Latent dynamical variables produce signatures of spatiotemporal criticality in large biological systems

Mia C. Morrell,^{1,*} Audrey J. Sederberg,^{1,2} and Ilya Nemenman^{1,3,2}

¹*Department of Physics, Emory University, Atlanta, GA 30322, USA*

²*Initiative in Theory and Modeling of Living Systems, Emory University, Atlanta, GA 30322, USA*

³*Department of Biology, Emory University, Atlanta, GA 30322, USA*

(Dated: August 19, 2020)

Understanding the activity of large populations of neurons is difficult due to the combinatorial complexity of possible cell-cell interactions. To reduce the complexity, coarse-graining had been previously applied to experimental neural recordings, which showed over two decades of scaling in free energy, activity variance, eigenvalue spectra, and correlation time, hinting that the mouse hippocampus operates in a critical regime. We model the experiment by simulating conditionally independent binary neurons coupled to a small number of long-timescale stochastic fields and then replicating the coarse-graining procedure and analysis. This reproduces the experimentally-observed scalings, suggesting that they may arise from coupling the neural population activity to latent dynamic stimuli. Further, parameter sweeps for our model suggest that emergence of scaling requires most of the cells in a population to couple to the latent stimuli, predicting that even the celebrated place cells must also respond to non-place stimuli.

A key problem in modern biological physics is extracting useful knowledge from massive data sets enabled by high-throughput experimentation. For example, now one can record simultaneous states of thousands of neurons [1–5] or gene expressions [6–8], or the abundances of species in microbiomes [9–11]. Inferring and interpreting the joint probability distributions of so many variables is infeasible. A promising resolution to the problem is to adapt the Renormalization Group (RG) [12] framework for coarse-graining systems in statistical physics to find relevant features and large-scale behaviors in biological data sets as well. Indeed, recently, RG-inspired coarse-graining showed an emergence of nontrivial scaling behaviors in neural populations [13, 14]. Specifically, the authors analyzed the activity of over 1000 neurons in the mouse hippocampus as the animal repeatedly ran through a virtual maze. Their coarse-graining scheme involved combining the most correlated neurons into neural clusters by analogy with Kadanoff’s hyperspins [15], while using cluster-cluster correlations as a proxy for locality. Various correlation functions of neural clusters exhibited self-similarity for different cluster sizes, suggestive of criticality. Further analysis inspired by Wilson’s momentum space approach to renormalization [16] revealed that the joint distribution of cluster activities flowed to a non-trivial, non-Gaussian fixed point. Mechanisms responsible for these behaviors remain unknown. Thus it is unclear which other systems may exhibit them.

Observation and interpretation of signatures of criticality in high-throughput biological experiments is a storied field [17–22]. As a specific example, one commonly observed signature is the Zipf’s law, which describes a power-law relation between the rank and the frequency of

a system’s states. It has been explained by the existence of stationary latent (unobserved) fields (such as stimuli or internal states) that couple neurons (spins) over long distances [23, 24]. Similarly, here we show that the observations of Ref. [14], including scaling properties of the free energy, the cluster covariance, the cluster autocorrelations, and the flow of the cluster activity distribution to a non-Gaussian fixed point can be explained, within experimental error, by a model of non-interacting neurons coupled to latent *dynamical* fields. This is the first model to explain such a variety of spatio-temporal scaling phenomena observed in large-scale biological data.

Below we introduce the model, implement the coarse-graining of Ref. [14] on data generated from it and compare our findings with experimental results. We conclude by discussing which other experimental systems may exhibit similar scaling relations under the RG procedure.

The model. — To understand how scaling relationships could arise from coarse-graining data from large-scale systems, we study a model of N binary neurons (spins) $s_i \in \{0, 1\}$, $i \in [1, N]$, where $s_i = 0$ or 1 corresponds to a neuron being silent or active. The neurons are conditionally independent and coupled only by N_f fields $h_m(t)$, $m \in [1, N_f]$ such that the probability of a population being in a certain state $\{s_i\}$ is

$$P(\{s_i\}|\{h_m\}) = \frac{1}{Z(\{h_m\})} e^{-H(\{s_i\}, \{h_m\})}, \quad (1)$$

where Z is the normalization, and H is the “energy”:

$$H = \eta \left[\sum_{i,m=1}^{N, N_f} h_m(t) W_{im} s_i + \epsilon s_i \right]. \quad (2)$$

Here ϵ is the bias toward silence, η controls the variance of individual neuron activity, and W_{im} are the coupling constants that link neurons to fields. The model includes two types of fields (place and latent), explained below.

* Current address: Los Alamos National Laboratory, XCP-8, Los Alamos, NM 87545, USA

FIG. 1. (A) Activity variance of coarse-grained variables at each coarse-graining iteration, fit to $\propto K^\alpha$, $\alpha = 1.37 \pm 0.01$. This is within the error of the experimental observation $\alpha = 1.4 \pm 0.06$ [14], shown in blue. (B) Average free energy, Eq. 5, at each coarse-graining iteration, fit to $\propto K^{\tilde{\beta}}$, $\tilde{\beta} = 0.84 \pm 0.01$, again close to the experimentally found $\tilde{\beta} = 0.88 \pm 0.01$ [14]. (C) Eigenvalue spectrum of cluster covariance for cluster sizes $K = 32, 64, 128$ against the scaled rank, averaged over clusters. We observe a scaling, as in Eq. (6), for about 1.5 decades with $\mu = 0.65 \pm 0.01$, within error of the experimental $\mu = 0.71 \pm 0.06$ [14]. For all panels, the error bars are standard deviations over randomly selected contiguous quarters of the simulation.

In the experiment analyzed in Ref. [14], a mouse ran on a virtual track repeatedly, while neural activity in a population of hippocampal neurons was recorded. A subset of these neurons, called *place cells*, are activated when the mouse is at certain points on the track. To capture this structure, we define place fields distributed along a virtual track of length X . We simulate 200 repetitions of a run along a track of length X with an average forward speed v . As in the experiments, at the end of each run, the mouse is transported instantaneously to the beginning of the track. Thus the mouse position is $x(t) = v(t \bmod T)$, where $T = X/v = 1$ is the time to run a track length. The place fields $h_m^{(\text{place})}(x)$ are modeled as Gaussians with centers $\mu_m \sim \text{unif}(0, X]$ and standard deviations $\sigma_m \sim \Gamma(4, X/40)$ drawn from the Γ -distribution with shape 4 and scale $X/40$. Coupling between a spin and its place field $W_{im}^{(\text{place})}$ is nonzero with probability q , with its value drawn from the standard Γ distribution, $\Gamma(1, 1)$. We include place fields in our model to match the observed data, but we reproduce the scaling results within error bars whether or not place cells are modeled (see *Discussion* and *Online Supplementary Materials*).

The second type of field is a latent field, which we interpret as processes, such as head position or arousal level,

known to modulate neural activity, but not directly controlled or measured by the experiment [25]. We model each latent field $h_m^{(\text{latent})}$ as an Ornstein-Uhlenbeck process with zero mean, unit variance, and the time constant τ . We model the couplings to the latent fields as

$$W_{im}^{(\text{latent})} = \phi \times \begin{cases} \sim \mathcal{N}(0, 1) & \text{if } i \text{ couples to latent fields,} \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

Here $\sim \mathcal{N}(0, 1)$ denotes sampling from the standard normal distribution, and ϕ controls the relative strength of the latent fields compared to the place fields in driving the neural activity. We present results with all latent fields $h_m^{(\text{latent})}$ possessing the same time constant τ (see Tbl. I for parameters), so that the temporal criticality cannot be attributed to the diversity of time scales in the fields driving the neural activity.

While we explored many different parameter choices (see Tbl. II), we present results largely with $N = 1024$ [14], and $N_f = 10$. Consistent with Ref. [14], we choose $p = 50\%$ of neurons to be place cells, each coupled to its own place field (μ_m, σ_m). Each latent field is coupled to every neuron. Thus in our typical simulations, about 512 neurons respond to place and latent stimuli, and about 512 are exclusively latent-stimuli neurons.

Results. — In the following, we simulate random neural activity according to Eq. (1) and then we replicate the real-space and momentum-space coarse-graining schemes of Ref. [14], while tracking the distributions of variables within clusters as we iterate the coarse-graining algorithms. Briefly, in each iteration of the real-space coarse-graining scheme, pairs of highly correlated neurons are combined into clusters. The cluster activity is the sum of the activity of the pair. At each iteration step, the population size is therefore halved. In the momentum-space coarse-graining scheme, neural activity fluctuations are projected onto the eigenvectors of the covariance matrix of the population activity, selecting the K eigenvectors

TABLE I. Simulation parameters for Figures 1-3.

Parameter	Description	Value
ϕ	latent field multiplier	$\phi = 1.0$
ϵ	bias towards silence	$\epsilon = -2.67$
η	variance multiplier	$\eta = 6.0$
q	probability of coupling to latent field	$q = 1.0$
N_f	number of latent fields	$N_f = 10$
τ	latent field time constant	$\tau = 0.1$
$h_m^{(\text{place})}$	presence or absence of place fields	all cells couple to latent fields, half couple to place fields

FIG. 2. (A) Average autocorrelation function for cluster sizes $K = 2, 4, \dots, 256$ as a function of time. (B) Same data, but with time rescaled by the appropriate τ_c for each cluster size. (C) Time constants τ_c extracted from each curve in (A) obey $\tau_c \propto K^{\tilde{z}}$, $\tilde{z} = 0.27 \pm 0.01$, for roughly 1 decade. Experimentally found $\tilde{z}$ is shown in blue [14]. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

with the largest eigenvalues, and then projected back to the original system size, N . All results of Ref. [14] can be *quantitatively* reproduced by our model, and we include corresponding experimental results in blue on each figure when appropriate. Several scaling exponents were not included or were only reported for a single recording in Ref. [13], and therefore we refer to Ref. [14].

1. *Scaling of the activity variance.* Real-space coarse-graining of experimental data [14] reported that the variance of the cluster variables scaled with the cluster size K as K^α , $\alpha = 1.40 \pm 0.06$, in one experiment. In our simulations, the coarse-grained activity variance scales as K^α , $\alpha = 1.36 \pm 0.01$, over more than two decades in K (Fig. 1A), within error bars of the experimental value. This indicates that the microscopic variables are not fully independent (which would be $\alpha = 1$), nor are they fully correlated (which would be $\alpha = 2$).

2. *Scaling of the free energy.* The effective free energy is related to the probability of silence in a cluster, and is expected to scale as a power of cluster size, $K^{\tilde{\beta}}$ [14]. Specifically, we marginalize Eq. (1) over all fields:

$$P(\{s_i\}) = \int d\{h_m\} P(\{h_m\}) P(\{s_i\}|\{h_m\}) \quad (4)$$

and compute $\ln P(\{s_i = 0\}) = \ln P(\{s_i = 0\}|\{h_m\}) + \ln \sum_{\{h_m\}} P(\{h_m\})$, where $P(\{s_i = 0\})$ is the probability that all neurons $\{s_i\}$ are silent. This defines

$$F(\{s_i\}) = -\ln P(\{s_i = 0\}|\{h_m\}), \quad (5)$$

where $F(\{s_i\})$ is effective free energy. In Fig. 1B, we observe that the average free energy at each coarse-graining scales, with a scaling exponent of $\tilde{\beta} = 0.84 \pm 0.01$, within error bars of experimental results, 0.88 ± 0.01 [14].

3. *Scaling of the eigenvalue spectra.* We expect the eigenvalues of the covariance matrix of microscopic variables within each cluster to scale as a power law of the scaled eigenvalue rank [14]. Thus there are two scalings: the rank by the cluster size, and the eigenvalue by the scaled rank. Specifically, the R^{th} eigenvalue λ_R of a clus-

ter of size K was shown in [14] to follow

$$\lambda_R \propto \left(\frac{K}{R}\right)^\mu. \quad (6)$$

In Fig. 1C, we plot the average eigenvalue spectrum of the covariance matrix for each coarse-grained variable for cluster sizes $K = 16, 32, 64, 128, 256$. We observe scaling according to Eq. 6 for roughly 1.5 decades, with the scaling exponent $\mu = -0.65 \pm 0.01$, within error bars of the experimental value of $\mu = -0.71 \pm 0.06$.

4. *Scaling of the correlation time.* Another signature of critical systems is that the timescale of cluster autocorrelation τ_c is a power law of length scale (cluster size K) with exponent $\tilde{z}$. In Fig. 2A we plot the average autocorrelation function for $K = 4, 8, \dots, 256$. In Fig. 2B, we show the same data as a function of the rescaled time, t/τ_c , where τ_c is calculated by fitting the correlation function to the exponential form. The collapse shown in Fig. 2B suggests that $C(t/\tau_c)$ is scale invariant. We then observe a power law relation between the time constant τ_c and the cluster size K for roughly 1.5 decades in Fig. 2C, with a scaling exponent $\tilde{z} = 0.27 \pm 0.01$. For the recording reported in Ref. [14], the exponent was somewhat different, $\tilde{z} = 0.16 \pm 0.02$, but the value over three different recordings, $\tilde{z} = 0.22 \pm 0.08 \pm 0.10$ (mean, individual recording rms error, standard deviation across recordings) again matches our result.

5. *Flow to a non-Gaussian fixed point.* We replicated the momentum space coarse-graining analysis of Ref. [14]. For this, we first calculated the covariance matrix Γ_{ij} of the neural activity fluctuations matrix $\Phi_{it} = S_{it} - \langle S_{it} \rangle_t$, where i indexes neurons and t indexes time step. We then calculated the eigenvalues and eigenvectors of Γ_{ij} and constructed a matrix $\tilde{S}_{ij}$ containing the eigenvectors in its columns, ordered by the corresponding eigenvalues, from largest to smallest. Summing over the first k modes, we calculated the coarse-grained variable

$$S_{it}^{(k)} = z_i \sum_{l,j'}^{N,k} \tilde{S}_{ij'} \Phi_{lt} \tilde{S}_{lj'}, \quad (7)$$

FIG. 3. Distribution of coarse-grained variables for $k = N/16, N/32, N/64, N/128$ modes retained under momentum-space coarse-graining, with a Gaussian distribution (gray dashed line) shown for comparison. The distribution of coarse-grained variables approaches a non-Gaussian limit as k decreases. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

where we set z_i such that $\langle [S_{it}^{(k)}]^2 \rangle_t = 1$ [14].

In Fig. 3, we follow the distribution of $S_{it}^{(k)}$ over coarse-graining cut-offs k . As the coarse-grained variables are linear combinations of the original variables, if the correlations between the original variables are weak, the distribution will approach a Gaussian due to the central limit theorem. However, close to criticality, the system may flow to a non-Gaussian fixed point. We show these distribution of coarse-grained variables $S_{it}^{(k)}$ for $k = N/16, N/32, N/64, N/128$ modes retained, observing the flow to a non-Gaussian limit as k decreases: the limit distribution retains a sharp peak at 0 and a heavy positive tail, similar to the experiments [14].

Experimental agreement. To investigate which parameter regimes give rise to scaling in our model, we vary the parameters η , ϕ , and ϵ in Eq. (2), the latent field correlation time τ , the number of latent fields N_f , and the probability that a neuron couples to a latent field q . We vary them one at a time, while keeping other parameters at values in Tbl. I. We also run simulations with only nonplace fields $h_m^{(\text{latent})}$ included, or with only place fields $h_m^{(\text{place})}$. We record parameters whose simulations display eigenvalue spectra collapse for at least 1.5 decades, as in Fig. 1D, and activity variance scaling for over 2 decades, as in Fig. 1A. Parameter regimes leading to scaling behaviors are summarized in Tbl. II, with detailed plots shown in *Online Supplementary Materials* [26]. We also provide scatter plots of pairs of scaling exponents (if scaling is observed) in Fig. 4, compared to the values from three different experiments as reported in Ref. [14], highlighting the experiment we used as a benchmark in the previous figures. Our simulations show that a broad range of parameters lead to scaling exponents in a quantitative agreement with the experiments.

Discussion. — When the number of activity variables

is large, working with their joint probability distributions is infeasible, and one need to coarse-grain to develop interpretable models of the data. We have shown that, under two different coarse-graining schemes, a model of a neural population in which neurons (spins) are randomly coupled to a few slowly varying latent stimuli or fields (certainly fewer than would be needed to overfit the data) replicates power law scaling relationships as well as the flow of activity distributions to a non-Gaussian fixed point, reported for the mouse hippocampus experiments [13, 14]. Other models, such as a randomly connected rate network [27], or a spiking Brunel neural network in the synchronous irregular regime [28], cannot reproduce these results [14]. In the latter case, one can approximate the network by a population of uncoupled neurons driven by a single common time-varying input [20], but we show that the scaling does not appear for fewer than about five latent processes, explaining why these previous models failed to match experiments.

Our parameter sweeps show that emergence of scaling in the model is robust to parameter changes. The existence of scaling is most sensitive to nearly all cells having significant latent field coupling, irrespective of whether they additionally couple to place fields. This is especially clear in Fig. S7, where only simulations with widespread latent field coupling reproduce the autocorrelation time collapse [26]. This allows us to make an interesting biological prediction that even place cells in hippocampus must be driven not solely by the animal's position. This is consistent with the observations that place cells carry information about activity of other cells in the population [29]. Further, since it is difficult to reproduce temporal scaling over many decades using latent fields with a single time constant, we suggest that this may be easier with latent fields with diverse time scales.

More broadly, we have shown that the surprising spatio-temporal scaling results of Ref. [14] can be explained by the presence of multiple unknown, time-varying driving fields (possibly with just a single time constant). Further, these latent fields *necessarily* result in scale free activity. To our knowledge, our mechanism is the first one to explain these results. While here we have focused on neural data, our results show that the signatures of criticality discussed in this *Letter* will emerge from *any* sparsely active multivariate system (whether biological, inanimate, social, or human-made) driven by several latent dynamical processes.

ACKNOWLEDGMENTS

We thank L. Meshulam and W. Bialek for helping us to understand their work, and S. Boettcher and G. Berman for valuable feedback. This work was supported in part by NIH Grants R01NS084844 (AS and IN), R01EB022872, and R01NS099375 (IN), and by NSF Grant BCS-1822677 (IN).

FIG. 4. (A)-(F) Scatter plots of scaling exponents $\alpha, \tilde{\beta}, \mu, \tilde{z}$ generated by simulations with varying parameters. Experimental results are in blue [14], with the result highlighted in previous figures in bold. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

TABLE II. Results from parameter sweeps over η, ϕ , and ϵ in Eq. (2). We vary these parameters one at a time while keeping all others at default values.

Param.	Sweep range	Critical values	Comments
ϕ	[0.8, 1.5]	$\phi \in [0.9, 1.2]$	Weak latent fields: damaged variance scaling (Fig. S53).
ϵ	[-2.91, -1.33]	$\epsilon < -1.92$	$\epsilon > -1.92$: damaged eigenvalue scaling (Fig. S27); $\epsilon > -1.5$: flow to a Gaussian fixed point (Fig. S33)
η	[2.8, 6.6]	$\eta \in [2.8, 6.6]$	Does not impact existence of scaling (Fig. S43-S49)
q	$q = [0.25, 1.0]$	$q \geq 0.5$	$q < 0.5$ is deleterious to variance scaling (Fig. S20)
N_f	[1, 20]	$N_f \geq 5$	≥ 5 latent fields needed for scaling (Fig. S11)
τ	[0.05, 1.2]	$\tau \in [0.05, 1.2]$	No significant impact on scaling (Fig. S35-S41)
$h_m^{(\text{place})}$	presence / absence	—	Place fields only: no scaling behavior (Fig. S3-S9)

- [1] R. Segev, J. Goodhouse, J. Puchalla, and M. J. Berry, *Nature Neuroscience* **7**, 1155 (2014).
- [2] J. Nguyen, F. Shipley, A. Linder, G. Plummer, J. Shaevitz, A. Leifer, and S. Setru, *Proceedings of the National Academy of Sciences* **113** (2015), 10.1073/pnas.1507110112.
- [3] J. L. Gauthier and D. W. Tank, *Neuron* **99**, 179 (2018).
- [4] D. A. Schwarz, M. A. Lebedev, T. L. Hanson, D. F. Dimitrov, G. Lehew, J. Meloy, S. Rajangam, V. A. Subramanian, P. J. Ifft, Z. S. Li, A. Ramakrishnan, A. J. Tate, K. Z. Zhuang, and M. A. L. Nicolelis, *Nature methods* **11**, 670 (2014).
- [5] Q. Lin, J. Manley, M. Helmreich, F. Schlumm, J. M. Li, D. N. Robson, F. Engert, A. Schier, T. Nöbauer, and A. Vaziri, *Cell* **180**, 536 (2020).
- [6] G. X. Y. Zheng, J. M. Terry, P. Belgrader, P. Ryvkin, Z. W. Bent, R. Wilson, S. B. Ziraldo, T. D. Wheeler, G. P. McDermott, J. Zhu, M. T. Gregory, J. Shuga, L. Montesclaros, J. G. Underwood, D. A. Masquelier, S. Y. Nishimura, M. Schnall-Levin, P. W. Wyatt, C. M. Hindson, R. Bharadwaj, A. Wong, K. D. Ness, L. W. Beppu, H. J. Deeg, C. McFarland, K. R. Loeb, W. J. Valente, N. G. Ericson, E. A. Stevens, J. P. Radich, T. S. Mikkelsen, B. J. Hindson, and J. H. Bielas, *Nature communications* **8**, 14049 (2017).
- [7] J. Cao, J. S. Packer, V. Ramani, D. A. Cusanovich, C. Huynh, R. Daza, X. Qiu, C. Lee, S. N. Furlan, F. J. Steemers, A. Adey, R. H. Waterston, C. Trapnell, and J. Shendure, *Science* **357**, 661 (2017).
- [8] T. M. Gierahn, M. H. Wadsworth, T. K. Hughes, B. D. Bryson, A. Butler, R. Satija, S. Fortune, J. C. Love, and A. K. Shalek, *Nature methods* **14**, 395–398 (2017).
- [9] H. G. Martín and N. Goldenfeld, *Environmental Microbiology* **8**, 1145 (2006).
- [10] C. Palmer, E. M. Bik, D. B. DiGiulio, D. A. Relman, and P. O. Brown, *PLOS Biology* **5**, 1 (2007).
- [11] N. M. Vega and J. Gore, *PLOS Biology* **15**, 1 (2017).
- [12] N. Goldenfeld, *Lectures on Phase Transitions and the*

Renormalization Group (2018) pp. 1–394.

- [13] L. Meshulam, J. L. Gauthier, C. D. Brody, D. W. Tank, and W. Bialek, *Phys. Rev. Lett.* **123**, 178103 (2019).
- [14] L. Meshulam, J. L. Gauthier, C. D. Brody, D. W. Tank, and W. Bialek, “Coarse-graining and hints of scaling in a population of 1000+ neurons,” arXiv:1812.11904. (2018).
- [15] L. P. Kadanoff, *Physics Physique Fizika* **2**, 263 (1966).
- [16] K. G. Wilson, *Rev. Mod. Phys.* **55**, 583 (1983).
- [17] T. Mora, A. M. Walczak, W. Bialek, and C. G. Callan, *Proceedings of the National Academy of Sciences* **107**, 5405 (2010).
- [18] J. E. S. Socolar and S. A. Kauffman, *Phys. Rev. Lett.* **90**, 068702 (2003).
- [19] M. Nykter, N. D. Price, M. Aldana, S. A. Ramsey, S. A. Kauffman, L. E. Hood, O. Yli-Harja, and I. Shmulevich, *Proceedings of the National Academy of Sciences* **105**, 1897 (2008), <https://www.pnas.org/content/105/6/1897.full.pdf>.
- [20] J. Touboul and A. Destexhe, *Physical Review E* **95**, 1 (2017), arXiv:1503.08033.
- [21] J. P. Barton, M. Kardar, and A. K. Chakraborty, *Proceedings of the National Academy of Sciences of the United States of America* **112**, 1965 (2015).
- [22] D. R. Chialvo, *Nature Physics* **6**, 744 (2010).
- [23] L. Aitchison, N. Corradi, and P. E. Latham, *PLOS Computational Biology* **12**, 1 (2016).
- [24] D. J. Schwab, I. Nemenman, and P. Mehta, *Phys. Rev. Lett.* **113**, 068102 (2014).
- [25] M. J. McGinley, M. Vinck, J. Reimer, R. Batista-Brito, E. Zagha, C. R. Cadwell, A. S. Tolias, J. A. Cardin, and D. A. McCormick, *Neuron* **87**, 1143 (2015).
- [26] See Supplemental Material at [URL will be inserted by publisher] for details concerning pairwise correlations, place cell activity, and parameters sweeps in our simulations.
- [27] C. V. Vreeswijk and H. Sompolinsky, *Science* **274**, 1724 (1996).
- [28] N. Brunel, *Journal of Computational Neuroscience* **8**, 183 (2000).
- [29] L. Meshulam, J. L. Gauthier, C. D. Brody, D. W. Tank, and W. Bialek, *Neuron* **96**, 1178 (2017).

Appendix A: Pairwise correlations and place cell activity

References [13, 14] reported two additional observations: the first and second moments of the cell activity were recorded and the effect of coarse-graining on place cell activity was tracked. We did not address these observations in the Main Text, but we report similar results here.

FIG. S1. (A) Distribution of pairwise correlation coefficients. (B) Rate of firing vs. rank of neuron. For (A) and (B), error bars are standard deviations over randomly selected contiguous quarters of the simulation.

In Fig. S1A, we plot the probability distribution of the pairwise correlation coefficients of our simulated neurons. This is qualitatively similar to the experimental results [14]: the distribution has a sharp peak in density just to the right of 0 (small positive correlations), a short left tail, and a long right tail, ending at correlation coefficients greater than 0.6. In Fig. S1B we plot neuron firing rate vs. its rank. Again, this is similar to the experiments [14], including a maximum firing rate of less than 0.2, and a slight elbow in the otherwise near-straight rate vs. rank curve.

Further, place cells are visible in our simulations across the coarse-graining, as in Ref. [14]. Indeed, Fig. S2 shows the activity of coarse-grained variables vs. position on the track, which shows the characteristic localized place cell bump over many coarse-graining iterations.

Appendix B: Parameter sweeps

Most plots in the Main Text use the default parameters listed in Tbl. I. In order to further investigate the behavior of our model, we perturb each parameter around its default value, while holding all others fixed. Below we include detailed results of each of these parameter sweeps. A summary is tabulated in Tbl. II.

a. Varying how cells couple to stimuli and latent fields

In the model in the Main Text, all cells coupled to latent fields, and half of the cells coupled to place fields. Here we show the effects of changing this. We complete the following 4 simulations, analyzed as in the Main Text:

1. $N/2$ cells couple only to latent fields, $N/2$ cells couple to both latent fields and place fields (Main Text).
2. $N/2$ cells couple only to place fields, $N/2$ cells couple only to latent fields.
3. N cells couple only to latent fields.
4. N cells couple only to place fields.

We refer to these simulations as “all”, “place + latent”, “latent only”, and “place only”, respectively. Note that for the “place only” simulation, we increased ϵ from default value $\epsilon = -2.67$ to $\epsilon = -1.33$ to compensate for the omission of latent fields and the resulting decrease in the activity. In Figs. S3-S9 we show that including place fields in our simulations together with latent fields does not significantly alter free energy scaling (Fig. S5), correlation time scaling (Figs. S6-S8), or approach to a non-Gaussian fixed point (Fig. S9). However, including place fields in simulations with latent fields creates slight deviation from power law scaling in variance at large cluster size, Fig. S4) and is damaging to the eigenvalue collapse, Fig. S3. In contrast, omitting latent fields from simulations has a disastrous effect on scaling.

By examining Figs. S3-S9, we conclude that the presence of scaling behavior does not depend on the presence of place fields, but does depend on whether *all* (or, at least, nearly all) cells also couple to latent stimuli. In fact, the presence of place fields is deleterious to scaling and does not yield scaling behavior without the inclusion of latent fields. Thus existence of scaling in experimental data suggests that most cells (including place cells) in the mouse hippocampus, in fact, are also coupled to latent fields.

b. Varying the number of latent fields N_f

We will now consider the effects of varying the number of latent fields N_f in our simulation. We perform

FIG. S2. Average activity at spatial location x for each neuron or coarse-grained variable. Simulation parameters are tabulated in Tbl. I. Coarse-graining steps 0,1,2,3 and 4 are displayed in (A), (B), (C), (D), and (E), respectively. Observe that place cell activity remains strong even over many coarse-graining iterations.

simulations with the default parameters sweeping over values of $N_f = 1, \dots, 20$. In Fig. S10, we note that our simulations include a regime that quantitatively matches experimental results [14].

We find that for $N_f < 5$, eigenvalue scaling, Fig. S11, and variance scaling, Fig. S12, are damaged. However, free energy scaling, Fig. S13, and correlation time scaling, Fig. S14-S16, are not significantly affected by variation in N_f . Figure S11 through Fig. S17 suggest that 5 or more latent fields are required to observe scaling.

There are hints of an upper limit of N_f for a simulation to display critical behavior. As N_f increases, distributions of coarse-grained activity become increasingly short-tailed (Fig. S17). In addition, Fig. S10A shows variance scaling exponent α approaching 1.2, and the autocorrelation starts having large negative lobes, Fig. S15. It is thus possible that, for some $N_f > 20$, the system will stop exhibiting nontrivial scaling, but additional analysis is needed to confirm this.

d. Varying the penalty term ϵ

We perform simulations sweeping over values of the penalty term ϵ , which controls the sparseness of activity. In Fig. S26, we note that our simulations include a regime which quantitatively matches experimental results [14].

Several scaling results are sensitive to ϵ , with damaged scaling for large ϵ , which corresponds to higher overall levels of activity. We find that varying ϵ significantly damages eigenvalue scaling for $\epsilon > -1.92$ (Fig. S27). We observe that coarse-grained distributions of activity from simulations with $\epsilon > -1.5$ approach but do not reach a Gaussian fixed point (Fig. S33). However, free energy scaling (Fig. S29), variance scaling (Fig. S28), and correlation time scaling (Fig. S30-S32) are not significantly affected by variation in ϵ . We conclude that highly active simulations do not display a clear eigenvalue spectra collapse or approach a non-Gaussian fixed point upon coarse-graining.

e. Varying the latent field time constant τ

c. Varying the probability of coupling to a latent field q

We vary the probability q of coupling to a latent field in our simulation. We perform simulations with all other parameters set to the default values while sweeping over $q = 0.25, \dots, 1.0$. Our simulations include a regime (Fig. S18) which quantitatively matches experimental results [14].

We find that varying q causes slight deviations in variance scaling for $q < 0.5$ (Fig. S20). Eigenvalue scaling (Fig. S19), free energy scaling (Fig. S21), approach to a non-Gaussian fixed point (Fig. S25), and correlation time scaling (Fig. S22-S24) are not significantly affected by variation in q from 0.5 to 1.0. We conclude that varying the probability of coupling to a latent field does not have a significant impact on scaling for $q \geq 0.5$, but is deleterious to scaling for $q < 0.5$.

We will now consider the effects of varying τ in our simulation while fixing the other parameters to default values. As in the Main Text, all latent fields have the same value of τ . We vary τ from 0.05 to 1.2, where the time for one track length to be run in simulations is 1. In Fig. S34, we note that our simulations include a regime which quantitatively matches experimental results [14].

We find that varying τ changes the exponent $\tilde{z}$, with larger τ corresponding to larger $\tilde{z}$ and smaller τ corresponding to small $\tilde{z}$ (Fig. S34). However, free energy scaling (Fig. S37), variance scaling (Fig. S36), eigenvalue scaling (Fig. S35), and approach to a non-Gaussian fixed point (Fig. S41) are not significantly affected by variation in τ . Figure S38 through Fig. S41 suggest that dynamic scaling is robust to an increase in τ , but that better quantitative agreement with the experimental $\tilde{z}$ is achieved with smaller τ .

f. Varying the multiplier η

We perform simulations with the default parameters sweeping over values of η , which is an overall multiplier for the “energy” (Eq. 2). In Fig. S42, we note that our simulations include a regime which quantitatively matches experimental results [14].

Free energy scaling (Fig. S45), variance scaling (Fig. S44), approach to a non-Gaussian fixed point (Fig. S49), and dynamic scaling are not significantly affected by variation in η . The quality of scaling of eigenvalues (Fig. S43) is high across all values of η , although the scaling exponent μ decreases with η . Thus, adjusting η has little effect on the quality of scaling.

FIG. S3. Average eigenvalue spectrum of cluster covariance for cluster sizes $K = 32, 64, 128$. Cell types present in the simulation are labeled. No significant difference in quality of scaling is observed for panels (A)-(C), while simulations with only place cells have a lower quality eigenvalue collapse. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

g. Varying the latent fields multiplier ϕ

Finally, we perform simulations sweeping over values of ϕ , which multiplies the latent field term in the energy (Eq. 2). We vary ϕ from 0.8 to 1.5, with all other parameters fixed to default values. In Fig. S51, we note that our simulations include a regime which quantitatively matches experimental results [14].

In Fig. S50 we show that the presence of place cells remains stable over coarse-graining over the full range $\phi \in [0.8, 1.5]$, but as ϕ increases, the relative strength of place cells compared to background activity is decreased.

We find that varying ϕ does not significantly affect the quality of scaling for eigenvalues (Fig. S52, free energy (Fig. S54), or correlation time (Fig. S57), and it does not affect the approach to a non-Gaussian fixed point (Fig. S58). However, setting $\phi < 1.0$ creates slight deviation from power law scaling of activity variance at large cluster size (Fig. S53). In summary, Fig. S50 through Fig. S58 show that weak latent fields are deleterious to scaling behavior.

FIG. S4. Activity variance over coarse-grained variables at each coarse-graining iteration. Cell types present in the simulation are labeled. No significant difference in quality of scaling is observed for panels (A)-(C), while simulations with only the place cells have a lower quality variance scaling. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S6. Average autocorrelation function for cluster sizes $K = 2, 4, \dots, 256$ as a function of time, cluster size indicated by legend. Cell types present in the simulation are labeled. Panel (D) with place cells only substantially differs from the rest. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S5. Average free energy at each coarse-graining iteration. Cell types present in the simulation are labeled. No significant differences in quality of scaling between simulations is observed. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S7. Average autocorrelation function for cluster sizes $K = 2, 4, \dots, 256$, where time is rescaled by the appropriate τ_c for that coarse-graining iteration. Panel (D) with place cells only substantially differs from the rest. The quality of collapse and agreement with experiments are worse in panels (B) and (D), where some cells have no latent field coupling. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S9. Distribution of coarse-grained variables for $k = N/16, N/32, N/64, N/128$ modes retained. Cell types present in the simulation are labeled. No significant differences in approach to a non-Gaussian fixed point are observed between simulation types except for panel (D), where latent fields are present. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S8. Time constants τ_c extracted from each curve in Fig. S6, and fits to $\tau_c \propto K^{\tilde{z}}$. Cell types present in the simulation are labeled. Panel (D) with place cells only shows no scaling. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S10. The critical exponent, $\alpha, \tilde{\beta}, \tilde{\alpha}, \mu$ vs. number of latent fields N_f . Results from three different experiment [14] are shown in gray, pink, and blue bands. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S12. Activity variance over coarse-grained variables at each coarse-graining iteration for different N_f . While changing N_f changes the value of α , the scaling persists for all explored N_f . Error bars (too small to be seen) are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S11. Average eigenvalue spectrum of cluster covariance for cluster sizes $K = 32, 64, 128$ for different N_f . Too few latent fields are insufficient to reproduce the scaling. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S13. Average free energy at each coarse-graining iteration for different N_f . While changing N_f changes the value of $\tilde{\beta}$, the scaling persists for all explored N_f . Error bars (too small to be seen) are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S14. Average autocorrelation function for cluster sizes $K = 2, 4, \dots, 256$ for different N_f . Error bars (too small to be seen) are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S16. Time constants τ_c extracted from each curve in Fig. S15, fitted to $\tau_c \propto K^{\tilde{z}}$, for different N_f . While changing N_f changes the value of $\tilde{z}$, the scaling persists for all explored N_f . Default simulation parameters with labeled values of N_f .

FIG. S15. Average autocorrelation function for cluster sizes $K = 2, 4, \dots, 256$, where time is rescaled by the appropriate τ_c for that coarse-graining iteration, for different N_f . Too many scaling fields produce negative lobes in the correlation functions, which are not observed experimentally. However, it is not clear how significant this is. Error bars (too small to be seen) are standard deviations over randomly selected contiguous quarters of the simulation. Default simulation parameters with labeled values of N_f .

FIG. S17. Distribution of coarse-grained variables for $k = N/16, N/32, N/64, N/128$ modes retained, for different N_f . Panel (A), with just a single latent field, shows substantial deviations from the rest. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S18. Each critical exponent, $\alpha, \tilde{\beta}, \tilde{z}, \mu$ vs probability of coupling to a latent field q . Results from [14] marked and shaded in gray, pink, and blue. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S20. Activity variance over coarse-grained variables at each coarse-graining iteration for different q . For $q = 0.25$, deviations from a scaling relationship appear at large cluster sizes. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S19. Average eigenvalue spectrum of cluster covariance for cluster sizes $K = 32, 64, 128$ for different q . Varying q produces no noticeable differences in scaling quality. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S21. Average free energy at each coarse-graining iteration for different q . While changing q changes the value of $\tilde{\beta}$, the scaling persists for all explored q . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S22. Average autocorrelation function for cluster sizes $K = 2, 4, \dots, 256$ for different q . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S24. Time constants τ_c extracted from each curve in Fig. S22, and observe behavior obeying $\tau_c \propto K^z$ for roughly 1 decade. Varying q yields no significant difference in quality of temporal scaling. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S23. Average autocorrelation function for cluster sizes $K = 2, 4, \dots, 256$ where time is rescaled by the appropriate τ_c for that coarse-graining iteration, for different q . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S25. Distribution of coarse-grained variables for $k = N/16, N/32, N/64, N/128$ modes retained for different q . Convergence to a non-Gaussian fixed point is unaffected by varying q . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S26. Each critical exponent, α , $\tilde{\beta}$, $\tilde{z}$, μ vs penalty term ϵ . Results from [14] marked and shaded in gray, pink, and blue. Error bars are standard deviations over randomly selected contiguous quarters of the simulation

FIG. S28. Activity variance over coarse-grained variables at each coarse-graining iteration for different ϵ . Varying ϵ yields no significant difference in scaling quality. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S27. Average eigenvalue spectrum of cluster covariance for cluster sizes $K = 32, 64, 128$ for different ϵ . For $\epsilon = -1.33$ and $\epsilon = -1.67$, the scaling relationship degrades. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S29. Average free energy at each coarse-graining iteration for different ϵ . Varying ϵ yields no significant change in quality of free energy scaling. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S30. Average autocorrelation function for cluster sizes $K = 2, 4, \dots, 256$ for different ϵ . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S32. Time constants τ_c extracted from each curve in Fig. S30. Observe behavior obeying $\tau_c \propto K^z$ for roughly 1 decade. Varying ϵ results in no significant changes in temporal scaling. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S31. Average autocorrelation function for cluster sizes $K = 2, 4, \dots, 256$ where time is rescaled by the appropriate τ_c for that coarse-graining iteration, for different ϵ . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S33. Distribution of coarse-grained variables for $k = N/16, N/32, N/64, N/128$ modes retained for different ϵ . Coarse-grained distributions of activity approach a non-Gaussian fixed point for $\epsilon > -1.5$. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S34. Each critical exponent, $\alpha, \tilde{\beta}, \tilde{z}, \mu$ vs latent field time constant τ . Results from [14] marked and shaded in gray, pink, and blue. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S36. Activity variance over coarse-grained variables at each coarse-graining iteration for different values of τ . Varying τ does not have a significant effect on activity variance scaling. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S35. Average eigenvalue spectrum of cluster covariance for cluster sizes $K = 32, 64, 128$ for different values of τ . Varying τ does not have a significant effect on eigenvalue scaling. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S37. Average free energy at each coarse-graining iteration for different τ . Varying τ does not appear to affect free energy scaling. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S38. Average autocorrelation function for cluster sizes $K = 2, 4, \dots, 256$ as a function of time, for different τ . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S40. Time constants τ_c extracted from each curve in Fig. S38, and observe behavior obeying $\tau_c \propto K^{\tilde{z}}$ for roughly 1 decade. Varying τ does not have a significant effect on scaling quality, but results in different scaling exponents $\tilde{z}$, with larger τ resulting in larger $\tilde{z}$. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S39. Average autocorrelation function for cluster sizes $K = 2, 4, \dots, 256$, where time is rescaled by the appropriate τ_c for that coarse-graining iteration, for different τ . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S41. Distribution of coarse-grained variables for $k = N/16, N/32, N/64, N/128$ modes retained for different τ . Varying τ has no significant impact on approach to a non-Gaussian fixed point. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S42. Each critical exponent, α , $\tilde{\beta}$, $\tilde{\alpha}$, μ vs multiplier η . Results from [14] marked and shaded in gray, pink, and blue. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S44. Activity variance over coarse-grained variables at each coarse-graining iteration for different η . Varying η does not impact activity variance scaling. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S43. Average eigenvalue spectrum of cluster covariance for cluster sizes $K = 32, 64, 128$ for different η . Varying η does not appear to impact quality of eigenvalue scaling. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S45. Average free energy at each coarse-graining iteration for different η . Varying η does not affect quality of free energy scaling. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S46. Average autocorrelation function for cluster sizes $K = 2, 4, \dots, 256$ as a function of time, for different η . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S48. Time constants τ_c extracted from each curve in Fig. S46, and observe behavior obeying $\tau_c \propto K^{\tilde{z}}$ for roughly 1 decade. Varying η does not significantly affect the quality of temporal scaling or the value of exponent $\tilde{z}$. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S47. Average autocorrelation function for cluster sizes $K = 2, 4, \dots, 256$ where time is rescaled by the appropriate τ_c for that coarse-graining iteration, for different η . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S49. Distribution of coarse-grained variables for $k = N/16, N/32, N/64, N/128$ retained for different η . Approach to a non-Gaussian fixed point is not affected by varying η . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S50. Average activity at spatial location x for each neuron at coarse-graining step 4. Simulation parameters are those tabulated in Tbl. I with $\phi = 0.8$ (A), $\phi = 1.0$ (B), $\phi = 1.2$ (C), $\phi = 1.4$ (D), and $\phi = 1.5$ (E). Increasing the latent field multiplier ϕ decreases the relative strength of place cells compared to cells only coupled to latent fields.

FIG. S51. Each critical exponent, α , $\tilde{\beta}$, $\tilde{z}$, μ vs latent field multiplier ϕ . Results from [14] marked and shaded in gray, pink, and blue. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S52. Average eigenvalue spectrum of cluster covariance for cluster sizes $K = 32, 64, 128$ for different ϕ . Note that quality of eigenvalue collapse and scaling and value of the exponent μ is unaffected by varying ϕ . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S53. Activity variance over coarse-grained variables at each coarse-graining iteration for different ϕ . Note that for $\phi < 1.0$, variance scaling is damaged. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S54. Average free energy at each coarse-graining iteration for different ϕ . Quality of free energy scaling is unaffected by varying ϕ . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S55. Average autocorrelation function for cluster sizes $K = 2, 4, \dots, 256$ as a function of time, for different values of ϕ . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S56. Average autocorrelation function for cluster sizes $K = 2, 4, \dots, 256$ where time is rescaled by the appropriate τ_c for that coarse-graining iteration, for different values of ϕ . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S57. Time constants τ_c extracted from each curve in in FIG S55, and observe behavior obeying $\tau_c \propto K^{\tilde{z}}$ for roughly 1 decade. Error bars are standard deviations over randomly selected contiguous quarters of the simulation.

FIG. S58. Distribution of coarse-grained variables for $k = N/16, N/32, N/64, N/128$ modes retained for different ϕ . Note convergence to non-Gaussian fixed point regardless of value of ϕ . Error bars are standard deviations over randomly selected contiguous quarters of the simulation.