant parameters uniquely represent the polarization state of a wave [11], one can for example choose the inclination angle of the polarization ellipse, ψ, and the ellipticity angle, χ, which can be easily calculated from measured co- and cross-polarized responses. If the measured co-polarized field at a given frequency is $E_{co}e^{i\delta_{co}}$ and the measured cross-polarized field at the same frequency is $E_{cx}e^{i\delta_{cx}}$, where δ_{co} and δ_{cx} are the phase angles of the co-polarized and the cross-polarized fields, respectively, refered to some fixed reference. then the two polarization angles are calculated as follows

$$tan2\psi = \frac{2E_{co}E_{cx}}{E_{co}^2 - E_{cx}^2}cos\delta \tag{1}$$

and

$$sin2\chi = \frac{2E_{co}E_{cx}}{E_{co}^2 + E_{cx}^2}sin\delta \tag{2}$$

where $\delta = \delta_{co} - \delta_{cx}$. The above two parameters, ψ and χ, together with the complex amplitude of the scattered wave, all plotted against the frequency parameter contain complete information about the scattering object for a given aspect. Note that the amplitude information is already present in the range profile and hence addending the frequency variation of both ψ and χ to it would produce a complete signature of a given aspect of the target.

How such signatures are used to recognize targets and the importance of comprehensive signatures, which include polarization information, in enhancing the cognitive ability of a neural processing system is illustrated Section 6.

3.2 The Principle of Electromagnetic Similitude

It is usually not easy to acquire range profile data of desired aspects of an actual airborne target over some solid angle of encounter. It is much easier to obtain range profiles at different desired aspects of scale models of the real targets in an anechoic chamber environment. The question is whether an equivalence can be established between the range profiles of actual targets as opposed to scale models of these targets. This involves consideration of such factors as dimension and frequency scaling, and electromagnetic parameters of the object and the medium and constitutes what is called the problem of electromagnetic similitude [12]. Here we are paying the price of the convenience of having complete control over the range of aspects over which data can be acquired.

Assume that the permittivity and permeability of the material of actual and scale targets (whose dimensions are in the ratio $n : 1$) are the same. Then it can be shown [12] that the conductivity and measurement frequencies of the smaller model be n times that of the larger. The first requirement is very difficult to meet since the conductivities of metals used for real targets and scale models fall within a rather limited range. However, since the conductivities of metals are very high, increasing the conductivity further would hardly effect the fields in the smaller scale model, and in practice one is able to establish similitude by simply using correspondingly higher frequencies for the smaller scale models. Assume that a radar system is designed to operate in the range of frequencies f_1 to f_2 looking at a target of size L. Then to produce data that obeys similitude (with respect to the actual radar data) we need to use a frequency range of nf_1 to nf_2 in an anechoic chamber environment when using a scale model of size L/n.

4 Learning

Biological neural networks can learn to identify concepts, patterns or objects from examples and appropriately generalize from what they learn. By generalization we mean that learning is not rote. A child learns the concept of *dog* from few examples (encounters) and from there on recognizes all varieties of dogs when encountered. Moreover, biological networks are also known to perform amazingly well on information that is incomplete, noisy or distorted in different ways, what we call *sketchy information.* It is also recognized that these networks are especially adept at solving problems of a different nature than those which yield to parametrization and programmed numerical solutions on digital computer. These problems are usually based upon "natural" data, and sometimes termed random problems [13] because of their lack of structure (actually, quite complex or rich structure). Hence they defy an effective concise definition which could be transformed into an algorithm fit for a digital computer.The term "natural" refers to information that stimulates our senses and that of other species and emanates from the respective surroundings. As a clarification, one should note that not all difficult problems (from the computing perspective) are natural for neural networks. For example the problem of decryption (decoding an encrypted message) is hard but generally not natural in the biological learning sense. Another important attribute of many biological organisms is their ability to distinguish between what is familiar and what is novel. In other words, when confronted with a new object or concept, there is awareness about its novelty. This ability is infact crucial to continued learning in human beings and other species. It is also observed that different organisms can perform specific tasks relevant to the needs of the organism. Thus a given neural network is not expected to perform like a general purpose machine. The tasks that a system is expected to perform determine or are determined by the size and architecture of the network. It is difficult to say much in detail about the intricacies of biological computation, but there is substantial evidence of computing with different types of attractors in biological networks, suggesting that they behave like nonlinear dynamical systems, [14] and [15].

The need to propose and study models of how learning occurs is twofold. One is to be able to understand and explain the learning phenomenon in humans and animals. Second, which is most important from the engineering point of view, is to build systems that can learn. Most learning models proposed to date focus on mimicking some of the properties of biological learning and may be adequate for certain applications. For example, most models of associative memory or learning focus only on recognizing a limited number of possibly complex patterns from incomplete and/or distorted inputs. The environment is assumed to be secure or controlled in the sense that these are the only possible patterns that will appear. In statistical pattern recognition and inductive inference the aim is to infer a rule, e.g. some probability distribution, that could explain some given data well, [16] and [17]. In this paradigm, one requires (in the limit) that the *hypothesis* become equal to the actual underlying *target concept* that generated the data. Informally, a target concept can be an actual object or process from which the data originated in the first place, for example letters of English alphabet. The data is then different examples of these letters written by possibly different people. All the different examples form what is known as the *concept space.* Learning is then seen as a process that uses examples of the target concept to produce a hypothesis, an approximation of the concept. For example, a neural network (which is the physical implementation of the hypothesis), suitably trained on examples of alphabet, can classify new examples of the same alphabet by correctly outputting symbols for different letters.

Relatively recently (1984), Valiant [18] has proposed a more general framework to construct a mathematical model of the learning process. The model is variously known as the *distribution-free model* or the model of *probably approximately correct* learning. In this model, a learning algorithm attempts to learn a concept (or target) belonging to some known class of concepts (or targets). The algorithm is assumed to have access to the concept only through positive and negative examples of the concept. For example, all handwritten versions of the number "5" are positive examples of the concept "five". All handwritten versions of numbers (0-9) other than "5" are negative examples of the concept "five". The examples are thought to be generated randomly according to some unknown probability distribution, which may be arbitrary but fixed. Three realistic requirements are placed on the performance of the learning algorithm [19]. First, it is required to identify the unknown concept only approximately (*probably approximately correct*). The more accurate the approximation, the better it is. Second, it should learn in reasonable time, i.e. the learning algorithm should be computationally efficient, in the standard *polynomial time* sense. By polynomial time we mean that the time required to process the data is at most a polynomial function of the amount of input data [20]. Third, the learning algorithm should be general enough to perform well against any probability distribution on the examples[2] (*distribution-free learning*). Regarding the last point, it should be pointed out that not all biological systems are geared to perform equally well on all different problems. Hence the ability to process different types of data (in other words, data with different underlying probability distributions) should generally be linked to the types of tasks that a system is expected to achieve. The term probability distribution on examples in the context of object recognition can be seen as assigning values of the probability measure μ to different objects that belong to the sample space of all possible objects occuring in a particular application. For example, in zip code recognition in the U.S. only arabic numerals can occur with some probability, but the probability of occurence of Chinese characters is negligible.

Some comments are in order at this stage. First, that computational learning theories only guarantee existence of learning automata that are capable of doing certain tasks but are still lacking to a large extent on proposing constructive and adaptive procedures to achieve these goals. Second is the assumption that a certain number of examples (positive and negative) are available to the learning algorithm so that it can learn a certain concept rather well. It is shown by Kearns [19] for example, that certain concepts or representation classes can be learnt by negative-only or by positive-only examples. However certain concepts, for example, polynomially learnable representation classes like kCNF $\vee$ kDNF (i.e. the disjunction of the k conjunctive normal form and the k

[2] A probability space is a set X (of objects or elements), together with a family $\mathbf{A}$ of subsets of X and a function μ, the ***probability distribution*** or ***probabilty measure***, from $\mathbf{A}$ to the unit interval [0,1]. An element A of $\mathbf{A}$ is known as an event, and the value $\mu(A)$ is known as the probability of A [21]. As a simple example, X can contain all 26 letters of the English alphabet, and all possible combinations of letters constitute the family of sets $\mathbf{A}$. To each element A of $\mathbf{A}$ we assign a number $\mu(A)$, which is the probability of its occurence in a normal english text. See [21].

disjunctive normal form) and kCNF $\wedge$ kDNF (i.e. the conjunction of the k conjunctive normal form with the k disjunctive normal form), require both positive and negative examples for polynomial learnability[3] [19]. This problem is also related to the arbitrary but fixed probability distribution condition which sounds very strong indeed. However, the catch is that the same probability distribution that generated the examples used in training is the one that is used to test the system. This may be a reasonable assumption in applications where the system is not exposed to novel stimuli which disturb the probability distribution substantially. An example is the post office environment where a machine is required to sort mail by Zip code, where the number of possible characters is restricted and their frequencies of occurence are within certain limits. We will discuss this problem in greater detail soon.

Third, is the problem of using *a priori* knowledge about the problem. In many practical situations we do have some knowledge of the *target function* f, for example the shape of a certain object. In these cases it would be inefficient to take random examples without taking advantage of what is known about f, for example the object is symmetrical about a certain axis. Hence if one could use such hints in learning from examples, it may considerably help in reducing the hypothesis space from which functions may be chosen to approximate the unknown concept, or reduce the number of steps or examples needed to learn the concept [22]. For example, the range profiles or signature vectors used to represent radar targets in this work vary gradually, at a rate depending on the complexity of the target, as the aspect of the target is changed. Hence, one can use this information in selecting range profiles to train a system to recognize radar targets from their range profiles. In the absence of this information, one would choose the angular resolution criterion to calculate the number of range profiles needed to characterize a radar target over a given angular window. The number of range profiles required in this case can be very large depending on the bandwidth of illuminating radiation. *A priori* knowledge of the "angular correlation" of range profiles

[3] The *conjunctive normal form* or CNF is a conjuction of monomials. A monomial is itself a conjunction of literals. A literal indicates a variable (feature) or its negation. Similarly, the *disjunctive normal form* or DNF is a disjunction of monomials. Literals define the simplest (atomic) concepts. Monomials are conjunctions of literals and therefore define concepts which are more complex. The DNF and CNF define even more complex sets. If we want to describe concepts of even greater complexity, we can form conjunctions (or disjunctions) of DNFs (or CNFs). The kCNF and kDNF are obtained by restricting the number of literals in each monomial that makes up these functions to k.

helps determine how range profiles could be selected to achieve good generalization.

4.1 Complexity Theory and Efficient Learning: Brief Background

The difficulty with which different concepts can be learned from examples forms the subject of complexity of learning. Complexity theory deals with the relationship between the number of examples needed by the algorithm or machine to learn a concept to be able to perform valid generalization and the time required by the algorithm to learn that concept. The issue is whether the algorithm can achieve its goal "efficiently", i.e. in reasonable time (in polynomial time sense).

The problem can be formalized in the following way. We have closely followed the treatment in [21]. Suppose H is a hypothesis space defined on the example space X, from which a set $\mathbf{x}$ of N examples of a target concept is available. To make the discussion complete we will interweave with a simple example. The concept can be a black and white picture which partitions the two dimensional space into black (or positive) and white (or negative) regions. N examples of this concept (picture) consist of coordinates of N points in the picture plane. The N examples could be N shots of the same scene taken at different times of the day. The example space can be seen as a manifestation of the target concept f which is to be approximated by some $h \in H$ from N examples in X. For example, h may be the partition achieved by a certain feedforward network of the type we discuss below. H is then the class of all feedforward networks which implement different partitions. If H is restricted to all straight lines (or networks implementing 1-D hyperplanes), the types of partitions and therefore the concepts that can be learned are restricted to those that are linearly separable. H can classify the N examples (as positive or negative, i.e., binary classification) in at most 2^N ways. $\mathbf{x}$ is said to be shattered by H if this maximum possible value is attained by H. For example, the hypothesis space consisting of straight lines (1-D hyperplanes) can shatter three non-collinear points in a plane, i.e. the three points can be partitioned in all eight possible ways. If the set $\mathbf{x}$ contains examples which are not all distinct (therefore, not separable by any surface), then it cannot be shattered by any H. More formally, when the examples are distinct, $\mathbf{x}$ is shattered by H iff for any subset S of the examples, there is a hypothesis h in H such that for $1 \leq i \leq N$,

$$h(x_i) = 1 \Longleftrightarrow x_i \in S \tag{3}$$

S is then the set of positive examples of $\mathbf{x}$, the remaining being negative examples.

Let us assume that a possibly unknown function $h(f) \in H$ approximates the concept f well on all examples in X, i.e. $h(f)$ is the best possible approximation to f. Since $h(f)$ is not known, we can consider a function $h_N(f)$ as an approximation for $h(f)$ and therefore for f, and expect it approach $h(f)$ as N becomes very large. However the function $h_N(f)$ may be biased by the N examples used to obtain it. We want to know how bad is the estimate in the worst case. The key result is a bound given by Vapnik and Chervonenkis [23]

$$Pr(max_f|h_N(f) - h(f)| > \epsilon) \leq 4g(2N)e^{-\epsilon^2 N/8} \tag{4}$$

Unless the function $g(N)$ grows exponentially, the right side will approach zero as N increases.[4] The growth function $g(N)$ is the maximum number of different binary functions on the set of examples $x_1, \ldots, x_N$. It is either identically equal to 2^N for all N (VC-D is infinite since it keeps increasing with N) or else is bounded above by $N^d + 1$ for a constant d (VC-D = d is finite). The VC-D (VC dimension) of a hypothesis may be defined as the maximum number of samples N_{max} that are shattered by H. Finite VC-D implies a polynomial $g(N)$ and guarantees generalization. In this case, as the number of examples increases beyond VC-D, the concept is better learnt (number of valid hypothesis from H decreases) and generalization improves. It would be helpful to give an example at this point. We want to find the VC dimension of the hypothesis space H consisting of all 1-D hyperplanes that may be used to partition a plane. As already stated in the foregoing paragraph, H shatters any three non-collinear points in a plane, i.e. it can partition them in any of the eight possible ways. Hence, H has a VC dimension of at least three. It can be easily shown that any four points lying in a plane are not shattered by H. Therefore VC-D(H) = 3. It has been shown by Baum and Haussler [24] that the VC-D of a feedforward network with one hidden layer is proportional to the number of its nodes and adaptable weights

[4] 2N is used in the argument of $g(.)$ because in deriving the equation (2), two samples of length N each are used. This is needed to see whether the maximum difference between the relative frequencies of a certain event in these two samples uniformly converges to some value as the number of examples N in each sample is increased.

Within Valiant's framework, one wants to learn from examples of a Boolean function $f \in F$. The choice of the hypothesis (or representation) class H is crucial in the learnability of F [25]. A class H of representations is defined as a *p-time representation* if for all x and for all $h \in H$, $h(x)$ may be computed in time polynomial in n (dimension of the feature vector x) and the size of h. Baum proves that:

> For any class of concepts F and any p-time representation H, if F is learnable by H, then F is learnable by feedforward neural nets.

However, there are functions that are not learnable by neural networks. For example, Goldreich et al. [26] have constructed classes of *poly-random* functions not learnable by any representation (or hypothesis) and hence, in particular, not learnable by feedforward nets. Goldreich calls a function poly-random if any polynomial-time algorithm, given values of the function at arguments of its choice, cannot distinguish a computation during which it receives the true values of the function from a computation during which it receives the outcome of independant coin flips. Also Kearns and Valiant (1988) have shown under cryptographic hypothesis that the class of feedforward nets, even when restricted to be logarithmically deep (i.e. if the size of the input is n, then the number of layers is of the order of *log*n), with each node connected to a constant number of others, are still not learnable by any p-time representation. It is evident that human learning in natural world as well as a lot of practical problems are not concerned with solving the general decryption problem. The number of concepts that are learnable from examples (n-dimensional) in polynomial time are an exponentially small subset of possible concepts. According to this assumption, since people are capable of learning in the real world, there must exist a small set of concepts that are both rapidly learnable and adequate for accurately describing the world.

Independantly, Hornik et al. [27] have shown that standard multilayer feedforward networks with as few as one hidden layer using arbitrary squashing functions are capable of approximating any Borel measurable function[5] from one finite dimensional space to

[5] Let S_x and S_y be a system of subsets of any two sets X and Y, respectively. Then an abstract function $f(x)$ defined on X and taking values in Y is said to be (S_x, S_y)-*measureable* if $A \in S_y$ implies $f^{-1}(A) \in S_x$. If S_x and S_y are chosen to be system of all *Borel* sets, the the function defined above is called a *Borel-measureable* function. Put simply, B is a *Borel* set if B can be obtained by a countable number of operations on some given sets, starting from open sets and each operation consisting of taking unions, intersections, or complements.

another to any desired degree of accuracy, provided sufficiently many hidden units are available. This result thus establishes the class of concepts that can be learned by multilayered feedforward networks.

4.2 Learning and Working in Different Environments

In many applications, a machine is required to work in the same environment in which it was trained in. A robot working in an auto factory, a handwritten Zip Code recognition machine in the post office, and most classification tasks are examples of tasks confined in secure or controlled environments.

However, in other important applications, a machine is required to work in environments other than that it was trained in. This may be desireable when one is interested in identifying a small number of objects among a very large number of possible objects, or when training in the actual environment is practically not possible, as is the case with radar target identification. Another important consideration is that of the capacity of finite sized networks to learn.

The original PAC learning framework, as proposed by Valiant in 1984, assumes that the training and working environments are identical. In a modified PAC framework, Shvaytser [28] considers cases when the two environments can be different. For binary classification of examples, one can characterize an environment e by the probability distribution functions D_e^+ and D_e^- of the positive and negative examples[6]. The training environment is denoted by $e = 0$. $e = i$ where $i = 1, 2, \ldots, E$ represents other environments that could be encountered by the machine trained in environment $e = 0$. There are three possible cases that can occur in practice (Shvaytser [28])

1. The environment is unchanged during training and working (testing), i.e. $e = 0$ all along. A simple example of this is when a network trained to classify only two different objects or patterns is expected to encounter these two objects, to the exclusion of all other objects. Hence it operates in a controlled environment.

2. The working environment $e = i$ is completely unknown during the training, which is done in environment $e = 0$. In this case a common strategy is to take D_0^+ and

[6] Examples are instances of some concept(s), e.g. a tree. D_e^+ could be the probability distribution function of instances of trees in the environment, and D_e^- the pdf of instances of other objects that are not trees

D_0^- as uniform distributions.

3. D_1^+ is known and is used for D_0^+, but D_1^- is unknown. An example of this can be when a network is trained to recognize a letter "A" in different environments. The negative examples can all be other alphabets and/or some other random patterns, Chinese letters, etc. The two subcases that arise in this instant are: (a) negative examples are not used at all during training, and (b) D_0^- is assumed to be a uniform distribution over the negative examples.

In Valiant's framework, it is shown that in case (1), a polynomial number of examples is always sufficient for reliable training [29]. Also in this framework, it can be shown that reliable training is impossible in case 2. Shvaytser shows that it is impossible to reliably train feedforward networks to handle both subcases 3(a) and 3(b). We will illustrate this in the next section using radar target identification as an example. Specifically we will show why it is not possible to achieve sufficient cognition by only using feedforward networks. Also we propose and describe a novel composite network that has the ability to solve this problem.

In this and other similar applications, the reasons for using this approach can be enumerated thus: Given a set of all possible objects that can occur within the environment of a neural network or a cognitive system, one cannot practically think about learning all possible objects from their different manifestations. There are some important reasons for this. First, the amount of information available to the network to learn may be so great that the size of the network required to learn the environment in detail becomes horrendous and learning in reasonable time becomes improbable. Second, the objects of interest that are to be classified may be relatively small, and it would be inefficient to learn in detail, information about all other objects (i.e. negative instances) that is not directly useful. Third, information about all possible objects or concepts possible within the environment is rarely available in practical situations. Hence, although positive examples of the target concept are available, negative examples are either too numerous or expensive to come by. Also a reasonable number of negative examples brings us to the first point: namely the size of the network and the time required to learn (i.e. train it). Fourth, the network could be trained in a controlled environment and then required to operate in a different environment. This may be seen as changing the probability ditribution of the sample space of the examples used to test the network as compared to the probability distribution of the sample space used for training.

One can think of building a network that learns only a subset of the set of all possible objects to the exclusion of all the other objects in the set, i.e. to build a network that can distinguish between familiar objects belonging to its learning set and novel objects belonging to the set of all objects the net has not or could not be taught. We call this capability cognition. The inability of a network to distinguish independantly, i.e. on its own, between familiar and novel objects or its lack of cognition is one of the major outstanding issues in pattern recognition that is not widely appreciated. The second major issue is how to achieve distortion invariant recognition which is often referred to as displacement, rotation, scale, and SNR (signal-to-noise ratio) independant recognition. Both issues assume crucial importance in remote sensing and in autonomous systems that are meant to operate in a complex uncontrolled environment, and have consistently resisted attempts at their solution for a long time. The radar recognition problem, which presents itself as a marvellous example in illustrating these issues is used in the next two sections to highlight these issues.

5 Learning to solve the ATR problem

In section 3 we argued why applying the neural paradigm to the problem of ATR held promise because one can learn complex relationships through examples when it is difficult or impossible to arrive at them analytically (and therefore algorithmically). However, one has to consider issues of a rather different nature that emerge as a result of taking this route. For example in ATR problem, it is not practically possible to teach the system with all possible targets that can happen in its environment, both because of the limited capacity of the system and difficulty in acquiring data about all possible targets. Even the number of targets required to be classified may be large enough to be efficiently learned by a single network. This confronts us with the question of whether imparting cognition to the network can resolve these issues. In addition, targets of interest (e.g. certain class of airplanes) produce generally quite similar signatures, specially from certain aspects. Hence the recognition task requires making fine distinctions between similar echoes. These concerns are enhanced by the presence of noise and the signal level, which may vary depending on the distance of the target from the radar. Finally, there is the practical need to learn and identify targets in reasonable time so that information does not lose its value.

A neural network designed to solve the radar problem must therefore fulfill certain requirements. First, it should exhibit good generalization by performing well on new examples of the known targets and at the same time be able to discriminate against examples belonging to novel targets, i.e. have cognitive ability. The nature of the application also requires robust operation in the face of external and internal noise and imperfections. Also, the network should be able to perform its task in real time.

It is logical to start by examining existing neural net techniques to see how they relate to these characteristics. For example forming simple heteroassociations of target echoes with target labels (see for example, [30] and [3]) does not provide an answer to the problem of cognition. Not only are different target labels evoked by some echoes belonging to other targets but also by spurious inputs. As another different example, Ans's self-organizing network [31] requires long training times and also lacks cognition, i.e. it is unable to distinguish between familiar and unfamiliar targets. As a more interesting example of this difficulty in keeping these crucial properties together we discuss our experience with a high threshold version of the feedforward network, a possible candidate for providing cognition. This network is a simple feedforward network trained by error back-propagation in which the internal threshold of neurons is used to control the response region of neurons.

In this high threshold network we found that the generalization and cognitive performance of a feedforward network can be tuned by varying the internal thresholds of neurons. As a simple example, the network can be taught to associate selected range profiles of two targets with labels assigned to the two targets. To test the generalization of the network one tests it with novel range profiles of the familiar targets. To test its cognition, one can test it on range profiles from some novel targets as well as spurious inputs or signals. We found that as the threshold rises, the ability of the network to distinguish between familiar and novel targets increases at the expense of its generalization ability. For low to moderate values of thresholds, the generalization ability is quite good. However, when the threshold is quite high, the network becomes only a memorizer of training examples and can be seen as an example of rote learning. Also, with increasing levels of noise, the network rapidly loses its generalization performance and is no longer able to recognize the known targets most of the time. Generalization is important because as it will be seen later it is the mechanism with which recognition of the object or target from single echos independant of their aspect (aspect or rotation invariant recognition) is achieved.

Some results on the high threshold networks will be helpful to explain its performance better. The architecture of the feedforward network is as follows. The number of neurons at the input is fixed by the number of data points in the range profile at 128. The number of neurons in the hidden layer is chosen to be 24, which seems to be a good choice for this data set. The number of ouput neurons is chosen to be 32. The network is trained on 25 percent of the available range profiles from the B52 and the Space Shuttle. During testing all the range profiles from the three test objects, the B52, B747 and the Space Shuttle are used. When a zero threshold is used in training and testing, the network classifies the known objects correctly in all cases, but misclassifies the B747 as either a B52 or a Space Shuttle from 97 percent of the views. When the threshold (during training and testing) is raised to $\theta_i = 1$ the misclassification rate on the novel target, i.e. the B747, goes down from 97 percent to 78 percent. The remaining 22 percent of the range profiles from the B747 result in sparse activity at the output of the network, which can be taken as an indication of discrimination against novel targets by the network. The performance on known targets is almost unaffected. The behaviour of the network as a function of progressively increasing the neural threshold is shown in Table 1 and Figure 1. Beyond a certain value, raising the threshold further only marginally decreases misclassification of unknown targets at the expense of deterioration of performance on known targets.

Also the network is trained rather rapidly by increasing the threshold rather gradually. Training the network at higher thresholds directly either requires longer times or the network does not converge. Therefore we trained the higher threshold networks in stages to facilitate rapid learning. For example, if a network is to be trained to operate with a neural threshold of $\theta_i = 3.0$, it is trained with a neural threshold of $\theta_i = 1$ until the mean-squared error between the actual and the desired outputs has dropped below a given value. In the next stage, the threshold is raised to $\theta_i = 2.0$, for example, and the training is continued until the error between the actual and desired outputs has again dropped below the given value. Finally the threshold is raised to $\theta_i = 3.0$ and training is continued until the error between actual and desired outputs again drops below the given value.

Using different sets of targets to train the network influences the performance of the

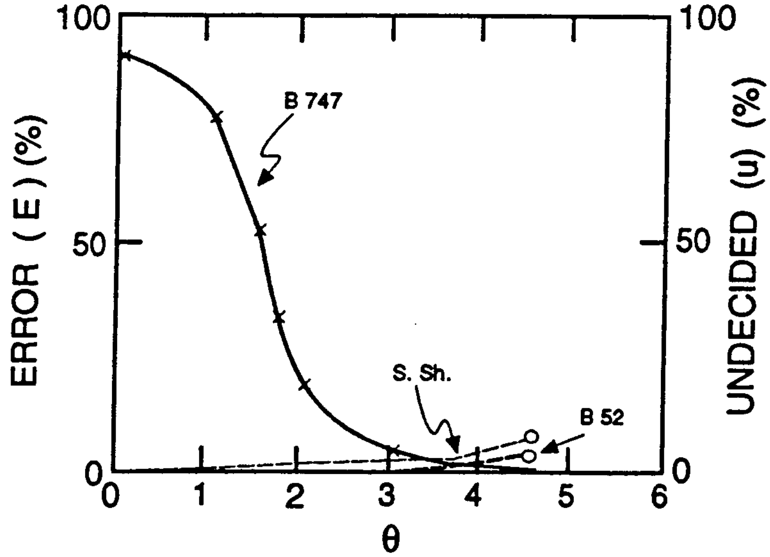

Figure 1: The effect of using different neural thresholds on the performance of feedforward networks. The network is trained on 25 percent data from the B52 and Space Shuttle scale models and tested on all data from these two models as well as a novel target (B747).

	Error (E) or Undecided (U)		
θ_i	B52	S.Sh.	B747
0	0	1 (U)	91 (E)
1.0	0	1 (U)	78 (E)
2.0	0 (U)	2 (U)	19 (E)
3.0	0 (U)	2 (U)	4 (E)
4.0	2 (U)	5 (U)	2 (E)
4.5	3 (U)	7 (U)	1 (E)

Table 1: The effect of using different neural thresholds on the performance of feedforward networks. The network is trained on 25 percent data from the B52 and Space Shuttle scale models and tested on all data from these two models as well as a novel target (B747).

	Error (E) or Undecided (U)		
θ_i	B747	S.Sh.	B52
0	0	0	98 (E)
1.0	0	0	94 (E)
2.0	0	0	63 (E)
3.0	0	0	40 (E)
4.0	0	0	17 (E)
5.0	0	0	12 (E)
6.0	0	1 (U)	12 (E)

Table 2: The effect of using different training targets on the performance of feedforward networks using various internal neural thresholds. This network is trained on 25 percent data from the B747 and Space Shuttle scale models and tested on all data from these two models as well as a novel target (B52).

high threshold network rather strongly. For example when we used the B747 and the Space Shuttle as the known targets and the B52 as the unknown target, with the network parameters same as those for the net described in detail above, the misclassification rate for the unknown target (the B52, in this case) was as high as 40 percent at $\theta_i = 3.0$, down from 98 percent at $\theta_i = 0$. The cognitive performance of the network in this case as a function of the internal neuron threshold is tabulated in Table 2 and plotted in Figure 2 for this case.

When the B52 and the B747 are used as known targets and the Space Shuttle as the unknown target, the misclassification rate on the Space Shuttle drops from 63 percent at $\theta_i = 0$ to only 3 percent at $\theta_i = 3$. This behaviour is tabulated in Table 3 and plotted in Figure 3.

The asymmetrical behaviour of the network *vis-a-vis* the training set is not the only problem with high threshold networks. Other critical properties such as the dynamic range and robustness against noise are far from satisfactory. Since the nets trained and operated at high neuron thresholds form tighter phase spaces only a small amount of Gaussian noise is tolerated before the network fails to recognize a given target either by classifying it as one of the other targets or by sparse activity at the output layer of the feedforward network as a signal of its inability to make a decision. Even a signal to noise ratio of 10dB is usually sufficient to cause such a failure. Note that the signal to

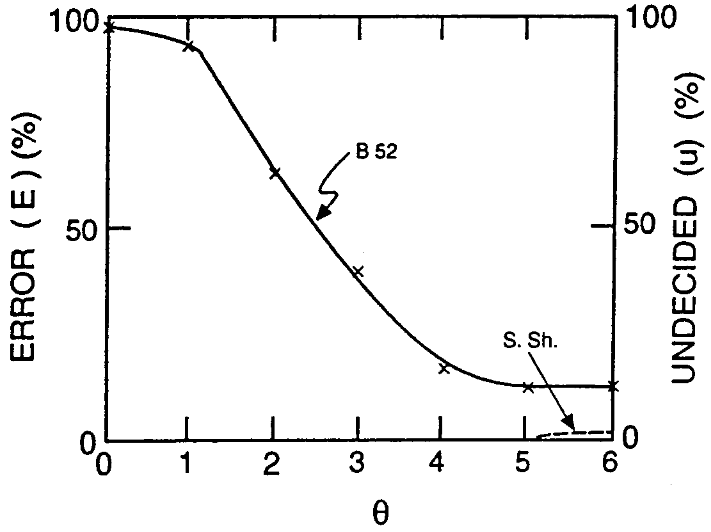

Figure 2: The effect of using different training targets on the performance of feedforward networks using various internal neural thresholds. This network is trained on 25 percent data from the B747 and Space Shuttle scale models and tested on all data from these two models as well as a novel target (B52).

	Error (E) or Undecided (U)		
$\theta_{train}=\theta_{test}$	B52	B747	Space Shuttle
0	0	0	63 (E)
1.0	4 (U)	0	40 (E)
2.0	4 (U)	0	10 (E)
3.0	2 (U)	0	3 (E)
4.0	1 (U)	1 (U)	2 (E)
5.0	2 (U)	1 (U)	1 (E)
6.0	0	1 (U)	1 (E)

Table 3: The effect of using different training targets on the performance of feedforward networks using various internal neural thresholds. This network is trained on 25 percent data from the B52 and B747 scale models and tested on all data from these two models as well as a novel target (Space Shuttle).

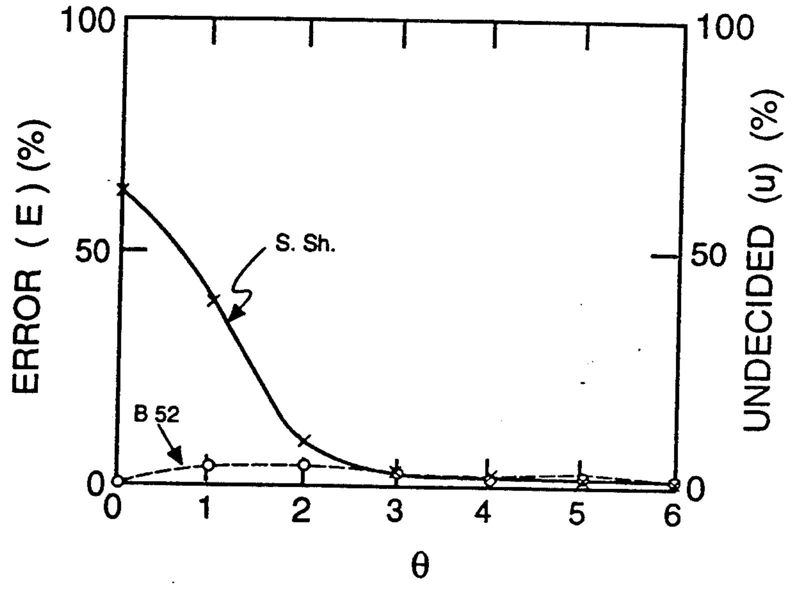

Figure 3: The effect of using different training targets on the performance of feedforward networks using various internal neural thresholds. This network is trained on 25 percent data from the B52 and B747 scale models and tested on all data from these two models as well as a novel target (Space Shuttle).

noise ratio of the original range profiles collected in the experimental facility is about 15 to 20 dB. Also the dynamic range decreases as the threshold is raised. Hence raising the threshold makes the network more and more inflexible to changes in the signal level.

A deep look at how the network operates tells us why it cannot be trained reliably with only positive examples (i.e. examples of the objects to be recognized), when it is expected to perform in a different environment. The knowledge of the network is only based upon the patterns from the targets used to train it. When some pattern is presented to the input of a network, a neuron in the following layer sees a weighted sum of the pattern inputs (depending on the relevant inter-layer weights). The threshold θ of the neurons serves as a gauge [32]. When the weighted sum is greater than θ, the particular neuron identifies the pattern as similar to some examplar pattern which produces the greatest value of the weighted sum. For some high value of threshold $\theta = \theta_H$, only one pattern will be classified as familiar. This corresponds to rote learning. For some lower value $\theta = \theta_L$, all possible patterns will be classified as familiar. In between, some patterns will be classified as familiar and some as non-familiar. The response region of the neuron can be seen as the mechanism that provides approapriate generalization and cognition in feedforward networks. The problem is how to choose an appropariate value of θ using only positive examples in the radar target recognition case. One can choose a reasonable threshold by observing performance on unknown targets, but this violates our condition that information about other than training targets is not available. This problem is not solvable using high threshold feedforward networks in the radar case, because some unknown targets have some echoes which are more similar to some of the known target's echoes than other echoes of that known target. Also, making the response region tight has the effect of making the network fragile to noise. The question is whether one can come up with a different scheme that would introduce cognition but not at the expense of sacrificing other desireable characteristics such as generalization and robustness.

6 Blueprint for Cognition

As we now explain, the nonlinear dynamical systems approach to computing offers an interesting opening into the problem. The biological plausibilty of such an approach is evidenced by the fact the higher level cortical circuits are nonlinear and exhibit rich feedback [33]. The behavior of such circuits can be macroscopically described in terms of the types of attractors they can exhibit, namely, *point*, *periodic* and *chaotic attractors*. There is evidence that a plausible mechanism for achieving cognition lies in the ability to bifurcate between different attractors depending on the input to the network. For example bifurcation between periodic and chaotic attractors in the rabbit olfactory bulb provides a mechanism for differentiation between familiar and novel odors as shown by Skarda and Freeman [14] and Baird [15]. How this is actually done is still difficult to comprehend, partly because of our limited understanding of chaos and chaotic attractors. Because it is easier to consider bifurcation between point and periodic attractors, we will explore the cognitive potential that can be tapped by bifurcating between these two types of attractors.

Here is where periodic attractor networks enter the picture, as agents for providing cognition. The periodic attractor network (PAN) is briefly described in the appendix. Here we summarize some of the important features of a PAN.

- The PAN is a fully connected feedback network in which highly correlated vectors can be stored in one or more non-intersecting open or closed trajectories in the phase space of the network.

- A relatively large number of vectors (of the order of N) can be stored on prescribed trajectories.

- These trajectories can be formed with a high degree of isolation in the sense that if the network is initiated by a stored vector or one close to it in the Hamming sense it triggers the periodic attractor, otherwise it goes to a point attractor. This is the mechanism for providing cognition.

- We have found that robustness of these networks to imperfections in weights is reasonable, in that they can withstand 6-10 percent weight imperfections without appreciable loss in isolation properties. This is an important when hardware implementations of nets are considered.

- The PAN are however intolerant to element failure but in practical nets this can be remedied by neuron redundancy.

- The PAN requires synchronous update and the implications of this are also discussed in the section on conclusions and discussion.

6.1 Integrating Diverse Attractors.

The problem is how one can combine the desireable properties of the feedforward networks with those of the PANs. At this moment it seems fruitful to see what hints neurobiology can provide about a possible mechanism for cognition in the brain. The current view (which we present in very simplified terms) can be condensed in the following way. The different modalities of information that impress on our various sensors end up as separate cortical maps on our cortex. As an example, the different sections in the somatosensory cortex can be related to associated areas on the body surface. These cortical maps are presumably integrated by intercortical circuits which connect different areas. Unfortunately this mechanism seems to be quite complex and it is not known exactly how the integration takes place. One can thus hypothesize a preliminary blueprint for cognition. Feedforward networks process segments of information and map them onto the cortical surface, and other networks which use feedback somehow bind these cortical features together to provide a mechanism for cognition.

How the integration (binding) takes place is quite difficult to answer. For example, Eckhorn *et al.* [34] based upon their discovery of feature linking of cell assemblies in cat primary visual cortex by mutual synchronization, suggest a neural model to explain this phenomenon. Their model net consists of two layers of neurons coupled by feedforward connections as well as lateral and feedback connections. The idea is that temporal correlations may be the means of achieving binding.

With this we may venture to propose this simple (engineering) model for achieving distortion-invariant recognition of radar targets. The idea is to process target signatures in segments with feature forming modules and then bind the features formed by these segments depending on the compatibility (consistency) of the features. The sub-spatial features may be formed by feedforward trainable networks which process segments of radar signatures. The composite features or labels formed are then processed by a periodic attractor network which either binds these features by a periodic attractor or, if they are not compatible, bifurcates to a point attractor. Note that in this case, the thresholds of the feature-forming networks are taken to be zero, and hence the problem of determining thresholds appropriate for generalization does not occur. This is replaced by the easier problem of choosing the threshold in the feedback PAN in order to ensure sufficient isolation of the periodic trajectories from the rest of the network phase space. In such composite networks, the feedforward feature forming modules lack cognition but furnish robust learning and generalization, while the PAN which lacks generalization furnishes the mechanism for cognition through its bifurcating ability.

6.2 Performance of Simple Composite Networks

As a simple test of the performance of such an architecture we use range profile segments as inputs to a composite network. The case of one segment corresponds to using one multi-layered feed-forward network, and therefore one does not have the mechanism of comparing different sub-spatial features of the echo for compatibility. With two segments comparison of different spatial features becomes possible through the PAN. The key point is that as the number of segments increases, the chance that an unknown target responds on all segments in exactly the same fashion as one of the known targets decreases rapidly and the PAN makes use of this to provide cognition.

Some results will help to make the discussion clear. The architecture of the composite network is scematically shown in Figure 4. To test the potential of such a scheme, we used simple perceptron networks at the front end to process segments of data. The target data used to evaluate our cognitive network is the same as that used to evaluate the high threshold networks described in the previous section.

We first used the B52 and the B747 as the training targets and the Space Shuttle as the novel target. When whole range profiles were used in training one fully connected network, the known targets are recognized with almost hundred percent certainty. However the performance on the unknown target (Space Shuttle) was undesireable since 80 percent of the time it was classified erroneously as one of the known targets and only 20 percent of the time did the net indicate its ignorance by going to a ground state (i.e. all neurons in the output layer are in the low state represented by 0). We then divided each range profile into two equal segments and each segment was used with two separate feedforward networks as shown in Figure 5. During testing, if both segments give the correct answer the target is recognized unambiguously since one of the periodic attractors is triggered. Otherwise, the network indicates its reservation about making a decision by going to a point atrractor. With two feedforward nets at the front end, the known targets were again recognized almost perfectly. The performance on the unknown target improved greatly since it was misclassified as one of the other two 29 percent of the time by both nets simultaneously. In the remaining 71 percent of the cases, the

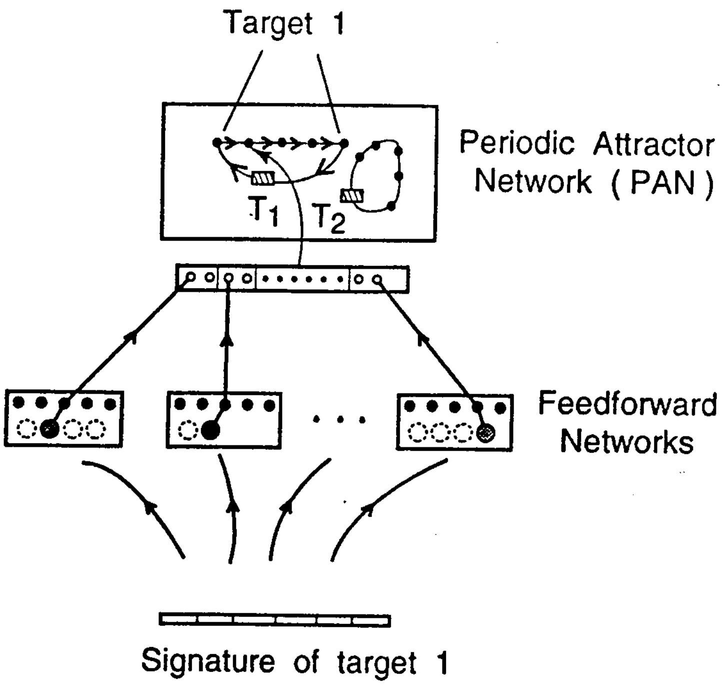

Figure 4: A schematic of a composite feedforward networks and periodic attractor network (PAN) that can be used to achieve controlled generalization. Each feedforward network outputs a certain label when initiated by an example from a certain region of object space. All feedforward networks cover the total desired space of examples from the object. The periodic attractor network binds the response labels of an object with its master label. Two master labels T_1 or T_2 for two different objects are shown.

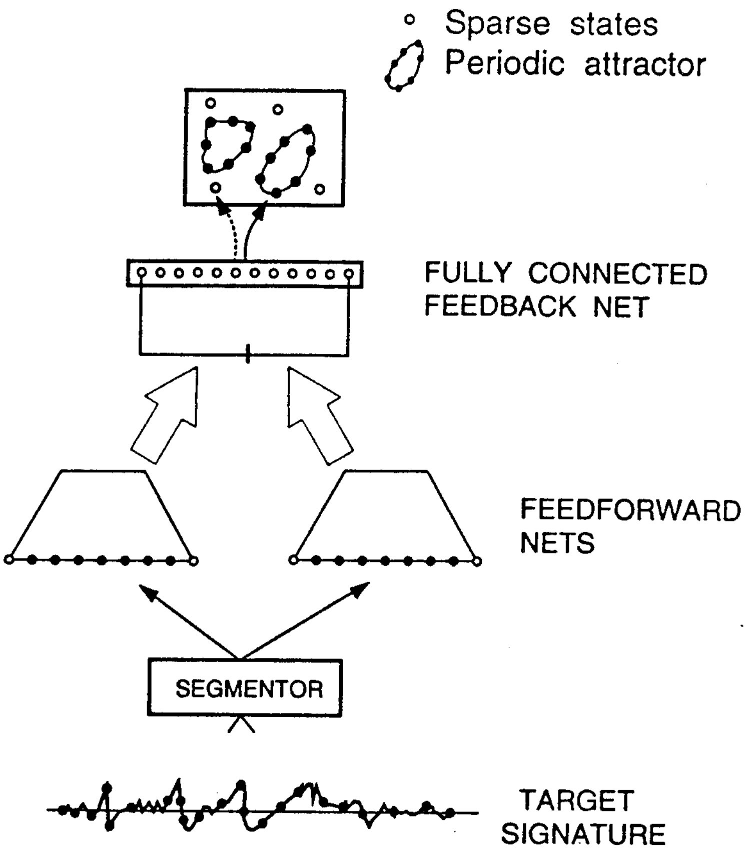

Figure 5: The simplest composite network. The range profile is divided into two equal segments and fed to two identical single layer feedforward networks. The outputs of both these networks are concatenated and used to trigger the periodic attractor network

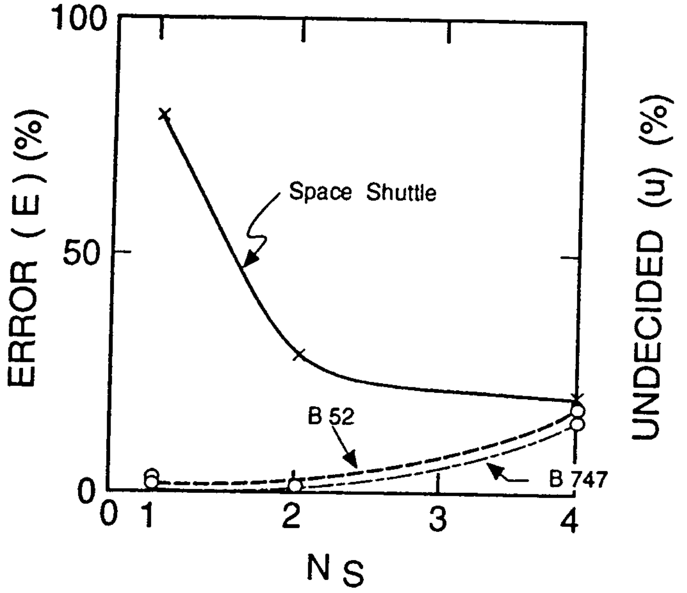

Figure 6: Cognition of the composite network as a function of the number of segments of range profiles (or modules of single layer feedforward networks). The network is trained on 25 percent of the data from the B52 and B747 scale models and tested on all data from these two models as well as a novel target (Space Shuttle).

networks indicated their undecidedness by outputting contradictory or unknown labels. Using four similar networks on four equal segments of a range profile further decreased the rate of incorrect classification to 20 percent, in which case all four nets misclassified the Space Shuttle as a B52. However the performance on range profiles from known targets also deteriorated since about 17 percent of the range profiles from both the B52 and B747 triggered ambiguous responses since one out of four networks misclassified the target or output an unknown label. With 8 equal segments used (each 16 data points) to train 8 networks, some of the networks did not converge. This might be used to indicate the minimum length of a segment required for containment of relavent target features. The dynamic range and noise robustness of the segmented network are still quite good although with decreasing segment size the effect of noise becomes more pronounced. A summary of simulation results is plotted in Figure 6. It is seen that as the number of segments, N_s, increase, the network discriminates against the unknown target better. However, its ability to recognize the known targets deteriorates to some extent.

We also tested the network with different training target sets to analyse its asymmetry with respect to known and unknown target sets. When the B52 and Space Shuttle are used to train the single layered nets we observed that the net does not converge when four segments are used, i.e. not all of these segments are now linearly separable.

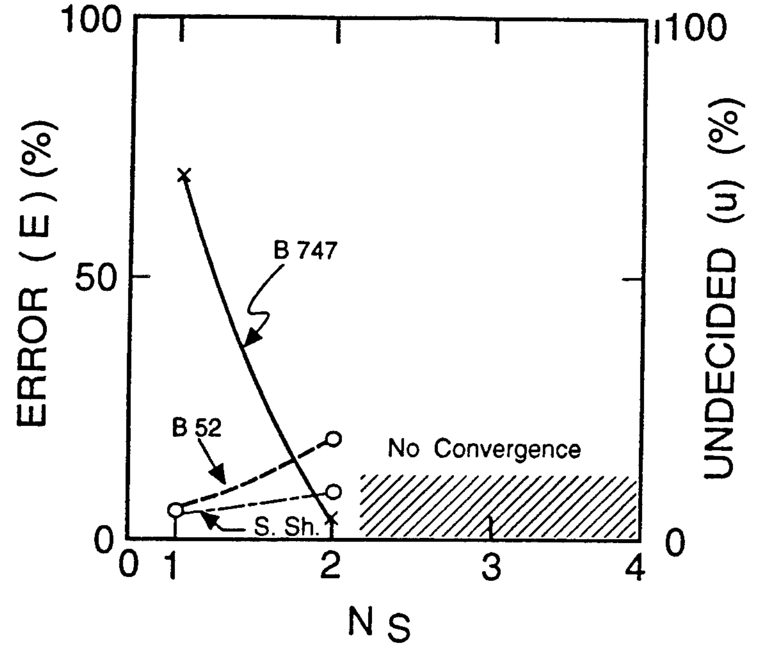

Figure 7: Cognition of the composite network as a function of the number of segments of range profiles (or modules of single layer feedforward networks). The network is trained on 25 percent of the data from the B52 and Space Shuttle scale models and tested on all data from these two models as well as a novel target (the B747).

However even with two segments, the misclassification of the unknown target (the B747) is only 4 percent, down from 70 percent when only one network is used. However the undecidibility rate with two segments is rather high at 20 percent for the B52 and 11 percent for the Shuttle. These results are plotted in Figure 7.

One may ask if one needs to know in advance or can determine the number of segments needed to achieve maximum cognition. The advantage of our cognitive scheme is that the extent of segmentation possible is determined by increasing the number of segments, until learning is not possible, i.e. the feedforward networks cannot extract features from segments smaller than a certain length. One way to do this is to start by training networks on a small segment containing first n_s points of the echoes of known targets, and progressively increasing the number of points, until the segment length becomes large enough to contain relavant features, and hence can be learnt. The length of the second segment can be similarly determined by using that portion of the echoes that were not used in the first segment. The extent of segmentability of the echoes of the known radar targets will depend on the structural complexity as well as the similarity of these targets. This process is a way to achieve maximum differentiabilty (cognition) between different possible targets (known and unknown) based upon information from a finite number of known targets. This differentiabilty or cognition can be increased by using additional information about the known targets, which would enable one to generate more segments with additional features. Hence the chance that an unknown target matches a known target on all these segments is further reduced. In the limit when complete information is available about known targets, one can say that one can distinguish them from all other different unknown targets, even if information about these unknown targets is not available during training. As a corollary, the lesser the similarity between known and unknown targets, the lesser is the information required about the known targets to achieve this goal.

6.3 The Need for Multisensory Information

Consider a simple example that illustrates the effect of the amount of information made available to the network on its ability to differentiate between objects that are similar, i.e. have some similar characteristics. We are required to differentiate between different shapes of different colors, say red, blue and green balls, cubes and pyramids. Using a black and white camera (i.e. color information is not available), we can identify balls from among balls, cubes and pyramids by training someone on balls only. It is obvious that with only black and white information differentiating between balls of different colors is not possible. On the other hand, if one has only a device to measure the color (wavelength of radiation) of the objects, then one cannot differentiate between different shapes but can recognize a particular color from other colors if that color is among the colors used to train the network. In order to recognize a particular shape of a particular color, one needs to use both the black and white camera that provides shape information and the color measuring device.

The above discussion and example also illustrates the role of multisensory information in reducing ambiguity between similar targets as well as imparting greater cognition to the composite network against unknown targets. We observed in the case of composite networks trained on the range profiles of some targets that their ability to discriminate against a novel target increases as the number of segments were increased provided the segment lengths are fixed, i.e. the amount of data within each segment is fixed. In the results cited, the novel target could be discriminated against from 80 to 96 percent of the aspects, depending on which targets were used in training the network. This is the maximum performance achievable using only range profile data. To increase the cognition ability of the network would require more information on the known targets, which would help by providing additional segments to make finer comparisons of the targets possible. Another important reason for using additional information is to improve the noise immunity of the system through larger signature segments.

We have conjectured that using multisensory information should greatly improve the cognition of the radar target recognition system. To get a general idea of the type of behaviour expected, we concatenated uncorrelated range profiles to simulate multisensory data. The composite signal formed by such concatenation was constructed as follows. For a given target the available range profiles are divided into two equal groups. In our case, the 50 range profiles from 0 to 10 degrees from headon towards broadside constitute the first group, and the 50 range profiles over the adjacent 10 degree angle form the second group. The n-th range profile from the first group is then concatenated with the n-th range profile of the second group to form the n-th composite signal. Hence from the 100 original 128 point range profiles we form 50 composite profiles, each with 256 discrete samples. The halves of the composite range profiles are uncorrelated because of the angular separation of the range profiles from which they were formed. Hence these composite range profiles can be taken to represent loosely multisensory information as when for example range profile data would be concatenated with polarization information to form a multisensory target representation. The lack of correlation between the polarization response and the range profile is a central assumption here. This lack of correlation helps also separate the target representations in the multisensory target signature space and this is desireable for enhancing cognition.

The results of simulations with these composite signals are tabulated in Table 4 and plotted in Figure 8 where the performance of networks trained on different targets is shown. The composite networks in this case had multilayered feedforward networks at the front end. Multilayered networks with one hidden layer of neurons were used since they are known to have more flexibility in partitioning the phase space that the simpler perceptrons [35]. The feedforward networks process non-overlapping segments (overlap $d = 0$) of the composite range profiles obtained by the process described in the foregoing paragraph. The learning parameter α and the momentum parameter β are fixed at 0.75 and 0.5 respectively. The internal neural threshold is fixed at zero in all the simulations, since working at higher thresholds makes the networks more sensitive to noise and hence the undecidibility about known targets increases as noise in the system

				Error (E) or Undecided (U)		
N_s	N_i	N_h	N_o	B52	B747	S.Sh.
1	256	32	32	0	0	100(E)
2	128	24	32	0	0	98(E)
4	64	16	32	4(U)	0	50(E)
8	32	10	16	12(U)	6(U)	0

Table 4: Cognitive performance as affected by processing the composite range profiles in segments by multiple feedforward networks. The networks are trained on 50 percent composite range profiles of the B52 and B747 and tested on all composite profiles from these two targets as well as the unknown target (the Space Shuttle). There is no overlap between segments, and training parameters α and β are fixed at 0.75 and 0.5, respectively.

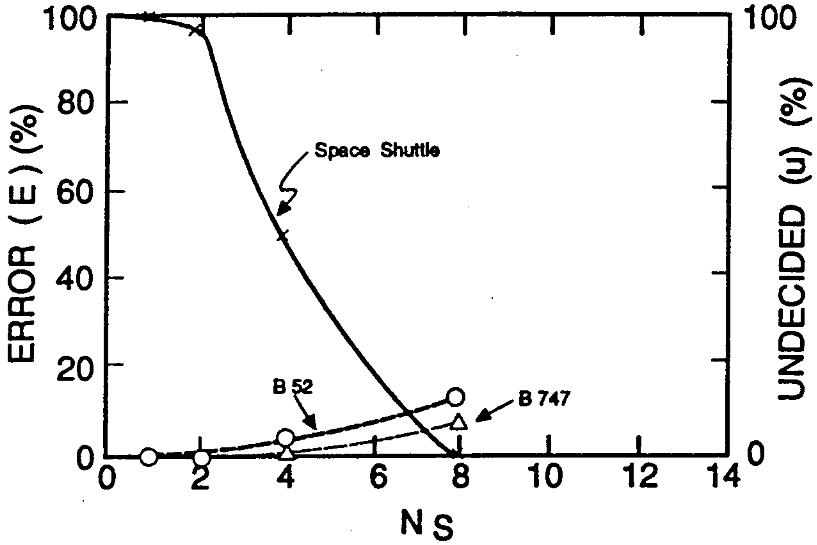

Figure 8: Cognitive performance as affected by processing the composite range profiles in segments by multiple feedforward networks. The networks are trained on 50 percent composite range profiles of the B52 and B747 and tested on all composite profiles from these two targets as well as the unknown target (the Space Shuttle).

increases. One strong trend is evident from Table 4 and the corresponding plots: that *the cognition capability of the system dramatically improves as the number of segments is increased.* This is independant of the targets used in training and testing the networks. We note that with eight segments of 32 data points and therefore 32 input neurons each, the recognition capability of the system is very good although performance on known targets deteriorates to some extent, depending on which targets were used to train the networks. The test statistics are obtained by testing the network with one signature vector instead of the majority vote technique which we use later in this section. For example, when $\eta = 50$ percent of available composite range profiles of only the B52 and B747 are used to train the network, the Space Shuttle (unknown target) is classified erroneously from all its composite profiles when a single network is used. Using four segment networks reduces this misclassification rate by 50 percent with negligible effect on network performance on known targets. Doubling the number of segments to eight, the misclassification rate on the unknown target drops to zero. The undecidibility on known targets (the B52 and the B747, in this case) rises moderately: 12 percent for the B52 and 6 percent for the B747.

When the B52 and Space Shuttle are used as known targets and the B747 as the unknown target (see Table 5 and Figure 9), the misclassification rates on the unknown target with 4 and 8 equal segments are 36 and 2 percent respectively. Increasing the number of segments is not possible since the segment length becomes too small for reasonable features to exist or be extracted and hence the net does not converge. The undecidibility on the known targets in this case, with 8 segments, is 20 percent for the B52 and 6 percent for the Shuttle. The maximum number of segments for which the nets converged is 9 and there was an overlap of 4 points between the segments in this case. This suggests that long composite signature vectors are desired and that is why multisensory information is important to consider.

In the final combination, with the B747 and the Shuttle as the known targets and the B52 as the unknown target, the misclassification rates on the unknown target with 4 and 8 equal segments are 22 and 8 percent respectively (see Table 6 and Figure 10).

The undecidibility on the known targets, with 8 segments, is 8 percent for the B747 and 6 percent for the Shuttle. In this case we were able to increase the number of segments to 14 without problems of convergence. However, the 14 segments had an overlap of four points and a length of 22 points. The misclassification error on the

				Error (E) or Undecided (U)		
N_s	N_i	N_h	N_o	B52	S.Sh.	B747
1	256	32	32	0	2(U)	94(E)
2	128	24	32	0	0	60(E)
4	64	16	32	12(U)	0	36(E)
8	32	10	16	20(U)	6(U)	2(E)
9	32	10	16	24(U)	8(U)	2(E)

Table 5: Cognitive performance as affected by processing the composite range profiles in segments by multiple feedforward networks. The networks are trained on 50 percent composite range profiles of the B52 and the Space Shuttle and tested on all composite profiles from these two targets as well as the unknown target (the B747).There is no overlap between segments, and training parameters α and β are fixed at 0.75 and 0.5, respectively. For the case of 9 segments, the overlap is 4, and α and β are fixed at 0.6 and 0, respectively.

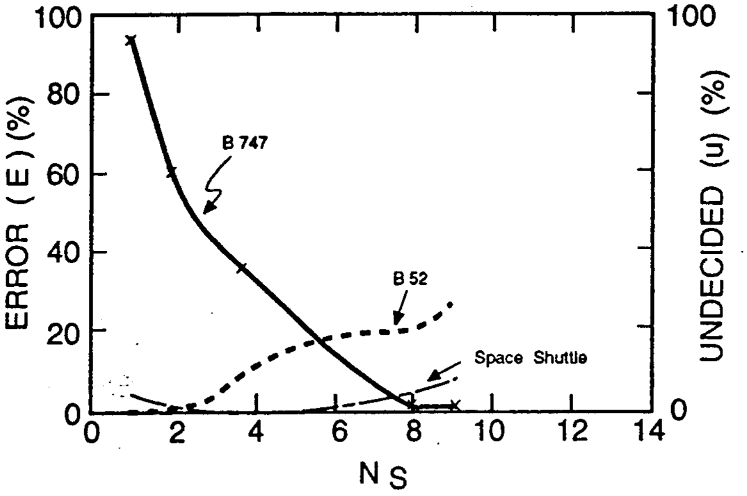

Figure 9: Cognitive performance as affected by processing the composite range profiles in segments by multiple feedforward networks. The networks are trained on 50 percent composite range profiles of the B52 and Space Shuttle and tested on all composite profiles from these two targets as well as the unknown target (the B747).

				Error (E) or Undecided (U)		
N_s	N_i	N_h	N_o	B747	S.Sh.	B52
1	256	32	32	0	0	50(E)
2	128	24	32	0	0	44(E)
4	64	16	32	0	0	22(E)
8	32	10	16	8(U)	6(U)	8(E)
9	32	10	16	0	0	10(E)
11	26	8	16	2(U)	6(U)	4(E)
12	25	8	16	8(U)	6(U)	2(E)
14	22	8	12	8(U)	6(U)	2(E)

Table 6: Cognitive performance as affected by processing the composite range profiles in segments by multiple feedforward networks. The networks are trained on 50 percent composite range profiles of the B747 and the Space Shuttle and tested on all composite profiles from these two targets as well as the unknown target (the B52). The overlap and learning parameters for 1 to 9 segments is the same as Table 8.12. For 11 segments, overlap is 3 and α and β are 0.5 and 0 respectively. For 12 and 14 segments, overlap is 4 and α and β are 0.4 and 0.6 respectively.

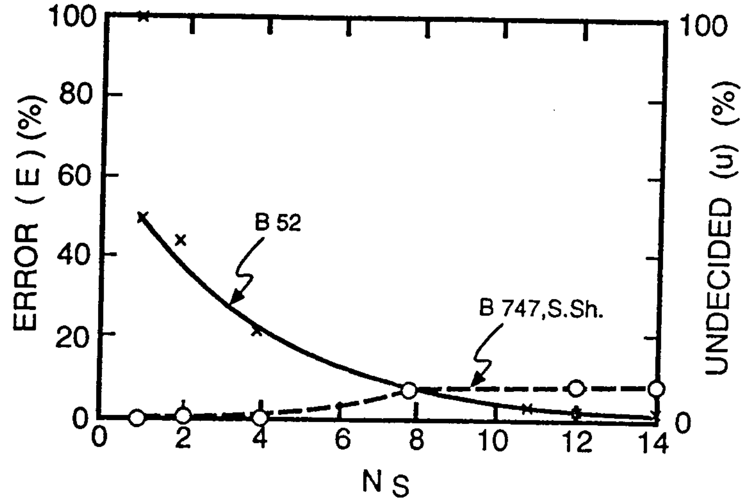

Figure 10: Cognitive performance as affected by processing the composite range profiles in segments by multiple feedforward networks. The networks are trained on 50 percent composite range profiles of the B747 and Space Shuttle and tested on all composite profiles from these two targets as well as the unknown target (the B52).

unknown target was reduced to 2 percent (with 14 segments) without any increase in the undecidibility on the known targets. The degree of segmentation possible is hence seen to be a function of the known targets, used to train the networks. It can be seen that with more complex targets, the maximum number of segments possible is smaller than with less complex targets. It is intuitive that separating more complex features is more difficult and hence a greater number of sample points per segment are required to define them i.e. more complex features require a broader context. An analogy can be seen with the problem of extracting a generating rule from a given series of numbers. If the series is a simple one, such as $1, 2, 3, 4, \ldots$, one can immediately see from a few numbers that the i-th element is simply gotten by adding 1 to the $(i-1)$-th element. A more difficult sequence may require many more elements before a generating rule can be extrapolated.

If the maximum number of segments is fixed at eight, we see that the misclassification error is reduced to zero for some cases but for other training sets it has a small positive value. The majority vote technique in which one decides on the basis of responses to three aspect queries from the target, can then be used with advantage once the

	Error (E) or Undecided (U)								
	B52			B747			S.Sh		
N_s	No	R	A	No	R	A	No	R	A
1	0	0	0	0	0	0	100	100	100
2	0	0	0	0	0	0	92	98	98
4	0	0	0	0	0	0	46	43	41
8	16	6.5	15.2	8	1.8	0.9	6	1	0

Table 7: The effect of majority vote on cognitive performance of the multiple segment network. "No" indicates that no vote is taken, "R" indicates a vote of 3 randomly selected composite profiles, and "A" indicates a vote of adjacent profiles separated by an angular distance of 0.2°. The network was trained with 50 percent of the data from B52 and B747 and tested with all profiles from these two targets and an unknown target (Space Shuttle). The learning parameters used are $\alpha = 0.4$ and $\beta = 0.6$ except when $N_s = 8$, in which case $\alpha = 0.4$ and $\beta = 0$ are used.

misclassification rate has been reduced to less than five or six percent, as illustrated by Table 7. The majority vote technique is applied in the following manner. The networks are initiated by three radar signatures in succession and the outputs are recorded. If the network responds at least twice identifying a given target, positive identification is indicated. The three target signatures can be selected randomly over a given angle or can be adjacent. Both cases are shown, and give similar results. We used 1000 trails in each case. Note that in a practical situation the adjacent range profile case would be much more appropriate, as when the radar tracks a moving target and target signatures of adjacent aspect angles are available to the network to make a decision.

6.4 Performance in Noise

The performance of networks was also evaluated when noisy signals with varying levels of zero mean Gaussian noise corrupted the composite signal. If P_s is the signal power and P_n is the noise power in the original signal then the signal to noise ratio SNR is defined as

$$SNR(dB) = 10log(\frac{P_s}{P_n}) \tag{5}$$

	SNR in dB								
	B52			B747			S.Sh		
N_s	8.5	2.85	1.1	8.5	2.85	1.1	8.5	2.85	1.1
1	0	2(U)	4(U)	0	0	0	98(E)	92(E)	96(E)
2	0	0	10(U)	0	0	6(U)	88(E)	74(E)	72(E)
4	6(U)	6(U)	30(U)	0	4(U)	10(U)	36(E)	26(E)	34(E)
8	34(U)	52(U)	76(U)	34(U)	50(U)	68(U)	6(E)	8(E)	2(E)

Table 8: The performance of multiple segment networks with different levels of noise. "E" indicates erroneous decisions and "U" indicates that the networks are undecided. The networks were trained by using 50 percent of the composite profiles from the B52 and B747 and tested on all profiles from these two targets as well as an unknown target (Space Shuttle).

The SNR of the original signal varies between 15 dB and 22 dB; the mean value is about 17 dB. If zero mean Gaussian noise, whose probability density function $g(x)$ is given by

$$g(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2} \qquad (6)$$

is used to contaminate the signal, the new SNR is given by

$$SNR(dB) = 10 log \frac{P_s}{P_n + \sigma^2} \qquad (7)$$

In the above equations x is a random variable, σ is the standard deviation of x about zero mean, and σ^2 is the variance and also the Gaussian noise power. The results of network performance with various signal-to-noise ratios are tabulated in Table 8 and plotted in Figure 11.

We see that as the length of one segment decreases, the system becomes more prone to be affected by high levels of noise. The effect of noise is less severe on performance of unknown targets than on recognition of known targets. For example with 4 segments, each of length 64, moderate levels of noise (upto SNR = 2.85) have little effect on network performance. With 8 segments of length 32 each, the performance on known targets deteriorates appreciably since the net cannot decide about their presence from an increasing number of the target aspects. The general conclusion we can draw from these results is that for good performance we need a reasonable number of segments of sufficient length. One way to achieve this is to include polarization information in the

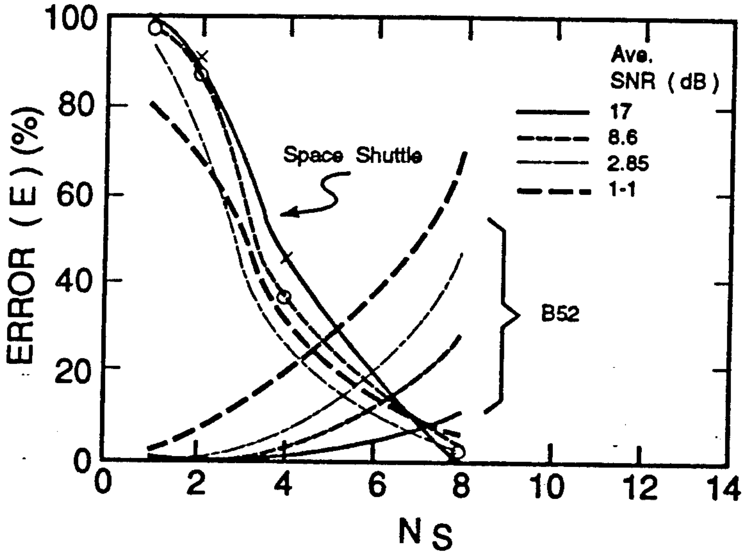

Figure 11: The performance of multiple segment networks with different levels of noise. "E" indicates erroneous decisions and "U" indicates that the networks are undecided. The networks were trained by using 50 percent of the composite profiles from the B52 and B747 and tested on all profiles from these two targets as well as an unknown target (Space Shuttle).

signature of the targets, i.e. to work with signature vectors consisting of concatenation of range profile information with polarization response (χ vs. frequency and ψ vs. frequency) of the target. We will elaborate on this point in section 7.1.

7 Designing a Radar Recognition System

Having described the basic aspects of a radar recognition system based on models of neural networks, we can tie our results together to propose a practical and autonomous system. Such a system is shown schematically in Figure 12, and can be described as a feature binding and cognitive hierarchial network. The system acquires interesting properties from processing partial spatial representations of a given object followed by an integration of partial decisions at the end. This approach offers some attractive benefits, such as

1. Modularity is introduced naturally, and hence the scaling problem of learning is considerably reduced. The problem of scaling can be explained by saying that neural net models are tested on toy problems do not always translate linearly to real (bigger) problems in terms of network size and/or learning time.

2. Reduces or eliminates ambiguities by making the cognition process dependent on the simultaneous occurence of a set of events at different locations, for example, like hitting a jackpot in a gambling machine, in that the correct window symbols must occur simultaneously in order that a winning condition (cognition in our case) does occur. See the one armed bandit analogy in Figure 13.

3. Enables the introduction of hierarchial processing, i.e., different levels of attractors, each level reducing the dimensionality of data but increasing the probability of correct recognition.

In the following subsections we will elaborate on different aspects of the system and how they complement each other to achieve excellent cognitive performance.

7.1 Signature Representations of Targets

Suppose one obtains representations for all possible manifestations of an object, that can occur in a practical setting. In the context of our ATR work, this means we have samples

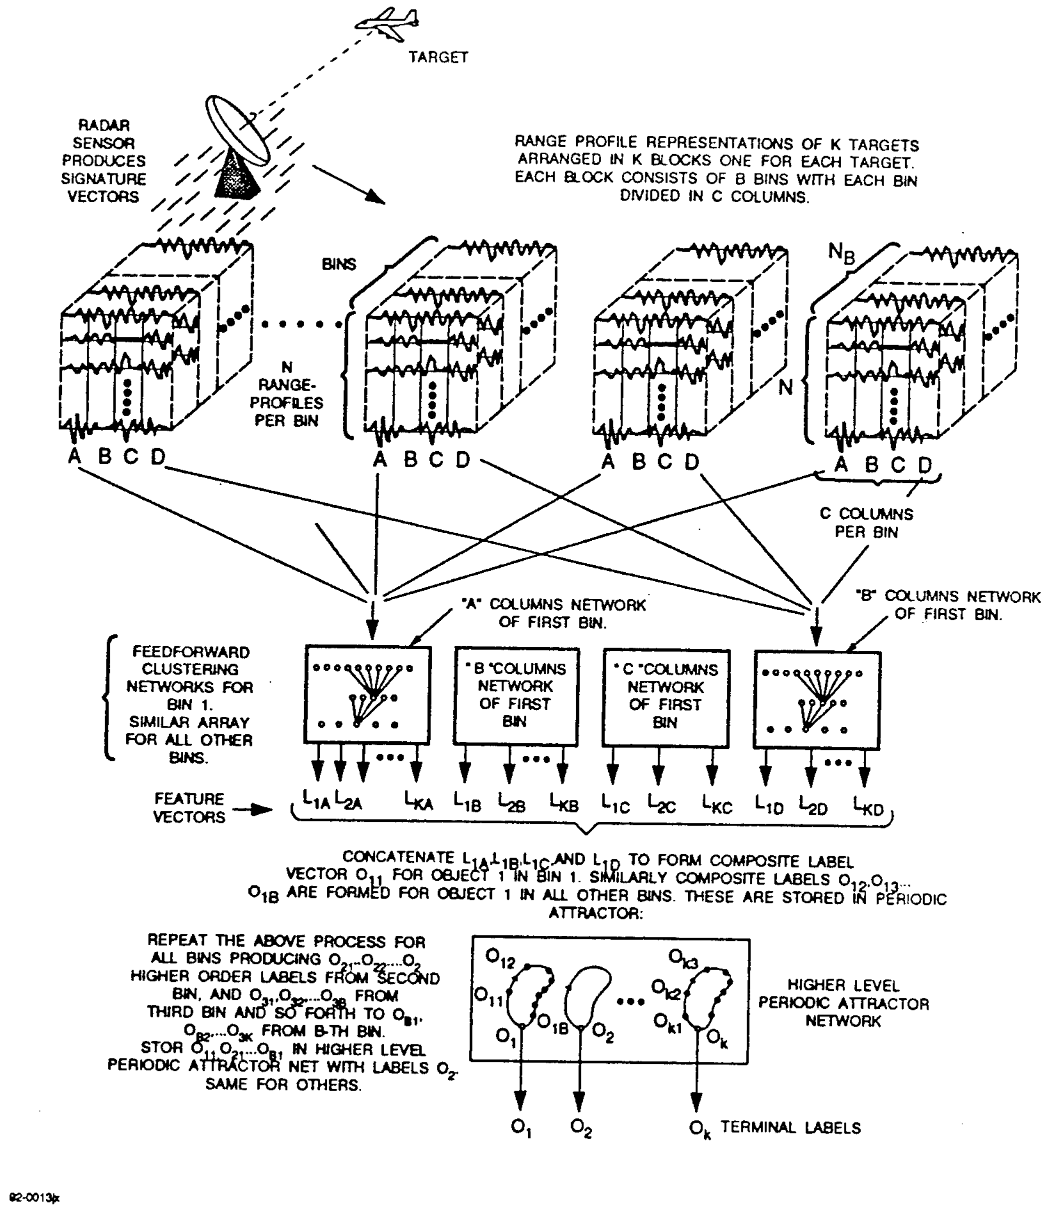
TARGET
RADAR SENSOR PRODUCES SIGNATURE VECTORS
RANGE PROFILE REPRESENTATIONS OF K TARGETS ARRANGED IN K BLOCKS ONE FOR EACH TARGET. EACH BLOCK CONSISTS OF B BINS WITH EACH BIN DIVIDED IN C COLUMNS.
BINS
N_B
N RANGE-PROFILES PER BIN
N
A B C D
A B C D
A B C D
A B C D
C COLUMNS PER BIN
"A" COLUMNS NETWORK OF FIRST BIN.
"B" COLUMNS NETWORK OF FIRST BIN.
FEEDFORWARD CLUSTERING NETWORKS FOR BIN 1. SIMILAR ARRAY FOR ALL OTHER BINS.
"B "COLUMNS NETWORK OF FIRST BIN
" C "COLUMNS NETWORK OF FIRST BIN
FEATURE VECTORS
L_{1A} L_{2A} L_{KA} L_{1B} L_{2B} L_{KB} L_{1C} L_{2C} L_{KC} L_{1D} L_{2D} L_{KD}
CONCATENATE L_{1A},L_{1B},L_{1C},AND L_{1D} TO FORM COMPOSITE LABEL VECTOR O_{11} FOR OBJECT 1 IN BIN 1. SIMILARLY COMPOSITE LABELS O_{12},O_{13}... O_{1B} ARE FORMED FOR OBJECT 1 IN ALL OTHER BINS. THESE ARE STORED IN PERIODIC ATTRACTOR:
REPEAT THE ABOVE PROCESS FOR ALL BINS PRODUCING O_{21},.O_{22},...O_2 HIGHER ORDER LABELS FROM SECOND BIN, AND O_{31},O_{32},...O_{3B} FROM THIRD BIN AND SO FORTH TO O_{B1}, O_{B2},...O_{3K} FROM B-TH BIN. STOR O_{11},O_{21}...O_{B1} IN HIGHER LEVEL PERIODIC ATTRACTOR NET WITH LABELS O_2. SAME FOR OTHERS.
O_{12}
O_{11}
O_1
O_{1B}
O_2
O_{k3}
O_{k2}
O_{k1}
O_k
HIGHER LEVEL PERIODIC ATTRACTOR NETWORK
O_1
O_2
O_k TERMINAL LABELS
92-0013k

Figure 12: A feature binding hierarchical cognitive network

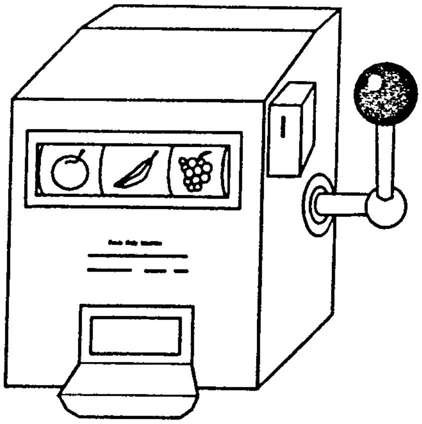

Figure 13: One armed bandit analogy of the coincidence of events in a cognitive system to signal positive cognition.

of normalized range profiles and/or all depolarization signatures of a given target scale model, falling within an expected solid angle of encounter for that target. Note that the representations are independent of the range to the target by virtue of the sensor characteristics used to produce them. Both types of signatures are otherwise influenced by noise and clutter and hence from the outset any cognitive system must be robust. Schematic depictions of the different types of range-independant target representations are shown in Figure 14(a). The range-profiles basically contain amplitude and phase information while the plots of ellipticity and inclination angles of the polarization ellipse of the echo versus frequency give the polarization information. One can concatenate these representations (see Figure 14(b)) to produce composite multisensory signatures which are characteristic of given targets. Assume that in the solid angle of encounter of interest for a given target, there are N_r such signatures. Let all these signatures be partitioned into N_b bins or groups, each containing $N = N_r/N_b$ signatures. In Figure 15, all members of one group are arranged in one plane, and correspond to the signatures within a small solid angle. The signatures are then partitioned into N_c columns. In the figure, $N_c = 4$ and the columns are labeled A,B,C and D. The data in each representation is viewed as containing specific features of the object. Some segmentation will be natural, for example the plots of the two polarization parameters, while others are somewhat

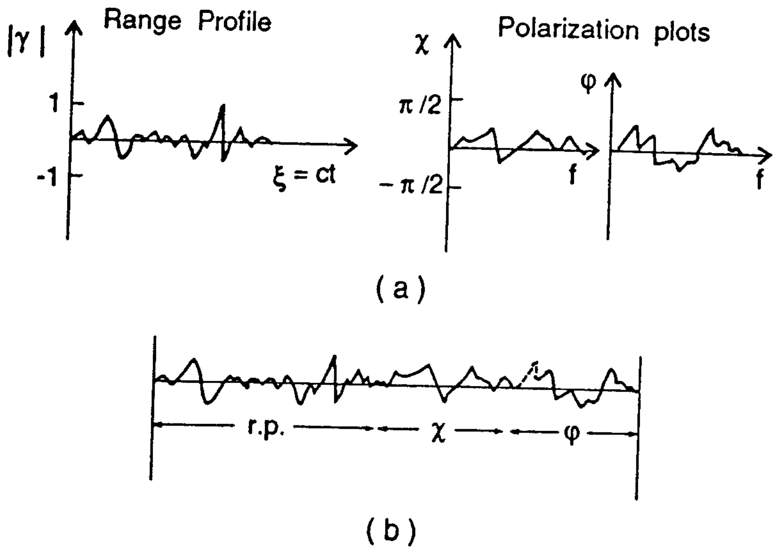

Figure 14: (a) Range independant Target Representations and, (b) A composite signature of a target obtained by concatenating it's range profile and polarization responses.

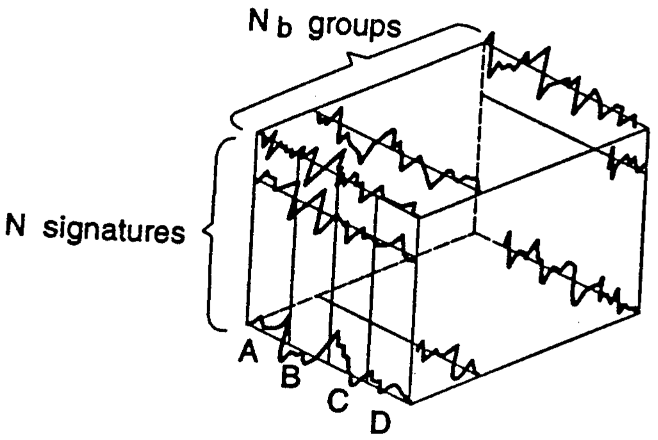

Figure 15: Data representation of one target or object. A total of N_r multisensory representations partitioned into N_b groups or bins each containing $N = N_r/N_b$ representations.

arbitrary divisions. To make the segments larger and the transition smoother, one can use overlapping segments. As illustrated by the simulation results given in the previous section the number of columns and bins is an important design parameter of the system and ultimately influences the performance of the network. Figure 12 shows feedforward clustering networks for bin 1. There would be a total of $N_b \times N_c$ such networks. Each network associates the data in its column with a given binary label and hence there would be $N_b \times N_c$ labels, which are not necessarily all distinct. Of course a given signature vector would trigger one label from each network.

7.2 Operational Principle of the System

The operation of the network can be visualized in the following manner. Figure 16 shows the expected angles of encounter of two known targets. For training the network, each solid angle of encounter is subdivided in smaller solid angles called bins. For example, five bins are shown in the figure for each target. A signature vector of a given target within an expected angle of encounter would then lie in one of these bins. Each signature vector is divided into a certain number of segments, labeled $A, B, C, \ldots$. As an example, each bin which contains a certain number of signature vectors is shown divided into three segments in Figure 16. The signature segments are fed into banks of feedforward networks, each trained to recognize a given target over a small angle of encounter by associating mini-labels of the target with its corresponding signature vector segment. For example, if the target signature belongs to bin 1 of target 1, its segments are fed into all the banks of networks, each bank containing 3 networks in our case. Then the networks in the bank shown on the right will output mini-labels L_{1A}, L_{1B} and L_{1C} when initiated by segments A, B and C of the signature vector. These mini-labels are concatenated to form a larger composite object label, C_{11} in this case, representing the particular solid aspect angle of the target. If the target is seen at another aspect angle contained in another solid angle, another bank of networks forms the corresponding label of the target associated with that solid angle. With proper design of the system, the probability that another bank will output C_{11} or another target's label is negligible. All the binary composite labels belonging to one target are stored with a master label for that target in a periodic attractor network shown in Figure 16. For two targets we would have two isolated, i.e., non-intersecting periodic trajectories stored in the same network. For

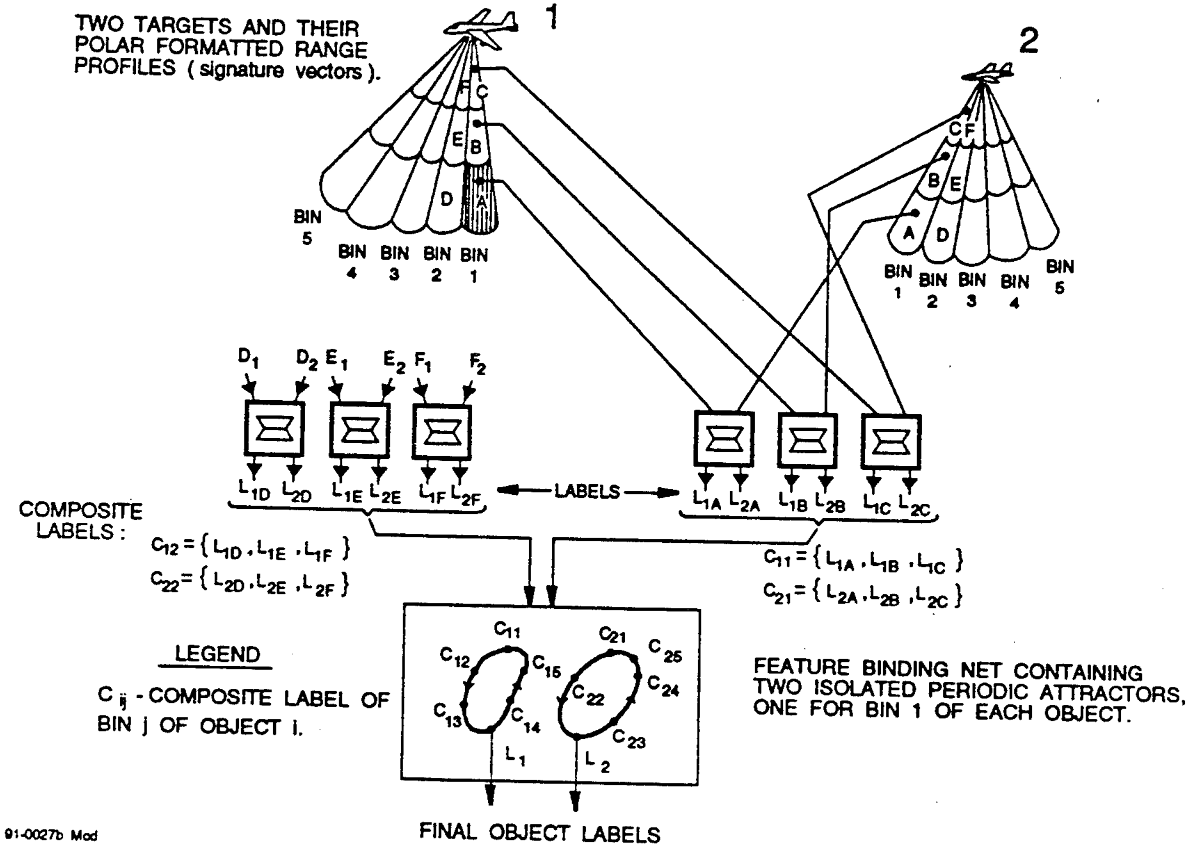
TWO TARGETS AND THEIR POLAR FORMATTED RANGE PROFILES (signature vectors).
1
2
F C
E B
D A
BIN 5
BIN 4
BIN 3
BIN 2
BIN 1
C F
B E
A D
BIN 1
BIN 2
BIN 3
BIN 4
BIN 5
D1 D2 E1 E2 F1 F2
L1D L2D L1E L2E L1F L2F
LABELS
L1A L2A L1B L2B L1C L2C
COMPOSITE LABELS :
C12 = { L1D , L1E , L1F }
C22 = { L2D , L2E , L2F }
C11 = { L1A , L1B , L1C }
C21 = { L2A , L2B , L2C }
LEGEND
C ij - COMPOSITE LABEL OF BIN j OF OBJECT i.
C11 C12 C15 C13 C14 L1
C21 C25 C22 C24 C23 L2
FEATURE BINDING NET CONTAINING TWO ISOLATED PERIODIC ATTRACTORS, ONE FOR BIN 1 OF EACH OBJECT.
FINAL OBJECT LABELS
91-0027b Mod

Figure 16: Organization of Cognitive Autonomous Target Recognition System

example, the composite labels C_{11}, C_{12}, C_{13}, C_{14} and C_{15} which represent the response of banks of networks to target signatures from the five bins belonging to target 1, are stored in a closed trajectory with the master label L_1. If a known target appears then it will trigger mini-labels, say L_{1A}, L_{1B} and L_{1C} representing the target and when concatenated together will form one of the vectors stored on the trajectory of the given target, C_{11} in this case. C_{11} will then trigger the trajectory containing the trajectory containing the master label L_1. We call this event "Jackpot" because of the similarity of what happens in hand operated gambling machines : alignment of certain labels in parallel rotating wheels signifies a jackpot (see Figure 13). If the composite representation is of a novel object, the chances of it erroneously producing a composite label vector stored in one of the two periodic attractors and thereby triggering a "Jackpot" in the same bin will be very remote and this furnishes the basis for robust cognition. The recognition process outlined above is neatly summarized in Figure 17. We have in this argument rested on the assumption that the periodic attractor trajectories representing different objects are appropriately isolated.

7.3 The Role of the Periodic Attractor

The periodic attractor network serves to provide a binding mechanism by which the feature outputs from the feedforward network banks are bound in the final step of the cognition process. It complements the function of the feature forming feedforward networks which furnish generalization and provide robustness against noise and scaling of the signal level, i.e. have a wide dynamic range. Although the fibres of cognition lie in the multiple local decisions arriving in parallel at the feedforward net outputs, the task of selectively weaving them into a substantial cognitive fabric is done by the PANs. Since one ultimately wants to implement the networks in hardware some comments about the robustness of the periodic attractor nets against setting weights with a given imprecision as well as element failure are in order.

M binary vectors (N-dimensional) can be stored as stations on a sequential trajectory in a fully connected N neuron network in the following manner. We denote the synaptic strength from neuron j to i denoted by w_{ij}. For simplicity, consider storing only one trajectory. The m-th vector on the trajectory is denoted by $V^{(m)} = (v_1^{(m)}, v_2^{(m)}, \ldots, v_i^{(m)}, \ldots, v_N^{(m)})$. The w_{ij}s thus form an $N \times N$ square matrix W. Dur-

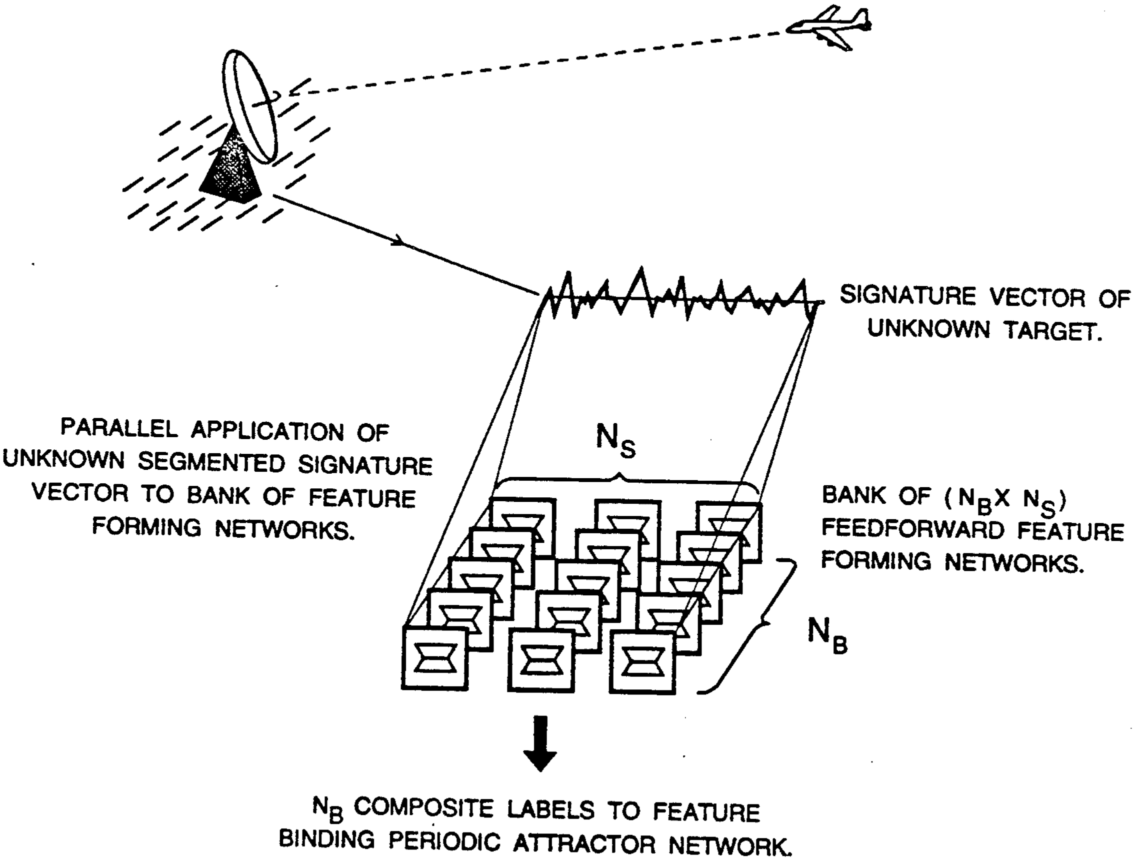
SIGNATURE VECTOR OF UNKNOWN TARGET.
PARALLEL APPLICATION OF UNKNOWN SEGMENTED SIGNATURE VECTOR TO BANK OF FEATURE FORMING NETWORKS.
N_S
BANK OF (N_BX N_S) FEEDFORWARD FEATURE FORMING NETWORKS.
N_B
N_B COMPOSITE LABELS TO FEATURE BINDING PERIODIC ATTRACTOR NETWORK.

Figure 17: Query and recognition of some target

ing learning, the net updates its state vector synchronously according to $v_i = B(u_i - \theta_{neuron}, \theta_{high}, \theta_{low})$, where $u_i = \sum_j w_{ij}v_j$ is the action potential of the i-th neuron, θ_{neuron} is its internal threshold, and θ_{high} and θ_{low} are the upper and lower limits of a neural band gap function $B(.)$ used during learning to ensure "good" learning. During the recognition phase, θ_{high} and θ_{low} are set at mean level, 0.5, which gives rise to zero band gap during recall. It is observed that using a zero band gap during the training phase produces a network with negligible tolerance to setting weights with some imprecision. Also filling the network more and more ($M \approx N$) reduces this tolerance whereas with less filled networks ($M \ll N$) the trajectory can be triggered by more and more vectors. Hence it is necessary to tailor the periodic attractor net to provide the desired isolation of the trajectory from the rest of the phase space and to allow some imprecision in setting weights in hardware.

8 A Design Example

It is best to illustrate the operation of the envisioned target recognition system with a simple design example. Consider a situation where one needs to recognize only two targets from their signature vectors, i.e. to be able to tell from a given signature vector whether it comes from these known objects or not, and if yes, then which one. To simplify the analysis, assume that the most probable target aspects lie in a solid angle of 40° in elevation, extending from 20° to 60° in elevation, and 70° in azimuth extending from head-on to both broadsides of the target. Note that for symmetric targets this translates to an azimuth angle of 140 degrees. If we choose the bin size to be 20 by 20 degrees, then the number of bins for one target is $N_b = 70 \times 40/20 \times 20 = 7$. The number of signatures required per bin for teaching the feedforward networks would depend on the complexity of the targets in the set of targets to be encountered. For the scale model targets we have used a choice of 0.5° and 1° as the angular distance between adjacent samples in azimuth and in elevation is appropriate. This estimate is based upon the variation of range profile correlations as a function of the difference in the aspect angles of the three targets (B52, Boeing 747, and Space Shuttle). For example, the range profiles of the B52 scale model have useful correlation (which is above the cross-correlation between different targets) over an angle of 0.8° in azimuth. The angles, over which useful correlation exists, for other two targets are greater. Since, the target extent in the elevation direction is less, the angle over which useful correlation exists will be greater. With this consideration in mind we can calculate the number of samples required per bin per target to be about $N_{sig} = 20 \times 20/0.5 \times 1.0$ or about 800 samples. In this example, one would need $7 \times 800 = 5600$ equally spaced samples per target to provide a library of echoes from which the training set to teach the networks can be chosen. If the number of segments is chosen to be $N_s = 4$, as shown in the schematic of the cognitive system, the total number of feedforward networks in the system is 28, arranged in 7 banks of 4 each. With the output neurons of each feedforward network chosen to be 8, the integrated label at the output would be 32 bits. For each bank a different output label can be selected and hence there are 7 different possible labels belonging to one target and associated with its different aspect regions. These 7 labels can be stored with a master label for the target for a total of 8 labels on a closed trajectory in a periodic attractor network. Similarly for the other target, we can store 8 labels on a different trajectory that does not intersect the first trajectory in the same periodic attractor network.

The system works as follows. The signature of an unknown target is input in segments to all the banks in parallel. The outputs from the networks of each bank are concatenated into a composite label or feature vector and used to interrogate the periodic attractor network. This can be done serially by applying the output of each bank, observing the behaviour of the periodic attractor network before applying the output of the next bank and so on. One can also do this operation in parallel by using seven identical periodic networks, but as the number of banks increases, this may be impractical and one can have several banks sharing a periodic attractor network. If all outputs from a bank are consistent, i.e. correspond to a given target, then the concatenated output lies on that target's periodic attractor and hence will trigger it. If the output labels are not consistent, i.e. do not belong to the same target, the concatenated output has a certain Hamming distance from the vectors stored on the two trajectories and will not trigger any one, provided the trajectories are well isolated. A feed-forward network in a bank outputs a mini-label belonging to one of the two objects depending on the similarity of its input to the signature segments of the targets used to teach the network. Hence, by choosing the length of these labels and their Hamming distances from each other one can determine the minimum Hamming distance of any concatenated label vector not on one of the closed trajectories, from the trajectory. What we need then is to have the

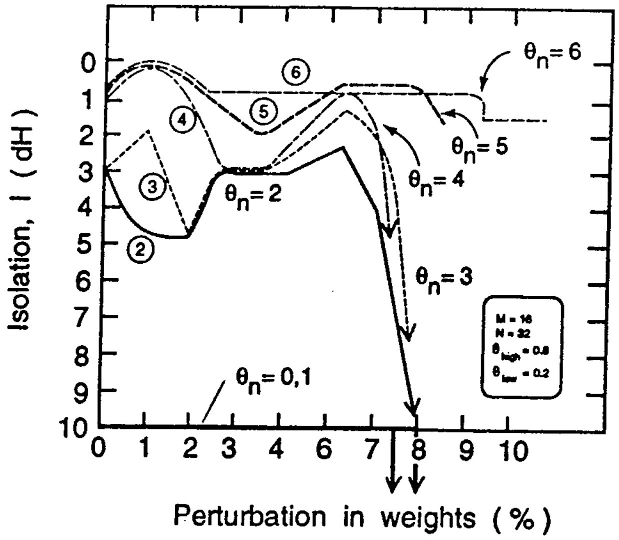

Figure 18: The isolation I of the periodic attractor network as a function of weight perturbation. The isolation-weight perturbation curves for different values of θ_n (encircled) are drawn. The network is half-filled, i.e. $M = 0.5N$.

trajectories isolated enough so as not to be triggered by one of these "lurking" vectors, which lie at a minimum distance from them. These ideas are illustrated quantitatively below.

We have chosen the length of the outputs from the segments to 8. Let the minimum Hamming distance, dH, between the different mini-labels be 5. The concatenated output will be $N = 32$ bits in length, and the minimum distance of a composite output not on the trajectory from any label on the trajectory will be $dH_{min} = 5$. Hence we need a periodic attractor with as isolation such that no vector at a Hamming distance greater than $dH = 4$ triggers one of the trajectories. Based on this value of the minimum isolation required, and a given tolerance in setting weights in hardware, we proceed to find out the other parameters of the net, namely, θ_{neuron}, θ_{high} and θ_{low}. θ_{neuron}, the internal threshold of a neuron, mainly controls the degree of isolation of the periodic attractors. θ_{high} and θ_{low} during training, mainly determine the tolerance of setting weights with a given imprecision in the periodic attractor.

Figure 18 shows how the isolation and tolerance in weight values change with increasing the internal neural threshold, θ_{neuron}. The values of θ_{high} and θ_{low} are fixed during training at 0.8 and 0.2 respectively, while during the recall phase both are fixed at 0.5.

At low values of the internal neural threshold, the isolation of the trajectories is very low, although the tolerance to weight imprecision is reasonably good. As the threshold is increased, isolation improves with a corresponding reduction in the tolerance to weight imprecision. We find that the required value of the internal neural threshold to achieve required isolation of $dH = 4$ is $\theta_{neuron} = 4$. The tolerance allowed in setting weights in hardware is about 6%. If a greater tolerance in weights is desired, one can increase the length of the segment output lables and hence increase the minimum Hamming distance of vectors outside the trajectories from the trajectories.

9 Summary and Discussion

This paper addresses the issue of how the neural paradigm can be applied to an electromagnetic scattering problem. Traditionally, the inverse scattering problem has been a central issue in electromagnetics. The approach is to invert the measured data. This problem is known to be ill-posed and therefore difficult to solve. Regularization methods are applied to facilitate solution.

Inverse scattering requires use of *a priori* knowledge of the mechanism involved in creating the measured data. Living organisms seem to be adept at solving inverse problems. The neural paradigm of information processing is therefore important. The approach adopted in this paper is applicable to other problems in inverse scattering and not only to ATR.

Cognition, which is an important attribute of biological systems, has been generally neglected in most of ANN research. We have explained in detail why it is crucial to success in many applications. We have also argued in support of the hypothesis that to make a neural network cognitive, it must be nonlinear, dynamical and computing with with diverse attractors. Also it must be capable of bifurcating between them depending on the nature of the objects being presented to the network. Our results also indicate why multisensory information may be of great importance in enhancing cognition and reducing ambiguities between similar objects. It is worth noting that the composite hierarchial network we describe handles multisensory information, in the form of concatenated multisensory signature vectors, in a natural way.

Usually neural net architectures and learning methods are adapted to tasks that a system is required to perform, as evidenced by many biological systems. The radar identification problem also has its own peculiarities. Different neural architectures with their associated modes of operation offer diverse potential, not always well quantified. and it is a tricky business to put them together to have the desired effect. In this work we found it necessary to couple different types of attractor networks as a means to achieving performance that is not achievable by simpler networks alone.

A practical concern in neural network resaerch is that of scalibility. A legitimate complaint is that neural net models are generally tested on toy problems and that translating them to real problems is not feasible or straightforward in terms of time and network size. One of the advantages of gained using composite networks, as proposed in this paper, is that the problem of scalibility is addressed efficiently. A certain composite network, comprising the feature forming and feature binding networks, divides its environment into two disjoint domains: that of objects known to it and that of all the other objects and signals. Hence if new objects have to be learned, one does not need to retrain this network to include the information about new objects in addition to relearning the "already known" objects. The new objects can be added through additional composite network modules. This feature is very helpful in enabling one to train the networks only once. This means that even somewhat prolonged training times maybe acceptable. The feature of data (signature vector) segmentation ensures that each modular label forming network is relatively small and hence its training time is not as protracted as compared to the traditional approach when all the data is to be taught to a single large network.

We have worked with very simplified models of neural networks in trying to realize certain characteristics which are critical to the solution of the recognition problem. We feel that more realistic networks, for example those incorporating the temporal dimension would provide one with increased power and flexibility to approach such problems. For example in the feedback periodic attractor network we have used we have to ensure synchronous operation (update of state vector), possibly by external means. Biological evidence suggests that groups of neurons can become sychronized under certain conditions to operate as a synfire chain [36], or that trajectories can be realized in the state space of a network with asynchronous operation. Work along these lines is already being pursued in our work.

The use of multisensory information in facilitating recognition has been indirectly demonstrated. This concept needs to be explored further, specially in finding the effects of using increasingly diverse modalities of information in increasing not only the performance of the cognitive system but also the cost of acquiring that additional information. Other related aspects such as how uncorrelated information is used and can be beneficial need to be looked into to arrive at some knowledge of recognition phenomena used by various biological systems.

10 Acknowledgement

This work was supported by SDIO/IST, the Office of Naval Research, JPL/ASAS Program Office and the Army Research Office.

11 Appendix

We here describe and present results obtained with a simple training algorithm to learn and recall arbitrary sequences of pattern vectors in a fully connected artificial neural network, i.e., feedback network, and synchronous update. Note that no requirement about the orthogonality of patterns is made. We are given K sets of M_k N_d-dimensional pattern vectors to be stored as K different sequences in the N-neuron network. Let us start with a blank memory so that $w_{ij}^{(0)}=0$ for all $i, j = 1, 2, \ldots, N$. Consider the case when only one sequence is to be stored.

When an m-th pattern, $v^{(m)}$, is presented to the network it produces an output, $o^{(m)}$, which is compared with the desired output, $v^{(m+1)}$, and based on this the weights are updated as follows:

$$w_{ij}^{(t+1)} = w_{ij}^{(t)} + \lambda(v_i^{(m)}v_j^{(m+1)} - v_i^{(m)}o_j^{(m)}) \tag{8}$$

where λ is a positive learning parameter. The training process is continued until all the patterns have been stored in the desired order.

The method we have described above produces a network with the given states stored on open or closed trajectories. However, the isolation of the trajectories from the rest of the network phase space is uncontrolled. For the radar cognition scheme described in this paper, the isolation of the trajectories has to be controlled within prescribed limits. We would also need to be able to set the weights of the network with some tolerance, when implementing the network in hardware is contemplated. There are two parameters of interest that determine these characteristics of the network: the internal threshold of the neurons and

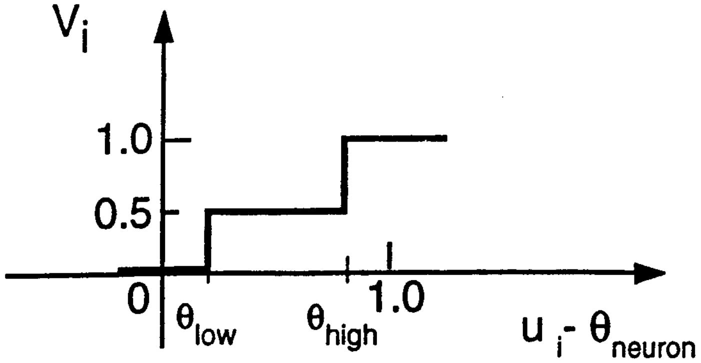

Figure 19: The neural bandgap function used during the training phase of the periodic attractor neural network.

the output function of the neuron. Raising the threshold of the neurons improves the isolation of the trajectories learned by the network, but also makes the performance of the network more sensitive to perturbations of synaptic weights. The flexibility to set weights with some tolerances is important when one needs to implement the network in hardware. This flexibility can be achieved by training the network with a bandgap or deadzone neuron output function as described below.

Assume that the net updates its state vector according to the neuron function $v_i = B(u_i, \theta_{neuron}, \theta_{high}, \theta_{low})$, shown in Figure 19, where $u_i = \sum_j w_{ij}v_j$ is the action potential or activation of the i-th neuron, θ_{neuron} is its internal threshold, and θ_{high} and θ_{low} are the upper and lower limits of a band gap used during learning to ensure "good" learning. Learning is then continued until the responses of all individual neurons to all patterns to be stored are either above $\theta_{neuron} + \theta_{high}$ or below $\theta_{neuron} + \theta_{low}$. During the recognition phase, both θ_{high} and θ_{low} can be set at mean level, 0.5, and a certain distortion in the input or synaptic weights can be rectified due the two buffer zones above and below the mean level, which were created during the learning phase. We observe that using a zero band gap during the training phase produces a network with negligible tolerance to setting weights with some imprecision. Also filling the network with more and more patterns ($M \approx N$) reduces this tolerance whereas with less filled networks ($M \ll N$) the trajectory can be triggered by more and more vectors. This illustrates the need to tailor the periodic attractor net to the situation at hand, *vis-a-vis* the amount of isolation of the trajectory desired and the variation allowed in setting weights in hardware.

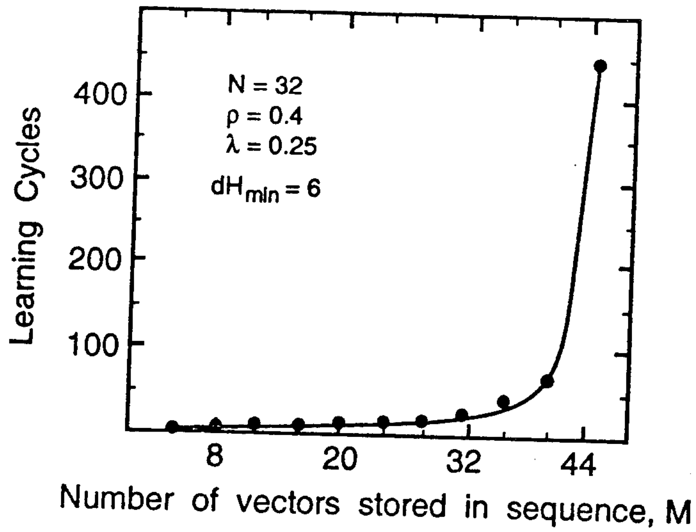

Figure 20: The number of learning cycles needed by a 32 neuron network to learn a sequence as a function of the number M of patterns in the sequence.

The number of learning cycles needed to learn M patterns or vectors using the above procedure is shown in Figure 20 for a neural network of $N = 32$ neurons. The stored sequences consisted of pattern vectors whose density ρ was about 0.4. The density ρ of a pattern vector is defined as the ratio of the number of 1's in the vector to the total number of its elements. The minimum Hamming distance between any pair of vectors in a given sequence was $dH_{min} = 6$. Similar behaviour is obtained for $N = 64$ and 128. It is seen from Figure 20 that learning is rapid as long as $M < N$. As M increases beyond N the number of learning cycles, and hence the learning time, required to learn the sequence of given pattern vectors increases exponentially.

As the internal threshold of the neurons increases the isolation of the trajectories learned by the network increases until it becomes a true filamentary trajectory, i.e., any vector which is not designed to lie on the trajectory does not trigger it. However, it might be desireable to allow a trajectory with a more or less controlled narrow region of attraction around it, so that an initiating vector lying in that region can also trigger the periodic attractor and the one outside it does not. An important point to note here is that unfamiliar states end in a sparse phase space of the network, most of them going to a ground attractor. Hence the network's response to unknown inputs is one of very low neural activity, whereas familiar states trigger a cyclic response. Those states that are partially familiar, initially elicit some firing before going to the sparse regions.

The isolation of the trajectories to be stored is mainly controlled by the internal threshold of the neurons. We define the isolation, I, of the network as the maximum Hamming distance, D_{max}, of vectors that can trigger the trajectory, from the stored trajectory, for given values of internal neuron threshold and weight perturbation. That is, any vector with a Hamming distance, $D > D_{max}$, from the trajectory will not trigger the trajectory. The Hamming distance, D, of an arbitrary vector from a stored trajectory is the minimum distance from the vectors on the trajectory. The perturbation P of the network is the percentage error in setting weights of the synapses between neurons.

The boundary of the periodic attractor is fractal in that not all vectors at a given Hamming distance from the trajectory will trigger the periodic attractor. Therefore in designing the network for isolation from a region beyond some Hamming distance D we actually design for isolation beyond a Hamming distance $D_s = D/f$ where f is a safety factor greater than one. Once the design Hamming distance D_s is decided, we have to find by experiment the values of three network parameters, namely θ_{neuron} the internal threshold of the neurons, and θ_{high} and θ_{low} for the neural bandgap function, which is used during the training phase.

In Figure 21, the isolation I (in dH: Hamming distance) of the network is plotted as a function of perturbation in weights as a measure of the robustness of the network. Each original learned weight is perturbed by randomly increasing or decreasing it by a fraction of its value, depending on the amount of perturbation P desired. In this network of $N = 32$ fully-connected neurons, $M = 16$ pattern vectors are stored on one closed trajectory. The internal neuron threshold is held fixed at $\theta_{neuron} = 4$, and was arrived at experimentally, assuming that an isolation of $D = 6$ is required. A safety factor of $f = 2$ was used to design for an isolation of $D_s = 3$.

As can be seen, the isolation fluctuates for various values of perturbation until for greater perturbations of weights, the trajectory is completely lost. Also, as evident from Figure 18 the minimum isolation increases as the internal threshold of the neurons is raised. For our application, it is the minimum isolation that we will use in the region where the the trajectory is recognized. As noted earlier initiating the network with unfamiliar states results in sparse or no activity. Hence the network's response to unknown inputs is one of very low neural activity, whereas familiar states trigger a cyclic response. Those states that are partially familiar, initially elicit some firing before going to the sparse regions.

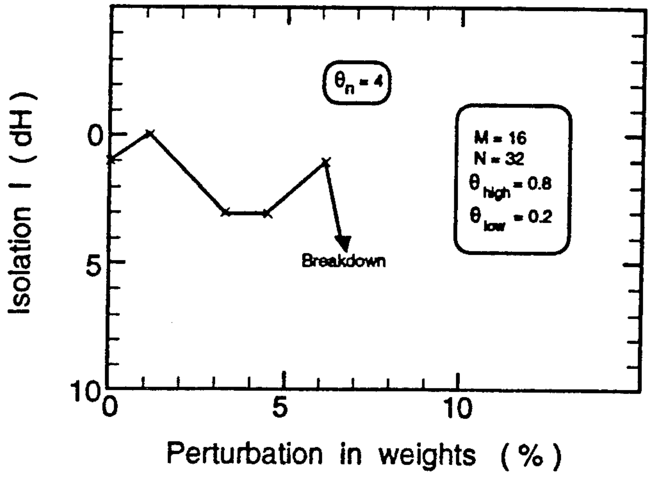

Figure 21: The effect of weight perturbations on the isolation of the periodic attractor network. Note that perturbing the weights more than 6 percent results in the loss of the trajectory in this case.

References

[1] A.N. Tikhonov and V.Y. Arsenin, *Solutions of Ill-Posed Problems*, J. Wiley, New York, 1977.

[2] N.H. Farhat and B. Bai, "Echo Inversion and Target Shape Estimation by Neuromorphic Processing," *Neural Networks*, Vol. 2, pp. 117-125, 1989.

[3] N.H. Farhat, "Microwave Diversity Imaging and Automated Target Identification Based on Models of Neural Networks," *Proc. IEEE*, Vol. 77, pp. 670-681, May 1989.

[4] N.H. Farhat, "Cognitive Neural Networks: Computing with Diverse Attractors," JPL Seminar, August 1989.

[5] A.J. Poelman, *et al.*, "Final Report of Working Discussion Group W-C on Polarization Utilization in High Resolution Imaging," in *Inverse Methods in Electromagnetic Imaging* ed. W.-M. Boerner *et al.*, Part 2, pp. 1277-1280, D. Reidel Publishing Company, 1985.

[6] A.A. Ksienski, *et. al.*, "Low Frequency Approach to Target Identification," *Proc. IEEE*, vol. 63, pp. 1651-1660, December 1975.

[7] C.E. Baum, "The Singularity Expansion Method," in *Transient Electromagnetic Fields*, L.B. Felsen (Ed.), Springer-Verlag, Heidelberg 1976.

[8] M.L. Blaricum and R. Mittra, "Problems and Solutions Associated with Prony's Method for Processing Transient Data," *IEEE Trans. Antennas Propagat.*, vol. AP-26, pp. 174-182, January 1978.

[9] D.H. Dudley, "Progress in Identification of Electromagnetic Systems," *IEEE Antennas Propagat. Newsletter*, pp. 5-11, August 1988.

[10] Bai, B., "Image Reconstruction and Target Identification Based on Neural Network Models," Ph.D. Dissertation in Electrical Engineering, The Moore School of Electrical Engineering, University of Pennsylvania, 1990.

[11] P. Beckmann, *The Depolarization of Electromagnetic Waves*, The Golem Press, Boulder, Colorado, 1968.

[12] J.D. Stratton, *Electromagnetic Theory*, McGraw-Hill, New York, pp 480-490, 1941.

[13] Y.S. Abu-Mostafa, "The Complexity of Information Extraction," *IEEE Trans. Inform. Theory*, vol. IT-32, pp 513-525, July 1986.

[14] C. Skarda and W. Freeman, "How Brains make Chaos in Order to make Sense of the World," Behav. and Brain Sciences, Vol. 10, pp. 161-195, Cambridge U. Press, 1987.

[15] Baird, B., "Bifurcation Analysis of Oscillating Network Model of Pattern Recognition in Rabbit Olfactory Bulb," *Neural Networks for Computing*, Denker, J.S., (Ed.), Am. Inst. of Phys., pp. 29-34, New York, 1986.

[16] D. Angluin and C. Smith, "Inductive inference: theory and methods," *A.C.M. Computing Surveys*, 15, pp. 237-269, 1983.

[17] R. Duda and P. Hart, *Pattern Classification and scene analysis*, John Wiley and Sons, 1973.

[18] L.G. Valiant, "A Theory of the Learnable," *Communications of the ACM*, 27(11), pp. 1134-1142, 1984.

[19] M.J. Kearns, *The Computational Complexity of Machine Learning*, An ACM Distinguished Dissertation 1989, The MIT Press, Cambridge, 1990.

[20] H.S. Wilf, *Algorithms and Complexity*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1986.

[21] M. Anthony and N. Biggs, *Computational Learning Theory, An Introduction*, Cambridge University Press, Cambridge, 1992.

[22] Y.S. Abu-Mostafa, "Learning from Hints in Neural Networks," *Journal of Complexity*, **6**, pp. 192-198, 1990.

[23] V.N. Vapnik and A. Chervonenkis, "On the uniform convergence of relative frequencies of events to their probabilities," *Theory Prob. Appl.* **16**, pp. 264-280, 1971.

[24] E.B. Baum and D. Haussler, "What size net gives valid generalization?," *Neural Computation*, **1**, pp. 151-160, 1989.

[25] E.B. Baum, "A Proposal for More Powerful Learning Algorithms," *Neural Computation*, **1**, pp. 201-207, 1989.

[26] O. Goldreich, et. al., "How to Construct Random Functions," *Journal of ACM*, Vol. 33, No. 4, pp. 792-807, October 1986.

[27] K. Hornik, et. al., "Multilayer Feedforward Networks are Universal Approximators," *Neural Networks*, Vol. 2, pp. 359-366, 1989.

[28] H. Shvaytser, "Even Simple Neural Nets Cannot be Trained Reliably with a Polynomial Number of Examples," *Proc. IJCNN*, Vol. II, pp. 141-145, Washington, D.C., June 1989. IEEE Catalog No. 89CH2765-6.

[29] A. Blumer *et al.*, "Classifying learnable geometric concepts with the Vapnik-Chervonenkis dimension," *Proc. ACM Symp. Theory Computing*, **18**, pp. 273-282, 1986.

[30] N.H. Farhat *et al.*, "Optical analogs to two-dimensional neural networks and their application in recognition of radar targets," in *Neural Networks for Computing*, J.S. Denker, Ed., New York, Am. Inst. of Phys., pp. 146-152, 1986.

[31] B. Ans, "Learning arbitrary associations: a neuromimetic model," *Advances in Cognitive Science Vol. 2: Theory and Applications*, G. Tiberghien (Ed.), pp. 62-85, Ellis Horwood LImited, Chichester, England, 1989.

[32] P. Kanerva, *Sparse Distributed Memory*, The MIT Press, Cambridge, 1988.

[33] Shepherd, G.M., *Neurobiology*, 2nd edition, Oxford University Press, Inc., New York, 1988.

[34] R. Eckhorn, H.J. Reitboek, M. Arndt and P. Dicke, "Feature Linking Via Stimulus-Evoked Oscillations: Experimental Results from Cat Visual Cortex and Functional Implications from a Network Model," Proc. IJCNN'89, Vol. I, pp. I-723-730, IEEE Cat. No. 89CH2765-6 (1989).

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

[36] Abeles, M., *Local Cortical Circicuits: an Electrophysiological Study*, Springer-Verlag, Berlin; New York, 1982.