

Understanding Neural Complexity: A Role For Reduction

JOHN BICKLE

Department of Philosophy and Neuroscience Graduate Program, University of Cincinnati, Cincinnati, OH 45221, USA; E-mail: bicklejw@email.uc.edu

Abstract. Psychoneural reduction is under attack again, only this time from a former ally: cognitive neuroscience. It has become popular to think of the brain as a *complex system* whose theoretically important properties *emerge* from dynamic, non-linear interactions between its component parts. “Emergence” is supposed to replace reduction: the latter is thought to be incapable of explaining the brain *qua* complex system. Rather than engage this issue at the level of theories of reduction versus theories of emergence, I here emphasize a role that reductionism plays – and will continue to play – even if neuroscience adopts this “complex systems” view. In detailed investigations into the function of complex neural circuits, certain questions can only be addressed by moving down levels and scales. This role for reduction also finds a place for approaches that dominate mainstream neuroscience, like the widespread use of experimental techniques and theories from molecular biology and biochemistry. These are difficult to reconcile with the anti-reductionist sentiments of the “complex systems” branch of cognitive neuroscience.

1. Dueling Images

Contemporary neuroscience is a battleground. Two “global” views are at war. Ecumenical reconciliation or shared governance ultimately is unacceptable. Although each side rarely comes right out and says it, both imply that their view of the brain is *right*. And their views are jointly inconsistent: both *can't* be right, as a simple matter of logic.

This war in current neuroscience is just a re-enactment of the old holism-reductionism debate. Is the whole equal to, or more than, the sum of its parts? We've known for some time that “sum” is the wrong term here. Sum is a linear function and the interactions between components in even the simplest physical systems involve non-linear dynamics. The real question is thus: Is the whole equal to, or greater than, the dynamic, non-linear interactions between its structured components?

Within brain science, one branch says “greater than.” Increasingly this group adopts the technical notion of a “complex system” to clarify its holistic image. As a “working definition” for this paper, we can characterize this technical notion as

- A set of *interacting* components whose *collective behavior* cannot be deduced from the behavior of those components *in isolation*.

A non-reductive sense of “emergence,” difficult to define, is then offered as the relation between the brain's physical/chemical/molecular/physiological/anatomical and its cognitive features. On the other hand, the answer from mainstream neur-



oscience – from the Society for Neuroscience crowd – says “equal to.” A mature account of the brain’s physical-chemical-molecular structure *and dynamic interactions* will explain its physiology and anatomy, and in combination its cognitive features.

Holist-reductionist debates invigorate everybody, but are only of scientific value when they generate novel methodologies and fruitful experimental results. The debate in current neuroscience has succeeded in doing this. Cognitive neuroscience now employs a variety of methodologies for investigating whole brain function, and no one should doubt that these are returning important data. *Neuroimaging techniques* – positron emission tomography (PET), functional magnetic resonance imaging (fMRI), magnetoencephalography (MEG), and others – are now as prevalent in the experimental lab as in clinical practice. These techniques are generating images with increasingly fine spatial and temporal resolution of regional neural activity during well-controlled cognitive and behavioral tasks. Neuropsychologists continue to refine their traditional methods and supplement them with new techniques (Kolb and Whishaw, 1996). Computational neuroscientists construct massively parallel network architectures, borrowing ideas from neuron physiology and neuroanatomy. Often these networks can be simulated on digital computers, both to test hypotheses about network performance capacities and to implement machine intelligence. A powerful interpretive tool unites these “whole brain” methodologies. This is the theory of dynamical systems from mathematical physics, where states of the system are represented as points in high-dimensional vector spaces and state transitions are trajectories through these spaces. How these vector spaces are implemented physically/chemically/molecularly is often assumed to be of little importance at these cognitive-neuroscientific levels of theory.

At the same time, mainstream neuroscience (the Society for Neuroscience crowd) has grown increasingly reductionistic in both its central questions and its methodologies. The cellular and intracellular methods of traditional neurophysiology continue to be applied, developed, and refined. Increasingly one finds the methods of molecular biology being put to use, including receptor subtype-specific pharmacological and genomic manipulations. Even neuroanatomical track-tracing now uses specific molecular and immunological markers. Five minutes perusal of this year’s *Society for Neuroscience Abstracts* volume (cataloguing the 13,000+ slide and poster presentations at the year’s annual meeting) reflects how intracellular, molecular, and biochemical mainstream neuroscience has become. This attitude is also reflected in the discipline’s principal textbooks. Nearly one decade ago, Kandel et al. (1991) opened the 3rd Edition of their comprehensive *Principles of Neural Science* by stating: “The goal of neural science is to understand the mind, how we perceive, move, think, and remember. In the previous editions of this book we stressed that important aspects of behavior could be examined at the level of individual nerve cells . . . Now it is possible to address these questions directly on the molecular level” (1991, p. xii). By the 4th Edition, after another decade of

cell-biological and molecular neuroscience, their introductory statement was even bolder:

This book . . . describes how neural science is attempting to link molecules to mind – how proteins responsible for the activities of individual nerve cells are related to the complexity [!] of mental processes. Today it is possible to link the molecular dynamics of individual nerve cells to representations of perceptual and motor acts in the brain and to relate these internal mechanisms to observable behavior. (2000, pp. 3–4)

Here again, it would be folly to deny that fruitful data has resulted. Take a look at the 1400+ pages of Kandel et al. (2000) if you're doubtful (or any other comprehensive recent text).

So which view is right? Neural holism/emergence or reductionism? Brain function greater than or equal to the dynamic interactions between its highly structured parts? The typical philosopher's approach is to construct or elaborate "theories" about the key concepts involved (e.g., emergence and reduction), then "test" these "theories" against actual and imagined scientific practice and results. I propose a different strategy. I will examine a recent attempt my group made to hypothesize the function of a neural circuit that meets the working definition stated above for a complex system. Treating this circuit at the network-architectural level yielded a functional hypothesis for it. But one type of question that had to be answered to defend *the biological plausibility* of these modeling results and functional hypothesis required us to reorient our modeling efforts *down a level*: to adopt a reductive approach. This requirement is not specific to our particular project: it generalizes across the board in cognitive neuroscience. Treating the brain as a complex system is no escape from reductionist methodology. Nor should it be, since a central aspect of the reductionist intuition is correct. Or so I shall argue below.

2. A Functional Hypothesis for a Thalamocortical Circuit Suggested by Computer Simulation

Anatomical and physiological discoveries over the past three decades have revealed the complex circuitry between the relay nucleus in the thalamus and the primary cortical region for each sensory modality. The mammalian visual projection system, from retina to thalamic *lateral geniculate nucleus* (LGN) to *primary visual cortex* (V1), has been studied extensively. The following features are now well documented (Jones, 1985; Adams et al., 1997; Sherman and Koch, 1998):

- V1 ("feedback") projections to LGN outnumber LGN ("feedforward") projections to V1 by a ratio of at least 10:1.
- V1 feedback projections obey a "principle of reciprocity" so precise that LGN-V1 feedforward connections can be predicted from V1-LGN feedback connections. These feedback projections retain the retinotopic map (location of retinal stimulation) of the feedforward projections.




Figure 1. Schematic of the mammalian visual projection system, including afferents and efferents of the thalamic reticular nucleus (TRN).

- Both LGN feedforward and V1 feedback projections use glutamate as their predominant neurotransmitter, and hence are excitatory. (Activity in the pre-synaptic neuron tends to produce rather than inhibit membrane depolarization in the post-synaptic neuron.)
- Both LGN feedforward and V1 feedback projections drop off axon collaterals into specific regions of the thalamic *reticular nucleus* (TRN).
- TRN axons project intrathalamically, with a high degree of reciprocity between a specific TRN region and the thalamic nucleus from which it receives collaterals. Specific regions are dedicated to specific sensory modalities. For example, a map of the visual field exists in the region of the TRN receiving collaterals from LGN and V1.
- TRN projections to thalamic sensory relay neurons use GABA (γ -amino butyric acid) as their sole neurotransmitter, and hence are inhibitory. Evidence from a variety of *in vivo* and *in vitro* experiments suggest that these intrathalamic projections implement a kind of *lateral inhibition* within thalamic sensory nuclei (reviewed in Bickle et al., 1999).

Illustrate these details in the circuit diagram of Figure 1.

It is now clear from these and other findings that thalamic sensory nuclei are more than just passive relay stations from receptors to cortex. Increasingly, theorists hypothesize the role of thalamus and its connections with both brainstem and cortical structures in the neural mechanisms subserving such clinical and psychiatric phenomena as attention and arousal (Steriade et al., 1993; Sherman and Guillery, 1996). But the *specific* function(s) of this LGN-V1-TRN circuit remains puzzling (as do similar cell properties and connectivities in mammalian auditory and somatosensory projection systems). More than a decade ago Francis Crick and Christof Koch pointed out that “there are many back projections (and some cross-connections) between cortical areas and also from cortex back to the thalamus, *the function of all of which are unknown*” (1990, p. 268; emphasis added). A recent review paper reveals that the most straightforward functional questions remain unanswered: “we still cannot even determine if activation of these [corticothalamic] circuits excites or inhibits [thalamic] relay cells” (Sherman and Koch, 1998, p. 307). Note that the “working definition” of complex system presented in the previous section applies to this circuit. It is a set of neurons (“interacting components”) whose circuit function (“collective behavior”) cannot be deduced from the well-known anatomical, physiological, and even cellular and molecular properties of the neurons themselves (“behavior of those components in isolation”).

To explore these circuit-functional questions, we employed techniques from computational neuroscience and computer simulation (Bickle et al., 1999). First, we noticed that the abstract computational architecture of this LGN-V1-TRN circuit resembles that of a well-studied artificial neural network, the *Interactive Activation and Competition* (IAC) net. IAC networks contain segregated “pools” of units, with excitatory bidirectional connections between some units across pools and inhibitory connections between units within pools (Grossberg, 1978). Existing software for implementing IAC networks on digital computers permits modelers to specify connections across and within pools and choose values for a range of processing parameters (McClelland and Rumelhart, 1987). We specified these connections and chose these values to mimic known cell properties and circuit features of the LGN-V1-TRN system (Table I). We then ran a variety of experiments with this simulation. For example, we mimicked a strongly salient visual stimulus (i.e., in terms of luminance contrast) to one portion of the (simulated) visual field, simultaneously with a moderately salient stimulus to another portion. We tracked the activity dynamics in both simulated LGN relay and V1 neurons over time. Initially, both simulated LGN relay and cortical units displayed activity that reflected the relative strength of the stimulus in the units’ simulated visual fields. (This result is in keeping with the “rate law” of sensory physiology; see Carlson, 1994). But after simulated cortical excitatory feedback and intrathalamic inhibitory input began to occur, both simulated LGN relay and cortical cells whose receptive fields contained the stronger stimulus *maintained* their activity rates, while those whose receptive fields contained the moderate stimuli began declining toward “rest” values (Figures 2 and 3). By calculating activity dynamics in the actual LGN-V1-TRN circuit, we




Figure 2. Activity dynamics in simulated cortical (CVC) units over a number of processing cycles to a single visual stimulus. The simulated receptive fields of CVC1-4 received the strong stimulus; those of CVC5-8 received the moderate stimulus. Initially the units reflect the relative strengths of the stimuli (cycle numbers C2, C3). But as excitatory feedback projections from V1 and inhibitory projections from simulated TRN to LGN relay units begin to take effect in later processing cycles (C4-C7), units whose receptive fields contain the strong stimulus *maintain* their activity level while those whose receptive fields contain the moderate stimulus *decline* toward resting values. For full analysis and biological interpretation, see Bickle et al. (1999). (Adapted from Figure 3 in Bickle et al., 1999, p. 257.)

estimated that these global dynamics should begin appearing roughly 50 msec following stimulus onset: well within the saccade latency period that sets an absolute upper bound on neural processing of a specific retinal input to LGN.

Based on these modeling results, we concluded that the actual LGN-V1-TRN circuit functions to “prime” selective attention mechanisms further up in the cortical visual streams toward specific portions of a total visual stimulus (Bickle et al., 1999). Hence significant stimulus-driven, “bottom-up” selectivity and priming occurs at the earliest stage of cortical visual processing. This circuit appears to be one more among many thalamocortical mechanisms that subserve the neural basis of attention (McCormick, 1989). Our model also goes far beyond existing ones about the processing and function of thalamic relay-sensory cortical circuits. Given the known cell-physiological and anatomical similarities between the visual projection system and those of other sensory modalities, we expect a similar “pre-attentive priming” function to be performed by thalamocortical-TRN circuitry in both auditory and somatosensory projection systems.

Table I. Choice of values for parameters and variables in IAC network equations and their biological justification for our LGN-V1-TRN simulation

Parameter	Value	
<i>max</i> maximum activation value of unit	1.0	Set strictly for computational reasons based on the equations governing $\Delta \alpha_i$
<i>min</i> minimum activation value of unit	-0.1	Set strictly for computational reasons based on the equations governing $\Delta \alpha_i$
<i>rest</i> activation value of unit given if <i>exin</i> =0	0.0	Setting <i>rest</i> slightly higher than <i>min</i> reflects the biological fact that neurons discharge spontaneously
α scales excitatory input from other units within the network	0.25	
<i>estr</i> scales excitatory input (<i>exin</i>) from outside the network	0.4	We set <i>estr</i> slightly higher than α for the biological reason that the principle function of thalamic relay neurons is to relay information about <i>external stimuli</i> to primary sensory cortices.
<i>decay</i>	0.25	We set <i>decay</i> equal to α (0.25) based on trial and error with numerous IAC simulations.
For biological reasons γ (receives special treatment)		
γ scales inhibitory input to units from within the network	Initially 0.0 After 2nd processing cycle 0.125 After 3rd processing cycle 0.25	No input to TRN from simulated stimulus. This keeps initially active TPCs from inhibiting others immediately after the former are activated by <i>exin</i> . After the second processing cycle, when TPCs have processed a second burst of <i>exin</i> and CVCs have processed their first burst of excitatory input from initially activated TPCs, we set $\gamma=0.125$. This is one-half its fully activated value, mimicking the fact that at this point real TRN neurons have received stimulus driven input only from thalamocortical collaterals. After the third processing cycle to a simulated visual stimulus we set $\gamma=0.25$, its full value. This mimics the fact that at this point real TRN neurons have received their full stimulus driven input from both thalamocortical and corticothalamic collaterals.




Figure 3. Activity dynamics in two simulated LGN relay neurons with classic ON-center, OFF-surround receptive fields over a number of processing cycles to a single visual stimulus. The simulated receptive field of TPC53 contained the strong stimulus; that of TPC30 contained the moderate stimulus. A similar account of the activity dynamics in simulated cortical units holds here as well. (See Figure 2 above.) For full analysis and biological interpretation, see Bickle et al. (1999). (Adapted from Figure 4 in Bickle et al., 1999, p. 258.)

We were able to provide neurobiological justification for both connectivities between units within our computer-simulated model and our choices of most processing parameter values (Table I). Using existing cell-physiological data, we were also able to defend the biological plausibility of key temporal features of our results and our use of lateral inhibition within the LGN relay pool to mimic the effects of TRN inputs. However, there remained one important computational assumption that we could not discharge biologically. IAC networks require that the scaling values for excitatory and inhibitory inputs to a unit from other units within the network be (close to) equal; otherwise an IAC net quickly displays unstable performance. The question is whether this assumption, interpreted biologically, is even approximately true in the neural system we were modeling. Does a given excitatory input from V1 feedback projections exert roughly identical effects towards depolarizing an LGN relay neuron, compared to a given inhibitory TRN input toward hyperpolarizing it? There is no cell-physiological evidence, *in vivo* or *in vitro*, that speaks directly to this question. Nor does there seem to be any indirect evidence that permits a confident answer. Yet modeling the LGN-V1-TRN circuit on a computer-simulated IAC network requires this biologically unjustified computational assumption. Clearly, this presents a problem for any *neurofunctional* hypothesis derived from our results with this model.

3. Available Electron Microscope Data

Electron microscope studies of LGN relay neurons have revealed some structural details relevant to this worry. In the cat, Montero (1991) has shown that RSD terminals (*Round vesicles, Small terminals, Dark mitochondria*) from V1 afferents comprise 57–59% of terminals sampled on LGN relay neurons. GABAergic F1 terminals (pleomorphic vesicles, dark cytoplasmic matrix, dark mitochondria, symmetrical synaptic contacts) from TRN afferents comprised 23–25% of terminals sampled. This suggests a ratio of roughly 12:5 between excitatory cortical and inhibitory TRN synapses on a typical LGN relay dendritic tree. In the primate *Galago crassicaudatis* (bush baby) Feig and Harting (1994) and Harting et al. (1991) have shown the locations of these (and other) synapses on LGN dendritic trees. RSD terminals from cortical (V1) afferents were found distributed onto all LGN layers (parvo-, magno-, and koniocellular) as well as interlaminar zones. They were found exclusively on conventional dendrites, and never on soma or proximal dendrites. The dendritic diameter postsynaptic to these terminals was typically between 0.6–1.0 μm . GABAergic F1 terminals were found distributed onto all LGN layers and interlayer zones, but were found on soma and proximal dendrites in addition to conventional dendrites. F1 terminals were often found in close spatial proximity to RSD terminals. The resulting picture is of more terminals from excitatory cortical feedback projections, but with these synapses more distal to soma and axon hillock than their inhibitory intrathalamic TRN counterparts. Intriguingly, the latter seem placed to halt the spread of excitatory potentials down dendritic membranes. (See Harting and Feig, 1994, Figure 12.)

How might this subcellular detail be brought to bear upon the worry about the biological plausibility of our original model: that is, upon the relative effects of excitatory inputs from cortex and inhibitory inputs from TRN on activity in individual LGN relay cells? Does it legitimize or force us to revise our assumption of rough equality on parameters scaling excitatory and inhibitory inputs? Does it force us to give up on higher-level “network circuitry” modeling for our neurocomputational project? To answer these questions, we need some way to build this level of electron-microscopic anatomical and physiological detail into our model and computer simulation.

4. The Promise of Compartmental Modeling, Or: Reduction to the Rescue!

It is here that we turned to *compartmental modeling*. This approach is now dominant within computational *neuroscience*. There are few serious *neurocomputational* modelers who still work at the level of abstract network architectures. In the words of two of the approach’s principal developers, “when constructing detailed neuronal models that explicitly consider all of the potential complexities of a cell, the increasingly standard approach is to divide the neuron into a finite number of interconnected anatomical compartments... With the appropriate differential equations for each compartment, we can model the behavior of each compart-

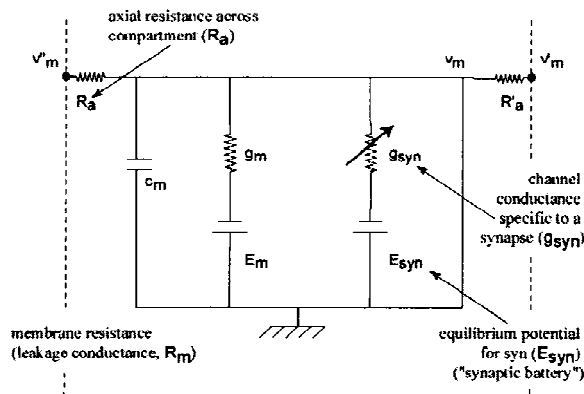


Figure 4. A GENESIS compartment Ohm's law computes the value of V_m based on the other values illustrated in the diagram. See text for details. (Adapted from Figure 6.1B in Segev, 1995, p. 65.)

ment as well as its interactions with neighboring compartments" (Bower and Beeman, 1995, p. 8). The unit of analysis is the patch of neuron membrane and its ion-specific channel conductances. The existence of general simulation software (NEURON, GENESIS) for implementing compartmental models on digital computers is partly responsible for its growing popularity. Modelers can simulate and manipulate a variety of intracellular factors that determine, e.g., rate of action potentials. These factors include the topology of membrane structure, variations in ion channels across membrane patches, and electrical field properties of post-synaptic potentials depending on synapse locations (i.e., diameter of dendritic membrane, distance from axon hillock). The GENESIS software is particularly attractive for our modeling effort. It allows us to "custom build" simulated LGN relay neurons according to the detailed intracellular and synaptic parameters revealed by electron microscope studies. But it also does not sacrifice the capacity to represent and study circuit properties and dynamics of simulated networks composed of these modeled neurons.

GENESIS computes *voltage potential across a membrane patch* V_m (a "compartment"). It uses a differential equation that expresses the rate of change to voltage potential across *membrane capacitance* C_m as proportional to the net current flowing into the compartment to charge the capacitance (Figure 4). Ohm's law calculates V_m based on

- *axial resistance* R_{a1} , R_{a2} across adjacent compartments,
- *ionic channel conductance* g_k or g_{syn} specific to an ion (or combination) k or synapse syn ,
- the *equilibrium potential* E_k or E_{syn} for that ion (or combination) or synapse (the "ionic" or "synaptic battery"), and
- the *membrane potential* R_m (leakage conductance across the membrane)




Figure 5. Our current 30-compartment model of an LGN relay neuron reflecting the relative number, location, and arrangements of excitatory cortical and inhibitory TRN synapses on its dendritic tree. The model contains twelve excitatory synaptic compartments (where $E_{syn} > E_{rest}$), five inhibitory synaptic compartments (where $E_{syn} \leq E_{rest}$), eleven “passive” compartments representing membrane distances between these synapses, and generic soma and axon compartments.

(Bower and Beeman, 1995, chapter 2; Segev and Burke, 1998). An E_{syn} value greater than resting membrane potential E_m ($E_{syn} > E_m$) represents an excitatory synapse. An E_{syn} value less than or equal to E_m ($E_{syn} \leq E_m$) represents an inhibitory synapse (Segev, 1995b). By the appropriate choices of values for the parameters defining specific compartments and for the axial resistances across adjacent compartments, modelers can construct multi-compartmental models of individual neurons to virtually any level of biological plausibility (De Shutter and Bower, 1994a, b). GENESIS even provides computational resources for calculating the cable properties of dendritic diameter (Segev, 1995a; Rall and Agmon-Snir, 1998).

Using GENESIS, we have constructed a thirty-compartment model of an LGN relay neuron that incorporates the electron-microscope discoveries presented above (Figure 5). Twenty-seven compartments represent the branching dendritic tree. Twelve represent excitatory cortical feedback synapses; four represent inhibitory TRN synapses. The inhibitory compartments lie in spatial proximity to excitatory compartments on distal dendrites. Eleven “passive” compartments represent membrane distances between these synapses. One additional compartment represents an inhibitory TRN synapse onto a proximal dendrite or soma. Two compartments represent, respectively, the cell’s soma and axon. (These two compartments receive no specialized treatment in our current model. We set processing values for the equations computing voltage potential at default values for generic soma and axon.) By choosing values judiciously in the dendritic compartments for $R_{d,s}$, $E_{syn,s}$, and g_k,s , we mimic the relative effects of these two types of synapses on LGN

relay cells. This yields a biologically-plausible test of the purely computational assumption of our original LGN-V1-TRN simulation.

Preliminary results with our single GENESIS neuron suggest that inhibitory synapse activation (from the simulated TRN projections) is very effective at inhibiting action potential generation in the cell's axon (Bickle et al., 2000). This in turn suggests that the equal scaling parameters for excitatory and inhibitory inputs to a unit in our original model probably must be revised. But we are still exploring the range of biologically plausible values for E_{syns} , g_k s, and R_a s and their ultimate effects on simulated action potential dynamics. As one would expect, different combinations of these values produce very different activity patterns. Hopefully, with enough results using different combinations of these values, patterns will begin to emerge that we can compare with existing electrophysiological data. It is here that collaboration is required between neurcomputational modelers and single-cell experimental neurophysiologists to push forward solutions to circuit-functional questions.

Our current goal is to reconstruct the entire thalamocortical network of our original study (Bickle et al., 1999), replacing the IAC units and equations within the thalamic pool with these more biologically plausible multi-compartmental models. (We will make the same replacement, albeit at a less biologically detailed level, of the cortical units of the original model with compartmental constructs.) GENESIS provides the capacity to create networks of interacting multi-compartmental neurons (Crook and Cohen, 1995; Protopapas and Bower, 1995; Koch and Segev, 1998, chapters 10 through 13). Implementing a parallelized version of this network (using PARALLEL GENESIS) on a high-speed multi-processor will enable us to repeat experiments from the earlier study. But now we will be able to track activity dynamics in a more biologically realistic LGN-V1-TRN circuit simulation. We expect results to bear on two important points.

- Results should subject the functional hypothesis for the LGN-V1-TRN circuitry suggested by our previous simulation to a sterner *neurocomputational* test. If a similar “maintenance versus decline” dynamics obtains in cells receiving the more versus the less salient stimuli in their simulated visual fields, we will have stronger evidence that this is a function of the neural circuitry we are modeling.
- Results should also yield increasingly detailed predictions about activity dynamics in actual LGN relay and cortical V1 neurons under a variety of experimental conditions. We expect not only to predict results in experimental situations not yet tested, but also to probe the activity of individual cells in our model under conditions that are difficult or even impossible to track electrophysiologically with existing experimental techniques.

Clearly, a reductionist approach came to the rescue in our specific project. Although the circuit under investigation meets the working definition of a complex system, we still had to move our modeling efforts down levels to justify the *biological plausibility* of our initial results and the ensuing functional hypothesis. The

sense of reduction we employ is quite robust. Our unit of analysis is no longer the (simplified) neuron and its anatomical connectivity. Now it is the patch of neuron membrane with its ion-specific conductance capacities and its electrochemical interactions with neighboring patches. Our model incorporates the intracellular detail revealed by electron microscopy, including relative synapse numbers and locations. Even the dendritic diameter post-synaptic to terminals gets incorporated into our multi-compartmental neuron models.

This lesson generalizes beyond our specific project. As noted above, compartmental modeling with this much and even greater amounts of electrochemical realism is *the* current standard in computational *neuroscience*: despite its virtually complete unfamiliarity among cognitive scientists and philosophers. There is a straightforward and utterly convincing reason for its popularity, which fits hand-in-glove with the reason behind the earlier shift to cellular and molecular approaches in mainstream neuroscience. The best place at present for seeking the *mechanisms* of neuronal complexity and cognition is at the level of cell and molecular biology and biochemistry, supplemented with the neuroanatomy that tells us who is talking to whom. Molecular biology and biochemistry are hardly “merely structural” branches of science. They are rife with dynamical mechanisms and interactions. Think of the biochemistry of three-dimensional protein folding (only some of which is now understood). That’s “mere structure?” Nonsense!

It is at these levels that current neuroscience’s best explanation resides about *how* the brain works, even in all its complex glory. We can learn a lot about the specific neural regions dedicated to particular cognitive and behavioral tasks using neuroimaging techniques. We can learn a lot about the capacity of certain computational operations implemented in networks that mimic simplified neurons and their anatomical connectivities. We can learn a lot by correlating careful behavioral measures with neural deficits and damage. But if we want to know *how* activity in these regions produces these effects, reductionist approaches are our only resource. Kandel, Schwartz and Jessell formulate this question nicely: “If specific mental processes are represented locally in different brain regions, what rules relate the anatomy and physiology of a region to its specific role in mentation?” (2000, p. 3). Our only answer to this question comes from cell- and molecular-level investigations: Reduction to the rescue in general, not just in selected projects. Denying the relevance of this wealth of mechanism and interaction for any project seeking comprehensive understanding of the cognizing brain is simply confused. That is the central “reductionist intuition.” It is as true about the brain today as it was when “complex systems” were just a gleam in a climatologist’s eye.¹

5. Note

¹This paper was presented originally as a talk at the International Conference for Complex Systems in Nashua, NH in May, 2000. Many thanks to John Symons, who arranged and chaired the session, fellow speakers William Bechtel, Alfredo Pereira, Joao Teixeira, and Alex Rueger, and to audience

participants. Thanks also to Marica Bernstein, who created or adapted the figures and commented on an earlier draft.

References

- Adams, N.C., Lozsati, D.A. and Guillery, R.W. (1997), 'Complexities in the thalamocortical and corticothalamic pathways', *European Journal of Neuroscience* 9/2, pp. 204–209.
- Bickle, J., Bernstein, M., Heatley, M., Worley, C. and Stiehl, S. (1999), 'A functional hypothesis for LGN-VI-TRN connectivities suggested by computer simulation', *Journal of Computational Neuroscience* 6, pp. 251–261.
- Bickle, J., Lee, D. and Bernstein, M. (2000), 'Experiments with a multi-compartment model of a thalamic relay neuron reveals the relative effects of cortical excitatory feedback and intrathalamic inhibitory synapses on membrane potential', *Society for Neuroscience Abstracts* 26, pp. 1471.
- Bower, J.M. and Beeman, D. (eds.) (1995), *The Book of GENESIS*, New York: Springer-Verlag.
- Crick, F. and Koch, C. (1990), 'Towards a neurobiological theory of consciousness', *Seminars in the Neurosciences* 2, pp. 263–275.
- Crook, S. and A. Cohen (1995). 'Central pattern generators', in J.M. Bower and D. Beeman, pp. 141–158.
- De Shutter, E. and Bower, J.M. (1994a), 'An active membrane model of the cerebellar Purkinje cell I. Simulation of current clamps in slice', *Journal of Neurophysiology* 71, pp. 375–400.
- De Shutter, E. and Bower, J.M. (1994b), 'An active membrane model of the cerebellar Purkinje cell II. Simulation of the synaptic responses', *Journal of Neurophysiology* 71, pp. 401–419.
- Feig, S. and Harting, J.K. (1994), 'Ultrastructural studies of the primate lateral geniculate nucleus: Morphology and spatial relationships of axon terminals arising from the retina, visual cortex (area 17), superior colliculus, parabrachial nucleus, and pretectum of *Galago crassicaudatus*', *Journal of Comparative Neurology* 343, pp. 17–34.
- Grossberg, S. (1978), 'A theory of visual coding, memory, and development', in E.J. Leeuwenberg and H.F.J.M. Buffart, eds., *Formal Theories of Visual Perception*, New York: Wiley.
- Harting, J.K., van Lieshout D.P. and Feig, S. (1991), 'Connectional studies of the primate lateral geniculate nucleus: Distribution of axons arising from the thalamic reticular nucleus of *Galago crassicaudatus*', *Journal of Comparative Neurology* 310, pp. 411–427.
- Jones, E.G. (1985), *The Thalamus*, New York: Plenum Press.
- Kandel, E.R., Schwartz J.H. and Jessell, T.M. (eds.) (1991), *Principles of Neural Science*, 3rd Ed, New York: McGraw-Hill.
- Kandel, E.R., Schwartz, J.H. and Jessell, T.M. (eds) (2000). *Principles of Neural Science*, 4th Ed, New York: McGraw-Hill.
- Koch, C. and Segev, I. (eds.) (1998), *Methods in Neuronal Modeling: From Ions to Networks*, 2nd Ed., Cambridge, MA: MIT Press.
- Kolb, B. and Whishaw, I.Q. (1996), *Fundamentals of Human Neuropsychology*, New York: Freeman.
- McClelland, J.L. and Rumelhart, D.A. (1987), *Explorations in Parallel Distributed Processing: A Handbook of Models, Programs, and Exercises*, Cambridge, MA: MIT Press.
- Montero, V.M. (1991), 'A quantitative study of synaptic contacts on interneurons and relay cells of the cat lateral geniculate nucleus', *Experimental Brain Research* 86, pp. 257–270.
- Protopapas, A. and Bower, J.M. (1995), 'Dynamics of cerebral cortical networks', in J.M. Bower and D. Beeman, eds., pp. 159–179.
- Rail, W. and Agmon-Snir, H. (1998), 'Cable theory for dendritic neurons', in C. Koch and I. Segev, eds., pp. 27–92.
- Segev, I. (1995a), 'Cable and compartmental models of dendritic trees', in J.M. Bower and D. Beeman, eds., pp. 53–81.

- Segev, I. (1995b), 'Temporal interactions between post-synaptic potentials', in J.M. Bower and D. Beeman, eds., pp. 83–101.
- Segev, I. and Burke, R.E. (1998), 'Compartmental models of complex neurons', in C. Koch and I. Segev, eds., pp. 93–136.
- Sherman, S.M. and Guillery, R.W. (1996), 'The functional organization of thalamocortical relays', *Journal of Neurophysiology* 76, pp. 1367–1395.
- Sherman, S.M. and Koch, C. (1998), 'Thalamus', in G. Shepherd, ed., *The Synaptic Organization of the Brain*, 4th Ed, New York: Oxford University Press, pp. 289–328.
- Steriade, M., McCormick, D.A. and Sejnowski, T.J. (1993), 'Thalamocortical oscillations in the sleeping and aroused brain', *Science* 262, pp. 679–685.