



HAL
open science

Pairwise Ising Model Analysis of Human Cortical Neuron Recordings

Trang-Anh Nghiem, Olivier Marre, Alain Destexhe, Ulisse Ferrari

► **To cite this version:**

Trang-Anh Nghiem, Olivier Marre, Alain Destexhe, Ulisse Ferrari. Pairwise Ising Model Analysis of Human Cortical Neuron Recordings. Frank Nielsen; Frédéric Barbaresco. Geometric Science of Information. Third International Conference, GSI 2017, Paris, France, November 7-9, 2017, Proceedings, 10589, Springer, pp.257-264, 2017, Lecture Notes in Computer Science, 978-3-319-68444-4. 10.1007/978-3-319-68445-1_30 . hal-02064133

HAL Id: hal-02064133

<https://hal.science/hal-02064133v1>

Submitted on 13 Dec 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Pairwise Ising model analysis of human cortical neuron recordings

Trang-Anh Nghiem¹, Olivier Marre² Alain Destexhe¹, and Ulisse Ferrari^{2,3}

¹ Laboratory of Computational Neuroscience, Unité de Neurosciences, Information et Complexité, CNRS, Gif-Sur-Yvette, France.

² Institut de la Vision, INSERM and UMPC, 17 rue Moreau, 75012 Paris, France.

³ ulisse.ferrari@gmail.com

Abstract. During wakefulness and deep sleep brain states, cortical neural networks show a different behavior, with the second characterized by transients of high network activity. To investigate their impact on neuronal behavior, we apply a pairwise Ising model analysis by inferring the maximum entropy model that reproduces single and pairwise moments of the neuron’s spiking activity. In this work we first review the inference algorithm introduced in Ferrari, *Phys. Rev. E* (2016) [1]. We then succeed in applying the algorithm to infer the model from a large ensemble of neurons recorded by multi-electrode array in human temporal cortex. We compare the Ising model performance in capturing the statistical properties of the network activity during wakefulness and deep sleep. For the latter, the pairwise model misses relevant transients of high network activity, suggesting that additional constraints are necessary to accurately model the data.

Keywords: Ising model, maximum entropy principle, natural gradient, human temporal cortex, multielectrode array recording, brain states

Advances in experimental techniques have recently enabled the recording of the activity of tens to hundreds of neurons simultaneously [2] and has spurred the interest in modeling their collective behavior [3,4,5,6,7,8,9]. To this purpose, the pairwise Ising model has been introduced as the maximum entropy (most generic [10]) model able to reproduce the first and second empirical moments of the recorded neurons. Moreover it has already been applied to different brain regions in different animals [3,5,6,9] and shown to work efficiently [11].

The inference problem for a pairwise Ising model is a computationally challenging task [12], that requires devoted algorithms [13,14,15]. Recently, we proposed a *data-driven* algorithm and applied it on rat retinal recordings [1]. In the present work we first review the algorithm structure and then describe our successful application to a recording in the human temporal cortex [4].

We use the inferred Ising model to test if a model that reproduces empirical pairwise covariances without assuming any other additional information, also predicts empirical higher-order statistics. We apply this strategy separately to brain states of wakefulness (Awake) and Slow-Wave Sleep (SWS). In contrast to

the former, the latter is known to be characterized by transients of high activity that modulate the whole population behavior [16]. Consistently, we found that the Ising model does not account for such global oscillations of the network dynamics. We do not address Rapid-Eye Movement (REM) sleep.

1 The model and the geometry of the parameter space

The pairwise Ising model is a fully connected Boltzmann machine without hidden units. Consequently it belongs to the exponential family and has probability distribution:

$$P_\eta(\mathbf{X}) = \exp(\mathbf{T}(\mathbf{X}) \cdot \boldsymbol{\eta} - \log Z[\boldsymbol{\eta}]) , \quad (1)$$

where $\mathbf{X} \in [0, 1]^N$ is the row vector of the N system's free variables and $Z[\boldsymbol{\eta}]$ is the normalization. $\boldsymbol{\eta} \in \mathcal{R}^D$ is the column vector of model parameters, with $D = N(N+1)/2$ and $\mathbf{T}(\mathbf{X}) \in [0, 1]^D$ is the vector of model sufficient statistics. For the fully-connected pairwise Ising model the latter is composed of the list of free variables \mathbf{X} and their pairwise products:

$$\{T_a(\mathbf{X})\}_{a=1}^D = \{ \{X_i\}_{i=1}^N, \{X_i X_j\}_{i=1, j=i+1}^N \} \in [0, 1]^D . \quad (2)$$

A dataset Ω for the inference problem is composed by a set of τ_Ω *i.i.d.* empirical configurations \mathbf{X} : $\Omega = \{\mathbf{X}(t)\}_{t=1}^{\tau_\Omega}$. We cast the inference problem as a log-likelihood maximization task, which for the model (1) takes the shape:

$$\boldsymbol{\eta}^* \equiv \underset{\boldsymbol{\eta}}{\operatorname{argmax}} \ell[\boldsymbol{\eta}] ; \quad \ell[\boldsymbol{\eta}] \equiv \mathbf{T}_\Omega \cdot \boldsymbol{\eta} - \log Z[\boldsymbol{\eta}] , \quad (3)$$

where $\mathbf{T}_\Omega \equiv \mathbf{E}[\mathbf{T}(\mathbf{X}) \mid \Omega]$ is the empirical mean of the sufficient statistics. As a consequence of the exponential family properties, the log-likelihood gradient may be written as:

$$\nabla \ell[\boldsymbol{\eta}] = \mathbf{T}_\Omega - \mathbf{T}_\eta , \quad (4)$$

where $\mathbf{T}_\eta = \mathbf{E}[\mathbf{T}(\mathbf{X}) \mid \boldsymbol{\eta}]$ is the mean of $\mathbf{T}(\mathbf{X})$ under the model distribution (1) with parameters $\boldsymbol{\eta}$. Maximizing the log-likelihood is then equivalent to imposing $\mathbf{T}_\Omega = \mathbf{T}_\eta$: the inferred model then reproduces the empirical averages.

Parameter space geometry In order to characterize the geometry of the model parameter space, we define the minus log-likelihood Hessian $\mathcal{H}[\boldsymbol{\eta}]$, the model Fisher matrix $J[\boldsymbol{\eta}]$ and the model susceptibility matrix $\chi[\boldsymbol{\eta}]$ as:

$$\chi_{ab}[\boldsymbol{\eta}] \equiv \mathbf{E}[T_a T_b \mid \boldsymbol{\eta}] - \mathbf{E}[T_a \mid \boldsymbol{\eta}] \mathbf{E}[T_b \mid \boldsymbol{\eta}] \quad (5)$$

$$J_{ab}[\boldsymbol{\eta}] \equiv \mathbf{E}[\nabla_a \log P_\eta(\mathbf{X}) \nabla_b \log P_\eta(\mathbf{X}) \mid \boldsymbol{\eta}] , \quad (6)$$

$$\mathcal{H}_{ab}[\boldsymbol{\eta}] \equiv -\nabla_a \nabla_b \ell[\boldsymbol{\eta}] , \quad (7)$$

As a property inherited from the exponential family, for the Ising model (1):

$$\chi_{ab}[\boldsymbol{\eta}] = J_{ab}[\boldsymbol{\eta}] = \mathcal{H}_{ab}[\boldsymbol{\eta}] . \quad (8)$$

This last property is the keystone of the present algorithm.

Moreover, the fact that the log-likelihood Hessian can be expressed as a covariance matrix ensures its non-negativity. Some zero Eigenvalues can be present, but they can easily be addressed by L_2 -regularization [14,1]. The inference problem is indeed convex and consequently the solution of (3) exists and is unique.

2 Inference algorithm

The inference task (3) is an hard problem because the partition function $Z[\boldsymbol{\eta}]$ cannot be computed analytically. Ref. [1] suggests applying an approximated natural gradient method to numerically address the problem. After an initialization of the parameters to some initial value $\boldsymbol{\eta}_0$, the natural gradient [17,18] iteratively updates their values with:

$$\eta_{n+1} = \eta_n - \alpha J^{-1}[\boldsymbol{\eta}_n] \cdot \nabla \ell[\eta_n] . \quad (9)$$

For sufficiently small α , the convexity of the problem and the positiveness of the Fisher matrix ensure the convergence of the dynamics to the solution $\boldsymbol{\eta}^*$.

As computing $J[\boldsymbol{\eta}_n]$ at each n is computationally expensive, we use (8) to approximate the Fisher with an empirical estimate of the susceptibility [1]:

$$J[\boldsymbol{\eta}] = \chi[\boldsymbol{\eta}] \approx \chi[\boldsymbol{\eta}^*] \approx \chi_\Omega \equiv \text{Cov}[\mathbf{T} \mid \Omega] . \quad (10)$$

The first approximation becomes exact upon convergence of the dynamics, $\boldsymbol{\eta}_n \rightarrow \boldsymbol{\eta}^*$. The second assumes that (i) the distribution underlying the data belongs to the family (1), and that (ii) the error in the estimate of χ_Ω , arising from the dataset's finite size, is small.

We compute χ_Ω of Eq. (10) only once, and then we run the inference algorithm that performs the following approximated natural gradient:

$$\eta_{n+1} = \eta_n - \alpha \chi_\Omega^{-1} \cdot \nabla \ell[\eta_n] . \quad (11)$$

Stochastic dynamics.⁴ The dynamics (11) require estimating $\nabla \ell[\boldsymbol{\eta}]$ and thus of \mathbf{T}_η at each iteration. This is accounted by a Metropolis Markov-Chain Monte Carlo (MC), which collects I_η , a sequence of τ_I *i.i.d.* samples of the distribution (1) with parameters $\boldsymbol{\eta}$ and therefore estimates:

$$\mathbf{T}_\eta^{\text{MC}} \equiv \mathbf{E}[\mathbf{T}(\mathbf{X}) \mid I_\eta] . \quad (12)$$

This estimate itself is a random variable with mean and covariance given by:

$$\mathbf{E}[\mathbf{T}_\eta^{\text{MC}} \mid \{I_\eta\}] = \mathbf{T}_\eta ; \quad \text{Cov}[\mathbf{T}_\eta^{\text{MC}} \mid \{I_\eta\}] = \frac{J[\boldsymbol{\eta}]}{\tau_I} , \quad (13)$$

where $\mathbf{E}[\cdot \mid \{I_\eta\}]$ means expectation with respect to the possible realizations I_η of the configuration sequence.

⁴ The results of this section are grounded on the repeated use of central limit theorem. See [1] for more detail.

Data: $\mathbf{T}_{\Omega, \chi_{\Omega}}$

Result: $\boldsymbol{\eta}^*, \mathbf{T}_{\eta^*}^{\text{MC}}$

Initialization: set $\tau_{\Gamma} = \tau_{\Omega}, \alpha = 1$ and $\boldsymbol{\eta}_0$; estimate $\mathbf{T}_{\eta_0}^{\text{MC}}$ and compute ϵ_0 ;

while $\epsilon > 1$ **do**

$\boldsymbol{\eta}_{n+1} \leftarrow \boldsymbol{\eta}_n - \alpha \chi_{\Omega}^{-1} \cdot \nabla l[\boldsymbol{\eta}_n]$;
 estimate $\mathbf{T}_{\eta_{n+1}}^{\text{MC}}$ and compute ϵ_{n+1} ;
if $\epsilon_{n+1} < \epsilon_n$ **then**
 | increase α , keeping $\alpha \leq 1$;
else
 | decrease α and set $\boldsymbol{\eta}_{n+1} = \boldsymbol{\eta}_n$;
end
 $n \leftarrow n + 1$;

end

Fix $\alpha < 1$ and perform several iterations.

Algorithm 1: Algorithm pseudocode for the Ising model inference.

For $\boldsymbol{\eta}$ sufficiently close to $\boldsymbol{\eta}^*$, after enough iterations, this last result allows us to compute the first two moments of $\nabla \ell_{\boldsymbol{\eta}}^{\text{MC}} \equiv \mathbf{T}_{\Omega} - \mathbf{T}_{\boldsymbol{\eta}}^{\text{MC}}$, using a second order expansion of the log-likelihood (3):

$$\mathbf{E}[\nabla \ell_{\boldsymbol{\eta}}^{\text{MC}} \mid \{I_{\boldsymbol{\eta}}\}] = \mathcal{H}[\boldsymbol{\eta}] \cdot (\boldsymbol{\eta} - \boldsymbol{\eta}^*) ; \quad \text{Cov}[\nabla \ell_{\boldsymbol{\eta}}^{\text{MC}} \mid \{I_{\boldsymbol{\eta}}\}] = \frac{J[\boldsymbol{\eta}^*]}{\tau_{\Gamma}} . \quad (14)$$

In this framework, the learning dynamics becomes stochastic and ruled by the master equation:

$$P_{n+1}(\boldsymbol{\eta}') = \int d\boldsymbol{\eta} P_n(\boldsymbol{\eta}) W_{\boldsymbol{\eta} \rightarrow \boldsymbol{\eta}'}[\boldsymbol{\eta}] ; \quad W_{\boldsymbol{\eta} \rightarrow \boldsymbol{\eta}'}[\boldsymbol{\eta}] = \text{Prob}(\nabla \ell_{\boldsymbol{\eta}}^{\text{MC}} = \boldsymbol{\eta}' - \boldsymbol{\eta}) , \quad (15)$$

where $W_{\boldsymbol{\eta} \rightarrow \boldsymbol{\eta}'}[\boldsymbol{\eta}]$ is the probability of transition from $\boldsymbol{\eta}$ to $\boldsymbol{\eta}'$. For sufficiently large τ_{Γ} and thanks to the equalities (8), the central limit theorem ensures that the unique stationary solution of (15) is a Normal Distribution with moments:

$$\mathbf{E}[\boldsymbol{\eta} \mid P_{\infty}(\boldsymbol{\eta})] = \boldsymbol{\eta}^* ; \quad \text{Cov}[\boldsymbol{\eta} \mid P_{\infty}(\boldsymbol{\eta})] = \frac{\alpha}{(2 - \alpha)\tau_{\Gamma}} \chi^{-1}[\boldsymbol{\eta}^*] . \quad (16)$$

Algorithm. Thanks to (8) one may compute the mean and covariance of the model posterior distribution (with flat prior):

$$\mathbf{E}[\boldsymbol{\eta} \mid P^{\text{Post}}(\boldsymbol{\eta})] = \boldsymbol{\eta}^* ; \quad \text{Cov}[\boldsymbol{\eta} \mid P^{\text{Post}}(\boldsymbol{\eta})] = \frac{1}{\tau_{\Omega}} \chi^{-1}[\boldsymbol{\eta}^*] \quad (17)$$

where τ_{Ω} is the size of the training dataset. From (14), if $\boldsymbol{\eta} \sim P^{\text{Post}}$ we have:

$$\mathbf{E}[\nabla \ell_{\boldsymbol{\eta}}^{\text{MC}} \mid \{I_{\boldsymbol{\eta} \sim P^{\text{Post}}}\}] = 0 ; \quad \text{Cov}[\nabla \ell_{\boldsymbol{\eta}}^{\text{MC}} \mid \{I_{\boldsymbol{\eta} \sim P^{\text{Post}}}\}] = \frac{2\chi[\boldsymbol{\eta}^*]}{\tau_{\Gamma}} . \quad (18)$$

Interestingly, by imposing:

$$\frac{1}{\tau_{\Omega}} = \frac{\alpha}{(2 - \alpha)\tau_{\Gamma}} \quad (19)$$

the moments (16) equal (17) [1]. To evaluate the inference error at each iteration we define:

$$\epsilon_n = \|\nabla \ell_{\eta_n}^{\text{MC}}\|_{\chi_\Omega} = \sqrt{\frac{\tau_\Omega}{2D} \nabla \ell_{\eta_n}^{\text{MC}} \cdot \chi_\Omega^{-1} \cdot \nabla \ell_{\eta_n}^{\text{MC}}}. \quad (20)$$

Averaging ϵ over the posterior distribution, see (18), gives $\epsilon = 1$. Consequently, if $\eta_n \neq \eta^*$ implies $\epsilon_n > 1$ with high probability, for $\eta_n \rightarrow \eta^*$ thanks to (19) we expect $\epsilon_n = \sqrt{\tau_\Omega/\tau_\Gamma/(2-\alpha)}$ [1]. As sketched in pseudocode 1, we iteratively update η_n through (11) with $\tau_\Gamma = \tau_\Omega$ and $\alpha < 1$ until $\epsilon_n < 1$ is reached.

3 Analysis of cortical recording

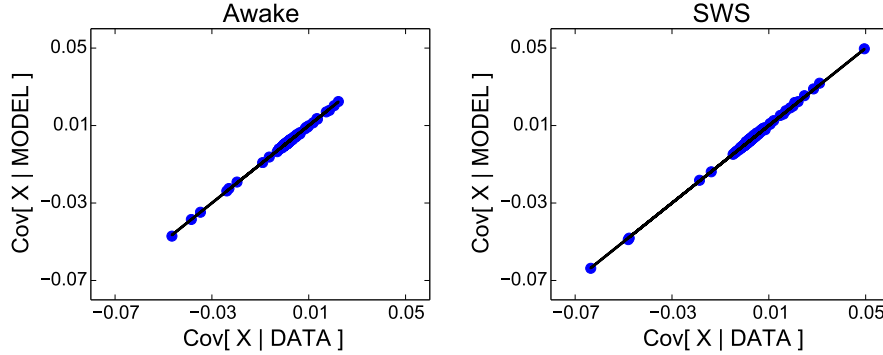


Fig. 1. Empirical pairwise covariances against their model prediction for Awake and SWS. The goodness of the match implies that the inference task was successfully completed. Note the larger values in SWS than Awake

As in [4,7], we analyze ~ 12 hours of intracranial multi-electrode array recording of neurons in the temporal cortex of a single human patient. The dataset is composed of the spike times of $N = 59$ neurons, including $N^I = 16$ inhibitory neurons and $N^E = 43$ excitatory neurons. During the recording session, the subject alternates between different brain states [4]. Here we focused on wakefulness (Awake) and Slow-Wave Sleep (SWS) periods. First, we divided each recording into τ_Ω short $50ms$ -long time bins and encoded the activity of each neuron $i = 1, \dots, N$ in each time bin $t = 1, \dots, \tau_\Omega$ as a binary variable $X_i(t) \in [0, 1]$ depending on whether the cell i was silent ($X_i(t) = 0$) or emitted at least one spike ($X_i(t) = 1$) in the time window t . We thus obtain one training dataset $\Omega = \{\{X_i(t)\}_{i=1}^N\}_{t=1}^{\tau_\Omega}$ per brain state of interest. To apply the Ising model we assume that this binary representation of the spiking activity is representative of the neural dynamics and that subsequent time-bins can be considered as independent. We then run the inference algorithm on the two datasets separately to obtain two sets of Ising model parameters η_{Awake} and η_{SWS} .

Thanks to (4), when the log-likelihood is maximized, the pairwise Ising model reproduces the covariances $\mathbf{E}[X_i X_j | \Omega]$ for all pairs $i \neq j$. To validate the inference method, in Fig. 1 we compare the empirical and model-predicted pairwise covariances and found that the first were always accurately predicted by the second in both Awake and SWS periods.

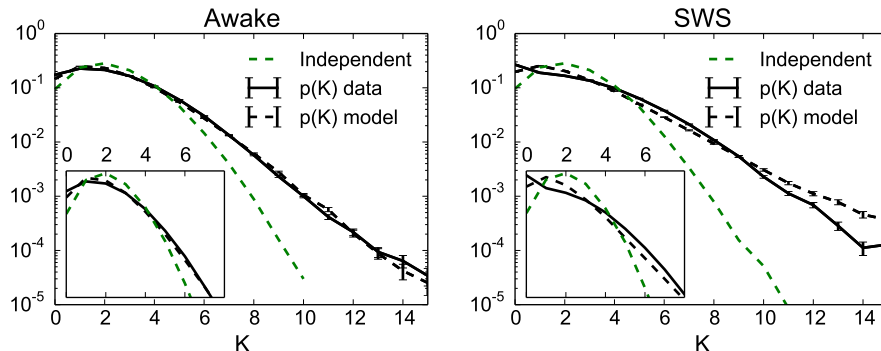


Fig. 2. Empirical and predicted distributions of the whole population activity $K = \sum_i X_i$. For both Awake and SWS periods the pairwise Ising model outperforms the independent model (see text). However, Ising is more efficient at capturing the population statistics during Awake than SWS, especially for medium and large K values. This is consistent with the presence of transients of high activity during SWS.

This shows that the inference method is successful. Now we will test if this model can describe well the statistics of the population activity. In particular, synchronous events involving many neurons may not be well accounted by the pairwise nature of the Ising model interactions. To test this, as introduced in Ref. [6], we quantify the empirical probability of having K neurons active in the same time window: $K = \sum_i X_i$. In Fig. 2 we compare empirical and model prediction for $P(K)$ alongside with the prediction from an *independent* neurons model, the maximum entropy model that as sufficient statistics has only the single variables and not the pairwise: $\{T_a(\mathbf{X})\}_{a=1}^N = \{X_i\}_{i=1}^N$. We observed that the Ising model always outperforms the independent model in predicting $P(K)$.

Fig. 2 shows that the model performance are slightly better for Awake than SWS states. This is confirmed by a Kullback-Leibler divergence estimate:

$$D_{\text{KL}}(P_{\text{Awake}}^{\text{Data}}(K) | P_{\text{Awake}}^{\text{Ising}}(K)) = 0.005; D_{\text{KL}}(P_{\text{SWS}}^{\text{Data}}(K) | P_{\text{SWS}}^{\text{Ising}}(K)) = 0.030.$$

This effect can be ascribed to the presence of high activity transients, known to modulate neurons activity during SWS [16] and responsible for the larger covariances, see Fig. 1 and the heavier tail of $P(K)$, Fig. 2. These transients are known to be related to an unbalance between the contributions of excitatory and inhibitory cells to the total population activity [7]. To investigate the impact of

these transients, in Fig. 3 we compare $P(K)$ for the two populations with the corresponding Ising model predictions. For the Awake state, the two contributions are very similar, probably in consequence of the excitatory/inhibitory balance [7]. Moreover the model is able to reproduce both behaviors. For SWS periods, instead, the two populations are less balanced [7], with the inhibitory (blue line) showing a much heavier tail. Moreover, the model partially fails in reproducing this behavior, notably strongly overestimating large K probabilities.

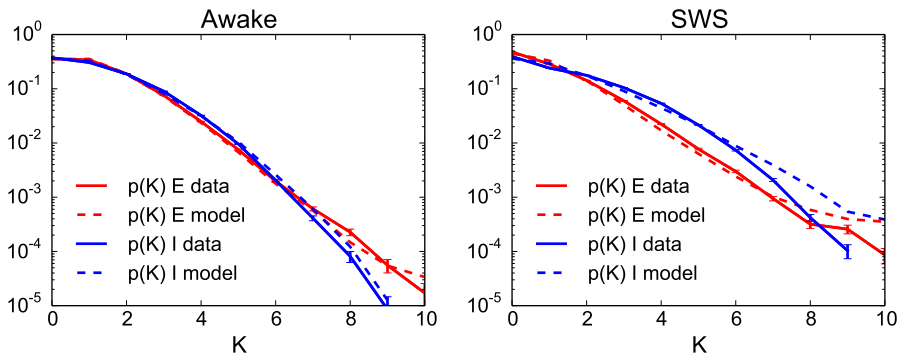


Fig. 3. Empirical and predicted distributions of excitatory (red) and inhibitory (blue) population activity. During SWS, the pairwise Ising model fails at reproducing high activity transients, especially for inhibitory cells.

Conclusions: (i) The pairwise Ising model offers a good description of the neural network activity observed during wakefulness. (ii) By contrast, taking into account pairwise correlations is not sufficient to describe the statistics of the ensemble activity during SWS, where (iii) alternating periods of high and low network activity introduce high order correlations among neurons, especially for inhibitory cells [16]. (iv) This suggests that neural interactions during wakefulness are more local and short-range, whereas (v) these in SWS are partially modulated by internally-generated activity, synchronizing neural activity across long distances [4,16,19].

Acknowledgments We thank B. Telenczuk, G. Tkacik and M. Di Volo for useful discussion. Research funded by European Community (Human Brain Project, H2020-720270), ANR TRAJECTORY, ANR OPTIMA, French State program Investissements dAvenir managed by the Agence Nationale de la Recherche [LIFESENSES: ANR-10-LABX-65] and NIH grant U01NS09050.

References

1. U. Ferrari. Learning maximum entropy models from finite-size data sets: A fast data-driven algorithm allows sampling from the posterior distribution. *Phys. Rev. E*, 94:023301, 2016.
2. I.H. Stevenson and K.P. Kording. How advances in neural recording affect data analysis. *Nature neuroscience*, 14(2):139–142, 2011.
3. E. Schneidman, M. Berry, R. Segev, and W. Bialek. Weak pairwise correlations imply strongly correlated network states in a population. *Nature*, 440:1007, 2006.
4. A. Peyrache, N. Dehghani, E.N. Eskandar, J.R. Madsen, W.S. Anderson, J.A. Donoghue, L.R. Hochberg, E. Halgren, S.S. Cash, and A. Destexhe. Spatiotemporal dynamics of neocortical excitation and inhibition during human sleep. *Proceedings of the National Academy of Sciences*, 109(5):1731–1736, 2012.
5. L. S. Hamilton, J. Sohl-Dickstein, A. G. Huth, V. M. Carels, K. Deisseroth, and S. Bao. Optogenetic Activation of an Inhibitory Network Enhances Feedforward Functional Connectivity in Auditory Cortex. *Neuron*, 80:1066–76, 2013.
6. G. Tkacik, O. Marre, D. Amodè, E. Schneidman, W Bialek, and Berry M.J. Searching for collective behaviour in a network of real neurons. *PloS Comput. Biol.*, 10(1):e1003408, 2014.
7. N. Dehghani, A. Peyrache, B. Telenczuk, M. Le Van Quyen, E. Halgren, S.S. Cash, N.G. Hatsopoulos, and A. Destexhe. Dynamic balance of excitation and inhibition in human and monkey neocortex. *Scientific Reports*, 6(23176), 2016.
8. C. Gardella, O. Marre, and T. Mora. A tractable method for describing complex couplings between neurons and population rate. *eneuro*, 3(4):0160, 2016.
9. G. Tavoni, U. Ferrari, S. Cocco, F.P. Battaglia, and R. Monasson. Functional coupling networks inferred from prefrontal cortex activity show experience-related effective plasticity. *Network Neuroscience*, MIT press, 2017.
10. E.T. Jaynes. On The Rationale of Maximum-Entropy Method. *Proc. IEEE*, 70:939, 1982.
11. U. Ferrari, T. Obuchi, and T. Mora. Random versus maximum entropy models of neural population activity. *Phys. Rev. E*, 95:042321, 2017.
12. D. H. Ackley, G. E. Hinton, and T. J. Sejnowski. A learning algorithm for boltzmann machines. *Cognitive Science*, 9:147–169, 1985.
13. T. Broderick, M. Dudik, G. Tkacik, R.E. Schapire, and W. Bialek. Faster solutions to the inverse pairwise Ising problem. *arXiv:0712.2437*, 2007.
14. S. Cocco and R. Monasson. Adaptive cluster expansion for inferring Boltzmann machines with noisy data. *Phys. Rev. Lett.*, 106:090601, 2011.
15. J. Sohl-Dickstein, P. B. Battaglino, and M. R. DeWeese. New Method for Parameter Estimation in Probabilistic Models: Minimum Probability Flow. *Phys. Rev. Lett.*, 107:220601, 2011.
16. A. Renart, J. De La Rocha, P. Bartho, L. Hollender, N. Parga, A. Reyes, and K.D. Harris. The asynchronous state in cortical circuits. *Science*, 327(5965):587–590, 2010.
17. S. Amari. Natural Gradient Works Efficiently in Learning, *Neural Computation*. *Neural Comput.*, 10:251–276, 1998.
18. S. Amari and S.C. Douglas. Why natural gradient? *Proc. IEEE*, 2:1213–16, 1998.
19. M. Le Van Quyen, L.E. Muller, B. Telenczuk, E. Halgren, S. Cash, N.G. Hatsopoulos, N. Dehghani, and A. Destexhe. High-frequency oscillations in human and monkey neocortex during the wake–sleep cycle. *Proceedings of the National Academy of Sciences*, 113, 33, 9363–68, 2016