# CRCNS Discussion Forum
## Biological Relevance in Computational Neuroscience—Four Perspectives

**What is biological relevance in computational neuroscience, and why is it hard?**

Nancy Kopell
Department of Mathematics and Center for BioDynamics
Boston University

Mathematics focuses on finding recurrent and reproducible patterns, and explaining how they come about. Biology deals with behavior that is ever changing and context dependent, in which there is robustness as well as very rich diversity; one of its central problems is to explain how the robustness comes about in the face of all the potential alternatives provided by the complexity of the nervous system. One structural reason why computational neuroscience is hard is that the central question of robustness vs diversity requires looking simultaneously at the largest possible picture, to get a sense of why diversity is necessary, and at the much more detailed levels, to understand how any part of that diversity is generated.

We are likely to be talking in this symposium about bottom-up and top down and even center-out, but all of us know that the real problem is going from one end to the other. I think we also agree that there is no one privileged place to start. So how can we judge if the work we or others do is providing useful pieces for the overall jig-saw puzzle? I do not think this question is answerable in generality, but I can offer some opinions that guide my own work.

- Look for some level of generality that also makes contact with some detailed observations. My favorite example in my own work is about the lamprey central pattern generator for swimming. In that work, the issue was the mechanism for the traveling wave of neural activity associated with the physical movement. The bottom line is that the activity was shown to be a consequence of the local dynamics and global anatomy: the spinal cord is essentially a chain of oscillators, which has some generic properties independent of the physiological details. This gave rise to testable (and successfully tested) implications.

- Be very critical about the hypotheses underlying the simulations and analysis of any given project. There is nothing wrong with leaving out most of what is in the full biological system motivating a given project - indeed, it is inescapable. But there may be a problem if the conclusions depend critically on specific hypotheses or parameters used, especially if these hypotheses are counter-factual. That is, try to evaluate a caricature by looking at more complex versions in which various hypotheses are relaxed. As an example, integrate-and -fire neurons have been much used to investigate network behavior. It is an open question when the results of such investigations generalize to more complex HH equations. For situations in which complex voltage-dependent currents are active between spikes, the network behavior can be very different. Sometimes it takes multiple investigators to do this: I like working with Roger Traub

since we can go back and forth between very detailed models that are hard to intuit and simpler descriptions that are vulnerable to being too simple in the absence this kind of control.

- Look for questions that deal with bridging levels of behavior, and that can be handled without knowing all the details of lower levels. This is what I mean by "center-out." My example of this is the investigation of rhythms in the nervous system. Using knowledge of the ionic currents of cells and their pharmacology, we can investigate the origin of different neural rhythms and their coherence behavior, without knowing, e.g., biophysical details of the channels or their genetic manipulations. However, ability to make those manipulations can (and have) then be used to test predictions from the network analysis. The network behavior can then also be used to investigate behavior at higher levels of organization: the timing and organization of rhythmic firing has implications for how sensory information can travel among substructures of the nervous system, and is relevant to, e.g., attention, learning and recall, motor planning.

I find the "center-out" strategy especially important for issues of structure and function. It is reasonably straightforward to guess at the basic function of an eye or a leg. But for structures such as the entorhinal cortex, many synapses away from either sensory input or motor output, there is a deeper mystery. It is well accepted that there is a relationship between structure and function; this "structure" is usually interpreted in terms of anatomy, physiology, pharmacology etc. However, there is also dynamical structure, associated with patterns of firing, and changes of firing in different behavioral and modulational states. By looking at network behavior in slices in different pharmacological situations  (not a low-level as genes or as high as behavior), one gets
baseline dynamical structure critical to understanding the different ways in which the network can process inputs, which in turn can be very suggestive about function. Indeed, I believe that this dynamical structure is a clue to "design principles," another theme that will come up in this discussion.

# Models are relevant if taken seriously by neuroscientists

Tomaso Poggio
McGovern Institute
Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology

Two points:

1) In some sciences, like physics, models are often taken seriously. In neuroscience, models will become relevant if and when some of them will be taken seriously.

2) In some sciences, say physics, there are top-down models/theories like thermodynamics. There are also bottom-up models, like statistical mechanics. They both play a fundamental and complementary role in the development of science and of our understanding.

Early in my scientific life, I was taught by Werner Reichardt that models in neuroscience should be developed in very close contact with experimental data. Their success is directly measured by how much they affected the design of new experiments. At about the same time, I worked on biophysics of neurons and synapses in the vertebrate retina and in the visual system of the fly, following a bottom-up approach to the brain, which is of course important and has an established role in computational approaches in neuroscience, at least since the work of Hodgkin and Huxley. Of course, I think that it cannot be the only approach, unless one believes that we should not worry about biology because it follows from quantum mechanics. I believe that models of a different type are also needed, models that are more top-down in the sense that they start from a computational problem/need like visual recognition and suggest possible ways of how the brain might solve it, taking into account not only computational constraints but also data we have from behaviour, physiology and anatomy. Models of this type cannot be as yet too detailed – simply because we do not know many of the details – but must be quantitative.

An example of an approach centred on models of this intermediate type is our NIH-funded Conte Centre, which is also funded by NSF-CRCNS. Its central focus is to study the neural computations underlying object recognition in visual cortex. The Centre's framework is based on the collaboration of labs working on monkey physiology, cat physiology and human psychophysics with a quantitative, "intermediate" computational model providing the main conduit through which experimental results in one lab affect experiments in another lab. The model represents a novel tool for driving a collaborative, highly multidisciplinary enterprise, providing a way to integrate the data, to check their consistency, to suggest new experiments and to interpret the results. One of the Centre's key motivations is in fact the belief that quantitative models of complex neural system, when developed in close cooperation with experimental labs, can be powerful tools in basic research that integrates across several levels of analysis – from synaptic, to cellular, to systems, to complex visual behaviour. In particular, I believe that quantitative models can be tools to a) think about the problems (some cognitive problems are too complex for the qualitative, simple models used so far); b) make predictions, suggest and plan new experiments; c) analyse and interpret data; d) integrate experimental findings of different types and from different labs, drawing implications for future work from multiple sources of evidence.

I believe that there are several problems in higher brain functions that are now at a stage in which quantitative models of this type are required just to make progress. For instance, the problems of visual recognition and motor control are computationally difficult – we do not know as yet how to make computers solve them at the level of human performance! -- and the size of the experimental data from different sources is growing rapidly. Quantitative models will increasingly replace the traditional qualitative mental models of the visual and motor physiologist and will become a more and more important tool for developing consistent explanations, for interpreting existing data and for planning and analysing new experiments.

**How relevant is biological relevance?**

Rob de Ruyter
Department of Physics
Indiana University

A fitting way to start this discussion is by quoting the opening line in Horace Barlow's contribution to Rosenblith's 1961 book on Sensory Communication: "A wing would be a most mystifying structure if one did not know that birds flew." Barlow makes the argument that we can analyze all kinds of features of a bird's wing in great detail. But, he notes, the work would very soon become unrewarding if we don't have a clue what the wing does for the bird or if we are ignorant of the physical principles underlying that function.

In the years since, there has been great progress along many fronts in neuroscience. Most of that progress, I think, has been in isolating more parts, describing them in ever greater detail, and unraveling how they work. But I wonder how much closer we have come to understanding *what* those parts really do, how they work together, and how they combine to help the animal function in its environment. In aiming to understand that, the field has been chasing a moving target. I think it is fair to say that today we realize that animal behavior is far more complicated than was thought in the early sixties. The classical ethological picture of stimuli acting as "releasers" of more or less stereotyped behaviors has made way for an understanding of animal behavior as a much more adaptable and plastic phenomenon, even in so called "simple" animals. Further circumstantial evidence that animals have more up their sleeves than we used to think comes from the failure, by and large, of artificial intelligence to produce autonomous robots that can go out and perform tasks in the real world. Ants and bees, to name an example, possess navigational capabilities far superior to those of artificial autonomous systems.

This evolution in our views on animal behavior has a parallel in present day thinking about neural processing: To name two examples, the original picture of cells in primary visual cortex as bar detectors has made way for a more nuanced understanding in which stimulation outside the classical receptive field plays an important role. And, as an example from "simple" animals that I am familiar with, motion sensitive neurons in the fly visual system used to be described as rather fixed in their properties, but have recently been found to adapt strongly to statistical parameters of the motion stimulus.

Viewed together, these findings suggest that we have but a very limited understanding of the tasks that animals perform when they are dealing with the real world. And of course this in itself is nothing new. The success of the scientific method in general hinges on the realization that the world at large is complicated, and that we must carefully choose simplified cases that allow us to isolate and study the interesting phenomena. But when we study the neural substrate of behavior within a rigorous regime of simplification, we may throw the baby out with the bathwater. I would argue that we need to study cases, simplified to be sure, but such that we nevertheless retain some connection with the ways in which animals process more or less natural stimuli. This requires studying the relevant properties of natural stimuli, as well as the behavior that our animal of choice displays in natural conditions. Difficult as it is already, this alone is not enough. We will also need organizing principles to guide our thinking and put our observations in

perspective. One such principle is optimality: We may, perhaps, understand the interactions of animals with their environment as the best possible solution given certain physical and biological constraints, and given the irreducible uncertainties presented by the environment. This clearly is in line with evolutionary views, in which one expects a process of "survival of the fittest" to drive animal design toward some form of optimal behavior.

In the following segment, Dan Margoliash will have more to say about evolution and neuroscience. I would like to conclude by noting that the study of animals in biologically relevant conditions could very well be highly relevant for deepening our understanding of neural information processing. That enterprise is not going to be easy. The alternative, however, is far worse, as we would have to feel the scorn of Horace Barlow, who concludes his paper by saying: "it is foolish to investigate sensory mechanisms blindly—one must also look at the ways in which animals make use of their senses."

**Are there "design principles" in computational neurobiology? How do we identify them?**

Daniel Margoliash
Department of Organismal Biology and Anatomy
University of Chicago

I would like to comment starting from the perspective described by Rob de Ruyter. One aspect missing from the discussion in the field is the relation to our organizing, central principle, evolution. It may be helpful to characterize the differences between systems neuroscience and computational neuroscience. At its highest level systems neuroscience attempts to figure out "how it works," typically from the perspective of a modality. At its highest level computational neuroscience attempts to figure out the design principles, typically from the perspective of the output (behavior). Our methods for assessing design principles are weak. We do not have a tool box of design principles, at least we have yet to attempt to create a tool box of computational design principles based on what we know about the evolution of the nervous system and the body plan. Thus, our approach is to work on a particular model (particular animal, particular organismal or circuit behavior). This is hard work for all the reasons mentioned in this discussion (intrinsic difficulty of working with neurons and behavior, multiple levels of analysis, collaboration across disciplines). Then we should do this for multiple species chosen on a phylogenetic basis. But the first steps are so difficult that to contemplate launching a computational efforts into a new species or set of species is daunting. Also comparative anatomy is dying (or regressed?) which is an acute loss for these efforts. All this makes addressing the big problems very hard indeed, which I think helps answer Nancy Kopell's question.

A case in point is the idea of the canonical cortical circuit. People typically study monkeys or cats or other derived animals (in the evolutionary sense) and try to make inference into general circuit properties. The cortex is conserved at some level but is that weak claim sufficient to guide an analysis of its design principles. A good discussion point. Another example, in my field of birdsong there was a fascinating period where song learning was examined across multiple species. We learned quite a bit about species diversity - how different organisms solved different but related problems. This continues with behavioral ecology of song birds but in neuroethology these days there is a strong tendency to gravitate towards studying one species, the zebra finch, and a tendency to pay less attention to behavior. Our white rat. I am sure we are learning a great deal about zebra finches, and quite a bit about vocal learning. But our approach would be strengthened if we re-embraced species diversity.

Other thoughts: There are at least two fairly distinct meanings of computational neuroscience. The first is very practical, and brings quantitative ideas and approaches to bear on difficult problems in experimental design and data analysis. The second is not so practical but no less valid, which might be call theoretical neuroscience. We need to support both.

I personally do not think the distinction between bottom up and top down is particularly helpful. Both work, often times wonderfully and in concert in the same system, and stated practitioners of each benefit from the other.