Chapter 9
Dynamics of Cerebral Cortical Networks

ALEXANDER PROTOPAPAS and JAMES M. BOWER

9.1 Introduction

Previous chapters in this volume have considered detailed models of single cells and small networks of cells. In this chapter we consider a large scale multicellular model of the mammalian olfactory cortex. The simulation consists of three distinct neuronal populations of 135 cells each for a total of 405 interconnected neurons. With this simulation, we will explore the possible physiological basis for experimentally recorded electroencephalographic patterns in this cortex.

When constructing a model of any neural system, one must always strike a balance between biological realism and computational efficiency. This is especially true in the case of network models. The real piriform cortex of a rat, for example, contains on the order of 10^6 neurons (Haberly 1990). Even when the complexity of individual neurons is reduced, it is still not possible to simulate all the neurons found in this network. Accordingly, the modeler is always faced with determining the level of detail necessary to explore and illuminate the particular physiological and computational properties of these networks. An important question becomes: how do the physiological properties or computational capabilities of a cortical network scale with the number of neurons? In this tutorial we demonstrate how at least a rough understanding of network behavior can be obtained with a quasirealistic model of cerebral cortex. First, however, we briefly introduce the piriform cortex.

9.2 Piriform Cortex

The piriform cortex is the primary olfactory cortical area in the mammalian brain. For the last several years we have been constructing realistic models of this network (Wilson and Bower 1992) with the ultimate objective of understanding its role in olfactory object recognition (Hasselmo and Bower 1993, Bower 1995). One motivation for this work is our assumption that this cortex computationally represents a kind of associative memory (Haberly 1985, Haberly and Bower 1989).

Piriform cortex receives its afferent input from the olfactory bulb which itself receives input directly from the nasal epithelium where the olfactory receptors are located. Thus, piriform cortex is quite close to the sensory periphery, and unlike other sensory cortical areas, it does not receive its afferent input through the thalamus. Piriform cortex sends its primary projections to the entorhinal cortex but also connects to the thalamus, olfactory tubercle, superior colliculus, peri-amygdaloid, and has strong reciprocal connections with the olfactory bulb (Haberly 1990). Figure 9.1 illustrates the connections between piriform cortex and related areas.

Piriform cortex is believed to be phylogenetically older than other sensory cortical areas and is commonly referred to as *paleocortex*. It also has a particularly well-defined and somewhat simpler anatomical organization than neocortical regions (Haberly 1985). For example, it has only three layers instead of the six normally found in neocortex. In this sense, it is similar to the hippocampus which also has a trilaminar structure. There are considerable physiological data available on its neurons, their interactions, and on network level responses (Haberly 1990). The detailed internal structure of the piriform cortex is discussed in the context of the model's implementation.

9.3 Structure of the Model

The simulation discussed in this chapter was originally constructed by Matt Wilson when he was a graduate student at the California Institute of Technology. In fact, this simulation served as the initial basis for the construction of GENESIS itself. Portions of the model description below were taken from a paper originally written by Wilson and Bower (Wilson and Bower 1989). The graphical interface was later added to make the simulation user-friendly. Although this model has been used to explore a wide range of cortical behavior (Wilson and Bower 1989, Wilson 1990), including associative memory function (Wilson and Bower 1988, Hasselmo, Wilson, Anderson and Bower 1990), the tutorial version has been simplified for the sake of computational speed and pedagogical ease. In its current manifestation, the model allows you to reproduce experimental EEG patterns and to explore their possible physiological basis. Using later chapters of this book, the user can expand this simulation as he or she wishes.

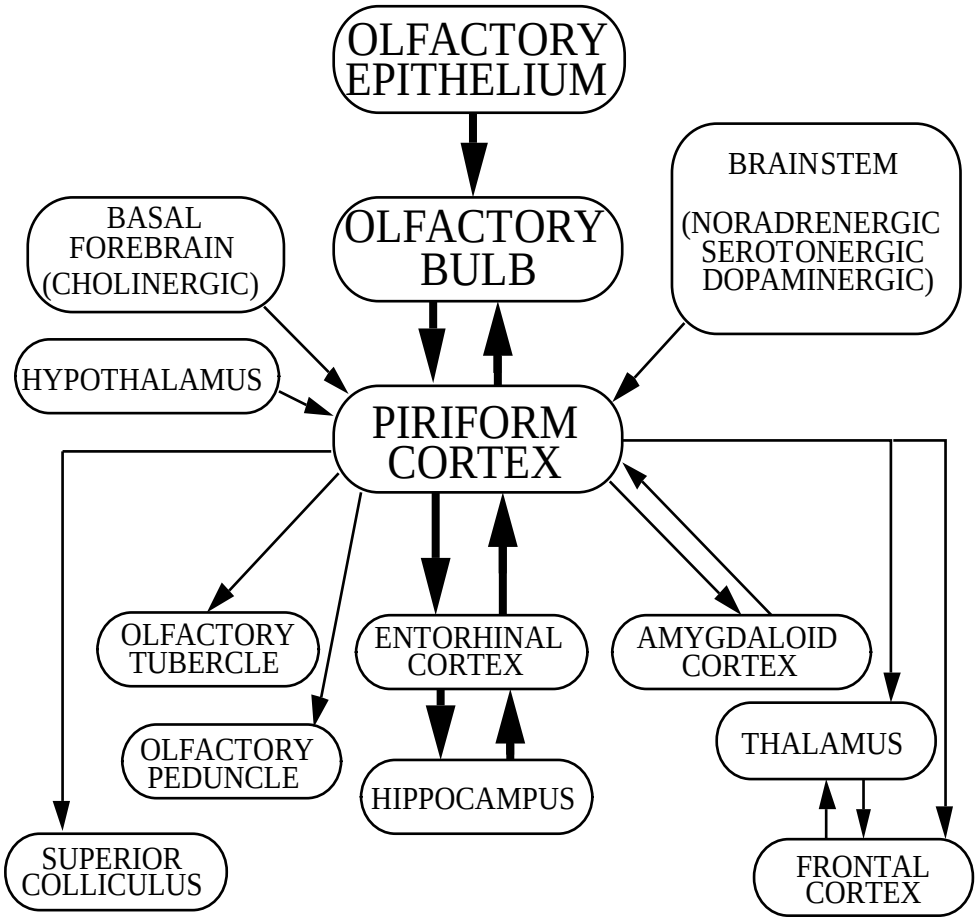

Figure 9.1 Piriform cortex and connected areas. Major pathways are indicated by thick arrows. In the model, only the olfactory bulb to piriform cortex connection is included. The feedback pathway is not incorporated.

9.3.1 Cellular Complexity

Pyramidal cells are the principal cell type in piriform cortex, and are believed to be exclusively excitatory (Haberly and Price 1978, Haberly and Bower 1984). Superficial and deep pyramidal cells form two distinct populations that differ significantly in their physiology (Tseng and Haberly 1989). There are also several populations of nonpyramidal cells or interneurons that can be distinguished on anatomical grounds (Haberly 1983). These neurons appear to be GABAergic and seem to mediate both feedback and feedforward inhibitory effects.

As mentioned previously, computational limitations require that realistic network models be constructed of fewer neurons than are actually found in the brain. The current model

is based on a single population of pyramidal cells plus two populations of inhibitory interneurons. A total of 135 neurons of each type are simulated, yet these neurons are intended to represent the full extent of the actual cortex (approximately 10 $mm \times 6$ mm). Accordingly, although we simulate neurons individually, the output of each neuron is taken to represent the average activity of a larger group of cells that would normally be in its region. This adjustment for scale is made in the strength of synaptic connections between the cells, and in the number of cells contacted by a particular neuron within the simulated network.

Network models also usually include much simpler representations of neurons than are found in realistic single cell simulations. In the current case, pyramidal cells are modeled using five electrical compartments, whereas interneurons are modeled using only one. We have chosen to model pyramidal cells with five compartments for several reasons. First, in order to accurately model field potentials, it is necessary to distribute synaptic inputs spatially along the dendritic processes of these cells. Secondly, the spatio-temporal distribution of synaptic input along the dendritic tree of the cell may be computationally relevant. Much more realistic multi-compartment models are currently being used to explore this possibility (Protopapas and Bower 1994). Little is lost in modeling inhibitory neurons as single compartments because little is known about the organization of their synaptic input and their small size and radially spanning dendrites preclude them from making significant contributions to the EEG.

In addition to the reduction in the number of compartments for each cell, the membrane properties of these neurons have also been simplified. For example, although experimental evidence indicates that there are a number of Ca^{2+} and K^{+} currents in piriform pyramidal cells (Constanti and Sim 1987, Constanti, Galvan, Franz and Sim 1985, Constanti and Galvan 1983) in addition to standard Hodgkin-Huxley Na^{+} and K^{+} currents, none of these are modeled in any of the cells. Rather, a simple threshold criterion is applied to the membrane potential to generate discrete spike events. The occurrence of spikes is indicated with a spike waveform "pasted" onto the actual membrane potential at the appropriate time. In this way, the computationally expensive details of spike generation are avoided but we are still able to generate the appropriate currents and membrane potentials associated with real action potentials. Similarly, synaptic conductances are modeled neglecting computationally expensive details such as the kinetics of ligand binding, neurotransmitter uptake, etc. Instead, changes in synaptic conductance are modeled as the difference of exponential functions that approximates the shape of EPSPs seen in experimental studies.

9.3.2 Network Circuitry

Primary afferent input enters piriform cortex via the lateral olfactory tract (LOT) projection from the mitral and tufted cells of the olfactory bulb. This is shown in Fig. 9.2. Experimental evidence suggests that this projection is exclusively excitatory (Biedenbach and

Stevens 1969a, Biedenbach and Stevens 1969b, Haberly 1973b, Haberly and Bower 1984) and extremely diffuse or non-topographic. These afferents make excitatory connections to pyramidal cells and feedforward inhibitory cells in layer Ia (Haberly 1985). LOT input is modeled as a set of independent fibers that make sparse connections with pyramidal cells and both types of inhibitory interneurons. In this tutorial, activity along the afferent pathway is random over time. In both the actual cortex and the model, conduction velocities along axons are finite and vary with the axonal type (Haberly 1978). Signals travel along the LOT rostrally to caudally, and are distributed across the cortex via many small collaterals (Devor 1976). In the model, as in the brain, signals proceed along the LOT towards the cortex at a speed of 7.0 m/s. Collaterals leave the main fiber tract at a 45° angle and travel across the cortex at a speed of 1.6 m/s (Haberly 1973a). In the biological cortex there is a diminution of afferent input to pyramidal cells moving rostrally to caudally that is reflected anatomically in the number of synaptic terminals (Price 1973, Schwob and Price 1978), and physiologically in the amplitude of shock-evoked potentials mediated by the afferent system (Haberly 1973a). To simulate this effect in the model, the strength of synaptic input due to afferent signals is exponentially attenuated with increased distance from the rostral site of stimulation.

In addition to excitatory input from the bulb, pyramidal cells within the piriform cortex make excitatory connections with other pyramidal cells across the entire cortex (Biedenbach and Stevens 1969a, Haberly and Bower 1984). The fibers appear to spread out radially from the originating cell and travel rostrally at a speed of 1.0 m/s, and caudally at a speed of 0.5 m/s (Haberly 1973a, Haberly 1978) making local connections on basal dendrites of other pyramidal cells and distant connections on apical dendrites. In the model, fibers originating from pyramidal cells follow the same pattern of interconnectivity and signals are propa-

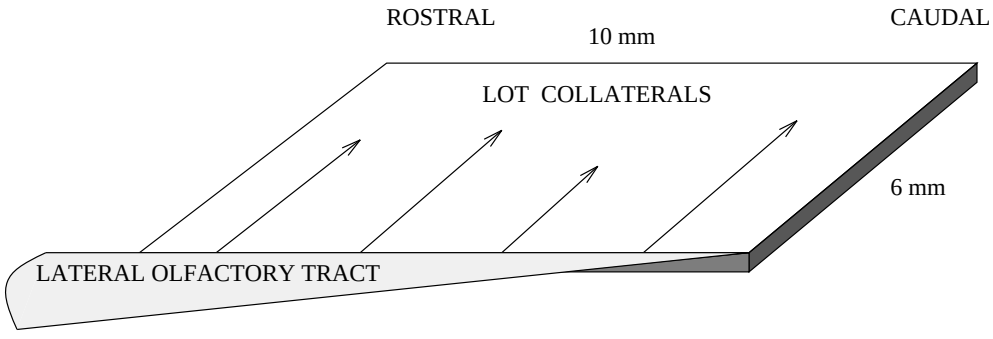

Figure 9.2 The diagram indicates the distribution of afferent input to real and simulated piriform cortex. Input from the olfactory bulb arrives via the lateral olfactory tract which then sends perpendicular collaterals into the cortex that make sparse connections with piriform cells. The number of connections between the LOT and the cortex decreases as one travels rostrally to caudally.

gated along each fiber with the corresponding delays. Simulation scaling considerations as described earlier require that association fiber interconnectivity be greatly increased as compared with that of the actual cortex. As with afferent input, intrinsic excitatory connections are attenuated exponentially with distance from the originating cell.

Experimental evidence suggests the existence of two types of inhibition in the piriform cortex, both of which are incorporated into the model. A well-documented Cl^- mediated feedback inhibition is thought to be generated by local interneurons that receive input primarily from local pyramidal cells as well as some afferent fibers (Biedenbach and Stevens 1969a, Biedenbach and Stevens 1969b, Haberly 1973b, Satou, Mori, Tazawa and Takagi 1982, Haberly and Bower 1984). The outputs of these inhibitory interneurons feed back into nearby pyramidal cells where they activate a significant Cl^- conductance increase at the cell body. A K^+ mediated inhibition is also present in the biological cortex and appears to be generated by local inhibitory interneurons receiving primarily direct afferent input from the LOT as well as some association pathway input from pyramidal cells (Satou et al. 1982, Tseng and Haberly 1986). The outputs of these interneurons generate a long-latency, long-duration hyperpolarizing inhibitory potential in nearby pyramidal cells in the most distal part of the pyramidal cell apical dendrite. In the model, the K^+ mediated inhibition is activated on the apical pyramidal cell dendrite by inhibitory neurons with both feedforward and feedback input. The network circuitry for the model is illustrated in Fig. 9.3.

9.4 Electroencephalography

Because the electroencephalogram (EEG) is the physiological measure that this tutorial attempts to simulate, it merits some explanation. Unfortunately, the EEG is something which has a long history in neurophysiology but whose origins are still debated. In general, physiological measurement techniques can be grouped into those that record the responses of individual neurons, and those that measure the more complex aggregate electrical activity of networks of cells. Single cell electrical recordings can be either intracellular or extracellular and reflect the actual output of single neurons. The origins and significance of aggregate recordings such as the EEG are more difficult to determine. The EEG is identical to extracellular single unit recording in that it measures the field potentials generated in the space around neurons. It differs because EEG recordings represent electrical activity over a wider area of the brain. Typically, EEGs are recorded from an array of electrodes placed on the surface of the brain or even the scalp. In this sense, one may think of the EEG as the average electrical activity of many neurons over a sizable area. The EEG is calculated in the model using an array of 40 evenly spaced electrodes on the surface of the simulated cortex. Recordings from the array are averaged to produce the EEG.

Understanding in detail how extracellular field potentials such as the EEG are related to the activity of collections of single cells is not a straightforward matter. A neuron can

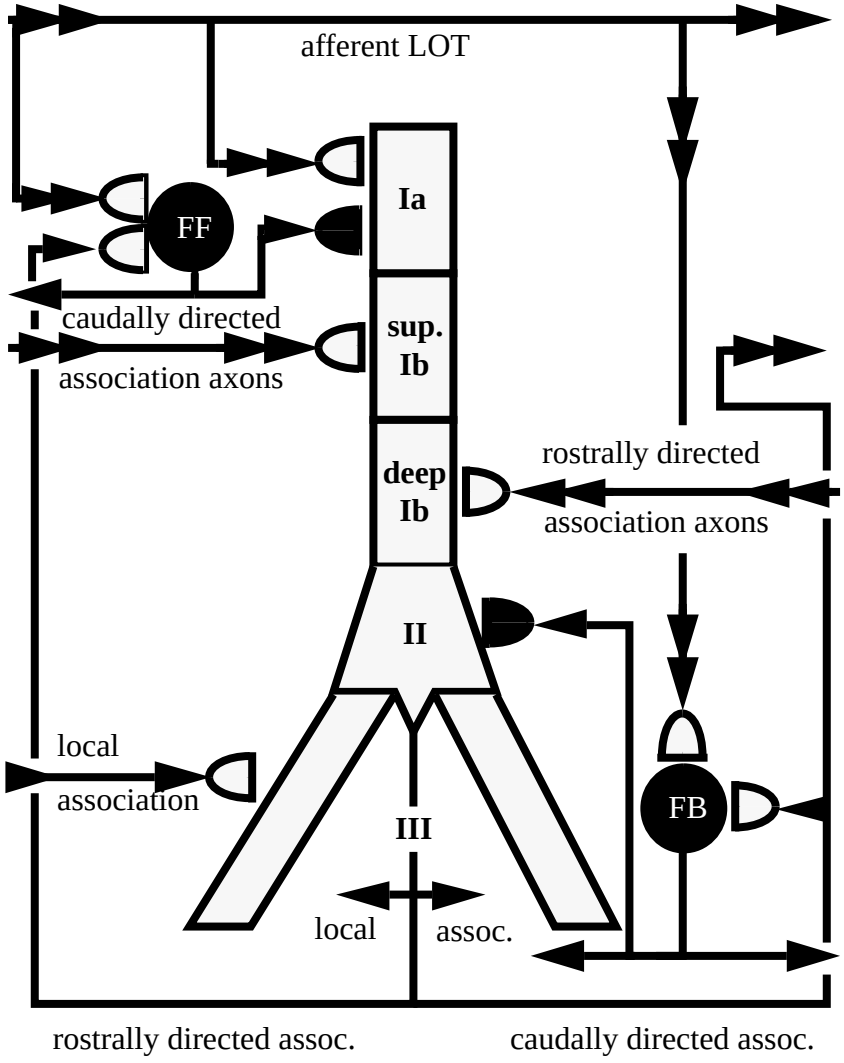

Figure 9.3 Diagram of local model circuitry. The large cell in the center represents a pyramidal neuron. Inhibitory cells are represented by black fi lled circles. The two types of inhibitory interneurons are distinguished as FF (feedforward) and FB (feedback). Single arrowheads signify local pathways. Double arrowheads are used to show distant pathways. Excitatory synapses are shown as lightly colored cones and inhibitory synapses are black. In the model, the basal dendrites are modeled as a single compartment.

be thought of as a very complicated circuit consisting of resistors, capacitors, and batteries (see earlier chapters). Like any circuit, it must obey Kirchoff's current law, which states that the sum of the total current entering and leaving a circuit node must equal zero. In the case of a neuron, this means that if synaptic current, for example, enters the cell at one point, it must leak from another, thereby generating an extracellular current (see Fig. 9.4). An area of the neuron where current is entering the cell is called a current sink. An area where current flows outward is called a current source. Here, we use the "physiologists' convention" (Sec. 4.3.2) which holds that inward current is negative and outward current is positive. When treating current sources and sinks discretely as we do in the model, the field potential is dependent on these extracellular currents according to the equation:

$$\Phi(\vec{r},t) = \frac{1}{4\pi\sigma}\sum_{i=1}^{n}\frac{I_i(t)}{R_i}, \tag{9.1}$$

where Φ is the field potential in volts, $I_i(t)$ is the total current (amperes) from the ith current source into the brain tissue of conductivity σ ($\Omega^{-1}m^{-1}$), and R_i (meters) is the distance of the ith current source from the field point $\vec{r}$ (Nunez 1981). Although this equation is used to calculate field potentials in the model, it is based on the assumption that the brain is a homogeneous conductor, which is only an approximation to biological reality. The important thing to note here is that the size of the field potential increases in amplitude with the magnitude of extracellular current and decreases with distance between current source (or sink) and electrode. This has certain important implications. For example, it suggests that the action potentials of individual neurons often make little contribution to the EEG. Because the extracellular currents produced during spike generation are generally small, the greater magnitude of synaptic currents makes more significant contributions. Since the EEG represents the averaged electrical activity of many neurons, the more synchronous the activity, the stronger the signal will be.

Although understanding the basis for EEG activity is by no means straightforward, many attempts have been made to correlate EEG patterns with certain types of animal and human behavior. EEGs are primarily distinguished on the basis of their frequency. The two frequencies most commonly discussed in the context of the olfactory system are the theta (4–7 *Hz*) and gamma (30–80 *Hz*). These types of EEG activity are common, not only in the olfactory system, but throughout the brain. They are also found across species (Ketchum and Haberly 1991). In the rat, hippocampal EEGs in the theta range are prominent during exploratory behavior (Ranck 1973). More recent experiments have indicated that theta-patterned stimulation in the hippocampus is optimal for the induction of hippocampal long term potentiation (Larson, Wong and Lynch 1986, Staubli and Lynch 1987). Theta stimulation has been used to successfully induce LTP in the piriform cortex as well (Kanter and Haberly 1990). Interestingly, the theta rhythm approximates the rate of exploratory sniffing in the rat (Macrides, Eichenbaum and Forbes 1982).

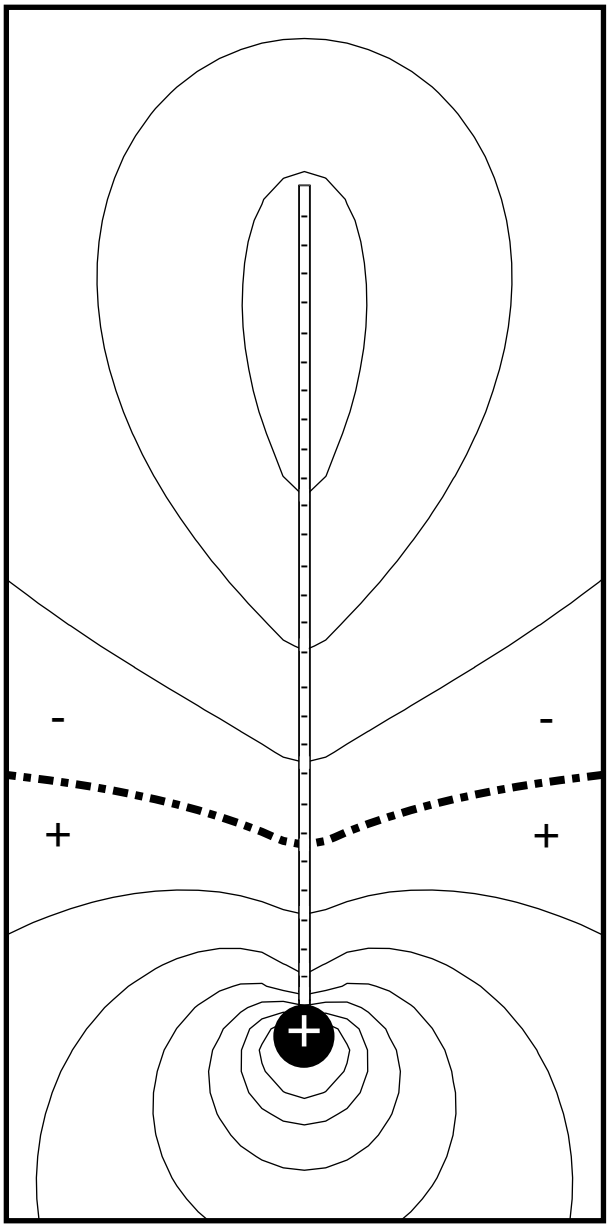

Figure 9.4 A model cell shows the distribution of current sources and sinks and extracellular isopotential contours at the moment when the cell is receiving excitatory input all along its apical dendrite. The thick contour shows the area of zero potential. Contour lines above the zero line give negative potentials. Below the zero line are positive potentials. Minus signs represent current sinks (current fbwing into the cell) and the plus sign depicts a current source (current fbwing out of the cell).

The gamma frequency (in the 40 *Hz* frequency range) has recently been the focus of intense experimental and modeling efforts. One of the earliest observations of the gamma frequency was in the olfactory system of the hedgehog in response to an odor stimulus (Adrian 1942). Since then, these oscillations have been found in a variety of cortical areas. Numerous theories have evolved to explain the significance of the gamma rhythm. Researchers have proposed that it is a solution to the binding problem (Gray, Konig, Engel and Singer 1989), a cortical information carrier (Bressler 1990), and even a hallmark of consciousness (Crick and Koch 1990). Although such statements are highly speculative, the ubiquity of these events suggests that the activity underlying the 40–60 *Hz* oscillations is related in some way to neural computation. The research model on which this tutorial is based was used to explore the possible physiological underpinnings of these oscillations (Wilson and Bower 1991) and led to the conclusion that cortical oscillations do not represent an information code, as suggested in some of the speculations above, but rather reflect the coordination of interneuronal communication within these networks (Bower 1995) A comparison of the EEGs generated by the model and real data is shown in Fig. 9.5.

9.5 Using the Tutorial

As should be clear from the introduction to this chapter, the piriform cortex model that is at the base of this tutorial is quite complex. In fact, this is the most complex tutorial in this book. We have designed an interface for the tutorial that is relatively intuitive to use, however, as you will see, one can quickly become inundated with graphical displays and flashing colors. Any effort to understand all the details of this simulation is certain to be a substantial undertaking. Accordingly, the rest of this chapter should be regarded only as an introduction to some of the features of the tutorial. We encourage you to make further explorations on your own.

9.5.1 Getting Started

If you type "`genesis Piriform`" after changing to the *Scripts/piriform* directory, you will execute the script for the *Piriform* tutorial. Because of the complexity of this model, it will take some time to initialize the simulation, especially on a slow machine. You will notice a multitude of messages appearing in the terminal window from which you started the simulation. You can ignore these for now. They simply provide information about which scripts have been executed. These messages can act as a powerful debugging tool if you have any intention of later modifying the scripts. Once these messages have ended, a large colorful window should appear in the upper left corner of your screen. This is the menu window that also serves as a circuit diagram for the modeled network. In order to decipher the various colors and shapes you see in the window, click on the `help` button at the bottom

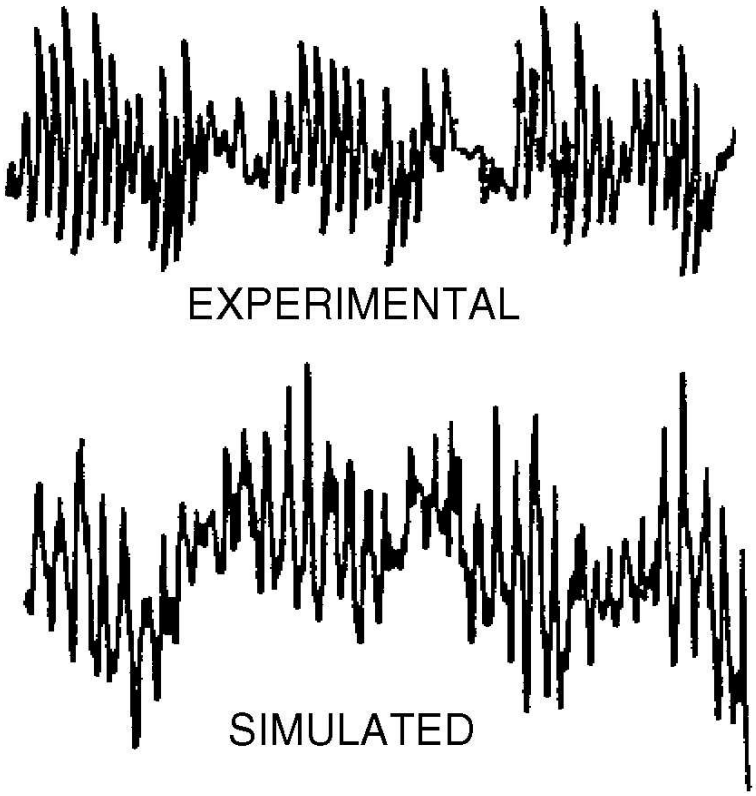

Figure 9.5 Comparison of an EEG measured from rat piriform cortex and data generated using this simulation (Wilson 1990).

of the window. A large text-filled window should appear on the right half of your screen. This window contains an on-line help file which you can scroll up and down using the scroll bar on the left side of the window. Click on the button at the bottom of the `help` window labeled `legend`. This should generate a window that explains the meaning of the colors and shapes you see in the menu. To hide any of these windows, simply click on the buttons labeled `cancel`.

9.5.2 Generating Simulated Data

In order to generate data using this tutorial, you must first go to a directory to which you have write access. To do this, type "`cd <directory>`" from the window that is running GENESIS, where "`<directory>`" is some directory to which you are allowed to write. Then copy the file called *test_dir* from the tutorial directory to the directory to which you

wish to write your data. The *test_dir* file contains a UNIX shell script that tests to see if the data directory you are trying to create already exists.

The tutorial operates in two modes: PLAYBACK and SIMULATE. In order to generate data, you must be in SIMULATE mode. On the lower left corner of the menu window is the mode toggle labeled either SIMULATE or PLAYBACK. If it reads PLAYBACK, click on it so that it reads SIMULATE. Now click on the simulate button next to the toggle. A window also labeled simulate should appear over the button. The dialog labeled step size allows the user to set the time step of the simulation. As discussed earlier in this text (Chapter 2), there is always a tradeoff in setting step size. A large step size will increase the speed of the simulation but decrease the accuracy of the results, whereas a small step will improve accuracy at the cost of speed. The optimal step size will depend on the speed of your machine, your desire for accuracy, and most importantly, your patience. For now, let's leave step size at 0.1 *ms*. This means that the simulation will generate ten data points for every one millisecond of simulated time. A millisecond of simulated time corresponds to a millisecond of time in the biological network.

The dialog labeled simulation duration indicates the length of the simulation run and data directory name is the name of the directory to which you will store your simulation data. For now, let the dialog read default. The dialog output rate contains the sampling frequency for stored data. If step size is 0.1 *ms* and output rate is 1.0 *ms*, this means that for every ten simulated steps, only one value will be stored. This parameter is just as important as step size because once you are finished generating data, you will only be able to look at your data with the temporal resolution indicated in this dialog.

Now it is time to run the simulation. In the simulate window, click on the button labeled reset. This must be done each time you run a new simulation. Now click on start simulation. Wait about ten seconds and then click on simulation status. A window should appear over the menu that will tell you what percent of the simulation run is complete. This will give you a good idea of how long it will take to generate your data. The simulation status is not automatically updated. It will only update whenever you click on the simulation status button in the simulate window or the update button in the sim_stat window. By clicking on the stop button, you can freeze the simulation run. By clicking on continue simulation, data generation will resume from the point at which it was halted. Now wait until the simulation has finished. The run is complete when the start simulation or continue simulation button is no longer highlighted and when the simulation status window reads 100%.

9.5.3 Initial Look at Simulated Activity

As you will see in later sections, the graphical interface for this tutorial allows you to observe many different aspects of the model's behavior. Before introducing you to the intricacies of the interface, it will be useful to give you a sense for the overall network

behavior. To do this, you should find the three buttons to the right of the pyramidal cell in the menu window labeled `field potential`, `spike count`, and `full cell view`. Each of the graphs associated with these buttons allows you to look at a different measure of global network activity. Click on the button titled `field potential`. Beneath the title bar labeled `fp's generated by synaptic currents` are toggles corresponding to the six termini of the six different synaptic pathways going to the pyramidal cells, and below that a title bar reading `eeg display`. For the moment, just click on the button labeled `eeg`. When you play back the simulation, this graph will show you the simulated EEG.

A second form of network behavior you can monitor is the total spike count over time for any neuronal class in the model. Click on the `cancel` button at the bottom of the `field potentials` window and then click on `spike count`. A menu window should appear over the button listing three options: `pyramidal`, `feedback inhibitory`, and `feedforward inhibitory`. Click on `pyramidal`. A graph titled `spike_activity_pyr_default` should appear next to the menu window. `Spike count` graphs show the total number of action potentials occurring during a playback step size for all the members of a particular cell class. In this case, we are looking at spike count for the pyramidal cells, but we also could have chosen to count spikes for feedforward and feedback inhibitory cells.

In addition to overall patterns of network activity, the interface allows you to examine the response properties of many individual network components. The details of this capacity are described in later sections of this chapter. For now, look at the pyramidal cell in the menu window and notice that at the center of every cell compartment there is a button indicating laminar position. These buttons allow you to visualize data associated with each compartment. Click on the button labeled `supIb`. The dialog box that appears has two options `network view` and `total synaptic conductance`. Click on `network view`. The new window includes several display options. Click on `synaptic conductance (Na)`. A window containing a two-dimensional array of rectangles should appear in the upper right corner of the screen. It should be labeled `net_pyr_gCA.default`. This window shows the synaptic Na^+ conductance for each superficial Ib layer compartment of each pyramidal cell.

We are finally ready to play back the data we have generated. In order to look at the simulation results, you must put the tutorial in `PLAYBACK` mode. Go to the toggle in the lower left part of the menu window and click on it if it reads `SIMULATE`. It should now read `PLAYBACK`, indicating that you are in `PLAYBACK` mode. Go to the `playback` button at the bottom of the menu window and click on it. The `playback` window that appears contains many of the same options you saw in the simulate window and should therefore be self-explanatory. The only thing to note is the clock dialog at the top of the window which indicates the number of milliseconds of simulated activity that have been played back. Click on the `reset` button and then on the `step` button. The `step` button advances the playback by one time step as specified in the `playback step size` dialog. In order to play everything from start to finish, click on `run`. Do not hide any of the display windows because you will need them for the next section.

9.5.4 Observing Network Behavior

If you followed all the instructions up to this point, you should be observing three different types of simulated results simultaneously. You will first notice the complex oscillatory pattern of the EEG (note the roughly 40 *Hz* oscillations), and the oscillatory spiking activity of the network. If you now pay close attention to the network window displaying synaptic conductance (labeled `net_pyr_gCA.default`), you will notice waves of synaptic activity propagating from the rostral to the caudal end of the network. This is what we would expect given that the superficial Ib compartment of the pyramidal cell receives synaptic input from rostral pyramidal cells.

In the real brain, the olfactory cortex is known to receive oscillatory inputs from the olfactory bulb (Bressler 1987). Furthermore, these bulbar oscillations occur in the same theta and gamma ranges seen in the olfactory cortex EEG. In this tutorial, input coming from the olfactory bulb is modeled as continuous and random, yet the piriform model still generates the oscillatory patterns found in the piriform cortex when recording from an awake behaving animal. This suggests that the intrinsic properties of the piriform cortex, rather than patterned input, are responsible for these oscillations in the real cortex (Wilson and Bower 1992).

9.5.5 Varying Network Parameters

Over the last several years we have published several papers on what we consider to be the physiological and computational implications of these simulations of the oscillatory structure of the piriform cortex (Wilson and Bower 1992, Bower 1995). We have also modified this simulation to make the model more like neocortex (Wilson and Bower 1991) and used it to replicate the recent oscillation results in visual cortex (Gray et al. 1989). Readers interested in our interpretations are encouraged to read these papers. Here, we briefly introduce you to the way in which you can modify model parameters to further explore the physiological basis of these oscillations.

When one wishes to study the behavior of a model, it is important to be able to vary its parameters. As you have learned, GENESIS was designed to make that relatively easy. Of course, one must still carefully choose which parameters to vary because a biologically realistic simulation as complex as this one is dependent on a multitude of parameters that would take many years to fully explore at random. Typically, in complex simulations of this sort, which parameters are varied is determined either by the level of experimental certainty of the parameter value, or by parameters that are known to vary in the biological system. For example, if a parameter is well defined experimentally, and well known to be invariant, it is less necessary to run a full parameter search (Bhalla and Bower 1993). In the current simulation, for example, the values for conduction velocities in piriform cortex were previously experimentally determined, and most likely remain constant in the context of normal

functioning; therefore, it may be less important to vary conduction velocities except to illuminate their role in normal activity. On the other hand, it is well known that synaptic efficacy varies greatly with the presence of certain neuromodulators. Experiments in piriform cortex have shown that synaptic transmission in the intrinsic excitatory pathways (layer Ib) is significantly diminished following the application of acetylcholine or norepinephrine (Hasselmo and Bower 1992, Vanier and Bower 1993). In addition, carbachol (a drug which mimics the effects of acetylcholine) has been shown to suppress gamma oscillations and intensify the theta rhythm in the piriform cortex of cat (Biedenbach 1966). In the following exercise, we will decrease the synaptic efficacy of excitatory inputs going to layer Ib and study the effects of this change on the EEG.

In this tutorial, the user is allowed to change the `weights` (or efficacy) of synaptic pathways in the model as well as the duration of `channel open times` (this can also be thought of as duration of synaptic current) in the pyramidal cell. To see how this is done, click on one of the synapses on the pyramidal cell. A window should appear over the synapse containing dialogs specifying `weight` and `channel open time`. Synapses on inhibitory cells can have their `weight`, but not their `channel open time` altered.

To examine the effects of changing the synaptic efficacy of the excitatory association pathways in layer Ib, first make sure the tutorial is in `SIMULATE` mode. You will notice that the previous `PLAYBACK` windows will disappear once the simulation is in `SIMULATE` mode. Don't be alarmed by this. It's supposed to happen. Now go to the excitatory synapse on layer `supIb` and click on it. A window titled `rostral pyr to caudal pyr synapse` should appear over the synapse. The weight value is `1.0` in the dialog labeled `from rostral pyr`. This indicates that the strength of this pathway is 1.0 times the default value. Default values for synaptic strengths are determined by finding synaptic strengths that permit the simulation to replicate physiological activity. Change the weight dialog to `0.05` in order to make the synaptic efficacy 20 times weaker than the default value. Click on the synapse going to layer deepIb. When the window titled `caudal pyr to rostral pyr synapse` appears, go to the weight dialog and enter `0.05`. Now click on the button labeled `apply`. Every window associated with a synapse has an `apply` button that incorporates the changes the user has made. This button has to be clicked in only a single window after all the appropriate changes have been made. You can also click on it every time you make changes in a synapse window, but clicking it only once will save you time.

Open the `simulate` window by clicking on the `simulate` button on the control bar at the bottom of the menu window. Go to the `data directory name` dialog and enter `ach` (for acetylcholine). Remember that in the previous simulation run, the name of the data directory was `default`. Click on `reset` and then `start simulation`.

When the simulation is complete, set the mode toggle to `PLAYBACK`. In the `playback` window, make sure that the `data directory` is set to `default`. Click on the `field potential` button, then click on `eeg`. Go to the `spike count` button and click on it. Then select `pyramidal`. Go back to the `supIb` button and click on it. In the windows that pop

up, choose `network view` and then `synaptic conductance (Na)`. These are the same displays you were looking at previously. Now click on the `playback` button at the bottom of the menu window and type `ach` in the `data directory` dialog. Go back to the `field potentials` window and click on `eeg` again. Then make the same `spike count` and `supIb` selections that you made previously for the default data. After you have done this, you will notice that you now have a duplicate set of windows. Note that the last word in each window title indicates the data directory name. This is the case with all display windows used in `PLAYBACK` mode. The tutorial allows you to display any number of windows from any number of directories simultaneously. Go back to the `playback` window and click on `reset` and then `run`. All the windows should start running simultaneously. Compare the way in which the simulation evolves in the default and acetylcholine case. Is the difference subtle or dramatic?

Now let's look at the two EEG graphs to determine how well the simulation was able to replicate the experimental results described earlier. The one corresponding to the `default` directory shows the EEG generated when weights were left at their default values. The `ach` EEG display is different in a number of ways. First, the amplitude of the gamma frequency oscillations is significantly attenuated. We also see high frequency activity riding on top of the remaining gamma oscillations. Although these EEG results differ significantly from what is seen in experimental preparations, the attenuation of the gamma rhythm is qualitatively similar. We may ask ourselves why the model was not able to replicate the full extent of the experimental results. Perhaps we neglected to incorporate an important parameter. For example, it is well known that acetylcholine and norepinephrine increase the excitability of pyramidal cells by blocking a slow K^+ current (Hasselmo and Bower 1992, Madison and Nicoll 1982). Yet, neither the current nor the increase in excitability is modeled. The ambitious student may try to better match model data to experimental EEGs by incorporating these details into the model using knowledge gleaned in later chapters.

9.6 Detailed Examination of Network Behavior

So far, you have observed network behavior using only three types of data windows. As already mentioned, the interface to this tutorial actually allows the user to observe many different aspects of network response. Because the possibilities are far too numerous to describe in detail, the following sections describe some of your options and some suggested exercises. In going through the sections below, remember to specify your playback data directory as `default`.

If you return to the `field potential` button and click, beneath the title bar labeled `fp's generated by synaptic currents` you will find toggles corresponding to the six termini of the six different synaptic pathways going to the pyramidal cells. These toggles allow you to look at the contributions of individual synaptic pathways to the overall EEG.

By clicking on several of these at once, you can add the field potentials generated by the synaptic currents in different pathways. Thus, if you click on the toggles labeled `supIb` and `deepIb` and then click on `display` a graph will appear on the screen whose title should indicate the combination of toggles you just clicked. If you play back the simulation, you will see the summed contribution of the synaptic currents in these two layers to the overall EEG. Note that if you sum the contributions from all synaptic pathways, you will not get a graph that is identical to the EEG. This is because you are neglecting the contribution to the EEG from voltage-gated currents in the soma.

At the total network level you can observe the activities of rows of neurons all at once by selecting the button labeled `full cell view` to the right of the pyramidal cell. A menu window containing two dialogs should appear over the button. The `full cell view` option allows you to display a row of cells in the network. The `parameter` dialog gives you the choice of viewing membrane potential (`Vm`), transmembrane current (`Im`), or synaptic conductance (`Gk`). The `row` dialog specifies the row number to be displayed. Valid row numbers are 1 through 9. Enter Gk in the `parameter` dialog and click on `display` to see row 5 and to look at synaptic conductance. Each column you see in the display window represents a single pyramidal cell in the network. Each cell in the display is composed of six visual compartments that represent conductance changes caused by each of the six synaptic pathways terminating on the pyramidal cell. `Vm` and `Im` `full cell view` displays show only five visual compartments for each cell which correspond to the five electrical compartments used to model the pyramidal cell.

In addition to overall patterns of network activity, the interface allows you to examine the response properties of many individual network components. For example, if you look at the pyramidal cell in the menu window, you will notice that at the center of every cell compartment there is a button indicating laminar position. These buttons allow you to visualize network data associated with that compartment only. For example, press the button labeled `deepIb`. A menu titled `deepIb_options` should appear over the button you just pressed. Click on `network view`. Another menu titled `deepIb_network` should appear next to the original `deepIb_options` window. The `network` option allows you to display a slice of piriform cortex that is parallel to the surface, but contains only the specified compartments from all the pyramidal cells in the network. The options `Vm, Im,` and `synaptic conductance` indicate the parameter to be visualized. To get a better idea of what this means, click on `Vm, Im,` and `synaptic conductance`. Three small windows each containing a 9×15 grid of rectangles should appear on the right side of the screen. Each rectangle represents a deepIb compartment from a single simulated pyramidal cell found in that position within the network. The value of parameters `Vm` (membrane potential), `Im` (transmembrane current), and `synaptic conductance` are indicated by the color of the rectangles. Hot colors like red and yellow represent high values and cold colors like blue represent low values. The parameter `synaptic conductance` requires some explanation since it is not something neurophysiologists are accustomed to observing di-

rectly. Each synaptic pathway induces a change in conductance in the compartment (or layer) on which it is acting. This change in conductance is measured as the `synaptic conductance`. Although this is not something a physiologist can yet measure, it is a useful way of keeping track of the total amount of synaptic activity. This in turn can help one to understand network dynamics. Click on `cancel` to get rid of the `deepIb_network` window and then click on the button labeled `total synaptic conductance (Na)` in the `deepIb_options` window. A graph should appear in the upper part of your screen next to the menu window. When you play back your data, this graph will show the summed synaptic conductance over the entire network for the deepIb compartment.

Other compartments contain the same visualization options and can be executed simultaneously with any combination of other compartment displays. One exercise you can perform is to determine which parameters are clearly oscillatory and which are not. For example, there is clear oscillatory behavior in the full cell view window displaying synaptic conductance (labeled `full_cell_Gk_5.default`). You should notice that the fourth row of compartments down displays a similar undulating pattern to that seen in the network display. This was to be expected since this row of compartments represents the synaptic conductance of deepIb compartments in the pyramidal cell, which is also oscillatory.

You can also use multiple output graphs to observe in detail the relationships among different components of the model. For example, you should compare the graphs for pyramidal cell spike count, EEG, supIb and deepIb field potential, and total synaptic conductance for the deepIb compartment (press the button labeled `total synaptic conductance (Na)` in the `deepIb_options` window). You will notice that all graphs display oscillatory waveforms of varying amplitudes. You might want to rescale some of the graphs by magnifying the waveforms. To do this, click on the `scale` button in any of the graph windows. Play with `ymin` and `ymax` until you have properly magnified the waveform. Do this for all the graphs making sure that `xmax` is `500` in each case. Look closely again and you will notice that there is a one-to-one correspondence between the peaks of all the graphs. This tells us a number of very important things. First it allows us to correlate the EEG to pyramidal cell activity. By comparing the graphs labeled `spike_activity_pyr_default` and `field_eeg_default` we can say that a peak in the EEG corresponds to a peak in pyramidal cell activity. To see this better, scale the EEG and spike count windows so that `xmax` is `100`. Now scale the other graphs to an `xmax` of `100` and compare the EEG to the added supIb and deepIb field potentials. The peaks and valleys of the waveforms correspond almost exactly; furthermore, both waveforms have similar amplitudes. This suggests that the field potentials generated by association pathway synaptic currents underlie a significant portion of the EEG. Now compare the conductance graph to the EEG. The valleys in the EEG correspond to the peaks in the synaptic conductance. The reason for this is that the inward currents generated by deepIb synaptic conductance create a negative field potential (see Eq. 9.1). Our last comparison is between the pyramidal cell spike count and synaptic conductance. You will notice that the synaptic conductance trace is positively phase shifted with respect

to the pyramidal cell spike count. This is because of the delay between the firing of caudal pyramidal cells and the arrival of synaptic input to layer deepIb.

9.7 Summary

The piriform model has been used to study a wide range of physiological and functional phenomena including associative memory function (Wilson and Bower 1988, Wilson and Bower 1992). The model attempts to retain biological plausibility, but simplifies the structure of the piriform cortex considerably. Despite the existence of a number of cell classes in the piriform cortex, only three were modeled: superficial pyramidal cells, feedforward and feedback inhibitory cells. Pyramidal neurons are represented with five electrical compartments and inhibitory cells with only one. Although the piriform area contains on the order of millions of cells, the model (in its tutorial manifestation) uses only 405. Axon conduction velocities and anatomical circuitry were fitted to experimental data in order to place realistic constraints on the network. Despite its simplicity, the model has been able to replicate a wide range of cortical behavior and to make experimentally testable predictions. The reader is encouraged to use information in the second part of this book to modify this network model as he or she sees fit in order to further explore cerebral cortical dynamics and function.

9.8 Exercises

1. Change the synaptic weight for the feedback interneuron to pyramidal cell pathway from `1.0` to `0.0` and then run the simulation. Change the pathway back to its default weight and then change the weight for the pyramidal cell to feedback inhibitory interneuron from `1.0` to `0.0`. Compare the data you obtain from each change to the default data you generated earlier. What happens to the EEG in each case? What role do the feedback inhibitory interneurons seem to play in generating 40 Hz oscillations?

2. The $GABA_A$ receptor activates the fast ligand-gated Cl^- permeable synaptic channel responsible for the shunting inhibition generated by the feedback inhibitory cells in the layer II area of the pyramidal cell. Barbiturates are known to act at a special site on the $GABA_A$ receptor that prolongs the burst time of the Cl^- channel (Hille 1992). Simulate the effect of applied barbiturates to the piriform cortex by changing the channel open time for the inhibitory synaptic input to layer II from `7.0` *ms* to `18.0` *ms*. What effect does this have on the EEG? What do you think is the physiological basis for this effect?

3. The olfactory bulb is known to generate both fast and slow EEG rhythms that match those found in the piriform cortex when recordings are done simultaneously (Bressler 1987). In the piriform model, bulb input is modeled by a random number generator with a flat frequency distribution; nonetheless, fast oscillations still occur at roughly 40 *Hz*. Change the synaptic weight of the afferent pathway (coming from the bulb via the LOT and going to the pyramidal cell) to `0.1`. Then generate a new batch of data where the synaptic weight of the afferent pathway is set to `5.0`. Compare the EEGs you just generated to the default case. Is there any significant difference? Do you think piriform cortex requires patterned stimulation from other brain areas in order to generate 40 *Hz* oscillations? If not, what purpose might patterned input serve?