Signatures of criticality in a maximum entropy model of the *C. elegans* brain during free behaviour

Miguel Aguilera[1,2], Carlos Alquézar[1,2,3] and Eduardo J. Izquierdo[4]

[1]Aragón Institute of Engineering Research, Zaragoza, Spain
[2]Dept. of Computer Science, Univ. of Zaragoza, Zaragoza, Spain
[3]Dept. of Biochemistry and Molecular and Cellular Biology, Univ. of Zaragoza, Zaragoza, Spain
[4]Cognitive Science Program, Indiana University, Bloomington, United States.
sci@maguilera.net

Abstract

A popular hypothesis suggests that the nervous system of different organisms, from neural tissue to whole brains, may operate at or near a critical point. During the last decade, maximum entropy techniques have allowed to go beyond merely finding statistical signatures of criticality, to models directly inferred from data recorded in neural cultures, providing stronger evidence of criticality in neural activity. Nevertheless, these modeling techniques are restricted to neural cultures and have not been extended to neural tissue in living organisms. In this paper, we extend this line of research by analyzing signatures of criticality in a pairwise maximum entropy model inferred from neural recordings of *C. elegans* during freely-moving locomotion. From the analysis of the inferred models we find some signatures of criticality, as a divergence of the heat capacity of the system. Other indicators, such as Zipf's distributions, were not found. However, inspecting a similar analysis based in a 2D lattice Ising model we suggest that this could be due to the restricted number of samples in our data set. The availability of larger recordings of the *C. elegans* neural system during free locomotion could provide more conclusive results.

Introduction

Several studies have reported evidence that neural systems operate at or near critical points, poised at transitions in their parameter space (Mora and Bialek, 2011; Beggs, 2008). Criticality refers to a distinctive set of properties found at the boundary separating regimes with different dynamics. A classical example is the transition between an ordered and a disordered phase. Some of these properties include power-law divergences of some quantities (such as the heat capacity or correlation-length of the system) described by critical exponents and fractal behaviour (Salinas, 2001).

Evidence of criticality in neural networks was first found in model simulations (Chialvo and Bak, 1999; Bak and Chialvo, 2001; Bertschinger and Natschlger, 2004; Lin and Chen, 2005). Since, signatures of criticality have been found in recordings of cortical neuron cultures, in the form of power law distributions of avalanche sizes (Beggs and Plenz, 2003; Beggs, 2008; Timme et al., 2016), as well as in recorded brain activity from humans and other animals (Chialvo, 2014; Scott et al., 2014), although the presence of signatures of criticality in the brain is still controversial and not fully resolved (Dehghani et al., 2012; Beggs and Timme, 2012). As well, criticality has been found in dynamical models using human or animal connectomes, including the *C. Elegans* connectome (Moretti and Muoz, 2013). Furthermore, during the last decade, a series of experiments inferring maximum entropy models reproducing the statistics of neural cultures have allowed a deeper characterization of criticality in neural networks, inferring statistical mechanics models displaying criticality in the thermodynamic limit (Schneidman et al., 2006; Yu et al., 2011; Tkacik et al., 2015). All these results suggest that biological neural networks are poised near critical points as a mechanism, and it is suggested that this may allow them to optimize information processing (Beggs, 2008). Highly correlated networks would generate a limited range of possible states, while weakly correlated dynamics would prevent information flow in large systems. In between, systems at criticality present unique features combining integration and segregation of information.

The pervasive finding of signatures of criticality in nervous systems supports the provocative hypothesis predicting that brains and organisms will self-organize to operate near critical points. Nevertheless, unlike the case of controlled cultures of a limited number of neurons, criticality has not been studied much in statistical mechanics models reproducing real data from an operative neural system. Thus, the hypothesis of criticality in the brain is largely unresolved since the models capturing these statistical signatures of criticality are substantially low-dimensional compared to the neural structure of interest and/or restricted to an unrealistic simple range of behaviours (Chialvo, 2014). Moreover, experimental evidence of criticality in the nervous system of humans and other animals is generally restricted to measurements in static situations where the subject is not freely moving nor interacting in normal conditions with its environment.

In this paper, we further explore the presence of criticality in the nervous system of freely moving organisms by applying inference of maximum entropy models to neural

Carole Knibbe et al, eds., Proceedings of the ECAL 2017, Lyon, France, 4-8 September 2017,
(Cambridge, MA: The MIT Press, ©2017 Massachusetts Institute of Technology).

activation data from the nematode *Caenorhabditis elegans*. We analyze an Ising model whose connectivity is inferred from data published by Nguyen et al. (2016b): whole-brain recordings of *C. elegans* during free locomotion. Up to date, publicly available neural recordings of the freely moving animal are still limited to a few minutes in duration and a few thousand samples. This makes impractical finding statistical indicators of criticality such as scaling laws. Instead, maximum entropy models deal with data scarcity by inferring models that capture measurable statistical quantities of the data (such as means and pairwise correlations of neural activation) without making additional assumptions. These models have been successfully applied to map the activity of networks of neurons (Schneidman et al., 2006; Tkacik et al., 2015) and other biological systems, such as antibody protein sequences (Mora et al., 2010) and flocks of birds (Bialek et al., 2012).

In this paper, we find evidence of criticality in the *C. elegans* nervous system. Due to the limited size of the dataset, the results are not conclusive. Our work suggests the usefulness of an analysis on datasets of larger duration and finer resolution.

Model

In order to inspect how *C. elegans* neurons are operating, we derive statistical mechanics models of the system from the published dataset. The objective of these models is to infer macroscopic properties from microscopic descriptions of the system that cannot be accessed directly from the data.

We use publicly available recordings of a freely-moving *C. elegans*, genetically modified to express the calcium indicator GCaMP6s (see Materials and Methods). The activation $a_i(t)$ of $N = 56$ neurons from the head of the worm (out of 302 in the whole neuron system) was captured for a period of approximately 5 minutes with a sampling rate of 6Hz, giving a total of $T = 1872$ samples. Since we are interested in macroscopic statistical properties of the system, we do not focus in which specific neurons are analyzed, but rather in the collective behaviour of groups of neurons.

Since our data is limited, we simplified the data to binary variables to maximize the accuracy of the statistical descriptions. Prior to that, we found the presence of low frequency transients combined with faster frequency oscillations in the data (see Figure 1 top). As Nguyen et al. (2016b,a) suggest, the movement of the worm could introduce some artifacts that the system might be interpreting as calcium transients. For example, the compression or expansion of the neurons could lead to changes to the density of the fluorescent protein. On the other hand, part of this low-frequency activity may correspond to neural activity at the frequency of the worm's movement. Thus, for the rest of the paper we use two versions of the data. The first version is the raw activity measured from GCaMP6s fluorescence. The second is a filtered version of the activity, where the original signal is taken through a high-pass Butterworth filter with a cutoff frequency of $0.06Hz$ (Figure 1).

Even with binary variables, the number of states of such a system is dauntingly large (2^N states). Thus, obtaining a good estimate of $P(s)$ from data for measuring power laws or other indicators of critical activity is unrealistic. The principle of maximum entropy is a strategy that addresses this problem by assuming a model that is as random as possible, but that agrees with some average observables of the data. As we will see, maximum entropy models can be used to naturally map the statistics of a given data onto known statistical mechanics models, which will ease the study of their macroscopic properties.

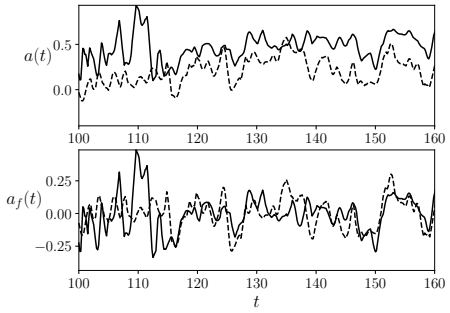

Figure 1: Example of the activation $a_i(t)$ of two arbitrary neurons unfiltered (top) and applying a high pass Butterworth filter with a cutoff frequency of $0.06Hz$ (bottom).

From discretized signals (see Materials and Methods), we extract the means and correlations present in the data and infer the pairwise maximum entropy models (i.e. Ising models for binary variables) that reproduce them. This is the least-structured model that is consistent with the mean activation rate and correlations of the nodes in the network. The maximum entropy distribution consistent with a known average energy is the Boltzmann distribution $P(s) = Z^{-1}e^{-\beta E(s)}$, where s is a state of the network, Z is the partition function and $\beta = 1/(Tk_B)$, being k_B Boltzmann's constant and T the temperature. Without loss of generality we can set the temperature $\beta = 1$. The energy of an Ising model with pairwise interactions is defined as $E(s) = -\sum_i h_i s_i - \sum_{i<j} J_{ij} s_i s_j$, where 'magnetic fields' h_i represent influences in the activation of individual nodes and 'exchange couplings' J_{ij} stand for the tendencies correlating the activity between nodes. The resulting distribution of the maximum entropy model is:

$$P(s) = \frac{1}{Z} exp\left[\beta \sum_i h_i s_i + \sum_{i<j} J_{ij} s_i s_j\right] \qquad (1)$$

where the h_i and J_{ij} are adjusted to reproduce the measured mean and correlation values between nodes in the network.

Carole Knibbe et al, eds., Proceedings of the ECAL 2017, Lyon, France, 4-8 September 2017,
(Cambridge, MA: The MIT Press, ©2017 Massachusetts Institute of Technology).

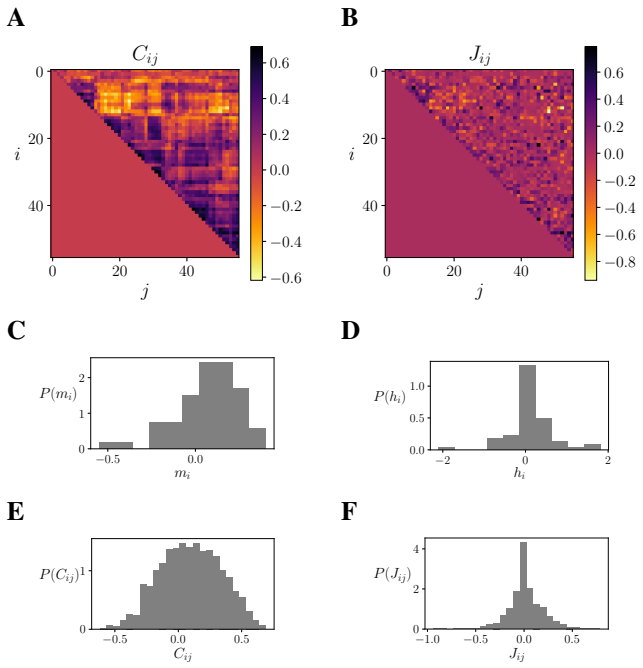

Figure 2: Ising model for the raw version of the neural activations. (**A**) Correlations $C_{ij} = \langle s_i s_j \rangle - \langle s_i \rangle \langle s_j \rangle$ present in the data. (**B**) Coupling constants, J_{ij}. (**C**) Mean activation values. (**D**) Neuron fields h_i. (**E**) Distribution of correlation coefficients, C_{ij}. (**F**) Distribution of coupling values, J_{ij}. We can observe how, although strong correlations are present in the system, connectivity is relatively sparser.

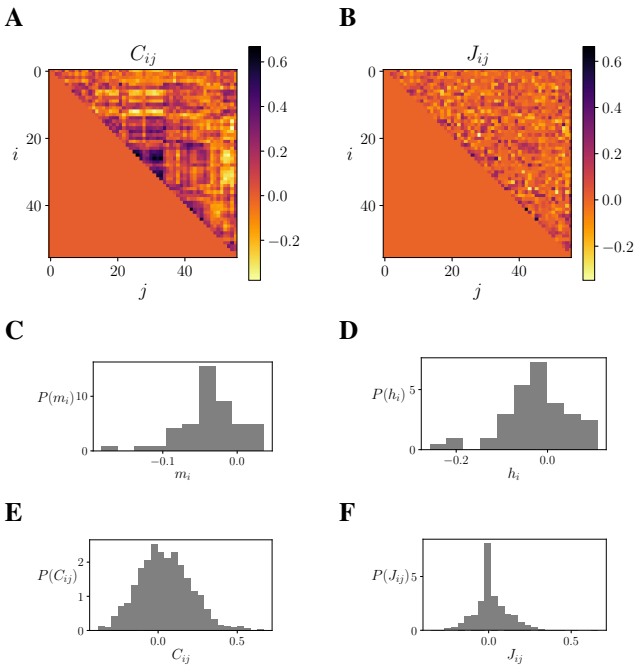

Figure 3: Ising model for the filtered version of the neural activations. (**A**) Correlations $C_{ij} = \langle s_i s_j \rangle - \langle s_i \rangle \langle s_j \rangle$. (**B**) Coupling constants, J_{ij}. (**C**) Mean activation values. (**D**) Neuron fields h_i. (**E**) Distribution of correlation coefficients, C_{ij}. (**F**) Distribution of coupling values, J_{ij}. In comparison with the unfiltered model, stronger correlation values are removed, as well as couplings with larger values and most bias.

From the neural activation data of the worm, we inferred an Ising model $P(s)$ that solves the corresponding inverse Ising problem. We do this using an adaptive cluster expansion (Barton et al., 2016) and Monte Carlo sampling to fit the parameters h_i and J_{ij} that reproduce the means and correlations found in the series of states s (see Materials and Methods).

In Figures 2 and 3 we can observe the inferred models. The models reproduce the correlations found in the neural activation within the margin of error given by the number of samples (see Figure 4). The description of the relative error can be found in the Materials and Methods section.

Results

Once we have extracted a model of the worm's neural activation, $P(s)$, we explore the thermodynamic (macroscopic) properties associated with it. The Ising model allows us to find evidence of the critical behaviour of the model by exploring divergences in quantities as the heat capacity of the system when it approaches the thermodynamic limit. A divergence of the heat capacity implies that the system is maximally sensitive to internal parametric changes. By introducing a fictitious temperature value, changing the inverse temperature parameter β (previously assumed to be equal to 1), we can explore the parameter space of the system computing its heat capacity of the system as:

$$C(T) = \frac{\partial \langle E(s) \rangle}{\partial T} = \beta^2 \langle E(s)^2 \rangle - \langle E(s) \rangle^2 \qquad (2)$$

In order to estimate whether the heat capacity diverges at the thermodynamic limit, we compute networks of different sizes $N = 4, 8, 16, 32$. For each size, we generate 10 different networks, select random combinations of nodes, and compute the heat capacity of the system at different values of β using a Metropolis Monte Carlo sampler. In Figure 5 we observe how the normalized heat capacity of the system changes as the system grows. For the filtered activity, it is not clear that there is a divergence of the heat capacity, since the curves for $N = 32$ and $N = 56$ are similar (Figure 5A). On the other hand, the heat capacity of the system increases with size for filtered neural activity, approaching the operating temperature at $\beta = 1$. This suggests that the system is poised near a critical point.

Another classical signature of criticality is found when the probability distribution of $P(s)$ follows Zipf's law. Due to

Carole Knibbe et al, eds., Proceedings of the ECAL 2017, Lyon, France, 4-8 September 2017,
(Cambridge, MA: The MIT Press, ©2017 Massachusetts Institute of Technology).

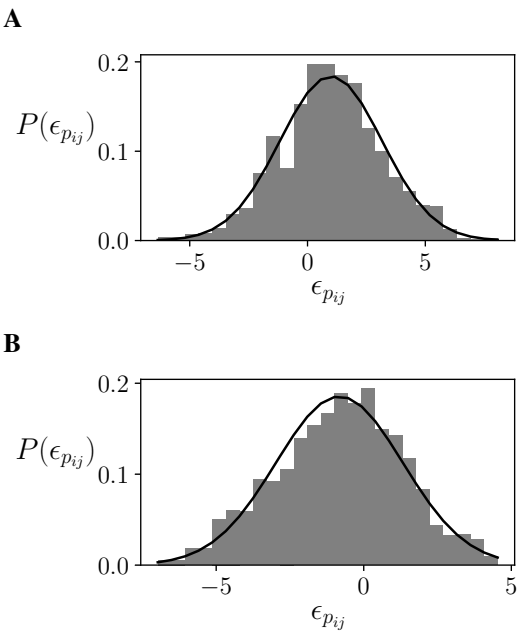

Figure 4: Characterization of the relative error in the pairwise correlations between the real data and that generated by the obtained Ising models. The relative error is defined as the ratio between the deviations of the predicted observables from the data and the expected statistical fluctuations due to finite sampling (see Materials and Methods). (**A**) The error from the model generated using the raw activity. (**B**) The error from the model generated using the filtered activity.

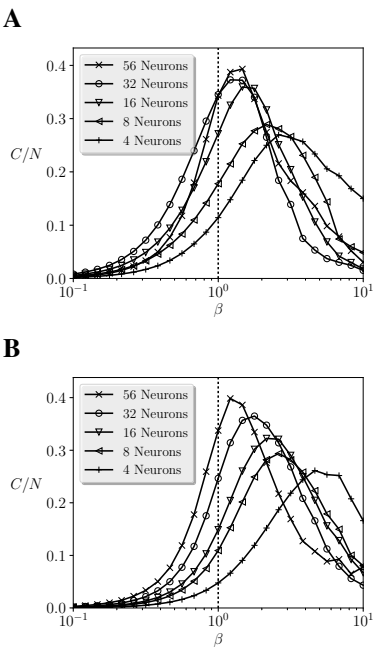

Figure 5: Normalized Heat Capacity C/N of the Ising model of the network of all neurons and subnetworks with a random subsets of neurons for the (**A**) raw and (**B**) filtered models. We observe that C/N shows a peak which approaches 1 as we consider larger networks. Similarly, the maximum value of the peak increases with larger networks. This suggests a divergence of the heat capacity in the thermodynamic limit.

the limited number of samples in the data, it is impossible to find power law scaling directly there since most samples appear only one or twice in the series. On the other hand, the inferred Ising models describe a huge probability distribution function which is not directly computable (with 2^{56} states). Instead, we first searched for power law scaling in the probability distribution in the unfiltered and filtered case by creating a probability distribution from limited sets of Metropolis Monte Carlo samples, which did not yield results compatible with the Zipf law (Figures 6, panels A and B).

Finally, we tested the presence of power law structures in the probability distribution of the samples present in the data, using equation 1. The results suggests that the model resembles Zipf's law for a limited range of samples (around two orders of magnitude) falling rapidly afterwards for both the unfiltered an filtered case (Figure 6C-D). This could be an indicator that the system is not at a critical point. Alternatively, the system could be near a critical point. However, the limited samples available might preclude the model to learn part of the long-range correlations necessary to display a consistent Zipf's law distribution. In the following section, we test this intuition using a 2D lattice Ising model.

Comparison with a 2D lattice Ising model

How does the number of samples influence in the detection of signatures of criticality? Is it possible that maximum entropy models inferred from a limited number of samples display signatures of criticality in the divergence of the heat capacity but are not able to capture the Zipf's distribution of the probability density function?

In order to test this idea we generate 2000 samples of a 7x7 lattice Ising model at the critical temperature. As in the case of *C. elegans* neurons, we extracted a model of the 49 units, as well as 10 models with N random neurons, with $N = 4, 8, 16, 32$. Repeating the same analysis we performed in the models inferred from the neurons of the *C. elegans*, we find similar results than in the filtered case.

First, we observe a similar divergence of the heat capacity, which becomes larger and with a peak closer to the operating temperature $\beta = 1$ as the number of units of the model increases (Figure 7).

Furthermore, if we analyze the probability distribution of the model, we do not find clear signs of power law distributions (Figure 8). This may suggest that, working on a limited number of samples, the divergence of the heat capacity is a

Carole Knibbe et al, eds., Proceedings of the ECAL 2017, Lyon, France, 4-8 September 2017,
(Cambridge, MA: The MIT Press, ©2017 Massachusetts Institute of Technology).

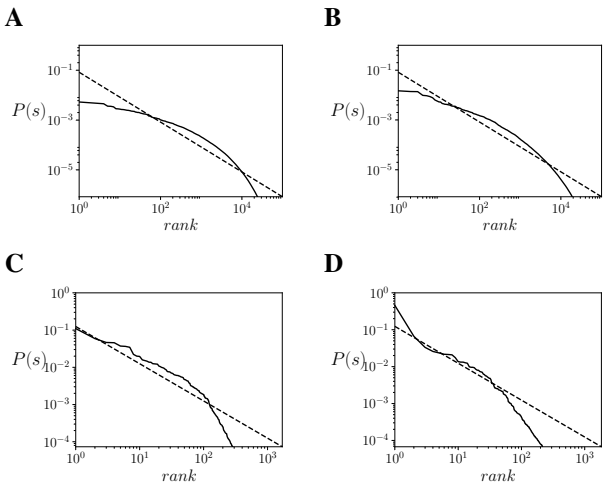

Figure 6: Probability distribution $P(s)$ (solid line)generated with 100000 Monte Carlo samples, in comparison with Zipf's distribution (dashed line) for the models of the (**A**) unfiltered and (**B**) filtered neural activity. Probability distribution $P(s)$ for the states s found in the recordings of the worm for the models of the (**C**) unfiltered and (**D**) filtered neural activity.

better indicator of criticality. Contrarily, long-range correlations necessary for generating power laws in the probability distribution might not be easily captured with limited samples. Nevertheless, this hypothesis should be further tested, both in theoretical models as well as in larger recordings of freely-moving *C. elegans* neural activity when those become technologically possible.

Discussion

In this paper we have inferred a pairwise maximum entropy model from neural activation data extracted from a freely moving *C. elegans*. As far as we are aware, this is the first such attempt.

We analyze the presence of criticality in the model under two conditions. The first extracts a model of binarized samples from the raw neural activity, and the second adds a high-pass filtering steps in order to remove low-frequency artifacts generated by worm motion (though this filtering may be removing activity from the neurons at that frequency).

From both sets of data, we extracted Ising models replicating the statistical structure of the nervous system of the worm. The models obtained displayed a divergence of the heat capacity of the system mapping filtered data, suggesting that the nervous system of the *C. elegans* is near a point of criticality. Results are not clear for the case without filtering, suggesting that the combination of low-frequency transients from the worm's body and discretization of the data may be distorting its statistical content (as well, comparing Figures

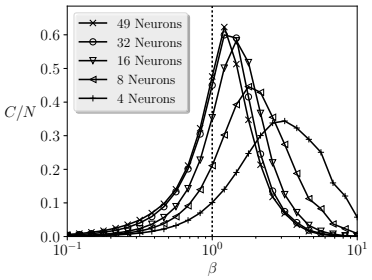

Figure 7: Normalized Heat Capacity C/N of the 7x7 lattice Ising model. As expected, we observe that the maximum value of the peak increases with larger networks, suggesting a divergence of the heat capacity in the thermodynamic limit.

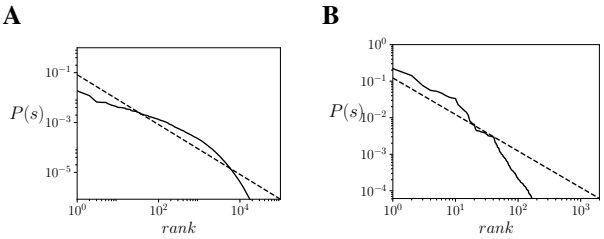

Figure 8: Probability distribution $P(s)$ (solid line) generated with 100000 Monte Carlo samples (**A**) and for the states s found in the recordings of 2000 samples of the Ising model (**B**), in comparison with Zipf's distribution (dashed line).

2 and 3 we can observe how correlations strongly increase throughout the system when filtering is not applied). Nevertheless, some of the slow fluctuations may be the product of criticality in the nervous system. Future analysis may explore more advanced filtering techniques to discriminate between external noise and intrinsic fluctuations in low-frequency contents of the signal. In addition, the generality of the results shall be tested to compare the behaviour of different individual worms. In any case, the presented results show some evidence that the brain of a freely-moving C elegans worm might be posed at criticality. Generally, the brain has been shown to be at criticality in rest state, but some evidence points that long-range correlations associated with criticality might disappear with movement (Botcharova et al., 2015). New available data in freely-moving animals offers the opportunity to study in detail the presence of criticality during movement, or even while interacting with the animal's environment.

Analyzing the probability distribution of the inferred model does not yield the presence of clear power laws in neither the filtered or unfiltered case. Thus not being conclusive for the assessment of criticality in the nervous system of *C. elegans*. However, we hypothesize that the lim-

Carole Knibbe et al, eds., Proceedings of the ECAL 2017, Lyon, France, 4-8 September 2017,
(Cambridge, MA: The MIT Press, ©2017 Massachusetts Institute of Technology).

ited length of the data (only presenting around 1900 samples) might prevent the model to learn the long-range correlations present in power law distribution. Analyses of a restricted number of samples generated from an Ising model in a 2D lattice at the critical temperature seem to support this intuition. Replicating this analysis with larger datasets or better methods for analyze the scaling of the data under limited samples might help to corroborate this idea (Levina and Priesemann, 2017). In addition, in this preliminary study only one worm was analyzed, so future analysis including recordings from other worms, further experiments and the availability of more extensive data of neural activity in freely moving *C. elegans* should validate or dismiss the analysis presented here. As well, a model of the *C. elegans* nervous system inferred directly from the data may offer other interesting possibilities, for example allowing a comparison between the worm's connectome and the inferred pairwise structure of the Ising model identifying coupling and inhibitory tendencies, or to better understand the relations between topology and dynamics in the neural activity of *C. elegans*.

Materials and Methods

Data acquisition. We used the data corresponding with worm 1 available from Nguyen et al. (2016b). In this study were analyzed four *Caenorhabditis elegans* modified genetically to express the calcium indicator GCaMP6s in order to register the calcium transients related to neural dynamics. The four worms were freely behaving on their environments and also the position and body-shapes were registered. Furthermore, the data from an additional transgenic worm capable to express calcium insensitive GFP protein was provided as a control worm. In our case, worm 1 was analyzed since it presented larger recordings of neural activity.

Discretization of neural activity. A calcium imaging technique was used to record changes in intracellular calcium levels of each neuron. In order to make the minimum assumptions analysing this non-spiking activity, we discretized the activation of each neuron into a set of two possible binary states according to a threshold θ. The threshold was determined by computing the mean entropy of clusters of 1, 2 and 3 neurons for different values of θ. For the three cluster sizes, the entropy peaked around the same value for each worm. We chose θ to maximize the entropy of the discretized data, in order to capture as much information as possible from the original data.

Inference of maximum entropy models. We used adaptive cluster expansion (Barton et al., 2016) for fitting the parameters h_i and J_{ij} that reproduce the means and correlations found in the discretized neural data mentioned above. We used $l2$ normalization with a value of $1/T$, being T the number of samples available. After applying adaptive cluster expansion, we adjusted the resulting values through Boltzmann Learning (Ackley et al., 1985) by computing the means and correlations of the model using 800000 samples.

Convergence of the model is defined in terms of the relative error found in the one-point and pairwise correlations between the model and the actual data, defined as:

$$\begin{aligned} \varepsilon_{p_i} &= \frac{|{p_i}^m - {p_i}^d|}{\delta_{p_i}} \\ \varepsilon_{p_{ij}} &= \frac{|{p_{ij}}^m - {p_{ij}}^d}{\delta_{p_{ij}}}| \end{aligned} \tag{3}$$

where $\delta_{p_i} = \sqrt{\frac{p_i(1-p_i)}{T}}$ and $\delta_{p_{ij}} = \sqrt{\frac{p_{ij}(1-p_{ij})}{T}}$ are the expected statistical fluctuations due to finite number of samples T. In Figure 4 it is shown the relative error $\varepsilon_{p_{ij}}$ of correlations in the model.

We define the normalized maximum error as:

$$\varepsilon_{max} = max_{\{i,j\}} \left[\frac{1}{\sqrt{2log\,(M)}} \left(\varepsilon_{p_i}, \varepsilon_{p_{ij}}\right) \right] \tag{4}$$

where M is the total number of one- and two-point correlations, i.e. the number of p_i and p_{ij} values.

We first apply adaptive cluster expansion until we reach a cluster size of 12 and afterwards Monte Carlo Boltzmann learning until $\varepsilon_{max} < 1$.

Acknowledgements

This work was supported by the Spanish National Programme for Fostering Excellence in Scientific and Technical Research project PSI2014-62092-EXP, Spanish MINECO project TIN2016-80347-R and NSF grant No. PHY-9723972.

References

Ackley, D. H., Hinton, G. E., and Sejnowski, T. J. (1985). A Learning Algorithm for Boltzmann Machines*. *Cognitive Science*, 9(1):147–169.

Bak, P. and Chialvo, D. R. (2001). Adaptive learning by extremal dynamics and negative feedback. *Physical Review E*, 63(3):031912.

Barton, J. P., De Leonardis, E., Coucke, A., and Cocco, S. (2016). ACE: adaptive cluster expansion for maximum entropy graphical model inference. *Bioinformatics*, 32(20):3089–3097.

Beggs, J. M. (2008). The criticality hypothesis: how local cortical networks might optimize information processing. *Philosophical Transactions. Series A, Mathematical, Physical, and Engineering Sciences*, 366(1864):329–343.

Beggs, J. M. and Plenz, D. (2003). Neuronal Avalanches in Neocortical Circuits. *Journal of Neuroscience*, 23(35):11167–11177.

Carole Knibbe et al, eds., Proceedings of the ECAL 2017, Lyon, France, 4-8 September 2017,
(Cambridge, MA: The MIT Press, ©2017 Massachusetts Institute of Technology).

Beggs, J. M. and Timme, N. (2012). Being Critical of Criticality in the Brain. *Frontiers in Physiology*, 3.

Bertschinger, N. and Natschlger, T. (2004). Real-time computation at the edge of chaos in recurrent neural networks. *Neural Computation*, 16(7):1413–1436.

Bialek, W., Cavagna, A., Giardina, I., Mora, T., Silvestri, E., Viale, M., and Walczak, A. M. (2012). Statistical mechanics for natural flocks of birds. *Proceedings of the National Academy of Sciences*, 109(13):4786–4791.

Botcharova, M., Berthouze, L., Brookes, M. J., Barnes, G. R., and Farmer, S. F. (2015). Resting state MEG oscillations show long-range temporal correlations of phase synchrony that break down during finger movement. *Frontiers in Physiology*, 6.

Chialvo, D. R. (2014). Critical Brain Dynamics at Large Scale. In Plenz, S. I. D. and Niebur, E., editors, *Criticality in Neural Systems*, pages 43–66. Wiley-VCH Verlag GmbH & Co.

Chialvo, D. R. and Bak, P. (1999). Learning from mistakes. *Neuroscience*, 90(4):1137–1148.

Dehghani, N., Hatsopoulos, N. G., Haga, Z. D., Parker, R., Greger, B., Halgren, E., Cash, S. S., and Destexhe, A. (2012). Avalanche Analysis from Multielectrode Ensemble Recordings in Cat, Monkey, and Human Cerebral Cortex during Wakefulness and Sleep. *Frontiers in Physiology*, 3.

Levina, A. and Priesemann, V. (2017). Subsampling scaling. *Nature Communications*, 8:ncomms15140.

Lin, M. and Chen, T. (2005). Self-organized criticality in a simple model of neurons based on small-world networks. *Physical Review E*, 71(1):016133.

Mora, T. and Bialek, W. (2011). Are biological systems poised at criticality? *Journal of Statistical Physics*, 144(2):268–302.

Mora, T., Walczak, A. M., Bialek, W., and Callan, C. G. (2010). Maximum entropy models for antibody diversity. *Proceedings of the National Academy of Sciences*, 107(12):5405–5410.

Moretti, P. and Muoz, M. A. (2013). Griffiths phases and the stretching of criticality in brain networks. *Nature Communications*, 4:2521.

Nguyen, J. P., Linder, A. N., Plummer, G. S., Shaevitz, J. W., and Leifer, A. M. (2016a). Automatically tracking neurons in a moving and deforming brain. *arXiv:1610.04579 [physics, q-bio]*. arXiv: 1610.04579.

Nguyen, J. P., Shipley, F. B., Linder, A. N., Plummer, G. S., Liu, M., Setru, S. U., Shaevitz, J. W., and Leifer, A. M. (2016b). Whole-brain calcium imaging with cellular resolution in freely behaving Caenorhabditis elegans. *Proceedings of the National Academy of Sciences*, 113(8):E1074–E1081.

Salinas, S. R. A. (2001). Scaling Theories and the Renormalization Group. In *Introduction to Statistical Physics*, pages 277–304. Springer New York.

Schneidman, E., Berry, M. J., Segev, R., and Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–12.

Scott, G., Fagerholm, E. D., Mutoh, H., Leech, R., Sharp, D. J., Shew, W. L., and Knpfel, T. (2014). Voltage Imaging of Waking Mouse Cortex Reveals Emergence of Critical Neuronal Dynamics. *Journal of Neuroscience*, 34(50):16611–16620.

Timme, N. M., Marshall, N. J., Bennett, N., Ripp, M., Lautzenhiser, E., and Beggs, J. M. (2016). Criticality Maximizes Complexity in Neural Tissue. *Frontiers in Physiology*, 7.

Tkacik, G., Mora, T., Marre, O., Amodei, D., Palmer, S. E., Berry, M. J., and Bialek, W. (2015). Thermodynamics and signatures of criticality in a network of neurons. *Proceedings of the National Academy of Sciences*, 112(37):11508–11513.

Yu, S., Yang, H., Nakahara, H., Santos, G. S., Nikoli, D., and Plenz, D. (2011). Higher-Order Interactions Characterized in Cortical Activity. *Journal of Neuroscience*, 31(48):17514–17526.

Carole Knibbe et al, eds., Proceedings of the ECAL 2017, Lyon, France, 4-8 September 2017,
(Cambridge, MA: The MIT Press, ©2017 Massachusetts Institute of Technology).