



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Effects of inhibitory and excitatory neurons on the dynamics and control of avalanching neural networks

Jacob Carroll, Ada Warren, and Uwe C. Täuber

Phys. Rev. E **99**, 052407 — Published 17 May 2019

DOI: [10.1103/PhysRevE.99.052407](https://doi.org/10.1103/PhysRevE.99.052407)

The effects of inhibitory and excitatory neurons on the dynamics and control of avalanching neural networks

Jacob Carroll, Ada Warren, and Uwe C. Täuber

*Department of Physics (MC 0435) and Center for Soft Matter and Biological Physics,
Virginia Tech, Robeson Hall, 850 West Campus Drive, Blacksburg, Virginia 24061, USA**

The statistical analysis of the collective neural activity known as avalanches provides insight into the proper behavior of brains across many species. We consider a neural network model based on the work of Lombardi, Herrmann, De Arcangelis et al. that captures the relevant dynamics of neural avalanches, and we show how tuning the fraction of inhibitory neurons in this model alters the connectivity of the network over time, removes exponential cut-offs present in the distributions of avalanche size and duration, and transitions the power spectral density of the network into an “epileptic” regime. We propose that the brain operates away from this power law regime of low inhibitory fraction to protect itself from the dominating avalanches present in these extended distributions. We present control strategies that curtail these power law distributions through either random or, more effectively, targeted disabling of excitatory neurons.

I. INTRODUCTION

Neurons are the ubiquitous cells found in all intelligent animals, whose operations and dynamics are responsible for cognition. Individually the dynamics of neurons are well known [1–5], but neurons are connected dynamically in ensembles of billions, and the collective behavior of these networks of neurons that gives rise to high-level cognitive functions such as memory and consciousness is only partially understood.

A common method for analyzing the collective behavior of biological neural networks is through the study of neural avalanches. An avalanche is a period of continuous neural activity in the network where signals are continually transmitted from neuron to neuron. Distributions of various quantities associated with these neural avalanches, such as the avalanche size, duration, and power spectral density, are observed to follow power laws, and the exponents of these power laws appear to govern the proper functioning of the network [6–8].

In this paper we discuss a model based on the work of Lombardi, Herrmann, De Arcangelis et al. [9–15] of an avalanching neural network, that correctly reproduces the power law behavior of the avalanche size distribution, duration distribution, and power spectral density of neural activity. We show how these power laws are affected by modifying the fraction of inhibitory neurons, neurons that serve to suppress signals in the network, and observe intriguing extended power law behavior indicative of criticality in the avalanche size and duration distributions, as well as exponents suggestive of epileptic behavior in the power spectral density at low inhibitory fractions. We also monitor how the outgoing connectivity distribution of the network evolves under the effects of different inhibitory fractions.

Finally, we present two distinct strategies to control and remove the extended tails of the avalanche size and

duration distributions in networks with low inhibitory fractions through the disabling of either random or highly connected neurons. Removing these extended tails serves to protect these networks from the extreme avalanches that occur in these extended distributions.

A. Neurons

Biological neural networks are composed of individual neurons connected to each other by synapses. Synapses are small gaps between neighboring neurons where neurons can release and receive neurotransmitters: chemicals that cause the receiving neuron to open or close ion gates and pumps in order to increase or decrease its membrane potential, depending on the neurotransmitter received. Some neurons solely release inhibiting neurotransmitters, and will be referred to as “inhibitory neurons.” These neurons serve to suppress signals in brain, and make up 20-30% of the neurons in the human cortex [15, 16]. As a matter of definition, the neuron releasing the neurotransmitters will be referred to as the “pre-synaptic” neuron, and the neuron receiving the neurotransmitters will be referred to as the “post-synaptic” neuron [1]. A schematic of a synapse is shown in Fig. 1.

If the membrane potential of the post-synaptic neuron is increased beyond a threshold value then the post-synaptic neuron generates a spiking potential that travels down the length of the post-synaptic neuron’s axon. Synapses that this signal reaches will release neurotransmitters of their own, to be picked up by other neurons. This signal is referred to as an “action potential,” and this process will summarily be called “firing.” It is important to note that the connections between neurons are not symmetric. It is not necessarily true that if neuron A can transmit to neuron B, then neuron B can transmit to neuron A [1].

Pre-synaptic neurons that have just fired lock down their ion channels for a period of 3-4 ms [2]. During this period, the neuron cannot respond to any received

* jac21934@vt.edu; tauber@vt.edu




FIG. 1: A schematic of a synapse. The upper half of the image represents the pre-synaptic neuron, while the lower half represents the post-synaptic neuron. The gap between the pre- and post-synaptic neurons is the synapse. Neurotransmitters (small red circles) are collected in the synaptic vesicles (circular cavities in the pre-synaptic neuron). When the neuron fires, the synaptic vesicles are transported to the surface of the neuron at the synapse and expel their neurotransmitters. The neurotransmitters propagate through the synapse and bind to receptors on the surface of the post-synaptic neuron (red-orange half-circles). These receptors activate ion channels (paired green rectangles) on the post-synaptic neuron which cause the membrane potential of the post-synaptic neuron to change as ions are transported into or out of the post-synaptic neuron’s cell body [1, 2].

neurotransmitter, and is for our purposes dormant. This period of quiescence is referred to as the “refractory period,” and serves (among other purposes) to enforce that action potentials propagate unidirectionally down axons [1].

The strength of signals transmitted across synapses is determined by the combination of several different factors from both the pre- and post-synaptic neurons, such as the quantity of neurotransmitters produced by the pre-synaptic neuron, and the number of neurotransmitter receptors available on the post-synaptic neuron’s surface. The strength of these signals can change over time [3–5].

One such mechanism for the change in synaptic strength is Hebbian plasticity, a rule which states that synapses which successfully perpetuate signals, i.e., when the pre-synaptic neuron causes the post-synaptic neu-

ron to fire, increase their strength [17]. This increase in synaptic strength and efficiency can be affected by a variety of processes, which include the increased production of neurotransmitters in the pre-synaptic neuron [3] or an increase in the number of receptors on the post-synaptic neurons [4]. Other types of neural plasticity can be present in synapses, but the model devised and described in detail in Refs. [9–15] only attempts to reproduce Hebbian plasticity in the pursuit of constructing a simplified, minimally complex model system that still reproduces many pertinent non-trivial observed features.

B. Avalanches

In the context of neural networks, avalanches are defined as a period of continuous neural activity. The name is chosen in analogy to Abelian sandpile models, where a column of sand can topple sending sand down onto lower columns causing more grains to topple, and so forth, creating a literal avalanche of sand. Sandpile avalanches are a classic example of “self-organized criticality,” where the balance between the addition to and cascading depletion of the sandpiles appears to drive the system to a critical point without any obvious external tuning or adjustment [18, 19]. It is however now well-understood that the assumed time scale separation in the model between instantaneous toppling and external driving constitutes a relevant control parameter in the system, and the corresponding ratio of rates in fact needs to be tuned to zero to observe dynamic critical behavior. The collective dynamics of the avalanching system are then characterized by a set of universal power laws, as is typical of critical systems. Among other system quantities, the avalanche size and duration distributions both follow algebraic decay with exponents $-3/2$ and -2 , respectively [18].

The same cascading process can happen among neurons, where at least one neuron fires, sending signals to other neurons and causing them to fire, etc. until the system reaches a time where no neurons in the network can exceed the firing threshold. This period of continuous neuron firing constitutes a neural avalanche, and the system’s dynamics displays (near-)critical behavior when the characteristic time scale for internal firing cascading is again assumed much shorter than that of external stimuli, as is indeed implemented in the model of Refs. [9–15].

The dynamics of neural avalanches were studied by Beggs and Plenz in rat cortex cultures, and various properties of avalanches in these cultures were found to follow power law distributions [6]. The observables investigated included the avalanche size: the sum of all signals sent by firing neurons; and the avalanche duration: the length of time that the avalanche persists. The avalanche size and duration distributions followed the same exponents of $-3/2$ and -2 that are expected in the Abelian sandpile models [6, 18]. These results for avalanche size have since been replicated in macaque monkeys [8].

Additionally, power law behavior has been recorded

in the power spectral density of neural activity in humans using electroencephalography and electrocorticography [7], and avalanche models have been shown to replicate this behavior, even matching exponents seen experimentally [12].

II. NEURAL NETWORK MODEL

Our numerical model is based on the work of Lombardi, Herrmann, De Arcangelis et al. [9–15], and in addition to modeling key features of biological neural networks; namely firing at a threshold potential, refractory periods, and Hebbian learning [17], the model accurately reproduces several experimentally determined distributions related to avalanches of neural activity in biological neural networks [6, 7].

We note that an extension to this model developed by Lombardi, Herrmann, De Arcangelis et al. also recreates the avalanche waiting time distributions is described in Appendix A. The following simplified model variant however does not attempt to address these waiting time distributions, and hence this extension is excluded in the present study to reduce the complexity of the system. Table I lists the various parameters used in our model. We

TABLE I: The network parameters used in the avalanching neural network model.

Parameter	Description	Value
N	Number of neurons in the network	64000
t	Time step	4-6 ms
p_{inh}	Fraction of inhibitory neurons in the network	—
$J_{ij}(t)$	Weight of connection from i^{th} to j^{th} neuron	—
$g_{ij}(t)$	Weight scaled by degrees of connectivity	—
J_{min}	Minimum weight strength	0.001
J_{max}	Maximum weight strength	2
$n_i(t)$	Potential of i^{th} neuron	—
$s_i(t)$	Action potential of i^{th} neuron	—
n_{max}	Threshold potential	-55 mV
$k_{\text{in}_i}(t)$	Number of incoming connections to i^{th} neuron	—
$k_{\text{out}_i}(t)$	Number of outgoing connections from i^{th} neuron	—

would like to emphasize that the following model is definitely a very simplified representation of true neuronal networks and certainly does not grasp the full complexity of such systems, but still reproduces crucial pertinent power law features that are our focus in this study.

A. Neuron dynamics

The model consists of a number of neurons N , each neuron i defined by its potential n_i . At the beginning of the simulation, each neuron is randomly designated as inhibitory, with probability p_{inh} , or excitatory with probability $1 - p_{\text{inh}}$. This will determine whether signals from

this neuron increase (excitatory) or decrease (inhibitory) the potentials of other neurons.

The neuron is then randomly assigned an out-degree k_{out_i} from a truncated power law distribution, formed such that $P(k_{\text{out}}) \sim k_{\text{out}}^{-2}$ for $k_{\text{out}} \in [2, 100]$. This range was chosen to mimic the distribution of connectivity experimentally observed in human cortices, which demonstrates power law behavior across two decades of connectivity, following an exponent of -2 [20]. Additionally, the truncated nature of the power law also allows this distribution to be normalized. The k_{out_i} neurons are then chosen from a uniform distribution, and connections between them and the i^{th} neuron are established by assigning each synapse an initial weight J_{ij} uniformly distributed on the interval $(0, 1)$. Connections from a neuron to itself are not allowed in the model. Once the initial network topology is established, the in-degree k_{in_i} is tabulated for each neuron.

The network is initialized such that the potential of each neuron is set to 90% of the threshold potential $n_{\text{max}} \sim -55$ mV to facilitate and accelerate the initial building-up of network activity.

The system dynamics are discretized into time steps t . A single time step is taken to be the length of time an individual neuron takes to fire, which is in the range of 4-6 ms. During each time step t , any neuron whose potential has increased past the system's firing threshold, $n_i(t) \geq n_{\text{max}}$, fires, sending an action potential $s_i(t)$ to each of the $k_{\text{out}_i}(t)$ connected neurons. If the potential of the i^{th} neuron is not above the threshold potential, the action potential is zero:

$$s_i(t) = \begin{cases} 0, & n_i(t) < n_{\text{max}} \\ n_i(t), & n_i(t) \geq n_{\text{max}} \end{cases} \quad (1)$$

After a neuron has fired, its potential is set to zero for one time step, during which it cannot receive signals from other neurons. This mimics the refractory period of real neurons [1]. The signal received by the post-synaptic neurons not in a refractory period is proportional to the action potential of the pre-synaptic neuron. This behavior is summarized in Eq. (2),

$$n_j(t+1) = \begin{cases} 0, & s_j(t) > 0 \\ n_j(t) \pm g_{ij}(t)s_i(t), & s_j(t) = 0 \end{cases} \quad (2)$$

where the upper and lower signs are for i excitatory and inhibitory, respectively, and $g_{ij}(t)$ controls how much the j^{th} neuron is affected by signals from the i^{th} neuron:

$$g_{ij}(t) = \frac{k_{\text{out}_i}(t)}{k_{\text{in}_j}(t)} \frac{J_{ij}(t)}{\sum_k J_{ik}(t)} \quad (3)$$

While the strength of the synaptic signal is proportional to the pre-synaptic neuron's potential, this signal is scaled by a factor $J_{ij}(t)/\sum_k J_{ik}(t)$ that determines the relative strength of the connection between the i^{th} and j^{th} neurons compared to all connections from the

i^{th} neuron. This is then in turn scaled by the factor $k_{\text{out}_i}(t)/k_{\text{in}_j}(t)$.

While $J_{ij}(t)/\sum_k J_{ik}(t)$ determines the strength of the $i \rightarrow j$ connection relative to all of the i^{th} neuron's connections, the factor $k_{\text{out}_i}(t)/k_{\text{in}_j}(t)$ rescales each of these connections relative to their importance to the network. This is necessary in order to properly compare the strength of signals from different neurons.




FIG. 2: An example network of six neurons showcasing the various possible interactions between neurons. The colored circles represent six different neurons, labeled n_i , while the arrows connecting them represent the synaptic connections between neurons with the corresponding elements of the weight matrix g centered in each line. The (blue) neurons n_2 , n_4 , and n_5 have zero action potential and send no signals through their connections (the black lines) but receive signals from firing neurons. The (red) neurons n_1 and n_6 are firing, sending their non-zero action potentials through their outgoing connections (the red dashed lines). The (grey) neuron n_3 is in a refractory period, and can neither send or receive signals through its synaptic connections (the grey dash-dotted lines).

A neuron with many outgoing connections will be more important to the network than a neuron with few, and so its outgoing signals will be scaled by the neuron's number of outgoing connections, $k_{\text{out}_i}(t)$. A neuron that receives signals from many different neurons will be less excited by any one connection, so any incoming signals to it will be scaled by its number of incoming connections, $k_{\text{in}_j}(t)$. Indeed, this factor $k_{\text{out}_i}(t)/k_{\text{in}_j}(t)$ is responsible for the power law distribution of the avalanche sizes, though it has no bearing on the waiting time or duration distributions. The change in potential in Eq. (2) is referred to

as the depolarization of j due to i for this time step.

A series of successive time steps during each of which at least one neuron fires constitutes an avalanche.

Figure 2 shows a schematic of a simple, six neuron network with examples of different possible connections and neuron behavior.

B. Hebbian learning and pruning

After each avalanche, the strength of connections between neurons is adjusted according to Hebbian-like rules [17], and then pruned (set to zero) if the strength of connection drops below a threshold J_{min} .

Due to the variable length of each avalanche, it is convenient to index avalanches on a separate variable τ where each avalanche has beginning and ending times $t_i(\tau)$ and $t_f(\tau)$. At the end of each τ^{th} avalanche, we implement Hebbian-like plasticity rules in parallel across the synaptic connections between all neurons. The strength of each synapse J_{ij} is increased proportional to the sum of all signals sent through the synapse during the avalanche, and decreased by the average increase in synaptic strength. We cap each J_{ij} to a maximum value of J_{max} in order to ensure stability in the network. The change in synaptic strength is summarized in Eq. (4):

$$J_{ij}(\tau) = J_{ij}(\tau - 1) + \frac{\delta n_{ij}(\tau)}{n_{\text{max}}} - \Delta J(\tau), \quad (4)$$

where $\delta n_{ij}(\tau)$ is the sum of magnitudes of all signals sent from neuron i to neuron j during the τ^{th} avalanche,

$$\delta n_{ij}(\tau) = \sum_{t=t_i(\tau)}^{t_f(\tau)} \frac{k_{\text{out}_i}(t)}{k_{\text{in}_j}(t)} \frac{J_{ij}(t)}{\sum_k J_{ik}(t)} s_i(t), \quad (5)$$

and $\Delta J(\tau)$ is the average increase in connection strength after the τ^{th} avalanche,

$$\Delta J(\tau) = \frac{1}{N_C(\tau)} \sum_{i=1}^N \sum_{j=1}^N \frac{\delta n_{ij}(\tau)}{n_{\text{max}}}. \quad (6)$$

Here, $N_C(\tau)$ is the number of non-zero connections in the network,

$$N_C(\tau) = \sum_{i,j} \Theta(J_{ij}(\tau)),$$

where Θ represents Heaviside's step function. The competition between the second and third terms of Eq. (4) causes used connections to increase in strength, while unused connections weaken. If any connection is lowered below a threshold J_{min} , the connection is permanently removed as that element J_{ij} is set to zero. This removal of weak connections is called "pruning."

To prevent over-pruning, we impose an upper bound of J_{max} on the strength of any given connection. This

strengthening procedure can saturate this bound, but not exceed it. This rule forces the network to prioritize those connections which are most often used, as in biological neural networks [17].

Between avalanches, the system is stimulated via small ($\sim 1\%$ of the threshold potential) constant noise applied to randomly chosen neuron potentials. This noise tends to drive the system towards another avalanche.

III. DISTRIBUTIONS OF AVALANCHE PARAMETERS

Several different parameters of neural avalanches can be studied through statistical analysis, and have been shown to be important to the proper operation of biological neural networks [6–8, 21]. These observables are the avalanche size, the avalanche duration, and the power spectral density of neural activity. In this section we describe each of these quantities in turn, and our methods for modeling and recording them.

A. Avalanche size distribution

The size of the τ^{th} avalanche is defined to be the sum of absolute values of all signals sent between neurons during the avalanche.

Biologically, this is the sum of all neuron action potentials. This has been recorded experimentally through the careful placement of electrodes on both *in-vivo* and *in-vitro* neural networks [6, 8]. The size of many different avalanches can be collected to form a probability distribution describing the likelihood of a given avalanche having a certain size P_S , where S is the size of a given avalanche. This distribution P_S has been found to follow a power law of $P_S(S) \sim S^{-1.5}$ [6, 8].

We calculate the total avalanche size by summing each signal sent between neurons during an avalanche. If each avalanche has beginning and ending times $t_i(\tau)$ and $t_f(\tau)$ respectively, we define the size of the avalanche as:

$$S(\tau) = \sum_{t=t_i(\tau)}^{t_f(\tau)} \sum_{i=1}^N \sum_{j=1}^N \frac{k_{\text{out}_i}(t)}{k_{\text{in}_j}(t)} \frac{J_{ij}(t)}{\sum_k J_{ik}(t)} s_i(t). \quad (7)$$

B. Avalanche duration distribution

The duration of an avalanche is the length of time that the avalanche persists. This has been recorded experimentally in the same manner as the avalanche size, and has been collected into a distribution P_D describing the likelihood of a given avalanche duration. Experimentally, this distributions has been observed to follow a power law, $P_D(D) \sim D^{-2.0}$ followed by an exponential cut-off [6].

The duration of the τ^{th} avalanche, $D(\tau)$, is taken to be the number of time steps that the avalanche persists for. This can be written as

$$D(\tau) = t_f(\tau) - t_i(\tau), \quad (8)$$

where $t_i(\tau)$ is the initial time step of the τ^{th} avalanche, and $t_f(\tau)$ is the final time step of the τ^{th} avalanche.

C. Avalanche power spectral density

The power spectral density (PSD) of a signal describes the distribution of power in the signal as a function of frequency. This is a common form of analysis done on the measurements of *in-vivo* neural activity via techniques such as electroencephalography and electrocorticography. Electroencephalography and electrocorticography both measure the action potentials of neurons firing in living brains through the placement of electrodes either outside (electroencephalography) or inside (electrocorticography) the skull. The time series measurement of electrical activity can be decomposed into their power spectral density and has been suggested as a means to diagnose epilepsy [7, 21]. Healthy non-epileptic brains display a PSD exponent in the range of $(-1.5, -0.8)$ (at large frequencies f) [15], while brains undergoing epileptic events have been recorded with PSD exponents in the range of $(-2.2, -1.8)$ [7, 15].

In our simulations the sum of all depolarizations is calculated after each time step in an avalanche. This sum is then appended to a time series as x_n , the n^{th} sum of depolarizations. We perform this summation for each time step in every avalanche until the series $\{x_n\}$ contains the sum of depolarizations for every time step of every avalanche.

The power spectral density of time series data can be determined by computing this series' discrete Fourier transform. The average of the square of the contribution of each frequency in the discrete Fourier transform gives the power spectral density of that frequency,

$$PSD(f) = \frac{1}{N} \left| \sum_{n=1}^N x_n e^{-ifn} \right|^2, \quad (9)$$

where PSD is the power spectral density as a function of the frequency f , x_n is the sum of depolarizations in the n^{th} time step, and N is the total number of time steps that these depolarizations were recorded for. We compute the PSD across the entire frequency range using Welch's method [22] and record it in a histogram using logarithmic binning to smooth out fluctuations that occur at the lower frequency ranges due to their limited occurrence in our data. This limited occurrence is due to the power law nature of the model's PSD (see Fig. 6).

IV. RESULTS

For each of the following results, with the exception of the finite-size effects plot shown in Fig. 4, we simulated 100 different networks of 64,000 neurons, each randomly initialized according to the methods described in Sec. II A. Each network was allowed to operate for 10,000 – 100,000 separate avalanches, after which the various distributions described in Sec. III were calculated and averaged across the 100 separate networks. The averaged distributions are recorded below.

A. Avalanche size distribution

Using the network described in Sec. II A and Eq. (7) we measured the distribution histogram of avalanche sizes. Figure 3 shows the avalanche size distributions for two different networks, each made of 64,000 neurons. The two networks differ only in their inhibitory fraction p_{inh} . The (blue) dots represent the avalanche size distribution of a network with $p_{inh} = 0.04$ and the (orange) triangles represent the avalanche size distribution of a network with $p_{inh} = 0.10$. The (green) dashed line indicates a power law with exponent -1.55 . In both cases the early avalanche size distributions ($S \in [10^1, 10^5]$) follow a power law of $P_S \sim S^{-1.55}$, which agrees well with the distributions found in rat and macaque monkey cortices [6, 8].

For larger inhibitory fractions, the power law behavior reaches an exponential cut-off at high avalanche sizes as the network is unable to sustain the activity necessary for massive avalanches. Again, as seen in the avalanche duration distributions, as the inhibitory fraction decreases the network becomes better able to sustain increasingly larger avalanches, and at an inhibitory fraction of $p_{inh} = 0.04$ we see the exponential cut-off disappear as the power law behavior is extended for several more decades. This extension of the power law behavior due to these massive avalanches suggests the system is approaching a critical regime as the inhibitory fraction is lowered.

The “hump” displayed in both sets of data is a finite-size effect related to the total number of neurons in the network. Because the network is stimulated by small constant noise between avalanches, every neuron in the network will, on average, be very close to firing when an avalanche begins. This allows avalanches to initially propagate more easily through the network, until every neuron has fired at least once. After this, this initial “supply” of neuron potential has been exhausted, and avalanches must be self-sustaining to continue past this point. Many avalanches are not strong enough to continue propagating without many neurons in the network having highly elevated potentials, and so many avalanches end with this value of size, creating a broad local maximum in the distribution. The dependence on the location of this hump to system size is shown in Fig.




FIG. 3: Avalanche size distributions for two 64,000 neuron networks with differing inhibitory fractions p_{inh} . All other network parameters are as described in Sec. II A. The x axis represents the avalanche size as defined in Sec. III A. The y axis represents the probability of an avalanche occurring with that size. The (blue) dots are the distribution of avalanche size for a network with $p_{inh} = 0.04$ inhibitory fraction. The (orange) triangles are the distribution of avalanche size at $p_{inh} = 0.1$. The (green) dashed line indicates a power law with exponent -1.55 . Note that for $p_{inh} = 0.04$ the exponential cut-off seen at $S = 5 \cdot 10^5$ in the $p_{inh} = 0.1$ data disappears, and the power law behavior extends to much larger avalanche sizes. Avalanches were recorded with sizes up to $S \sim 10^{10}$, but these were excluded from the figure due to poor statistics. The bump seen at $S \sim 10^6$ in both distributions is a finite-size effect that is proportional to the simulation value of the neuron firing threshold multiplied by the number of neurons in the network. This relation is shown in greater detail in Fig. 4.

4. This finite-size effect is also apparent in the avalanche duration distribution, Fig. 5.

B. Avalanche duration distribution

Using the network described in Sec. II A and Eq. (8) we measured the distribution histogram of avalanche durations. Figure 5 shows the avalanche duration distributions for two networks of 64,000 neurons. The (blue) dots represent the distribution of a network with inhibitory fraction $p_{inh} = 0.10$, while the (orange) triangles represent the distribution of a network with inhibitory fraction $p_{inh} = 0.04$. The (green) dashed line indicates a power law with exponent -2.1 .

For short avalanche durations both distributions agree well with the experimental results for the avalanche duration distribution which follows a power law with exponent of -2.0 followed by an exponential cut-off [6]. The




FIG. 4: The avalanche size distribution for four different system sizes, each with inhibitory fraction of 0.09. The data for each system size was averaged from fifty different randomly initialized instances of the network. The (blue) circles, (orange) inverted triangles, (green) squares, (red) diamonds each represent the distribution of avalanche size for networks with $N = 100$, $N = 1000$, $N = 10000$, and $N = 64000$ neurons, respectively. Each distribution follows a power law before reaching a “hump” that is followed by a cut-off. In each case the location of this hump is very close to the system size (number of neurons) multiplied by the simulation value of neuron threshold potential. These humps form because in between avalanches we repeatedly stimulate randomly chosen neurons with a small increase in the neuron’s potential until a neuron fires. This small stimulation guarantees that on average, when the first neuron in an avalanche fires a neuron fires many neurons will have a potential close to the threshold potential. These stimulated potentials will help sustain the avalanche until either the avalanche ends through some fluctuation in the system, or until all neurons have on average fired once.

distribution with higher inhibitory fraction, $p_{inh} = 0.10$, matches the experimental results very well, while the distribution from the network with low inhibitory fraction $p_{inh} = 0.04$ does not show the experimentally found exponential cut-off, and instead displays continued power law behavior at the same exponent for several more decades before statistics of the measured events becomes too poor. This increase in available avalanche durations is due to the inability of the network to suppress signals because there are few inhibitory neurons in the network. This generates very long lasting avalanches that in turn give rise to the extended power law seen in the avalanche size distribution shown in Fig. 3. These long-lasting avalanches ultimately dominate the dynamics of the network, because while they are rare, these avalanches can be up to 10^4 time steps longer than




FIG. 5: Avalanche duration distributions for two 64,000 neuron networks with different inhibitory fractions p_{inh} . The x axis represents the duration D of an avalanche in time steps of single neuron firing periods, which are in the range of 4 – 6 ms. The y axis represents the probability of this avalanche duration. The (blue) dots represent the distribution of avalanche durations with $p_{inh} = 0.04$. The (orange) triangles represent the distribution of avalanche durations with $p_{inh} = 0.1$. The (green) dashed line indicates a power law with exponent -2.1 . As with the avalanche size distributions (see Fig. 3) the exponential cut-off seen at $D = 3 \cdot 10^2$ time-steps when $p_{inh} = 0.1$ disappears and the power law behavior persists for several decades at the same exponent, before the statistics of the events become too poor. This extended tail for the lower inhibitory fraction ultimately dominates the dynamics of the system, as the avalanches in this regime last for orders of magnitude more time steps than the original regime. This tail continues for several more decades, but these data points have been trimmed from the plot due to poor statistics. The hump seen in both the high and low inhibitory fraction data around $D \sim 5 \cdot 10^2$ is a finite-size effect related to the duration necessary for all neurons in the network to fire, on average, once.

avalanches seen in a network with a higher inhibitory fraction.

The inhibitory fraction at which we see these long-lasting avalanches occur is much lower than the value observed in human cortices, which is around 0.2 – 0.3 [15, 16]. Milton et al. [23] suggest that the exponential cut-offs seen in the avalanche distributions exist to protect the brain from runaway avalanches; our results corroborate this idea, namely that the brain might operate away from this regime in order to not be dominated by the incredibly long-lasting avalanches present at low inhibitory fractions, so the brain actually ultimately avoids truly critical behavior, yet still maintains a highly cooperative and correlated state.




FIG. 6: Two power spectral densities of 64,000 neuron networks with different inhibitory fractions p_{inh} . The (blue) dots represent the power spectral density of the network with $p_{inh} = 0.04$. The (red) squares represent the power spectral density of the network with $p_{inh} = 0.30$. The (blue) dashed line indicates a power law with exponent -2.04 , an exponent in the regime of epileptic behavior seen in humans [7, 15]. The (orange) dash-dotted line represents a power law with exponent -1.0 , which is in the range of exponents observed in “healthy” human brains [7, 15]. While there appears to be a well defined power law for the $p_{inh} = 0.04$ data, the power law regime shrinks as we raise the inhibitory fraction to $p_{inh} = 0.30$, and we at most observe a brief period of power law behavior around $f \sim 10^{-1}$ where there is approximately a power law with an exponent close to -1 . To more effectively interpret this data, we calculated the local derivative between each consecutive data points and plotted the effective exponents of each of the two PSDs in Fig. 7(a) and 7(b).

Additionally, both the avalanche size and duration distributions were found to be very stable with respect to the noise strength, minimum and maximum weight strengths, as well as the threshold potential, with changes in these parameters resulting in little to no modifications in their dynamics.

C. Power spectral density

Using Eq. (9) we constructed the power spectral density of our networks of 64,000 neurons for different inhibitory fractions. Figure 6 shows two different network power spectral densities for inhibitory fractions 0.04 and 0.30.

The PSD data for the $p_{inh} = 0.04$ network is represented by the (blue) circles, while the data for the $p_{inh} = 0.30$ network is represented by the (red) triangles. Both PSDs displayed in Fig. 6 show two distinct

dynamical regimes: a power law regime at mid to low frequencies, and a semi-constant regime at high frequencies. The PSD’s power law regime is due to long-range temporal correlations present between neuron firings. It is an assumption of this model that neurons firings are uncorrelated across different avalanches, so this power law regime becomes more pronounced as the inhibitory fraction decreases, due to the increasing accessibility of long duration avalanches. This behavior corroborates results shown by Lombardi et al. [15]. We also observe the range of power law behavior to shrink as the inhibitory fraction is increased. While the network with a lower inhibitory fraction shows a very clear power law following an exponent around -2.04 , (shown in Fig. 6 as the blue dashed line), the higher inhibitory network seems to at most follow a power law with exponent around -1.0 (represented in Fig. 6 by the orange dash-dotted line) for only a small portion of its frequency range, if at all.

This decrease in the power law regime is most likely due to the suppression of highly correlated neuron firing events by the increased fraction of inhibitory neurons. To investigate the true extent of the power law regime we extract the effective exponents of each of the PSDs shown in Fig. 6 and plot the effective exponents as a function of frequency in Figs. 7(a) and 7(b).

The effective exponents shown in Fig. 7(a) showcase very clear power law behavior for the low inhibitory network with exponent around -2.0 . This value of the exponent is in the regime of experimentally observed epileptic behavior seen in humans [7, 15] and we observe this behavior for the same values of inhibitory fractions where we observed the extended avalanche size and duration distributions.

Figure 7(b) displays no true power law behavior, with only a small possible plateau of the effective exponent around -1.0 . While this value is in the regime of normal operating brain behavior for human PSDs [7, 15], and a transition of exponents from a value of -1.0 to -2.0 as the inhibitory fraction of these networks is lowered was reported by Lombardi et al. [15], we cannot claim to see true power law behavior at biologically relevant inhibitory fractions of $p_{inh} = 0.30$.

D. Neuron connectivity distribution

In addition to the various avalanche distributions, we can also observe how the connectivity of the network evolves over time for different inhibitory fractions.

As avalanches occur in the network, the Hebbian plasticity rules defined in Sec. IIB will cause the connections between neurons to change according to their use. Frequently used connection will be strengthened, and infrequently employed connections will be weakened or even “pruned” (i.e., removed) if the connections become too small. We are interested in observing how the distribution of outgoing connectivity evolves as function of time and inhibitory fraction under these effects.




FIG. 7: Double-logarithmic plot of the local derivative between consecutive pairs of points for the PSD of the (a) $p_{\text{inh}} = 0.04$ network and (b) $p_{\text{inh}} = 0.30$ network shown above in Fig. 6. The y axes display the effective exponent of any power law behavior observed across the two data sets, while the x axes represent the average frequency between the pairs of PSD data points that were used to calculate each local derivative. True power law behavior in the PSDs will then be displayed on this graph as a plateau in the effective exponent across many different frequencies. This behavior is clearly visible in (a), where the effective exponent converges to a value around ~ -2 for much of the frequency range between 10^{-2} and 10^{-1} . In contrast we see at most very short-lived power law behavior in (b). The effective exponent perhaps displays a short-lived plateau around an exponent of approximately ~ -1 , but this behavior ceases very quickly.

Figure 8 shows the both the initial distribution of outgoing degrees of connectivity for two networks with (a) $p_{\text{inh}} = 0.30$, and (b) $p_{\text{inh}} = 0.04$, as well as their respective distributions after 45,000 avalanches. The horizontal axis represents the degrees of outgoing connectivity k_{out} of individual neurons, while the vertical axis displays the number of observed neurons with a particular k_{out} . The (gray) squares indicate the initial distributions, while the (blue) circles represent the number of neurons measured with all degrees of connectivity except zero, and the (red) square represents the measured number of neurons with zero outgoing connections. The (red) square is placed with some positive offset on the horizontal axis in order to display this data on a double-logarithmic plot. The (orange) line represents the best fit of a power law to the distributions measured after 45,000 avalanches, which follows an exponent very close to -2.0 in both situations. The inset is a log-linear graph of the same data, replotted to highlight the maximum degree of connectivity in the network, $k_{\text{out}} = 75$ for $p_{\text{inh}} = 0.04$ (a), and $k_{\text{out}} = 91$ for $p_{\text{inh}} = 0.04$ (b).

The power law behavior constructed in the initial distributions for both (a) and (b) remains relatively unchanged after the networks' evolution over 45,000 avalanches, but there are noticeable differences not only between the evolved distributions and their initial states, but also between the two evolved distributions. Notably, the height of both evolved distributions is lower than their respective initial configurations, with the distribu-

tion taken from networks with higher inhibitory fractions of $p_{\text{inh}} = 0.30$ (a) being much lower than the evolved distribution taken from networks with the lower inhibitory fraction of $p_{\text{inh}} = 0.04$ (b). Additionally, there are striking differences in the head and tail of the distribution after the Hebbian rules have been applied for 45,000 avalanches. Many of the high degrees of connectivity have been pruned out of the distribution, and many more neurons retain either only one or zero outgoing connections; additionally, no neurons have more than 75 and 91 outgoing connections, respectively, for these two inhibitory fractions. The capability of the model to prune away all connections of a neuron is a noticeable departure from biological relevancy, and we find that the Hebbian rules detailed in Sec. II B have a tendency to “over-prune” due to the inability for new connections to be formed, or for old zero-strength “dead” connections to be reestablished. Future iterations of this model could be improved by allowing new or dead connections to become established between neurons through some probabilistic rate during the network updates that take place after each avalanche.

Even with the general decrease in connectivity, the lower inhibitory fraction network is inherently more strongly connected than the higher inhibitory fraction network after a long simulation time allowing for 45,000 avalanches. The neural network with smaller p_{inh} displays a higher maximum value of k_{out} , and the number of neurons with only one or zero outgoing connections is

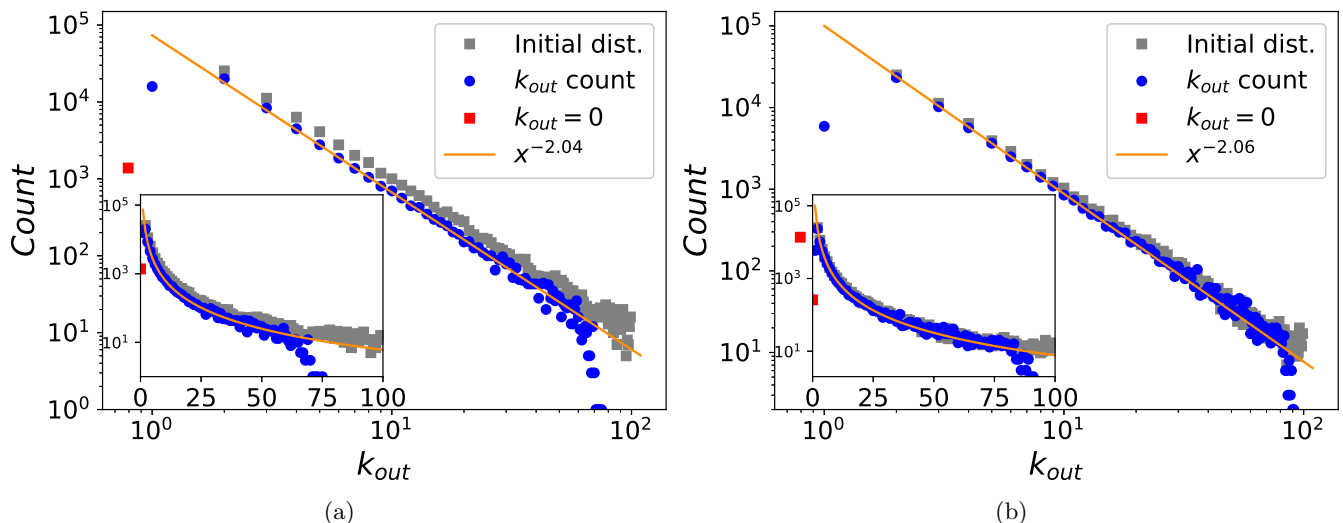


FIG. 8: The k_{out} degree distribution for networks with an inhibitory fraction of (a) $p_{inh} = 0.30$ and (b) $p_{inh} = 0.04$. The horizontal axis (k_{out}) represents the number of outgoing connections individual neurons have. The vertical axis represents the number of neurons that have a particular value of k_{out} . The (gray) squares in (a) and (b) represent the respective initial distributions, while the (blue) circles are the data points we measured and averaged from 100 networks with (a) $p_{inh} = 0.30$, and (b) $p_{inh} = 0.04$ after 45,000 avalanches, and represent the number of neurons with each degree of connectivity except zero. The (red) squares represents the data points for the number of neurons with zero outgoing connections. These data points are plotted with some offset along the horizontal axis in order to display it on a double-logarithmic plot. The (orange) line is the best fit to the power law regime of the evolved connectivity distribution, with exponents -2.04 (a) and -2.06 (b). The inset represents a log-linear graph of the same data, replotted to highlight the maximum value of k_{out} for this network, which is 75 (a) and 91 (b), respectively.

reduced by almost an order of magnitude as compared with the network with larger p_{inh} .

This difference in connectivity evolution results from the distinct inhibitory fractions of the networks as follows: A stronger inhibitory network will result in weaker signals being transmitted through it. Weaker signals in turn cause weaker connections between neurons, and hence an increase in the number of neuron connections pruned due to the Hebbian rules that govern the system. In comparison, the network with a lower inhibitory fraction cannot suppress the signals in the network as strongly as the network with a higher p_{inh} . The network with a lower inhibitory fraction will sustain stronger connections and will consequently not prune away neuron links as drastically. In general we observe that the pruning mechanism preserves the power law structure of the initial distribution, and the effects of the pruning are only visible in the very head and tails of the distribution.

E. Control of avalanche distributions

Figure 3 shows how the avalanche size distribution changes as the inhibitory fraction of the network is varied. The plot shows two regimes of activity: At higher inhibitory fractions the distribution follows a power law behavior terminating in an exponential cut-off; at very

low inhibitory fractions the algebraic decay extends further and the exponential cut-off is shifted to very high avalanche sizes. These extended power laws dominate the dynamics of the networks they occur in. In the following subsection we draw inspiration from work previously done on the robustness of scale-free networks [24, 25] and propose two different control strategies to remove these extended power law tails in networks with low inhibitory fractions through the disabling of either (1) randomly picked or (2) specifically selected highly connected excitatory neurons.

1. Control through disabling random excitatory neurons

The first strategy we implemented is disabling randomly chosen excitatory neurons. Disabling a neuron means that the potential of the neuron is permanently set to zero from the time step it is disabled. We choose excitatory neurons only because the goal is to prevent the large avalanches occurring in the extended tail of this network's normal avalanche size distribution. An active inhibitory neuron is more effective at stopping these avalanches than a disabled neuron, so we only pick excitatory neurons. The neurons are selected with equal probability from all excitatory neurons. We tested several different fractions of excitatory neurons to disable

randomly, and only observed the extended tail present in the avalanche size distribution to disappear for disabling fractions greater or equal to 0.30.

Figure 9(a) shows the avalanche size distribution of a network with inhibitory fraction of $p_{\text{inh}} = 0.04$ after 30% of the excitatory neurons were randomly selected and disabled, i.e. $n_i(t)$ was held at zero for all disabled neurons in the subsequent evolution. The network was allowed to evolve unperturbed for 60,000 avalanches before the excitatory neurons were disabled. The data shown was averaged over 100 realizations of the network. The extended tail seen in avalanche size distributions with this inhibitory neuron fraction has disappeared, due to the fragmentation of the network caused by disabling so many neurons. The network is no longer able to sustain the large network-wide avalanches necessary for the extended power law tail in the avalanche size distribution. Disabling fractions of random excitatory neurons less than 30% does not remove the extended tail. The system is thus quite robust against random disabling, which is not surprising given its scale-free structure. Scale-free networks are known to be quite stable against the random removal of nodes, which is analogous to our disabling of neurons. In order to prevent a signal from being propagated across a generic scale-free network of the same size as our network, more than 90% of the nodes must be randomly removed [24, 25]. We only need to disable a much lower fraction of random neurons to see the extended power law disappear because we are not attempting to disrupt the entire network activity, but only curtail the power law regime of these extended avalanches.

However, even with the network’s robustness, disabling this many neurons from the network is destructive to the normal dynamics of the network, and we observe a change in exponent of the power law behavior of the distribution from -1.5 to -1.74 . This lower exponent results in a much lower probability of strong avalanches, and we see the cut-off appear several decades below the threshold in networks with higher inhibitory fractions.

Disabling random neurons consequently is an ultimately successful strategy for removing the long power law tail in this avalanche distribution; however, it requires a significant portion of the network’s neurons be disabled, which is certainly not ideal, and very likely quite detrimental in any biological neural network.

2. Control through disabling highly connected neurons

The second strategy we implemented is the disabling of the most highly connected excitatory neurons. The neurons were chosen based on their degree of outgoing connectivity k_{out} . We tested several different fractions of the most connected excitatory neurons, and observed that only the top 1% of highly connected neurons need to be disabled in order to stop these incredibly large avalanches.

Figure 9(b) depicts the avalanche size of a network with

inhibitory fraction of $p_{\text{inh}} = 0.04$ with the top 1% of the most highly connected neurons disabled. The network was again allowed to evolve for 60,000 avalanches before the highly connected neurons were disabled. The (blue) circles are the averaged data points of the avalanche size distributions from 100 different network realizations. The (orange) line is the best fit of a power law to the data. This model is very sensitive to disabling highly connected neurons, which is also to be expected given its scale-free structure. In addition to being very robust against random disabling, scale-free networks are highly susceptible to “targeted” disabling, where the most highly connected nodes are removed [24, 25]. A generic scale-free network of the same size and exponent of the connectivity distribution needs approximately the top 3% of highly connected nodes disabled to completely fragment the network and destroy any long-range connectivity [24–26]. In our system, we need only disable the top 1% of the highly connected neurons because we do not aim to completely destroy the network dynamics, yet merely wish prevent the occurrence of exceedingly strong avalanches.

Permanently disabling these highly connected neurons does still considerably affect the network dynamics. The avalanche size distribution follows a power law with exponent -1.66 instead of the typical -1.5 . This is due to the signals being unable to propagate as strongly through the network after the highly connected neurons have been disabled. These weaker avalanches are more likely to die out earlier than their counterparts in an unsuppressed network, resulting in a lower exponent and an earlier cut-off in the avalanche size distribution.

Disabling only the highly connected neurons is hence demonstrably a very successful control strategy for removing these extended power law tails that dominate the network. However, this approach does require significant knowledge about the structure of the network and is still destructive to the network dynamics because, by definition, these highly connected neurons are very important to signal propagation through the system.

3. The disabled fraction of connections

We can also compare the fraction of connections between neurons that are effectively removed from our networks due to each of these two strategies, in addition to the fraction of neurons. We observed the process of randomly selecting a certain fraction of neurons in our networks to select the same fraction of the connections. Thus when 30% of the excitatory neurons are disabled we also disable approximately 30% of the network’s connections. In contrast, when we target the highly connected neurons we sample the tail of our connectivity distribution, and we observe that we disable approximately 10% of the network’s connections even though we disable only 1% of the network’s neurons.

The discrepancy in the fraction of disabled connections between these two strategies is due to the power




FIG. 9: The avalanche size distribution of a network with an inhibitory fraction of $p_{\text{inh}} = 0.04$ after (a) randomly selected 30% of the excitatory neurons, and (b) the top 1% of highly connected excitatory neurons have been disabled. In either case, the network was allowed to evolve naturally for the duration of 60,000 avalanches before the excitatory neurons were disabled. The (blue) circles represent the probability of an avalanche having a given avalanche size S , and the (orange) line is the best fit of a power law to the data. The extended tail present in a normal $p_{\text{inh}} = 0.04$ network has disappeared after disabling (a) a large fraction, (b) a minor fraction of excitatory neurons. For case (a), disabling 30% of inhibitory of neurons is very destructive to the dynamics of the network, and a marked change in the power law behavior of the avalanche size distribution is observed: This network follows a power law with exponent ~ -1.74 , lower than the normal exponent of -1.5 . This is the result of signals dying out more quickly in this heavily diluted network. Disabling random fractions less than 0.30 of the excitatory neurons retained the power law tails typically seen in the avalanche size distributions of these low inhibitory fraction networks; indeed, 0.30 is the lowest fraction of random excitatory neurons that must be disabled to curtail these extended power law tails. For (b), the extended tail present in a normal $p_{\text{inh}} = 0.04$ network has disappeared after disabling only a minor fraction of the excitatory neurons. Due to the power law structure of the network connectivity, these top 1% of the excitatory neurons are much more important to the network dynamics than the vast majority of all other neurons. Disabling these neurons is still destructive to the network dynamics, though less so than the random disabling case shown in (a), and a change in the power law behavior of the network is observed. This network follows a power law with exponent ~ -1.66 , again lower than the normal exponent of -1.5 . This is the result of signals being unable to propagate through the most heavily connected neurons, which play a large role in the network's transmission capability. 0.01 is the required lowest fraction of highly connected excitatory neurons that needs to be disabled in order to effectively remove these extended power law tails. Data averaged over 100 independent network realizations.

law nature of the network’s connectivity distribution, as well as the self reinforcement of connection strength that is driven by the Hebbian learning rules detailed in Sec. II B. With a sufficiently large network, randomly choosing neurons will access the same fraction of connections as neurons, but due to the power law connectivity distribution most of these neurons have a very small degree of connectivity and are relatively unimportant to the network dynamics. The Hebbian learning rules enforce that connections that are not often used will grow weaker than those that are used frequently, and neurons with low numbers of connections will not use those connections very often. Thus one needs to disable a large number of these weak connections that make up the majority of the network’s connectivity distribution in order to dramatically effect the network dynamics. Conversely, the connections from highly connected neurons will be utilized frequently and will hence remain inherently stronger than the connections to neurons with lower k_{out} . In addition to the highly connected neurons serving as “hubs” that link many otherwise poorly connected parts of the network together, their inherently stronger connections will sustain signals more easily in the network and will be more important to the propagation of extremely large avalanches. Thus only approximately 10% of the total connections need to be removed from the network when taken from the tail of the connectivity distribution through the disabling of the most highly connected neurons.

V. DISCUSSION

Operating our model with the parameters described in Sec. II A and an inhibitory fraction of $p_{\text{inh}} = 0.30$, we observe that the avalanche size distribution of our model follows a power law of $P_S(S) \sim S^{-1.55}$, and that the avalanche duration distribution of our model obeys a power law of $P_D(D) \sim D^{-2.1}$, both of which agree well with experimental results [6, 8] and reproduce the results shown by Lombardi, Herrmann, De Arcangelis et al. As we lower the inhibitory fraction of our network towards an inhibitory fraction of 0.04, we intriguingly find behavior suggestive of criticality as the exponential cut-off present previously in the avalanche size and duration distributions disappears, and these distributions continue to follow power laws for several more decades. At this value of inhibitory fraction, the network becomes dominated by long-lasting avalanches that persist for orders of magnitude more time steps than avalanches in the cut-off distributions. The particular value of inhibitory fraction at which we see this extension of the distributions is far below the fraction found in human cortices, which is closer to 0.2 – 0.3 [15, 16].

Additionally, the power spectral density of our network at low inhibitory fractions ($p_{\text{inh}} = 0.04$) behaves similarly to power spectral densities of epileptic humans by following a power law with exponent -2.0 . As the inhibitory fraction of the network is increased to a more biologically

relevant value of 0.3 the power law behavior in the PSD vanishes, and we see only a brief minima in the effective exponent of the PSD near a value of -1.0 , though this value of exponent matches the observed exponent of healthy human brains [7]. Lombardi, Herrmann, De Arcangelis et al. observed a transition of the exponent from an “epileptic” value of -2.0 to a “healthy” value of -1.0 [15]; however, we only observe the emergence of power law behavior at low inhibitory fractions.

Low inhibitory fractions allow the network to access much higher avalanche durations and corresponding avalanche sizes, because even though the underlying power law distribution of these quantities does not change, the exponential cut-off disappears allowing the power law distributions to extend into regimes of greatly increased duration and size. The incredibly large “black swan events” that the network can access have correspondingly low probabilities due to the power law distribution, but because they are so large, they dominate the network for many orders of magnitude more time steps than avalanches from a “healthy” distribution, once they occur.

The exponential cut-offs protect the network from these events, and human cortices may naturally operate at higher inhibitory fractions in order to avoid a truly critical point, yet still benefit from wide distributions at lower intensity avalanche events.

This corroborates the idea proposed by Milton et al. [23], that critical behavior in the brain, though long sought after, might be destructive as the long-range correlations introduced by approaching a critical point could destroy and dominate the short-range interactions necessary for the proper operation of the brain.

We also observe how the outgoing connectivity distribution changes as the network evolves under the Hebbian learning rules described in Sec. II B after 45,000 avalanches have run through the system.

Networks with a high inhibitory fraction ($p_{\text{inh}} = 0.30$) prune away many connections, as the system is unable to propagate avalanches strong enough to sustain all of the links. This results in the tail of the connectivity distribution being truncated with ultimately no neurons maintaining more than 75 outgoing connections, while the head of the distribution becomes inflated, as many neurons end up having only one or zero outgoing connections. Networks with a lower inhibitory fraction ($p_{\text{inh}} = 0.04$), prune away fewer connections than networks with a higher inhibitory fraction, as they are able to sustain stronger avalanches in the network. This results in an extended connectivity distribution tail, with some neurons having as many as 91 outgoing connections after 45,000 avalanches. Additionally these networks display an order of magnitude fewer neurons with zero or merely one connection than the networks endowed with a higher inhibitory fraction.

The combination of the inhibitory neurons and the Hebbian rules of the system cause networks with high inhibitory fractions to evolve into more sparsely connected

networks than networks with a low inhibitory fraction. These differences in connectivity reinforce the networks' ability to sustain or disrupt very large avalanches. Networks with high inhibitory fractions will display weaker avalanches, causing them to be less connected, which in turn further weakens them in their capability to sustain large-scale avalanches. Networks with lower inhibitory fractions will on occasion go through massive avalanches allowing them to remain more connected, which hence will assist these systems in permitting further strong avalanches.

Finally we investigate two different strategies to remove these exceedingly large avalanches from networks with low inhibitory fraction through either the disabling of randomly selected or carefully chosen highly connected excitatory neurons, respectively. In order to curtail these large events through random disablings, 30% of the networks excitatory neurons must be disabled. This strategy is therefore ultimately effective, but would be quite destructive to any biological neural networks. In contrast, switching off highly connected neurons proves to be a much more effective strategy, as only the top 1% of these prominently connected excitatory neurons need to be disabled in order to prevent such large-scale avalanche events. Both of these strategies provide a means to circumvent the inherent occurrence of incredibly large "epileptic" avalanches in systems with very low inhibitory neuron fraction.

ACKNOWLEDGMENTS

We thank Nuno Araujo for pointing us towards the work of Lombardi, Herrmann, De Arcangelis, et al. We would like to thank Michel Pleimling and Priyanka for helpful discussions, Alexandra Bosh for her contributions to this work in the course of her undergraduate research project, and Lucilla De Arcangelis, Hans J. Herrmann, and Fabricio Lombardi for valuable comments. Research was sponsored by the Army Research Office and was accomplished under Grant Number W911NF-17-1-0156. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

Appendix A: Extended model to reproduce waiting time distribution

The model described in section II accurately recreates the avalanche size distribution, the avalanche duration distribution, and the power spectrum distribution observed in rat and human cortices [6, 21]. In order to also generate experimentally detected waiting time dis-

tributions, the model must be extended. This extended model is detailed by Lombardi, Herrmann, de Arcangelis et al. [12–14] and is briefly summarized here.

The waiting time between avalanches is the time between the end of the last avalanche and the beginning of the next. Figure 10 shows experimentally determined waiting time of seven different slices of rat cortex. This figure was taken with permission from Lombardi, Herrmann, de Arcangelis et al.'s paper *The balance between excitation and inhibition controls the temporal organization of neuronal avalanches* [12].




FIG. 10: Figure reproduced with permission from Ref. [12]. This plot shows the avalanche waiting time distributions for seven different slices of rat cortex. The inset shows two examples of temporal neural sequences. The majority of the waiting time distributions display bimodal behavior, with an initial power law at low waiting times, followed by a bump at higher waiting times.

The experimentally determined waiting time distributions [12] shown in Fig. 10 display bimodal behavior with an initial power law regime followed by a “bump.”

This bimodal behavior requires that avalanches which occur within short waiting times should be highly correlated to previous avalanches in order to reproduce the power law behavior at low waiting times. In comparison, avalanches with long waiting times need to be uncorrelated to reproduce the exponential behavior seen at longer waiting times.

These distinct features can be reproduced by introducing two network-wide macro-states: “up” and “down.” The up state is defined as a period of high network activity, during which many neurons are close to firing potential. Anytime an avalanche occurs, the system transitions to or remains in the up state.

During an up state, the noise driving the system is drawn from the distribution $(0, S_{\max}/S(\tau)]$, where S_{\max} is the avalanche size threshold, and $S(\tau)$ is the size of the last avalanche. Additionally, when an avalanche ends in the up state, each neuron in the network is reset to be

close to the neuron firing threshold.

$$n_i \rightarrow n_{\max}(1 - S(\tau)/S_{\max}). \quad (\text{A1})$$

The resetting in Eq. (A1) ensures with high probability that avalanches in the up state will have short waiting times, and the correlations introduced from the up state noise distribution produces the power law behavior seen in Fig. 10.

If the size of the last avalanche exceeds the avalanche size threshold S_{\max} , the system transitions to the down state. The down state is characterized as a period of no activity in the network in which the system is slowly brought back to firing. When the network transitions from the up state to the down state, each neuron is dras-

tically polarized in opposition of the previous behavior,

$$n_i \rightarrow n_i - h\Delta n_i, \quad (\text{A2})$$

where Δn_i is the sum of depolarizations during the last avalanche in the up state, and h is a system parameter introduced to control the size at which the neuron is anti-polarized. During the down state, the network is driven by small ($\sim 0.01 \cdot n_{\max}$) constant noise.

The hyperpolarization of neurons coupled with the small constant noise ensures that the waiting times during the down state will be very long, and will produce an approximately Gaussian distribution due to the central limit theorem.

-
- [1] G. L. Fain, *Molecular and Cellular Physiology of Neurons*, 2nd ed. (Harvard University Press Cambridge, 2014) p. 178.
- [2] R. F. Schmidt and G. Thews, *Human Physiology*, 2nd ed. (Springer Berlin Heidelberg, 1989) p. 26.
- [3] J.-L. Gaiarsa, O. Caillard, and Y. Ben-Ari, *Trends Neurosci.* **25**, 564 (2002).
- [4] K. Gerrow and A. Triller, *Curr. Opin. Neurobiol.* **20**, 631 (2010), neuronal and glial cell biology New technologies.
- [5] R. C. Malenka and M. F. Bear, *Neuron* **44**, 5 (2004).
- [6] J. M. Beggs and D. Plenz, *J. Neurosci.* **23**, 11167 (2003), <http://www.jneurosci.org/content/23/35/11167.full.pdf>.
- [7] J. Yan, Y. Wang, G. Ouyang, T. Yu, Y. Li, A. Sik, and X. Li, *Nonlinear Dyn.* **83**, 1909 (2016).
- [8] S. Yu, A. Klaus, H. Yang, and D. Plenz, *PLoS One* **9**, 1 (2014).
- [9] L. De Arcangelis, C. Perrone-Capano, and H. J. Herrmann, *Phys. Rev. Lett.* **96**, 028107 (2006).
- [10] G. L. Pellegrini, L. de Arcangelis, H. J. Herrmann, and C. Perrone-Capano, *Phys. Rev. E* **76**, 016107 (2007).
- [11] L. De Arcangelis and H. Herrmann, *Front. Physiol.* **3**, 62 (2012).
- [12] F. Lombardi, H. J. Herrmann, C. Perrone-Capano, D. Plenz, and L. de Arcangelis, *Phys. Rev. Lett.* **108**, 228703 (2012).
- [13] L. De Arcangelis, *Eur. Phys. J. Spec. Top.* **205**, 243 (2012).
- [14] F. Lombardi and L. De Arcangelis, *Eur. Phys. J. Spec. Top.* **223**, 2119 (2014).
- [15] F. Lombardi, H. J. Herrmann, and L. De Arcangelis, *Chaos* **27**, 047402 (2017), <https://doi.org/10.1063/1.4979043>.
- [16] S. Sahara, Y. Yanagawa, D. D. M. O’Leary, and C. F. Stevens, *J. Neurosci.* **32**, 4755 (2012), <http://www.jneurosci.org/content/32/14/4755.full.pdf>.
- [17] S. J. Cooper, *Neurosci. Biobehav. Rev.* **28**, 851 (2005).
- [18] V. B. Priezzhev, D. V. Ktitarev, and E. V. Ivashkevich, *Phys. Rev. Lett.* **76**, 2093 (1996).
- [19] H. J. Jensen, *Self-Organized Criticality: Emergent Complex Behavior in Physical and Biological Systems*, Cambridge Lecture Notes in Physics (Cambridge University Press, 1998).
- [20] V. M. Eguíluz, D. R. Chialvo, G. A. Cecchi, M. Baliki, and A. V. Apkarian, *Phys. Rev. Lett.* **94**, 018102 (2005), [cond-mat/0309092](https://arxiv.org/abs/cond-mat/0309092).
- [21] M. H. Myers, E. Jolly, Y. Li, A. de Jongh Curry, and H. Parfenova, *Ann. Neurosci.* **24**, 12 (2017).
- [22] P. D. Welch, *IEEE Trans. Audio Electroacoust.* **15**, 70 (1967).
- [23] J. G. Milton, *Eur. J. Neurosci.* **36**, 2156 (2012).
- [24] R. Albert, H. Jeong, and A. Barábasi, *Nature* **406**, 378 (2000).
- [25] R. Albert and A.-L. Barabási, *Rev. Mod. Phys.* **74**, 47 (2002).
- [26] D. S. Callaway, M. E. J. Newman, S. H. Strogatz, and D. J. Watts, *Phys. Rev. Lett.* **85**, 5468 (2000).