Chapter 8

Netlet and Population Modeling

§ 1. Characterizing Proxy Neuron Models

In chapter 7, the idea of a proxy neuron model representing a population of physiological neurons was introduced. The key assumption behind the idea of a proxy neuron is the idea that it is meaningful to speak of an "average neuron" within the populace of neurons. Like the famous "average American family with 2.3 children," the "average neuron" is, to borrow one of William James' famous put-downs, "as mythical an entity as the jack of spades." What are we to mean by "average neuron," and how does one go about coming up with one?

One logical starting point is to ask: What do all neurons do so far as their basic signal processing behavior is concerned? and can this commonality be meaningfully represented by some sort of average descriptive function or functions? One rather obvious answer to the first question is: Neurons respond to synaptic stimuli, and if the stimulus is sufficiently large, and if the neuron communicates with other neurons via chemical synapses, the neuron fires an action potential. If, for now, we ignore non-spiking neurons that communicate via gap junctions, it seems reasonable to pick the neuron's responses to chemical synapse stimuli as a likely starting point. One of the most common measures of neuronal response is the peak amplitude of its excitatory postsynaptic potential (EPSP) in response to single-event synaptic stimulation.

Studies detailing distributions of EPSP responses to synaptic stimulation are not common, but

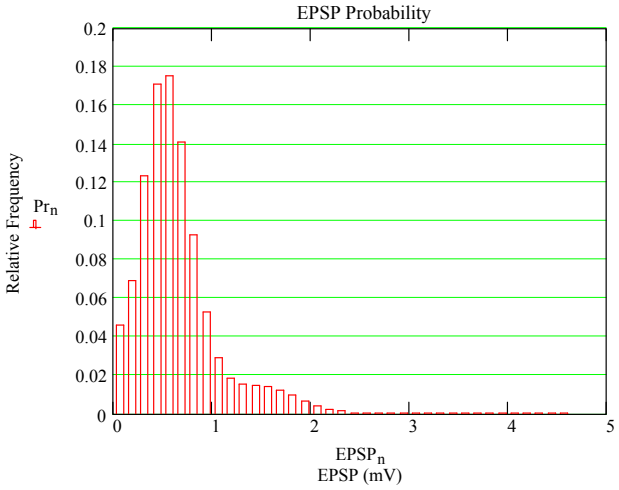

Figure 8.1: Model distribution of excitatory postsynaptic potentials (EPSPs) to single-event synaptic stimulation. The histogram bins are 0.125 mV wide. The model is a tri-modal distribution composed of three weighted binomial distributions. See the text for discussion of this model.

there have been a few reported findings, e.g. [SMET]. Taking Smetters' findings on layer V pyramidal cell neurons as representative, the first noticeable feature of EPSP distribution is that it tends to be skewed toward the smaller EPSP values, with occasional low-probability occurrences of much larger EPSPs. The shape of the distribution is not well-fitted by standard probability functions such as the binomial or Poisson distribution functions, suggesting that we are dealing with a multimodal population of responses.

Figure 8.1 shows a model distribution comprised of three separate populations, each modeled as a binomial distribution, $P(n, N; p)$. The model assumes a "unit EPSP" response such that each bin in the histogram represents a peak EPSP range from $0.125n$ to $0.125(n + 1)$ mV with the bin centered on $0.125n + 0.0625$ mV, $n = 0, 1, 2, \ldots N$. The model used $N = 40$ for all three populations. The mean value of the EPSP for each population is $\overline{n} = pN$, and the populations were weighted to give an overall distribution function

$$\Pr(n, N) = w_1 \cdot P(n, N; p_1) + w_2 \cdot P(n, N; p_2) + w_3 \cdot P(n, N; p_3) \ . \qquad (8.1)$$

The weighting factors and probabilities (w_i, p_i) were (0.05, 0.01), (0.85, 0.10), and (0.10, 0.29) respectively for the three model populations. These values were chosen to give a reasonable match to the mean value and skew for the experimental data reported by Smetters [SMET]. Smetters' actual data is skewed somewhat more to the left and also shows fairly high-voltage outliers as large as about 4.2 mV. However, the mean value, the probabilities in the region from 0.4 mV to 1 mV, and the flattened "back porch" from 1 mV to 2 mV agree reasonably well with Smetters' data (given the constraint of having three subpopulations in the model distribution). The binomial function that underlies $\Pr(n, N)$ is, of course,

$$P(n, N; p) = \binom{N}{n} \cdot p^n (1-p)^{N-n} \ .$$

Parameter p_i, along with N, determines the mean EPSP response for population i. Weighting factor w_i models the fraction of the overall population made up of members of subpopulation i.

Although, undoubtedly, closer empirical fits to the experimental data could be obtained by other models, figure 8.1 is nonetheless adequate for making an important first observation: There are multiple, statistically-distinct populations producing the distribution of EPSP responses. To the extent that figure 8.1 is an adequate approximation for the overall population, we are faced with three distinct subpopulations. By itself, this does not necessarily mean a unique proxy neuron model cannot be achieved. But it does alert us to the need to understand what mechanisms underlie the presence of observably distinguishable subpopulations.

The distribution of figure 8.1 does not tell us what these mechanisms are. It merely alerts us to the presence of something that systematically results in different neurons exhibiting statistically different response properties. There are many possible mechanisms that might produce a multimodal distribution of EPSP responses. Differences in EPSP from proximal vs. distal synapses in the dendritic arbor of one neuron, due to the "cable effect" (chapter 7), might be a contributor. If this is the dominant mechanism for producing the distribution, a single proxy neuron with multiple synaptic channel models (a multi-compartment model) might likely suffice to model the entire population. But if other significant mechanisms exist as well, differing from neuron to neuron rather than from synapse to synapse, then having only a single, unique proxy neuron model might not be adequate for modeling a netlet. Figure 8.1 presumes no specific neuron model; it only makes a statistical statement regarding the observable values of EPSPs exhibited by an overall neuron population (in the case of Smetters' data, a population of pyramidal cells). We must therefore examine more closely how parametric differences among neurons affect the EPSP response.

§ 1.1 Single-event EPSP Variations in the Single Compartment Model

Let us begin with neurons modeled by a single-compartment model such as the Wilson neuron model. Naturally, use of a full Hodgkin-Huxley model for this examination would be preferable. However, the experiment from which the model of figure 8.1 was deduced does not provide the data necessary to determine the details required for full H-H modeling of the neurons in the population. What is known is that the neuron population is comprised of pyramidal cells (more specifically, layer V pyramidal cells in rat visual cortex). Most pyramidal cells belong to the RS class of spiking responses, and so use of the Wilson model will reflect the extent of our actual knowledge of the system we are dealing with. However, because the Wilson model is based on fitting firing patterns and firing rates in response to injected test current stimuli, we have to ask a key question regarding the use of this model for this purpose: Is the subthreshold response of the Wilson RS model consistent with the EPSP data when physiological values are used for its synaptic conductance and membrane capacitance parameters? If it is, we may proceed with an analysis. If it is not, we can attach little significance to whatever we might learn from the Wilson model in regard to our present purpose.

Physiological range values for synaptic g_{max} and membrane capacitance C_m were presented in chapter 7. As we are dealing with subthreshold responses to stimulation, a modeling experiment can be run using the six parameter combinations (g_{max}, C_m) for AMPA channels. Statisticians refer to this as a ***factor level*** experiment. Table I summarizes the six factor level test cases.

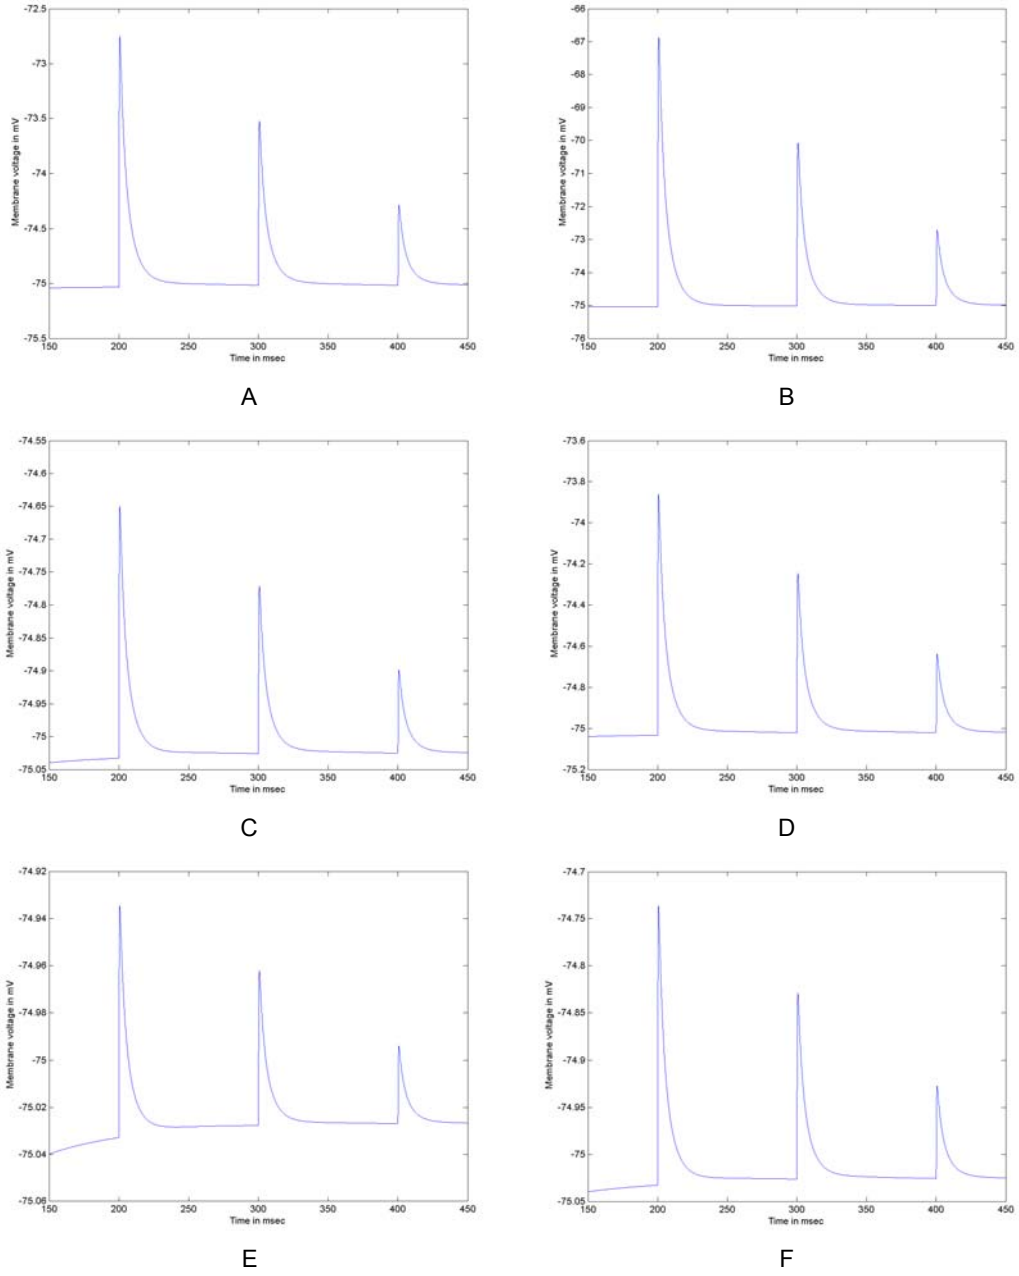

Figure 8.2: EPSP responses to simultaneous stimulation of one, two, and three synapses. (A) and (B): Class 3 (small neuron) parameters. (C) and (D): Class 2 (medium neuron) parameters. (E) and (F): Class 1 (large neuron) parameters. The order in which stimulus is applied is: 3-synapses, 2-synapses, 1-synapse.

Table I: EPSP Factor Level Experiment

Factor Class	Factor Designation	g_{max} (pS)	C_m (fF)
3	A	100	0.09
3	B	300	0.09
2	C	100	0.51
2	D	300	0.51
1	E	100	1.97
1	F	300	1.97

Table II: Peak EPSPs for Single Event Stimuli

Factor Class	Factor Designation	$EPSP_1$	$EPSP_2$	$EPSP_3$	
3	A	0.75	1.5	2.3	mV
3	B	2.25	5.0	8.0	mV
2	C	0.125	0.275	0.39	mV
2	D	0.39	0.80	1.35	mV
1	E	0.03	0.065	0.08	mV
1	F	0.10	0.20	0.285	mV

For discussion purposes the six test cases are divided into three classes according to C_m. We will call Class 1 the "large neuron" class, Class 2 the "medium neuron" class, and Class 3 the "small neuron" class since it can be presumed that C_m is correlated with the size of the neuron. Figure 8.2 shows the EPSP responses to stimulus events involving three, two, and one synapses, respectively. Peak EPSP responses in mV are summarized in Table II. The reason for testing single stimulus events involving multiple synapses is White's corollary number 2 from chapter 7. An axon projecting into a target area will make synaptic contact with all the dendrites of a neuron occurring in that target area. Therefore, we must account for the possibility that a single event stimulus can simultaneously stimulate more than one synapse. $EPSP_1$, $EPSP_2$, and $EPSP_3$ in Table II denote simultaneous stimulation of 1, 2, and 3 synapses, respectively. We will first discuss the correspondences between this model simulation and the distributions in figure 8.1. Afterwards, we will comment on the shortcomings of the model simulation experiment.

The first thing to note from analysis of the simulation results is that $EPSP_1$ values all fall within the distribution range of figure 8.1. This is encouraging because it implies that the Wilson model's subthreshold EPSP response is at the least physiologically reasonable, despite the fact that the Wilson model itself was not developed using any consideration of subthreshold responses. To this extent, we may say the Wilson model does have predictive power, since this model result goes beyond the conditions under which it was developed and is therefore not a mere curve fit function. This addresses one concern voiced earlier in the text when the Wilson model was first introduced. To be useful, an approximation model must have some predictive power, and this is our first serious test of whether or not this is true of the Wilson model.

At the same time, however, a second caution is merited. We are comparing the results of a simulation model with the results of a distribution *model*. A philosopher of science might tell us at this point that what we are doing is making comparisons entirely within the Platonic world of mathematics. Before we can celebrate clearing this first hurdle of testing the Wilson model, we must ask: How do the simulation results compare with the original experimental data from which the *approximation* provided by the model of figure 8.1 was deduced? Here it turns out that there is cause for some concern. Smetters' data shows *no* EPSP responses for amplitudes less than

about 0.125 mV. The experimental data shows a relative frequency of about 0.16 for responses in the range of 0.250 to 0.375 mV, an abrupt drop in frequency to only about 0.02 in the range from 0.125 to 0.250 mV, and no observed cases at all for EPSPs less than 0.125 mV. The first three histogram bins of figure 8.1 are those most seriously at odds with Smetters' data in the model distribution. We should therefore be very cautious in our interpretation of the $EPSP_1$ responses involving the Class 1 (large neuron) test cases. At the same time, however, one should also know that it is very difficult to accurately measure responses below 0.1 mV in a noisy environment, that it seems very peculiar that there should be so abrupt a drop in peak EPSPs between the $n = 1$ and $n = 2$ bins, and equally peculiar that there should be no heavily attenuated "cable effect" responses from distal dendrite synapses of a pyramidal cell in the sub-0.1 mV range. The point is that even experimental data, obtained at the edges of measurement capabilities, merits some caution in interpretation. Even experimental results are not delivered to us on stone tablets. This is not, and never should be used as, an excuse for ignoring experimental results. All it does is merely inject into the picture an idea of where the frontier of one's "knowledge of hard facts" lies. This is one of the realities of science that makes science such a fascinating and challenging area in which to work. The safest conclusion we can draw from this comparison is the following: Regardless of whether or not the $EPSP_1$ large neuron simulation results are truly physiological, it should be concluded that the pooled $EPSP_1$ results are *not typical* of the entire neuron population. (A "pooled" result is a set of results treated as if they all belong to a single distribution).

Next we should compare the three simulation test cases against the three putative population statistics from which the distribution of figure 8.1 is composed. Because we are taking the factor levels for g_{max} and C_m as parametric ranges, what one should expect is for the simulation results to "bracket" the distributions in a statistical sense. To see if this happens, we need the mean and standard deviations for the three binomial distribution subpopulations in figure 8.1. The mean index $\bar{n}_i$ and standard deviation $\sigma_n = \sqrt{(1-p_i)\cdot\bar{n}_i}$ for each of the three probabilities p_i for the subpopulations in figure 8.1 are shown in Table III, along with the corresponding EPSP values.

Comparing the statistics in Table III against the values for EPSP in Table II, the following observations can be made. The distribution for $p_1 = 0.01$ is most closely matched by the range for

Table III: Subpopulation Means and Standard Deviations

p_i	$\bar{n}_i$	σ_n	EPSP($\bar{n}_i$)	EPSP($\bar{n}_i+\sigma_n$)	EPSP($\bar{n}_i-\sigma_n$)
0.01	0.4	0.629	0.115 mV	0.193 mV	–
0.10	4.0	1.9	0.565 mV	0.803 mV	0.328 mV
0.29	11.6	2.87	1.515 mV	1.874 mV	1.156 mV

the Class 1 (large neuron) model for the $EPSP_2$ test case. The Class 1 $EPSP_2$ model responses bracket the mean EPSP of 0.115 mV with case F providing a fairly close match to the EPSP for one standard deviation above the mean. This implies that the small-EPSP population in the model of figure 8.1 is best represented by the class of large neurons with an average of two synaptic connections per presynaptic terminal over the typical physiological range of g_{max}. Although this alignment is nothing more than a matching-up of two model characteristics (the EPSP distribution and the Wilson RS-type model with the stated parameters), as an "average neuron" range model, this assignment does make some physiological sense. Large neurons are more difficult to stimulate because the concentration of inflowing synaptic channel charge is smaller for large cell bodies, thus implying less effect on membrane potential. Furthermore, it is reasonable to expect that a large pyramidal cell will present more dendrite surface area within a given axon's projection region, thus implying under White's second corollary a greater average number of synapses per axon terminal.

The best match for the $p_2 = 0.10$ subpopulation is provided by a somewhat complex mixture of the Class 2 (medium-sized) neuron cases. Specifically, model responses for case C $EPSP_2$ through $EPSP_3$ along with case D $EPSP_1$ through $EPSP_2$ fall more or less within the ± one standard deviation range about the mean EPSP for $p_2 = 0.10$. Such a mix is not physiologically unreasonable, but it resists a simple description similar to that for the Class 1 case, other than to say the $p_2 = 0.10$ population seems to represent a wide range of synaptic efficacies, with "effective" values for g_{max} ranging from 200 to 600 pS. This is perhaps not inconsistent with the fact that the p_2 subpopulation in figure 8.1 represents 85% of the overall neuron population. It is a modeling outcome that would tend to implicate a multimodal distribution with more than three subpopulations. It is also an outcome that suggests it is functionally possible to represent this subpopulation with a single "average neuron" model ($C_m = 0.51$ fF) with the variance in the subpopulation being taken into account by means of a distribution in the synaptic weights of the subpopulation.

Finally, the $p_3 = 0.29$ subpopulation is best matched by the Class 3 (small neuron) cases for $EPSP_1$. These responses bracket the mean EPSP in Table III, although the ranges somewhat exceed the ± one standard deviation range. The "midpoint" between the two values of $EPSP_1$ (about 1.5 mV) is reasonably close to the distribution model value of 1.515 mV. $EPSP_1$ implies one synapse per presynaptic axon projection, which would seem to be consistent with what one could expect from a small neuron.

In this analysis, "best match" does not denote "best fit." The model distribution (figure 8.1) and the Wilson model parameters have been determined independently of one another. The

analysis has merely pointed out what aspects of one model are most consistent with specific characteristics of the other (see Exercise 5). Put another way, what we have here is a *functional correspondence* rather than a physiological correspondence. This is consistent with the basic question we are asking here, namely "What might constitute an 'average neuron' model?" To go beyond a mere functional correspondence and attempt to state these conclusions as proof of or evidence for physiological implications can only be called speculative. Remember, the "average neuron" is "as mythical an entity as the jack of spades."

Papers in the neural network literature often invoke the notion of the average neuron, and speak of it as if to say, "If you take the arithmetic averages of the parameters of individual neurons, you will come up with a set of parameters describing behaviors that most of the neurons in a biological network will show." Whether an author means to leave this impression or not, the analysis conducted above should start to warn us away from this sort of thinking when our models are close to the mechanism level. This point will be further illustrated in the following section. In the present case, the problem of using a "one-size-fits-all" neuron model stems from the wide variance in the distribution of responses. The notion of a single "one-size-fits-all" average neuron is a notion owed to thinking about systems in terms of the Region-II paradigm we discussed in chapter 7. Large population variance is the deadly enemy of this kind of general systems thinking. The analysis just completed implies that capturing population dynamics in neural systems with any one "average" approximation model is not to be taken for granted.

The reason one uses an approximation model, such as Wilson's model, to construct small neural netlets is to tie the behavior of the netlet back to its biological substratum. If this tie-back is made successfully, then netlet behaviors can be seen as having important implications for brain function. But if the tie-back is illusory, then study of the netlet constructed from these kinds of "average neurons" might yield interesting and important mathematical findings of neurodynamical interest, but the imputation that such a netlet really is a model for an actual biological system is illusory as well. This is why the question, "What *is* an 'average' neuron?" is such an important issue for biological signal processing and computational neuroscience. It is as important to know what conclusions one may *not* draw from one's model as it is to know what conclusions one *may* draw.

When operating at any one "rung" on the modeling hierarchy "ladder" in neuroscience, it is a trivial if often tedious accomplishment to "torture your model's parameters until the model tells you what you want to hear." But what does it mean to say a model is "true"? The only objectively valid meaning science can have for the word "truth" is "congruence between the model and its object." Congruence, however, must be both "local" (found on the specific "rung" where one is

working) and also non-local (in the "rails" of the "ladder" connecting one rung to the next).

The analysis just concluded above implies that the notion of the "average neuron" might not be so easy to define as one might have thought. However, this analysis by itself does not rule out the possibility that perhaps a suitably-chosen distribution of synaptic weight parameters, even a functional "phenomenological" one, could capture all the pesky variance seen in the population responses to single-event stimuli. We have been working against only a single experimental observable, the single-event EPSP, and it is perhaps clear to you that if we are free to choose any g_{max} value we wish, we can force our model to agree with any subthreshold EPSP response. Would this suffice to give us an average neuron? To explore this question, we must turn to other neuronal behaviors within the netlet environment.

§ 1.2 Spatial Summation

Another commonality shared by all neurons *in vivo* is that they sum synaptic currents arising from simultaneous activation of their synaptic channels and they integrate the channel current responses over time for successive synaptic events. In this section we will examine the first of these behaviors using the neuron models of Table I.

We begin with spatial summation. A vector of action potential inputs converging on a single neuron within the same brief time period is called a ***volley***, and the neuron's response to a volley is called its ***spatial summation*** response. Let n_{syn} denote the number of synapses receiving the volley. We shall examine the response of our six neuron models to widely-spaced volleys for n_{syn} values of 10, 20, 40, and 60. The simulation responses are shown in Figure 8.3 below.

Action potential responses differ significantly among the neuron models. Cases A, C, D, E, and F all show one or more subthreshold responses for one or more of the volley stimulations. Only case B responds with action potentials for all volley strengths n_{syn}. The firing threshold for Wilson's RS-type neuron model is approximately –55 mV, which is approximately 20 mV above the resting potential of the model cell. In these simulations, as in those of the previous section, the excitatory synaptic channels are modeled by a $g^{(\beta)}$-function with time constants of 0.162 and 3.0 ms. By comparing the peaks of the subthreshold responses against the approximately –55 mV firing threshold of the Wilson model, one can see how closely the response for a subthreshold value of n_{syn} approaches the trigger point of the neuron.

Mathematically, the test conditions for these simulations assigns the same value of g_{max} to each of the synapses of the neuron. This is obviously an "average synapse" assumption built into the simulation. If the different synapses are given a distribution of g_{max} values (by assigning a synaptic weight vector W as discussed in the earlier chapters), $n_{syn} \cdot g_{max}$ would correspond to the

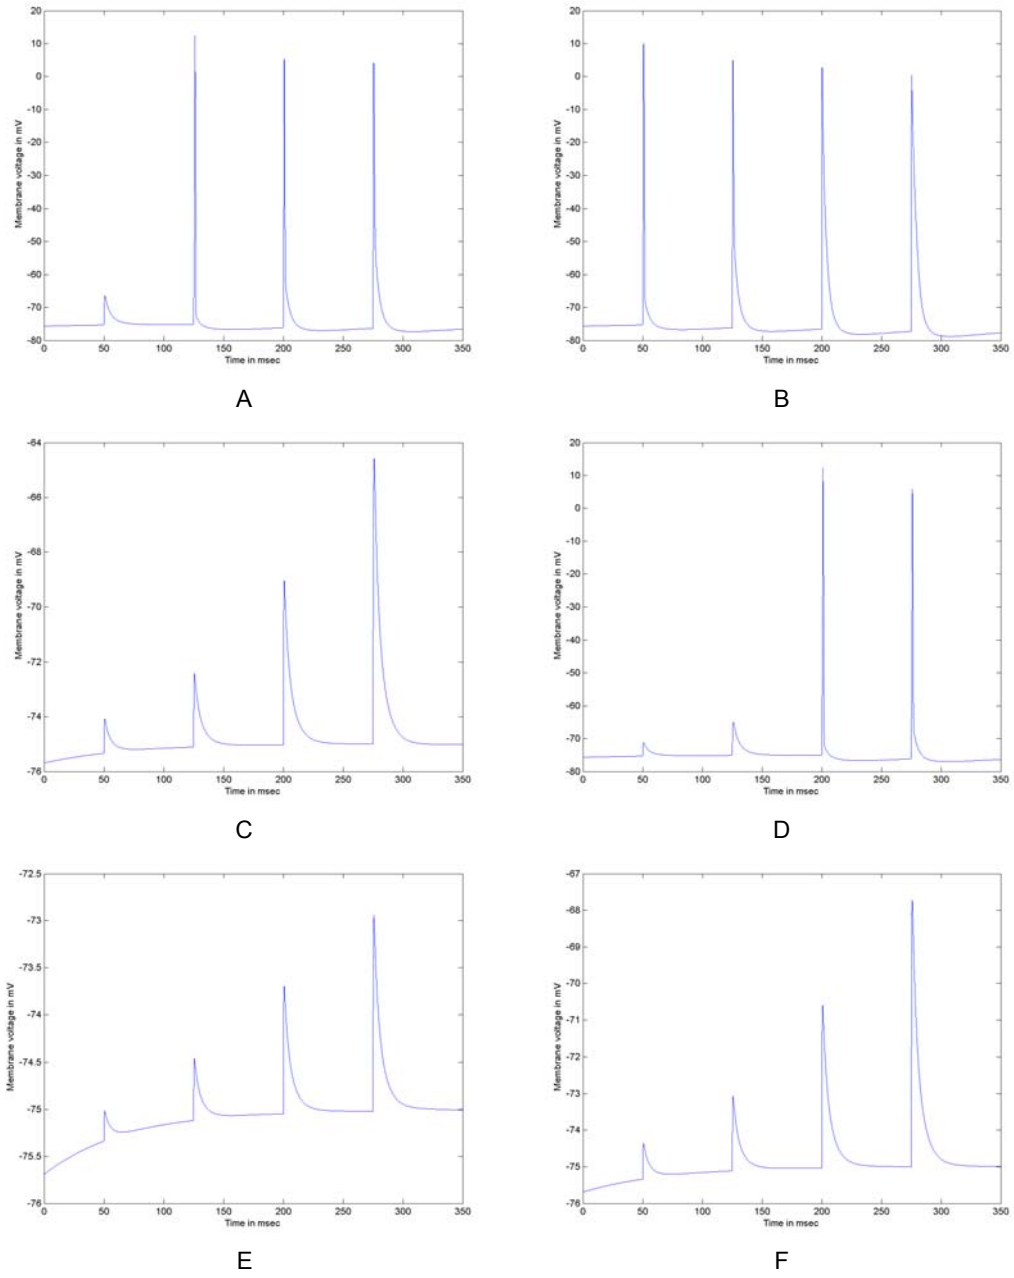

Figure 8.3: Volley responses for the six neuron models defined in Table I. The gently-arcing transient visible in subthreshold cases C, E, and F is caused by an initial condition transient in the model. This non-steady-state initial condition transient is less than 1 mV in range and does not affect whether or not an action potential response is generated during the test.

weighted sum of all synaptic inputs from the volley.

At first glance it might appear as if these response differences could likewise be accounted for by a phenomenological assignment, W, of synaptic weights. This is true to a first approximation. However, another interesting response property becomes evident through careful examination of the subthreshold volley responses. What we find is that the change in membrane voltage, ΔV_m, is

not a strictly linear function of $n_{syn} \cdot g_{\max}$. Rather, the peak subthreshold V_m response increases disproportionately faster as the firing threshold is approached. This property can also be seen quite clearly in the $EPSP_1$ through $EPSP_3$ responses in Table II. This property is actually observed in cortical cells (another success for Wilson's model). Abeles has dubbed this the ***coincidence advantage*** of synchronous signaling [ABEL1: 230-232]. The existence of coincidence advantage has led to a general interpretation that a principal functional role for cortical neurons is in acting as a ***coincidence detector*** in neural network systems [ABEL2]. Prior to the discovery of coincidence advantage, the neuron was generally looked upon as serving as a simple integrator in biological signal processing.

Can coincidence advantage be functionally captured by a single neuron model with a suitably chosen synaptic weight distribution W for modeling the response variances demonstrated by figures 8.2 and 8.3? To find out, we must conduct an ***analysis of variance*** for the peak subthreshold responses with $n_{syn} \cdot g_{\max}$ treated as an independent variable in the test. Analysis of variance (ANOVA) methods are standard tools in the application of statistics to experimental results, and a number of available statistics packages exist for carrying it out. For our present analysis, the easiest and most general of the available methods is based on linear regression models and is called "testing the mean-squared drop in the sum of squared errors" [OTT: 469-485].

Table IV provides the test data from the spatial summation simulations above. The analysis factors for the ANOVA are $n_{syn} \cdot g_{\max}/C_m$ and C_m, recalling that in Wilson's model synaptic conductances are scaled by the factor $1/C_m$. (Note that we cannot analyze case B in this modeling experiment because all its responses produced action potentials). To carry out the ANOVA by the mean-squared drop method we form two regression models. The first,

$$\Delta V_m = \beta_0 + \beta_1 \cdot C_m + \beta_2 \cdot \left(n_{syn} \cdot g_{\max}/C_m\right) + \beta_3 \cdot \left(n_{syn} \cdot g_{\max}/C_m\right)^2$$

is called the "complete model" because it contains all the parameters we wish to test. The coefficients β_i are model parameters obtained from a least-squares fit of the data. The β_3 term is included in order to try to fit the coincidence advantage phenomenon. The significance of the parameters in the complete model are tested by setting their respective β_i equal to zero, e.g.,

$$\Delta V_m = \beta_2 \cdot \left(n_{syn} \cdot g_{\max}/C_m\right) + \beta_3 \cdot \left(n_{syn} \cdot g_{\max}/C_m\right)^2,$$

and re-fitting the remaining parameters using a least-squares fit. The second model is called the "reduced model." The ANOVA consists of testing the significance of the change in the sum of squared errors between the complete model and the reduced model.

TABLE IV: Subthreshold Membrane Response Data

C_m (nF)	$(n_{syn} \cdot g_{\max}/C_m)$ (nS/nF)	ΔV_m (mV)
0.09	11.11	8.79
0.51	1.96	1.29
0.51	3.92	2.65
0.51	7.84	6.00
0.51	11.76	10.45
0.51	5.88	4.29
0.51	11.76	10.39
1.97	0.508	0.32
1.97	1.105	0.66
1.97	2.03	1.35
1.97	3.046	2.06
1.97	1.523	0.98
1.97	3.046	1.99
1.97	6.091	4.41
1.97	9.137	7.29

Carrying out a linear regression for the complete model, we obtain

$$\Delta V_m = -1.244 \cdot 10^{-3} + 0.043 \cdot C_m + 0.557 \cdot (n_{syn} \cdot g_{\max}/C_m) + 0.026 \cdot (n_{syn} \cdot g_{\max}/C_m)^2$$

with a total sum of squared errors $SSE_c = 0.534$ and mean-squared error $MSE_c = 0.0485$. The regression fit has 11 degrees of freedom. Comparing the regression coefficients against the square root of the MSE (= 0.22) immediately suggests that the β_0 term is probably negligible. Computing a reduced model with $\beta_0 = 0$ and testing the change in the total SSE which results verifies this. We do not require the β_0 term in the model. This same examination also suggests the β_1 and β_3 terms might also be set to zero, leaving only the β_2 term in the reduced model. However, when this is tested we find that the change in SSE is significant at the 0.05 level of statistical significance, which tells us that at least one of the terms β_1 or β_3 must be retained in the reduced model. Keeping the β_3 term (and noting that $0.026 = 0.161^2$) results in a reduced model

$$\Delta V_m = 0.582 \cdot (n_{syn} \cdot g_{\max}/C_m) + 0.024 \cdot (n_{syn} \cdot g_{\max}/C_m)^2 . \qquad (8.2)$$

The reduced sum of squared errors, SSE_r for this model is 0.558. The mean-squared drop in the comparison of the reduced model to the complete model is $MSdrop = (SSE_r - SSE_c)/2 = 0.012$, where the 2 in the denominator of MSdrop is the degrees of freedom for this statistic (equal to the difference in the number of terms in each model). This change is not statistically significant, and so the conclusion is: the reduced model does not have a variance different from that of the complete model. The largest error term in the reduced model is –0.618 mV and occurs at the C_m = 0.09 nF data point in Table IV.

We can also test the hypothesis $\beta_1 \neq 0$ and $\beta_3 = 0$. This is tantamount to making the hypothesis

that the neuron size (represented through C_m) is a separate factor in addition to synaptic strength factor $n_{syn} \cdot g_{max}/C_m$. In this case, however, the change in SSE measured by the mean-squared drop statistic turns out to be significant at the 0.05 level of significance, and therefore this reduced model must be rejected.

The conclusion drawn from this analysis is the following. The subthreshold stimulation response of the membrane voltage does not depend separately on neuron type (C_m) at least within the range of subthreshold responses examined in the simulations. (Because the largest ΔV_m in the data set is about half of what is required to trigger an AP, a cautious examination would add a few more stimulation test points to provide ΔV_m data closer to the trigger point, $\Delta V_m \approx 20$ mV). *This means that the subthreshold responses can be obtained using a single neuron model,* with the variability across the population taken into account by a distribution of "normalized" synaptic weight parameters $n_{syn} \cdot g_{max}/C_m$. Thus, we are a step closer to being able to answer the original question, namely: Can the population of (in this case) pyramidal cells in a netlet be represented by a single "average" pyramidal cell with phenomenologically-chosen synaptic conductance parameters? So far, it looks like the answer is turning out to be "yes." However, we have one more behavior to examine before finalizing this conclusion.

§ 1.3 Temporal Integration

Temporal integration refers to the build-up of membrane potential in response to closely-spaced volleys arriving sequentially in time, each of which alone would produce only a sub-threshold membrane response. Because we have just tentatively concluded that we can capture the variability in neuron responses by a suitably-chosen synaptic weighting, let us examine the temporal integration responses of our six test case neuron models by choosing n_{syn} values that produce subthreshold responses. We will apply four successive volleys at a spacing of 5 ms between each volley and simulate the membrane voltage responses of our six neuron models. The test data is tabulated in Table V.

Figure 8.4 illustrates the responses of the model neurons to the successive subthreshold volleys. In all cases except case E, the volley stimulus is sufficient to produce an AP response on

TABLE V: Temporal Integration Test

Neuron Test Case	n_{syn}
A	10
B	3
C	60
D	20
E	60
F	60

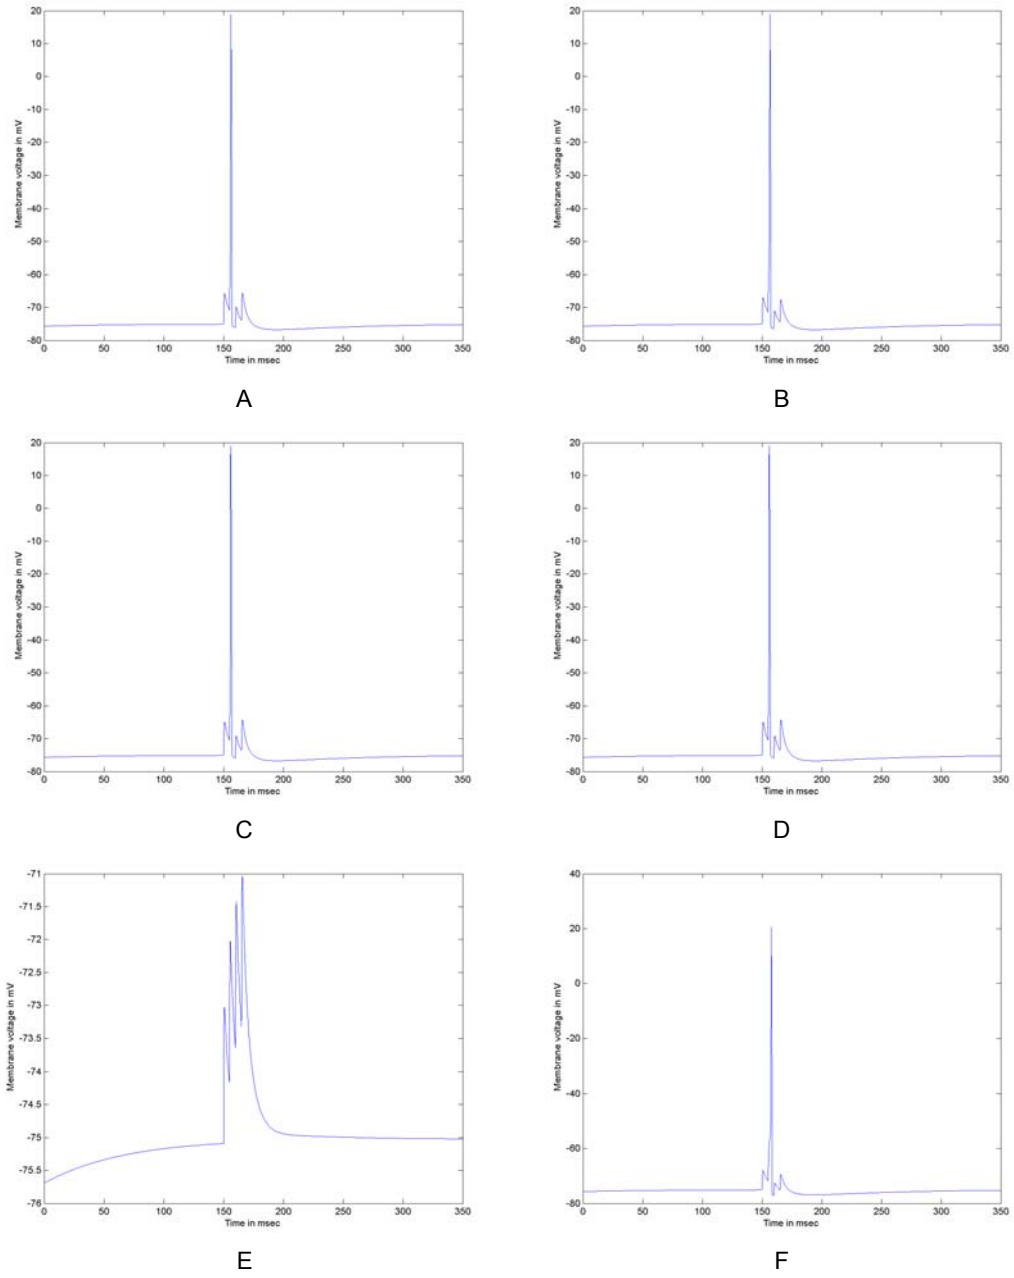

Figure 8.4: Temporal integration responses for the six model neurons in response to subthreshold synaptic stimulation.

the second volley. Integration action is clearly visible in all six test cases, but particularly so in case E. The absence of a second AP in response to the fourth volley is an indicator of the refractory behavior of the neuron model. It was noted in the chapter introducing the Wilson approximation model that Wilson's approximations remove the mechanisms responsible for the absolute refractory period characteristic of the neuron. The simulations shown above tell us that at least some relative refractory behavior is still maintained in the Wilson model if the stimulus is

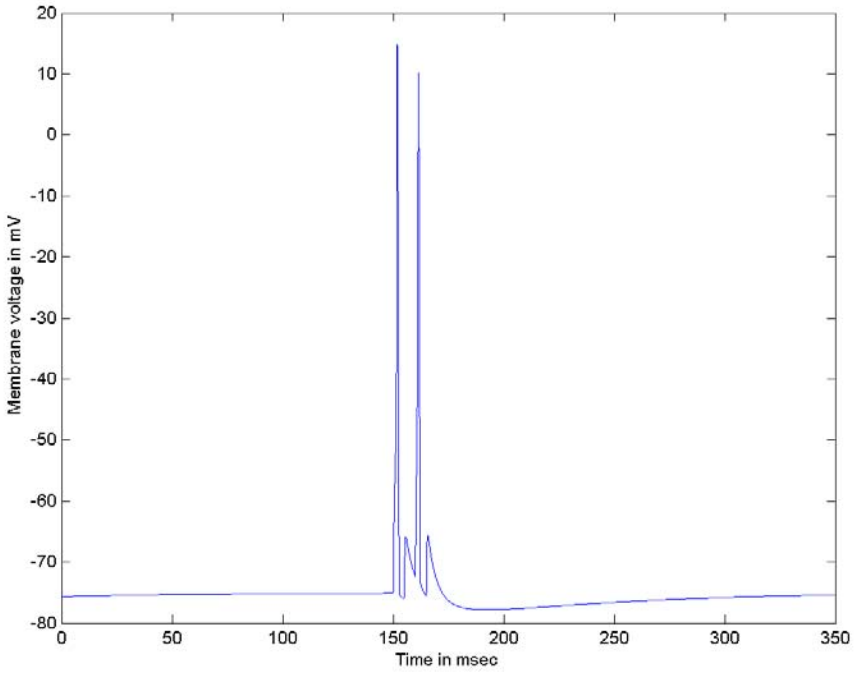

Figure 8.5: Test case B when n_{syn} = 5. This stimulus is sufficient to produce an action potential at the first volley input. The neuron has sufficient refractory-like behavior to suppress a second AP for the volley occurring 5 ms later, but the model does produce a second AP upon receipt of the third volley. This 10 ms spacing between AP responses is somewhat faster than is physiological for most RS-type pyramidal cells.

not too great. There are, however, limits to this. Figure 8.5 illustrates the temporal integration response for test case B when n_{syn} = 5 is used. At this level of stimulation, the model is capable of producing APs at a 10 ms spacing, which is somewhat faster than is physiologically realistic for most RS-type pyramidal cells.

The principal lesson we learn from examining the temporal integration responses is the following. Our first two studies from the previous sections provided us with some confidence that it is realistic to represent a population of neurons (RS-type pyramidal cells in our present case) by means of one "average" neuron model with suitably chosen synaptic weight parameters. But of the six different model cases we have examined, is there a way to decide which of the three "size" models would best serve as the "average" neuron? The temporal integration simulations illustrate that "size matters" in the context of modeling functional microcircuits and netlets. The "small neuron" cases (A and B) also have small thresholds for n_{syn}, which acts as a limitation on the number of converging inputs from other model neurons that will be realistically represented by the model. In computer engineering language, C_m determines the **fan-in** property of the neuron model. Likewise, the "large neuron" cases (E and F) have a likewise large fan-in for n_{syn}. Thus, the functional selection of an average neuron model must be made within the context not only of

the neuron signaling type (RS-, FS-, IB-, or CB-type), but also within the context of the netlet model being considered. Because all the neuron models in a functional microcircuit or a netlet will be "average neurons" of one type or another, if the modeler wishes to tie the netlet model to an identifiable biological system, such as a cortical functional column, the population of biological neurons for which one average neuron acts as proxy determines in part the number of proxy neurons in the functional microcircuit or netlet. Thus the fan-in to be represented for each proxy is tied to the anatomical system it is to represent, and this in turn affects what one should select as the "size model" for the proxy neurons.

To date there has been no general rule-of-thumb put forward to guide this selection process. Indeed, the number of proxy neurons used in functional microcircuit models often seems to be a rather arbitrary choice on the part of the modeler. We can, however, find a qualitative sort of consideration that ought to go into this decision. Is there any sort of "reference center point" that might usefully serve to "calibrate" this decision-making process? In the case of neocortical model networks, the answer is, "Yes, somewhat." It is provided by another of Abeles' ideas, which he calls the ***synchronous gain 50*** (SG50) characteristic for cortical neurons [ABEL1-2].

SG50 is a parameter emerging from Abeles' 1982 statistical model of the integrative properties of the neuron. Its mathematical development is found in the original citations, but as a statistical parameter SG50 has a simple meaning. It is the inverse of the number of synchronous synaptic inputs required for a 50% probability of generating an action potential for an "average" synapse. Using model parameters in the physiological range found in neocortex, Abeles found the average value for SG50 to be 1/37 (37 synchronized inputs to produce an AP with 50% probability).

Let us compare this statistic with the response of the Wilson RS-neuron using the "medium sized neuron" parameter value C_m = 0.51 nF with mid-range maximum conductance parameter $g_{\max}$ = 0.2 nS. Setting $x = g_{\max}/C_m = 0.392$, $n_{syn} = 1$, and using our reduced model equation (8.2) from the previous section gives us

$$\Delta V_m \approx 0.582 \cdot x + 0.024 \cdot x^2 = 0.23 \text{ mV} .$$

This estimate is consistent with the most commonly observed experimental EPSP responses for neocortical and other pyramidal neurons and, likewise, agrees with the "average" cortical neuron single-event EPSP of 0.22 mV used by Abeles [ABEL1: 121-123]. The simulated Wilson RS-type neuron response using these parameters returns a peak simulated EPSP of 0.255 mV, which is in quite good agreement with both Abeles' "typical" model values and with laboratory results.

Unfortunately, we cannot trust our reduced model formula to estimate the n_{syn} required to evoke an AP response. This is because an AP requires $\Delta V_m \approx 20$ mV for Wilson's model, and this

value is well outside the range of values used in finding our reduced model curve fit expression. (Regression models break down outside the parametric range from which they are determined). Setting $x = n_{syn} \cdot g_{max}/C_m$ with $n_{syn} = 37$ gives us $x = 14.5$ and $\Delta V_m \approx 13.5$ mV. As it happens to turn out, the Wilson simulation agrees fairly precisely with this value even though it is outside the range of the parameters for our reduced model formula. However, setting $\Delta V_m = 20$ mV and solving for the required value of x tells us we must have $x = 19.1855$, which implies $n_{syn} = 48.9$. Simulations demonstrate that the AP can be evoked for much smaller values of n_{syn}. For the neuron at steady-state at its resting potential, simulation results show that the fractional value $n_{syn} = 39.17$ produces $\Delta V_m \approx 20$ mV, and APs are evoked for higher values of n_{syn}. The curve fit model, in contrast, predicts $\Delta V_m \approx 14.6$ mV for $n_{syn} = 39.17$. This tells us that the fit rather badly underestimates the coincidence advantage near the neuron's firing threshold. To obtain an improved reduced model curve fit expression, one must simulate additional data points and re-fit, possibly adding a cubic term, x^3, to the curve-fit expression to boost its prediction of coincidence advantage near firing threshold (see Exercise 8).

If we take $n_{syn} = 39.18$ to represent a kind of statistical measure for SG50, we obtain a simulation value for SG50 of SG50 = 1/39.18 = 0.0255, compared to Abeles' value of 0.0270. The simulation result is therefore about 5.5% different than the value coming out of Abeles' statistical model. Given the variances in neural systems we have been discussing, this is reasonably good agreement. The conclusion to be drawn from this exercise is the following: The "medium-sized neuron" model ($g_{max} = 0.2$ nS; $C_m = 0.51$ nF) appears to represent a reasonable best choice for a base "average neuron" model to serve as proxy in functional microcircuit models of neocortex. *Thus, at least for the case of neocortical pyramidal cells of the RS-type, we have finally completed our first objective for this chapter, which was to determine how one can define an "average neuron" proxy with physiological validity*.

A moment's reflection will raise a rather obvious question pertaining to what we have just seen. Abeles' SG50 parameter is statistical and it carries the qualifying statement that the number of synchronized synaptic inputs it imputes has only a 50% chance of evoking an AP. The Wilson model, on the other hand, is not a stochastic model. Wouldn't its "SG50" n_{syn} evoke an AP every time?

The answer to this question is, "no." The reason lies with the qualifying phrase used above, namely, " for the neuron *at steady-state* at its resting potential." The n_{syn} value of 39.18 is a value just sufficient to reach the AP firing threshold when the neuron is at rest. Figure 8.6 illustrates what happens when this "at rest" condition is violated even slightly. The first stimulus arrives at t

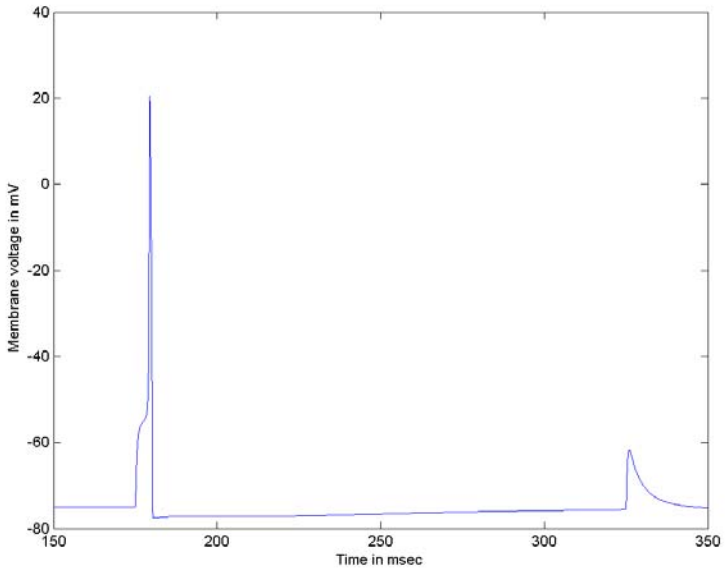

Figure 8.6: Response of the average RS-type neuron model to stimulations with n_{syn} = 39.18 spaced at 150 ms intervals. At the first stimulus the neuron is in its steady-state resting condition. The EPSP at the second stimulus reflects the very slight change from equilibrium produced by the aftereffects of the first AP.

= 175 ms with n_{syn} = 39.18. The wide, slow base of the action potential is typical for a stimulus that just barely reaches the firing threshold. The second stimulus, arriving 150 ms later, sees a neuronal state that has not quite returned to the rest condition (a relative refractory period effect). Even though the volley bombardment rate is only 6.67 Hz, the peak EPSP for the second stimulus reaches only about –60 mV (equivalent to an n_{syn} of about 37 for the neuron completely at rest). This means the coincidence advantage is a strong function of the neuron's state for stimuli in the vicinity of the firing threshold.

This phenomenon is one of the characteristics of cortical neurons Abeles studied with his statistical model. Indeed, variations about the steady-state condition of the neuron is a principal component of Abeles' statistical model and is why SG50 is defined in the way it is. This effect is further illustrated in Figure 8.7 below. The neuron receives three volleys with n_{syn} = 15, 38, and 39.18, respectively, at t = 175, 200, and 325 ms, respectively.

When the neuron is at rest in the steady-state, n_{syn} = 38 would be insufficient to evoke an action potential response. In this case, however, it does suffice even though the second volley is a full 50 ms later than the mild first volley. Examination of the figure reveals that V_m has actually returned to, and perhaps is even slightly below, its resting level at the time the second volley is applied. Looks, however, are deceiving in this case. Although V_m appears to have returned to its initial value before the second volley arrives, the voltage-gated channels have *not* returned to their respective resting conditions. The neuron sensitivity to stimulation is, accordingly, different

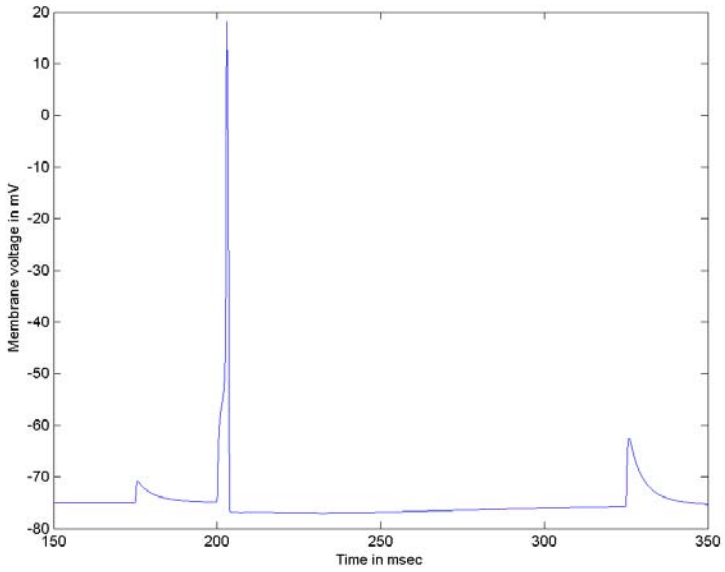

Figure 8.7: Illustration of effects of neuron state on coincidence advantage. The neuron receives three volleys of n_{syn} = 15, 38, and 39.18 at t = 175, 200, and 325 ms, respectively. Normally n_{syn} = 38 is insufficient to evoke an AP in the neuron at rest, but in this case the residual effect on the neuron's state due to the first volley makes n_{syn} = 38 sufficient to trigger an AP.

when the second volley arrives and produces an AP even though n_{syn} = 38 is less than that implicated by our "SG50 value" for n_{syn} = 39.18. We may also observe that the third volley, which is "at the SG50 value," fails to evoke an AP response.

When modeling functional microcircuits by means of proxy neuron models, the modeler must bear considerations of these nonlinear dynamics in mind for designing targeted fan-in properties in the proxy model. AP triggering is more dependent on the state of the underlying VGCs than it is on merely the membrane potential or n_{syn}. Coincidence advantage is more of a statistical than a mechanistic characterization of neuronal response.

§ 2. Proxy Neuron Mimic Models

The discussions in §1 remained very closely focused on maintaining a linkage between the proxy model and physiology in its use of Wilson's approximation model. We have seen that for at least the case of RS-type pyramidal cells of the neocortex, Wilson's model proves to be impressively accurate across an impressively large scope of neuronal behaviors it can successfully exhibit. Unfortunately, the Wilson approximation model does suffer the drawback of being computationally expensive, and this limits our ability to use it efficiently above the level of functional microcircuit modeling of neural netlets.

To proceed further – say to the level of functional minicolumn modeling – we need a practical

method to overcome the "Square Law of Computation" limits for the Wilson model. To ascend to this next rung in the modeling hierarchy of computational neuroscience and biological signal processing, we turn to the mimic models. Of the two principal mimic model schemas presently favored in active research (Izhikevich's schema and Rulkov's map schema), we will consider only Rulkov's schema. This is because its per-iteration computational cost is comparable to that of the Izhikevich schema but its Δt (0.495 ms) is much larger than the practical limit needed for the Izhikevich schema (0.01 ms) using Euler's method. It therefore achieves a nearly 50-fold computational cost advantage over differential equation based mimic models.

At the time this is being written, the Rulkov schema is still quite newly arrived on the scene. Consequently, our science has not yet had very much time in which to give proper consideration to how to best use this exciting new modeling schema, nor to develop the sort of research programs needed to establish and ensure its linkages with its lower neighboring rung on the modeling ladder (that of the approximation models). For this reason, the discussion in this section is considerably more qualitative than was the discussion given Wilson's model in §1.

In §3 we will be introducing even more abstract functional models for neuron population proxies, and so one pertinent question to raise in this section is: Why bother with this intermediate level of abstraction in neuron proxy modeling? Why not just jump ahead to the models in the next section? As you might by now suspect, the answer is: *linkage*. The models we will be discussing later are older than Rulkov's model, and they are very widely used in neural network research. However, the very nature of the abstractions employed in them quite effectively severs the direct link between them and physiology.

The neuronal world is a world filled with marvelous diversity: neurons with regular-spiking responses, neurons with fast-spiking and non-accommodating responses, neurons with bursting responses, and so on. The obvious question – and a question that as of yet has not been clearly settled at the present state of our science – is: What does this marvelous signal processing diversity "do" for the animal whose nervous system possesses it? One can argue – and this often *is* argued – that there must be some important evolutionary advantage gained from this diversity. We are still very far from having a clear understanding of what this advantage or advantages might be.

But there is another attendant question. We do not even know for a fact that such a putative advantage exists. Evolutionary changes that result in varieties and species are brought about by gene mutations. Many of these mutations are disadvantageous rather than advantageous (when other environmental factors are more or less constant), and some few are advantageous. However, these very same genetic mutation mechanisms also produce ***neutral variations*** – variations that

bring no survival advantages or disadvantages whatsoever. The genetic mutation that gives rise to the individuality of fingerprints is one example of this. The success of one's Cro-Magnon ancestors can hardly have been due even in the tiniest part to the individuality of fingerprints; the extinction of the Neanderthals could hardly have been due to the individuality of fingerprints (assuming that in Neanderthals this individuality did in fact exist) or to its lack (if Neanderthals did not have this trait; no one knows whether or not they did).

Are a few, many, most, or even all the varieties in neuronal signaling types functionally advantageous? Most neuroscientists assume that at least some of them are (but which ones?), yet the fact is we do not actually know for sure. Your author is inclined to think this variety is probably very, very important, but this is just a speculation on his part. He has no *convincing* evidence, yet, to back up his opinion. It is an area of computational neuroscience research that is queerly lacking in research investment.

Of course, such an investment in research time and effort has only recently become practically feasible. The Rulkov model is the development that makes it practical.

In its brief time on the computational neuroscience stage, Rulkov's mimic model has already been used to construct impressively large crystalline neural network models, comprised of up to hundreds of thousands of Rulkov proxy neurons. However, and this is not really a criticism of the work done to date, the network models that have so impressively demonstrated the computational power of the Rulkov schema are of more mathematical than biological interest. The crystalline neural network structures so far reported are *very* crystal-like in the sense that they are step-and-repeat structures with no identified biological model of what they supposed to represent. Some very tantalizing phenomena have emerged from these mathematical constructions – e.g. the finding of an interesting emergent wave propagation property that only shows up in simulation when the network exceeds a certain large number of neurons in its matrix [RULK1].

There is no shortage of fascinating research questions for the mathematics of nonlinear neuro-dynamics. There cannot be all that much doubt that the Rulkov neuron models will eventually prove to be important contributors to this research. There can be almost no doubt at all that some yet-to-be-reported mathematical discovery will one day prove to be an epic, groundbreaking discovery for the advancement of neuroscience. Yet all of this remains merely Platonic if it is not accompanied by research that forges the linkage between mimic-based netlets and neural physiology and anatomy. This, too, is a difficult and challenging research area, and we must scout out how this particular trail might be blazed.

As an illustration of why the question of neuron firing-type is at issue, let us look at a comparison of RS- and IB- type neuron responses for volley stimulation and co-localized NMDA

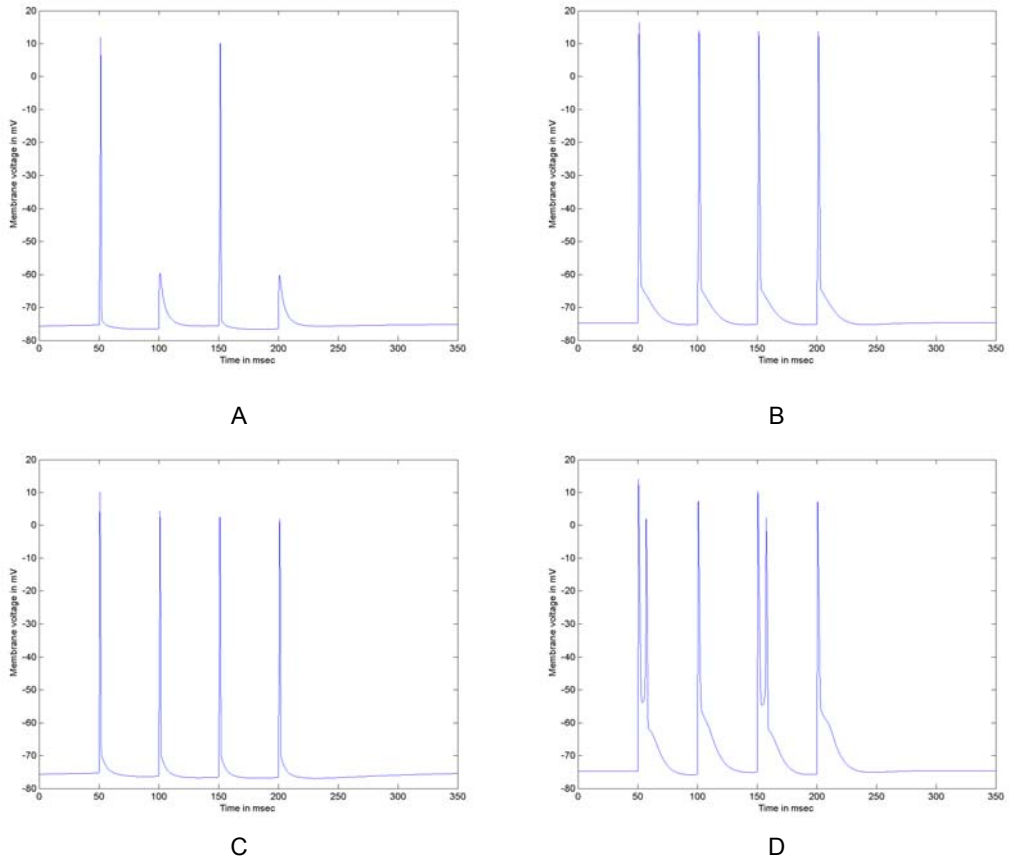

Figure 8.8: Comparisons of RS-type and IB3-type firing patterns for co-localized AMPA and NMDA channel synapses. (A) RS-type neuron with n_{syn} = 50 volley inputs. (B) IB3-type neuron with n_{syn} = 50 volley inputs. (C) RS-type neuron with n_{syn} = 100 volley inputs. (D) IB3-type neuron with n_{syn} = 100 volley inputs.

and AMPA channels present in the excitatory synapses. Co-localization of NMDA and AMPA channels is common in cortical neurons, although the ratio of NMDA to AMPA channels within a synapse is highly variable, ranging from NMDA-dominated synapses to AMPA-dominated synapses. Recall that NMDA channels are glutamate-enabled/membrane voltage activated channels. Thus, they contribute very little current flow until V_m is depolarized to about –50 mV.

IB-type neurons [McCO2] are found mainly in neocortex layers IV and V (along with RS-type neurons). The layer V IB-type pyramidal cells, like the RS-type pyramidal cells, make long distance white-matter projections to many different cortical and subcortical targets. Thus, both neuron types play an "output neuron" role in the functional column architecture of the neocortex. For this reason, any differences found in the firing patterns they produce in response to synaptic stimulation is clearly of interest for understanding neural network functioning.

It is now known that intrinsic bursting neurons come in a wide variety bursting responses. McCormick et al. noted,

> The vigor with which neurons could burst was quite variable. . . These data suggest that there is

> actually a continuous distribution of spiking modes among pyramidal neurons. The molecular mechanisms underlying the burst mode may thus be expressed in various degrees in cells with similar somadendritic structure. The distinct laminar distribution of bursting cells may imply that this trait is determined at an early stage of cortical development, since the laminar position of a neocortical cell is strongly determined by its birth date [McCO2].

The burst characteristics for the IB-cells found by McCormick et al. [McCO2], [GRAY1] are not particularly well matched by either Wilson's IB1 or IB2 model parameters. A fairly reasonable match can be obtained by modifying several of the model parameters. We will call this modified model the IB3-cell. Its unique parameters (in scaled units, as presented in chapter 6) are: $\tau_T = 9$, $\tau_H = 17$, $\tau_R = 5.7$ (all in ms), $g_T = 1.5$, and $g_H = 2.4$. These parameters were used for the simulation.

The Wilson model simulations are shown in Figure 8.8. The synaptic g_{max} is 0.2 nS for both the AMPA and NMDA channels. The AMPA channel time constants are $\tau_1 = 0.162$ ms, $\tau_2 = 3.0$ ms (which is the same as was used in §1 earlier). $C_m = 0.51$ nF for the RS-type neuron and 0.53 for the IB3-type neuron. The NMDA time constants are $\tau_1 = 3.0$ ms and $\tau_2 = 40.0$ ms. The NMDA channel model is the modified $g^{(\beta)}$-function described in chapter 4. Figures 8.8A and C are the RS-type responses; figures 8.8B and D are the IB3-type responses. Figures 8.8A and B are responses with $n_{syn} = 50$. Figures 8.8C and D are responses with $n_{syn} = 100$.

Comparing figures 8.8A and B, the first thing to note is the marked difference in output firing rates and patterns. Although both neurons receive an identical volley tetanus with an incoming firing rate of 20 Hz, the RS-type neuron response is a "divide by two" response producing an output firing rate of 10 Hz. In contrast, the IB3-type cell responds to all four volleys in a "relay-like" mode. When n_{syn} is increased to 100 for the input tetanus (figures 8.8C and D), the two output responses are altogether different. The RS-type neuron is able to "follow" the input tetanus in this case, passing along one AP spike for each volley input. The IB3-type neuron, on the other hand, responds to the first and third volleys with a two-spike burst, and responds to the second and fourth volleys with a single AP spike.

This very interesting contrast in neuron behaviors suggests the possibility that the two neuron types, working in concert and projecting to the same target areas, might be viewed as constituting a very sophisticated ***encoding system*** based on the intensity (and frequency) of the input tetanus. There is a fairly wide range, between $n_{syn} = 50$ and $n_{syn} = 90$, over which the IB3-cell does not burst, whereas by $n_{syn} = 75$ the RS-cell is following the input volleys in "relay" fashion. At $n_{syn} = 40$ the IB3 response is similar to figure 8.8A. In comparison, the RS-type neuron responds with only a single spike for $n_{syn} = 40$. It would appear, then, that the two neuron types encode the strength and the firing rate of the volley stimulus over fairly broad ranges of values, and represent

this encoding in the firing patterns they produce in response. The two neural responses, taken jointly, therefore appear to be capable of ***classifying*** the nature of their stimuli into different ***sets*** of stimulus conditions that are identifiable by looking at the spiking patterns of the neuron pair.

A scientific question of long-standing interest to computational neuroscience – one might almost call it the "holy grail" of the field – is the question of how the brain encodes the information that its neurons process. Historically, the two most-pursued hypotheses for "the neural code" have been the firing rate hypothesis ("information is carried by the firing rate") and the synchronized signaling hypothesis ("information is carried in vector form by the synchronized firing of specific 'feature fragment' netlets"). The firing-rate hypothesis was first proposed by von Neumann in the 1950s, and the thinking in regard to this hypothesis has always born a very noticeable analogy to the frequency-modulation methods of communication system theory. The synchronized-signaling hypothesis is one of the major cornerstones of Malsburg's correlation theory of brain function, and it presently is commanding a great deal of attention from both physiologists and neural network theorists. It is an hypothesis closely tied to the research interest in activity-wave propagation through neural networks, e.g. [RULK1].

Both these long-standing hypotheses owe much of their justification to the habit of thinking of network models in terms of the traditional, one-size-fits-all proxy neuron presupposition generally used in neuroscience for many years. However, the example just given suggests that it might be the case where "the neural code" follows a much more complicated and elegant plan. Does the coexistence of RS- and IB- type neurons in the primary output layer of the neocortex mean that information is being coded for transmission by an encoding scheme such as the one conjectured here? We do not know. Clearly the possibility exists in principle. The only way to find out is to study it with neural netlet and network models capable of dealing with this sort of signaling behavior. Almost no work, either mathematical or physiological, has to date been widely reported which addresses this question.

This brings us around to the role for mimic models in netlet modeling. If in fact there is an information encoding system operating in neural networks according to the diversity of firing patterns possible with differently-responding neuron signaling types, a root question is: How much information can such an encoding scheme represent? We find ourselves here at that fine gossamer boundary that separates "signal processing" from "information processing." It is good to recall Weaver's comment from chapter 1 that, "information relates not so much to what you *do* say, as to what you *could* say." Referring to our example above, the classification-encoding scheme (if that is what it is) made possible in principle by this pair of cortical output cells would be a rough classification indeed. There is a range of stimulus intensity, from approximately n_{syn} =

65 to n_{syn} = 90 at the tetanus rate of 20 pulses per second, where the joint output patterns of the RS- and IB3- cells do not change. Thus, at this tetanus frequency, there is a fairly wide range of stimulus activity levels that are "all the same to the output neurons." If different levels of stimulus activity within this range do in fact represent "different messages" for the neural system to process, then the encoding capability of just our two pyramidal cells is ***information lossy***, i.e. information being presented to them is not passed on to other cell groups.

Information loss is not necessarily a bad thing. Information theory's technical definition of information loss, in non-mathematical language, merely says "if you can't tell what the original information at the source was from the signals you receive, information loss has occurred." Under information theory's definition, the adder used in a computer is an information-lossy device because if its output is "7" you cannot tell if its input was "3 + 4" or "6 + 1" or etc. The adder inside a computer is useful *precisely because it is information lossy* in a very specific way.

A neural network that always exhibits the same response to a wide variety of different stimuli is said to have a ***stereotyped response***. The wider the variety of stimuli that evokes the same response, the more stereotyped the neural network is said to be. The more stereotyped is a neural network, the more information lossy it will be. Some parts of the central nervous system, e.g. the brain stem and many of the neural reflex circuits of the ventral horn of the spinal cord, are regarded as being highly stereotyped. This suits the specific and specialized nature of their functions. Other parts of the brain, e.g. the cerebral hemispheres, are wondrously diverse in their responses – the functional opposite of being stereotyped. This is both blessing and puzzle for us as human beings. William James wrote,

> The dilemma in regard to the nervous system seems, in short, to be of the following kind. We may construct one which will react infallibly and certainly, but it will then be capable of reacting to very few changes in the environment – it will fail to be adapted to all the rest. We may, on the other hand, construct a nervous system potentially adapted to respond to an infinite variety of minute features in the situation; but its fallibility will then be as great as its elaboration. We can never be sure that its equilibrium will be upset in the appropriate direction. In short, a high brain may do many things, and may do each of them at a very slight hint. But its hair-trigger organization makes of it a happy-go-lucky, hit-or-miss affair. It is as likely to do the crazy as the sane thing at any given moment. A low brain does few things, and in doing them perfectly forfeits all other use. The performances of a high brain are like dice thrown forever on a table. Unless they be loaded, what chance is there that the highest number will turn up oftener than the lowest? [JAME, vol. I: 140].

Can a "high brain" emerge as the consequence of vast numbers of interacting "low brain" neural netlets? Artificial intelligence and neural network pioneer Marvin Minsky thinks so:

> We think the difference in abilities comes from the fact that a brain is not a single, uniformly structured network. Instead, each brain contains hundreds of different types of machines, interconnected in specific ways which predestine that brain to become a large, diverse society of partially specialized agencies. . . Why did our brains evolve so as to contain so many specialized

> parts? Could not a single, uniform network learn to structure itself into divisions with appropriate architectures and processes? We think this would be impractical because of the problem of representing knowledge. In order for a machine to learn to recognize or perform *X*, be it a pattern or a process, that machine must in one sense or another learn to represent or embody *X*. Doing that efficiently must exploit some happy triadic relationship between the structure of *X*, the learning procedure, and the initial architecture of the network. It makes no sense to seek the "best" network architecture or learning procedure because it makes no sense to say that *any* network is efficient by itself; that makes sense only in the context of some class of problems to be solved. Different kinds of networks lend themselves best to different kinds of representations and to different sorts of generalizations. This means that the study of networks in general must include attempts, like those in this book, to classify problems and learning processes; but it must also include attempts to classify the network architectures [MINS: 273-274].

A computer's adder cannot function as its instruction decoder. Complex systems tend to be systems comprised of multiple "specialist" functions, and from a purely mathematical point of view such specialization often depends on the ability to classify multiple different input conditions into a smaller set of outcomes; and this is what the word "generalization" in neural network theory actually means.

However, it seems almost self-evident that Minsky's "hundreds of different types of machines" is something of an underestimation. If, as James' and Minsky's comments imply, high brain behavior involves the cooperative efforts of a legion of specialist netlets and networks, it is clear that developing a functional understanding of such a brain organization must involve the study of increasingly complex and interacting neural assemblies with increasingly large numbers of neurons. It does not take very long for Weinberg's Square Law of Computation to overwhelm the ability to model netlets and functional columns using approximation models such as Wilson's.

But to push back the onset of our computational limitations and yet retain the needed linkage to real biological function, it is necessary that mimic models, such as Rulkov's, must be *calibrated* not only against those approximation models at the neuron level, but against netlet *architectures* that are adequately representative of the biological substrate. Functional biological phenomena, not merely biological "plausibility," must be the standard of comparison. This cannot be demonstrated using the traditional "one-size-fits-all" proxy approach until and unless it is first shown that this highly abstract level of representation actually does correspond to outcomes emerging from intercourse among more minute biological signal processing netlets.

Chapter 6 provided Rulkov parameters for an RS-, FS-, and one type of IB- class neuron model. In view of what has just been said about the need to match Rulkov neuron responses against a physiological standard and the need to go beyond the "one-size-fits-all" approach to netlet modeling, one can anticipate that a broader suite of Rulkov model parameters is likely to be necessary. The discussion of Rulkov neurodynamics in chapter 6 provides some guidance for carrying out the calibration process, but no general "design procedure" has yet been presented.

Clearly, this rhetoric is more of a "go west, young man" type of dictum than a specific recipe for how one develops a practical research program. But, at the present state of knowledge, this vague theme is perhaps about the best one can do in advance of more specific findings – not from crystal networks involving hundreds of thousands of cells, but from less lofty explorations of netlet architectures capable of predicting effects the anatomists and physiologists can test in their laboratories. At the very dawn of what came to be known as "the Age of Reason," Francis Bacon wrote,

> 19. There are and can exist but two ways of investigating and discovering the truth. The one hurries on rapidly from the senses and particulars to the most general axioms, and from them, as principles and their supposed indisputable truth, derives and discovers the intermediate axioms. This is the way now in use. The other constructs its axioms from the senses and particulars, by ascending continually and gradually, till it finally arrives at the most general axioms, which is the true but unattempted way.
>
> 22. Each of these two ways begins from the senses and particulars, and ends in the greatest generalities. But they are immeasurably different; for the one merely touches cursorily the limits of experiment and particulars, whilst the other runs duly and regularly through them; the one at the very outset lays down some abstract and useless generalities, the other gradually rises to those principles which are really the most common in nature.
>
> 104. Nor can we suffer the understanding to jump and fly from particulars to remote and most general axioms . . . and thus prove and make out their intermediate axioms according to the supposed unshaken truth of the former. This, however, has always been done to the present time from the natural bent of the understanding, educated to, and accustomed to, this very method by the syllogistic mode of demonstration. But we can then only augur well for the sciences, when the ascent shall proceed by a true scale and successive steps, without interruption or breach, from particulars to the lesser axioms, thence to the intermediate (rising one above the other), and, lastly, to the most general. For the lowest axioms differ but little from bare experiments; the highest and most general (as they are esteemed at present), are notional, abstract, and of no real weight. The intermediate are true, solid, full of life, and upon them depend the business and fortune of mankind; beyond these are the really general, but not abstract, axioms, which are truly limited by the intermediate.
>
> We must not then add wings, but rather lead and ballast to the understanding, to prevent its jumping or flying, which has not yet been done; but whenever this takes place, we may entertain greater hopes of the sciences [BACO: 108, 128].

This is still as good a bit of advice today as it was in 1620.[1] Let us follow it and add some ballast to our discussion.

In chapter 7 the idea of the functional microcircuit model was introduced by means of the example provided by Douglas and Martin. Using a handful of approximation neuron models, each said to represent a whole population of neurons, they were able to successfully imitate certain features of neocortical response to particular patterns of thalamic inputs. Without prejudice to the

[1] Bacon himself made no important scientific discoveries, but he performed an even more valuable service as a sort of Field Marshal who issued the call to science and the rejection of authoritarianism in the 17th century. Although it is not widely remembered today, the founding of the Royal Society was inspired by another of Bacon's books, *New Atlantis*.

real accomplishments of this model, let us look at some of the things this model did not achieve. The model did not address cortico-cortical afferents (inputs coming in from elsewhere in the neo-cortex). The model did not simulate all the phenomena of interest to Douglas and Martin using *one single neural netlet structure* (cf. figures 7.4 and 7.5). It was not very specific in modeling the interactions between neighboring chain links of cortex even within the one functional micro-circuit, and it did not at all model interactions between adjacent functional microcircuits or adjacent functional minicolumns (refer to figure 7.5).

In fairness to Douglas and Martin, most of the general (if not overly specific) interconnection pathways in neocortex (figure 7.2) were not known in 1991. They did what could be done at that time. However, as the saying goes, that was then and this is now. If we still do not possess complete detailed knowledge of cortical (or thalamical, or hippocampal, or cerebellar, or etc.) anatomy and physiology, we do know much, much more than was known in 1991, and the power of our computational resources today is tremendously greater than it was then. What shall we do with all this knowledge and computational capability?

Douglas and Martin provide us with an example to follow, namely in the idea of *using physiological* (approximation) *models to represent well defined, small functional microcircuits*. General estimates are available for making approximate determinations of general organizational features of some specific cortical regions, e.g. [ABEL1: 50, 53, 59], [GIBS], [SCHÜ]. We also know a few details about what sort of neuron-type to neuron-type connections are factually present in cortex, e.g. [WHIT: 81]. From these we can at least begin to explore what sort of signal processing modalities are likely to exist in various netlet structures, based on approximation models and with recourse to anatomical and physiological findings whenever these are available.

We also have, from above, a speculation on what sort of signaling characteristics we might expect to be phenomena of interest – namely set classifications with stereotyped responses. (A small microcircuit most likely cannot do very much by itself). Simple classifier and logic function problems were the first signal processing tasks studied by neural network theory in its infancy half a century ago, and they are tasks neural networks do very well. The issue, of course, is "what sort of things are to be classified?" What we can do now, that the pioneers back then were in no position to do, is: classify according to signaling patterns more accurately representative of the physiology of real cortex.

Except for the bothersome amount of detail this task entails, this might seem a trite thing to spend one's time on. As was said, classifiers and classifier theory have been around since the early days of the science. But, and this is a point not to be under-appreciated, *classification is the first and most elementary step in understanding mathematical structure*. The three most basic

structures of pure mathematics – algebraic structure, order structure, and topological structure – all must presume a *classification structure* at the first step. To understand this, one must first know that modern mathematics is based on ***set theory***. But a set is nothing other than a classification, a subset nothing other than a sub-classification. *To understand neural classifiers is to better understand how to develop mathematical descriptions of neuronal function*.

Functional microcircuit classifiers are likely to be pretty elemental and limited in capability. If we liken a neuron to a single "logic gate" (borrowing from the language of computer engineers), a functional microcircuit modeled with approximation neuron models is likely to rise no higher than the level of a "medium scale" function (analogous to a basic adder). The limitation here is simply the cost of computing. To go beyond this, to explore the nature of *interactions* among such circuits in a functional column (Minsky's "different types of machines") we must be able to put many such functional microcircuits together in a larger system. *This* is the *real significance* of efficient mimic models such as Rulkov's. The mimic model enables us to put together functional microcircuits to form larger functional netlets, just as the "medium scale logic circuit" lets the computer engineer put together functionally useful computing subsystems. This, however, does presume that the mimic netlets correctly reproduce the behaviors of the approximation-model functional microcircuits and correctly account for the interactions among them. We must not merely *assume* they do so; we must *demonstrate* they do so.

This is the research in pursuit of Bacon's movement "from the lesser axioms to the intermediate axioms." Through this work we find the promising potential for building up a "library" of functional microcircuit modules, which then in turn can support even more complex neural netlets at the next rung up the ladder. Equally important, the functions held in this library will unavoidably pose new questions for the anatomists and physiologists to pursue. In a manner of speaking, the computational theorist will be able to walk down the hall to his colleagues in the biological arm of the science and say, "My model says such-and-such. Can you look to see if that's true, please?" Neural network theory is thereby made *predictive* rather than merely *descriptive*, and this is what every proper science should strive to be. There is no more fertile and untilled research territory in all of computational neuroscience than this today.

§ 3. Spiking Population Proxy Models

As we move up to the next rung of the modeling ladder, one can expect to encounter a new limitation – different in kind from Weinberg's Square Law of Computation limit – which may well prove to be more of a limiting factor than computational horsepower. Rulkov et al. have already successfully demonstrated that simulating large "crystal" neural networks containing

hundreds of thousands of mimic neurons is both possible and practical right now. It is one thing to stand and watch the intricate and beautiful activity patterns these networks show themselves capable of producing, and to appreciate the unexpected novel features these displays reveal. It is quite another thing altogether *to understand what these behaviors mean and what the model has to tell us*. Without comprehension of function, we can no more claim a scientific understanding of brain theory than the ancient Egyptian surveyors could claim to understand geometry.

§ 3.1 Populations and Local Functions

The great difficulty in understanding network function stems from this: The great majority of neurons in the brain have no direct contact with either sensory inputs or motor outputs. This makes it very, very difficult to understand what precisely is the functional role of most small regions of neurons in central systems. For some time now, much research attention has been focused on the phenomenon of ***activity wave propagation*** in one- and two-dimensional chains of crystalline neural network structures. An activity wave is a somewhat abstract concept, but in its simplest examples it is merely the totality of more or less synchronized neuron output activities across some smaller region of the network. The regions exhibiting firing activity move from one place to another in a more or less continual fashion reminiscent of waves rippling across a pond. Simulations of activity wave behavior in crystal-like neural networks typically focus on just the activities of the excitatory neurons, although the networks themselves usually also contain local inhibitory neurons as well. (We do not have evidence of *inhibitory* activity wave propagation, except as a localized response to propagating excitatory signaling).

What the spatial extent and functional significance of activity wave propagation may be is far from certain at present. We know that activity waves do exist in neocortex, and we do know they propagate for some distance, although how far this distance extends is less certain. Evidence suggests the propagation distance may not be very far (on the order of a few millimeters across the surface of the cortex). What the functional significance of these waves may be – and even the issue of whether there is any functional significance to them at all – is currently still somewhat controversial. Not the least of the factors involved in this controversy is due to the non-anatomical nature of the crystal-like neural networks used to study the mathematical properties of activity wave propagation. Neural network theorists have their reasons for selecting the one- and two-dimensional chain models they use, of course. But these reasons usually do not impress neurobiologists, who demand to know what any of this has to do with real networks and ignore the mathematical findings if the network theorists cannot answer this question to their satisfaction.

One reasonable question that can be asked in regard to any crystal-like chain model of neural networks is, "Chains of what? What are the links in the chain supposed to represent?" The usual answer is, "activity levels of the local neural population." But this is a rather vague answer and is difficult to connect to the idea that a "local population" has some "local function." How might this question be more specifically addressed?

There is reason to think local neuron populations *do perform* some kind of local function or functions. (It is pointless to ask what the local function *is* unless one thinks *there is* local function). The reason *par excellence* to think so is the documented existence of functional columns in the neocortex across a wide range of mammalian species. White tell us,

> Coexisting with the horizontal, laminar aspect of cortical organization is a vertical one envisioned by Lorente de Nó (1938) to consist of chains of interconnected neurons extending across all layers of the cortex. Confirmation of the notion that the cerebral cortex has a distinct vertical aspect to its organization was provided subsequently by the observation of functional columns in both sensory and motor areas . . . Functional columns, demonstrated initially by electrophysiological methods, are composed, in sensory areas, of neurons that share similar response properties. . . Columnar arrangements also have been identified in motor cortex where stimulation of discrete groups of neurons causes specific muscles to contract. . .
>
> These efforts have resulted in the following discoveries: (1) Functional columns in many areas of the cerebral cortex may be shaped more like bands or slabs than like columns; (2) columnar arrangements based on the distribution of afferent fibers are a general feature of cortical organization; and (3) in most instances, clear structural correlates for functional columns have yet to be identified [WHIT: 12-13].

Two important attributes of functional columns are the fact that they are populations that respond to the same or almost the same set of afferent inputs, and these populations project their responses to the same or almost the same set of target populations. These two attributes certainly suggest some kind of "input-output" organization of local cell populations, and any relationship between input signals they receive and output signals they produce would be, by definition, the *mathematical* function or functions of that population. A system theorist is obliged to think that where one finds a mathematical function, there must also be some sort of biological function underpinned by it. To understand the latter, we must first figure out the former.

Is there any common enough architectural "theme" we can exploit in pursuit of obtaining an improved understanding of population function? Many neurobiologists think there is, at least insofar as specific regions (neocortex, hippocampus, cerebellum, etc.) are concerned. (The functional organization almost certainly differs between brain regions, e.g. neocortex architecture differs from cerebellar architecture). Of course, the picture is made somewhat more complicated by observed evidence suggesting many functional columns can "re-wire themselves" in responding to different patterns of stimulating afferents. Nonetheless, qualitative models have been proposed as a first step in understanding functional column organization. Figure 8.9 presents

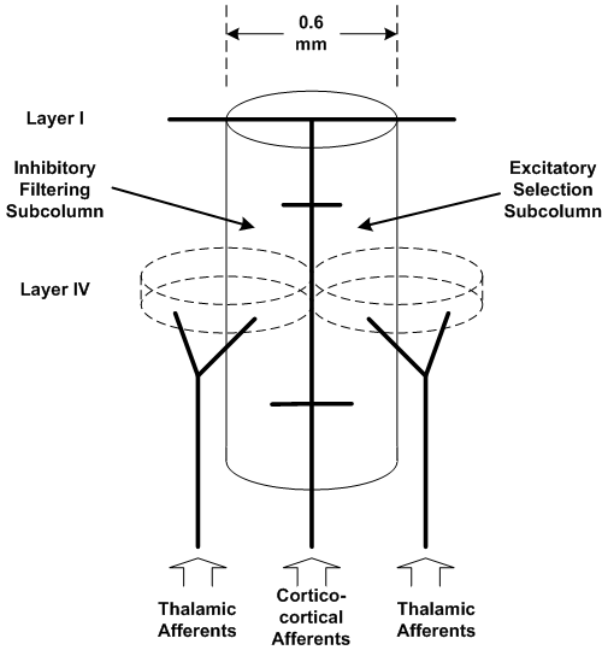

Figure 8.9: Functional column organization model proposed by Szentágothai.

one such model proposed by Szentágothai [SZEN]. It is instructive to compare this architectonic theme with the functional microcircuit models and the overall cortical model presented in chapter 7. The Szentágothai column is approximated as a cylinder with an approximate diameter of about 0.6 mm (600 microns). Cortico-cortical afferents, which make up the great majority of input signals in neocortex, project up through the center of the column, which is divided into an "inhibitory filtering" subcolumn and an "excitatory selection" subcolumn. Thalamic afferents project primarily into layer IV of the cortex, where a small, offset sub-cylinder structure is situated which overlaps adjacent columns. (It is instructive to compare this with Douglas' and Martin's functional microcircuit model in figure 7.5). The current findings on cortical signal flow organization [DOUG1] tell us layer IV makes its local output projections into layers II/III of the cortex (see figure 7.2). Roughly speaking, pyramidal cells receive about half of their inputs from white matter projection ("external inputs") and the other half from local circuits ("intra-columnar" and neighboring column inputs) [ABEL1: 58]. Thalamical afferents, again roughly speaking, make up less than 10% of the synaptic connections in layer IV [DOUG1] but appear to have synaptic strengths on the order of twice that of the synapses made with local pyramidal cells.

For neocortex the architectonic of figure 8.9 is the obvious candidate to serve as the "unit" of a local population model, with its connections to and from neighboring columns serving as the "links of the chain." Because the functional column structure is vertical, this implies that a two-dimensional population network model ought to be able to represent the neural network topology.

The obvious next question to ask is, "How many neurons per 'unit column' do we have to deal with?"

This number is going to be very species-dependent as well as dependent upon what specific area of the neocortex (somatosensory cortex, motor cortex, etc.) we are dealing with. A rough estimate of neuron density for human neocortex is 20,000 neurons per mm^3. Using an average cortex thickness of 3 mm and the diameter of the column in figure 8.9, this places the number of neurons per unit column at about 17,000 neurons, with roughly half allotted to each subcolumn.

That, as they say, is a lot of neurons. Rulkov et al. have demonstrated the practical capability of simulating a neural network of this size (although in [RULK1] the network was not organized around the functional column architectonic), but it is clear that current computational capabilities do not much exceed the capacity for modeling one column at the neuron-by-neuron level. To this, of course, is added the problem that we do not know the specific "wiring" within a column to very much level of detail beyond that of White's rules and corollaries.

We are, however, aided by the fact that some significant fraction of the column population will fire synchronously or nearly synchronously (when they fire) in response to stimuli. Again, it is not clear how large this fraction will be; this is why functional microcircuit and netlet modeling levels are important for understanding biological signal processing. Putting all this together illustrates why we need population models. It also illustrates why, as we move from the functional microcircuit and netlet level to the level of neural networks, the proxy models become more abstract. At this transition point, our interest moves from the behavior of individual and small groups of cells to the functional properties of *the synchronized population of cells*.

§ 3.2 The Integrate-and-fire Proxy Model

The first proxy model for a network-scale population model we shall look at is a very simple, and for that reason very popular, proxy model. It is called the integrate-and-fire (I&F) model. This model was first introduced in 1965 by Stein [STEIN]. Stein developed this model for the purpose of studying variability in neuronal firing rates rather than for the purpose of simulating neural networks, but within a decade it had been adopted by neural network simulation modelers.

Figure 8.10 illustrates the I&F model in diagram form. Variable V represented the synaptic stimulation in Stein's original model, although when the I&F model is used to represent neuron populations V takes on an abstract meaning, namely as a variable that represents a general *level* of stimulation regarded as being present in the population. It is computed as the weighted sum of input impulses p_n, $n = 1, \ldots, N$, so that

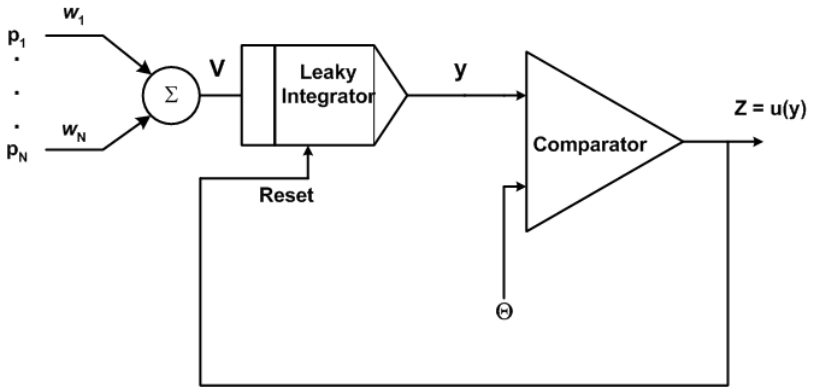

Figure 8.10: The Integrate-and-Fire Proxy Model

$$V(t) = \sum_{n=1}^{N} w_n \cdot p_n(t) \ . \tag{8.3}$$

For each input impulse function, $p_n(t) = 0$ if the input is inactive at time t, and equals a unit impulse if the input is active at time t. In a population model, each input line represents a tract of neuron inputs coming into the population from some other population. Weights w_n represent the weighting that is given to the signals represented by input tract n. Output $Z(t)$ represents the population output.

The heart of the model is the "leaky integrator" function (called an "integrate-and-dump" by electrical engineers). The "reset" signal is normally equal to zero, but when output $Z(t)$ is active (represented by $Z = 1$) the reset signal is activated (set equal to 1) and remains active for a period t_r called the refractory time of the I&F model. When reset = 1, the signal $y(t)$ is held equal to zero. Otherwise, when reset = 0, y and V are related by the differential equation

$$\frac{dy(t)}{dt} = -\frac{1}{\tau} \cdot y(t) + \frac{1}{\tau} \cdot V(t) \ . \tag{8.4}$$

Variable Θ is called the firing threshold. For $y < \Theta$, $Z = 0$. When $y \geq \Theta$, $Z = 1$ (an impulse) and the reset signal is set equal to one and held at that value for the next t_r seconds.

As always, equation (8.4) must be converted to a difference equation for computer simulation. If the $p_n(t)$ inputs are modeled as perfect impulses (technically called the "Dirac delta function"), then $V(t)$ is also a perfect impulse, and if no impulses are received between time t and time $t + \Delta t$ the solution of (8.4) is

$$y(t + \Delta t) = y(t) \cdot \exp(-\Delta t / \tau) + \frac{V(t)}{\tau} \ . \tag{8.5}$$

The key presupposition in deriving the difference equation (8.5) is the supposition that no new

impulsive inputs arrive in the interval between t and $t + \Delta t$. So long as this condition is met, (8.5) is an exact solution for (8.4). Note that the exponential term in (8.5) depends only on the fixed ratio $\Delta t/\tau$, and is therefore a constant and its value can be pre-computed before the simulation begins. This means we can set Δt to as large a value as we wish, subject to the constraint that no additional input pulses occur in the interval between simulation steps, and no discretization error will be introduced into (8.5). Because very, very few neurons are capable of firing action potentials at a rate exceeding 1000 pulses per second, a simulation step size of Δt = 1 ms is very popular in network models using the I&F proxy.

The I&F proxy has the lowest computational cost of any "neuron" model save only the original McCulloch-Pitts model (which has no differential equations at all in its description). Also, the fact that we can use large values for Δt reduces the computation costs even more. It is little wonder why the I&F proxy has been the most popular proxy model for many years. The weights w_n can mimic excitatory input tracts ($w_n > 0$) or inhibitory input tracts ($w_n < 0$). Further, the input impulses are typically obtained as the output of other proxy cells, and therefore they are binary-valued, i.e. $p_n(t)$ is either zero or one at each simulation time step.

Because (8.4) is a linear differential equation with constant coefficients, analytic solutions are obtainable for describing a variety of behaviors for various input conditions. One basic characteristic is the firing rate of the proxy in response to a constant level of stimulation V. Let V be a constant in (8.4). In this case, the exact solution for $y(t)$ for initial condition $y(0) = 0$ is

$$y(t) = V \cdot (1 - \exp(-t/\tau)) \ .$$

From this expression, we see that if $V < \Theta$ then y never reaches the firing threshold. On the other hand, if $V > \Theta$ then y increases until it reaches $y = \Theta$, whereupon the proxy fires an output pulse and resets y back to zero, holding it there for t_r seconds. After this, the process repeats and the result is a constant, steady output tetanus from the proxy. It is easily shown that the output firing rate of the proxy in pulses per second is given by

$$R = \frac{1/\tau}{(t_r/\tau) - \ln[1 - (\Theta/V)]} \ , \quad V > \Theta \ . \tag{8.6}$$

Another characteristic of interest is the response of the proxy to a synchronous input tetanus of rate R_{in}, where we will assume $R_{in} = 1/T$, $T = k \cdot \Delta t$ for some integer constant k. At each time $t = mT$ with $m = 0, 1, 2, \ldots$ the stimulation (8.3) will have some non-zero value $V = V_T$ and will be zero for all other times t. We define the geometric ratio $\rho = \exp[\ -T/\tau\ ]$ and assume $y = 0$ for $t < 0$. Then from (8.5) we have

$$y(0) = \frac{V_T}{\tau} \ .$$

We have two cases to consider here. If $V_T > \tau \cdot \Theta$, then the proxy fires immediately. Provided that the refractory interval $t_r < T$, the proxy will likewise fire at each subsequent incoming volley. We will call this condition the ***all-pass response mode*** of the I&F model. Otherwise, for $V_T < \tau \cdot \Theta$ we will have successive solutions

$$y(T) = \frac{V_T}{\tau} \cdot \rho + \frac{V_T}{\tau} = \frac{V_T}{\tau} \cdot (1 + \rho)$$

$$y(2T) = \frac{V_T}{\tau} \cdot (1 + \rho + \rho^2)$$

$$\vdots$$

$$y(M \cdot T) = \frac{V_T}{\tau} \cdot \sum_{m=0}^{M} \rho^m$$

for as long as $y < \Theta$. The summation term in the expression above is called a ***power series*** and its solution resolves to a well known expression,

$$y(M \cdot T) = \frac{V_T}{\tau} \cdot \frac{1 - \rho^{M+1}}{1 - \rho} \ . \tag{8.7}$$

We again have two cases to consider. Either y will build up and eventually surpass the firing threshold Θ, or else $y(t)$ will settle into a periodic signal where its peak value never reaches the firing threshold as $M \to \infty$. Since by definition $\rho < 1$, in the latter case we have

$$\lim_{M \to \infty} y(M \cdot T) = \frac{V_T}{\tau \cdot (1 - \rho)} < \Theta$$

from which we obtain the threshold condition

$$V_T > (1 - \rho) \cdot \tau \cdot \Theta \tag{8.8}$$

for firing in response to the input tetanus. Note that because ρ is a function of T, this firing threshold is a function of the tetanus firing rate. As T decreases, ρ increases in value, and so a particular value of V_T sufficient to trigger firing at one value of T might be insufficient for triggering the proxy at a larger value of T (a lower firing rate for the input tetanus). Thus, if the proxy is not operating in the all-pass response mode, it acts in what will be termed a ***high-pass filter mode***, suppressing the response to a tetanus at lower firing rates but responding to it at

higher firing rates. Likewise, for some fixed $V_T < \tau \cdot \Theta$, we can solve (8.8) to find the maximum period T_c for which a tetanus of sufficient duration will eventually evoke a firing response from the proxy. Calling $R_c = 1/T_c$ the ***cut-in rate*** for the input tetanus, it is easily found that

$$R_c = \frac{1/\tau}{\ln\left[\frac{1}{1 - V_T/(\tau \cdot \Theta)}\right]} \quad . \tag{8.9}$$

The proxy will fire for any input rate $R_{in} \geq R_c$.

If V_T is large enough to eventually trigger a firing response, but small enough that the proxy is not operating in all-pass response mode, the proxy will fire at the first integer value M for which $y(MT)$ equals or exceeds Θ. Applying this condition to (8.7) we obtain

$$1 - \rho^{M+1} \geq (1 - \rho)\frac{\tau \cdot \Theta}{V_T}.$$

Rearranging and taking the natural logarithm of each side, we have

$$\ln\left[\exp\left(-\frac{(M+1)T}{\tau}\right)\right] = -(M+1)\frac{T}{\tau} \leq \ln\left[1 - (1 - \rho)\frac{\tau \cdot \Theta}{V_T}\right].$$

M will therefore be the smallest integer satisfying the constraint

$$M + 1 \geq -\frac{\tau}{T}\ln\left[1 - (1 - \rho)\frac{\tau \cdot \Theta}{V_T}\right] \quad . \tag{8.10}$$

Note that the logarithm term in (8.10) is always a negative number and V_T is constrained to the range

$$(1 - \rho) \cdot \tau \cdot \Theta < V_T \leq \tau \cdot \Theta \,. \tag{8.11}$$

This implies the solution range is $M \geq 0$ for all V_T falling within this range, with equality if and only if $V_T = \tau \cdot \Theta$ (onset of the all-pass response mode). Recall that one of the conditions of this analysis is $t_r < T$, i.e. the input tetanus is not coming in faster than the refractory period of the I&F proxy. Therefore, the firing rate of the proxy's response is

$$R = \frac{R_{in}}{M + 1} \,. \tag{8.12}$$

The output firing rate of the proxy in high-pass filter mode is always less than the firing rate of the input tetanus. Because a smaller R is equivalent to a larger T, when this proxy projects its

output to another I&F proxy with the same time constant τ, that destination proxy will see a smaller value for *its* ρ parameter, thus implicating a larger minimum V_T required to evoke a firing response to a tetanus coming in at rate R. If the sum-total of its afferent inputs produces the same value for its V_T as the first proxy, its firing rate R_2 will be less than R, and etc. for each succeeding link in a chain of I&F proxies. Eventually at some point we will arrive at an I&F proxy for which V_T does not reach the minimum stimulation required to evoke a firing response and the chain will cease to propagate a response. We will call this situation an ***evanescent propagation mode***.

Finally, let us consider the case where the input tetanus is coming in with $T < t_r$. Assuming the I&F cell can respond with some $M \geq 0$, it will occasionally see pulses coming in from the tetanus that arrive while the cell is still in its refractory period. Since y is held equal to zero during the refractory period, the cell will, in effect, "miss" this stimulus event. Consequently, its response rate R will be less than $R_{in} = 1/T$. The difference in this case is that the output firing pattern will no longer consist of a series of equally-spaced output pulses. The output response can become very complicated-looking, consisting of ***packets*** of pulses spaced at intervals of $1/R$, with variable numbers of pulses in the different packets and various time spacing between packets. We will say such a response ***loses coherence*** with the input tetanus.

These various characteristics of the I&F cell have consequences for neural network models constructed with them. We will discuss some of these consequences later. For now we will merely note one thing: The consequence of successive firing rate reduction in chains of I&F cells operating in high-pass filter mode is one factor motivating the use of a widely used model neural network architecture first reported on by Abeles and known as the ***synfire chain*** [ABEL1, 3], [HAYO].

§ 3.3 Some Shortcomings of the I&F Proxy

The attractive simplicity of the I&F cell has made it a popular proxy model for many years. At the same time, this simplicity comes with a price tag. The I&F proxy has an absolute refractory period but no relative refractory period. It cannot generate a bursting response to an input stimulus, no matter how strong that stimulus may be. It is usually difficult to construct I&F networks in which a group of I&F cells can be made to fire synchronously unless these cell are operated in the all-pass response mode. It has only one way to respond to any stimulus: either it fires or it does not. Are these serious shortcomings?

Figures 8.4 through 8.7 all contained illustrations of cases where the relative refractory behavior of the neurons exercised an effect on the output firing pattern it produced. When one uses a proxy model to proxy for the behavior of a large group of synchronously-responding

neurons, it is reasonable to guess that relative refractory characteristic probably are not too terribly dissimilar across the population (although some dissimilarity should be expected just from the fact that neuron variability is so great in a biological network). If one finds at the netlet level of modeling that relative refractory effects are significant for the cell group's overall function, then the I&F cell's inability to produce such effects would then be a cause for concern; otherwise it would not be.

Figure 8.8 illustrated a hypothetical case for which the conjecture was made that bursting neurons might play a role in the functional encoding of information within a neural network. The I&F cell cannot respond with a bursting output, and this means it does not allow the modeler a way to test this encoding conjecture without adding additional I&F proxies specifically for this purpose. This is part and parcel with the fourth shortcoming listed above: it either responds or it does not, and when it does respond this response is very stereotyped. Of course, one could mimic the bursting effect by using a small netlet of I&F cells; one of the exercises at the end of this chapter will illustrate this. However, if one finds oneself adding more and more I&F proxies for the sole purpose of producing a richer suite of possible responses to stimuli, one also starts to lose the principal advantage of the I&F model, namely its low cost of computation.

Finally, in regard to the difficulty in obtaining I&F neural networks in which different I&F cells are made to respond synchronously, there are two points from our earlier discussions to bear in mind. First, let us recall that our primary motivation for using a proxy model is to allow us to replace netlets of many neuron mimic models with, preferably, one population proxy. Second, one of the experimentally observed properties of functional columns is what White called the "ephemeral" and "transitory" character of many functional columns:

> Clear structural correlates for functional columns are difficult to find. One possibility for the general lack of correspondence between anatomical structure and functional columns probably has to do with the likelihood that certain functional columns are rather ephemeral, owing their existence to the continued presence of a specific set of stimulus conditions [WHIT: 13].

> Mapping the different systems of functional columns, identified by anatomical and physiological techniques, has resulted in a situation where the visual cortex has more columns than it has space; that is, different columnar structures overlap . . . One way out of this dilemma is to presume that the formation of functional columns is a transitory phenomenon and so the cortex does not contain the same set of functional columns from one minute to the next. Because the different columns appear in sequence, the cortex would not at any time contain more columns than it has space [WHIT:200].

This situation is one we might have anticipated from looking at figure 8.9. Not only does the Szentágothai column share its layer IV with columns on either side; its cortico-cortical afferent tract rises through the column and then spreads out via layer I to neighboring columns. The hypothesis that functional columns dynamically "re-wire themselves" in response to different

stimuli is an hypothesis coming straight out of anatomical and physiological studies. It is interesting that this finding dovetails quite nicely with a similar hypothesis put forth by Malsburg from mathematical and psychological considerations:

> Time is divided into two scales, a psychological time-scale (some tenths of a second) which is characteristic of mental processes, and a fast time-scale (some thousandths of a second). Mean unit activity evolves on the psychological time-scale, but the activity fluctuates around this mean on the fast temporal scale. Units bind to each other by correlating their activity fluctuations. A set of units can be bound into a block by synchronizing their fast activity fluctuations. Several such blocks can coexist if their activity is desynchronized relative to each other: this is the solution to the superposition catastrophe. . .
>
> Correlations are shaped by connections. If correlations are to represent variable bindings, connection strengths must vary. This function is called synaptic modulation. The excitatory connection between synchronized units is increased in strength, up to a maximum strength which is characteristic of the connection. . . The excitatory connection between two desynchronized units is decreased in strength, down to the value zero. These changes take place on the psychological time scale. If there are no signals in the two units, the connection slowly sinks back, within times characteristic of short-term memory, to a resting state, in which it conducts with a fraction of its maximum strength [MALS4].

This idea, described in qualitative terms here, arose from Malsburg's correlation theory of brain function in the early 1980s. He calls a neural network architecture that implements this sort of dynamic "re-wiring" a ***dynamic link architecture*** or ***DLA***.

> The units of DLA play the role of symbolic elements. . . Units are endowed with structured signals changing in time. These signals can be evaluated under two aspects, intensity and correlation. Intensity measures the degree to which the meaning of the unit is alive in the mind of the animal. Correlations . . . quantify the degree to which the signal of one unit is related to that of others.
>
> Dynamic links constitute the glue by which higher data structures are built up from more elementary ones. Conversely, the absence of links (temporary or permanent) keeps mental objects separate from each other and prevents their direct interaction. . . More generally, mental objects are formed by binding together units representing constituent parts. The infinite richness and flexibility of the mind is thus made possible as a combinatorial game. . . Dynamic links are the means by which the brain specializes its circuit diagram to the needs of the particular situation at hand [MALS3].

How precisely all this *might* take place is far from settled. Malsburg's dynamic links are envisioned as fundamentally a synaptic-level mechanism, and this view is not without dispute by other prominent theorists. Regardless of the details of *mechanism*, one point does enjoy fairly widespread agreement: the basic *functional* idea is that cell groups or assemblies "change allegiance" and join (or leave) particular functional "units" (as Malsburg calls them) for others. At the level of the proxy population model, a key requirement for the possibility of this is an ability for different proxy cells to synchronize their activities with one another. This is not impossible to do with I&F proxies, but it is not particularly straight-forward either and tends to foster multiplication of the number of proxy cells used in the network for the explicit purpose of effecting this "dynamic link" capability. Again, as we increase the number of I&F cells in order to

get a structure to do what one I&F cell by itself cannot, the model starts to lose its attractive computational advantage.

§ 3.4 The Eckhorn Model

One can, of course, make changes to the basic I&F model to overcome the shortcomings just discussed. When this is done, the result is no longer an I&F model. We close this chapter by taking a look at one such modification that has enjoyed some degree of popularity since its introduction in 1990: the Eckhorn model [ECKH1].

Figure 8.11 illustrates the basic Eckhorn model. It is a multi-compartment model consisting of two main types of compartments, called the "soma" compartment and the "dendrite" compartment. An Eckhorn neural unit (ENU) consists of one "soma" and at least one "dendrite" compartment. It allows for multiple "dendrite" compartments by means of a summing node placed just before the "soma" compartment. The "dendrite" compartment is composed of two distinct input sections called the ***feeding field*** and the ***linking field***. Each field has its own leaky integrator function, called the feeding field leaky integrator (FFLI) and the linking field leaky integrator (LFLI), respectively. Each leaky integrator is mathematically defined by equation (8.4) with appropriate changes to the variable names as indicated in figure 8.11. Both LI functions convert to difference equation form as described by equation (8.5). Generally, the FFLI and the LFLI have different time constants, denoted τ_{ff} and τ_{lf}, respectively. If an ENU has more than one "dendrite" the second and subsequent "dendrites" might or might not have linking fields.

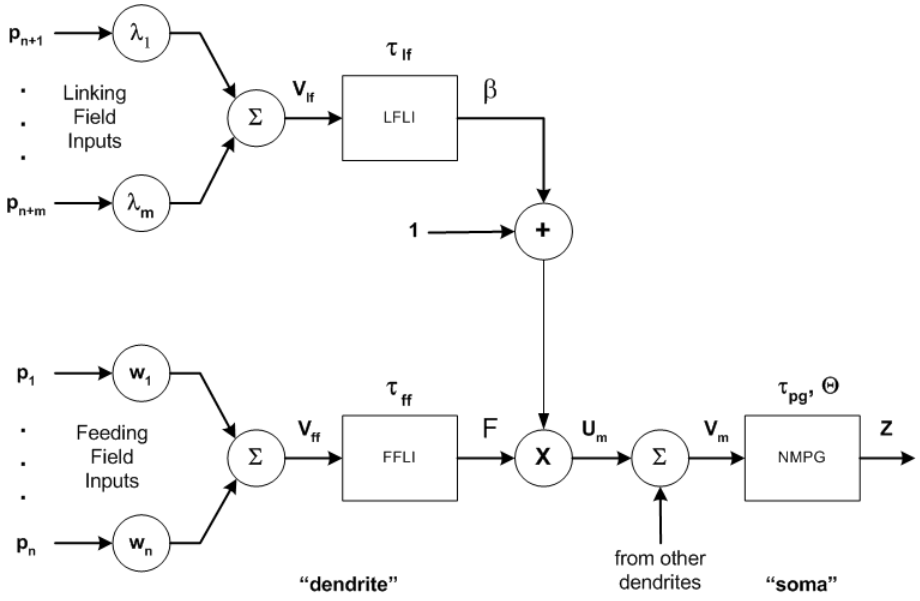

Figure 8.11: The basic Eckhorn model

The "soma" compartment consists of a variable firing threshold function, Θ, and a pulse generator, commonly called a "neuromime pulse generator" or NMPG. The output Z of the ENU is zero when the neuromime input $V_m < \Theta$. In difference equation form, letting Δt denote the simulation step size, the variable threshold function is defined by

$$\Theta_v(t+\Delta t) = \begin{cases} \Theta_v(t)\cdot\exp\left(-\Delta t/\tau_{pg}\right), & V_m(t) < \Theta(t) \\ A, & V_m(t) \geq \Theta(t) \end{cases} \qquad (8.13)$$
$$\Theta(t+\Delta t) = \Theta_0 + \Theta_v(t+\Delta t)$$

where Θ_0 is a constant establishing the minimum firing threshold, A is a constant establishing the maximum threshold value $\Theta_0 + A$, and τ_{pg} is a time constant determining the relative refractory interval of the ENU. $\Theta_v(t)$ is the time-varying component of the threshold.

The output signal for the "dendrite" compartment, U_m, is the product of the FFLI output $F(t)$ (as described by (8.5) with the appropriate changes of variables denoted in figure 8.11), and the offset output, $1 + \beta$, of the LFLI. $\beta(t)$, of course, is also given by (8.5) with appropriate change of variables. Since $U_m(t) = (\beta(t) + 1) \cdot F(t)$, it is easily seen that the "dendrite" output is zero if F is zero, which means linking field activity by itself can never cause the ENU to fire. There must be stimulation of the feeding field inputs to evoke a firing response. The NMPG input, V_m, is merely the sum of all the "dendrite" outputs in the ENU.

The effect of linking field activity is more or less equivalent to a short-term modulation of all the weights w_i, i = 1 to n, in the feeding field summing tree. As is the case for the I&F proxy, all the feeding field and linking field input signals are impulses. Linking field modulation of the feeding field pathway is therefore a means for implementing Malsburg's "synaptic modulation" on what he called the "fast time scale." Technically, this is not what Malsburg has in mind for his "synaptic modulation on the psychological time scale," and thus is not the same as his "dynamic link." Rather, the linking field signals provide a method for approximating a "measure of correlation" among ENUs interconnected through linking field connections. More precisely, the linking field tends to *produce* synchronized firing in groups of ENUs connected by linking fields based on the feeding field activities converging on each ENU in the group.

> To understand the general idea of temporal binding, the exact nature of the signal fluctuations is not relevant. . . Although much of the discussion of signal correlation focuses on the binary case, the correlation of just two neurons, it should be emphasized that the much more relevant and important type of event concerns correlations of higher order: the simultaneous firing of larger groups of neurons. The reason is this: for correlations to play an important role in the brain it must be possible for them to be evaluated quickly and reliably [MALS5].

Those who oppose and criticize the correlation theory of brain function often make the point that the mathematical evaluation of correlation functions, as these are defined in stochastic signal

processing theory, necessarily requires a significantly long record of firing patterns be evaluated. This mathematical requirement is seen as being at odds with psychological findings on human and animal reaction times. It is a tough argument to refute. However, merely because science found out about correlation functions long before anyone had the idea that synchronized firing patterns were important in brain function, this does not mean brain function necessarily has to use correlation functions. The Eckhorn linking field is an alternate way of looking at the problem, i.e. it provides a mechanism for focusing the signal processing on firing events occurring together within a brief time period. While this is not "correlation" in the mathematical sense, it may very well be an appropriate way of expressing the *functional end result* that gives the correlation theory of brain function its basic theoretical importance. Put another way, *linking fields promote binding* among ENUs.

Another key difference between the Eckhorn model and the I&F model is this: the Eckhorn model has no reset function. The FFLI and LFLI are not summarily "discharged" by the firing of an output pulse by the NMPG. Refractory effects are produced instead by the NMPG function. The lack of a reset function in the ENU is not biological; when a neuron fires an AP this action does lead to the opening of VGCs in the neuron which *do* result in the flushing of excess charge in the soma (the equivalent of discharging the capacitor in the H-H and Wilson models). The Eckhorn model is not a neuron model (under the definition used in this text for neuron models; this has not kept authors from calling the ENU a "neuron" in the literature, but it is bad terminology nonetheless). The Eckhorn model is a population model acting as proxy for large groups of neurons.

The absence of a reset function in the Eckhorn model changes its firing pattern characteristics in some important ways from those of the I&F model. One of the important changes is this: the absence of a reset function *makes it possible for the ENU to fire in bursts*. We saw earlier that bursting simply cannot occur in the I&F model because of its reset function. The absence of the reset function can lead to bursting by the ENU in the following way.

Let us suppose that at $t = t_0$ the NMPG input $V_m(t_0) = V_0$ exceeds threshold and the NMPG fires an output pulse. In the differential equation (continuous time) form of the Eckhorn dynamics this action sets $\Theta(t) = \Theta_0 + A$. Assuming no further pulses are applied to the input of the ENU for $t > t_0$, we have

$$\begin{aligned} V_m(t) &= V_0 \cdot \exp\left[-(t - t_0)/\tau_{ff}\right] \\ \Theta(t) &= \Theta_0 + A \cdot \exp\left[-(t - t_0)/\tau_{pg}\right] \end{aligned} \tag{8.14}$$

where the linking field dynamics have been neglected for the sake of simplicity. Right at $t = t_0$ we

may assume $\Theta(t) > V_m(t)$. If this condition maintains for all $t > t_0$ the NMPG will not fire again and a single output pulse will be produced. However, if for some $t = t_0 + T$ we have

$$V_0 \cdot \exp\left(-T/\tau_{ff}\right) = \Theta_0 + A \cdot \exp\left(-T/\tau_{pg}\right) \tag{8.15}$$

the NMPG will re-trigger and fire a second output pulse even in the absence of new input signals.

(8.15) is a nonlinear equation with no known closed-form expression for finding T given the values of the other parameters and initial input V_0. However, numerical solutions are easily obtained for it using either numerical or graphical methods. If a second output pulse is produced at some time spacing T, the analysis can be repeated to find out if there is some new time spacing T_2 at which the NMPG will produce a third output pulse (again, in the absence of new input signals to the ENU), and so on. By appropriate choices of the NMPG parameters and the time constant of the FFLI, it is possible to design a NMPG model capable of producing V_0-dependent burst firing patterns. This is, in principle, an important capability the ENU possesses that the I&F model does not.

The absence of the reset function also has a significant effect on the output firing rate of the ENU. With the I&F model, the output firing rate in response to a steady tetanus is always less than the firing rate of the tetanus unless the I&F neuron is operating in the all-pass response mode. The same is not necessarily true for the ENU. Like the I&F, the ENU has an all-pass response mode defined by the condition $V_{ff} > \tau_{ff} \cdot \Theta_0$, where V_{ff} has the same role as V_T in our earlier discussion and where the effect of the linking field is being ignored. For V_{ff} below this value and in the absence of linking field modulation, the ENU responds in a high-pass filter mode for the first M pulses in a tetanus with period T, responding with an output pulse on the $(M + 1)^{st}$ pulse of the tetanus. The value of M for this *first* output pulse is given by (8.10) with appropriate changes in the variable names.

However, the ENU response to the $(M + 2)^{nd}$ and later pulses in the tetanus is different from that of the I&F because the FFLI is not reset. Letting $\rho = \exp[\, -T/\tau_{ff}\,]$, the FFLI output for the m^{th} succeeding pulse in the tetanus is

$$U\left[(M+m)T\right] = \frac{V_{ff}}{\tau_{ff}} \frac{1-\rho^{M+m+1}}{1-\rho}$$

In the absence of other "dendrites," this will also be the value of V_m. Assuming the ENU is not responding in a bursting mode, the simulation threshold at this time is

$$\Theta\left((M+m)T\right) = \Theta_0 + A \cdot \exp\left(-\left(mT - \Delta t\right)/\tau_{pg}\right), \quad m > 0.$$

If V_m exceeds this value of Θ, the ENU will fire again. An important special case occurs when τ_{pg} is several times smaller than T but still large enough to prevent bursting mode operation. In this case, Θ will have returned to approximately Θ_0 for $m = 1$, and since $V_m[(M + 1)T]$ will be larger than, or at worst equal to, $V_m(MT)$, the ENU will fire again upon arrival of this input volley and will continue to fire for every successive input volley (unless inhibited by the arrival of inhibitory input pulses). Therefore, after suppressing the first M pulses in the tetanus, *the ENU thereafter responds with an output pulse for each succeeding input pulse volley and with the same firing rate as the input tetanus*. It can *relay* the tetanus (except for its first M volleys), which is something the I&F model is incapable of doing except in all-pass response mode.

Up to this point in the discussion, the contribution of the linking field has been neglected. Now let an input volley arrive at the linking field inputs at time $t = t_0$. In the absence of further linking field volleys, the LFLI response is

$$\beta(t) = \frac{V_{lf}}{\tau_{lf}} \cdot \exp\left[-(t - t_0)/\tau_{lf}\right]$$

and $U_m(t) = F(t) \cdot (1 + \beta(t))$. Thus, the feeding field signal is modulated by an exponentially decaying linking field signal, which produces what is called the ***linking window***. It is obvious that the linking field contribution is capable of boosting an otherwise sub-threshold value of F above the threshold value of the NMPG. This is the means by which ENU synchronization in a neural network is produced by the Eckhorn model. In typical ENU networks reported in the literature, the linking field time constant is usually about one order of magnitude smaller than the feeding field time constant, and so the linking field window is highly selective to feeding field volleys occurring within just a few milliseconds on either side of the linking field volley.

Because the ENU supports multiple "dendrites," it is possible to assign different time constants for different feeding field and linking field pathways (at the cost of additional computation for the model). In particular, it is often the case that one desires *inhibitory* feeding field inputs to have a different feeding field time constant from that of the excitatory feeding field signals. A second "dendrite" is therefore often used (with negative values assigned to the w_i) to implement the inhibitory pathways. It is also common for the "inhibitory dendrite" to omit the linking field, which saves half the incremental computational cost of adding this "dendrite."

§ 3.5 The Eckhorn Model in Comparison with the Rulkov Model

The per-iteration-step computational cost of a one-dendrite ENU is approximately the same as that of a Rulkov neuron when the cost of computing synapse signals in the latter is included. This

naturally raises the question: Why would one use one of these models rather than the other in a given modeling situation? Several considerations play a role in making this decision.

Although the Rulkov model is a mimic model separated from physiology by the abstract nature of its parameters and variables, it is nonetheless closer to biological mechanism than is the Eckhorn model. Although calibrating a Rulkov model to an approximation model (e.g. Wilson's schema) or to the H-H model is by no means a trivial undertaking, when this is done it is quite apparent what parametric decisions were made in support of which biological mechanism playing through to input-output function.

In contrast, there are functional features of the Eckhorn model that have no presently known and clearly understood linkage to physiological mechanism. The absence of a reset function in the ENU is not matched by anything at the neuron level, although it might be something characteristic of a netlet as a whole. (This really does remain to be seen at our present state of knowledge). There is nothing presently established about a neuron mechanism that corresponds to the linking field of the Eckhorn model. Typical values for τ_{lf} employed in Eckhorn-based neural network models are far too small to be accounted to known metabotropic mechanisms. It might be the case that a linking field function could result from the interplay of AMPA-dominated synapses and NMDA-dominated synapses [WELL1], but this is presently no more than a conjecture. Despite its use of such suggestive names as "dendrite" and "soma," the ENU remains an abstract model at the level of population proxy modeling.

Although the per-iteration computational cost of the Eckhorn and Rulkov models are quite comparable, the Eckhorn model lends itself easily to larger values for iteration step size Δt because an exact conversion (within very mild constraints) between its differential and difference equation forms can be made. In contrast, a map-based model such as Rulkov's has a fixed equivalent Δt. This is not to say that other Δt equivalences cannot be produced by map-based techniques, but it *is* to say this is not easily accomplished. Consequently, the network modeling cost of computation using the Eckhorn model has a slight edge, all else being equal, over the Rulkov model. On the other hand, the distance the Eckhorn model stands from biological mechanism produces a natural niche for applying the Rulkov model, namely in the modeling of netlets and more complex functional microcircuits and linking them to underlying models that capture the functional properties of the physiological neuron. It does not seem at all unlikely that when an explanation for the mechanistic underpinning of the Eckhorn linking field is finally obtained, the Rulkov model will play a central role in uncovering this explanation.

Finally, the Eckhorn model better lends itself to the *design* of signal processing functions at the neural network level than does the Rulkov model. Although it is true that some of the design

equations (e.g. for producing burst firing, or for determining a desired *T:M* response relationship for the ENU) require numerical solution methods, the expressions involved in this are cast in terms of time constants, thresholds, and other intuitively appealing parameters one can lean upon to guide the design development. In contrast, the Rulkov parameters are a level of abstraction removed from this *functional* aspect of neural network design. Put another way, design equations are relatively easy to derive for the Eckhorn model, relatively difficult to derive for the Rulkov model. Because experimental findings descriptive of network-level function are frequently reported in the literature (e.g. through such techniques as microprobe arrays) a significant amount of the literature data first implicates network-level function, and this makes *scientific* reduction rather than model order reduction (i.e., going from an Eckhorn network to Rulkov netlets) an important research direction for computational neuroscience.

§ 4. Chapter Summary

This chapter has covered much ground, so a few closing remarks in summation are in order. Model systems have been examined here ranging from approximation modeling very close to the level of physiological mechanism, through mimic modeling at the functional microcircuit and netlet level, to, finally, population proxy modeling at the neural network level.

At the lower levels of the modeling hierarchy, the question, "What is meant by 'average' neuron?" was raised. Considerable discussion was presented regarding the considerations that go into deciding what sort of modeling can be carried out to achieve an "average" dynamical response. The principal finding here is that it is possible to use a single signaling type model to capture many dynamical effects, and to use phenomenological synaptic weight factors to introduce variability into the neuron population.

The major research issue of neural encoding was introduced, and the possibility was raised that different neuron types may be key to understanding the long-sought "neural code." Although this research area is very, very active – and has been for a very long time – it is still true that neuroscience is not in possession of a paradigm accepted by all researchers for "cracking the code" as this putative code is actually carried out in brain function.

The difference between stereotyped and non-stereotyped function was introduced. A brief discussion was given over to the idea that perhaps "higher" brain function is a consequence of interactions among many much more stereotyped "neural machines." This is a view that tends to be opposed by an older, more traditional school of thought first established in the 1960s and later resurrected by what is widely known as the "parallel distributed processing" (PDP) school. PDP theorists, who tend to work at the psychophysical and psychological levels of neuroscience,

promote the view that neural networks should be seen as "universal function approximators," a point of view that can be traced back to the very early work in automaton theory by Alan Turing. The notion of regarding neural networks as universal function approximators is strongly opposed by Minsky and Papert, and also by dynamic link architecture theorists such as Malsburg:

> There is a widespread opinion that classical neural networks are a universal medium with no limits to their abilities and that consequently they are not subject to the binding problem. . . But what does universality buy? It can be compared to the universality of a pen and sufficiently many sheets of white paper as a universal medium for formulating novels. You still have to write them. Over time, the field of Artificial Intelligence discovered that it is not a practical task at all to write a program that emulates the capabilities of the brain. It is becoming increasingly clear that the only goal we can hope for is to establish a system that constitutes a basis for self-organization and learning . . .
>
> Brain theorists realized this in the late 1950s and modified the McCulloch and Pitts' networks to accommodate self-organization and learning. . . However, these changes may have come at a price: it is not clear whether neural networks are universal in any sense, although the community seems to have inherited the implicit belief that they are and that any brain function can be modeled on the basis of those few abstractions from the real nervous system that went into the formulation of neural networks. It is not even clear how to formulate a new universality theorem. . . . But how can we characterize brain problems in any general and satisfactory way? It would be foolish to argue that "this is a particular problem I have solved on the basis of classical neural networks, which proves that all of them can be solved this way" [MALS5].

Finally, two network-level models, serving as proxies for populations of neurons, were introduced. Properties of these models, and where they stand in regard to lower level netlet models (particularly Rulkov's models) and higher network signal processing functions (synchronized signaling, correlations), were discussed. The overriding theme covering all these various models has been the theme of ***hierarchical modeling***, which is made necessary by the ever-present practical reality of computational costs and Weinberg's "Square Law of Computation."

Something touched upon very briefly in this chapter was the notion that a description of mathematical ***structure*** *for* neuronal function (rather than merely a mathematical description *of* a neural network) is needed for theory to really comprehend neural systems. The main idea brought up here was the notion of characterizing signals and networks as *sets* and *transformations of sets*. This topic has received very, very little attention in the present day computational neuroscience literature (compared to the other topics covered), but it is a key and fundamental topic in its own right. We will make a brief digression in chapter 10 and say a bit more about this.

Exercises

1. There are many different kinds of averages. Three of the most common are called the mean, the median, and the mode. The mean is also called the "arithmetic average" and is the one you probably use the most. The median is the value for which half the cases have a smaller value and half the case have a larger value. The mode is the value most

frequently encountered. Given the series of numbers {1, 1, 2, 2, 2, 2, 3, 4, 4, 4, 5, 5, 6}, find the mean, median, and mode averages of this set.

2. The values of the mean, median, and mode averages come out pretty close to each other when the population follows a normal distribution (i.e., the well known bell-shaped curve you are probably very familiar with; the normal distribution is also called the Gaussian probability distribution). However, when the distribution is skewed, the mean, median, and mode averages can be very different. Suppose a population is given that is characterized by the following numbers: {2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 3, 3.7, 3.7, 3.7, 3.7, 5, 5, 5, 5.7, 10, 10, 15, 45}. Find the mean, median, and mode averages. If these numbers represent the monthly pay (in thousands of dollars) for the people who work in a particular small company, what do the three different averages tell you?

3. Data was given in chapter 6 describing the population of different kinds of neurons found in the neocortex. For purposes of modeling the neocortex, what kind of average (mean, median, or mode) do you think would be the most useful in describing an "average" cortical neuron? Do you think that any <u>one</u> average is adequate for describing this? Explain why or why not. How would <u>you</u> use the idea of "averages" to set up a model of the neocortex?

4. A curve fit model (called a "regression model" by statisticians) fits a mathematical function (often, but not always, a polynomial function) to experimental data. The variables in the model are "factors" that were varied to produce different experimental outcomes. A good regression model manages to produce accurate "predictions" of what outcomes one will see for other combinations of factors provided these all fall within the range of values that was used to obtain the experimental curve fit. It is also well known that regression models usually "break down" (fail to accurately predict) outcomes for factors that exceed the range of values used in obtaining the fit. In contrast, a good phenomenological model is said to have "predictive power" better than that of a curve fit. Discuss what this distinction is meant to imply about the difference between phenomenological models and regression models.

5. Section §1.1 provided an analysis for matching up the three Wilson model test cases against the model distribution statistics in Table III. This analysis might be called "semi-quantitative" because, while it used numbers to compare ranges, the decisions on how to "match up" one model to the other were qualitative and "feature-based" (i.e., "this one looks the most like that one"). A semi-quantitative analysis is always open to counter-arguments simply because the decision is not quantitatively "hard and fast." Other interpretations are almost always possible. Carry out your own analysis for this model-to-model comparison and state your own conclusions on how you would "match up" the two models. Explain your reasoning.

6. Explain in your own words what "coincidence advantage" is. What is the evidence in §1.1 and §1.2 that points to the existence of coincidence advantage?

7. Let Y be an $n\times1$ vector of experimental data and let B be an $m\times1$ vector of model parameters. Let X be an $n\times m$ matrix of experimental factors (such as $n_{syn} \cdot g_{\max}/C_m$ and C_m in §1.2). A linear regression model is a model of the form $Y = X\cdot B$. Y and X are known and B is to be determined. Show that the solution is $B = (X^T X)^{-1} X^T Y$. This solution is called the least-squares fit. Find X for the complete and reduced-order models of §1.2.

8. Modify your computer simulation program for the RS-type Wilson model so that you can repeat the simulation exercise in §1.2. Then obtain additional sub-threshold data points so that you can obtain a regression fit to peak sub-threshold EPSP responses up to about 20

mV. Your goal is to obtain an adequate reduced model similar to that of (8.2), possibly with the addition of a cubic term if you find such a term statistically significant. How closely is your model able to predict the n_{syn} firing threshold?

9. In your own words, describe what the term "fan-in" denotes and discuss how fan-in might be relevant for modeling functional microcircuits using "average" neurons.

10. Why would the "wide, slow base" of an action potential, as mentioned in regard to figure 8.6, be indicative of a stimulus that just barely reaches the firing threshold of the neuron?

11. Discuss in what way or ways the SG50 from Abeles' statistical model and an "SG50" for Wilson's deterministic model are similar, and in what way or ways they are different. Hint: You are being asked to think about what the difference is between a statistical model and a deterministic model, and to reflect on how hard-to-predict nonlinear effects in a deterministic model might be regarded as "being like" random variables.

12. Modify your computer simulation model of the Wilson RS-type neuron to co-localize both AMPA and NMDA channels together in the same synapse. Your model should allow for variable ratios of AMPA to NMDA receptors, and allow for different $g_{\max}$ values to be assigned to each. Simulate three test cases: (1) AMPA channels only; (2) NMDA channels only; and (3) equal AMPA and NMDA channel densities with equal $g_{\max}$ values of 200 pS per synapse. Find and compare the "SG50" value of n_{syn} for cases (1) and (3) for $C_m = 0.51$ fF.

13. The spinal cord contains neural networks that produce what are known as reflex responses. Examples include kicking your leg in response to a gentle tap from the doctor's mallet and snatching your hand away when you touch something very hot. Would these sorts of reflex responses constitute stereotyped responses? Explain your reasoning.

14. Here is a conjecture for you to consider: "The more stereotyped the response of a netlet or network is, the more model order reduction can be applied to it." Decide whether or not you agree with this conjecture and write a short essay presenting and defending your view.

15. The existence of vertically-organized functional columns in the neocortex implies it is meaningful and correct to model neocortical function by means of two-dimensional chains of functional columns. Explain why this is so.

16. Explain what "columnar arrangements based on the distribution of afferent fibers are a general feature of cortical organization" means.

17. What sort of two-dimensional chain model might be suggested by Szentágothai's model of figure 8.9? (You do not need to assume the functional column is perfectly cylindrical). Draw a two-dimensional sketch of your conjectured chain organization model.

18. Show that equation (8.5) is an exact solution of (8.4) under the specified conditions.

19. Derive equation (8.6). Plot R as a function of V/Θ for the interval $V/\Theta = 1.1$ to 10 in steps of 0.1 using $\tau = 1$ for $t_r = 0$, 0.5, 1, 2, and 5. Describe the effect of t_r on R.

20. Write a computer simulation program for the I&F proxy model and verify that the model responds to tetanus inputs in the way described in the text.

21. Derive equation (8.9).

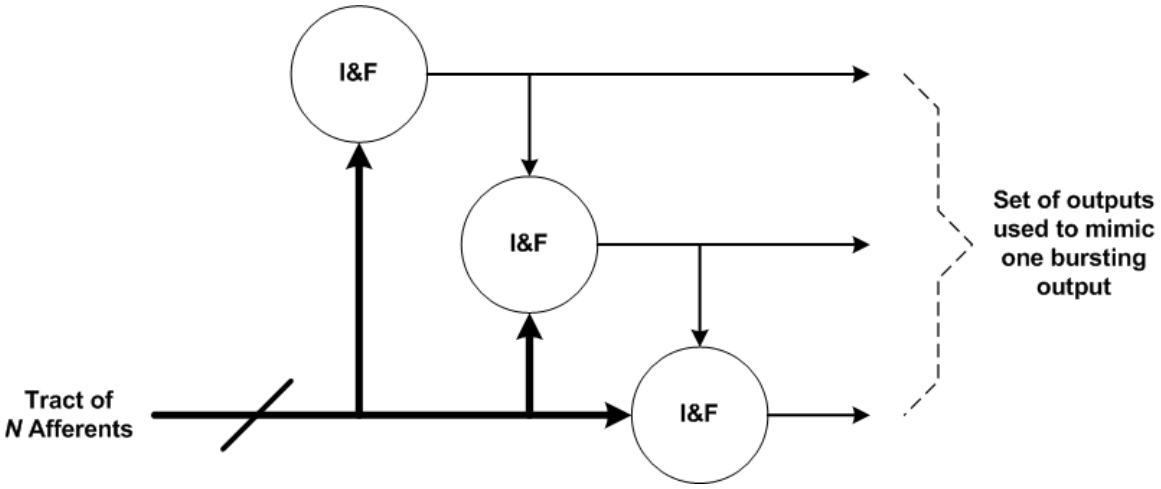

Figure P8.22

22. Bursting responses can be mathematically imitated by using a small network of I&F proxies such as that shown in Figure P8.22. Each I&F cell receives the same set of N afferent inputs. Different weightings are applied by each I&F cell so that successive I&F cells in the chain require stimulation from the previous I&F cell before they can fire. Design a set of weightings and time constants so that the network can generate one, two, or three pulse responses depending on the level of stimulation from the afferent tract. Demonstrate that your design works by simulating it. Comment on any shortcomings or disadvantages this method has.

23. One of the simplest methods to evaluate equation (8.15) is to plot its left-hand and right-hand sides on the same graph and find where they intersect. Let the ENU have the following parameters: τ_{ff} = 15 ms, τ_{lf} = 7.5 ms, Θ_0 = 0.5, $A = 20 \cdot \Theta_0$. Determine if re-triggering of the NMPG occurs, and, if so, when, for (a) $V_0 = 10 \cdot \Theta_0$ (b) $V_0 = 5 \cdot \Theta_0$. Use time steps of 0.1 ms and evaluate the functions over the range from $T = 0$ to $T = 50$ ms.

24. Typical values for Eckhorn model parameters used in ENU network models reported in the literature are on the order of the following: $\tau_{ff} \approx 10$ ms, $\tau_{lf} \approx 1$ ms, $\tau_{pg} \approx 7.5$ ms, $A \approx 50$, $\Theta_0 \approx 0.5$. Typically the linking field weights are on the order of about 10 times greater than the feeding field weights. These values are merely examples and different modeling situations will often call for different parameter settings. Write a computer program to simulate the ENU and verify that the model responds to tetanus inputs in the manner described in the text.

25. Inhibitory inputs to a one-dendrite ENU can be implemented by using negative-valued weights for the inhibitory inputs. However, this means that the inhibitory signals have the same time constant as the excitatory signals. Add a second inhibitory dendrite to your Eckhorn simulation program from exercise 24. The second dendrite does not require a linking field pathway. Set $\tau_{ff} \approx 20$ ms for the inhibitory dendrite and evaluate the effect of inhibitory inputs on the tetanus response of the ENU.