J. Indian Inst. Sci., Sept.-Oct., 1991, **71**, 491–502.

Neural networks: An overview *

CHANDAN DASGUPTA
Department of Physics, Indian Institute of Science, AND Jawaharlal Nehru Centre for Advanced Scientific Research, Bangalore 560 012.

Abstract

A brief overview of the newly emerging field of neural networks is presented from the perspective of a condensed matter physicist. The basic ingredients of neural network modeling are described and the principles governing the functioning of symmetric and feed-forward networks are outlined. A survey of some of the important results obtained from neural network research in the fields of neurobiology, physics, computer science and engineering is given.

Key words: Neural networks, symmetric and feed-forward networks, condensed matter physics, neuro-biology, computer science and engineering.

1. Introduction

The subject of neural networks has received a great deal of attention during the last few years. Researchers from a wide variety of disciplines such as neurobiology, cognitive science, computer science, electrical engineering, physics and mathematics are currently working in this area. In this article, I shall attempt to provide a brief overview of this rapidly developing field. The interdisciplinary nature of this field makes the task of providing a comprehensive review rather difficult. I, therefore, mention at the beginning that this review is from the point of view of a physicist interested in the statistical mechanical aspects of neural network modeling, and my notion of what is important and interesting in this field would probably differ considerably from that of neural network researchers from other disciplines.

The rest of this article is organised as follows. The basic ingredients of neural network modeling are described in Section II. In Section III, the working principles of symmetric (Hopfield-type) and feed-forward networks are explained. Section IV is devoted to a description of some of the important developments which have taken place in this field during the last few years.

* Text of lecture delivered at the Annual Faculty Meeting of the Jawaharlal Nehru Centre for Advanced Scientific Research, on November 23, 1990.

2. Basic ingredients of neural network modeling

A neural network may be defined as a large, highly interconnected assembly of simple computing elements (*model neurons*). These networks exhibit non-trivial computing abilities emerging from the collective dynamics of the constituent neurons. The model neurons used in most neural networks are highly simplified versions of their biological counterparts. They are characterized by a sigmoid input-output relation similar to the one shown in Fig. 1. In some models, the input-output relation is simplified further by approximating the sigmoid function by a step function shown by the dotted line in Fig. 1. A model neuron with such input-output relation may be considered to be a two-state threshold device. If the net input (u_i) to the ith neuron exceeds a threshold value u_{io}, then the neuron is in the 'on' state for which the output $v_i = 1$; if the input u_i is less than u_{io}, the neuron is in the 'off' state with $v_i = 0$. Model neurons of this type, first introduced by McCulloch and Pitts[1] in 1943, are called *digital* neurons and those with a continuous input-output relation are called *analog* neurons. Each neuron in the network receives inputs from and provides outputs to a large number of other neurons. The connectivity of the network is specified by the *synaptic interaction matrix J*. The element J_{ij} of this matrix represents the input provided by the jth neuron (if it is in the 'on' state) to the ith neuron. Thus, the net input to the ith neuron is given by

$$u_i = \Sigma_j J_{ij} v_j. \tag{1}$$

Connections with positive values of J_{ij}, which act to increase the input to the receiving neurons, are called *excitatory* and those with negative values are called *inhibitory*. A common feature of all neural network models is the property that the

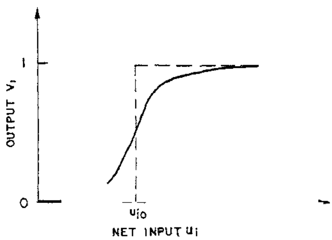

FIG. 1. Input-output relation of model neurons. continuous line: analog neurons; broken line: digital neurons.

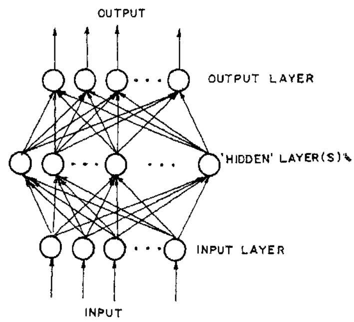

FIG. 2. Typical structure of feed-forward neural networks.

computations to be performed by the network are encoded in the synaptic matrix J. The time evolution of such a network is governed by an assumed dynamics which specifies how the state of each neuron, denoted by the value of the output variable v, is updated. The updates, which take place in parallel, can be synchronous or asynchronous. In most models, the update rule is assumed to be deterministic, *i.e.*, the output state of each neuron is uniquely determined in terms of the net input received by it. In some models, a stochastic (probabilistic) update rule is used. The network is started off from an initial configuration in which the output state of each neuron is specified. The network then evolves in time in accordance with the assumed dynamics until a *time-persistent* state (a fixed point or a limit-cycle attractor of the underlying dynamics) is reached. This final state of the network represents the result of the computation performed by the network.

There are several reasons behind the current interest in models of this sort. At a highly simplified level, these systems serve as models of some of the collective computational properties of biological networks. The simple two-state neurons used in most neural network models resemble, although in a highly schematic way, biological neurons which fire (*i.e.*, send a pulse of electric potential along the axon) when the sum of the membrane potentials received from the synaptic inputs exceeds a 'firing' threshold. Also, the 'parallelism' inherent to neural network models is known to be an essential element of biological computation. Thus, it is hoped that studies of simple neural network models, some of which are amenable to mathematical analysis, may shed some light on the principles underlying computations in biological systems. At a more practical level, neural networks provide a new paradigm of parallel computing, with numerous applications in pattern recognition and combinatorial optimisation problems. The parallel processing performed by neural networks may also be used in designing fast general-purpose computers. Another attractive feature of these systems is their robustness. Since the computation to be performed is coded in a large number of interconnections, the performance of the system is not degraded much if a small fraction of the interconnections are corrupted. This fault-tolerance capability arising from the distributed nature of the coding in neural networks is not present in conventional digital computers.

3. Working principles of neural networks

The most commonly studied neural network models may be divided into two classes: networks with symmetric interconnections and feed-forward networks. The basic principles underlying the functioning of these two classes of networks are described in this section.

3.1. *Networks with symmetric interconnections*

In this class of models, the synaptic interaction matrix J is symmetric ($J_{ij} = J_{ji}$ for all $i \neq j$). The Hopfield model[2] is the most well-known example of this class. It is a model of *associative* or *content-addressable* memory. A system behaves as an associative memory if it can retrieve patterns stored in it from 'hints' representing a knowledge of the stored information. To understand how a neural network with symmetric connections may act as an associative memory, it is useful to define an *energy function* E in the following way:

$$E = -\Sigma_{i>j} J_{ij} v_i v_j + \Sigma_i u_{io} v_i. \tag{2}$$

The most commonly used dynamics, which is asynchronous and deterministic, corresponds to updating the neurons one at a time in a random sequence. The ith neuron is updated according to the rule

$$v_i(t+1) = 1 \text{ if } u_i(t) = \Sigma_j J_{ij} v_j(t) - u_{io} \geqslant 0,$$
$$v_i(t+1) = 0 \text{ if } u_i(t) < 0, \tag{3}$$

where t represents a discrete 'time' label. It is easy to see that this update scheme corresponds to the rule that the state of a neuron is changed only if the energy defined in eqn (2) is decreased in the process. It then follows that every configuration corresponding to a *local minimum* of E (*i.e.*, every configuration of the set of binary variables $\{v_i\}$ with the property that a change of any one of these variables increases the energy E) represents an attractive limit point of the assumed dynamics. Any configuration close to such a local minimum converges to it under the dynamics and the system remains in that configuration for all subsequent times. This property of a network of this type allows it to be used as an associative memory if the interaction matrix is chosen so as to make each pattern to be stored and associatively recalled (which is represented in this scheme as a binary string) correspond to the configuration at a local minimum of the energy function E. An initial state close to such a local minimum, representing partial knowledge of the stored information, converges to the local minimum under the collective dynamics of the network, thus retrieving the complete information. Such a network also functions as a pattern classifier. All input configurations lying within the basin of attraction of one of the stored patterns are classified by the network as belonging to the same category.

To construct such a model of associative memory, it is necessary to find a prescription for the construction of the synaptic matrix J which ensures that the states representing the memories to be stored in the network correspond to local minima of the associated energy function defined in eqn (2). Such a prescription for constructing the synaptic matrix is called a *learning rule*. Many different learning rules have been proposed over the years, the most frequently used among them being the Hebb rule[3] and its variants[4] and the pseudo-inverse rule[5]. Analytic studies and numerical simulations of models using these or other learning rules have demonstrated that networks of this type are indeed capable of storing and retrieving a large number of patterns.

3.2. *Feed-forward networks*

Most neural networks used in pattern recognition belong to this class. These networks have a layered structure (*see* Fig. 2) in which neurons in the first 'hidden layer' receive their inputs from the neurons in the input layer, and in turn, provide inputs to the neurons in the second 'hidden layer', and so on. The information thus flows forward and the computation performed by the network is a mapping of the state of the input layer on to that of the output layer. The single-layer perceptron[7], consisting of only the input and the output layer, is the simplest example of such a network. Although interest in the perceptron diminished in the 1960s after demonstration by Minsky and Papert[8] that the class of problems which can be solved by such a network is severely restricted, interest in feed-forward networks, containing one or more hidden layers, has been revived in recent years as new studies have demonstrated that these systems are much more powerful (*see* Rumelhart *et al*[9] for an introduction to the theory and applications of feed-forward networks). One of the most interesting features exhibited by these networks is the ability to learn from examples and to generalize what it has learnt. To illustrate this aspect, let us consider a pattern classification problem in which the network is supposed to correctly classify a group of binary patterns $\{\xi_i^{j,k}\}$, $i = 1,2, \ldots, N$; $j = 1,2, \ldots, M$; and $k = 1,2 \ldots, p$. Here, N is the number of neurons in the input layer, M the number of distinct classes which is also the number of neurons in the output layer, and p the number of patterns in each class. The problem is to find a set of synaptic couplings which ensure that when a pattern belonging to the jth class is presented to the input layer, the corresponding output has only the jth neuron in the 'on' state and all other neurons in the 'off' state. In most cases, it is not possible to give an explicit prescription for finding the right couplings. However, in many cases, it is possible to 'teach' the network to carry out the required task by using a suitable 'learning algorithm'. Many such algorithms have been proposed (*see* Hinton[10] for a review), the simplest among them being the so-called *supervised learning scheme*. In this scheme, a fraction of the patterns to be classified is used to teach the network the rule underlying the classification process. The training process consists of presenting patterns which are known to belong to particular classes to the network one by one. If the network classifies the input pattern correctly, then the couplings are left unchanged. If the pattern is not classified correctly, then the couplings are changed according to a prescribed algorithm (such as the popular *back propagation algorithm*[9]). This procedure is continued until all the patterns in the training set are classified correctly, or until the average error in classifying these patterns drops below a specified limit. In most cases, there is no guarantee that the learning procedure will converge to a solution. However, it is found from experience that by adjusting the number of hidden layers and the number of neurons in each hidden layer, it is usually possible to design a network for which the training procedure converges to a solution. It is also found that if a pattern not belonging to the training set is presented to the trained network, then the network can classify it correctly with a high probability. These networks, thus, exhibit the ability to learn from examples and the capacity to generalize what it has learnt to perform non-trivial computational tasks. Networks of this type have been used during the last few years in a variety of practical applications, some of which are described in the next section.

4. Survey of recent developments

This section contains a brief survey of some of the interesting results obtained during the last few years from neural network research in the fields of neurobiology, physics and computer science and engineering. This survey is certainly not complete; the intention here is to provide some indication of the kind of developments which have taken place in neural network research in these fields during recent years.

4.1. *Neurobiology*

The question of whether neural network models of the type described above are relevant to real biological networks such as the brain has been the subject of much debate in recent years. Some researchers, mostly neurobiologists, argue that these models are much too simple to be of use to the study of brain functions. It is indeed true that the model neurons used in most neural network models are highly simplified versions of their biological counterparts, which are highly complicated cells requiring a large number (~ 40) of variables for a realistic description of their behavior. Thus, neural network models using simple model neurons are not expected to provide much insight into neurobiological phenomena in which properties of individual neurons play an important role. However, many other researchers in this field, mostly physicists, believe that details of the functioning of individual neurons may not be crucial in understanding some of the *collective computational properties* of biological networks. These researchers have concentrated on the development of models which mimic some of the simpler aspects of the functioning of the brain. Some of the developments taking place in this line of research during the last few years are listed below.

a) *Models of short-term memory*: The Hopfield model described above is the simplest one in this class. The behavior of this model is now almost fully understood (see below). More recently, many researchers have investigated the effects of incorporating some of the known neurobiological facts on the functioning of the Hopfield model. The features which have been incorporated in the modeling include the presence of static synaptic noise, less-than-full connectivity of the network, limited analog depth of the synaptic connections, low average level of the activity of the network, asymmetry of the synaptic connections and synaptic specificity (*see* Amit[11] for a description of these models). A very interesting result has emerged from the study of these models. It has been established that the basic functional features of the Hopfield model remain, to a large extent, unaffected by the incorporation of these neurobiological details. This result lends some support to the viability of this class of models. Another interesting development in this field has been the construction of models of 'memory palimpsests' [12,13]. These are networks exhibiting a selective erasure ('forgetting') of old patterns stored in the memory as new information is memorised. Nearly all available models of memory use the Hebbian learning scenario of synaptic modifications for the storage of the memorised information. Although there exists a lot of evidence indicating that learning causes synaptic modifications in biological networks, the details of this process are not understood yet.

b) *Central pattern generators*: Many biological systems exhibit *central pattern generators* which are neural groups repeatedly generating a specific temporal sequence of patterns. These neural circuits control the muscles involved in a wide variety of rhythmic functions such as breathing, locomotion, swimming and chewing (*see e.g.,* Kristan[14]). These systems have been modeled by neural networks which can store and recall temporal sequences of patterns. These models make use of either a time-delay mechanism[15,16] or the presence of dynamic synaptic noise[17] to generate a passage of the network through a specified periodic sequence of patterns. Numerical simulations of these networks have reproduced several features exhibited by biological central pattern generators. Networks of this type can also be used for the recognition of temporal sequences of patterns[16] and for performing abstract computations such as counting the number of chimes of a clock[18].

c) *Neurological disorders*: An interesting example of research in this area is the modeling of a mental disorder known as *prosopagnosia*. Persons suffering from this disorder can recognise only the generic class to which an object belongs, not the specific object itself. For example, a patient may know that a face is a face, but would not be able to recognise whose face it is even if it is the face of a familiar person (*see e.g.,* Damasio *et al*[19]). This disorder has been modeled by a hierarchical network which stores correlated memories belonging to different categories, and their 'ancestor' patterns representing various categories. It is found that if this network is subjected to random synaptic corruption or to a high level of synaptic noise, then it can recall the ancestor patterns representing different categories, but not the individual patterns belonging to the selected category[20]. This behavior is quite similar to the symptoms of prosopagnosia. Another neurological disorder for which neural network models are being developed is epilepsy[21].

Although the models described above look quite promising, the real test of their viability will come from experiments probing neural processes at the individual neuron level. Work in this area is just beginning, and many new and interesting developments are expected in the coming years.

4.2. *Physics*

At a basic level, neural networks described above are dynamic systems involving a large number of interacting variables. Such systems are rather common in condensed matter physics, and the branch of physics known as statistical mechanics has been developed specifically for the study of such systems. Thus, it is perhaps not surprising that a large number of physicists have become involved in neural network research. Many interesting results about the properties of neural network models have been obtained by physicists from the application of the concepts and methods of statistical mechanics. Some of these developments are summarised below.

a) *Equilibrium statistical mechanics*: Methods of equilibrium statistical mechanics are readily applicable to the study of neural network models with a symmetric synaptic interaction matrix J. As mentioned before, one may define an energy function for such networks. The binary variables $\{v_i\}$ which specify the states of a network consisting of digital neurons can be readily transformed to variables $\{\sigma_i\}$, each of which can have the values $+1$ and -1 [$\sigma_i = 2\,v_i - 1$]. Thus, the energy function can be written in terms of these new variables $\{\sigma_i\}$, which are called Ising spins in studies of magnetism. Neural network models with digital neurons and a symmetric J may, therefore, be looked upon as Ising models with a Hamiltonian (energy function) defined in eqn (2). A similar formulation for models with analog neurons is also possible. It can be shown that the most commonly used stochastic update rule for the neurons generates a set of network configurations which are distributed according to the Boltzmann distribution in which the probability of occurrence of a state $\{\sigma_i\}$ is proportional to exp $(-E\{\sigma_i\}/T)$, where T is a parameter (the temperature) that specifies the degree of stochasticity in the dynamics (the $T = 0$ limit corresponds to the deterministic case). The average properties of the states generated by the dynamics may, therefore, be calculated by using methods of equilibrium statistical mechanics which have been developed for the calculation of average properties of systems described by the Boltzmann distribution. The statistical mechanics of Hopfield-type models turns out to be quite similar to that of a class of models used to describe certain disordered magnetic systems known as *spin glasses*. These spin glass models have been the subject of intense study during the last two decades, and a number of analytic methods have been developed by physicists for dealing with the equilibrium statistical mechanics of these systems (*see* Mezard *et al*[22] for a review of spin-glass physics). The application of these methods to the analysis of neural network models was pioneered by Amit *et al*[23]. Since then, a wealth of information about the behavior of Hopfield-type models has been obtained by physicists from similar studies. The issues addressed in these studies include the storage capacity of models of associative memory, the average fraction of errors in the retrieval of stored information, the effects of fast synaptic noise on the performance of the network, the relative merits and demerits of different learning rules, the effects of incorporating neurobiological facts in the models and the performance of models using hierarchical structures for the storage of correlated memories. Extensive numerical simulations have established the validity of the results obtained from these analytic calculations. A comprehensive account of the work in this area performed during the last few years may be found in the book of Amit[11].

b) *Dynamics*: The methods of equilibrium statistical mechanics cannot be used in the study of models with asymmetric connections or time-delayed interactions. A full analysis of the dynamics is necessary for understanding the behavior of such models. Neural network models form a class of general dynamical systems known as *cellular automata* which have been used extensively by physicists as models of complex systems. Some of the methods developed for the study of these systems and the insight gained from such studies may be used in the analysis of the dynamics of neural networks. However, the methods for the study of the dynamics of complex systems are not as well developed as those of equilibrium statistical mechanics. Consequently, not many analytic calculations on the dynamics of neural network models have been carried out. There exist a few specially constructed models (*see e.g.*, Derrida *et al*[24]) whose dynamics can be solved exactly. One has to rely on numerical simulations for information about the dynamic behavior of nearly all other models. This is an area where more research is needed.

c) *Learning*: The theory of learning involves methods of finding the synaptic connections appropriate for a specified computational task. In many cases, the problem of finding the right synaptic matrix may be formulated as an optimisation problem by defining an appropriate *cost function* which is a function of the elements of the synaptic matrix with the property that every solution of the learning problem corresponds to a minimum of this function. Recently, a set of analytic methods known as 'statistical mechanics in the space of interactions' has been developed[25] to deal with various aspects of this optimisation process. The issues which have been addressed in studies using these methods include the maximum number of patterns which can be classified by a perceptron[25], the convergence time of the learning process[26], and the ability of a perceptron to generalise what it has learnt[27]. This line of research is very promising, with many potential applications to the theory of feed-forward networks.

4.3. *Computer science and engineering*

Neural network research has obvious implications in the area of artificial intelligence. Recent developments in neural network modeling have led to the formation of the so-called 'connectionist' school of artificial intelligence and a lively debate is currently going on between researchers belonging to this school and the proponents of the more conventional, rule-based methods of artificial intelligence (*see e.g.*, Graubard[28]). This issue will not be settled before the capabilities of either approach are fully understood.

Another area of computer science in which neural network research has had considerable impact is parallel processing. The parallelism inherent in the functioning of neural nets provides a new paradigm for the development of parallel processing computers. A lot of research is going on in this area, and neural networks have already found applications in obtaining near-optimal solutions of hard combinatorial optimisation problems such as the travelling salesman problem[29].

Neural networks have also been successfully used in a large number of applications involving pattern-recognition problems. Most of these applications use feed-forward networks and the back propagation algorithm for training. The number of such applications is growing day by day, and it is not possible to list all of them here. A representative sampling is given below.

NETtalk, a network that can pronounce English text[30] given a moving 7-character window as input and a phoneme code as output.

Signal prediction: Given a signal $x(t)$ for times $t \leqslant t_0$, the network predicts its value for times $t > t_0$[31]. This has many applications, from stock market to weather. This method is found to perform considerably better than all standard algorithms for signal prediction.

Protein secondary structure determination: This network takes as input a moving window of 13 amino acids in the primary structure and produces as output a prediction of the secondary structure[32].

Explosive detection at airports: In this device, the baggage to be examined is bathed in neutrons and the gamma-ray signal produced is fed through a neural network which is trained to recognise the signatures of common explosives. This device has been endorsed by the Federal Aviation Administration of the USA.

A neural network that tests the quality of products from their acoustic signatures has been installed on a production line of electric motors by Siemens in a pilot project.

Recognition of handwritten characters and spoken words: Several companies have come up with neural network-based systems for these pattern-recognition tasks. There are several defence applications such as target recognition and tracking, detection of submarines from sonar signals, etc.

The success of neural networks in these and other pattern-recognition/classification problems of practical interest has sparked a lot of research in hardware implementation of neural nets. Several chip manufacturers (AT&T, Intel, Philips) have already come up with VLSI chips which implement small neural nets. Other possibilities such as optical implementation are currently being studied by many groups.

Although neural networks have been successfully used in a number of practical pattern-recognition problems, the theory of the working of feed-forward nets is still not well developed. For this reason, the designing of a network for a particular task is usually done by a trial-and-error procedure and there is not much understanding of what a particular network can do and what it cannot do. More research on the development of a theory of the working of these networks would be extremely useful.

5. Concluding remarks

In the preceding pages, I have attempted to provide an introduction to the rapidly developing field of neural network research. This field is still in its infancy and it is not clear at this stage what this line of research will eventually lead to. Thus, it is too early to ask questions like whether research in this area will some day produce a theory of the working of the human brain. This subject, however, shows a lot of promise and I believe that an interdisciplinary effort in this field involving researchers from neurobiology, physics, computer science and engineering will lead to many interesting and useful developments in the future.

References

1. McCulloch, W. S. and Pitts, W. A. — *Bull. Math. Biophys.*, 1943, **5**, 115–133.

2. Hopfield, J. J. — *Proc. Natn. Acad. Sci. USA*, 1982, **79**, 2554–2558, and 1984, **81**, 3088–3092.

3. HEBB, D. O. *The organisation of behavior*, 1949, Wiley.

4 SOMPOLINSKY, H. In *Heidelberg Colloquium on glassy dynamics* (eds, I. Morgenstern and J. L. van Hemmen), 1987, Springer-Verlag.

5. KOHONEN, T. *Self organisation and associative memory*, 1984, Springer-Verlag.

6. PERSONNAZ, L. GUYON, I. AND DREYFUS, G. *J. Phys. Lett.*, 1985, **46**, 359–365.

7. ROSENBLATT, F. *Principles of neurodynamics*, 1961, Spartan Press.

8. MINSKY, M. AND PAPERT, S. *Perceptorons, An introduction to computational geometry*, 1969, MIT Press.

9. RUMELHART, D. E., MCCLELLAND, J. L. AND the PDP Group *Parallel distributed processing*, 1986, MIT Press.

10. HINTON, G. E. *Artif. Intell.*, 1989, **40**, 185.

11. AMIT, D. J. *Modeling brain function: The world of attractor neural networks*, 1989, Cambridge University Press.

12. NADAL, J. P., TOULOUSE, G., CHANGEUX, J. P. AND DEHAENE, S. *Europhys. Lett.*, 1986, **1**, 535–542.

13. PARISI, G *J. Phys. A*, 1986, **19**, L617–L623.

14. KRISTAN, W. B. In *Information processing in nervous systems* (eds, Pinsker, H. M. and Wills, W. D.), 1980, Raven Press.

15. SOMPOLINSKY, H. AND KANTER, I. *Phys. Rev. Lett.*, 1986, **57**, 2861–2864.

16. KLEINFELD, D. *Proc. Natn. Acad. Sci. USA*, 1986, **83**, 9469–9473.

17. BUHMANN, J. AND SCHULTEN, K. *Europhys. Lett.*, 1987, **4**, 1205–1209.

18. AMIT, D. J. *Proc. Natn. Acad. Sci. USA*, 1988, **85**, 2141–2145.

19. DAMASIO, A. R., DAMASIO, H. AND VAN HOESEN, G. W. *Neurology*, 1982, **32**, 331.

20. VIRASORO, M. A. *Europhys. Lett.*, 1988, **7**, 293–298.

21. KNOWLES, W. D., TRAUB, R. D., WONG, R. K. S. AND MILES, R. *Trends Neurosci.*, 1985, **8**(2), 73–79.

22. MEZARD, M., PARISI, G. AND VIRASORO, M. A. *Spin glass theory and beyond*, 1987, World Scientific.

23. AMIT, D. J., GUTFREUND, H., AND SOMPOLINSKY, H. *Phys. Rev. Lett.*, 1985, **55**, 1530–1533 and *Ann. Phys.*, 1987, **173**, 30–67.

24. DERRIDA, B., GARDNER, E. AND ZIPPELIUS, A. *Europhys. Lett.*, 1987, **4**, 167–173.

25. GARDNER, E., *J. Phys. A*, 1988, **21**, 257–269.

26. HERTZ, J. A., KROGH, A. AND THORBERGSSON, G. I. *J. Phys. A*, 1989, **22**, 2133–2150.

27. HANSEL, D. AND SOMPOLINSKY, H. *Europhys. Lett.*, 1990, **11**, 687–692.

28. GRAUBARD, S. J. (ed.) *The artificial intelligence debate*, 1988, MIT Press.

29. HOPFIELD, J. J. AND TANK, D. W. *Science*, 1986, **233**, 625–633.

30. SEJNOWSKI, T. J. AND ROSENBERG, C. R. *Complex Systems*, 1987, **1**, 145

31. LAPIDES, A. AND FARBER, R. *Non-linear signal processing using neural networks: Predictions and systems modeling*, Technical Report LA-UR-87-2662, Los Alamos National Laboratory, 1987.

32. BRUNEK, S., BOHR, H. AND LAUTRUP, B. Preprint, Neils Bohr Institute, Copenhagen, 1988.

Published by N. M. Malwad, Executive Editor, Journal of the Indian Institute of Science, Bangalore 560 012; Typeset at Eastern Fototype and printed at The Eastern Press, 4, Bannerghatta Road, Bangalore 560 029.