Computational Neuroscience of Neuronal Networks

76

Samuel A. Neymotin, AnnMary Mathew, Cliff C. Kerr, and William W. Lytton

Brief History

History: Missing Links

Through multiscale modeling, computer simulation typically utilizes multiple simplifications, each providing a different view. The various models then provide complementary insights into a system that is too large, and too complex, to grasp directly. Each simplification will match, albeit always imperfectly, some important aspect of the system under study. In the story of the blind men and the elephant, each blind man takes hold of a different part of the elephant's anatomy. Each, from his own limited observation, declares the elephant to be something quite different: long and tufted (tail), long and muscular (trunk), flat and broad (ear), thick like a tree (leg), etc. Each of these models of the elephant is accurate but limited. Taking these models together, adding in additional information as to the location of each blind man, one could begin to build a preliminary idea of the elephant as a whole. While multiscale modeling emphasizes the use of different models at different scales, multiple views are also often useful at a single scale. The multiple scales of multiscale may be spatial (microns to millimeters to centimeters to meters) or temporal (milliseconds to seconds to minutes up to years).

Whether multiscale or not, a major challenge in modeling is the same as it is in the elephant story: make the connections between models so as to build a coherent

S.A. Neymotin (✉)
Biomedical Engineering Program, SUNY Downstate / NYU-Poly, Brooklyn, NY, USA
e-mail: samn@neurosim.downstate.edu

A. Mathew • C.C. Kerr • W.W. Lytton
Physiology and Pharmacology, State University of New York Downstate Medical Center, Brooklyn, NY, USA
e-mail: annmary1985@gmail.com, cliffk@neurosim.downstate.edu, billl@neurosim.downstate.edu

D.W. Pfaff (ed.), *Neuroscience in the 21st Century*,
DOI 10.1007/978-1-4614-1997-6_87,

story. In the case of multiscale modeling, the connections are typically effected via an embedding: models at each scale include highly simplified versions of one or more models at the scale below. Unlike the blind men, we know in advance that we want to connect models, so we try to do our modeling with a view upward toward future simplifications as well as a view downward at the models to be included.

Neuronal modeling, dating back to the 1940s, is an early example of multiscale modeling. Here the scales to be connected range from the cellular level covered in the last chapter to the network level covered in this chapter. The major problem in neural network modeling is that scientists still do not know enough about how individual neurons work. Therefore, modelers must make various simplifying assumptions without being confident that these are the correct assumptions. For example, the prior chapter ended with a brief discussion of activity spread in the apical dendrite of a pyramidal cell. These large cells are a major component of the neocortex and hippocampus and are highly conserved through mammalian evolution, with similar sizes seen from shrews to elephants. It is likely that there is some special processing occurring in these cells, but this remains unknown. As a result, this special processing is at present not included in network models.

Other key attributes are also missing from the knowledge and understanding of neuronal networks. One of the most important of these is chemistry. Most of neural modeling, as highlighted in the prior chapter, has focused on the electrical properties of cells. The complex chemical machinery of neurons has not been well explored. In fact, the information processing aspects of the chemical machinery have not yet been disentangled from the metabolic and reproductive chemical machinery that is part of almost every cell in the body. This separation is made more difficult by the fact that the brain has undoubtedly co-opted metabolic machinery to take on additional functions during the course of evolution. For example, the major neurotransmitters discussed below are close analogues of foods: you think with what you eat.

This chapter is divided into two parts. First, it is noted that simplifications are both required and desirable in order to model the brain. This section is partly a methods section, covering some of the large number of decisions that have to be made when choosing how and even what to model. The simplification section has itself been greatly simplified by focusing on simulation of spiking networks and omitting most of the alternatives: coupled oscillators, artificial neural networks, etc. The second part of the chapter describes a set of themes that recur in much of neuronal network modeling. Again, this section necessarily omits a number of these themes, both for reasons of space and due to the fact that some of these modeling themes are best characterized utilizing modeling techniques that have not been covered. As noted in the prior chapter, many of these themes and techniques are covered in three textbooks: those by Hertz et al., by Dayan and Abbott, and by Lytton. A theme of these themes, a meta-theme, is the introduction of computer simulation as an experimental pursuit. The complexity of the systems being investigated is such that the entire system cannot be grasped in its totality. This leads to the notion of in silico experimentation as an analogue of the more familiar in vivo and in vitro techniques.

Simplifications

A central theme in all modeling is simplification. It is often said that model simplification is an end in itself, since reduction is a necessary component of science. While this was largely true in the era of paper-and-pen modeling, the computer era has opened up new opportunities and new styles of modeling beyond reductionism. Much modern computer modeling is done to simply provide a more readily accessible version of the thing itself. The flight simulator for a 747, based on a computer model, is judged based on its ability to accurately include everything. It is used to test pathologies that one hopes will never occur in the real thing. Similarly, a complete computer model of the brain, though unachievable, would be of great use.

Reductionism is sometimes derided as the last refuge of the scientific mind. Reductionism, however, comes in many forms. Reductionism to a single cause, the smoking gun, is often impossible in complex systems. However, conceptual reductionism is another word for explanation, which is a necessary end of science. A computer model may embody a conceptual reduction by doing a direct simplification. Alternatively, as in the case of the 747 model, a simulation may simply be used as a tool to enable concepts to emerge. Often, these concepts will emerge via other models which do provide the simplification. One facet of multiscale modeling is the building of computer models of computer models of computer models.

How and where are simplifications made? A major reason for simplifications is that we do not know what is there. We often must exclude things that we believe to be important. Leaving out the things that we know we do not know is vexing; leaving out the things that we do not know we do not know is even more problematic. For example, there may be some remarkable phenomenology hidden in the use of gaseous communication (ethylene in bananas; nitric oxide in human brains). Another reason for simplification is that we do not have the computer power, or computer techniques, necessary to simulate all the things that we do know. This has become increasingly obvious with the rise of the "omes": massive amounts of information are coming out of genetic and protein assays. These data cannot yet be accommodated in our models.

Simplification #1: Dichotomies

There are two kinds of people in the world: those who dichotomize and those who do not. Modelers generally fall into the former group. There are thousands of different kinds of neurons in the brain. These are definable into different subsets by using different criteria: cell morphology, neurotransmitters, proteins produced, projection patterns, electrical properties, dynamic properties, etc. For most of these properties, we attempt to simplify through dichotomies, lumping things together into only two groups instead of 10s or 100s or 1,000s.

For most models, we recognize two ways in which a cell will influence other cells: excitation or inhibition. In this case, the cell type is defined based on the transmitter that is being produced. A key concept here is Dale's law: a given neuron

will produce glutamate (the major excitatory transmitter) or GABA (a major inhibitory transmitter) but not both. A model cell is therefore strictly defined as being either an excitatory cell or an inhibitory cell, and synapses are defined as strictly excitatory or inhibitory. This simplification ignores the fact that transmitters are often coexpressed and coreleased with other transmitters, for example, peptides with GABA. Excitatory (inhibitory) cells have axons that end with excitatory (inhibitory) synapses. Excitatory (inhibitory) cells produce excitatory (inhibitory) postsynaptic potentials (EPSPs/IPSPs). The reversal potential of the ion channels associated with excitation is above the resting potential of the cell so that activation produces depolarizing potentials. This brings the postsynaptic cell closer to firing its own action potentials. By contrast, when an inhibitory neuron fires an action potential, the voltage of a postsynaptic target is generally moved in a hyperpolarizing direction, below resting potential. As a result, the postsynaptic cell is typically less likely to fire an action potential. Inhibitory effects can also be achieved without hyperpolarization by simply increasing conductance and thereby making it harder to subsequently depolarize the cell.

For both EPSP and IPSP, we typically model a further dichotomy among the postsynaptic receptors. For glutamate (excitatory), we model fast AMPA responses and slow NMDA responses. For GABA (inhibitory), we typically model $GABA_A$ (fast) and $GABA_B$ (slow) receptors. The different receptors are also associated with different reversal potentials based on the mixture of ions that flow through their associated ion channels. We generally ignore metabotropic receptors and their chemical responses which produce communication with internal transmitters and with the cell nucleus. It is likely that some receptors encompass both metabotropic and ionotropic (current passing) effects.

An anatomical projection dichotomy is the distinction made between interneurons and principal cells. Interneurons are those which only project locally in their own immediate area. Principal neurons are generally larger and have outside as well as local projections. In the neocortex, principal cells are excitatory and often pyramidal, in the shape of their cell bodies. Hence, the terms excitatory, pyramidal, and principal are sometimes used as shorthand synonyms. Cortical interneurons are largely inhibitory, and the use of interneuron as a synonym for inhibitory cell is very common. This synonym should be used with caution, since in other brain areas (e.g., cerebellum and striatum), the principal cells are inhibitory and local organization via interneurons is mixed. Another set of anatomical distinctions involves the location of projections onto the postsynaptic cells: distal vs. proximal. For neocortical pyramidal cells, excitatory synapses tend to project onto more distal locations and onto spines (the "twigs" of the dendritic "tree") at these locations. By contrast, inhibitory projections arrive at proximal dendrites, at the soma itself, or at the beginning of the axon.

A dynamical dichotomy is made between tonic and bursting cells. Tonic means that the activation of the neuron produces spiking with fairly regular interspike intervals. Generally, the firing will continue for the duration of an input. This is often tested by artificially applying a current clamp (constant current) through an electrode. Burst firing is the tendency to produce clumps of spikes together,

followed by a longer interburst interval. Bursting can be either repetitive or single. It can also be intrinsic or extrinsic. Intrinsic bursting can be elicited by a prolonged input. Extrinsic bursting is elicited by a brief input that defines the burst duration. Note that extrinsic bursting still requires that the cell have particular intrinsic properties that permit it to fire at a high frequency. Tonic responses are sometimes referred to as regular and burst responses as phasic.

The above dichotomies are simplifications on the scale of cells and synapses that we use in order to build networks. There are also a number of popular dichotomies that concern networks, which could in turn be used to build more complex systems. These will come up as we go through the chapter but are worth foreshadowing here. First, a network or subnetwork can be regarded primarily as an input (e.g., one that receives primary sensory information) or an output (e.g., one that projects to muscles). Another dichotomy that we will explore further is the notion that a system may be a generator of signals or a filter of signals. Another dichotomy will be the notion of single-cell coding vs. ensemble coding. As with all the prior dichotomies, these either/or propositions are oversimplifications that are nonetheless useful for improving our understanding.

Simplification #2: Biophysical Omission

In the prior section, we simplified by lumping or categorizing. There might be 100 kinds of neurons and 1,000 kinds of receptors, but we will say that there are 2 (or sometimes 3 or 4) of each. In this section, we will simplify by omission. In the next section, we will simplify by redefinition.

The main thing that we will leave out is the axon. The axon, ironically, was the first part of the neuron that was modeled, with the discovery of the Na^+ and K^+ channels in the squid axon by Hodgkin and Huxley. In neurons, the axon is a relatively thin, very long projection that arises from the soma and that synapses on the dendrites and soma of other neurons. Most cells appear to conform to Cajal's hypothesis of functional polarity: synaptic inputs propagate from dendrites toward the soma and are summed there. Signals then are communicated to the axon initial segment and thence out via the axon to other cells. Sufficient depolarization of the axon initial segment leads to the generation of an action potential, which travels down the length of the axon until it hits the end of the axon. At this point, its depolarization causes the opening of voltage-dependent Ca^{++} channels, allowing Ca^{++} to enter the axon. Next, a cascade of molecular events causes the release of a neurotransmitter into the synaptic cleft. The neurotransmitter then diffuses across the synapse and binds to the postsynaptic receptor. This receptor may in turn trigger another molecular cascade with many potential consequences. One possible consequence will be an EPSP or IPSP.

We could simulate the axon, but simulating a lot of axons would be very computationally expensive. To simplify by omission, the computational function of an axon can be thought of as a common outflow path with transmission of information to postsynaptic targets. Consequently, in many compartmental models of neurons, an explicit compartment for the axon is simply left out, and signal integration takes place in the somatic compartment instead of the axon initial

segment. With the absence of the axon, simulating electrical propagation down the axon is no longer possible. The molecular events leading to neurotransmitter release and the neurotransmitter diffusion and binding to postsynaptic receptors are also left out. Instead, a delay is simply added before the effects of the EPSP/IPSP are felt to represent the axonal transmission delay. This delay is typically on the order of several milliseconds (2–4 ms), but can vary by model type.

In addition, the synapse itself is half missing. We leave out the presynaptic mechanism and just trigger a postsynaptic response in response to the arrival of the presynaptic signal following the delay. This postsynaptic response may be represented by a standardized function, called the alpha function, whose shape approximately matches the shape of a postsynaptic potential. Alternatively, a slightly more complicated postsynaptic modeling choice is to use a biexponential conductance alteration. To this may be added a square wave which grossly represents the duration of action of the unmodeled neurotransmitter. Notice that in this sequence of models, each adds a little more detail.

We mentioned at the beginning that we are leaving out all of the internal chemical signaling in the cell. As a result, we also left out dendritic spines from our models. One hypothesis is that the spine provides a small volume where reactants can come together. It is not clear what, if any, electrical functionality spines would have, so they are left out. Leaving out these details, which are undoubtedly important in biology, allows the modeler to focus attention on the electrophysiological activity in the cells of the network.

Perhaps the major omission arises from the fact that the brain is immensely large in terms of numbers of cells. There is no reliable estimate of how many cells there are; 100 billion is a round figure that is sometimes used. One cubic millimeter of cortex may contain about 50,000 cells; 50,000 cells would be manageable in our simulations were it not for the large numbers of interconnections, perhaps as high as 10,000 per cell, or up to 500 million synapses (50,000 $\times$ 10,000). In general, for n cells where every cell connects to every other cell except itself, there will be $n^2 - n$ synapses, often an overwhelming number. In order to run simulations within the limited computer memory available, and to get results within a reasonable amount of time, we tend to design networks that are considerably smaller than the actual networks that we want to model.

Simplification #3: Dynamic Replacement

The multicompartment model discussed in the prior chapter is computationally intensive. A single neuron can have hundreds of compartments. It is possible to run networks with these models or with reduced models with only tens of compartments. However, for many large simulations, a far more drastic simplification is used, collapsing the entire cell into a small computational element whose computational load is no greater than that of a single compartment of the multicompartment model.

The basic version of this minimalist model is the integrate-and-fire model (I-F). These models do not simulate currents associated with ion channels explicitly. Instead, I-F models work by using a threshold for firing and summing synaptic

inputs until that threshold is reached. These models also collapse the neuron so that all synaptic inputs it receives are treated equally. Such a spatially collapsed neuron, whether using a Hodgkin-Huxley model or an I-F model, is referred to as a *point neuron*, all of the spatial complexity of the neuron has been collapsed down to a point. An I-F neuron therefore has a single state variable, v, which represents the voltage level of the cell (note that V is used for voltage in compartmental models; v is often used for the quasi-voltage in simplified models). When a synaptic input arrives, v is incremented (or decremented), depending on the synapse type. The voltage change may be effected by current injection or by conductance change. The voltage then decays exponentially to a resting level with a time constant t, an analogue of membrane time constant, t_{memb}. When v passes a threshold, the cell produces an action potential. To add some realism to the model, the cell is prevented from firing again for a few milliseconds, representing an absolute refractory period. As with compartmental models, numerical integration is required to simulate I-F models. However, since the equations involved are simple and not linked, I-F models are much computationally cheaper to run.

Innovations in neuronal models combine features of I-F with more complex models. For example, extended I-F models were developed independently by MacGregor and by Friesen. These models had two compartments that combined a few details of the parallel conductance model with a thresholding mechanism from the I-F model. The parallel conductance model is used for subthreshold signal summation. Once a threshold is reached, computational savings are generally achieved by stopping the numerical integration and simply issuing an action potential as a signal to other cells. The threshold may be an adaptive value driven by voltage or may be calculated from the partial integration of the active currents. The action potential may also produce a resetting of some of the other state variables.

Another variant of the I-F neuron was introduced by Izhikevich. This model uses two state variables and four parameters. The two state variables represent membrane voltage v and a combined ion-channel state variable u. The four parameters are (a) the time constant for return to RMP following an action potential, (b) the sensitivity of the model to subthreshold voltage variation, (c) the value of the resting membrane potential (RMP), and (d) the magnitude of a post-spike afterhyperpolarization. Different choices for the four parameters allow the model to replicate a wide variety of biological dynamics, including regular firing, burst firing, and subthreshold firing, as shown in Fig. 76.1. This is a simple dynamical system with only two state variables governed by the four aforementioned parameters. There are two reasons such a simple model can show so many different kinds of dynamics. First, it is *nonlinear*, meaning that small changes in one state variable can produce a large effect on the rate of change in the other state variable. Second, even a small number of parameters can give a large variety of dynamics if the model is sensitive to them. For example, assuming each parameter can be small, medium, or large in value, four parameters yield $3^4 = 81$ possible combinations. Although not all of these 81 combinations yield unique or useful dynamics in this model, about 20 do.

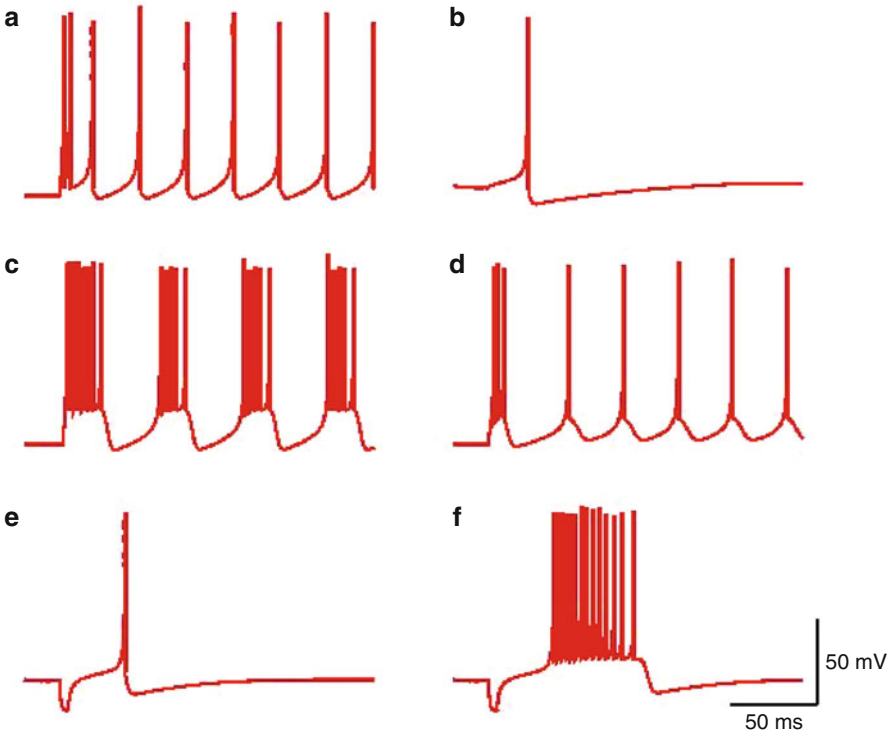

Fig. 76.1 Single-cell voltage traces from the Izhikevich (2003) integrate-and-fire model showing a subset of the possible dynamics. (**a**) Burst followed by regular firing. (**b**) Phasic single spike. (**c**) Repetitive bursting. (**d**) Mixed-mode bursting/spiking. (**e**) Rebound spike following hyperpolarization. (**f**) Rebound burst following hyperpolarization

Another class of I-F neurons was developed by Hines and Carnevale. This model eliminated numerical integration entirely by predicting the timing of subsequent firing based on the time at which inputs are received. There are computational cost savings resulting from rule evaluation only occurring at input event times, although these savings may not be great if the frequency of input signals is high. These models can also be extended by including a set of rules to simulate more complex activity, such as bursting and depolarization blockade, features not usually incorporated into simplified models. Extended versions of these models may also have several state variables comparable to those of an extended I-F model. In our version, the update rule for the models is a fixed-step augmenting (decrementing) of a particular synaptic state variable (e.g., AMPA, NMDA, $GABA_A$) in response to an excitatory (inhibitory) input. The changed state variable then decays back to baseline exponentially with a specified time constant. Once the state variables from the synapses are set, their values are added to produce the overall membrane voltage V_m.

The computational time savings associated with these various simplified cell models allow scaling up network size significantly, which in turn allows the investigation of network dynamics associated with large networks.

Simplification #4: Plasticity

To store memories and respond optimally to ever-changing environmental demands, changes take place in the network. Such changes are referred to as plasticity. Plasticity is responsible for a variety of adaptations to the environment. Of these, learning and memory have received the greatest interest. It has been demonstrated that there are many different kinds of memory: episodic, procedural, semantic, and others. These different forms of memory are located in different parts of the brain and presumably derive from different forms of learning, involving various patterns of plasticity.

Although plasticity occurs at multiple locations in neurons, plasticity at the synapse has received the greatest attention. The leading hypothesis is that many forms of learning are accomplished via alterations in the strength of synapses between neurons. In 1949, Donald Hebb hypothesized that neurons in an ensemble that are functionally coupled form connections via a growth process. Although a simplification, this idea has been popularized as "cells that fire together, wire together." Since that time, long-term potentiation (LTP) was demonstrated both in vitro and in vivo: a prolonged alteration in synaptic efficacy in response to activity across synapses. There is now additional evidence connecting LTP to learning in behaving animals. Meanwhile, many variants of Hebb's hypothesis have been applied to computational models. These algorithms are most clearly described in the context of artificial neural network models where they are used for various tasks including optical character recognition, speech processing, and stock market prediction.

In the context of detailed neuronal network simulations, a number of algorithms have been considered as possible sources of synaptic change. Much interest has focused on the phenomenon of spike-timing-dependent plasticity (STDP). This learning rule augments synaptic weights when a presynaptic neuron fires before its postsynaptic target and decrements when a postsynaptic neuron fires before its presynaptic source. In this way, presynaptic cells that are effective in causing their targets to fire are strengthened, while ineffective sources are weakened. The learning rules employed can take several different forms. In the simplest, the weight of the synapse is incremented using an exponential function: $A \exp(-d/t)$, where d is the difference in time between the pre- and postsynaptic spike, t is the time constant, and A is the scaling coefficient. STDP learning tends to enhance the synchrony of neurons in a network. Initially, a pre- and postsynaptic neuron may be weakly coupled. However, STDP will gradually strengthen the connection and hence the synchrony between the neurons. In the extreme, the network may be prone toward epileptic activity. To balance this, there are modifications to the basic rule that balance the synaptic strength in inverse proportion to the firing rate of the postsynaptic cell.

Simplification #5: Whence Activity?

Before any discussion of intrinsic neuronal network dynamics, it is important to understand how activity is generated in a network. Although some network models include spontaneously spiking cells, in most models, the network is silent (i.e., no cells fire action potentials) in the absence of external inputs. This is comparable to the situation in a brain slice which will typically be silent unless provided with shocks or pharmacological manipulation, because most of the synaptic inputs to the cell have been cut off.

There are two approaches to initiating network activity: (a) provide a brief shock stimulus and (b) provide ongoing random synaptic inputs. The first approach provides an excitatory stimulus over a short duration, typically 1–10 ms, to a proportion of the excitatory and/or inhibitory cells in the network, similar to what is done in vitro in brain slice experiments. Modeling this type of activity is useful for observing relatively brief responses from the network. In certain network models, it is possible to generate self-sustaining activity after a brief shock. This type of activity involves high cell firing rates with strong connectivity between excitatory cells or disinhibited models comparable to an "epileptic" slice after pharmacological manipulation.

These in vitro models do not replicate the relatively slow firing rates of cortical neurons in vivo, where many excitatory neurons fire at 1 Hz or less when not stimulated. Thus, the second approach is to provide ongoing synaptic inputs that are meant to represent ongoing activity being received by neurons from other brain areas that are not being simulated. These would be the other brain areas that are cut off when a brain is sliced up. This approach allows for more realistic in vivo modeling of neuronal networks. In these types of simulations, it is important to distinguish activity from external random inputs from activity generated by the network. At one extreme, an individual cell will simply fire at a brief delay from the times when stimulation is received from the random inputs. In this case, there is no value added by the network, the network is just following the random activity that sustains it. At the other extreme, there would be no relation between these inputs and cell firing. In order to assess this, we can compare activity patterns when cells in the network are connected to patterns of firing when cells in the network are disconnected.

One widely cited model by Izhikevich and Edelman used plasticity as a means of maintaining activity. In their simulations, network activity died out during the first second of activation, even if the network was initially strongly stimulated and had high strength of synaptic connections. Using STDP, they found that the network rewired in a way that would allow it to continue activity without the need for continuous exogenous stimulation.

Simplification #6: Network Replacement

A drastic simplification that can be made is to simulate an entire network as a set of equations, doing away with individual cells altogether. These models, called lumped, continuum, or mean-field models, were pioneered by Beurle in the 1950s and developed by Wilson, Cowan, and Nunez in the 1970s and 1980s. This type of

model generally depends on one of the dichotomies: that between excitatory and inhibitory neurons. Each of these populations is assumed to be homogeneous so that individual properties cancel out, leaving only bulk, average properties that affect the entire ensemble. In some ways, this is analogous to the relationship between statistical mechanics and thermodynamics in physics. In statistical mechanics (network models), the position and momentum (position and firing rate) of each particle (neuron) are nominally considered. However, at large scales, these properties blur into macroscopic quantities like temperature and pressure, which are the domain of thermodynamics.

Likewise, instead of considering the firing rate of each neuron independently, one can consider the average firing rate across a population or ensemble of neurons. Whereas a single neuron either fires an action potential in a given time interval or does not fire, in a lumped model, the population firing rate is a continuous variable. Similarly, instead of discrete connections between individual neurons, connections are between populations and are again represented by continuous variables. Note that some neuron-level properties are preserved even in lumped models: for example, the temporal integration of inputs in the dendrites results in low-pass filtering of input signals. Thus, even if the individual dendrites are not modeled, this low-pass filtering effect may still be taken into account.

With their low dimensionality, these systems are easier to manipulate than large neuronal networks and can be used to obtain very close approximations to physiological spectra (Fig. 76.2). Furthermore, surprisingly complex dynamics can arise from lumped models due to the complex patterns of feedback both within and between populations. For example, these models can give rise to both standing and traveling waves. A big advantage of lumped models is that they can be solved or partially solved analytically. This allows the investigation of mathematical properties that would be impossible to calculate for a network model. Lumped models typically include parameters that correspond to the strength of projections between populations or from one brain area to another. Therefore, they can be used to look at, or predict, how changes in the dynamics of brain areas might spread across the brain. Since lumped models are only valid approximations on "large" spatial scales ($>$1 mm), they are applied to large-scale dynamics such as EEG spectra (whole brain) or the sleep-wake cycle (whole day).

Themes and Measures

We now look at a number of themes that arise in the study of neuronal networks. These themes necessarily arise based on what we can look for, on what we know how to measure. We look where the light is brightest. The distance between what we want to know about and what we can actually measure is a recurring tension in the history of science. Before there is a way to measure something, there are many concepts or theories that turn out to be untenable. For example, phlogiston was long believed to be the essential element for combustion, until it was discovered that it could not be measured. It could not be measured because it did not, in fact, exist.

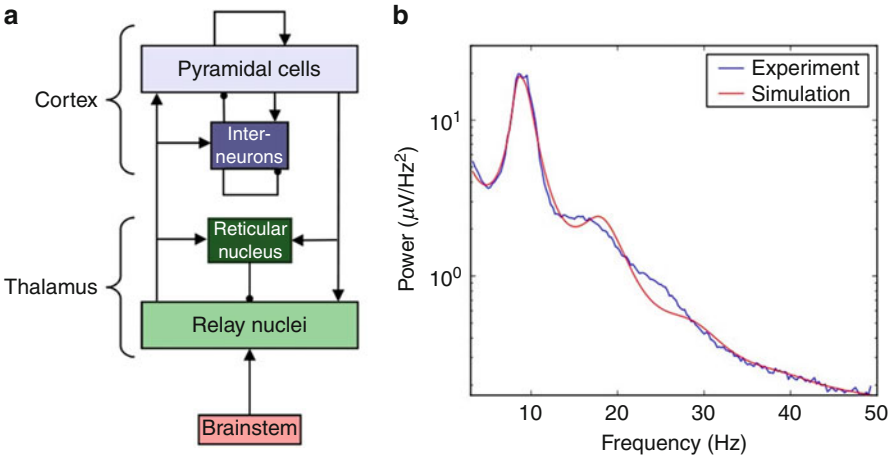

Fig. 76.2 Structure and output of a lumped model developed by Robinson et al. (**a**) Schematic diagram of the model. *Boxes* represent key neuronal populations in the thalamocortical circuit (*light blue*/*green*/*pink*, excitatory; *dark blue*/*green*, inhibitory). *Arrows* represent connections between populations (*triangle*, excitatory; *circle*, inhibitory). (**b**) Comparison of experimental eyes-closed EEG data from a single subject (*blue*) and the model (*red*). Note that a logarithmic y-axis is used to emphasize the 100-fold power difference between low (<10 Hz) and high (>40 Hz) frequencies (Experimental data courtesy of the Brain Resource International Database, www.brainnet.net)

Conversely, the ether, long believed to be a requirement for transmission of light and other electromagnetic waves, was measurable. Ether was sought, it was measured, but it was not there. In brain science, there is currently much interest in consciousness. There is also interest, and concern, about the possibility of building simulations, and machines, that achieve consciousness. However, consciousness remains underdefined and unmeasurable. The concept of consciousness will probably end up being split into a number of measurable subconcepts. For the time being, the origin of consciousness is left as an exercise for the reader.

Visualization

The brain is a complex system, perhaps the most complex that is known (some might argue that the full earth ecosystem is more complex). *Complex systems* is not only a descriptor, it has become a field of study. The proper approach to understanding complex systems remains a subject for debate, recapitulating the discussion of reductionism at the beginning of this chapter. Some insist that true understanding can only come from a reduction of complex models or aspects of complex models to simpler versions that can then be more fully analyzed, while others believe that one can learn much from direct analyses of the complex systems. Both perspectives have validity. The approach of reductionism through

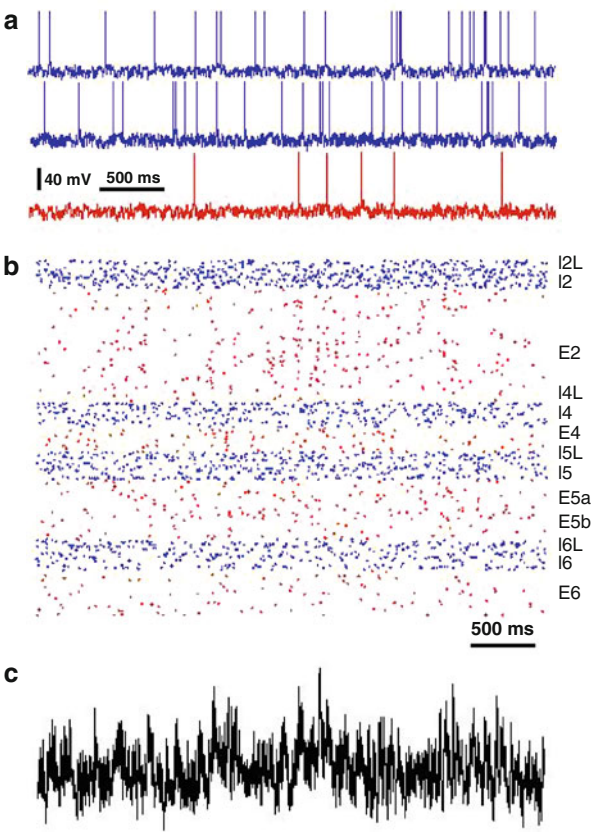

Fig. 76.3 Activity in a large multilayer cortical network model. (**a**) Representative spike trains from cells of different types (*blue* inhibitory; *red* excitatory). (**b**) Raster plot of spiking. Each cell is represented by a vertical position. Each *dot* represents a spike in a cell at that position, at the given time (x-axis). Excitatory cells are represented by *red* and inhibitory by *blue*. The cells are grossly arranged by cortical layer so that the cells toward the top would be more in the superficial cortex and those toward the *bottom* in the deep cortex. (**c**) Simulated local field potential (LFP) (Figure adapted from Neymotin et al. (2011a))

dimensional reduction (phase plane analysis) and other techniques generally requires sophisticated mathematical tools covered in the textbook by Izhikevich (see Further Reading). In this chapter, the focus is on visualizations and activity assessments that are done directly on large network models.

The main technique is to look. This seems absurdly obvious, but looking is overlooked surprisingly often. Investigators, caught up in the excitement of their latest clever statistical tools, neglect to look back at the raw data to confirm their observations. Experimental scientists are not immune to this syndrome, sometimes running Student's t-tests of questionable validity while neglecting to take a careful look at the data itself.

In a network of 1,000s, or even 100s, of cells, it is difficult to look at the dynamics of each cell individually. However, it is valuable to look at the voltage plots of at least some representative cells. This allows checking if the general pattern of firing is reasonable. For example, in Fig. 76.3a, the three cells show a reasonable pattern of firing with no evidence of depolarization blockade, as would

occur if a cell received excessive stimulation. In addition, the two inhibitory cells (blue) are firing faster than the excitatory cell (red), a common pattern in cortical networks. Since it is not practical to look individually at each cell in the network, data mining can be used to ensure that none of the cells are behaving badly. For example, a program that would go through all of the cell traces and check them for depolarization blockade, by looking for any periods of prolonged depolarization, would be useful. It would also be important to check all of the average rates for all cells and check whether there are some cells in the network showing unrealistically high firing rates. Similarly, we might also look at the maximum firing rate for each cell. If we found such overdriven cells, we would then check their connectivity to see whether they had perhaps inadvertently been overconnected and received too much excitatory drive.

Another typical way to view network activity is the raster plot (Fig. 76.3b). In this representation, each cell has its own horizontal row. Time is represented from left to right. Each time a cell spikes, the spike is represented with a small symbol. One can customize the raster plot in order to show more about the network. For example, in Fig. 76.3b, the excitatory cells are represented by red dots and the inhibitory cells are represented by blue. The cells are also organized from top to bottom to approximate the location that they would have in the neocortex. In a network of 1,000 cells or fewer, we might look at a raster plot that shows all of the cells. With much more than a thousand, the amount of information in the raster can become overwhelming. We would then only want to look at a subset of the cells at a given time. In the example of Fig. 76.3, we might want to only look at the cells from a particular layer, and then look at the cells from another layer. Alternatively we might randomly choose 10% of the cells in each layer to look at.

Once there are too many cells to represent readily in a raster plot, we utilize lumped representations that summarize a lot of data in a relatively simple display. Two typical lumped representations are local field potential (LFP) and the multiunit activity (MUA). Both summarize activity across many units by doing an activity summation. Both also have the advantage of corresponding directly to physiological measurements. Physiologically, LFPs are generated due to differences in potential along long dendrites, typically the apical dendrites of large cells in the neocortex or hippocampus. Essentially, these dendrites are acting as antennas that send out a signal that can be read locally as LFPs or from outside the head as EEG. Since the voltage differences are set up by synaptic inputs, the LFP can be considered a measure of inputs into the cells. Many or most of these inputs will be coming from the network itself, so a physiological LFP is *not* a measure of inputs to the network. In a simulation, one can isolate just inputs and create an LFP that does represent a network input. By contrast, the MUA is the summation or superposition of action potentials from many neurons. This then can be considered as outputs from the cell, but again not as outputs from the network unless specifically isolated in the simulation.

Although lumped representations necessarily leave out a lot of detail, they can be used to see things that are not readily apparent to the unaided eye. This is properly the art and science of *data visualization*, an area that concerns itself with

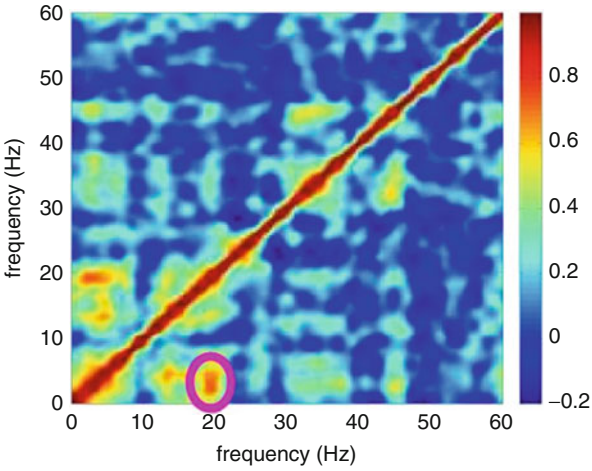

Fig. 76.4 Frequency relationships in the network of Fig. 76.3. The color scheme uses hotter (*red*) colors to indicate where oscillations in MUA frequencies occur at the same time. The *red spot* circled in *purple* indicates strong co-occurrence of 20 and 5 Hz activity (Figure adapted from Neymotin et al. (2011a))

using appropriate graphical representations to reveal patterns that are hidden within large amounts of data. The books by Edward Tufte provide field guides to a great variety of graphical representations (see Further Reading). Visualization is in turn one aspect of knowledge discovery and data mining (KDD), which concerns itself with finding patterns in data. An example of this kind of complex analysis is shown in Fig. 76.4. Here, we are re-representing a MUA, which is spiking activity summed across neurons (over a small time interval) from the raster in Fig. 76.3b. The frequency relationships shown in Fig. 76.4 are present in the MUA. These are essentially two representations of the same data (except that Fig. 76.4 has left out the dimension of time). However, one would have to spend a lot of time looking at the details of waves and waves-within-waves in the MUA and raster in order to begin to see the features that immediately hit the eye when looking at Fig. 76.4. Figure 76.4 is a reduction of the data, but it is data illumination rather than simple data summary. The red spot circled in purple shows comodulation of 20 Hz (beta) activity with 5 Hz (theta) activity in this model.

Figure 76.4 is just one example of the many kinds of analyses that can be done to extract information about a network. Another large class of data-mining visualizations involves assessment of correlations among firing cells in order to identify groups of cells that may be working together in neural ensembles. Another set of visualizations are those utilizing spectrograms.

Emergent Properties

Extremely simple networks can produce remarkably complex patterns of activity. This is most clearly seen in the case of *cellular automata*. Cellular automata are highly simplified network models with memoryless individual units which do

thresholding. They also generally only have nearest-neighbor connectivity. Although not developed to be models of neural systems, they are sometimes used as such. Cellular automata produce complex activity patterns which *emerge* from their simple rules. An emergent property is one that could not have been predicted based on knowledge of the elemental properties of the constituent parts of a system. Weather would be an example of an emergent system that can ultimately be traced to differential energy absorption across the earth's surface. Emergent properties can also occur in very simple dynamical systems, sometimes producing effects that are opposite to expectations. An interesting example dates from the early days of space exploration. An astronaut attempting to catch up with another spacecraft ahead would naturally accelerate forward toward the other vehicle. However, the more he accelerates, the further behind he falls. It was discovered that the way to catch up was to accelerate backward, effectively braking. This placed the spacecraft into a tighter orbit with a shorter orbital period. Another set of examples comes from the regime of chaotic dynamics where sensitivity to initial conditions allows small perturbations to produce big effects.

Complex patterns of activity also emerge from neuronal networks. This means that it is generally impossible to predict the patterns of activity that will occur when connecting neurons together. Varying connections, competing effects of excitation and inhibition, and the mix of various cellular dynamics all contribute to this unpredictability. Here again, dynamical patterns are sometimes the opposite of expected, due to the many recurrent loops with different effective time constants and different overall excitatory or inhibitory influence. For example, increasing synaptic weights from excitatory to inhibitory cells increases the firing rate of inhibitory cells, which then feeds back and reduces the firing rates of excitatory cells. This, in turn, decreases the firing rate of inhibitory cells, leading to an increase of excitatory cell firing rates. Round and round it goes.

One heavily studied class of emergent properties is oscillation. Observing local field potentials from electrodes placed in the brain or electroencephalogram (EEG) from electrodes on the scalp reveals voltage oscillations. Brain oscillations are generated by groups of neurons receiving inputs at or near the same time.

Various computer models have been used to investigate the origin of synchronized oscillations, with particular emphasis on the gamma band since gamma rhythms have been suggested to have relevance for cognitive processing. Early work showed that a network of interneurons is capable of generating gamma oscillations spontaneously when driven with external excitatory inputs. This type of gamma is called ING, for interneuron network gamma. In the models used, all cells may be initialized to have random membrane voltages. Since the cells are coupled, after a short period of time, all the cells synchronize and fire together. Because all the neurons in the network are inhibitory, when a presynaptic cell fires, the GABA synapses produce a hyperpolarizing change in the postsynaptic cell's voltage level. This results in the postsynaptic cell becoming less likely to fire. This raises a question as to how the synchronization emerges: if each interneuron makes

it less likely for its postsynaptic target to fire, then how do the two cells synchronize? Wouldn't the presynaptic firing simply prevent the postsynaptic cell from firing altogether?

The synchronizing effects of the mutually coupled interneurons can be explained with reference to neuronal ensembles. Although at the beginning of the simulation most of the neurons are firing asynchronously, there is a chance that several of the neurons will fire with some level of synchrony by chance. These randomly synchronous neurons can be said to belong to the initial ensemble, and all other neurons are excluded from this ensemble. All the postsynaptic targets of the ensemble are inhibited after a synaptic delay. The ensuing hyperpolarization delays the firing time of the postsynaptic targets. In other words, the postsynaptic target firing probability sharply decreases after the initial ensemble fires and gradually increases to a maxima before the ensemble fires on its next cycle. This promotes neurons outside of the ensemble to fire synchronously with the ensemble. In this way, more and more of the neurons that are initially outside of the ensemble are gradually recruited into the ensemble. Eventually, all neurons in the network oscillate synchronously. This type of emergent synchronization depends on the synaptic decay time constant being significantly larger than the gamma oscillation period. In addition, a minimum number of synapses per neuron is required that does not depend on network size.

The synchronized oscillations in the interneuron networks demonstrate that interneuron gamma can potentially drive and synchronize excitatory cell activity in neuronal networks in vivo. However, this type of model leaves out excitatory cells, which make up approximately 80% of all neurons, from being modeled explicitly. More complex models have been developed to address these shortcomings and to determine how gamma oscillations emerge from the interactions between pyramidal and interneuron populations. The gamma in these networks is called PING, for pyramidal interneuron network gamma.

These studies use simulations of networks containing both excitatory and inhibitory neurons and demonstrate that under certain conditions, inhibitory cells induce phase locking of excitatory cell firing to gamma oscillations (Fig. 76.5). In these simulations, the populations of inhibitory and excitatory cells are mutually coupled, leading to synchronization of the activity of pyramidal cells via feedback from the inhibitory cells.

For gamma oscillations to emerge, it is important that the inhibitory cells are sufficiently interconnected. This allows for the production of gamma oscillations, using a mechanism similar to that used in ING networks. However, in PING networks, since the excitatory cells both drive and receive feedback from inhibitory cells, important changes to the dynamics occur. Specifically, the excitatory cell population tends to fire together, driving the smaller inhibitory cell population to fire at high rates. This produces prolonged inhibition which then delays the firing of excitatory cells in the next gamma cycle. In this way, inhibitory-inhibitory connections can be viewed as both speeding up and sharpening the oscillation by shortening the time window for excitation of inhibitory cells after an excitatory impulse from excitatory cells.

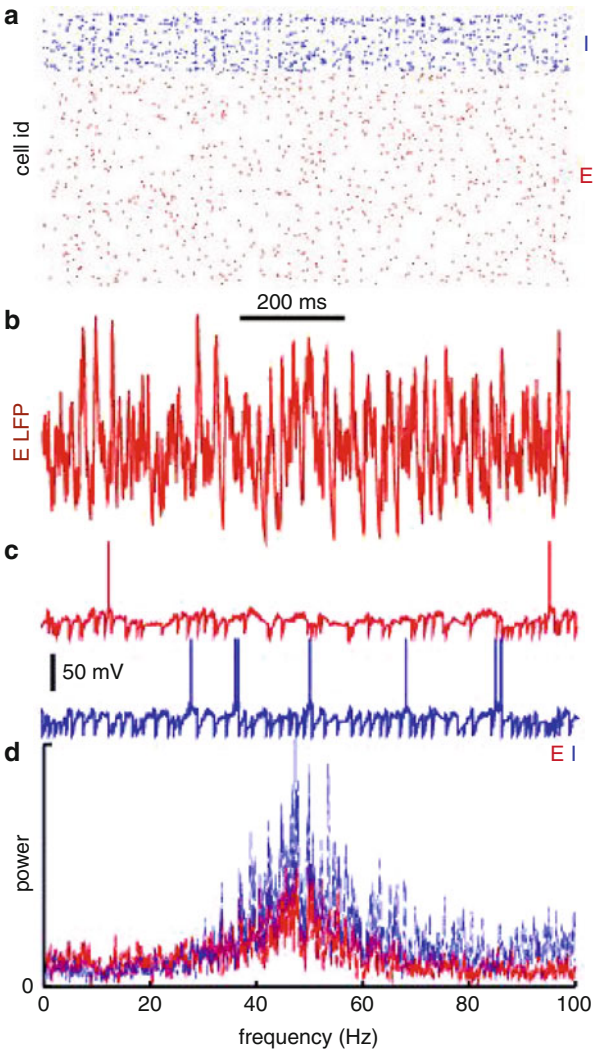

Fig. 76.5 Pyramidal interneuron network gamma emerges through interactions of mutually coupled excitatory (*E*) and inhibitory (*I*) neurons. (**a**) Raster plot (**b**) LFP from E cells. (**c**) single-cell voltage traces from one *E* and one *I* cell. (**d**) Peak in multiunit power spectrum in 40–60-Hz gamma range for both *E* (*red*) and *I* (*blue*) cells

Multiple Interneuron Classes Generate Multiple Embedded Rhythms

Studies of simulations of hippocampal CA3 networks demonstrate that different types of interneurons, differing in pattern of connectivity and in synaptic time constants, can have different roles in generating distinct oscillation frequencies in a network (Fig. 76.6). At least two classes of interneurons are known to exist in CA3: fast-spiking basket cells and slower-spiking oriens-lacunosum moleculare (OLM) cells. The basket cells synapse on each other's somata and the soma of pyramidal cells and have fast synaptic time constants at their $GABA_A$ synapses.

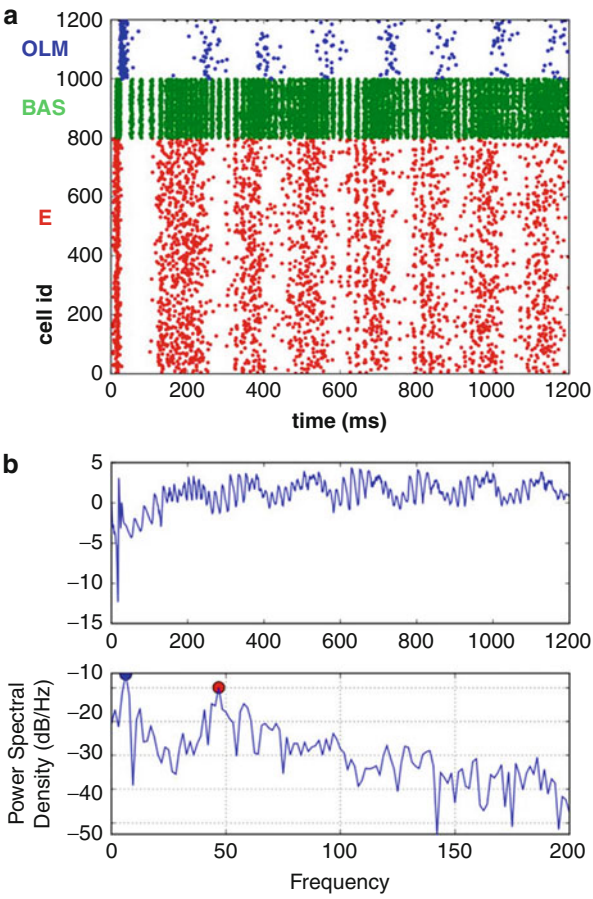

Fig. 76.6 Simulation of hippocampal CA3 with three cell populations: excitatory (*E*), basket interneurons (BAS), and oriens-lacunosum moleculare (OLM) interneurons. Spike raster is shown in (**a**). As seen, the pyramidal cells fire with moderate rates (2–4 Hz) and are synchronized by the fast-spiking basket cells, which oscillate in the gamma range (~30 Hz). The OLM interneurons generate theta activity (~4–8 Hz). Interactions between the different populations produce coupling between the theta and gamma oscillations, as seen in the local field potential in (**b**). Here, the gamma oscillation amplitude is periodically modulated by theta activity in the network. (**c**) shows that there are multiple rhythms present in the local field potentials, as seen by a peak in the theta band and a peak in the gamma band (This simulation was adapted from Neymotin et al. (2011b))

As a result, these fast-spiking basket cells are involved in the generation of gamma oscillations, utilizing the ING mechanism described above. In contrast, OLM cells synapse on the dendrites of pyramidal and basket cells, and the $GABA_A$ synapses they utilize have longer time constants. In addition, due to dendritic filtering, the effect of OLM cells are felt for longer durations. The longer timescale over which these interneurons operate allows them to contribute to the generation of the theta rhythm. Pyramidal cells synapse on both types of interneurons as well as on each other, utilizing fast-acting AMPA and longer-lasting NMDA synapses. When pyramidal cells fire at heightened rates, they increase the drive to basket cells, which in turn generates increases in gamma oscillation power. These gamma oscillations are then fed back onto the pyramidal cells, as can be seen in local field potential recordings, which are generated largely by pyramidal cells (Fig. 76.6b). In addition, the pyramidal cells provide a significant source of excitatory drive to the OLM cells, hence also playing a role in the production of the slow rhythm.

OLM cells also project directly onto basket cells. Due to the longer-lasting effects of OLM cells, this allows them to modulate the gamma oscillation strength at the theta frequency. In vivo, in addition to intrinsic connectivity, OLM and basket cells receive periodic inhibitory inputs at a rate between 4 and 8 Hz from the medial septum (MS), an external pacemaker for hippocampal theta rhythm. As a result of this action, the inhibitory cells in the network are periodically inhibited. When the interneurons are inhibited, the pyramidal cells are disinhibited and fire with higher rates, creating higher power in the gamma oscillations via the feed-forward connections to basket cells.

Connectomics and Graph Theory

A major area of interest in neurobiology is connectomics, the study of how neurons and areas are connected to one another. As a broad term, connectomics is often used to refer to how brain areas connect to one another. In a network context, connectivity may be referred to as microconnectomics, or simply as network architecture. Specific choices of network architecture, varying wiring layout between different cells and between cell classes, will lead to particular network dynamics.

Biological connection densities across and within cell type groupings are believed to lie somewhere in the wide range of 1–50%, depending on the area. In complex multilayer networks such as the neocortex, exact numbers are not available and probably will not be for some time. For example, one experimental technique for estimating connectivity is serial testing of intracellular recordings of multiple neurons simultaneously. This allows a researcher to stimulate a single neuron and record changes in the membrane potential of possible postsynaptic cells. Cells that tend to have their voltage change a short time after presynaptic neuron stimulation are assumed to be postsynaptic targets. This procedure is extremely difficult and generally results in relatively few pairings from which density is grossly estimated. By contrast, connections can be visualized by anatomical techniques at the electron microscope level or inferred from appearances at the light microscope level. However, the existence of an anatomical synapse does not guarantee the presence of a functional connection.

For lack of information, neuronal network wiring must start from whatever limited data is available. Data for a particular model may come from different species and different brain areas. Being widely studied, cat and primate visual cortex and rat barrel cortex are favorite reference areas when constructing neocortical models. However, these particular sensory areas are probably not particularly typical of all cortical areas, given the many anatomical feature differences across cortical regions, as originally described by Brodmann over a century ago. Different cortical areas feature different thicknesses of layers and numbers of cells in each subpopulation and are likely to have different connectivity densities. These anatomical differences will then be associated with differences in network dynamics.

The study of connectivity in neuronal networks has its antecedents in graph theory, a branch of discrete mathematics that originated with the work of Euler in 1736. Graph theory treats graphs, or networks, as mathematical objects, defined by

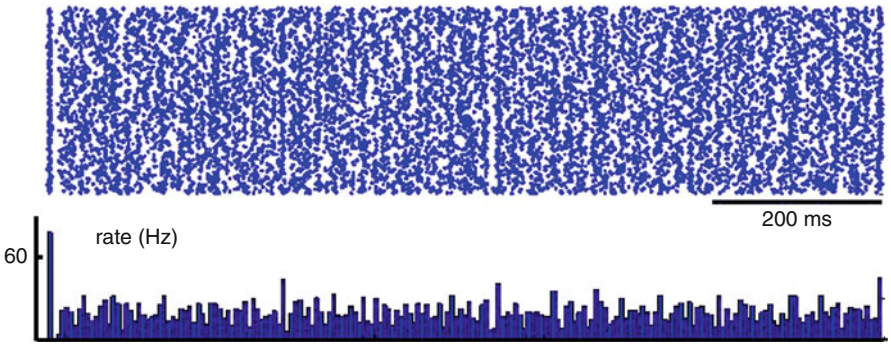

Fig. 76.7 Randomly wired network of excitatory and inhibitory cells developed by Brunel. Network has 12,500 cells; 400 cells are shown in the raster plot at the top (each *dot* represents a spike occurrence for each cell; cells arrayed from top to bottom). This network has 15,639,500 synapses. Below the raster is a representation of average excitatory cell firing frequencies (Simulation adapted from Brunel (2000))

vertices and edges. Vertices (also called nodes) are connected by edges (also called lines or connections). In neuroscience, the vertices may represent cells or may represent brain areas. Correspondingly, the edges may represent synaptic connections or axonal tracts. Although much of graph theory works with nondirectional edges, directed graphs are generally used in neuroscience, corresponding to the directionality of information flow through axons and synapses.

The use of graph theory in neuroscience is a direct consequence of our ignorance. In the absence of precise knowledge of connectivity, it is important that we systematize our network architectures so as to be able to define or categorize one network relative to another. This then enables us to develop network taxonomies and to relate structure to function. With the absence of information, our networks are random graphs, a field of study developed by Erdos and Renyi in the mid-twentieth century. Random graph theory provides descriptions of the properties and statistics of random graphs and algorithms to generate or to categorize these graphs. Random graphs are easy to generate in several ways. For example, one way, simple though not optimally efficient, is as follows: First, one decides on the size of the graph, given by the number of vertices n. Then, there is a fixed probability, p, of connections between vertices. Next, all pairs of vertices are traversed, and a random number generator selecting numbers uniformly from the interval [0,1] is used to determine whether the given cell pair connects, using p. If the number drawn from the random number generator is less than or equal to p, a connection between the given pair is formed.

In many cases, a neuronal network to be explored will be fully random. A p is chosen to represent the density of connectivity between neurons. Often, a p higher than the estimated density is used, because the network being simulated is so small compared to the network being represented. The higher p in this case allows for higher convergences and divergences and greater interactions between cells in the network. Figure 76.7, a network by Brunel, shows the dynamics of a randomly connected network with excitatory and inhibitory cells.

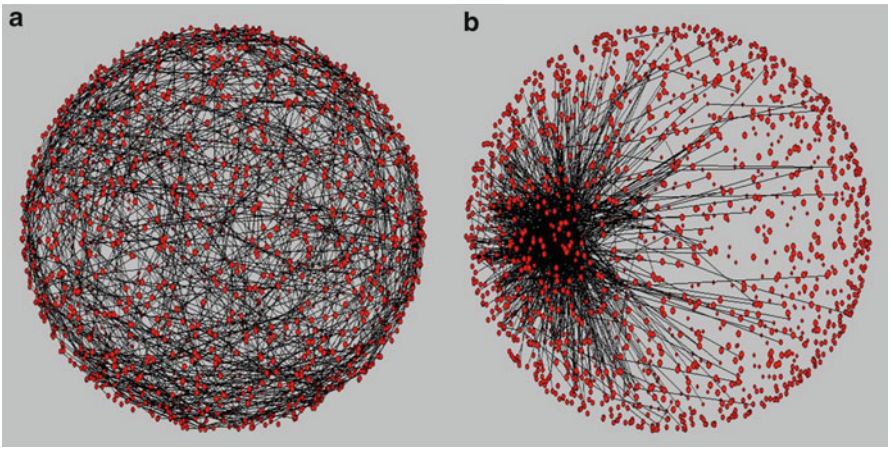

Fig. 76.8 Two randomly wired networks drawn with different degree distributions. Only selected connections are shown since drawing the full graph would completely obscure the picture. (**a**) Normal (Gaussian) degree distribution, with all cells having similar degree. (**b**) Heavily skewed degree distribution, with a few cells (hubs) having many more connections than average and most cells sharing a small number of connections (Figure drawn with Networks/Pajek program)

Graph Theoretic Measures

Graph theory has a multitude of measures useful for characterizing the properties of graphs. The simplest measures include a cell's convergence (number of inputs) and divergence (number of outputs). If a cell's convergence is high, this can signify its role as an integrator of multiple streams of information. Similarly, cells with large divergence may be more effective at spreading information throughout a network. Convergence and divergence can be further subdivided based on the types of inputs from and outputs to the different cell classes in the network. These quantities can be used to determine if a cell is more or less likely to fire; for example, if the cell has more excitatory inputs than inhibitory inputs, it may be hyperexcitable and biased toward firing.

The graph theory terms for convergence and divergence are in-degree (or simply *degree*) and out-degree, respectively. Note that although in-degree and out-degree can differ for an individual cell, the network's average in-degree must exactly equal the average out-degree. To understand network dynamics, it is useful to count the number of cells in a network that have a specified degree. The distribution of the degrees of all cells in the network is referred to as the network's degree distribution. Random graphs tend to have a distribution that appears as a bell curve, with most cells having a similar degree around the average value (Fig. 76.8a). In these networks, the probability of a cell having very many more or less than average degree trails off exponentially with distance from the mean. Some real-world graphs have been shown to have degree distributions significantly different from

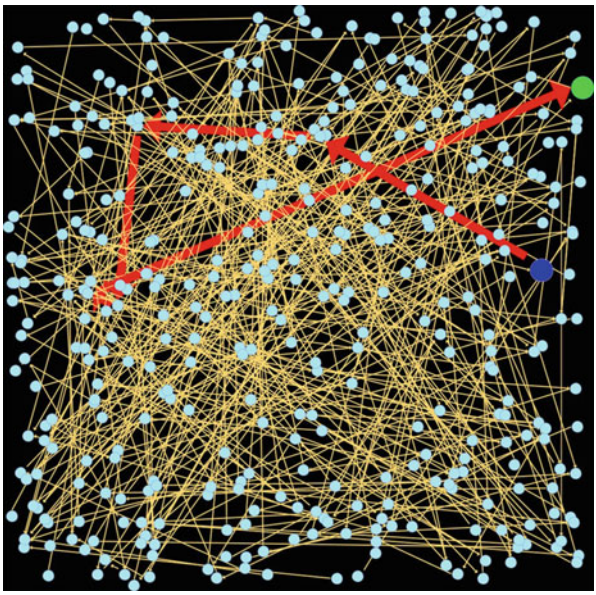

Fig. 76.9 Network of 400 cells with random connectivity and low connection density. Cells drawn in light *blue* and synapses between cells drawn with *yellow* lines, with *arrowheads* representing direction of connection from pre- to postsynaptic cell. Synaptic pathway from *blue* (B) to *green* (G) cell is highlighted by the *red* path. This pathway utilizes five synapses and therefore has a path length of 5. Although the cells are spatially close to each other, due to random connectivity, the pathway traverses the graph in a roundabout fashion. In order for information to get from B to G, all five synapses must be traversed (Figure drawn with GUESS: The Graph Exploration System)

those of Erdos-Renyi random graphs. An example of this is shown in Fig. 76.8b, where a small set of cells act as hubs which receive and give a large number of projections.

Average path length is one of several graph theoretic measures that is useful for network analysis (Fig. 76.9). Average path length gives the average distance that you need to travel to get from any neuron to any other neuron along edges. A low value of average path length indicates closeness between any two cells and provides for efficient communication within a network. Graphs with short average path lengths tend to spread activity more readily. Hubs serve as focal points. They receive activity widely and then transmit it widely, giving many node pairs that are connected via a 2-hop path, path length 2. Note the comparison to airport hubs which are used because they allow one to readily travel between any two points in the country. Other common graph theoretic measures include clustering and centrality, the latter measuring the number of shortest paths between any two cells that a given cell lies on.

The presence of hubs is typical of small-world networks. Much of this work has derived from studies of social networks, used to model relationships between

people. Each vertex in such a graph would correspond to a person. A nondirected edge between two people indicates that they know each other or are connected via some activity. For example, the movie actor Kevin Bacon has become (erroneously) famous as a hub; in these analyses, the nodes are actors and the edges are provided where two actors appear in the same movie. Small-world networks have short average path lengths and high clustering: it is easy to find a path between two people, and people who are acquainted with one person tend to know each other. A small-world network can be generated readily as follows: First, create a regular graph in a square grid with each cell connecting only to its four nearest neighbors. Thus, there is high clustering, but the path length between cells tends to be high, as it takes many steps to get from one side of the graph to the other. To make a small-world network, the average path length must be reduced. This is accomplished by taking existing edges and randomly rewiring them to connect a random pair of cells. After some rewiring, a small-world network emerges.

Graph theoretic analyses of brain data, ranging from human fMRI to mammalian electrophysiology to electron micrographic analysis of tiny invertebrates, have suggested that brains and neuronal networks may have small-world properties. This suggests the existence of hubs which serve as well-connected modules that provide anatomical/functional segregation. The short path lengths suggest efficient information transmission between modules. The presence of such modules may also be important in setting neuronal dynamics and aiding in synchronization.

A very thorough combination of simulation, electrophysiology, and anatomical analysis, performed by Soltesz and coworkers, demonstrated that granule cells in the hippocampal dentate gyrus might serve as hubs that pathologically promote hypersynchrony and epilepsy. In the simulations, granule cells were highly effective in promoting seizures after only a few of them were converted into hubs through the development of new connections through sprouting. These studies also showed that pathological hubs had to have both high convergence and high divergence. A hub with only high in-degree (convergence) or only high out-degree (divergence) would not predispose to epilepsy. The existence of hubs in the epileptic dentate gyrus was explained anatomically by noting that these cells had both large dendritic and large axonal arborizations.

At the larger spatial scale of networks of brain regions, hubs have been found to play an important role in functional integration. Other recent work on fMRI-based functional networks has demonstrated the presence of hubs in regions of unimodal or multimodal association cortex, suggesting these types of hubs can integrate different forms of information effectively.

Information Flow

In 1948, Claude Shannon introduced information theory to rigorously quantify the information content of a signal. This theory has been applied to the problem of neuronal coding in order to try to understand how neurons in the brain represent and transmit information. Unfortunately, information is intuitively a somewhat

nebulous thing. As we will see below, the information content of a signal depends not only on the signal itself but also on the preexisting knowledge of a particular receiver or on the set of similar signals that that receiver has access to. Because of this, one person's information is another person's nonsense. This conundrum is related to one of the many great unanswered questions about the brain: is the brain filled with uninterpreted information, or is it just filled with noise? Many would argue that neurons and synapses are intrinsically noisy components through which an information-containing signal must scream in order to be heard.

Information theory rests on the foundation of probability theory. The key insight is that the probability of an event (a particular signal) is inversely proportional to the information it conveys. This is the notion of surprise: the unexpected signal has more information. Surprise is captured by the equation for information gain, $I = \log(1/p)$, where p is the probability of an event occurring. The inversion of p means that the lower the probability, the higher the information. A weather prediction of 90^{o} F for mid-afternoon, mid-summer, in a mid-latitude location would carry little information; the information gained by hearing that communication is minimal. By contrast, the prediction of snow would be a major piece of information. The logarithm is used to make information additive – if you read half a book and read the other half later, the two information segments will add up to be the same as the value associated with reading the whole book at one sitting.

Typically, one is dealing with a large number of signals rather than just one signal. In this case, rather than dealing with the probability of a single event p, we deal with probability distributions involving large numbers of events, p_i. These events can be studied by measuring the entropy of a probability distribution. The entropy of a system X is the average information gained from the occurrence of events in the distribution, $H(X) = p_1 \cdot \log(1/p_1) + p_2 \cdot \log(1/p_2) + p_3 \cdot \log(1/p_3) + \ldots$. Those who have studied physical chemistry will recognize the concept of entropy from statistical mechanics. Entropy will be maximal when all events in the distribution are equally likely to occur, which occurs for a uniform probability distribution. In this situation, the ability to predict future events from knowledge of past events is minimal.

Conversely, for low entropy distributions, new information is of low value since predictions may be based on prior knowledge – knowledge of the relatively few possibilities for this particular information system. In the probability distribution, most events cluster around a particular value, so it is likely that future events will occur close to that value. In order to compute probability, we count events. Then $p(f) = [(n_f)/n]$, where n_f is the number of times the event f has occurred and n is the total number of observations for all events. In the case of a continuous signal, we must first discretize. In a digital computer, for example, continuous voltage is divided up into low, the 0 signal (clear bit), and high, the 1 signal (set bit). Any voltage in between is thresholded so as to count as one of these two possibilities.

There are a number of related information theoretic measures that have to do with transfer or transmission of information. Mutual information (MI) uses the dependence between two probability distributions, X and Y, to measure the information gained about X, when information is received about Y. MI is useful for quantifying how much information about a stimulus variable a neuron represents.

For example, MI has been used to quantify the amount of information neurons in the visual cortex use to encode the different patterns of light hitting the retina.

MI is a useful measure but lacks any notion of direction, as to whether information is flowing from X to Y or from Y to X. In neuroscience, we would like to know how neurons and networks influence one another. When successful, knowledge of the direction of information transfer between neurons can allow inference of functional or synaptic connectivity. However, we find that we are often unable to distinguish the leader from the follower. In many cases, there will be no consistent causality since two areas are hopelessly intertwined through multiple paths of mutual interconnection. This consideration further motivates the notion of computer simulation as an experimental pursuit. Even though everything is laid out by the investigator, and everything is visible to the investigator, there are still these mysteries that must be explored experimentally. The complexity of neural systems, with their many recurrent dynamical loops, means that sophisticated tools, such as those of information theory, are needed to determine activity flow in silico, just as they are in vivo.

Several information theoretic measures have been developed to assess the direction of information flow. These include Granger causality, the directed transfer function, and normalized transfer entropy (nTE). The key to these methods is that the results differ depending on the direction: knowing X's past may tell you something about Y's future, whereas knowing Y's past may not tell you anything about X's future. The key thing to remember about all of these methods is that they are weak. First, they can be fooled: a signal that appears to flow B → C might be emitted by A but received by B before being received by C. Second, they are complex statistical estimates where inaccuracy may be introduced due to chance. They generally require complex normalization to try to avoid such statistical artifact: this is the "n" in nTE. We will now describe the use of nTE as an exemplar of this class of measures.

In order to use nTE, spike trains from neurons or inputs are typically binned to give the number of spikes per time interval. The bin size for nTE calculations can be set similar to a synaptic delay and propagation of activity from dendrite to soma/axon, where the postsynaptic spike is generated. This can be in the order of several milliseconds. Varying the size of the bin can have a large impact on the value of nTE and may even determine whether a significant value will be detected.

Figure 76.10a shows two networks, with excitatory cells in red and inhibitory cells in blue. We use nTE to determine the direction of information flow across the two networks when a weak unidirectional connection is created to connect them from left to right. nTE was measured in both directions using 18-s periods selected during the course of a several-minute simulation (Fig. 76.10b). Each point shown is an estimate of influence based on assessing the influence of one excitatory cell on one side to one excitatory cell on the other side. Only a few cell pairs are represented; most cell pairs cannot be shown to influence each other at all. However, the results show a fairly reliable directional difference above about 1% connectivity. For example, at 1.6% connectivity, 22 (2.2%) cells are found with presumptive influence in a left-right direction. This suggests that the method is detecting disynaptic as well as monosynaptic projections.

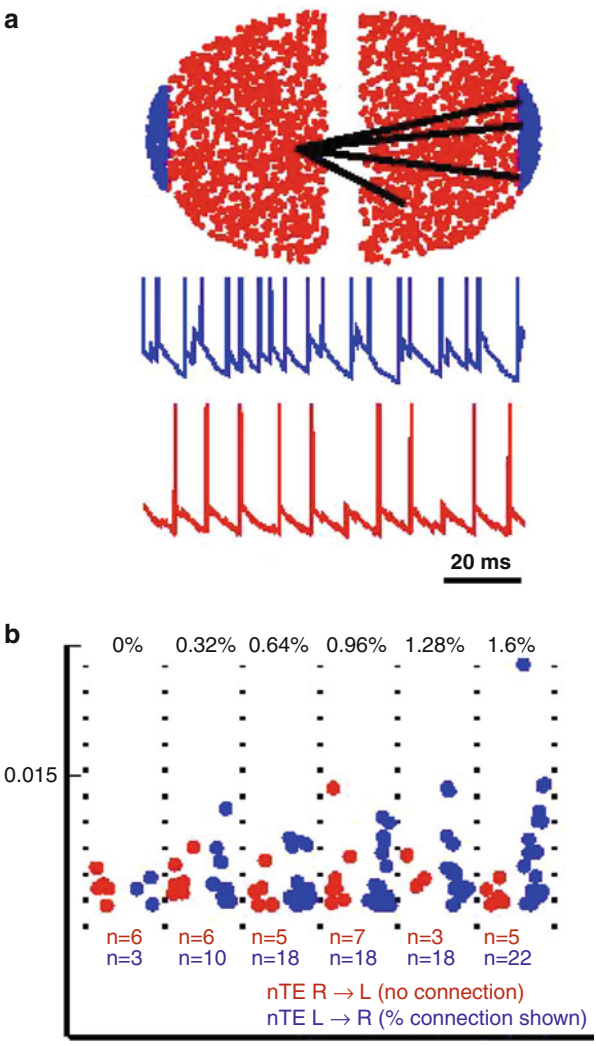

Fig. 76.10 Two neuronal networks with two types of cells each, inhibitory (*blue circles*) and excitatory (*red circles*), shown in **a**. Schematic of simulation layout and example voltage traces shown from inhibitory (*blue*) and excitatory (*red*) cells. (**b**) Connections are added incrementally from the *left* to the *right* network, and nTE is measured between all cell pairs in both directions. Individual cell pairs that display significant nTE are shown (*blue*: left to right; *red*: right to left). As connectivity from the *left* to the *right* side increases (shown horizontally along the x-axis of **b**), the number of cells with significant L → R nTE increases (*blue dots*). There are also false positives (*red dots*), wrongly indicating R → L projections

However, at the same time, the method appears to be spuriously detecting five pairs with evidence of projection in the wrong direction, a direction where there are in fact no projections. This is the influence of randomness on the measure – cell B firing before cell A a number of times randomly has given the mistaken impression that cell B projects to cell A. This problem could be partly solved by basing the calculations on a larger amount of data, thus averaging out more of the random coincidences. However, the assessments shown already required an enormous amount of data. Although not a problem in our in silico sandbox, similar methods applied experimentally would also run into the problem of nonstationarity – the system is changing during the long periods of time that we are recording from – and

cell A may only be strongly influencing cell B during part of the time. There are also technical limitations on physical recordings; current physiological technology does not allow for reliable recording from more than about 50 cells simultaneously.

Figure 76.10 illustrates that the nTE measure can be very sensitive to small effects. Each cell is receiving a large number of inputs from many presynaptic cells, each of which is only contributing a small subthreshold activation. The detection of a presynaptic effect from a single cell is the proverbial problem of the needle in a haystack: the effect of any single presynaptic cell on a single postsynaptic cell is weak and hard to detect. Detection threshold depends not only on the strength of the interaction but also on the amount of data available – information theory methods usually require a lot of data. Also, providing sensitivity to weak connections means being vulnerable to false positives: the red data in Fig. 76.10b which erroneously suggest the presence of projections in a right to left direction.

Using the techniques described in this section, it is possible to quantify how information transfer is modulated by a network. Simulations of neuronal networks have been used to measure how the activity and synaptic weights in networks alter information flow related to experimental data from animal models of epilepsy and schizophrenia.

Activity Spread: Avalanches and Traveling Waves

How does activity in neuronal networks evolve over time and space? One phenomenon studied with neuronal networks is the avalanche. In this context, an avalanche is defined as runaway activation that spreads far from its origin. In order to create an avalanche, the average number of cells recruited to the avalanche at each step must exceed one. Otherwise, the avalanche will die out. The balance between excitation and inhibition in a network may be important parameters controlling the spread of avalanches. If all neurons in a network are excitatory, and sufficient interconnections exist, avalanches are likely to occur, but these avalanches will be brief and will self-terminate due to running out of further cells to excite; previously excited cells will go into a refractory state. This is the situation of an interictal spike, a phenomenon typically seen in epileptic tissue. Inhibition can prevent this rapid runaway excitation but may actually worsen avalanche pathology if it leads to a more prolonged abnormal activity state. The relation of avalanches to the excitation-inhibition balance suggests that neuronal networks may be sensitively poised on the verge of large-scale transitions in dynamics.

A recent large-scale thalamocortical model from Izhikevich showed wavelike patterns of activity similar to those observed in vivo. This model also showed miniature avalanche-like effects; adding an extra spike anywhere in the model changed global cortex activity after just a few hundred milliseconds. The model ended up displaying a different activity pattern than it would have without the spike. One extra spike triggered numerous extra spikes or missed spikes. In this way, a little local activity spread over the entire network, and every neuron was affected. Clearly, damage to single neurons does not dramatically degrade brain

function, the phenomenon of graceful degradation. If a single spike can significantly alter spike patterns of other individual neurons, then this suggests that either the single spike or the single neuron or both are not the proper level of analysis for understanding information processing in the network. This meshes with the notion of ensemble codes and distributed neural processing, which suggest a relatively minor role for each neuron when taken individually but a contribution of many neurons through coordinated activity.

Homeostasis

In order to maintain stability, animals sense changes in the environment and then adjust internal processes of internal states, such as temperature and blood pressure. This is homeostasis. In the brain, homeostasis needs to adjust activity not only in response to external changes to the environment but also in response to internal changes elsewhere in the brain. For example, in retrieving one set of associations, the brain may have to suppress the activity of other neural ensembles which provide related but irrelevant information, actively coordinating the different ensembles. These mechanisms, which are part of the selective activations of memory, have been hypothesized to share underlying mechanisms with homeostatic functions.

One form of brain homeostasis entails maintenance of baseline activity patterns across varying levels of inputs and neuromodulators. For example, frequency homeostasis is the maintenance of the physiological spectrum of frequencies and frequency interrelationships despite alternation of drive. This drive can take the form of enhanced sensory information or modulation of activity via attentional faculties. As an example of alteration of intensity of sensory information, consider the problem of visually identifying an object in conditions of variable-intensity light. When a room is dimly lit, the visual system makes optimal use of the minimal information present in the environment. When viewing a brightly lit outdoor scene, the high intensity inputs entering the nervous system could potentially cause overactivation of cells, preventing normal processing from occurring. Although some of this mechanism could be described as simple adaptation, the broader term homeostasis encompasses the many dynamical readjustments that must be made downstream to maintain the entire system in a state of equilibrium.

Attentional modulation of perception, demonstrated physiologically, has been hypothesized to arise from augmented drive from higher to lower sensory systems. This pattern of added driving can be produced in silico by adding feedback into supragranular layers (layer 2/3) in neocortical models. A major effect of adding such drive is that this extra activity spreads to the rest of the network, increasing firing rates of all cell populations. In our neocortical model, these large changes in individual cell firing rates did not produce substantial changes in the spectral profile, an example of frequency homeostasis (Fig. 76.11). Spectral power increased across all frequencies with no appreciable shift in peak, demonstrating that the network compensated internally for the increased activation. This homeostatic effect in the simulation appeared to be provided by a balance between

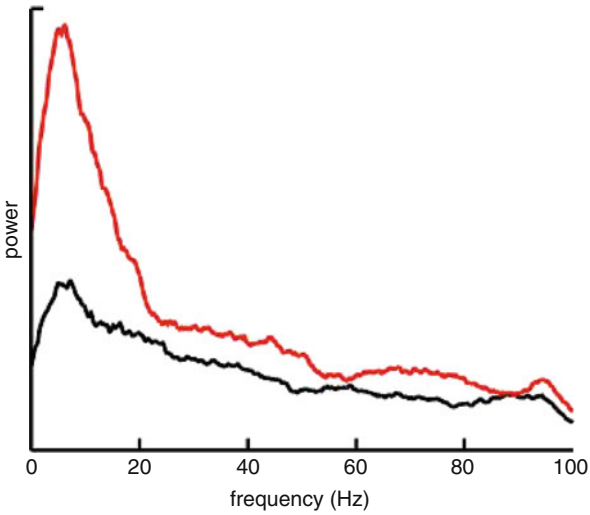

Fig. 76.11 Frequency homeostasis: no appreciable shift in excitatory cell spectral peak with increased external inputs to layer 2/3 excitatory cells, which may be caused by enhanced attention in vivo (*black*: baseline; *red*: 10% increase in AMPA activation onto excitatory cells) (Figure adapted from Neymotin et al. (2011a))

inhibitory cell domination of the high-end of the spectrum (gamma frequencies) and excitatory cell domination of the low-end of the spectrum (theta/alpha). The increased inputs pushed both excitatory and inhibitory populations which led to countervailing tendencies to move toward higher network frequencies (due to direct effects of the excitatory cells) and toward lower network frequencies (due to indirect effects of inhibitory cells which sculpt out these higher frequencies).

When homeostasis and feedback between excitatory and inhibitory populations fail, oversynchronization between neural elements can occur, leading to the occurrence of population spikes (when many cells discharge over a short duration) or seizures. One homeostatic mechanism built into neural architecture that may prevent seizure spread is the pattern of feed-forward inhibition seen connecting adjacent cortical areas. Overexcitation within one area may then inhibit surrounding areas and reduce the spread of pathological activation.

Outlook

One of the great frontiers of brain science in the coming decades involves the assessment of meaning in the brain. Part of this task will involve the discovering of the neural code, or, more likely, neural codes. These codes somehow provide the representations on which meaning, cognition, and behavior depend, just as the genetic code provides the representations that determine structure and mechanism in both body and brain. Due to the complexity of neural systems, computational neuroscience will be a necessary partner in this endeavor. In this, we see the same patterns of computational assistance that have become vital for studying and understanding weather and other large systems. In all of these cases, computers are necessary both to manage massive amounts of information and to put the information into context using simulation technology.

For those who are interested in doing this type of research, a diverse background in neuroscience, computers, and engineering is valuable. Fields of applied mathematics/engineering that are particularly valuable include signals and systems, numerical calculus, linear algebra, and probability theory. Most of these topics are covered in the field that used to be called *Scientific Computing* but is now referred to as *Computational Science*. It is also important to become a sophisticated computer user. Ironically, computer skills are deteriorating in this era of ubiquitous computing, now that all everyday computing tasks have been simplified to the click of a button.

Glossary

Dynamics The study of how a process changes over time.

ING Interneuron network gamma – gamma oscillations generated by the activity of interneurons.

PING Pyramidal interneuron network gamma – gamma oscillations generated by the interactions of pyramidal and inhibitory neurons.

Raster plot Plot showing spiking of cells – vertical axis is cell identity and horizontal axis is time.

Further Reading

Anderson JA (ed) (1995) An introduction to neural networks. MIT Press, Cambridge

Brunel N (2000) Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. J Comput Neurosci 8:183–208

Hertz J, Krogh A, Palmer RG (1991) Introduction to the theory of neural computation. Addison-Wesley, Redwood City

Izhikevich EM (2003) Simple model of spiking neurons. Neural Networks, IEEE Transactions on 14(6):1569–1572

Izhikevich EM (2010) Dynamical systems in neuroscience: the geometry of excitability and bursting. MIT Press, Cambridge, MA

Koch C (2004) The quest for consciousness: a neurobiological approach. Roberts and Company Publishers, Englewood, CO

Koch C, Segev I (1998) Methods in neuronal modeling: from synapse to networks, 2nd edn. MIT Press, Cambridge, MA

Lytton WW (2002) From computer to brain. Springer, New York

Neymotin SA, Lee HY, Park EH, Fenton AA, Lytton WW (2011a) Emergence of physiological oscillation frequencies in a computer model of neocortex. Frontiers in Computational Neuroscience 5:19

Neymotin SA, Lazarewicz MT, Sherif M, Contreras D, Finkel LH, Lytton WW (2011b) Ketamine disrupts theta modulation of gamma in a computer model of hippocampus. J Neurosci 31(32):11733–11743

O'Reilly R, Munakata Y, McClelland JL (2000) Computational explorations in cognitive neuroscience. MIT Press, Cambridge, MA

Sporns O (2010) Networks of the brain. MIT Press, Cambridge, MA

Tufte ER (1990) Envisioning information. Graphics Press, Cheshire, CT