Brain Advance Access published July 24, 2012

BRAIN
A JOURNAL OF NEUROLOGY

BOOK REVIEW

Hodgkin and Huxley's legacy: the science of neural supercomputing

MATHEMATICS FOR NEUROSCIENTISTS
By *Fabrizio Gabbiani* and *Stephen Cox* 2010.
New York: Academic Press
ISBN: 978-0-12-374882-9 (Hardcover)
Price: €71.95/$US99.95

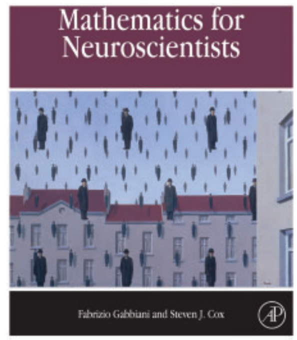

MATHEMATICAL FOUNDATIONS OF NEUROSCIENCE
By *G. Bard Ermentrout* and *David H. Terman* 2010.
New York: Springer
ISBN: 978-0-387-87707-5 (Hardcover)
Price: £53.99/$US74.95

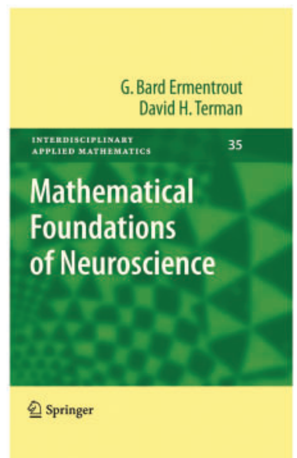

The small sub-tropical island of Okinawa, some 400 miles south of Japan, is home to a unique colony of creatures. About the size of large rabbits, they possess strange translucent heads into which a single eye is embedded. They spend most of their time foraging and mating and communicate with each other using a luminescent bulb located towards the small of their backs. Their predilection for reproduction has allowed their primitive genome to evolve rapidly to adapt to their unusual circumstances. For their food is a power supply to charge their battery packs, and mating involves infrared data transfer from their hard drives. They are cyber-rodents: mobile, autonomous robots confined only to their home in the Okinawa Institute of Science and Technology by the doors of the laboratory.

There are two remarkable things about the cyber-rodents. First, they demonstrate that remarkably little computational power can sustain 'awake', behaving organisms. Their brain is a small four-CPU onboard computer, not a giant supercomputer as you might expect. Second, many algorithms used by the rodents are based on those thought to be implemented by live rats and indeed humans. These algorithms control basic behaviours of self-preservation, mimicking the functions thought to be controlled by neuromodulators such as serotonin and dopamine.

This second fact underpins a basic fact about neuroscience: if we want to understand the brain, we need to understand how it actually processes information. Like it or not, the brain is a computer, and in the same way, as we cannot hope to understand the brain of the cyber-rodent without a fair knowledge of computer science, so too we are not going to understand the human brain without some sort of neural computer science framework.

It is easy to feel a bit daunted by this prospect. With 100 billion neuronal processing units (not to mention those involving glia), each with thousands of synaptic connections, the astronomical complexity that the brain harbours is abundantly clear. If anything, this sentiment is aggravated by the sheer scale of recent big-money projects popular with the media. For instance, the 'Blue Brain' project is attempting to build a full-scale model of the human organ based on assembling biologically accurate simulations of single neurons and exploring behaviour of the assembled network using the formidable Swiss BlueGene supercomputer. In much the same way, The Human Connectome project is attempting to map all the connections in the brain, by painstakingly identifying neuronal connections using fibre-tracing studies. But despite the many insights and tools that these projects will provide, one thing they seem unlikely to yield when completed and 'turned on' is a behaving and talking android.

This, then, still leaves us with the problem of how to understand behaviour in terms of information processing. David Marr, the brilliant computer scientist-turned-theoretical neuroscientist, famously proposed three distinct levels of understanding how the brain works: characterizing the computational problems that the brain solves at one level; specifying the algorithmic solution used at the next and probing its implementation by neurons at the third level.

Consider, as an illustration, the problem of object recognition given hazy visual sensory input. The computational problem can be formalized as data classification given noisy information. A putative algorithmic solution to this problem would be Bayesian decision theory, which would require representing and integrating probability distributions relating to the prior probability of an object identity, and its likelihood given the sensory input and computing the true (posterior) probability using Bayes' rule. This is the 'software' solution. If such an algorithm was used by the brain, then these computations need somehow to be physically implemented by the 'hardware' of the brain. One such possibility might be to code the probabilities with the spike trains of populations of cortical neurons and integrate their activities with a simple neural network (e.g. Beck *et al*., 2008).

This simplified example illustrates the goal of traversing different levels of understanding that potentially take one from the activity of individual neurons to the behaviour of an organism. But understanding the brain at such vastly different scales—the fundamental aim of systems neuroscience—is a formidable challenge. The complexity of neurons means that this is only going to be possible with quantitative models of the brain at each level—a somewhat intimidating prospect for many biologists who lack a graduate training in mathematics, computer science or physics.

Mathematics for neuroscientists by Gabbiani and Cox shows that the sophisticated mathematical methods involved in understanding and modelling the brain are not necessarily out of reach for the 'ordinary' neurobiologist. Based on courses taught by the authors, the book is a superb introduction to quantitative methods that take the reader on a journey from ion channel to network. In doing so, it follows the history of many pivotal neuroscience discoveries along the way, asking the same questions and using the same analysis as was used by the pioneering neuroscientists that discovered them.

Gabbiani and Cox start by developing dynamical models of ion channels to understand passive membrane potentials. The mathematics needed to do this involves specifying the properties of different ion channels as a series of partial differential equations. Having introduced the basic problem here, the authors then add a chapter dealing specifically with the problem of how to solve those equations, using algebraic, numerical and approximate solutions. This sort of chapter characterizes the approach of the book—to use the biological problem to introduce and frame the mathematical problem, and then slowly to take the reader through the solution, incorporating all the intermediate analytical steps and equations that are usually skipped over in more concise texts. Thus, the reader is guided through the properties of active membranes, and in Chapter 4, to the most famous equation in neuroscience—Hodgkin and Huxley's model of the action potential. Their mechanistic model marks the historically defining point at which the basic computing unit of the brain can be said to have been discovered, and with it they founded a new discipline in neuroscience—neural computation.

Quantitative models of neuronal signalling developed rapidly after Hodgkin and Huxley. The basic membrane ion channel equations considered the voltage at a single point on the neuronal membrane, but to understand information transfer, one needed to know how current travels from one point to the next along dendrites and axons. One of the pioneers of modelling this process was Wilfred Rall, who among others applied Cable Theory (originally developed for telegraphic cables) to understand the dendrite as an electrical conducting cylinder. As described in Chapters 6–9, this can be extended to a long length of dendrite with varying diameter by treating it as a series of adjacent, connected compartments. Dendrites are of course more complicated than that. They are typically highly branched and often have far more active membranes than is usually appreciated, but the theory can be extended to deal with much of this complexity.

As matters start to get complicated, Chapter 10 introduces a simplified model neuron—the leaky integrate-and-fire neuron—which has been invaluable to reductionist single cell systems aimed at modelling the properties of networks involving multiple neurons (such as CA3 hippocampal networks), nicely illustrating the continuity of theory as one scales up from ion channel, to single neuron, to network. In a further drive for simplification, Chapter 14 also describes a neat statistical approach to dimensionality reduction of complex neuronal models using singular value decomposition.

One of the most mysterious things about the brain is the constant hum of background neural noise and the randomness apparently inherent in neural signals. Appreciating the importance of this intrinsic biological uncertainty turns out to be at the heart of many mathematical aspects of neural information processing. Gabbiani and Cox provide a limpid exposition of the core concepts of probability theory necessary for the diversity of areas to which it has now proved critical. Perhaps, the clearest early demonstrations of this were in models of synaptic transmission, embodied by Katz's seminal work on quantal release, uncovering another fundamental component on the brain's microprocessing unit.

It is now abundantly clear that grasping the mathematics of stochastic processes is essential for understanding neuronal signals, and these methods turn out to be critical for understanding correlation and coherence across neural spike trains, the mainstay of so much in modern neurophysiological data analysis. Gabbiani and Cox beautifully show how this can be applied to early visual processing, providing clear systems-level insight of how real sensory information can be coded and processed by networks of real neurons founded on Hubel and Wiesel's remarkable discovery of simple and complex cells in the cat primary visual cortex. Their work provided one of the most striking insights into how a neural code can represent key steps in the information processing of a sensory input and critical insight into the mapping of neurophysiological responses to behaviour.

Mathematics for neuroscientists is remarkable in combining both authority and accessibility in describing transition from the dynamics of neuronal signalling to the statistics of information processing. It appeals to, and deserves, a wide audience by functioning both as a textbook on neural computation for biologists and on neurophysiology for mathematicians. Indeed it realizes this aim slightly better than as a primary textbook of mathematics for neuroscientists, since it is likely that many classically trained neuroscientists will have to resort to other more basic mathematical texts to get a firm grounding in areas such as dynamical systems, linear algebra and matrix theory and probability theory. But that does not detract from the clarity by which the authors apply these concepts to neuroscience.

Much of the same motivation underlies *Mathematical foundations of neuroscience* by Ermentrout and Terman. They set out in a similar, albeit more concise, fashion covering the fundamentals of the Hodgkin–Huxley equations and Cable Theory. However, they soon lay the groundwork for a specific approach where they have been at the forefront—applying dynamical systems theory at multiple levels of analysis. They start by introducing phase planes—a geometric representation of the solution to the membrane ion channel equations, which depict the dynamic evolution of the voltage trace following an injection of current. They then describe how fixed points and limit cycles emerge, reflecting periodic solutions that result in oscillatory firing of the model neuron. Within this framework, transitions between different dynamic states can then be studied with powerful analytical tools such as bifurcation theory.

As the authors remark, Hodgkin and Huxley's decision to study the squid axon was inspired, or just plain lucky, given the small number of ion channels it has in comparison with mammalian neurons. Ultimately, however, some of these other channels are necessary to understand more complex patterns of activity; and Ermentrout and Terman elegantly illustrate how their incorporation predicts much of the variety of bursting patterns seen in certain brain regions. Well-known examples are the sleep rhythm generators in the thalamus and respiratory rhythm generators in the pre-Bötzinger complex in the medulla.

But neurons do not burst in isolation, and the authors extend the dynamical systems approach to the spatial propagation of action potentials and spike trains down the axon, and the generation of postsynaptic currents, by considering neurotransmitter sensitive ion channels. This is central to the transition from the single neuron to networks of neurons, approached head-on in the second part of the book.

Neuronal dynamics gets much more interesting when one starts to connect neurons together. One of the first clear examples came from the study of locomotion in the lamprey eel, and the discovery of networks of neurons that could sustain rhythmic behaviour in the absence of any external input—so called central motor pattern generators. Such networks of neurons function as primitive 'mini-brains' in their own right, as they continue to function when severed from the rest of the nervous system. This also makes them a perfect target to test the validity of dynamic models based on coupling single neuron models together. The subsequent success of such simulations in yielding stable oscillatory patterns of activity was therefore a major milestone in neural computation.

However, understanding the behaviour of large-scale populations of neurons requires scaling up from the dynamics of individual neurons, and their synaptic coupling, to embrace the complexities that arise from different network topologies. Clearly, the state space becomes extremely large when one considers the diversity of neural architectures found in real brains. The authors review some of the key parameters likely to be important, such as the sparseness and homogeneity of connections, the proportions of excitatory and inhibitory connections and their speed. With such a large parameter space, however, one of the most challenging problems is knowing what to simplify or approximate without losing too much biological validity. This is an open question, because experimentally it remains difficult to study large numbers of individual neurons simultaneously. Instead, many experimental and clinical methods are stuck with the aggregate behaviour of large numbers of neurons, such as local field potentials, magneto- and electroencephalography and functional MRI.

But does this matter, or can we adequately understand large neuronal populations by considering their average firing rates? Ermentrout and Terman convey a diplomatic agnosticism on this point. But they highlight one of the most influential network models, the Wilson–Cowan model, which uses simplified excitatory and inhibitory neurons to study how neuronal populations interact. This yields a basic dynamic systems analysis that can be applied to a range of competitive and stimulus-evoked behaviours, such as binocular rivalry and perceptual decision making. Indeed many systems level models of more complex cognitive behavioural phenomena, such as working memory, are direct descendants of this dynamical approach (for instance neural mass models and neural field theory; Deco *et al*., 2008).

Although *Mathematical foundations of neuroscience* provides a thorough survey of the field, it is perhaps less of a textbook than *Mathematics for neuroscientists* but rather a more comprehensive analysis of the application of dynamical systems theory to neuroscience. Notably, the approach in both books is founded on understanding information processing in terms of the neurophysiological properties of neurons—reflecting the authors' considerable contributions in this area. Neither, however, considers top–down approaches to information processing. That is they ask the question 'What can neurons compute given what we know about their physiological properties?' rather than 'What must neurons compute to produce behaviour?' For this reason, large areas of theory that stem from reverse-engineering behaviour are not covered. This includes control theoretic models of motor learning, reinforcement learning models of decision making and statistical and information theoretic models of representational learning. Thus neglected are a number of fascinating areas such as fear conditioning in the amygdala (McNally *et al.*, 2011) and spatial encoding in the hippocampus (Burgess *et al*., 2007), which have been so successful in highlighting tangible links between the multiple levels along the scale from molecular mechanisms to behaviour.

From a clinicians' perspective, the world of single neuron computation and non-linear neuronal dynamics might at first seem somewhat distant from the practical reality of neurological disease. But this is an illusion, and indeed some of the most common diseases in neurology can easily be classified as disorders of neural dynamics. An obvious example is epilepsy, and indeed there is a long history of non-linear systems analysis in the study of seizure initiation and propagation. A popular example is in the use of bifurcation analysis of neural fields to predict phase transition to chaotic dynamics (Grimbert and Faugeras, 2006). Perhaps less well known is that symptoms in Parkinson's disease are closely correlated with abnormal synchronous oscillatory activity recordable in multiple levels of the basal ganglia loop, and efficacy of treatment with deep brain stimulation of the nucleus seems tied to modulation of this abnormal firing (Hammond *et al*., 2007). Finally, and most aesthetically illustrative, is migraine: it has long been known that disturbed cortical synchronization coincides with the onset of migrainous attacks. Work by Ermentrout and Cowan (1979) has suggested that abnormal synchronous activity might account for the beautiful and complex geometric patterns that characterize the visual aura experienced by many patients.

All these examples should precipitate new avenues for treatment, in particular in one of the most exciting frontiers in neurology—brain-machine interface technologies. Future neural prostheses are likely to involve simultaneous decoding of complex neurophysiological signals coupled with selective, temporally sophisticated modulation of neuronal activity using techniques such as optogenetics. At some point in the future, these therapies will be in the hands of clinicians to manage, and 'computational neurology' will have come of age as a new clinical discipline.

There are many fascinating questions that lie ahead. Is the background neural activity merely noise, or does it reflect computation of some mysterious unknown function; does the dynamic synchronization of activity across spatially distant areas underlie the binding of the sensory world into a coherent and meaningful conscious perceptual experience; to what extent does the specific timing of individual spikes matter and are we missing some crucial neural code when considering only the average rates in neural populations? This last question—the true nature of the neural code—is perhaps the greatest unanswered question in biology. And it is the field of neural computation, and no other, which will provide the answer; and with it fulfil the legacy provided by the pioneering work of its founding fathers, Alan Hodgkin and Andrew Huxley.

It is easy to get lazy about neural computation and become sidetracked by philosophical conversations about whether the brain can ever understand itself, and can we understand consciousness in terms of computation. But this should not detract from the simple truth that the brain computes to behave, and we will not understand behaviour, and its pathologies, until we understand the underlying computations.

Ben Seymour[1,2]
[1]*Centre for Information and Neural Networks, National Institute of Communications Technology, Osaka University, Osaka, Japan*
[2]*Department of Engineering and Department of Clinical Neurology, Cambridge University, Cambridge, UK*
E-mail: bjs49@cam.ac.uk

References

Beck JM, Ma WJ, Kiani R, Hanks T, Churchland AK, Roitman J, et al. Probabilistic population codes for Bayesian decision making. Neuron 2008; 60: 1142–52.

Burgess N, Barry C, O'Keefe J. An oscillatory interference model of grid cell firing. Hippocampus 2007; 17: 801–12.

Deco G, Jirsa VK, Robinson PA, Breakspear M, Friston K. The dynamic brain: from spiking neurons to neural masses and cortical fields. PLoS Comput Biol 2008; 4: 1–35.

Ermentrout GB, Cowan JD. A mathematical theory of visual hallucination patterns. Biol Cybern 1979; 34: 137–50.

Hammond C, Bergman H, Brown P. Pathological synchronization in Parkinson's disease: networks, models and treatments. Trends Neurosci 2007; 30: 357–64.

Grimbert F, Faugeras O. Bifurcation analysis of Jansen's neural mass model. Neural Comput 2006; 18: 3052–68.

McNally GP, Johansen JP, Blair HT. Placing prediction into the fear circuit. Trends Neurosci 2011; 34: 283–92.