Published in final edited form as:

Epilepsy Behav. 2008 October ; 13(3): 511–522. doi:10.1016/j.yebeh.2008.05.007.

An open hypothesis: is epilepsy learned, and can it be unlearned?

David Hsu¹, Wei Chen², Murielle Hsu¹, and John M. Beggs²

¹Department of Neurology, University of Wisconsin, Madison WI

²Department of Physics, Indiana University, Bloomington IN

Abstract

Plasticity is central to the ability of a neural system to learn and also to its ability to develop spontaneous seizures. What is the connection between the two? Learning itself is known to be a destabilizing process at the algorithmic level. We have investigated necessary constraints on a spontaneously active Hebbian learning system and find that the ability to learn appears to confer an intrinsic vulnerability to epileptogenesis on that system. We hypothesize that epilepsy arises as an abnormal *learned response* of such a system to certain repeated provocations. This response is a network level effect. If epilepsy really is a learned response, then it should be possible to reverse it, i.e., to *unlearn* epilepsy. Unlearning epilepsy may then provide a new approach to its treatment.

Keywords

epileptogenesis; critical connectivity; homeostasis; Hebbian learning

Introduction

The primary functions of the brain are to transmit, process and store information about the body and the environment. Higher order functions such as problem solving and adaptation also exist in some animals. We will refer to all these functions loosely as components of the learning process. Plasticity of neurons and of the connections between neurons is central to these capabilities. Plasticity is also central to epileptogenesis [1]. Is this simply an unhappy coincidence? Or is there a deeper reason why plasticity plays a role in both learning and epileptogenesis?

How pathologic plasticity leads to epileptogenesis is a subject of intense interest. There is an enormous number of ways in which plasticity can go wrong at all levels of description [2–9]. Particularly at the genetic level, the process of epileptogenesis is a bewildering complex with many contributory factors. Indeed, so intricately is normal brain function dependent on the proper mix of receptors, channels, chemical environment and other factors that it may appear a miracle that so many animals are able to function at all, and that more of us are not subject to epilepsy.

In principle, any abnormality in gene expression or molecular environment associated with clinical or electrophysiological evidence of epileptogenesis may represent one of three

Corresponding author: David Hsu, **Address:** Department of Neurology, 600 Highland Av H6/526, University of Wisconsin, Madison WI, 53792, **Phone:** 608-263-8551, **Fax:** 608-263-0412, **Email:** hsu@neurology.wisc.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

possibilities: each abnormality may be a causative factor, a response to epileptogenesis or a noncontributory, incidental finding. It is often difficult to classify each abnormality to one of these three classes. Prinz et al. have cautioned that a great variety of combinations of receptor and ion channel types can result in the same electrophysiological behavior [10]. Thus it may be necessary to understand how all the constitutive parts come together to understand the behavior of the whole, at least at the whole neuron level and likely also at the network level.

It would be useful to have organizing principles for the functional behavior of biological neural systems. In this paper, we review some of the recent advances in this regard. As a consequence of some of these organizing principles, the ability of a neural system to learn appears to confer an intrinsic vulnerability to epileptogenesis on that system. We hypothesize that epilepsy arises as an abnormal *learned response* of such a system to certain repeated provocations. This response is a network level effect. If, as we hypothesize, epilepsy really is a learned response, then it should be possible to reverse it, i.e., to *unlearn* epilepsy. Unlearning epilepsy may then provide a new approach to its treatment.

We also propose that there must be at least three conditions for the development of spontaneous seizures, or epilepsy. These are (1) neuronal hyperexcitability due to an imbalance between excitatory and inhibitory influences, (2) overconnectivity in space leading to abnormally wide spatial spread of neuronal activity, and (3) overconnectivity in time leading to abnormally persistent activity. All three conditions must exist for spontaneous seizures to appear, although all three conditions do not have to be present continuously in an epileptic brain. These three conditions are distinct, and each is a potential target for treatment.

Background: Hebbian learning, neuronal avalanches and critical connectivity

Here we review the recent literature that bears on learning and the self-organizing properties of biological neural systems. The key points are: (1) learning destabilizes both the activity and connectivity of a neural system, (2) there is a certain level of connectivity, called critical connectivity, that optimizes brain performance, and (3) there are homeostatic mechanisms that maintain stable levels of activity and critical connectivity in the face of the destabilizing effects of learning.

Neural systems are thought to obey Hebbian learning rules [11], summarized by the phrase “cells that fire together, wire together.” In the simplest formulation, if two neurons consistently fire consecutively, then the connection from the first-firing neuron to the second-firing neuron is strengthened by a small factor. Such a learning rule represents long-term potentiation (LTP). Conversely, if firing of one neuron is not followed by firing of a second neuron, then the connection of the first-firing neuron to that of the second neuron is weakened. Such a learning rule represents long-term depression (LTD). A third learning rule combines both LTP and LTD and is known as spike-timing dependent plasticity (STDP). These learning rules are all competitive associative rules.

Hebbian rules are destabilizing

However, computational models utilizing simple Hebbian learning rules eventually result in either runaway excitation or global silence [12–15]. At the same time, connectivity tends either to rise to an overconnected state, such that excitation at one neuron is immediately followed by global or near-global activation, or else to decrease to a severely underconnected state, such that activation of one neuron is not followed by activation of any other neuron. Connectivity at either extreme is not useful for information processing. The most useful level of connectivity lies somewhere between these two extremes, allowing for a wide variety of spatial patterns extending from the smallest possible clusters of a few neurons discharging simultaneously to the largest possible extent, covering macroscopic portions of the brain. Is there a level of

connectivity that yields the widest possible repertory of spatial activation patterns? The answer appears to be yes. The clearest results have come from the study of acute cortical slices and slice cultures on microelectrode arrays. We review these results next.

Network activity

With the advent of microelectrode arrays, it has become common to record activity from acute cortical slices and cortical slice cultures on tens of channels [16,17]. Figure 1 shows a representative recording of population spikes on many electrodes recorded simultaneously from a slice of rat parietal cortex. Activity in such slices is typically characterized by brief bursts of population spikes lasting tens of milliseconds, separated by periods of quiescence lasting several seconds [18–20]. Each population spike represents the near simultaneous discharge of a group of local principal neurons. The exact number of neurons contributing to a given population spike is not known and is not essential to our discussion, but based on lengthscale and neuronal density estimates, this may be as high as 1000 [21]. In practice, population spikes are identified as excursions of electrode potentials that are more than 3 standard deviations above or below the mean.

How does activity propagate in these networks? Figure 1 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. If a sequence is taken to mean a series of consecutively active frames bracketed by inactive frames, and the size of a sequence is the total number of electrodes activated in that sequence [18,19], then the example sequence in Figure 1 is 3 frames long and it has a size of 10.

Neuronal avalanches

Figure 2 shows that cortical slice networks produce cascades of activity such that the distribution of sizes follows a power law [18]. Because avalanches in critical sand pile models also follow a power law [22,23], the term “neuronal avalanche” was chosen to describe these cascades of neural activity [18,19]. Although first discovered in cortical slice networks, power law distributions of sequence sizes have now been reported in awake behaving monkeys [24, 25], in the isolated leech ganglion [26], and in dissociated cultures of neurons [26]. These findings suggest that neuronal avalanches are a general phenomenon, reflecting a fundamental property of neuronal networks.

It has also been shown that specific sequences of network activation patterns can persist in a given tissue culture for many hours at a time [19,27]. This finding suggests that these tissue cultures are capable of storing information as spatiotemporal patterns of activation, and that such storage even in tissue cultures is stable on a time scale of hours.

The power law behavior is not sensitive to the occasional malfunctioning electrode. In our microelectrode arrays, there are 60 microelectrodes arranged on an 8 by 8 grid with the four corner electrodes removed. If avalanches are analyzed by restricting attention to any combination of 8, 15 or 30 microelectrodes out of the total of 60 microelectrodes, the same power law behavior is observed with the same exponent [18]. Such insensitivity to scale is seen in other *scale-free* systems with power law behaviors.

The critical point

The power law distribution of avalanche sizes suggests that these networks are operating near a “critical point,” so named because the power laws are reminiscent of the critical point of phase transitions of matter [21,27–29]. Although not identical to the critical point of phase transitions, because it lacks the property of universality, the critical point of neural systems has attracted increasing theoretical and computational interest because of its importance to neural

systems function [21,30–40]. A critical neural system is balanced between a phase where activity is damped and a phase where activity is expanding. Such a state can be characterized by a simple branching ratio, here denoted by σ , which gives the average number of “Descendant” neurons that would be activated by a single “Ancestor” neuron in the previous time step:

$$\sigma = \frac{\text{Descendants}}{\text{Ancestors}} \quad \text{Eq (1)}$$

Essentially, the branching ratio σ tells us that if one neuron fires an action potential, it will cause, on average, σ other neurons to fire in response.

A key experimental finding is that the branching ratio as defined in Eq (1) tends to hover very near $\sigma = 1$ [18]. This finding has also attracted increasing theoretical and computational interest. Simple computational models have demonstrated that a branching ratio of $\sigma = 1$ indeed gives a power law size distribution, and it has the correct power of -1.5 [21,30,31]. Thus we refer to systems with $\sigma = 1$ as being at the critical point, while systems with $\sigma < 1$ are subcritical and systems with $\sigma > 1$ are supercritical. We speak of systems at criticality as having *critical connectivity*. As shown in Figure 3, activity in subcritical systems quickly dies out (in fact, exponentially fast); in supercritical systems, when $\sigma > 1$, activity escalates (also exponentially); and when $\sigma \approx 1$, activity propagates on average in a stable way without expanding or contracting.

A power law of avalanche sizes has now been reported in vitro [18,19] and in vivo [41], in organisms from leeches [26], to rats [18] and primates [24,25], suggesting that criticality may be a fundamental state of neuronal networks that is maintained in the face of perturbations. Indeed, when the branching ratio is tracked over time, it hovers around a mean value near criticality, similar to fluctuations observed in blood pressure or body temperature, parameters that are known to be homeostatically regulated. In addition, the power law distribution of avalanche sizes is maintained in cortical cultures over a period of four weeks during in vitro development [42], indicating that synaptogenesis and axonal sprouting do not disrupt criticality. The question now arises, *why*? What function is so important to biological neural systems that criticality is maintained over such diverse environments?

Criticality optimizes information processing

Computational modeling studies suggest that networks operating at the critical point can simultaneously optimize information transmission [18,37,43], information storage [30], computational power [43] and stability [30,43]. When the network deviates from the critical point, information processing and stability will be compromised [21,27]. A neural network whose job it is to process information, to learn and to adapt, must therefore maintain itself near the critical point, even when synaptic weights change strengths during the process of learning. This then is the reason why biological neural systems wish to maintain criticality: in the face of the destabilizing effects of learning, maintaining criticality re-stabilizes the system and allows the system to continue to learn.

This principle is so important to learning systems that we believe it *must* exist in biological neural systems [21]. We have suggested that the branching ratio may fluctuate away from criticality for periods of time but that it must not remain too far from criticality for too long; we have also suggested that components of a neural system may be tuned away from criticality for specific purposes but that at a large enough lengthscale, the system as a whole must be tuned to be near-critical. Failure to maintain criticality negates the usefulness of a neural system, and should, from an evolutionary perspective, be strongly selected against.

Homeostatic mechanisms

If critical homeostasis is so important and if it does exist, then what is its mechanism? Turrigiano and colleagues [44–46] have shown that homeostasis of activity (or firing rate) is maintained at a target firing rate in neuronal networks even when they are perturbed by agents that block excitatory or inhibitory transmission. Synaptic strengths are homeostatically scaled by up- or down-regulation of post-synaptic AMPA receptors to either strengthen or weaken connections when a network is perturbed. This scaling mechanism, furthermore, is multiplicative, meaning that all incoming connections arriving at a given neuron are scaled up or down by the same factor, so that the overall excitability of the neuron can be regulated without erasing information stored in individual synaptic weights. Multiplicative scaling does not erase what Hebbian learning has written.

In addition, Royer and Pare [47] have demonstrated conservation of synaptic weights after the induction of either long-term potentiation or depression. If by electrical stimulation, a synapse is made either stronger or weaker, the other synapses onto the same cell are concomitantly made weaker or stronger, so that the total incoming connectivity to a neuron is maintained constant. This mechanism appears to depend on intracellular calcium currents. In this manner, learning can occur without altering the overall level of connectivity of a system. Other candidate mechanisms may exist, including local interneuronal circuitry, and possibly backpropagation of axonal activity into the dendritic tree of the same cell.

Algorithmic consequences of criticality: connectivity in space

In this section, we explain why computational models that incorporate learning and maintain stable levels of activity and connectivity have a tendency to enter the overconnected or supercritical state under certain provocations. These provocations include status epilepticus and acute deafferentation, the latter being a model of post-traumatic brain injury. We argue that supercritical connectivity should be epileptogenic.

Recent theoretical and computational work on homeostasis of criticality and activity have done so both in the absence of Hebbian learning [31,32,48], and in its presence [21]. Perhaps as expected, critical homeostasis is more difficult to achieve in the presence of Hebbian learning, since as we have discussed, learning is a destabilizing force. Figure 4 illustrates our simplest computational model. In this model, a local grouping of neurons is represented by a “node” that can either fire a population spike spontaneously, with no input from any other node, or it can fire in response to activity at one or more other nodes. The latter is referred to as stimulated activity. Most of our simulations were performed with a total of $N = 64$ nodes, because our experimental system consists of an 8 by 8 microelectrode array with the four corner electrodes removed (thus, 60 electrodes). At any given time t , the probability that node i fires spontaneously within the next time window of 4 ms is given by $S(i;t)$, which can be different for each node and which can also vary in time. At any given time t , the conditional probability that a prior population spike at node j causes a population spike at node i , within the next time window of 4 ms, is given by $P(i,j;t)$. This conditional probability can be different between every pair of nodes and it can also vary in time. The branching ratio of each node can then be defined as the sum of outputs to all other nodes:

$$\sigma(i,t) = \left\{ \sum_{j=1}^N P(j,i;t) \right\} \quad \text{Eq (2)}$$

A corresponding measure of excitatory input into a given node i is given by the input ratio, defined as:

$$\eta(i,t) = \left\{ \sum_{j=1}^N P(i,j;t) \right\} \quad \text{Eq (3)}$$

The branching ratio is a pre-synaptic attribute while the input ratio is post-synaptic one; they are equivalent measures of connectivity. Critical connectivity occurs when the branching and input ratios are 1. To achieve critical and firing rate homeostasis, the current firing rate and input ratio for each node i are compared against the pre-set target firing rate and the critical input ratio of 1. The values of $S(i;t)$ and all the $P(i,j;t)$'s associated with that node i are scaled by small constant factors either up or down so as to approach the target firing rate and critical connectivity. A different factor is used for $S(i;t)$ and the $P(i,j;t)$'s. For example, if the current firing rate at node i is too low compared to the target rate, then $S(i;t)$ and all the $P(i,j;t)$'s associated with node i are scaled up by their respective scaling factors at every time step of the simulation, until the target firing rate is reached or exceeded. If the target firing rate is exceeded, then $S(i;t)$ and all the $P(i,j;t)$'s associated with node i are scaled down. Similar scaling is also performed for the connectivity, with two different scaling factors, for a total of four different scaling factors [21].

In the presence of Hebbian learning, we found that critical and firing rate homeostasis are independent principles and *both* must exist for a neural system to be algorithmically stable. For instance, scaling the $S(i;t)$'s and $P(i,j;t)$'s so as to achieve firing rate homeostasis alone will not guarantee that critical homeostasis is maintained. In addition, there are certain other constraints that must be satisfied for the system to be stable [21]. One important constraint is that the rate of scaling of the $P(i,j;t)$'s must be fast enough to keep up with changes due to Hebbian learning. Another important constraint is that scaling of the $P(i,j;t)$'s must operate more quickly than scaling of the $S(i;t)$'s. The faster that scaling of the $P(i,j;t)$'s takes place, relative to scaling of the $S(i;t)$'s, the more stable the system. This constraint has important consequences for how networks respond to provocations.

It is important to distinguish spontaneous from stimulated or connectivity-related activity, that is, activity due to the $S(i;t)$'s vs the $P(i,j;t)$'s. The importance arises because these two types of activity often do not change in parallel, and in fact, they often change so as to counterbalance each other. Below we give two examples:

Example 1: status epilepticus

Figure 5 shows that forced increased activity of a subset of neurons during a simulated seizure triggers homeostatic mechanisms to scale down all the $S(i;t)$'s and $P(i,j;t)$'s to very small values. When the simulated seizure stops, homeostasis causes the $S(i;t)$'s and $P(i,j;t)$'s to be scaled back up; they recover to baseline values. However, since scaling of the $P(i,j;t)$'s must operate more quickly than scaling of the $S(i;t)$'s, in fact the total connectivity as measured either by the branching or input ratio can overshoot steady state values for a time until the spontaneous firing probabilities, the $S(i;t)$'s, return to steady state values. Therefore, in the post-ictal state, the overall activity is decreased compared to baseline but the level of connectivity is *supercritical*. As a result, if and when a population spike occurs, in the post-ictal period, there is an increased chance of abnormally wide spatial spread of this excitation. The significance of activating spatially hyperextended states in a learning system is that if one such hyperextended state occurs frequently enough, the system will “learn” it and “burn” it into memory. If such a state is burned into memory, then there is an increased likelihood that that state will be reactivated again at some random time in the future. The reactivation of a spatially hyperextended state is a *necessary condition for epilepsy*, as seizures in epilepsy tend to start from the same focus in a stereotypic way, and each seizure focus must involve a macroscopic number of neurons in order to generate clinical symptomatology. Thus we claim

that prolonged post-ictal states are epileptogenic, while shorter seizures with no post-ictal state are not as epileptogenic.

Example 2: acute deafferentation

When a small section of cortex is suddenly deprived of inputs from the rest of the network, epilepsy can gradually develop [49–56]. This kind of acute deafferentation is an experimental [53,54] and computational model [55,56] for post-traumatic epilepsy. In post-traumatic epilepsy in humans, the appearance of epilepsy can be delayed for as long as 10–20 years. What is the mechanism of epileptogenesis and why does it take so long?

We simulated acute deafferentation by allowing 100 nodes to come into equilibrium with each other, and then we suddenly disconnected 10 of these nodes from the others. We then focused attention on the smaller subset of 10 nodes (see Figure 6). Our computational model predicts that acute deafferentation causes an immediate drop in firing rate and connectivity in the smaller subset of 10 nodes, because of a loss of excitatory drive from the other 90 nodes. Homeostasis will then cause both connectivity and spontaneous firing probabilities to increase. Connectivity will again adjust more quickly, and indeed it may overshoot. This situation is similar to what happens after prolonged status epilepticus; both status epilepticus and acute deafferentation provoke a hypoactive, hypersynchronous state which can result in the burning into memory of an epileptogenic spatially hyperextended pattern of activation.

In our simulation, the supercritical state is maintained for 9 hours. The time to recovery is longer if the relative magnitude of deafferentation is larger, for instance, by cutting off 20 nodes from a network of one million nodes. In addition, it may be that in real systems, the capacity for increasing the spontaneous activity is limited, such that there is a ceiling above which the spontaneous activity cannot rise. If this ceiling is below what is needed to return connectivity to the critical level, then the deafferentated nodes will remain supercritical *indefinitely*. Thus it may be that the reason epileptogenesis takes so long is that it is a learned process. It takes time for accidental repetitions of a hyperextended state to burn it into memory.

Our computational result also offers a potential explanation for the results of Graber and Prince [53,54], who found that total blockade of all activity by focal application of tetrodotoxin for a period of 3 days, if applied within 3 days after acute deafferentation, can prevent epileptogenesis in rats. Based on our simulations, we suggest that if there is *absolutely no activity* for an extended period after acute deafferentation, then no hyperextended states can actually be activated and thus no hyperextended states will be burned into memory, even though connectivity is at supercritical levels. As the spontaneous firing probability gradually approaches or exceeds the new steady state value, if activity is allowed to return to the system at that time, then the connectivity will relax back to critical levels. Once the spontaneous firing probability is near or above its new steady state value, the danger for epileptogenesis has passed.

Therefore, in the prevention of epileptogenesis after acute deafferentation, if one chooses to suppress activity prophylactically, it may be important to suppress activity to a *profound* degree, because in this time period the connectivity rises to supercritical levels and any activity at all is likely to produce a transient spatially hyperextended state. If one allows these hyperextended states to occur too frequently, by insufficient suppression of activity, then one may actually promote epileptogenesis by teaching the system to burn these hyperextended states into memory. Such a situation would be counterproductive. It would also be important that the method chosen to block activity does not suppress the *drive* towards higher spontaneous activity. If one suppresses the homeostatic drive to ramp up spontaneous activity, then one may again promote epileptogenesis, in this case by prolonging the epileptogenic period. Such a situation would also be counterproductive.

On the other hand, if the spontaneous firing probabilities can be artificially and more rapidly boosted to near steady state values, then the connectivity should decline more rapidly to its steady state value, because the drive towards higher levels of connectivity has been relieved. *That is, boosting spontaneous activity should be protective against epileptogenesis.* This counterintuitive idea arises from analysis of our computational model, and would not have been apparent without such analysis. We have conjectured that this may be one possible mechanism by which electrical brain stimulation works in the treatment of refractory epilepsy. *Electrical brain stimulation may work by boosting spontaneous activity and suppressing the supercritical state* [21].

To return to the question of homeostatic mechanisms and their relevant timescales, the synaptic scaling mechanism discussed by Turrigiano and colleagues is a slow process, with a timescale of hours to days [44–46]. Since critical homeostasis must occur on the same timescale as Hebbian learning, in order to prevent Hebbian learning from destabilizing the system, we must look elsewhere for a fast biomolecular mechanism for critical homeostasis. In addition to being fast, such a mechanism must also be nonlocal, i.e., not restricted to the level of individual synapses, because the input and branching ratios are nonlocal properties requiring simultaneous knowledge of total input and output weights across an entire node. There exists at least three candidate fast, nonlocal mechanisms: (a) When homosynaptic LTP (or LTD) is induced in the intercalated neurons of the amygdala, compensatory heterosynaptic depression (or facilitation) is observed such that the total synaptic weight of a given neuron remains constant [47]. The counterbalancing heterosynaptic response is suggestive of critical homeostasis. This mechanism depends on the release of intracellular calcium stores and has a timescale of minutes. (b) The phenomenon of backpropagation of action potentials into the dendritic tree [57–59] allows widely distributed numbers of synapses to receive nearly simultaneous information about neuronal output. This information is conjectured to play a role in LTP, LTD and STDP, but might conceivably also be used for critical homeostasis. (c) A third mechanism may involve the interaction of principal output neurons with local interneurons. It may be that a certain subset of local interneurons can sense both the total input into and total output out of a local community of output neurons. This information might then be used to modulate either the input or branching ratio of that group of output neurons. In support of this possibility, blocking interneurons with bicuculline can produce a dramatic increase in the branching ratio within minutes (Beggs unpublished).

In summary for this section, the study of the normal function of the brain can provide useful testable hypotheses regarding epileptogenesis. Epileptogenesis is usually discussed in terms of neuronal hyperexcitability, caused by an imbalance between excitatory and inhibitory drives within a neural system or subsystem, although excessive synchronization is also known to play a role [6]. Within our framework, synchronization is a normal part of brain function, wherein distinct spatial patterns of neuronal activation are activated simultaneously or near-simultaneously to represent an item of stored information. The ability of the brain to maintain a wide repertory of spatial patterns of activation requires an active homeostatic mechanism to maintain neuronal connectivity at the optimal, critical level. *Excessive* synchronization can arise when the homeostatic mechanism is provoked in certain ways so as to drive connectivity to supercritical levels. In a learning system, such provocations are epileptogenic because they can lead to the learning and burning into memory of hyperextended spatial patterns of neuronal activation. Thus the role of plasticity in epileptogenesis is not just an unhappy coincidence, a misapplication of the tools of plasticity; rather, the algorithmic requirements of a stable learning system builds in an intrinsic functional vulnerability to epileptogenesis which can be unmasked with repeated provocations.

After neuronal hyperexcitability, the formation of spatially hyperextended states represents the second necessary condition for epileptogenesis. This condition is neither more important nor

less important than the condition of neuronal hyperexcitability, but it suggests new ways of thinking about and treating epileptogenic states. For instance, one may wish to maintain connectivity in people at risk for epilepsy more closely about criticality by testing and designing electrical brain stimulation protocols that boost spontaneous activity whenever the brain enters a supercritical state. Suppressing the supercritical state in this way would not only help prevent epileptogenesis but it should also improve brain performance. In contrast, brain stimulation protocols that do not monitor connectivity may not improve brain performance and may not suppress the supercritical state. Similarly, pharmacological suppression of neuronal hyperexcitability also does not guarantee improved brain performance; rather, it more often degrades it.

Non-Markovian connectivity: connectivity in time

Are there other necessary conditions for epileptogenesis?

This question arose because even though we provoked our computational model by subjecting it to prolonged post-ictal states, we failed to detect spontaneous seizure-like events in the computational model. The system was too stable. In our original definition of the conditional probability $P(i,j;t)$, we took $P(i,j;t)$ to mean the probability that node i will fire if node j fired *in the immediately preceding time step*, 4 ms earlier. Inputs from times earlier than one time step back are “forgotten”. Such connectivity is referred to as Markovian; we chose this type of connectivity for the sake of simplicity. Although we did detect an increased activation of spatially hyperextended states, such activations did not result in what we expect to see in a seizure, namely, a rhythmic, hyperactive and repetitive activation of a spatially hyperextended state. We reasoned that there must be at least one other condition for the occurrence of spontaneous seizures, namely, that the system is able to enter an abnormally prolonged, temporally recurrent sequence of activations. Thus, we have proposed that there are at least three necessary conditions for epileptogenesis [21]:

1. Neuronal hyperexcitability: neuronal discharges occur above baseline levels.
2. Overconnectivity in space: spatially hyperextended states are learned and burned into memory.
3. Overconnectivity in time: neuronal activations enter a recurrent temporal loop and persist for an abnormally long period of time.

There are examples of brain states that satisfy one or two of the conditions of epileptogenesis, but not all three; these states do not represent true seizures. For example, the well-known phenomenon of isolated interictal spikes seen on clinical scalp EEG's represents a state of local supercritical connectivity, spanning brain areas from millimeter to centimeter lengthscales. These isolated spikes do not represent seizures because they do not persist, and patients exhibit minimal to no clinical manifestations during them. Similarly, the recent observation of persistent, oscillatory discharges recorded from intracranial micro-wire electrodes in human epileptic brain [60,61] is evidence of abnormal temporal recurrence of patterns of activity that fail to spread spatially. These “microseizures”, or more appropriately, micro-oscillations do not represent true seizures in that electrical activity remains confined to the submillimeter lengthscale, and patients show no clinical manifestation during these events. Both interictal spikes and these subclinical micro-oscillations are important markers of epileptic risk, as brain tissue in which such spikes and micro-oscillations are found is also tissue which exhibits potentiality for spontaneous seizures. But a true seizure must combine all three conditions above.

Markovian vs non-Markovian connectivity

What is necessary for a neural system to enter a prolonged temporally recurrent state? One way to do this is to relax the condition of Markovian connectivity, and to allow $P(i,j;t)$ to depend on activity at node j from more than one time step back. This kind of connectivity is known as non-Markovian connectivity. Markovian connectivity is adequate for coding static memory but is not reliable for coding temporal sequences. The reason is that the temporal links induced by Markovian connectivity are fragile. To see this, let A, B, C, D and E represent five distinct spatial patterns of nodal activation, and consider trying to learn the temporal sequence $A \rightarrow B \rightarrow C \rightarrow D \rightarrow E$. A Markovian brain would learn this sequence as 4 separate links: $A \rightarrow B$, $B \rightarrow C$, $C \rightarrow D$ and $D \rightarrow E$. If any one link is disrupted by chance, then the rest of the sequence is lost. If pattern C misfires, there is no way to look further back in time, to see that pattern B was the pattern one time step back, pattern A the pattern two time steps back, and therefore the current pattern is probably pattern C and the next pattern should be pattern D. To be able to look further back than the immediately preceding time step requires non-Markovian connectivity. In a non-Markovian brain where $P(i,j;t)$ is allowed to remember what happened 4 time steps back, the temporal sequence $A \rightarrow B \rightarrow C \rightarrow D \rightarrow E$ can then be learned entire. In this case, if patterns A and B fire correctly in sequence but pattern C misfires, it may yet be possible that the correct firing of patterns A and B contains enough information to “skip over” pattern C and to cause pattern D to fire and then E next.

It is possible to generate fairly reliable temporal sequences in a purely Markovian system by the use of many Markovian recurrent feedback loops to reinforce each temporal link [62], but such a network bears the cost of maintaining the many connections and neurons making up these feedback loops. That we did not observe this in our simulations may be that we did not incorporate enough nodes; we used only 64, because our experimental system was of approximately this size. A non-Markovian mechanism provides for reliable temporal sequencing even for small numbers of nodes, and as such is much more energy efficient.

Is there experimental evidence for non-Markovian connectivity?

Non-Markovian connectivity can best be detected by studying the cross-correlogram in the time domain. In a purely Markovian network, the cross-correlogram will fall off quickly and monotonically in time -- in fact, exponentially. In networks with both Markovian and non-Markovian connectivity, however, the cross-correlogram is expected to extend beyond this period of quick, exponential decay, and may have a long-time shoulder beyond the exponential delay or even distinct bumps that occur at more substantial delays. We have observed all 3 types of correlograms in data from cortical slice networks. Out of 60 such slices, 3 showed distinct bumps with maxima at about 120 ms, and the remainder either showed a rapid exponential decay, a hint of a shoulder or more clear evidence of a shoulder. In Figure 7, we show one example demonstrating a non-Markovian shoulder and another example with a distinct non-Markovian peak at about 120 ms. All samples were prepared in the same way (see Methods section).

The distinct bump with maximum at about 120 ms is strong evidence of non-Markovian connectivity. The only way that a purely Markovian system can produce such a peak is if neuronal excitation migrates to a region that is not being monitored by the experimental electrodes, so that the apparent interval of relative neuronal silence in the correlogram (between approximately 50–80 ms) is one during which neuronal activity is localized in the unmonitored region. However, because of our experimental setup, there is no such neural tissue that is not within range of one or more recording micro-electrode at any time.

Is the non-Markovian bump plastic?

That the non-Markovian bump appeared in only 3 out of 60 slices suggests that it was acquired through some vagary in preparation. That is, the non-Markovian bump may have been learned. *Is the bump plastic?* We stimulated 2 out of 60 electrodes at 10 Hz for 5 minutes and compared recordings consisting of two 10 minute samples chosen before stimulation and one 20 minute sample after stimulation. Figure 8 shows that a new peak appears at about -90 ms, indicating that the neurons underlying most of the 60 electrodes have become followers behind certain other neurons with a time delay of firing of 90 ms. Similar results were obtained on 3/3 trials. That we were able to change the non-Markovian bump by electrical stimulation suggests (although it does not prove) that this bump may be plastic. The time delay is far outside of the usual LTP, LTD and STDP time delays, which are on the order of 20 ms or less. Potter and coworkers have found analogous results [63,64].

The new peak in the cross-correlogram, induced by electrical stimulation, is evidence of plasticity. Experiments are underway in our lab and other labs aimed at improving control over this non-Markovian plasticity. The ultimate goal is to be able to “write” arbitrary spatiotemporal patterns onto the neural network. For instance, if all the keys of a computer keyboard are assigned an arbitrary but distinct spatial pattern of neuronal activation, then we wish to be able to write any “sentence” onto the neural network of our choosing. If there are 60 active electrodes, then there are $2^{60} \approx 10^{18}$ different possible spatial patterns, so there are more than enough spatial patterns to choose from.

What is the functional significance of non-Markovian connectivity?

The significance is that many higher cognitive skills require reliable and efficient temporal linking and sequencing [21]. That is, if A, B, C, D and E each represent a particular spatial pattern of nodal activation, then each may represent a basic unit of information. These bits of information may be linked together in a temporal sequence, to form a more complex unit of information. With more temporal links, more and more complex information can be coded, processed and stored [65–69]. In particular, the temporal dimension frees the brain from spatial limitations. The complexity of a thought can be successively extended in the temporal domain by increasing the number of temporal links in that thought, which in principle can extend to the lifetime of the animal.

To appreciate the value of non-Markovian connectivity, consider two animals, one capable only of Markovian connectivity, and the other capable of non-Markovian connectivity as well. The animal capable of strictly Markovian connectivity may learn the spatial features of a particular scene, as in a photographic still, but a non-Markovian animal would in addition be able to link these photographic stills together into a temporal sequence, in essence, into a movie. A Markovian animal may be trained to recognize the letters of the alphabet, but a non-Markovian animal can in addition link these letters together into words, sentences, paragraphs, even a novel. An external stimulus presented to a Markovian animal will appear the same each time it is presented, and the animal will then respond in the same way, no matter how many times the stimulus has been presented. A non-Markovian animal, however, will be aware that it has seen the same stimulus before, and its response may be modified accordingly. That is, a non-Markovian animal is capable of *historical context*. These ideas are summarized in Table 1.

Gamma and theta-alpha frequency timescales

It is also worth remarking on the timescales of Markovian vs non-Markovian connectivity that appear spontaneously in our samples. Figure 7 suggests that Markovian connectivity operates on a timescale of 20 ms while non-Markovian connectivity operates on a timescale of 80 to 160 ms. These timescales correspond to 50 Hz Markovian and 6–12 Hz non-Markovian

frequencies, i.e., gamma and theta-alpha band frequencies. Gamma and theta-alpha band frequencies are widespread in the intact brain [66], and often appear to be phase-locked such that gamma oscillations appear on the crests of theta oscillations. Lisman has interpreted this type of coupling in terms of temporal linkage and delimitation, with each gamma crest representing a “letter” of a “word” and with each theta-alpha crest delimiting the beginning and end of a word [65]. This coding scheme is attractive in that it implies plasticity at both gamma and theta-alpha frequencies, in order to allow the coding and linkage of different letters and different words.

Another area of research in which theta-alpha frequencies figure prominently is that of the design of stimulation protocols that optimize traditional LTP. This field has recently been reviewed [70]. In short, LTP is most strongly induced when electrical stimulation is given in bursts of gamma frequency pulses with an interburst interval at theta-alpha frequencies. Workers in this field test for LTP at the traditional gamma frequency timescales, but we suspect LTP is also occurring at theta-alpha frequencies in these experiments.

In summary, non-Markovian connectivity may not only improve temporal sequencing of neuronal spatial activation patterns, but there may also be an interplay between non-Markovian and Markovian connectivity in the construction of more complex spatiotemporal patterns. Future experiments should test for LTP not only at gamma frequencies but also at longer timescales.

The discussion above on non-Markovian connectivity may be of interest to the basic cognitive neuroscientist. Now we turn to the key question of interest for the clinician caring for patients with epilepsy.

What is the relationship to epileptogenesis?

If non-Markovian connectivity exists and if it is plastic, what happens if a temporal sequence is created (or accidentally learned) that comes back on itself, and an endlessly recurring sequence is created, e.g., $A \rightarrow B \rightarrow C \rightarrow A \rightarrow B \rightarrow C \rightarrow \text{etc?}$ If such a recurring sequence were to persist for an abnormally long time, then it would represent the microoscillations that have been observed in epileptic brain [60,61]. Such an endlessly recurrent loop would constitute the third condition for epileptogenesis, related to overconnectivity in time. When combined with the other two conditions of epileptogenesis, an electrographic seizure should result.

Of the three conditions we have proposed as necessary for epileptogenesis, the third condition is the most intriguing. If a given seizure circuit (e.g., $A \rightarrow B \rightarrow C \rightarrow A \rightarrow B \rightarrow C \rightarrow \text{etc}$) is *learned*, then it should also be possible to *unlearn it*. If we can learn how to write onto a neural network using specific spatiotemporal patterns of electrical stimulation, then we should be able to rewrite or to erase seizure circuits.

A great deal of work remains to be done before we will know if such rewriting or erasing is possible. We can imagine at least three electrical stimulation protocols to try. Protocol 1 consists of regular high frequency stimulation of a subset of neighboring electrodes at regular intervals, to mimic standard brain stimulation protocols [71–78]. If a given seizure circuit is given by $A \rightarrow B \rightarrow C$, then Protocol 2 consists of repetitively presenting the sequence $B \rightarrow A$ at random intervals. Protocol 3 consists of repetitively activating the sequence $A \rightarrow B \rightarrow C \rightarrow X$ where X is a random pattern, different for each presentation.

Our predictions are: Protocol 1 will result in suppression of seizures by lowering connectivity to subcritical levels, but the seizure circuit will not be erased and there will still be occasional activations of the seizure circuit; Protocol 2 will disrupt the seizure circuit and result in the

learning of the sequence $B \rightarrow A$, erasing not only the full seizure circuit but also the sequence $B \rightarrow C$; however, if applied too frequently, Protocol 2 may create a new, shorter seizure circuit consisting of $A \rightarrow B \rightarrow A \rightarrow B \rightarrow \text{etc}$; Protocol 3 will disrupt repetition of the seizure circuit but it will not lead to the learning of a new seizure circuit. If our predictions are verified, in slices and in animal models, then it would be of interest to apply these protocols to the treatment of human epilepsy.

Summary and future directions

We have proposed that there are at least 3 necessary conditions for epileptogenesis. The first one, neuronal hyperexcitability, is the textbook standard. Treatments aimed at suppressing neuronal hyperexcitability typically involve pharmacological suppression of excitatory drive within a circuit or boosting inhibitory drive. Such interventions do not “cure” epilepsy in that the seizure circuit remains intact and can still resurface at unpredictable times. Future interventions may involve genetic manipulations to control more precisely the component parts of the neural network. Such an undertaking would be incredibly complex, as the behavior of the whole system almost certainly cannot be predicted from control of a few component parts. We would also suggest that any intervention that targets neuronal hyperexcitability also take into account interictal baseline activity and network connectivity, as our computer simulations predict that suppression of neuronal firing rates to levels that are below set point levels may result in compensatory supercritical connectivity, which may then promote further epileptogenesis and result in the generation of new seizure circuits. In addition, perturbation of connectivity away from critical connectivity in either direction will degrade brain performance.

The second condition for epileptogenesis, overconnectivity in space, states that supercritical states are epileptogenic and that boosting spontaneous activity should relieve the drive towards supercritical connectivity. If connectivity can be restored to critical levels, then not only will seizures be suppressed but the brain may be returned to optimal performance. The third condition for epileptogenesis, overconnectivity in time, requires the greatest further theoretical and experimental development before it can be useful, but it also holds out the greatest hope for a possible *cure* of epilepsy. This hope rests on the supposition that what has been learned can be unlearned.

The second and third conditions for epileptogenesis are both built-in vulnerabilities of a system designed to learn, and both require network connectivity. As such, they are independent conditions distinct from neuronal hyperexcitability, because hyperexcitable neurons that are not connected to each other are not capable of generating seizures, no matter how hyperexcitable they are.

We emphasize that we have not proven that epilepsy involves the learning of seizure circuits, nor that epileptogenesis requires a period of supercritical connectivity. However, in exploring the constraints of stable Hebbian learning systems, we have found that there are intrinsic vulnerabilities of such systems to the development of epilepsy. Our goal in writing this paper is to encourage consideration of these vulnerabilities as candidate mechanisms of epileptogenesis. The mechanisms we have discussed are particularly relevant to the design of interventions for suppressing and possibly reversing epileptogenesis.

Future issues

There are many future issues to be resolved. First and foremost, it would be highly desirable to obtain branching ratio estimates from multiple parts of the brain under multiple conditions including the awake and asleep states, in disease and in health. However, in the presence of non-Markovian connectivity, *how does one extract the branching ratio from experiment?* The

influence of one neuron on another can be delayed for 120 ms in our experiments, and there are hints that there may be even longer timescales [Beggs, unpublished]. The simple formulas of Eq (1) and Eq (2) are still valid, but a direct causal relationship between events separated by long times can be difficult to discern. In particular, one can no longer assume that causation between neuronal firing events is restricted to the duration of an avalanche. Each avalanche likely represents only the Markovian part of the total connectivity, and misses the non-Markovian contribution. The firing of neurons in one avalanche may contribute to firing of neurons in a different avalanche, and indeed neuronal discharges in live animals may be more complex and may not even occur in the form of avalanches.

The presence of neuronal avalanches in our experimental system helped lead us to the idea of critical connectivity, but critical connectivity itself may not depend on the presence of identifiable neuronal avalanches. Indeed there have been in vivo experiments that have failed to detect power law avalanche behavior in unit neuronal activity in the live cat, observing instead an exponential dependence more typical of subcritical systems [79]. Because of the importance of critical connectivity to brain performance, we believe that it must exist even in systems where neuronal avalanches are absent. If one estimates connectivity assuming only Markovian connectivity, then one is bound to underestimate true connectivity and, consequently, one easily may mistake a critical system for a subcritical one. A reliable method for estimating the branching ratio in systems with non-Markovian connectivity, however, is still lacking. This problem promises to be challenging.

The pairwise cross-correlogram can be used to suggest when there is a causal relationship between unit potentials or population spikes occurring at two different electrodes, but it cannot be used directly to calculate a branching ratio, because besides the direct correlation it also implicitly contains cross-correlations between many intermediate interactions. That is, the cross-correlogram $C(i,j;t)$ contains interactions not just from the interaction $j \rightarrow i$ (by which is meant that a population spike at node j directly later causes a spike at node i) but it also contains interactions from “chain diagrams” such as $j \rightarrow k_1 \rightarrow i$, $j \rightarrow k_1 \rightarrow k_2 \rightarrow i$, plus all higher order chain diagrams, as well as all possible “branching” diagrams. Figure 9 shows examples of a simple chain and branching diagram.

How can the direct correlation function be extracted from the total correlogram? A similar diagrammatic expansion occurs in the theory of the radial distribution function in condensed matter physics [80]. An exact solution, whereby the exact direct correlation function is extracted from the total cross-correlogram, is not generally possible, but an approximate solution is possible, for instance through use of the Ornstein-Zernicke equation [80]. If we let $D_{OZ}(i,j;t)$ be the Ornstein-Zernicke estimate of the direct correlation function, then it is related to the total correlogram by:

$$C(i,j;t) = D_{OZ}(i,j;t) + \sum_{k=1}^N \sum_{t'=0}^t D_{OZ}(i,k;t-t')C(k,j;t') \quad \text{Eq (4)}$$

The direct correlation function can be extracted from Eq (4) by Fourier transforming Eq (4) into frequency space, solving for $\hat{D}_{OZ}(i,j,f)$ in frequency space, then backtransforming into the time domain. The branching ratio can then be estimated using:

$$\sigma_{OZ} = \frac{1}{N} \sum_{i=1}^N \sum_{j \neq i}^N \sum_{t=0}^T D_{OZ}(i,j;t) \quad \text{Eq (5)}$$

Here T is a time beyond which one does not expect any direct correlations; in practice, we find that beyond 2–3 s, Eq (5) is not sensitive to its choice [Hsu and Hsu, unpublished]. More experience using Eq (5) is required.

Once we have an estimate of the branching ratio, we will wish to know not only the mean value of the branching ratio, averaged over time, but also how widely it fluctuates over time. Are there wider fluctuations in brains that are more susceptible to epilepsy? How about the mean firing rate, how wide are those fluctuations? It may turn out that in epileptic foci, that the mean firing rate is normal but that the interictal baseline rate is subnormal, to compensate for the higher firing rates during seizures. In such brain tissue, it may be counterproductive to further suppress neuronal activity in the interictal period.

What are the timescales for LTP?

A second issue is, at what timescales is LTP possible? The traditional time window extends to 20 ms but we believe there is also a time window centered near 120 ms. It would be of great interest to determine the full range of timescales at which LTP can be induced. A similar interest also holds for the LTD and STDP time windows.

Reading the brain

Third, in spontaneously active neural tissue, how do we recognize and classify spatiotemporal patterns into a unique set of fundamental elements? This problem arises because the activation of one neuron or one set of neurons does not always result in exactly the same response from the other neurons. How do we recognize the “intended” pattern? Many solutions to this “clustering” problem have been proposed [81,82]. Since the number of possible patterns that the brain makes is enormous in number, this problem is likely to be challenging as well.

Writing on the brain

Finally, how do we “write” onto neural tissue through electrical stimulation? How do we selectively strengthen or weaken individual connections? The most obvious approach is to apply high frequency bursts consecutively first at one electrode and then at another, with interburst intervals within a known LTP time window. More complex spatiotemporal patterns may be written onto a neural network by the simultaneous application of electrical stimulation, in the spatial pattern desired, followed by a second spatial pattern delivered within a known LTP time window. Once we learn how to write onto a neural system, then we can try to create a “seizure circuit” by writing a recurrent loop; next we would try to suppress or erase this seizure circuit by the protocols discussed above. If we can, one day, do all of this, then we will have learned how to unlearn epilepsy.

Methods

Acute slices were prepared from Sprague-Dawley rats 14–35 days old (Harlan). Rats were deeply anesthetized with Halothane and then decapitated. Brains were removed and immediately placed for 3 mins in ice-cold artificial cerebrospinal fluid (ACSF) containing in mM: Sucrose 125, KCl 3, $\text{NaH}_2\text{PO}_4 \cdot \text{H}_2\text{O}$ 1.25, NaHCO_3 26, $\text{MgSO}_4 \cdot 7\text{H}_2\text{O}$ 2, $\text{CaCl}_2 \cdot 2\text{H}_2\text{O}$ 2, D-glucose 10, saturated with 95% O_2 /5% CO_2 . After cooling, brains were blocked into ~5 mm³ sections containing somatosensory cortex, striatum and thalamus. Blocks were then sliced into coronal sections with a thickness of 250 μm using the tissue slicer. After cutting, slices incubated for ~1hr at room temperature in ACSF with the same ingredients as listed above, but with 125 mM NaCl substituted for 125 mM sucrose to restore Na^+ and allow cells to fire action potentials again. After incubation, slices were adhered to microelectrode arrays with a solution of 0.1% polyethelamine that had been previously applied and let to dry for 2 hrs [16]. We attempted to place the tissue so that neocortical layers I-V covered the array. Slices were maintained thermostatically at 37° C and were perfused at 1.0 ml/min with ACSF solution containing 5 mM KCl and 0 mM Mg^{+2} during recording, which typically lasted 5 hrs. These external ionic concentrations are known to produce local field potential activity in cortical brain slices (e.g., [83,84]).

Stimulation was applied through the electrodes of a 60-channel microelectrode array. We used arrays purchased from Multichannel Systems, Reutlingen, Germany. Voltage pulses of 500 μ V in amplitude, 50 μ s duration, were delivered in a 20 Hz train for 1 s through four selected electrodes.

Acknowledgements

DH was supported by a grant from the National Institutes of Health, grant number 1KL2RR025012-01. JMB was supported by a grant from the National Science Foundation.

References

1. Scharfman HE. Epilepsy as an example of neural plasticity. *Neuroscientist* 2002;8(2):154–173. [PubMed: 11954560]
2. Williams PA, et al. Development of spontaneous seizures after experimental status epilepticus: implications for understanding epileptogenesis. *Epilepsia* 2007;48:157–163. [PubMed: 17910596]
3. Bender RA, Baram TZ. Epileptogenesis in the developing brain: what can we learn from animal models? *Epilepsia* 2007;48:2–6. [PubMed: 17910574]
4. Dube CM, et al. Fever, febrile seizures and epilepsy. *Trends Neurosci* 2007;30(10):490–496. [PubMed: 17897728]
5. Chen JW, Naylor DE, Wasterlain CG. Advances in the pathophysiology of status epilepticus. *Acta Neurol Scand Suppl* 2007;186:7–15. [PubMed: 17784531]
6. Scharfman HE. The neurobiology of epilepsy. *Curr Neurol Neurosci Rep* 2007;7(4):348–354. [PubMed: 17618543]
7. Crino PB. Gene expression, genetics, and genomics in epilepsy: some answers, more questions. *Epilepsia* 2007;48:42–50. [PubMed: 17571352]
8. Pitkanen A, et al. Epileptogenesis in experimental models. *Epilepsia* 2007;48:13–20. [PubMed: 17571349]
9. Traub RD, Contreras D, Whittington MA. Combined experimental/simulation studies of cellular and network mechanisms of epileptogenesis in vitro and in vivo. *J Clin Neurophysiol* 2005;22(5):330–342. [PubMed: 16357637]
10. Prinz AA, Bucher D, Marder E. Similar network activity from disparate circuit parameters. *Nat Neurosci* 2004;7(12):1345–1352. [PubMed: 15558066]
11. Hebb, D. The organization of behavior: a neuropsychological theory. New York: Wiley; 1949.
12. Miller KD. Synaptic economics: competition and cooperation in synaptic plasticity. *Neuron* 1996;17(3):371–374. [PubMed: 8816700]
13. Marder E, Prinz AA. Modeling stability in neuron and network function: the role of activity in homeostasis. *Bioessays* 2002;24(12):1145–1154. [PubMed: 12447979]
14. Abbott LF, Nelson SB. Synaptic plasticity: taming the beast. *Nat Neurosci* 2000;3(Suppl):1178–1183. [PubMed: 11127835]
15. Turrigiano GG. Homeostatic plasticity in neuronal networks: the more things change, the more they stay the same. *Trends Neurosci* 1999;22(5):221–227. [PubMed: 10322495]
16. Wirth C, Luscher HR. Spatiotemporal evolution of excitation and inhibition in the rat barrel cortex investigated with multielectrode arrays. *J Neurophysiol* 2004;91(4):1635–1647. [PubMed: 14627664]
17. Gholmieh G, et al. A biosensor for detecting changes in cognitive processing based on nonlinear systems analysis. *Biosens Bioelectron* 2001;16(7–8):491–501. [PubMed: 11544043]
18. Beggs JM, Plenz D. Neuronal avalanches in neocortical circuits. *J Neurosci* 2003;23(35):11167–11177. [PubMed: 14657176]
19. Beggs JM, Plenz D. Neuronal avalanches are diverse and precise activity patterns that are stable for many hours in cortical slice cultures. *J Neurosci* 2004;24(22):5216–5229. [PubMed: 15175392]
20. Tang A, et al. A maximum entropy model applied to spatial and temporal correlations from cortical networks in vitro. *J Neurosci* 2008;28(2):505–518. [PubMed: 18184793]

21. Hsu D, et al. Simple spontaneously active Hebbian learning model: Homeostasis of activity and connectivity, and consequences for learning and epileptogenesis. *Phys Rev E Stat Nonlin Soft Matter Phys* 2007;76(4 Pt 1):041909. [PubMed: 17995028]
22. Bak P, Tang C, Wiesenfeld K. Self-organized criticality: An explanation of the 1/f noise. *Physical Review Letters* 1987;59(4):381–384. [PubMed: 10035754]
23. Paczuski M, Maslov S, Bak P. Avalanche dynamics in evolution, growth, and depinning models. *Phys Rev E Stat Phys Plasmas Fluids Relat Interdiscip Topics* 1996;53(1):414–443. [PubMed: 9964272]
24. Priesemann, V.; Wibral, M.; Munk, MHJ. Detection of neuronal avalanches under incomplete sampling conditions in models of self-organized criticality and the macaque brain; Society for Neuroscience Annual Meeting; 2007.
25. Hahn, G., et al. Neuronal avalanches in vivo and in spiking activity; Society for Neuroscience Annual Meeting; 2007.
26. Mazzoni A, et al. On the dynamics of the spontaneous activity in neuronal networks. *PLoS ONE* 2007;2(5):e439. [PubMed: 17502919]
27. Beggs JM. The criticality hypothesis: how local cortical networks might optimize information processing. *Philos Transact A Math Phys Eng Sci*. 2007
28. Plenz D, Thiagarajan TC. The organizing principles of neuronal avalanches: cell assemblies in the cortex? *Trends Neurosci* 2007;30(3):101–110. [PubMed: 17275102]
29. Beggs, JM. Scholarpedia. 2007. Neuronal avalanche. www.scholarpedia.org/article/Neuronal_Avalanche
30. Haldeman C, Beggs JM. Critical branching captures activity in living neural networks and maximizes the number of metastable States. *Phys Rev Lett* 2005;94(5):058101. [PubMed: 15783702]
31. Hsu D, Beggs JM. Neuronal avalanches and criticality: a dynamical model for homeostasis. *Neurocomputing* 2006;69:1134–1136.
32. Levina A, Herrmann JM, Geisel T. Dynamical synapses causing self-organizing criticality in neural networks. *Nat. Phys* 2007;3:857–860.
33. Abbott LF, Rohrkemper R. A simple growth model constructs critical avalanche networks. *Prog Brain Res*, to be published. 2006
34. Beggs JM. How to build a critical mind. *Nat. Phys* 2007;3:834–835.
35. de Arcangelis L, Perrone-Capano C, Herrmann HJ. Self-organized criticality model for brain plasticity. *Phys Rev Lett* 2006;96(2):028107. [PubMed: 16486652]
36. Chialvo DR. Psychophysics: Are our senses critical? *Nat. Phys* 2006;2:301–302.
37. Kinouchi O, Copelli M. Optimal dynamical range of excitable networks at criticality. *Nat. Phys* 2006;2:348–352.
38. Juanico DE, Monterola C, Saloma C. Dissipative self-organized branching in a dynamic population. *Phys Rev E Stat Nonlin Soft Matter Phys* 2007;75(4 Pt 2):045105. [PubMed: 17500949]
39. Teramae J, Fukai T. Local cortical circuit model inferred from power-law distributed neuronal avalanches. *J. Comput. Neurosci* 2007;22:301–312. [PubMed: 17226088]
40. Lee KE, Lee JW. Avalanche dynamics of idealized neuron function in the brain on an uncorrelated random scale-free network. *The European Physical Journal B* 2006;50:271–275.
41. Petermann, T., et al. Neuronal avalanches in vivo; Society for Neuroscience Annual Meeting; 2006.
42. Stewart CV, Plenz D. Homeostasis of neuronal avalanches during postnatal cortex development in vitro. *J Neurosci Methods*. 2007
43. Bertschinger N, Natschlager T. Real-time computation at the edge of chaos in recurrent neural networks. *Neural Comput* 2004;16(7):1413–1436. [PubMed: 15165396]
44. Turrigiano GG, Nelson SB. Homeostatic plasticity in the developing nervous system. *Nat Rev Neurosci* 2004;5(2):97–107. [PubMed: 14735113]
45. Turrigiano GG, Nelson SB. Hebb and homeostasis in neuronal plasticity. *Curr Opin Neurobiol* 2000;10(3):358–364. [PubMed: 10851171]
46. Wierenga CJ, Ibata K, Turrigiano GG. Postsynaptic expression of homeostatic plasticity at neocortical synapses. *J Neurosci* 2005;25(11):2895–2905. [PubMed: 15772349]

47. Royer S, Pare D. Conservation of total synaptic weight through balanced synaptic depression and potentiation. *Nature* 2003;422(6931):518–522. [PubMed: 12673250]
48. Abbott LF, Rohrkemper R. A simple growth model constructs critical avalanche networks. *Prog Brain Res* 2007;165:13–19. [PubMed: 17925237]
49. Jacobs KM, Prince DA. Excitatory and inhibitory postsynaptic currents in a rat model of epileptogenic microgyria. *J Neurophysiol* 2005;93(2):687–696. [PubMed: 15385597]
50. Li H, Prince DA. Synaptic activity in chronically injured, epileptogenic sensory-motor neocortex. *J Neurophysiol* 2002;88(1):2–12. [PubMed: 12091528]
51. Jacobs KM, et al. Postlesional epilepsy: the ultimate brain plasticity. *Epilepsia* 2000;41:S153–S161. [PubMed: 10999537]
52. Salin P, et al. Axonal sprouting in layer V pyramidal neurons of chronically injured cerebral cortex. *J Neurosci* 1995;15(12):8234–8245. [PubMed: 8613757]
53. Graber KD, Prince DA. A critical period for prevention of posttraumatic neocortical hyperexcitability in rats. *Ann Neurol* 2004;55(6):860–870. [PubMed: 15174021]
54. Graber KD, Prince DA. Tetrodotoxin prevents posttraumatic epileptogenesis in rats. *Ann Neurol* 1999;46(2):234–242. [PubMed: 10443889]
55. Houweling AR, et al. Homeostatic synaptic plasticity can explain post-traumatic epileptogenesis in chronically isolated neocortex. *Cereb Cortex* 2005;15(6):834–845. [PubMed: 15483049]
56. Frohlich F, Bazhenov M, Sejnowski TJ. Pathological effect of homeostatic synaptic scaling on network dynamics in diseases of the cortex. *J. Neuroscience* 2008;28:1709–1720.
57. Johnston D, et al. Active dendrites, potassium channels and synaptic plasticity. *Philos Trans R Soc Lond B Biol Sci* 2003;358(1432):667–674. [PubMed: 12740112]
58. Stuart G, et al. Action potential initiation and backpropagation in neurons of the mammalian CNS. *Trends Neurosci* 1997;20(3):125–131. [PubMed: 9061867]
59. Waters J, Schaefer A, Sakmann B. Backpropagating action potentials in neurones: measurement, mechanisms and potential functions. *Prog Biophys Mol Biol* 2005;87(1):145–170. [PubMed: 15471594]
60. Schevon, C.; Emerson, R. Microphysiology of human epileptogenic cortex; American Epilepsy Society Annual Meeting presentation; Philadelphia, PA: 2007.
61. Worrell, G. The spatial scale of ictogenesis; American Epilepsy Society Annual Meeting presentation; Philadelphia, PA: 2007.
62. Levy WB, Hocking AB, Wu X. Interpreting hippocampal function as recoding and forecasting. *Neural Netw* 2005;18(9):1242–1264. [PubMed: 16269237]
63. Madhavan R, Chao ZC, Potter SM. Plasticity of recurring spatiotemporal activity patterns in cortical networks. *Phys. Biol* 2007;4:181–193. [PubMed: 17928657]
64. Rolston JD, Wagenaar DA, Potter SM. Precisely timed spatiotemporal patterns of neural activity in dissociated cortical cultures. *Neuroscience* 2007;148:294–303. [PubMed: 17614210]
65. Lisman J. The theta/gamma discrete phase code occurring during the hippocampal phase precession may be a more general brain coding scheme. *Hippocampus* 2005;15(7):913–922. [PubMed: 16161035]
66. Buzsaki, G. *Rhythms of the Brain*. New York: Oxford University Press; 2006.
67. Buzsaki G. Theta rhythm of navigation: link between path integration and landmark navigation, episodic and semantic memory. *Hippocampus* 2005;15(7):827–840. [PubMed: 16149082]
68. Skaggs WE, et al. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 1996;6(2):149–172. [PubMed: 8797016]
69. Maurer AP, et al. Organization of hippocampal cell assemblies based on theta phase precession. *Hippocampus* 2006;16(9):785–794. [PubMed: 16921501]
70. Albensi BC, et al. Electrical stimulation protocols for hippocampal synaptic plasticity and neuronal hyper-excitability: are they effective or relevant? *Exp Neurol* 2007;204(1):1–13. [PubMed: 17258711]
71. Morrell M. Brain stimulation for epilepsy: can scheduled or responsive neurostimulation stop seizures? *Curr Opin Neurol* 2006;19(2):164–168. [PubMed: 16538091]

72. Theodore WH, Fisher R. Brain stimulation for epilepsy. *Acta Neurochir Suppl* 2007;97(Pt 2):261–272. [PubMed: 17691312]
73. Andrade DM, et al. Long-term follow-up of patients with thalamic deep brain stimulation for epilepsy. *Neurology* 2006;66(10):1571–1573. [PubMed: 16540602]
74. Boon P, et al. Deep brain stimulation in patients with refractory temporal lobe epilepsy. *Epilepsia* 2007;48(8):1551–1560. [PubMed: 17726798]
75. Velasco AL, et al. Electrical stimulation of the hippocampal epileptic foci for seizure control: a double-blind, long-term follow-up study. *Epilepsia* 2007;48(10):1895–1903. [PubMed: 17634064]
76. Benifla M, et al. Vagal nerve stimulation for refractory epilepsy in children: indications and experience at The Hospital for Sick Children. *Childs Nerv Syst* 2006;22(8):1018–1026. [PubMed: 16816981]
77. Crumrine PK. Vagal nerve stimulation in children. *Semin Pediatr Neurol* 2000;7(3):216–223. [PubMed: 11023179]
78. Murphy JV, et al. Vagal nerve stimulation in refractory epilepsy: the first 100 patients receiving vagal nerve stimulation at a pediatric epilepsy center. *Arch Pediatr Adolesc Med* 2003;157(6):560–564. [PubMed: 12796236]
79. Bedard C, Krogen H, Destexhe A. Does the 1/f frequency scaling of brain signals reflect self-organized critical states? *Phys. Rev. Lett* 2006;97:118102. [PubMed: 17025932]
80. McQuarrie, DA. *Statistical Mechanics*. Sausalito, CA: University Science Books; 2000.
81. Theodorakis, S.; Koutroumbas, K. *Pattern recognition*. 2nd ed.. Amsterdam: Academic Press; 2003.
82. Montgomery EB Jr, Huang H, Assadi A. Unsupervised clustering algorithm for N-dimensional data. *J Neuroscience Methods* 2005;144:19–24.
83. Schiff SJ, et al. Controlling chaos in the brain. *Nature* 1994;370(6491):615–620. [PubMed: 8065447]
84. Wu JY, Guan L, Tsau Y. Propagating activation during oscillations and evoked responses in neocortical slices. *J Neurosci* 1999;19(12):5005–5015. [PubMed: 10366633]

Figure 1.

Population spikes recorded at the network level. Upper left, cortical slice on the 60-channel microelectrode array. Electrodes appear as small black circles at the ends of lines. Upper right, local field potential (LFP) signals on electrodes. Large LFPs are caused by the synchronous spiking of many neurons near the electrodes, as seen in interictal spikes. Bottom, suprathreshold LFPs represented by small black squares. A sequence of 3 active frames is shown. The sequence has a size of 10, as this is the number of electrodes driven over threshold. Adapted from Beggs [27].

Figure 2.

Power law of neuronal avalanches. Distribution of sequence sizes 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 an independent Poisson model is shown in red. The actual data follow a nearly straight line for sequence sizes from 1–35; after this point there is a cutoff induced by the electrode array size. The nearly straight line is indicative of a power law. From Beggs [27].

Figure 3.

How the branching ratio, σ , influences the spread of network activity. The left column shows the three ways activity could spread in a network, characterized by the three values of σ . The far right column shows avalanche size distributions produced by a network model with these different values of σ . Note that the power law distribution (middle plot; dashed line) only occurs when $\sigma = 1$. Adapted from Beggs [27].

Fig 4.

Our computational model. Represented are nodes i and j with respective spontaneous firing probabilities $S(i)$ and $S(j)$. The connection strengths are represented by the conditional firing probabilities $P(i,j)$ and $P(j,i)$.

Fig 5.

Recovery after simulated seizure. A, the spontaneous firing probability declines with seizure onset at arrow, and gradually recovers after seizure stops. B, Connectivity recovers faster than the spontaneous firing probability, and overshoots for 50 million time steps (or 55 hours) in this example. From Hsu et al [21].

Fig 6.

Simulated acute deafferentation. A system of 100 nodes is suddenly reduced (arrow) to 10 nodes. A, Steady state spontaneous firing probability is higher for the smaller system. B, Acute deafferentation is accompanied by an immediate drop in the relative firing rate, which slowly recovers as the spontaneous firing probability attains its new steady state value. In the meantime, connectivity rises to supercritical levels and remains there for 8 million time steps (9 hours) in this example. From Hsu et al [21].

Fig 7.

Cross-correlograms of Markovian and non-Markovian connectivity. A, Network simulation can produce either Markovian (dark line, rapidly falling) or non-Markovian (light line, with bumps) connectivity. The number of bumps can be varied by varying the non-Markovian parameters of the model (see Methods Aim 3). That is, through our simulations we can produce many bumps, one bump or no bumps. B, Actual data from two cortical slices. The dark curve is rapidly falling and indicates Markovian connectivity. There is a broad shoulder with peak at 40–50 ms suggestive of possible non-Markovian connectivity as well. There are also examples where no shoulder is apparent (not shown). The lighter curve with a distinct, separate bump peaking near 120 ms is strong evidence of non-Markovian connectivity. In this sample, there is only one non-Markovian bump. Note different time scales in the two plots.

Figure 8.

Plasticity induced by electrical stimulation at 10 Hz. Each plot shows cross-correlograms for all electrodes with respect to one reference electrode. All correlograms are superimposed to give a summary of temporal relationships in the entire network. Plots A and B were taken from the first and second 10 minutes of the baseline period. Note minor differences in correlograms, suggesting relative stability. Plot C is the overall correlogram for the 20 min pre stimulation period. Plot D is the correlogram for the 20 min period that began 10 min after stimulation. Note the appearance of a new bump (at arrow) near -100 ms. This bump is probably not the result of drift, as there was no evidence of it in the pre stimulation baseline plots A or B. Evidence for long term plasticity of this sort has been seen in 3/3 experiments so far.

Figure 9. Top diagram: simple chain diagram with one intermediate node. Bottom diagram: simple branching diagram with two parallel intermediate nodes.

Table 1

Markovian connectivity	Non-Markovian connectivity
Individual letters Individual snapshots Stereotyped response to stimuli Static memory	Words, sentences, ... novels Movies Response modified by history Sequential memory

Table 2

Conditions for epileptogenesis	Intervention
Hyperexcitability	Decrease excitatory or boost inhibitory drive, e.g., pharmacologically
Overconnectivity in space	Boost spontaneous activity → suppress supercritical connectivity
Overconnectivity in time	Recognize and unlearn seizure circuit