


A lighthouse for neural modeling



Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems

by Peter Dayan and Larry Abbott
MIT Press, Cambridge, Massachusetts, 2001. \$50.00
hardcover, pp 460
ISBN 0-262-04199-5

Reviewed by Christof Koch

A famous theorem in astrophysics states, “A black hole has no hair.” It implies that, under steady-state conditions, a macroscopic black hole, such as the monster engulfing the center of our galaxy, can be exactly described by three variables: its mass, angular rotation and charge. This is not an approximation under highly idealized conditions but describes the real universe. Contrast this with a single chemical synapse characterized by a myriad of pre- and postsynaptic properties—the probability of release of a vesicle, its ionic reversal potential, number and time constants of activation and inactivation of the underlying channels and so on—that change in complex and ill-understood ways as a function of the prior usage of the synapse. No wonder, then, that theoreticians have had a vastly more difficult time making sense of the nervous system than physicists have had in rendering the structure and evolution of the universe comprehensible. It is therefore immensely satisfying to open the pages of *Theoretical Neuroscience* to see how much progress has occurred over the past 20 years in our emerging understanding of computation, coding, representation and learning in brains.

As the authors, Dayan and Abbott, two well known computational neuroscientists, point out, there are different reasons for wanting to acquire such a tome. Some neurobiologists seek to link the firing activity of one or more neurons to the behavior of the animal or to quantify the amount of information that a spike train contains. Others are interested in capturing the behavior of their favorite system with the aid of a computer model that tracks intracellular calcium, membrane potential or fir-

ing rate using biophysically or phenomenologically justified differential equations. Finally, because the brain processes information, theoreticians seek to understand the computational steps performed by neural networks, how information is represented—implicitly or explicitly—in the processing hierarchies found in all but the simplest brains, and how these representation can be learned. To all, this book will serve as guidepost, as a lighthouse.

Dayan and Abbott’s textbook reflects these three distinctive interests. The first four chapters cover the all-important question of how the nervous system encodes sensory information and how this information can be decoded by a mathematically astute or ‘ideal’ observer. This section contains the most concise discussion of spike-based codes that I have seen anywhere. The mathematical formalism used is justified by frequent reference to relevant experimental data. The middle part focuses on the quantitative description of single neurons and networks of them. At its core are the nonlinear differential equations that describe the biophysics of synapses, membrane conductances, dendrites and axons, masterfully exploited half a century ago by Hodgkin and Huxley. The dynamics of feedforward and feedback nets and their computational abilities are analyzed with the help of continuous firing-rate models and are applied to associative memory, oscillations, orientation selectivity in visual cortex and processing in the olfactory bulb. The last section places the most demand on the reader’s mathematical abilities. Its focus is on supervised and unsupervised learning algorithms based on variants of Hebb’s rule. The range of applications is catholic, ranging from synaptic plasticity and LTP/LTD to development of ocular dominance and orientation columns, perceptron and delta learning rules, classical conditioning and reinforcement learning, as well as the learning of

receptive fields and a thorough discussion of sparse coding. The book culminates in a final chapter on how the statistics of the input can be used to construct optimal representations for natural scenes or sounds.

The book reflects well the strengths and limitations of theorizing about the nervous system around the turn of the millennium. (More than 80% of citations are to papers from the last two decades.) Although much still remains to be discovered, the field has a solid grasp on the biophysics of single neurons—how synapses work and how dendritic and somatic conductances contribute to the generation and patterning of action potentials—and on the dynamics of small networks or largely feedforward structures, such as the retina. The same claim—of understanding the basic guiding principles and expressing them in this book—can also be made for the encoding and representation of sensory information in spike trains in the periphery and for the development of columns and other structures in cortex and elsewhere based on unsupervised learning rules. But tremendous challenges remain when it comes to thinking about large, adaptive and extremely heterogeneous networks of spiking neurons. In central structures, such as the medial temporal lobe in mammals or the mushroom body of insects, cells signal a significant event using one or a few action potentials that fire at a particular phase. Because an average firing rate doesn’t really make much sense under these conditions, what is the right way to think about such dynamical systems? And how can such sparse representations arise and adapt to a continual changing environment? We still don’t know how to answer these questions, as the field doesn’t even have the right mathematical tools yet to deal with them. Even so, given the many young, competent and highly interdisciplinary scientists attracted to these questions, the outlook for the future is very bright.

Theoretical Neuroscience succeeds admirably as a textbook for a course on computational neuroscience, collective computation or sensory coding. Its scope is universal, with no significant elision. To the extent possible, the mathematics is collected into clearly organized appendices. Web-based exercises illustrate concepts with well chosen examples of relevance to neuroscience. The well designed book, with large margins emphasizing the introduction of key terms, is the ideal companion to a class taught at the senior undergraduate or graduate level to computationally minded neuro-, cognitive or computer scientists. I will certainly use it in this capacity and warmly recommend it to anybody else.

Christof Koch is in the Division of Biology, 139-74, California Institute of Technology, 1200 E. California Blvd., Pasadena, California 91125, USA.
email: koch@klab.caltech.edu