VOL. 4 NO. 1 of Seminars In The Neurosciences (Feb. 1992):
The Use Of Models In The Neurosciences
Guest Editor: Kenneth D. Miller

CONTENTS:

Introduction: The Use Of Models In The Neurosciences Kenneth D. Miller

A Quantitative Model of Phototransduction and Light Adaptation in Amphibian Rod Photoreceptors V. Torre, M. Straforini and M. Campani

Realistic Single-Neuron Modeling William W. Lytton and John C. Wathey

Modeling Hippocampal Circuitry Using Data From Whole Cell Patch Clamp and Dual Intracellular Recordings In Vitro
Roger D. Traub and Richard Miles

Models of Central Pattern Generators as Oscillators: The Lamprey Locomotor CPG Karen A. Sigvardt and Thelma L. Williams

Realistic Neural Network Models Using Backpropagation: Panacea or Oxymoron? S.R. Lockery

Models of Activity-Dependent Neural Development Kenneth D. Miller

Neural Modeling of a Visual Perceptual Task — A Case Study
Gerald Westheimer

An Introduction to Silicon Neural Analogs
M.A. Mahowald, R.J. Douglas, J.E. LeMoncheck, and C.A. Mead

Seminars in the Neurosciences is a review journal. Each issue is devoted to a single topic or theme of current interest to neuroscientists, and is edited by a guest editor appointed for that issue. The aim of each issue is to provide a co-ordinated, topical review of a selected subject that will be of interest to all levels of reader from the senior undergraduate to the researcher. Forthcoming topics (with guest editor): Dopamine (Trevor W. Robbins); Immune Responses in the Nervous System (Cedric S. Raine); Building the Architecture of the Nervous System (Andrew Lumsden). The Editorial Advisory Board for the journal is: Marianne Bonner-Fraser, J.P. Changeux, Klaus Peter Hoffman, John S. Kelly, Peter Kennedy, Motoy Kuno, Dale Purves, and Janis C. Weeks.

Six issues are published annually. SUBSCRIPTION RATES: Personal: $65, or 39 pounds (UK); Institutional: $130, or 78 pounds (UK). SINGLE ISSUES: $34, or 18 pounds (UK). Canada: add GST at current rate of 7%.

SEND ORDERS TO: Academic Press Limited, Foots Cray, Sidcup, Kent DA14 5HP, UK (telephone: 081-300-3222).

gone above and beyond her role as Executive Editor of this series, taking on serious reconstruction of two of the articles in addition to her normal thorough editing of all of the articles, and providing valuable guidance at many stages as this issue emerged. I thank her for her involvement.

References

[1] Atick JJ, Redlich AN (1991) What does the retina know about natural scenes? *Neural Comp.*, in press

[2] Bialek W, Rieke F, de Ruyter van Steveninck RR, Warland D (1991) Reading a neural code. *Science* **252**, 1854-1857

[3] Marr D (1982) *Vision.* W.H. Freeman, NY

[4] Bizzi E, Mussa-Ivaldi M, Giszter S (1991) Computations underlying the execution of movement: a biological perspective. *Science* **253**, 287-291.

[5] Brooks, RA (1991) New approaches to robotics. *Science* **253**, 1227-1232.

could not be significantly addressed in this issue, a theory may suggest computational principles that are embodied by the nervous system. Some recent elegant examples offer explanations of retinal receptive fields in terms of efficient and noise-resistant encoding of the visual information reaching the retina [1], or characterize the information about the world carried by real-time neural responses [2].

Today, as reflected in the articles in this issue, modeling and theory tend to begin from the biological structure being studied. This represents something of a shift from the influential program suggested by David Marr a decade ago [3]. He proposed that problems should be decomposed into three hierarchical elements: the computational task, the algorithm that implements the task, and the hardware that implements the algorithm. The emphasis was on understanding the task as a prerequisite for understanding neural function, much as understanding the aerodynamics of flight is a prerequisite for fully understanding the function of feathers. While these ideas remain important and fruitful (as one example, see [4]), there is a growing awareness that the computational task is often not cleanly separable from the tissue that carries it out (a point emphasized by Mahowald *et al.*). The computational tasks carried out by a massively parallel nervous system made up of large numbers of noisy and variable neurons, organized into many interacting but independent subsystems that each can directly link sensing and action [5], may differ fundamentally — not just at the level of algorithm or implementation — from the tasks we would assign a digital computer coping with the same environment.

Thus far I have avoided the term 'computational neuroscience.' Like all jargon, this is a useful term but should not be taken literally. Though we study the computations done by the nervous system, we also study many other aspects — the process of neural development, the biophysics of cellular activation, the conditions under which a piece of tissue generates an oscillation — without immediate reference to their role in a computation. We use computers heavily, and the tremendous growth of computing power in the last decade is one of the major contributors to the effort, but we also use pen and paper, mathematics, silicon. Perhaps the term should be 'theoretical neuroscience,' but this revives memory of failed efforts to transplant theory from physics directly into biology, rather than growing a theory in our own soil. 'Computational neuroscience' captures the excitement of the present wave of theoretical work in neuroscience, and names our effort; but let us not forget that the real scope of the project is larger.

I believe in this issue of *Seminars in the Neurosciences* you will find some of the excitement, promise, and scope of computational neuroscience, a field already far too large to be encompassed in a single volume or by any set of eight authors. It is a new field whose contributions to the tasks of neuroscience are growing rapidly; may you discover here a sense of its potential and power.

In assembling this issue, I have been aided greatly by Jennifer Altman. She has

'beginner's guides' to each type or level of model. They address questions such as: when is a model useful or not; how can a model aid a neuroscientist and what should one watch out for; what is essential for determining the correct modeling approach to a new problem or for interpreting modeling papers? While there is no simple way to characterize an effective model, a number of criteria are suggested: does the model correctly and simply represent our knowledge and hypotheses about the system? Does it discover important implications of those ideas that were not known at the outset? Does it express these discoveries in a testable framework, in terms of biologically identifiable and measurable quantities? Does it clearly distinguish between results that are implied by the biological knowledge or ideas being studied and results that depend on arbitrary details of implementation or *ad hoc* assumptions? Does it allow exploration of alternative assumptions and choices of parameters, enabling understanding of a class of relationships rather than some accidental or special example? Are the results robust to the vagaries of biological variability? Do the results give us insights that will guide or alter experimental work, or do they sit idly and speculatively on the sidelines of the scientific enterprise?

A model is necessarily an abstraction. There is more in nervous tissue than any representation can fully embody. One point emphasized here is that different levels of abstraction are appropriate for different problems. Detailed biophysical models can help us understand signal transduction (Torre *et al.*) or integration (Lytton and Wathey) in single cells, or the behavior of a hippocampal network (Traub and Miles). But simple models lacking in biophysical detail can capture and extend our understanding of large-scale phenomena such as the generation of oscillations (Sigvardt and Williams), the development of large-scale patterns of connectivity (Miller) or the computational significance of observed responses in the cortex (Lockery). Several authors discuss the strengths and weaknesses of alternative levels of abstraction for approaching a single problem (Traub and Miles; Sigvardt and Williams; Lockery; Miller). The modeling presented here is bottom-up: we are looking to the biology to inform us. But bottom-up does not necessarily mean biophysically detailed; excess detail, if not informed by measurement, can be as harmful to useful modeling as insufficient detail (discussed in Sigvardt and Williams; Miller).

A fuzzy line distinguishes models from theories. A model is a representation of what is known or believed about a piece of neural tissue, while a theory is a statement of underlying principles. At one level, a model is a representation that is understood by exploration of parameter space through simulation, while a theory is a representation that is understood in general terms. Thus, a model of lamprey central pattern generators leads to the development of a theory of the dynamics of coupled neural oscillators (Sigvardt and Williams); a model of neural development leads to a theory of the outcome of correlation-based developmental rules (Miller). Theory at this level gives us a complete picture of a model, the whole elephant rather than two legs and a tail. At another level, which

The Use of Models in the Neurosciences

Introduction for Seminars in the Neurosciences, Vol. 4, No. 1, Feb. 1992 (pp. 1-3)

Kenneth D. Miller
Division of Biology
Caltech 216-76
Pasadena, CA 91125
ken@cns.caltech.edu

Few things are easier to speculate about than the workings of the brain and the mind, but useful theories are far harder to develop. Yet neuroscience needs theory. Without some implicit and tentative conceptual framework, experimental results would be only disconnected phenomena. The need for explicit and formal models arises as the complexity of our knowledge and theoretical ideas about the nervous system grows: it becomes difficult to simply intuit the consequences of our knowledge and assumptions.

A model is a formal representation of some aspect of the complex tissue that is the nervous system. It embodies knowledge and assumptions at levels ranging from the molecular to the perceptual or behavioral. The medium may be mathematical equations, computer programs, silicon chips, or some combination. A model may, for example, explore the range of subcellular or channel properties that can account for a single cell's physiological responses (Torre *et al.*; Lytton and Wathey; Traub and Miles; Mahowald *et al.*); the patterns of connectivity that can account for the dynamical behavior of a piece of neural tissue (Traub and Miles; Sigvardt and Williams; Lockery; Mahowald *et al.*) or that can implement an observed or hypothesized input-output relationship (Lockery); the range of developmental outcomes that can result from postulated developmental rules and measurable patterns of connectivity and neural activity (Miller); or the ways in which measured neural responses could account for observed perceptual responses (Westheimer).

Although neuroscience needs models and theories, there is no guarantee that useful ones can be found. A model may be too speculative to have any obvious connection to something that can be measured. It may be too particular to have much relevance, with the outcome dependent on implementation details or on obscure and perhaps unbiological assumptions. It may seem plausible, but neither unique nor testable. Clothed in the mysteries of mathematics and unreadable or unavailable programs, revealing only the beautiful face of computer-generated color graphics, even the silliest model may appear alluring. How can a neuroscientist distinguish models possessing the depth and substance to make a serious contribution?

In this issue we present examples of the effective use of models in the neurosciences at a variety of levels. The articles illustrate the achievements of modeling, but also serve as