Manuscript Template

FRONT MATTER

Title

Low dimensional criticality embedded in high dimensional awake brain dynamics

Short Title

Critical subspace in awake cortex

Authors

Antonio J. Fontenele,[1†] J. Samuel Sooter,[1†] V. Kindler Norman,[1] Shree Hari Gautam,[1] Woodrow L. Shew[1]*

Affiliations

[1]UA Integrative Systems Neuroscience Group, Department of Physics, University of Arkansas, Fayetteville, AR, USA, 72701

*Corresponding email: shew@uark.edu

[†]These authors contributed equally to this work.

Abstract

Whether cortical neurons operate in a strongly or weakly correlated dynamical regime determines fundamental information processing capabilities and has fueled decades of debate. Here we offer a resolution of this debate; we show that two important dynamical regimes, typically considered incompatible, can coexist in the same local cortical circuit by separating them into two different subspaces. In awake mouse motor cortex, we find a low-dimensional subspace with large fluctuations consistent with criticality – a dynamical regime with moderate correlations and multi-scale information capacity and transmission. Orthogonal to this critical subspace, we find a high-dimensional subspace containing a desynchronized dynamical regime, which may optimize input discrimination. The critical subspace is apparent only at long timescales, which explains discrepancies among some previous studies. Using a computational model, we show that the emergence of a low-dimensional critical subspace at large timescale agrees with established theory of critical dynamics. Our results suggest that cortex leverages its high dimensionality to multiplex dynamical regimes across different subspaces.

Teaser

Temporal coarse-graining reveals a low-dimensional critical subspace coexistent with a desynchronized subspace in awake cortex.

Introduction

The ongoing dynamics of neuronal populations in cerebral cortex are high dimensional and the nature of the dynamics differs across dimensions. Some dimensions contain large amplitude fluctuations, while other dimensions tend to have smaller fluctuations. Accumulating evidence shows that a relatively small fraction of the total possible dimensionality - i.e. a low-dimensional subspace - contains much of the large amplitude collective fluctuations (*1–6*). This fact is commonly used to justify simplified views of population activity using just a few dominant dimensions defined by principal component analysis (PCA), for example (*7*, *8*). Less commonly discussed are the smaller amplitude, more desynchronized fluctuations in the higher dimensional subspace, beyond the first few principal components. These basic observations lead to an interesting possible resolution of a long-standing debate in systems neuroscience.

The debate concerns the answer to a fundamental question: what is the dynamical regime of cortical neural networks in awake animals? Some researchers argue in favor of a desynchronized dynamical regime with weak collective fluctuations. This view is supported by measurements of low correlations of spikes recorded from pairs of neurons. Also, the desynchronized regime comes with potential computational benefits, like low noise and high input discrimination (*9–11*). Others argue for the importance of synchronized dynamics, based on observed oscillatory brain activity and potentially improved long range signal transmission (*12–14*). Of particular importance for our work here, a third camp argues in favor of criticality - a dynamical regime poised at the boundary between uncorrelated and strongly coordinated regimes (*15–19*). Criticality manifests with moderate correlations, highly diverse spatiotemporal fluctuations, and is thought to confer multiple functional advantages, like optimal information transmission and dynamic range (*20–26*).

How might consideration of the high-dimensional geometry of population dynamics help resolve the debate about cortical dynamical regime? Here we hypothesize that cortical populations could support more than one dynamical regime at the same time by separating different regimes into different subspaces. Criticality could occupy the low-dimensional subspace with large fluctuations (first few principal components) while a more desynchronized regime could occupy the higher-dimensional subspace with weak fluctuations (principal components beyond the first few). If this possibility is correct, then a cortical circuit need not choose between weak and strong correlations; it can have both simultaneously.

Is this hypothesis consistent with previous studies? In particular, how could this hypothesis be consistent with the multiple previous studies that have analyzed recordings of spikes from awake cortex and concluded that the activity was not consistent with criticality (*27–32*). These studies found activity that appeared to be weakly correlated, more consistent with a desynchronized regime. At first glance, these studies seem to contradict our hypothesis here of a low-dimensional critical subspace with large, correlated fluctuations; why might these studies miss the critical subspace? A likely answer to this question comes from a comprehensive comparison of all attempts to seek evidence for criticality based on recordings with single neuron resolution in awake animals (22 cases in all including the 13 negative cases in the papers cited above, a complete list to our knowledge at the time of writing this paper, Supplementary Table S1). A distinguishing feature of the 9 cases that reported positive evidence for criticality was that they focused on relatively long time scale fluctuations; they considered coarse-grained

spike counts in time bins larger than about 10 ms. This temporal coarse-graining was either a deliberate choice in the data analysis (*31*, *33*) or was due to limited time resolution of experimental measurements, which is typical for calcium imaging (*32*, *34–38*). The most strongly negative reports were based on analyses at the millisecond time scale, with little temporal coarse-graining. Based on this small meta-analysis of previous work and theoretical support for the importance of temporal coarse-graining (more on this in the results below), we refine our hypothesis. We hypothesize that desynchronized and critical subspaces coexist, but that the critical subspace is not apparent without temporal coarse-graining. To test this hypothesis we recorded spike activity from motor cortex of awake mice and compared our measurements to a ground-truth theoretical model of critical dynamics.

Results

We performed extracellular spike recordings of up to 247 units in motor cortex of awake, behaving mice (Fig 1A, 4 mice, 19 recordings, n=104±43 single units, 44±18 multi-units per recording, 44±9 minutes recording duration, more details in Methods). Our analysis of each recording begins with generating an N × M spike count matrix (Fig 1B), where N is the number of neurons and M is the number of time bins (M = recording duration divided by time bin duration ΔT). The entry in the ith row and jth column is the number of spikes fired by the ith unit during the jth time bin. We performed principal component analysis on each spike count matrix, and found that the activity is high dimensional, but much less than N dimensional; 45±0.05% of principal components (PCs) were needed to explain 95% of variance (Fig 1C).

Next, we performed avalanche analysis following previously established methods (*26*, *33*, *39*, *40*). We first created a 1-dimensional time series by summing spike counts across all N neurons. Then, we defined an avalanche as a period when this summed population activity exceeds a threshold (Fig 1D, see Methods). According to theory (*40*, *41*), if the neural system operates near criticality, then avalanches should be diverse, with sizes (S) and durations (T) distributed according to power-laws, $P(S) \sim S^{-\tau}$ and $P(T) \sim T^{-\tau t}$. We found good agreement with this prediction; avalanche sizes and durations were power-law distributed over a wide range of scales (Fig. 1E,F, purple). In addition, theory predicts that the exponents, τ and τ_t, are not independent; they should be related to each other and to a third exponent ($1/\sigma\upsilon z$) according to the scaling law $1/\sigma\upsilon z = (1 - \tau_t)/(1- \tau)$ (*40*, *41*), sometime called the "crackling noise" scaling law. Moreover, according to theory, avalanche size depends on duration as $S \sim T^{1/\sigma\upsilon z}$. The avalanches we measured conformed well to all these predictions (Fig 1G). Thus, we conclude that the awake spike activity we observed here is in good agreement with predictions for a system operating at criticality. (As alluded to in the Introduction, this finding depends on the choice of time scale for the time bins ΔT, which we will investigate further below.)

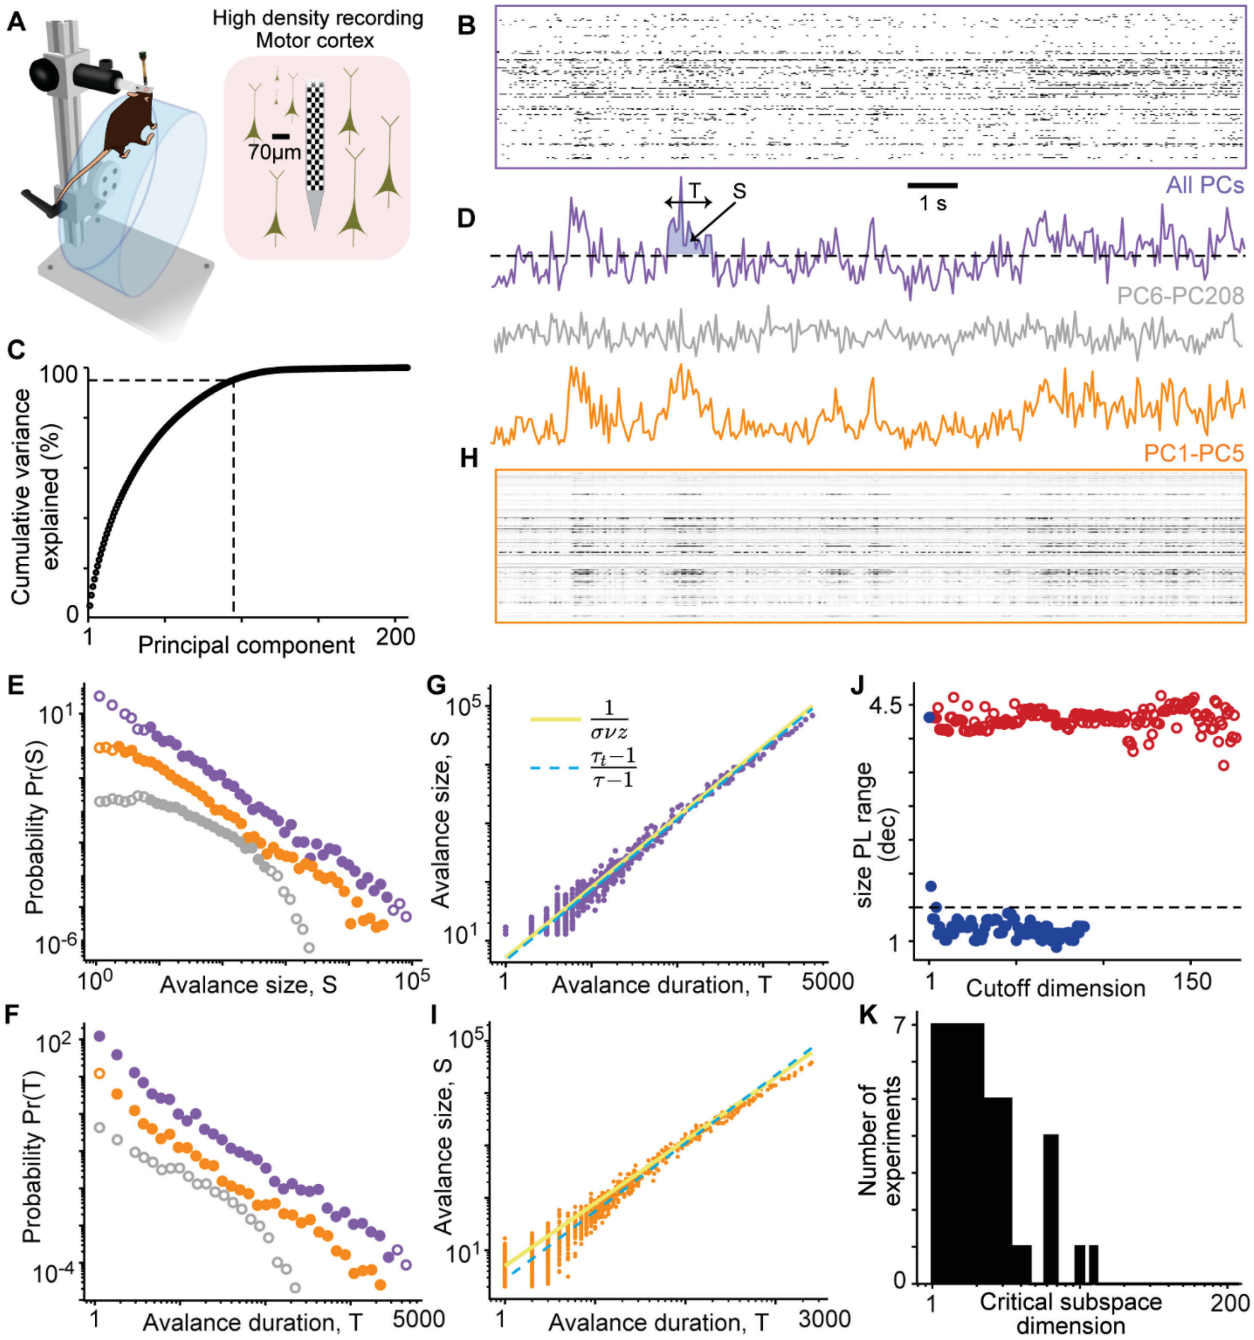

Fig. 1. Low dimensional critical subspace. (A) Mice were head-fixed and placed on top of a wheel, free to run, rest, groom, etc. High density recordings were performed in motor cortex, using a Neuropixels 1.0 probe. **(B)** Example spike count raster plot for 208 units. **(C)** Cumulative variance of original activity explained by increasing number of principal components. Dashed lines marks the number of PCs needed to explain 95% of the variance. **(D)** Spike count time series for entire population (ΔT = 50 ms). Dashed line represents the avalanche threshold. The shaded suprathreshold excursion represents one avalanche. **(E)** Distributions of avalanche size (S) for the original data (purple), reconstructed data using the first five principal components (orange) and after removing the first five components (grey). Distributions are shifted vertically for comparison. **(F)** Same as in E, but for avalanche durations (T). **(G)** Avalanche sizes and durations follow predicted scaling law for original data. **(H)** Raster and population sum for reconstructed data

based on the first five principal components (orange) or excluding the first five principal components (gray). **(I)** Scaling laws obeyed for PC1-5 reconstructed data. **(J)** Power-law range for avalanche size distributions as a function of PCs removed ascending/descending order (blue/red). Dimension of critical subspace is defined as the number of PCs removed before power-law range drops below 1.5 decades (dashed line). **(K)** Histogram of critical subspace dimension for all 19 recording sessions.

Critical subspace is spanned by first few principal components

Traditional avalanche analysis (Fig 1D-G purple) is based on the 1-dimensional population summed activity, which, by definition, precludes insight on how the dynamics may be distributed across the multiple dimensions revealed by PCA (Fig 1C). Do the scale-free avalanche dynamics exist in a subspace of the high dimensional population dynamics? If so, then the results of avalanche analysis should be robust to removal of PCs that are not in the avalanche subspace. We tested this possibility first with an example case (Fig 1 orange). We generated a new N × M activity matrix, reconstructed using only the first five PCs, excluding the other 203 dimensions (Methods). The population sum time series based on the reconstructed data was very similar to the original (Fig 1D,H). Moreover, the avalanche statistics for this 5-dimensional reconstruction remained in good agreement with predictions for criticality (Fig 1 E,F,I orange). In contrast, if we reconstructed the data using PCs 6-208, the avalanche sizes and durations were not power-law distributed (Fig 1E,F gray). Thus, for this example, we conclude that the critical dynamics are primarily contained within the subspace defined by the first 5 PCs.

This example raises interesting questions. Is the dimensionality of this "critical subspace" exactly 5; is it higher or lower? Is the critical subspace always spanned by the first several PCs? To address these questions, we repeated the avalanche analysis using dimensionally-reduced reconstructed data, but systematically varying the cutoff dimension d_c from 1 to N ($d_c = 5$ in the example above). For each cutoff dimension we performed avalanche analysis on two data sets; one reconstructed using PCs 1 through d_c, the other reconstructed using PCs d_c+1 through N. For each case, we quantified the range of the avalanche size distribution that was well-fit by a power law. When reconstructed using PCs 1 through d_c, the power-law range remained high, largely independent of d_c (Fig 1J red). Consistent with the $d_c = 5$ example above, this suggests that the avalanche statistics are not impacted by the activity in the dimensions defined by the high PCs. Indeed, when we reconstructed the data using PCs d_c through N, the power-law range dramatically dropped when d_c exceeded a relatively small number (Fig 1J blue). Thus, the critical subspace is strongly dependent on the first few PCs. This observation suggests a convenient quantitative definition for the dimensionality of the critical subspace - the number of PCs that can be removed (starting from PC 1) before the power-law range drops below 1.5 decades. Using this definition, we found that the critical subspace rarely had a dimensionality larger than 3 (13 at most, Fig 1K). Thus, we conclude that the critical subspace is always low dimensional and is always spanned by the first few PCs.

Temporal coarse-graining is required to reveal critical subspace

The activity in the critical subspace manifests as large amplitude fluctuations, coordinated across many neurons. Previous studies suggest that the spatiotemporal structure of such population activity can depend on the time scale of observation (*31*, *42–44*). Moreover, theory of critical phenomena suggests that temporal coarse-graining, i.e. excluding details at the small time scales, may be required to reveal universal properties of critical dynamics

(*45*). Therefore, we next sought to determine how the critical subspace depends on the time scale of observation ΔT (for the results in Fig 1, ΔT = 50 ms). In many previous studies of neuronal avalanches based on spike recordings (*30*, *31*, *46*, *47*), a common approach has been to set ΔT to the average interspike interval ⟨*ISI*⟩ for the entire population of recorded neurons, following the approach pioneered by (*48*). We note, however, that this approach was originally developed for LFP-peak events, not spikes. For our recordings here, the average single neuron spike rate was about 3 Hz. Thus, for a typical recording of 200 neurons, the ⟨*ISI*⟩ was about 1.5 ms. Obviously, ⟨*ISI*⟩ will be even smaller in experiments with more recorded neurons. Here we systematically investigated a range of ΔT between 1 ms and 500 ms.

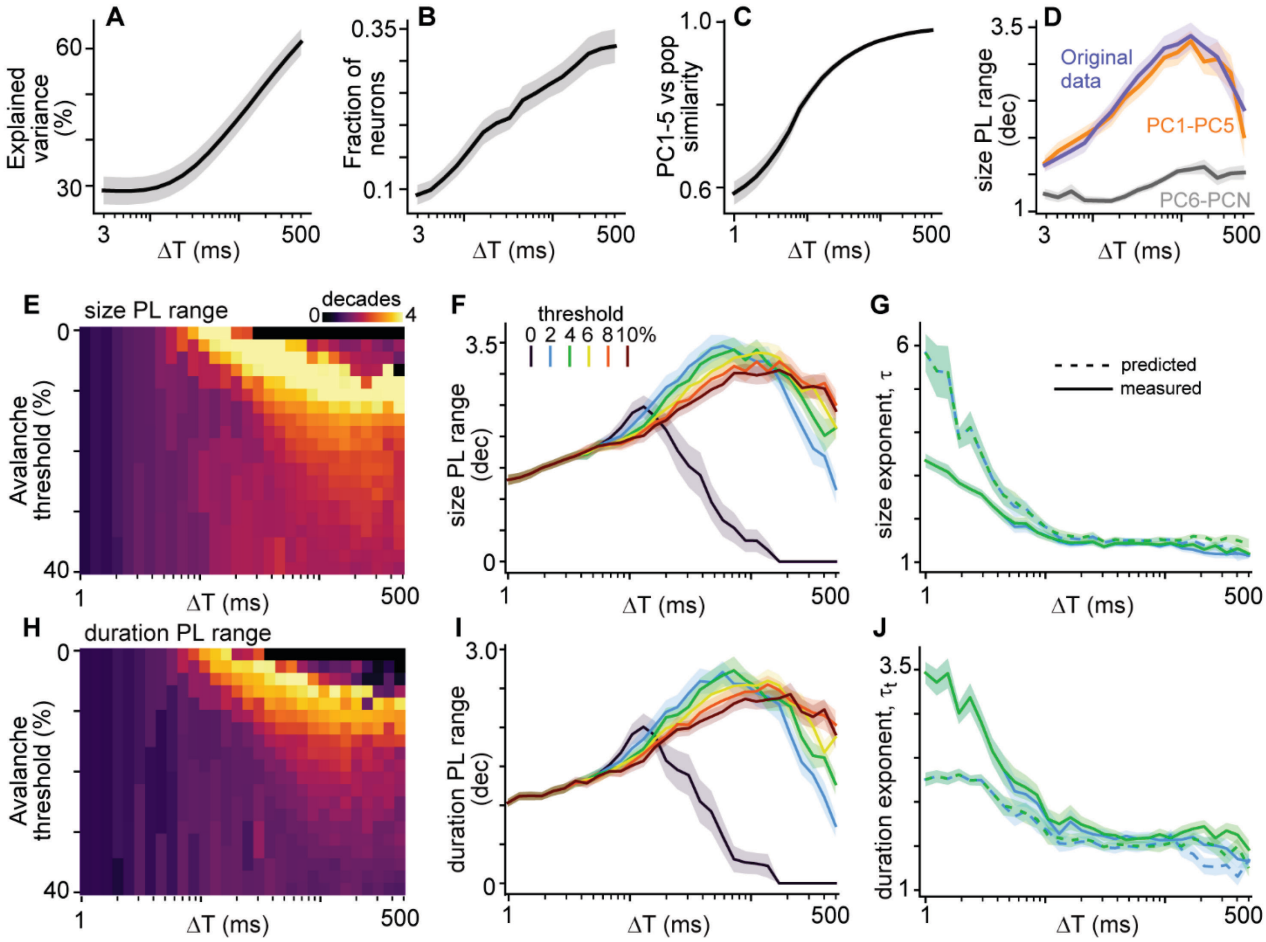

Fig 2. Temporal coarse-graining reveals critical subspace. (A) Total explained variance for the first five principal components is sensitive to temporal coarse-graining, sharply rising for ΔT > 10 ms, indicating the emergence of a low-dimensional subspace. **(B)**, Fraction of neurons engaged with (loading > 0.1) at least one of the first five principal components increases with ΔT. **C)** Similarity between the original summed population activity and that reconstructed using the first five principal components increases with ΔT. **(D)** Large power-law range is less than 2 decades for ΔT < 10 ms for the original data and PCs 1-5. For PCs 6-N, power-law range never exceeds 2 decades. **(E)** One example recording showing how power-law range (for population sum) depends on ΔT and the threshold used for defining avalanches. For large ΔT, the traditional approach (zero threshold) precludes finding a good power law. **(F)** Summary of how avalanche size power-law range depends on threshold and ΔT for all recordings.

(G) The measured exponent τ (solid lines) of the size distribution power-law decreases as ΔT increases for ΔT < 10 ms, but is relatively steady for ΔT > 10 ms. Only for ΔT > 10 ms do we find agreement between the measured τ and that predicted by the crackling noise scaling law (dashed). **(H-J)** Same as panels E-G, but for avalanche duration statistics. For all panels, solid lines represent the mean across all recordings and shaded areas represent standard error.

First, since the critical subspace coincides with the first several PCs, we quantified how the importance of these PCs depends on ΔT. We found that the variance explained by the first 5 PCs is relatively small for small ΔT, but rises sharply around ΔT ~ 10 ms (Fig 2A). Next, we asked how many neurons are involved in the first 5 PCs. In principle, it is possible that as few as 5 neurons are fully responsible for the first 5 PCs. We measured the fraction of neurons with strong engagement (loading > 0.1, Methods) with at least one of the first five PCs. We found that, as we increased ΔT, the fraction of neurons engaged with this low-dimensional subspace triples, reaching 30% (Fig 2B); note that 30% is around 30 to 60 neurons for our recordings. Thus, the importance of the first 5 PCs is hidden for small time scales, and emerges only after temporal coarse-graining. In our initial example with ΔT = 50 ms we saw that the population sum activity of the full population was very similar to that reconstructed from PCs 1-5 (Fig 1D,H). Next, we asked how the similarity between these two signals depends on ΔT. We found that they were not strongly correlated for small ΔT; this correlation sharply increased around ΔT ~ 10 ms (Fig 2C).

Next, we determined how ΔT impacts evidence for criticality based on avalanche analysis. We quantified the range of good power-law fit (number of decades, Methods) for the avalanche size distribution; we interpret a larger power-law range as better evidence for criticality. For both the full population and the PC 1-5 subspace, we found that at small time scales power-law range is small (Fig 2D) and the avalanche distribution is better fit by an exponential distribution (Supplementary Fig S1). The power-law range rises around ΔT ~ 10 ms. If we consider avalanches based on the PC 6-N subspace, power-law range is small for all ΔT (Fig 2D, grey). We note that the emergence of large power-law range at larger ΔT also depends on the threshold used for defining avalanches. Some previous studies have used a median threshold (*26*, *34*), some have used the 35th percentile (*32*, *33*), and many have used a zero threshold (*27*, *29–31*, *46*, *47*, *49*). Here we found that if the threshold is set to zero, the emergence of robust power-law statistics at larger ΔT will be missed (Fig 2 E,F,H,I). Finally, we examined the exponents of the power-laws and found that for ΔT < 10 ms, the exponents did not agree with the crackling noise scaling laws predicted at criticality and varied dramatically as ΔT changes (Fig 2G,J). For larger ΔT, the exponents were robust, less dependent on ΔT, and were in good agreement with crackling noise scaling laws (Fig 2G,J). Taken together, the results in Fig 2 show that evidence for criticality is robust in the critical subspace, but will be missed if spike data are not sufficiently coarse-grained in time. The critical subspace emerges for timescales above about 10 ms.

Theory of critical dynamics confirms experimental results

Does the necessity of temporal coarse-graining to find the critical subspace agree with theory? Could the temporal coarse-graining we employ mislead us, producing apparent critical dynamics from a system that is not actually at criticality? Here we address these questions, showing that our results agree with a simple, model – the multivariate Hawkes process (*50*, *51*) – but only if the model is tuned to operate close to criticality. Our model

is similar to branching processes and random walks (*40*, *45*, *52*), providing an established theoretical case of critical dynamics, but unlike many of these classic models, the Hawkes process treats time continuously. This is helpful for our goals here of studying how the dynamics depend on the timescale ΔT. Most discrete time models can have $\langle ISI \rangle$ well below the duration of one time step, which precludes quantitative comparison to previous studies which often set ΔT = $\langle ISI \rangle$.

We set up our model with N = 10^4 units with fixed, random connectivity. In a multivariate Hawkes process, each unit's spikes are drawn from an inhomogeneous Poisson process. At time t, the ith unit fires with rate $\lambda_i(t)$ which depends on the recent history of spikes from other connected units,

$$\lambda_i(t) = \mu + \sum_{j=1}^{N} C_{ij} \sum_{t_j<t} \phi(t - t_j),$$

where C is the connectivity matrix (C_{ij} = 1 with probability 0.1, otherwise C_{ij} = 0). The parameter μ determines a baseline rate of noisy or spontaneous firing, independent of interactions among units. The second sum accounts for the spikes fired by the jth unit (at times t_j) before time t and with weight that decays exponentially $\phi(t) = e^{-t}$. The magnitude of the largest eigenvalue of C, called Λ here, determines whether the system operates at criticality. For Λ well below 1, the units fire relatively asynchronously, the system is subcritical. As Λ approaches 1, firing rates and fluctuations become large and correlated and the autocorrelation diverges (Fig 3B,C); Λ = 1 at criticality.

We performed avalanche analysis on the model population activity, as we did for the experiments. We analyzed the dynamics of a subset of only 200 units, neglecting the rest, to account for subsampling effects that are certainly present in our experiments and may be important for assessing critical dynamics (*53*, *54*). When the model was near criticality (Λ = 0.99), the dynamics closely matched our experimental results. Avalanche sizes and durations were power-law distributed and obeyed the crackling noise scaling law (Fig 3D-F), but only after sufficient temporal coarse-graining (approximately ΔT > 10$\langle ISI \rangle$, Fig 3D,E,G,H, Supplementary Fig S1). When the model was subcritical (Λ = 0.9), the model dynamics were not power-law distributed for any ΔT, demonstrating that temporal coarse-graining does not produce evidence for criticality unless the system is in fact near criticality (Fig 3D,E,G,H, Supplementary Fig S1). When the noise parameter μ was increased, the importance of temporal coarse-graining increased (Supplementary Fig S1). We note that when we coarse-grained our experimental data using a ΔT set to a specified factor of $\langle ISI \rangle$ (instead of a specified number of milliseconds), the results of different experiments were more consistent with each other and in good quantitative agreement with the model (light colored lines in Fig 3D,E,G,H).

Like the experiments, critical dynamics in the model are confined to a low-dimensional subspace. In fact, the model is essentially 1-dimensional. Removing the first principal component fully removes the collective fluctuations in firing rates (Fig 3B). Moreover, removing the first PC completely abolishes the power-law statistics of avalanches (Fig 3I). In Supplementary Fig S2, we show that the model can generate higher dimensional dynamics if we implement spatially-dependent near-neighbor connectivity, similar to critical phenomena in other systems with spatial translational invariance (*55*). However, the neurons recorded from one high-density shank electrode like in our experiments are all within a few hundred micrometers from each other and, thus, are expected to have more

spatially independent connectivity. Thus, we conclude that our experimental measurements are well-explained by theory of low-dimensional critical dynamics.

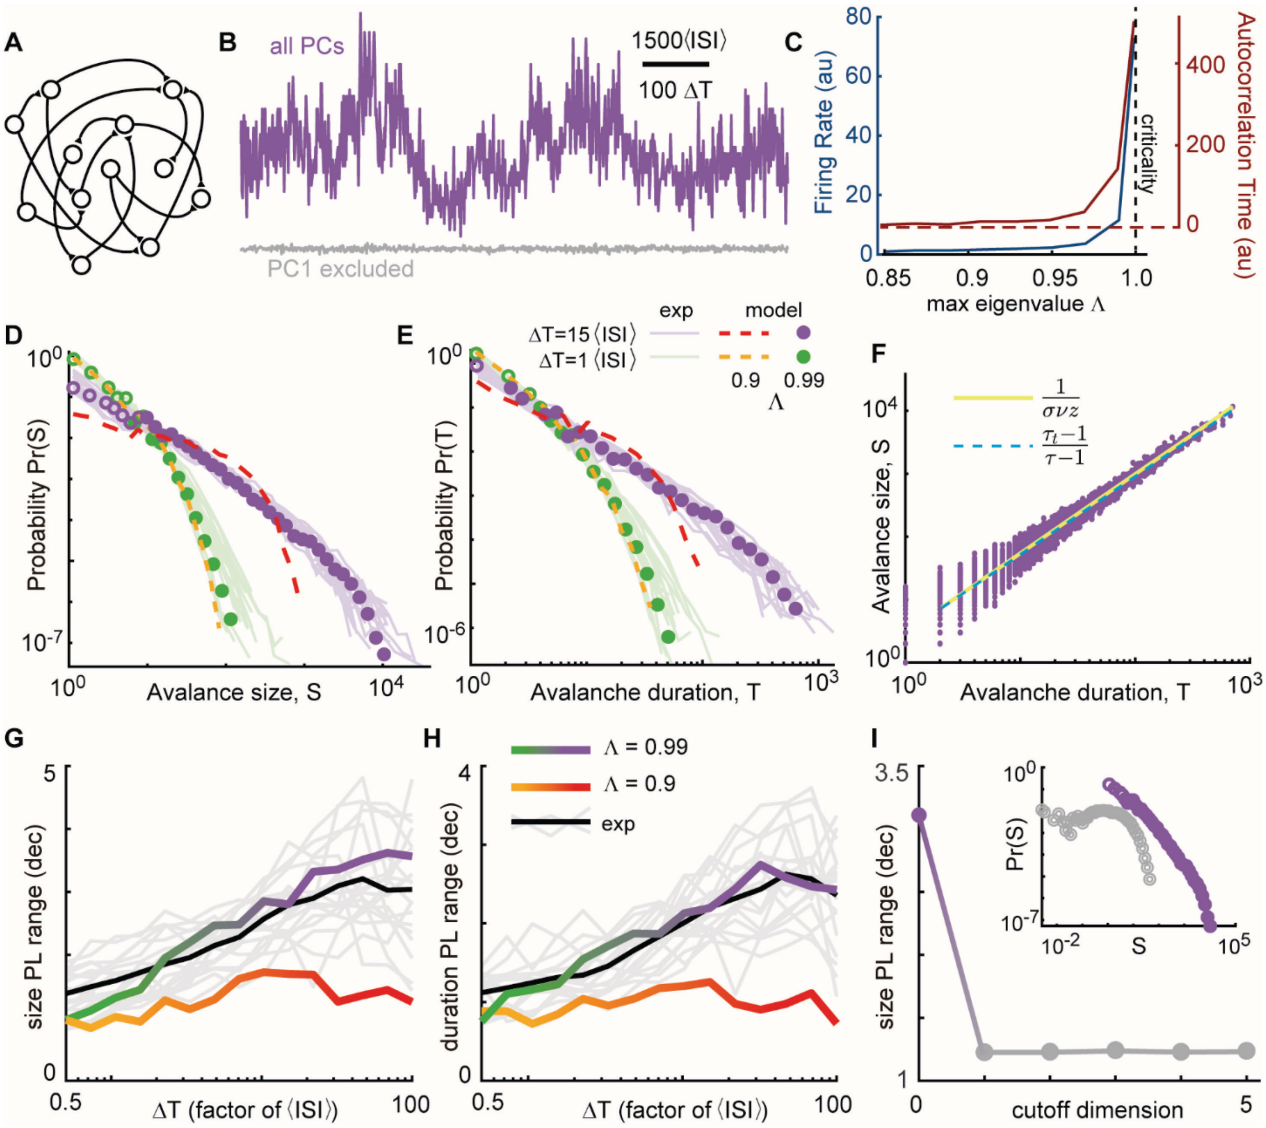

Fig 3. Low-dimensional model of critical dynamics agrees well with experiments. (A) We study a multivariate Hawkes process with random connectivity. **(B)** When the model is near criticality $\Lambda = 0.99$), the model generates population activity with large fluctuations (purple), which are abolished if the first principal component is projected out (gray). **(C)** As Λ approaches 1 (criticality), the firing rate and the timescales of fluctuations diverge. **(D)** Near criticality and with sufficient temporal coarse-graining, the model avalanche sizes were power-law distributed ($\Delta T = 15\langle ISI \rangle$, purple points), but the power-law was obscured for insufficient temporal coarse-graining ($\Delta T = 1\langle ISI \rangle$, green points). Subcritical model dynamics were not power-law distributed, with or without temporal coarse-graining (dashed lines). Experiments (light colored lines) agree well with the model. **(E)** Same as panel D, but for avalanche duration. **(F)** With $\Delta T = 15\langle ISI \rangle$, the model agrees with the crackling noise scaling law. **(G,H)** At criticality, the power-law range becomes large only for large ΔT. The subcritical model never exhibits more than 2 decades of power-law range for all ΔT. Individual experiments (gray) and mean

across experiments (black) agree best with model with Λ=0.99. **(I)** The model critical subspace is 1-dimensional; power-law range collapses to less than 1.5 decades after excluding PC1. Inset shows size distributions with (purple) and without (gray) PC1.

Desynchronized subspace coexists with the critical subspace

As discussed in the introduction, the standard point of view is that criticality is not compatible with desynchronized dynamical regimes; it is thought that the cortex must "choose" between criticality and a desynchronized regime. However, the fact that the critical subspace is low-dimensional raises an interesting question that challenges the standard view. Could there be a desynchronized dynamical regime that coexists with the critical subspace, in the same neural circuit, but in a different subspace?

As demonstrated in Figs 1C,J, and K, the number of dimensions needed to explain a substantial fraction of the total variance is much greater than the dimension of the critical subspace. This suggests that the additional dimensions beyond the critical subspace are "important", in the sense that they are needed to explain much of the measured variance. What type of dynamics is contained in these additional dimensions? The basic mathematical facts of PCA require that these dimensions outside the critical subspace must have smaller fluctuations, because the critical subspace spans the first few PCs. But PCA does not tell us whether the dynamics in these dimensions has the characteristic features of a desynchronized regime: weak correlations among units, short autocorrelation timescales, and Gaussian distributed population activity. To answer these questions we reconstructed our measured spike count dynamics using a small number of PCs (d_c+1 through 2d_c) just beyond the cutoff dimension d_c of the critical subspace, as defined in Fig 1J,K. Fig 4A shows an example recording where each subspace is 6-dimensional (i.e. d_c = 6).

First, we summed the reconstructed dynamics across units to obtain a single population summed activity for each subspace (Fig 4A). As expected, this population activity fluctuated with much larger amplitude (larger variance, Fig 4B,C) in the critical subspace compared to the desynchronized subspace. In addition, we found that the distribution of population activity in the desynchronized subspace was close to Gaussian (low skewness), while the critical subspace had heavy tailed, skewed distributions (Fig 4D). Another defining feature of desynchronized regimes is a lack of long-range temporal correlations. We compared autocorrelation functions for the population activity in the two subspaces and found that the critical subspace had much longer temporal correlations (Fig 4E). Finally, we assessed pairwise correlations among units, using the reconstructed spike count rasters (Fig 4A). Here, some caution is needed; the dimensionality of the subspace impacts the basic nature of pairwise correlations for such reconstructed unit activity. Nonetheless, defining the dimensionality of the two subspaces to be equal allows for meaningful relative comparisons and we found that the critical subspace always had much higher pairwise correlations than the desynchronized subspace. Thus we conclude that the dynamics in this subspace with the greatest variance, after removing the critical subspace, manifested all the defining features of desynchronized regimes: weak pairwise correlations, short-range temporal correlations, and Gaussian distributed population activity (*9*, *56*).

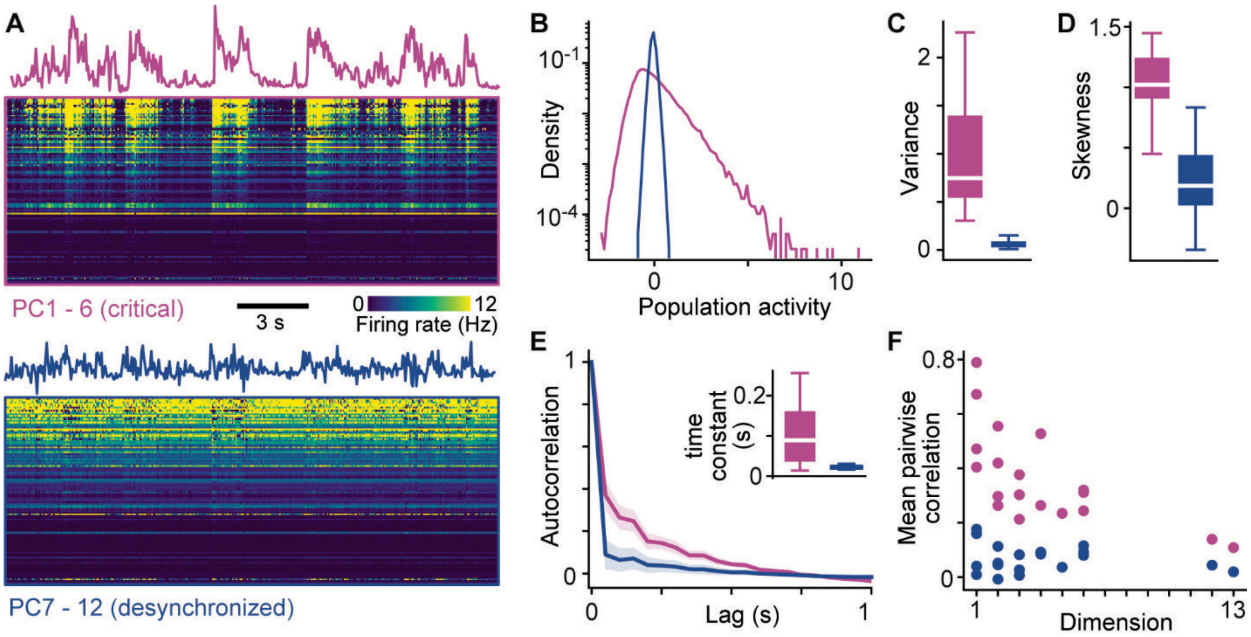

Fig 4. Desynchronized subspace coexists with critical subspace (A) Population summed activity and raster reconstructed from the critical subspace (purple, d_c=6 in this example) and the desynchronized subspace (blue). **(B - D)** The distribution of population summed activity is broad, heavy-tailed, and skewed for the critical subspace, but narrow and near gaussian for the desynchronized subspace. Box-plots summarize all experiments (box spans quartiles, white line is median). **(E)** For the desynchronized subspace, the autocorrelation function drops sharply indicating short-range temporal correlations compared to the critical subspace. Inset: time constants from an exponential fit for all experiments. **(F)** Mean pairwise correlations among reconstructed single unit activities are closer to zero for the desynchronized subspace.

Discussion

We have shown that population spiking activity in awake mouse motor cortex can be partitioned into different subspaces, each containing fundamentally different kinds of coordinated dynamics. The most prominent subspace – defined by the first few principal components – is home to critical dynamics with long-range temporal correlations, heavy-tailed distributions of activity and multifaceted agreement with scaling laws predicted at criticality. The next most prominent dimensions contain desynchronized dynamics with weak spatiotemporal correlations and Gaussian distributed population activity.

These observations open several interesting new questions and offer answers to several long-standing questions about the fundamental operating regime of cortical neuronal networks. One such long-standing question asks whether cortical circuits operate near criticality or not in awake cortex. Many previous experiments with strong agreement with criticality theory were based on techniques that measure collective brain signals, like local field potential (LFP) (*30*, *48*, *57–60*), wide field imaging (*25*, *61*), and human brain imaging (*57*, *62*). These collective signals represent an aggregate of the underlying spike activity of many individual neurons. However, when spikes recorded in awake animals have been analyzed directly, results have been less clear - some studies report support for criticality (*31*, *33–38*, *49*, *63*) while others do not (*27–32*). Considering that spikes are the fundamental information carriers underlying brain function, the equivocal support for

criticality at the level of spike measurements has created skepticism and confusion surrounding the hypothesis (*29*, *30*). Why is evidence for criticality clear in collective signals, but unclear in spike data? As we showed here, one important factor explaining discrepancies among previous spike recordings is that the existence of the critical subspace is clear only at larger timescales, beyond about 10 $\langle ISI \rangle$. Our results suggest that the studies based on spikes with negative reports about criticality missed the evidence for criticality because they did not sufficiently coarse-grain in time (Supplementary Table 1). What about collective signals like LFP? One possibility is that collective signals, by their nature, may carry some degree of overlapping signals (LFP at two nearby locations can reflect signals from the same source neurons), which results in spurious correlations that could be mistaken for criticality (*64*). Considering our results here, another possibility is that the coordinated activity in the critical subspace is apparent in measurements of collective signals like LFP, but the desynchronized subspace is hidden from collective signals because it is weakly coordinated. Thus, measuring collective signals may effectively filter out the desynchronized subspace, leaving only the critical subspace signals.

Several recent studies have proposed that critical phenomena in neural systems might be fundamentally high dimensional (*65–69*). Do these studies contradict our work here? Several of these previous studies focused on edge-of-chaos (EOC) criticality, which, unlike our results here, is thought to occur without large fluctuations at the population level (*66*, *68*, *69*). In other words, EOC criticality has no avalanches. In this basic sense, our experimental results are inconsistent with EOC criticality – we observe prominent large amplitude fluctuations and power-law distributed avalanches. In other studies (*55*, *65–67*), criticality was hypothesized to be high-dimensional based on data analysis and concepts adapted from traditional renormalization group in systems with spatial translational invariance (similar to the Ising model, for example). Our experimental results do not directly confirm or deny this hypothesis, but our model shows a clear case of low-dimensional criticality; thus, criticality in neural systems is not required to be high dimensional. We further clarify this point in Supplementary Fig S2; we show how a transition from low dimensional to high dimensional criticality can result from tuning the model connectivity from global to local and introducing approximate spatial translational invariance. These theoretical considerations suggest that we observed a low dimensional critical subspace in our experiments, because we measured neurons from a local patch of cortex (approximately one cortical column) using a single shank Neuropixels probe. Within the population of recorded neurons connectivity is approximately global. Thus, we might expect to find a higher dimensional critical subspace based on more spatially wide-spread recordings, as found in recent fMRI measurements, for example (*55*).

As discussed in the introduction, another controversy that our results may help resolve is the question of cortical state: Are cortical population dynamics weakly or strongly correlated? Our results show that this question may be a false dichotomy; we show that weakly correlated dynamics can coexist with relatively strongly dynamics by separating them into different subspaces. However, this observation raises important follow up questions. Do these two different subspaces perform different brain functions? Can the different subspaces really be used by the brain in a way that confers the different functional advantages of weakly versus strongly correlated dynamics? For example, desynchronized dynamics are thought to be beneficial for low-noise discrimination and representation of sensory input. Whereas criticality is thought to be beneficial for multi-scale information transmission, to name one example. Do previous studies support the

possibility that the critical subspace and the desynchronized subspace might execute different functions? Stringer et al (2019) showed that visual input is encoded in one subspace with relatively small fluctuations while some body movements (e.g. running and whisking) are encoded in a different subspace with larger fluctuations (first few PCs). Moreover, Jones et al (2023) reanalyzed the data from Stringer et al (2019) and found that subsets of neurons with the strongest correlations to body movements also exhibited power-law fluctuations similar to the avalanches we observed in our critical subspace here. Taken together, these results suggest that sensory input might be encoded in the desynchronized subspace we found here, while behavior may be encoded in the critical subspace. Another possible use for the critical subspace might be exchange of signals among cortical areas (*70*, *71*). These studies used ongoing activity in visual cortex (i.e. they studied the stimulus independent "noise" fluctuations, consistent with our critical subspace). Further studies will be needed to carefully test these interesting possible functional implications.

Materials and Methods

Animals

All procedures followed the Guide for the Care and Use of Laboratory Animals of the National Institutes of Health and were approved by University of Arkansas Institutional Animal Care and Use Committee (protocol 21022). We studied adult male C57BL6/6J mice (Jackson Labs). After acclimatization to handling, a small aluminum plate (0.5 g), was attached to the skull with dental cement. Then, mice were trained for head fixation for 20 sessions, gradually increasing in duration. At the time of recordings, the mice weighed ≈ 28 g and were 21-23 weeks old. 1-2 days before the first recording for each mouse, a craniotomy (2 mm diameter) was performed over right motor cortex (Anterior-Posterior = 0 mm, Medial-Lateral = 1 mm). Each recording day began with a brief period of isoflurane anesthesia to expose the craniotomy and head fix the mouse. The mice were free to run, sit, groom, and walk for the entire duration (45 minutes) of each recording. During recordings, after inserting the electrode array, the craniotomy was covered with gel-foam pieces soaked in sterile phosphate buffer solution.

Electrophysiology

The extracellular voltage was recorded using Neuropixels probes (NP version 1.0, IMEC) consisting of an electrode shank (width: 70 µm, length: 10 mm, thickness: 100 µm) of 960 total sites laid out in a checkerboard pattern with contacts (18 µm site-to-site separation), enabling up to 384 recording channels. On the recording day, following head fixation, the Neuropixels probe was inserted to a tip-depth of approximately 1.2 mm, ensuring that the active recording sites spanned all cortical layers. An Ag/AgCl pellet was used as ground, placed in the saline-soaked gel foam covering the craniotomy. The ground pellet wire was soldered to the Neuropixel midway along the ribbon cable. Electrophysiological data were collected (30 kHz) using SpikeGLX software. Spike sorting was performed using Kilosort 3.0 (https://github.com/MouseLand/Kilosort) and then manually curated using phy (https://github.com/cortex-lab/phy) (*72*).

Data analysis

PCA - We performed principal component analysis (PCA) in Python using the function 'decomposition.PCA' from package 'sklearn'. Let Z be a spike count matrix with M rows (number of time bins) and N columns (number of neurons). Then, PCA generates V which contains the principal components, i.e. the eigenvectors of the covariance matrix of Z. V has N rows and N columns (one column for each eigenvector, i.e. one column for each

principal component). We calculated % variance explained by a set of PCs as the sum of their corresponding eigenvalues of V (reported in Fig 1C and Fig 2A). To reconstruct data based on a subset of PCs (e.g. PCs 1-K), we first define the M × K projection matrix B by B=Z$\hat{V}$, where $\hat{V}$ is N × K, including the first K columns of V. Then the reconstructed activity is $\hat{Z}$=B$\hat{V}^{T}$=Z$\hat{V}\hat{V}^{T}$.

Avalanche analysis - The first step in avalanche analysis was to create a spike count matrix Z, where Z_{tj} is the number of spikes fired by unit j during time bin t. Next, the population sum spike count time series X was created by summing spike counts over all neurons at each time bin, $X_t = \sum_j Z_{tj}$. The threshold Θ used for avalanche detection was defined as some percentile of X. By definition, an avalanche begins when X exceeds the threshold and ends when X returns below threshold. The size S of an avalanche is defined as $S = \sum_{t_i}^{t_f}(X_t - \Theta)$, where the start and end times of the avalanche are t_i and t_f, respectively. Avalanche duration is defined as T = t_f - t_i.

In order to assess whether avalanche sizes and durations were distributed according to a power-law and to obtain power-law exponents and ranges, we built on previously developed maximum likelihood methods (*34*, *58*, *73*, *74*). In brief, the fitting algorithm identifies the best fit truncated power-law that meets a pre-defined goodness-of-fit criterion. There are three fitting parameters: the minimum avalanche size x_m, the maximum avalanche size x_M, and the power-law exponent τ. The following steps summarize the algorithm. First, events with size/duration less than x_m and larger than x_M were excluded. Second, the maximum likelihood power-law exponent was calculated. Third, we assessed the goodness-of-fit. We repeated these four steps for all the possible pairs of x_m and x_M values, in the end, identifying the largest power-law range that passed the goodness-of-fit criterion. We define power-law range as the number of decades of power-law scaling $\log_{10}(x_M/x_m)$. We note that this algorithm is independent of any choice of bins used to create the PDF plots in the paper.

The primary improvement we made compared to our most recently published methods (*34*) was to make our goodness-of-fit criterion less sensitive to sample size (number of avalanches) and more computationally efficient. For a given x_m, x_M, and τ, goodness-of-fit was quantified as follows. First, we created a cumulative distribution function (CDF) of the real data (excluding samples below x_m and x_M). Second, we define a theoretical CDF for a truncated power-law with the same range and exponent. Third, we define a region delimited by upper and lower bounds defined as the theoretical CDF +0.03 and -0.03, respectively. Fourth, we resample the real CDF at 10 logarithmically spaced values per decade. Fifth, we calculated the fraction F of resampled points in the CDF of the real data that fell within ±0.03 bounds of the theoretical CDF. F is our goodness-of-fit measure. F=1 means that the entire range of the real data varies less than 3% from a perfect power-law. We sought the fit with largest power-law range that meets the goodness-of-fit criterion F ≥ 0.8.

Decorrelated subspace – The analysis presented in Fig 4, began with making a spike count matrix with time bin size ΔT = 50 ms. Next, the activity was reconstructed using different dimensions as described above. The third step was to perform a band pass filter including 0.1 – 100 Hz on the reconstructed activity time series for each unit. Finally, the activity distributions, activity variance and skewness, autocorrelation functions, and pairwise correlations were analyzed. The time constant for the autocorrelation functions was the time constant of a best-fit exponential function.

Computational model

The numerical implementation of the multivariate Hawkes process described in Fig 3 and corresponding text was done according to ref (*75*).

References

1. C. Stringer, M. Pachitariu, N. Steinmetz, C. B. Reddy, M. Carandini, K. D. Harris, Spontaneous behaviors drive multidimensional, brainwide activity. *Science (80-.).* **364**, eaav7893 (2019).
2. P. Gao, E. Trautmann, B. Yu, G. Santhanam, S. Ryu, K. Shenoy, S. Ganguli, A theory of multineuronal dimensionality, dynamics and measurement. *bioRxiv* (2017), doi:10.1101/214262.
3. S. Recanatesi, S. Bradde, V. Balasubramanian, N. A. Steinmetz, E. Shea-Brown, A scale-dependent measure of system dimensionality. *Patterns*. **3**, 100555 (2022).
4. J. Li, P. A. Kells, A. C. Osgood, S. H. Gautam, W. L. Shew, Collapse of complexity of brain and body activity due to excessive inhibition and MeCP2 disruption. *Proc. Natl. Acad. Sci.* **118**, e2106378118 (2021).
5. L. Huang, P. Ledochowitsch, U. Knoblich, J. Lecoq, G. J. Murphy, R. C. Reid, S. E. J. de Vries, C. Koch, H. Zeng, M. A. Buice, J. Waters, L. Li, Relationship between simultaneously recorded spiking activity and fluorescence signal in gcamp6 transgenic mice. *Elife*. **10** (2021), doi:10.7554/ELIFE.51675.
6. D. Kobak, W. Brendel, C. Constantinidis, C. E. Feierstein, A. Kepecs, Z. F. Mainen, X. L. Qi, R. Romo, N. Uchida, C. K. Machens, Demixed principal component analysis of neural population data. *Elife*. **5**, 1–36 (2016).
7. M. M. Churchland, J. P. Cunningham, M. T. Kaufman, J. D. Foster, P. Nuyujukian, S. I. Ryu, K. V. Shenoy, K. V. Shenoy, Neural population dynamics during reaching. *Nature*. **487**, 51–56 (2012).
8. P. Gao, S. Ganguli, On simplicity and complexity in the brave new world of large-scale neuroscience. *Curr. Opin. Neurobiol.* **32**, 148–155 (2015).
9. A. Renart, J. de la Rocha, P. Bartho, L. Hollender, N. Parga, A. Reyes, K. D. Harris, The Asynchronous State in Cortical Circuits. *Science (80-.).* **327**, 587–590 (2010).
10. C. van Vreeswijk, H. Sompolinsky, Chaotic balanced state in a model of cortical circuits. *Neural Comput.* **10**, 1321–71 (1998).
11. A. S. Ecker, P. Berens, G. A. Keliris, M. Bethge, N. K. Logothetis, A. S. Tolias, Decorrelated Neuronal Firing in Cortical Microcircuits. *Science (80-.).* **327**, 26–31 (2010).
12. A. Palmigiano, T. Geisel, F. Wolf, D. Battaglia, Flexible information routing by transient synchrony. *Nat. Neurosci.* **20**, 1014–1022 (2017).
13. C. M. Gray, W. Singer, Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc. Natl. Acad. Sci. U. S. A.* **86**, 1698–702 (1989).
14. A. M. Bastos, J. Vezoli, C. A. Bosman, J.-M. Schoffelen, R. Oostenveld, J. R. Dowdall, P. De Weerd, H. Kennedy, P. Fries, Visual Areas Exert Feedforward and Feedback Influences through Distinct Frequency Channels. *Neuron*. **85**, 390–401 (2015).
15. J. M. Beggs, *The Cortex and the Critical Point* (The MIT Press, Cambridge, MA, 2022; https://direct.mit.edu/books/book/5372/The-Cortex-and-the-Critical-PointUnderstanding-the).
16. M. A. Muñoz, Colloquium: Criticality and dynamical scaling in living systems. *Rev. Mod. Phys.* **90**, 31001 (2018).
17. W. L. Shew, D. Plenz, The Functional Benefits of Criticality in the Cortex. *Neuroscientist* (2012), doi:10.1177/1073858412445487.
18. J. Wilting, V. Priesemann, 25 years of criticality in neuroscience — established results, open controversies, novel concepts. *Curr. Opin. Neurobiol.* **58**, 105–111 (2019).

19. L. Cocchi, L. L. Gollo, A. Zalesky, M. Breakspear, Criticality in the brain: A synthesis of neurobiology, models and cognition. *Prog. Neurobiol.* **158**, 132–152 (2017).
20. W. L. Shew, H. Yang, S. Yu, R. Roy, D. Plenz, Information Capacity and Transmission Are Maximized in Balanced Cortical Networks with Neuronal Avalanches. *J. Neurosci.* **31**, 55–63 (2011).
21. W. L. Shew, H. Yang, T. Petermann, R. Roy, D. Plenz, Neuronal avalanches imply maximum dynamic range in cortical networks at criticality. *J. Neurosci.* **29**, 15595–600 (2009).
22. O. Kinouchi, M. Copelli, Optimal dynamical range of excitable networks at criticality. *Nat. Phys.* **2**, 348–351 (2006).
23. T. Ezaki, E. Fonseca dos Reis, T. Watanabe, M. Sakaki, N. Masuda, Closer to critical resting-state neural dynamics in individuals with higher fluid intelligence. *Commun. Biol.* **3** (2020), doi:10.1038/s42003-020-0774-y.
24. W. P. Clawson, N. C. Wright, R. Wessel, W. L. Shew, Adaptation towards scale-free dynamics improves cortical stimulus discrimination at the cost of reduced detection. *PLOS Comput. Biol.* **13**, e1005574 (2017).
25. E. D. Fagerholm, G. Scott, W. L. Shew, C. Song, R. Leech, T. Knöpfel, D. J. Sharp, Cortical Entropy, Mutual Information and Scale-Free Dynamics in Waking Mice. *Cereb. Cortex*, 1–8 (2016).
26. S. H. Gautam, T. T. Hoang, K. McClanahan, S. K. Grady, W. L. Shew, Maximizing Sensory Dynamic Range by Tuning the Cortical State to Criticality. *PLOS Comput. Biol.* **11**, e1004576 (2015).
27. C. Bédard, H. Kröger, A. Destexhe, Does the 1/f frequency scaling of brain signals reflect self-organized critical states? *Phys. Rev. Lett.* **97**, 1–4 (2006).
28. T. L. Ribeiro, M. Copelli, F. Caixeta, H. Belchior, D. R. Chialvo, M. a L. Nicolelis, S. Ribeiro, Spike avalanches exhibit universal dynamics across the sleep-wake cycle. *PLoS One*. **5**, e14129 (2010).
29. J. Touboul, A. Destexhe, Can power-law scaling and neuronal avalanches arise from stochastic dynamics? *PLoS One*. **5**, e8982 (2010).
30. N. Dehghani, N. G. Hatsopoulos, Z. D. Haga, R. a Parker, B. Greger, E. Halgren, S. S. Cash, A. Destexhe, Avalanche Analysis from Multielectrode Ensemble Recordings in Cat, Monkey, and Human Cerebral Cortex during Wakefulness and Sleep. *Front. Physiol.* **3**, 302 (2012).
31. V. Priesemann, M. Wibral, M. Valderrama, R. Pröpper, M. Le Van Quyen, T. Geisel, J. Triesch, D. Nikolić, M. H. J. Munk, Spike avalanches in vivo suggest a driven, slightly subcritical brain state. *Front. Syst. Neurosci.* **8**, 108 (2014).
32. Z. Ma, H. Liu, T. Komiyama, R. Wessel, Stability of motor cortex network states during learning-associated neural reorganizations. *J. Neurophysiol.* **124**, 1327–1342 (2020).
33. Z. Ma, G. G. Turrigiano, R. Wessel, K. B. Hengen, Cortical Circuit Dynamics Are Homeostatically Tuned to Criticality In Vivo. *Neuron*. **104**, 655-664.e4 (2019).
34. S. A. Jones, J. H. Barfield, V. K. Norman, W. L. Shew, Scale-free behavioral dynamics directly linked with scale-free cortical dynamics. *Elife*. **12** (2023), doi:10.7554/eLife.79950.
35. Z. Bowen, D. E. Winkowski, S. Seshadri, D. Plenz, P. O. Kanold, Neuronal Avalanches in Input and Associative Layers of Auditory Cortex. *Front. Syst. Neurosci.* **13** (2019), doi:10.3389/fnsys.2019.00045.
36. T. Bellay, A. Klaus, S. Seshadri, D. Plenz, Irregular spiking of pyramidal neurons organizes as scale-invariant neuronal avalanches in the awake state. *Elife*. **4**, 1–25 (2015).
37. Y. Karimipanah, Z. Ma, J.-E. K. Miller, R. Yuste, R. Wessel, Neocortical activity is stimulus- and scale-invariant. *PLoS One*. **12**, e0177396 (2017).

38. D. Curic, V. E. Ivan, D. T. Cuesta, I. M. Esteves, M. H. Mohajerani, A. J. Gruber, J. Davidsen, Deconstructing scale-free neuronal avalanches: behavioral transitions and neuronal response. *J. Phys. Complex.* **2**, 045010 (2021).
39. S.-S. Poil, R. Hardstone, H. D. Mansvelder, K. Linkenkaer-Hansen, Critical-state dynamics of avalanches and oscillations jointly emerge from balanced excitation/inhibition in neuronal networks. *J. Neurosci.* **32**, 9817–23 (2012).
40. L. Laurson, X. Illa, M. J. Alava, The effect of thresholding on temporal avalanche statistics. *J. Stat. Mech. Theory Exp.* **2009** (2009), doi:10.1088/1742-5468/2009/01/P01019.
41. J. P. Sethna, K. a Dahmen, C. R. Myers, Crackling noise. *Nature*. **410**, 242–50 (2001).
42. D. A. Leopold, N. K. Logothetis, Spatial patterns of spontaneous local field activity in the monkey visual cortex. *Rev. Neurosci.* **14**, 195–205 (2003).
43. B. Doiron, A. Litwin-Kumar, R. Rosenbaum, G. K. Ocker, K. Josić, The mechanics of state-dependent neural correlations. *Nat. Neurosci.* **19**, 383–393 (2016).
44. M. Okun, N. A. Steinmetz, A. Lak, M. Dervinis, K. D. Harris, Distinct Structure of Cortical Population Activity on Fast and Infraslow Timescales. *Cereb. Cortex*. **29**, 2196–2210 (2019).
45. U. C. Tauber, *Critical Dynamics: A Field Theory Approach to Equilibrium and Non-equilibrium Scaling Behavior* (Cambridge University Press, Cambridge, 2014).
46. T. L. Ribeiro, M. Copelli, F. Caixeta, H. Belchior, D. R. Chialvo, M. a L. Nicolelis, S. Ribeiro, Spike avalanches exhibit universal dynamics across the sleep-wake cycle. *PLoS One*. **5**, e14129 (2010).
47. G. Hahn, A. Ponce-Alvarez, C. Monier, G. Benvenuti, A. Kumar, F. Chavane, G. Deco, Y. Frégnac, Spontaneous cortical activity is transiently poised close to criticality. *PLOS Comput. Biol.* **13**, e1005543 (2017).
48. J. M. Beggs, D. Plenz, Neuronal avalanches in neocortical circuits. *J. Neurosci.* **23**, 11167–77 (2003).
49. A. J. Fontenele, N. A. P. De Vasconcelos, T. Feliciano, L. A. A. Aguiar, C. Soares-Cunha, B. Coimbra, L. Dalla Porta, S. Ribeiro, A. J. Rodrigues, N. Sousa, P. V. Carelli, M. Copelli, Criticality between Cortical States. *Phys. Rev. Lett.* **122**, 208101 (2019).
50. F. Y. K. Kossio, S. Goedeke, B. Van Den Akker, B. Ibarz, R. M. Memmesheimer, Growing Critical: Self-Organized Criticality in a Developing Neural System. *Phys. Rev. Lett.* **121**, 58301 (2018).
51. D. Notarmuzi, C. Castellano, A. Flammini, D. Mazzilli, F. Radicchi, Percolation theory of self-exciting temporal processes. *Phys. Rev. E*. **103**, L020302 (2021).
52. J. Marro, R. Dickman, *Nonequilibrium phase transitions in lattice models* (Cambridge University Press, 2005).
53. A. Levina, V. Priesemann, Subsampling scaling. *Nat. Commun.* **8**, 1–9 (2017).
54. A. Levina, V. Priesemann, J. Zierenberg, Tackling the subsampling problem to infer collective properties from limited data. *Nat. Rev. Phys.* **4**, 770–784 (2022).
55. A. Ponce-Alvarez, M. L. Kringelbach, G. Deco, Critical scaling of whole-brain resting-state dynamics. *Commun. Biol.* **6**, 1–11 (2023).
56. K. D. Harris, A. Thiele, Cortical state and attention. *Nat. Rev. Neurosci.* **12**, 509–23 (2011).
57. A. Zhigalov, G. Arnulfo, L. Nobili, S. Palva, J. M. Palva, Relationship of Fast- and Slow-Timescale Neuronal Dynamics in Human MEG and SEEG. *J. Neurosci.* **35**, 5385–5396 (2015).
58. W. L. Shew, W. P. Clawson, J. Pobst, Y. Karimipanah, N. C. Wright, R. Wessel, Adaptation to sensory input tunes visual cortex to criticality. *Nat. Phys.* **11**, 659–663 (2015).
59. S. Yu, T. L. Ribeiro, C. Meisel, S. Chou, A. Mitz, R. Saunders, D. Plenz, Maintained

avalanche dynamics during task-induced changes of neuronal activity in nonhuman primates. *Elife*. **6** (2017), doi:10.7554/eLife.27119.
60. T. Petermann, T. C. Thiagarajan, M. A. Lebedev, M. A. L. Nicolelis, D. R. Chialvo, D. Plenz, Spontaneous cortical activity in awake monkeys composed of neuronal avalanches. *Proc. Natl. Acad. Sci. U. S. A.* **106**, 15921–6 (2009).
61. G. Scott, E. D. Fagerholm, H. Mutoh, R. Leech, D. J. Sharp, W. L. Shew, T. Knopfel, Voltage Imaging of Waking Mouse Cortex Reveals Emergence of Critical Neuronal Dynamics. *J. Neurosci.* **34**, 16611–16620 (2014).
62. E. Tagliazucchi, P. Balenzuela, D. Fraiman, D. R. Chialvo, Criticality in large-scale brain FMRI dynamics unveiled by a novel point process analysis. *Front. Physiol.* **3**, 15 (2012).
63. A. Ponce-Alvarez, A. Jouary, M. Privat, G. Deco, G. Sumbre, Whole-Brain Neuronal Activity Displays Crackling Noise Dynamics. *Neuron*. **100**, 1446-1459.e6 (2018).
64. J. P. Neto, F. Paul Spitzner, V. Priesemann, Sampling effects and measurement overlap can bias the inference of neuronal avalanches. *PLoS Comput. Biol.* **18**, 1–21 (2022).
65. S. Bradde, W. Bialek, PCA Meets RG. *J. Stat. Phys.* **167**, 462–475 (2017).
66. G. B. Morales, S. di Santo, M. A. Muñoz, Quasiuniversal scaling in mouse-brain neuronal activity stems from edge-of-instability critical dynamics. *Proc. Natl. Acad. Sci. U. S. A.* **120**, 1–12 (2023).
67. L. Meshulam, J. L. Gauthier, C. D. Brody, D. W. Tank, W. Bialek, Coarse Graining, Fixed Points, and Scaling in a Large Population of Neurons. *Phys. Rev. Lett.* **123**, 178103 (2019).
68. D. Dahmen, S. Grün, M. Diesmann, M. Helias, Second type of criticality in the brain uncovers rich multiple-neuron dynamics. *Proc. Natl. Acad. Sci. U. S. A.* **116**, 13051–13060 (2019).
69. J. O'Byrne, K. Jerbi, How critical is brain criticality? *Trends Neurosci.* **45**, 820–837 (2022).
70. J. D. Semedo, A. I. Jasper, A. Zandvakili, A. Krishna, A. Aschner, C. K. Machens, A. Kohn, B. M. Yu, Feedforward and feedback interactions between visual cortical areas use different population activity patterns. *Nat. Commun.* **13**, 1–14 (2022).
71. J. D. Semedo, A. Zandvakili, C. K. Machens, B. M. Yu, A. Kohn, Cortical Areas Interact through a Communication Subspace. *Neuron*. **102**, 249-259.e4 (2019).
72. M. Pachitariu, N. Steinmetz, S. Kadir, M. Carandini, K. D. Harris, Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. *bioRxiv*, 061481 (2016).
73. A. Clauset, C. R. Shalizi, M. E. J. Newman, Power-Law Distributions in Empirical Data. *SIAM Rev.* **51**, 661–703 (2009).
74. D. Langlois, D. Cousineau, J. P. Thivierge, Maximum likelihood estimators for truncated and censored power-law distributions show how neuronal avalanches may be misevaluated. *Phys. Rev. E*. **89**, 012709 (2014).
75. A. Dassios, H. Zhao, Exact simulation of Hawkes process with exponentially decaying intensity. *Electron. Commun. Probab.* **18** (2013), doi:10.1214/ECP.v18-2717.

Acknowledgments

Funding:
National Institutes of Health grant R15NS116742 (AJF, VKN, SHG, WLS)
National Science Foundation grant R01AB123456 (AJF, SHG, WLS)
Arkansas Biosciences Institute (SHG, WLS)

Author contributions:
Conceptualization: AJF, JSS, WLS

Experimental recordings: AJF, VKN, SHG, WLS
Data analysis: AJF, JSS, WLS
Computational modeling: JSS
Figure creation: AJF, JSS, WLS
Writing: WLS, AJF, JSS

Competing interests: All authors declare that they have no competing interests.

Data and materials availability: Spike data and analysis code are publicly available without restriction on Figshare at doi: *TBD upon paper acceptance*. All data are available in the main text or the supplementary materials.

Supplementary Materials

Supplementary Materials include Supplementary Text, two Supplementary Figures, and one Supplementary Table.

Supplementary Materials for

Low dimensional criticality embedded in high dimensional awake brain dynamics

Fontenele *et al.*

*Corresponding author. Email: shew@uark.edu

This PDF file includes:

Supplementary Text
Figs. S1 to S2
Table S1

Supplementary Text

Temporal coarse-graining reveals power-laws

In the main text, we presented power-law distributed avalanches after sufficient temporal coarse-graining and claimed that avalanche distributions without coarse-graining were not power laws. Here directly support this claim, quantitatively showing that when ΔT is too small, the avalanche distributions are better fit by an exponential distribution than a power-law. We compared the quality of the maximum likelihood exponential fit ($P_{exp}(\mathrm{S}) \propto \mathrm{e}^{-\mathrm{S}/\mathrm{a}}$, fitting parameter a) to that of the maximum likelihood power-law fit ($P_{pl}(\mathrm{S}) \propto \mathrm{S}^{-\tau}$, fitting parameter τ) using the Akaike information criterion (AIC) (Fontenele et al., 2019; Hirotugo, 1974).

$$AIC \coloneqq 2k - \hat{l} + \frac{2k^2 + 2k}{N - k - 1}$$

where k = 1 is the number of fitting parameters, N is the number of samples, and $\hat{l}$ is the log-likelihood at its maximum. Unlike in the power-law fitting algorithm used in the main text, here we included all observed avalanches in the fit, without imposing an upper or lower cutoff. Smaller AIC means a better fit, so $\Delta \coloneqq AIC_{pl} - AIC_{exp}$ is greater than zero when the exponential fit is better than the power-law fit, and less than zero when the reverse is true. As shown in Fig S1, the exponential was a better fit for all experiments except 1 when ΔT=$\langle ISI \rangle$ and the power-law was a better fit for all experiments when ΔT=30$\langle ISI \rangle$.

Dimensionality of critical dynamics depends on connectivity of units

In the main text Discussion, we mentioned that some previous studies have hypothesized that critical dynamics in neural systems is high dimensional, unlike the low-dimensional critical subspace we presented. Here we show using our computational model, that the dimensionality of critical dynamics depends strongly on the nature of connectivity among neurons in the model. The main results in Fig 3 are based on global random connectivity; each pair of neurons is connected with probability p = 0.1, independent of the distance between the two neurons. In this case, the critical dynamics are 1-dimensional as we showed in Fig 3. However, if we consider neurons on a square lattice (50 x 50) with local connectivity (connection probability is strongly dependent on distance between neurons), then we find that the critical dynamics of the model become higher dimensional (Fig S2). The spatial reach of local connectivity is inversely related to the dimensionality of the critical subspace (computed in the same way as in Fig 1 of the main text).

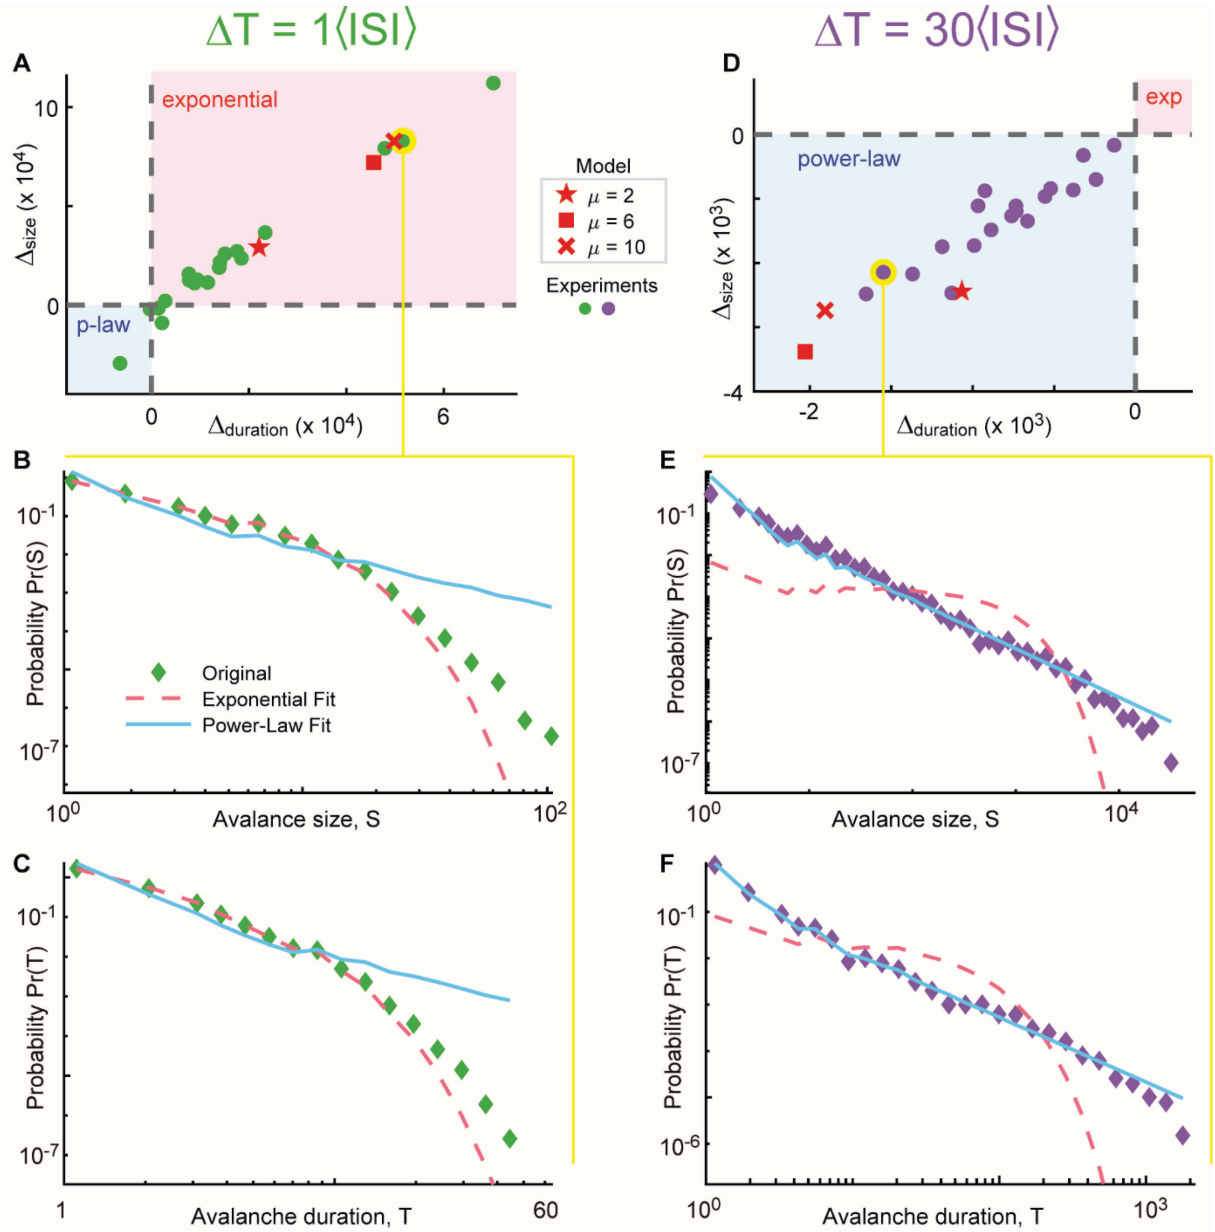

Fig. S1. Without temporal coarse graining, avalanches are exponentially distributed. (**A**) Each point represents one experiment (green) or one run from the model (red) analyzed using ΔT=⟨ISI⟩. The vertical position of each point represents the comparison $\Delta \coloneqq AIC_{pl} - AIC_{exp}$ for the avalanche size distribution; the horizontal position represents the same comparison, bur for the duration distribution. A point that falls in the pink shaded area is best fit by an exponential distribution for both size and duration. Note that, for the model, increasing uncorrelated noise (μ) results in better exponential fits, suggesting that uncorrelated firing can be one reason that temporal coarse graining is needed. (**B,C**) The size and duration distributions for the example point highlighted in yellow in panel A. Clearly, the exponential fit (dashed) is better than the best power-law fit (blue line). (**D-F**) Same as panels A-c, but ΔT= 30⟨ISI⟩, resulting in all experiments being better fit by a power-law for both size and duration distributions (blue shaded region). Note: the blue lines are perfect power laws. The slight deviations from a straight line on log-log axes are due to logarithmic binning to create the PDF plots.

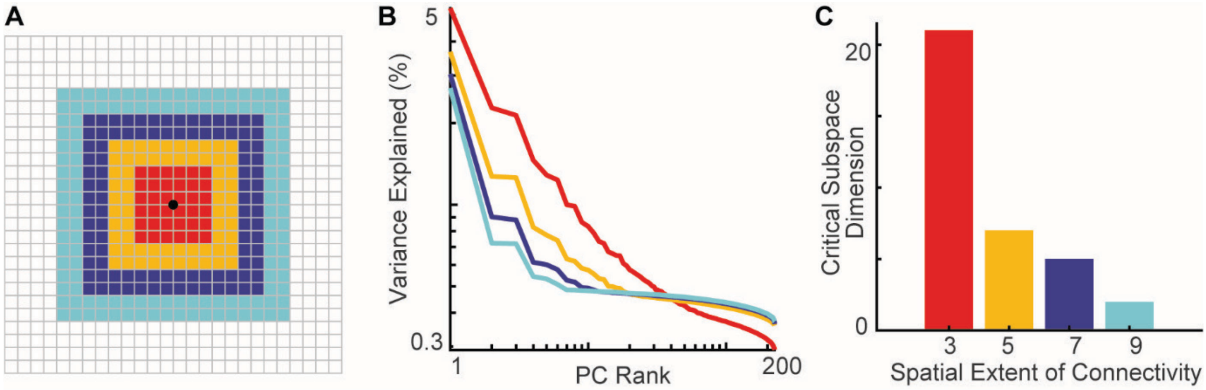

Fig. S2. Spatial translational invariance and local connectivity results in higher dimensional critical subspace.

(**A**) We simulated a multivariate Hawkes Process on a 50 x 50 lattice, with outgoing connections from each site to the sites in an L x L square centered on it. All nonzero connections are excitatory and of equal strength. As in the main text, we scaled the connectivity matrix to have largest eigenvalue $\Lambda = 0.99$, poising the network close to criticality, and set the baseline spiking rate to $\mu = 6$. For the avalanche analysis and PCA, we took a 15 x 15 localized subsample and chose a time bin of length $\Delta T = 30\langle ISI \rangle$. Increasing the spatial extent L of the connectivity decreased both the dimension of the critical subspace (**C**) and the number of PCs over which the PC eigenspectrum decays as an approximate power-law (**B**).

Authors	Journal	Year	Orgaanism	Measurement device	Time bin duration (ms)	Avalanche threshold (%)	high density	Temporal resolution (ms)
Ponce-Alvarez et al	Neuron	2018	zebra fish	2p Ca imaging	470	n/a	X	470
Jones et al	biorXiv	2022	mouse	2p Ca imaging	300	50	X	300
Karimipanah et al	Plos One	2017	mouse	2p Ca imaging	250	50	X	250
Curic et al	J Phys Complexity	2021	mouse	2p Ca imaging	50	0*	X	50
Ma et al	Neuron	2019	mouse	16 ch array	40	35	X	<1
Ma et al	J Neurophys	2020	mouse	2p Ca imaging	30	35	X	30
Bowen et al	Frontiers Sys Neuro	2019	mouse	2p Ca imaging	30	0*	X	30
Bellay et al	eLife	2015	mouse	2p Ca imaging	30	0*	X	30
Hahn et al	Plos Comp Biol	2017	monkey	96 ch MEA (utah)	14 - 31	0		<1
Priesemann et al	Frontiers Sys Neuro	2014	rat	32 or 64 ch Buzsaki probes	16	0	X	<1
			cat	54 ch polytrodes	16	0	X	<1
			monkeys	16 single wires (4x4 grid)	16	0		<1
Dehghani et al	Frontiers Physiol	2012	cat	96 ch MEA (utah)	1 - 16	0		<1
			monkey	96 ch MEA (utah)	1 - 16	0		<1
			humans	96 ch MEA (utah)	1 - 16	0		<1
Ribiero et al	Plos One	2010	rat	16-32 ch micro wire	1 - 20	0		<1
Fontenele et al	Phys Rev Lett	2019	mouse	64 ch array	3	0	X	<1
Priesemann et al	Frontiers Sys Neuro	2014	rat	32 or 64 ch Buzsaki probes	0.5	0	X	<1
			cat	54 ch polytrodes	0.5	0	X	<1
			monkeys	16 single wires (4x4 grid)	0.5	0		<1
Touboul & Destexhe	Plos One	2010	cat	MEA (8 ch)	n/r	0		<1
Bedard & Destexhe	Phys Rev Lett	2006	cat	MEA (8 ch)	n/r	0		<1

Table S1. Insufficient temporal coarse-graining explains previous discrepancies.

We summarize all previous experiments that performed avalanche analysis based on spikes recorded from awake animals. Experiments that reported poor evidence for power-law distributed avalanches are shaded yellow. The rows of the table are ordered from top to bottom in order of largest to smallest amount of temporal coarse graining. These results suggest that previous negative reports are due to insufficient temporal coarse-graining.