The criticality hypothesis: how local cortical networks might optimize information processing

By John M. Beggs*

Department of Physics, Indiana University, Bloomington, IN 47405, USA

Early theoretical and simulation work independently undertaken by Packard, Langton and Kauffman suggested that adaptability and computational power would be optimized in systems at the 'edge of chaos', at a critical point in a phase transition between total randomness and boring order. This provocative hypothesis has received much attention, but biological experiments supporting it have been relatively few. Here, we review recent experiments on networks of cortical neurons, showing that they appear to be operating near the critical point. Simulation studies capture the main features of these data and suggest that criticality may allow cortical networks to optimize information processing. These simulations lead to predictions that could be tested in the near future, possibly providing further experimental evidence for the criticality hypothesis.

Keywords: edge of chaos; neural network; power law; criticality; cortex; avalanche

1. Introduction

One of the great scientific challenges of our time is to understand how the brain processes information. Current work suggests that individual neurons are relatively modest in their computational power, but that networks of neurons can collectively perform extremely complex operations. Any system that sheds light on emergent properties may therefore potentially advance our understanding of neural function.

Fortunately, emergent properties have received much attention in a wide variety of systems. Simulations of evolving agents led Packard (1988) to suggest that adaptability was optimized when agents approached the 'edge of chaos' (but see Mitchell *et al.* 1993). Langton (1990) studied cellular automata that could be tuned into ordered, critical or chaotic regimes, and concluded that computations were performed best by systems near the critical point between order and chaos. Kauffman and colleagues (e.g. Kauffman & Johnsen 1991; Kauffman 1993) modelled cellular protein interactions with random Boolean networks and showed that under selection pressure these networks would self-organize into a nearly critical state where perturbations propagated in the form of avalanches. Neural models have also picked up on this idea, and several authors have suggested that neural networks should operate close to the chaotic regime (Chialvo & Bak 1999;

*jmbeggs@indiana.edu

Bak & Chialvo 2001; Bertschinger & Natschlager 2004; Lin & Chen 2006; Legenstein & Maass 2007). Collectively, these simulations predict that living systems will self-organize to operate near a critical point. Hereafter, we will refer to this prediction as the 'criticality hypothesis'.

Despite the abundance of simulations, experimental evidence supporting the criticality hypothesis had been relatively sparse. Recent experiments, however, have begun to provide support for this idea in the area of neural networks (Beggs & Plenz 2003, 2004; Petermann *et al.* 2006; Stewart & Plenz 2006).

The organization of the rest of this paper is as follows. First, we will review experiments with living neural networks that support the criticality hypothesis. Second, we will discuss models of these networks which capture relevant features of the data. Third, we will explore implications of these models on information processing in living neural networks. Fourth, we will conclude with a discussion of how future experiments could test these models.

2. Experiments

Here we will describe recordings from small cultures and slices of rat neocortex mounted on top of multielectrode arrays. To prepare a brain slice, a rat is first fully anaesthetized and then the intact brain is quickly dissected from the skull. After being placed in a cold, oxygenated solution that contains all the salts and sugars found in cerebrospinal fluid, the brain is sliced into thin sections, each approximately 300 μm thick. These sections allow oxygen and nutrients to diffuse into the tissue and waste products to diffuse away. In such a way, these slices may be kept alive for up to 12 hours even after the circulatory system is removed. For still longer recording sessions, slices may be bathed in culture medium, which contains hormones and blood serum, and thus kept alive for weeks at a time. Cultured slices afford the opportunity to make extremely long-term recordings of neural activity, thus enhancing statistical power during analysis. Such slices are called *in vitro* preparations, meaning 'in a dish'. Both acute slices and slice cultures are vastly simplified versions of the larger intact brain and may be easier to understand. As with any reduced system, however, *in vitro* preparations have the potential drawback of not including some relevant features of the intact system.

To record activity in slices, they are placed on multielectrode arrays, as shown in figure 1. These arrays are now manufactured commercially by both Multichannel Systems in Germany and Panasonic in Japan, and were pioneered over 25 years ago by Guenter Gross (Gross *et al.* 1982) and Jerome Pine (Pine 1980). Today, several groups are developing their own custom arrays, often with larger numbers of electrodes and smaller interelectrode spacing (e.g. Litke 1998).

Activity in acute slices and neuronal cultures is typically characterized by brief bursts of activity lasting tens of milliseconds (Corner *et al.* 2002; Wagenaar *et al.* 2006), separated by periods of quiescence lasting several seconds (figures 1 and 2). This activity is similar in many respects to that seen in slabs of cortex from *in vivo* animals that have been isolated from surrounding tissue (Timofeev *et al.* 2000). Activity may take the form of local field potentials (LFPs), which are caused by the synchronous activation of many neurons located near the electrode tip (Bove *et al.* 1996, 1998). Alternatively, the voltage deflections caused by the activation of individual neurons, called spikes, may also be detected.
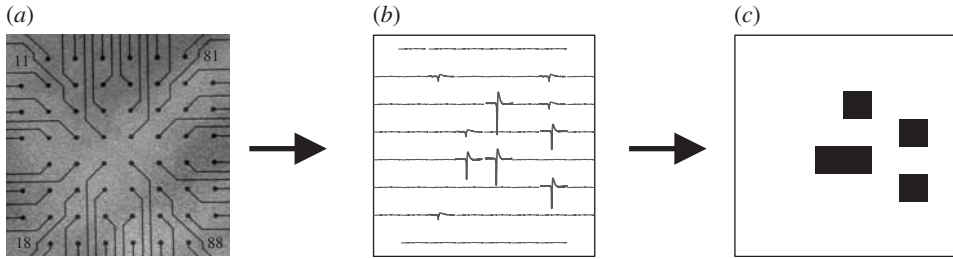
(*a*)   (*b*)   (*c*)

Figure 1. Schematic of data representation. (*a*) Tissue on array: photo of the 60-channel microelectrode array. The slice is placed on top of this and can only be seen as a grey background. Electrodes appear as small black circles at the ends of lines. Interelectrode distance is 200 μm and electrode diameter is 30 μm. (*b*) LFP signals on electrodes. LFPs are caused by the synchronous spiking of many neurons near the electrodes. (*c*) Thresholded signals: suprathreshold LFPs represented by small black squares. LFPs that exceed three standard deviations of the signal are considered suprathreshold.
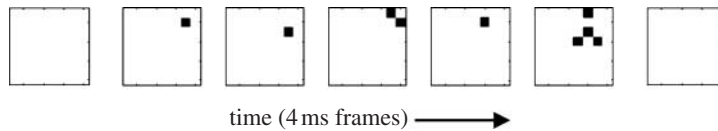
time (4 ms frames) ⟶

Figure 2. Example of a sequence. Seven frames are shown, where each frame represents activity on the electrode array during one 4 ms time step. Suprathreshold LFP activity at electrodes is shown by small black squares. A sequence is a series of consecutively active frames that is preceded by and terminated by blank frames. Sequence length is given by the number of active frames, while sequence size is given by the total number of active electrodes. The sequence shown here has a length of 5 and a size of 9.

Figure 2 illustrates that multielectrode data can be broken down into frames where there is no activity and frames where there is at least one active electrode. We assume that the data are binned at a temporal resolution that is appropriately matched to the spatial resolution of the electrode array. Here, the data are binned at 4 ms, and this is roughly the time it takes for activity to propagate over a distance of 200 μm, the interelectrode distance of the array (Maeda *et al.* 1995; Jimbo & Robinson 2000). We consider several consecutively active frames that are preceded by and terminated by inactive frames. Let us call these consecutively active frames, bracketed by inactivity, a *sequence*. The example sequence shown in figure 2 is five frames long, and the total number of electrodes activated during the sequence is nine. We may thus say that this sequence has a *length* of 5 and a *size* of 9.

The distribution of sequence *lengths* shown in figure 3*a* has a fat tail, suggesting that many sequences are longer than would be expected by chance. While this distribution is not clearly fitted by a power law, it cannot be accurately modelled by an ensemble of independent units (grey curve in figure 3*a*). This suggests that somehow the units in the actual network are interacting to produce sequences with long lifetimes.

A look at the distribution of sequence *sizes* further indicates that the independent model will not work (figure 3*b*). Here, sequence sizes are distributed in a manner that is nearly fitted by a power law (Beggs & Plenz 2003). Owing to
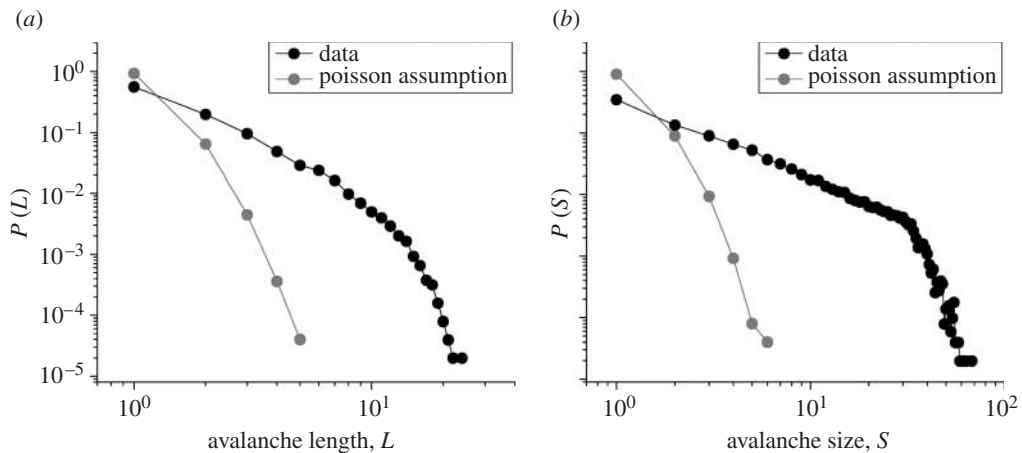
Figure 3. Sequence length and size distributions for LFPs. (*a*) Distribution of sequence lengths taken from acute slice LFPs recorded with a 60-electrode array, plotted in log–log space. Actual data are shown in black, while the output of a Poisson model is shown in grey. In the Poisson model, each electrode fires at the same rate as in the actual data, but independently of other electrodes. Note the large difference between the two curves, indicating that a Poisson model cannot account for the observed sequence lengths. (*b*) Distribution of sequence sizes from the same acute slice, plotted in log–log space. Here, the actual data follow a nearly straight line for sequence sizes from 1 to 35; after this point there is a cut-off induced by the electrode array size. The straight line is indicative of a power law, suggesting that the network is operating near the critical point. Note again the large difference between the Poisson model and the actual data. Because critical sand-pile models also demonstrate power-law distributions (Bak *et al.* 1987), these sequences of LFPs were called neuronal avalanches. Data here recorded by W. Chen, C. Haldeman, S. Wang, A. Tang and J. M. Beggs. Very similar results are obtained with organotypic cortical cultures (see Beggs & Plenz 2003).

the limited size of the array (60 electrodes), the power law begins to bend downward in a cut-off before 60. But for larger arrays (512 electrodes), the power law is seen to extend much further (figure 4). The equation of a power law is

$$P(S) = kS^{-\alpha},$$

where $P(S)$ is the probability of observing a sequence of size S; α is the exponent that gives the slope of the power law in a log–log graph; and k is a proportionality constant. For experiments with slice cultures, the size distribution of sequences of LFPs has an exponent $\alpha \approx 1.5$, but this may depend significantly on the spacing between electrodes in the array. In the example shown in figure 3, the array has an interelectrode distance of 200 μm. When the size distribution of sequences of spikes from slice cultures is plotted, it has a slightly larger exponent $\alpha \approx 2.1$ (figure 4), but these data are collected from an array with a smaller interelectrode spacing (60 μm). The factors that determine the observed value of the exponent α are still being actively researched. Because avalanches in critical sand-pile models also follow a power law (Bak *et al.* 1987; Paczuski *et al.* 1996), the term 'neuronal avalanche' was chosen to describe these cascades of local field potential activity (Beggs & Plenz 2003, 2004).
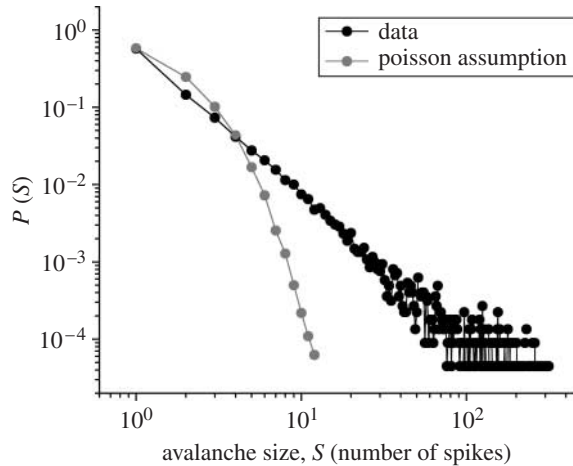
Figure 4. Avalanche size distribution for spikes. The log–log plot shows that the avalanche size distribution of spikes can be approximated by a straight line over three orders of magnitude in probability, without a sharp cut-off as shown in figure 3. This indicates that the data would be well fitted by a power law. Note again that the Poisson model cannot account for the observed distribution. Data here were collected by A. Litke, S. Sher, M. Grivich, D. Petrusca, S. Kachiguine and J. M. Beggs with a 512-electrode array from an acute cortical slice bathed in high potassium and zero magnesium. This large electrode array was used to verify that the power law continued beyond the cut-off point of 60 observed in data obtained from smaller, 60-electrode arrays. Similar results are obtained from cortical slice cultures. Spikes were thresholded at three standard deviations and were not sorted. Data were binned at 1.2 ms to match the short interelectrode distance of 60 μm.

The power-law distribution of avalanche sizes also suggests that these networks of neurons may be operating near a critical point. This is because many tunable systems display power-law distributions at the critical point: magnets at the Curie point (Yeomans 1992), water molecules at a second-order phase transition (Stanley 1987) or a lattice at the percolation threshold (Stauffer 1994). Strictly speaking, only infinitely large systems can operate at the critical point, but here we will use the term 'critical' to describe behaviour in finite systems that would approach criticality if they were extended to unlimited sizes.

It is important to note that, while power laws have been reported for many years in neuroscience in the temporal correlations of single time-series data (e.g. the power spectrum from the electroencephalogram (Linkenkaer-Hansen *et al.* 2001; Worrell *et al.* 2002), Fano or Allan factors in spike count statistics (Teich *et al.* 1997), neurotransmitter secretion times (Lowen *et al.* 1997), ion channel fluctuations (Toib *et al.* 1998), interburst intervals in neuronal cultures (Segev *et al.* 2002)), they had not been observed from the interactions seen in multielectrode data. Thus, power-law distributions from multielectrode data suggest that distributed neuronal networks operate near criticality. This conclusion could not come from single time-series data alone.

Power-law distributions of sequence sizes have also been observed from the distribution of spikes in the isolated leech ganglion (Mazzioni *et al.* 2007) and in spikes from dissociated cortical cultures (Alessio *et al.* 2006), suggesting that the phenomenon of criticality may be quite general. Preliminary reports
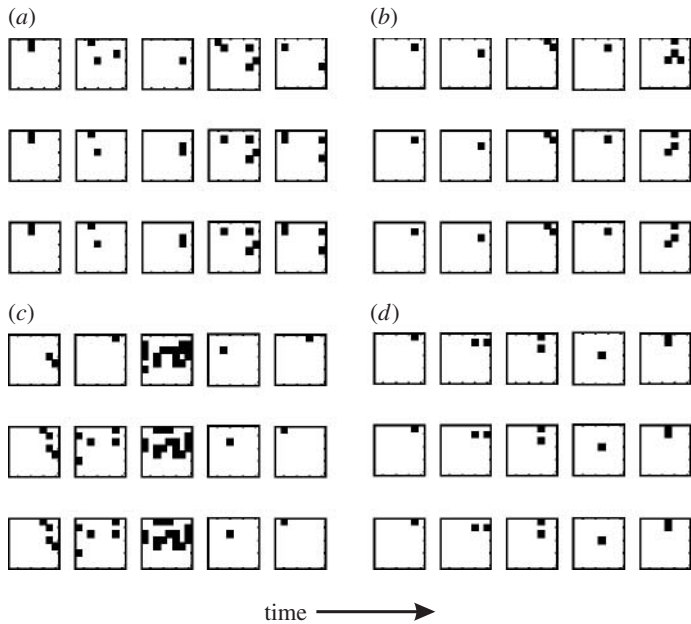
J. M. Beggs

Figure 5. Families of repeating avalanches. Each family (*a–d*) shows three similar avalanches. Similarity within each group was higher than expected by chance when compared with 50 sets of shuffled data. On average, 30 ± 14 (mean $\pm$ s.d.) distinct families of reproducible avalanches, each containing approximately 23 avalanches, occur in a given slice culture (Beggs & Plenz 2004). Repeating avalanches are stable for 10 hours and have a temporal precision of 4 ms, suggesting that they could serve as a substrate for storing information in neural networks.

indicate that power-law distributions are also present in LFPs from superficial cortical layers of awake primates (Petermann *et al.* 2006).

While avalanches in critical sand-pile models are stochastic in the patterns they form, this is not the case with the neuronal avalanches reported so far. Avalanches of local field potentials can occur in spatio-temporal patterns that repeat more often than expected by chance (Beggs & Plenz 2004). Figure 5 shows several such patterns from an acute cortical slice. These patterns are reproducible over periods of as long as 10 hours, and have a temporal precision of 4 ms (Beggs & Plenz 2004). The stability and precision of these patterns suggest that neuronal avalanches could be used by networks as a substrate for storing information. In this sense, avalanches appear to resemble sequences of spikes observed from the brains of animals performing cognitive tasks (e.g. Abeles *et al.* 1993; Dave & Margoliash 2000; Hahnloser *et al.* 2002).

3. Models

Computational work suggested criticality in neuronal networks several years before the phenomenon was observed experimentally. Bak *et al.* (1987) proposed a cellular automaton simulation of a sand pile. A two-dimensional lattice has cells that can hold up to three grains of sand before toppling. When a fourth grain is added, the site is driven over threshold and topples, distributing one

grain to each of its four neighbouring cells. If any of these neighbouring cells is over threshold, it too will re-distribute sand to its neighbours. If single grains of sand are slowly added to this system, over time the distribution of cascade sizes will approach a power law, indicating that the system is at a critical point. The critical point is characterized by avalanches with infinite spatial and temporal correlations. The essential ingredients of the model responsible for this critical behaviour seem to be the following (Jensen 1998):

(i) many nonlinear threshold units connected to each other and
(ii) a separation of driving and relaxation time scales, such that the addition of sand (driving) occurs very slowly relative to the duration of avalanches (relaxation).

In both these respects, the sand-pile model is similar to *in vitro* neural networks. Neurons within a network are coupled, nonlinear threshold units; the time between spontaneous bursts of activity (tens of seconds) is much longer than the time it takes for activity to propagate within the network (tens of milliseconds).

Models that explicitly predicted avalanches of neural activity include the work of Herz & Hopfield (1995), which connects the reverberations in a neural network to the power-law distribution of earthquake sizes. Also notable is the work of Eurich *et al.* (2002), which predicted that the avalanche size distribution from a network of globally coupled nonlinear threshold elements should have an exponent of $\alpha \approx 1.5$ at the critical point. Remarkably, this exponent turned out to match that reported experimentally (Beggs & Plenz 2003). A recent paper also arrives at the same exponent by invoking neural field theory (Buice & Cowan 2007).

Here, we will describe in more detail a critical branching process model (Harris 1989; Beggs & Plenz 2003; Haldeman & Beggs 2005; reviewed by Vogels *et al.* 2005), because it captures the power-law distribution of avalanche sizes and the reproducible activity sequences seen in the data. In a neural branching model, a neuron that is active at one time step will produce, on average, activity in σ neurons in the next time step. The number σ is called the *branching parameter* and can be thought of as the expected value of this ratio,

$$\sigma = \frac{\# \text{ descendants}}{\# \text{ ancestors}},$$

where # ancestors is the number of neurons active at time step t and # descendants is the number of neurons active at time step $t+1$. There are three general regimes for σ, which are as follows.

(i) *Subcritical.* If $\sigma < 1$, then activity is damped and will quickly die out.
(ii) *Critical.* If $\sigma \approx 1$, then activity is nearly sustained, but because σ is an average measure, the number of descendants sometimes will be zero and propagation will terminate. The distribution of sequence lengths will approach a power law.
(iii) *Supercritical.* When $\sigma > 1$, then activity is amplified over time.
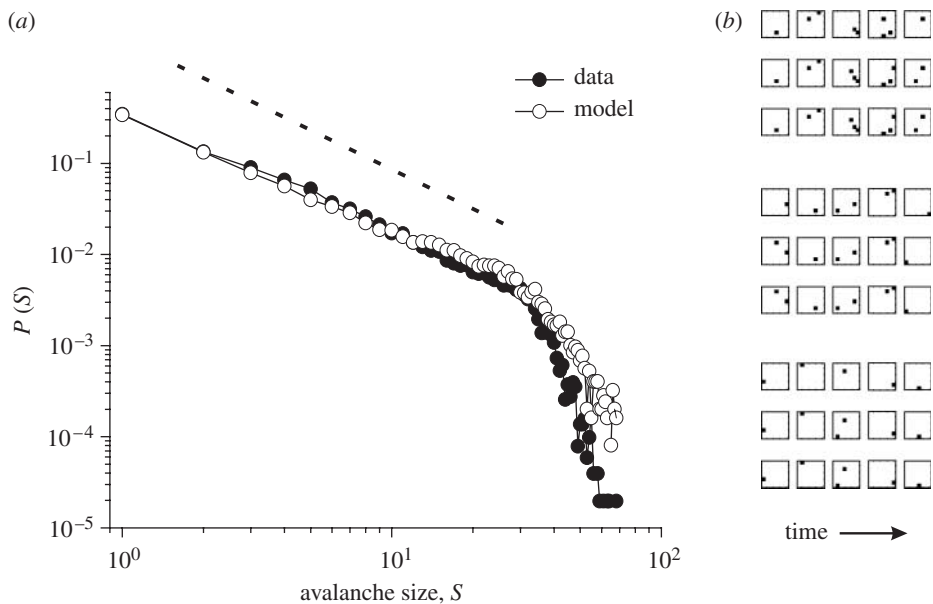
Figure 6. A branching model captures the two main features of the data. (*a*) Avalanche size distribution from data and model compared, showing fairly close correspondence. Note that both show a straight line portion in log–log space, extending over avalanche sizes 1–35. Model was tuned to the critical point such that the branching parameter, σ, equalled unity. (*b*) Three families of significantly similar avalanches produced by the model. Note similarity to avalanche families produced by actual data in figure 5.

At the level of a single neuron in the network, the branching parameter σ is set by the following relationship:

$$\sigma_i = \sum_{j=1}^{M} p_{ij},$$

where σ_i is the expected number of descendant neurons activated by neuron i; M is the number of neurons that neuron i connects to; and p_{ij} is the probability that activity in neuron i will transmit to neuron j. By setting the sum of transmission probabilities p at every neuron, the value of the branching parameter σ may be tuned for the entire network (Beggs & Plenz 2003; Haldeman & Beggs 2005). Because some transmission probabilities are greater than others, preferred paths of transmission may occur, leading to reproducible avalanche patterns.

Using this model, Haldeman & Beggs (2005) found that both the power-law distribution of avalanche sizes and the repeating avalanches can be qualitatively captured when σ is tuned to the critical point ($\sigma=1$), as shown in figure 6. When the model is tuned either above ($\sigma>1$) or below ($\sigma<1$) the critical point, it fails to produce a power-law distribution of avalanche sizes.

An alternative model relies on the fact that power laws can be constructed from a summation of many exponential processes with different time constants (Fusi *et al.* 2005) as shown in figure 7. This model is appealing because it can account for the power law as well as the high levels of memory storage and long retention times observed in experiments. As with the branching model, it is too early to tell whether this model accurately accounts for other aspects of the data as well.
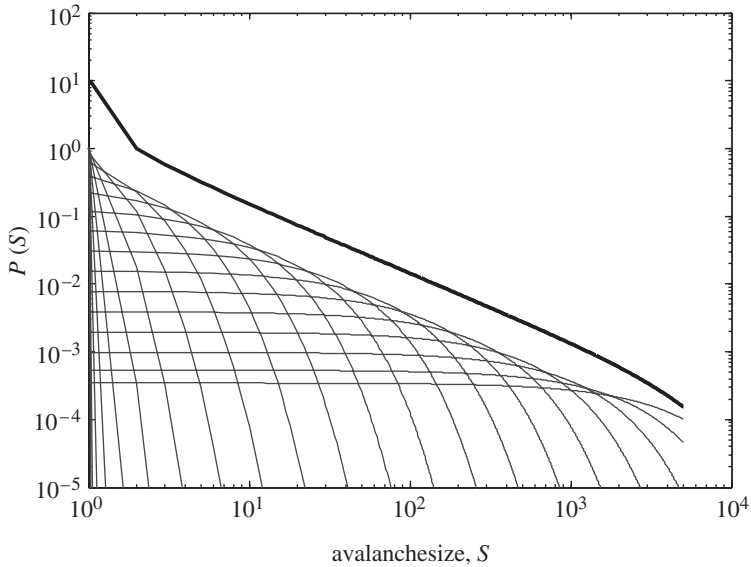
Figure 7. A sum of exponentials can also produce a power law. Thin curves are normalized exponential distributions, each with a different exponent. The thick curve is the sum of these exponential curves, producing a close approximation to a power law. This type of model might also account for observed neuronal avalanches.

Because so many physiological systems demonstrate homeostatic regulation (heart rate, respiration rate, body temperature, etc.), it is natural to suppose that neural networks would also approach the critical point through some self-organizing processes. Candidate mechanisms for leading the network into this state include competitive activity-dependent attachment and pruning (e.g. Bornholdt & Rohlf 2000; de Arcangelis *et al.* 2006), short-term synaptic plasticity (Levina *et al.* 2005) or some combination of homeostatic regulation of excitability and Hebbian learning (Hsu & Beggs 2006).

4. Information processing

What are the implications of criticality for information processing in neuronal networks? To clearly describe this, it is first necessary to explain in more detail the implications of power laws. There are two features of power laws that are relevant to neuronal avalanches, and are as follows.

(i) An important property of many power laws is that their mean values *diverge*. Let us consider a distribution of avalanche sizes: $P(S) = S^{-\alpha}$. The expected value of this distribution is given by

$$\langle S \rangle = \int_{S=1}^{S=\infty} S \cdot P(S) \, \mathrm{d}S,$$

which reduces to

$$\langle S \rangle = \int_{S=1}^{S=\infty} S \cdot S^{-\alpha}\, \mathrm{d}S = \int_{S=1}^{S=\infty} S^{1-\alpha}\, \mathrm{d}S.$$

For power laws with exponents $\alpha \leq 2$, this integral will diverge. Because the observed exponent for avalanches of local field potentials is $\alpha \approx 1.5$, this means that the average avalanche size will be infinite. Although it will be much more common for small avalanches to occur, the fact that some extremely large avalanches appear with a moderate probability makes the expected value blow up. Infinitely large avalanches can be related to information transmission, and will be described in detail below.

(ii) A power law is also a description of a *scaling relationship*. In the case of avalanches, it says that a given ratio of avalanche sizes will have a corresponding ratio of avalanche probabilities. This ratio of probabilities will remain constant for all avalanche sizes. This can also be explained by equations. If the ratio of avalanche sizes is $S'/S = a$, then we have $P(S) = kS^{-\alpha}$ for an avalanche of size S, and $P(S') = k(aS)^{-\alpha}$ for an avalanche of size S', since $S' = aS$. Therefore, the ratio of avalanche probabilities will be $P(S')/P(S) = ka^{-\alpha}S^{-\alpha}/kS^{-\alpha} = a^{-\alpha}$. The ratio of avalanche probabilities $a^{-\alpha}$ does not depend on avalanche size S. Owing to this, the power-law relationship has been called 'scale free', 'scale invariant' and 'fractal'. Regardless of the magnification, the patterns appear to be similar. Because avalanches occur at all sizes, scale invariance is related to diversity (Corral *et al.* 1997) and also to complexity (Langton 1990; Bak 1996). Note that only power-law distributions are scale invariant. For non-power-law distributions, most avalanche sizes would occur near one mean value. The diversity of power laws can be related to the information storage capacity and the computational power of a network, as elaborated below.

These properties of the avalanche size distribution have implications for information processing in neural networks in the following four areas.

(i) *Information transmission.* Infinitely large avalanches might seem like a problem for a neural network, but they open up the possibility of communication between all neurons in the network. Indeed, when a feed-forward neural network based on a branching process is tuned to the critical point, it has optimal information transmission (figure 8; Beggs & Plenz 2003; Bertschinger & Natschlager 2004; Kinouchi & Copelli 2006). Thus, the fact that the avalanche size distribution diverges suggests that neuronal avalanches are efficient for transmitting information.

(ii) *Information storage.* When a recurrent branching network is tuned to the critical point, the number of significantly repeating avalanche patterns is maximized, as shown in figure 9 (Haldeman & Beggs 2005). The diversity of avalanches at the critical point leads to a wider range of stable patterns in the network, and this may lead to optimal information storage.

(iii) *Computational power.* By changing the variance in synaptic weights in a spiking network model, Bertschinger and colleagues (Maass *et al.* 2002; Bertschinger & Natschlager 2004) produced networks with damped,
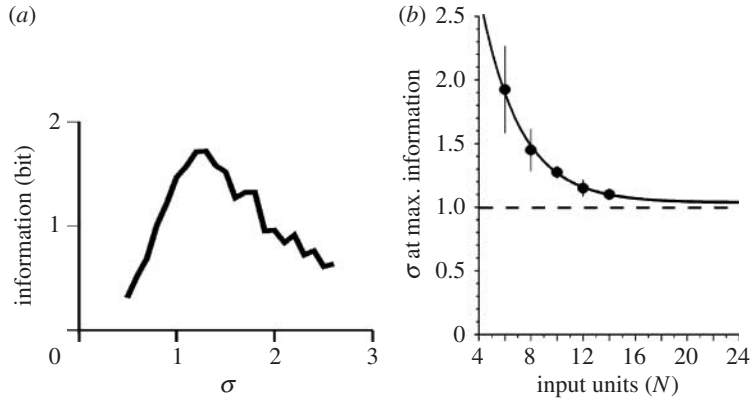
Figure 8. Information transmission is optimized at the critical point in a branching model. (*a*) Information transmission through a feed-forward network peaks when the branching parameter σ is between 1 and 2. This suggests that there is a relatively narrow range in which information may flow efficiently. (*b*) As network models are made larger, the value of σ at which information transmission peaks asymptotically approaches 1. This suggests that, for a very large network, information transmission would be optimized at the critical point. Note that at the critical point, when $\sigma = 1$, the model produces a power-law distribution of avalanche sizes and repeating avalanches, as shown in the data in figure 3*b*. Reproduced with permission from Beggs & Plenz (2003).
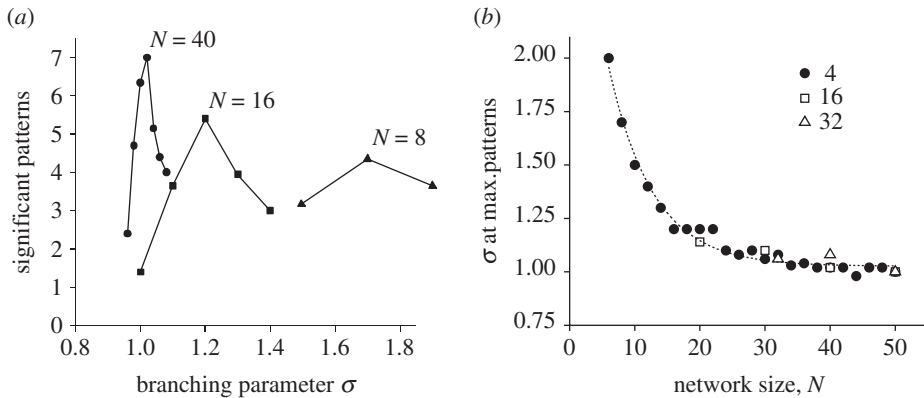

Figure 9. Number of statistically significant repeating avalanche patterns is maximized at the critical point. (*a*) Number of significant repeating avalanche patterns plotted against σ for three network sizes. Larger networks peak closer to $\sigma = 1$. (*b*) Curve asymptotically approaches $\sigma = 1.03 \pm 0.01$, and large networks peak exactly at $\sigma = 1$, regardless of the number of connections per unit (4, 16, 32), suggesting universality. Reproduced with permission from Haldeman & Beggs (2005, fig. 3; http://prl.aps.org/). Copyright © American Physical Society.

sustained and expanding activity. When synaptic strengths have low variance, activity coming from a given neuron will produce relatively similar downstream effects. This will lead to stable, repeatable activity patterns. In contrast, when synaptic strengths have high variance, activity from one neuron tends to activate downstream targets differently each time. This generates highly variable activity patterns. It was found

that computation depends on a delicate balance between order and variability. If a network is boringly ordered, it cannot perform many mappings between inputs and outputs. This limits the range of computations that a network could perform. On the other hand, if a network is too random, it has a wider range of mappings, but the mappings will be unreliable, thus undermining computation. Networks operating between these two extremes, near the critical point, performed more effectively on a variety of computational tasks than networks that were tuned to have either subcritical or supercritical dynamics (see also Latham & Nirenberg 2004).

(iv) *Stability.* When a recurrent, branching network model is tuned to the critical point, it produces largely parallel trajectories in state space, meaning that the network is at the edge of stability (Bertschinger & Natschlager 2004; Haldeman & Beggs 2005). Trajectories are still stable and controllable with minor corrective inputs.

Optimization of all of these information processing tasks may occur simultaneously when a network operates near the critical point.

5. Testable predictions

Emerging work suggests that pharmacological agents can be used to tune living neural networks through subcritical, critical and supercritical regimes. For example, reducing the effect of inhibitory transmission through the application of bicuculline seems to cause larger avalanches, while reducing excitatory transmission with agents like CNQX or APV seems to reduce avalanche size. These results are still preliminary and are sometimes inconsistent, but a general trend seems to be emerging that supports the above statements.

Which predictions should then be tested? First, information transmission between randomly chosen ensembles of recording sites in a slice network should be optimal when the tissue is operating near the critical point. Second, the number of statistically significant repeating avalanche patterns should also be maximal when the network is operating at the critical point. Third, the computational power, measured in a manner similar to that proposed by Bertschinger & Natschlager (2004), should also peak near the critical point. Fourth, trajectories in state space should be neutral in critical networks, attracting in subcritical networks and chaotic in supercritical networks, as predicted by Haldeman & Beggs (2005).

However, a couple of notes of caution should be mentioned about these future experiments. True criticality can only occur in systems of infinite size, while most of these experiments will be performed with only tens, or at most hundreds, of recording sites. Care should be taken to correct for system size when interpreting the results. For example, the number of repeating avalanche patterns in *small* networks should actually peak at *supercritical* values of σ, and will only peak at the critical point in the asymptotic limit as system size approaches infinity. This is clearly shown in figure 9b, where a 10×10 network with 100 units is expected to have maximum storage capacity when $\sigma \approx 1.5$, not when $\sigma = 1$. Another point of caution concerns studies of information transmission. All of the data presented

above were collected from networks that were spontaneously active, and all of the conclusions above apply to this condition. It is also possible to stimulate neural networks with electrical pulses and observe their responses. In this case, care must be taken to use stimulation intensities that are approximately equal to the stimulation caused by neurons in the network. If an overwhelmingly strong stimulus is used, one might erroneously conclude that signals can propagate without much loss even in subcritical networks. Conversely, if an extremely weak stimulus is used, one might conclude that only supercritical networks favour information transmission. For these reasons, it is probably best to perform tests of the criticality hypothesis first using spontaneous activity of neural networks. If these tests produce promising results, then they could help to move the criticality hypothesis closer towards biological reality.

References

Abeles, M., Bergman, H., Margalit, E. & Vaadia, E. 1993 Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J. Neurophysiol.* **70**, 1629–1638.

Alessio, R., Cozzi, L., D'Angelo, P. & Sanguinetti, V. 2006 Cultures of dissociated neurons display a variety of avalanche behaviors. In *ESANN'2006 Proc.—European Symp. on Artificial Neural Networks, Bruges, Belgium, 26–28 April 2006.*

Bak, P. 1996 *How nature works: the science of self-organized criticality.* New York, NY: Copernicus.

Bak, P. & Chialvo, D. 2001 Adaptive learning by extremal dynamics and negative feedback. *Phys. Rev. E* **63**, 031912. (doi:10.1103/PhysRevE.63.031912)

Bak, P., Tang, C. & Wiesenfeld, K. 1987 Self-organized criticality: an explanation of the $1/f$ noise. *Phys. Rev. Lett.* **59**, 381–384. (doi:10.1103/PhysRevLett.59.381)

Beggs, J. M. & Plenz, D. 2003 Neuronal avalanches in neocortical circuits. *J. Neurosci.* **23**, 11 167–11 177.

Beggs, J. M. & 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. (doi:10.1523/JNEUROSCI.0540-04.2004)

Bertschinger, N. & Natschlager, T. 2004 Real-time computation at the edge of chaos in recurrent neural networks. *Neural Comput.* **16**, 1413–1436. (doi:10.1162/089976604323057443)

Bornholdt, S. & Rohlf, T. 2000 Topological evolution of dynamical networks: global criticality from local dynamics. *Phys. Rev. Lett.* **84**, 6114–6117. (doi:10.1103/PhysRevLett.84.6114)

Bove, M., Genta, G., Verreschi, G. & Grattarola, M. 1996 Characterization and interpretation of electrical signals from random networks of cultured neurons. *Technol. Health Care* **4**, 77–86.

Bove, M., Martinoia, S., Verreschi, G., Giugliano, M. & Grattarola, M. 1998 Analysis of the signals generated by networks of neurons coupled to planar arrays of microtransducers in simulated experiments. *Biosens. Bioelectron.* **13**, 601–612. (doi:10.1016/S0956-5663(98)00015-3)

Buice, M. & Cowan, J. D. 2007 Field-theoretic approach to fluctuation effects in neural networks. *Phys. Rev. E* **75**, 051919. (doi:10.1103/PhysRevE.75.051919)

Chialvo, D. & Bak, P. 1999 Learning from mistakes. *Neuroscience* **90**, 1137–1148. (doi:10.1016/S0306-4522(98)00472-2)

Corner, M. A., van Pelt, J., Wolters, P. S., Baker, R. E. & Nuytinck, R. H. 2002 Physiological effects of sustained blockade of excitatory synaptic transmission on spontaneously active

developing neuronal networks—an inquiry into the reciprocal linkage between intrinsic biorhythms and neuroplasticity in early ontogeny. *Neurosci. Biobehav. Rev.* **26**, 127–185. (doi:10.1016/S0149-7634(01)00062-8)

Corral, A., Perez, C. J. & Diaz-Guilera, A. 1997 Self-organized criticality induced by diversity. *Phys. Rev. Lett.* **78**, 1492. (doi:10.1103/PhysRevLett.78.1492)

Dave, A. S. & Margoliash, D. 2000 Song replay during sleep and computational rules for sensorimotor vocal learning. *Science* **290**, 812–816. (doi:10.1126/science.290.5492.812)

de Arcangelis, L., Perrone-Capano, C. & Herrmann, H. J. 2006 Self-organized criticality model for brain plasticity. *Phys. Rev. Lett.* **96**, 028107. (doi:10.1103/PhysRevLett.96.028107)

Eurich, C. W., Herrmann, J. M. & Ernst, U. A. 2002 Finite-size effects of avalanche dynamics. *Phys. Rev. E* **66**, 066137. (doi:10.1103/PhysRevE.66.066137)

Fusi, S., Drew, P. J. & Abbott, L. F. 2005 Cascade models of synaptically stored memories. *Neuron* **45**, 559–611. (doi:10.1016/j.neuron.2005.02.001)

Gross, G. W., Williams, A. N. & Lucas, J. H. 1982 Recording of spontaneous activity with photoetched microelectrode surfaces from mouse spinal neurons in culture. *J. Neurosci. Methods* **5**, 13–22. (doi:10.1016/0165-0270(82)90046-2)

Hahnloser, R. H., Kozhevnikov, A. A. & Fee, M. S. 2002 An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* **419**, 65–70. (doi:10.1038/nature00974)

Haldeman, C. & Beggs, J. M. 2005 Critical branching captures activity in living neural networks and maximizes the number of metastable states. *Phys. Rev. Lett.* **94**, 058101. (doi:10.1103/PhysRevLett.94.058101)

Harris, T. E. 1989 *The theory of branching processes.* New York, NY: Dover.

Herz, A. V. & Hopfield, J. J. 1995 Earthquake cycles and neural reverberations: collective oscillations in systems with pulse-coupled threshold elements. *Phys. Rev. Lett.* **75**, 1222–1225. (doi:10.1103/PhysRevLett.75.1222)

Hsu, D. & Beggs, J. M. 2006 Neuronal avalanches and criticality: a dynamical model for homeostasis. *Neurocomputing* **69**, 1134–1137. (doi:10.1016/j.neucom.2005.12.060)

Jensen, H. J. 1998 *Self-organized criticality: emergent complex behavior in physical and biological systems.* Cambridge, UK: Cambridge University Press.

Jimbo, Y. & Robinson, H. P. 2000 Propagation of spontaneous synchronized activity in cortical slice cultures recorded by planar electrode arrays. *Bioelectrochemistry* **51**, 107–115. (doi:10.1016/S0302-4598(99)00083-5)

Kauffman, S. A. 1993 *The origins of order: self-organization and selection in evolution.* New York, NY: Oxford University Press.

Kauffman, S. A. & Johnsen, S. 1991 Coevolution to the edge of chaos—coupled fitness landscapes, poised states, and coevolutionary avalanches. *J. Theor. Biol.* **149**, 467–505. (doi:10.1016/S0022-5193(05)80094-3)

Kinouchi, O. & Copelli, M. 2006 Optimal dynamic range of excitable networks at criticality. *Nature Phys.* **2**, 348–351. (doi:10.1038/nphys289)

Langton, C. G. 1990 Computation at the edge of chaos—phase-transitions and emergent computation. *Physica D* **42**, 12–37. (doi:10.1016/0167-2789(90)90064-V)

Latham, P. E. & Nirenberg, S. 2004 Computing and stability in cortical networks. *Neural Comput.* **16**, 1385–1412. (doi:10.1162/089976604323057434)

Legenstein, R. A. & Maass, W. 2007 Edge of chaos and prediction of computational performance for neural circuit models. *Neural Netw.* **20**, 323–334. (doi:10.1016/j.neunet.2007.04.017)

Levina A., Hermann, J. M. & Geisel, T. 2005 Dynamical synapses give rise to a power-law distribution of neuronal avalanches. In *2005 NIPS Proc.* See http://books.nips.cc/papers/files/nips18/NIPS2005_0646.pdf.

Lin, M. & Chen, T. 2006 Self-organized criticality in a simple model of neurons based on small-world networks. *Phys. Rev. E* **71**, 016133. (doi:10.1103/PhysRevE.71.016133)

Linkenkaer-Hansen, K., Nikouline, V. V., Palva, J. M. & Ilmoniemi, R. J. 2001 Long-range temporal correlations and scaling behavior in human brain oscillations. *J. Neurosci.* **21**, 1370–1377.

Litke, A. M. 1998 The retinal readout system: an application of microstrip detector technology to neurobiology. *Nucl. Instrum. Methods Phys. Res. A* **418**, 203–209. (doi:10.1016/S0168-9002(98)00734-7)

Lowen, S. B., Cash, S. S., Poo, M. M. & Teich, M. C. 1997 Quantal neurotransmitter secretion rate exhibits fractal behavior. *J. Neurosci.* **17**, 5666–5677.

Maass, W., Natschlager, T. & Markram, H. 2002 Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.* **14**, 2531–2560. (doi:10.1162/089976602760407955)

Maeda, E., Robinson, H. P. C. & Kawana, A. 1995 The mechanisms of generation and propagation of synchronized bursting in developing networks of cortical neurons. *J. Neurosci.* **15**, 6834–6845.

Mazzioni, A., Broccard, F. D., Garcia-Perez, E., Bonifazi, P., Ruaro, M. E. & Torre, V. 2007 On the dynamics of the spontaneous activity in neuronal networks. *PLoS ONE* **2**, e439. (doi:10.1371/journal.pone.0000439)

Mitchell, M., Hraber, P. T. & Crutchfield, J. P. 1993 Revisiting the edge of chaos: evolving cellular automata to perform computations. *Complex Syst.* **7**, 89–130.

Packard, N. 1988 Adaptation towards the edge of chaos. In *Dynamic patterns in complex systems* (eds J. A. S. Kelso, A. J. Mandell & M. F. Shlesinger), p. 293. Singapore: World Scientific.

Paczuski, M., Maslov, S. & Bak, P. 1996 Avalanche dynamics in evolution, growth, and depinning models. *Phys. Rev. E* **53**, 414–443. (doi:10.1103/PhysRevE.53.414)

Petermann, T., Lebedev, M., Nicolelis, M. & Plenz, D. 2006 Neuronal avalanches *in vivo*. *Soc. Neurosci. Abstr.* **531**, 1.

Pine, J. 1980 Recording action potentials from cultured neurons with extracellular microcircuit electrodes. *J. Neurosci. Methods* **2**, 19–31. (doi:10.1016/0165-0270(80)90042-4)

Segev, R., Benveniste, M., Hulata, E., Cohen, N., Palevski, A., Kapon, E., Shapira, Y. & Ben Jacob, E. 2002 Long term behavior of lithographically prepared *in vitro* neuronal networks. *Phys. Rev. Lett.* **88**, 118 102. (doi:10.1103/PhysRevLett.88.118102)

Stanley, H. E. 1987 *Introduction to phase transitions and critical phenomena*. New York, NY: Oxford University Press.

Stauffer, D. 1994 *Introduction to percolation theory*. London, UK: Taylor and Francis.

Stewart, C. V. & Plenz, D. 2006 Inverted-U profile of dopamine–NMDA-mediated spontaneous avalanche recurrence in superficial layers of rat prefrontal cortex. *J. Neurosci* **26**, 8148–8159. (doi:10.1523/JNEUROSCI.0723-06.2006)

Teich, M. C., Heneghan, C., Lowen, S. B., Ozaki, T. & Kaplan, E. 1997 Fractal character of the neural spike train in the visual system of the cat. *J. Opt. Soc. Am. A* **14**, 529–546.

Timofeev, I., Grenier, F., Bazhenov, M., Sejnowski, T. J. & Steriade, M. 2000 Origin of slow cortical oscillations in deafferented cortical slabs. *Cereb. Cortex.* **10**, 1185–1199. (doi:10.1093/cercor/10.12.1185)

Toib, A., Lyakhov, V. & Marom, S. 1998 Interaction between duration of activity and time course of recovery from slow inactivation in mammalian brain Na channels. *J. Neurosci.* **18**, 1893–1903.

Vogels, T. P., Rajan, K. & Abbott, L. F. 2005 Neural network dynamics. *Annu. Rev. Neurosci.* **28**, 357–376. (doi:10.1146/annurev.neuro.28.061604.135637)

Wagenaar, D. A., Pine, J. & Potter, S. M. 2006 An extremely rich repertoire of bursting patterns during the development of cortical cultures. *BMC Neurosci.* **7**, 7–11. (doi:10.1186/1471-2202-7-11)

Worrell, G. A., Cranstoun, S. D., Echauz, J. & Litt, B. 2002 Evidence for self-organized criticality in human epileptic hippocampus. *NeuroReport* **13**, 2017–2021. (doi:10.1097/00001756-200211150-00005)

Yeomans, J. M. 1992 *Statistical mechanics of phase transitions*. New York, NY: Oxford University Press.