Problem 94: Analysis of DNA sequences. The file `seqs.txt` contains data from the experiments of Kinney et al. (2010); focus on the data labeled crp-wt. There are a large number of sequences for the promoter region of a gene, and each sequence is associated with a batch B0, B1, . . . , B9. B1 through B9 are those sequences that generate expression levels in some specified ranges, and B0 is a sample of all sequences used in the experiment. At each site along the sequence, you should have enough samples to estimate the joint distribution of the base at that site (four possibilities) and the batch number (nine possibilities). Using the ideas in Chapter 6, compute the amount of information that the base gives about the expression level (or its surrogate in this experiment, the batch number). Can you identify regions of the sequence that are particularly informative? Does the information ever go to zero, suggesting that some bases are completely irrelevant to the control of expression?

Now that we have some confidence in our description of the binding energies, we can go back and ask once more about the statistics of sequences and the problem of robustness versus fine tuning. The essential point, I think, is that although proteins bind to segments of DNA that are ∼10 bases long, to first approximation all that matters is the binding energy, which is one number. Out of $\sim 4^{10}$ sequences, there are many ways of achieving the same value of the binding energy. The natural hypothesis is that evolution allows nearly random wandering among the isoenergetic sequences; comparing the genomes of closely related organisms gives results consistent with this hypothesis.

Our discussion has focused on the interaction of (mostly) single proteins with single DNA sequences. The real problem is in a larger context—there are more than 100 transcription factors, even in relatively small bacteria, and several thousands of functional binding sites. It is not enough that one protein bind to one site with a sensible energy; it must also be true that this protein does not bind to the wrong sites, nor that other proteins bind to its intended site. If we can take seriously the representation of a protein's sequence specificity by a linear model, then we can start to ask how different transcription factors distribute themselves in the space parameterized by the matrix elements $W_{i\mu}$, and whether this distribution provides maximum discriminability among the alternative regulatory signals. As explained in the references (see the Annotated Bibliography for this section), there are some interesting efforts in this direction, but much remains to be done.

5.2 Ion Channels and Neuronal Dynamics

The functional behavior of neurons involves the generation and processing of electrical signals, voltage differences, and current flows across the cell membrane. As noted in our discussion of the rod photoreceptor cell (Section 2.3), the membrane itself is insulating, and hence there would be no interesting electrical dynamics if not for the presence of specific conducting pores or channels. These pores are protein molecules that can change their structure in response to various signals, including the voltage across the membrane, which means that the system of channels interacting with the voltage constitutes a potentially complex, nonlinear dynamical system. We can also think of the ion channels in the cell membrane as a network of interacting protein molecules, with the interactions mediated through the transmembrane voltage. In contrast to many other such biochemical networks, we actually know the equations that describe the

network dynamics, and as a result the questions of fine tuning versus robustness can be posed rather sharply.

When we move from thinking about individual neurons to thinking about circuits and networks of neurons, which really do the business of the brain, it is easy to imagine that the neurons are circuit elements with some fixed properties. We enhance this tendency by drawing circuit diagrams in which we keep track of whether neurons excite or inhibit one another, but nothing else about their dynamics is made explicit. Indeed, we will take this point of view in Section 5.4, just to keep things tractable. Despite our hopes for simplicity, our genome encodes $\sim 10^2$ different kinds of channels, each with its own kinetics. This range is expanded even further by the fact that many of these channels have multiple subunits, and cells can splice together the subunits in different combinations. On the one hand, this creates enormous flexibility and presumably adds to the computational power of the brain. On the other hand, this range of possibilities raises a problem of control. A typical neuron might have eight or nine different kinds of channels, and we will see that the dynamics of the cell depend rather sensitively on how many of each kind of channel is present. In keeping with the theme of this chapter, it might seem that cells need to tune their channel content very precisely, yet this needs to happen in a robust fashion.

To explore the trade-off between fine tuning and robustness in neurons, we need to understand the dynamics of the channels themselves.[7] For simplicity, let's neglect the spatial structure of the cell and assume we can talk about a single voltage difference V between the inside and the outside of the cell. Then because the membrane acts as a capacitor with capacitance C, we can write, quite generally,

$$C\frac{dV}{dt} = I_{\text{channels}} + I_{\text{ext}}, \tag{5.32}$$

where I_{ext} is any external current that is being injected (perhaps by us as experimenters) and I_{channels} is the current flowing through the channels. Each open channel acts more or less as an Ohmic conductance, and the structure of the channel endows it with specificity for particular ions. Because the cell works to keep the concentrations of ions different on the inside and outside of the cell, the thermodynamic driving force for the flow of current includes both the electrical voltage and a difference in chemical potential; it is conventional to summarize this by the reversal potential V_{i} for the currents flowing through channels of type i, which might involve a mix of ions. Because current only flows through open channels, we can write

$$I_{\text{channels}} = -\sum_{\text{i}} g_{\text{i}} N_{\text{i}} f_{\text{i}} \cdot (V - V_{\text{i}}), \tag{5.33}$$

where g_{i} is the conductance of one open channel of type i, N_{i} is the total number of these channels, f_{i} is the fraction that are open, and V_{i} is the reversal potential. If each

7. As with many of the topics discussed in this text, we could spend an entire semester on ion channels and not exhaust the subject. I admit that in some sections of the book I feel that I am providing a good guide to potentially complex matters, whereas in other sections I feel very strongly the weight of the things I am leaving out. As always, I encourage you to dig into the references in the Annotated Bibliography for this section.

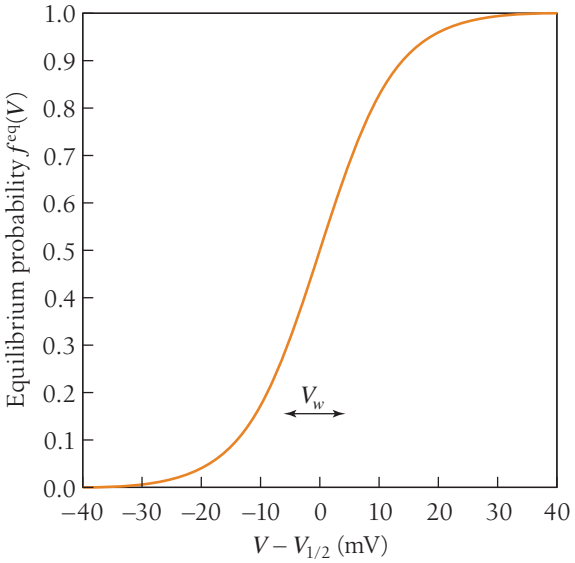

FIGURE 5.21
Activation curve for an ion channel, from Eq. (5.36), with $Q = 4$.

channel has just two states, open and closed, then the dynamics would be described by

$$\frac{df_i}{dt} = -\frac{1}{\tau_i(V)}\left[f_i - f_i^{eq}(V)\right]. \tag{5.34}$$

The equilibrium fraction of open channels as a function of voltage, $f_i^{eq}(V)$, often is called the activation curve, and $\tau_i(V)$ is the time constant for relaxation to this equilibrium.

What is a reasonable shape for the activation curve? We are describing a protein molecule that can exist in two states, and the equilibrium between these two states depends on voltage. This is possible only if the transition from closed to open rearranges the charges in the protein. In the simplest model, the opening of the channel effectively moves a charge Q across the membrane, and so the free-energy difference between open and closed states will be $\Delta F = F_0 - QeV$. Then the equilibrium probability of a channel being open will be given by

$$f^{eq}(V) = \frac{1}{1 + \exp\left[(F_0 - QeV)/k_BT\right]} \tag{5.35}$$

$$= \frac{1}{1 + \exp\left[-(V - V_{1/2})/V_w\right]}, \tag{5.36}$$

where the point of half maximal activation is $V_{1/2} = F_0/(Qe)$, and the width of the activation curve is $V_w = k_BT/Qe$, as shown in Fig. 5.21. The charge Q is referred to as the gating charge. At room temperature we have $k_BT/e = 25$ mV, so that even with relatively small values of Q we expect channels to make the transition from closed to open in a window of $\sim$10 mV. The location of the midpoint $V_{1/2}$ depends on essentially all aspects of the protein structure in the open and closed states, so it is harder to develop intuition for this parameter.

In the absence of external inputs, all currents through the channels have to cancel. If only one kind of channel dominates, it will drive the voltage across the membrane

to V_i, the reversal potential for that particular channel. Cells have pumps that maintain differences in the concentration of ions between the inside and outside of the cell, but these differences are not infinitely large. If, for example, there is a ratio of 100 between the internal and external concentrations of an ion with charge 1, then the reversal potential for the flow of this ion will be $\pm(k_B T/e)\ln(100) \sim 115$ mV, where the sign depends on whether the higher concentration is inside or outside the cell. This estimate suggests that full dynamic range of voltage changes across the membrane will be limited to ± 100 mV, which is correct.

To get started on the dynamics, it is useful to identify the "resting potential" $V = V_0$, and study small perturbations around this steady state. The full dynamics are

$$C\frac{dV}{dt} = -\sum_i g_i N_i f_i \cdot (V - V_i) + I_{ext}, \tag{5.37}$$

$$\frac{df_i}{dt} = -\frac{1}{\tau_i(V)}\left[f_i - f_i^{eq}(V)\right], \tag{5.38}$$

and the linearization is

$$C\frac{d\delta V}{dt} = -\sum_i g_i N_i f_i^{eq}(V)\delta V - \sum_i g_i N_i (V_0 - V_i)\delta f_i + I_{ext}, \tag{5.39}$$

$$\frac{d\delta f_i}{dt} = -\frac{1}{\tau_i(V_0)}\left[\delta f_i - \left.\frac{df_i^{eq}(V)}{dV}\right|_{V=V_0}\delta V\right]. \tag{5.40}$$

Fourier transforming, we can solve for the channel dynamics, then substitute and collect terms to find

$$\left[-i\omega C + \frac{1}{R_0} + \sum_i \frac{g_i N_i (V_0 - V_i)[df_i^{eq}(V)/dV]_0}{1 - i\omega\tau_i(V_0)}\right]\delta\tilde{V}(\omega) = \tilde{I}_{ext}(\omega). \tag{5.41}$$

The resting resistance of the membrane is defined by

$$\frac{1}{R_0} = \sum_i g_i N_i f_i^{eq}(V). \tag{5.42}$$

The term in brackets in Eq. (5.41) is the inverse impedance (or admittance) of the system.

Problem 95: Details of membrane impedance. Fill in the steps leading to Eq. (5.41).

To understand what is going on here, it is helpful to think about channels with fast ($1/\tau_i \gg \omega$) or slow ($1/\tau_i \ll \omega$) responses. The fast channels renormalize the resistance:

$$\frac{1}{R_0} \to \frac{1}{R_0} + \sum_{i\in fast} g_i N_i (V_0 - V_i)\left.\frac{df_i^{eq}(V)}{dV}\right|_{V=V_0}. \tag{5.43}$$

Importantly, the correction to the resistance can be either positive or negative. Suppose that, as in Fig. 5.21, the channels tend to open in response to increasing voltage, as

most channels do. Then $[df_i^{eq}(V)/dV]_0 > 0$. But if this channel is specific for an ion with a reversal potential above the resting potential ($V_i > V_0$), then opening the channel creates a stronger tendency to pull the voltage toward this higher potential, which is a regenerative effect—a negative resistance. The power supply for this negative resistor is provided by the pumps that maintain the reversal potentials.

If the channels are slow, they make a contribution to the imaginary part of the admittance, along with the capacitance,

$$-i\omega C \to -i\omega C + \frac{1}{-i\omega}\sum_{i\in slow}\frac{g_i N_i}{\tau_i(V_0)}(V_0 - V_i)\frac{df_i^{eq}(V)}{dV}\bigg|_{V=V_0}. \tag{5.44}$$

Again the sign depends on details. If the channels are opened by increasing voltage and the reversal potential is below the resting potential, then their contribution is (almost) like an inductance and can generate a resonance by competing with the capacitance. This resonance is at a frequency

$$\omega_* = \left[\frac{1}{C}\sum_{i\in slow}\frac{g_i N_i}{\tau_i(V_0)}(V_0 - V_i)\frac{df_i^{eq}(V)}{dV}\bigg|_{V=V_0}\right]^{1/2}. \tag{5.45}$$

Problem 96: Equivalent circuits. Equation (5.41) shows that each type of channel contributes a parallel path for current flow through the membrane. The impedance of this path is defined by

$$\frac{1}{\tilde{Z}_i(\omega)} = g_i N_i f_i^{eq}(V) + \frac{g_i N_i (V_0 - V_i)[df_i^{eq}(V)/dV]_0}{1 - i\omega\tau_i(V_0)}. \tag{5.46}$$

Without resorting to the fast/slow approximations above, draw an equivalent circuit using the standard lumped elements (capacitance, resistance, inductance) that realizes this impedance. Show how the parameters of the lumped elements relate to the parameters of the channels.

So, we have seen that even in response to small signals, the dynamics of ion channels generate an interesting complement of electronic parts: resistors, inductors, and negative resistors. Certainly these elements together can make a filter, playing the effective inductance of the channels against the intrinsic capacitance of the membrane, as noted above. The negative resistor can sharpen the resonance and even generate an instability; on the other side of the instability is a genuine oscillator.

Problem 97: Oscillations. Construct a minimal model for ion channels in the cell membrane that supports a stable limit-cycle oscillation of the voltage.

The negative resistance alone means that we can have (without oscillations) an instability of the steady state around which we were expanding, presumably because

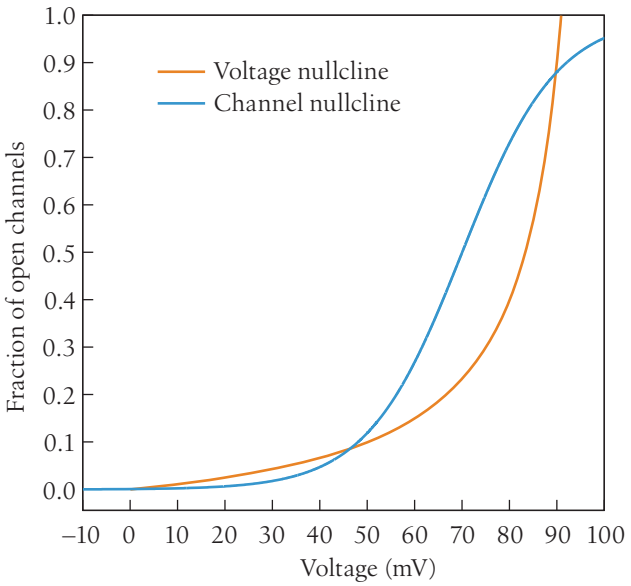

FIGURE 5.22

Bistability in a simple model of a neuron. The channel nullcline is Eq. (5.49), and the voltage nullcline is Eq. (5.50). To be explicit, we choose $f_{eq}(V)$ from Eq. (5.36), with $V_{1/2} = 70$ and $V_w = 10$, and $G_{leak}/gN = 0.1$. Note that there are three crossing points, corresponding to steady states. The low- and high-voltage states are stable; the intermediate-voltage state is unstable.

the real system is multistable. To see this behavior more clearly, consider just two types of channels—a leak channel, which has a total conductance G_{leak} independent of the voltage and has a reversal potential of zero, and some other channel, which opens in response to increasing voltage and has a reversal potential V_r. Then the dynamics are

$$C\frac{dV}{dt} = -G_{leak}V - gNf \cdot (V - V_r), \tag{5.47}$$

$$\frac{df}{dt} = -\frac{1}{\tau(V)}[f - f_{eq}(V)]. \tag{5.48}$$

The steady state solutions are determined by solving two simultaneous equations, usually called the nullclines, obtained by setting the time derivatives equal to zero:

$$f = f_{eq}(V), \tag{5.49}$$

$$V = V_r\frac{f}{f + G_{leak}/gN}; \tag{5.50}$$

these nullclines are shown schematically in Fig. 5.22 for some reasonable choice of parameters. We can see that there are three solutions to the two simultaneous equations; two are stable and one is unstable. The two stable states correspond, roughly, to one state in which all channels are closed and the voltage is zero (the reversal potential of the leak), and one state in which all channels are open and the voltage is near the reversal potential for these channels. The bistability means that, if the cell starts in the low-voltage state, injection of a relatively small, brief current can drive the system across a threshold (separatrix), so that it falls into the high-voltage state after the current pulse is complete. This behavior constitutes a form of memory (interesting, although not very realistic), but it also substantially amplifies the incoming signal, especially if the parameters are tuned so that the difference in voltage to the unstable state is small.

Problem 98: Bistability. Work through a concrete example of the ideas in the previous paragraphs, perhaps using the detailed model from Fig. 5.22. Verify analytically the claims about stability of the three different steady states. Explain how these analytic criteria can be converted into a test for stability of each steady state that can be read off directly from the plots in Fig. 5.22. Analyze the response to brief pulses of current, showing that there is a well-defined threshold for switching from one stable state to the other.

All the different kinds of dynamics we have seen thus far—filtering, oscillation, and bistability—can be generated by just one kind of channel with only two states. Real neurons are more complex. One important class of dynamics that we can't quite see in the simplest models is excitability. In this case, a small pulse again drives the system across a threshold, but what would have been a second stable state is destabilized by relaxation of some other degrees of freedom. The result is that the system takes a long, and often stereotyped, trajectory through its phase space before coming back to its original steady state after the input pulse is over. The action potential is an example of such excitable dynamics.

Our understanding of ion channels goes back to the classic work of Hodgkin and Huxley in the 1940s and 1950s. They studied the giant axon, a single cell, visible to the naked eye, which runs along the length of a squid's body and along which action potentials are propagated to trigger the squid's escape reflex. Passing a conducting wire through the interior of the long axon, they short-circuited the propagation, ensuring that the voltage across the membrane was spatially uniform, as in our idealization above.[8] They then studied the current that flowed in response to steps of voltage. If the picture of channels is correct, then with the voltage held constant, there should be an (Ohmic) flow of current through the open channels. If we step suddenly to a new value of the voltage, Ohm's law states that the current through the open channels will change immediately, but there will be a prolonged time dependence that results from the open or closing of channels as they equilibrate at the new voltage. In the simple model with two states, this changing current should relax exponentially to a new steady state; in particular, the initial slope of the current should be finite.

Hodgkin and Huxley found that the relaxation of the current at constant voltage has a gradual start, as if the channels had not one closed state but several, and the molecules had to go through these states in sequence before opening. They chose to describe these dynamics of the currents by imagining that, for the channel to be open, there were several independent molecular gates that all had to be open. Each gate could have only two states and would obey simple first-order kinetics, but the probability that the channel is open would be the product of the probabilities that the gates were open. In the simple case that the multiple gates are identical, the probability of the channel being open is just a power of the gating variable describing the probability that one gate is open. Hodgkin and Huxley also discovered that at least one important class of channels opens in response to increased voltage and then closes over time. They

8. There is a video of Hodgkin himself (along with a colleague) recreating some of these experiments in the 1970s: http://youtu.be/k48jXzFGMc8. You can also find the great anatomist J. Z. Young dissecting a squid, and taking us along the rather astonishing path to realizing that there are single nerve axons more than 1 mm in diameter: http://youtu.be/pw6_Si5jOpo.

described this mechanism by postulating that in addition to activation gates that were opened by increasing voltage, there were inactivation gates that closed in response to increasing voltage but had slower kinetics. Putting the pieces together, they described the fraction of open channels as

$$f_i = m_i^{\alpha_i} h_i^{\beta_i}, \tag{5.51}$$

where m and h are activation and inactivation gates, respectively, and the powers α and β count the number of these gates that contribute to the opening of one channel. The kinetics are then described by

$$\frac{dm_i}{dt} = -\frac{1}{\tau_i^{(m)}(V)}\left[m_i - m_i^{\text{eq}}(V)\right], \tag{5.52}$$

$$\frac{dh_i}{dt} = -\frac{1}{\tau_i^{(h)}(V)}\left[h_i - h_i^{\text{eq}}(V)\right], \tag{5.53}$$

and finally the voltage (again neglecting spatial variations) obeys

$$C\frac{dV}{dt} = -\sum_i g_i N_i m_i^{\alpha_i} h_i^{\beta_i} \cdot (V - V_i). \tag{5.54}$$

Problem 99: Two gates. Suppose that each channel has two independent structural elements (gates), each of which has two states. Assuming that the two gates are independent of each other, fill in the steps showing that the dynamics of the channels are as described above. In particular, show that after a sudden change in voltage, the fraction of open channels starts to change as $\propto t^2$, not $\propto t$ as expected if the entire channel only has two states.

Problem 100: Hodgkin and Huxley revisited. The original equations written by Hodgkin and Huxley are:[9]

$$C\frac{dV}{dt} = -\bar{g}_L(V - V_L) - \bar{g}_{\text{Na}} m^3 h (V - V_{\text{Na}}) - \bar{g}_{\text{K}} n^4 (V - V_{\text{K}}) + I(t), \tag{5.55}$$

$$\frac{dn}{dt} = \frac{0.01(-V + 10)}{e^{(-V+10)/10} - 1}(1 - n) - 0.125e^{-V/80} n, \tag{5.56}$$

$$\frac{dm}{dt} = \frac{0.1(-V + 25)}{e^{(-V+25)/10} - 1}(1 - m) - 4e^{-V/18} m, \tag{5.57}$$

$$\frac{dh}{dt} = 0.07e^{-V/20}(1 - h) - \frac{1}{e^{(-V+30)/10} + 1} h, \tag{5.58}$$

where Na and K refer to sodium and potassium channels, respectively; time is measured in milliseconds and V is measured in millivolts. These equations are intended to describe a small

9. The only difference from the original paper is that we use the modern sign convention for the voltage. Notice that this original formulation is in terms of a "maximal conductance" for each type of "current," whereas in modern language we could talk about the number of each type of channel. In fact, the more phenomenological description persists, because it corresponds more directly to what is measured, but it allows us to forget that such parameters as $\bar{g}_{\text{K}}$ actually measure the number of copies of a protein that have been inserted into the membrane.

patch of the membrane, and so many parameters are given per unit area: $C = 1\,\mu\text{F/cm}^2$, $\bar{g}_L = 0.3\,\text{mS/cm}^2$, $\bar{g}_{\text{Na}} = 120\,\text{mS/cm}^2$, and $\bar{g}_{\text{K}} = 36\,\text{mS/cm}^2$; the reversal potentials are $V_L = 10.613\,\text{mV}$, $V_{\text{Na}} = 115\,\text{mV}$, and $V_{\text{K}} = -12\,\text{mV}$, all measured from the resting potential.

(a) Rewrite these equations in terms of equilibrium values and relaxation times for the gating variables, for example,

$$\frac{dm}{dt} = -\frac{1}{\tau_m(V)}\left[m - m_{\text{eq}}(V)\right]. \tag{5.59}$$

Plot these quantities. Can you explain, intuitively, the form of the curves?

(b) Simulate the dynamics of the Hodgkin-Huxley equations in response to constant current inputs. Show that there is a threshold current above which the system generates periodic pulses. Explore the frequency of the pulses as a function of current.

(c) Suppose that the injected current consists of a mean (less than the threshold you identified in part (b)) plus a small component at frequency ω. By some appropriate combination of analytic and numerical methods, find the impedance $Z(\omega)$ for different values of the mean injected current. Show that the membrane has a resonance, and explore what happens to this resonance as the mean current is increased toward threshold. How do your results connect to the frequency of pulses above threshold?

(d) Real axons are essentially long thin cylinders. Show that, if we allow the voltage to vary along the length of the axon, a current per unit area should flow across the membrane of

$$I = \frac{a}{2R}\frac{\partial^2 V}{\partial z^2}, \tag{5.60}$$

where z is the coordinate along the cylinder, a is its radius, and R is the resistivity of the fluid filling the axon, assuming that resistance outside the axon is negligible. For the squid giant axon, $a \sim 250\,\mu\text{m}$ and $R \sim 35\,\text{ohm}\cdot\text{cm}$. Use this result to write equations for the voltage and gating variables along the axon. Note that only the equation for the voltage has spatial derivatives. Why?

(e) Simulate the response of a long segment of the axon to a current pulse injected at one end. Show that small pulses result in spatially restricted voltage responses, whereas larger inputs produce a propagating pulse. Confirm that these pulses become more stereotyped as they propagate and have a velocity that is independent of the input current. What is this velocity? How does it compare to the observed speed of action potentials, $v \sim 20\,\text{m/s}$?

Problem 101: Simplification. It is difficult to make analytic progress in understanding the dynamics of a system with five variables. There is a history of trying to approximate the Hodgkin-Huxley model by exploiting the fact that the different variables have very different time scales. See how far you can go along this path. I have left this problem deliberately open ended. For some approaches, see the references in the Annotated Bibliography for this section.

It is good to pause here and review how we know that the Hodgkin-Huxley description of ion channels is correct. The initial triumph, which you are asked to reproduce in problem 100, is the prediction of the propagating action potential itself, as in Fig. 5.23, with the correct speed. The model also predicts that, as the action potential passes, there is a net flux of potassium and sodium across the membrane. On long time scales,

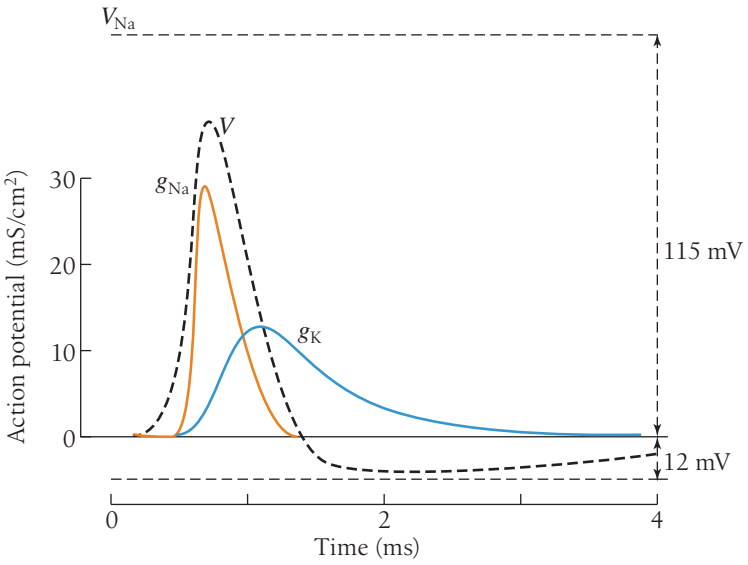

FIGURE 5.23

The action potential that emerges from the Hodgkin-Huxley model. The voltage trace is shown by the heavy dashed line, and the time dependence of the underlying conductances for sodium and potassium are shown by solid lines. Note that the rise of the action potential is associated with the opening of sodium channels, whereas the fall is a combination of these channels closing and the opening of potassium channels. The dynamic range of the electrical signal is bounded by the reversal potentials for these two ions (V_{Na} and V_K), as shown.

this flux must be balanced by the action of pumps that maintain the concentration differences between the inside and outside of the cell. But either by looking quickly or by poisoning the pumps, one should be able to detect the flux (e.g., using radioactive tracers); this works, quantitatively.

Nature provides a variety of toxins that block the action potential in different ways,[10] and we can also use artificial blockers, for example, by using ions with very large radius that can literally plug the hole in open channels. It is striking that these agents act selectively on different channels, and one can verify that this way of isolating the dynamics of sodium and potassium channels matches the Hodgkin-Huxley description. If we can arrange for the channels to open but be blocked, then the structural change of the channel molecule upon opening should still move the gating charge across the membrane. In addition, if we are careful, this movement should be measurable essentially as a delayed capacitive response to changes in the applied voltage. These gating currents have indeed been detected, and the magnitude of the gating current matches quantitatively what is predicted from the voltage dependence of the activation curve. In some cases this agreement can be extended to genetically engineered channels, where one can show that changes in the activation curve and gating currents track one another.

If individual channels are independent of one another, then their opening and closing events should be independent. If we look at a small patch of the membrane, there will not be that many channels present, and we might be able to see that the discrete events in the individual molecules do not quite average out—there should be noise from the random opening and closing of the single channels. This channel noise has been detected and has the spectral properties expected from the Hodgkin-Huxley model. Finally, if we look at even smaller patches of the membrane, and have

10. The most famous of these toxins might be tetrodotoxin, produced by the puffer fish. This molecule blocks the sodium channel and hence eliminates action potentials. It is worth remembering that these toxins serve a positive function for the organisms that produce them.

proportionately more sensitive amplifiers, we should be able to see the opening and closing of single channels. Again, this has been done. Most importantly, we can look at the distribution of times that individual channels spend in the open and closed states and connect this distribution to the kinetics predicted by the Hodgkin-Huxley model and its generalizations. Although these more-detailed measurements have revealed new features of channel kinetics even in well-studied examples, in outline the picture given to us by Hodgkin and Huxley has stood the test of time.

Now that we have confidence in our mathematical description of neurons, it is time to appreciate just how many parameters are involved. A typical cell expresses eight or nine different kinds of channels. Each channel is described by the dynamics of two gating variables. If activation and inactivation curves have a simple sigmoidal form as in Fig. 5.22, then there are two parameters for each such curve—the voltage at half activation and the slope or width—and at least one more parameter to set the time scale of the kinetics. Finally, there is a parameter to count the total number of channels, or equivalently the maximal conductance achieved if all channels are open. All together, then, this is ∼7 parameters per channel type, or roughly 50 parameters for the entire neuron, conservatively. Importantly, to a large extent the cell actually has control over these parameters and, in a meaningful sense, can adjust them almost continuously.

How do these adjustments occur? Most obviously, the total number of open channels is controlled in the same way that all other protein copy numbers are controlled. Sometimes, because of a clear connection to experiment, one speaks of the maximal conductance associated with a particular type of channel ($G_i^{max} = g_i N_i$), but this terminology obscures the fact that this parameter really is the total number of copies of the protein that the cell has expressed and inserted into the membrane. The parameters of the activation curves and the time constants are intrinsic properties of the proteins, but they too can be adjusted in several ways. Like all proteins, ion channels can be covalently modified by phosphorylation or other actions. More importantly, the genome encodes a huge number of different ion-channel proteins; the human genome has 90 different potassium channels alone. Although they do form classes based on their dynamics, considerable variation exists within classes, and because many of these genes have multiple alternative splicings, there is the potential for almost continuous parameter variation. These different mechanisms of variation interact; as an example, different splicing variants can exhibit different sensitivities to phosphorylation.

Problem 102: Continuous adjustment of electrical dynamics. To illustrate the possibility of nearly continuous adjustments in the electrical dynamics of neurons, consider the case of the hair cells in the turtle ear. In these cells (see Section 4.5), one contribution to frequency selectivity comes from a resonance in the electrical response of the hair cell itself. This resonance is driven by a combination of voltage-gated calcium channels and calcium-activated potassium channels. A detailed model of this system is described by Wu and Fettiplace (2001). Try to understand what they have done, and reproduce the essential theoretical results. In particular, what is the role of "details" (e.g., the building of channels out of combinations of different subunits) in generating the correct qualitative behavior?

One well-studied example of channel dynamics is in the stomatogastric ganglion of crabs and lobsters, schematized in Fig. 5.24. This network of ∼30 neurons generates a

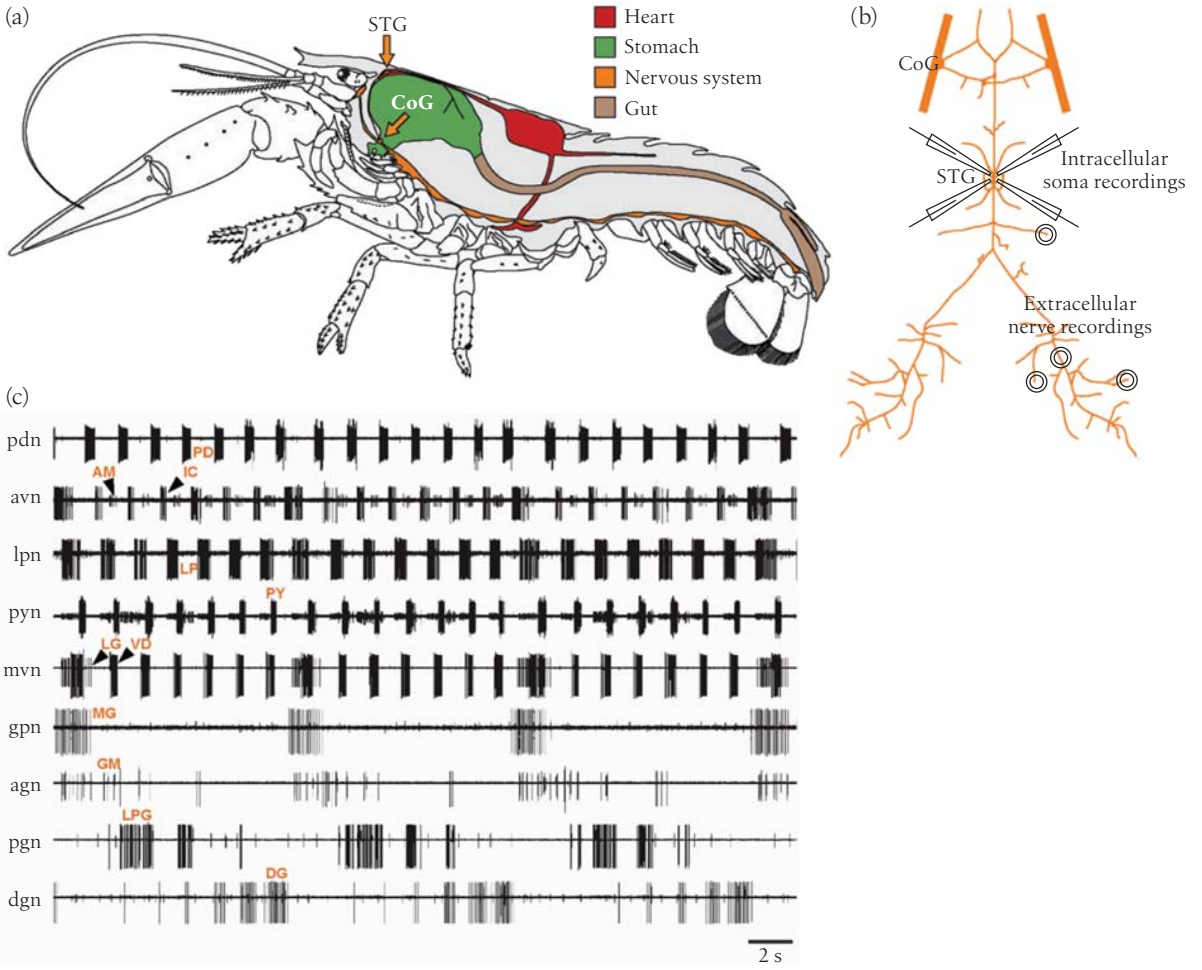

FIGURE 5.24

The stomatogastric ganglion (STG) in crustaceans. Reprinted with permission of Annual Reviews, Inc., from Marder and Bucher (2007); permission conveyed through Copyright Clearance Center, Inc. (a) The location of the STG and the commissural ganglion (CoG) in a lobster. (b) Schematic of the ganglion dissected out of the animal, and the opportunities for recording the activity of the neurons. (c) Simultaneous extracellular recordings from nine motor nerves at the output of this network. Names indicate particular neurons that can be identified in each individual (as with the named neurons in the fly visual system discussed in Section 4.4), and in some cases (e.g., avn, mvn) we can identify spikes from several individual neurons in the recording from one nerve. There are two main rhythms, the faster pyloric rhythm in cells PD, LP, PY, VD and IC, and the slower gastric-mill rhythm in cells MG, DG, GM, LPG, and LG.

rhythm, and this rhythm in turn drives muscles that actuate teeth in the crab stomach, grinding its food. Evidently getting the correct rhythm is important in the life of the organism. Records of the electrical signals from individual neurons show that several of the cells produce periodic bursts of action potentials, and a handful of cells are pacemakers that can generate this periodic pattern without input from the other cells. In one such cell (the lateral pyloric neuron), experiments show that there are seven different channel types. An important feature of this cell, shared by many other cells, is the presence of voltage-gated calcium channels. As action potentials occur, they trigger

calcium flux into the cell. Because some channels are also directly affected by the calcium concentration, a complete model must include a description of the calcium buffering or pumping that counterbalances this flux.

It is worth being explicit about all these ingredients in the dynamics of the lateral pyloric neuron, not least to get a sense for the state of the art in such analyses. This will, however, take us to a level of detail that I have largely tried to avoid until this point in the text. It is essential, however, because this level of detail is where the problems become apparent. Once we have identified the problem, we can zoom back out to a more schematic view.

As before, we neglect the spatial structure of the cell, so there is just one relevant voltage difference V between the inside and outside of the cell, which obeys Eq. (5.54),

$$C\frac{dV}{dt} = -\sum_{\mathrm{i}} g_{\mathrm{i}} N_{\mathrm{i}} m_{\mathrm{i}}^{\alpha_{\mathrm{i}}} h_{\mathrm{i}}^{\beta_{\mathrm{i}}} (V - E_{\mathrm{i}}) + I_{\mathrm{ext}}, \tag{5.61}$$

where I_{ext} is any externally injected current, and E_{i} is the reversal potential for channel type i. The kinetics of the gating variables m_{i} and h_{i} are governed by Eq. (5.52) and Eq. (5.53), respectively. For most channels, we can take the equilibrium values of the gating variables to be given by the generalization of Eq. (5.36),

$$m_{\mathrm{i}}^{\mathrm{eq}}(V) = \frac{1}{1 + \exp[-(V - V_{1/2}^{m_{\mathrm{i}}})/V_w^{m_{\mathrm{i}}}]}, \tag{5.62}$$

$$h_{\mathrm{i}}^{\mathrm{eq}}(V) = \frac{1}{1 + \exp[-(V - V_{1/2}^{h_{\mathrm{i}}})/V_w^{h_{\mathrm{i}}}]}, \tag{5.63}$$

and the time constants for relaxation of the gating variables are phenomenologically,

$$\frac{1}{\tau_{\mathrm{i}}^{(m)}(V)} = \frac{k_{\mathrm{i}}^{(m)}}{1 + \exp[-\gamma_{\mathrm{i}}^{(m)}(V - V_{\mathrm{i}}^{(m)})]}, \tag{5.64}$$

$$\frac{1}{\tau_{\mathrm{i}}^{(h)}(V)} = \frac{k_{\mathrm{i}}^{(h)}}{1 + \exp[-\gamma_{\mathrm{i}}^{(h)}(V - V_{\mathrm{i}}^{(h)})]}. \tag{5.65}$$

As shown in Table 5.1, this description works for several channel types, one selective for potassium (the delayed rectifier), two for calcium, and one mixed (the inward rectifier), plus a leak that exhibits no significant time or voltage dependence of its conductance.

Two of the important channel types allow calcium to flow into the cell. As we will see, this current is big enough to change the concentration of calcium inside the cell, which has a variety of effects on other processes, including one of the channels that does not fit the simple description we have given so far. So, we need to describe the dynamics of the calcium concentration itself. The simplest model is that the calcium relaxes back to some internally determined steady state, $[\mathrm{Ca}]_0 = 0.05\ \mu\mathrm{M}$, with a rate $k_{\mathrm{Ca}} = 360\ \mathrm{s}^{-1}$, and the current through the open calcium channels is driving an increase in the intracellular calcium concentration. In this case,

$$\frac{d[\mathrm{Ca}]}{dt} = -k_{\mathrm{Ca}}\left([\mathrm{Ca}] - [\mathrm{Ca}]_0\right) + AI_{\mathrm{Ca}}, \tag{5.66}$$

TABLE 5.1

Subset of channels in the lateral pyloric neuron

Channel type	$g_i N_i$ (μS)	E_i (mV)	Midpoint (mV)	Width (mV)	Rate (s^{-1})
i = 1: delayed rectifier	0.35	$E_K = -80$			
Activation equilibrium ($\alpha_1 = 4$)			$V_{1/2}^{m_1} = -25$	$V_w^{m_1} = 17$	
Activation kinetics			$V_1^{(m)} = 10$	$1/\gamma_1^{(m)} = 22$	$k_1^{(m)} = 180$
i = 2: Ca^{++} current 1	0.21	E_{Ca}			
Activation ($\alpha_2 = 1$)			$V_{1/2}^{m_2} = -11$	$V_w^{m_2} = 7$	50
Inactivation ($\beta_2 = 1$)			$V_{1/2}^{h_2} = -50$	$V_w^{h_2} = -8$	16
i = 3: Ca^{++} current 2	0.047	E_{Ca}			
Activation ($\alpha_3 = 1$)			$V_{1/2}^{m_3} = -22$	$V_w^{m_3} = 7$	10
i = 4: inward rectifier	0.037	−10			
Activation equilibrium ($\alpha_4 = 1$)			$V_{1/2}^{m_4} = -70$	$V_w^{m_4} = -7$	
Activation kinetics			$V_4^{(m)} = -110$	$1/\gamma_1^{(m)} = 13$	$k_1^{(m)} = 0.33$
i = 5: leak	0.1	-50			
i = 6: A-current	2.2	$E_K = -80$			
Activation equilibrium ($\alpha_6 = 3$)			$V_{1/2}^{m_6} = -12$	$V_w^{m_6} = 26$	
Activation kinetics					$k_6^{(m)} = 140$
Inactivation equilibrium ($\beta_{6a} = 1$)			$V_{1/2}^{h_{6a}} = -62$	$V_w^{h_{6a}} = 6$	
Inactivation kinetics					$k_{6a}^{(h)} = 50$
Inactivation equilibrium ($\beta_{6b} = 1$)			$V_{1/2}^{h_{6b}} = -40$	$V_w^{h_{6b}} = -12$	
Inactivation kinetics					$k_{6b}^{(h)} = 3.6$

For the delayed rectifier and the second type of calcium channel, there is no evidence for inactivation. The negative value of $V_w^{(h_2)}$ means, from Eq. (5.36), that the probability of the inactivation gate being open decreases with increasing voltage. For calcium channels, the reversal potential varies, depending on the calcium concentration inside the cell, as in Eq. (5.67), and the relaxation times do not have a detectable voltage dependence. The voltage dependence of the inward rectifier kinetics is opposite to Eq. (5.64), that is, $1/\tau \propto 1 + \exp[-\gamma_i^{(m)}(V - V_i^{(m)})]$. The leak current, by convention, is the current that exhibits no voltage or time dependence of its conductance. From Buchholtz et al. (1992).

where I_{Ca} is the total calcium current ($I_{Ca} = I_2 + I_3$ from Table 5.1). The constant $A = 300\ \mu$M/nC is inversely proportional to the volume into which the current flows, which experimentally comes out to be much smaller than the total volume of the cell body. As the concentration of calcium changes, the reversal potential for the calcium currents also changes,

$$E_{Ca} = \frac{k_B T}{2e} \ln\left(\frac{[Ca]_{out}}{[Ca]}\right), \tag{5.67}$$

where the calcium concentration outside the cell is $[Ca]_{out} = 13$ mM.

We are still missing three channel types in this cell. First, there is another potassium channel that is almost described by our standard model, but the inactivation seems to

involve two processes that occur on different time scales. This behavior can be captured by replacing

$$h_6 \to x(V)h_{6a} + [1 - x(V)]h_{6b}, \tag{5.68}$$

where the weighting function is

$$x(V) = \frac{1}{1 + \exp[-(V - 7)/15]}, \tag{5.69}$$

with V measured in mV, as before.

Second, there is a fast sodium channel not unlike the ones that Hodgkin and Huxley found in the squid giant axon, with $\alpha_7 = 3$ and $\beta_7 = 1$. The activation is sufficiently fast that it can be approximated as instantaneous, so that m_7 is always at its equilibrium value, which varies with voltage in a slightly more complicated way than for the other channels,

$$m_7 = m_7^{\text{eq}}(V) = \frac{1}{1 + 136\exp[-(V + 34)/13]\left(1 - \exp[-(V + 6)/20]\right)/(V + 6)}, \tag{5.70}$$

where V again is measured in mV. The inactivation gates obey

$$\frac{dh_7}{dt} = a_7(V)(1 - h_7) - b_7(V)h_7, \tag{5.71}$$

where the rates,

$$a_7(V) = 40\exp[-(V + 39)/8], \tag{5.72}$$

$$b_7(V) = \frac{500}{1 + \exp[-(V + 40)/5]}, \tag{5.73}$$

are measured in s^{-1}. The total conductance contributed by these channels is large, $g_7 N_7 = 2300\ \mu\text{S}$, although they are only open briefly.

The last type of channel, like the first two in Table 5.1, is selective for potassium ions, but the probability of the channel being open is modulated by the intracellular calcium concentration. This channel has $\alpha_8 = \beta_8 = 1$, and the equilibrium state of the inactivation gate depends only on the calcium concentration:

$$h_8^{\text{eq}} = \frac{1}{1 + [\text{Ca}]/(0.6\ \mu\text{M})}. \tag{5.74}$$

The equilibrium state of the activation gate, in contrast, depends both on voltage and on calcium,

$$m_8^{\text{eq}} = \frac{1}{1 + \exp[-(V + f[\text{Ca}])/23]} \cdot \frac{1}{1 + \exp[-(V + 16 + f[\text{Ca}])/5]} \cdot \frac{[\text{Ca}]}{2.5\ \mu\text{M} + [\text{Ca}]}, \tag{5.75}$$

where $f = 0.6\ \text{mV}/\mu\text{M}$. The relaxation rates, $k_8^{(m)} = 600\ \text{s}^{-1}$ and $k_8^{(h)} = 35\ \text{s}^{-1}$, show little if any voltage dependence. This seems like a complicated model, but it fits the experimental results very well, as shown in Fig. 5.25.

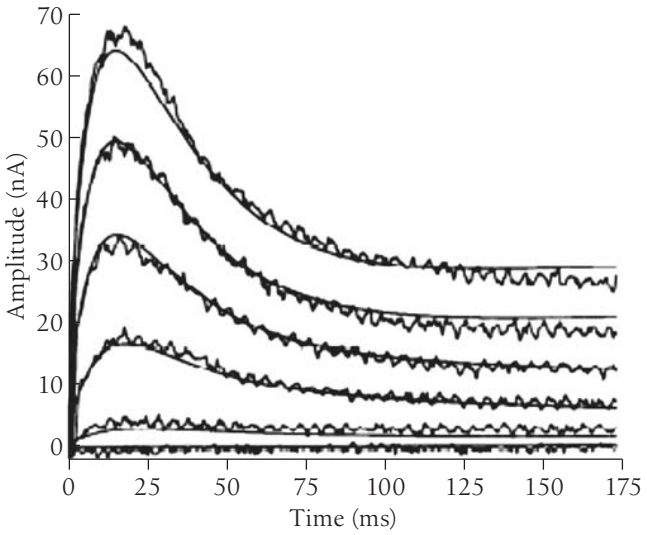

FIGURE 5.25

Dynamics of the calcium-dependent potassium current in response to voltage steps of varying size. Redrawn from Buchholtz et al. (1992). Experimental data (noisy traces) from Golowasch and Marder (1992); solid lines are from the model including Eq. (5.74) and Eq. (5.75).

Problem 103: Calcium-dependent potassium conductances. Develop a microscopic picture to explain the combination of voltage and calcium dependences seen in Eq. (5.74) and Eq. (5.75). Remember that these equations describe the equilibrium fractions of molecules in particular states, so you need to relate these fractions back to the free energies of the different states. Connect your discussion with the MWC models introduced in the discussion of the cGMP gated channels in rod cells (Section 2.3).

The model of the lateral pyloric neuron described here represents the culmination of many years of effort, both in experiments on this particular system and in the exploration of these fully realistic generalizations of the Hodgkin-Huxley model to what seems the more typical case, with many different channel types functioning together. This model also represents a level of detail and complexity that I have tried to avoid so far, so some explanation is called for. First, the complexity consists largely of variations on a theme. Many channels are known to be described by the general picture of multiple activation and inactivation gates, so this picture provides a framework in which each new type of channel can be fit. Second, the complexity is justified by a large body of data. Independent experiments have been done on other systems, exploring quantitatively each type of channel seen in this neuron, and detailed experiments on this one cell have teased out the contributions of each channel type.

Problem 104: Justifying complexity. Go through Buchholtz et al. (1992), Golowasch and Marder (1992), and Golowasch et al. (1992), and explain the justification for each channel type in the model discussed above.

Indeed, the program of describing the electrical dynamics of single neurons in terms of generalized Hodgkin-Huxley models, usually with many different channel types functioning together, became a small industry. It really worked. In some cases one could go so far as to characterize the kinetics of particular channel types through measure-

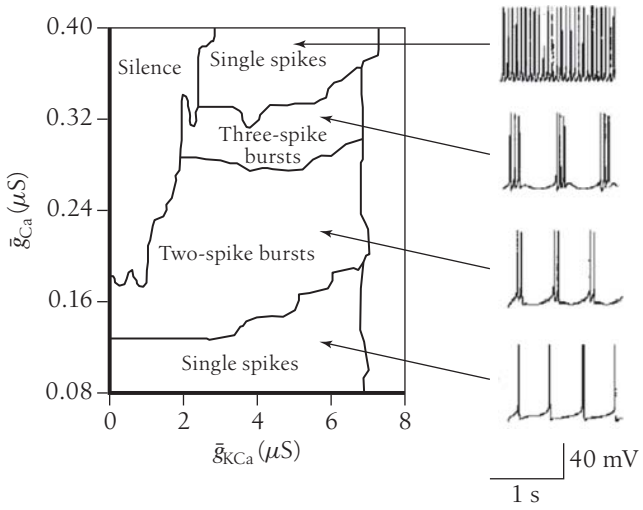

FIGURE 5.26

Simulations of a detailed model, with seven types of channel, for the lateral pyloric neuron in the stomatogastric ganglion of the crab. Changes in the pattern of activity are shown as a function of the numbers of two different kinds of channel, where channel number is expressed as the maximal conductance when all channels are open ($\bar{g}_{Ca}$ and $\bar{g}_{KCa}$). Note that relatively small changes in these parameters can result in both quantitative and qualitative changes in the pattern of electrical activity, running the full range from silence to single spike firing to bursting. Redrawn from Le Masson et al. (1993).

ments on single molecules and then put these single-molecule properties together to reproduce the functional behavior of the cell as a whole. This body of work is beautiful and implements what many people would like to do in other systems, building from measured properties of individual molecular events up to macroscopic biological function. As emphasized above, we can think of the ion channels in the cell membrane as a network of interacting proteins, where the interaction is mediated by the voltage across the membrane rather than direct protein-protein encounters, and where the equations for the dynamics of individual channels have a firm foundation. It is not unreasonable to claim that ion channels in the cell membrane are in fact the best-understood examples of biochemical networks, although the language typically used in describing these systems obscures this connection.

Despite their success, it came to be known, though not widely commented on, that these models of coupled ion-channel dynamics had a problem. Although experiments often characterize the activation curves and kinetics of the individual channels, it is hard to make independent measurements of the total number of channels (or equivalently, the maximum conductance when all channels are open). Thus, one is left adjusting these parameters, trying to fit the overall electrical dynamics of the neuron—for example, the rhythmic bursting of the pyloric neuron. This fitting turns out to be delicate; as one adjusts the (many) parameters, one finds bifurcations to qualitatively different behaviors in response to relatively small changes. An example of this behavior is shown in one two-dimensional slice through the seven-dimensional space of channel numbers in the pyloric model (Fig. 5.26).

From a physicist's point of view, this all seems a mess. There are many details one has to keep track of and many parameters to adjust.[11] One might be tempted just to walk away, and count this model as a part of biology we don't want to know about. But

11. As in the case of kinetic proofreading, I think there is a tendency to remember the original papers as having proposed mechanisms that solve problems. But in many ways, it was a much deeper contribution to formulate the problems. Even if the solutions turn out not to be precisely the ones chosen by Nature, the problems are important.

FIGURE 5.27

Mean calcium concentration follows the pattern of electrical activity. Main figure gives a coarse map of mean calcium concentration as a function of the same two variables shown in Fig. 5.26. Regions are labeled A, B, . . . , E in order of increasing mean concentration. Region A corresponds to near-zero concentration, region B to $c < 0.1\,\mu$M, up to region E in which $0.4 < c < 0.5\,\mu$M. Small figure at right shows that the region of bursting activity corresponds almost perfectly to the region of parameter space in which the mean calcium concentration is between 0.1 and $0.3\,\mu$M, so that holding the calcium level fixed will stabilize bursting. Redrawn from Le Masson et al. (1993).

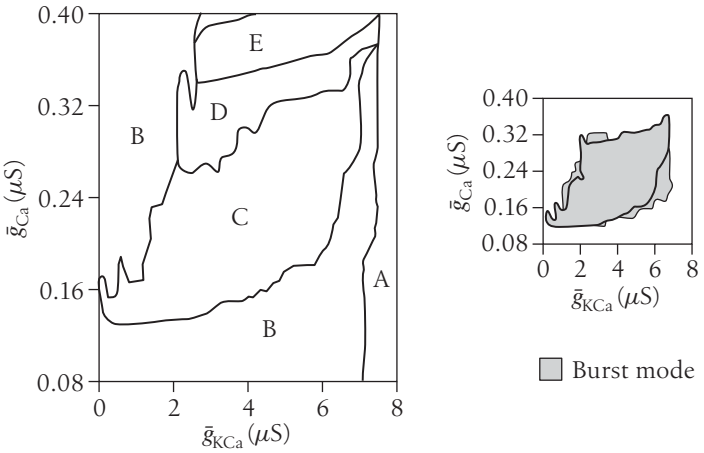

there is a deep question here. If we have trouble adjusting the parameters of models to reproduce the observed functional behaviors of particular cells, how do the cells themselves adjust these parameters to achieve their correct functions? How does a cell choose the correct number of each type of channel to express? One could imagine that the cell has some sort of lookup table—I am a cell of type α, so I should express N_1^α molecules of channel type 1, N_{37}^α molecules of channel type 37, and so on. This is a bit implausible. More likely would be that the cell has some way of monitoring its activity, asking "how close am I to doing the right thing?", and generating an error signal that could be used to drive changes in the expression of the channels or perhaps their insertion into the membrane.

How can a neuron "know" whether it is exhibiting the desired pattern of electrical activity? It would need some signal that couples voltage changes across the membrane, which are quite fast, to the biochemical events regulating gene expression, which are quite slow. One idea is to use the intracellular calcium concentration as an intermediary. We know that many cellular processes are regulated by calcium, so one end of this mechanicsm is easy to imagine. But in the models described above the calcium concentration is an explicit part of the dynamics, so we can calculate, for example, the time-averaged calcium concentration as function of the parameters of the model. What we see in Fig. 5.27 is that $[Ca^{++}]$ does an excellent job of tracing the pattern of electrical activity in this cell. Thus, if the system needs to stabilize a pattern of rhythmic bursting, it can do so by feedback mechanisms that try to hold the calcium concentration near a target value of $C_0 \sim 0.2\,\mu$M.

Let us suppose that the expression of each channel protein is regulated by calcium, so that

$$\tau_i \frac{dN_i}{dt} = N_i^{max} f_i([Ca^{++}]/C_0) - N_i, \tag{5.76}$$

where $f_i(x)$ is a sigmoidal function, such as

$$f_i(x) = \frac{1}{1 + x^{\pm n}}. \tag{5.77}$$

These equations have their steady state at $N_i = N_i^{max} f_i([Ca^{++}]/C_0)$, but the calcium concentration must be determined self-consistently through the full dynamics of the channels and voltage. We should choose the signs of the calcium dependences to ensure stability: channels that allow excitatory currents to flow will tend to drive increases in $[Ca^{++}]$, and so they should be opposed by a decreasing function $f_i(x)$, whereas channels that allow inhibitory currents to flow should be controlled by an increasing function $f_i(x)$. Once the signs are chosen, if the regulation functions are steep (large value of n in Eq. (5.77)), and the maximum possible numbers of channels (N_i^{max}) are large, then the dynamics will always be pulled into regimes where $[Ca^{++}] \approx C_0$.

Problem 105: A simple example of a self-tuning neuron. Imagine a neuron with three types of channels. Two of these are always open, and they have different reversal potentials V_+ and V_-; there are N_+ and N_- copies of these two types of channel, and they have the same single-channel conductance. Further, there are channels that allow calcium to flow into the cell, and these channels have a probability of opening that depends on voltage. Calcium flow into the cell is opposed by a pump that would cause the internal calcium concentration to relax exponentially if there were no influx.

(a) Write out the equation for the dynamics of the voltage in the cell. In the approximation that calcium currents have a negligible impact on the voltage, show that there is a steady state in which the voltage is a weighted average of the two reversal potentials.

(b) Write out equations for the dynamics of the internal calcium concentration and the probability of the calcium channel being open. Feel free to hypothesize some simple activation curve for the channel, or look in the literature for inspiration. Introduce calcium-dependent dynamics for the $N_\pm$ as in Eq. (5.76) and Eq. (5.77). For simplicity, assume that the time constants for relaxation of the channel numbers are the same for the two channels, as are the maximal expression levels.

(c) What does your model predict? In particular, you have the choice of signs in Eq. (5.77); try the case where N_+ has the $+$ in this equation, and conversely for N_-. Is there a stable steady state voltage, even when the numbers of channels are allowed to fluctuate? How does this stability depend on the reversal potentials $V_\pm$?

This example, where the voltage is stabilized at a steady state, is quite simple and misses some crucial features of the dynamics of real neurons. Still, you can explore generalizations. Can you, for example, solve a similar model with many different kinds of channels, still neglecting their gating, and show that there is again a stable steady state voltage if the signs of the calcium regulation functions in Eq. (5.77) are chosen in correct relation to the reversal potential? How important are details, such as the assumption that the maximum numbers of the different channels are the same? This is a deliberately open-ended problem; you should try to see how far you can go without having to do simulations, which seem essential once we include realistic channel dynamics.

How can we tell whether something like this sort of self-tuning really is happening? If neurons "knew" how many of each kind of channel to make, then they would try to make this number no matter what the conditions were. For example, inputs from other neurons would drive changes in the electrical activity but not changes in channel expression. However, if the cell is maintaining some mean calcium concentration, or some other measure of activity, then changing the environment in which the neuron

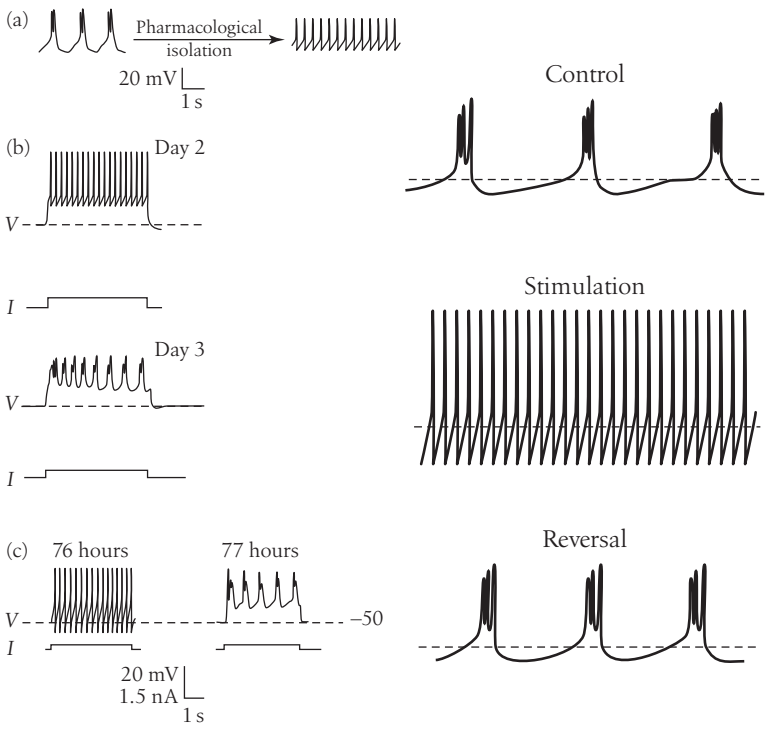

FIGURE 5.28

Changing intrinsic properties of the stomatogastric ganglion neurons. At left, an experiment in which one cell is ripped from the network and placed in isolation. At first (top) the electrical activity shifts from rhythmic bursts to repeated ("tonic") firing of single action potentials. After 2 days in culture, the cell is silent but responds to small positive currents with tonic firing; after 3 days the response consists of bursts not unlike those in the native network environment. At bottom, continuous recordings demonstrate that this switch from tonic firing to bursting can occur within an hour. At right, 1 hr of stimulation with negative current pulses drives a shift from bursting to tonic firing, which is reversed after 1 hr of no stimulation. All these changes in activity reflect changes in the numbers of different types of ion channels in the cell membrane, as predicted from the models discussed in the text. From Turrigiano et al. (1994). Reprinted with permission from AAAS.

operates will change channel expression. As an extreme example, if we rip the neuron from its network and put it in a dish, the normal pattern of rhythmic bursting will go (wildly) wrong, but the calcium-sensitive dynamics of the channel expression levels will eventually bring the system back into something close to the original pattern. In this new state, the channels are playing different roles in the dynamics, because the driving forces for ionic current flow are different, but the final pattern of activity is the same. A literal version of this rather dramatic scenario actually works experimentally, as shown in Fig. 5.28.

We have noted already that, in invertebrates, such as flies and crabs, the structure and function of neurons in many circuits is sufficiently reproducible from individual to individual within a species that it makes sense to give these cells names and numbers. This discussion of stabilizing patterns of activity rather than expression levels suggests that such reproducibility of function can be achieved without exactly reproducing the number of copies of each channel protein. Further, although the slice through parameter space shown in Fig. 5.27 suggests that the region compatible with normal function is convex, this geometry in fact is not general, and real models often have banana-shaped volumes in parameter space that are consistent with particular patterns of electrical activity. Again, such results are consistent with experiments, most impressively in experiments that measure directly the number of copies of mRNA for several channel types in single cells.

You might worry that we have replaced the tuning of channel copy numbers with a fine tuning of the regulatory mechanisms on all channels. In fact, it is not plausible that calcium acts directly on expression of genes. It is more likely that calcium binds

to some protein, and when its binding sites are occupied, the protein can act, directly or indirectly, as a transcription factor. Then the fact that all the genes have the same calcium dependence of their steady state values reflects their common regulation by the same calcium-binding protein. Exploring this scenario in more detail, we see that the kinetics of binding and unbinding of calcium to the sensitive protein can span the time scales of action potentials, bursts, and even the basic rhythm itself. By combining signals from calcium-binding proteins with different kinetics, the more subtle details in the pattern of electrical activity can be stabilized.

A model that explains the behavior of cells only when parameters are finely tuned provokes suspicion that we are missing something. One possibility—often the most plausible—is that the model simply is wrong. The models that we have for biological systems are not like the Navier-Stokes equations for fluids or the standard model of particle physics; we have many reasons to suspect that we are simply solving the wrong equations. But the electrical dynamics of neurons are a special case. Our mathematical models of channel dynamics emerged as accurate summaries of a huge body of data and are nearly exact on the time scales that are experimentally accessible; it is for this reason that we have gone into rather more detail here than in other sections of the text. Rather than rejecting the models, we therefore must conclude that we are missing something, presumably on time scales longer than the experiments that are used to characterize the channel kinetics. In particular, what look like constant parameters must become dynamical variables on long time scales. The simplest implementation of this idea seems to work, and it generates several dramatic experimental predictions that have subsequently been confirmed. Indeed, this theoretical work on the problem of parameter determination has launched a whole subfield of experimental neurobiology, investigating the activity-dependent regulation of the intrinsic electrical properties of neurons.

5.3 The States of Cells

Cells have internal states. Sometimes these states are expressed in a very obvious way, even to external observers, as when we see the alternating black and white stripes of a zebra. In other cases, the states are hidden, as when a neuron stops responding to a constant external stimulus but then rebounds when the stimulus is removed; the amplitude of the rebound reflects the initial amplitude of the stimulus, which must have been stored in some internal state, separate from the output. In these two examples, we also see that these internal states can be (approximately) discrete or continuous. In many cases, the states of cells are known to be encoded by the concentrations of particular, identifiable molecules, and these concentrations in turn reflect a balance of multiple kinetic processes. If we try to transcribe these qualitative ideas simply into quantitative models, we will find that the states of cells depend on parameters. Most obviously, these states will depend on absolute concentrations, and there is a widespread suspicion that absolute concentrations are highly variable, making them poor candidates for the markers of cellular state. More generally, it would seem that, unless we are careful, states will depend sensitively on parameters, providing another example of the problem of fine tuning versus robustness that we have been discussing.

When you tie your shoes in the morning, you can feel the pressure against the skin of your foot, but very quickly this sensation dissipates. When you step outside on a