MODELING FRICTION FORCE AVALANCHES IN MOLYBDENUM DISULFIDE AND SPIKING NEURONAL AVALANCHES IN MOUSE CORTEX

BY

TYLER SALNERS

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Physics
in the Graduate College of the
University of Illinois Urbana-Champaign, 2023

Urbana, Illinois

Doctoral Committee:

Professor Yann R. Chemla, Chair
Professor Karin A. Dahmen, Director of Research
Professor Michael Stone
Professor Emeritus Richard L. Weaver

Abstract

In nature we find widely varying systems can sometimes be described by a single set of simple laws. For example, intermittent jumps of both sliding tectonic plates (during an earthquake) and spontaneous magnetization (the Barkhausen effect in magnets) follow the same set of "crackling" laws. Crackling refers to the quick jumps between metastable configurations that are facilitated by slow external forcing in a disordered system. The cascades of jumps (referred to as 'avalanches') follow simple scaling laws for quantities like their size and duration. In this work, we highlight two systems that follow scaling laws and compute predictions about the system using avalanches. Building on a simple slip model we show that friction force fluctuations of molybdenum-disulfide may result from slip-avalanches. We show how experimental conditions can affect the observed scaling behavior, and thus provide a blueprint for the interpretation of data in a range of other systems which might have similar experimental conditions. Importantly, we find that building the experimental conditions into the model yields good agreement with the experimental results. Next, we study fluctuations of spiking neurons in mice cortex. They too can be described as avalanches and follow similar statistics as the slips in the friction experiments. Here we use the same basic avalanche model used for friction, and to capture the essential physics of inhibition and recurrent firing, we add mean field connections to an external inhibitory population and a threshold reduction mechanism, respectively. We find that this simple model can mimic many statistical properties of the data. We show that model simulations with a few biological tweaks can even reproduce empirical signatures of recurrent firing. Finally, we show that though the inhibitory species are the minority population, they make up a majority of the activity in neuronal avalanches. We demonstrate the models ability to reproduce this inhibitory dominance and conclude with future directions.

This dissertation is dedicated to my grandparents, Jack & Kathleen Wilkerson, and Glenn & Luanne Salners

Acknowledgments

I would like to acknowledge first and foremost my advisor, Professor Karin Dahmen, for without her unwavering patience and keen insights, this dissertation would not be possible. I would also like to acknowledge my soon to be wife, Katelynn, for constantly being by my side, not only to support me, but to push me to be the best version of myself. I'd also like to thank my family, especially my mother, Tricia and my twin sister, Shelby, for without them I would not have ever made it to graduate school. I especially owe my long time friends, Patrick Kaner, Erik Lemmerman, Jonathon Berry, Kent Williams, Clayton Morris, Randy Harp and Brian Vienne, for they're the reason I ever developed a drive and determination to reach new heights. I would also like to acknowledge all of the members of the Dahmen research group, with special thanks to William McFaul, Li Shu, Alan Long, Siddharth Mansingh and Jordan Sickle for any and all help they provided me along the way. Lastly, I'd especially like to thank the graduate department advisor, Professor Lance Cooper, who made my life, and those of all UIUC physics graduate students, so much easier than they would have been without him.

List of Abbreviations

PDF	Probability Distribution Function
CDF	Cummulative Distribution Function
CCDF	Complement of the Cummulative Distribution Function
MFT	Mean Field Theory
RG	Renormalization Group
MoS_2	Molybdenum Disulfide
STZ	Shear Transformation Zone
MD	Molecular Dynamics
UHV	Ultra High Vacuum
FFM	Friction Force Microscopy
AFM	Atomic Force Microscopy
N_2	Nitrogen Gas
SiO_2	Silicon Dioxide
MEA	Micro Electrode Array
TE	Transfer Entropy
SOC	Self Organized Criticality

List of Symbols

u	position of interfacial element in the direction of applied force
F_{ext}	applied force
η	viscosity/damping coefficient
σ_{int}	2D elastic interaction kernel
f_w	random pinning force
$\mathcal{P}(X)$	probability distribution function of the random variable X
J	interaction strength
N	system size
L	perpendicular depth of the bulk material
v	external velocity
K_{L}	effective spring constant
S	avalanche size
T	avalanche duration
$\mathcal{G}_S$	scaling function for $\mathcal{P}(S)$
$\mathcal{G}_T$	scaling function for $\mathcal{P}(T)$
Γ	force drop rate
Γ_{max}	maximum force drop rate
$\mathcal{H}_S$	critical scaling function for size averaged avalanche profiles
$\mathcal{H}_T$	critical scaling function for duration averaged avalanche profiles
τ	critical scaling exponent for $\mathcal{P}(S)$
α	critical scaling exponent for $\mathcal{P}(T)$
$1/\sigma\nu z$	avalanche fractal dimension
Δu_{i}	local slip distance of cell 'i'
Δt	simulation timestep
τ_{i}	local stress of cell 'i'

$\tau_{\mathrm{arrest,i}}$	local arrest stress of cell 'i'
τ_{thresh}	threshold stress
g	distance to failure of most unstable cell
c	conservation parameter
Θ	total friction force
$\mathcal{C}(X)$	complement of the cummulative distribution function of random variable X
$\mathcal{T}_{\mathrm{i}}$	duration bin center
$\mathcal{S}_{\mathrm{i}}$	size bin center
μ	friction coefficient
R	probe sphere radius
C_{z}	normal force calibration
k_{z}	normal spring constant
V_{z}	photodiode voltage
F_{n}	normal force
V_{xt}	lateral photodiode voltage during trace
V_{xr}	lateral photodiode voltage during retrace
W	friction loop width
Δ	friction loop offset
F_{f}	friction force
S_0	avalanche size rescaling parameter
T_0	avalanche duration rescaling parameter
G_0	avalanche force drop rate rescaling parameter
δt	temporal bin width
t_{start}	avalanche start time
t_{end}	avalanche end time
d	maximal delay parameter
$\%DGB$	percentage of durations longer than the shortest burst duration
V_{i}	potential of neuron 'i'
V_{thresh}	threshold potential
ΔV_{i}	potential released when neuron 'i' spikes
$V_{\mathrm{r,i}}$	reset potential of neuron 'i'
α	strength of inhibition
ϵ	strength of weakening

ΔT	seizure duration/persistence time
S_{cr}	critical avalanche nucleation size
U	number of unique neurons
N_{inh}	number of inhibitory neurons in a sample/simulation
N_{exc}	number of excitatory neurons in a sample/simulation
S_{inh}	number of inhibitory spikes in an avalanche
S_{exc}	number of excitatory spikes in an avalanche
U_{exc}	number of active inhibitory neurons in an avalanche
U_{exc}	number of active excitatory neurons in an avalanche
γ	inhibitory to excitatory synaptic strength reduction factor
ΔJ	width of lognormal distribution used to define synaptic strengths

Table of contents

Chapter 1

Introduction

1.1 Forcing Disordered Systems out of Equilibrium: Avalanches

In this thesis we'll highlight an interesting and ubiquitous phenomena – the 'avalanche' – which is prevalent in many varying systems that are all united by the fact that they are heterogeneous, or 'disordered'. These avalanches, just like the ones most of us are familiar with, are triggered at some origin, and they proliferate through far reaching parts of the system. What 'trigger' means here depends on the system at hand; for instance, we could be talking about soft magnets, whose microscopic dynamics consists of separate spins flipping from the influence of an external magnetic force – here the spin flipping would be the trigger. How one trigger leads to another (via. an 'interaction') is also system dependent. In the example of the magnet this interaction is afforded by the spin-spin coupling (i.e., an Ising-like interaction) that causes a flipping spin to change the local magnetic field of the other spins.

This change in local magnetic field can induce another spin to flip, which then may induce more spins to flip, and so on and so forth. We aptly refer to the cascade of spin flips as an avalanche. Avalanches can be quite mundane in an ordinary system, but if you add heterogeneous forces locally throughout the system (i.e., 'disorder'), the resulting avalanche dynamics becomes much richer. In the magnet analogy, this disorder is achieved by including a random on-site magnetic field influencing the direction of the spin at the site. This can be from, for instance, a local magnetic impurity obscuring the state of a particular spin.

A subtle but essential point is that avalanches are not a property of e.g., a single spin, but instead require the cooperative action of many spins in the system – we say avalanches are a 'collective behavior'. To characterize an avalanche, we can attribute to it properties such as the number of total spin flips (referred to as its 'size'), the total space these spin flips spanned (its 'area') or the total time elapsed during the avalanche (its 'duration'), to name a few. Thus, we are packaging the many different spin flips across the system into one collective mode – the avalanche – and characterizing it with a single set of properties. This way of compartmentalizing a system at different scales has found immeasurable success in physics, particularly in the field of condensed matter where it is often used to understand the behavior of systems on large length and time scales [1]. We'll find that, what's more interesting than the avalanches themselves, is what they imply about the physics that dominates the system at these coarse length and time scales.

The avalanches reflect an underlying (non-equilibrium) critical point that facilitates a sharp increase in correlations among large parts of the system. We say that the correlation length has diverged, and with it, all smaller scales (for instance, the typical area of an avalanche) have washed away. In other words, the long-range correlations lead to an emergent scale invariance, enabling us to make powerful predictions for experimentally accessible measurements.

For example, the scale invariance of the avalanche properties manifests in their probability distributions, which will be 'power-law' distributed. The specific set of exponents for the different power-laws form what's known as the universality class of the critical point. These avalanches will also have a characteristic temporal profile that, when appropriately scaled, will look the same (on average) for avalanches of different sizes and durations. This characteristic profile is what's known as a 'critical scaling function' and is an example of a more stringent prediction of criticality than a mere power-law measurement, since it's showing invariance of an entire functional form and not just of a single number. A further prediction of criticality is the fulfillment of critical scaling relations, relating critical exponents in certain ways dictated by the universality class.

1.2 Finding Avalanches in Nature

In practice, it's not so easy to *directly* measure the number of flipping spins in a magnet during an avalanche. Fortunately, there is a way to *indirectly* measure the number of flipped spins by considering instead the total change in magnetization. After all, any net change in the number of spins pointing up (for instance) is reflected in the magnetization since the magnetization is just the spatial average of the spins. Therefore, cascades of spin flips correspond to jumps in magnetization and one can obtain avalanche sizes by simply measuring the size of these jumps. In similar spirit, many other systems with a clear microscopic definition of an avalanche can be analyzed in terms of a coarse-grained definition.

For instance, a microscopic system describing a discretized rough interface being sheared has many experimental realizations describing deformation in materials, earthquakes in faults etc., (see [2] and references therein). Taking the interface to be, for instance a localized shear band that has formed between the two pieces of a column of material under compression, our microscopic degrees of freedom can be thought of as little cells that tile the interface. Each cell has its own local stress and is coupled to other cells by an elastic interaction such that when sheared, a cell will slip and cause other cells to slip as well. The cascades of slipping cells will be seen in the experiments as the sharp drop in total *stress* which can be measured to obtain the avalanche size and duration etc.

In this thesis, we'll highlight the essence of universality by using one simple model to describe two completely different systems. We'll first explore avalanches in friction force measurements made via., a tribometer sliding over a steel coupon coated in molybdenum disulfide (also referred to as MoS_2, a common solid lubricant). Here, we can imagine the sliding of the tribometer on the sample emulates a situation similar to interface sliding described in the paragraph above. It is natural then, to attribute drops in friction force to avalanches at the frictional interface between the tip of the tribometer and the layer of MoS_2 coating the coupon. After, we'll seemingly switch gears and study the spiking signals from a slice of mouse cortex grown in-*vitro*. There, to extract avalanches, a new definition is necessary. We can think of an avalanche as a burst of spiking activity with a period of quiescence directly before and after. If we imagine temporally coarse graining the spike signals from our data into one signal that represents the number of spikes per short bin of time (i.e., a spiking rate), one can define an avalanche as the sum of any consecutive nonzero segments of such a signal. In other words, once the spiking rate is greater than zero, the avalanche has begun, and it only ends once the rate returns to zero again. Here the size of the avalanche is just the area under this excursion of the spiking rate (equivalent to the total number of spikes during the avalanche). Importantly, with these two definitions out of the way, we find both systems obey predictions of some form of the same simple avalanche model.

1.3 Outline of Thesis

First, in Chapter 2, we will survey the results of model simulations for varying parameter regimes. Then, in Chapter 3 we will use the model to cast friction force fluctuations in MoS2 within the framework of avalanches described above. We will compare to data from both micro tribometer experiments at the macro scale and molecular dynamic simulations at the nano scale. Next, in Chapter 4, we will describe the analysis performed on spiking neuronal data in order to extract avalanches. In Chapter 5 we will discuss how the boundary dissipation in our model is similar to inhibition in neuronal avalanches and show that we can tune it to capture neuronal avalanche statistics decently well with our simple slipping model. To capture neuronal avalanche statistics even better, in Chapter 6 we will upgrade the bare model to not only include excitatory neurons but inhibitory neurons as well. We will discuss how simulations of this amended model capture recurrent spiking activity as seen in experiments. Lastly, in Chapter 7 we will illustrate how the inhibitory neurons dominate the avalanche activity despite being the minority population, and we show how the model reproduces this inhibitory dominance.

Chapter 2

A Survey of the Simple Slip Model

2.1 Modeling Avalanches: Driven Interface Model

Up until now, we have used an analogy with spins and magnets to illustrate avalanches. Moving forward, we will play to the strength of universality and adopt a realization of avalanches in driven elastic interfaces with disorder. This has been described in great detail in [3], where the interface could in general be in any dimension. Here for simplicity of visualization, we discuss a one-dimensional interface moving through a two-dimensional medium; you can think of the contact line of a water droplet sliding down sandpaper [4]. There, since the water droplet is not a rigid body, it will respond to an applied external force with differing advancements of the interface at different positions due to the heterogeneity, or 'disorder', from e.g., the varying heights of the grains of sand on the paper. Since the water droplet also experiences elastic forces (i.e., surface tension) this is an example of a driven elastic interface with disorder.

With these ingredients, an equation of motion can be written down for the position of the interface, $u(\mathbf{r}, t)$, as measured along the direction of an applied external force, F_{ext}:

$$\frac{\eta\, \partial u(\mathbf{r}, t))}{\partial t} = F_{\text{ext}} + \sigma_{\text{int}}(\mathbf{r}, t) - f_w[u(\mathbf{r}, t), \mathbf{r}, \{u(\mathbf{r}, t' < t)\}]. \tag{2.1}$$

Here, η is a constant describing the viscosity or damping, $\sigma_{\text{int}}(\mathbf{r}, t)$ represents the elastic interaction between cells in the interface, and f_w represents the quenched disorder (see Figure 2.1(a)). When the external force causes a cell to slip, its local force will relax while all other cells experience an increase in force from the elastic interactions. F_{ext} serves as a tuning parameter between two characteristic behaviors of the system (Figure 2.1(b)); With F_{ext} large, the interface will slide with nonzero average interface velocity, $\langle du/dt \rangle$. Conversely, with F_{ext} exceedingly small, the heterogeneities keep the interface completely stuck (or 'pinned').

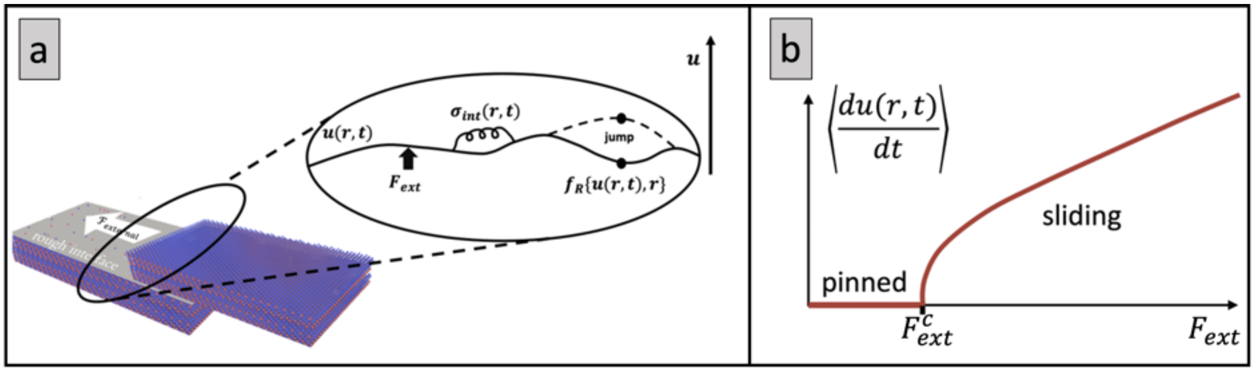

Figure. 2.1: Driven disordered systems and continuous non-equilibrium phase transitions. a) A schematic of a driven interface, taking the example of a two-dimensional interface made by two pieces of materials being sheared against one another. A caricature picture of the interface is shown, for simplicity, as only one-dimensional. $u(\mathbf{r}, t)$ measures the displacement of the interface in the direction of the applied force, F_{ext}. The interface also experiences forces from 1) random heterogeneities on the interface (f_w) and 2) elastic forces from the interactions with the other cells in the interface (σ_{int}). **b)** A non-equilibrium phase transition is illustrated with a plot of average interfacial velocity ($\langle du(\mathbf{r}, t)/dt \rangle$) vs. applied external force (F_{ext}). The interface is 'pinned' when the average interfacial velocity is zero and will transition to 'sliding' in a continuous fashion as the force is increased through the critical point, $F_{\text{ext}}^{\text{c}}$.

Right as the system crosses through the critical force (or 'critical point', $F_{\text{ext}}^{\text{c}}$) it transitions from pinned to sliding in a remarkably interesting way. The abrupt change in the systems properties as the force is moved through the critical point can be understood with the framework of phase transitions.

Near the critical point, with $F_{\text{ext}} < F_{\text{ext}}^{\text{c}}$, the interface is neither stuck nor sliding, but instead moving forward in quick intermittent jumps. By quick we mean with respect to the timescale of the external drive. These jumps are interesting because they aren't characterized by a particular size, but rather, by a broad range of sizes instead. Specifically, the probability distribution (PDF) of sizes takes the form of a 'power-law.' A power-law describes a function that depends on a power of its argument (i.e., $f(x) \approx x^{-a}$) and can be plotted on double-logarithmic axes as a straight line with slope $-a$. This scaling gives us a way to categorize the systems behavior via. the fingerprint made from the set of power-law exponents and the relationships that exist between them.

Now, we can convince ourselves that due to the long-range correlations and corresponding scale invariance, there must be certain microscopic details that are irrelevant to these macroscopic power-laws. We would like to know what details *do* matter for the avalanche scaling and if we can make useful predictions using them. The first detail that concerns us is the explicit form of the elastic interaction term, $\sigma_{\text{int}}(\mathbf{r}, t)$. We will assume the range of the interactions is sufficiently large so that we can replace the spatially dependent kernel

$\sigma_{\text{int}}(\mathbf{r},t)$ with a constant

$$\sigma_{\text{int}}(\mathbf{r},t) \equiv \int_{-\infty}^{t} dt' \int d^2\mathbf{r}' \frac{J(\mathbf{r}-\mathbf{r}', t-t')}{N}[u(\mathbf{r}',t') - u(\mathbf{r},t)] \rightarrow \int d^2\mathbf{r}' \frac{J}{N}[u(\mathbf{r}',t') - u(\mathbf{r},t)], \tag{2.2}$$

an approximation known as 'mean-field theory' (MFT). In this last step we have also assumed the interactions are Dirac delta-correlated in time i.e., the neurons at time step $t + \Delta t$ only feel the effects of spiked emitted at time t.

We also need to consider the explicit form of the external force, F_{ext}. We will drive the interface with a spring from a perpendicular distance L away from the interface, emulating shear in two pieces of material with bulk width L causing slip at an interface. The spring is pulled at velocity v i.e., $F_{\text{ext}} = K_{\text{L}}(vt - u(\mathbf{r},t))$, where $u(\mathbf{r},t)$ again represents the interface displacement and K_{L} is an effective spring constant. For simplicity we will also assume the force is slow enough so as to not increase during the course of an avalanche. We say that v drives the system 'adiabatically slowly', which just means that we drive the system slow enough so that the avalanches are well separated in time.

Under these conditions, predictions for the avalanche statistics can be computed analytically (see [2, 5] and references therein) using the Renormalization Group (RG). In the table below we catalogue the predictions that will be highlighted in this thesis (see [6] for a more detailed version of Table 2.1).

Table 2.1: Analytical predictions of the mean field slip model. $\mathcal{G}_S$, $\mathcal{G}_T$, $\mathcal{H}_S$ and $\mathcal{H}_T$ are universal scaling functions (see Eq. (2.10) and Eq. (2.11)) and k sets the distance to criticality. Adapted from [6]

Quantity	**Scaling Form**	**MFT Prediction**
Avalanche Size Distribution	$\mathcal{P}(S) \sim S^{-\tau}\, \mathcal{G}_S(\sqrt{k}S)$	$\tau = 3/2$
Avalanche Duration Distribution	$\mathcal{P}(T) \sim T^{-\alpha}\, \mathcal{G}_T(kT)$	$\alpha = 2$
Average Size vs. Duration	$\langle S \vert T\rangle \sim T^{1/\sigma\nu z}$	$\sigma\nu z = 1/2$
Average Maximum Velocity vs. Size	$\langle \Gamma_{\max} \vert S\rangle \sim S^{\sigma\rho}$	$\sigma\rho = 1/2$
Average Maximum Velocity vs. Duration	$\langle \Gamma_{\max} \vert T\rangle \sim T^{\rho/\nu z}$	$\rho/\nu z = 1$
Size Averaged Avalanche Temporal Profile	$\langle \Gamma(t) \vert S\rangle \sim S^{1-\sigma\nu z}\, \mathcal{H}_S(t/S^{\sigma\nu z})$	$1 - \sigma\nu z = 1/2$
Duration Averaged Avalanche Temporal Profile	$\langle \Gamma(t) \vert T\rangle \sim T^{1/\sigma\nu z - 1}\, \mathcal{H}_T(t/T)$	$1/\sigma\nu z - 1 = 1$

2.2 Simulations

In this thesis, we'll use these analytical predictions to help understand avalanches arising from several types of systems. We inevitably run into situations that do not strictly adhere to the conditions laid out by the model outlined above; in these cases we can turn to simulations to make further progress. Starting with the ground truth (the simple slip model in the mean field limit), we can simulate avalanches to confirm they agree with the analytical prediction above. In later chapters, we'll make amendments to the simulation to more appropriately model situations not described within the constraints of the analytical prediction outlined above.

2.2.1 Outline of Simulation

To simulate the model, we discretize the interface equation of motion Eq. (2.1) in space and time ($\mathbf{r} \to i, \mathbf{r}' \to j$):

$$\Delta u_{\mathrm{i}} = \frac{K_{\mathrm{L}}}{\eta}(vt - u_{\mathrm{i,t}}) + \frac{J}{\eta N}\Sigma_{\mathrm{j=1}}^{N}(u_{\mathrm{i,t}} - u_{\mathrm{j,t}}) - \frac{f_w[u_{\mathrm{i}}, i, \{u_{\mathrm{i,t'<t}}\}]}{\eta} \tag{2.3}$$

where we have defined the local change in position after one time step $\Delta u_{\mathrm{i}} \equiv u_{\mathrm{i,t+\Delta t}} - u_{\mathrm{i,t}}$. Eq. (2.3) provides us with a map of the cell positions from time t to time $t + \Delta t$. We are, however, going to be more interested in following the local *stress*, τ_{i}, as the system evolves in time. We can write down an equation that represents the local stress while no cells are slipping (i.e., when there is not an ongoing avalanche) as [7]

$$\tau_{\mathrm{i}} = K(vt - u_{\mathrm{i,t}}) + \frac{J}{N}\Sigma_{\mathrm{j=1}}^{N}(u_{\mathrm{i,t}} - u_{\mathrm{j,t}}) \tag{2.4}$$

Notice here that the disorder term is absent. This is because we have translated the random pinning-sites and their effect on local cell positions into a picture of cells slipping with random *arrest stresses* instead. We prescribe that cells slip when their local stress τ_{i} exceeds a threshold $\tau_{\mathrm{thresh}} \equiv 1$, relieving their stress to the rest of the system through the elastic interactions, and settling at an arrest stress, $\tau_{\mathrm{arrest,i}}$. This arrest stress could in principle be set to zero for all cells, but here we see that we can incorporate the disorder of the pinning sites by simply adding some disorder to the arrest stresses. Unless specified otherwise, in this thesis, we will focus on simulations where these arrest stresses are sampled from a uniform distribution centered around zero with width w. This disorder in the arrest stresses is held fixed for the duration of the simulations (usually referred to as 'quenched' disorder). We reiterate that the stress release ($\Delta\tau_{\mathrm{i}} \equiv \tau_{\mathrm{thresh}} - \tau_{\mathrm{arrest,i}}$) for a single site is fixed in *time* and thus disorder is introduced only because $\Delta\tau_{\mathrm{i}}$ is varying in *space*, i.e., for each cell of the interface.

Imposing the external velocity to be adiabatically slow means we only drive the system to initiate the avalanche by forcing the most unstable cell to slip. During the avalanche, in this adiabatic limit, Eq. (2.4) simplifies.

$$\tau_{\mathrm{i}} = -(K_{\mathrm{L}} + J)u_{\mathrm{i,t}} + \frac{J}{N}\Sigma_{\mathrm{j=1}}^{N}u_{\mathrm{j,t}}. \tag{2.5}$$

With these definitions in place, our simulation can be defined quite simply. First, we initialize our system, which is represented by a vector, $\vec{\tau}$, whose N components represent the local stresses of each cell in the interface. We first sample $\tau_{\mathrm{arrest,i}}$ from a uniform distribution centered on zero with width w. Then we

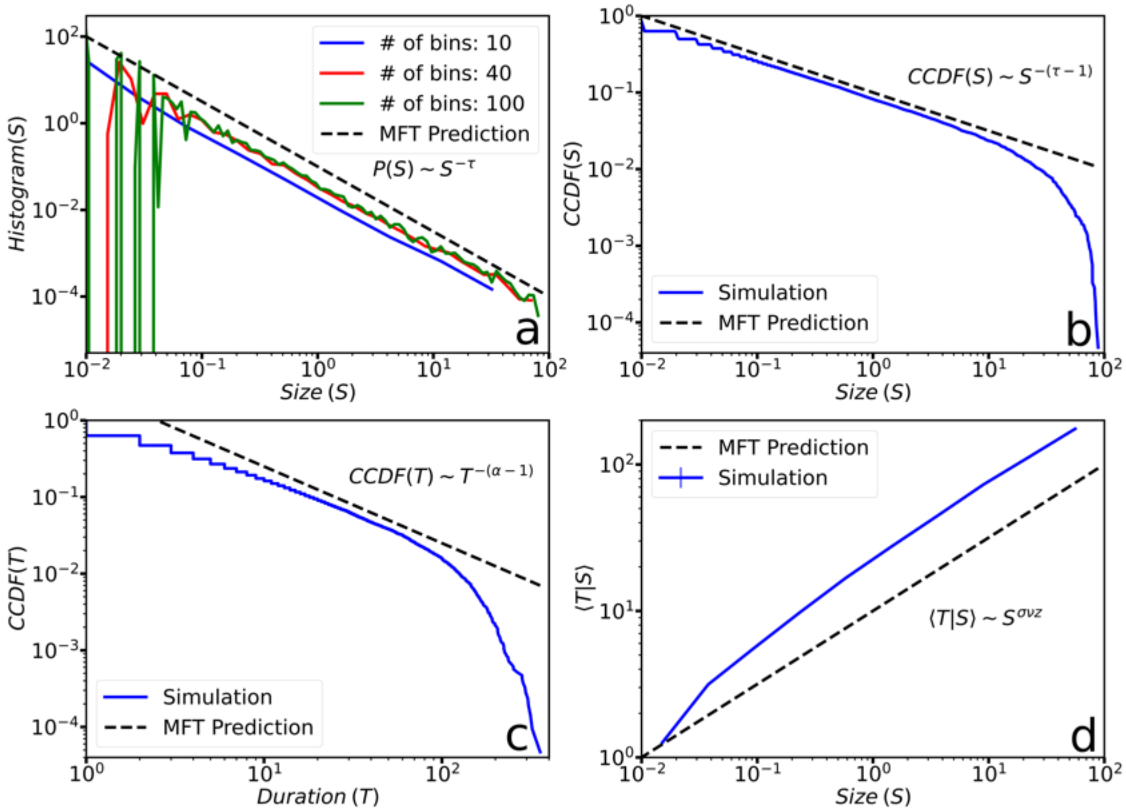

Figure. 2.2: Comparing avalanche statistics: PDF **vs.** $CCDF$**.** **a)** A histogram of avalanche sizes, constructed using 10 (blue) 40 (red) and 100 (green) logarithmically spaced bins. The MFT prediction is shown as a black dashed line. **b)** The complement of the cumulative distribution function ($CCDF$) is shown for simulated avalanche sizes, with a black dashed line showing the MFT prediction. **c)** The $CCDF$ for simulated avalanche durations is shown with a black dashed line plotted to show the MFT prediction. The $CCDF$ of size and duration have power-law slopes of one less than that predicted for the PDF (i.e., $\tau \rightarrow \tau - 1$). **d)** The duration of the avalanche, averaged on size, is plotted for the simulations with the MFT prediction shown as a black dashed line. Ten logarithmically spaced bins were used for the averaging, with error bars showing a 95% confidence interval of each bin. Simulations were run for $N = 10000$, with uniformly distributed arrest stresses with width $w = 0.1$, conservation parameter $c = 1 - 1/\sqrt{N}$, for 200000 time steps.

initialize the stress of each cell to a different random number, $\tau_{\mathrm{i}} = \tau_{\mathrm{arrest,i}} + \xi(\tau_{\mathrm{thresh}} - \tau_{\mathrm{arrest,i}})$ where ξ is uniformly distributed on $[0, 1]$. This just means that the stresses are uniformly distributed between the arrest and threshold stress. To begin the first avalanche, we find the cell closest to threshold, needing an amount of stress g to slip. We raise the stress of the entire system by this amount g, causing the unstable cell to slip, reducing its stress, and redistributing it to the rest of the cells.

Eq. (2.4) tells us precisely how much stress will be lost by a cell when it slips

$$\Delta\tau_{\mathrm{slipping,i}} = -(K_{\mathrm{L}} + J)\Delta u_{\mathrm{i}} \tag{2.6}$$

as well as how much of that stress will be received by all other cells, $'j'$

$$\Delta\tau_{\mathrm{incoming,j}} = \frac{J}{N}\Delta u_{\mathrm{i}} \tag{2.7}$$

Using Eq. (2.6) we can write Eq. (2.7) in terms of the stresses only

$$\Delta\tau_{\mathrm{incoming,j}} = -\frac{c\Delta\tau_{\mathrm{j}}}{N} \tag{2.8}$$

where we have defined the 'conservation parameter' $c \equiv \frac{J}{J+K_{\mathrm{L}}}$, which determines how much of the stress is dissipated to the boundary spring when a cell slips; $c = 1$ means all of the stress is conserved, leading to its name. We simulate the system by tracking the stresses for all times. At each time step, we subtract stress from any cells that slipped, and we take that subtracted stress and redistribute it to the system. We will typically set the conservation parameter to $c = 1 - \frac{1}{\sqrt{N}}$ so that enough stress is dissipated that avalanches do not go on forever. When no cells are beyond the threshold, the avalanche terminates. With these definitions in place, we outline the simulation results below.

2.2.2 Simulation Results

From the simulation we track the time trace of total system stress $\Theta_{\mathrm{t}} \equiv \Sigma_{\mathrm{i=1}}^{N}\tau_{\mathrm{i}}$, from which we can extract avalanches by measuring drops in stress. The total drop in stress is the avalanche size, the total time spent dropping is the avalanche duration. We will also consider the time derivative of this total stress time trace, which we will define by capital gamma $\Gamma_{\mathrm{t}} \equiv \frac{d\Theta}{dt}$. This time trace Γ_{t} will start at zero at the beginning of an avalanche (t_{start}), making an excursion to negative values during the avalanche, and then crossing over zero again to end the avalanche (t_{end}). We will call this temporal profile of Γ_{t} during $t_{\mathrm{start}} \leq t \leq t_{\mathrm{end}}$ the avalanche shape, and we will be interested in the average profile of these shapes for given sizes or durations, as well as the statistics of the peak values of these shapes, Γ_{max}. Technically, considering stress *drops*, our avalanche sizes would be negative, so we will multiply Γ_{t} by -1 to instead consider stress *jumps*, which have a positive size. Γ_{t} gets flipped about the time axis, accordingly, meaning the shapes are always *above* zero now.

Statistics

As mentioned in Chapter 1, the model predicts power-law distributions for the avalanche size, $\mathcal{P}(S) \sim S^{-\tau}$, and duration, $\mathcal{P}(T) \sim T^{-\alpha}$. By simulating many avalanches, we can estimate the PDF by constructing a normalized histogram (see Appendix A). We use these histograms to verify analytical predictions for avalanches. We show in Figure 2.2(a) the normalized histogram of avalanche size as simulated by our model

(N=10000). We can see that it follows a power-law with the same slope as predicted by MFT, $\tau = 3/2$ (see Table 2.1). Because the procedure for constructing the histogram requires you to pick a bin width, there are technically several different normalized histograms for a given data set. We illustrate in Figure 2.2(a) how the normalized histogram changes as you adjust the bin width.

Though these adjustments don't seem to make much of a difference for the simulation data, for experimental data, we find the distributions can be highly sensitive to this bin width. To avoid future discrepancies, moving forward we opt instead to use the complement of the cumulative distribution function ($CCDF$) of avalanche properties. The $\mathcal{P}(S)$ tells you the probability to have an avalanche of size S, while the $CCDF(S)$ ($\mathcal{C}(S)$) tells you the probability of having an avalanche size *greater than or equal to* S. The $CCDF$ is $1 - CDF$ where CDF stands for cumulative distribution function. Since the CDF is just the integral of the PDF, our prediction for the distribution of avalanche size, for instance, $\mathcal{P}(S) \sim S^{-\tau}$ translates to a prediction for the $CCDF$, $\mathcal{C}(S) \sim S^{-(\tau-1)}$, i.e., the same function, just with an exponent of $\tau - 1$, instead of τ. The $CCDF$ is much easier to construct since it does not require to define a bin width. This also means there is one, and only one, $CCDF$ for a given data set (see Appendix A). In light of these details, we will use the $CCDF$ for the remainder of this thesis. Indeed, the simulated $CCDF$ of avalanche sizes is shown to follow a power-law with exponent $\tau - 1$ in Figure 2.2(b), with rather indisputable comparison.

Similarly, the $CCDF$ of avalanche duration, T, is shown in Figure 2.2(c), where it clearly follows a power-law with exponent $-(\alpha - 1)$, $\mathcal{C}(T) \sim T^{-(\alpha-1)}$. The average avalanche duration (see Appendix A) for a given size is also predicted to follow a power-law $\langle T|S\rangle \sim S^{\sigma\nu z}$ with the MFT prediction $\sigma\nu z = 1/2$ (see Figure 2.2(d)). The three scaling exponents mentioned above are related to each other at the critical point by the formula

$$\frac{1-\alpha}{1-\tau} = \frac{1}{\sigma\nu z} \tag{2.9}$$

known as a scaling relation (see [8] for details on how one arrives at this relation).

Dynamics

Apart from the predictions outlined above, we can also make predictions about the entire temporal shape of the avalanches. By temporal shape, we mean the time trace of the stress-drop rate, $\Gamma(t)$. Figure 2.3(a) shows a sample of $\Gamma(t)$ from the simulations with the stress trace, $\Theta(t)$, in the inset. Though the model cannot predict the jagged shape of an individual avalanche, it can predict the maximum stress-drop rate ($\Gamma_{\max}$, see Figure 2.3(a)) averaged on size $\langle\Gamma_{\max}|S\rangle$, see Figure 2.3(b). The model also predicts what an *average* temporal shape will look like. Specifically, it predicts that the average shape will be self-similar– average avalanche shapes at differing sizes and durations will look the same under a simple rescaling. The

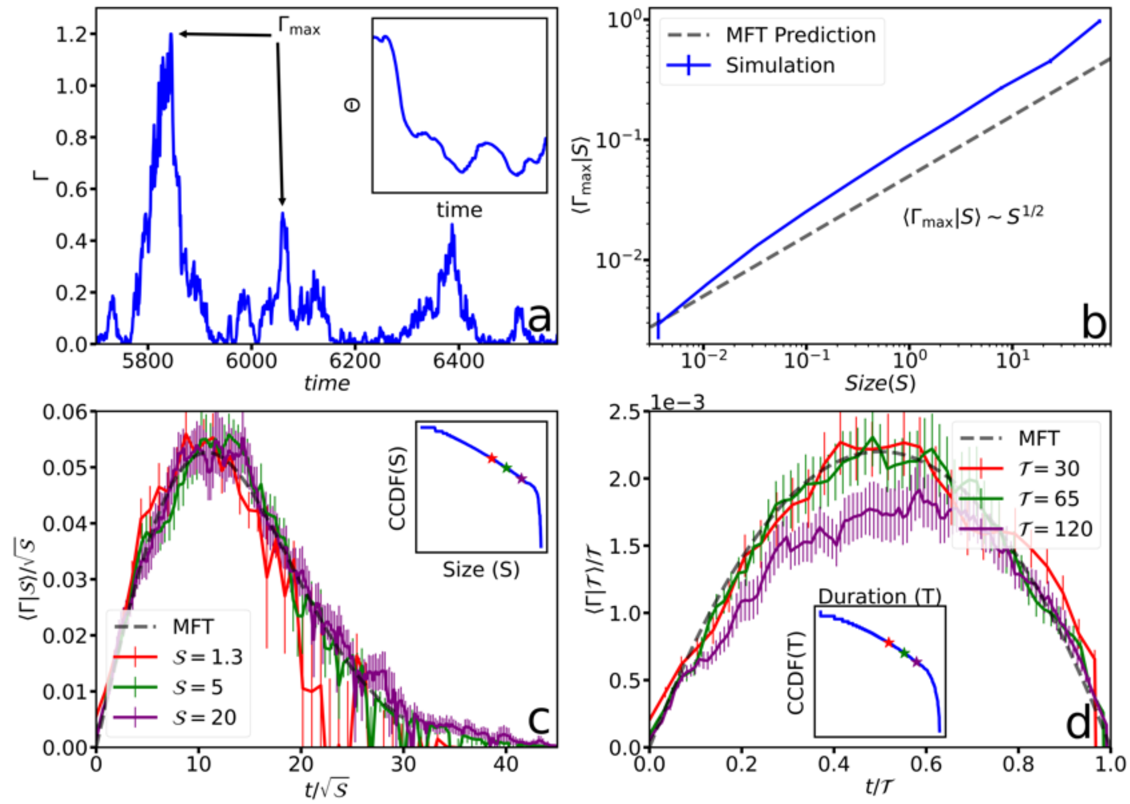

Figure. 2.3: Comparing avalanche statistics: ***PDF*** **vs.** ***CCDF*****. a)** (inset) The time-trace of total system stress θ_t for a portion of the simulation, showing (main) the time derivative of total stress Γ_t. **b)** The simulated maximum value of Γ_t for a given avalanche, or $\Gamma_{\max}$, averaged on avalanche size, with the MFT prediction shown as a dashed black line. Ten logarithmically spaced bins were used for the averaging, with error bars showing the confidence interval of each bin. **c)** The average avalanche profiles shown for three different sizes in the scaling regime. The x and y axis are rescaled according to the MFT prediction to collapse the average shapes onto one predicted critical scaling function (black dashed line, see Table 2.1). Averages for each bin were taken using 51 avalanches, with error bars showing a 95% confidence interval. In the inset we recreate the *CCDF* of avalanche sizes to highlight the place in the scaling regime where we sampled avalanches from for shape averages in the main figure (marked with colored stars). **d)** The same thing as **c)**, but for bins in avalanche *duration*. Simulations were run for $N = 20000$, with uniformly distributed arrest stresses with width $w = 0.1$, conservation parameter $c = 1 - 1/\sqrt{N}$, for 200000 time steps. See Appendix A for details on how these quantities were extracted and/or computed.

profile that these average avalanche shapes all collapse onto after rescaling is called a critical scaling function, and our model predicts (at least) two of them. Averaging avalanches of similar *sizes* predicts the scaling function (Figure 2.3(c))

$$\frac{\langle \Gamma_{\mathcal{S}_\mathrm{i}}(t_\mathrm{j})\rangle}{\sqrt{\mathcal{S}_\mathrm{i}}} = A(\frac{t_\mathrm{j}}{\sqrt{\mathcal{S}_\mathrm{i}}})e^{-B(\frac{t_\mathrm{j}}{\sqrt{\mathcal{S}_\mathrm{i}}})^2} \equiv \mathcal{H}_S(\frac{t_\mathrm{j}}{\sqrt{\mathcal{S}_\mathrm{i}}}). \tag{2.10}$$

Averaging avalanches of similar *durations* predicts instead (Figure 2.3(d))

$$\frac{\langle \Gamma_{\mathcal{T}_\mathrm{i}}(t_\mathrm{j})\rangle}{\mathcal{T}_i} = 4(\frac{t_\mathrm{j}}{\mathcal{T}_\mathrm{i}})(1 - (\frac{t_\mathrm{j}}{\mathcal{T}_\mathrm{i}})) \equiv \mathcal{H}_T(\frac{t_\mathrm{j}}{\mathcal{T}_\mathrm{i}}). \tag{2.11}$$

See Appendix A for details on how average shape profiles are computed.

Disorder

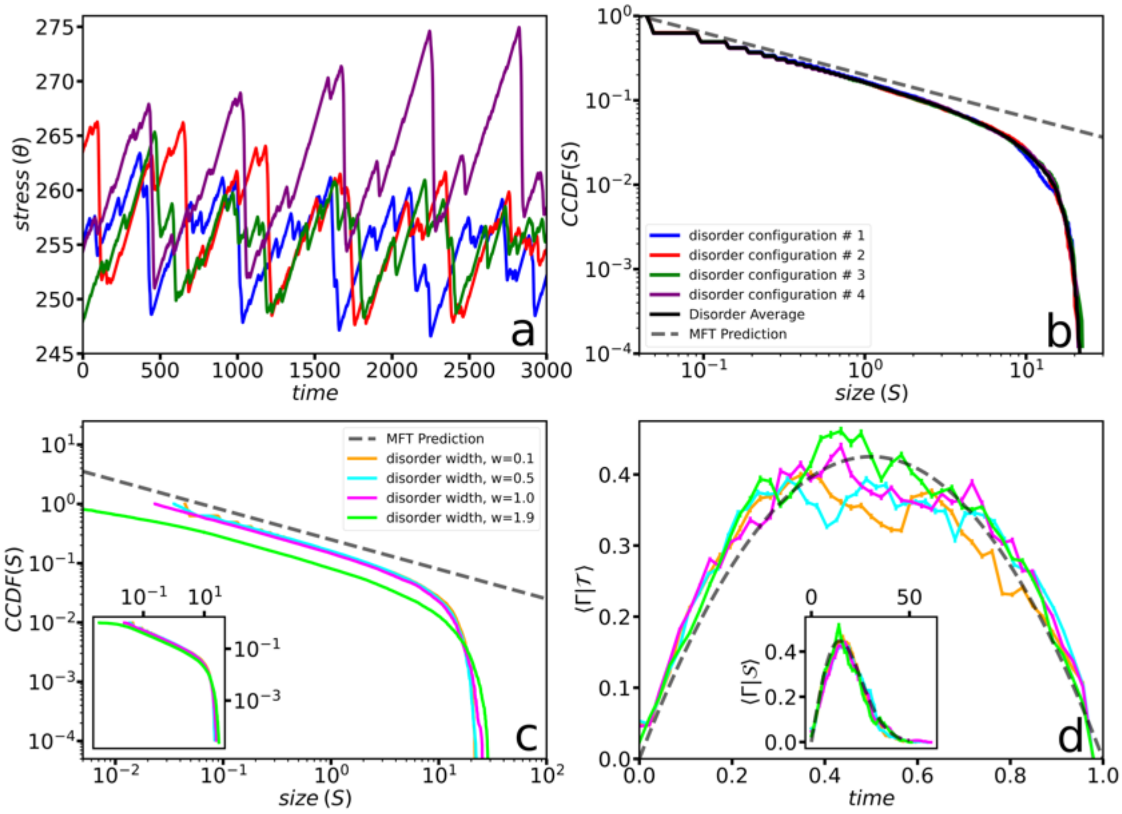

Figure. 2.4: Disorder. a) A portion of the simulated stress trace for four different disorder configurations. **b)** The $CCDF$ of avalanche size, for each of the four simulated disorder configurations. The black line shows the disorder average of the four simulation trials and the dashed grey line shows the MFT prediction. **c)** The $CCDF$ of avalanche sizes for four different disorder strengths. Each line represents a disorder average over four simulation trials. The grey dashed line shows the MFT prediction. We show the same thing in the inset, but for a parabolic distribution of arrest stresses. **d)** The averaged avalanche shape for a given duration, $\mathcal{T}$, shown for the corresponding disorder strengths in (c). The predicted critical scaling function is shown with a dashed grey line. In the inset we show the equivalent averaged shapes for a given *size*, $\mathcal{S}$, with the corresponding MFT predicted scaling function as a grey dashed line. Error bars show the standard error of the mean. Simulations were run for $N = 500$, with conservation parameter $c = 1 - 1/\sqrt{N}$, for 200000 time steps. See Appendix A for details on how these quantities were extracted and/or computed.

We introduced disorder to the simulations by using random arrest stresses. That is, each cell has an arrest stress that was chosen *randomly* by sampling a uniform distribution. We gave the uniform distribution a width, w, and centered it on zero. In this section we illustrate how the simulated avalanche properties depend on 1) the particular realization of disorder for a given simulation trial and 2) the width of the distribution, w. We shall find that the statistical and dynamical features highlighted in the next chapter are not significantly dependent on the disorder.

For each particular disorder realization, the system stress evolves uniquely in time. We illustrate this in Figure 2.4(a) where we plot the total stress (θ) as a function of time (we refer to this as the "stress time-trace" or "stress-trace" for short) for four different disorder configurations. Similar features can be seen between all configurations; the stress changing erratically over time but will occasionally build up stress to an unusually large value followed by an unusually large drop (or avalanche in our language). These large drops also seem to happen roughly periodically in time. Despite the similarities, the four different stress-traces are all unique since they take on very different values of stress at any given time. It will turn out that the similarities pointed out between disorder configurations emerge as an invariance (between disorder configurations) of the scaling properties. We show as an example the avalanche size $CCDF$ in Figure 2.4(b). We can see that the slope of the power-law is consistent between all configurations and that the disorder average of all for configurations gives a result that is indeed representative of the configurations themselves. The cutoff for the size distribution, $S_{\max}$, also stays fixed from one configuration to the next.

The disorder configuration does not change the relevant features of the model results and we show below that the disorder strength only has an effect on the avalanche size cutoff, $S_{\max}$. In Figure 2.4(c) we again plot the avalanche size $CCDF$, where each line represents the average of four different disorder configurations. While the slope of the power-law seen is consistent between different disorder strengths, the cutoff of the $CCDF$ does seem to change. Specifically, we see increasing disorder strength increases the avalanche size cutoff. For completeness, we compared results with a parabolic distribution of the arrest stresses instead of the uniform one. We show those results in the inset of Figure 2.4(c), where we see no noticeable differences in behavior; the slope of the $CCDF$ doesn't change but the cutoff increases with increasing disorder strength.

Lastly, in Figure 2.4(d) we plot the simulated avalanche shapes, averaged around a certain duration, $\mathcal{T}$, for the same disorder averaged simulation results of Figure 2.4(c). We show similar results for shapes averaged on size, $\mathcal{S}$, in the inset. One thing to note is the maximum height of the average avalanche shape increases with increasing disorder. This is consistent with the corresponding increase in avalanche size cutoff mentioned above. The rough functional form of the shapes, however, does seem to be consistent between different disorder strengths.

System Size

To close, we turn briefly to the system size dependence of our model simulations, highlighting subtle deviations from the analytical predictions which technically only hold for infinite system size. Recall that we have adopted the total stress drop as our definition of avalanche size. We do this since many experiments of interest in avalanche research are related to some kind of deformation (see [9], and references therein), and thus the stress signal is experimentally accessible; meaning we can measure avalanche sizes by measuring the stress drops. It is important to understand that the spring driving the system is responsible for absorbing any lost stress during an avalanche and thus is directly related to the avalanche size. Specifically, a slipping cell relieves to the spring a fraction $1 - c$ of its total released stress. In this thesis, unless otherwise specified, we set $c = 1 - 1/\sqrt{N}$. This means the smallest avalanche (one cell slipping) has on average a size of $1/\sqrt{N}$ and the largest avalanche (N cells slipping) has on average a size of $\sqrt{N}$. In Figure 2.5(a), we confirm this with the avalanche size $CCDF$ from simulations of four different system sizes.

We see a similar trend of increasing avalanche duration cutoff with increasing system size (Figure 2.5(a), inset). It is perhaps unsurprising that the smallest avalanche duration, on the other hand, does not depend on the system size (Figure 2.5(a), inset) since the minimum avalanche duration is fixed by the simulation time step to be $\Delta t = 1$ for all N. Importantly, neither the power-law slope of avalanche size nor duration changes with system size.

This trend of avalanche sizes translates of course to the size averaged on duration, shown in Figure 2.5(b), i.e., the smallest avalanche is proportional to $1/\sqrt{N}$. For small durations, there is a clear deviation from the predicted MFT slope, until $\sim T = 10$, when $\langle S|T \rangle$ kinks upward to a slope much more comparable to MFT. We again see that the slope of these power-laws for larger durations is not a function of system size, and the cutoff in avalanche size increases like $\sqrt{N}$. We point out that for the smallest simulated system size ($N = 100$) the power-law slope is always shallower than the prediction of MFT. Said another way, the scaling regime shrinks with decreasing system size until it is non-existent for system sizes smaller than $N \sim 100$.

The maximum velocity vs. avalanche size also shows an invariance of the power-law slope to the system size, comparing excellently to the MFT prediction (Figure 2.5(c)). Here, instead of a power-law cut off beyond some largest avalanche size, we find the power-law is obeyed only above a minimum avalanche size. This is because the smallest avalanches consist of only one slip, but, because of the disorder, the stress released by this one slip (and therefore the maximum stress-drop rate in the avalanche) varies according to the strength of the disorder. That is, this minimum cutoff does not show up if we take the avalanche size to instead be equal to the number of total slips, a definition we will make use of in future chapters (see, Chapter 4). For now, to best prepare us for describing friction force avalanches in the next chapter, we adopt the former definition. Finally, in Figure 2.5(d) we plot the first moment of avalanche size (main) and duration (inset) vs. system size. Both moments roughly follow a power law, with slope 0.1 and 0.25 for size and duration, respectively. We will compare this scaling behavior to experimentally extracted avalanches in samples of varying system sizes (see Chapter 7).

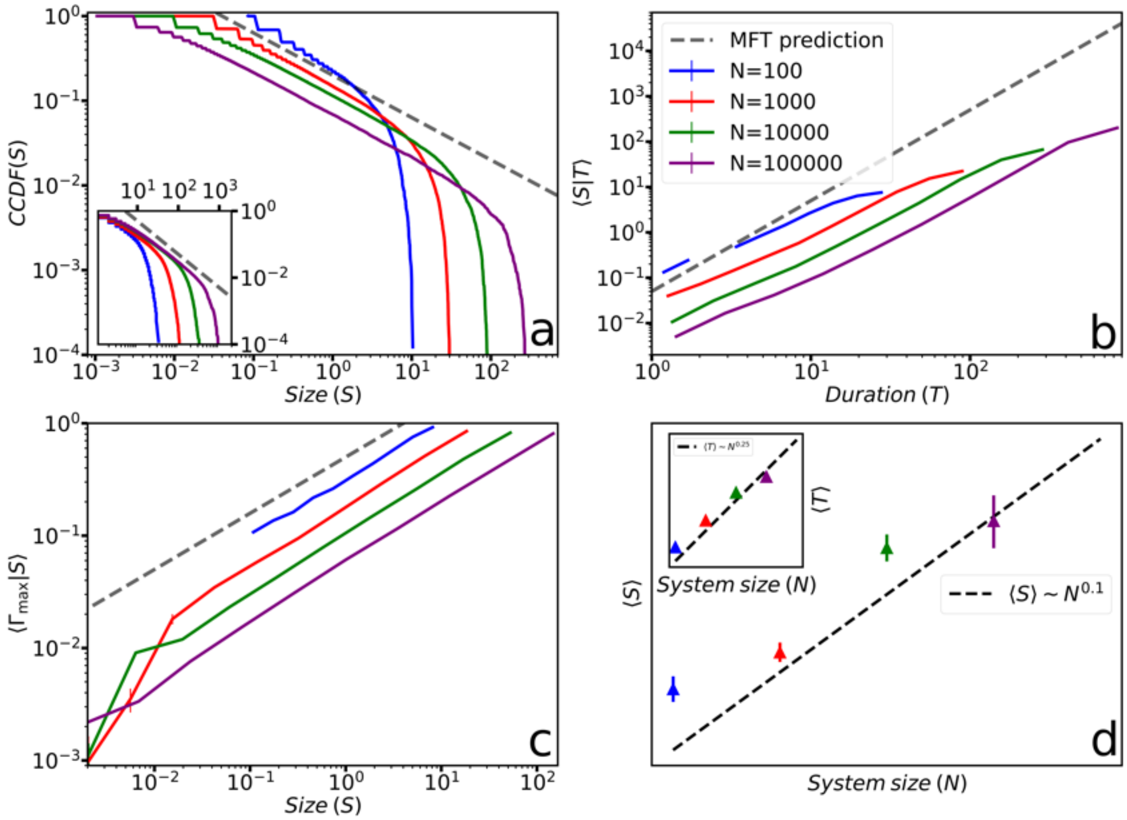

Figure. 2.5: System size. a) The $CCDF$ of avalanche sizes, shown for varying system sizes, with the MFT prediction shown as a black dashed line. The inset shows the corresponding $CCDF$ of avalanche durations. **b)** The avalanche size, averaged on duration, with the MFT prediction shown as a black dashed line, with error bars showing a 95% confidence interval. **c)** The maximum velocity for a given avalanche ($\Gamma_{\max}$) averaged on size, with the MFT prediction shown as a black dashed line, with error bars showing a 95% confidence interval. **d)** The first moment of avalanche size plotted vs. system size, with error bars showing a 95% confidence interval. The black dashed line is shown as a guide for the eye. The first moment of avalanche duration vs. systems size is plotted in the inset, with error bars showing a 95% confidence interval. A black dashed line is shown as a guide for the eye. See Appendix A for details on how these quantities were extracted and/or computed.

Chapter 3

Linking Friction Scales in Molybdenum Disulfide ($\mathrm{MoS_2}$) from Nano to Macro via Avalanches

3.1 Description of Collaborative Team

The contents of this chapter in its entirety are published in the journal of Tribology Letters (Tribol Lett 70, 82 (2022). https://doi.org/10.1007/s11249-022-01619-x). Several authors contributed to this work; John Curry, Tomas Babuska and Nicolas Argibay ran tribometer experiments and post-processed the data. Frank DelRio ran AFM/FFM experiments and post-processed the data. Adam Hinkle and Michael Chandross performed molecular dynamics simulations and post-processed the data. Karin Dahmen helped with interpretation of the data as well providing overall guidance for the project.

3.2 Abstract

Steady-state fluctuations in the friction force of molybdenum disulfide ($\mathrm{MoS_2}$), a prototypical lamellar solid, were analyzed experimentally for newton-scale forces and computationally via molecular dynamics (MD) simulations for nanonewton-scale forces. A mean field model links the statics and the dynamics of the friction behavior across these eight orders of magnitude in friction force and six orders of magnitude in friction force fluctuations (i.e., avalanches). Both the statistics and dynamics of the avalanches match model predictions, indicating that friction can be characterized as a series of avalanches with properties that are

predictable over a wide range of scales.

3.3 Introduction

Understanding the shear strength of lamellar solids like MoS_2 and graphene, materials that are widely used for their exceptional lubricity (i.e., low shear strength), is important for many engineering applications where traditional fluid lubricants cannot be used. Commonly used in mechanisms that operate in inert environments, such as in hermetic enclosures or exposed to the vacuum of space, MoS_2 films in the form of engineered coatings can undergo a transition to low friction coefficients (μ) during use, a phenomenon commonly referred to as "run-in." This transition in friction behavior is accompanied by structural changes of the coating surface [10–12] and the formation of a transfer film [13–15] on the mating counterface, generating a self-mated contact. The first few cycles of sliding cause shear-induced reorientation of the constituent atoms, leading to the formation of large, basally-oriented MoS_2 sheets at the shearing interface parallel to the direction of sliding. These self-mated sulfur-terminated lamellae interact through weak van der Waals forces that easily shear and provide low steady-state friction coefficients ($\mu < 0.05$). In cases where the MoS_2 films are sufficiently pure, devoid of adsorbate contaminants, and basally oriented, superlubricity can be achieved, which is associated with friction coefficients $\mu < 0.01$ [16]. Additionally, in this low friction regime, wear rates are also low, corresponding to the removal of as little as one molecular layer per sliding pass [17–19]. When in the steady-state friction regime, shearing planes of highly-ordered MoS_2 can display persistent fluctuations in friction force (as shown in Figure 3.1,a(b)) [17, 20] that have been attributed to stick-slip behavior [2, 21]. In this regime, atoms in adjacent lamella periodically become stuck in mutual local potential wells.

Here, we try to understand these low friction coefficients and wear rates from the statistics of fluctuations in the friction force that arise from underlying slip avalanches. An avalanche describes the cascading nature of failure events, experimentally observable as sharp drops in friction force. To understand the statistics of these avalanches, we model friction of MoS_2 lamellae via the slipping of elastically-coupled weak spots between material layers [2, 22, 23]. These weak spots arise from the natural variations at the interface between the two sliding layers. Weak spots are regions that are most prone to slipping. Without thresholds to sliding, sliding would proceed smoothly in response to an applied external force. Heterogeneities at the interface, for example via local asperities, introduce a local threshold force to sliding, implying that the interface will remain stuck (or 'pinned') for small external forces. In response to a small force increase the interface slips in avalanches until is pinned again by new asperities. Only for forces larger than a critical force, i.e. large enough to overcome all asperities, will the interface keep slipping. That critical force constitutes a critical point [24] that separates the pinned regime from the sliding regime. If we drive the system with a soft loading spring the system self organizes to operate near the critical force. In that case it responds with avalanches that are widely distributed in size with statistical distributions that are characterized by scaling exponents and scaling functions. These only depend on a few basic properties such as symmetries, dimensions, interaction ranges etc., but not on the microscopic details of the system. Below we describe a simple microscopic model that captures the basic properties and makes predictions for these universal exponents and scaling functions that can then be compared to experiments.

The statistics of these friction force avalanches resemble those in a chaotic regime of friction described in [19], as well as sudden stress drops arising from internal friction during deformation of a wide range of solids across a broad range of length scales [2, 25–30]. The model suggests that this behavior is similar because the underlying mechanism in all cases is slip avalanche weak spots. Predictions from this model have been confirmed by recent experiments on internal friction during slow compression of bulk metallic glass specimens; in these materials, the weak spots are shear transformation zones (STZs) [25] that predominantly slip in shear bands. Here, we show the statistics and dynamics of friction force avalanches measured in macroscale tribological experiments on sputter-deposited MoS_2 thin films. The results are compared with data from nanoscale molecular dynamics simulations of shear between MoS_2 lamella and predictions from a simple analytic model for the slip statistics during friction. The application of renormalization group methods to this simple analytic model enables the determination of aspects of the slip statistics and dynamics that are independent of microscopic details and therefore apply across a wide range of scales. We find good agreement between the experiments, simulations, and theoretical mean-field theory (MFT) model.

In plastic deformation experiments, avalanches manifest as sudden drops in stress [2, 25–30] with statistics and dynamics that are described well by a simple MFT model of coupled weak spots undergoing slip [2]. For friction, avalanches are observed as sudden drops in friction force with weak spots corresponding to local asperities at the interface. We define $u(\mathbf{r}, t)$ (sometimes referred to as the displacement field) as the interfacial offset since time $t = 0$ of a weak spot at location $\mathbf{r}$. With an applied shearing force F_{ext}, elastic forces from all other asperities $\sigma_{\mathrm{int}}(\mathbf{r}, t)$, and asperity pinning due to pinning forces $f_w[u, \mathbf{r}, \{u(\mathbf{r}, t' < t)\}]$, the displacement field $u(\mathbf{r}, t)$ evolves in time according to (see Chapter 2 and [2])

$$\frac{\eta\, \partial u(\mathbf{r}, t))}{\partial t} = F_{\mathrm{ext}} + \sigma_{\mathrm{int}}(\mathbf{r}, t) - f_w[u, \mathbf{r}, \{u(\mathbf{r}, t' < t)\}], \tag{3.1}$$

where η is an effective viscosity-like constant that sets the time scale of the slips and depends on the microscopic details of the system. Renormalization group calculations show that elastic interactions are sufficiently long range so that mean field theory, which approximates $\sigma_{\mathrm{int}}(\mathbf{r}, t)$ to be a constant, gives the correct scaling behavior in the physical dimension for slips along a shear band or frictional surface [5]. Thus, we take the mean field limit by assuming each site couples to the average force of all other sites $\bar{u} = \int dr\, u(\mathbf{r}, t)$ with $\sigma_{\text{int}}(\mathbf{r}, t) = J[\bar{u} - u(\mathbf{r}, t)]$, where J characterizes how strongly the slipping weak spots interact, i.e. it really is an "elastic coupling coefficient". To make contact with the physical system of interest, the change in local friction of a weak spot slipping at position $\mathbf{r}'$ is defined as $\Delta f(\mathbf{r}', t) = (J + K_{\text{L}})\Delta u(\mathbf{r}', t)$, where K_{L} represents the effective bulk loading stiffness. The friction force at any asperity $f(\mathbf{r}', t)$ can surpass a failure threshold value, f_{fail}, and slip to reduce its friction force to an arrest force, $f_{\text{arr}}(\mathbf{r}', t)$. The elastic interactions mediate the transfer of this released energy to other asperities and may trigger more slips in what we refer to as an avalanche. To mimic periodic noise during an experiment (e.g., from noise in the electrical circuits to measure forces or ringing of the force sensors), we overlay friction-force traces produced by our model with a signal resembling the periodic noise.

3.4 Methods

Friction force fluctuations of MoS_2 were measured in ultra-high vacuum (UHV; 10^{-9} torr) with a linear-reciprocating micro-tribometer at a 10 kHz acquisition rate [31] (Figure 3.1,a(c)). Force measurements were achieved via capacitance probes that measure the deflection of a custom dual-axis cantilever. On the end of the cantilever, a spherical probe (the counterface) is held fixed while the substrate is displaced parallel to the contact plane using a motorized stage. To establish a reference dataset for system noise, measurements were performed while the probe was out of contact and the stage was moving, and also while in contact with the stage stationary. More information regarding measurement uncertainty for force measurements can be found in previous manuscripts [32]. Experiments consisted of bidirectional sliding cycles at a normal force of 0.2 N and sliding speed of 1 mm/s. The 3.2 mm diameter spherical probes, made of $440C$ stainless steel, were brought into sliding contact against $440C$ stainless steel polished coupons with an approximately 100 nm thick pure MoS_2 film deposited using an N_2 spray impingement process [20]. Each experiment consisted of 500 bidirectional cycles with a 2 mm long stroke length. MoS_2 coatings under shear typically show a run-in period where friction coefficients and wear rates evolve from high ($\mu > 0.1$) to low ($\mu < 0.05$) as the surface structure transitions from amorphous (or near-amorphous, depending on the deposition method) to ordered lamellae [33]. To capture the physics of the steady state regime, only the final 400 cycles of each experiment were used to extract magnitudes S, lengths T, and force-drop rate Γ of the friction force avalanches (Figure 3.1(d)). The method to extract the avalanche sizes and durations from the data is as follows: Zero crossings of the time derivative of the friction signal ($\Gamma(t)$) from positive to negative indicate when avalanches (i.e. force drops) start ($t = t_{\text{start}}$). At the following zero crossing from negative to positive the avalanche ends (t=tend)., see (Figure 3.1,a(b)). The avalanche duration T is then defined as $T = t_{\text{end}} - t_{\text{start}}$, and the avalanche size S

is equal to the amount by which the friction force dropped during the duration, $S = F(t_{\text{start}}) - F(t_{\text{end}})$.

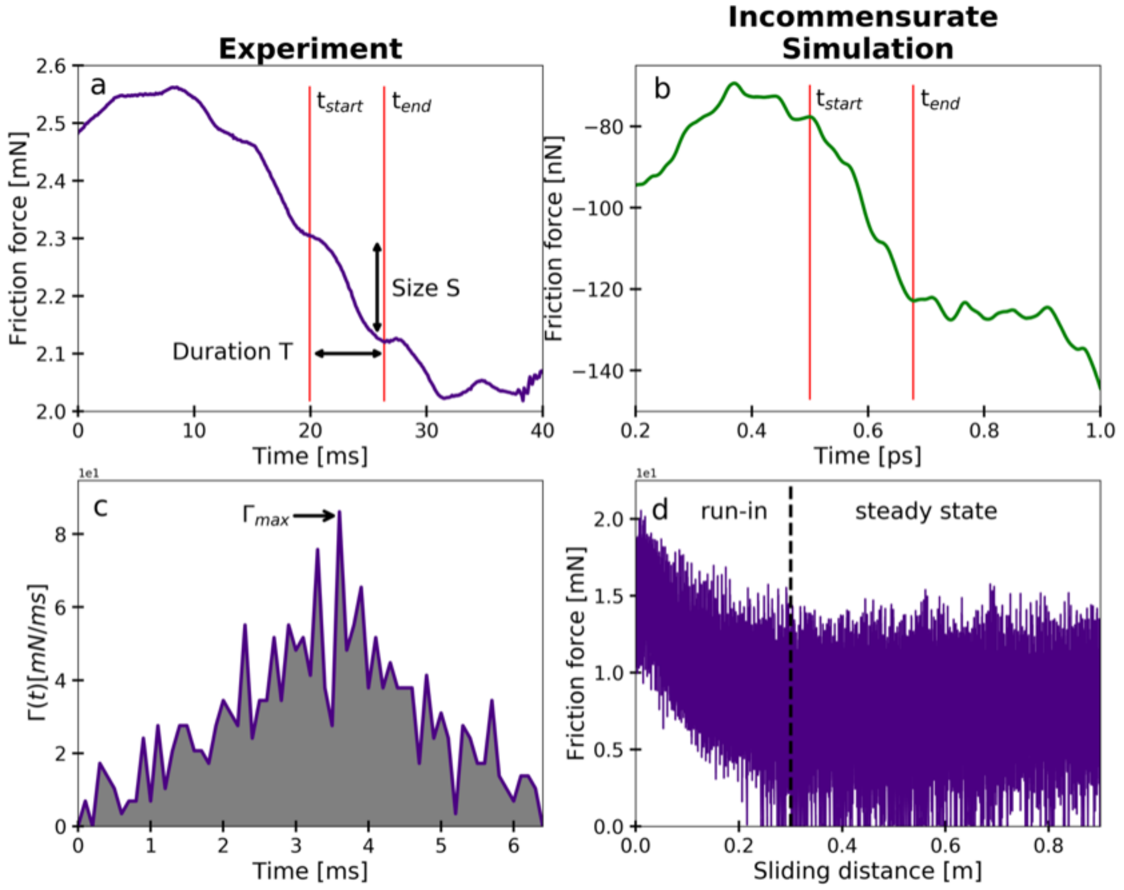

Figure. 3.1: Friction force fluctuations. a) Friction force measured in a UHV tribometer. Vertical red lines mark the start and end of a single avalanche. S is the total drop in force, while T is the difference between start and stop times. **b)** Friction forces from MD simulations show similar fluctuations. **c)** The force-drop rate profile of the single experimental avalanche and the maximum force drop rate Γ_{max}. We use only experimental data from the steady state regime, as shown in **d)**. Note in (b), the flat sections are not actually flat, but slightly curved downward, and thus the region between the red lines is one single avalanche.

Friction force microscopy (FFM) measurements were performed on an Asylum MFP-3D atomic force microscope (AFM). MoS_2-coated silicon substrates were mounted inside a humidity sensing cell on the sample stage. The humidity cell was attached to high purity nitrogen (N_2), such that the cell was maintained at a positive gas pressure and the relative humidity was close to that of the dry N_2 source. Rectangular quartz-like cantilevers with SiO_2 colloidal spheres were used to facilitate a well-defined contact geometry (Nanosensors CP-qp-SCONT-SiO probes with a sphere radius of $R = 1\ \mu\text{m}$). The normal force calibration C_{z} and spring constant k_{z} for a single probe were measured with the thermal fluctuation method [34]; multiple C_{z} and k_{z} values resulted in means and standard deviations of $C_{\text{z}} = 1.48 \pm 0.02$ nN/V and $k_{\text{z}} 14.6 \pm 0.2$ pN/nm. The lateral sensitivity S of the probe was assessed with the extended wedge method [35]; multiple S values resulted in $S = 12.8 \pm 1.9$. Following calibration, the probe was brought into contact with the MoS_2 until a normal photodiode voltage $V_z = 5.26\,\mathrm{V}$ was reached, translating to a normal force $F_n = 7.8\,\mathrm{nN}$. Repetitive, single-line scans were performed over a 1 μm distance at a frequency of 0.10 Hz. The lateral photodiode voltage during the trace V_{xt} and retrace V_{xr} were measured for each cycle. For each V_{xt} and V_{xr} pair, the friction loop width $W = (V_{xt} - V_{xr})/2$ and offset $\Delta = (V_{xt} + V_{xr})/2$ were calculated to determine friction force $F_f = C_z(S)(W)$ and coefficient of friction $\mu = (S)(W)/V_z$ [35]. The F_f data from the steady-state μ region (final 400 cycles) were analyzed for avalanches.

MD simulations were performed using a reactive force field [36] at a fixed shear velocity of 1 m/s. Two sets of simulations were performed: (1) one with MoS_2 single crystals (15 square lamellae of $MoS_2 \approx 9\,\mathrm{nm}$ on a side) to study commensurate (high friction) contact and (2) another with a single sheet rotated by 30° to study incommensurate (low friction) contact; details of the procedure are in [37]. Friction forces from the simulations showed persistent fluctuations like those in the experiments (Figure 3.1(b)). This enabled a determination of sizes, durations, and force-drop rates for the avalanches in the steady-state sliding regimes (after $\approx 100\,\mathrm{ps}$ and $\approx 35\,\mathrm{ps}$ of sliding for the commensurate and incommensurate cases, respectively).

3.5 Results

The distribution $\mathcal{P}(S)$ of force drop sizes S (Figure 3.1,a(b)) scales as $\mathcal{P}(S) \sim S^{-\tau}$, with universal (i.e., detail-independent) exponent τ. The prediction for τ depends on whether there is slip weakening ($\tau = 2$) or not ($\tau = 3/2$). This scaling holds for S within a power-law scaling interval $S \in [S_{\min}, S_{\max}]$. The interval is bounded by an upper cutoff $S_{\max}$ that depends on experimental tuning parameters such as static and dynamic friction, machine stiffness, sample size, and driving rate. The region $[S_{\min}, S_{\max}]$ is often referred to as the scaling regime, and the exponent τ is a critical exponent. To normalize scaling regimes with the macroscale experiments (for better comparison in a single figure) we scaled the size of MFT+noise simulations, commensurate MD simulations, incommensurate MD simulations and AFM/FFM experiments by $S_0 = 2.75 \times 10^{-6}\,\mathrm{N}, 2275\,\mathrm{N}, 9000\,\mathrm{N}$, and $1.75 \times 10^{-3}\,\mathrm{N}$, respectively. Similarly, the durations were rescaled by $T_0 = 1.4 \times 10^{-5}\,\mathrm{s}, 49 \times 10^{9}\,\mathrm{s}, 87.5 \times 10^{9}\,\mathrm{s}$ and $15\,\mathrm{s}$, respectively, along with and the maximum force-drop rates by $G_0 = 0.54, 4.64 \times 10^{-8}\,\mathrm{N/s}, 1.03 \times 10^{-7}\,\mathrm{N/s}$ and $1.5 \times 10^{-4}\,\mathrm{N/s}$, respectively. We extracted several universal scaling properties from this regime and compared them to the critical exponents and scaling functions predicted by the MFT model (Figure 3.2).

To avoid ambiguities from varying histogram bin widths (see Appendix A and [38]), we focus on power-laws in the complement of the cumulative distribution function ($CCDF$), or one minus the integral of the probability distribution function (PDF) $\mathcal{P}(S)$. The $CCDF$ of avalanche size $\mathcal{C}(S)$ is the probability to have

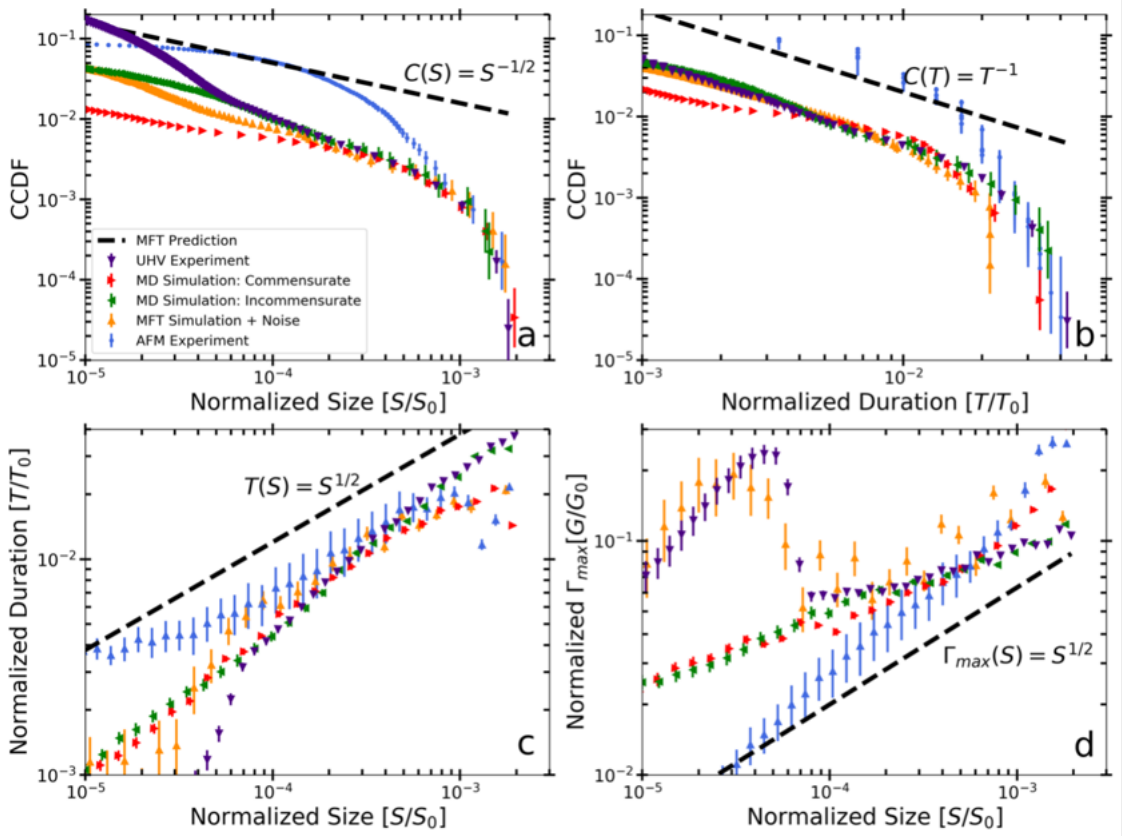

Figure. 3.2: Comparing to simple slip model. *CCDF* of **a)** size $\mathcal{C}(\mathbf{S})$ and **b)** duration $\mathcal{C}(\mathbf{T})$. (a,b) are plotted with error bars calculated with a Bayesian technique with 95% confidence interval. Distributions of (c) fluctuation duration and (d) maximum force-drop rate vs. size. (c,d) are plotted with error bars from the standard error in a log binned average over size. The power-law slopes predicted by the MFT model are plotted as black dashed lines. All units were normalized to the UHV experiments via three parameters (S_0, T_0, G_0). The power-laws are predictions for an infinite size system without noise or oscillations and overlapping blue data points in (b) stem from the discrete durations of AFM avalanches.

an avalanche of size S or larger, and scales as $\mathcal{C}(S) \sim S^{-(\tau-1)}$ (Figure 3.2(a), and Appendix A) with $\tau = 3/2$ in the absence of slip weakening. Similarly, the *CCDF* of avalanche durations $\mathcal{C}(T)$ scales as $\mathcal{C}(T) \sim T^{-(\alpha-1)}$ (Figure 3.2(b)) with $\alpha = 2$. Avalanches are predicted to proceed in time such that the distribution of sizes scales as $\langle S\rangle(T) \sim T^{1/\sigma\nu z}$ (Figure 3.2(c)), where angle brackets denote the average (see Appendix A).

The size-averaged maximum force-drop rate scales as $\langle\Gamma_{\max}\rangle(S) \sim S^{\sigma\nu z}$ (Figure 3.2(d)). Historically, $1/\sigma\nu z$ is the fractal dimension of the avalanches in the scaling regime, with σ describing the stress dependence of the largest avalanche in the scaling regime, z describing the scaling of the duration of avalanches with their diameter, and ν describing the scaling of the correlation length near criticality. Within error bars, all critical exponents obtained in experiments and simulations were consistent with the MFT simulations with emulated machine noise (see Appendix B, Section 2). To strengthen agreement, we compared to stricter predictions. Two of the most stringent validations are the time profiles of the average force-drop rate as a function of size and duration, or Γ-profiles [11, 15, 32]. The Γ-profiles of $N = 21$ avalanches within a bin size centered at $\mathcal{S}_{\mathrm{i}}$ were averaged to obtain the size averaged Γ-profile (see Appendix A)

$$\langle \Gamma_{\mathcal{S}_{\mathrm{i}}}(t_{\mathrm{j}}) \rangle = \Sigma_{\mathrm{n}-1}^{N} \frac{\Gamma_{\mathrm{n}}^{i}(t_{\mathrm{j}})}{N}. \tag{3.2}$$

$\Gamma_{\mathrm{n}}^{\mathrm{i}}(t_{\mathrm{j}})$ is the force-drop rate of the n^{th} avalanche in the i^{th} bin, where N is the number of avalanches in the bin $\mathcal{S}_{\mathrm{i}}$. Time was offset such that all avalanches start at $t = 0$ ($t_{\mathrm{j}} = t - t_0$, where t_0 is the avalanche start time). The MFT model predicts that the averaged force-drop rates, drawn from different size bins in the scaling regime of $\mathcal{C}(S)$, have the same form as the universal scaling function (see Appendix A and [2, 22])

$$\frac{\langle \Gamma_{\mathcal{S}_{\mathrm{i}}}(t_{\mathrm{j}}) \rangle}{\sqrt{\mathcal{S}_{\mathrm{i}}}} = A(\frac{t_{\mathrm{j}}}{\sqrt{\mathcal{S}_{\mathrm{i}}}}) e^{-B(\frac{t_{\mathrm{j}}}{\sqrt{\mathcal{S}_{\mathrm{i}}}})^2}. \tag{3.3}$$

Figure 3.3,a,c,e (left column) show that plotting $\frac{\langle \Gamma_{\mathcal{S}_{\mathrm{i}}}(t_{\mathrm{j}}) \rangle}{\sqrt{\mathcal{S}_{\mathrm{i}}}}$ versus $(\frac{t_{\mathrm{j}}}{\sqrt{\mathcal{S}_{\mathrm{i}}}})$ collapsed the profiles. The experimental noise caused deviation from the analytical prediction at small values of $\frac{t}{S^{0.5}}$, however when the experimentally-determined periodic noise signal was included in the MFT model, the deviations were nearly the same (Figure 3.3(c) and Appendix B, Section 2, Figure B.1, top right). Matching the functional form of scaling functions was a more stringent test of the theory than merely matching power-law exponents. As shown below, we find good agreement between experiments and theory for two different scaling functions and four exponents.

In addition, $\mathcal{C}(T)$ followed a power-law in the duration scaling regime (Figure 3.2(b)). Averaging the force drop-rate time profiles of avalanches of about the same duration $\mathcal{T}_{\mathrm{i}}$ gave the duration averaged Γ-profiles $\langle \Gamma_{\mathcal{T}_{\mathrm{i}}}(t_{\mathrm{j}}) \rangle$. MFT predicts that avalanches in this regime obey the scaling law

$$\frac{\langle \Gamma_{\mathcal{T}_{\mathrm{i}}}(t_{\mathrm{j}}) \rangle}{\mathcal{T}_{\mathrm{i}}} = 4(\frac{t_{\mathrm{j}}}{\mathcal{T}_{\mathrm{i}}})(1 - (\frac{t_{\mathrm{j}}}{\mathcal{T}_{\mathrm{i}}})). \tag{3.4}$$

Figure 3.3,b,d,f (right column) show that this prediction also matches experiments and MFT with added noise. The noise in the experiments again obscured the shapes, here for both small and large $\frac{t}{S^{0.5}}$, and the MFT with noise showed similar deviations (Figure 3.3(d)). In summary, we see strong agreement between the experiments and the model predictions for (1) the avalanche statistics, as seen in the scaling exponents of the statistical distributions, and (2) the avalanche dynamics, described by the average Γ-profiles that match the predicted critical scaling functions. This agreement between theory and experiments indicates that the fluctuations in the friction force indeed reflect the underlying critical point predicted by our simple MFT model. That is an important result because the model predicts that the statistics and dynamics of these fluctuations should be universal and independent of the microscopic details and the scale of the system. In other words, the MFT predicts that friction experiments on different materials should yield similar results, and that the results from the lab scales should be transferrable to systems on much larger or much smaller scales, as long as the symmetries, dimensions and interaction ranges are the same.

3.6 Discussion

System oscillations paired with periodic noise limited our analysis to regimes away from spikes in the power spectrum (see Appendix B, Section 4 and Figure B.4 - B.6). Although the scaling regime in the experiments spans less than two orders of magnitude, the agreement of our data with more than six measures including scaling collapses with universal scaling functions suggests the data strongly support model predictions. It is the robust scale invariance predicted by the model that explains why both UHV experiments at mN scales and MD simulations at nN scales follow the same MFT scaling behavior. As predicted by MFT, both experiments and simulations showed a distinct scale invariance, reflecting an underlying critical point. The agreement in friction force fluctuations across six orders of magnitude suggests a common thread of slipping elastically-coupled weak spots, as described by MFT.

The MFT model is not necessarily expected to describe experimental results for small event sizes and durations, as it is designed to describe only long wavelength and low frequency regimes where universal behavior is found. At small scales, microscopic details come into play and deviations between MFT and experiments are unsurprising. Such deviations can be seen in the $CCDF$ (Figure 3.2,a(b)) for small sizes and durations. Interestingly, in the small size and duration regime, experimental results more closely matched MD simulations with incommensurate contacts. While the registry of experimental systems is unknown, the two simulation cases (i.e., commensurate and incommensurate) represent bounding cases and suggest that macro-scale steady-state low friction consists of primarily incommensurate shear (see Appendix B, Figure B.7).

Moreover, there was no evidence for critical avalanches in the AFM/FFM data (Figure 3.2). However, this was again expected, considering the single-asperity nature of FFM [15, 17, 18]: Avalanches in FFM involve slipping of the entire contact area, i.e., the entire tip slides on the surface until it is caught by a strong enough asperity. The model presented here predicts that power-law scaling is only seen when the slipping size is much smaller than the size of the contact area. If that is not the case then the model predicts that finite size effects prevent the emergence of power-law scaling, because the tip size sets a length scale that destroys scale invariant behavior. We estimated a ratio of slipping area to contact area, for the largest avalanche, and found it to be of order 1000 to 1 for the AFM experiments. Thus these dramatic finite size effects explain the absence of scaling of avalanches in AFM data. On the other hand, for the MD simulations we find this ratio to be 1 and for the UHV experiments it is 2.5. If the largest slip is comparable to the contact area, then all smaller slips in the experiment are smaller than the contact area, and power-law scaling can emerge as predicted. We note that the MD simulations do not have multiple, simultaneous slip locations, but instead the whole interface slides at once [39]. Importantly, there are layers beyond the interface that also slip, so in this case our model can be understood as describing the slipping of elastically coupled layers.

Importantly, we point out that the AFM data can still be described by our model, albeit in a parameter regime where the slips are large compared to the contact area (or system size). Similar effects can be achieved in the model, for example for large slip weakening in the system, and in a somewhat different way also for effectively soft loading springs. Either way this type of behavior is seen in model parameter regimes where the system is not close to a critical point and thus the avalanches are not power-law distributed. Thus these model predictions match the observations from the AFM experiments also.

Finally, new modeling work [40] further supports the idea that friction force fluctuations should match our model with weakening [5, 7], based on comparison to simulations of a rate and state model [41]. We note importantly here that, though the slope of the size and duration $CCDF$ in the mentioned work are steeper than the MFT prediction without weakening (instead being a closer match to the prediction of the model with weakening) our MFT+noise simulations suggest this is simply a consequence of the periodic machine noise. It will be interesting to see if the exponents resembling those of the weakened model will persist in experiments without machine noise.

3.7 Conclusion

In summary, the results demonstrated that friction behavior of lamellar solids, and possibly materials in general, can be accurately described by our MFT model, and that this framework can be used to understand the similarities and differences between macroscale and nanoscale friction measurements. The framework can also be used to transfer friction results between scales, predict future friction behavior, and diagnose and quantify potential changes in friction over time. As friction is linked to plastic deformation, the broad agreement with predictions from the MFT model suggests the possibility of linking friction force statistics with internal friction and wear rates as a future application.

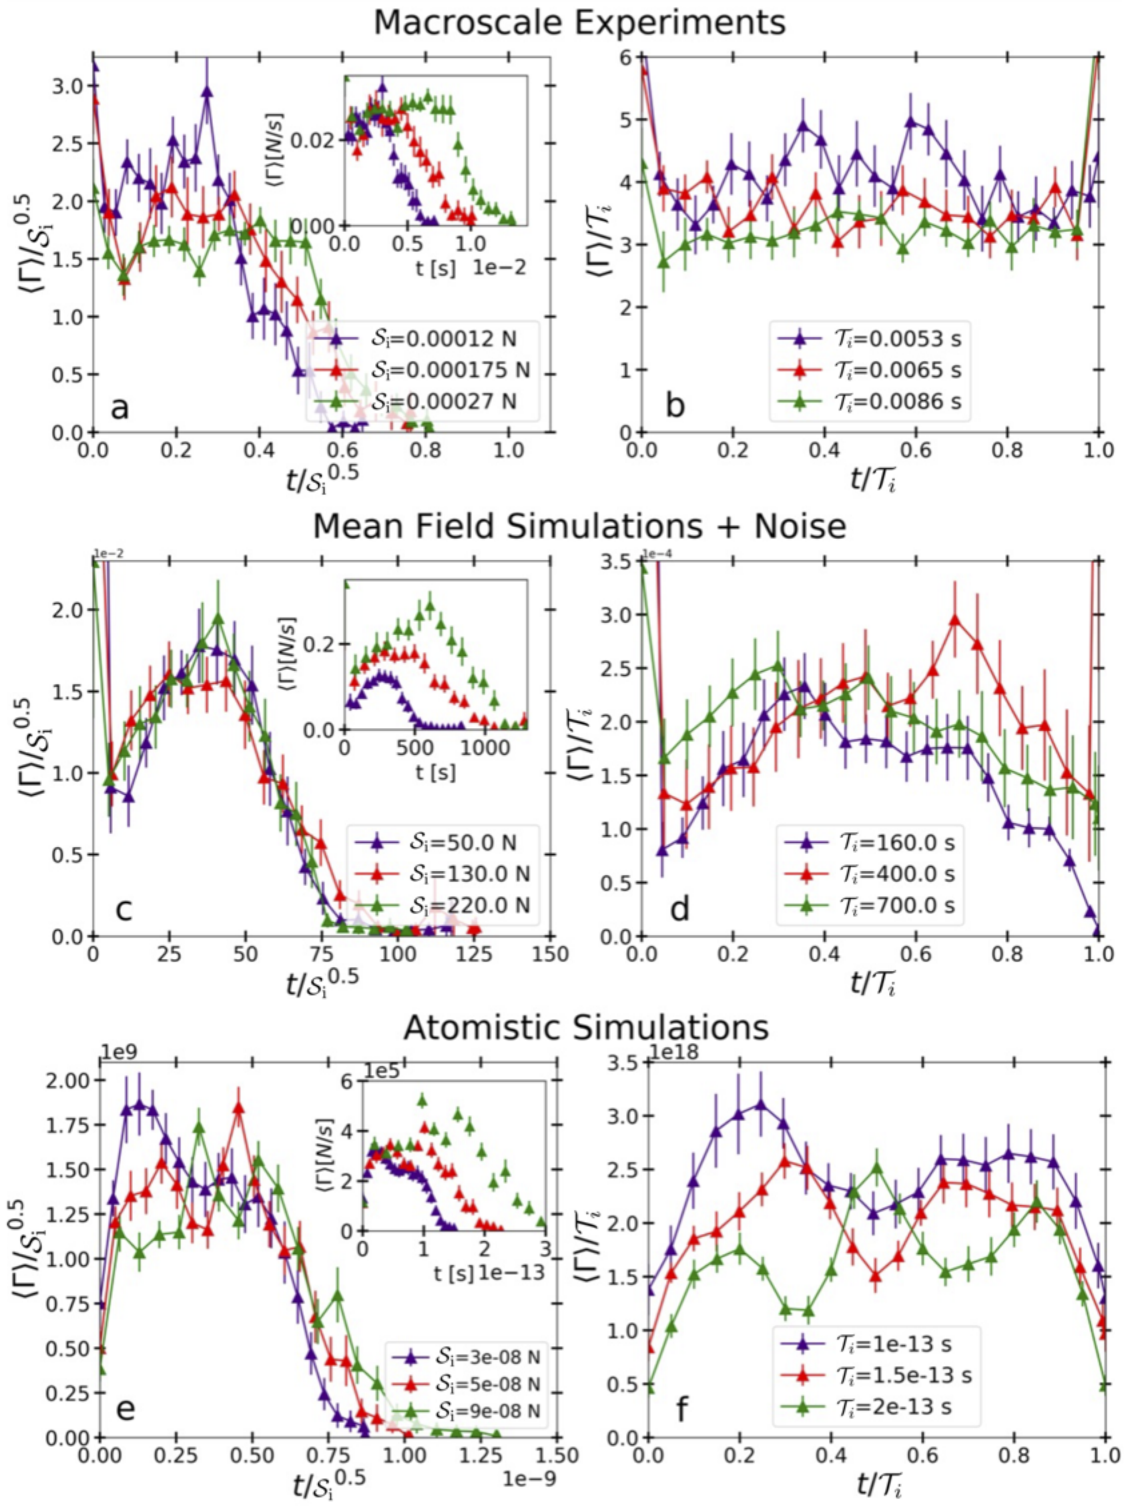

Figure. 3.3: Average force-drop rate profiles. Average force-drop rate profiles for **(a,b)** UHV macroscale experiments, **(c,d)** MFT calculations with added noise, and **(e,f)** MD simulations. Profiles computed from n avalanches found in a narrow size bin centered on $\mathcal{S}_\mathrm{i}$ are shown here in the left column insets for three bins in the scaling regime. Markers are shown with error bars calculated from the standard error in a log binned average over time. The size averaged profiles are predicted to collapse onto each other when rescaled by $\mathcal{S}_\mathrm{i}^{-1/2}$ (left column) and $\mathcal{T}_\mathrm{i}^{-1}$ (right column). The results for experiments and simulations are consistent with MFT.

Chapter 4

Analyzing Neuronal Spiking Data

4.1 Introduction

Now we transition from discussing friction force drops in materials to in-*vitro* neuron spiking signals from cultured mouse cortex. We'll introduce the analysis necessary to extract and characterize avalanches from spiking neuron activity, with a focus on how this characterization changes as we vary input parameters of the analysis.

To record spiking activity from many neurons, its typical to use an electrode which can be used to record voltage fluctuations from nearby neurons. A Micro electrode array (MEA) is an array of several (hundreds of) electrodes, spaced only micrometers apart. A very common way to record neuronal spiking activity is to culture a system of neurons on top of an MEA. Once the culture becomes spontaneously active, the electrical signals at each electrode inform us when a neuron spikes. With many details left to references [42], the data can be processed from continuous voltage signals at each *electrode* to discrete spiking signals at each *neuron*. In this data form, the neurons can be in one of two states; either they are spiking (we'll represent this with a 1) or they are not spiking (represented by a 0). Now it should be clear that our data set consists of a binary list of zeros and ones for each of the N neurons in the system.

From these lists we wish to learn something about neural dynamics, specifically, in this thesis we are going to be concerned with neuronal avalanches. We illustrate two techniques used for extracting avalanches from the kind of spiking data described above. We'll refer to the first technique as the 'Traditional' technique since it was introduced when the first neuronal avalanches were discovered [43]. The second technique was introduced under the name 'Causal Webs' over a decade after the discovery of neuronal avalanches [44]. After comparing results from each method, we close by highlighting how a possible artifact may be introduced to

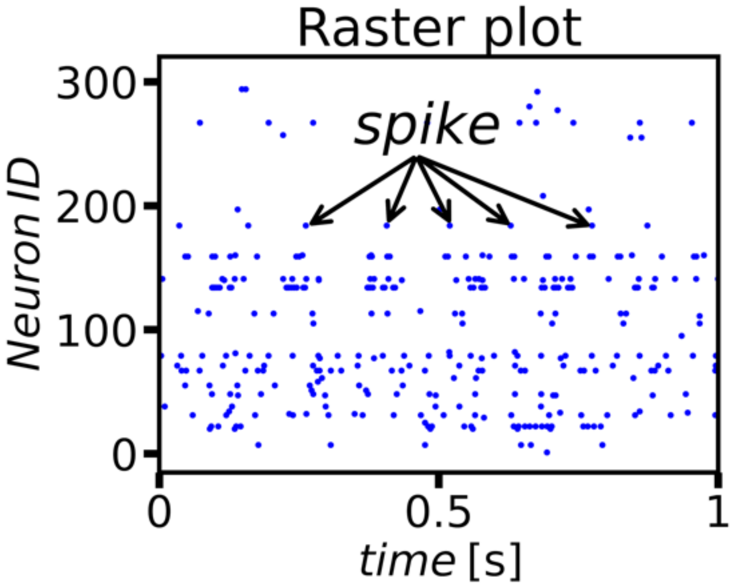

Figure. 4.1: Raster plot. The $N \sim 300$ neurons in the data set are labelled with an ID which is an integer between 1 and N. One second of spiking data from mouse cortex is plotted vs. time for each neuron. To indicate that neuron $'n'$ spiked at time $'t'$, we place a blue dot at the coordinate $(x = t, y = n)$.

the data from the causal webs processing.

4.2 Traditional Neuronal Avalanche Analysis

Starting with N binary lists representing the state of all N neurons of the system in discrete time, we can visualize the spiking behavior with a Raster Plot. To construct the raster plot, it will be useful to transform our data from a list of zeros and ones to a list of discrete times that a neuron spiked. We'll call this list $t_{\mathrm{i,n}}$, where $'i'$ indexes time (i.e., the discrete times that the neuron spiked) and $'n'$ indexes the neuron. A raster plot is built out of the set of $t_{\mathrm{i,n}}$ as follows; 1) Take the first neuron in the system, label it with the number $'1'$ and place a dot at the coordinates $(t_{\mathrm{i,1}}, 1)$ for all times $t_{\mathrm{i,1}}$ that neuron 1 spiked. 2) Repeat this for all N neurons at the corresponding y-value i.e., place dots at the coordinates $(t_{\mathrm{i,n}}, n)$ for all times $t_{\mathrm{i,n}}$ that neuron $'n'$ spiked. Figure 4.1 shows the raster plot for an exemplary 1-second long portion of a spiking data set.

From the raster plot we notice that there are not only correlations of a single neuron in time (a neuron firing in periodic fashion, for instance), but also correlations between different neurons in the system (for instance the quasi-synchronous activation of many different neurons). The goal of this section is to try and access these system wide correlations from the individual neurons spike-traces through a spatio-temporal average. This coarse-grained spiking rate will yield information about the systems collective dynamics, not just the behavior of a single neuron. We will extract avalanches from this signal to compare their statistical

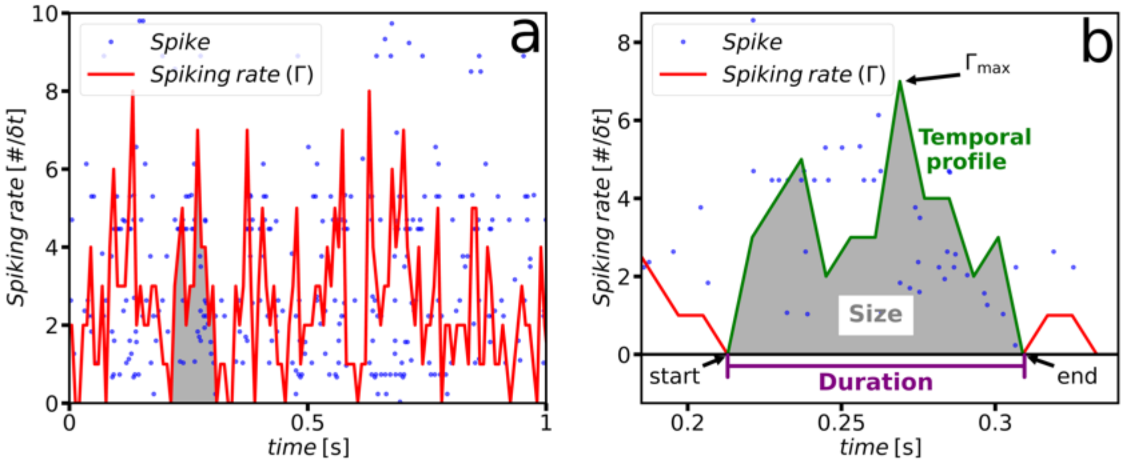

Figure. 4.2: Traditional neuronal avalanche extraction. a) The same raster plot seen in Figure. 4.1, with the spiking rate (Γ) overlayed in red. A temporal width of $\delta t = 8$ ms was used to construct the spiking rate signal. An example avalanche is shaded in grey. **b)** A zoom in on the shaded avalanche from Figure. 4.2(a). The start and end of the avalanche (when $\Gamma = 0$) are labelled with arrows, with the difference between the two times defining the avalanche duration (purple). The spiking rate signal itself between the start and end times is the temporal profile (green) and the total area under the spiking rate signal is the avalanche size (grey). The maximum value of Γ in the example avalanche, Γ_{max}, is marked with an arrow.

properties with a candidate theory of neuronal avalanches.

We'll refer to the system wide spiking rate as Γ since it will turn out to play an analogous role to the Γ defined in Chapter 2. We can construct Γ easily once we have the list of single neuron spike times, $t_{\mathrm{i,n}}$. First, we simply concatenate all $'n'$ lists together into one list representing the spike times of the entire system. The resolution of our data set is 20 kHz, but we would like to know what the averaged number of spikes are in a small window of time δt, otherwise known as the spiking-rate. By dividing the time axis into bins of width δt, we add up all spikes occurring within each bin to construct the spiking-rate at a temporal width δt (Figure 4.2(a)). The corresponding effects of δt on the estimated avalanche properties will be discussed in Section 4.4 below.

We define an avalanche as any excursion of the spiking-rate above zero. An example avalanche is shaded in grey in Figure 4.2(a), which we zoom in on in Figure 4.2(b). To be explicit, the avalanche starts once Γ becomes nonzero and ends once Γ returns to zero (marked by an arrow labelled $'$start$'$ and $'$end$'$, respectively, in Figure 4.2(b)). The time difference between these points, $t_{\mathrm{end}} - t_{\mathrm{start}}$, defines the avalanche duration, marked in purple in Figure 4.2(b). The total number of spikes between t_{end} and t_{start} defines the avalanches size (grey, Figure 4.2(b)) and the temporal profile is defined as the portion of the spiking-rate between t_{end} and t_{start} (green, Figure 4.2(b)). We will also be interested in statistics of the peak value of this profile, marked with an arrow labelled Γ_{max} in Figure 4.2(b).

Before we explore the resulting avalanche statistics, we describe a separate technique for extracting

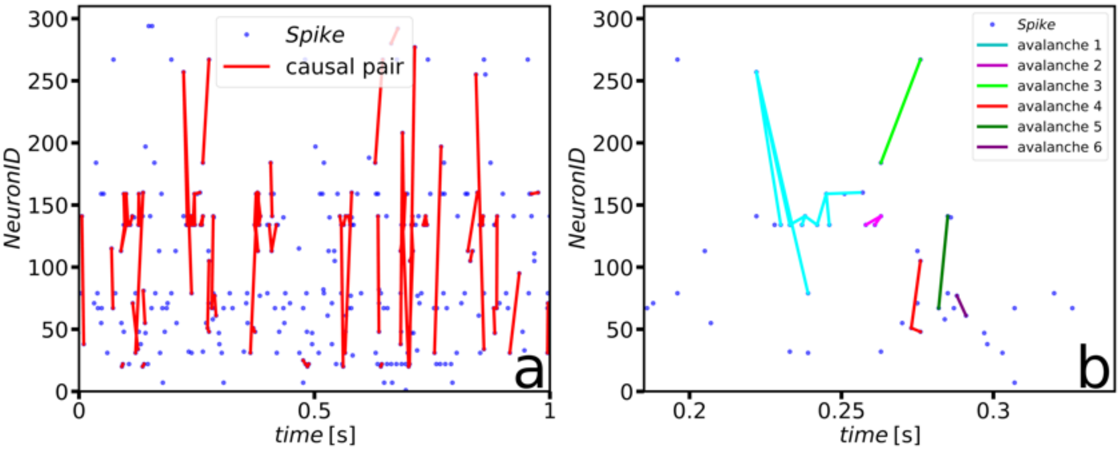

Figure. 4.3: Causal web avalanche extraction. a) The same raster plot seen in Figure. 4.1, with causal pairs drawn as red lines connecting any two spikes. **b)** A zoom in on the same section of spike trace as Figure. 4.2(b). Causal web avalanches are defined as any connected cluster of causal pairs. We show pairs belonging to the same avalanche in the same color. The size of the avalanche is the number of total neurons in a connected cluster, and the duration is the difference between start and end times.

neuronal avalanches known as 'Causal Webs' [44].

4.3 Causal Webs Avalanche Analysis

The causal webs technique, unlike the traditional one, does not involve directly averaging over neurons or time. Instead, there, one uses the estimated functional connectivity between neurons, paired with the spiking data and a measure of transfer entropy (TE) to extract causally connected pairs of spiking neurons. Briefly, when neuron n_1 spikes at time t_1 and neuron n_2 to spikes at time t_2, the causal webs method would determine whether or not neuron n_1 caused neuron n_2 to spike. The estimated functional connectivity between the two neurons can be used in part to determine this—if two neurons are not strongly connected, the spike pairs may not be causally connected. However, spike pairs between two neurons that are strongly connected are a candidate for a causal pair. The transfer entropy, an averaged measure of information transfer between neurons in the time series, will either increase or decrease from including this proposed causal connection. If the TE increases, the pair of spikes is deemed to be 'causal' *as long as* the time separating the two spikes isn't *too* long. We call the longest amount of time permitted between two spikes in order for them to be deemed causal, the 'maximal time delay parameter', or just 'd' for short (see [44] for more details).

Our raw data is a list of the tuples representing causal pairs. From this list of tuples, avalanches are defined as any connected cluster of causal spike pairs. In Figure 4.3(a), we show the same section of data as in Figure 4.2(a), but with causal spike pairs identified by a red bar connecting the two dots. We zoom in on the same subset of data as Figure 4.2(b) in Figure 4.3(b) and show causal pairs that are part of the same avalanche with bars of the same color. Here the avalanche size is again equated to the number of spikes in the cluster, with the duration again being the difference between the start and end times.

The first striking feature we see is that, where the traditional method finds a single avalanche, the causal webs method may find several. This not only leads to obscured avalanche statistics, but a deeper issue; which one is the *true* avalanche? It may be that there is not a clear answer to that question at this time, but we proceed, nevertheless.

4.4 Compare and Contrast: Traditional vs. Causal Webs Analysis

In this section we will compare and contrasts the estimations of the empirical avalanche properties, from the traditional vs. causal webs methods. We focus for the time being on the statistics of avalanche sizes and durations as well as their temporal profiles and the peak value of these profiles Γ_{max}. We will compare the results to the analytical predictions of our MFT model (see Table C.1) and also empirical estimations taken from neuronal avalanches in rat cortex [45].

4.4.1 Size and Duration Distributions

In Figure 4.4(a), we plot the CCDF of avalanche sizes as estimated by the traditional method (shown in greens) and the causal webs method (shown in reds). For the traditional analysis, we show results using three different values of the temporal bin width, δt, with darker green lines corresponding to larger values of δt. For the causal webs analysis, we show results using three different values of the maximal delay parameter, d, with darker red lines corresponding to larger d. We can see increasing δt or d has the effect of making the power-law slope *shallower*. The decrease of slope is particularly pronounced when δt is adjusted from 1.5 ms to 5 ms in the traditional method of analysis. Interestingly, this shallowest slope best matches the empirical estimates made in rat cortex with the same experimental setup (dashed purple line in Figure 4.4(a), see [45]). There, they too used the traditional method of avalanche detection and with a bin width of $4ms$, so this result is not surprising, but nevertheless confirms the statistics of these avalanches across these two species (rat and mouse). In the inset of Figure 4.4(a) we show the avalanche duration $CCDF$ for the corresponding methods and parameter values used in the main figure. The same general trend of decreasing slope with increasing δt or d seen for avalanche sizes is seen in avalanche durations as well. Again, using a $\delta t = 5$ms in the traditional analysis produced a slope closest to the empirical estimations in rat cortex (dashed purple line, Figure 4.4(a), inset). We do see, however, that the range of the power-law scaling is much smaller in the inset than it is in the main figure. This is of course consistent with our expectations from the theory, since the durations scale like the square root of the size (see Chapter 2), and thus the range of durations (and accordingly the range of durations that scales) is significantly reduced in comparison to the corresponding range of sizes.

4.4.2 $\langle T|S\rangle$ vs. Size

From the traditional analysis, the duration averaged in logarithmically spaced bins of size follows a power-law with a steeper slope than predicted by the MFT, but which is consistent with the empirical estimation found in rat cortex [45] (Figure 4.4(b), dashed purple line). This slope is consistent (within error bars) for all three values of δt used in the analysis. The mean has a small kink upward that appears for larger sizes in the dark green ($\delta t = 5\,\mathrm{ms}$) line. In contrast, for avalanches not too large, the causal webs analysis produces a shallower power-law than the traditional analysis, one that is much closer to the prediction of MFT (Figure 4.4(b), reds). However, for larger avalanche sizes, the slope kinks *slightly* upwards for all values of d. We note that for all values of d and for the largest value of δt shown, duration averaged on size kinks upwards to the same slope (which itself is steeper than the MFT prediction).

4.4.3 $\Gamma_{\max}$ vs. Size and Duration

Using the traditional analysis, the maximal value of the spiking-rate, $\Gamma_{\max}$, averaged in logarithmically spaced bins of size is shown in Figure 4.4(c). We can see that for $\delta t = 0.9$ and $1.5\,\mathrm{ms}$, the slope for small avalanches is slightly shallower than the MFT prediction, though still matching it within error bars. For larger avalanches, the power-law begins to fall off. For the largest value of δt ($5\,\mathrm{ms}$), the slope for the smaller avalanches is not as clear, and the largest avalanches kink downward instead of gradually falling off as was seen with the traditional analysis. The same general trends are true for $\Gamma_{\max}$ as estimated via. causal webs (Figure 4.4(c), reds), though the power-law slope for smaller avalanches seems to match the MFT predictions better than the traditional analysis (greens).

In the inset of Figure 4.4(c), we again show $\Gamma_{\max}$, but this time averaged on duration. Here the error bars are larger than in the main figure and the scaling regime extends over a smaller region of durations, making a power-law slope tough to define. Still, some sections obeying power-laws can be discerned. Specifically, for smaller avalanches found using the traditional analysis, a power-law shallower than the MFT prediction is followed for a very small range for durations. This range broadens for increasing δt. We see again for longer duration avalanches that the power-law falls off at some maximal duration that increases with increasing δt.

Similar trends are seen for the causal webs analysis, except the width of the scaling regime is slightly larger than in the traditional analysis, and this width does not vary significantly with d, as it did with δt. We also note that d sets the minimum duration for the power-law, below which $\langle \Gamma_{\max}|T\rangle$ is significantly shallower. After a maximal duration, the power-law slope begins to fall off. In the next section we discuss

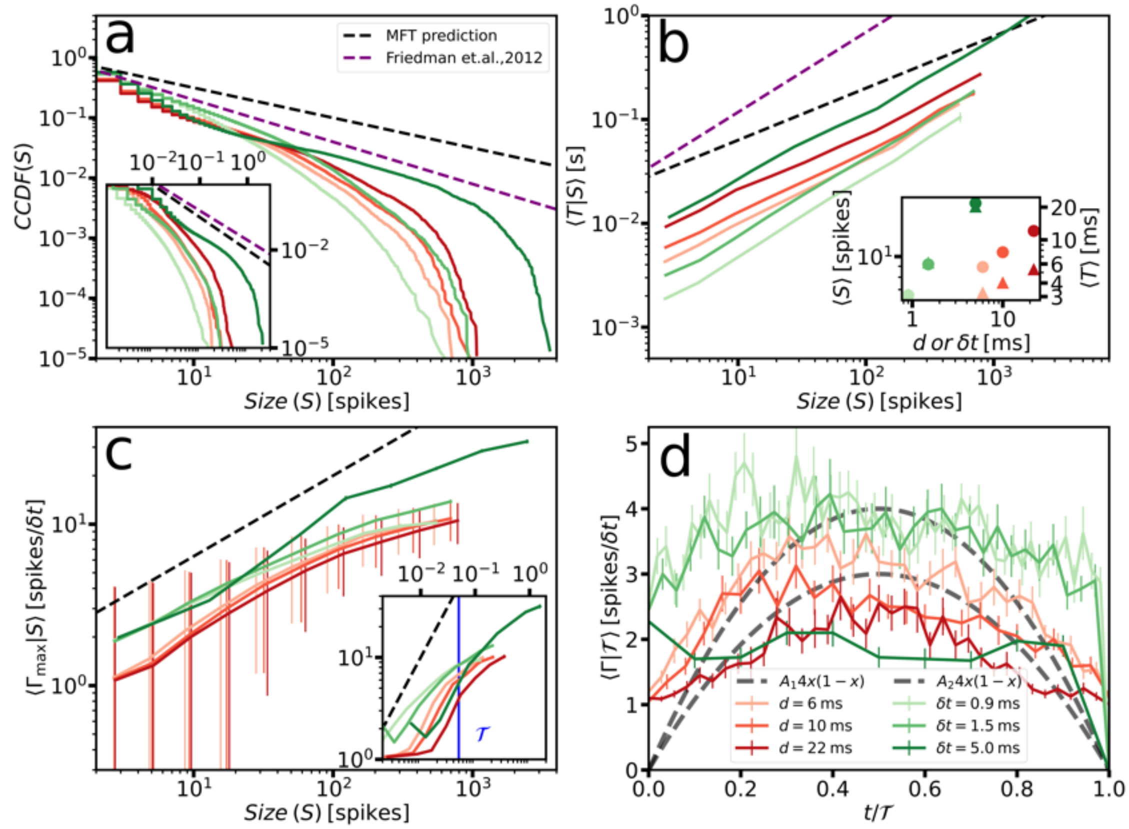

Figure. 4.4: Compare and contrast avalanches from both techniques. **a)** *CCDF* of avalanche size (main) and duration (inset), as estimated by the traditional avalanche analysis (greens) and causal webs avalanche analysis (reds). Darker lines correspond to increasing the temporal bin width (δt) and maximal delay parameter (d) for the traditional and causal webs avalanches, respectively. **b)** Avalanche duration averaged in ten logarithmically spaced bins of size, shown for causal webs and traditional avalanches with reds and greens, respectively. Error bars are showing a 95% confidence interval. **c)** $\Gamma_{\max}$ averaged in ten logarithmically spaced bins of size (main) and duration (inset), for the causal webs and traditional avalanches, with reds and greens, respectively.Error bars are showing a 95% confidence interval. **d)** The averaged avalanche shape for a given duration, $\mathcal{T}$, shown in reds and greens for the causal webs and traditional avalanches, respectively. Error bars represent the standard error of the mean. To reduce clutter, only every tenth bin is plotted for each of the averaged shape profiles shown in d). In all figures, we show the prediction of MFT with a black dashed line, and the empirical estimates found in [45] with a purple dashed line. See Appendix A for more details on how these quantities were computed.

the averaged temporal profiles for a given duration, $\mathcal{T}$, which is marked in blue in Figure 4.4(c) for reference.

4.4.4 Temporal Profiles

In Figure 4.4(d) we show the avalanche profile averaged in a bin of duration $\mathcal{T} = 50\,\mathrm{ms}$ (see Figure 4.4(c), vertical blue line). For the traditional analysis we again find agreement between $\delta t = 0.9$ ms and $\delta t = 1.5$ ms, where both profiles roughly follow a parabola with maximum height around 4 [spikes/δt]. In contrast, for $\delta t = 5$ ms, we see a relatively flat profile at the y value 2 [spikes/δt] for most of the avalanche. This is simply because at this largest value of δt, durations of $\mathcal{T} = 50$ ms are actually quite short in comparison to the longest avalanches, meaning they are also quite small and thus have a smaller maximum height. For the causal webs analysis (see Appendix C for details on how these shapes were computed), the maximum height consistently decreases with increasing d, and the shapes become more localized around $t/\mathcal{T} = 1/2$. We note there is a slight leftward skew in the averaged profiles estimated by causal webs with $d = 6$ ms and $d = 10$ ms, as well as the averaged profile estimated with the traditional analysis at $\delta t = 0.9$ ms.

4.5 Post-Processing Artifacts and Bursts

In this section, we take a closer look at the effects of the maximal delay parameter, d, used in the causal webs method. We have already found that the slope of the $CCDF$ of avalanche size decreased with increasing d, which corresponds to a decreasing value of τ. In Figure 4.5(a), we show the $CCDF$ of avalanche sizes for five different values of d, clearly showing the trend of decreasing τ. For a more quantitative look, we show the value of τ as extracted by a Bayesian maximum likelihood estimator (see Appendix C and [46]) as a function of d in the inset. In the inset we not only show the slope of the avalanche *size* $CCDF$ with triangular markers, but also the slope of the avalanche *duration* $CCDF$ and the slope of $\langle T|S\rangle$ vs. size, with circular and square markers, respectively. From this we notice that the slope of the avalanche duration $CCDF$ decreases with increasing d, just as τ does, but at a faster rate. In contrast, we see the slope of the power-law of $\langle T|S\rangle$ vs. size actually *increases* with increasing d.

We recall that the MFT not only predicts values for the corresponding scaling exponents, but it also predicts a relation that should be fulfilled by the exponents as well, Eq. (2.9). In Figure 4.5(b) we illustrate this relation by plotting the left- and right-hand sides of Eq. (2.9) on the x and y axes, respectively. If the scaling relation is held, then the dot should lie on the line $y = x$, shown in Figure 4.5(b) as a black dashed line. We can see that, increasing d roughly leads to a decrease in $(\alpha - 1)/(\tau - 1)$ and an increase in $1/\sigma\nu z$. Thus, from Figure 4.5(b) it seems there may be a hypothetical value of d that leads to this scaling relation being satisfied (somewhere between $d = 6$ ms and $d = 10$ ms).

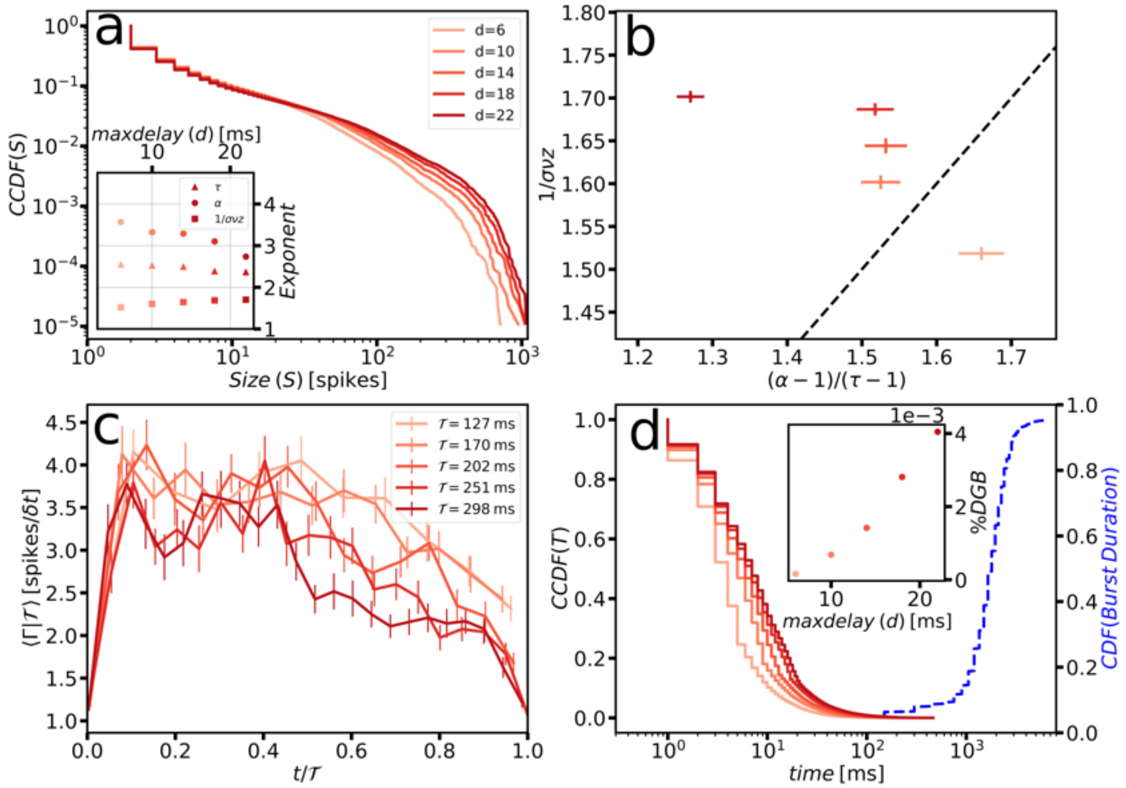

Figure. 4.5: Effects of changing maximal delay time, d, Sample 1. a) $CCDF$ of avalanche size, shown for five different values of d, the longest time allowed between two spikes if they are to be deemed causally connected. The inset shows the estimated values for the critical exponents τ, α and $1/\sigma\nu z$, with triangular, circular, and square markers, respectively. The error bars are insignificant enough to be hidden behind the markers (see Appendix C for details on how the error was estimated). **b)** For each value of d, the left and right hand of the scaling relation Eq. (2.9) is plotted on the the x and y axes, respectively. A black dashed line is shown on the line $y = x$ to mark where the scaling relation is fulfilled. **c)** The thirty longest avalanches are used to construct an averaged temporal profile, $\langle\Gamma|\mathcal{T}\rangle$, for each value of d. The median duration of the thirty events is shown in the legend and, to reduce clutter, data are only plotted for every tenth bin and the error bars show a 95% confidence interval. **d)** The $CCDF$ of avalanche durations, shown for each of the five values of d. The blue dashed line shows the CDF of burst durations (see Appendix D). The inset shown the percentage of avalanches in the sample that had durations larger than the minimum burst duration, '$\%DGB$', seen in the same sample, for each of the values of d. Error bars in all plots represent the standard error of the mean. See Appendix A for more details on how these quantities were computed.

In further investigation, we've sought out to reconcile a slower timescale phenomena seen in the data, known as 'bursting' (see Appendix C) with the extracted avalanche properties. We've averaged together the temporal profiles of the thirty largest avalanches and plotted them on top of each other in Figure 4.5(c). In these averaged profiles, in stark contrast to the downward facing parabola predicted by the MFT (see Figure 2.3(d)), we see a profile that is skewed to the left quite significantly. We hypothesis that this is due to the influence of bursts, which carry a similar profile that decays in time [47], though we return to this feature in Chapter 7.

To validate this hypothesis, we checked for overlap between the avalanche durations (Figure 4.5(d), reds) and the burst durations (Figure 4.5(d), black dashed line, see Appendix C). We indeed find that the percentage of avalanche durations that overlap with burst durations (i.e., the percentage of avalanches with durations larger than the minimum burst duration, what we define as $\%DGB$, or 'the percent of avalanches with **D**urations **G**reater than the minimum **B**urst duration', see Appendix C) increases with increasing d (Figure 4.5(d) inset). Aside from the median avalanche duration in the bin increasing with increasing d, we do not see any qualitative changes in the five separate average profiles (see Figure 4.5(a)).

Finally, we check to see how robust these results are when we change to a new sample. For bookkeeping purposes, we'll call the first sample (what we have analyzed until now) 'Sample 1' and this new sample 'Sample 2'. In Sample 2 we find, in contrast to Sample 1, the avalanche profiles become more and more skewed to the left with increasing d (see Figure 4.6(c)). This would suggest that the value of d is too large and gives credence to the idea that d should be adjusted for each sample. To add to that, we find in Sample 2 that increasing d decreases τ and α at a much faster rate than it did for Sample 1(Figure 4.6(a), inset). This manifests in a scaling relation that does not seem to be satisfied for any hypothetical d. Lastly, we also find a larger percentage in Sample 2 than Sample 1 of avalanches with durations greater than the minimum burst duration (Figure 4.6(d)).

4.6 Discussion

In this chapter we introduced techniques for extracting avalanches from spiking neuron data and illustrated their resemblance to avalanches produced by the MFT model of Chapter 2. Though we find clear evidence of scaling in the extracted neuronal avalanches, the exact values of the scaling exponents are not consistent with the MFT model. While we find that both traditional and causal webs analysis produce power-law distributed avalanche sizes and durations, the exponents were steeper than $\tau = 3/2$ and $\alpha = 1$, respectively. Both τ and α were found to decrease linearly with increasing temporal input parameter (δt or d) and in the most extreme case of the traditional analysis ($\delta t = 5$ ms) the slopes even decreased enough to be consistent with the MFT

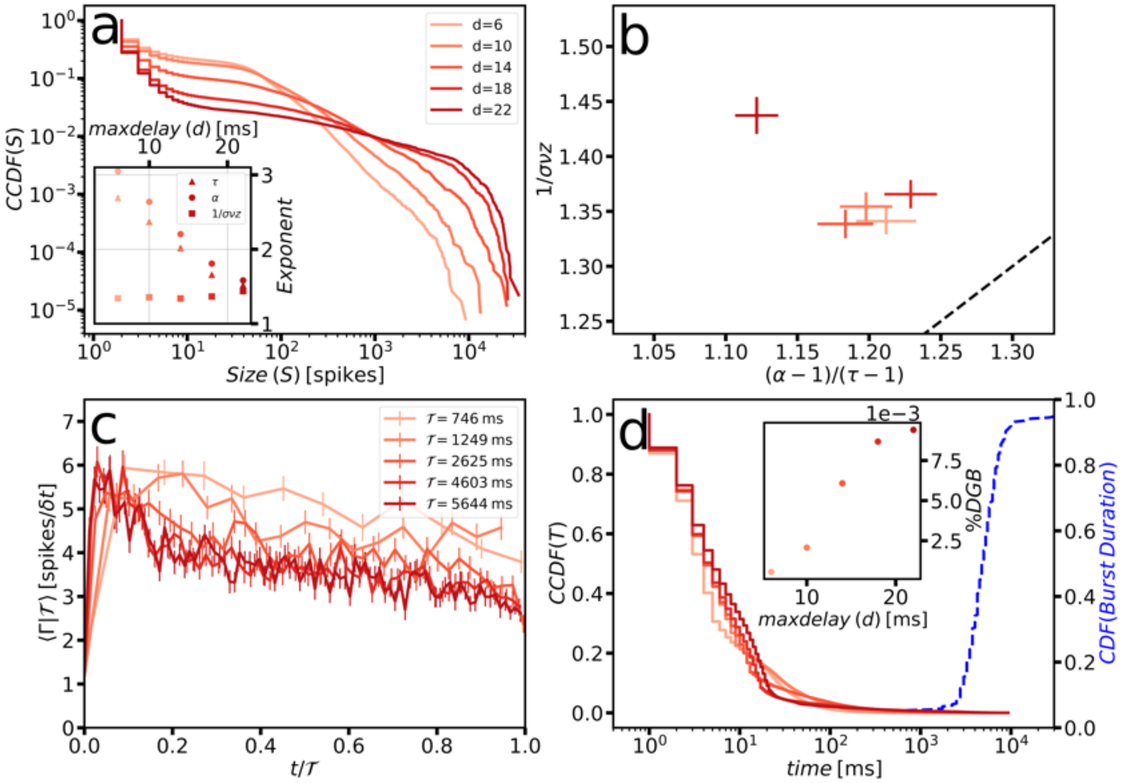

Figure. 4.6: Effects of changing maximal delay time, d, Sample 2. a) $CCDF$ of avalanche size, shown for five different values of d, the longest time allowed between two spikes if they are to be deemed causally connected. The inset shows the estimated values for the critical exponents τ, α and $1/\sigma\nu z$, with triangular, circular, and square markers, respectively. The error bars are insignificant enough to be hidden behind the markers (see Appendix C for details on how the error was estimated). **b)** For each value of d, the left and right hand of the scaling relation Eq. (2.9) is plotted on the the x and y axes, respectively. A black dashed line is shown on the line $y = x$ to mark where the scaling relation is fulfilled. **c)** The thirty longest avalanches are used to construct an averaged temporal profile, $\langle \Gamma | \mathcal{T} \rangle$, for each value of d. The median duration of the thirty events is shown in the legend and, to reduce clutter, data are only plotted for every tenth bin and the error bars show a 95% confidence interval. **d)** The $CCDF$ of avalanche durations, shown for each of the five values of d. The blue dashed line shows the CDF of burst durations (see Appendix D). The inset shown the percentage of avalanches in the sample that had durations larger than the minimum burst duration, '$\%DGB$', seen in the same sample, for each of the values of d. Error bars in all plots represent the standard error of the mean. See Appendix A for more details on how these quantities were computed.

prediction. Knowing that the power-law exponents estimated for a data set may depend on the given input parameters of the analysis is important in its own right, but what is perhaps more interesting is making a connection between two seemingly unrelated phenomena; the effect of increasing d merges avalanches together, just as what happens in magnet avalanches when you increase the rate at which you sweep the external magnetic field [48].

We found that α changed with changing maximal delay parameter, d, in a linear fashion i.e., $\alpha \sim \alpha_0 - \phi d$, with ϕ a constant. What makes this interesting is past work on avalanches in magnets [48, 49], where an analogous equation is seen, describing the change of α. There, the avalanches were composed of spin flips in a magnet (see Chapter 1) and the avalanche duration exponent α increased linearly with the rate that the external magnetic current was swept. That work proposed that the increasing sweep rate caused more than one avalanche to nucleate at a time, and, since the experiment only measures changes in magnetization, all simultaneous avalanches appear as one merged avalanche. Multiple avalanches merging shifts the statistics of the avalanche durations in the magnet. Here, in analogy, the linear increase of α with d suggests that multiple avalanches are merging together as you increase d. This aligns with our intuition since, as you increase d, new causal connections appear and link together previously separate clusters of causal pairs (i.e., multiple avalanche merge into one).

The causal webs method for all d, and the traditional method with $\delta t = 5\,\mathrm{ms}$ found agreement (within error bars) with the MFT predicted slope of $\langle \Gamma_{\max} | S \rangle$ and $\langle \Gamma_{\max} | T \rangle$ vs. S and T, respectively. The duration averaged temporal profiles for the sampled duration bin $\mathcal{T} = 50\,\mathrm{ms}$ also roughly obeyed the parabolic MFT prediction. For both methods, the power-law in $\langle \Gamma_{\max} | S \rangle$ and $\langle \Gamma_{\max} | T \rangle$ eventually decayed for the largest sizes and durations. This is unsurprising since we saw in Figure 4.5(c) and Figure 4.6(c) that in general, the larger duration shapes are flatter than the parabolic prediction of MFT, providing a probable explanation for the decay of power laws in $\langle \Gamma_{\max} | T \rangle$ for the largest durations.

Both techniques also admit scaling of $\langle T | S \rangle$ with S, and the causal webs technique even produces an exponent that is nearly consistent with, but slightly steeper than, MFT. When using the traditional analysis, however, the predicted slope is much steeper than anticipated by the MFT model. A steeper slope here means the avalanche durations are longer than expected, for a given size. An avalanche like this would be long but flat; again, resembling the abnormally flat duration averaged temporal profiles (see Figure 4.5(c) and Figure 4.6(c)).

Lastly, we raise concern for the possibility of mixing avalanches with bursts. We found that using too large a value of the maximal delay parameter, d, made duration average avalanche profiles skew leftward for the longest durations—resembling the exponential decay of a burst. Indeed, we found the durations of these longer avalanches were longer than the minimum burst duration measured in the same experiment. This affect was more prevalent in some samples than others (see Appendix C) suggesting that we may need to choose d separately for each sample, or that d is *non-universal*.

4.7 Conclusion

We showed that neuronal avalanches are consistently found to have scaling regardless of which technique is used or what the temporal input parameter is set to. In general, the avalanches extracted with the causal webs technique seemed to have larger scaling regimes with critical exponents and averaged shape profiles that matched the MFT predictions most closely. When extracting avalanches with the causal webs method, we illustrated that care must be taken to not overlap avalanches with bursts and shed light on the idea that the maximal delay parameter may need to be tuned for each sample.

Chapter 5

A First Step at Modeling Neuronal Avalanches: Weakening and Global Inhibition

5.1 Description of Collaborative Team

The contents of this chapter in its entirety is under review for publication in Physical Review E. Several authors contributed to this work; John Beggs ran the experiments and performed post-processing and also helped edit the manuscript. Karin Dahmen supervised the project and also helped edit the manuscript. I ran model simulations, analyzed experiments and wrote the manuscript.

5.2 Abstract

A simple model for a population of spiking excitatory neurons is used to model seizures (see Chapter 2). Two key features are introduced into the model: 1) an external inhibitory population of neurons that drain potential out of the system and 2) a threshold weakening mechanism that adds memory to the dynamics. A seizure-prone state corresponds to a phase where the system can suddenly switch from a "healthy" state with small scale-free avalanches, to a "seizure" state with almost periodic large avalanches ("seizures") where all neurons in the system spike. Simulations of the model predict the dependence of the average time spent in each state on the amounts of inhibition and threshold weakening. Our results predict that increasing inhibition will make seizures less frequent and shorter lived. The results are compared to empirical and

theoretical examples of similar behavior in neuronal systems.

5.3 Introduction

In this chapter we redefine α as the strength of the inhibition instead of the avalanche duration critical exponent (as it was defined in earlier chapters). This won't cause any confusion since we will not discuss the duration critical exponent in this chapter.

In the simplest sense, the characteristic feature of a neuron is the sudden release of an action potential (or 'spike') after the membrane potential surpasses a threshold. After a neuron spikes, the released membrane potential is transferred through synaptic connections to other neurons in the network. Spiking neurons can trigger other neurons to spike also, leading to collective behavior such as synchronous oscillations [50–53], spiral waves [54] and cascades of spiking activity (or 'neuronal avalanches') [43, 45, 55–57]. In a neuronal avalanche a single neuron spikes and triggers one or more neurons to spike afterwards. Tools from statistical physics, such as the theory of phase transitions and criticality, provide simple predictions for the statistical properties of these avalanches that can be validated by experiments [43, 45, 55, 56].

Criticality refers to a system that has large fluctuations because it is going through a 'continuous phase transition'. In the following, it will suffice to know that spiking neuronal systems have been hypothesized to be operating near a continuous phase transition between the so called 'absorbing' state, where the avalanches are small and the 'active' state where, avalanches proliferate through the entire system (i.e., grow macroscopically large). The critical point defines the parameter values where this phase transition occurs. Near the phase transition we expect avalanches of all sizes.

The total number of spikes emitted during an avalanche (or avalanche 'size'), S, can be recorded for many avalanches in an experiment to obtain a histogram, or the probability distribution of S, $\mathcal{P}(S)$ (see Appendix A). Statistical physics predicts that, near criticality, this distribution will look like a power-law over a broad range of sizes (see Chapter 2 for more details). In finite sized systems or for parameters slightly away from the critical point, there will be a maximum avalanche size S_{max} that depends on the system size or the distance of the parameters from the critical point. This maximum avalanche size is reflected by a cutoff of the probability distribution (e.g., $\mathcal{P}(S)\ e^{-S/S_{\mathrm{max}}}S^{-\tau}$, see Table B.1). The regime where the avalanche size distribution can be described by such simple scaling laws is referred to as the scaling regime. Indeed, empirically observed distributions of avalanche sizes and durations resemble power-laws, with the power-law extending over 3 or more decades in size (S) [43, 45, 56–58].

Further, these scaling predictions are equivalent for entire classes of models that have differing microscopic details—a feature known as 'universality'. Many biologically motivated models thus predict the same universal scaling of the avalanche statistics [45, 59–66]. These models are said to be in the same 'universality class'. The renormalization group (RG) makes it possible to identify which properties are relevant for the scaling behavior on large scales, and which ones are irrelevant. Therefore, the RG can be used to predict which biologically motivated amendments of existing models alter the universality class [67–69]. These tools are now commonly used to study neuronal activity and will inform our model.

Here, to include memory in an avalanche, we introduce a mechanism that, once a neuron has spiked, dynamically reduces the spiking threshold for the remainder of the avalanche. This threshold reduction (or 'weakening') changes the avalanche statistics [70], and so is a relevant parameter (in the sense outlined in the previous paragraph). Specifically, with weakening, periodically occurring system spanning (or 'runaway') avalanches occur, that are composed of a significant amount of cyclic recurrent spiking [71]. The recurrent spiking was validated by in-vitro cortical spiking data that too showed a large amount of cyclic recurrent activity in large avalanches. These results point towards a memory mechanism like weakening that has a significant effect on the size distribution of the spiking avalanches.

We wish to mention that similar recurrent large avalanches can be obtained by overshoots that may be caused for example by temporary increases in the coupling between neurons [5]. Here we use the weakening as a representative of a class of many different mechanisms that can cause such large avalanches. The results discussed here should apply in similar ways to other mechanisms, like overshoots or temporary coupling increases in the same way. In either case, the large recurrent avalanches may be similar to seizures, so it is of great interest to understand how their size can be reduced.

To this end, we also include inhibitory neurons, which send negative ionic currents when they spike, preventing more neurons from spiking soon after [72, 73]. In this work we test the hypothesis that the size of seizure-like large avalanches can be reduced by the inhibitory neurons. For simplicity, we model an excitatory population of neurons that feel an averaged effect of an external population of inhibitory neurons. The effect of an inhibitory neuron spike amounts to a reduction in the potential of the whole system. In the following, we survey the implications of our model for our understanding of the epileptic state and neuronal dynamics in general.

5.4 Model

To study the dynamics of a system of N spiking neurons, it's typical to define the state of the network as $\vec{V} = V_{\mathrm{i}}$ where V_{i} is the potential of neuron 'i'. We note that if the local cell stress, τ_{i}, is replaced by the neuron potential, V_{i}, the simulation outlined below is identical to that of Chapter 2. The network state evolves when a neuron reaches its threshold potential, $V_{\mathrm{thresh}} = 1$, whence it emits a spike, reducing its own

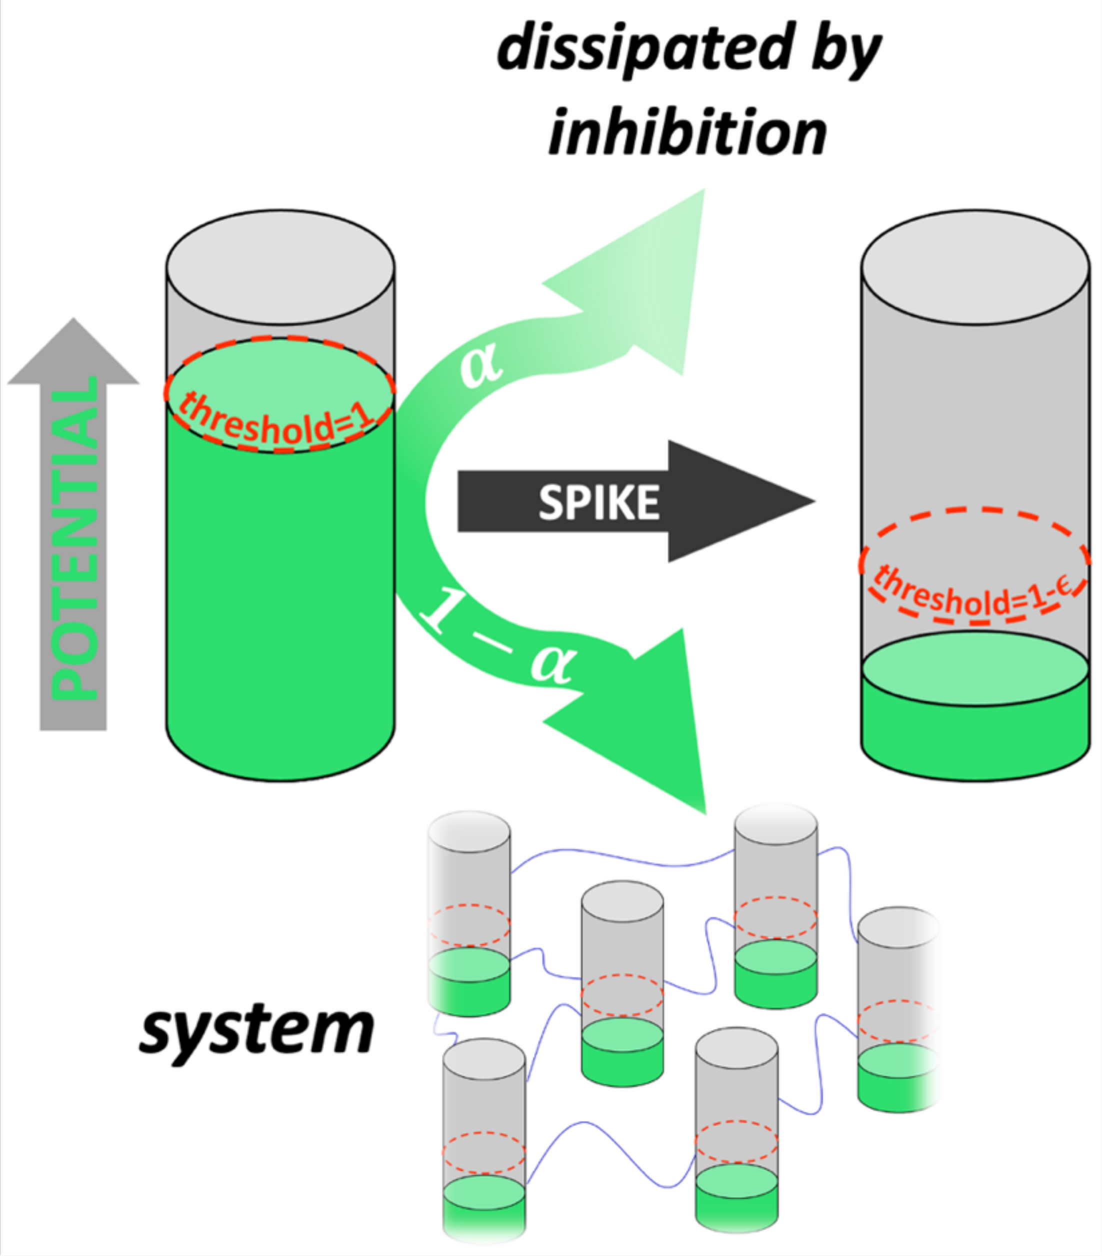

Figure. 5.1: Graphic of main changes to typical neuron models. When a neuron reaches threshold (left) it spikes, sending a portion of the potential released back into the system. The rest of the potential is dissipated by inhibition. The threshold is reduced after spiking until the avalanche is complete (right).

potential by $\Delta V_\mathrm{i} = 1 - V_\mathrm{r,i}$. Here, $V_\mathrm{r,i}$ is the neuron-dependent reset potential, which is distributed randomly to capture the inherent disorder in neuronal systems. The released potential is dispersed evenly to the rest of the system via the synaptic connections, increasing the potential of all other neurons by $\Delta V_\mathrm{i,system} = \Delta V_\mathrm{i}/N$.

To include realistic neuronal effects, here we will introduce two features: 1) A separate population of inhibitory spiking neurons whose effect will be to suck potential out of the system with strength α. This amounts to a change in $\Delta V_\mathrm{i,system} \rightarrow (1-\alpha)\Delta V_\mathrm{i,system}$. This assumption is justified by experiments that show that the number of inhibitory neurons that spike in each avalanche is proportional to the number of excitatory neurons that spike in the same avalanche (see Appendix D, Section 1). We also point out that the inhibition defined this way is related to the effective spring constant, K_{L}, of the model used in Chapter 2 and Chapter 3 describing slip at an interface. Specifically, $\alpha = 1 - c = 1 - J/(J + K)$. 2) A threshold reduction mechanism has been shown to be useful in replicating experimentally resolved recurrent spiking in large neuronal avalanches [71]. After a neuron spikes, the threshold necessary for a subsequent spike is lowered $V_{\mathrm{thresh}} \to V_{\mathrm{thresh}} - \epsilon \Delta V_{\mathrm{i}}$ for the remainder of the avalanche. $\epsilon \in [0, 1]$ is the strength of the weakening. We can see this effect is equivalent to instead fixing the threshold and increasing the gain function (or strengthening the couplings) of a neuron after it has spiked; either way the neuron reaches threshold faster. Though the latter may be more biologically motivated, we use the former for ease of comparison with previous work [5, 7]. In the following we'll therefore discuss the behavior of our system as a function of inhibition and weakening (i.e., in $\{\alpha, \epsilon\}$-space). A graphical representation of the model described above is shown in Figure 5.1.

5.5 Results

The model described above produces dynamics in the form of avalanches. There are two predicted steady states, determined by the position of the system in $\{\alpha, \epsilon\}$-space. For $\alpha > \alpha^* = \epsilon/(1 + \epsilon)$, in the region of Figure 5.2(a) labeled 'stable', the system response consists only of scale-invariant avalanches (Figure 5.2(b), $\epsilon = 0$) up to a maximum size that broadens with increasing alpha as $S^n_{\max} \sim 2\alpha^{-2} + \mathcal{O}(\alpha^2)$ [7]. We will refer to this as the normal state. For $\alpha < \alpha^*$, in the region labeled 'bi-stable', a second steady state is possible, consisting of temporally quasi-periodic system spanning avalanches, with smaller scale free avalanches in the times between. We refer to these smaller avalanches happening in between the system spanning avalanches as the "background activity". The size distribution thus has a power-law for small sizes (from the background activity), up to a cut off $S^s_{\max} = 2(1 - \epsilon)^2/(\epsilon - \alpha)^2 + O((\epsilon - \alpha)^2) < S^{\mathrm{n}}_{\max}$ [7], followed by the system-spanning avalanches, together yielding a bimodal distribution (Figure 5.2(b), $\epsilon = 0.4$, red). These large avalanches are referred to as runaway events, and we refer to this state as the seizure state.

Since the switching times grow with the system size, for $N \to \infty$, the initial conditions (within the bi-stable region) determine whether the steady state system operates in the normal or seizure state. For finite N, however, the probability of spontaneously switching between the two states is nonzero and switching from one to the other is expected. This switching from the normal state into the seizure state is facilitated by a nucleation mechanism, where an avalanche of size larger than the nucleation size S_{cr} $(1 - \epsilon)N/(1 - \alpha)$ grows into a system spanning avalanche [7]. The system can also transition back into the normal state from the seizure state if there are not enough neurons available to participate in a runaway event (see [7]).

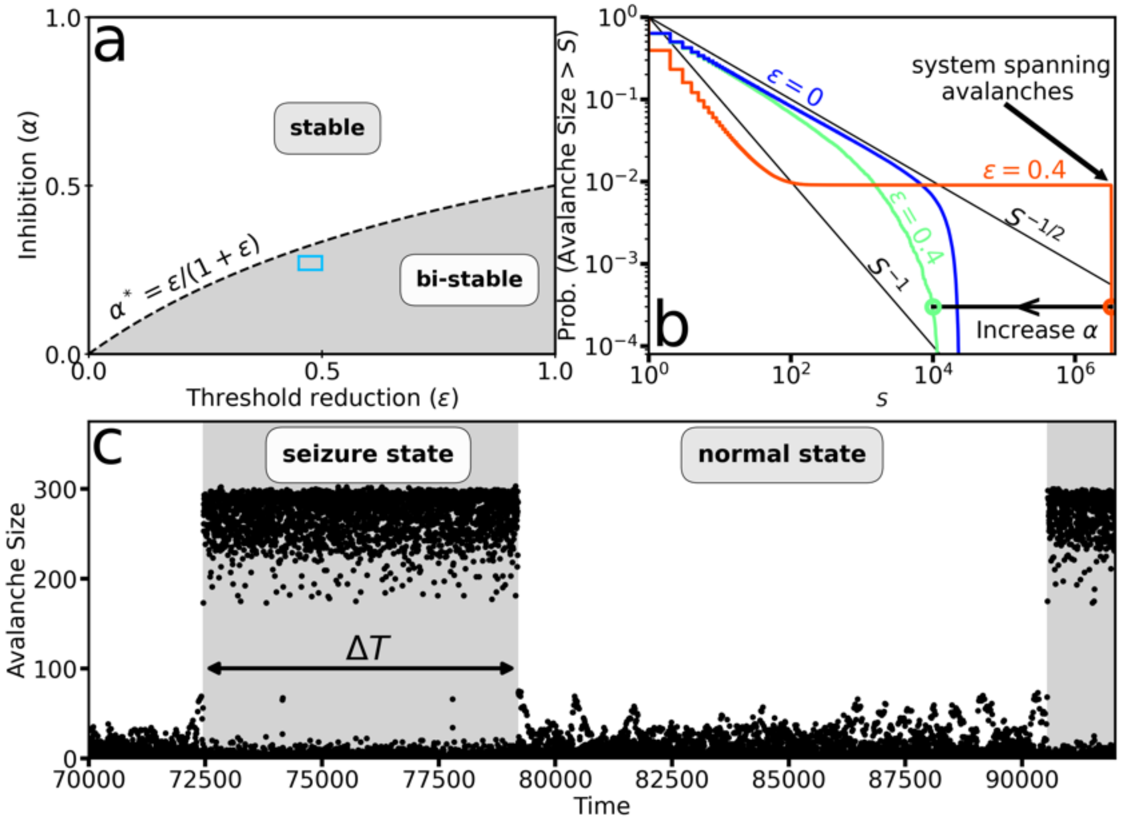

Figure. 5.2: Bi-stability between normal and seizure states in simulations. a) Phase diagram for $\alpha, \epsilon \in [0, 1]$. Phase boundary α^* (black dashed line) separates the stable and bi-stable regimes. A teal perimeter highlights the region of phase space sampled for switching statistics. **b)** Avalanche size distribution for weakened ($\epsilon > 0$) and non-weakened ($\epsilon = 0$) simulations. Black lines show the two predicted critical exponents. System spanning avalanches are evident by the sharp vertical lines in the weakened case. Simulations parameters ($N = 25000, w = 0.1, \alpha = 0.005$ [green line: $\alpha = 0.008$]). **c)** Time trace of avalanche sizes in model simulations. Seizure states are highlighted in grey and ΔT marks the seizure duration. Simulation parameters ($N = 100, w = 2/19, \epsilon = 0.46, \alpha = 0.25$).

This switching is visible in a typical time trace of avalanche sizes from model simulations (shown in Figure 5.2(c)). We will focus on the mode switching process from seizure to normal state (describing the duration of the seizure state) which has a persistence time ΔT, depending on α, ϵ, N and the characteristic time needed to load a neuron to threshold, T_0. The average persistence time is predicted by our model to have the approximate form [7]:

$$\langle \Delta T \rangle \approx \frac{T_0 N^{3/2}}{(S^s_{\max})^2} e^{(\chi+\chi^2)N/S^s_{\max}} \tag{5.1}$$

respectively, where $\chi = (\alpha^* - \alpha)/[(1 - \alpha^*)(1 - \alpha)]$ [7].

For a portion of $\{\alpha, \epsilon\}$-space in the unstable region (red box Figure 5.2(a)) we extracted the first moment of time spent in the seizure state $\langle \Delta T \rangle$ from model simulations. In Figure 5.3(a) we plot the $\langle \Delta T \rangle$ vs.

weakening, for varying values of inhibition. We see inhibition reduces the average time spent in the seizure state (i.e., $\langle \Delta T \rangle$ decreases with increasing inhibition). We also see this reduction competes with the weakening which acts to increase $\langle \Delta T \rangle$ (see Figure 5.3(a)).

As an initial comparison of the model with data, we have examined spiking signals from many hundreds of neurons in a slice culture of mouse cortex [74]. A feature of most all in-vitro spiking neuron systems is the presence of slow oscillations between a state of low activity and one of high activity, a phenomena usually referred to as bursting [57]. We extracted times spent in each of the states (see Appendix D) and constructed complementary distribution functions to compare with the model's switching times (e.g., by identifying the high activity state with the seizure state and the low activity state with the normal state, Figure 5.3(b)).

5.6 Discussion

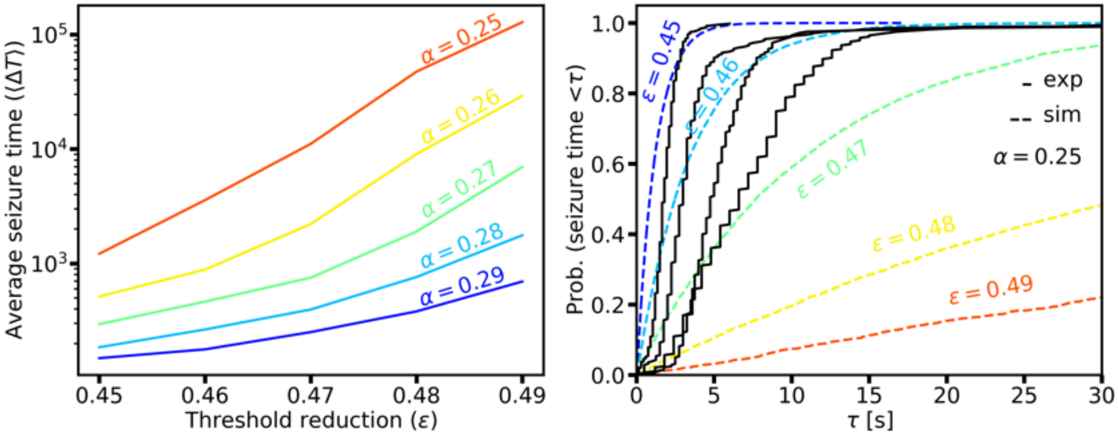

Figure. 5.3: Switching times in model simulations and experiments. a) average seizure time as a function of threshold reduction ϵ, for varying values of inhibition α. The error bars show a 95% confidence interval. **b)** Probability of seizure times being larger than τ shown for varying values of threshold reduction and fixed inhibition $\alpha = 0.25$. Three different exemplary spiking data sets are shown in black lines.

The results predict collective behavior in large systems of excitatory and inhibitory spiking neurons. Specifically, modeling both memory effects (through the dynamic threshold weakening) and inhibition (through the interaction with the inhibitory population) leads to the prediction of two possible outcomes for the neuronal dynamics. Without inhibition, the network is predicted to be in the bi-stable phase, susceptible to runaway avalanches. However, the results indicate that increasing the inhibition will tune the network into a stable phase absent of these runaway avalanches. Therefore, this model prediction can potentially help estimate the strength of inhibition needed in a network to prevent runway avalanches. Most importantly, if these runaway avalanches are related to seizures [75–78] our results predict a way to tune global inhibition (for instance through the addition of inhibiting drugs [79, 80]) to relax the epileptic system back to stability

(Figure 5.2(b)).

To understand seizures, which are temporary episodes of bursting behavior, we focus on the bistable phase. Here, the system switches between the normal state or the seizure state. In the seizure state the avalanche size distribution follows a power-law with a bump in the tail reflecting the runaway avalanches [81, 82], while in the normal state the avalanches are distributed according to a power-law with a smaller cutoff.

The switching phenomena provides a possible mechanism responsible for excessive brain activity [80, 82–84]. This is because the seizure state nucleates out of the normal state at random—that just means that if the avalanche size surpasses a critical value, it turns into a runaway avalanche, transitioning the network into the seizure state. This random occurrence would explain why it is difficult to find reliable precursors to seizures.

The nucleation of a seizure state is ultimately a consequence of the dynamic threshold weakening. The nucleation size, i.e., the critical avalanche size needed to trigger this switching, $S_{\rm cr} \sim 1-\epsilon)N/(1-\alpha)$ depends on N, α and ϵ. By measuring this size, it may be possible to estimate where a real system lies in α, ϵ-space. More importantly, by changing the various parameters through appropriate drugs it may be possible to increase the nucleation size so that seizure episodes are less likely and thus greatly suppressed.

These results show how inhibition reduces hyperactivity [75, 85]. We find that introducing more inhibition to the system will help reduce the average time spent in the seizure state (Figure 5.3(a)). Thus, increasing inhibition (α) is predicted to make seizures end more quickly. We found the effects of inhibition to be in competition with the dynamic threshold weakening—enhancing the amount of weakening (ϵ) made seizures longer, while a reduction in the weakening also reduces occurrence rate of seizures.

We compared the statistics of burst times in the experimental (in-*vitro*) data (e.g., time spent in the high or low activity states) to see if the trends resemble the predictions of our model. Our simple model has only a few limited ingredients, but we are still able to capture several features of the data. In particular we compare the high activity state in experiments to our model's seizure state. We find that the average time spent in the high activity state roughly follows a Poissonian distribution, consistent with the prediction of our model for the seizure state [7] (see Figure 5.3(b) and Appendix C). Further, comparing three separate data sets (i.e., from different brain slices on different MEA's) we find a possible explanation for differing burst statistics; that the effective memory from MEA to MEA is differing, as exemplified by Figure 5.3(b).

It's also important to note that studying these switching times (rather than avalanche properties directly) is not vulnerable to the massive problem of subsampling in neuron data sets [86]. To be clear, we mean that a subsampled data set will give the same results for switching times as the fully sampled set (since we see roughly the same burst outline even after subsampling, see Appendix C).

Our work also encourages a discussion about which critical point is controlling the system. Inhibition induced desynchronization is seen in experiments [75, 79, 80] and indeed many models suggest a synchronization transition is controlling the system [62, 78, 84, 87–89]. On the other hand, several other models suggest the critical point is that of directed percolation, describing the transition between the active and absorbing states [63, 65, 66, 81, 90, 91]. Even combinations of the two have been suggested [62, 90, 92]. We are instead suggesting a close cousin to directed percolation, the weakened-depinning transition, that resembles the problem of an elastic interface moving in a disordered environment [3] that includes dynamic threshold weakening [5]. This weakening is the key ingredient out of which the quasi-periodic runaway events emerge, whose connections to oscillations in neuronal systems warrants further investigation.

We would also like to draw some comparisons to other work. First, we note that the runaway avalanches in our model compare closely to dragon king avalanches [82] as well as giant depolarizing events [80, 83] reported in previous work. The various approaches corroborate that neuro-dynamics lie in the bistable phase.

Lastly, the question of self-organized criticality (SOC) that arises in much of the neuroscience literature [43, 57, 59, 61, 63, 68, 82, 93] can be studied in the context of domain wall motion in magnets, where global dipolar effects called "demagnetization fields" can self-organize the wall to the center of the system [70, 94, 95]. These demagnetization fields are equivalent to inhibition in our system, and an important prediction of [94] is that SOC is only realized with infinite ranged (global) weak demagnetization fields. This would suggest that real neuronal systems (i.e., those with an underlying network structure) are not expected to be self-organized critical, but rather near an ordinary critical point, and that they can be tuned closer to or farther away from criticality by tuning parameters such as inhibition strength and weakening. Future directions include extending this work beyond MFT models to include e.g., these underlying network structures.

5.7 Conclusion

We have shown that global inhibition plays the role of preventing large periodic avalanches in the collective activity of neuronal systems. Our results predict a regular (normal) and epileptic (seizure) state, and the possibility of switching from one to the other in any finite sized system. We showed that average time spent in the normal state was enhanced through inhibition (decreasing the likelihood that a seizure happens) and the seizure state was impeded by inhibition (preventing long seizures). This seizure phase is only possible when inhibition is below a critical value, suggesting that inhibition in real systems is an important ingredient for preventing epileptic activity. The model makes predictions that can be tested experimentally. For example, it predicts how the dwell time in these various dynamical phases depend on the experimental tuning parameters, such as inhibition strength and weakening.

Chapter 6

Modeling Recurrent Cortical Spiking Activity from Large Neuronal Avalanches in Mice

6.1 Description of Collaborative Team

The contents of this chapter in its entirety is published in the journal "Scientific Reports". Several authors contributed to this work; John Beggs ran the experiments and performed post-processing. Benjamin Nicholson and Christopher Meyers post processed the data from raw spikes into causal webs and performed an analysis of the cycles in these causal webs. Karina Avila wrote the simulation code and helped with preliminary modeling and data analysis. Karin Dahmen supervised my role in the project, specifically as it pertained to modeling. I ran model simulations and analyzed experiments. All authors contributed to writing the manuscript.

6.2 Abstract

A new statistical analysis of large neuronal avalanches observed in mouse and rat brain tissues reveals a substantial degree of recurrent activity and cyclic patterns of activation not seen in smaller avalanches. To explain these observations, we adapted a model of structural weakening in materials. In this model, dynamical weakening of neuron firing thresholds closely replicates experimental avalanche size distributions, firing number distributions, and patterns of cyclic activity. This agreement between model and data suggests

that a mechanism like dynamical weakening plays a key role in recurrent activity found in large neuronal avalanches. We expect these results to illuminate the causes and dynamics of large avalanches, like those seen in seizures.

6.3 Introduction

Neuronal firing signals measured from the cortex of animals and humans are known to occur in quick intermittent bursts, involving complex cascades of electrical activity [43, 45, 55–57, 68, 96–101]. Here we show that while the smaller cascades typically are not recurrent and occur temporally clustered, the largest cascades or avalanches are periodically occurring and highly recurrent. Understanding the origin of large avalanches and predicting their timing is important for preventing their potentially catastrophic impact for instance during a seizure. The hope is that understanding the underlying mechanism provides new insights for preventing serious problems such as epilepsy, which affects roughly one percent of the world's population. Possible mechanisms include pre-existing feedback loops in the underlying neuronal network, temporary strengthening of active connections between neurons, or firing threshold weakening. Here we appeal to a biological threshold weakening mechanism that has been reported [102]. We find that including this mechanism in an established model for firing avalanches qualitatively reproduces the experimental data. We also discuss connections to other models.

The results of this study are relevant not just for neurons but also for many other systems: similar collective bursts of cascading activity – or "avalanches"– occur in a broad range of experimental and natural systems, including magnets, fluids in porous media, and earthquakes [5, 24, 103]. The power-law scaling behavior of avalanche statistics and dynamics can often be described in the framework of second order phase transitions [24]. Since the smaller avalanches in healthy brain tissue are often also power-law distributed, the brain has been conjectured to operate near a critical point [43]. In this work, we expect the weakening parameter to increase the avalanche size scaling exponent from 1/2 to 1, describing the universality class of de-pinning with dynamical weakening [5]. We will refer these avalanches in the power-law scaling regime of the size distribution as "small avalanches", and those that are larger than the scaling regime (i.e., in the exponential tail of the distribution) as "large avalanche". Several statistical and dynamical model predictions have been investigated through experiments [45], but the interplay between the dynamically changing neural network and the avalanches and especially the dynamics of catastrophically large avalanches, have yet to be fully understood.

Electrical activity in neurons travels along networks defined by neuronal connectivity. Broadly speaking, activity can propagate through a network of interconnected neuronal nodes in either a feedforward or in a feedback manner. In feedforward propagation, activity visits each neuron only once in a given cascade. Feedforward cascades are relatively brief in duration, typically stable, and can be used to perform computations that rapidly converge [104]. In feedback propagation, some nodes are reactivated, possibly many times, during a cascade. Feedback cascades can be relatively long in duration, can amplify and become unstable, and can be used to perform iterative computations that take longer to converge [105]. This type of propagation could easily be harmful as well, since it can promote self-generated activity that may not be desirable. Some experiments have found recurrent activity in neurons [75, 85, 106–108], usually as the driving force behind self-generated sustained activity, where it was seen to arise during seizures in rats [75]. It was also shown to generate slow oscillations in cat neocortex [106] and suggested to play a role in synchronization in Guinea pig hippocampal slices [85]. Despite this past work, limitations of the spatial and/or temporal resolution of most methods for recording the activity of a sufficiently large number of neurons has prevented recurrent activity within neuronal avalanches from being experimentally resolved. Here we leverage a novel technique [44] to reexamine neuronal firing data [45, 74], allowing us to disentangle neuronal avalanches into causally connected webs of firing activity. This affords us a high enough spatial and temporal resolution to reliably (see Appendix E, sec. 2) characterize signatures of recurrent activity for the first time within avalanches, both qualitatively and quantitatively.

We note here that firing frequency adaptation [109] is normally thought to suppress recurrent firing (within the timescale T). However, the recurrent firing seen in the data suggests additional mechanisms are at play. A mechanism that weakens thresholds dynamically (and temporarily) is known to affect the dynamics of many systems such as earthquakes [110], shear deformation in materials [111] and even the formation of river basins [112, 113]. Here we invoke an analogous neuronal weakening mechanism for three main reasons: (1) Inhibitory neurons have a regime where their firing thresholds decrease as they are driven more strongly [102]. In addition, the hyperpolarizing effects of inhibition have been shown to lower firing thresholds and reduce refractory periods of excitatory neurons [114]. (2) The effects of a weakening mechanism are equivalent to the effects of temporary coupling strength increase immediately after a firing event [5]. Neuronal data support this idea as synapses onto inhibitory neurons in the cortex have short-term facilitation [115]. Thus, the literature suggests a biological tuning parameter that controls recurrent (such as "epileptic" [75]) activity in neuronal networks. The new parameter could be either temporary threshold weakening or, equivalently, a coupling strength increase. (3) Here we use the language of weakening simply to connect the description to a much broader class of systems with similar recurrence effects, including earthquakes and brittle behavior, where weakening is the more natural description. Yet, both mechanisms equally promote recurrent activity within large avalanches because they are analogous to each other, and both interpretations can be used interchangeably.

6.4 Methods

6.4.1 Data from microelectrode arrays

The data that we analyze here were collected using a 512-microelectrode array and previously published [45, 74]. This array has an electrode spacing (60 μm) and sampling rate (20 kHz) that let us detect individual neuron firing events and track their transmission within a densely connected cortical network. The tissue we used in these experiments came from slices of mouse and rat cortex that were cultured for two weeks until they showed mature patterns of neuronal activity. These patterns allow for extraction of avalanche size and avalanche duration distributions. For several of these samples, those distributions can be approximated by power-laws, indicating that these networks may be operating near a critical point [56]. We thus focused our attention on those samples that reflect the most relevant state seen in-*vivo* for active and anesthetized animals [116–118].

6.4.2 Causal webs

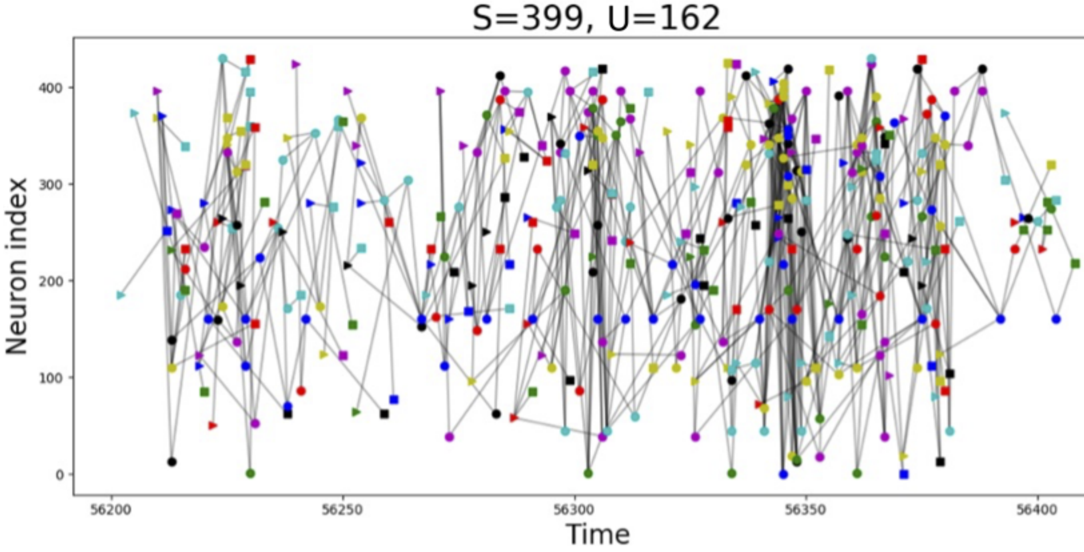

Figure. 6.1: Causal webs. Time trace of a recurrent avalanche in mouse. For a single large avalanche with size $S = 399$ and number of neurons $U = 162$, the causal web of firing activity. Individual neurons are labeled by an arbitrary neuron index (y-axis, not all neurons participate in this avalanche), which spike at particular times (x-axis, in ms since start of recording). The causal web is a directed graph whose nodes are indexed by the pair (neuron, time). Nodes with no incoming edges are indicated by right-pointing triangles, nodes with no outgoing edges are indicated by squares, and nodes with both incoming and outgoing edges are indicated by circles. Patterns of recurrent neural activity are seen as particular neurons spiking repeatedly during the avalanche, leading to horizontal trace of like-colored symbols.

From firing activity, we need to infer neuronal avalanches. Most previous work has relied on temporal binning of the data, to identify temporally connected bins of activity [43]. One problem with this method is that it can merge avalanches that occur at the same time but at disconnected locations. In this work, we have used instead a recently developed method of "causal webs" [44]. Briefly, this assumes that avalanches can only propagate along connections between neurons and at delays that are consistent with those connections. We used transfer entropy [119] to estimate statistically significant functional connections between neurons and their corresponding delays [120–123]. We then mapped the observed network activity onto avalanches that were consistent with the connections and delays we found previously in the network. Activity that could not be accounted for by these connections and delays was excluded from avalanches. From the raw firing data, we extracted causal webs for several mouse and rat neural tissue recordings (a single avalanche is shown in Figure 6.1). Generally, each recording consisted of many tens to hundreds of thousands of avalanches involving hundreds of neurons. Not all samples show evidence of power-law distributed avalanche statistics, so we use such statistics as a quality control filter, resulting in six mouse samples and two rat samples which we use for further analysis (see Appendix E, Figure E.1).

6.4.3 Bi-Test

The Bi-test extracts temporal information about avalanches in the form of a distribution of interevent times [124]. These times can be broadly classified as either clustered or periodic: information that can suggest different mechanisms underlying recurrent avalanche activity. Given a set of M avalanches one can observe the set of starting times $T = t_{\mathrm{i}}, i = 1, \ldots, M$ to construct the values $H_{\mathrm{i}} = \Delta t_{\mathrm{i}}/(\Delta t_{\mathrm{i}} + \Delta\tau_{\mathrm{i}}/2)$. Here Δt_{i} is taken as the minimal interevent time, $\min(|t_{\mathrm{i}} - t_{\text{i-1}}|, |t_{\mathrm{i}+1} - t_{\mathrm{i}}|)$, and $\Delta\tau_{\mathrm{i}}$ is the subsequent interevent time ($|t_{\mathrm{i}-1} - t_{\mathrm{i}-2}|$ or $|t_{\mathrm{i}+2} - t_{\mathrm{i}+1}|$). Once $H = \{H_{\mathrm{i}}\}$ is found, the form of the temporal correlations can be inferred from the shape of the cumulative distribution function (CDF) of H (see Figure E.13) [124, 125].

6.4.4 Model

The broad model describing systems like earthquakes and magnets [5, 24, 103] is adapted to include not only excitatory (i.e. positively coupled), but inhibitory (i.e. negatively coupled) cells as well, to better represent real neuronal systems. We further replace the uniform coupling constant J by a weight matrix J_{ij} with an all-to-all topology and the strengths randomly distributed. This represents the connections between a pre (i) and post (j) synaptic neurons. In view of experimentally observed functional connectivity [122, 126, 127], the elements of this matrix are sampled from a log-normal distribution

$$\mathcal{P}(J_{\mathrm{ij}}) = \frac{1}{J_{\mathrm{ij}}\sqrt{2\pi N}} e^{\frac{-\ln(J_{\mathrm{ij}})^2}{2N}} \tag{6.1}$$

Simulations start by defining for all neurons an arrest potential $V_{\mathrm{i,arrest}}$ that is sampled from a uniform

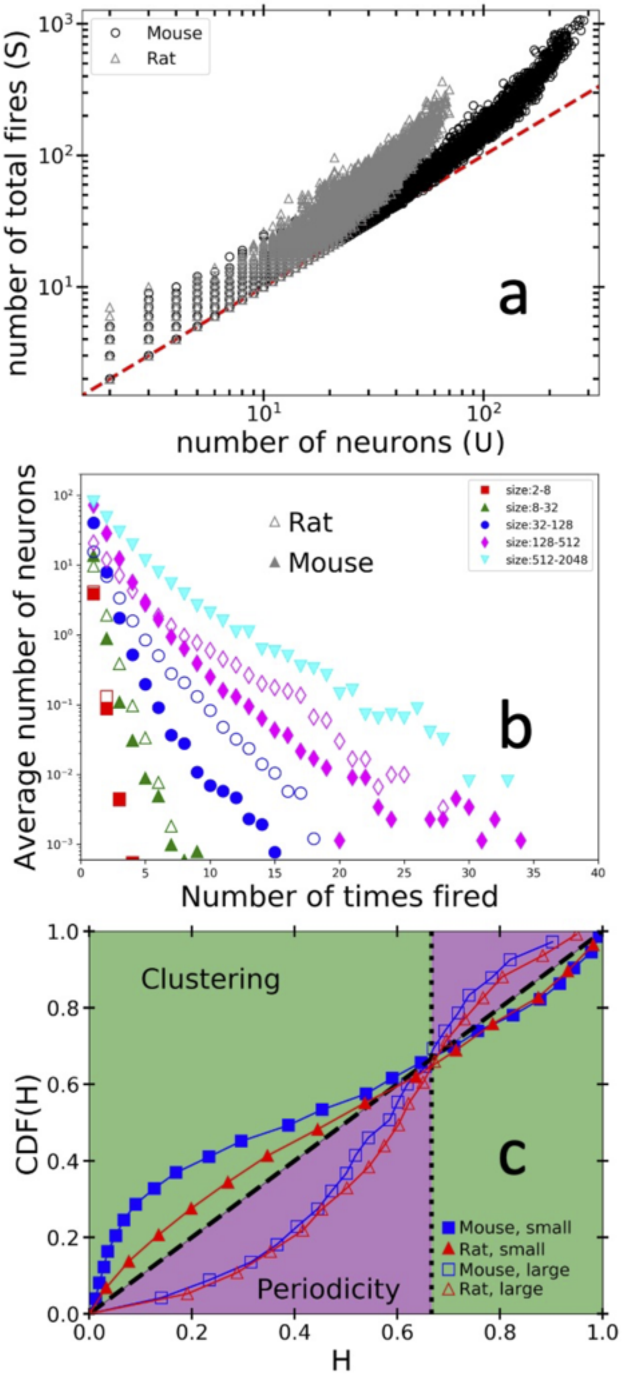

Figure. 6.2: Recurrent activity in mouse cortical neurons. a) Scatter-plot of avalanche size S as a function of the number of individual neurons U participating in each avalanche, as estimated by the causal-webs method [44]. For the mouse sample 434 neurons participated in 86936 avalanches. For rat 106 neurons participated in 77648 avalanches. **b)** Distribution of the mean number of times each neuron spikes in an avalanche, for avalanches within specified size bins. Note the rat data had no avalanches in the largest (size=512-2048) bin. **c)** Avalanche interevent time correlations. The cumulative distribution function of $H_\mathrm{i} = \Delta t_\mathrm{i}/(\Delta t_\mathrm{i} + \Delta\tau_\mathrm{i}/2)$, where Δt_i is taken as the minimal interevent time, $\min(|t_\mathrm{i} - t_{\mathrm{i}-1}|, |t_{\mathrm{i}+1} - t_\mathrm{i}|)$, and $\Delta\tau_\mathrm{i}$ is the subsequent interevent time ($|t_{\mathrm{i}-1} - t_{\mathrm{i}-2}|$ or $|t_{\mathrm{i}+2} - t_{\mathrm{i}+1}|$). Poisson distributed random interevent times fall directly on the dashed line, while perfectly periodic interevent times fall on the dotted line. The bi-test in the green regime is a signature of temporal clustering while in the purple regime is a signature of temporal quasi-periodicity.

distribution on the interval $[-w/2, w/2]$. We initialize all N neurons with a random potential value

$$V_{\mathrm{i}} = V_{\mathrm{i,arrest}} + \xi \Delta V_{\mathrm{i}} \tag{6.2}$$

where $0 \leq \xi \leq 1$ is uniformly distributed and for simplicity we have defined $\Delta V_{\mathrm{i}} = V_{\mathrm{thresh}} - V_{\mathrm{i,arrest}}$. This simply initializes the system in a mixed up state so the steady state is reached quickly. Next, the inclusion of inhibitory neurons amounts to designating N_{exc} neurons to be excitatory and N_{inh} neurons to be inhibitory ($N = N_{\mathrm{exc}} + N_{\mathrm{inh}}$).

When an inhibitory neuron i spikes, the synaptic connections with all other excitatory type neurons j are less than zero ($J_{\mathrm{ij}} < 0$) so that inhibitory neurons do indeed inhibit excitatory neurons from firing. We use four times as many excitatory neurons as inhibitory ($N_{\mathrm{exc}} = 4N_{\mathrm{inh}}$) as typically reported in cortical tissue [128].

To start an avalanche, we raise the potential V_{i} of all N neurons by the amount $V_{\mathrm{thresh}} - \max(V_{\mathrm{i}})$ such that the "origin" neuron i which is closest to the firing threshold $V_{\mathrm{thresh}} = 1$ is caused to spike. When this neuron i spikes, the simulation resets its potential to a resting value $V_{\mathrm{i}} \to V_{\mathrm{i}} - \Delta V_{\mathrm{i}}$. In the subsequent time step (we have checked that delaying this does not change our results, see Appendix E, sec. 4) the firing event from neuron i then contributes to all other $N - 1$ neurons the potential $\Delta V_{\mathrm{i}}/(N-1)$.

The new failure threshold for an already spiked neuron is weakened to a dynamical value $V_{\mathrm{dynamical}} = V_{\mathrm{thresh}} - \epsilon_{\mathrm{E/I}} \Delta V_{\mathrm{i}}$, with $0 \leq \epsilon_{\mathrm{E/I}} \leq 1$. Note that this parameter is not necessarily the same for excitatory neurons as it is for inhibitory ones. Thresholds for the firing of a neuron during an avalanche are diminished due to this weakening process but are reset to their original default after an avalanche terminates. Thus, the model assumption relies on a separation of time scales between intra-avalanche and inter-avalanche firing (verified in the data, see Appendix E sec.3) and assumes that some unspecified relaxation process allows for the resetting of firing thresholds on sufficiently long timescales after an avalanche terminates. As discussed below the separation of time scales does not have to be perfectly fulfilled to observe the predicted trends.

Avalanches emerge in this model when the origin firing event triggers one or more firing events in a cascading fashion. We identify the avalanche size S to be the number of total firing events, and U as the total number of neurons involved. All simulations were performed for $200,000$ time steps each. To match both mouse and rat data sets we simulated system sizes from 450 to 2000 neurons and disorder strength $w = 1.9$.

6.4.5 Sub-sampling

In the experiment, the cortical slices are laid atop a 1 by 2 mm rectangular electrode array. Because only roughly one quarter of the total slice is being recorded by the MEA, there is chance that neurons outside of the field of view affect the recorded neurons i.e., by causing one of them to spike. To emulate this in our simulations, we have up-sized our system by a factor of four and then analyzed only one fourth of the neurons when constructing avalanches. We outline the sub-sampling procedure in more detail in Appendix E. We find that sub-sampling the data has the effect of increasing both the level of recurrent activity as well as the robustness of the clustered bi-test signature for small avalanches. In the Appendix E, Figure E.9, we show that the sub-sampling alone cannot reproduce the data, and that weakening is indeed necessary for our simulations to reproduce the data.

6.5 Results

6.5.1 S vs U: evidence of recurrent activity

Having reconstructed avalanches, we examine the relationship between the overall size S of an avalanche (total number of firing events) and the number of unique neurons U firing during that avalanche. For both mouse and rat samples, Figure 6.2(a), shows the relationship between the size S and number U participating in each avalanche. Smaller avalanches tend to involve approximately only a single firing event per neuron during an avalanche ($S \approx U$), whereas the largest avalanches in this recording involve some neurons firing more than once ($S \gg U$). Note that the mouse sample (black) shows more recurrent activity than the rat sample (S/U greater for the largest avalanches). Appendix E shows similar plots for other samples and recordings (Figure E.2).

6.5.2 Distribution of recurrent activity

We examine further the nature of these repeated firing events by considering the distribution of the number of times each neuron spikes during an avalanche. Representative data are shown in Figure 6.2(b), for the rat and mouse samples, where we bin avalanches by their total size S and compute the mean number of neurons $F(\mathrm{m})$ that spike m times during an avalanche. For any given avalanche bin (range of avalanche sizes), $F(\mathrm{m})$ falls off roughly exponentially, but for larger avalanches, this exponential decay becomes slower and slower, with some neurons firing dozens of times in a single avalanche. In some samples (see Appendix E Figure E.3), some neurons spike hundreds of times during an avalanche.

6.5.3 Cycles: the nature of recurrent dynamics

While the distributions shown in Figure 6.2(b) demonstrate that some neurons spike many times during an avalanche, they do not reveal the structure of these avalanches or the role played by neurons firing many

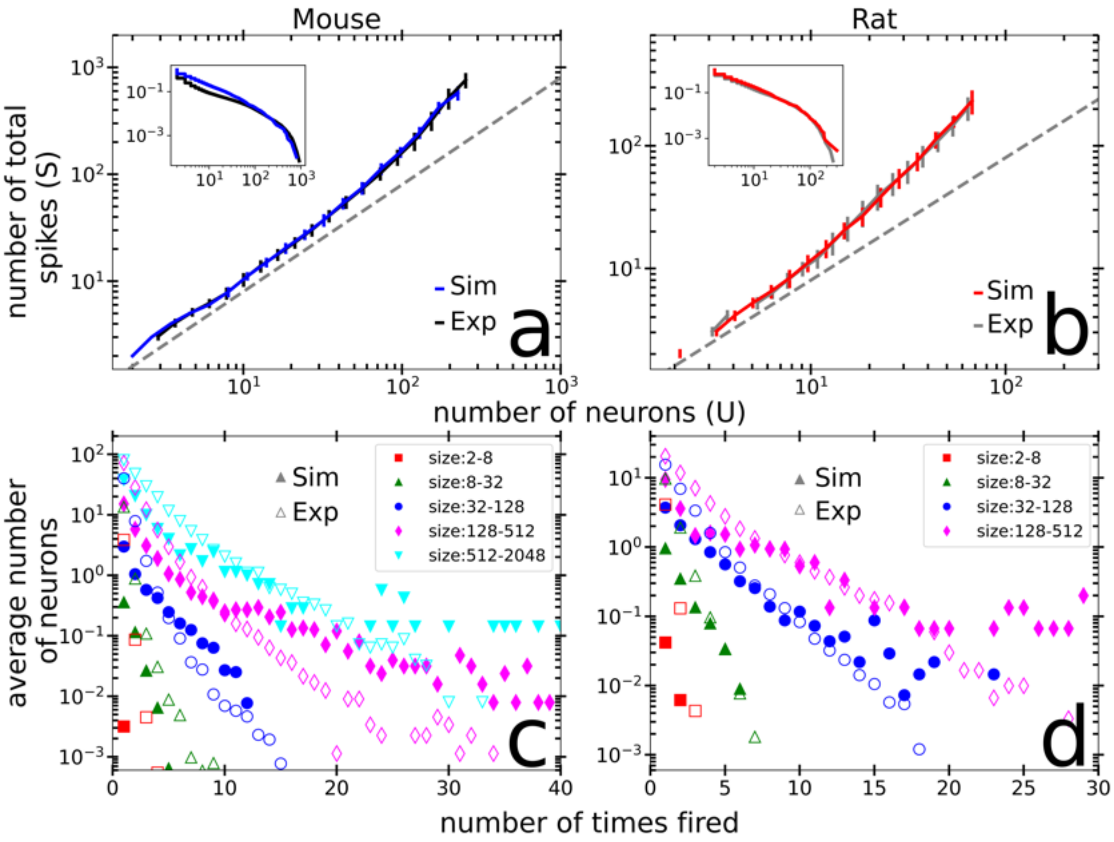

Figure. 6.3: Model comparison to data. Causal web analysis of neuronal avalanches in both mouse (black line, **a)** and rat (grey line, **b)** are shown in comparison to model simulations (blue and red lines, respectively). Lines represent the average value of size, S, within bins of the number of neurons, N. We use twenty logarithmically spaced bins and error bars show the standard deviation of the bin. Using $\epsilon_e = 0.5$ $\epsilon_i = 0.8$ the model captures the recurrent activity and matches the S vs. U distribution in the mouse. Adjusting the number of neurons in the simulation and the weakening parameters ($\epsilon_e = 0.6$ $\epsilon_i = 0.7$) the model fits the rat distribution as well. We reduced the inhibitory-excitatory coupling strengths by a factor of $\gamma = 4.8$ and $\gamma = 2.5$ for the mouse and rat simulations, respectively. The complement of the cumulative distribution for avalanches size S ($\mathcal{C}(S)$) is shown in the inset compared to the simulated distribution, which qualitatively (power-law) and quantitatively (-1 slope) match the experiments, for mouse and rat both. Note that the avalanche sizes involved in the mouse experiments were larger than those in the rat, so the axes in the right span a smaller range for qualitative comparison of the distributions in different species. Mouse **c)** and rat **d)** neurons spike multiple times in an avalanche. The model simulations (empty markers) qualitatively follow the broadening of the experimental distributions (filled markers) signifying recurrent activity. Importantly, the simulations also broaden more for larger size avalanches, which is seen in the experiments facilitates the bending in Figure. 6.1. No avalanches existed in the largest size bin ($128 - 512$) for the rat, and the axes are different between left and right for better qualitative comparison. All simulations were run for $200,000$ time steps with a disorder strength, $w = 1.9$.

times. Are large avalanches characterized by many spontaneous (untriggered) firing events that propagate through the system in a feedforward manner, or are there cycles of feedback in which activity loops back to trigger previously active neurons? To address this question, we have probed the net- work structure of these neuronal avalanches to identify cycles of activity. Each avalanche consists of a directed graph whose nodes are particular neurons firing at particular times (i.e., (neuron, time) pairs), and whose edges connect nodes if there is an inferred causal link from a source node to a destination node (i.e., source contributes to the firing of destination). A time trace of one such large avalanche with recurrent activity is shown in Figure 6.1. For our purposes here, a cycle is any directed path in an avalanche on which a given neuron spikes at some initial time and then spikes again at one or more subsequent times on the same directed path, implying a causal path connecting a neuron to itself at a later time.

In general, we find a complex, interwoven set of cycles, with multiple paths of activity leading into and out of recurrent loops, shorter cycles sometimes embedded in longer cycles, and injections of new activity from spontaneous activation events with no apparent triggering event. Characterizing the detailed structure of these cyclic patterns of activity is beyond the scope of the current work, but we can summarize the statistics of cycles as shown in Figure 6.4, for both the mouse (a) and rat (b) samples examined previously. In particular, for any avalanche, we have calculated what fraction of (neuron, time) nodes in a causal web lie on at least one cycle, as defined above. Most small avalanches are completely feedforward with no cycles (the lower branch for small S), although some of the small avalanches do involve small cycles among a few neurons (the upper branch). As avalanches grow larger, they tend to have a larger fraction of firing events occurring within these cycles, and the mouse sample reveals a greater fraction of activity taking place within recur- rent cycles for the largest avalanches. Interestingly, recent work examined the same data set [129], and although defining "cycles" in a different manner (more akin to our definition of "recurrence"), showed in simulations of networks that cascade duration increased with the number of embedded cycles, thus preserving information longer. In our analysis, working back from the data, we see that larger avalanches are associated not just with a greater amount of recurrent activity, but also a greater amount of cyclical/feedback activity, as we have defined it.

6.5.4 Evidence of weakening

The Bi test extracts interevent time correlations that are used to investigate the presence of dynamic weakening. Here we use the Bi-test to show that the addition of dynamical weakening to our model produces a specific temporal signature: quasi-periodic interevent times between the large avalanches in the tail of the distribution Figure 6.2,(c) (see [125]). We note that a further signature is predicted; clustered interevent times between the small avalanches in the scaling regime, for our bare MFT model. We find, however, that using both inhibitory neurons as well as increased disorder strength causes this temporal signature to vanish. We outline in Appendix E when the clustered signature does and does not show up in our model simulations (Figure E.10) and show how sub-sampling in the data could be a reason that we see bi-test signatures in both small and large avalanches in the data (See Figure E.11 in Appendix E). See Figure E.5 for results of the Bi-test on other samples.

6.5.5 Model fit to S vs U

Adding dynamic weakening to the model enhances the possibility of recurrent activity, by successively lowering thresholds to activation and facilitating neuronal firing. We find this activity in both mouse and rat cortex, which manifests as an upward bending of the S vs. U curve away from the purely feedforward line (Figure 6.2(a)). Model simulations without weakening produce S vs. U curves that follow this feedforward line (see Figure E.14, bottom, blue).Those with weakening however produce curves bent above the feedforward line (Figure E.14, bottom, red). We allow the inhibitory and excitatory weakening values to differ and find the optimal parameters for mouse and rat to be ($\epsilon_{\mathrm{E}} = 0.5$, $\epsilon_{\mathrm{I}} = 0.8$) and ($\epsilon_{\mathrm{E}} = 0.6$, $\epsilon_{\mathrm{I}} = 0.7$), respectively. With these parameters alone, we find the S vs. U distributions bend, although there are larger avalanches in the experiment, and thus the tails of the scatter plot do not match (See Figure E.14, top, red).

By reducing the coupling strength between excitatory and inhibitory neurons by a factor of γ (Figure E.14, top, green), we effectively enhanced the network excitability, leading to larger avalanches for a better fit to the tail of the experimental scatter plot Figure 6.3,(a-b). For the mouse and rat data, we found good agreement with $\gamma = 4.8$ and $\gamma = 2.5$, respectively.

6.5.6 Model fit to firing distribution

With the addition of dynamic weakening, the model is able to qualitatively capture the phenomena of multiple neuronal firing. To further investigate the structure of recurrent firing we have constructed the distribution of single neuron firing times for every avalanche which we then averaged together using bins in size, analogous to the data presented in Figure 6.2(b). As with the experimental data, we see that larger avalanches in the model are associated with a widening of the distribution of neural firing (i.e., neurons are more likely to spike repeatedly). With these parameters, however, the simulations somewhat overestimate the probabilities in the tails of the distributions (Figure 6.3,(c-d)). We note that this widening of the distribution of neural firing in our model is very slight unless we add weakening (see Appendix E, Figure E.15).

6.5.7 Recurrent cycles in model simulations

We analyzed the dynamical trajectories generated by the simulation model to assess the presence of cycles, as for the experimental data. In the model, a given set of neurons spike at one discrete time step, leading to the firing of a different set of neurons firing at the next time step. Therefore the “causal webs” that we infer from the model trajectories are more densely connected than those we infer from the experimental data, since each neuron firing at one time step contributes in principle to the activity of all firing neurons at the next time step. Nonetheless, we have computed the fraction of model firing events that lie on a cycle as a function of avalanche size, and plotted those results in Figure 6.4,(c-d) for mouse and rat parameters. These results on top (experiments) are qualitatively similar to the analogous plots on bottom (model simulations), revealing the merging of distinct feed-forward and recurrent branches of activity for small avalanches, and an increase in cyclic activity as avalanches grow larger.

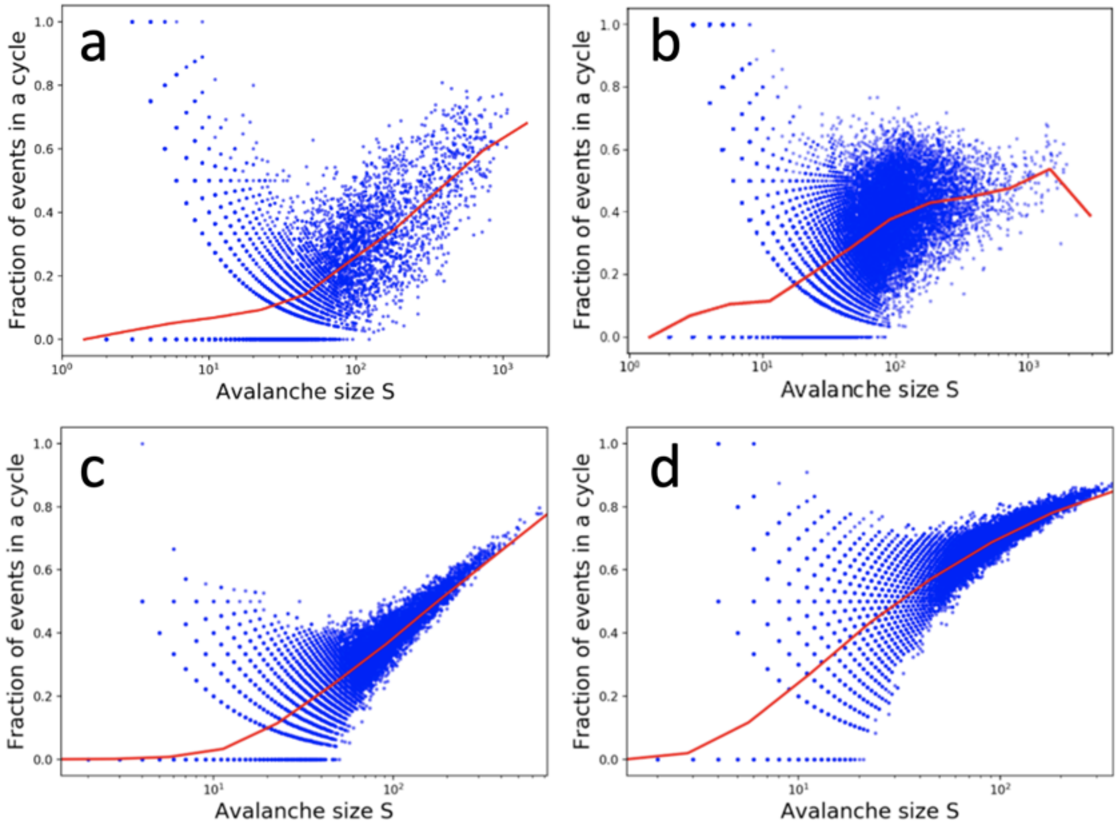

Figure. 6.4: Cycles in recurrent avalanches. Fraction of firing events that lie on a cycle of activity, as a function of overall avalanche size S, for mouse experiments **a)**, rat experiments **b)**, mouse simulations **c)** and rat simulations **d)**. By detecting cyclical activity of firing events on directed causal paths (see text), all firing events are classified as either lying on a cycle or not. For each model simulated avalanche, the fraction of firing events in a cycle for that avalanche is plotted as a function of avalanche size. Large avalanches are characterized by more recurrent firing, and a greater fraction of firing lying on cycles.

6.6 Discussion

Here we have demonstrated that dynamic threshold weakening can induce recurrent firing patterns in cortical neurons. This dynamic reduction of firing thresholds is reminiscent of structural weakening in materials models [5] and has been inferred from parameter estimation algorithms applied to current clamp recordings of neurons [102]. Importantly we note that in addition to dynamic threshold weakening, there is also evidence for a firing frequency adaptation mechanism [102]. Since recurrent activity emerges at the onset of a seizure [75] we could use this effect, whatever its origins, to understand and combat dangerous pathologies such as epileptic activity.

These basic patterns were observed in multiple tissue samples, from both rat and mouse. The similarity of these two example cases suggests much broader applicability of the results. For both mouse and rat data, the same model qualitatively reproduces the bending of the S vs U curves (Figure 6.3,(a-b)), the broadening of the distribution of multiple neuron firing events (Figure 6.3,(c-d)) and the propagation of activity around cycles (Figure 6.4,(c-d)). Note that the relationship between avalanche size and the time interval between avalanches (previously studied in [68]), is separate from the relationship between avalanche size and the intervals of recurrent firing events within them (here called cycles). With very minimal tuning of the weakening parameter between species, the model comparison can even be quantitative (e.g., it predicts the maximum values of avalanche sizes for a given network size, maximum number of times one neuron will spike in a given avalanche, etc.). This robust confirmation across data sets indicates that this weakening phenomenon is likely a general feature of neuronal avalanches in cortex.

Interestingly, other systems in nature also display weakening effects, which might help provide intuition about the dynamics of neuronal systems. Weakening has been observed in a wide variety of systems, and it is often associated with dynamical instabilities, spatio-temporal pattern formation, and spatial localization. Velocity-weakening friction in models of earthquakes leads to localized pulses of slip that propagate along faults [110]. Shear localization in materials involves spatially concentrated bands of shear deformations triggered by weakening [111]. Formation of river basins proceeds through a feedback process of canalization whereby locally increased water flow results in increased soil erosion, thereby supporting greater flow through emergent channels [112, 113]. We therefore speculate that weakening in neuronal networks might show similar dynamics, with channels of recurrent neural activity carved out by the avalanche dynamics itself, a scenario that can be studied in future work.

To strengthen the model, different biologically inspired amendments could be considered. For instance, the topology of neuronal systems is not all-to-all as we have assumed in this work for simplicity. Whether large avalanches and their recurrent activity are affected by topology in the model presented here could be the subject of some interesting future work. Further, our instant re-healing of a firing threshold after an avalanche terminates was ultimately used to limit the timescales introduced in the problem, but could instead in the future be modeled by a slower re-healing process. We lastly point out that results in-*vivo*, where the separation of timescales assumed here is not likely to be true, have never the less been captured by models that require such a separation in timescales [56, 100, 101, 116], supporting the claim that our results may too be applicable in-*vivo*.

6.7 Conclusion

Here we have shown that the larger avalanches in neuronal networks involve distributions of recurrent electrical activity, with multiple feedback cycles permeating the avalanche dynamics. We have demonstrated that a dynamic weakening of thresholds for neuronal activation is a plausible explanation for the experimental neuronal firing data.

Our model was purposefully designed as simply as possible to reproduce the trends in the data. That way we could extract the key components of the underlying mechanism for the recurrent behavior, and make predictions for universal aspects of the dynamics, that do not depend on the microscopic details. Future work could include the microscopic details of specific neuronal systems and refine the model for detail dependent predictions.

Our simple coarse-grained model with weakening reproduces patterns of recurrent activity observed in the data without the need to include intrinsic neural dynamics. A Bi-test applied to the data, and the fits of the avalanche model with weakening, both support the notion of a dynamical weakening mechanism. This weakening mechanism leads to the recurrent operation of neurons and the formation of large "runaway" neuronal avalanches. For large weakening these runaway avalanches sweep most of the neurons – reminiscent of pathologies such as epileptic seizures.

Chapter 7

Exploring Dual Species Neuronal Avalanches

7.1 Introduction

This chapter will take a closer look at the data and model of the previous chapter, this time with information not only about which neuron spiked and at what time, but also whether a neuron is inhibitory or excitatory. We seek answers to two questions: 1) Could we differentiate between the two types of neurons within an avalanche? The fraction of neurons that are activated within an avalanche could be compared between the two species, for instance. Then, we could even compare the recurrent activity in an avalanche between the two species. 2) Piggy backing on this question, if we looked at a system with one ratio of number of inhibitory to number of total neurons (N_{inh}/N, we'll refer to this as the 'inhibitory ratio' from now on), how would we expect avalanche properties– size, duration, averaged temporal profiles, etc.,– look as compared to a system with another inhibitory ratio? We'll first tease out signatures in the data that could help us to answer these questions, and then we compare them with the model simulations. We will find that both are highly sensitive to the inhibitory ratio, with the inhibitory neurons making up most of the avalanche activity.

7.2 Experiments

In this section, we reinvestigate the experiments from Chapter 6 with a valuable new piece of information– each neuron in the data set is now labeled as 'excitatory', 'inhibitory', or 'unidentified'. We'll first use this information to look for correlations between various avalanche properties and the percentage of inhibitory neurons in the sample. Next, we'll take a closer look at the internal structure of neuronal avalanches by comparing the abundance of inhibitory activity relative to the excitatory activity.

7.2.1 Sample to Sample Variation

We start by making some definitions. Since there are now two species of neurons, we define the number of excitatory (inhibitory) neurons as N_{inh} (N_{exc}). We define the total number of inhibitory and excitatory spikes with S_{inh} and S_{exc}, respectively, as well as the number of unique excitatory (inhibitory) neurons active in an avalanche, U_{exc} (U_{inh}). In Figure 7.1(a) we show an example of a cartoon avalanche before and after species labeling).

For the sample focused on in Chapter 6, of the $N = 434$ neurons, 95 were identified as inhibitory, 247 identified as excitatory, and 92 were not able to be identified as either inhibitory or excitatory. This of course is inconsistent with the estimated ratio used for simulations in Chapter 6 (20%) as it severely underestimates the prevalence of inhibitory neurons in the sample. In fact, according to the species labels, not one of the 14 examined data sets had an inhibitory ratio as low as 20% Figure 7.1(b). We do, however, point out that the inhibitory ratio roughly decreases with increasing system size (i.e., number of total neurons, N), hinting that the discrepancy between previously reported ratios [128] and those found here might resolve with larger sample sizes.

What's even more obscure is the first moment of the avalanche size distribution has no visible correlation with system size here. This is in stark contrast to the expectation from our simulations without inhibitory neurons (c.f. Figure 2.5(d)), where we see instead see the first moment of the avalanche size distribution scales with system size, N. Here, the average avalanche size does, however, roughly increase exponentially with increasing inhibitory ratio, though there are a handful of outliers to this trend (Figure 7.1(c)). Similarly, the average amount of recurrent activity in a sample, $\langle S/U \rangle$, roughly increases exponentially with increasing inhibitory ratio, with no obvious trend in system size (Figure 7.1(c), inset). In summary, we observe that the inhibitory ratio changes from sample to sample with no obviously predictable consequences on the reflected avalanche properties.

7.2.2 Inhibitory vs. Excitatory

Now we take a closer look at the internal structure of the neuronal avalanches in hopes of clarifying the aforementioned observations. Since the avalanches are composed of two different species, we first look at what fraction of neurons from each underlying species are activated in an avalanche. In Figure 7.2(a) we plot the fraction of activated excitatory (inhibitory) neurons, $U_{\text{exc}}/N_{\text{exc}}$ ($U_{\text{inh}}/N_{\text{inh}}$) as a function of the total

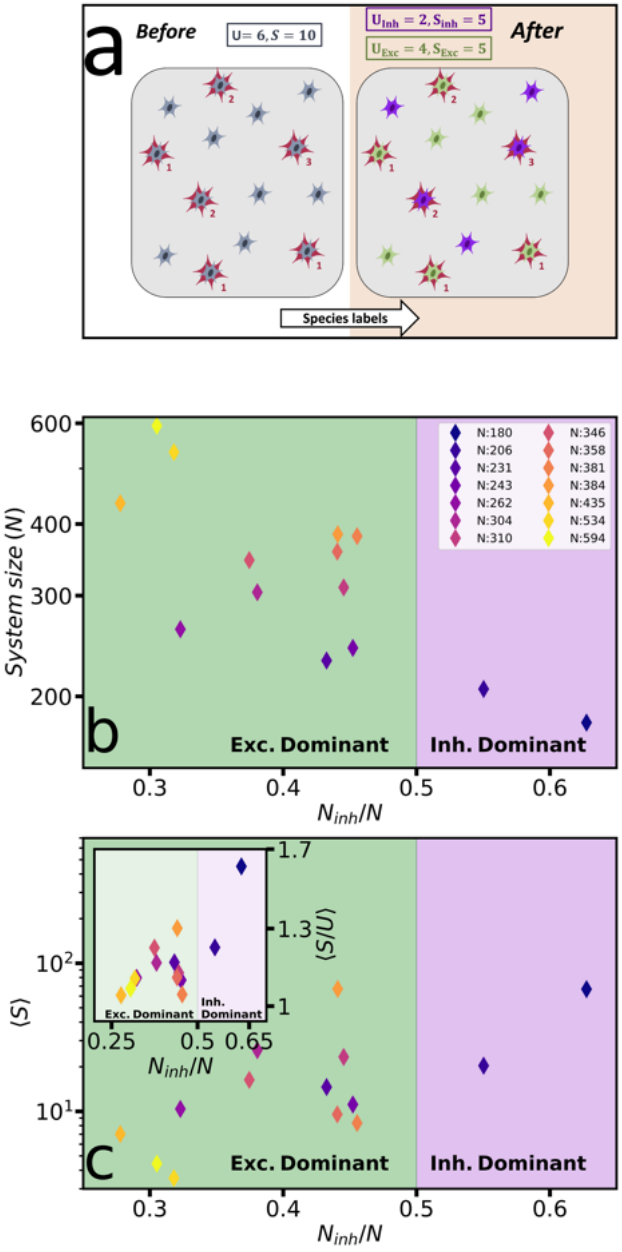

Figure. 7.1: Differentiating between inhibitory and excitatory neurons. a) A schematic of avalanche characterization with species (inhibitory or excitatory). A red outline indicates that a neuron has spiked, and a label 'n'' indicates that the neuron spiked n times. (Left) The size is the number of total spikes ($S = 10$) and the number of unique neurons is ($U = 6$). (Right) Avalanches now are characterized by the species differentiated number of active neurons (U_{inh} and U_{exc}) i.e., the number of neurons that spiked at least once, as well as the species differentiated avalanche size (S_{inh} and U_{exc}). **b)** The system size is plotted as a function of the percentage of inhibitory neurons, (N_{inh}/N), for all fifteen samples. Marker color corresponds to system size, and we have highlighted the background of the figure in green where the inhibitory ratio is less than 50% (excitatory dominant) and purple where the inhibitory ratio is greater than 50% (inhibitory dominant). **c)** The first moment of the avalanche size distribution, $\langle S \rangle$, plotted as a function of percentage of inhibitory neurons. Markers and background are colored the same as in b). Error bars show a 95% confidence interval.

level of recurrent activity in the avalanche (S/U). On average, around four percent of excitatory neurons are active in a feedforward (i.e., $S/U = 1$) avalanche and this percentage monotonically increases until saturating for the most recurrent avalanches at approximately $\sim 60\%$ average activated excitatory neurons (see green lines and markers in Figure 7.2(a)).

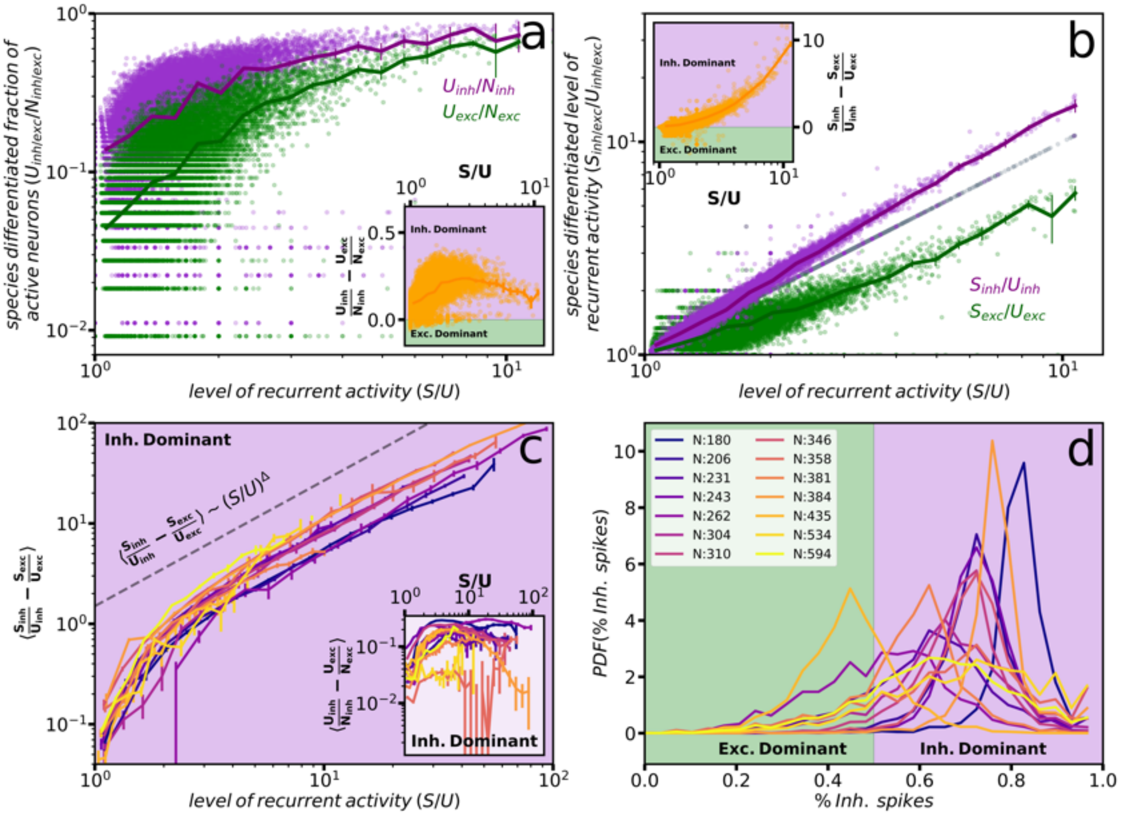

Figure. 7.2: Inhibitory neurons dominate avalanche activity. **a)** The species-differentiated fraction of active neurons is shown for each avalanche (circle markers) as a function of the level of total recurrent activity. Markers are color coded green and purple for excitatory and inhibitory neurons, respectively. The average, taken in twenty logarithmically spaced bins, is shown with solid lines and the error bars represent a 95% confidence interval. We highlight the difference of the species-differentiated fraction of active neurons $(\frac{U_{\rm inh}}{N_{\rm inh}} - \frac{U_{\rm exc}}{N_{\rm exc}})$ in the inset. **b)** The same as in a) but for the species-differentiated level of recurrent activity $(\frac{S_{\rm inh}}{U_{\rm inh}} - \frac{S_{\rm exc}}{U_{\rm exc}})$. **c)** The difference of the species-differentiated levels of recurrent activity is plotted on double-logarithmic scale to highlight the scaling with level of total recurrent activity. The lines represent the average, taken in twenty logarithmically spaced bins with error bars representing a 95% confidence interval. The fifteen samples are shown with line colors corresponding to different system sizes (see legend in d)). The background is highlighted in purple to indicate the recurrent activity is dominated by the inhibitory neurons. **d)** The probability distribution of the percentage of inhibitory spikes in an avalanche ('% inh. spikes′). The distribution is constructed for avalanches larger than size $S = 10$ and the background is highlighted green and purple to indicate excitatory and inhibitory dominance, respectively.

Qualitatively, the fraction of inhibitory neurons behaves like the excitatory neurons with increasing level of recurrent activity, however, we can see that the inhibitory neurons consistently have a higher average activated fraction for a given level of recurrent activity. Specifically for the feedforward avalanches, there are on average roughly 15% activated fraction of inhibitory neurons, and this fraction increases monotonically with the level of total recurrent activity in the avalanche, saturating at roughly $\sim 75\%$ average activated inhibitory neurons (see purple lines and markers in Figure 7.2(a)). We highlight the difference between these activated fractions ($U_{\text{inh}}/N_{\text{inh}} - U_{\text{exc}}/N_{\text{exc}}$) in the inset of Figure 7.2(a). There, it seems that the average value has a peak for some intermediate level of recurrent activity. By that, we just mean that this difference initially increases with increasing recurrent activity, until some intermediate value of S/U, beyond which the difference gradually decreases.

We have seen that there is, on average, a larger fraction of inhibitory neurons involved in an avalanche than excitatory ones. Now we'll disentangle the recurrent activity coming from the inhibitory vs. excitatory neurons by looking at the difference of species differentiated recurrent activity, $S_{\text{inh}}/U_{\text{inh}} - S_{\text{exc}}/U_{\text{exc}}$. In Figure 7.2(b) we show the species differentiated level of recurrent activity ($S_{\text{inh}}/U_{\text{inh}}$ in purple and $S_{\text{exc}}/U_{\text{exc}}$ in green) vs. the total level of recurrent activity (S/U). As a guide for the eye, in grey we plot the line S/U vs. S/U, to emphasize that the inhibitory neurons lie above this line and the excitatory neurons lie below this line. This means that the majority of the recurrent activity is emitted from inhibitory neurons, and this becomes truer for avalanches with higher levels of activity. This is quantified in the inset of Figure 7.2(b), where we plot the difference of species differentiated recurrent activity as a function of total recurrent activity. Plotted with double logarithmic axes, the difference of species differentiated recurrent activity even appears to scale with the total recurrent activity (Figure 7.2(c))

$$\left\langle \frac{S_{\text{inh}}}{U_{\text{inh}}} - \frac{S_{\text{exc}}}{U_{\text{exc}}} \right\rangle \sim \left(\frac{S}{U}\right)^{\Delta} \tag{7.1}$$

with $\Delta \sim 5/4$. As important as the scaling itself, is the consistency with which it shows up among the 14 samples (shown in different colors, Figure 7.2(c)). Nearly all samples at least begin to scale with the level of recurrent activity, and for a majority of the samples, the scaling is undeniable. We show in the inset of Figure 7.2(c) the average difference of the species differentiated fraction of active neurons. Again, all samples show a dominance of the inhibitory species, and the qualitative features are roughly consistent between samples.

To further confirm the dominance of inhibitory spikes in neuronal avalanches, we wanted to characterize what percentage of an avalanche is composed of inhibitory spikes. We extracted a histogram of this inhibitory percentage ($'$%inh.$'$) and normalized the counts to estimate the probability distribution, PDF(%inh.) Figure 7.2(d). Here, we are interested in the larger, and more recurrent, avalanches, where the inhibitory dominance is most clear. To extract a PDF resembling the statistics of just these larger avalanches, we only sample avalanches that are greater than size 10. This number is chosen mostly arbitrarily, just large enough to suppress the very small avalanches causing the spikes at zero- and one-hundred percent inhibitory spikes, and instead highlight the statistics of the most recurrent avalanches. We have checked that the results that follow do not depend on the exact value chosen for this minimal avalanche size (see Appendix F).

When focusing just on the larger avalanches, an indisputable preference towards inhibitory spikes is seen in all but one sample Figure 7.2(d). What we mean is that the distribution is sharply peaked around a value that is larger than one-half, signifying the avalanches are composed of a majority of inhibitory spikes rather than excitatory spikes. The results of this section point to a clear dominance of activity by the inhibitory neurons, despite the fact that they are the minority population, which should be considered the most important takeaway from this section.

7.3 Model

In this section, we test whether these results outlined above can be recovered by our model that includes both inhibitory and excitatory neurons. First, we'll reiterate the details of the model, emphasizing the parameter combinations that will be most important to us when fitting the data. We'll show that the model is indeed able to produce qualitative features seen in the difference of species differentiated fraction of active neurons as well as the difference of species differentiated recurrent activity. We also show the percentage of spikes in the model avalanches are on average inhibitory dominant. We'll close with a brief investigation of the average avalanche shapes produced by the model simulations.

We employ the model from the previous chapter (Chapter 6) to simulate two species of spiking neurons, inhibitory and excitatory. The underlying process that is simulated is the spiking of one ('pre-synaptic') neuron, triggering the release of, and corresponding redistribution of membrane potential to the neurons that are synaptically connected to this pre-synaptic neuron (i.e., the 'post-synaptic' neurons). Excitatory and inhibitory neurons both behave the same way in our model– once their membrane potential surpasses a threshold, they spike. However, the species which was the presynaptic neuron determines the resulting redistribution of membrane potential. If the pre-synaptic neuron is excitatory, when it spikes, it *increases* the potential of all post-synaptic neurons, regardless of species type. In contrast, if the pre-synaptic neuron is *inhibitory* when it spikes it will increase the membrane potential of the inhibitory post-synaptic neurons and *decrease* the potential of all post-synaptic excitatory neurons. This is where the name 'inhibitory' comes from since they are inhibiting the excitatory neurons from spiking.

As outlined in the previous chapter, when a neuron spikes, it loses potential in the amount ΔV_{i}

$$\Delta V_{\mathrm{i}} = V_{\mathrm{dynamical}} - V_{\mathrm{arrest,i}} \tag{7.2}$$

where this dynamical threshold, arising from the weakening, takes on two values during an avalanche

$$V_{\mathrm{dynamical}} = \begin{cases} 1 & \text{before spiking} \\ 1 - \epsilon_{\mathrm{E/I}} & \text{after spiking, for remainder of avalanche} \end{cases} \tag{7.3}$$

When the neuron is not spiking, it is just receiving incoming potential from all neurons that spiked in the previous time step. Specifically, a neurons potential increases by an amount

$$\Delta V_{\mathrm{incoming,i}} = \sum_{\mathrm{j}} \frac{J_{\mathrm{ij}} \Delta V_{\mathrm{j}}}{\sum\limits_{\mathrm{k}} |J_{\mathrm{kj}}|} \tag{7.4}$$

where the sum is over all neurons that spiked in the previous time step. Given the way we defined the interaction strengths between any two pairs of pre- and post-synaptic neurons, it will be more illuminating to rewrite this equation separately for each species. First, an inhibitory neuron experiences the same increase in potential whether an excitatory or an inhibitory neuron had spiked. The increase in potential would then be identical to Eq. (7.4), where the sum over $'j'$ is over all spiking neurons in the system. In contrast, excitatory neurons experience an increase in potential from other spiking excitatory neurons, but a *decrease* in potential from spiking inhibitory neurons. We then define any connection strength between a pre-synaptic inhibitory neuron and a post-synaptic excitatory neuron as $-J/\gamma$ where, as explained in Chapter 6, γ is introduced to control the level of inhibition in the system. The two important equations defining our model simulations then simplify to

$$\Delta V_{\mathrm{incoming,i}} = \begin{cases} \sum\limits_{\mathrm{j} \in \mathcal{E}} \frac{J_{\mathrm{ij}} \Delta V_{\mathrm{j}}}{\sum\limits_{\mathrm{k}} |J_{\mathrm{kj}}|} + \sum\limits_{\mathrm{j} \in \mathcal{I}} \frac{J_{\mathrm{ij}} \Delta V_{\mathrm{j}}}{\sum_{\mathrm{k}} |J_{\mathrm{kj}}|} & \text{for i inhibitory} \\ \sum\limits_{\mathrm{j} \in \mathcal{E}} \frac{J_{\mathrm{ij}} \Delta V_{\mathrm{j}}}{\sum\limits_{\mathrm{k}} |J_{\mathrm{kj}}|} - \frac{1}{\gamma} \sum\limits_{\mathrm{j} \in \mathcal{I}} \frac{J_{\mathrm{ij}} \Delta V_{\mathrm{j}}}{\sum_{\mathrm{k}} |J_{\mathrm{kj}}|} & \text{for i excitatory} \end{cases} \tag{7.5}$$

where $\mathcal{E}$ and $\mathcal{I}$ are the sets of excitatory and inhibitory neurons that spiked in the previous time step, respectively. The denominator of this term represents the sum of the absolute value of $'j'^{th}$ row in the synaptic weight matrix J.

When simulating this model, we must fix the number of excitatory and inhibitory neurons, N_{exc} and N_{inh}, respectively. The next free parameter is γ which we can use to adjust the strength with which a spiking inhibitory neuron reduces the potential of an excitatory neuron. We'll find that γ and the inhibitory ratio have similar effects on the avalanches, and thus we can increase γ to compensate for an increase in N_{inh}.

With the addition of dynamical weakening, there are two more parameters,ϵ_{E} and ϵ_{I}, representing the strength of this weakening for the excitatory and inhibitory neurons, respectively. Finally, the width of the disorder distribution, w, and the width of the coupling matrix distribution, ΔJ, are the last free parameters. We'll discuss below how these parameters can be adjusted to find excellent agreement with the empirical results outlined in the previous section.

7.3.1 Comparing with Data

For the simulation results below, we will fix the inhibitory ratio to 20% (previously reported in mouse cortex [128]) as well as the widths of the two disorder distributions, $w = 1.9$ and $\Delta J = 1$, leaving only $\gamma, \epsilon_I, \epsilon_E$ to vary. The first question we have is whether our estimate of 20% inhibitory neurons was skewing our simulations. We show in Figure 7.3(a) the comparison between avalanche size CCDF of the experiment and simulations. As we saw in the previous chapter, with 20% inhibitory neurons, there is a set of parameters that matches the data very well ($\gamma = 2.25$,$\epsilon_I = 0.85$,$\epsilon_E = 0.55$). When we adjust this ratio to the empirically measured value ($\sim 28\%$), we see the power-law decays faster than when using a 20% inhibitory ratio, and to smaller avalanche cutoff (Figure 7.3(a), green line). Similarly, the level of recurrence decreases slightly, descending slightly from the maximum value of S vs. U we had earlier (Figure 7.3(a)). Importantly, by simply reducing the strength of the inhibitory to excitatory connections, $\gamma = 2.25 \rightarrow 3$, we can get the simulations to converge back to the experiment even while using the empirically observed ratio of inhibitory to excitatory neurons ($\sim 28\%$, Figure 7.3(a), red line).

We've seen the results of Chapter 6 are robust when we increase the inhibitory ratio in the simulations to the empirical value. The next question is, does the simulation also reproduce the dominance in the inhibitory species that we saw in the previous section? First, we show the difference of species differentiated recurrent activity (c.f. Figure 7.2(b)) for our simulations before (blue) and after (red) increasing this inhibitory ratio. We can see that in both cases the simulations reproduce the scaling law found in the experiments (c.f. Eq. (7.4)). Interestingly, increasing the inhibitory ratio from 20% to $\sim 28\%$ did not change the scaling produced by the simulations, but did increase the maximal level of recurrent activity (the red line extends further than the blue in Figure 7.3(b)). In the inset of Figure 7.3(b) we confirm the fraction of activity inhibitory neurons is greater than the excitatory ones, just as seen in the experiments.

A clear picture of inhibitory dominance is starting to appear, but could it really be that avalanches are predominately composed of inhibitory spikes? After all, the excitatory neurons are certainly the majority species (except for two samples). To verify this, we extract the fraction of total spikes in an avalanche that were emitted from an inhibitory neuron and characterize with a PDF that we can compare to experiments (c.f. Figure 2.2(d)). We see that the simulations with only 20% inhibitory neurons were not able to reproduce the experiment, and instead, the PDF is skewed to the left, signaling excitatory dominance. In contrast, once we corrected the inhibitory ratio to the empirical value of $\sim 28\%$ and adjusted γ from $2.25 \rightarrow 3$, the PDF shifts to the right, signaling more inhibitory dominance, and, lining up more closely with experiments.

Lastly, we wanted to see if this inhibitory dominance has any noticeable signature in the averaged avalanche shape profiles. As we pointed out in Chapter 4, these profiles, averaged in bins of duration, start to skew leftward for the longest duration avalanches (Figure 7.3(d), black). Amazingly, this is indeed a feature that our simulations produce when we have inhibitory neurons. We can see that when adjusting the ratio of inhibitory neurons from 20% (blue) to the empirical value ($\sim 28\%$, red) the duration averaged spiking-rate profile skews further to the left, suggesting the inhibition may be at the root of the leftward skew. We leave further investigation of the shapes to future work but mention this additional agreement between the model and simulations here for completeness.

7.4 Discussion

The first striking feature we observed in this chapter was that the inhibitory ratio was larger than the typically reported [128] and simulated [2] value of 20%. This was true of all samples investigated, though as the sample sizes (N) got larger, the inhibitory ratio roughly tended towards 20%, bottoming out at $\sim 28\%$ amongst the fourteen data sets (Figure 7.1(b)). This trend may be suggesting that for larger system sizes this ratio may truly be closer to 20%, as typically reported [128]. Nevertheless, the observation we point out here (that the inhibitory ratio is larger than 20%) certainly should likely be considered when trying to model in-*vitro* cortical slices with order hundreds of neurons.

Though the inhibitory ratio was larger than anticipated for all samples, it was almost always less than 50% (excluding two samples) i.e., the excitatory population tends to be the majority. It is then very interesting that the majority of the spikes measured in avalanches are in fact coming from the minority population; that of the inhibitory neurons. We confirmed this inhibitory dominance with several measures, indeed finding consistency across all of them—the inhibitory neurons are more active. Not only this but since there are typically fewer inhibitory than excitatory neurons in a sample, this smaller population of inhibitory neurons must have been firing several times each, i.e., they were highly recurrent, which is exactly what we observe.

We also point out the need to introduce γ into our model in order to produce large enough avalanches to match the experiment. This was introduced in [71], when we used the conservative inhibitory ratio of 20%. Importantly, we found that increasing the inhibitory ratio (to the empirically estimated value of $\sim 28\%$) in our simulations required that we also increased γ from $2.25 \rightarrow 3$ in order to still match the experiments. This provides a useful insight—that the strengths of the synapses may be able to adjust in order to account for the specific inhibitory ratio in a system. Synapses that adjust themselves dynamically have long been discussed

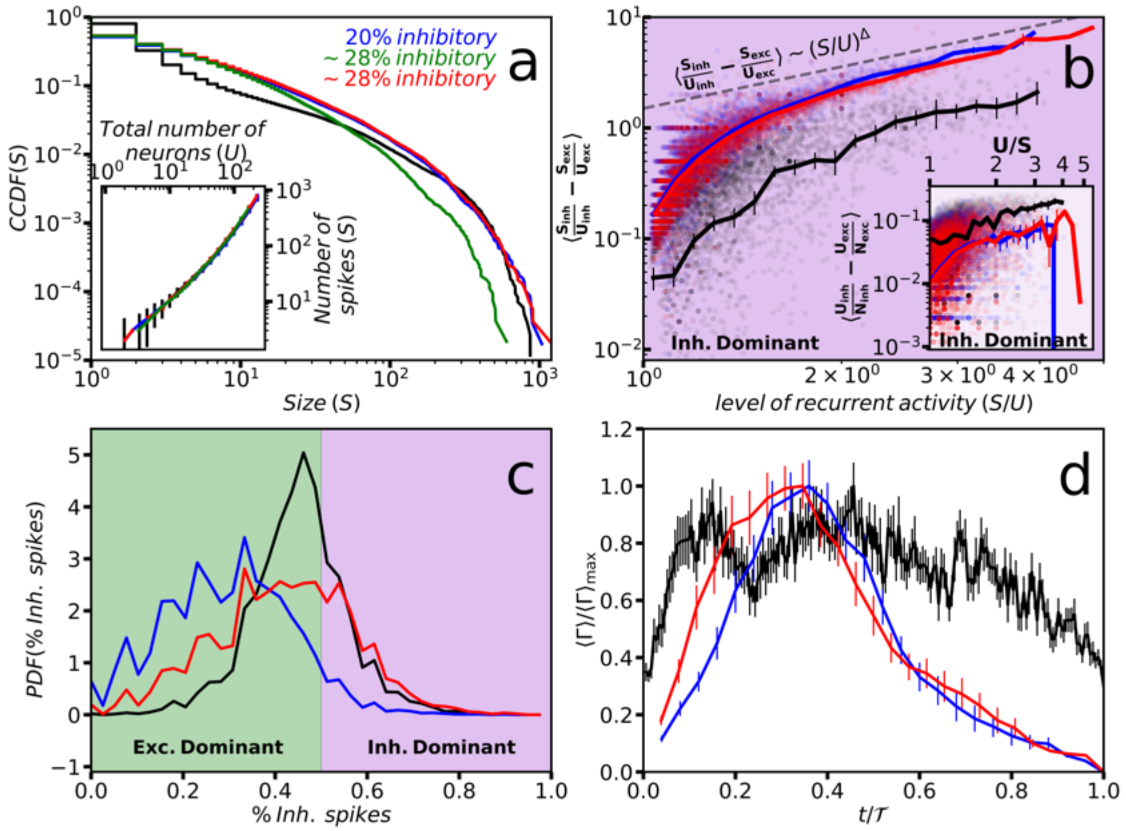

Figure. 7.3: Reproducing inhibitory dominance with model simulations. a) The CCDF of avalanche size is shown for the experiment (black) and model simulations with parameter values found in Chapter 6 (blue). The ratio of inhibitory neurons is adjusted from 20% to the empirically estimated value of $\sim 28\%$, with all other parameters held fixed (green). The simulation parameter γ is then adjusted from 2.25 (green) to 3 (red). The inset shows the averaged number of spikes (S) given the total number of neurons (U), with error bars showing the standard deviation. **b)** The averaged difference in species-differentiated recurrent activity for each avalanche is plotted as a function of total level of recurrent activity (circle markers) for the experiment (black) and the simulations before (blue) and after (red) adjusting the inhibitory ratio and γ from 2.25 to 3, respectively. The average, taken in twenty logarithmically spaced bins, is shown with a solid line. Error bars represent a 95% confidence interval, and the background is highlighted purple to indicate inhibitory dominance. The inset shows the average difference of species differentiated fraction of active neurons for each avalanche (circle markers) plotted as a function of total level of recurrent activity. Solid lines represent the average, taken in twenty logarithmically spaced bins. Error bars represent a 95% confidence interval, and the background is highlighted purple to indicate inhibitory dominance. **c)** The probability distribution of the percentage of inhibitory spikes in an avalanche ('% inh. spikes$'$) for the experiment (black) and the simulations before (blue) and after (red) adjusting the inhibitory ratio from 20% to $\sim 28\%$ and γ from 2.25 to 3, respectively. The distribution is constructed for avalanches larger than size $S = 10$ and the background is highlighted green and purple to indicate excitatory and inhibitory dominance, respectively. **d)** The averaged temporal profile for forty of the longest duration avalanches (i.e., those outside of the duration scaling regime) is shown for experiment (black) and the simulations before (blue) and after (red) adjusting the inhibitory ratio and γ. The y-axis is rescaled by the maximum value of each curve for ease of comparison.

in the literature of both experimental and theoretical neuroscience, under the name of synaptic plasticity [130]. Our work here may be suggesting that some aspect of synaptic plasticity, specifically involving those synapses connecting inhibitory to excitatory neurons, may be driven by changes in the inhibitory ratio. If future experiments could replicate several samples with precise inhibitory ratios, one could test this prediction by comparing the fits to γ between the different samples (i.e., different inhibitory ratios). In this scenario, our work suggests you would find γ would increase with increasing inhibitory ratio.

We have started to build up a sketch of inhibitory-excitatory avalanche dynamics, where the inhibitory neurons are spiking much more than the excitatory ones, and they are coupled more weakly to excitatory neurons than to other inhibitory ones. This image is consistent with the idea that the inhibitory neurons act as a sea of constant spiking [131], taking several inhibitory spikes per excitatory spike to calm down the behavior. Crucially, this observation would be consistent with the idea that neurons act like the 'brakes of the brain' [132, 133]. Interestingly, we also found the average avalanche temporal profiles displayed a leftward skew of the averaged temporal profiles for fixed duration for the largest duration avalanches, and this feature was not just seen in the experiments, but in the model as well. It is certainly suggestive that this leftward skew could originate from a situation where, at first, many excitatory neurons spike quickly, and then, with some temporal delay, the inhibitory neurons came on afterward to quell the activity. Indeed, this might cause the decay in the temporal profile observed in Figure 7.3(d). This is certainly work that should be pursued in the future, and while it is beyond the scope of this thesis, we have included a brief investigation of this idea in Appendix F.

To close, we briefly mention that our work connects to avalanches outside of the realm of neuroscience as well; anywhere that negative couplings arise (for example in antiferromagnets). Recall that we found the difference of species differentiated recurrent activity scaled with total recurrent activity (c.f. Eq. (7.4)). The traditional scaling reported in avalanches (neuronal avalanches and otherwise) is usually attributed to an underlying critical point, which tells you about the the behavior of the system on large length and time scales. It would be important as a next step then to figure out if the scaling law observed in this chapter also stems from the underlying critical point, or if it stems from something more mundane. If it is the former, our results could be connected to other systems using the principle of universality.

7.5 Conclusion

In this chapter, we showed how the additional information of the species type helped to uncover previously unknown features of neuronal avalanches. We found that, though the inhibitory neurons typically were the minority population, they tended to make up a majority of the avalanche activity. This of course implies the inhibitory neurons must have been much more recurrently active, a feature we predicted with our model simulations. We also showed that the average avalanche size, as well as the total recurrent activity (i.e., the total number of spikes in an avalanche (S) divided by the number of unique neurons that participated in the avalanche (U)), increased roughly exponentially with the inhibitory ratio. Model simulations found best comparison to data when reducing the inhibitory to excitatory connection strengths (by a factor of $\gamma = 3$), evidence in favor of the idea that the inhibitory spikes are like a weak sea of activity surrounding the excitatory neurons. Lastly, with the leftward skewed avalanche shapes (i.e. the average firing rates for avalanches of a given duration) we provided preliminary experimental and simulation evidence that the inhibitory neurons turn on to stop long avalanches, i.e., that they are the ‘brakes of the brain’.

Appendix A

Appendix to Chapter 2

A.1 Constructing a Probability Distribution Function (PDF)

The probability distribution function, PDF, can be constructed from data via. a histogram. In Figure A.1(a) we show the counts (i.e., the number of occurrences) for a list of random integers X; this is what we mean by a histogram. To construct the PDF from such a histogram, one just has to scale the y-axis by the number of samples, N_{samples}. We show in orange the PDF for the list of random numbers X, highlighting that the PDF is normalized to unity. In Figure A.1(b) we highlight the PDF estimated from a histogram of avalanche sizes (neuronal avalanches from Chapter 6). You can see that, depending on the number of bins used to construct the histogram from the data, you can get varying results for the PDF. This discrepancy can be avoided completely by instead using the complement of the cummulative distribution function, $CCDF$.

A.2 Constructing the Complement of the Cummulative Distribution Function (CCDF)

The $CCDF(x)$ of a random variable $\mathcal{X}$ tells you the probability of having x take on a value greater than or equal to $\mathcal{X}$. This means, for $x < \min(\mathcal{X})$, the $CCDF$ should equal unity and for $x > \max(\mathcal{X})$, the $CCDF$ should equal zero, i.e., there is one-hundred percent chance of finding a value of x that is in $\mathcal{X}$, and zero percent chance of finding a value of x that is not in $\mathcal{X}$. To construct the $CCDF$, simply order the list of random numbers $\mathcal{X}$ from smallest to largest. This ordered list is the x value of the $CCDF$ graph. If the number of elements in the list is N, then the 'y' value assigned to the smallest value in the list is 1, and every

$i'th$ element after the first one is assigned a 'y' value of $1 - i(1/N)$. That is, the 'y' value assigned to the largest value in the list is just $1/N$ (i.e., approaching zero for large N). In Figure A.1(c), we plot the $CCDF$ corresponding to the random list X from Figure A.1(a).

A.3 Averaging in Bins

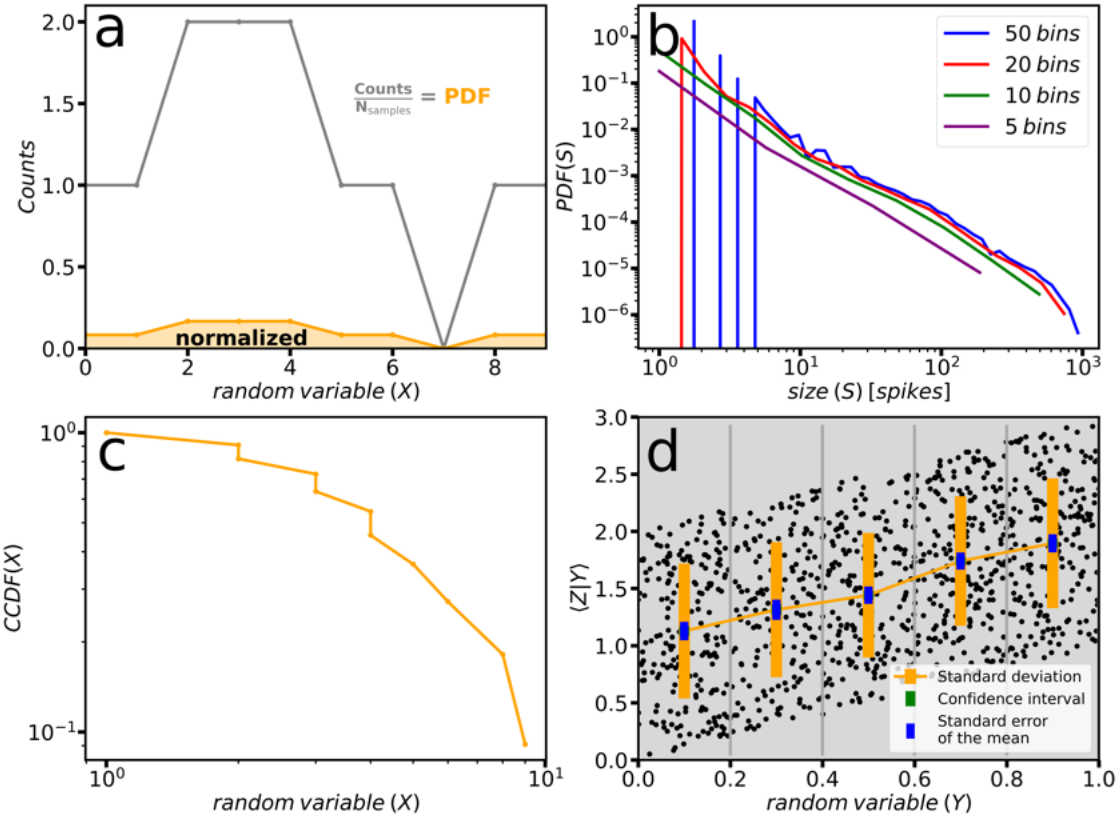

Figure. A.1: Statistical analysis of data a) For an example set of random numbers (X) we plot the counts (number of times of occurrence of a number in the set) in grey. The PDF, shown in orange, is constructed by normalizing the counts by the number of samples, N_{samples}. **b)** We show the PDF of avalanche size, S, for the neuronal avalanches found in Sample 1 (see Chapter 4). We construct the PDF using 50 (blue), 20 (red), 10 (green) and 5 (purple) bins, to highlight that the number of bins can cause slight differences in the estimated PDF. **c)** The complement of the cummulative distribution function ($CCDF$) is plotted for the same set of random numbers in (a), X. **d)** Two sets of random variables, Y and Z are plotted vs. each other as black dots. To find the average value of Z for a given value of Y, we take bins in the Y-variable (shown with grey vertical lines) for which we assign a Z value equal to the mean of the Z value of all black dots that lie within the bin (orange line connects these means). To show the error of the bin average, we can use the standard deviation (orange), the confidence interval (green) or the standard error of the mean (blue).

We can test predictions for expected values of avalanche size given duration, for instance, by 'bin-averaging' the scatter plot of size vs. duration. In Figure A.1(d), we illustrate the procedure of bin-averaging with some fabricated random data, Y and Z. First, we make a scatter plot of Z vs. Y (shown as black dots in Figure A.1(d)), then we divide the Y axis into bins; evenly spaced segments of the Y range of the data (shown with grey lines in Figure A.1(d)). The bin-average for a given bin is simply the mean of the data points that lie within the bin. The error bar for the bin-average is typically chosen to be the standard deviation of the data in the bin (S.T.D), or the standard error of the mean (S.E.M=S.T.D/$\sqrt{N}$, where N is the number of events in the bin) or even the confidence interval (C.I, 2(S.E.M)). In Figure A.1(d), we illustrate these three different error bars in orange, blue and green, respectively.

A.4 Constructing Average Avalanche Temporal Profiles

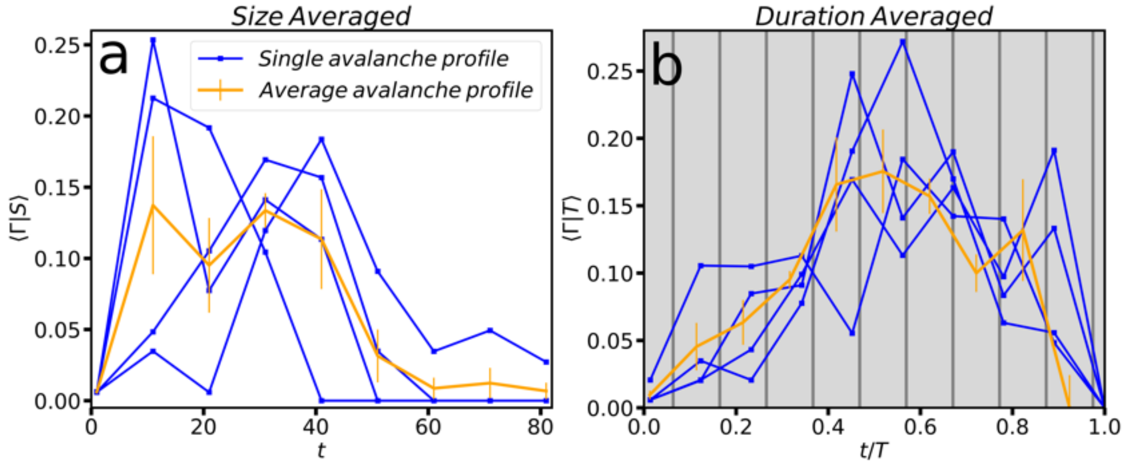

Figure. A.2: Constructing average avalanche temporal profiles a) Individual avalanche profiles ($\Gamma(t)$ during the duration of the avalanche) are shown in blue for four different avalanches with similar sizes. The average avalanche profile is shown in orange, with error bars representing the standard error of the mean. **b)** Individual profiles ($\Gamma(t)$ during the duration of the avalanche) are shown in blue for four different avalanches with similar durations. The average avalanche profile is shown in orange, with error bars representing the standard error of the mean.

To compare data with the predictions of the critical scaling function, we need to form an averaged avalanche profile, and we have to do this for different sizes (and/or durations), in order to check for a scaling collapse. To collapse average profiles is easy enough; just follow the scaling laws in Table 2.1. It is maybe not obvious, however, how to construct the average avalanche profile from a set of individual profiles. To average avalanche profiles of similar size, we first rescale all the profiles to start at $t = 0$ (Figure A.2(a)). Avalanches with similar size will not necessarily have the same duration, so we pad all avalanches with zeros at the end until they are all equal to the length of the longest duration avalanche in the bin. Now, the average shape is taken simply by averaging the individual shapes point-wise. The error bar shows the standard error of the mean, since we would ultimately compare this to a prediction of the critical scaling function, which predictions how the *average* values of these profiles will look.

Similarly, to average profiles of similar duration, we first rescale them to start at $t = 0$ (Figure A.2(b)). Here, in contrast to the size averaging, we do not pad the avalanches with zeros, but instead rescale the time axis by the avalanches duration, T, so that all of the individual avalanche profiles start and end at zero and one, respectively. In contrast to the size averaged profiles, now we see that there does not exist a unique data point in each shape for us to take a point wise average with. By this, we just mean that there are not the same number of data points in each avalanche profile. To average, we simply divide the time axis into bins and take the average of any data points that lie within the bins. Here, we again use the standard error of the mean for the error bars, and, the number of bins is set by the number of data points in the shortest of the individual (blue) avalanche profiles.

Appendix B

Appendix to Chapter 3

Table B.1: Analytical predictions of simple slip model.

System	$\boldsymbol{\tau}$	$\boldsymbol{\alpha}$	$\boldsymbol{\sigma\nu z}$	$\boldsymbol{\frac{\alpha-1}{\tau-1}}$	$\boldsymbol{\frac{1}{\sigma\nu z}}$
Micro-tribometer	2.0139 ± 0.0223	2.1285 ± 0.0088	0.7125 ± 0.0117	1.1130 ± 0.0124	1.4034 ± 0.0117
MD Simulations	1.9448 ± 0.0638	2.8535 ± 0.1775	0.5554 ± 0.0322	1.9617 ± 0.1032	1.8004 ± 0.0322
MFT $\epsilon = 0$	1.5	2	0.5	2	2
WMFT $\epsilon > 0$	2	3	0.5	2	2
MFT+Noise	1.9063 ± 0.0164	1.8336 ± 0.0065	0.4892 ± 0.0121	0.9197 ± 0.0089	2.0438 ± 0.0121

The mean-field model for material deformation was solved [2, 134] and the critical exponents cataloged [135]. **Table B.1** summarizes the six model predictions that were compared in this paper.

B.1 Simulating the MFT Model

With the appropriate pinning forces (see [5]) the dynamics of the model are completely described by the local friction force at each asperity site 'i'. When the local friction force surpasses the (uniform) failure force (i.e., $f_{\text{fail}} = 1$), its friction force is reduced to $f(r_{\text{i}}, t_{\text{j+1}}) = f(r_{\text{i}}, t_{\text{j}}) - [f_{\text{fail}} - f_{\text{arr}}(r_{\text{i}})]$. Here, we chose to include the effects of the quenched disorder (f_w) in the form of disorder in this arrest force. Specifically, the arrest force is sampled from a uniform distribution of width w, centered on $f_{\text{arr}} = 0$. In the mean field theory, the energy lost due to this reduction of force at site 'i' is redistributed by increasing the friction force at all other local asperities by $J/N[f_{\text{fail}} - f_{\text{arr}}(r_{\text{i}})]$ at time step $t_{\text{j+1}}$.

To account for the externally applied force, F, we apply a uniform velocity v a distance L away from the friction interface, which emulates a spring force with constant K_{L} such that $F = vK_{\text{L}}$ [5, 7]. Thus, at all time steps we uniformly raise the friction force at each asperity by the amount vK_{L}, which could in principle

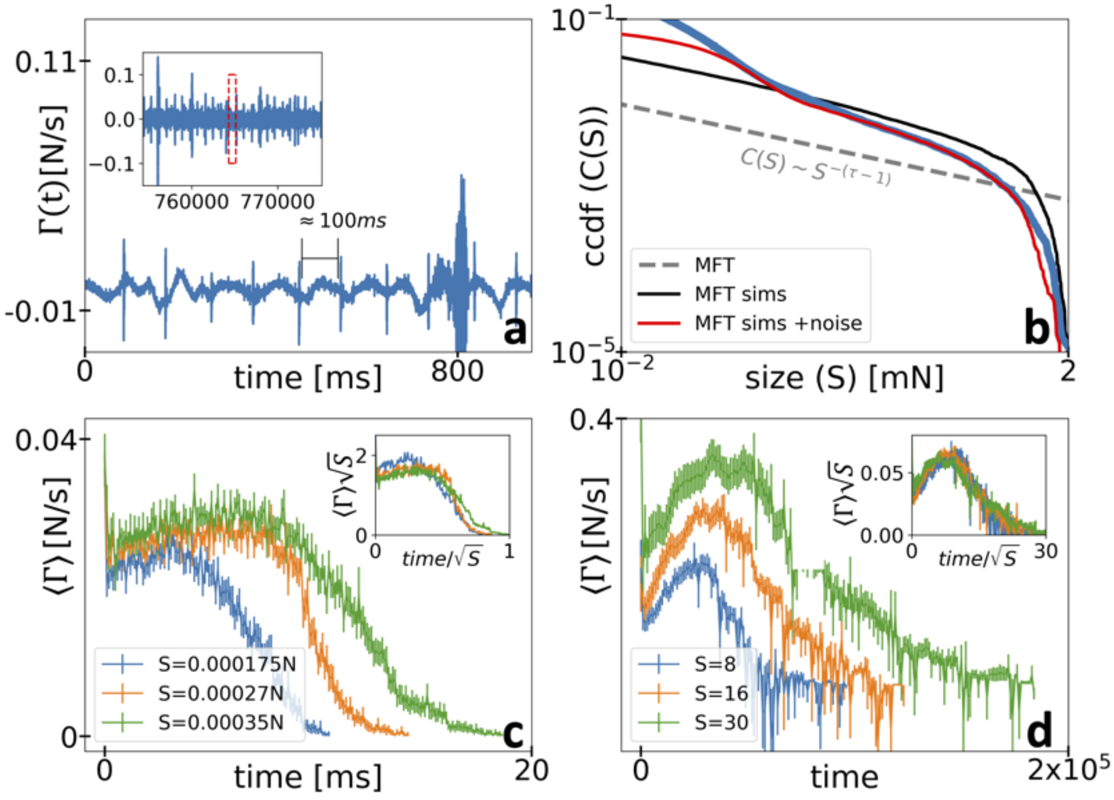

Figure. B.1: Emulating periodic machine noise in model simulations reproduces artifacts from tribometer experiments. a) The derivative with respect to time of friction force (main) and a zoom-out (inset). **b)** ccdf of avalanche size, for micro-tribometer experiments (blue) MFT prediction (dashed grey) MFT simulation (black) and MFT+noise simulation (red).**c)** Average force-drop rate profiles for micro-tribometer experiments. The profiles were computed from n avalanches found in a narrow size bin centered on S_i, for three bins in the scaling regime. Markers are shown with error bars calculated from the standard error in a log binned average over time. The size averaged profiles are predicted to collapse onto each other when both the x and y-axis are rescaled by $S_i^{-1/2}$ (inset). **d)** Same as (c), but for MFT+noise simulations.

help to persist avalanches longer if this extra amount of friction force allows a new asperity to slip. Here, however, we are interested in the dynamical regime where the asperities interact much more quickly than the external force changes, i.e., $vK_\mathrm{L} \ll \langle f_\mathrm{fail} - f_\mathrm{arr}(r_\mathrm{i}) \rangle$ [7]. Here the average is taken over all asperities, and so represents the characteristic amount of force needed to raise an arrested asperity to failure. With the system driven slowly enough, the effect of driving force F on the avalanche statistics vanishes.

In our simulations, without the external driving force, we simply raise the friction uniformly at all sites until one fails. From this failure, we distribute the friction force to the rest of the system, which causes other asperities to surpass threshold and fail, in a chain reaction or 'avalanche'. After an avalanche, we randomize the friction forces, raise the force uniformly until the next site fails, and repeat until we have constructed an ensemble of avalanches. The size of the avalanche is taken to be the number of cites that slipped, and duration of the avalanche is the time of the final slip minus the time of the initial slip

B.2 Periodic Noise (Experiments)

The experiments have periodic noise (Figure B.1(a), downstream from the measurement. Every 100 μs, there is an 18 time-step (1.8×10^{-3} s) burst of fluctuations in the force-drop rate, with average amplitudes close to twice the average maximum velocity seen in the extracted avalanches. We therefore expect that simply overlaying the equivalent periodic noise signal over a simulated time trace of the mean field theory (MFT) will show results similar to experiments (Figure B.1(b)). Specifically, we add Gaussian noise (with a mean of 0 and a standard deviation of 1) to the velocity time trace from the MFT simulations, in the same periodic manner of the experiments. The effect of the noise is to cut off the avalanches at the beginning, resulting in nonzero starting velocities, as seen in the size averaged shapes from both experiment and MFT simulations with noise (Figure B.1(c) and (d), respectively). The periodic noise populates the avalanche size distribution with multiple fast events that have a size proportional to the noise amplitude. The MFT simulations with an amplitude of 2 yielded results that qualitatively resemble the experimental size distribution (Figure B.1(b)).

B.3 Unavoidable Oscillations (MD Simulations)

The MD simulations have oscillations upstream from the measurement that are an artifact of the simulation method. We therefore expect that modifying the MFT simulation to include oscillations will show results those from the molecular dynamics simulations [125]. These effects are most clearly seen in averaged velocity profiles ($\Gamma(t)$) where 3 clear size regimes emerge (Figure B.2). The smallest avalanches are unaffected by the oscillations (Figure B.2(a-b)) until a maximum size, after which the profiles of the avalanches themselves are the oscillation profiles (Figure B.2(c-d)). In avalanches that have an inverse duration T greater than the frequency of these applied oscillations the duration and size averaged shapes resemble the MFT predicted scaling function, but with a sinusoidal gain added (Figure B.2(e-f)).

B.4 Avoidable Oscillations (Experiments)

The experiments can also be affected by oscillations when, for instance, the cantilever is near a resonant frequency. It was important to use a cantilever with a soft normal stiffness to minimize these resonance effects (Figure B.3). The blue lines show the avalanche results for the stiffest cantilever. The signature in the duration distributions of oscillations in experiments with stiff cantilevers is a shoulder at $T = T_{\text{osc}}$. This implies that the chance for an event of such a duration is large (Figure B.3,bottom right). The oscillation frequency can be seen as a clear peak in the power spectrum (Figure B.3, bottom left). Using a softer cantilever yields a lower resonant frequency, implying that for softer and softer cantilevers the peak seen in Figure B.3(c) (blue) should shift further (orange) and further (green) to the left. The shoulder in the CCDF will become less sharp (Figure B.3(d), orange) and ultimately disappear (green). The analysis regimes for experiment and MD simulations were chosen by determining the regimes in the power spectra that were not affected by oscillations Figure B.4- B.6.

B.5 Run-In Behavior (Experiments)

As mentioned in the main paper, the experimental analysis was performed in the steady state regime only (Figure 3.1(c)), corresponding to the last 400 of 500 total cycles for each experimental dataset. These results matched best with the purely incommensurate simulations. However, it is interesting to note that when the run-in regime is analyzed, the results compare best with a mixture of MD simulations, 4 parts incommensurate per 1 part commensurate (Figure B.7).

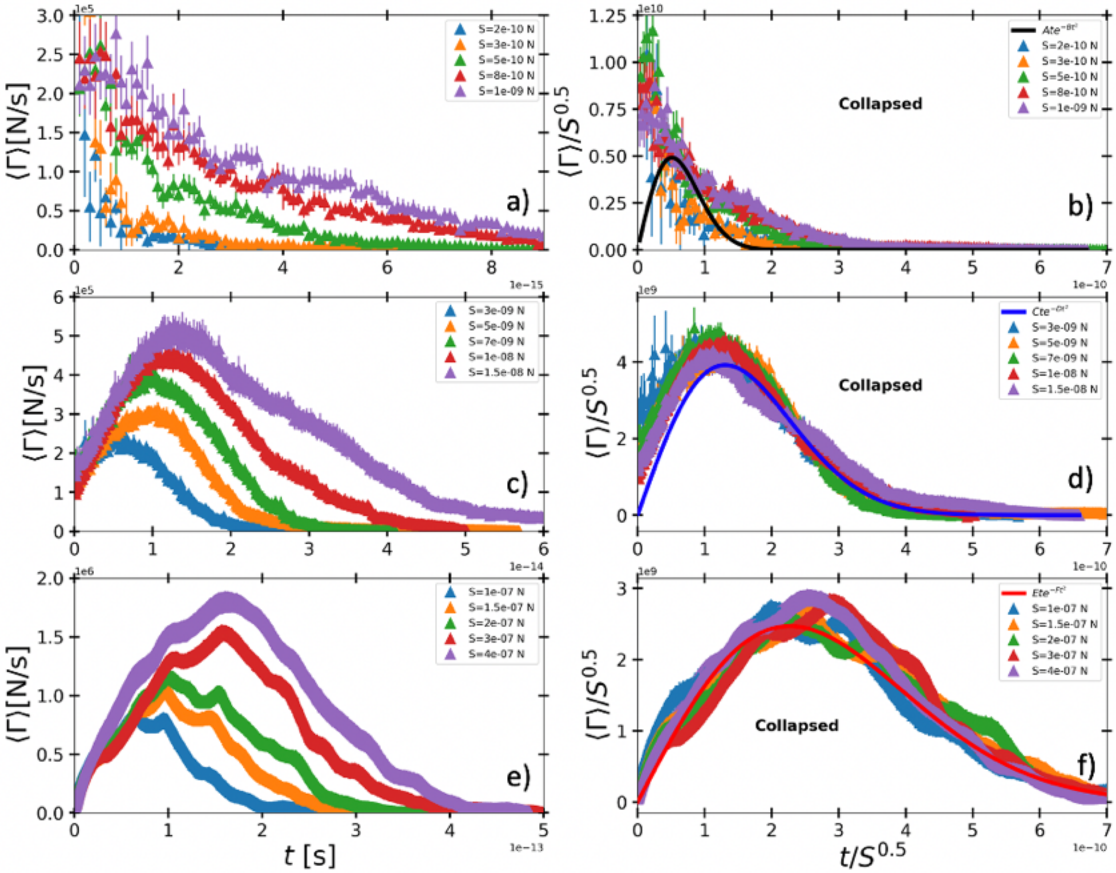

Figure. B.2: The three qualitatively distinct size regimes seen in the average velocity profiles. The analysis in this paper used avalanches from the largest bins shown in **(e)** and **(f)**, with velocity profiles that relaxed back to the MFT form. This is in contrast to those in the second regime where oscillations dominate the profile **(c, d)**. (See also [125]).

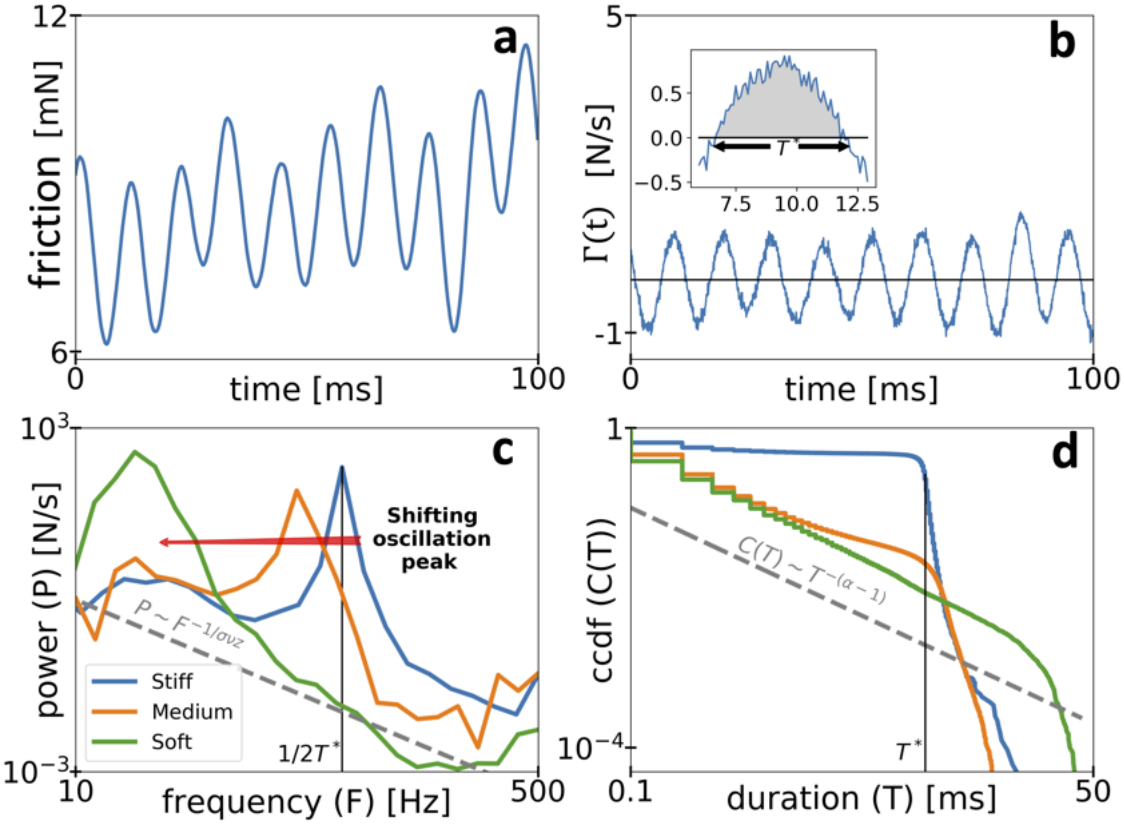

Figure. B.3: Reducing cantilever stiffness lowers oscillation frequency to uncover avalanches in friction force fluctuations. a) Friction force measurements of MoS_2 in a micro-tribometer. **b)** The derivative with respect to time of the friction force, Γ. A black line at $d\Gamma/dt = 0$ highlights the zero crossings at the beginning and end of an avalanche. The inset shows a zoom in of one avalanche with a characteristic duration T^*. **c)** The power spectrum of the friction force trace for cantilevers with lateral stiffness's 3794 N/m ('Stiff'), 1571.5 N/m ('Medium') and 42.8 N/m ('Soft'). **d)** The ccdf of avalanche duration for the same three cantilevers. The MFT predictions are shown with a dashed grey line.

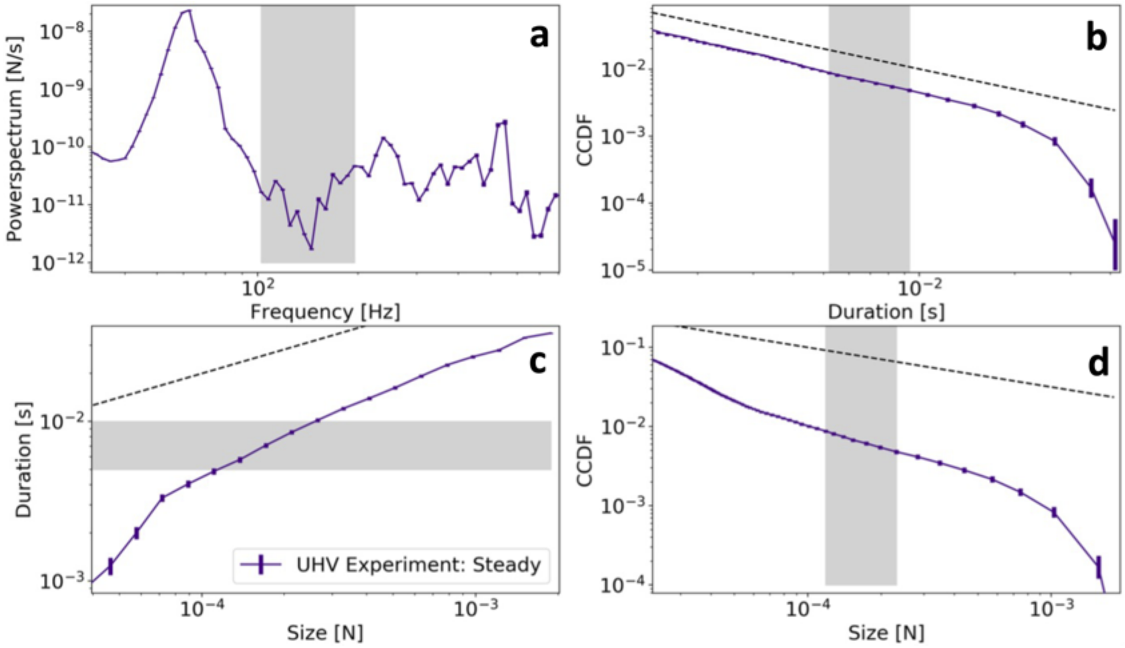

Figure. B.4: Defining the scaling regime in experiments. The power spectrum of the experimental friction force time trace **a)** and the avalanche duration vs. size **c)** are shown with error bar markers calculated from the standard error in a log-binned average over frequency and size, respectively. The duration **b)** and size **d)** CCDFs are plotted with error bar markers calculated with a Bayesian technique with 95% confidence interval [38]. Using the highlighted frequency regime in the power spectrum $[f_{\min}, f_{\max}]$, we then define the duration regime as $[1/f_{\max}, 1/f_{\min}]$ (highlighted in the duration CCDF), that is used in conjunction with the size vs. duration plot to fix the size regime $[S_{\min}, S_{\max}]$ (highlighted in the size CCDF). Mean field power-law predictions for infinite systems size with zero oscillations are plotted as black dashed lines.

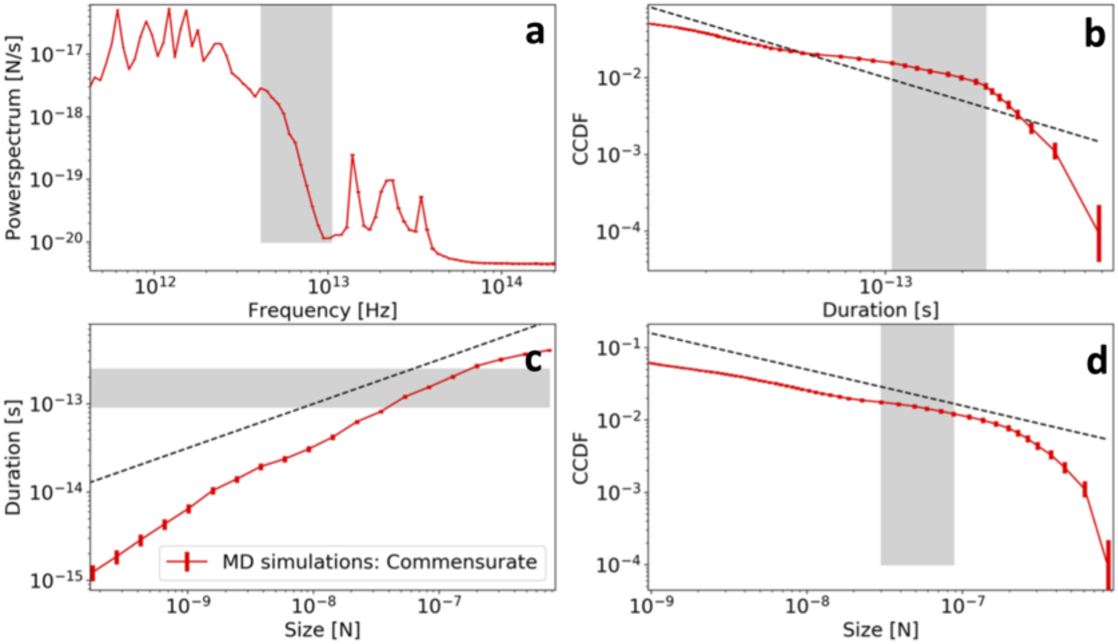

Figure. B.5: Defining the scaling regime in commensurate MD simulations. The power spectrum of the experimental friction force time trace **a)** and the avalanche duration vs. size **c)** are shown with error bar markers calculated from the standard error in a log-binned average over frequency and size, respectively. The duration **b)** and size **d)** CCDFs are plotted with error bar markers calculated with a Bayesian technique with 95% confidence interval [38]. Using the highlighted frequency regime in the power spectrum $[f_{\min}, f_{\max}]$, we then define the duration regime as $[1/f_{\max}, 1/f_{\min}]$ (highlighted in the duration CCDF), that is used in conjunction with the size vs. duration plot to fix the size regime $[S_{\min}, S_{\max}]$ (highlighted in the size CCDF). Mean field power-law predictions for infinite systems size with zero oscillations are plotted as black dashed lines.

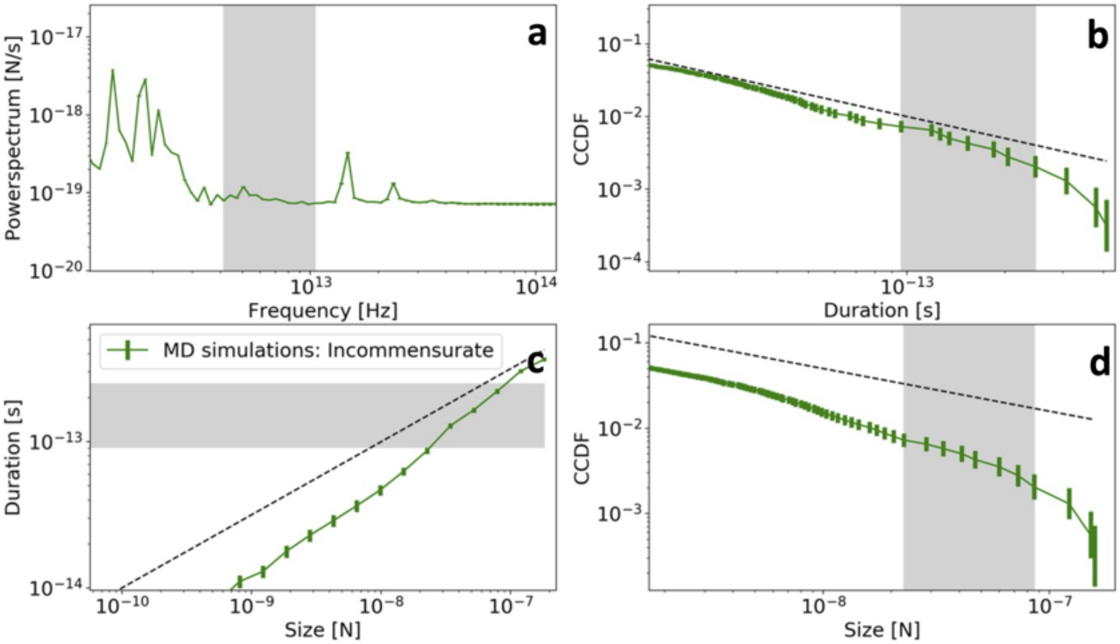

Figure. B.6: Defining the scaling regime in incommensurate MD simulations. The power spectrum of the experimental friction force time trace **a)** and the avalanche duration vs. size **c)** are shown with error bar markers calculated from the standard error in a log-binned average over frequency and size, respectively. The duration **b)** and size **d)** CCDFs are plotted with error bar markers calculated with a Bayesian technique with 95% confidence interval [38]. Using the highlighted frequency regime in the power spectrum $[f_{\min}, f_{\max}]$, we then define the duration regime as $[1/f_{\max}, 1/f_{\min}]$ (highlighted in the duration CCDF), that is used in conjunction with the size vs. duration plot to fix the size regime $[S_{\min}, S_{\max}]$ (highlighted in the size CCDF). Mean field power-law predictions for infinite systems size with zero oscillations are plotted as black dashed lines.

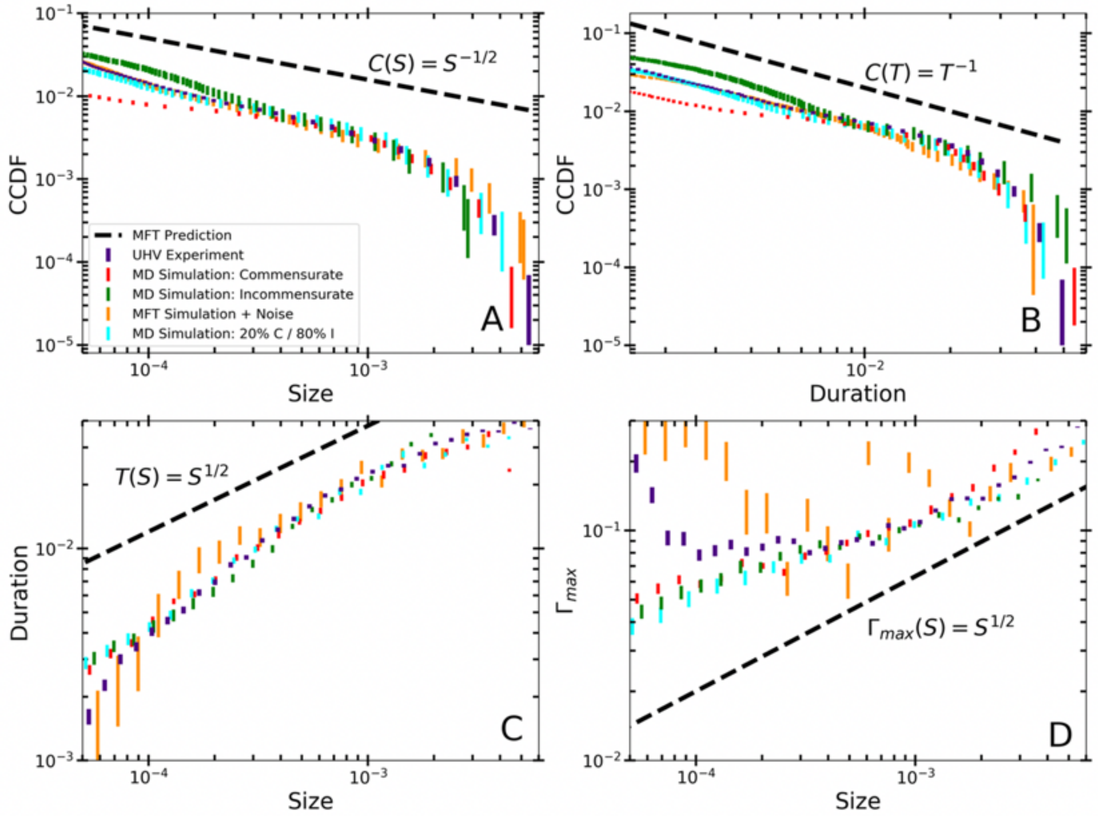

Figure. B.7: Demonstration of the comparison between results from a mixture of commensurate and incommensurate simulations and the run-in regime from the experiments. a) The complementary cumulative distribution function (CCDF; probability of having slip size S or larger) for UHV experiments during run-in (indigo), commensurate MD simulation (red), incommensurate MD simulations (green), mixed MD simulation with 20% commensurate / 80% incommensurate (cyan), and MFT simulations that were amended with noise (orange). **b)** CCDF of durations, $C(T)$. The cumulative distributions (a, b) are plotted with error bar markers calculated with a Bayesian technique with 95% confidence interval [38]. **c)** Distribution of fluctuation size vs. duration. **d)** Distribution of maximum force-drop rate ($\Gamma_{\max}$) during avalanches vs. avalanche size. (c, d) are shown with error bar markers calculated from the standard error in a log binned average over size. The power-laws predicted by the mean field model are plotted as black dashed lines (see Table B.1 in Appendix B). The x and y axis of the three simulations were rescaled to lie on the experiment, for comparison of the functional form only. Note that the power-laws are predictions for an infinite size system without noise or oscillations.

Appendix C

Appendix to Chapter 4

C.1 Constructing Shapes from Causal Web Data

To construct avalanche shape profiles, we usually start with the spiking rate (Γ), but Γ is not technically defined within the causal webs framework. However, each causal web can be condensed into a list of spike times. From this list of times, we can create something resembling a spiking rate by simply adding up all of the spikes that happened at each specific time. We show an example avalanche profile found using this method in Figure C.1,(a). We wish to point out here one unusual feature, which is that the time between two consecutive points in the avalanche profile is *not* constant (see Figure C.1,(a), inset). This is because in the causal web's method, unlike the traditional analysis (see Chapter 4), avalanches do not terminate when zero neurons spike in a given time step; the avalanche only terminates once there are no more future causal pairs connected to the cluster of causal pairs that is the avalanche. This results in uneven spaced points in the avalanche shape (see Figure C.1,(a), red), which at first sight seems concerning.

We'll explain now, why for the results of this thesis, at least, the aforementioned issue is of no concern. First, in this thesis, we focused mostly on the profiles of the largest duration avalanches. We show in Figure C.1,(b) that the average percentage of time steps in an avalanche that are larger than the minimum time step ($\Delta t_0 = 1\,\mathrm{ms}$) becomes insignificant for the longest avalanches. We confirm also that this is true for all samples (see Figure C.1,(b), inset). Lastly, even if this discrepancy was more prevalent in the data, the *averaged* avalanche profiles (the only profiles we have investigated in this thesis) would be unaltered by the occasional absence of a data point. For these reasons we are confident the conclusions on the avalanche shape profiles resulting from the causal webs analysis are valid.

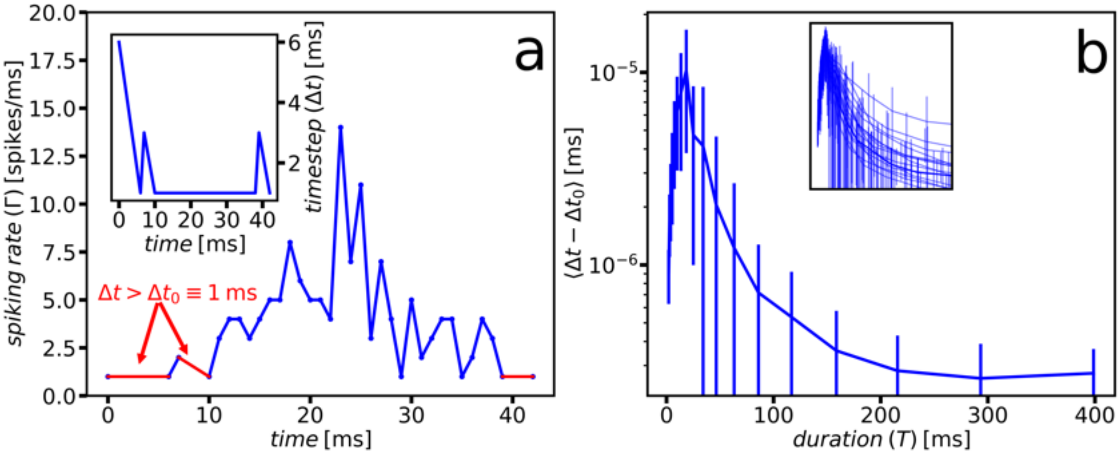

Figure. C.1: Constructing shapes from causal web data. a) A single avalanche profile as extracted by the causal webs method is shown in blue. We highlight segments of the shape with temporal spacing (Δt) greater than the acquisition time $\Delta t_0 = 1$ms. In the inset, we plot the temporal spacing (Δt) between each data point in the shape, to highlight that there are portions where the temporal spacing is larger that Δt_0. **b)** In each avalanche, we plot the average difference between the time step and the acquisition time, $\langle \Delta t - \Delta t_0 \rangle$, versus the corresponding avalanches duration. In the inset, we show the equivalent figure for all fifteen samples.

C.2 Estimating Critical Exponents in Data

To estimate the value of the critical scaling exponents τ and α, we use a Bayesian maximum likelihood estimator in the python library 'Power-law' (see [46] for more details). When the entire range of the data is used for the estimation (dashed grey line, Figure C.2(a)), the estimation is not very accurate (i.e., not representative of the potential power-law seen in the data). This is because the estimation considers sizes outside of the scaling regime. The 'Power-law' package has the ability to locate the scaling regime (Figure C.2(a), orange highlight), however, after which the estimation is much more accurate (Figure C.2(a), orange dashed line). To compare with the prediction of the scaling law of average size given duration, $\langle S|T \rangle$, vs. duration T, we do not need to use the Bayesian maximum likelihood estimator, and instead, linear regression works fine. However, when we limit the linear regression to the scaling regime, we again find a more accurate measurement (Figure C.2(b), grey vs. orange).

C.3 Bursting

We used the data sets openly provided in [74] to make empirical observations on the switching time distributions. The raw data (n = 15, mouse cortical slices above 512 electrode array) consists of labeled spike times (i.e., the seventh neuron spiked at time $t = 400 \rightarrow (7, 400)$. We could visualize this, for instance, as a 'raster plot' (Figure C.3(a)) where a single dot is shown for each labeled spike-time-tuple in the data set,

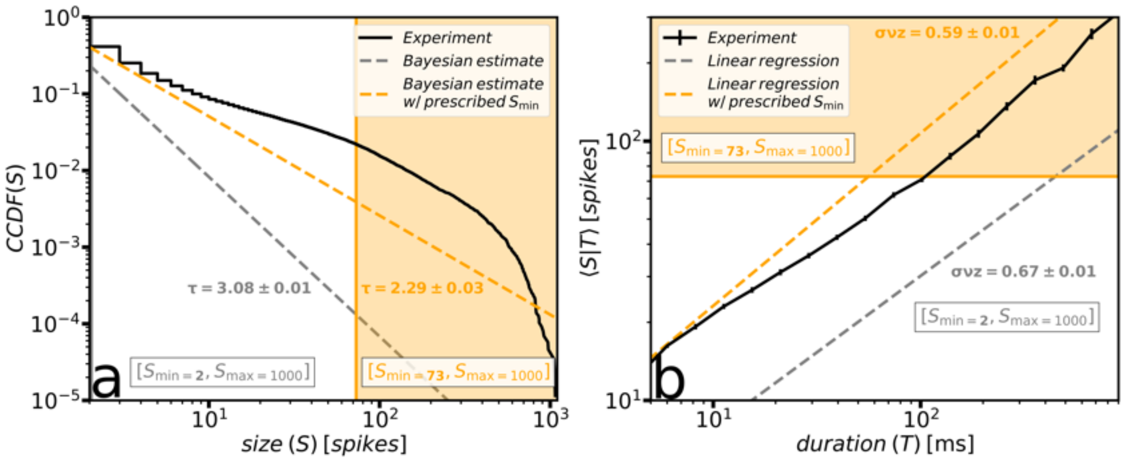

Figure. C.2: Estimating critical exponents. a) The $CCDF$ of avalanche size is shown for the experimental Sample 1 in black. The estimated power-law slope $\mathcal{C}(S) \sim S^{-(1-\tau)}$ is shown with a grey dashed line, as measured using the power-law package [46]. The orange dashed line shows the power-law estimated while imposing a minimum size on the scaling regime ($[S_{\mathrm{min}} = 73, S_{\mathrm{max}} = 1000]$, highlighted in orange). **b)** The avalanche size, averaged on duration ($\langle S|T \rangle$) is plotted in black versus duration for the experimental Sample 1. The estimated power-law slope $\langle S|T \rangle \sim T^{\sigma\nu z}$ is shown with a grey dashed line, as measured using linear regression. The orange dashed line shows the power-law estimated while imposing a minimum size on the scaling regime ($[S_{\mathrm{min}} = 73, S_{\mathrm{max}} = 1000]$, highlighted in orange).

located at x = spike time (e.g., 400) and y = neuron id (e.g., 7). In doing so, we observe (in all data sets) synchronous activity, known as 'bursting', where almost all the neurons in the system are spiking. These regions of activity are usually referred to as 'up-states' and the regions of relative inactivity between them are referred to as 'down-states'. We will be interested in characterizing how long the system dwells in this 'up-state' (i.e., 'seizure times').

To this end, we extract the start and end time of each labeled 'up-state' in the 1-hour long recording. Using a method derived from [136], we will first remove tonically spiking neurons (i.e., neurons that never seem to stop spiking, no matter what state the rest of the system is in). We perform this by simply removing neurons that spike more than a prescribed number of times (SpikeMax) in the 1-hour long recording. We note that this simply makes our algorithm more robust while not obstructing the true start vs. end times of an up-state. This is because the synchronous activity we seek to characterize is relatively insensitive to spatial subsampling, by this we just mean that since the synchronous 'up-state' is a global feature of the system, it is still very apparent even after subsampling (see Figure C.3(b)).

With this pre-processing of the data complete, we can cook up quite a robust algorithm for identifying the start and end times of the 'up-states'. Starting from the set of tuples, we can construct a global list of spike times t_i (e.g., ignoring neuron labels for now) and order it $t_{\mathrm{i}} \rightarrow \{t_{\mathrm{i}} | \forall i, j > i, t_{\mathrm{i}} < t_{\mathrm{j}}\}$. The next step is to coarse-grain this list (in time with bin-width Δt) such that we construct the sequence representing the global spiking-rate $\Gamma_{\tau}^{\Delta t} = \Sigma_{i=\tau}^{i=\tau+\Delta t} \delta_{\mathrm{i},\mathrm{t_i}}$ (Figure C.3(b), black line). We can use a simple algorithm (outlined in pseudo-code in Figure D.1) that picks out a start time (Figure C.3(b), vertical green lines) wherever the signal rises above some threshold, $\Gamma_\tau^{\Delta t} > T_{\mathrm{start}}$, and equivalently picks out the end time (Figure C.3(b), vertical red lines) when the signal subsequently crosses below a (separate) threshold, $\Gamma_\tau^{\Delta t} > T_{\mathrm{end}}$.

We point out that the algorithm isn't perfect, for instance in the second up-state of Figure C.3(b), we see the end time is overestimated. This means the threshold T_{end} wasn't perfect. The parameters $\mathrm{SpikeMax}, \Delta t, T_{\mathrm{start}}, T_{\mathrm{end}}$ are chosen to get the most accurate estimation of start and end times. We catalog the chosen parameters for each data set in Table C.1.

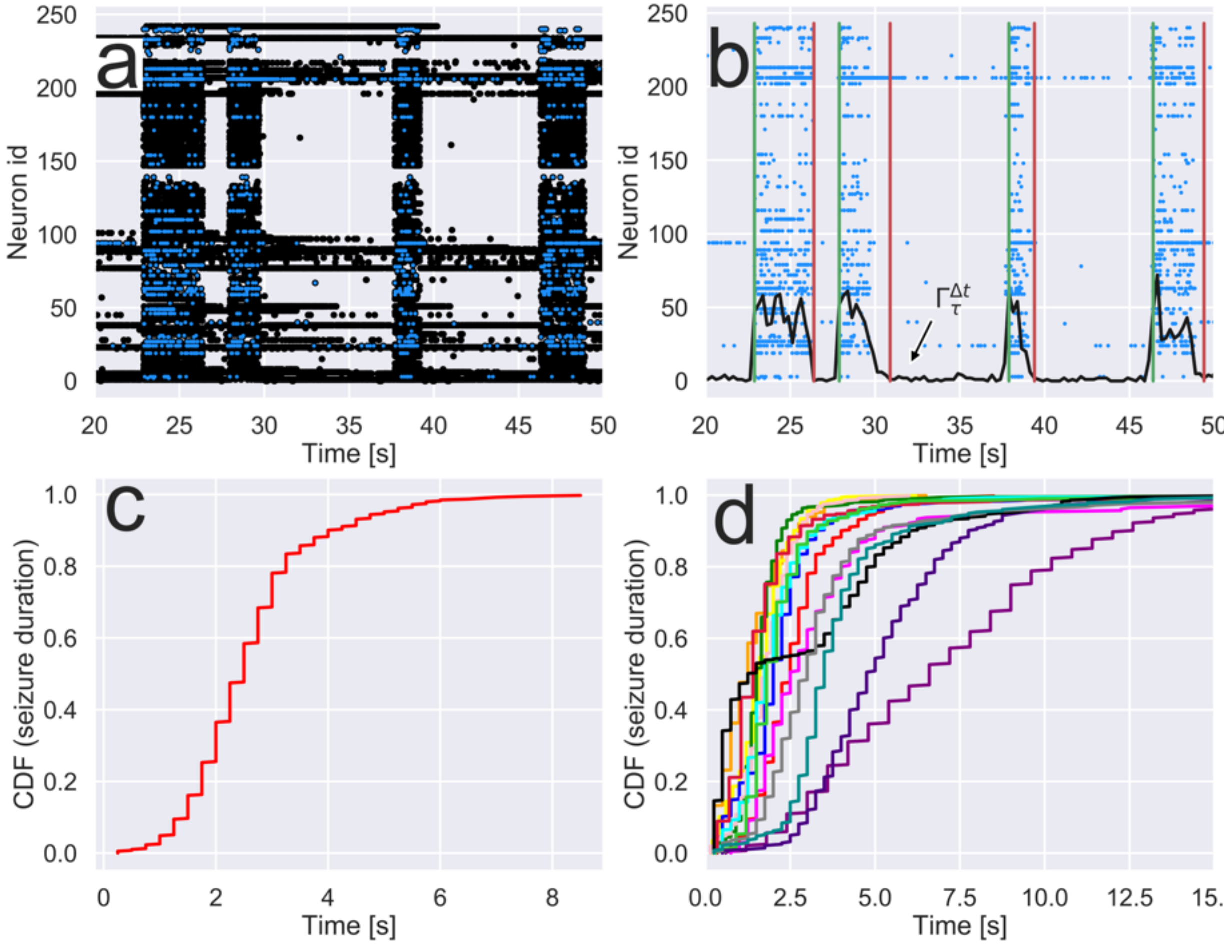

Figure. C.3: Extracting seizure times from data. a) With black dots we show the raster plot for spiking data taken in mouse cortex [120]. Blue dots show spikes only from neurons below some maximum average firing rate. **b)** Blue dots from (a) are reproduced, and the spiking rate for a given temporal bin width, Δt, is plotted as a black line. Green and red vertical lines show the start and end of the bursts, respectively, as identified by the algorithm described in Appendix C.3. **c)** The cumulative distribution function (CDF) of burst/seizure duration shown for a single sample. **d)** The cumulative distribution function (CDF) of burst/seizure duration shown for all fifteen samples, each shown in a different color.

Table C.1: Catalogue of parameters used in burst detection algorithm outlined in Appendix C.3

Data set	**Spike Max** [$\# \times 10^3$]	$\mathbf{\Delta t}$[s]	$\mathbf{T}_{\text{start}}$[#]	$\mathbf{T}_{\text{end}}$[#]
1	2.5	0.25	20	3
2	5	0.25	20	15
3	2.5	0.2	10	3
4	5	1.5	20	3
5	5	0.25	10	3
6	4	0.6	145	45
7	4	0.15	55	2
8	10	0.25	50	6
9	4	0.35	40	6
10	5	0.25	30	24
11	5	0.25	100	3
12	5	0.25	50	3
13	8	0.25	15	3
14	3	0.25	40	3
15	5	0.25	40	3

C.4 Burst/Avalanche Overlap

We show the avalanche duration $CCDF$ plotted in comparison to the burst/seizure duration CDF, in order to highlight overlap. The inset zooms in on this overlap, highlighted in orange, which quantifies how much the bursts are interfering with the avalanches. We measure the percentage of the avalanches that have durations larger than the minimum burst duration (i.e., those lying in the highlighted region of Figure C.5(a), '$\%DGB$') for each of the fifteen samples. In Figure C.5(b), we plot $\%DGB$' versus the average avalanche duration in the corresponding sample, to highlight that not all samples are effected by bursts in the same way.

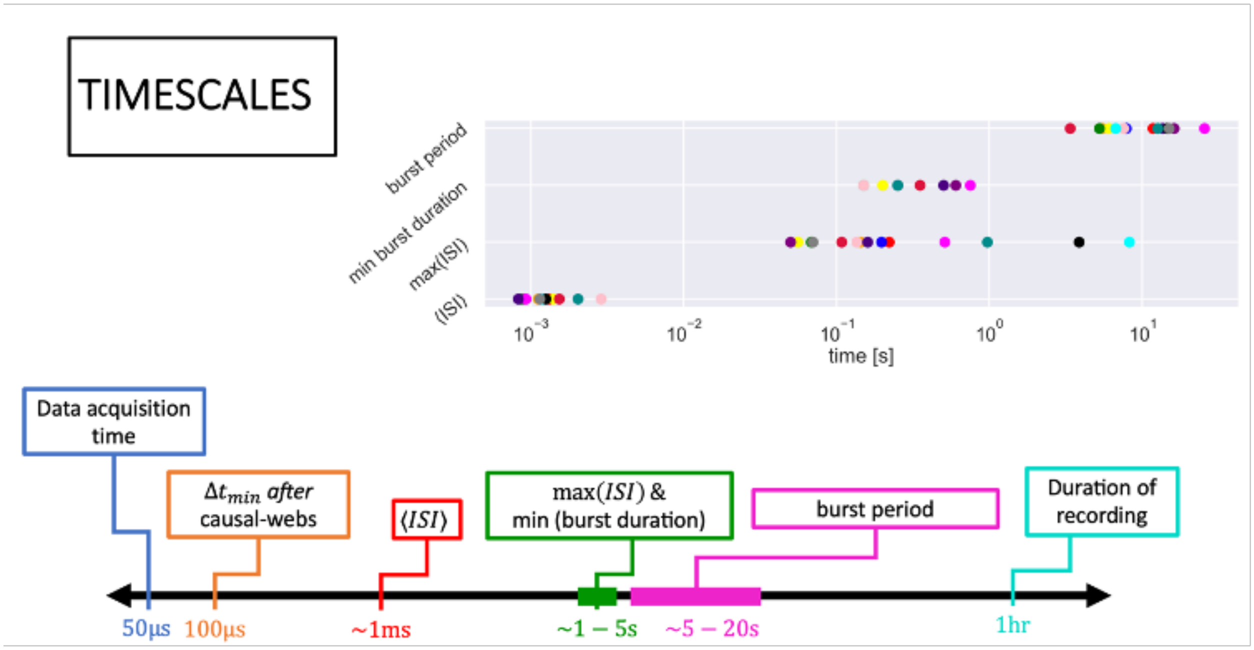

Figure. C.4: Timescales of the cortical neuronal spiking data used in Chapter 4, Chapter 5, Chapter 6 and Chapter 7.

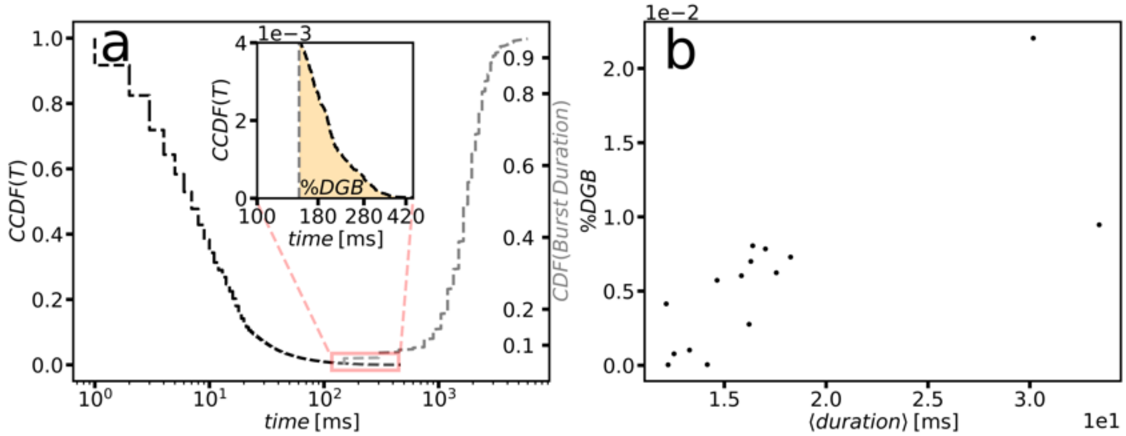

Figure. C.5: Avalanches overlap with bursts. a) The $CCDF$ of avalanche durations, T, is shown with a dashed black line for the experimental Sample 1. The CDF of burst durations is plotted with a dashed grey line. In the inset we zoom in on the portion outlined in red, showing the overlap with an orange highlight. **b)** The percentage of avalanches with durations greater than the minimum burst duration, $\%DGB$, for all fifteen samples versus the samples corresponding average duration.

Appendix D

Appendix to Chapter 5

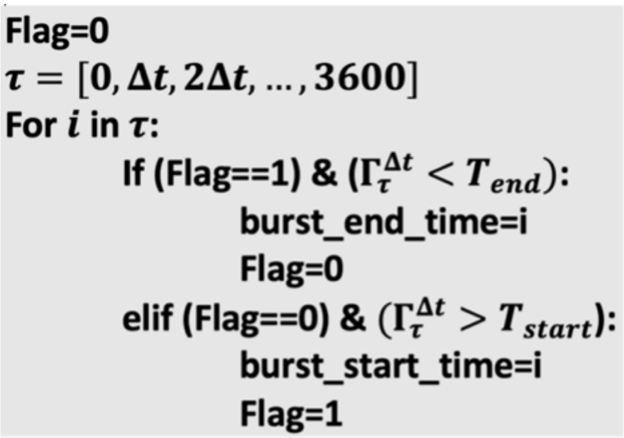

```
Flag=0
$\tau = [0, \Delta t, 2\Delta t, \ldots, 3600]$
For $i$ in $\tau$:
    If (Flag==1) & ($\Gamma_{\tau}^{\Delta t} < T_{end}$):
        burst_end_time=i
        Flag=0
    elif (Flag==0) & ($\Gamma_{\tau}^{\Delta t} > T_{start}$):
        burst_start_time=i
        Flag=1
```

Figure. D.1: Pseudo-code for algorithm.

D.1 Inhibitory Neurons Spike Proportionally to Excitatory Ones

In this section, we briefly discuss the validity of an assumption made in Chapter 5–that the averaged spiking rate of the inhibitory neurons scales linearly with the excitatory spiking rate. With recently discerned information about species labels in the data studied above, we compare the number of excitatory spikes in an avalanche to the number of inhibitory spikes in an avalanche (Figure D.2, left, 3 data sets used in main paper). Since there is a different number of excitatory than inhibitory neurons, we normalize the spike count in each avalanche by the number of excitatory neurons and inhibitory neurons in the system, respectively. We see a striking trend towards linearity and further confirm this with a bin average over the x-axis, for all n = 15 data sets (Figure D.2, right).

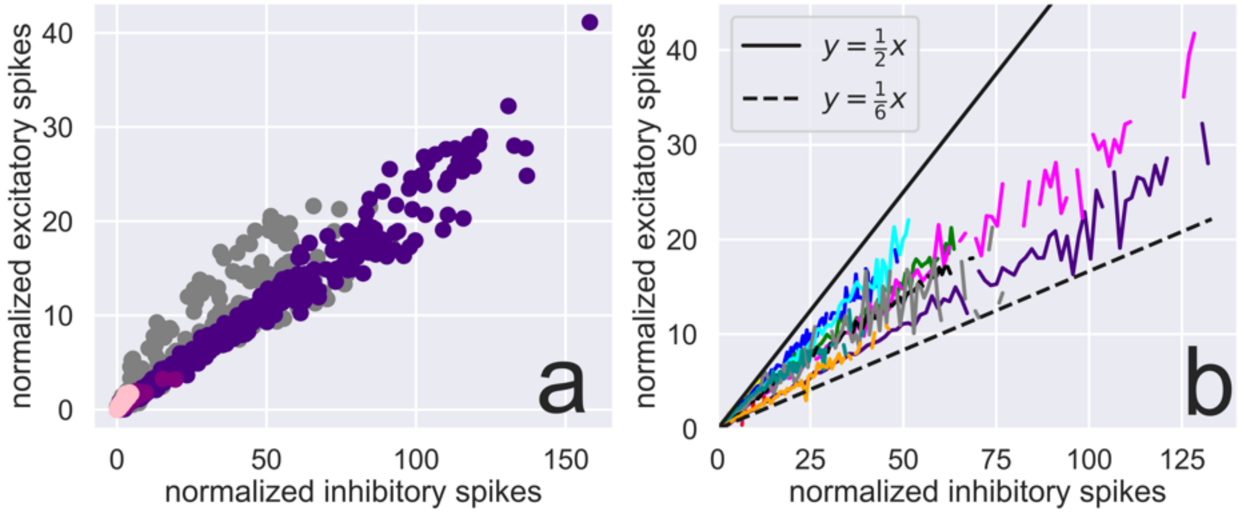

Figure. D.2: Inhibitory neurons spike roughly proportional to excitatory neurons on average. **a)** A scatter plot of normalized excitatory spikes vs. normalized inhibitory spikes, for the four data sets used in the main paper (Figure. 6.3). **b)** A recreation of the left with the x-axis binned to show an averaged trend), where different colors correspond to the different (n = 15) experiments.

Appendix E

Appendix to Chapter 6

E.1 Filtering for Criticality and Experimental Robustness

The experiments that show agreement with the predicted -1 power-law decay of avalanche size CCDF were used as critical data sets. Figure E.1 shows the critical data sets not used in the main paper. S vs. N curves for all critical mouse and rat data show the characteristic bending (Figure E.2). Firing time distributions for all experiments show multiple firings (recurrent activity) in both rat and mouse (Figure E.3). Cycle (Figure E.4) and Bi-test (Figure E.5) plots looks the same for all experiments.

E.2 Artifacts of Spike Signal Processing

Activation of a neuron in separate avalanches can be falsely identified as recurrent activation if caution is not taken during spike signal post processing. Traditionally, spike signals of continuous time are coarse grained into bins of width δt, after which an avalanche is defined as a set of consecutive nonempty bins (bounded by empty bins) [43]. With this definition, two simultaneous yet separate avalanches involving neuron i can be processed into a single merged avalanche. Therefore, the recurrent activation (of neuron i) is only an artifact of this merging. Causal webs processing avoids this ambiguity [44] whereby only causally connected strings of spikes are inferred to be avalanches. Accordingly, for a spike train analyzed via the traditional time binning method, we would expect to see more recurrent behavior as compared to the same signal analyzed via causal webs. Using a time bin width of 5 milliseconds we confirmed this hypothesis in the mouse data set (Figure E.6). Thus, this method helps to assure that we are reporting something physical and not just an artifact of post processing.

E.3 Inter and Intra-Avalanche Firing Timescales

Our model assumes that there is sufficient time between avalanches to allow for firing thresholds to be reset to their initial value after an avalanche ends, implying that typical intra-avalanche firing times are more tightly clustered than are inter-avalanche firing times. Here we investigate that assumption. From the mouse data set we extracted the inter-spike times of each neuron and calculated both the average time between spikes within the same avalanche and the average time between spikes that were not in the same avalanche. In Figure E.7 We plot the mean of each neuron below, with error bars showing the standard error of the mean. For each neuron, the average time between intra-avalanche spikes (red) is orders of magnitudes smaller than the average time between inter-avalanche spikes (green). This separation of the means was verified for each neuron with a hypothesis test that the two means were the drawn from the same normal distribution. We found the hypothesis to be rejected ($p \sim 0$) for all neurons in the system. Thus, the assumption of these separated timescales has been justified for our data. We note that causal webs' do allow for simultaneous avalanches– yielding the possibility of a neuron having negative inter-avalanche firing times. We show in magenta the averages with these negatives values removed and see that the change to the results are negligible.

E.4 Synaptic Time Delay in Simulations

Our model assumes that activity happening during an avalanche is much quicker than the time between consecutive avalanches i.e., that the neurons effectively communicate instantaneously. Here we us the model simulations to understand how our results change with the addition of a more biologically accurate synaptic time delay. In the original simulation protocol, the i^{th} neuron firing at time t affected all other neurons j at time $t+1$. We show in Figure E.8 the effect of delaying this until time $t+\tau$ for $\tau = 1, 2, 4, 20$ simulation steps. Though we find that increasing time delay decreases the avalanche size cutoff, the effect of recurrent activations persists as an upward bending of the S vs. N curve.

E.5 Sub-sampling

The experimental setup suggests that roughly one quarter of the total number of neurons in the cortex slice are being recorded by the MEA. This is not because the electrode density is too sparse to see all of the neurons; in fact the electrode density is so high, all of the neurons above it can be identified with very high accuracy through triangulation [120]. The other three quarters of the neurons actually come from the other three quarters of the cortex slice that is not directly above the MEA. So, we have neurons outside of the view ('outside' neurons) and neurons inside the view ('recorded' neurons). Not recording the entire cortical slice allows the possibility of an 'outside' neuron firing and causing a 'recorded' neuron to spike as well.

To account for this in our model, we performed a simple sub-sampling procedure on our simulation output that allowed us to emulate this experimental detail; namely that the experiment involved an entire slice but only recorded from a quarter of it. If our experimental system consisted of N 'recorded' neurons, we would simply simulate $4N$ neurons, but only consider firing events from every fourth neuron in the system. This way, when we construct the avalanches, neurons are truly receiving information from 'outside' neurons, which causes them to spike seemingly unprovoked. We have thus emulated the experimental situation of neurons from outside the field of view influencing neurons within the field of view, some times referred to as a 'sub-sampling' effect.

The results in the main paper (Figure 6.3,(a-b)) are shown for simulations that had this sub-sampling procedure performed and we can see again consistency between the data and experiment. We have also verified that this sub-sampling was not affecting the presence of the recurrent activity. We show in Figure E.9,(a) how the S vs. U curves change as we sub-sample. We find that the sub-sampling procedure enhances the recurrent activity in our simulations. We point out however that this enhancement through sub-sampling is not sufficient to account for the extent of the recurrent activity seen in the experiments alone (Figure E.9,(b)). We can reproduce experimental avalanche size $CCDF$'s and $S\,vs.\,U$ curves with our simulations with weakening alone or with weakening including sub-sampling, but not with sub-sampling alone. We also find the sub-sampling procedure can induce a clustered bi-test signature in simulations where it once was not present without sub-sampling, this trend is described in detail in the next section.

E.6 Bi-test for Model Simulations

Our simple slip model with weakening has been shown through large system- size simulations to admit two unique signatures related to temporal correlations between avalanches (see chapter 2 in [137]). Specifically, predictions about the intervals between consecutive avalanche start times can be made which admit a signature that differs between small and large avalanches. This signature, the bi-test of inter-avalanche start times (outlined in the manuscript), for small avalanches is either Poissonian when there is no weakening (Figure E.10, Blue), or temporally clustered if there is weakening (Figure E.10, Red). Though these signatures are seen in the bare MFT model, we can make the clustering vanish by making the width of the disorder distribution large enough (Figure E.10, Green). Further, once we add inhibitory neurons, the clustering signature is almost certain to vanish (Figure E.10, magenta). With both inhibitory neurons and high disorder width, as we need for our simulation to match the model, we also fail to recover clustered small avalanche inter-start times (Figure E.10, cyan).

Since the clustered signature seen in the data is not predicted by our simulations (i.e., to be a result of weakening), what causes it? We have investigated the effects of sub-sampling the simulations to more closely replicate the experiments. We find that, under sub-sampling of our simulations, the clustered bi-test signature for small avalanches is recovered (Figure E.11,). We were not able to get a sub-sampled simulation to all at once match the size distribution, S vs. U and the Bi-test, but we propose that sub-sampling is a possible reason for the appearance of the clustering signature in the data. We leave a perfect implementation of a sub-sampling scheme to future work, for this paper, we simply show that it is a possible explanation.

E.7 Disorder Configurations

Each time we run the simulation, a new set of arrest stresses is drawn from a uniform distribution. This means that each simulation will not have exactly the same evolution over time. Specifically, the stress configuration of the system can stay relatively uniform for the duration of the simulation, or it can evolve into a bunched state; see [5]. This can result in a handful of large events (significantly larger than the rest) obscuring the fit to experiment (Figure E.12, disorder realization trials $1, 3-6, 8-10$). Importantly the statistical features we are concerned with, for instance, the power-law slope of the avalanche size CCDF, do not change with these different realizations of disorder (Figure E.12). With that being said, the experiments have one realization of disorder, which we are a priori unaware of. For this reason, in the main paper (Figure 6.4), we highlight one disorder configuration that compares best with the experiment (disorder realization trial 2, Figure E.12).

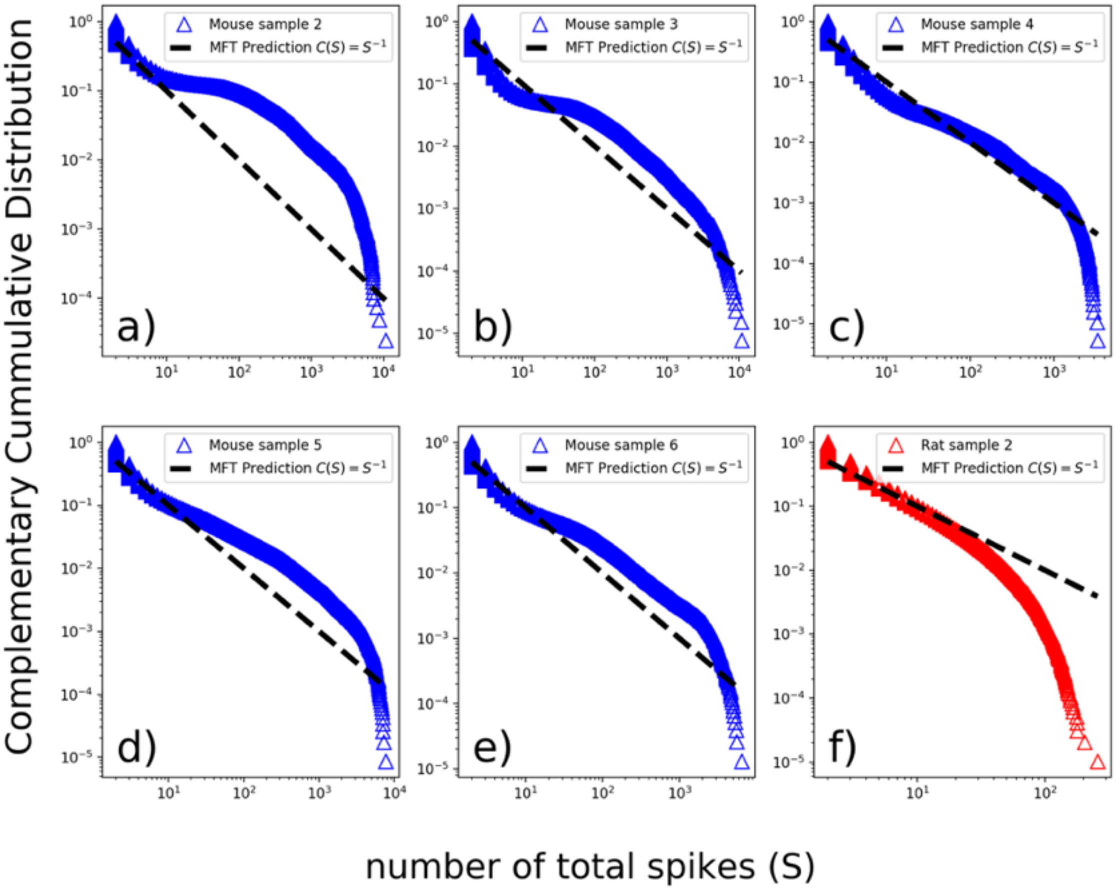

Figure. E.1: Complementary cumulative distribution function (CCDF) of avalanche size S. spike signals of cultured cortex recorded with a 512-electrode array and processed into causal webs were used to extract avalanche sizes (S). The $CCDF$ in size S (number of total spikes) is shown here for five critical mouse sets **(a-e)** and one critical rat set **(f)** not shown in the main paper. The critical distributions follow a power-law on a log-log scale with exponent minus one, shown here as a dashed line.

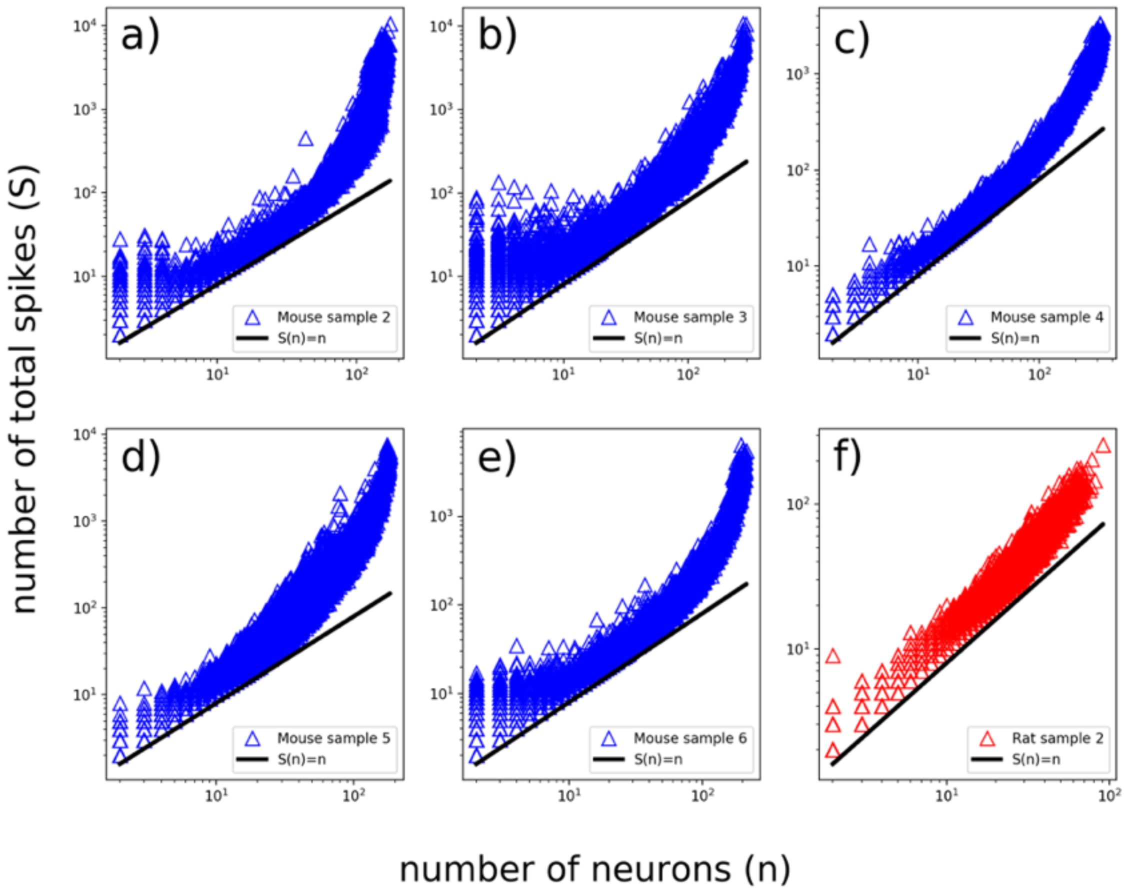

Figure. E.2: Bending in S vs. U. spike signals of cultured cortex recorded with a 512-electrode array and processed into causal webs were used to extract from each avalanche the number of unique neurons (U) and size (S). The two are plotted vs. each other above on a log-log scale, with the black line $S(U) = U$, shown here for five critical mouse data sets **(a-e)** and one critical rat data set **(f)** not shown in the main paper. Bending in S vs. U. spike signals of cultured cortex recorded with a 512-electrode array and processed into causal webs were used to extract from each avalanche the number of unique neurons (U) and size (S). The two are plotted vs. each other above on a log-log scale, with the black line $S(U) = U$, shown here for five critical mouse data sets (a-e) and one critical rat data set (f) not shown in the main paper.

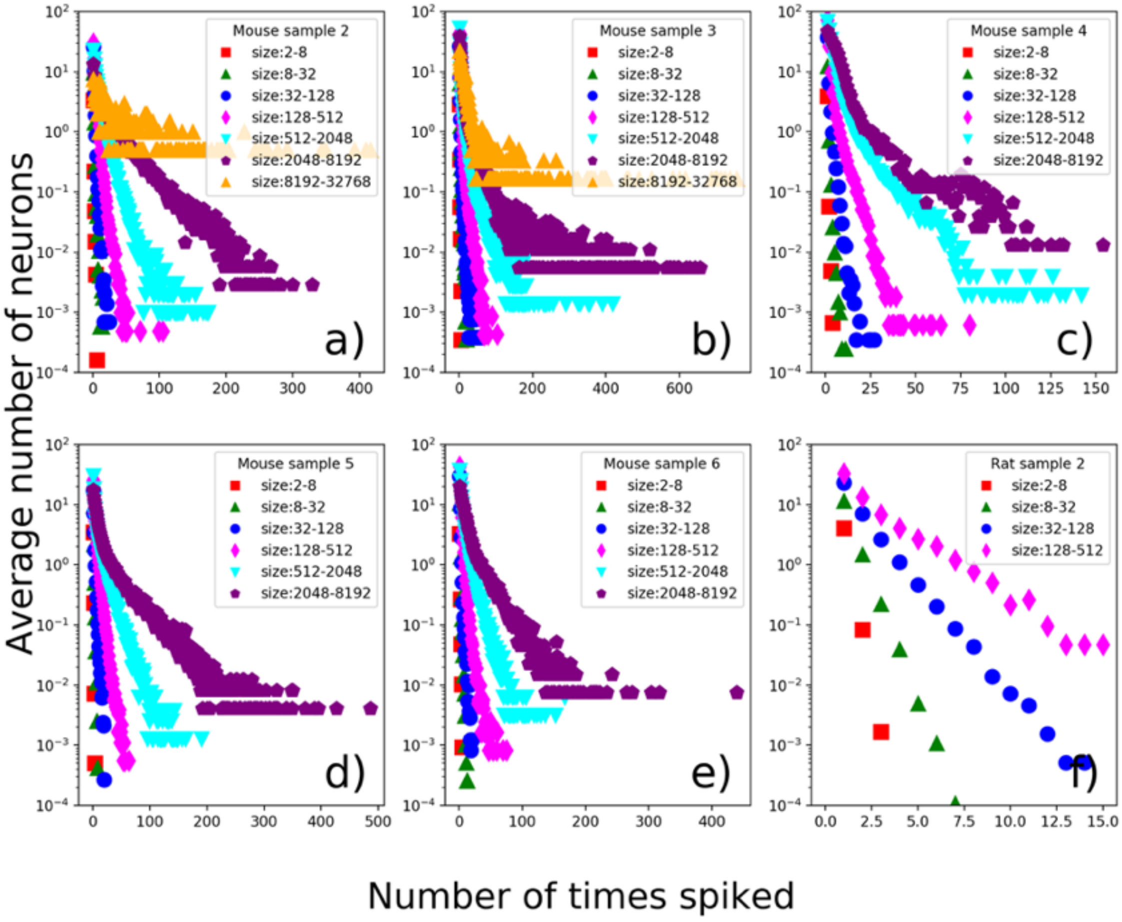

Figure. E.3: Distribution of multiple firing events. spike signals of cultured cortex recorded with a 512-electrode array and processed into causal webs were used to extract the distribution of the mean number of times each neuron spikes in an avalanche, for avalanches within specified size. They are shown here for five critical mouse data sets **(a-e)** and one critical rat data set **(f)** not shown in the main paper.

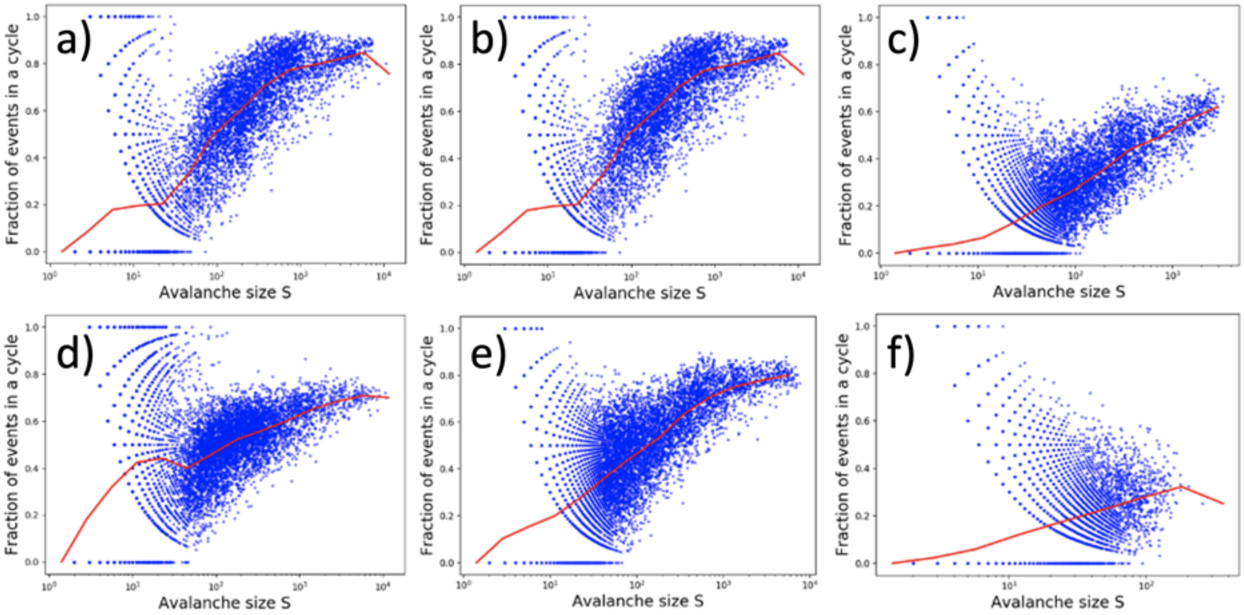

Figure. E.4: Cycles in avalanches. spike signals of cultured cortex recorded with a 512-electrode array and processed into causal webs were used to extract cycles of recurrent activity shown here for five critical mouse data sets **(a-e)** and one critical rat data set **(f)** not shown in the main paper.

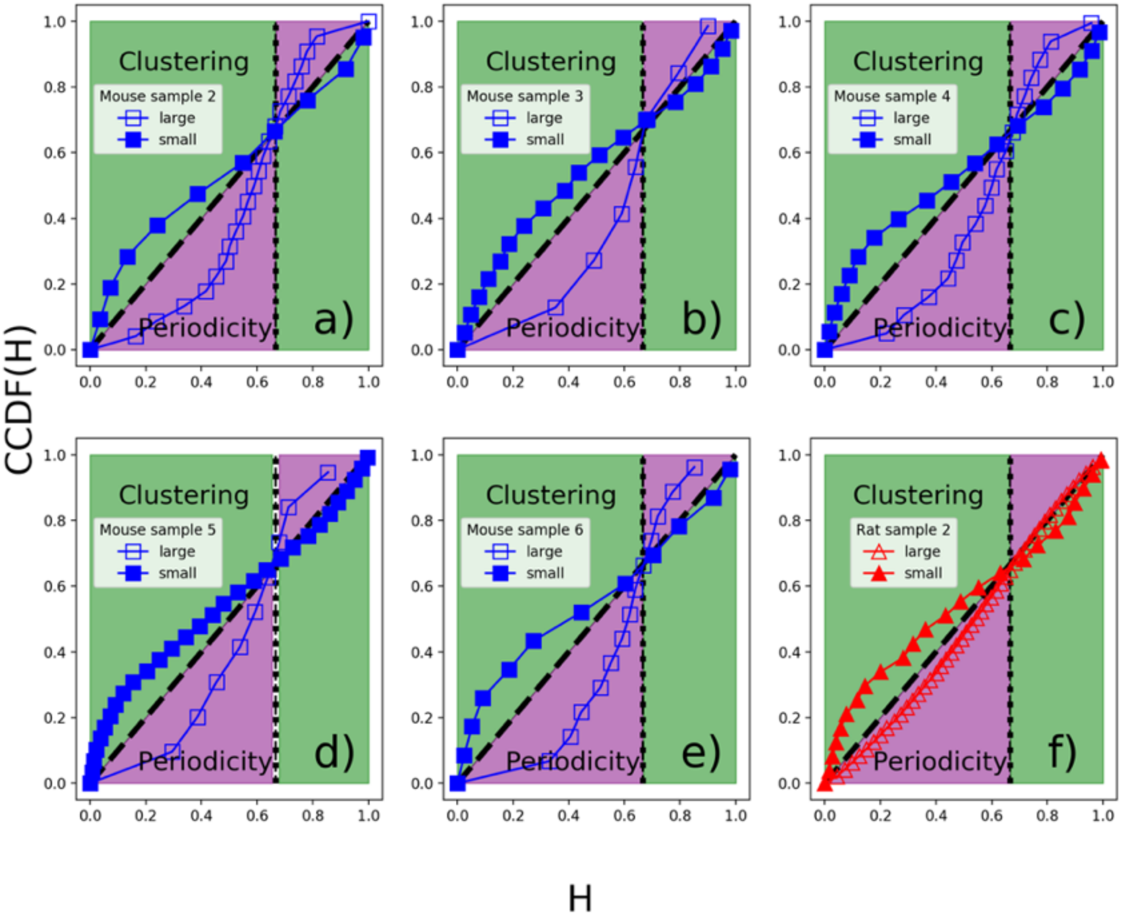

Figure. E.5: Bi-test. spike signals of cultured cortex recorded with a 512-electrode array and processed into causal webs were used to extract avalanche interevent-times. The cumulative distribution function of $H_{\mathrm{i}} = \Delta t_{\mathrm{i}}/(\Delta t_{\mathrm{i}} + \Delta\tau_{\mathrm{i}}/2)$ is shown here for large (empty markers) and small (filled markers) avalanches. Δt_{i} is taken as the minimal interevent time, $\min(|t_{\mathrm{i}} - t_{\mathrm{i}-1}|, |t_{\mathrm{i}+1} - t_{\mathrm{i}}|)$, and $\Delta\tau_{\mathrm{i}}$ is the subsequent interevent time ($|t_{\mathrm{i}-1} - t_{\mathrm{i}-2}|$ or $|t_{\mathrm{i}+2} - t_{\mathrm{i}+1}|$). Poisson distributed random interevent times fall directly on the dashed line, while perfectly periodic interevent times fall on the dotted line. The green regime is a signature of temporal clustering while in the purple regime is a signature of temporal quasi-periodicity. They are shown here for five critical mouse data sets (**a-e**, blue, square) and one critical rat data set ((**f**), red, triangle) not shown in the main paper.

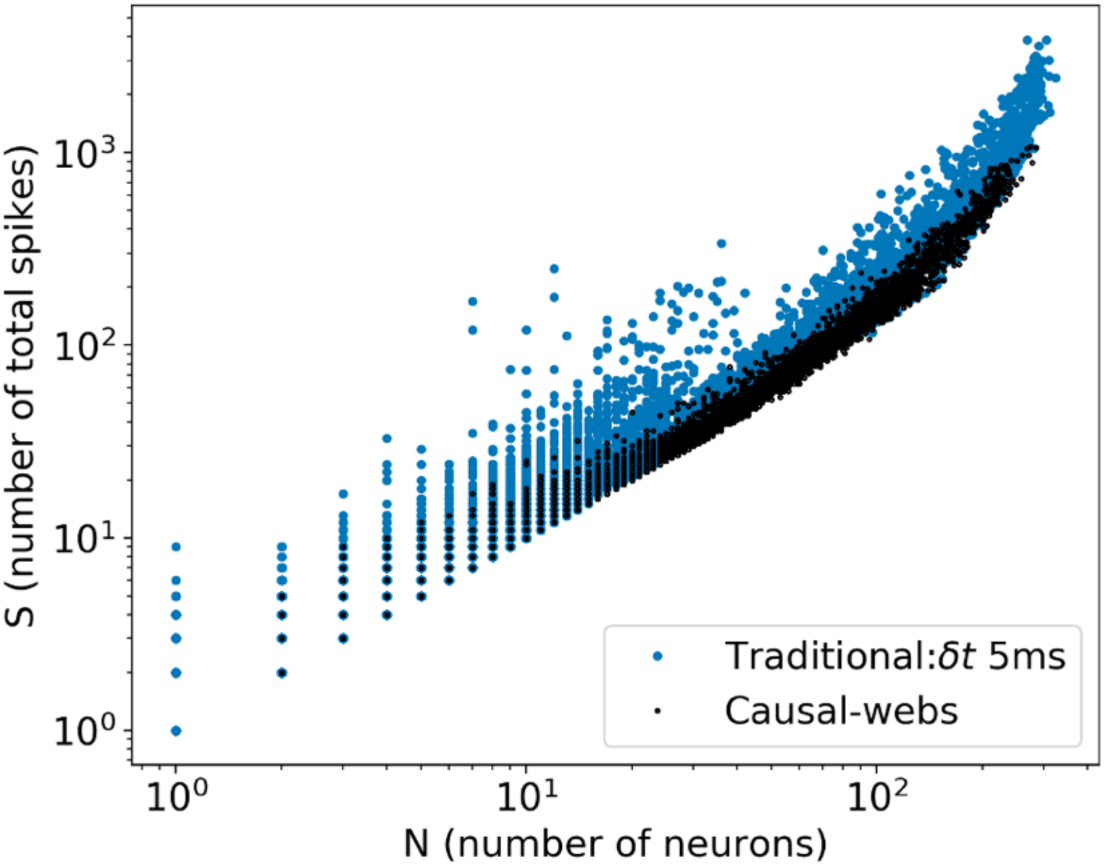

Figure. E.6: Traditional avalanche definitions over-emphasize recurrent firing. spike signals of cultured mouse cortex recorded with a 512-electrode array processed two ways: (blue) sorting spikes into time bins of width δt, defining an avalanche as any consecutive set of non-empty bins bounded by empty bins [43]. (black) sorting spikes into causally connected space (neuron)-time pairs, defining an avalanche as any string of causally connected spikes [44]. Avalanche sizes are plotted versus the number of unique neurons that spiked in the avalanche.

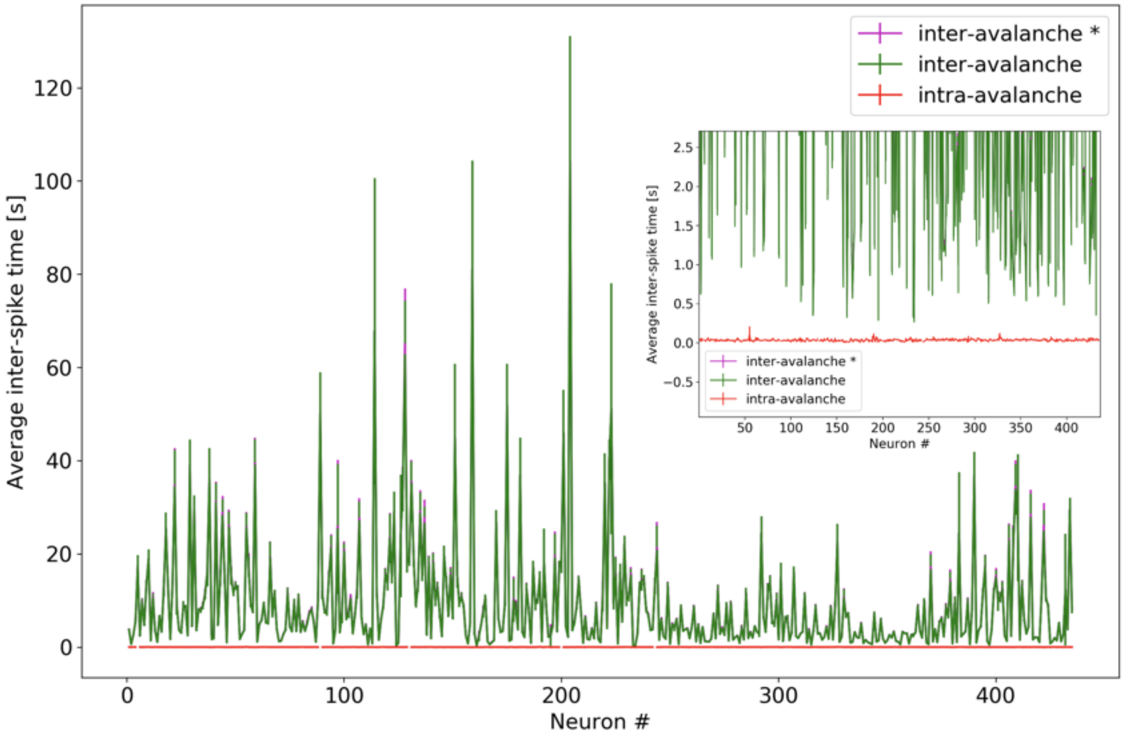

Figure. E.7: Neurons spike more frequently within an avalanche. Single neuron average inter-spike time within avalanches (intra, red) and between consecutive avalanches (inter, green) are plotted for all 434 neurons in the mouse cortex sample. Error bars show the standard error of the mean. The average with negative values removed is plotted in magenta and labelled as 'inter-avalanche*'. (inset) y-axis zoom in to show separation of timescales.

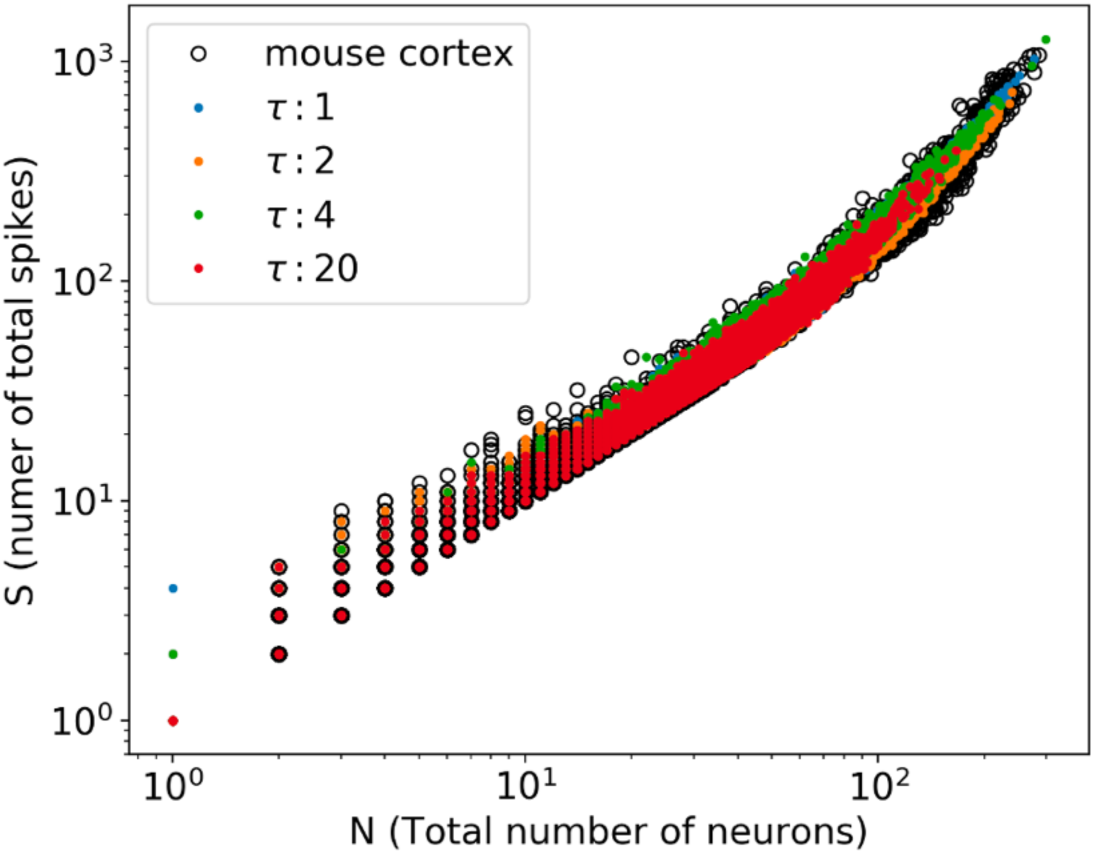

Figure. E.8: Synaptic time delays do not hinder recurrent activity in model simulations. The number of spikes (S) is plotted versus total number of neurons (U) for different synaptic time delays $\tau = 1, 2, 4$ and 20.

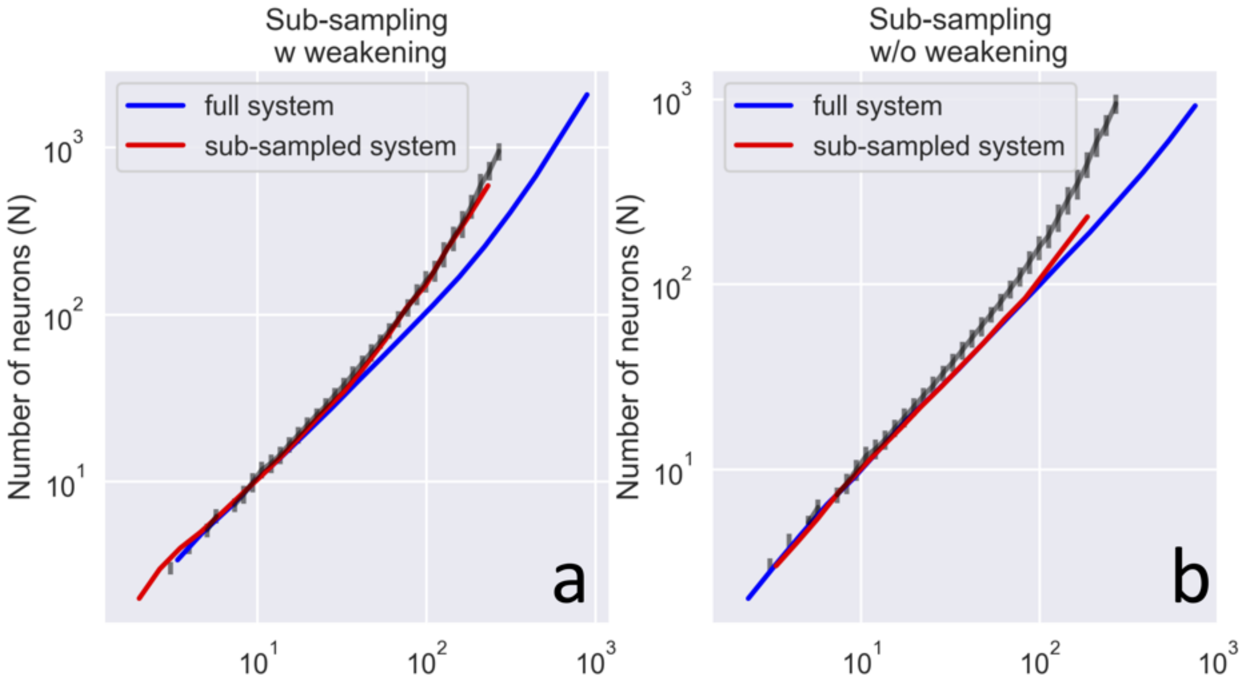

Figure. E.9: Clustered Bi-test signature is restored by sub-sampling. Simulations of the mouse system ($N = 434$) were ran with $w = 1.9$, $\epsilon_{\mathrm{e}} = 0.5$ $\epsilon_{\mathrm{i}} = 0.7$, $\gamma = 4.8$. **a)** simulations with nonzero weakening show the sub-sampled system has good agreement with experiment (red line). **b)** the sub-sampled system without weakening does not replicate the data nearly as well as the simulations with weakening (red line). From this figure we conclude that sub-sampling alone cannot account for the recurrent activity seen in the data.

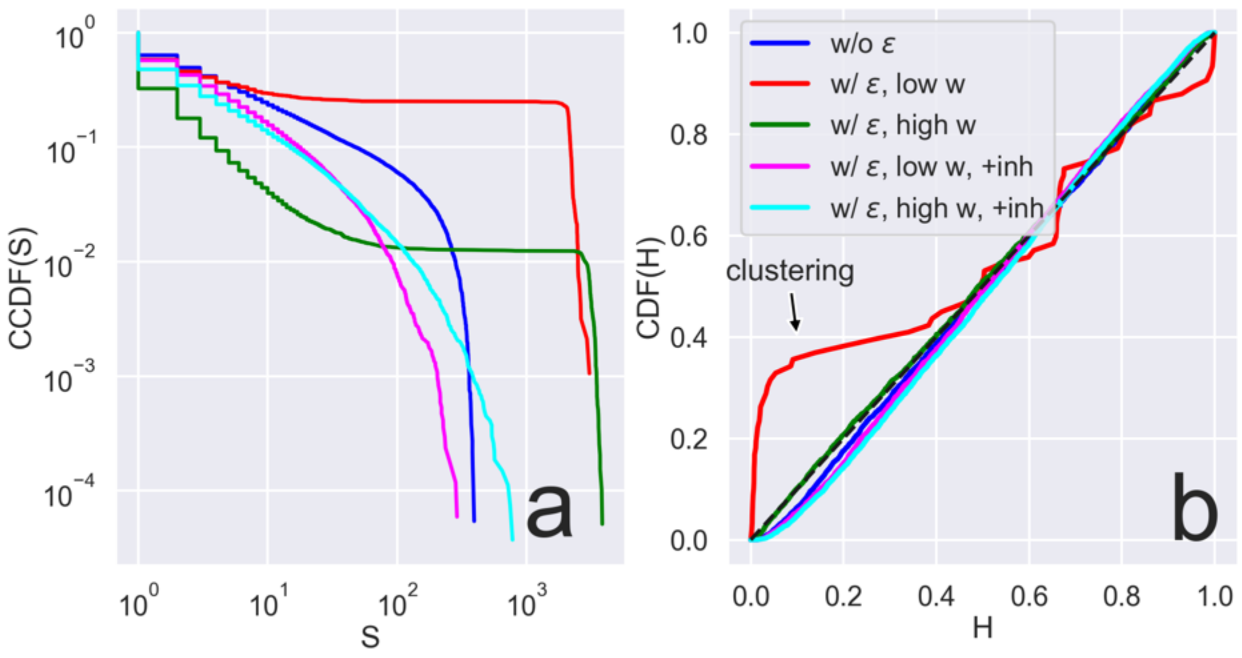

Figure. E.10: Clustered Bi-test signature is destroyed by disorder and inhibition. Simulations of the mouse system (N=434) were run with low disorder ($w = 0.1$), $\epsilon_e = 0.5$ $\epsilon_i = 0.7$, $\gamma = 4.8$. **a)** With out any weakening, the small avalanches show a near Poissonian bi-test signature (blue). **b)** With the addition of weakening, the signature resembles clustering (red). However, the addition of strong disorder ($w = 1.9$, green) or inhibitory neurons (magenta) or both (cyan) leads to the destruction of this clustered signature. From this investigation we provide a possible explanation for why the clustered signature is not present in our model simulations that include inhibitory neurons and large disorder.

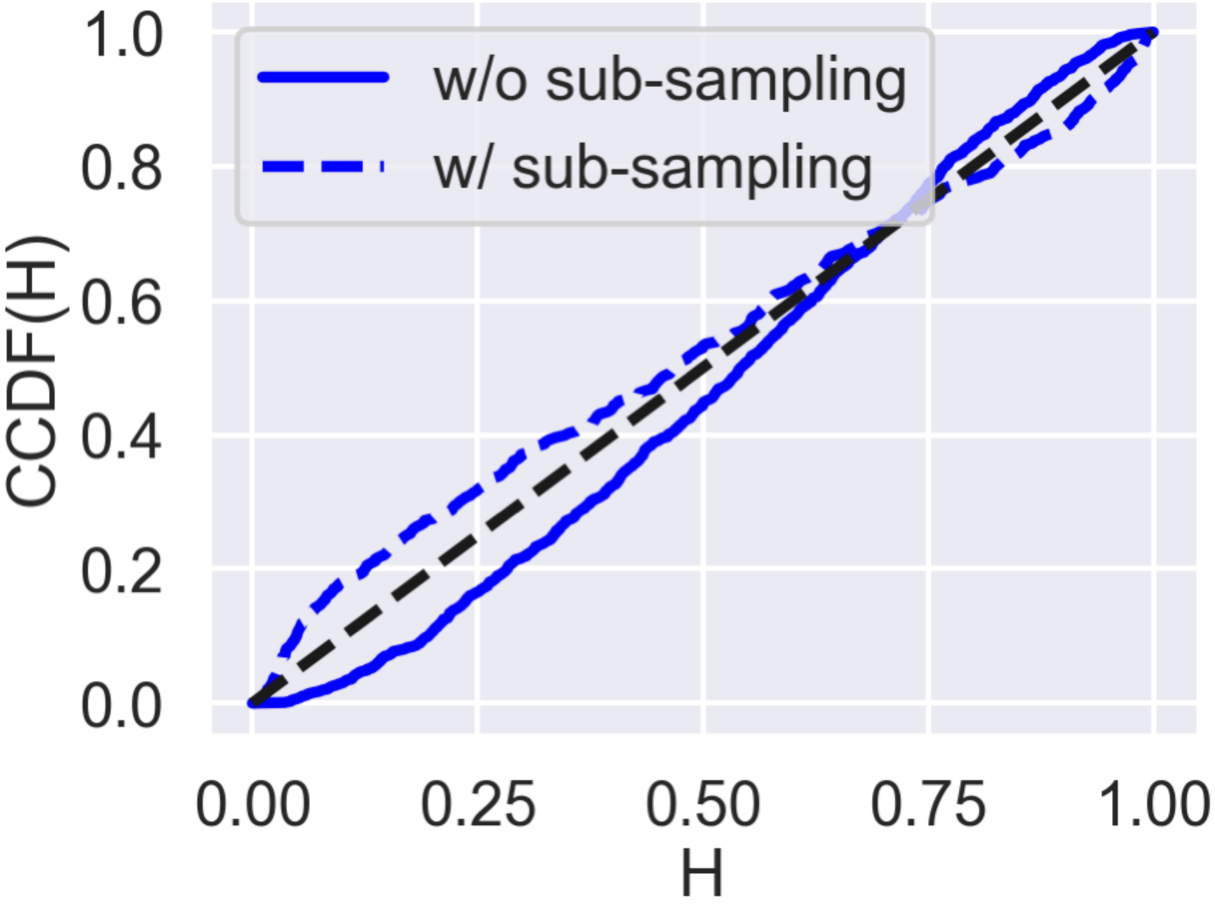

Figure. E.11: Clustered Bi-test signature is restored by sub-sampling. Simulations of the mouse system ($N = 434$) were run with $w = 1.9$, $\epsilon_\mathrm{e} = 0.5$ $\epsilon_\mathrm{i} = 0.7$, $\gamma = 4.8$. The solid line shows the full simulation (with $4N$ neurons) and the dashed line that same simulation after being sub-sampled by a factor of 4. From this figure we conclude that sub-sampling could be partially responsible for the clustered signature we see for small avalanches in the data.

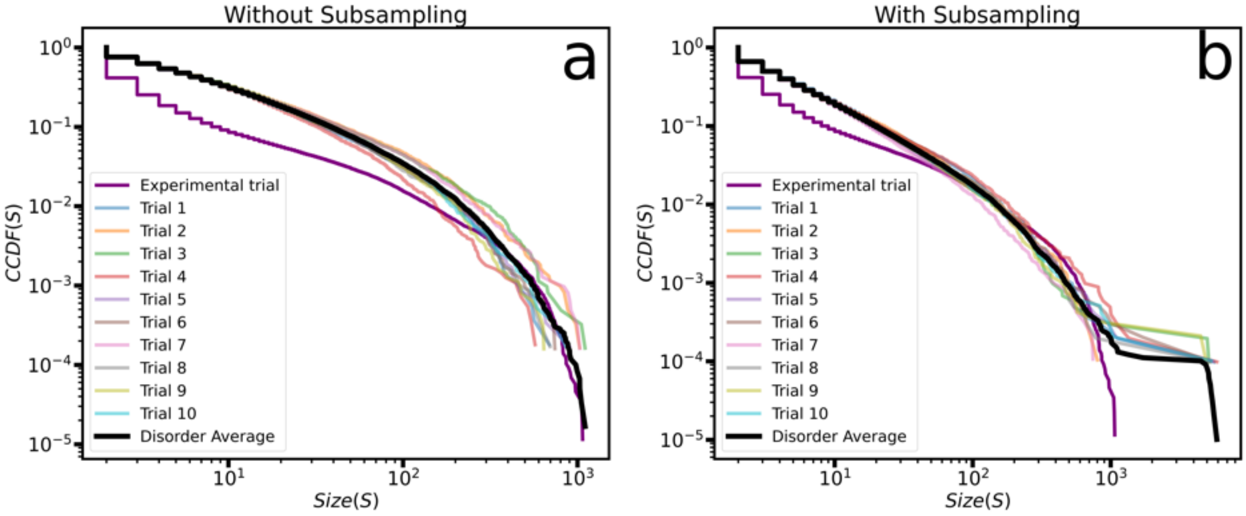

Figure. E.12: Disorder average. We show the $CCDF$ of avalanche size for 10 different realizations of disorder (or 'trials') in varying colors. The experimental trial from the main paper, mouse, is shown in purple and the disorder averaged $CCDF$ is shown in black. **a)** shows the $CCDF$ with out sub-sampling, **b)** shows the $CCDF$ with sub-sampling. All configurations on the left match the tail of the distribution fairly well, along with the disorder average. There is discrepancy for smaller sized avalanches though, and this discrepancy seems to relax with sub-sampling (right figure). The trade off is that there are sometimes large avalanches in the sub-sampled (and accordingly quadrupled in size) system, as explained in the Appendix E. Still, an individual simulation trial (trials 2 and 7, for instance) compare quite well to one experimental trial.

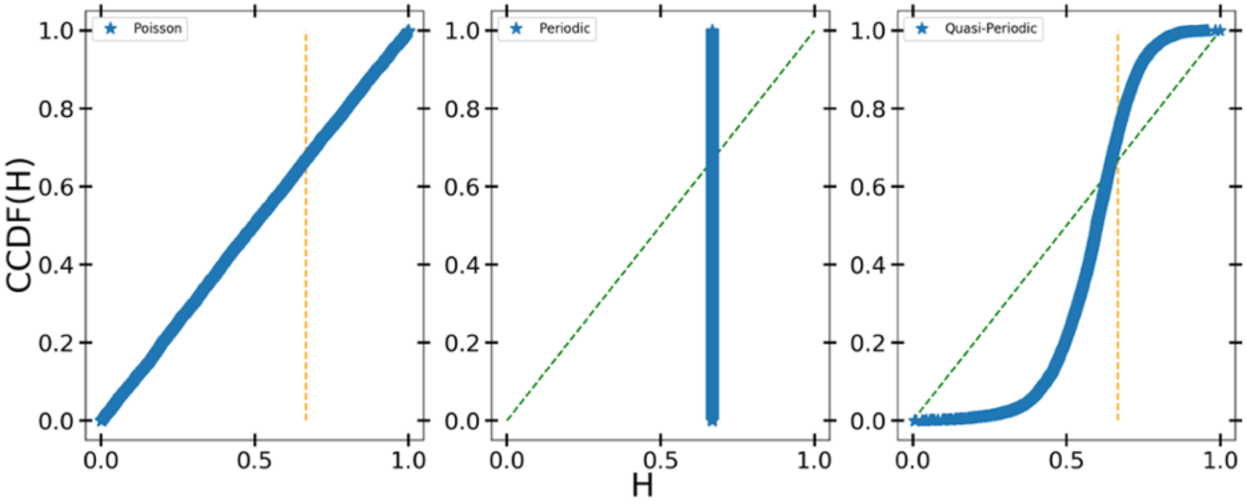

Figure. E.13: Calibrate the Bi-test. The cumulative distribution function of $H_i = \Delta t_{\rm i}/(\Delta t_{\rm i} + \Delta\tau_{\rm i}/2)$, where $\Delta t_{\rm i}$is taken as the minimal interevent time, $\min(|t_{\rm i} - t_{{\rm i}-1}|, |t_{{\rm i}+1} - t_{\rm i}|)$, and $\Delta\tau_{\rm i}$ is the subsequent interevent time ($|t_{{\rm i}-1} - t_{{\rm i}-2}|$ or $|t_{{\rm i}+2} - t_{{\rm i}+1}|$) shown here for example Poissonian (left), periodic (center) and quasi-periodic (right) spike-trains. For applications of the Bi-test to plasticity see [125].

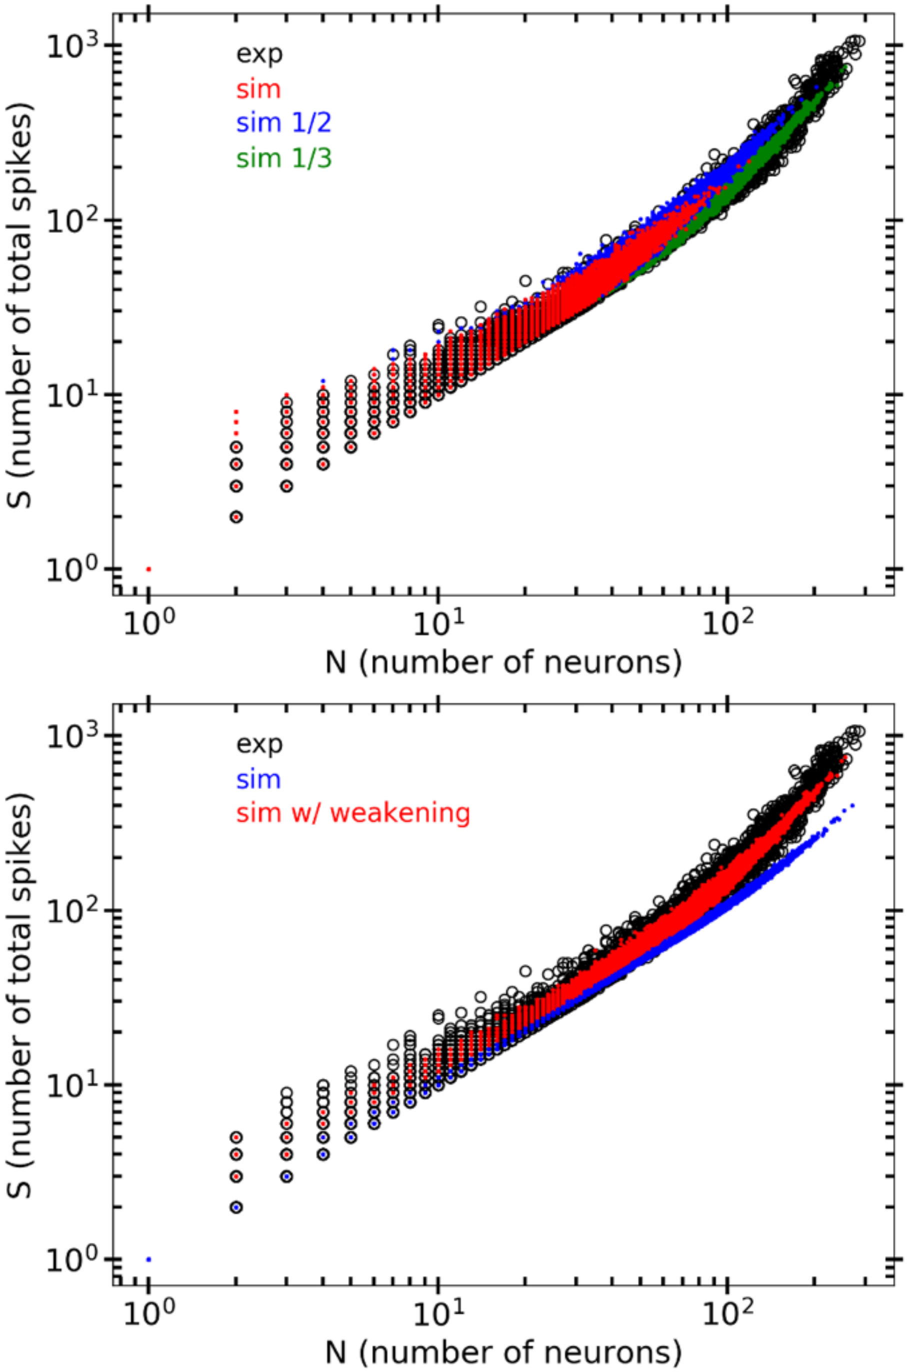

Figure. E.14: Reducing the strength of inhibitory couplings and adding dynamic weakening. Simulations of the mean field model were run with 108 inhibitory neurons and 326 excitatory neurons connected by a weight J_{ij} drawn from a log normal distribution $P(J_{\mathrm{ij}}) = \frac{1}{J_{\mathrm{ij}}\sqrt{2\pi N_{\mathrm{t}}}} e^{\frac{-ln(J_{\mathrm{ij}})^2}{2N_{\mathrm{t}}}}$, where N_{t} is the total number of neurons and a factor of -1 is added when neuron i and j are not of the same type. Avalanche sizes are calculated, along with total number of participating neurons, shown here plotted against each other on a log-log scale. Spike signals of cultured cortex recorded with a 512-electrode array and processed into causal webs were used to extract from each avalanche the number of unique neurons (n) and size (S) shown in black circles.(Top) Non-symmetric coupling: The excitatory to inhibitory weights were reduced by a factor of one (red dots) one half (blue dots) and one third (green dots).(Bottom) Weakening needed for bending in model: The model simulations are plotted with (red dots) and without (blue dots) dynamical weakening.

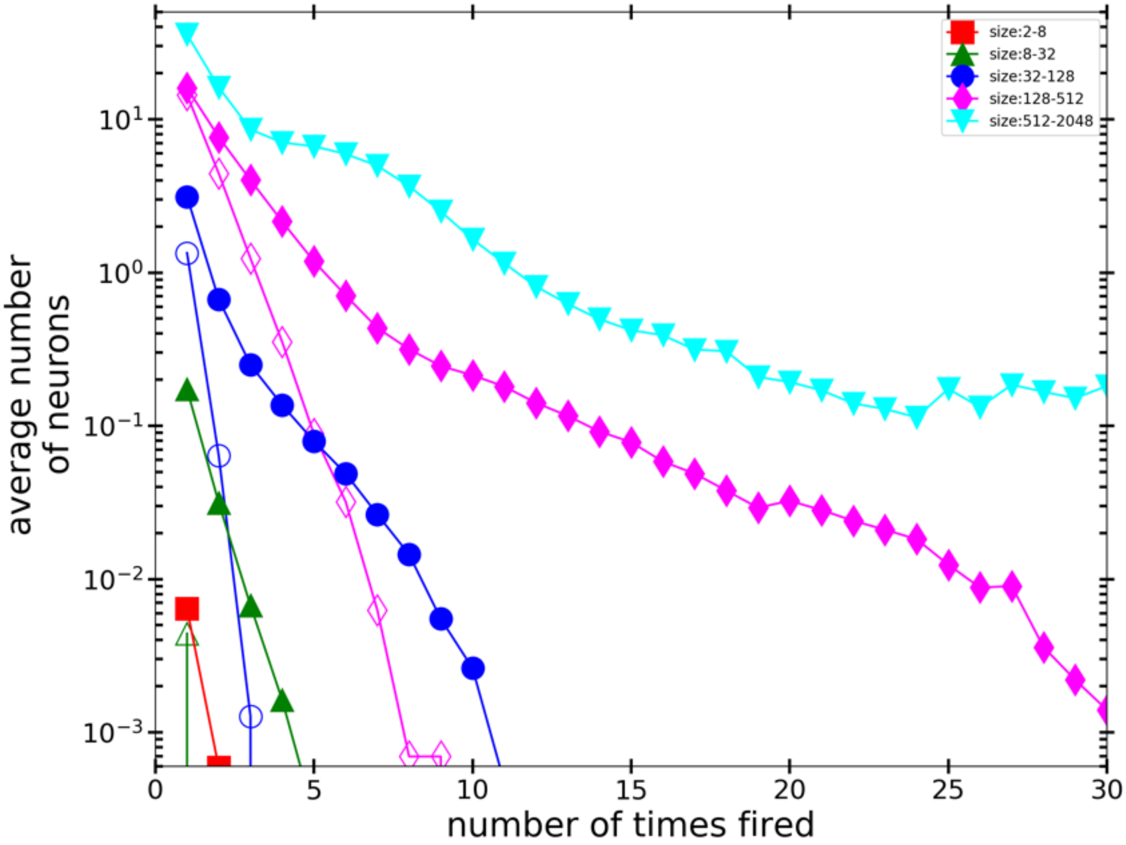

Figure. E.15: Weakening in simulated firing time distributions. Firing time distributions from simulations of the mouse sized system are shown with (filled markers) and without (empty markers) dynamical weakening.

Appendix F

Appendix to Chapter 7

F.1 Percentage of Inhibitory Spikes in an Avalanche

To investigate whether the avalanches of Chapter 7 were really dominated by inhibitory activity, in each avalanche we measured the percentage of spikes that were emitted from inhibitory spikes, % inh. spikes. We plot a distribution of these values, for experimental Sample 2, in Figure F.1(a), where we see that the distribution is dominated by the values zero, one-half, and one. This is coming from very small avalanches, of only a few spikes, such that there can be no inhibitory spikes, all inhibitory spikes, or exactly one-half inhibitory spikes. To parse out the features of the larger (and more recurrent) avalanches, we will construct the PDF again, but ignoring avalanches of size S less than some prescribed amount. Specifically, in Figure F.1(a), we do this for five particular values of S_{min}, showing that the larger avalanches do indeed have a bias toward inhibitory domination.

F.2 Inhibitory Neurons are the 'Brakes of the Brain'

As a first test of the idea that the inhibitory neurons could be turning on *after* the excitatory neurons– resembling a 'brake' of the brain– we plotted an average avalanche profile for the longer duration avalanches (Figure F.1(b), black). We separated the spikes coming from inhibitory and excitatory neurons and plotted their separated profiles in purple and green, respectively. From this first glance, it seems that the inhibitory neurons do seem to be responsible for the left-ward skew, but it is hard to tell if this is just because they are overwhelmingly active in the avalanche compared to the excitatory species. More investigation will need to be done in order to parse out the details of this hypothesis, but they are left as future work.

F.3 Unidentified Species

The method used to identify the species of each neuron in the cortical mouse data was not always successful. There were some neurons in each data set that were not able to be identified as excitatory or inhibitory, and were labeled as 'unidentified'. We show in Figure F.1(c-d) that results of Chapter 7 are not significantly different when we ignore the unidentified neurons (black) compared to when we assume all unidentified neurons were actually inhibitory (red) or even when we assume the unidentified neurons were all actually excitatory (blue). Quantitative changes may be present between the three options, but qualitatively the story is identical; scaling of the species differentiated the difference of recurrent activity (signaling inhibitory dominance) and a rough increase of the average total recurrent activity versus inhibitory ratio for the fifteen samples.

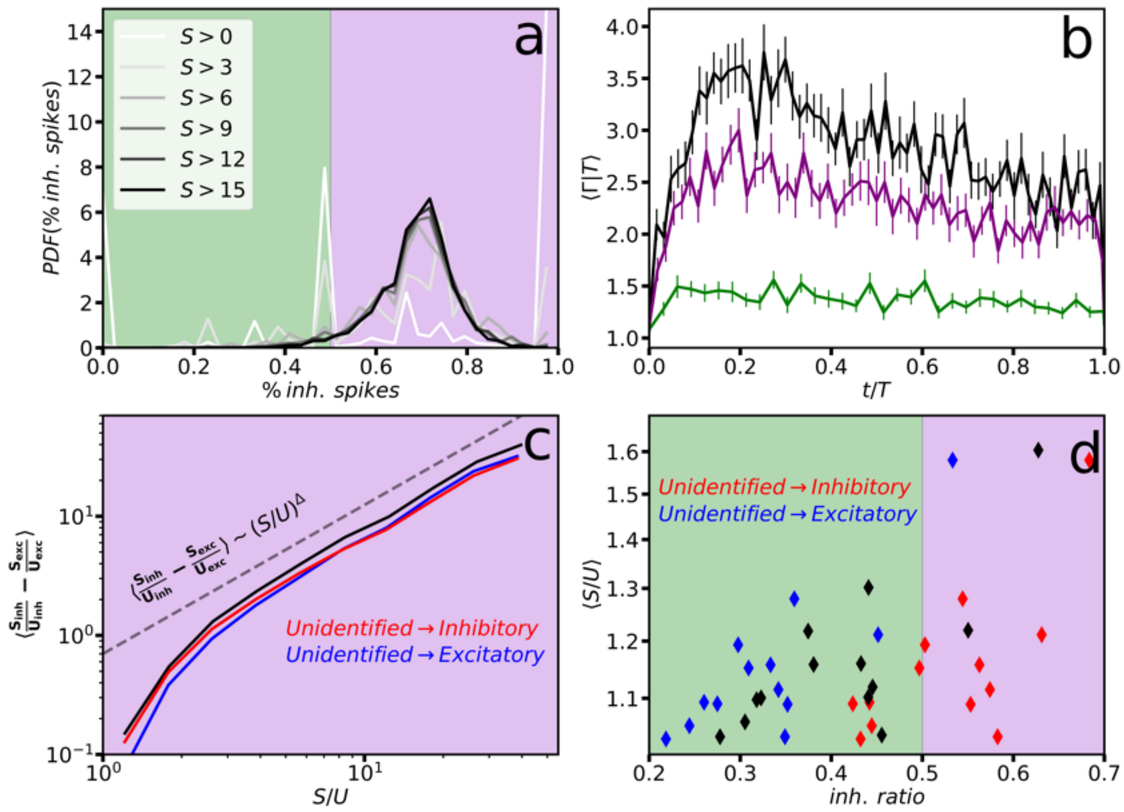

Figure. F.1: Verifying miscellaneous features in the data. **a)** The PDF of the percentage of spikes in an avalanche emitted from an inhibitory neuron, S_{inh}/S. The PDF is constructed, omitting avalanches smaller than a minimum, S_{min}, for $S_{\text{min}} = 0, 3, 6, 9, 12$ and 15. **b)** An average profile of a longer duration avalanche in experimental Sample 1, shown in black. The avalanche is separated into excitatory and inhibitory spikes, and the separate profiles are plotted in green and purple, respectively. Error bars show the standard error of the mean. **c)** The average difference of species differentiated recurrent activity, $\langle \frac{S_{\text{inh}}}{U_{\text{inh}}} - \frac{S_{\text{exc}}}{U_{\text{exc}}} \rangle$, is plotted vs. the total recurrent activity, S/U. The black line shows the experimental Sample 1, with all unidentified neurons not considered. The red and blue lines illustrate how this changes when you take all of the unidentified neurons to be inhibitory or excitatory, respectively. **d)** The average total recurrent activity, $\langle S/U \rangle$, is plotted for all fifteen samples versus the sample's corresponding inhibitory ratio. Black markers denote $\langle S/U \rangle$ with the unidentified neurons not considered, red markers denote $\langle S/U \rangle$ with the unidentified neurons taken to be inhibitory and blue markers denote $\langle S/U \rangle$ with the unidentified neurons taken to be excitatory.

References

[1] Nigel Goldenfeld. *Lectures on Phase Transitions and the Renormalisation Group.* Addison-Wesley, 1992.

[2] Karin A. Dahmen, Yehuda Ben-Zion, and Jonathan T. Uhl. “Micromechanical model for deformation in solids with universal predictions for stress-strain curves and slip avalanches”. In: *Physical Review Letters* 102.17 (Apr. 2009), p. 175501. ISSN: 00319007. DOI: `10.1103/PhysRevLett.102.175501`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.102.175501`.

[3] Daniel S. Fisher. “Collective transport in random media: From superconductors to earthquakes”. In: *Physics Report* 301.1-3 (1998), pp. 113–150. ISSN: 03701573. DOI: `10.1016/S0370-1573(98)00008-8`. arXiv: `9711179v1 [arXiv:cond-mat]`.

[4] Deniz Ertaş and Mehran Kardar. “Critical dynamics of contact line depinning”. In: *Physical Review E* 49.4 (1994), pp. 2532–2535. ISSN: 1063651X. DOI: `10.1103/PhysRevE.49.R2532`. arXiv: `9401027 [cond-mat]`.

[5] Daniel S. Fisher et al. “Statistics of Earthquakes in Simple Models of Heterogeneous Faults”. In: *Physical Review Letters* (1997). ISSN: 10797114. DOI: `10.1103/PhysRevLett.78.4885`. eprint: `9703029` (cond-mat).

[6] Michael LeBlanc. “AVALANCHES IN PLASTIC DEFORMATION: MAXIMUM VELOCITY STATISTICS, FINITE TEMPERATURE EFFECTS, AND ANALYSIS OF LOW TIME RESOLUTION DATA”. PhD thesis. University of Illinois at Urbana Champaign, 2016.

[7] Karin Dahmen, Deniz Ertaş, and Yehuda Ben-Zion. “Gutenberg-Richter and characteristic earthquake behavior in simple mean-field models of heterogeneous faults”. In: *Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics* 58.2 (1998), pp. 1494–1501. ISSN: 1063651X. DOI: `10.1103/PhysRevE.58.1494`. arXiv: `9803057 [cond-mat]`.

[8] J. Patrick Coleman, Karin A. Dahmen, and Richard L. Weaver. “Avalanches and scaling collapse in the large- N Kuramoto model”. In: *Physical Review E* 97.4 (2018), pp. 4–11. ISSN: 24700053. DOI: `10.1103/PhysRevE.97.042219`.

[9] Alan Long. “NUCLEATION OF LARGE AVALANCHES”. PhD thesis. University of Illinois at Urbana Champaign, 2022.

[10] J. Moser and F. Levy. “Crystal reorientation and wear mechanisms in MoS2 lubricating thin films investigated by TEM”. In: *Journal of Materials Research* 8.1 (1993), pp. 206–213. ISSN: 08842914. DOI: `10.1557/JMR.1993.0206`.

[11] J. J. Hu et al. “Transmission electron microscopy analysis of Mo-W-S-Se film sliding contact obtained by using focused ion beam microscope and in situ microtribometer”. In: *Tribology Letters* 32.1 (2008), pp. 49–57. ISSN: 10238883. DOI: `10.1007/s11249-008-9360-z`.

[12] T. W. Scharf et al. “Synthesis of Au-MoS2 nanocomposites: Thermal and friction-induced changes to the structure”. In: *ACS Applied Materials and Interfaces* 5.22 (2013), pp. 11762–11767. ISSN: 19448244. DOI: `10.1021/am4034476`.

[13] J. K. Lancaster. “Anisotropy in the mechanical properties of lamellar solids and its effect on wear and transfer”. In: *Wear* 9.3 (1966), pp. 169–188. ISSN: 00431648. DOI: `10.1016/0043-1648(66)90142-6`.

[14] P. D. Fleischauer and R. Bauer. “The influence of surface chemistry on MoS SUB 2 transfer film formation.” In: *in: Technical Preprints Presented At Asle 41St Annual Meeting, (Toronto, Canada: May 12-15, 1986), Park Ridge, U.S.a., Am. So* 8197.86 AM-5G-2 (1986). ISSN: 0569-8197. DOI: `10.1080/05698198708981744`.

[15] T. W. Scharf and I. L. Singer. “Monitoring transfer films and friction instabilities with in situ Raman tribometry”. In: *Tribology Letters* 14.1 (2003), pp. 3–8. ISSN: 10238883. DOI: `10.1023/A:1021942830132`.

[16] J. M. Martin et al. “Superlubricity of molybdenum disulphide”. In: *Physical Review B* 48.14 (1993), pp. 10583–10586. ISSN: 01631829. DOI: `10.1103/PhysRevB.48.10583`.

[17] Mohammad R. Vazirisereshk et al. “Solid lubrication with MoS2: A review”. In: *Lubricants* 7.7 (2019). ISSN: 20754442. DOI: `10.3390/LUBRICANTS7070057`.

[18] Michael Urbakh et al. “The nonlinear nature of friction”. In: *Nature* 430.6999 (2004), pp. 525–528. ISSN: 00280836. DOI: `10.1038/nature02750`.

[19] Jacob Israelachvili. *Intermolecular and Surface Forces.* Vol. 46. 6. Elsevier, June 2011, pp. 1349–1352. ISBN: 9780123751829. DOI: `10.1016/C2009-0-21560-1`. URL: `https://linkinghub.elsevier.com/retrieve/pii/C20090215601`.

[20] J. F. Curry et al. "Temperature-Dependent Friction and Wear of MoS2/Sb2O3/Au Nanocomposites". In: *Tribology Letters* 64.1 (2016), pp. 1–5. ISSN: 10238883. DOI: `10.1007/s11249-016-0748-x`.

[21] Dansong Zhang, Karin A. Dahmen, and Martin Ostoja-Starzewski. "Scaling of slip avalanches in sheared amorphous materials based on large-scale atomistic simulations". In: *Physical Review E* 95.3 (2017), pp. 1–12. ISSN: 24700053. DOI: `10.1103/PhysRevE.95.032902`.

[22] Yehuda Ben-Zion, Karin A. Dahmen, and Jonathan T. Uhl. "A Unifying Phase Diagram for the Dynamics of Sheared Solids and Granular Materials". In: *Pure and Applied Geophysics* 168.12 (2011), pp. 2221–2237. ISSN: 00334553. DOI: `10.1007/s00024-011-0273-7`.

[23] O M Braun et al. "Collective Effects at Frictional Interfaces". In: *Tribology Letters* 48.1 (2012), pp. 11–25. ISSN: 1573-2711. DOI: `10.1007/s11249-012-9913-z`. URL: `https://doi.org/10.1007/s11249-012-9913-z`.

[24] James P. Sethna, Karin A. Dahmen, and Christopher R. Myers. "Crackling noise". In: *Nature* (2001). ISSN: 00280836. DOI: `10.1038/35065675`.

[25] James Antonaglia et al. "Bulk metallic glasses deform via slip avalanches". In: *Physical Review Letters* 112.15 (Apr. 2014), p. 155501. ISSN: 10797114. DOI: `10.1103/PhysRevLett.112.155501`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.112.155501`.

[26] Ekhard K.H. Salje and Karin A. Dahmen. "Crackling noise in disordered materials". In: *Annual Review of Condensed Matter Physics* 5.1 (Mar. 2014), pp. 233–254. ISSN: 19475462. DOI: `10.1146/annurev-conmatphys-031113-133838`. URL: `http://www.annualreviews.org/doi/10.1146/annurev-conmatphys-031113-133838`.

[27] Jonathan T. Uhl et al. "Universal Quake Statistics: From Compressed Nanocrystals to Earthquakes". In: *Scientific Reports* 5 (2015), pp. 1–10. ISSN: 20452322. DOI: `10.1038/srep16493`. URL: `http://dx.doi.org/10.1038/srep16493`.

[28] D. V. Denisov et al. "Universality of slip avalanches in flowing granular matter". In: *Nature Communications* 7.1 (Apr. 2016), p. 10641. ISSN: 20411723. DOI: `10.1038/ncomms10641`. URL: `http://www.nature.com/articles/ncomms10641`.

[29] Dmitry V. Denisov et al. “Universal slip dynamics in metallic glasses and granular matter - Linking frictional weakening with inertial effects”. In: *Scientific Reports* 7.January (2017), pp. 1–8. ISSN: 20452322. DOI: `10.1038/srep43376`. URL: `http://dx.doi.org/10.1038/srep43376`.

[30] James P. Sethna et al. “Deformation of Crystals: Connections with Statistical Physics”. In: *Annual Review of Materials Research* 47.1 (2017), pp. 217–246. ISSN: 15317331. DOI: `10.1146/annurev-matsci-070115-032036`. arXiv: `1609.05838`.

[31] Adam R. Hinkle et al. “Low friction in bcc metals via grain boundary sliding”. In: *Physical Review Materials* 4.6 (2020), p. 63602. ISSN: 24759953. DOI: `10.1103/PhysRevMaterials.4.063602`. URL: `https://doi.org/10.1103/PhysRevMaterials.4.063602`.

[32] Brandon A. Krick et al. “Optical in situ micro tribometer for analysis of real contact area for contact mechanics, adhesion, and sliding experiments”. In: *Tribology Letters* 45.1 (2012), pp. 185–194. ISSN: 10238883. DOI: `10.1007/s11249-011-9870-y`.

[33] John F. Curry et al. “Highly Oriented MoS2 Coatings: Tribology and Environmental Stability”. In: *Tribology Letters* 64.1 (2016), pp. 1–9. ISSN: 10238883. DOI: `10.1007/s11249-016-0745-0`.

[34] Jeffrey L. Hutter and John Bechhoefer. “Calibration of atomic-force microscope tips”. In: *Review of Scientific Instruments* 64.7 (1993), pp. 1868–1873. ISSN: 00346748. DOI: `10.1063/1.1143970`.

[35] H. S. Khare and D. L. Burris. “The extended wedge method: Atomic force microscope friction calibration for improved tolerance to instrument misalignments, tip offset, and blunt probes”. In: *Review of Scientific Instruments* 84.5 (2013). ISSN: 00346748. DOI: `10.1063/1.4804163`.

[36] Alex Vasenkov et al. “Reactive molecular dynamics study of Mo-based alloys under high-pressure, high-temperature conditions”. In: *Journal of Applied Physics* 112.1 (2012). ISSN: 00218979. DOI: `10.1063/1.4731793`.

[37] John F. Curry et al. “Atomistic Origins of Temperature-Dependent Shear Strength in 2D Materials”. In: *ACS Applied Nano Materials* 1.10 (2018), pp. 5401–5407. ISSN: 25740970. DOI: `10.1021/acsanm.8b01454`.

[38] Pieter Vermeesch. “Statistical uncertainty associated with histograms in the earth sciences”. In: *Journal of Geophysical Research: Solid Earth* 110.2 (2005), pp. 1–15. ISSN: 21699356. DOI: `10.1029/2004JB003479`.

[39] Joseph M. Monti and Mark O. Robbins. “Sliding Friction of Amorphous Asperities on Crystalline Substrates: Scaling with Contact Radius and Substrate Thickness”. In: *ACS Nano* 14.12 (2020), pp. 16997–17003. ISSN: 1936086X. DOI: `10.1021/acsnano.0c06241`.

[40] Tom W. J. de Geus and Matthieu Wyart. "Scaling theory for the statistics of slip at frictional interfaces". In: 065001 (2022), pp. 1–11. ISSN: 24700053. DOI: `10.1103/PhysRevE.106.065001`. arXiv: `2204.02795`. URL: `http://arxiv.org/abs/2204.02795`.

[41] Tom W.J. de Geus et al. "How collective asperity detachments nucleate slip at frictional interfaces". In: *Proceedings of the National Academy of Sciences of the United States of America* 116.48 (2019), pp. 23977–23983. ISSN: 10916490. DOI: `10.1073/pnas.1906551116`. arXiv: `1904.07635`.

[42] A.M. Litke et al. "What does the eye tell the brain?: Development of a system for the large-scale recording of retinal output activity". In: *IEEE Transactions on Nuclear Science* 51.4 (2004), pp. 1434–1440. DOI: `10.1109/TNS.2004.832706`.

[43] John M. Beggs and Dietmar Plenz. "Neuronal Avalanches in Neocortical Circuits". In: *The Journal of Neuroscience* 23.35 (Dec. 2003), pp. 11167–11177. ISSN: 0270-6474. DOI: `10.1523/JNEUROSCI.23-35-11167.2003`. URL: `http://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.23-35-11167.2003`.

[44] Rashid V. Williams-García, John M. Beggs, and Gerardo Ortiz. "Unveiling causal activity of complex networks". In: *EPL (Europhysics Letters)* (2017). ISSN: 0295-5075. DOI: `10.1209/0295-5075/119/18003`. eprint: `1603.05659`.

[45] Nir Friedman et al. "Universal critical dynamics in high resolution neuronal avalanche data". In: *Physical Review Letters* 108.20 (2012), pp. 1–5. ISSN: 00319007. DOI: `10.1103/PhysRevLett.108.208102`.

[46] Jeff Alstott, Ed Bullmore, and Dietmar Plenz. "powerlaw: A Python Package for Analysis of Heavy-Tailed Distributions". In: *PLOS ONE* 9.1 (Jan. 2014), pp. 1–11. DOI: `10.1371/journal.pone.0085777`. URL: `https://doi.org/10.1371/journal.pone.0085777`.

[47] Damian Berger et al. "Power Laws in Neuronal Culture Activity from Limited Availability of a Shared Resource". In: *Emergent Complexity from Nonlinearity, in Physics, Engineering and the Life Sciences*. Ed. by Giorgio Mantica, Ruedi Stoop, and Sebastiano Stramaglia. Cham: Springer International Publishing, 2017, pp. 209–220. ISBN: 978-3-319-47810-4.

[48] Robert A. White and Karin A. Dahmen. "Driving Rate Effects on Crackling Noise". In: *Phys. Rev. Lett.* 91 (8 Aug. 2003), p. 085702. DOI: `10.1103/PhysRevLett.91.085702`.

[49] Robert Allen White. "SWEEP RATE AND TEMPERATURE AFFECTS ON CRACKLING NOISE". PhD thesis. University of Illinois at Urbana Champaign, 2005.

[50] Elad Schneidman et al. "Weak pairwise correlations imply strongly correlated network states in a neural population". In: *Nature* 440.7087 (Apr. 2006), pp. 1007–1012. ISSN: 0028-0836. DOI: `10.1038/nature04701`. URL: `http://www.nature.com/articles/nature04701`.

[51] Andreas Draguhn and György Buzsáki. "Neuronal Oscillations in Cortical Networks". In: *Science* 304.June (2004), pp. 1926–1930.

[52] Vikaas S. Sohal et al. "Parvalbumin neurons and gamma rhythms enhance cortical circuit performance". In: *Nature* 459.7247 (2009), pp. 698–702. ISSN: 00280836. DOI: `10.1038/nature07991`.

[53] Andreas K. Engel, Pascal Fries, and Wolf Singer. "Dynamic predictions: Oscillations and synchrony in top–down processing". In: *Nature Reviews Neuroscience* 2.10 (2001), pp. 704–716. ISSN: 14710048. DOI: `10.1038/35094565`.

[54] Xiaoying Huang et al. "Spiral waves in disinhibited mammalian neocortex". In: *Journal of Neuroscience* 24.44 (2004), pp. 9897–9902. ISSN: 02706474. DOI: `10.1523/JNEUROSCI.2705-04.2004`.

[55] Antonio J Fontenele et al. "Criticality between Cortical States". In: *Physical Review Letters* 122.20 (May 2019), p. 208101. ISSN: 0031-9007. DOI: `10.1103/PhysRevLett.122.208101`. URL: `https://doi.org/10.1103/PhysRevLett.122.208101%20https://link.aps.org/doi/10.1103/PhysRevLett.122.208101`.

[56] Thomas Petermann et al. "Spontaneous cortical activity in awake monkeys composed of neuronal avalanches". In: *Proceedings of the National Academy of Sciences of the United States of America* 106.37 (2009), pp. 15921–15926. ISSN: 00278424. DOI: `10.1073/pnas.0904089106`.

[57] V. Pasquale et al. "Self-organization and neuronal avalanches in networks of dissociated cortical neurons". In: *Neuroscience* 153.4 (2008), pp. 1354–1369. ISSN: 03064522. DOI: `10.1016/j.neuroscience.2008.03.050`.

[58] Woodrow L. Shew et al. "Neuronal avalanches imply maximum dynamic range in cortical networks at criticality". In: *Journal of Neuroscience* 29.49 (2009), pp. 15595–15600. ISSN: 02706474. DOI: `10.1523/JNEUROSCI.3864-09.2009`. arXiv: `0906.0527`.

[59] Mikail Rubinov et al. "Neurobiologically realistic determinants of Self-Organized criticality in networks of spiking neurons". In: *PLoS Computational Biology* 7.6 (2011). ISSN: 1553734X. DOI: `10.1371/journal.pcbi.1002038`.

[60] Anirban Das and Anna Levina. "Critical Neuronal Models with Relaxed Timescale Separation". In: *Physical Review X* 9.2 (2019), p. 21062. ISSN: 21603308. DOI: `10.1103/PhysRevX.9.021062`. arXiv: `1808.04196`. URL: `https://doi.org/10.1103/PhysRevX.9.021062`.

[61] Jun Nosuke Teramae and Tomoki Fukai. "Local cortical circuit model inferred from power-law distributed neuronal avalanches". In: *Journal of Computational Neuroscience* 22.3 (2007), pp. 301–312. ISSN: 09295313. DOI: `10.1007/s10827-006-0014-6`.

[62] Serena Di Santo et al. "Landau–Ginzburg theory of cortex dynamics: Scale-free avalanches emerge at the edge of synchronization". In: *Proceedings of the National Academy of Sciences of the United States of America* 115.7 (2018), E1356–E1365. ISSN: 10916490. DOI: `10.1073/pnas.1712989115`. arXiv: `1801.10356`.

[63] Ludmila Brochini et al. "Phase transitions and self-organized criticality in networks of stochastic spiking neurons". In: *Scientific Reports* 6.October (2016), pp. 1–15. ISSN: 20452322. DOI: `10.1038/srep35831`. arXiv: `1606.06391`.

[64] Zhixin Lu et al. "Inhibitory neurons promote robust critical firing dynamics in networks of integrate-and-fire neurons". In: *Physical Review E* 94.6 (Dec. 2016), p. 062309. ISSN: 2470-0045. DOI: `10.1103/PhysRevE.94.062309`. URL: `https://link.aps.org/doi/10.1103/PhysRevE.94.062309`.

[65] Mauricio Girardi-Schappo et al. "Self-organized critical balanced networks: A unified framework". In: *arXiv* 2.1 (2019), p. 12042. ISSN: 23318422. DOI: `10.1103/PhysRevResearch.2.012042`. arXiv: `1906.05624`. URL: `https://doi.org/10.1103/PhysRevResearch.2.012042`.

[66] Michael A. Buice and Jack D. Cowan. "Field-theoretic approach to fluctuation effects in neural networks". In: *Physical Review E* 75.5 (May 2007), p. 051919. ISSN: 1539-3755. DOI: `10.1103/PhysRevE.75.051919`. URL: `https://link.aps.org/doi/10.1103/PhysRevE.75.051919`.

[67] W. L. Shew et al. "Information Capacity and Transmission Are Maximized in Balanced Cortical Networks with Neuronal Avalanches". In: *Journal of Neuroscience* 31.1 (Jan. 2011), pp. 55–63. ISSN: 0270-6474. DOI: `10.1523/JNEUROSCI.4637-10.2011`. URL: `http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.4637-10.2011`.

[68] F. Lombardi et al. "Balance between excitation and inhibition controls the temporal organization of neuronal avalanches". In: *Physical Review Letters* (2012). ISSN: 00319007. DOI: `10.1103/PhysRevLett.108.228703`.

[69] Mohammad Yaghoubi et al. "Neuronal avalanche dynamics indicates different universality classes in neuronal cultures". In: *Scientific Reports* 8.1 (Dec. 2018), p. 3417. ISSN: 2045-2322. DOI: `10.1038/s41598-018-21730-1`. URL: `http://www.nature.com/articles/s41598-018-21730-1`.

[70] J. S. Urbach, R. C. Madison, and J. T. Markert. "Interface depinning, self-organized criticality, and the Barkhausen effect". In: *Physical Review Letters* 75.2 (1995), pp. 276–279. ISSN: 00319007. DOI: `10.1103/PhysRevLett.75.276`.

[71] Tyler Salners et al. "Recurrent activity in neuronal avalanches". In: *Scientific reports* 13.1 (2023), p. 4871. ISSN: 20452322. DOI: `10.1038/s41598-023-31851-x`. URL: `https://doi.org/10.1038/s41598-023-31851-x`.

[72] B Y J S Coombs, J C Eccles, and P Fatt. "Potential Change". In: *October* (1955), pp. 326–373.

[73] D. Whitteridge. *Studies in Physiology.* Ed. by D.R. Curtis and A.K. McIntyre. Vol. 51. 2. 1966, pp. 166–166. ISBN: 9781119130536. DOI: `10.1113/expphysiol.1966.sp001840`.

[74] Shinya Ito et al. "Spontaneous spiking activity of hundreds of neurons in mouse somatosensory cortex slice cultures recorded using a dense 512 electrode array". In: *https://www.jneurosci.org/content/16/14/4438* (2016). DOI: `10.6080/K07D2S2F`. URL: `https://www.jneurosci.org/content/16/14/4438`.

[75] Jean-Pierre Wuarin and F. Edward Dudek. "Electrographic Seizures and New Recurrent Excitatory Circuits in the Dentate Gyrus of Hippocampal Slices from Kainate-Treated Epileptic Rats". In: *Journal of Neuroscience* 16.14 (1996), pp. 4438–4448. ISSN: 0270-6474. DOI: `10.1523/JNEUROSCI.16-14-04438.1996`. eprint: `https://www.jneurosci.org/content/16/14/4438.full.pdf`. URL: `https://www.jneurosci.org/content/16/14/4438`.

[76] Colin Deransart et al. "Single-unit Analysis of Substantia Nigra Pars Reticulata Neurons in Freely Behaving Rats with Genetic Absence Epilepsy". In: *Epilepsia* 44.12 (2003), pp. 1513–1520. ISSN: 00139580. DOI: `10.1111/j.0013-9580.2003.26603.x`.

[77] Jon P. Hobbs, Jodi L. Smith, and John M. Beggs. "Aberrant neuronal avalanches in cortical tissue removed from juvenile epilepsy patients". In: *Journal of Clinical Neurophysiology* 27.6 (2010), pp. 380–386. ISSN: 07360258. DOI: `10.1097/WNP.0b013e3181fdf8d3`.

[78] Bethany Percha et al. "Transition from local to global phase synchrony in small world neural network and its possible implications for epilepsy". In: *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 72.3 (2005), pp. 1–6. ISSN: 15393755. DOI: `10.1103/PhysRevE.72.031909`. arXiv: `0411126 [physics]`.

[79] Olga Garaschuk, Eric Hanse, and Arthur Konnerth. "Developmental profile and synaptic origin of early network oscillations in the CA1 region of rat neonatal hippocampus". In: *Journal of Physiology* 507.1 (1998), pp. 219–236. ISSN: 00223751. DOI: `10.1111/j.1469-7793.1998.219bu.x`.

[80] Y. Ben-Ari et al. “Giant synaptic potentials in immature rat CA3 hippocampal neurones.” In: *The Journal of Physiology* 416.1 (1989), pp. 303–325. ISSN: 14697793. DOI: `10.1113/jphysiol.1989.sp017762`.

[81] Ariadne Costa, Ludmila Brochini, and Osame Kinouchi. “Self-Organized Supercriticality and Oscillations in Networks of Stochastic Spiking Neurons”. In: *Entropy* 19.8 (Aug. 2017), p. 399. ISSN: 1099-4300. DOI: `10.3390/e19080399`. URL: `http://www.mdpi.com/1099-4300/19/8/399`.

[82] L. de Arcangelis. “Are dragon-king neuronal avalanches dungeons for self-organized brain activity?” In: *The European Physical Journal Special Topics* 205.1 (May 2012), pp. 243–257. ISSN: 1951-6355. DOI: `10.1140/epjst/e2012-01574-6`. URL: `http://link.springer.com/10.1140/epjst/e2012-01574-6`.

[83] L. Menendez De La Prida and J. V. Sanchez-Andres. “Heterogeneous populations of cells mediate spontaneous synchronous bursting in the developing hippocampus through a frequency-dependent mechanism”. In: *Neuroscience* 97.2 (2000), pp. 227–241. ISSN: 03064522. DOI: `10.1016/S0306-4522(00)00029-4`.

[84] Liset Menendez De La Prida, Sonia Bolea, and Juan V. Sanchez-Andres. “Analytical characterization of spontaneous activity evolution during hippocampal development in the rabbit”. In: *Neuroscience Letters* 218.3 (1996), pp. 185–187. ISSN: 03043940. DOI: `10.1016/S0304-3940(96)13095-0`.

[85] R. Miles and R. K. Wong. “Inhibitory control of local excitatory circuits in the guinea-pig hippocampus.” In: *The Journal of Physiology* 388.1 (1987), pp. 611–629. ISSN: 14697793. DOI: `10.1113/jphysiol.1987.sp016634`. URL: `https://doi.org/10.1113/jphysiol.1987.sp016634`.

[86] A. Levina, J. M. Herrmann, and T. Geisel. “Dynamical synapses causing self-organized criticality in neural networks”. In: *Nature Physics* 3.12 (2007), pp. 857–860. ISSN: 17452481. DOI: `10.1038/nphys758`. arXiv: `0712.1003`.

[87] J. Lehnert et al. “Loss of synchronization in complex neuronal networks with delay”. In: *Epl* 96.6 (2011), pp. 1–6. ISSN: 02955075. DOI: `10.1209/0295-5075/96/60013`. arXiv: `arXiv:1107.4195v2`.

[88] Michael G. Rosenblum and Arkady S. Pikovsky. “Controlling Synchronization in an Ensemble of Globally Coupled Oscillators”. In: *Physical Review Letters* 92.11 (2004), pp. 1–4. ISSN: 00319007. DOI: `10.1103/PhysRevLett.92.114102`.

[89] Yu Shen, Zhonghuai Hou, and Houwen Xin. “Transition to burst synchronization in coupled neuron networks”. In: *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 77.3 (2008), pp. 2–6. ISSN: 15393755. DOI: `10.1103/PhysRevE.77.031920`.

[90] Yahya Karimipanah, Zhengyu Ma, and Ralf Wessel. "Criticality predicts maximum irregularity in recurrent networks of excitatory nodes". In: *PLoS ONE* 12.8 (2017), pp. 1–15. ISSN: 19326203. DOI: `10.1371/journal.pone.0182501`.

[91] Felipe Yaroslav et al. "Interface depinning, self-organized criticality, and the Barkhausen effect". In: *Physical Review Letters* 6.1 (Dec. 2010). Ed. by Karl J. Friston, pp. 1–4. ISSN: 00319007. DOI: `10.1371/journal.pcbi.1001013`. arXiv: `1801.10356`. URL: `http://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.23-35-11167.2003%20https://dx.plos.org/10.1371/journal.pcbi.1000846%20https://link.aps.org/doi/10.1103/PhysRevE.75.051919%20http://link.springer.com/10.1007/s00422-006-0068-6%20http://www.mdpi.com/1099-4300/19`.

[92] S.-S. Poil et al. "Critical-State Dynamics of Avalanches and Oscillations Jointly Emerge from Balanced Excitation/Inhibition in Neuronal Networks". In: *Journal of Neuroscience* 32.29 (July 2012), pp. 9817–9823. ISSN: 0270-6474. DOI: `10.1523/JNEUROSCI.5990-11.2012`. URL: `http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.5990-11.2012`.

[93] Manfred G Kitzbichler et al. "Broadband Criticality of Human Brain Network Synchronization". In: 5.3 (2009). DOI: `10.1371/journal.pcbi.1000314`.

[94] Onuttom Narayan. "Self similar Barkhausen noise in magnetic domain wall motion". In: *Physical Review Letters* 77.18 (Oct. 1996), pp. 3855–3857. ISSN: 10797114. DOI: `10.1103/PhysRevLett.77.3855`. arXiv: `9610053 [cond-mat]`. URL: `http://arxiv.org/abs/cond-mat/9610053%20http://dx.doi.org/10.1103/PhysRevLett.77.3855`.

[95] Gianfranco Durin and Stefano Zapperi. "Universality and size effects in the Barkhausen noise". In: *Journal of Applied Physics* 87.9 III (2000), pp. 7031–7033. ISSN: 00218979. DOI: `10.1063/1.372921`.

[96] Elakkat D Gireesh and Dietmar Plenz. "Neuronal avalanches organize as nested theta- and beta/gamma-oscillations during development of cortical layer 2/3". In: *Proceedings of the National Academy of Sciences* 105.21 (May 2008), 7576 LP –7581. DOI: `10.1073/pnas.0800537105`. URL: `http://www.pnas.org/content/105/21/7576.abstract`.

[97] Benjamin R. Cowley et al. "Stimulus-Driven Population Activity Patterns in Macaque Primary Visual Cortex". In: *PLoS Computational Biology* (2016). ISSN: 15537358. DOI: `10.1371/journal.pcbi.1005185`.

[98] Yuta Senzai, Antonio Fernandez-Ruiz, and György Buzsáki. "Layer-Specific Physiological Features and Interlaminar Interactions in the Primary Visual Cortex of the Mouse". In: *Neuron* 101.3 (2019), 500–513.e5. ISSN: 10974199. DOI: `10.1016/j.neuron.2018.12.009`.

[99] Woodrow L. Shew et al. “Adaptation to sensory input tunes visual cortex to criticality”. In: *Nature Physics* (2015). ISSN: 17452481. DOI: `10.1038/nphys3370`.

[100] Jiaqing Yan et al. “Analysis of electrocorticogram in epilepsy patients in terms of criticality”. In: *Nonlinear Dynamics* 83.4 (2016), pp. 1909–1917. ISSN: 1573269X. DOI: `10.1007/s11071-015-2455-9`.

[101] Oren Shriki et al. “Neuronal avalanches in the resting MEG of the human brain”. In: *Journal of Neuroscience* (2013). ISSN: 02706474. DOI: `10.1523/JNEUROSCI.4286-12.2013`.

[102] Skander Mensi et al. “Parameter extraction and classification of three cortical neuron types reveals two distinct adaptation mechanisms”. In: *Journal of Neurophysiology* 107.6 (2012), pp. 1756–1775. ISSN: 00223077. DOI: `10.1152/jn.00408.2011`.

[103] Karin Dahmen and James P. Sethna. “Hysteresis, avalanches, and disorder-induced critical scaling: A renormalization-group approach”. In: *Physical Review B* 53.22 (June 1996), pp. 14872–14905. ISSN: 0163-1829. DOI: `10.1103/PhysRevB.53.14872`. URL: `https://link.aps.org/doi/10.1103/PhysRevB.53.14872`.

[104] Victor A.F. Lamme, Hans Supèr, and Henk Spekreijse. “Feedforward, horizontal, and feedback processing in the visual cortex”. In: *Current Opinion in Neurobiology* 8.4 (1998), pp. 529–535. ISSN: 09594388. DOI: `10.1016/S0959-4388(98)80042-1`.

[105] Rodney J. Douglas et al. “Recurrent excitation in neocortical circuits”. In: *Science* (1995). ISSN: 00368075. DOI: `10.1126/science.7638624`.

[106] Maria V. Sanchez-Vives and David A. McCormick. “Cellular and network mechanisms of rhytmic recurrent activity in neocortex”. In: *Nature Neuroscience* 3.10 (2000), pp. 1027–1034. ISSN: 10976256. DOI: `10.1038/79848`.

[107] Jason C. Wester and Diego Contreras. “Columnar interactions determine horizontal propagation of recurrent network activity in neocortex”. In: *Journal of Neuroscience* 32.16 (2012), pp. 5454–5471. ISSN: 02706474. DOI: `10.1523/JNEUROSCI.5006-11.2012`.

[108] Garrett T. Neske, Saundra L. Patrick, and Barry W. Connors. “Contributions of diverse excitatory and inhibitory neurons to recurrent network activity in cerebral cortex”. In: *Journal of Neuroscience* 35.3 (2015), pp. 1089–1105. ISSN: 15292401. DOI: `10.1523/JNEUROSCI.2279-14.2015`.

[109] Go Eun Ha and Eunji Cheong. “Spike frequency adaptation in neurons of the central nervous system”. In: *Experimental Neurobiology* 26.4 (2017), pp. 179–185. ISSN: 20938144. DOI: `10.5607/en.2017.26.4.179`.

[110] Christopher R. Myers and J. S. Langer. "Rupture propagation, dynamical front selection, and the role of small length scales in a model of an earthquake fault". In: *Physical Review E* 47.5 (1993), pp. 3048–3056. ISSN: 1063651X. DOI: `10.1103/PhysRevE.47.3048`.

[111] Karin A. Dahmen and Wendelin J. Wright. "Avalanches in Solids, Theory and Experiments". In: *Encyclopedia of Continuum Mechanics* (2018), pp. 1–19. DOI: `10.1007/978-3-662-53605-6_73-1`.

[112] J. I. Díaz et al. "Mathematical analysis of a model of river channel formation". In: *Pure and Applied Geophysics* 165.8 (2008), pp. 1663–1682. ISSN: 00334553. DOI: `10.1007/s00024-004-0394-3`.

[113] Subhasish Dey and Sk Zeeshan Ali. "Fluvial instabilities". In: *Physics of Fluids* 32.6 (2020). ISSN: 10897666. DOI: `10.1063/5.0010038`. URL: `https://doi.org/10.1063/5.0010038`.

[114] Na Chen et al. "Afterhyperpolarization improves spike programming through lowering threshold potentials and refractory periods mediated by voltage-gated sodium channels". In: *Biochemical and Biophysical Research Communications* 346.3 (2006), pp. 938–945. ISSN: 0006291X. DOI: `10.1016/j.bbrc.2006.06.003`.

[115] Alex D. Reyes. "Synaptic short-term plasticity in auditory cortical circuits". In: *Hearing Research* 279.1-2 (Sept. 2011), pp. 60–66. ISSN: 03785955. DOI: `10.1016/j.heares.2011.04.017`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S0378595511001304`.

[116] Zhengyu Ma et al. "Cortical Circuit Dynamics Are Homeostatically Tuned to Criticality In Vivo". In: *Neuron* 104.4 (2019), 655–664.e4. ISSN: 10974199. DOI: `10.1016/j.neuron.2019.08.031`. URL: `https://doi.org/10.1016/j.neuron.2019.08.031`.

[117] Zhengyu Ma et al. "Stability of motor cortex network states during learning-associated neural re-organizations". In: *Journal of Neurophysiology* 124.5 (2020), pp. 1327–1342. ISSN: 15221598. DOI: `10.1152/jn.00061.2020`.

[118] Benedetta Mariani et al. "Beyond resting state neuronal avalanches in the somatosensory barrel cortex". In: *Frontiers in Systems Neuroscience* 0 (1), p. 89. ISSN: 1662-5137.

[119] Thomas Schreiber. "Measuring information transfer". In: *Physical Review Letters* (2000). ISSN: 00319007. DOI: `10.1103/PhysRevLett.85.461`. eprint: `0001042` (nlin).

[120] Shinya Ito et al. "Extending Transfer Entropy Improves Identification of Effective Connectivity in a Spiking Cortical Network Model". In: *PLoS ONE* 6.11 (Nov. 2011). Ed. by Michal Zochowski, e27431. ISSN: 1932-6203. DOI: `10.1371/journal.pone.0027431`. URL: `https://dx.plos.org/10.1371/journal.pone.0027431`.

[121] Michael Wibral et al. "TRENTOOL: an open source toolbox to estimate neural directed interactions with transfer entropy". In: *BMC Neuroscience* 12.S1 (2011), pp. 4–5. ISSN: 1471-2202. DOI: `10.1186/1471-2202-12-s1-p200`.

[122] Masanori Shimono and John M. Beggs. "Functional clusters, hubs, and communities in the cortical microconnectome". In: *Cerebral Cortex* 25.10 (2015), pp. 3743–3757. ISSN: 14602199. DOI: `10.1093/cercor/bhu252`.

[123] Sunny Nigam et al. "Rich-Club Organization in Effective Connectivity among Cortical Neurons". In: *The Journal of Neuroscience* 36.3 (Jan. 2016), pp. 670–684. ISSN: 0270-6474. DOI: `10.1523/JNEUROSCI.2177-15.2016`. URL: `http://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.2177-15.2016`.

[124] G BÖRNER. "Correlations in the absorption lines of the quasar Q0420-388". In: *Astronomy and astrophysics (Berlin. Print)* 218.1 (1989), pp. 19–23. ISSN: 0004-6361. URL: `http://library1.nida.ac.th/termpaper6/sd/2554/19755.pdf`.

[125] Louis W. McFaul et al. "Applied-force oscillations in avalanche dynamics". In: *Physical Review E* 101.5 (2020), pp. 1–10. ISSN: 24700053. DOI: `10.1103/PhysRevE.101.053003`.

[126] Tomáš Hromádka et al. "Sparse Representation of Sounds in the Unanesthetized Auditory Cortex Figure S7 Neuronal responses are heterogeneous". In: *PLoS Biology* 6.1 (2008), pp. 4–5. ISSN: 1545-7885. DOI: `10.1371/journal.pbio.0060016`.

[127] Sen Song et al. "Highly nonrandom features of synaptic connectivity in local cortical circuits". In: *PLoS Biology* 3.3 (2005), pp. 0507–0519. ISSN: 15457885. DOI: `10.1371/journal.pbio.0030068`.

[128] S. H. Hendry et al. "Numbers and proportions of GABA-immunoreactive neurons in different areas of monkey cerebral cortex." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 7.5 (1987), pp. 1503–1519. ISSN: 02706474. DOI: `10.1523/jneurosci.07-05-01503.1987`.

[129] Harang Ju et al. "Network structure of cascading neural systems predicts stimulus propagation and recovery". In: *Journal of Neural Engineering* 17.5 (2020). ISSN: 17412552. DOI: `10.1088/1741-2552/abbff1`. eprint: `1812.09361`.

[130] Jeffrey C. Magee and Christine Grienberger. "Synaptic Plasticity Forms and Functions". In: *Annual Review of Neuroscience* 43 (2020), pp. 95–117. ISSN: 15454126. DOI: `10.1146/annurev-neuro-090919-022842`.

[131] William A. Liberti et al. "Unstable neurons underlie a stable learned behavior". In: *Nature Neuroscience* 19.12 (2016), pp. 1665–1671. ISSN: 15461726. DOI: `10.1038/nn.4405`.

[132] Joana Lourenço, Fani Koukouli, and Alberto Bacci. "Synaptic inhibition in the neocortex: Orchestration and computation through canonical circuits and variations on the theme". In: *Cortex* 132 (2020), pp. 258–280. ISSN: 19738102. DOI: `10.1016/j.cortex.2020.08.015`.

[133] Robert Kim and Terrence J. Sejnowski. "Strong inhibitory signaling underlies stable temporal dynamics and working memory in spiking neural networks". In: *Nature Neuroscience* 24.1 (2021), pp. 129–139. ISSN: 15461726. DOI: `10.1038/s41593-020-00753-w`. URL: `http://dx.doi.org/10.1038/s41593-020-00753-w`.

[134] Karin A. Dahmen, Yehuda Ben-Zion, and Jonathan T. Uhl. "A simple analytic theory for the statistics of avalanches in sheared granular materials". In: *Nature Physics* 7.7 (2011), pp. 554–557. ISSN: 17452481. DOI: `10.1038/nphys1957`. URL: `http://dx.doi.org/10.1038/nphys1957`.

[135] Michael LeBlanc et al. "Universal fluctuations and extreme statistics of avalanches near the depinning transtition". In: *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 87.2 (Mar. 2013), pp. 1–14. ISSN: 15393755. DOI: `10.1103/PhysRevE.87.022126`. arXiv: `1303.1781`. URL: `http://arxiv.org/abs/1303.1781%20http://dx.doi.org/10.1103/PhysRevE.87.022126`.

[136] Douglas J. Bakkum et al. "Parameters for burst detection". In: *Frontiers in Computational Neuroscience* 7.JAN (2014), pp. 1–12. ISSN: 16625188. DOI: `10.3389/fncom.2013.00193`.

[137] William McFaul. "TEMPORAL PROPERTIES OF AVALANCHE DYNAMICS". PhD thesis. University of Illinois at Urbana Champaign, 2019.