Part II

Creating Simulations with GENESIS

Chapter 11
Constructing New Models

JAMES M. BOWER

You have now explored several tutorials intended to both provide some insights into fundamental concepts in neurobiology and to introduce you to the use of neural simulations. The remainder of this book is intended to guide you in the creation of your own simulations. In general, we have found that the most efficient way to develop new simulations is to modify one that already exists and that you understand well. We would encourage you to consider starting with a tutorial most similar to the simulation you would like to build, after having learned a few basics of GENESIS programming from the following chapters.

As discussed in Chapter 3, GENESIS has been specifically designed to allow user modification of existing simulations as well as the easy incorporation of new biological components. At the time of its development, no simulator existed that could easily support models based on the anatomical and physiological details of real nervous systems, or that could be easily modified to do so. Instead, those interested in building more realistic models of neural structures (primarily single neurons) either had to write their own special purpose code (Getting 1989, Pellionisz and Llinás 1977) or use a simulator designed for another purpose. For example, some neural modelers used electric circuit simulators originally intended for use by the electronics industry (Segev, Fleshman and Burke 1989, Segev, Fleshman, Miller and Bunow 1985). In other cases, modelers tried to adopt simulators constructed to model abstract "neural networks" of the connectionist type (Goddard, Fanty and Lynne 1987). In each case, including descriptions of real neural components was tedious, time consuming and limited.

GENESIS was designed in response to this situation. We specifically set out to build a simulator capable, in principle, of supporting any level of biological detail. Furthermore, we wanted to assure that as new properties and details of biological systems were discovered, they could be easily incorporated into the GENESIS system. Accordingly, this effort was fundamentally based on the assumption that the structural and physiological details of the nervous system matter. In fact, we believe that an eventual understanding of the way nervous systems compute will be very closely dependent on understanding the full details of their structure. For this reason, our own approach to modeling involves constructing computer simulations that are very closely linked to the detailed anatomical and physiological structure of the particular neuronal system being studied (Bower 1995). We refer to the resulting simulations as “structurally realistic” models to distinguish them from the more abstract types of models that have formed the principal basis for most neural modeling efforts in the past. Before beginning to describe the process of building and modifying GENESIS simulations, it will be useful to first discuss several general and specific issues relevant to our philosophy of model construction. To a considerable extent, these issues and philosophies have guided the construction of GENESIS.

11.1 Structurally Realistic Modeling

Every indication suggests that structurally realistic modeling will represent one of the largest areas of growth in computational neurobiology (Bower 1992). Continuing rapid advances in computer software and hardware technology are making it possible for neurobiologists themselves to build models. Typically, when neurobiologists are in charge of modeling efforts, they include more of the fine biological details. In addition, experimental technology has advanced to the point that the necessary information to construct realistic models can be more readily obtained.

Although technology is clearly enabling the development of more and more complex models, we do not believe that the simple capacity to make such models is the reason they should be built. Instead, as neurobiologists and modelers, we are more and more convinced that understanding, or “reverse engineering” (Bower 1995) the nervous system will critically depend on the construction of anatomically and physiologically realistic simulations.

Several practical arguments can be made in support of constructing realistic models. First, simply from the point of view of model construction and use, these simulations have several advantages over more abstract modeling efforts: For example, biological realism allows known neuroanatomical and neurophysiological data to be used as constraints for establishing model components. This is an important feature for models of complex systems that can otherwise posit an almost limitless number of components and interactions. In addition, the relationship between the parameters in a structural model and real biological measurements limits the parameter space of the models that must be explored (Bhalla and Bower 1993, Jaeger, De Schutter and Bower 1997). Because most of the time spent working with a model involves determining the dependence of its activity on model parameters, this greatly increases the chances that modeling will yield something interesting. By generating biologically relevant outputs, modeling results are also more readily comparable to data from actual experiments, making predictions testable (Hasselmo, Barkai, Horwitz and Bergman 1994, Jaeger et al. 1997). Testability is one of the most serious problems facing all models of complex systems. Finally, the process of building such a simulation itself, in effect, highlights what information still must be obtained to build the model. This often requires the modeler to more carefully quantify exactly what is known about a particular neuron or circuit. In fact, in our experience the early stages of model building often point out more vividly what is not known about a network than what is. In this way, structural models lead naturally to ideas for additional experiments as well as provide a context in which to interpret the data once they are obtained.

Beyond the practical advantages of building realistic simulations, we believe there are also very important pedagogical reasons for constructing simulations of this kind. Specifically, we believe that, when used properly, these models have an increased chance of generating unanticipated functional insights based on emergent properties of neuronal structure. Traditionally, most models, especially those that are structurally abstract, have been less concerned with the details of neural organization than in testing a particular pre-existing theory. Those that have been concerned with neural structure still have often been primarily intent on proving biological plausibility for a particular pre-existing idea rather than using the structure of the nervous system itself to suggest new functional ideas. Although there may be nothing, in principle, wrong with these "theory demonstration" approaches, we believe that models which replicate the structure of the nervous system as a basis for exploring its computational features are more likely to uncover features that had been previously overlooked or unsuspected (cf., Hasselmo and Bower 1992, Hasselmo, Anderson and Bower 1992, Bower 1995, Bower 1997a,b). Once a new idea has emerged, abstract models can always be built to more completely explore parameter space (Hasselmo et al. 1992) and/or to cast new ideas into more traditional theoretical forms (Hasselmo 1993). However, if a realistic model comes first, it is easier to make specific physiological predictions. In our experience, it is also easier to make a realistic model abstract than it is to make an abstract model realistic. In any case, the hardest test of any model should be how well the modeler can answer both, "What do you know now that you did not know before?" as well as "How can you determine its likelihood?" On both counts, we believe biologically realistic models have a distinct advantage.

Although realistic modeling has an increasing number of converts and successes, one still often hears the criticism that the modeler in these cases is simply substituting the problem of a complex neuron or network in an animal for a complex neuron or network in a computer (cf., Churchland and Sejnowski 1988). This criticism, however, does not adequately take into account the fact that the construction of realistic models is really a process rather than an end in itself (Bower 1995, Bower 1997a,b). As mentioned, one of the principal values of this type of modeling is that it makes very clear what you do not know and what you need to find out. Furthermore, such criticisms of realistic modeling are also often based on the assumption, or hope, that the specific details of the nervous system might not matter. As the neural theorist David Marr proposed (Marr 1982), in this view, any particular biological neuron or network should be thought of as just one implementation of a more general computational algorithm. Following a physics model, understanding the algorithm is thought to represent a more general form of understanding than can be provided by considering the details of the nervous system itself. However, it is not clear that the "technology of understanding" in the simple systems typically studied by physicists will readily apply to exceedingly complex systems like brains, which are also built as a consequence of biological evolution. For example, it is at least possible that evolutionary pressure has forced the elimination of "unnecessary components" resulting in a very tight relationship between the structure of the nervous system and its function. If correct, then a particular neural system might not represent one of many ways to perform a particular computation, but instead, could represent a much more limited solution set, or even the only solution. This is especially true if the magnitude of the computational problems solved by nervous systems are much greater than we even now realize. The dominating trend in the history of artificial intelligence is arguably the discovery of how hard these problems really are. Similarly, the recent history of engineering efforts like "neural networks" and "connectionism" has involved the development of ever more complex computing structures, despite the predisposition of those in the field toward simplicity.

For all of these reasons, we see no evidence yet to suggest that we can avoid considering all the anatomical and physiological details in our pursuit of a full understanding of nervous systems. As biologists, given the apparent complexity of these systems, coupled with our profound ignorance about how they work, or even what kind of machine they are (Nelson and Bower 1990), it seems prudent to err on the side of biology and include the details in our models, in the hope that structure will guide us to function.

11.2 The Modeling Process

Although understanding the brain may involve an eventual consideration of all its details, model building need not await a thorough description of the neural system in question. In fact, the construction of realistic simulations may be essential to the process of experimental design, aiding in the choice of which data, and therefore which experiments are currently most relevant to our evolving understanding of neural function. Neurobiology is approaching the point where the massive amount of data already obtained may more impede progress than promote it. We believe that modeling can serve as an important tool to determine which experimental investigations will yield the most useful data at a particular time. Realistic models may also represent an extremely efficient way to store the information obtained. In this regard, we are currently exploring the use of GENESIS as an actual data base for neurobiological information. Over the next several years, we will release software to allow neurobiologists to "mine" the information about neural structures that modelers have incorporated in GENESIS.

From the point of view of the modeler, these issues come down to the question of how and where to start constructing a model. It is important to realize that this approach to modeling does not simply involve packing all known features of a particular network into a computer simulation. Even if all the information about a neuron or network were known, as it never is, including all the details at the beginning of a modeling effort is often not practical, even with today's computers. Instead, model building is a stepwise process that we believe should start by including the minimal principal features of the network of interest. These are then added to, as necessary, to replicate specific experimental results (Bower 1995, Jaeger et al. 1997).

11.2.1 Single Neurons or Networks

From a practical point of view, the best place to start a new modeling project is with an existing model, but which model and at which level? Often the first decision that must be made is whether to begin with a detailed model of a single cell (Chapter 7) or a more general model of a network of cells (Chapters 8 or 9). As all neurons operate as parts of networks, this question comes up even in those animals with relatively small nervous systems (Selverston 1985, Harris-Warrick, Marder, Selverston and Moulins 1992).

The initial appropriate answer to this question is dependent on a number of different considerations. From the previous chapters, you should already be familiar with some of the tradeoffs in different levels of modeling. Detailed models of single cells can provide specific information about the dynamics of cellular responses to synaptic input, but do not necessarily provide any clue as to how the cell is normally synaptically activated. The more complex the cellular model, the more are the possibilities of different patterns of synaptic activation. On the other hand, models of networks usually employ more simplified neurons without the detailed channel properties that have a profound effect on cellular output and therefore on network activity. So whatever level of modeling one starts with, there are necessary abstractions leading to important limitations. For this reason, the decision about where to start usually revolves around the type of experimental data most readily available to the modeler. This is particularly true if the modeler is also an experimentalist using particular techniques, which we regard as the ideal situation.

To some extent, the question of where to start is somewhat mitigated by the likelihood that whatever level of realistic modeling one starts with, it is likely that one will eventually end up doing the other kind of modeling as well. For example, our own studies of the olfactory cortex began with network modeling (Wilson and Bower 1992), but soon required us to develop more detailed models of single pyramidal cells (Protopapas and Bower 1994). In contrast, our detailed cellular modeling of olfactory bulb mitral cells (Bhalla and Bower

1993) has now required that we develop a network model of the bulb to more completely interpret the data. Similarly, our detailed modeling of cerebellar Purkinje cells (De Schutter and Bower 1994a,b,c, Jaeger et al. 1997) is leading us to develop a model of the cerebellar cortical circuitry to better understand the pattern of synaptic input to these neurons (Santamaria and Bower 1997). Accordingly, no matter where one starts building realistic models of a particular structure, eventually the realism of all components of the model will probably need to be enhanced. The same trend can be seen with respect to the interactions between neural structures. Thus, one motivation for originally constructed detailed simulations of the olfactory bulb's mitral cells was to provide more realistic input to the olfactory cortex model (Bhalla and Bower 1993). This single cell modeling, however, led to electrophysiological experiments (Bhalla and Bower 1997) whose interpretation now requires an increase in the realism of our olfactory cortex model to understand the effects of cortical feedback on mitral cell behavior. Although this may start to sound like a never-ending process, few biologists ever said that understanding the brain would be easy. Such statements have usually been made by engineers and physicists. Probably the most important advice we can give is to simply get started and not worry about it.

11.2.2 Modeling Steps

Retreating from the somewhat daunting prospect of eventually needing to model the entire system, there are several practical suggestions that can be made about how to start modeling, regardless of the level at which one starts. In our approach (Bower 1995, De Schutter and Bower 1994a,b), the first stage of model building usually involves establishing sufficient structural detail in the simulation to be able to replicate a set of well-characterized physiological responses. This serves to ensure, at the outset, that modeling predictions are testable experimentally. We have found that it is actually most useful to start with what we refer to as "non-physiological responses" that nevertheless reflect the overall organization of the structure in question. Responses of this type are typically experimentally generated, activating a neuron or network in some completely unnatural way. For example, in our network modeling effort with piriform cortex, we initially sought to replicate the oscillatory cortical responses obtained experimentally by artificially shocking the input pathway to the cortex (Wilson and Bower 1992). Replicating responses from these direct electrical shocks turned out to be a good initial test of the model because such massively evoked cortical activity was less dependent on the more detailed network properties that were not yet included in the simulation. In our detailed models of single cells, the initial criterion for the accuracy of a model is often its ability to replicate voltage or current clamp data (De Schutter and Bower 1994a, Jaeger et al. 1997). In this case the experimentalist has artificially injected current into a neuron with an electrode rather than with more natural synaptic activation. Although the stimulus is artificial, the responses can reflect very global features of the cell. For example, in Chapter 7 we observed a transition from bursting to single spike firing in the

Traub pyramidal cell model as the injection current was increased. Although this behavior is only seen under laboratory conditions, it is an important initial test of the model.

A second important feature of this initial stage of model building involves simulating results from as many different recording methods as possible. Thus, our piriform cortex network simulations were designed from the start to generate intracellular potentials from simulated neurons, extracellular spike train activity, and extracellular bulk electrical responses such as evoked potentials or EEGs. At the level of single cells, our Purkinje cell models not only replicate transmembrane voltage potentials, but also levels of intracellular calcium (De Schutter and Bower 1994a,b,c). The more types of responses a model generates, the more rigorously it can be tested. This effort at the beginning also ensures that the results of later simulations can be tested using diverse types of real experimental data. This is important because it is often not clear at the beginning of a modeling effort which experimental technique will provide the best means to test modeling predictions.

Once a model has been tuned on these presumably "function neutral" measures, the modeler has more confidence that the model itself includes enough of the essential features to begin exploring functional properties of the model. In practice, for single cell models, this usually involves adding synaptic conductances (De Schutter and Bower 1994b), whereas in network models it often involves introducing more complex patterns of synaptic activation or synaptic plasticity (Protopapas and Bower 1994). However, it is important to mention that even though we expected that exploring function would depend on adding these features to the models, we found that, in each of our modeling efforts, new, functionally interesting features of neural structure became apparent even in the process of tuning the model (Bower 1995, Bhalla and Bower 1993, De Schutter and Bower 1994a,b, Bower 1997a,b). This repeated experience is one basis for our claim that paying attention to the anatomical and physiological details of real neural systems is very likely to lead to important clues to their function. It seems quite likely that the relationship between the structure of nervous systems and their function may be tighter than for any other machine we know. Time and modeling will tell.