A Connectionist/Neural Network Bibliography, Volume II

Dan Hammerstrom
Oregon Graduate Center
19600 N.W. von Neumann Dr.
Beaverton, Oregon 97006

Technical Report No. CS/E-87-008
October 1987

Abstract

This technical report contains a bibliography of connectionist and neural modelling papers and books as well as related references in the areas of computer architecture and brain science. Abstracts are included for many of the references.

BIBLIOGRAPHY

Alspector, Joshua and Robert B. Allen, "A Neuromorphic VLSI Learning System," Advanced Research in VLSI: Proceedings of the 1987 Stanford Conference, MIT Press, Cambridge, MA, 1987.

We have studied the properties of a model of neural nets based on a modified Boltzmann machine approach that can be implemented using VLSI technology. Such a VLSI system can use physical relaxation and massively parallel processing to achieve approximately a million-fold increase in speed of computation over an equivalent simulation run on a serial, digital computer. The proposed integrated circuit chip uses differential amplifiers as the analog neurons and simple digital processors as the dendrites and synapses whose connection strengths may change. The system achieves relaxation asynchronously using analog voltages rather than computer cycles. The system is trained rather than programmed, changing connection strengths digitally and synchronously in the process. We have performed functional simulations of such a chip to specify the architecture and we report our findings. We describe the important neuron-like features of our model and point out directions for future work in neuromorphic electronic systems.

Anderson, C.A., "Learning and Problem Solving with Multilayer Connectionist Systems," Ph.D. Dissertation, Amherst, MA, Sept. 1986.

The difficulties of learning in multilayered networks of computational units has limited the use of connectionist systems in complex domains. This dissertation elucidates the issues of learning in a network's hidden units, and reviews methods for addressing these issues that have been developed through the years. Issues of learning in hidden units are shown to be analogous to learning isseus for multilayer systems emplying symbolic representations.

Comparisons of a number of algorithms for learning in hidden units are made by applying them in a consistent manner to several tasks. Recently developed algorithms, including Rumelhart, et al.'s error back-propagation algorithm and Barto, et al.'s, reinforcement-learning algorithms, learn the solutions to the task much more successfully than methods of the past. A novel algorithm is examined that combines aspects of reinforcement learning and a data-directed search for useful weights, and is shown to out perform reinforcement-learning algorithms.

A connectinist framework for the learning of strategies is described which combines the error back-propagatrion algorithm for learning in hidden units with Sutton's AHC algirhtm to learn evaluation functions and with a reinforcement-learning algirithm to learn search heuristics. The generality of this hybrid system is demonstrated through successful applications to a numerical, pole-balancing task and to the Tower of Hanoi puzzle. Features developed by the hidden units in solving these task are analyzed. Comparisons with other approaches to each task are made.

Arbib, M.A. and J. Szentagothai, "Conceptual Models of Neural Organization," Neurosciences Res. Prog. Bull., vol. 12, no. 3, pp. 313-510, October 1972.

A worksession that covered the concept of hierarchy and its place in the structure of neural systems.

Arbib, Michael A. and Allen R. Hanson, Vision, Brain, And Cooperative Computation, The MIT Press, Cambridge, Massachusetts, 1987.

Athale, Ravindra A. and John N. Lee, "Optical Processing Using Outer-Product Concepts," Proceedings of the IEEE, vol. 72, no. 7, pp. 931-941, IEEE, 1984.

A row vector when left-multiplied by a column vector produces a two-dimensional rank-one matrix in an operation commonly called an outer product between the two vectors. The outer product operation can form the basis for a large variety of higher order algorithms in linear algebra, signal processing, and image processing. This operation can be best implemented in a processor having two-dimensional (2-D) parallelism and a global interaction among the elements of the input vectors. Since optics is endowed with exactly these features, an optical processor can perform the outer product operation in a natural fashion using orthogonally oriented one-dimensional (1-D) input devices such as acoustooptic cells. Algorithms that can be implemented optically using outer-product concepts include matrix multiplication, convolution/correlation, binary arithmetic operations for higher accuracy, matrix decompositions, and similarity transformations of images. Implementation is shown to be frequently tied to time-integrating detection techniques. These and other hardware issues in the implementation of some of these algorithms are discussed.

Bakoglu, H.B. and J.D. Meindl, "Optimal Interconnection Circuits for VLSI," IEEE Transactions of Electron Devices, vol. ED-32, no. 5, p. 903, May 1985.

The propagation delay of interconnection lines is a major factor in determining the performance of VLSI circuits because the RC time delay of these lines increases rapidly as chip size is increased and cross-sectional interconnection dimensions are reduced. In this paper, a model for interconnection time-delay is developed that includes the effects of scaling transistor, interconnections, and chip dimensions. The delays of aluminum, WSi2, and polysilicon lines are compared, and propagation delays in future VLSI circuits are projected. Properly scaled multilevel conductors, repeaters, cascaded drivers, and cascaded driver/repeater combinations are investigated as potential methods for reducing propagation delay. The model yields optimal cross-sectional interconnection dimensions and driver/repeater configurations that can lower propagation delay by more than an order of magnitude in MOSFET circuits.

Ballard, D., "Modular Learning in Neural Networks," Proc. of the AAAI Conf., pp. 279-284, Seattle, WA, August 1987.

In the development of large-scale knowledge networks, much recent progress has been inspired by connections to neurobiology. An important component of any "neural" network is an accompanying learning algorithm. Such an algorithm to be biologically plausible, must work for very large numbers of units. Studies of large-scale systems have so far been restricted to systems without internal units (units with no direct connections to the input or output). Internal units are crucial to such sysstems as they are the means by which a system can encode high-order regularities (or invariants) that are implicit in its inputs and outptus. Computer simulations of learning using internal units have been restricted to small-scale systems. This paper describes a way of coupling autoassociative learning modules into hierarchies that should greatly improve the performance of learning algorithms in large-scale systems. The idea has been tested experimentally with positive results.

Barlow, H.B., D. Rose, and V.G. Dobson, "Cerebral cortex as model builder," in Models of the Visual Cortex, pp. 37-46, Wiley Interscience, New York, 1985.

The cerebral cortex is supposed to be responsible for humanity's dominance of the natural world, and in particular for the intellectual preeminence that underlies this position. With the hope that the function of the whole cortex may illuminate the role of the parts devoted to vision, and vice versa, I have asked the following five questions:
1. Why does the cortex everywhere possess a similar structure?
2. What sort of new behavior does the cortex make possible?
3. What types of operation are required for such new behaviour?
4. Do neurons in visual cortex perform such operations?
5. Can we suggest how the cortex does it?
In this article I attemp to five the simplest and most straightforward answer to each question in turn, and a unified view emerges which makes sense of a wide assortment of facts about the cortex.

Barto, Andrew G., Richard S. Sutton, and Charles W. Anderson, "Neuronlike Adaptive Elements That Can Solve Difficult Learning Control Problems," IEEE Transactions of Systems, Man, And Cybernetics, vol. SMC-13, no. 5, pp. 835-846, September/October 1983.

It is shown how a system consisting of two neuronlike adaptive elements can solve a difficult learning control problem. The task is to balance a pole that is hinged to a movable cart by applying forces to the cart's base. It is assumed that the equations of motion of the cart-pole system are not known and that the only feedback evaluating performance is a failure signal that occurs when the pole falls past a certain angle from the vertical, or the cart reaches an end of a track. This evaluative feedback is of much lower quality than is required by standard adaptive control techniques. It is argued that the learning problems faced by adaptive elements that are components of adaptive networks are at least as difficult as this version of the pole-balancing problem. The learning system consists of a single associative search element (ASE) and a single adaptive critic element (ACE). In the course of learning to balance the pole, the ASE constructs associations between input and output by searching under the influence of reinforcement feedback, and the ACE constructs a more informative evaluation function than reinforcement feedback alone can provide. The differences between this approach and other attempts to solve problems using neuronlike elements are discussed, as is the relation of this work to classical and instrumental conditioning in animal learning studies and its possible implications for research in the neurosciences.

Baum, E.B., J. Moody, and F. Wilczek, "Internal Representations for Associative Memory," Technical Report NSF-ITP-86-138, Institute for Theoretical Physics, Santa Barbara, CA, 1986.

We describe a class of feed forward neural network models for associative content addressable memory (ACAM) which utilize sparse internal representations for stored data. In addition to the input and output layers, our networks incorporate an intermediate processing layer which serves to label each stored memory and to perform error correction and association. We study two classes of internal lable representations: the unary representation and various sparse, distributed representations. Finally, we consider storage of sparse data and sparsification of data. These models are found to ahve advantages in terms of storage capacity, hardware efficiency, and recall reliability when compared to the Hopfield model, and to possess analogies to both biological neural networks and standard digital computer memories.

Cleary, John G., A Simple VLSI Connectionist Architecture, Proceedings of the ICNN, San Diego, CA, June 1987.

A connectionist architecture based on a VLSI chip, the sigma-chip, is described. In it single bit values are communicated between units by heavily multiplexing them. Individual units perform a weighted sum of their inputs and compare the result to a threshold in order to compute the (single bit) output. Potential applications, including rule based reasoning, constraint satisfaction, picture processing, and the game of life are briefly described and their space and time requirements analyzed.

Cohen, Michael A. and Stephen Grossberg, Masking Fields: A Massively Parallel Neural Architecture for Learning, Recognizing, and Predicting Multiple Groupings of Patterned Data, Boston, 1986.

A massively parallel neural network architecture, called a masking field, is characterized through systematic computer simulations. A masking field is a multiple scale, self-similar, automatically gain-controlled cooperative-competitive feedback network F2. Network F2 receives input patterns from an adaptive filter F1 - F2 that is activated by a prior processing level F1. Such a network F2 behaves like a content-addressable memory. It activates compressed recognition codes that are predictive with respect to the activation patterns flickering across the feature detectors of F1, and competitively inhibits, or masks, codes which are unpredictive with respect to the F1 patterns. In particular, a masking field can simultaneously detect multiple groupings within its input patterns and assign activation weights to the codes for these groupings which are predictive with respect to the contextual information embedded within the patterns and the prior learning of the system. A masking field automatically rescales its sensitivity as the overall size of an input pattern changes, yet also remains sensitive to the microstructure within each input pattern. In this way, a masking field can more strongly activate a code for the whole F1 pattern than for its salient parts, yet amplifies the code for a pattern part when it becomes a pattern whole in a new input context. A masking field can also be primed by inputs from F1: it can activate codes which represent predictions of how the F1 pattern may evolve in the subsequent time interval. Network F2 can also exhibit an adaptive sharpening property: repetition of a familiar F1 pattern can tune the adaptive filter to elicit a more focal spatial activation of its F2 recognition code than does an unfamiliar input pattern. The F2 recognition code also becomes less distributed when an input pattern contains more contextual information on which to base an unambiguous prediction of which F1 pattern is being processed. Thus a masking field suggests a solution of the credit assignment problem by embodying a real-time code for the predictive evidence contained within its input patterns. Such capabilities are useful in speech recognition, visual object recognition, and cognitive information processing.

An absolutely stable design for a masking field is disclosed through an analysis of the computer simulations. This design suggests how associative mechanisms, cooperative-competitive interactions, and modulatory gating signals can be joined together to regulate the learning of compressed recognition codes. Data about the neural substrates of learning and memory are compared with these mechanisms.

Cohen, Neal J., Irene T. Abrams, Walter S. Harley, and Lisa Tabor, "Skill Learning and Repetition Priming in Symmetry Detection: Parallel Studies of Human Subjects and Connectionist Models 1," in Eighth Annual Conference of the Cognitive Science Society, pp. 23-44, Lawrence Erlbaum Associates, Hillsdale, August, 1986.

The present paper is a preliminary report of our work exploring skill learning and repetition priming in parallel studies of mirror symmetry detection in humans and network models. The memory mechanisms supporting the acquisition of skill and repetition priming in humans have been the subject of much speculation. On one account, drawing on the distinction between procedural and declarative learning, these learning phenomena grow out of experience-based tuning and reorganization of processing modules engaged by performance in a given domain, in a manner that molules engaged by performance in a given domain, in a manner that is intimately tied to the operation of those modules. Such learning appears similar to that suggested by the incremental learning algorithms currently being explored in massively-parallel connectionist models (e.g., the Boltzmann machine). In the present work, both learning phenomena were observed in the behavioral data from human subjects and the simulation data from the network models. The network models showed priming effects from the start of de novo learning despite being designed to handle generalization to new materials - the essence of skill learning - and without additional mechanisms designed to provide a temporary advantage for recently presented material. Priming occurred for the human subjects despite the use of novel materials for which pre-existing representations cannot already be present in memory. These findings support the notion that skill learning and repetition priming are linked to basic incremental learning mechanisms that serve to configure and reorganize processing modules engaged by experience.

Conwell, Peter R., Effects of Connection Delays in Two State Model Neural Circuits, Proceedings of the ICNN, San Diego, CA, June 1987.

A two state model utilizing a modified version of the McCulloch-Pitts decision rule is presented. Simulations confirm that fixed delays introduce a noise-like component into a units output function. Surprisingly, simulations of the evolution of the network revealed that stable states are

still found. The corresponding output distribution was determined by ensemble averaging. It was found to be similar to distributions introduced explicitly in other probabalistic models. Histograms depicting the distribution while the system is in various stages of convergence are presented. Accordingly, it was discovered that the shape of the output distribution narrowed and became more deterministic as the network approached stability. In effect the network cools as it converges. The shape of the distribution was also found to be dependent on update rate (the time a unit waits before it evaluates its input). Lower rates reflect a narrow, more deterministic distribution. Implications for the application of such networks to optimization problems are discussed.

Cormen, Thomas H. and Charles E. Leiserson , A Hyperconcentrator Switch for Routing Bit-Serial Messages, pp. 721-728, IEEE, 1986 . The authors are located at the Computer Science Lab at MIT

In highly parallel message routing networks, it is sometimes desirable to concentrate relatively few messages on many wires onto fewer wires. We have designed a VLSI chip for this purpose which is capable of concentrating bit-serial messages quickly. This hyperconcentrator switch has a highly regular layout using ratioed nMOS and takes advantage of the relatively fast performance of large fan-in NOR gates in this technology. A signal incurs exactly 2log2n gate delays through the switch, where n is the number of inputs to the circuit. The architecture generalizes to domino CMOS as well.

Cover, T.M., "Nearest Neighbor Pattern Classification," Transactions on Information Theory, vol. IT-13, no. 1, pp. 21-27, IEEE, January, 1967.

The nearest neighbor decision rule assigns to an unclassified sample point the classification of the nearest of a set of previously classified points. This rule is independent of the underlying joint distribution on the sample points and their classifications, and hence the probability of error R of such a rule must be at least as great as the Bayes probability of error R*-the minimum probability of error over all decision rules taking underlying probability structure into account. However, in a large sample analysis, we will show in the M-category case that R* < (or equal to) R < (or equal to) R*(2 - MR*/(M - 1)), where these bounds are the tightest possible, for all suitable smooth underlying distributions. Thus for any number of categories, the probability of error of the nearest neighbor rule is bounded above by twice the Bayes probability of error. In this sense, it may be said that half the classification information in an infinite sample set is contained in the nearest neighbor.

Cover, T. M. and P. E. Hart, "Nearest Neighbor Pattern Classification," IEEE Transactions on Information Theory, vol. IT-13, no. 1, pp. 21-27, January, 1967.

The nearest neighbor decision rule assigns to an unclassified sample point the classification of the nearest of a set of previously classified points. This rule is independent of the underlying joint distribution on the sample points and their classifications, and hence the probability of error R of such a rule must be at least as great as the Bayes probability of error R*--the minimum probability of error over all decision rules taking underlying probabililty structure into account. However, in a large sample analysis, we will show in the M-category case that $R^* < R < R^*(2 - MR^*/(M-1))$, where these bounds are the tightest possible, for all suitably smooth underlying distributions. Thus for any number of categories, the probability of error of the nearest neighbor rule is bounded above by twice the Bayes probability of error. In this sense, it may be said that half the classification information in an infinite sample set is contained in the nearest neighbor.

Cruz-Young, C. A. and J. Y. Tam, "NEP: An Emulation-Assist Processor for Parallel Associative Networks," IBM Palo Alto Scientific Center, June 1985.

This document describes the Network Emulation Processor, or NEP, which is a special-purpose computer designed and built at the Palo Alto Scientific Center. The NEP was created in support of the Parallel Associative Networks project at this Center. It is designed to allow the real-time emulation of large associative networks, a novel information-processing formalism which is briefly described here. The processor performs its task by efficiently executing a discrete-time digital simulation of a network's constituent elements. Each NEP can emulate a network containing 4k nodes and 16k links, with up to 30 network update cycles per second. In addition, up to 256 NEPs can be cascaded, allowing emulation of networks consisting of 1M nodes and 4M links.

Csernai, L. P. and J. Zimanyi, "Mathematical Model for the Self-Organization of Neural Networks," Biological Cybernetics, vol. 34, pp. 43-48, Springer-Verlag, 1979.

Mutual inhibition between neurons combined with a learning principle similar to that proposed by Hebb is shown to secure a powerful self-organizing property for neural networks. Numerical analysis reveals that the system investigated always organizes itself into the same final state from any arbitrarily chosen initial state.

Dally, William J., "A VLSI Architecture for Concurrent Data Structures," Thesis, Pasadena, CA, 1986.

Concurrent data structures simplify the development of concurrent programs by encapsulating commonly used mechanisms for synchronization and communication into data structures. This thesis developes a notation for describing concurrent data structures, presents examples of concurrent data structures, and describes an architecture to support concurrent data structures. Concurrent Smalltalk (CST), a derivative of Smalltalk-80 with extensions for concurrency, is developed to describe concurrent data structures. CST allows the programmer to specify objects that are distributed over the nodes of a concurrent computer. These distributed objects have many *constituent objects* and thus can process many messages simultaneously. They are the foundation upon which concurrent data structures are built. The *balanced cube* is a concurrent data structure for ordered sets. The set is distributed by a balanced recursive partition that maps to the subcubes of a binary *n*-cube using a Gray code. A search algorithm, VW search, based on the distance properties of the Gray code, searches a balanced cube in *O*(log *N*) time. Because it does not have the root bottleneck that limits all tree-based data structures to *O*(1) concurrency, the balanced cube achieves *O*(Considering graphs as concurrent data structures, graph algorithms are presented for the shortest path problem, the max-flow problem, and graph partitioning. These algorithms introduce new synchronization techniques to achieve better performance than existing algorithms. A message-passing, concurrent architecture is developed that exploits the characteristics of VLSI technology to support concurrent data structures. Interconnection topologies are compared on the basis of dimension. It is shown that minimum latency is achieved with a very low dimensional network. A deadlock-free routing strategy is developed for this class of networks, and a prototype VLSI chip implementing this strategy is described. A message-driven processor complements the network by responding to messages with a very low latency. The processor directly executes messages, eliminating a level of interpretation. To take advantage of the performance offered by specialization while at the same time retaining flexibility, processing elements can be specialized to operate on a single class of objects. These *object experts* accelerate the performance of all applications using this class.

Dell, Gary S., "A Spreading-Activation Theory of Retrieval in Sentence Production," *American Psychological Association, Inc.*, vol. 93, no. 3, pp. 283-321, 1986.

This article presents a theory of sentence production that accounts for facts about speech errors--the kinds of errors that occur, the constraints on their form, and the conditions that precipitate them. The theory combines a spreading-activation retrieval mechanism with assumptions regarding linguistic units and rules. Two simulation models are presented to illustrate how the theory applies to phonological encoding processes. One was designed to produce the basic kinds of phonological errors and their relative frequencies of occurrence. The second was used to fit data from an experimental technique designed to create these errors under controlled conditions.

Derthick, Mark and David C. Plaut, "Is Distributed Connectionism Compatible with the Physical Symbol System Hypothesis?," in Program of the Eighth Annual Conference of the Cognitive Science Society, pp. 639-644, Lawrence Erlbaum Associates, Hillsdale, NJ, August, 1986.

All existing intelligent systems share a similar biological and evolutionary heritage. Based on the conviction that cognition is computation, artificial intelligence researchers are investigating computational models as a means of discovering properties shared by all intelligent systems.

One property that has been proposed as central to intelligence is the ability to construct and manipulate symbol structures. If intelligence may be described completely in terms of symbol processing, then cognitive science need not be concerned with the particular physical implementation details of either artificial or biological examples; neuroscience would no longer be part of cognitive science. On the other hand, if important aspects of intelligence evade symbolic explanation, it may prove necessary to consider phenomena below the symbol level. The connectionist approach to artificial intelligence is founded on the conviction that the structure of the brain critically constrains the nature of the computations it performs. However, if the symbolic position is correct and neural networks only implement symbol systems, then connectionism contributes little to cognitive science. The notion of intelligence as symbol processing was made explicit by Newell and Simon with the Physical Symbol System Hypothesis (PSSH) (Newell & Simon, 1976, Newell, 1980) and the Knowledge Level Hypothesis (KLH) (Newell, 1982). Taken together, these hypotheses have significant implications for the nature of any system capable of general intelligence. We examine a number of connectionist systems in light of the hypotheses and distinguish three kinds: (1) rule-based systems, which are symbol systems; (2) rule-following systems, which are symbol systems only under a weakened version of the PSSH; and (3) systems which are not rule-following, and thus are not symbol systems even in a weak sense. According to the PSSH, non-symbolic connectionist systems must be incapable of general intelligence. There are strong arguments both for and against this conclusion. On the one hand, such connectionist systems may provide more parsimonious accounts of certain cognitive phenomena than do symbolic approaches. On the other hand, these connectionist systems have significant limitations, relating to universality, not shared by symbol systems. We conclude that a comprehensive theory of intel-

ligence may require a hybrid model that combines the strengths of both approaches.

Desimone, R., S.J. Schein, J. Moran, and L.G. Ungerleider, "Contour, Color and Shape Analysis Beyond the Striate Cortex," Vision Res., vol. 25, no. 3, pp. 441-452, Pergamon Press Ltd., Great Britain, 1985.

The corticocortical pathway from striate cortex into the temporal lobe plays a crucial role in the visual recognition of objects. Anatomical studies indicate that this pathway is mainly organized as a serial hierarchy of multiple visual areas, including V1, V2, V3, V4, and inferior temporal cortex (IT). As expected from the anatomy, we have found that neurons in V4 and IT, like those in V1 and V2, are sensitive to many kinds of information relevant to object recognition. In the spatial domain, many V4 cells exhibit length, width, orientation, direction of motion and spatial frequency selectivity. In the spectral domain, many V4 cells are also tuned to wavelength. Thus, V4 is not specialized to analyze one particular attribute of a visual stimulus; rather, V4 appears to process both spatial and spectral information in parallel. A special contribution of V4 neurons to visual processing may lie in specific spatial and spectral interactions between their small excitatory receptive fields and large silent suppressive surrounds. Thus, although the excitatory receptive fields of V4 neurons are small, the responses of V4 neurons are influenced by stimuli throughout a much larger portion of the visual field. In IT, neurons also appear to process both spatial and spectral information throughout a large portion of the visual field. However, unlike V4 neurons, the excitatory receptive fields of IT neurons are very large. Many IT neurons, for example, are selective for the overall shape, color, or texture of a stimulus, anywhere within the central visual field. Together these qualities are processed in parallel, but the type of analysis may become more global at each stage of processing.

Don, Hon-Son and King-Sun Fu, "A Parallel Algorithm for Stochastic Image Segmentation," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. PAMI-8, no. 5, pp. 594-602, September, 1986.

A parallel algorithm for syntactic image segmentation is introduced. Stochastic tree grammar is used as a context-generating model. It is shown that when this context-generating process is in the equilibrium state, a matched filter can be designed and applied in parallel to the image. This process can be used for image segmentation in a syntactic pattern recognition system to enhance the performance of the succeeding recognition process.

Easton, Paul and Peter E. Gordon, "Stabilization of Hebbian Neural Nets by Inhibitory Learning," Biological Cybernetics, vol. 51, pp. 1-9, Springer-Verlag, 1984.

In Hebbian neural models synaptic reinforcement occurs when the pre- and post-synaptic neurons are simultaneously active. This causes an instability toward unlimited growth of excitatory synapses. The system can be stabilized by recurrent inhibition via modifiable inhibitory synapses. When this process is included, it is possible to dispense with the non-linear normalization or cut-off conditions which were necessary for stability in previous models. The present formulation is response-linear if synaptic changes are slow. It is self-consistent because the stabilizing effects will tend to keep most neural activity in the middle range, where neural response is approximately linear. The linearized equations are tensor invariant under a class of rotations of the state space. Using this, the response to stimulation may be derived as a set of independent modes of activity distributed over the net, which may be identified with cell assemblies. A continuously infinite set of equivalent solutions exists.

El-Leithy, N., R. W. Newcomb, and M. Zaghioul, "A Basic MOS Neural-Type Junction / A Perspective on Neural-Type Microsystems," Proceedings of the ICNN, San Diego, CA, June 1986.

This paper presents a perspective on neural-type electronics as a novel device-circuit design approach inspired by neural structures. Such an approach attempts to provide a well-engineered, theory-tested substrate on which the emerging computational, adaptive, learning, massively-parallel architectures can be built. Further, the paper proposes an MOS neural-type junction as an example of the adopted design philosophy. The proposed circuit captures the temporal and spatial activity of neural elements through N input structures. Each of these structures comprises an initial adaptive weighting stage feeding into a charge-accumulating system. The net excitatory drive resulting from space and time effects stimulates a dynamic firing output stage. The circuit permits the implementation of information processing elements in which pulse frequency modulation is the coding mechanism. Some modifications to the design, implications and future prospects are outlined.

Farmer, J.D., S.A. Kauffman, and N.H. Packard, "Autocatalytic replication of peptides," Phys. Briefs, vol. 8, no. 23, 1986. Presented at Conference on Evolution, Games, and Learning: Models for Adaptation in Machines and Nature, Los Alamos NM, 20-24 May 1985.

We construct a simplified model for the chemistry of molecules such as polypeptides or single stranded nucleic acids, whose reactions can be restricted to catalyzed cleavage and condensation. We use this model to study the spontaneous emergence of autocatalytic sets from an initial set of simple building blocks, for example short strands of amino acids or nucleotides. When the initial set exceeds a critical diversity, autocatalytic reactions generate large molecular species in abundance. Our results suggest that the critical diversity is not very large. Autocatalytic sets formed in this way can be regarded as primitive connected metabolisms, in which particular species are selected if their chemical properties are advantageous for the metabolism. Such autocatalytic sets may have played a crucial role in the origin of life, providing a bridge from simple molecular species to complex proteins and nucleic acids. Many of our results are experimentally testable.

Gabor, D., Theory of Communication, British Thomson-Houston Co., Ltd., Research Laboratory, 1945.

The purpose of these three studies is an inquiry into the essence of the "information" conveyed by channels of communication and the application of the results of this inquiry to the practical problem of optimum utilization of frequency bands.

In Part I, a new method of analysing signals is presented in which time and frequency play symmetrical parts, and which contains "time analysis" and "frequency analysis" as special cases. It is shown that the information conveyed by a frequency band in a given time-interval can be analyzed in various ways into the same number of elementary "quanta of information," each quantum conveying one numerical datum.

In Part II, this method is applied to the analysis of hearing sensations. It is shown on the basis of existing experimental material that in the band between 60 and 1,000 c/s that this efficiency of nearly 50% is independent of the duration of the signals in a remarkably wide interval. This fact, which cannot be explained by any mechanism in the inner ear, suggests a new phenomenon in nerve conduction. At frequencies above 1,000 c/s the efficiency of discrimination falls off sharply, proving that sound reproductions which are far from faithful may be perceived by the ear as perfect, and that "condensed" methods of transmission and reproduction with improved waveband economy are possible in principle.

In Part 3, suggestions are discussed for compressed transmission and reproduction of speech or music, and the first experimental result obtained with one of these methods are described.

Gallant, Stephen I. and Donald Smith, Random Cells: An Idea Whose Time Has Come and Gone ... And Come Again?, College of Computer Science, Northeastern Univ., Boston, MA, January 1987.

Frank Rosenblatt proposed using random functions in a neural

(or connectionist) network in order to improve the learning capability of other cells. This proposal was not successful, for reasons reviewed here. A modification of Rosenblatt's idea is to: 1)use random linear discriminants, each of which is connected to all inputs, for fixed cells, and 2) use the pocket algorithm, a well-behaved modification of perceptron learning, to generate coefficients for cells that learn. The main idea is to have a sufficiently rich distributed representation in the activations of the random discriminants rather than having individual random cells compute individual features. To emphasize this distinction, the procedure has been named the Distributed Method. This revised approach overcomes or sidesteps the main theoretical problems with Rosenblatt's original proposal.

Experimental results for some of the most difficult classical learning problems are presented. The Distributed Method appears to be a practical way to construct network models for complex problems.

Gallant, Stephen I., Sequential Associative Memories, College of Computer Science, Northeastern Univ., Boston, MA, March 1987.

Humans are very good at manipulating sequential information, but sequences present special problems for connectionist models. This paper examines subnetworks of connectionist models called Sequential Associative Memories (SAM's) that have feedback and high connectivity. Gallant The coefficients for SAM cells are unmodifiable and are generated at random. SAM's serve two functions.

1. Their activations determine a state for the network which permits previous inputs and outputs to be recalled, and

2. They increase the dimensionality of input and output representations to make it possible for other (modifiable) cells in the network to learn difficult tasks.

The second function is similar to what has been called the distributed method, a way of generating intermediate cells for problems not involving sequences. To illustrate the underlying difficulties of learning various tasks we examine several ways of learning to add, the most difficult of which serves as an indication of how well SAM's can manipulate general sequences. The question of network macrostructure is also addressed. Our conclusion is that SAM's represent a promising approach for sequence manipulation by connectionist models.

Gallant, Stephen I., Automated Generation of Connectionist Expert Systems for Problems Involving Noise and Redundancy, College of Computer Science, Northeastern Univ., Boston, MA, March 1987.

When creating an expert system the most difficult and expensive task is constructing a knowledge base. This is particularly true if the problem involves noisy data and redundant measurements.

This paper shows how to modify the MACIE process for generating connectionist expert systems from training examples so that it can accommodate noisy and redundant data. The basic idea is to dynamically generate appropriate training examples by constructing both a `deep' model and a noise model for the underlying problem. The use of winner-take-all groups of variables is also discussed.

These techniques are illustrated with a small example that would be very difficult for standard expert system approaches.

Geller, Ronnie, "A VLSI Architecture for a Neurocomputer using Higher-Order Predicates.," Tech. Report CS/E-87-004, Dept. of Computer Science/Engineering, Oregon Graduate Center, Beaverton, Oregon, May 1987.

Some biological aspects of neural interactions are presented and used as a basis for a computational model in the development of a new type of computer architecture. A VLSI microarchitecture is proposed that efficiently implements the neural-based computing methods. An analysis of the microarchitecture is presented to show that it is feasible using currently available VLSI technology. The performance expectations of the proposed system are analyzed and compared to conventional computer systems executing similar algorithms. The proposed system is shown to have comparatively attractive performance and cost/performance ratio characteristics. Some discussion is given on system level characteristics including initialization and learning.

Genis, C. Torras i, "Neural Network Model with Rhythm-Assimilation Capacity," IEEE Transactions on Systems, Man, and Cybernetics, vol. 16, no. 5, pp. 680-693, Sept/Oct 1986.

Assimilation of a stimulus rhythm by certain nervous structures of vertebrate animals in a conditioning situation has been reported. A lateral-inhibition-type network of plastic pacemaker neurons is proposed to model such behavior. Through simulation, the network exhibits capacity to assimilate and encode in separate groups of neurons two successively presented frequencies. Learning of the second frequency does not disrupt memory of the first one. From the exploration of the effect of varying several factors upon the learning process - related to the connectivity, the intraneuronal functioning, the initial state, and the simulation conditions - it follows that the most influential factors are the proportion of excitatory connections over the total, the ratio between the ranges of excitatory and the inhibitory connectivity, and the degree of intraneuronal randomness.

Gluck, Mark A. and Richard F. Thompson, "Modeling the Neural Substrates of Associative Learning and Memory: A Computational Approach," Psychological Review, vol. 94, no. 2, pp. 176-191, Stanford, CA, 1987.

We develop a computational model of the neural substrates of elementary associative learning, using the neural circuits known to govern classical conditioning of the gill-withdrawal response of Aplysia. Building upon the theoretical efforts of Hawkins and Kandel (1984), we use this model to demonstrate that several higher order features of classical conditioning could be elaborations of the known cellular mechanisms for simple associative learning. Indeed, the current circuit model robustly exhibits many of the basic phenomena of classical conditioning. The model, however, requires a further assumption (regarding the form of the acquisition function) to predict asymptotic blocking and contingency learning. In addition, if extinction if mediated by the nonassociative mechanism of habituation--rather than the associative process postulated by Rescorla and Wagner (1972)--then we argue that additional mechanisms must be specified to resolve a conflict between acquisition and maintenance of learned associations. We suggest several possible extensions to the circuit model at both the cellular and molecular levels that are consistent with the known Aplysia physiology and that could, in principle, generate classical conditioning behavior.

Golden, R.M., A Probabilistic Computational Framework for Neural Network Models, Learning and Development Center, University of Pittsburgh, Pittsburgh, PA. Submitted for Publication July, 1987

Information retrieval in a "connectionist" or neural network is viewed as computing the most probable value of the information to be retrieved with respect to the probability density function, P. With a minimal number of assumptions, the "energy" function that a neural network minimizes during information retrieval is shown to uniquely specify P. Inspection of the form of P indicates the class of probabilistic environments that can be learned. Learning algorithms can be analyzed and designed by using maximum likelihood estimation techniques to estimate the parameters of P. The large class of nonlinear auto-associateive networks analyzed by Cohen and Grossberg, nonlinear associative multi-layer back-propagation networks, and certain classes of nonlinear multi-stage networks are analyzed within the proposed computational framework.

Golden, Ricahrd M., "Representing Casual Schemata in Connectionist Systems," in Eighth Annual Conference of the Cognitive Science Society, pp. 13-22, Lawrence Erlbaum Associates, Hillsdale, August, 1986.

The connectionist approach to human memory is based upon the idea that knowledge can be stored implicitly in the form of real-valued interconnections among a set of simple "neuron-like" computing elements (Hinton & Anderson, 1981). The schema system approach (Rumelhart, 1980; Schank & Abelson, 1977) considers human memory to be organized in terms of many small packets of knowledge called schemata. If a knowledge packet is defined as some sequence of causally related events, then it is referred to as a "casual schema" or "script". Although these two seemingly different approaches to the problem of modelling human memory might seem incompatible, they are actually intimately related (Rumelhart, Smolensky, McClelland & Hinton, 1986; also see Touretzky & Hinton, 1985). In this paper, a conectionist model of how causal schemata are used in the recall of actions from simple stories is described. The paper is organized in the following manner. First, an explicit procedure for representing complex causal schemata as "neural activation patterns" is discussed in detail. Next, the fundamental neural mechanisms that are used to process and learn information are described and motivated from a probabilistic viewpoint. Finally, the resulting system is used to model some experimental data obtained by Bower, Black and Turner (1979) in their studies of human memory for written test.

Golden, Richard M., "A Developmental Neural Model of Visual Word Perception," Cognitive Science 10, pp. 241-276, 1986.

A neurally plausible model of how the process of visually perceiving a letter in the context of a word is learned, and how such processing occurs in adults is proposed. The model consists of a collection of abstract letter feature detector neurons and their interconnections. The model also includes a learning rule that specifies how these interconnections evolve with experience. The interconnections between neurons can be interpreted as representing the spatially redundant, sequentially redundant, and transgraphemic information in letter string displays. Anderson, Silverstein, Ritz, and Jones's (1977) "Brain-State-in-a-Box" (BSB) neural mechanism is then used to implement the proposed model. The resulting system makes explicit qualitative predictions using both letter recognition accuracy and reaction time as dependent measures. In particular, the model offers an integrated explanation of some experiments involving manipulations of orthographic regularity, masking, case alterations, and experience with words. The similarities and differences between the model and models proposed by Adams (1979) and McClelland and Rumelhart (1981) are also discussed.

Graf, Hans P. and Paul de Vegvar, "FAM 22.1: A CMOS Associative Memory Chip Based on Neural Networks," IEEE International Solid-State Circuits Conference, pp. 304-305,437, 1987.

None

Gray, Paul R., David A. Hodges, and Robert W. Brodersen, Analog MOS Integrated Circuits, IEEE Selected Reprint Series, IEEE Press, New York, NY, 1980.

Grossberg, S., The Adaptive Brain (I and II), Elsevier/North-Holland, Amsterdam, 1986.

Grossberg, Stephen, The Adaptive Brain I, Advances in Psychology, 42, North Holland, Amsterdam, The Netherlands, 1987.

Grossberg, Stephen, The Adaptive Brain II, Advances in Psychology, 43, North Holland, Amsterdam, The Netherlands, 1987.

Gullichsen, Eric and Ernest Chang, Pattern Classification by Neural Network: An Experimental System for Icon Recognition, Proceedings of the ICNN, San Diego, CA, June 1987.

Recent advantages in the field of connectionism have led to a resurgence in popularity of neural models of computation [Rumelhart86]. Connectionism represents a "low level" (as opposed to a symbolic) approach to the problems of Artificial Intelligence, with knowledge represented globally in the real-valued strengths of connections between independent neuron-line processing elements, and not in discrete statements in a symbolic language.
We present information about a connectionist system for the recognition of handprinted icons on hydrographic charts, implemented in APL on a Sun workstation. A layered back-propagation network which learns to correctly classify icons constitutes the core of the system. Discussed are: icon propagation nets which proved useful in practice. A mechanism for the visualization of the extent of network learning, useful in a wide range of pattern recognition tasks, is discussed.

Hammerstrom, D., "A Connectionist/Neural Network Bibliography," Tech. Report CS/E-86-010, Dept. of Computer Science/Engineering, Oregon Graduate Center, Beaverton, Oregon, August 1986.

This technical report contains a bibliography of connectionist and neural modelling papers and books as well as related in the areas of computer architecture and brain science. Abstracts are included for many of the references.

Harth, E. M., T. J. Csermely, B. Beek, and R. D. Lindsay, Brain Functions and Neural Dynamics, pp. 93-120, 1970.

Anatomical and physiological evidence is cited for the existence in the CNS of more or less discrete populations of interconnected neurons. These are given the term netlets. A model based on these observations is presented, in which it is assumed that the netlets are the fundamental building blocks out of which nets of considerable complexity may be

assembled. The connectivity within each netlet is assumed to be random. Neuronal macrostates are defined in which the fractions of neurons active in each netlet are the dynamical variables. Thus the temporal and spatial fine structure of neuronal activity are considered to be of secondary significance and are disregarded. These assumptions bring about an enormous reduction in complexity. Thus calculations and computer simulation studies become possible for systems hitherto inaccessible to quantitative description. It is hoped that the features retained in the model play a sufficiently significant role in the functioning of real neural nets to make these results meaningful. The mathematical formalism and detailed numerical results appear in another paper of this issue (Anninos, 1970). Some of these results are anticipated in this paper and their implications for our model are discussed. The study proceeds from a treatment of isolated probabilistic netlets to the dynamics of interacting netlets. Of particular interest are the conditions under which a netlet will go into sustained activity and the often extremely delicate control exerted by afferent excitatory or inhibitory biases. Hysteresis effects are common and may represent a type of short-term memory. A variety of neural functions are listed to which some of these mechanisms may be applied. Among these are the modulating effects of the brain stem reticular formation on cortical and spinal neuron populations and the "energizing" of cortical centers by spontaneous activity in sensory systems. Finally the concepts of netlet interaction are applied in conjunction with the principle of synaptic facilitation to information processing in the cortex. Examples given are sensory-sensory cortical conditioning and the formation of the classical conditioned reflex.

Hecht-Nielson, Robert, "Performance Limits of Optical, Electro-Optical, and Electronic Neurocomputers," TRW Rancho Carmel AI Center, San Diego.

The performance limits of optical, electro-optical, and electronic artificial neural systems (ANS) processors (also known as neurocomputers) are discussed. After a brief introduction, an overview is provided of the recently revived field of ANS. Next, ANS performance measures are defined and a neurocomputer taxonomy is presented. Finally, the designs and performance limits of the various types of neurocomputers are discussed.

Hecht-Nielson, Robert, "Nearest Matched Filter Classification of Spatiotemporal Patterns," TRW Rancho Carmel AI Center, San Diego, 6/23/86.

Recent advances in massively parallel optical and electronic neural network processing technology have made is plausible to consider the use of matched filter banks containing large numbers of individual filters as pattern classifiers for

complex spatiotemporal pattern environments such as speech, sonar, radar, and advanced communications. This paper begins with an overview of how neural networks can be used to approximately implement such multidimensional matched filter banks. The "nearest matched filter" classifier is then formally defined. This definition is then reformulated to show that the classifier is equivalent to a nearest neighbor classifier in a separable infinite-dimensional metric space that specifies the local-in- time behavior of spatiotemporal patterns. The result of Cover and Hart is then applied to show that, given a statistically comprehensive set of filter templates, the nearest matched filter classifier will have near-Bayesian performance for spatiotemporal patterns. The combination of near-Bayesian classifier performance with the excellent performance of matched filtering in noise yields a powerful new classification technique. This result adds additional interest to Grossberg's hypothesis that the mammalian cerebral cortex carries out local-in-time nearest matched filter classification of both auditory and visual sensory inputs as an initial step in sensory pattern recognition - which may help explain the almost instantaneous pattern recognition capabilities of animals.

Hinton, Geoffrey E., "Connectionist Learning Procedures," CMU Tech Report #CMU-CS-87-115, Carnegie-Mellon University, Pittsburgh, PA, 1987.

A major goal of research on networks of neuron-like processing units is to discover efficient learning procedures that allow these networks to construct complex internal representations of their environment. The learning procedures must be capable of modifying the connection strengths in such a way that internal units which are not part of the input or output come to represent important features of the task domain. Several interesting gradient-descent procedures have recently been discovered. Each connection computes the derivative, with respect to the connection strength, of a global measure of the error in the performance of the the network. The strength is then adjusted in the direction that decreases the error. These relatively simple gradient-descent learning procedures work well for small tasks and the new challenge is to find ways of improving the speed of learning so that they can be applied to larger, more realistic tasks.

Hinton, Geoffrey E. and David C. Plaut, "Using Fast Weights to Deblur Old Memories," Proceedings of the 9th Cognitive Science Society Meeting, pp. 177-194, Seattle, Washington, July 1987.

Connectionist models usually have a single weight on each connection. Some interesting new properties emerge if each connection has two weights: A slowly changing, plastic weight which stores long-term knowledge and a fast-changing,

elastic weight which stores temporary knowledge and spontaneously decays towards zero. If a network learns a set of associations and then these associations are "blurred" by subsequent learning, all the original associations can be "deblurred" by rehearsing on just a few of them. The rehearsal allows the fast weights to take on values that temporarily cancel out the changes in the slow weights caused by the subsequent learning.

Hirai, Yuzo, "A New Hypothesis for Synaptic Modification: An Interactive Process between Postsynaptic Competition and Presynaptic Regulation," Biological Cybernetics, vol. 36, pp. 41-50, Springer-Verlag, 1980.

Based on recent physiological observations, a new hypothesis for the algorithm for synaptic modification is proposed: an interactive process between postsynaptic competition and presynaptic regulation controls the synaptic modification, and both functions obey the same competitive rule. A fundamental algorithm for the competitive rule is proposed and analysed. The interactive process is decomposed into two procedures. The first one is a postsynaptic competition, in which all synapses making contact with the same postsynaptic cell compete for their terminal sites, and the postsynaptic cell makes demands for synaptic modification upon the presynaptic cells. The second one is a presynaptic regulation, in which the presynaptic cell regulates the postsynaptic demands and determines the net amount of the modification for each synapse originating from it. The performance of the algorithm has been simulated on a digital computer. The initial structure of the simulated neural network consists of randomly-connected presynaptic and postsynaptic cells, with no other structures assumed. After repetitive presentations of a uniform stimulus, the basic configurations of neural connections (one-to-one, one-to-many, many-to-one and many-to-many connections) are organized in dependence on the parameters of the algorithm.

Hunsperger, R.G., Integrated Optics: Theory and Technology, Springer Series in Optical Sciences, 33, Springer-Verlag, Berlin, Germany, 1985.

Hwang, Kai and Joydeep Ghosh, "Hypernets for Parallel Processing with Connectionist Architectures," Technical Report CRI-87-03, Los Angeles, CA, 1/30/87.

A new class of modular networks is proposed for hierarchically constructing massively parallel multicomputer systems for distributed supercomputing and AI applications. These networks are called hypernets. They are constructed incrementally with identical cubelets, treelets and buslets that are well suited for VLSI implementation. Hypernets integrate positive features of both hypercubes and binary trees and maintain a constant node degree when the network

size is increased. This paper presents the principles of constructing hypernets and analyzes their architectural potentials in terms of message or marker routing complexity, cost-effective support for global as well as localized communication, I/O capabilities and fault tolerance. Several numeric and AI algorithms are mapped onto hypernets to illustrate their abilities in emulating hypercube connections while using less hardware, as well as in supporting parallel processing in a hierarchically structured or data-dependent environment. The potential of hypernets for efficient support of connectionist models of computation is also explored.

Jordan, Michael I., "Attractor Dynamics and Parallelism in a Connectionist Sequential Machine," in _Program of the Eighth Annual Conference of the Cognitive Science Society_, pp. 531-546, Lawrence Erlbaum Associates, Hillsdale, NJ, August, 1986.

Fluent human sequential behavior, such as that observed in speech production, is characterized by a high degree of parallelism, fuzzy boundaries, and insensitivity to perturbations. In this paper, I consider a theoretical treatment of sequential behavior which is based on data from speech production. A network is discussed which is essentially a sequential machine built out out of connectionist components. The network relies on distributed representations and a high degree of parallelism at the level of the component processing units. These properties lead to parallelism at the level of the component processing units. These properties lead to parallelism at the level at which whole output vectors arise, and constraints must be imposed to make the performance of the network more sequential. The sequential trajectories that are realized by the network have dynamic properties that are analogous to those observed in networks with point attractors (Hopefield, 1982): learned trajectories generalize, and attractors such as limit cycles can arise.

Judd, J. Stephen, Complexity of Connectionist Learning with Various Node Functions, Dept. of Comp. and Inf. Sci., Univ. of Massachusetts, Amherst, MA, 1987.

We formalize a notion of learning in connectionist networks that characterizes the training of feed-forward networks. Considering different families of node functions, we prove the learning problem NP-complete and thus demonstrate that it has no efficient general solution. One family of node functions studied is the set of logistic-linear functions, as used by the popular back-propagation algorithm. Several implications of the theorem are discussed, including why this result is actually helpful for connectionist learning research.

Kahan, Simon, Theo Pavlidis, and Henry S. Baird, "On the Recognition of Printed Characters of Any Font and Size," Transactions on Pattern Analysis and Machine Intelligence, vol. PAMI-9, no. 2, pp. 274-287, IEEE, 1987.

We describe the current state of a system that recognizes printed text of various fonts and sizes for the Roman alphabet. The system combines several techniques in order to improve the overall recognition rate. Thinning and shape extraction are performed directly on a graph of the run-length encoding of a binary image. The resulting strokes and other shapes are mapped, using a shape-clustering approach, into binary features which are then fed into a statistical Bayesian classifier. Large-scale trials have shown better than 97 percent top choice correct performance on mixtures of six dissimilar fonts, and over 99 percent on most single fonts, over a range of point sizes. Certain remaining confusion classes are disambiguated through contour analysis, and characters suspected of being merged are broken and reclassified. Finally, layout and linguistic context are applied. The results are illustrated by sample pages.

Kandel, E.R. and J.H. Schwartz, in Principles of Neural Science, Elsevier/North-Holland, New York, 1981.

Kauffman, S.A. and R.G. Smith, "Adaptive automata based on Darwinian selection," Phys. Briefs, vol. 8, 1986. Presented at Conference on Evolution, Games, and Learning: Models for Adaptation in Machines and Nature, Los Alamos NM, 20-24 May 1985.

The principle of natural selection in general, and it is natural to assess its implications for achieving automata with desired dynamical or structural properties. The following issues arise: 1) Appropriate definition of the analogue of "genotype" and "phenotype." 2) Definition of the ensemble of "possible" automata in which mutational search for desired properties is occurring. 3) Kinematic properties of the "fitness landscape" in the ensemble governing the statistical features of connected walks through filter variants. 4) Optimal mutation selection strategies given a particular landscape. We assess these questions in two ensembles of automata under selection for one attractor which matches a predetermined "target pattern." The results are: 1) The continuity of desired dynamical properties differs in the two ensembles. 2) When the best automaton seeds each generation, selection follows a characteristic curve, and asymtotes at automata which approach but fail to achieve the desired attractors. 3) Designation of a subset of the variables as hidden from fitness estimation as part of the target pattern does not improve approach to the target pattern. We discuss limitation in the capacity of mutation selection procedures, and approaches to overcoming them.

Klopf, A. H., "A Drive-Reinforcement Model of Single Neuron Function: An Alternative to the Hebbian Neuronal Model," in Neural Networks for Computing, ed. J.S. Denker, pp. 265-269, American Institute of Physics, Snowbird, UT, 1986.

A neuronal learning mechanism is porposed that accounts for the basic animal learning phenomena that have been observed. Among the classical conditioning phenomena predicted by the neuronal model are delay conditioning, trace conditioning, and simultaneous conditioning, stimulus amplitude effects, interstimulus interval effects, second and higher order conditioning, conditioned inhibition, habituation and extinction, reacquisition effects, backward conditioning, blocking overshadowing and serial compound conditioning. The proposed neuronal model and learning mechanisms offer a new building block for constructing neural network-like computer architectures for artificial intelligence.

Koch, C., T. Poggio, and V. Torre, S. Amari, and M. Arbib, "Micronetworks in Nerve Cells," in Competition and Cooperation in Neural Nets, Lecture Notes in Biomathematics, vol. 45, pp. 105-110, Springer-Verlag, Berlin, 1982. Chapter 6

A common belief is that the dendritic tree of neurons whose average length (l) is less than its estimated electronic space constant (L) is virtually equipotential. As a consequence, the geometry of the dendritic tree just merely reflects the spatial convergence of synaptic inputs. The above reasoning is strictly correct for a single cylindrical cable. In branched structures, even those satisfying the equivalent cylinder condition, the argument can be totally wrong. We have analyzed the electrical properties of the different types of retinal ganglion cells in the cat retina on the basis of passive cable theory. It is concluded that these neurons need not be equipotential despite their small dimensions, mainly because of their extensive branching. It is suggested that their dendritic architecture reflects characteristically different electrical properties, which are likely to be relevant for their physiological function and their information processing role.

Kohonen, T. and K. Makisara, "Representation of Sensory Information in Self-Organizing Feature Maps," in Neural Networks for Computing, ed. J.S. Denker, pp. 271-276, American Institute of Physics, Snowbird, UT, 1986.

This paper shows how the internal representations necessary for memorization can be automatically formed in neural networks or artificial systems.

Lapedes, Alan and Robert Farber, "Programming a Massively Parallel, Computation Universal System: Static Behavior," Technical Report LA-UR-86-1179, Los Alamos, NM, 3/4/86.

Massively parallel systems are presently the focus of intense interest for a variety of reasons. A key problem is how to control, or "program" these systems. In previous work by the authors, the "optimum finding" properties of Hopfield neural nets were applied to the nets themselves to create a "neural compiler." This was done in such a way that the problem of programming the attractors of one neural net (called the Slave net) was expressed as an optimization problem that was in turn solved by a second neural net (the Master net). The procedure is effective and efficient. In this series of papers we extend that approach to programming nets that contain interneurons (sometimes called "hidden neurons"), and thus we deal with nets capable of universal computation. Our work is closely related to recent work of Rummelhart et al. (also Parker, and LeChun), which may be viewed as a special case of this formalism and therefore of "computing with attractors." In later papers in this series, we present the theory for programming time dependent behavior, and consider practical implementations. One may expect numerous applications in view of the computation universality of these networks.

Legendy, Charles R., "Mathematical Link Between the width of Cortical Columns and the Range of Intracortical Inhibition," in <u>Models of the Visual Cortex</u>, ed. D. Rose and V. G. Dobson, pp. 473-478, John Wiley & Sons Ltd, 1985.

A little over ten years ago, von der Malsburg (1973) discovered, in computer simulation, that a straightforward set of assumptions leads to the creation of a column-like arrangement of preferred trigger feature orientations in a network of idealized striate cortex neurons. Subsequently I have succeeded in formulating the corresponding problem for mathematical treatment and solving it analytically (Legendy, 1978). Obtaining an analytic solution in this case has the advantage over computer simulation that it yields formulae, indicating what depends on what and how, and that it gives us a hint as to what is the simplest set of assumptions under which a certain result appears. The computer only gives us numbers and graphs, without comment. In this chapter the resulting model is described without the mathematics, its properties listed and the experimental evidence discussed.

Levine, Daniel S., "A Neural Network Theory of Frontal Lobe Function," in <u>Eighth Annual Conference of the Cognitive Science Society</u>, pp. 716-727, Lawrence Erlbaum Associates, Hillsdale, August, 1986.

The frontal cortex is six-layered only in primates and is the neocortical area best connected to the hypothalamus. For these reasons and many others, Fuster (1980, p. 144) stated: "The central notion...is that the prefrontal cortex plays a role in the temporal structuring of behavior. The prefrontal cortex is thought to be essential for the synthesis of cognitive and motor acts into purposive sequences." This article attempts to integrate this qualitative notion with existing neural network theories of motivation and cognition. Grossberg (1975) discusses the striving for balance between two subsystems in a network. The attentional system seeks stable response to fluctuating sensory cues by focusing attention on important subclasses of cues. The arousal system enables adaptation to unexpected events and new reinforcement contingencies. Frontal lesions often change the balance between attention and arousal.

Linsker, Ralph, "From Basic Network Principles to Neural Architecture: Emergence of Orientation-selective Cells," Proc. Natl. Acad. Sci. USA, vol. 83, pp. 8309-8394, November 1986.

This paper is number 2 in a series.

Linsker, Ralph, "From Basic Network Principles to Neural Architecture: Emergence of Orientation Columns," Proc. Natl. Acad. Sci. USA, vol. 83, pp. 8779-8783, November 1986.

This paper is number 3 in a series.

Linsker, Ralph, "From Basic Network Principles to Neural Architecture: Emergence of Spatial-opponent Cells," Proc. Natl. Acad. Sci. USA, vol. 83, pp. 7508-7512, October 1986.

The functional architecture of mammalian visual cortex has been elucidated in impressive detail by experimental work of the past 20-25 years. The origin of many of the salient features of this architecture, however, has remained unexplained. This paper is the first of three (the others will appear in subsequent issues of these Proceedings) that address the origin and organization of feature-analyzing (spatial-opponent and orientation-selective) cells in simple systems governed by biologically plausible development rules. I analyze the progressive maturation of a system composed of a few layers of cells, with connections that develop according to a simple set of rules (including Hebb-type modification). To understand the prenatal origin of orientation-selective cells in certain primates, I consider the case in which there is no external input, with the first layer exhibiting random spontaneous electrical activity. No orientation preference is specified to the system at any stage, and none of the basic developmental rules is specific to visual processing. Here I introduce the theory of "modular self-adaptive networks," of which this system is an example, and explicitly demonstrate the emergence of a layer of spatial-opponent cells. This sets the stage for the emergence, in succeeding layers, of an orientation-selective cell population.

Loizou, George and Stephen J. Maybank , "The Nearest Neighbor and the Bayes Error Rates," Transactions on Pattern Analysis and Machine Intelligence, vol. PAMI-9, no. 2, pp. 254-262, IEEE, 1987.

The (k,l) nearest neighbor method of pattern classification is compared to the Bayes method. If the two acceptance rates are equal than the asymptotic error rates satisfy the inequalities:
See the Paper for Algorithm
where d is a function of k, l, and the number of pattern classes, and 2 is the reject threshold for the Bayes method. An explicit expression for d is given which is optimal in the sense that for some probability distributions
again, see the paper for equations are equal.

Lynch, G., Synapses, Circuits, and the Beginnings of Memory, MIT Press, Cambridge, MA, 1986.

Maxwell, T., C.L. Giles, Y.C. Lee, and H.H. Chen, "Nonlinear Dynamics of Artificial Neural Systems," in Neural Networks for Computing, American Institute of Physics, 1986.

Now that significant progress has been made in developing algorithms for training hidden units, we suggest that it is time to reevaluate the nonlinear discriminate approach, which once fell into disfavor due to the problem of proliferation of high order terms. We show that there are many powerful techniques for reducing the number of spurious terms, and that the high order approach has many advantages over a cascaded slab approach in certain problem areas. Advantages include increased expressive ability, decreased architectural complexity, and dramatically increased learning rates.

Maxwell, T., C.L. Giles, Y.C. Lee, and H.H. Chen, "Transformation Invariance Using High Order Correlations in Neural Net Architectures," Proceedings International Conf. on Systems, Man, and Cybernetics, 1986.

In this paper we examine the properties of high order neuron-like adaptive learning units whose output is invariant under an arbitrary finite group of transformations on the input space. The transformation invariance is imposed by averaging the input of each unit over a transformation group, thus eliminating the capacity of the units to detect features which are incompatible with the imposed group invariance. This averaging process also generates equivalence classes of interactions among the units, and thus allows a collapse of the interaction weight matrix, reducing the number of high order terms. Massively parallel networks of these elements are then constructed for distributed computation. Simulations are presented for pattern recognition and associative memory (both spatial and temporal) applications.

These simulations indicate that translation and other invariances can be implemented in neural networks by the process of averaging over transformation groups.

McCulloch, W.S., Embodiments of Mind, MIT Press, Cambridge, MA, 1965.

McEliece, R.J., E.C. Posner, E.R. Rodemich, and S.S. Venkatesh, "The Capacity of Hopfield Associative Memory," IEEE Transactions on Information Theory, vol. IT-33, no. 4, pp. 461-482, July 1987.

This paper applies techniques from coding theory to rigorously study the capacity of the Hopfield Associative Memory. Such a memory stores n-tuples of +-1's. The components change depending on a hard-limited version of linear functions of all other components. With symmetric connections between components, a stable state is ultimately reached. By building up the connection matrix as a sum of outer products of m fundamental memories, one hopes to be able to recover a certain one of the m memories by using an initial n-tuple probe vector less than a Hamming distance n/2 away from the fundamental memory. If m fundamental memories are chosen at random, the maximum asymptotic value of m in order that most of the m original memories are so recoverable exactly is shown to be n/(2logn). With the added restriction that every one of the m fundamental memories be recoverable exactly, m can be no more than n/4logn asyptotically as n approaches infinity. Extensions are also considered, in particular to capacity under quantization of the outer product connection matrix. This quantized memory capacity problem is closely related to the capacity of the quantized Gaussian channel.

McKay, D.M., D. Rose, and V.G. Dobson, "The significance of 'feature sensitivity'," in Models of the Visual Cortex, pp. 47-53, Wiley Interscience, New York, 1985.

Mead, C., Analog VLSI and Neural Systems, California Institute of Technology, Pasadena, CA 91125, 1986. (Preliminary Version)

Moody, John, Perspectives on Associative Memories, Proceedings of the ICNN, San Diego, CA, June 1987.

It is shown that the outer product associative memory models of Hopfield (1982 & 1984), the higher order correlation memory models of Chen et al (1986), and the Unary Memory Model of Baum, Moody, and Wilczek (1986) and Graf and de Vegvar (1987) can be described as special cases of a general associative memory architecture (Moody 1987a). This architecture not only uses the hardware in an optimally efficient fashion, but allows an understanding of the comparative performance of the different models.

Munro, Paul, "A Dual Back-Propagation Scheme for Scalar Reward Learning," Proceedings of the 9th Cognitive Science Society Meeting, pp. 165-176, Seattle, Washington, July 1987.

Explicit supervised learning rules [e.g. the delta rule] require that each of the output units in a network receive a training signal indicating the "correct" response value; the unit can then adjust its parameters, so that its future response to the same stimulus is closer to the desired value. A much more realistic assumption for the nature of a supervisory signal is a single scalar "goodness-of-response" or "reward" signal. This credit assignment problem is handled here by a supervisory network which monitors the activities of both the sensory and effector units, and learns to predict the value of the reward signal using the generalized delta rule of Rumelhart, Hinton, and Williams (1986). The activity of a particular "predictor unit" thus comes to be associated with the expected reward. Having learned to mimic the environment's reward criteria, the supervisory network can provide each effector unit, by way of a back-propagation scheme, with an individualized correction signal that will lead to increased activity in the predictor. The actual reward is hence enhanced to the extent that the predicted reward is reliable.

Nielsen, Dan E., "Retinocortical Wiring of the Simple Cells of the Visual Cortex," in Models of the Visual Cortex, ed. D. Rose and V. G. Dobson, pp. 374-379, John Wiley & Sons Ltd, 1985.

Palmer, Larry A., Judson P. Jones, and Walter H. Mullikin, "Functional Organization of Simple Receptive Fields," in Models of the Visual Cortex, ed. D. Rose and V. G. Dobson, pp. 273-280, John Wiley & Sons Ltd, 1985.

Mammalian striate cortex has for some twenty years now been one of the main foci of progress in neuroscience. Testaments to this fact include the Nobel prize awarded to David Huble and Torsten Wiesel and the publication of this book. This situation is more true today than ever and the continued application of a host of new techniques promises to yield yet further understanding of this subtle and truly impressive piece of neural machinery. In the pages to follow, we will briefly summarize our own views on some rather narrow but important issues concerning the organization of simple receptive fields.

Parker, D.B., "Learning-Logic," Technical Report TR-47, Center for Computational Research in Economics and Management Science, MIT, Cambridge, MA, April 1985.

A Learning-Logic circuit is a nework of cells that can learn. This paper presents algorithms for several types of cells, and also presents an algorithm for connecting the

cells into networks. These algorithms may be implemented as electronic circuits or simulated in software. Learning-Logic is designed to work best at tasks which require massive amounts of parallel processing, such as pattern recognition. Learning-Logic cells and networks bear a strong resemblance to the neurons in the cortex of the human brain. A distinguishing feature of Learning-Logic is that each of the most important Learning-Logic cells performs on the order of p*p parallel arithmetic operations, where p is the number of inputs to the cell; most similar cells proposed by other authors only perform on the order of p operations. These other cells typically cannot be guaranteed to learn unless their signals are randomized in some manner, which leads to poor performance in situations where, for example, there isn't time to randomize the signals (such as real-time learning applications) or where randomizing the signals destroys information (such as when simultaneously learning several continuous functions). Learning-Logic eliminates the need for any form of randomization.

Parker, D.B., "Second Order Backpropagation: An Optimal Adaptive Algorithm," Proc. of the First International Conference on Neural Networks, San Diego, CA, June 1987.

An adaptive network consists of one or more adaptive cells that are connected together. The pattern of connections between the cells is called the architecture of the network. Many different types of adaptive cells and architectures are described in teh literature. This paper demonstrates that there is a single algorithm - called second order back-propagation - that can be used to give any of these adaptive cells or architectures optimal performance.

Parker, David B., "Second Order Backpropagation: An Optimal Adaptive Algorithm for Any Adaptive Network," Proceedings of the ICNN, San Diego, CA, June 1986.

An adaptive network consists of one or more adaptive cells that are connected together. The pattern of connections between the cells is called the architecture of the network. Many different types of adaptive cells and architectures are described in the literature. This paper demonstrates that there is a single algorithm - called second order backpropagation - that can be used to give any of these adaptive cells or architectures optimal performance.

We will proceed in two steps. First, we will derive an ideally optimal algorithm. For several reasons, this ideally optimal algorithm will be impossible to implement. >From it, however, we can then derive a realistically optimal algorithm, which will be second order backpropagation.

Parunak, H. Van Dyke, James Kindrick, and Bruce Irish, "Material Handling: A Conservative Domain for Neural Connectivity and Propagation," Proceedings of the AAAI, Seattle, Washington,

July 1987.

Two important components of connectionist models are the connectivity between units and the propagation rule for mapping outputs of units to inputs of units. The biological domains where these models are usually applied are nonconservative, in that a single output signal produced by one unit can become the input to zero, one, or many subsequent units. The connectivity matrices and propagation rules common in these domains reflect this nonconservatism in both learning and performance.
CASCADE is a connectionist system for performing material handling in a discrete parts manufacturing environment. We have described elsewhere the architecture and implementation of CASCADE [PARU86a] and its formal correspondence [PARU86c], [PARU87a] with the PDP model [RUME86]. The signals that CASCADE passes between units correspond to discrete physical objects, and thus must obey certain conservation laws not observed by conventional neural architectures.
This paper briefly reviews the problem domain and the connectionist structure of CASCADE, describes CASCADE's scheme for maintaining connectivity information and propagating signals, and reports some experiments with the system.

Patterson, P.H. and D. Purves, Readings in Developmental Neurobiology, Cold Spring Harbor Laboratory, 1982.

Pearlmutter, Barak A. and Geoffrey E. Hinton, "G-Maximization: An Unsupervised Learning Procedure for Discovering Regularities," in Neural Networks for Computing, American Institute of Physics, 1986.

Hill climbing is used to maximize an information theoretic measure of the difference between the actual behavior of a unit and the behavior that would be predicted by a statistician who knew the first order statistics of the inputs but believed them to be independent. This causes the unit to detect higher order correlations among its inputs. Initial simulations are presented, and seem encouraging. We describe an extension of the basic idea which makes it resemble competitive learning and which causes members of a population of these units to differentiate, each extracting different structure from the input.

Pineda, Fernando J., "Generalization of Backpropagation to Recurrent Neural Networks," Memo S1A-63-87, Johns Hopkins Applied Physics Laboratory, Laurel, MD, July 14, 1987.

An adaptive neural network with asymmetric connections is introduced. This network is related to the Hopfield network with graded neurons and uses a recurrent generalization of the _-rule of Rumelhart, Hinton and Williams to adaptively modify the synaptic weights. The new network bears a resemblance to the master/slave network of Lapedes and Farber but it is computationally more efficient. This adaptive network is well suited to implementation in VLSI.

Pollack, Jordan B., "Cascaded Back-Propagation on Dynamic Connectionist Networks," Proceedings of the 9th Cognitive Science Society Meeting, pp. 391-404, Seattle, Washington, July 1987.

The Back Propagation algorithm of Rumelhart, Hinton, and Williams (1986) is a powerful learning technique which can adjust weights in connectionist networks composed of multiple layers of perceptron-like units. This paper describes a variation of this technique which is applied to networks with constrained multiplicative connections. Instead of learning the weights to compute a single function it learns the weights for a network whose outputs are the weights for a network which can then compute multiple functions.
The technique is elucidated by example, and then extended into the realm of sequence learning, as prelude to work on a connectionist induction of grammars. Finally, a host of issues regarding this form of computation are raised.

Prazdny, K., "Similitude-Invariant Pattern Recognition Using Parallel Distributed Processing," Proceedings of the AAAI, Seattle, Washington, July 1987.

Translation, rotation-, and scale-invariant recognition of multiple, superimposed, partially specified or occluded objects can be accomplished in a fast, simple, distributed and parallel fashion using localizable features with intrinsic orientation. All known objects are recognized, localized, and segmented simultaneously. The method is robust and efficient.

Psaltis, D., "Two-Dimensional Optical Processing Using One-Dimensional Input Devices," Proceedings of the IEEE, vol. 72, no. 7, pp. 962-974, IEEE, July 1984.

Two-dimensional optical processing architectures that are implemented with one-dimensional input spatial light modulators are reviewed. The advanced state of the art of available one-dimensional devices and the flexibility that exists in the design of two-dimensional architectures with one-dimensional transducers leads to the implementation of the most powerful and versatile optical processors. Signal and image processing architectures of this type are discussed.

Raffell, Jack, James Mann, Robert Berger, Antonio Soares, and Sheldon Gilbert, "A Generic Architecture for Wafer-Scale Neuromorphic Systems," Proceedings of the ICNN, San Diego, CA, June 1986.

Wafer-scale integration is ideally matched to the extraordinary interconnect requirements of neuromorphic systems. The very high fanout associated with connectionist architectures results in designs which will ultimately be constrained by the pinout limitations of conventional packages. The Restructurable VLSI approach to wafer-scale integration uses a laser to both form and remove connections between conductors on a silicon wafer after fabrication and test. This technique accomplishes two important objectives: the incorporation of redundancy in order to circumvent defects and provide acceptable yields in very large monolithic systems, and the flexible configuration of interconnect between modular cells in order to program the functionality of the wafer-scale device for rapid turnaround customization. Whole wafer systems developed previously at Lincoln Laboratory using this technology include a 16-point-FFT, a Hough Transform Processor, and a system for implementing the Dynamic Time Warping algorithm for a speech recognizer.

Rose, D. and V.G. Dobson, in Models of the Visual Cortex, Wiley Interscience, New York, 1985.

Rosenberg, Charles R. and Terrence J. Sejnowski, "The Spacing Effect of NETtalk, A Massively-Parallel Network," in Eighth Annual Conference of the Cognitive Science Society, pp. 72-88, Lawrence Erlbaum Associates, Hillsdale, August, 1986.

NETtalk is a massively-parallel network that learns to convert English text to phonemes. In NETtalk, the memory representations are shared among many processing units, and these representations are learned by practice. In humans, distributed practice is more effective for long-term retention than massed practice, and we wondered whether learning in NETtalk had similar properties. NETtalk was tested on cued paired-associate recall using nonwords as stimuli. Retention of these target items was measured as a function of spacing, or the number of interspersed items between successive repetitions of the target. A significant advantage for spaced or distributed items was found for spacings of up to forty intervening items when tested at a retention interval of 64 items. Conversely, a significant advantage for massed items was found if testing immediately followed study. These results are strikingly similar to the results of many experiments using human subjects and suggest an explanation based on distributed representations in massively- parallel network architectures.

Rosenfeld, Edward, Neurocomputing - A New Industry, Intelligence Newsletter, New York, NY, June 1987.

For the second time in the past forty years, scientists and business people are forming companies to make practical, commercial applications out of discoveries and theories from the neurosciences, the cognitive sciences, and psychology.

By combining these theories and discoveries with advances in semiconductor fabrication and computing hardware and software technologies, new computing systems are evolving. These new systems, variously called neural networks, neurocomputers, adaptive and learning systems, associative and content-addressable memories, artificial neural networking simulation systems, connectionists and optical computing architectures, are now, once again, at the forefront of computing and artificial intelligence.

Ryan, T.W. and C.L Winter, Variations on Adaptive Resonance, Proceedings of the ICNN, San Diego, CA, June 1987.

In the process of implementing Adaptive Resonance Circuits (ARCs) [1] %A T.W. Ryan %A C.L Winter %C San Diego, CA %D June 1987 %I Proceedings of the ICNN %K ryan winter 1987 %T Variations on Adaptive Resonance %X In the process of implementing Adaptive Resonance Circuits (ARCs) [1] for a particular application [2], several circuit modifications and alternative processing conditions were considered. The objective of this paper is to report on some of these variations. First, we consider an adaptive thresholding technique that prevents inadvertent recoding of recognition nodes which can occur when novel patterns are presented. Next, we examine the behavior of an ARC when patterns are iteratively presented for relatively short periods of time. Finally, we consider the case of continuous ARC operation in which "neural" activity is not reinitialized with each new pattern presentation (as done in [1]). The adaptive thresholding technique provides a novel clustering algorithm applicable to both binary and multilevel data.

Segalowitz, S.J., in Language Functions and Brain Organization, Academic Press Inc., New York, New York, 1983.

Shannon, C.E. and W. Weaver, Mathematical Theory of Communication, The University of Illinois Press, 1949.

Smolensky, P., D.E. Rumelhart, and J.L. McClelland, Information Processing in Dynamical Systems: Foundations of Harmony Theory, Parallel Distributed Processing: Explorations in the Microstructure of Cognition, 1, pp. 194-281, Bradford Books/MIT Press, Cambridge, MA, 1986.

Somani, A.K. and N. Penla, "Compact Neural Networks," Proc. of the First International Conference on Neural Networks, San Diego, CA, June 1987.

A neural network model for content addressable memory (CAM) applications as suggested by Hopfield ahs a completely connected graph, i.e., each neuron (processing element or PE) is connected to every other neuron. This requires O(L**2) interconnections where L is the number of neurons in the system. Such as system, when implemented on a VLSI silicon chip, complicates the layout tremendously. Even in optical implementation, any reduction in the number of interconnections will ease the situation. In this paper, we explore this problem and suggest two neural networks, Compact Neural Network (CNN) and Reduced Interconnections Neural Network (RINN), which require far less interconnections than those of Hopfield's model without significant loss of performance.

Somani, Arun K. and Narong Penla, "Compact Neural Network," Proceedings of the ICNN, San Diego, CA, June 1986.

A neural network model for content addressable memory (CAM) applications as suggested by Hopfield has a completely connected graph, i.e., each neuron (processing element or PE) is connected to every other neuron. This requires O(L/squared) interconnections where L is the number of neurons in the system. Such a system, when implemented on a VLSI silicon chip, complicates the layout tremendously. Even in optical implementation, any reduction in the number of interconnections will ease the situation. In this paper, we explore this problem and suggest two neural networks, Compact Neural Network (CNN) and Reduced Interconnections Neural Network (RINN), which require far less interconnections than those of Hopfield's model without significant loss of performance.

Soulie, Francoise Fogelman and Gerard Weisbuch, "Random Iterations of Threshold Networks and Associative Memory," SIAM J. Comput., vol. 16, no. 1, pp. 203-220, 2/87.

A threshold network that can memorize and retrieve a given set of patterns Xs is presented. A theoretical framework is given to allow expression of this property in terms of the dynamics of the network: the patterns in Xs must be attractive fixed points. We give sufficient conditions for this, using a concept of energy. Scaling laws for the storing capacity of random patterns are provided and tested in the case of random as well as nonrandom patterns, such as the letters of the alphabet.

Sutton, Richard S., "Two Problems with Backpropagation and Other Steepest-Descent Learning Procedures for Networks," in Eighth Annual Conference of the Cognitive Science Society, pp. 823-831, Lawrence Erlbaum Associates, Hillsdale, August, 1986.

This article contributes tot he theory of network learning procedures by identifying and analyzing two problems with the backpropagation procedure of Rumelhart, Hinton, and Williams (1985) that may slow its learning. Both problems are due to backpropagation's being a gradient- or steepest-descent method in the weight space of the network. The first problem is that steepest descent is a particularly poor descent procedure for surfaces containing places which curve more sharply in some directions than others-- and such ravines are common and pronounced in performance surfaces arising from networks. The second problem is that steepest descent results in a high level of interference between learning with different patterns, because those units that have so far been found most useful are also those most likely to be changed to handle new patterns. The same problems probably also arise with the Boltzmann machine learning procedure (Ackley, Hinton and Sejnowski, 1985) and with reinforcement learning procedures (Barto and Anderson, 1985), as these are also steepest- descent procedures. Finally, some directions in which to look for improvements to backpropagation based on alternative descent procedures are briefly considered.

Sutton, Richard S. and Andrew G. Barto, "A Temporal-Difference Model of Classical Conditioning," Proceedings of the 9th Cognitive Science Society Meeting, pp. 355-378, Seattle, Washington, July 1987.

Rescorla and Wagner's model of classical conditioning has been one of the most influential and successful theories of this fundamental learning process. The learning rule of their theory was first described as a learning procedure for connectionist networks by Widrow and Hoff. In this paper we propose a similar confluence of psychological and engineering constraints. Sutton has recently argued that adaptive prediction methods called temporal-difference methods have advantages over other prediction methods for certain types of problems. Here we argue that temporal-difference methods can provide detailed accounts of aspects of classical conditioning behavior. We argue that it is an improvement over the Rescorla-Wagner model in its handling of within-trial temporal effects such as the ISI dependency, primacy effects, and the facilitation of the remote associations in serial-compound conditioning. The new model is closely related to the model of classical conditioning that we proposed in 1981, but avoids some of the problems with that model recently identified by Moore et. al. We suggest that the theory of adaptive prediction on which our model is based provides insight into the functionality of classical conditioning behavior.

Takeda, Mitsuo and Joseph W. Goodman, "Neural Networks for Computation: Number Representations and Programming Complexity," Dept. of Electrical Engineering, Stanford University, Stanford.

Methods for using neural networks for computation are considered. The success of such networks in finding good solutions to complex problems is found to be dependent on the number representation schemes used. Redundant schemes are found to offer advantages in terms of convergence. Neural networks are applied to the combinatorial optimization problem known as the "Hitchcock problem", and signal processing problems, such as matrix inversion, and Fourier transformation. The concept of programming complexity is introduced. It is shown that for some computational problems, the programming complexity may be so great as to limit the utility of neural networks, while for others the investment of computation in programming the network is justified. Simulations of neural networks using a digital computer are presented.

Tanenbaum, Andrew S. and Nancy Milnamow, in Computer Networks, Prentice-Hall, Englewood Cliffs, NJ, 1981.

As computers have become smaller, cheaper, and more numerous, people have become more and more interested in connecting them together to form networks and distributed systems. At first these connections were made in an ad hoc way, typically with each computer regarding the other ones as terminals. However, in the past decade, a substantial body of knowledge has developed on the subject of computer networking, so that future networks and distributed systems can be designed systematically. The key to designing a computer network was first enunciated by Julius Caesar: Divide and Conquer. The idea is to design the network as a sequence of layers, or abstract machines, each one based upon the previous one. By reducing the study of the whole to the study of its parts, the subject becomes more manageable. This book uses a model in which networks are divided into seven layers. The structure of the book follows the structure of the model to a considerable extent.

Tank, D.W. and J.J. Hopfield, Neural Computation by Concentrating Information in Time.

An analog model neural network that can solve a general problem of recognizing patterns in a time-dependent signal is presented. The networks use a patterned set of delays to collectively focus stimulus sequence information to a neural state at a future time. The computational capabilities of the circuit are demonstrated on tasks somewhat similar to those necessary for the recognition of words in a continuous stream of speech. The network architecture can be understood from consideration of an energy function that is being minimized as the circuit computes. Neurobiological mechanisms are known for the generation of appropriate delays.

Tesauro, G., "Simple Neural Models of Classical Conditioning," in Biological Cybernetics, Springer-Verlag, 1986.

A systematic study of the necessary and sufficient ingredients of a successful model of classical conditioning is presented. Models are constructed along the lines proposed by Gelperin, Hopfield, and Tank, who showed that many conditioning phenomena could be reproduced in a model using non-

trivial distributed representations of the sensory stimuli. The additional phenomena of extinction and blocking are found to be obtainable be generalizing the Hebbian learning algorithm, rather than by additional complications in the hardware. The most successful algorithms have a minimal number of adjustable parameters, and require only local-time information about the level of postsynaptic activity. The proper behavior of these algorithms is verified by both simple analytic arguments and by direct numerical simulation. Certain detailed assumptions concerning the distributed sensory representations are also found to have a surprising degree of importance.

Thagard, Paul, "Parallel Computation and the Mind-Body Problem," Cognitive Science 10, pp. 301-318, 1986.

The position in the philosophy of mind called functionalism claims that mental states are to be understood in terms of their functional relationships to other mental states, not in terms of their material instantiation in any particular kind of hardware. But the argument that material instantiation is irrelevant to functional relationships is computationally naive. This paper uses recent work on parallel computation to argue that software and hardware are much more intertwined than the functionalists allow. Parallelism offers qualitative as well as quantitative advantages, leading to different styles of programming as well as increased speed. Hence hardware may well matter to the mental: only by further empirical investigations of the relation between the mind and brain and between artificial intelligence software and underlying hardware will we be able to achieve a defensible solution to the mind-body problem. The major disadvantage of parallel systems is the need to coordinate their subprocesses, but recent proposals that consciousness provides a serial control for parallel computation are implausible.

Thompson, Richard F., "The Neurobiology of Learning and Memory," Science, vol. 233, pp. 941-947, 8/29/86.

Study of the neurobiology of learning and memory is in a most exciting phase. Behavioral studies in animals are characterizing the categories and properties of learning and memory; essential memory trace circuits in the brain are being defined and localized in mammalian models; work on human memory and the brain is identifying neuronal systems involved in memory; the neuronal, neurochemical, molecular, and biophysical substrates of memory are beginning to be understood in both invertebrate and vertebrate systems; and theoretical and mathematical analysis of basic associative learning and of neuronal networks is proceeding apace. Likely applications of this new understanding of the neural bases of learning and memory range from education to the treatment of learning disabilities to the design of new ar-

tificial intelligence systems.

Touretzky, David S. and Geoffrey E. Hinton, CMU Tech Report #CMU-CS-86-172, Carnegie-Mellon University, Pittsburgh, PA, 1986.

DCPS is a connectionist production system interpreter that uses distributed representations. As a connectionist model it consists of many simple, richly interconnected neuron-like computing units that cooperate to solve problems in parallel. One motivation for constructing DCPS was to demonstrate that connectionist models are capable of representing and using explicit rules. A second motivation was to show how "course coding" or "distributed representations" can be used to construct a working memory that requires far fewer units than the number of different facts that can potentially be stored. The simulation we present is intended as a detailed demonstration of the feasibility of certain ideas and should not be viewed as a full implementation of production systems. Our current model only has a few of the many interesting emergent properties that we eventually hope to demonstrate: it is damage resistant, it performs matching and variable binding by massively parallel constraint satisfaction, and the capacity of its working memory is dependent on the similarity of the items being stored.

Touretzky, David S., "Representing and Transforming Recursive Objects in a Neural Network, or "Trees Do Grow On Boltzmann Machines"," Proceedings of the 1986 IEEE International Conference on Systems, Man, and Cybernetics, Atlanta, October 14-17, 1986.

BoltzCONS is a nerual network that manipulates symbolic data structures. The name reflects the system's mixed representational levels: it is a Boltzmann Machine in which Lisp cons cell-like structures appear as an emergent property of a massively parallel distributed representation. BoltzCONS is controlled by an attached neural network production system interpreter also implemented as a Boltzmann Machine. Gated connections allow the production system and BoltzCONS to pass symbols back and forth. A toy example is presented where BoltzCONS stores a parse tree and the production system contains a set of rules for transforming parse trees from active to passive voice. The significant features of BoltzCONS are its ability to represent structured objects and its generative capacity, which allow it to create new symbol structures on the fly.

Waibel, A., "Prosody and Speech Recognition," CMU Tech Report #CMU-CS-86-162, Carnegie-Mellon University, Pittsburgh, PA, 1986.

Although numerous studies have demonstrated that prosody is critical to human speech perception, many automatic speech recognition systems process only spectral/phonetic cues. They ignore or deliberately remove prosodic cues such as pitch, intensity, rhythm, temporal relationships, and stress. Extending speech recognition systems to human performance levels, however, will require exploiting all available cues and sources of knowledge.

This work demonstrates the power of prosodic constraints in computer speech recognition systems. We first show theoretically that prosodic patterns can discriminate between words in large vocabularies (vocabularies the size an adult typically commands). We then introduce several novel algorithms to extract prosodic parameters reliable. These parameters include segmentation algorithms for detecting syllable boundaries and major segment boundaries and algorithms for measuring pitch and intensity contours, and lexical stress levels. Extensive performance evaluation of these algorithms is presented. We then implement and evaluate prosodic knowledge sources that apply the extracted parameters at appropriate processing levels including the lexical, syntactic and sentential levels. To permit large vocabulary capability, the knowledge source designs emphasize a concern for minimizing lexical search, exploiting parallelism and speaker-independent and/or template-independent operation.

Watanabe, S., Pattern Recognition: Human and Mechanical, John Wiley and Sons, Inc., New York, 1985.

Watrous, Raymond L. and Lokendra Shastri, "Learning Phonetic Features Using Connectionist Networks: An Experiment in Speech Recognition," Technical Report MS/-CIS-86-78, Philadelphia, PA, 10/86.

Weste, N. and K. Eshraghian, Principles of CMOS VLSI Design: A Systems Perspective, Addison-Wesley, 1985.

Widrow, B. and M.E. Hoff, "Adaptive switching circuits," WESCON Convention Record Part IV, pp. 96-104, 1960.

Widrow, B. and S.D. Stearns, Adaptive Signal Processing, Prentice-Hall, Engelwood Cliffs, NJ, 1985.

Williams, Ronald J., "Reinforcement - Learning Connectionist Systems," Technical Report NU-CCS-87-3, Boston, MA, 2/87.

This report begins by discussing a number of issues concerning learning algorithms to be used in training the behaviors of connectionist systems, or networks of neuron-like processing units. In particular, the need for algorithms which train temporally extended behaviors and which are capable of on-line implementation is argued. Then this report presents an intersting general class of associative reinforcement learning algorithms, which apply to systems where the pro-

cessing units operate stochastically and the training feedback consists of scalar reinforcement signals. What makes these algorithms interesting is that: (1) they statistically follow the gradient of a natural performance measure for reinforcement learning problems, making them analogues for such problems of the back-propagation learning procedure (Rummelhart, Hinton, & Williams, 1986) for supervised learning problems; (2) they admit a simple on-line implementation even when applied to the problem of training temporally extended behaviors, which is not true for the back-propagation algorithm; and (3) certain algorithms studied elsewhere belong to this class, including the reward-inaction variant of Barto's reward-penalty algorithm (1985), some stochastic learning automata algorithms, and variants of these in which Sutton's reinforcement comparison methods (1984) are used. Furthermore, while Sutton's promising adaptive heuristic critic algorithm (1984) is not a member of this class, ways the theory presented here might nevertheless help provide needed mathematical insight into its behavior are discussed. Finally, some particular implications of the theory presented here for the development of novel learning algorithms are noted.