

neurocomputer can implement a particular specified neural network to the necessary degree of accuracy, or it cannot do so.

Generality

Some neurocomputers are suitable for implementing only certain neural-network architectures. Other neurocomputers can implement essentially any architecture. Given a specific list of architectures, generality is the binary measurement of whether the neurocomputer can or cannot implement those network types.

Software interface provisions

For software programs to be able to call neural networks as subroutines, it is necessary to have software routines that can be linked with the user software and then called whenever needed to control the neurocomputer. The simplicity and ease of use of these software interface provisions are of great importance.

Configuration provision

For a computer to be used, it must be configured to run the desired neural networks. There are two basic approaches to this problem. First, most of the commonly used neural networks should be available in highly efficient (in other words, microcoded or otherwise optimized) prepackaged form. This coding efficiency

ensures that the networks used most often execute at maximum possible speed on the hardware. As with software interface provisions, configuration provisions are best judged on a relative-merit basis. The availability of a general-purpose neural network description language capable of describing a large percentage of neural-network architectures efficiently for use with the neurocomputer is an important consideration

Table 2 shows the different neurocomputers built so far, their performance in terms of speed and capacity, and the year introduced.

Conclusions

We discussed here the different approaches for implementing neurocomputers. The future development of neurocomputers will be governed by the development of accurate mathematical models for artificial neurons. Since it is predicted that the models will involve more of analog computations, it is hoped that more emphasis will be laid on analog hardware design and photonic computing for the design of efficient neurocomputers.

1. Lippman, R., *IEEE ASSP Mag.*, April 1987, pp. 4–22.
2. Hecht-Nielsen, R., *IEEE Spectrum*, March 1988, 36–41.
3. Mahadevan, Indu and Patnaik, L. M., *J. Parallel Comput.*, 1992, **18**, 401–413.
4. Mohan Kumar, J., Ph D Dissertation, Indian Institute of Science, Bangalore, 1992.
5. Patnaik, L. M. and Mohan Kumar, J., *J. Comput. Electr. Eng.* (to appear).

Mathematical modelling of neurons and neural networks

Chandan Dasgupta

Department of Physics, Indian Institute of Science, Bangalore 560 012, India and Jawaharlal Nehru Centre for Advanced Scientific Research, Bangalore 560 012, India

The basic concepts and techniques involved in the development and analysis of mathematical models for individual neurons and networks of neurons are reviewed. Some of the interesting results obtained from recent work in this field are described. The current status of research in this field in India is discussed.

THE development of an understanding how the human brain functions is one of the most important and

challenging tasks faced by modern science. During recent years, a great deal of progress has been made in experimental neurobiology. At the same time, efforts towards the development of mathematical models for some of the experimentally observed phenomena have gained momentum. Neurosciences cover a vast range of phenomena extending from the molecular to the behavioural level. In this paper, I have made an attempt to provide a brief overview of the concepts and techniques involved in the development and analysis of

mathematical models of individual neurons and neural networks. The molecular biophysical aspects of neurobiology are not dealt with here. To begin with, the basic framework for a mathematical description of the functioning of individual neurons is described and illustrated with a number of results obtained from recent studies. The next section is devoted to modelling of networks consisting of a large number of neurons. This kind of modelling has attracted a great deal of attention in the recent years. A large number of researchers from widely varying disciplines, such as neurobiology, computer science, electrical engineering, physics and mathematics, are currently working in this general area. Instead of trying to provide a comprehensive review of this newly emerging interdisciplinary field, I have concentrated in the section on neural networks only on those aspects which are relevant to neurobiology. This section contains a description of the basic ingredients of neural network modelling, a rough classification of the existing models on the basis of the degree of 'biological realism' incorporated in the modelling and a review of some of the important results obtained from recent studies in this area. Finally, the status of research in this field in India is discussed briefly.

Modelling of the individual neurons

Although neurons in the central nervous system show a wide range of variation in their morphology and function, the basic mechanism underlying their working is essentially the same for all of them. Neurons function electrochemically through the passage of different ions (Na^+ , K^+ , Ca^{2+} , Cl^- , etc.) through the cell membrane. It is, therefore, possible to model a neuron as an electrical device by constructing an equivalent electrical circuit. A segment of a cell membrane is represented in this circuit by several electrical elements. These are: (i) an electromotive force (or in other words, a battery) associated with each active ion channel, arising from the difference in ion concentration maintained across the membrane; (ii) a membrane conductance associated with each active ion channel, describing the passage of specific ions across the membrane; (iii) a leakage resistance that describes the passive leakage of unspecified ions across the membrane; (iv) a transmembrane capacitor describing the capacitance between the inner and outer walls of the membrane; and (v) internal and external resistances describing the transport of charge along the inner and outer walls of the membrane. Thus, a number of parameters are needed to characterize the electrical properties of a membrane element. Since these parameters have different values in different regions of the cell, an equivalent electrical circuit is usually obtained by dividing the cell into a

number of compartments and using different sets of parameters to describe the electrical properties of the membrane in different compartments. The appropriate set of circuit equations are then solved to determine the electrochemical properties of the cell. An important feature which complicates the analysis of the circuit equations arises from the fact that the values of some of the membrane parameters (most importantly, the membrane conductances) are themselves dependent on the membrane potential. An example of the kind of dependence to be incorporated in the model may be found in the classic work of Hodgkin and Huxley¹ on the mechanism of generation of the action potential. Such dependences make the circuit equations nonlinear and, therefore, difficult to solve analytically. These equations are usually solved numerically by using computers. A good example of this kind of computer modelling may be found in the work of Traub and Miles² on the pyramidal cells of the hippocampus. They have developed a computer model of hippocampal CA_3 neurons, in which ionic, ohmic and capacitive currents are calculated explicitly in 41 compartments representing the soma, the initial segment and the dendrites. The dendritic tree is represented in this model as a collection of equivalent cylinders. Somatic active currents are modelled using Hodgkin-Huxley-like Na^+ and fast K^+ conductances. Dendritic conductances include a Ca^{2+} conductance and a slow calcium-activated potassium conductance. This model is found to reproduce quite accurately the complex slow and fast spikes observed experimentally during the intrinsic bursting of the CA_3 cells. It is clear from this work and other studies of a similar nature that the basic processes involved in the electrochemical behaviour of individual neurons are fairly well understood now. Packages of computer programmes for simulating the electrical properties of isolated neurons are now available³ and integrated circuits which reproduce quite accurately the electrical characteristics of neurons have recently been fabricated⁴.

The models described above are deterministic, in the sense that they always produce the same outcome for a given set of parameter values and initial conditions. In reality, however, the behaviour of neurons is not completely predictable. The experimentally observed response of a neuron to a particular set of externally specified conditions is not always the same. It is, therefore, necessary to include an element of stochasticity (noise) in the dynamics of neurons. An interesting example that illustrates the importance of noise in neural dynamics may be found in a recent work of Longtin *et al.*⁵. They consider an experiment in which a single auditory nerve fibre of a squirrel monkey is stimulated with a sinusoidal sound stimulus of period T_0 and the time intervals between successive spikes generated by the neuron are recorded. The histogram of

the distribution of the interspike interval shows sharp peaks at integral multiples of T_0 . The probability of the interspike interval having a value nT_0 ($n = 1, 2, 3, \dots$) is found to fall off exponentially with increasing n . In order to explain these observations, Longtin *et al.*⁵ consider a very simple model that involves the motion of a particle in a one-dimensional bistable potential in the presence of a periodic driving force and a Gaussian white noise. The two minima of the bistable potential used in the model are supposed to represent the quiescent and firing states of the neuron. Their calculations show that the distribution of the time interval between successive transitions between the two minima of the potential exhibits all the substantive features of the experimentally observed interspike interval distribution. This simple calculation, thus, illustrates nicely the importance of including a noise term in the description of neural dynamics. It also serves the important purpose of showing that simple models in which many microscopic details are left out may indeed be adequate for explaining certain features of the experimentally observed behaviour of neurons.

It is clear from the discussion above that a lot of progress has been made in the modelling of the electrochemical properties of isolated neurons. This, however, is only a small part of the full story because neurons in the nervous system do not function in isolation. They receive inputs from and provide outputs to many other neurons through synaptic junctions. It is, therefore, necessary to look at models of the collective behaviour of networks of neurons in order to develop an understanding of brain functions. Certain properties of the individual neurons which govern the mechanism of communication between different neurons play an important role in determining the network behaviour. Although the basic processes involved in synaptic transmission of signals from one neuron to the other are fairly well understood now, many details are yet to be worked out. In particular, many questions related to the development of the dendritic tree, formation and subsequent modification (through learning) of synaptic junctions, specific neurotransmitters/receptors involved in synaptic plasticity, and so on, remain unanswered at the present time. This is one of the most active areas of current research and a lot of new results are expected in the near future.

Modelling of neural networks

The subject of neural network modelling has witnessed a spectacular growth in the recent years. It is a truly interdisciplinary field in which researchers from a wide range of backgrounds are involved. Since the term neural network carries different meanings for researchers from different disciplines, it would be useful

to begin with a fairly general definition of this term. A neural network may be defined as a large, highly interconnected assembly of *model neurons* which are computing elements with a specified input-output relation. Each neuron in the network receives inputs from and provides output to a large number of other neurons. The collective dynamics of the network endows it with nontrivial computing abilities. The connectivity of the network is specified by a synaptic interaction matrix. The computations to be performed by the network are encoded in this matrix through an appropriate learning process. The time evolution of the network is governed by an assumed dynamic rule which specifies how the state of a model neuron depends on the net input it receives from the other neurons. A time-persistent state reached by the network in the course of its time evolution represents the result of the computation performed by it.

There are several reasons for the current interest in models of this kind. In the neurobiological context, these systems serve as models (highly simplified in most cases) of some of the collective computational properties of the brain. The properties of the model neurons and their interconnections assumed in neural network models do resemble, although in a highly schematic way in most cases, some of the characteristics of their biological counterparts. Also, these models exhibit a number of features (such as parallel processing, fault tolerance, ability to learn, etc.) which are believed to be essential elements of biological computation. Physicists are interested in these models mainly because they are examples of interacting systems with complex dynamic behaviour. These models form an important part of a newly emerging subfield of physics that involves studies of complexity. At a more practical level, neural networks provide a new paradigm of parallel computation with numerous applications in practical problems involving pattern recognition and multivariable optimization. This is the primary reason for the interest of electrical engineers and computer scientists in this subject.

In this paper, I will concentrate on neural network modelling of neurobiological phenomena. A question that arises at the very beginning of any attempt at developing a model of this kind is that how much neurobiological detail should one incorporate in the model. Opinions on this issue vary a great deal. Many researchers in neurobiology argue that models in which many biological details are left out are much too simple to be of any use in the study of real neurobiological phenomena. On the other hand, many other researchers in this field feel that some of the microscopic details left out in the simpler neural network models may not be crucial in understanding some of the collective properties of biological networks. The debate on this issue is continuing. As a result, the existing neural network models show a great deal of variation in the

degree of 'biological realism' incorporated in the modelling. Using this aspect as a criterion, one may roughly classify the existing models into the following three categories.

Abstract models

In models of this kind, all microscopic details about the working of the individual neurons are left out. The state of a neuron in these models is described by a single variable v which may be interpreted as the firing rate (the rate at which an action potential is generated). The value of v for a particular model neuron is supposed to be a simple sigmoid function of the net input (membrane potential in the biological interpretation) it receives from the other neurons in the network. In some models, a further simplification is obtained by approximating the sigmoid function by a step function. The variable v in these models takes only two values, 0 and 1, representing, respectively, the quiescent and firing states of the model neuron. The pattern of synaptic connections and the values of the elements of the synaptic matrix used in models of this kind bear little resemblance to what is observed in biological networks. These network characteristics are usually chosen to maximize the efficiency of the network in performing the computational task it is designed for. The main advantage of networks of this kind is their relative simplicity, which often makes them amenable to analytic study. Physicists have played an important role in the development of analytic theories of such networks. A number of analytic methods developed by theoretical physicists in studies of statistical mechanics of disordered systems have been extremely useful in analysing the behaviour of these networks⁶. In cases where analytic treatments are too difficult, numerical simulations of the behaviour of such models can be performed rather easily. Analog implementations of networks of this kind are also possible. Most of the work carried out by physicists and engineers in the field of neural networks involves models of this kind. However, the connection of these models with neurobiology is rather tenuous. The only thing common between these models and real neuronal networks existing in the brain is a possible similarity in their basic working principles. For this reason, abstract models of this kind are sometimes called *artificial neural networks* to distinguish them from biological networks.

Models of this kind may be loosely classified into two categories: feed-forward networks (sometimes called multilayer perceptrons) and multiply connected networks. Feed-forward networks have a layered structure in which neurons in the m th layer receive inputs from those in the $(m-1)$ th layer and provide inputs to those in the $(m+1)$ th layer. Networks of this kind are primarily used for tasks involving recognition and classification of

patterns. A subset of the data to be classified is used first to 'train' the network. The trained network can then generalize the information contained in the training set and use this information to classify new data with a high degree of accuracy⁷. Multiply connected networks, which have feedback loops in the synaptic connections, are used mostly as models of associative memory. A system behaves as an associative memory if it can retrieve patterns stored in it from 'hints', representing a partial knowledge of the stored information. Human memory is known to be associative in this sense. In neural network models of associative memory, the connection matrix is chosen in such a way that network states representing the stored memories become locally stable attractors of the underlying dynamics. An initial state close to one of these attractors, representing partial knowledge of the stored information, is driven to the attractor itself in the course of the collective dynamics of the network. The complete information is thus retrieved. The Hopfield model⁸ is perhaps the most well-known neural network model of this kind. Networks of a similar nature are also being used⁹ with a considerable amount of success in finding near-optimal solutions of complex optimization problems involving many variables. It is interesting to note that computations which are performed more efficiently by neural networks than by conventional digital computers are also the ones at which the human brain excels. This observation suggests that the basic principles underlying the functioning of these models may indeed have some similarities with those used by the brain for carrying out various computational tasks.

Intermediate level models

In models belonging to this category, each neuron is still treated as a simple input-output device with no internal structure. However, a number of experimentally observed features are incorporated in the architecture and synaptic connectivity pattern of the network. The main advantage of studying models of this kind lies in the fact that due to their relative simplicity, analytic treatments are possible in some cases and numerical simulations can be carried out rather easily. At the same time, one can make some contact with experiments because these models are at least plausible biologically. Of course, comparisons with experimental data can be carried out only at the qualitative level because these models are not realistic enough. Also, models of this kind are obviously not appropriate for descriptions of neurobiological phenomena, in which properties of individual neurons play an important role. Nevertheless, many researchers in this field believe that studies of models of this kind do provide a lot of information and insight about the collective behaviour of biological networks.

Examples of models of this kind may be found in a large number of recent studies¹⁰ of biologically plausible models of associative memory. These models go beyond the simple Hopfield model in trying to incorporate different experimentally observed features in the structure of the network. The features which have been incorporated in the modelling include:

1. Neuronal specificity, i.e. the so-called Dale hypothesis, which states that in most cases the efferent synapses of a particular neuron are either all excitatory or all inhibitory.
2. The presence of static and dynamic synaptic noises.
3. Limited analog depth of synaptic strengths.
4. Hierarchical (tree-like) structures in the network architecture and in the stored data.
5. Less-than-full connectivity of the network.
6. Low average (spatial and temporal) level of activity of the network.

Models which exhibit a selective erasure of old patterns stored in the memory (forgetting) as new information is learnt have also been constructed¹¹. A very interesting result has emerged from these studies. It has been established that the basic functional features of the simple Hopfield model remain, to a large extent, unaffected by the inclusion of these biological details. This observation lends support to the expectation that studies of relatively simple models may indeed lead to useful insights into some aspects of the functioning of the brain.

Another class of models which have received a lot of attention in the recent years involve the storage and recall of temporal sequences of patterns. These models use either a time delay mechanism or the presence of dynamic synaptic noise to generate a passage of the network through a specified sequence of patterns. Numerical simulations of these models have reproduced several features exhibited by biological central pattern generators¹². These are neural groups which control the muscles involved in a variety of rhythmic activities (such as locomotion, swimming and chewing) by repeated generation of a specific sequence of patterns.

We have recently carried out a study¹³ which may be classified in this category of modelling. This work involves the development and simulation of a neural network model for kindling, which describes the process of generation of epilepsy in laboratory animals by repeated electrical stimulation of certain parts of the forebrain. The starting point of this study is a biologically plausible neural network that generates complicated limit cycles in the absence of any external stimulation. The resulting small-amplitude oscillations of the network activity are assumed to model the resting EEG. The effect of the external shocks is represented by changes in the postsynaptic potentials of neurons close to the implanted electrodes. These changes cause a large number of these neurons to fire. The excess activity of

the network is assumed to lead to the formation of new excitatory synaptic connections through a Hebbian mechanism¹⁴ of synaptic plasticity. The formation of these new synapses causes more neurons to fire and this 'snowballing' process leads to a rapid increase in the network activity, culminating in the epileptic state. Simulations of this model reproduce and provide some understanding of a large number of experimental observations.

Other neurobiological phenomena for which models of a similar nature have been developed include synchronization of oscillations in the visual cortex¹⁵ and oscillatory and chaotic EEG patterns observed in the olfactory system¹⁶.

Realistic models

In modelling of this kind, all known biophysical details of the individual neurons and synaptic connections among them are incorporated in the model. Due to their extreme complexity, models of this kind are not amenable to analytic treatments. Numerical simulations provide the only way of studying their properties. Such simulations require an extensive amount of computing resources. Consider, for example, a relatively small network consisting of 1000 neurons. As discussed earlier, one needs about 50 variables for a realistic description of the behaviour of a single neuron. Thus, a simulation of the properties of this network would involve solving a set of about 50,000 coupled nonlinear differential equations. This is a difficult task, even for the most sophisticated supercomputers available today. The complexity of these models sometimes makes interpretations of the simulation results difficult. Due to this reason, some researchers have raised questions about how much insight is actually gained from simulations of models which are almost as complicated as the physical system being studied. Nevertheless, the fact remains that modelling of this kind is necessary if one wants to make direct quantitative comparisons with experimental data. Work of this nature, thus, serves the important purpose of providing a way of validating theoretical ideas about how networks of neurons actually function. An example of modelling of this kind may be found in the work of Traub and collaborators^{2, 17} on rhythmic population oscillations observed in both the hippocampus (the hippocampal theta rhythm) and *in vitro* hippocampal slice preparations. Their model consists of 9000 excitatory pyramidal cells and 900 inhibitory cells. The parameters describing the properties of the pyramidal cells are chosen such that when they are excited, they generate intrinsic bursts with long, intrinsic hyperpolarization after the burst. Excitatory synaptic connections among the pyramidal cells are assumed to be sparse and random. These connections are spatially restricted, in the sense that the

probability of two pyramidal cells being connected decreases exponentially with the distance between the cells. Half of the population of inhibitory cells is assumed to produce fast inhibition and the remaining half produces slow inhibition. Axons of the inhibitory cells are assumed to be spatially more restricted than those of the pyramidal cells. Each neuron in the network has, on an average, 20 excitatory inputs and 20 inhibitory inputs. All these features incorporated in the model are consistent with the existing anatomical and physiological data. Simulations of the collective behaviour of the network are found to reproduce quite accurately a large number of experimentally observed results. From studies of how the network behaviour is affected by changes in the values of the various parameters appearing in the model, these researchers are able to determine the dependence of the emergent properties of the network on intrinsic cellular characteristics and on the connectivity and strength of both excitatory and inhibitory synapses. This study also leads to a number of specific predictions which can be tested experimentally.

Until recently, work of this nature could be carried out only at a few places in the world. Due to recent advances in computer technology, computing resources necessary for carrying out simulations of this kind have become more readily available. As a result, a large number of groups are currently getting involved in the development of models of this type¹⁸. This appears to be one of the fastest growing areas of neuroscience.

Status of research in India

Several groups in India are currently working in the field of neural networks. These groups, consisting mostly of physicists, mathematicians, engineers and computer scientists, are primarily working on the so-called artificial neural networks. The problems being studied include designing networks for specific computational tasks, testing the performance of different networks and analysis of the collective properties of the

various network models. Apart from a few isolated attempts, essentially no research is being carried out in India on modelling of neurobiological phenomena. This is probably due to the fact that there is hardly any interaction between the people involved in the modelling and those doing experimental research in neurobiology. A close collaboration between these two groups of researchers is essential for the development of a successful research programme in the area of neurobiological modelling. It would be worthwhile to initiate a programme which would foster the development of such collaborations.

1. Hodgkin, A. L. and Huxley, A. F., *J. Physiol. (London)*, 1952, **117**, 500.
2. Traub, R. D. and Miles, R., *Neuronal Networks of the Hippocampus*, Cambridge University Press, Cambridge, 1991.
3. *Trends in Neuroscience*, 1992, **15**, No. 12.
4. Mahowald, M. and Douglas, R., *Nature*, 1991, **354**, 515.
5. Longtin, A., Balsara, A. and Moss, F., *Phys. Rev. Lett.*, 1991, **67**, 656.
6. Amit, D. J., *Modelling Brain Functions*, Cambridge University Press, Cambridge, 1989.
7. Rumelhart, D. E. and McClelland, J. L. (eds.), *Parallel Distributed Processing*, MIT Press, Cambridge, Massachusetts, 1986.
8. Hopfield, J. J., *Proc. Natl. Acad. Sci. USA*, 1982, **79**, 2554; 1984, **81**, 3088.
9. Hopfield, J. J. and Tank, D. W., *Science*, 1986, **233**, 625.
10. Amit, D. J. and Treves, A., *Proc. Natl. Acad. Sci. USA*, 1989, **86**, 7871.
11. Nadal, J. P., Toulouse, G., Changeux, J. P. and Dehaene, S., *Europhys. Lett.*, 1986, **1**, 535.
12. Kleinfeld, D. and Sompolinsky, H., in *Methods in Neuronal Modelling* (eds. Koch, C. and Segev, I.), MIT Press, Cambridge (Massachusetts), 1989.
13. Mehta, M. R., Dasgupta, C. and Ullal, G. R., *Biol. Cybernet.* 1993, **68**, 335.
14. Hebb, D. O., *The Organization of Behaviour*, John Wiley, New York, 1949.
15. Sompolinsky, H., Golomb, D. and Kleinfeld, D., *Proc. Natl. Acad. Sci. USA*, 1990, **87**, 7200.
16. Freeman, W. J., *Int. J. Bifurcation and Chaos*, 1992, **2**, 451.
17. Traub, R. D., Miles, R. and Wong, R. K. S., *Science*, 1989, **243**, 1319.
18. Wilson, M. A. and Bower, J. M., in *Methods in Neuronal Modelling* (eds. Koch, C. and Segev, I.), MIT Press, Cambridge, Massachusetts, 1989.