P. Cisek, T. Drew & J.F. Kalaska (Eds.)
Progress in Brain Research, Vol. 165
ISSN 0079-6123

CHAPTER 2

A simple growth model constructs critical avalanche networks

L.F. Abbott[1,*] and R. Rohrkemper[2]

[1]*Department of Physiology and Cellular Biophysics, Center for Neurobiology and Behavior, Columbia University College of Physicians and Surgeons, New York, NY 10032-2695, USA*
[2]*Physics Department, Institute of Neuroinformatics, Swiss Federal Institute of Technology, Zürich CH-8057, Switzerland*

Abstract: Neurons recorded from electrode arrays show a remarkable scaling property in their bursts of spontaneous activity, referred to as "avalanches" (Beggs and Plenz, 2003, 2004). Such scaling suggests a critical property in the coupling of these circuits. We show that similar scaling laws can arise in a simple model for the growth of neuronal processes. In the model (Van Ooyen and Van Pelt, 1994, 1996), the spatial range of the processes extending from each neuron is represented by a circle that grows or shrinks as a function of the average intracellular calcium concentration. Neurons interact when the circles corresponding to their processes intersect, with a strength proportional to the area of overlap.

Keywords: network activity; homeostasis; plasticity; network development

Introduction

Theoretical (also known as computational) neuroscience seeks to use mathematical analysis and computer simulation to link the anatomical and physiological properties of neural circuits to behavioral and cognitive functions. Often, researchers working in this field have a general principle of circuit design or a computational mechanism in mind when they start to work on a project. For the project to be described here, the general issue concerns the connectivity of neural circuits. For all but the smallest of neural circuits, we typically do not have a circuit diagram of synaptic connectivity or a list of synaptic strengths. How can we model a circuit when we are ignorant of such basic facts about its structure? One answer is to approach the problem statistically, put in as much as we know and essentially average over the rest. Another approach, and the one that inspires this work, is to hope that we can uncover properties of a neural circuit from basic principles of synapse formation and plasticity. In other words, if we knew the rules by which neural circuits develop, maintain themselves, and change in response to activity, we could work out their architecture on the basis of that knowledge. To this end, we need to uncover the basic rules and principles by which neural circuits construct themselves.

When neurons are removed from the brain and grown in culture, they change from disassociated neurons into reconnected networks or, in the case of slice cultures, from brain slices to essentially two-dimensional neural circuits. These re-development processes provide an excellent opportunity for exploring basic principles of circuit formation. Using slice cultures from rat cortex (and also acute slices),

*Corresponding author. Tel.: +1 212-543-5070;
Fax: +1 212-543-5797; E-mail: lfa2103@columbia.edu

DOI: 10.1016/S0079-6123(06)65002-4

Beggs and Plenz (2003, 2004) uncovered an intriguing property of networks of neurons developed in this way. By growing neural circuits on electrode arrays, they were able to record activity over long periods of time and accumulate a lot of data on the statistical properties of the activity patterns that arise spontaneously in such networks. Of particular interest are the observations of scaling behavior and criticality. These results provide the inspiration for the model we construct and study here.

The networks recorded by Beggs and Plenz (2003, 2004) are often silent, but silent periods are punctuated by spontaneous bursts of activity observed on variable numbers of electrodes for different periods of time. Beggs and Plenz called these bursts avalanches. To define and parameterize neural avalanches, they divided time into bins of size t_{bin} through a procedure that selects an optimal size. Here, we simply use $t_{\text{bin}} = 10\,\text{ms}$, typical of the values they used. An avalanche is defined as an event in which activity is observed on at least one electrode for a contiguous sequence of time bins, bracketed before and after by at least one bin of silence on all electrodes. We use an identical definition here, except that electrode activity is replaced by neuronal activity, because our model has no electrodes and we can easily monitor each neuron we simulate.

The results of Beggs and Plenz (2003, 2004) of particular importance for our study are histograms characterizing both the durations and sizes of the avalanches they recorded. Duration was determined by counting the number of consecutive bins within an avalanche. Size was measured either in terms of the number of electrodes on which activity was recorded during an avalanche, or by a measure of the total signal seen on all electrodes during the course of an avalanche. In our modeling work, we measure the size of an avalanche by counting the total number of action potentials generated during its time course.

The histograms of duration and size constructed from the data revealed a fascinating property (Beggs and Plenz, 2003, 2004; Fig. 1); both were of a power-law form. The number of events of a given size fell as the size to the $-3/2$ power, and the number of events of a given duration fell as the duration to the -2 power. Power-law distributions are interesting because they contain no natural scale. For example, in this context we might expect the typical size of a neuronal dendritic tree or axonal arbor (around 100 μm) to set the spatial scale for avalanches. Similarly, we might expect a typical membrane time constant of around 10 ms to set the scale for avalanche durations. If this were true, the distributions should be exponential rather than power-law. Power-law distributions indicate that these networks can, at least occasionally, produce activity patterns that are much larger and much long-lasting that we would have expected. This is what makes power-law distributions so interesting. Another intriguing feature is that power-law behavior typically arises in systems

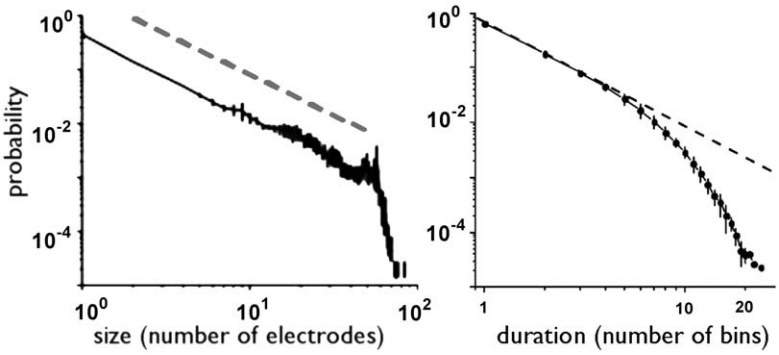

Fig. 1. Results of Beggs and Plenz on avalanche distributions. Left: probability of avalanches of different spatial sizes. The dashed line corresponds to a $-3/2$ power. Right: probability of avalanches of different durations. The dashed line corresponds to a -2 power. (Adapted with permission from Beggs and Plenz, 2004).

when they are critical, meaning that they are close to a transition in behavior. Thus, power laws arise when systems are specially configured.

Beggs and Plenz (2003, 2004) went on to note that the powers they observed, $-3/2$ and -2, are the same as those that arise in a very simple model (Zapperi et al., 1995). In this model, each neuron connects to n other neurons and, if it fires an action potential, causes each of its targets to fire with probability p. If $p<1/n$, activity in this model tends to die out, and if $p>1/n$ it tends to blow up. If $p=1/n$, on the other hand, this simple model produces distributions with the same power-law dependence and same powers as those observed in the data. The condition $p=1/n$ implies that every neuron that fires an action potential causes, on average, one other neuron to fire. This is critical in the sense discussed above that smaller values of p tend to induce patterns of activity that die out over time, and larger values of p tend to produce exploding bursts of activity. Thus, the results from these array recordings lead to the puzzle of how networks develop and maintain patterns of connectivity that satisfy this criticality condition. Do neurons somehow count the number of other neurons they project to and adjust the strengths of their synapses in inverse proportion to this number? If so, what would be the biophysical substrate for such a computation and adjustment (Teramae and Fukai, 2007)?

To address these questions, we made use of a model of neuronal circuit growth due to Van Ooyen and Van Pelt (1994, 1996). The model is simple, but here simplicity is exactly the point. We ask, in place of the above questions, whether a simple, biophysically plausible mechanism could account for the power-law behavior seen in the avalanche histograms without requiring any counting of synapses or criticality calculations. We are not proposing that the model we present is realistic, but rather use it to show that adjusting a network to be critical may not be as difficult as it would first appear.

The model

Following the work of (Van Ooyen and Van Pelt (1994, 1996); for reviews, see Van Ooyen, 2001, 2003), our model consists of N neurons positioned at random locations within a square region. The length and width of this square defines 1 unit of length. We can think of each location as the position of the soma of a neuron. The axonal and dendritic processes of each neuron are characterized by a circle drawn around its location. The size of this circle represents the extent of the processes projecting from the centrally located soma. Neurons interact synaptically when the circles representing their processes overlap, and the strength of the coupling is proportional to the area of overlap between these two circles. This is reasonable because synapses form in areas where neuronal processes intersect, and more intersections are likely to result in more synapses. All synaptic connections are excitatory.

The critical component of the model is the growth rule that determines how the process-defining circles expand or contract as a function of neuronal activity. The rule is simple: high levels of activity, which signify excessively strong excitation, cause the neuronal circle to contract, and low levels of activity, signifying insufficient excitation, cause it to grow. The initial sizes of the circles are chosen randomly and uniformly over the range from 0 to 0.05, in the units defined by the size of the square ''plating'' region.

Each neuron in the model is characterized by a firing rate and a radius, which is the radius of the circle defining the extent of its processes. Neuronal activity is generated by a Poisson spiking model on the basis of a computed firing rate. The firing rate for neuron i, where $i=1,\ 2,\ 3,\ldots,\ N$, relaxes exponentially to a background rate r_0 with a time constant τ_r according to

$$\tau_r \frac{\mathrm{d}r_i}{\mathrm{d}t} = r_0 - r_i. \tag{1}$$

We took $r_0 = 0.1$ Hz and $\tau_r = 5$ ms. The low background firing rate of 0.1 Hz is important to prevent the network from simply remaining silent. At every time step Δt, neuron i fires an action potential with probability $r_i \Delta t$. We took $\Delta t = 1$ ms. After a neuron fires an action potential, it is held in a refractory state in which it cannot fire for 20 ms.

Whenever another neuron, neuron j, fires an action potential, the firing rate of neuron i is

incremented by

$$r_i \rightarrow r_i + gA_{ij} \tag{2}$$

where A_{ij} is the overlap area between the two circles characterizing the processes of neurons i and j. In our simulations, the constant g, which sets the scale of synaptic strength in the model, is set to $g = 500\,\mathrm{Hz}$. This number is large because the overlap areas between the neurons are quite small in the units we are using.

The average level of activity of neuron i is monitored by a variable C_i that represents the internal calcium concentration in that neuron. C_i decays to zero exponentially,

$$\tau_C \frac{\mathrm{d}C_i}{\mathrm{d}t} = -C_i \tag{3}$$

and is incremented by one unit ($C_i \rightarrow C_i + 1$) whenever neuron i fires an action potential. This step size defines the unit of calcium concentration. The value of the time constant τ_C is not critical in what follows, but we took it to be 100 ms.

Two features make calcium a useful indicator of neuronal activity. First, resting calcium concentrations inside neurons are very small, but calcium enters the cell whenever the neuron fires an action potential. Because of this, the calcium concentration acts as an integrator of the action potential response and, for this reason, imaging calcium concentrations is a common way to monitor neuronal activity. Second, many molecules in a neuron are sensitive to the internal calcium concentrations, so this indicator can activate numerous biochemical cascades, including those responsible for growth.

The remaining equation in the model is the one that determines the contraction or growth of the radius a_i characterizing neuron i. This is

$$\frac{\mathrm{d}a_i}{\mathrm{d}t} = k(C_{\mathrm{target}} - C_i) \tag{4}$$

where k determines the rate of growth. We used a variety of values for k, but growth was always slow on the time scale of neuronal activity. We often started a run with a larger value of k ($k = 0.02\,\mathrm{s}^{-1}$) to speed up growth, but as an equilibrium state was reached we lowered this to $k = 0.002\,\mathrm{s}^{-1}$. The parameter C_{target} plays the dominant role in determining the behavior of the model. This sets a target level of calcium, and therefore a target level of activity, for the neurons. If activity is low so that $C_i < C_{\mathrm{target}}$, the above equation causes the processes from neuron i to grow (a_i increases) leading to more excitatory connections with other neurons and hence more activity. If activity is high so that $C_i > C_{\mathrm{target}}$, the processes will retract (a_i decreases) lowering the amount of excitation reaching neuron i. In this way, each neuron grows or contracts in an attempt to maintain the target level of calcium concentration ($C_i = C_{\mathrm{target}}$), which implies a certain target level of activity. We discuss the value of C_{target} more fully below, but $C_{\mathrm{target}} = 0.08$ was used to obtain the results in the figures we show.

Results

The left panel of Fig. 2 shows a typical configuration at the beginning of a run. In this case, 100 neurons have been located randomly with various radii, also chosen randomly. At this initial point, many of the neurons are disconnected or, at most, connected together in small clusters. Each neuron has a spontaneous firing rate of 0.1 Hz, even when isolated, so this network exhibits activity, but at a low level. Fig. 2 (left) shows a typical initial state of the model, but the results of running a model simulation are independent of the initial state unless a highly unlikely initial configuration (such as many neurons at the same position) limits the possibilities for developing connections through growth. The target calcium level we use, $C_{\mathrm{target}} = 0.08$, is larger than the average calcium level attained by the neurons in this initial configuration. Thus, when the simulation starts, the neurons (the circles in Fig. 2, left) grow larger.

As the neurons grow, they begin to form more and stronger connections, which causes the level of activity in the network to increase. Growth continues until the neurons are active enough to bring their average calcium concentrations near to the value C_{target}. At this point, the average rate of growth of the network goes to zero, but there are still small adjustments in the sizes of individual neurons. As neurons adjust their own radii, and

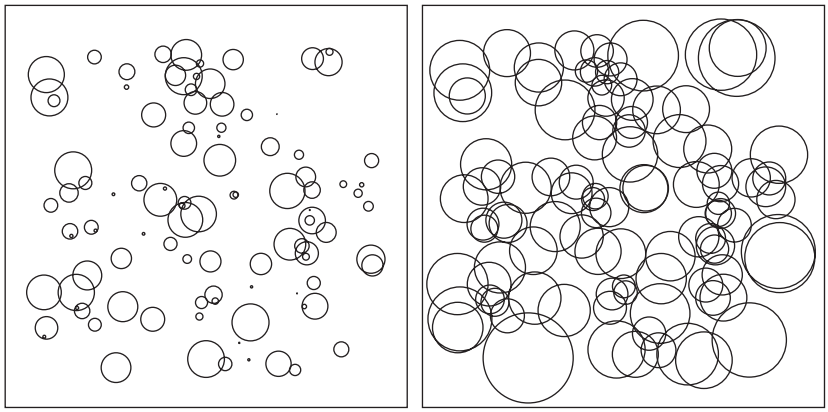

Fig. 2. Configuration of the model network before (left) and after (right) activity-dependent growth. Each circle represents the extent of the processes for one neuron. Neurons with overlapping circles are connected. Initially (left), the neurons are either uncoupled or coupled in small clusters. At equilibrium (right), the network is highly connected.

react to the adjustments of their neighbors, they eventually achieve a quasi-equilibrium point in which their time-averaged calcium concentrations remain close to C_{target}, with small fluctuations in their radii over time. From this point on, the network will remain in the particular configuration it has achieved indefinitely. This growth process has been described previously (Van Ooyen and Van Pelt, 1994, 1996; Abbott and Jensen, 1997). Our only modification on the original growth model of Van Ooyen and Van Pelt (1994, 1996) was to add Poisson spikes to their firing-rate model. The right panel of Fig. 2 shows the equilibrium configuration that arose from the initial configuration shown in the left panel.

The size of the small fluctuations in neuronal size about the equilibrium configuration is determined by the magnitude of the growth rate, k. Because growth processes are much slower than the processes generating activity in a network, we chose k to be as small as we could without requiring undue amounts of computer time to achieve equilibrium. The results we report are insensitive to the exact value of k.

Once the network has achieved an equilibrium configuration, we analyze its patterns of activity using the same approach as Beggs and Plenz (2003, 2004). In other words, we constructed histograms of the duration and total number of action potentials in periods of activity that were bracketed by 10 ms time bins in which no activity was observed. To assure that the resulting histograms reflect the dynamics of the network and not of the growth process, we shut off growth (set $k = 0$) while we accumulated data for the histograms, although for the small growth rate we use, this did not make any noticeable difference to the results.

Histograms of the durations and number of action potentials for the avalanches seen in the model at equilibrium are shown in Fig. 3. These are log-log plots, and the straight lines drawn indicate $-3/2$ (Fig. 3, left) and -2 (Fig. 3, right) power-law dependences. Over the range shown, the histograms follow the power-law dependences of a critical cascade model. As in the data (Beggs and Plenz, 2003, 2004), there are deviations for large, rare events due to finite-size effects.

Changing the initial size of the circles representing the neuronal processes in these simulations has no effect, because the growth rule simply expands small circles or shrinks large circles until they are in the equilibrium range. The model is, however, sensitive to the value of the target calcium concentration. The most sensitive result is the exponent of the power function describing the distribution of spike counts, as shown in the left panels of Figs. 1 and 3. The exponent for the distribution of durations is less sensitive. Fig. 4 shows how the spike count distribution exponent depends on C_{target} over a range of values from 0.04 to 1.2, with the value used for the previous figures, 0.08, in the middle of this range.

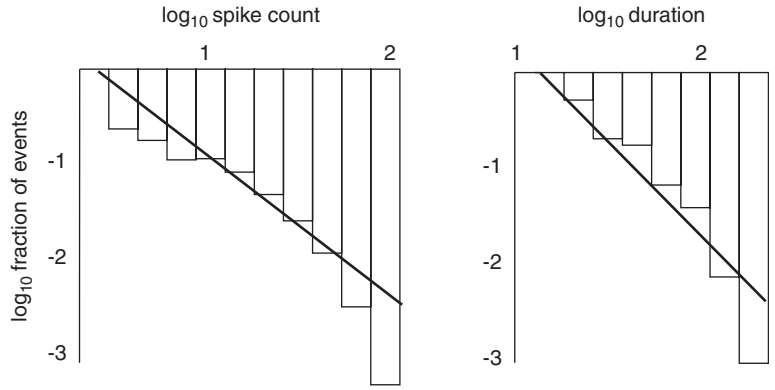

Fig. 3. Histograms of the fraction of avalanches with different numbers of spikes (left) and different durations (right). The plots are log-log and the lines indicate $-3/2$ (left) and -2 (right) powers.

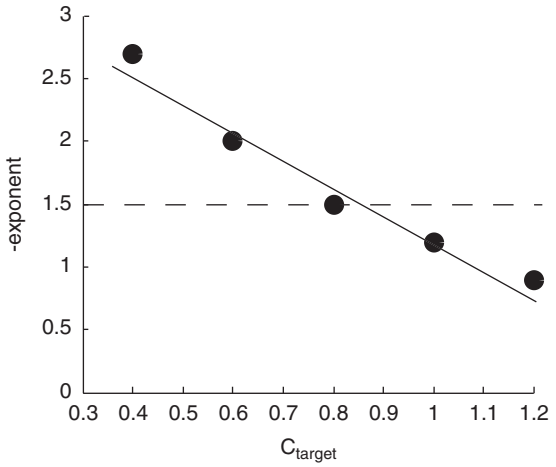

Fig. 4. Value of minus the exponent of the power function describing the spike count distribution as a function of the target calcium concentration. The value seen in the experiments, indicated by the dashed line, is 1.5, corresponding to $C_{\text{target}} = 0.08$. The solid line is drawn only to guide the eye.

Discussion

In our network model, the spontaneous level of activity for each neuron, 0.1 Hz, is insufficient to allow the internal calcium concentration to approach the target level we set. Therefore, disconnected neurons grow, and they can only reach an equilibrium size if they ''borrow'' activity from other neurons. Even the activity in small clusters is insufficient to halt growth. However, the target calcium concentration was set so that all-to-all connections or excessive large-scale firing over the entire network would produce internal calcium concentrations that exceed the target level and therefore induce process withdrawal. Therefore, the network is forced to find a middle ground in which individual neurons share activity in variable-sized groups, drawing excitation from both nearby and faraway neurons. This is what provides the potential for critical, power-law behavior.

The power-laws shown in Figs. 3 and 4 occur over a range of values of C_{target}, but they are not an inevitable consequence in the model. Values of C_{target} significantly higher than those we have used lead to an essentially flat distribution (over the finite range) of event sizes and durations. Smaller values lead to a shortage of large, long-lasting events.

The model we have considered warrants studying in more depth, and it can be extended in a number of ways. Obviously, inhibitory neurons should be added. In addition, it would be of interest to provide each neuron with two circles, one representing the extent of dendritic outgrowth and the other axonal. Separate growth rules would be needed for the two circles in this case. Finally, the axonal projections could be given both local extension, represented by a circle around the somatic location, and distal projections, represented by additional circles located away from the soma.

The fact that a simple growth rule can generate circuits with critical, power-law behavior suggests

that it could be the basis for developing interesting network models. We have only explored uncontrolled spontaneous activity, but the fact that this can occur over such a large range of sizes and durations makes the functional implications of these networks quite intriguing. If we can learn to grow circuits like this in which we can control the size and time scale of the activity, this could form a basis for building functional circuits that go beyond spontaneous activity to perform useful tasks.

Acknowledgments

Research supported by the National Science Foundation (IBN-0235463) and by an NIH Director's Pioneer Award, part of the NIH Roadmap for Medical Research, through grant number 5-DP1-OD114-02. We thank Tim Vogels and Joe Monaco for valuable input.

References

Abbott, L.F. and Jensen, O. (1997) Self-organizing circuits of model neurons. In: Bower J. (Ed.), Computational Neuroscience, Trends in Research 1997. Plenum, NY, pp. 227–230.

Beggs, J.M. and Plenz, D. (2003) Neuronal avalanches in neocortical circuits. J. Neurosci., 23: 11167–11177.

Beggs, J.M. and Plenz, D. (2004) Neuronal avalanches are diverse and precise activity patterns that are stable for many hours in cortical slice cultures. J. Neurosci., 24: 5216–5229.

Teramae, J.n. and Fukai, T. (2007) Local cortical circuit model inferred from power-law distributed neuronal avalanches. J. Comput. Neurosci., 22: 301–312.

Van Ooyen, A. (2001) Competition in the development of nerve connections: a review of models. Network, 12: R1–R47.

Van Ooyen, A. (Ed.). (2003) Modeling Neural Development. MIT Press, Cambridge, MA.

Van Ooyen, A. and Van Pelt, J. (1994) Activity-dependent outgrowth of neurons and overshoot phenomena in developing neural networks. J. Theor. Biol., 167: 27–43.

Van Ooyen, A. and Van Pelt, J. (1996) Complex periodic behaviour in a neural network model with activity-dependent neurite outgrowth. J. Theor. Biol., 179: 229–242.

Zapperi, S., Baekgaard, I.K. and Stanley, H.E. (1995) Self-organized branching processes: mean-field theory for avalanches. Phys. Rev. Lett., 75: 4071–4074.