

# Neuroscience Year

Supplement 1 to the  
Encyclopedia of Neuroscience

Edited by  
George Adelman

With 77 Figures

*A Pro Scientia Viva Title*

BIRKHÄUSER  
Boston · Basel · Berlin



First printing, 1989

Library of Congress Cataloging-in-Publication Data  
Neuroscience year. 1989: supplement 1 to the Encyclopedia of  
neuroscience / George Adelman, editor.

p. cm.

"A Pro scientia viva title."

Bibliography: p.

Includes index.

ISBN 0-8176-3383-9

1. Neurology. 2. Neuropsychology. I. Adelman, George, 1926-

RC329.N48 1989

616.8—dc19

88-39373

CIP

© 1989 by Birkhäuser Boston, Inc.

Copyright is not claimed for works by U.S. Government employees.

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Birkhäuser Boston, Inc., 675 Massachusetts Avenue, Cambridge, Massachusetts 02139, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use of general descriptive names, trade names, trademarks, etc. in this publication, even if the former are not especially identified, is not to be taken as a sign that such names, as understood by the Trade Marks and Merchandise Marks Act, may accordingly be used freely by anyone.

While the advice and information in this book are believed to be true and accurate at the date of going to press, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Typeset by Arcata Graphics/Kingsport, Kingsport, Tennessee.

Printed and bound by Arcata Graphics/Halliday, West Hanover, Massachusetts.

Printed in the United States of America.

9 8 7 6 5 4 3 2 1

ISBN 0-8176-3383-9

ISBN 3-7643-3383-9

# Computational Neuroscience

Terrence J. Sejnowski, Christof Koch, and Patricia S. Churchland

The ultimate aim of computational neuroscience is to explain how electrical and chemical signals are used in the brain to process information. This goal is not new, but much has changed in the last decade: more is known about the brain because of advances in neuroscience, much more computing power is available for performing simulations of neural systems, and new insights are available from simplified models of large networks of neurons. Brain models connect the microscopic level accessible by molecular and cellular techniques with the systems level accessible by studying behavior.

Understanding the brain is a challenge that is attracting a growing number of scientists from many disciplines. While there has been an explosion of discoveries over the last several decades concerning the structure of the brain at the cellular and molecular levels, we do not yet understand how the nervous system enables us to see and hear, to learn skills and remember events, to plan actions and make choices. Modeling of the brain represents a valuable technique for generating hypotheses about brain function.

## Computation

Although the digital computer in our age has become the prototypical example of a system that computes, many analog devices such as slide rules, optical Fourier analyzers, and even marbles rolling downhill can also compute. In each of these examples, the states of the physical system can be mapped onto the states in a more abstract algorithm that solves a class of computational problems (Sejnowski, et al., 1988). States of the nervous system represent events and states of affairs in the world. Computational explanations for states of the nervous system differ from mechanical or causal explanations in that a mapping can be made from the states of the brain onto information-bearing states of an abstract algorithm that solves a computational problem.

Existing computational devices invented by man may not be good guides to the computational solutions evolved by nature at least for the reason that evolutionary changes are made within the context of a design and architecture that already is in place. Evolution cannot start from scratch, even when the optimal design would require that course. As Francois Jacob has remarked in his book, "The Possible and the Actual," evolution is a tinkerer, and it fashions its modifications out of available materials, limited by earlier decisions. Moreover, any given capacity, such as binocular depth perception, is part of a much larger package subserving sensory-motor control and survival in general.

## Levels of analysis

An influential framework for analyzing information processing was articulated by David Marr in 1982. Three levels were characterized: (1) The *computational* level of abstract problem analysis, wherein the task (e.g. determining depth from binocular images) is decomposed into its fundamental constituents; (2) the level of the *algorithm*, which specifies a formal procedure by which, for a given input, the correct output could be given, and the task thereby performed; and (3) the level of physical *implementation* of the computation. In a programmable digital computer the computational level is independent of the algorithmic and implementation levels, in the sense

that the same computational problem can be solved by many different algorithms, and in turn each algorithm can be implemented with digital circuits constructed with many different technologies.

Unlike a digital computer, which is general purpose and can be programmed to run any algorithm, the brain appears to be a collection of special purpose systems that are limited in their flexibility, but very efficient at performing their tasks. In contrast to the doctrine of independence of computation from implementation, current research in neuroscience suggests that considerations of implementation play a vital role in the kinds of algorithms that are devised and the kinds of computational insights that are available. Even Marr, who advocated a top-down approach, was highly influenced by neurobiological considerations.

Another consideration is the relationship between levels of analysis and structural levels. At each structurally specified stratum from synapses to systems we can raise the question: What does it contribute to the wider, computational organization of the brain? Thus, the range of implementation levels in the brain is probably accompanied by an equally rich range of algorithms and task descriptions. Rather than seeking a single computational explanation we must expect a spectrum of explanations depending on the spatial and temporal scale of the phenomenon.

## Realistic brain models

Most of our information about the representation of sensory information and motor commands is based on recording from single neurons. This technique is revealing but it is also confining insofar as it biases us toward thinking about the cellular level rather than the subcellular or circuit levels. We are especially in need of techniques that would allow us to monitor populations of neurons. Modeling promises to be an important adjunct to these experimental techniques and is essential in addressing the conceptual issues that arise when studying information processing in the brain. The advantages of brain models are varied: (1) A model can make more accessible to scrutiny the consequences of a complex, nonlinear brain system with many interacting components. (2) New phenomena may be discovered by comparing the predictions of a simulation to experimental results and new experiments can be designed based on these predictions. (3) Experiments that are difficult or even impossible to perform in living tissue, such as the selective lesion of particular channels, synapses, neurons, or pathways, can be simulated using a model.

As knowledge in neuroscience accumulates at the cellular and molecular levels, it is tempting to incorporate all that is known into a model that aims to reproduce as much of the nervous system as possible. One problem with this approach is that a genuinely perfect model, faithful in every detail, is likely to be as incomprehensible as the nervous system itself. Another problem is that a model based on incorrect or incomplete knowledge may give misleading results. Nonetheless, this class of models, which we call realistic models, can provide valuable insights into the emergent properties of the nervous system, such as rhythmic pattern generation in neural circuits.

The range of spatial scales over which the nervous system has been explored spans over eight orders of magnitudes from

molecular dimensions measured in angstroms to fiber tracts that span many centimeters. Physiological time scales span over ten orders of magnitude from fractions of a millisecond in the case of gating of single ion channels to days or weeks for biophysical and biochemical events underlying memory (McNaughton and Morris, 1987). Organizational principles emerge on a hierarchy of spatial levels and time scales that are directly relevant to the function of the nervous system (Table 1). A few of these principles will be summarized here to serve as a framework for discussing concrete examples.

No single neural model can be expected to span all levels, and an essential feature at one level of organization may be an insignificant detail at another. The multiplicity of levels of organization is a feature not only of neuroscience, but also of physics and chemistry, where explanations of phenomena on distinct levels or organization are more developed. Assuming a comprehensive "theory of the brain" does emerge, it will involve establishing a successive and overlapping chain of explanations from the lowest levels to the highest, encompassing the various spatial, temporal, structural, and computational levels.

It is rare for a model to jump over many levels, and the most successful models typically link only neighboring levels. For example, the Hodgkin-Huxley ionic model of the initiation and propagation of action potentials links measurements of ionic currents in the whole axon to the kinetics of single ion channels, although such channels were only hypothetical at that time in 1952 when the model was proposed. Assumptions beyond the data available at one level are sometimes needed in order to reach a better understanding at that level. It is only with the introduction of single channel patch clamp recording techniques developed decades after the Hodgkin-Huxley model that their assumptions regarding the nature of ion channels could be verified by Bert Sackmann and Erwin Neher in 1976. Wilfrid Rall in the 1970s studied the detailed spread of current in dendrites using cable models, and more recently the effects of voltage-dependent processing have been included in these models by a number of investigators.

Sensory information tends to be organized in spatial maps that are topographically organized in brain structures that are often laminated, such as the retina and cerebral cortex. Topographic maps and laminae are special cases of a more general principle, which is the exploitation of geometry in the design of information processing systems. Spatial proximity may be an efficient way for biological systems to get together the information needed rapidly to solve difficult computational

problems. For example, it is often important to compute the differences between similar features of a stimulus at nearby locations in space. By maintaining neighborhood relationships the total length of the connections needed to bring together the signals is minimized. Maps also make it easier to synchronize the timing of information flow through an array of neurons. Perhaps the most successful realistic model of spatial processing in maps is the Hartline-Ratliff model of lateral inhibition in the *Limulus* lateral eye (Ratliff, 1974).

### Simplifying brain models

Models of the brain based on simplifying assumptions have a role in computational neuroscience that is very different from realistic models. Textbook examples in physics and other simplified problems that admit exact solutions are usually unrealistic, but they often illustrate important principles. In neuroscience, the study of simple models can provide a conceptual framework for studying the complex organizations in nervous systems. They may also help us to isolate the crucial issues and to understand the limitations of these kinds of systems.

The class of models that is currently being investigated under the general headings of connectionist models, parallel distributed processing models, and "neural networks" is of this second type, which we shall hereafter refer to as simplifying brain models. These models abstract from the complexity of individual neurons and the patterns of connectivity in exchange for analytic tractability (Hopfield and Tank, 1986). These models are being investigated as prototypes of new computer architectures or as models for psychological phenomena. Nonetheless, many of the results are applicable to the brain.

One of the best studied architectures is the class of layered feedforward networks. In this architecture, information is coded as a pattern of activity in an input layer of model neurons and is transformed by successive layers receiving converging synaptic inputs from preceding layers. Three findings are of significance for brain models: (1) Even systems with only a few intermediate layers have enormous power in representing complex nonlinear functions. (2) The performance of a network in specific problem domains (such as visual and speech processing) depends critically on how the incoming information is represented by the neurons (such as the type of preprocessing) and the symmetries in the pattern of connections. (3) For difficult problems the processing units in the middle or "hidden" layers generally encode many different combinations of input variables using a semidistributed type of representation (Sejnowski and Rosenberg, 1987). By combining the power of these models with further constraints from neurophysiology and neuroanatomy it may be possible to interpret some of the properties that have been observed from single-unit recordings, as we illustrate in the next section.

These simplifying brain models also make an important bridge to computer science and other disciplines that study information processing. Issues such as convergence of the network to a stable solution, the amount of time needed for the network to achieve a solution, and the capacity of networks to store information are being investigated in simplifying models in ways that are not at present feasible with realistic models. The scaling of these properties with the size of the network is crucially important to the practical feasibility of the model and its plausibility as a brain model.

### Technology for brain modeling

Computational brain models are almost always simulated on digital computers. Computers are getting faster, but they must

**Table 1.** Levels of Investigation of the Nervous System.\*

Structural	Analysis	Measurement
Channels	Implementation	Single channel
Synapses	Algorithmic	Single cell
Dendrites	Computational	Multiple cell
Local circuits		EEG/ERP/PET/MRI
Glomeruli		Psychophysics
Ganglia		Behavior
Nuclei/Maps		
Systems		

\* *Structural* elements in the nervous system can be studied at many spatial scales. *Measurements* can be made with many physiological techniques at each of these structural levels. *Analysis* of the information processing aspects of the nervous system can occur on three levels of abstraction (see text). Models typically link neighboring structural levels and are based on measurements made with several techniques.

perform the many parallel operations in the brain one at a time and are many orders of magnitude too slow. Parallel computers with thousands of processors are being developed, but are still inadequate. A new approach towards simulating biological circuitry is being pioneered by Carver Mead, who is constructing hardware devices with components that directly mimic the circuits in the brain. The severe physical restrictions imposed on the density of wires and the cost of communications in electronic circuits are similar to the constraints that are imposed on biological circuits. Fast hardware can deliver the computing power necessary to evaluate the performance of models in real time. This approach, which he terms "synthetic neurobiology," allows for the rapid determination of the strengths and limitations of a theory.

Mead uses analog subthreshold CMOS VLSI (very large scale integrated circuit) technology. Several chips that implement simplifying models of visual information processing have already been produced that are highly efficient. A "retina" chip computes the spatial and temporal derivative of arbitrary images projected onto an hexagonal array of 48 by 48 phototransistors—which are approximately logarithmic over 5 orders of magnitude of light amplitude—coupled via a horizontal resistive grid and injecting current into model "amacrine" cells that compute a temporal derivative. Similar circuits can be designed for computing optical flow in real time.

These VLSI chips and new techniques in optical information processing may lead to a new computing technology, sometimes called artificial neural systems, or neurocomputing. This technology for performing massively parallel computation could have a major influence on the next generation of research in computational neuroscience. For example, an analog VLSI model of a neuron that included conductance mechanisms, synaptic apparatus, and dendritic geometry could be produced in great quantities. These chips could be used as coprocessors in a conventional digital computer to greatly increase the speed of realistic simulations. If this technology is developed now, it should be possible to simulate our visual system in real time by the 21st century.

## Conclusions

A scientific field is defined primarily by its problem-space and its successful large-scale theories. Until there are such theories in computational neuroscience, the field is defined mostly by the problems it seeks to solve, and the general methodology and specific techniques it hopes will yield successful theories. Models of brain function driven primarily by functional considerations can provide only the most general guidance about what might be happening in the brain; conversely, models driven primarily by signal measurements and anatomy can easily miss those aspects of the signals that are relevant for information processing. Both realistic and simplifying models are being used to explore information processing by brain mechanisms at many different levels of structural organization, as indicated in Table 2.

Realistic and simplifying brain models have been distinguished to reveal their separate strengths and weakness. Neither type of model should be used uncritically. Realistic models require a substantial empirical database; it is all too easy to make a complex model fit a limited subset of the data. Simplifying models are essential but are also dangerously seductive; a model can become an end in itself and lose touch with nature. Ideally these two types of models should complement each other. For example, the same mathematical tools and techniques that are developed for studying a simplifying model could well be applied to analyzing a realistic model, or even

the brain itself. More accurately, the two types of models are really end points of a continuum, and any given model may have features of both. Thus, we expect future brain models to be intermediate types that combine the advantages of both realistic and simplifying models.

At this stage in our understanding of the brain, it may be fruitful to concentrate on models that suggest new and promising lines of experimentation, at all levels of organization. In this spirit, a model should be considered a provisional framework for organizing possible ways of thinking about the nervous system. The model may not be able to generate a full range of predictions owing to incompleteness, some assumptions may be unrealistic simplifications, and some details may even be demonstrably wrong. Nevertheless, if the computational model is firmly based on the available experimental data, it can evolve along with the experimental program and help to guide future research directions.

**Acknowledgments.** We are indebted to Drs. Francis Crick and Tomaso Poggio for many helpful comments on this paper. TJS is supported by a Presidential Young Investigator Award from the NSF and by grants from the Sloan Foundation, the Seaver Institute, the Lounsbery Foundation, the Air Force Office of Scientific Research, and the General Electric Founda-

**Table 2.** Other Models of Information Processing in the Brain

### Models in Vision:

- Hildreth EC, Hollerbach JM. In: *Handbook of Physiology*. Vol 3, Pt 2, *The Nervous System*. Plum F, ed. Washington, D.C.: American Physiological Society, 1987  
 Ullman S. *Ann Rev Neurosci* 9:1–26, 1986  
 Crick F. *Proc Nat Acad Sci (USA)* 81:4586–4590, 1984  
 Lehky SR, Sejnowski TJ. *Nature* 1988 333:452–454

### Models of Memory:

- Brown TH, Ganong AH, Kariss EW, Keenan Ch, Kelso SR. In: *Neural Models of Plasticity*. Byrne JH, Berry WO, eds. New York: Academic Press, 1988  
 Kohonen T. *Self Organization and Associative Memory*. New York: Academic Press, 1988  
 Donegan NH, Gluck MA, Thompson RF. In: *Models of Learning in Simple Neural Systems*. Hawkins RD, Bower GH, eds. *Psychology of Learning and Motivation*. Vol. 22. New York: Academic Press, 1988

### Models of Olfaction:

- Freeman WJ. *Mass Action in the Nervous System*. New York: Academic Press, 1975  
 Lynch G. In: *Perspectives in Memory Research* Gazzaniga MS, ed. Cambridge, Mass.: MIT Press, 1988  
 Bower JM, Wilson M, Nelson M. In: *Methods in Neuronal Modeling*. Koch C, Seger I, eds. Cambridge, Mass.: MIT Press, 1989

### Models of Sensorimotor Integration:

- Robinson DA. *Ann Rev Neurosci* 4:463–503, 1981  
 Arbib M. In: *Computational Neuroscience*. Schwartz E, ed. Cambridge, Mass.: MIT Press, 1989  
 Pellionisz P, Llinas R. *Neuroscience* 16:245–273, 1985.  
 Grossberg S, Kuperstein M. *Neural Dynamics of Sensory-Motor Control*. Amsterdam: North Holland, 1986

### Models of Development:

- Bear M, Cooper LN, Ebner FF. *Science*. 237:42–48, 1987  
 Linsker R. *Proc Nat Acad Sci (USA)* 83:8779–8783, 1986  
 von der Malsburg C, Willshaw D. *TINS* 4:80–83, 1981  
 Finkel LH, Edelman GM. In: *Synaptic Function*. Edelman GM, Gall WE, Cowan WM, eds. New York: Wiley, 1987

tion; CK is supported by a Presidential Young Investigator Award from the NSF and by grants from the Office of Naval Research, the Advanced Engineering program at the NSF, the Sloan Foundation, and by the Powell Lee Foundation; PSC is supported by a grant from the NSF and the James S. McDonnell Foundation.

**Further reading**

Sejnowski TJ, Koch C, Churchland PS (1988): Computational neuroscience. *Science* 241:1299-1306.

Hopfield JJ, Tank DW (1986): Computing with neural circuits: A model. *Science* 233:625-633

McNaughton BL, Morris RG (1987): Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends Neurosc* 10:408-415

Mead C (1989): *Analog VLSI and Neural Systems*. Reading, Mass.: Addison-Wesley

Ratliff F (1974) *Studies on Excitation and Inhibition in the Retina*. New York: Rockefeller Univ. Press

Sejnowski TJ, Rosenberg CR (1987): Parallel networks that learn to pronounce English text. *Complex Systems* 1:145-168