



HAL
open science

Parameters estimation for spatio-temporal maximum entropy distributions: application to neural spike trains.

Hassan Nasser, Bruno Cessac

► To cite this version:

Hassan Nasser, Bruno Cessac. Parameters estimation for spatio-temporal maximum entropy distributions: application to neural spike trains.. [Research Report] 2014. ⟨hal-00927080⟩

HAL Id: hal-00927080

<https://inria.hal.science/hal-00927080v1>

Submitted on 10 Jan 2014

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.



HAL Authorization

Parameters estimation for spatio-temporal maximum entropy distributions: application to neural spike trains.

Hassan Nasser, Bruno Cessac *

January 10, 2014

Abstract

We propose a numerical method to learn Maximum Entropy (MaxEnt) distributions with spatio-temporal constraints from experimental spike trains. This is an extension of two papers [10] and [5] who proposed the estimation of parameters where only spatial constraints were taken into account. The extension we propose allows to properly handle memory effects in spike statistics, for large sized neural networks.

1 Introduction

With the evolution of Multi-Electrode Arrays (MEA) acquisition techniques, it is currently possible to simultaneously record the activity of a few hundred of neurons up to a few thousand [12]. Stevenson et al [40] reported that the number of recorded neurons doubles approximately every 8 years. However, beyond the mere recording of an increasing number of neurons, there is a need to extract relevant information from data in order to understand the underlying dynamics of the studied network, how it responds to stimuli and how spike train response encodes these stimuli. In the realm of spike trains analysis this means having efficient spike sorting techniques [27, 18, 26, 34], but also efficient methods to analyze spike statistics. The second aspect requires using canonical statistical models whose parameters have to be tuned ("learned") from data.

The Maximum Entropy method (MaxEnt) offers a way to selecting canonical statistical models from first principles. Having its root in statistical physics, MaxEnt consists of fixing a set of constraints, determined as the empirical average of features measured from the spiking activity. Maximizing the statistical entropy given those constraints provides a unique probability, called a Gibbs distribution, which approaches at best data statistics in the following sense: among all probability distributions which match the constraints this is the one which has the smallest Kullback-Leibler divergence with the data. Equivalently, it satisfies the constraints without adding additional assumption on statistics [21].

*NeuroMathComp team (INRIA), Sophia Antipolis, France.
INRIA, 2004 Route des Lucioles, 06902 Sophia-Antipolis, France.
email: hassan.nasser@inria.fr

Most studies have focused on describing properly the statistics of *spatially* synchronized patterns of neuronal activity without considering time-dependent patterns and memory effects. In this setting pairwise models [39, 33] or extensions with triplets and quadruplets interactions [14], [13], [43] were claimed to correctly fit ≈ 90 to 99% of the information. However, considering now the capacity of these models to correctly reproduce *spatio-temporal* spike patterns, the performances drop-off dramatically, especially in the cortex [41, 28] or in the retina [44].

Taking into account spatio-temporal patterns requires to introduce memory in statistics, described as a Markov process. MaxEnt extends easily to this case (see section 2.2 and references therein for a short description) producing Gibbs distributions in the spatio-temporal domain. Moreover, rigorous mathematical methods are available to fit the parameters of the Gibbs distribution [44]. However, the main drawback of these methods is the huge computer memory they require, preventing their applications to large scale neural networks. Considering a model with memory depth D (namely, the probability of a spike pattern at time t depends on the spike activity in the interval $[t - D, t - 1]$), there are $2^{N(D+1)}$ possible patterns. The method developed in [44] requires to handle a matrix of size $2^{N(D+1)} \times 2^{N(D+1)}$. So, it becomes intractable for $N(D+1) > 20$.

In this paper, we propose an alternative method to fit the parameters of a spatio-temporal Gibbs distribution with larger values of the product $N(D+1)$. We have been able to go up to $N(D+1)$ (~ 120) on a small cluster (64 processors AMD Opteron(tm) 2300 MHz). The method is based on [10] and [5] who proposed the estimation of parameters in spatial Gibbs distributions. The extension in the spatio-temporal domain is not straightforward, as we show, but it carries over to the price of some modifications. Combined with parallel Montecarlo computing developed in [31] this provides a numerical method allowing to handle Markovian spike statistics with spatio-temporal constraints.

The paper is organized as follow. In section 2, we recall the theoretical background for spike train with Gibbs distribution. We discuss both spatial and spatio-temporal case. In the next section, 3, we explain the method to fit the parameters of MaxEnt distributions. We demonstrate that the extension from the spatial to the spatio temporal case is possible and we also discuss sequential and parallel update algorithms. As we mathematically show, the convex criterion used by [10] still applies for spatio-temporal constraints. However, the method used by [5] to avoid recomputing the Gibbs distribution at each parameters change cannot be directly used and has to be adapted using a Linear Response scheme. In the last section, 4, we show benchmarks evaluating the performance of this method and discuss the computational obstacles that we encountered. We made tests with both synthetic and real data. Synthetic data were generated from known probability distributions using a Montecarlo method. Real data corresponds to spike trains obtained from retinal ganglion cells activity (courtesy of M.J. Berry and O. Marre). The method shows a satisfying performance in the case of synthetic data. The performance on real data, although satisfying, is affected by the large number of parameters in the distribution, consequence of the choice to work with canonical models (Ising, pairwise with memory). This effect is not related to our method, as we show, but is in fact a standard problem in statistics.

2 Gibbs distributions in the spatio-temporal domain

2.1 Spike trains and observables

2.1.1 Spike trains

We consider the joint activity of N neurons, characterized by the emission of action potentials ("spikes"). We assume that there is a minimal time scale, δ , set to 1 without loss of generality such that a neuron can at most fire a spike within a time window of size δ . This provides a time discretization labeled with an integer time n . Each neuron activity is then characterized by a binary variable¹ $\omega_k(n) = 1$ if neuron k fires at time n and $\omega_k(n) = 0$ otherwise. The state of the entire network in time bin n is thus described by a vector $\omega(n) \stackrel{\text{def}}{=} [\omega_k(n)]_{k=1}^N$, called a *spiking pattern*. A *spike block* is a consecutive sequence of spike patterns $\omega_{n_1}^{n_2}$, representing the activity of the whole network between two instants n_1 and n_2 .

$$\omega_{n_1}^{n_2} = \{ \omega(n) \}_{\{n_1 \leq n \leq n_2\}}.$$

The *time-range* (or "range") of a block $\omega_{n_1}^{n_2}$ is $n_2 - n_1 + 1$, the number of time steps from n_1 to n_2 . Here is an example of a spike block with $N = 4$ neurons and range $R = 3$:

$$\begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

A *spike train* or *raster* is a spike block ω_0^T from some initial time 0 to some final time T . To alleviate notations we simply write ω for a spike train. We note Ω the set of spike trains.

2.1.2 Observables

An *observable* is a function \mathcal{O} which associates a real number $\mathcal{O}(\omega)$ to a spike train. In the realm of statistical physics common examples of observables are the energy or the number of particles (where ω would correspond to a spin configuration). In the context of neural networks examples are the number of neuron firing at a given time n , $\sum_{k=1}^N \omega_k(n)$, or the function $\omega_{k_1}(n_1)\omega_{k_2}(n_2)$ which is 1 if neuron k_1 fires at time n_1 and neuron k_2 fires at time n_2 and is 0 otherwise.

Typically, an observable does not depend on the full raster, but only on a sub-block of it. The *time-range* (or "range") of an observable is the minimal integer $R > 0$ such that, for any raster ω , $\mathcal{O}(\omega) = \mathcal{O}(\omega_0^{R-1})$. The range of the observable $\sum_{k=1}^N \omega_k(n)$ is 1; the range of $\omega_{k_1}(n_1)\omega_{k_2}(n_2)$ is $n_2 - n_1 + 1$. From now on, we restrict to observables of range R , fixed and finite. We set $D = R - 1$.

An observable is *time-translation invariant* if, for any time $n > 0$ we have $\mathcal{O}(\omega_n^{n+D}) \equiv \mathcal{O}(\omega_0^D)$ whenever $\omega_n^{n+D} = \omega_0^D$. The two examples above are time-translation invariant. The observable $\lambda(n_1)\omega_{k_1}(n_1)\omega_{k_2}(n_2)$, where λ is a real function of time, is not time-translation invariant. Basically, time-translation

¹We use the notation ω to differentiate our binary variables $\in \{0, 1\}$ to the notation σ or S used for "spins" variables $\in \{-1, 1\}$.

invariance means that \mathcal{O} does not depend explicitly on time. We focus on such observables from now on.

2.1.3 Monomials

Prominent examples of time-translation invariant observables with range R are products of the form:

$$m_{p_1, \dots, p_r}(\omega) \stackrel{\text{def}}{=} \prod_{u=1}^r \omega_{k_u}(n_u). \quad (1)$$

where $p_u, u = 1 \dots r$ are pairs of spike-time events (k_u, n_u) , $k_u = 1 \dots N$ being the neuron index, and $n_u = 0 \dots D$ being the time index. Such an observable, called *monomial*, takes therefore values in $\{0, 1\}$ and is 1 if and only if $\omega_{k_u}(n_u) = 1, u = 1 \dots r$ (neuron k_1 fires at time n_1, \dots , neuron k_r fires at time n_r). A monomial is therefore a binary observable that represents the logic-AND operator applied to a prescribed set of neuron spikes events. We allow the extension of the definition (1) to the case where the set of pairs p_1, \dots, p_r is empty and we set $m_\emptyset = 1$. For a number N of neurons and a time range R there are thus 2^{NR} such possible products. Any observable of range R can be represented as a linear combination of products (1). Monomials constitute therefore a canonical basis for observable representation.

2.1.4 Potential

Another prominent example of observable is the function called "energy" or *potential* in the realm of the MaxEnt. Any potential of range R can be written as a linear combination of the 2^{NR} possible monomials (1):

$$\mathcal{H}_\lambda = \sum_{l=1}^{2^{NR}} \lambda_l m_l, \quad (2)$$

where some coefficients λ_l in the expansion may be zero. Therefore, by analogy with spin systems, monomials somewhat constitute spatio-temporal interactions between neurons: the monomial $\prod_{u=1}^r \omega_{k_u}(n_u)$ contributes to the total energy $\mathcal{H}_\lambda(\omega)$ of the raster ω if and only if neuron k_1 fires at time n_1, \dots , neuron k_r fires at time n_r in the raster ω . The number of pairs in a monomial (1) defines the degree of an interaction: degree 1 corresponds to "self-interactions", degree 2 to pairwise, and so on. Typical examples of such potentials are the Ising model [39, 33, 38]:

$$\mathcal{H}_{Ising}(\omega(0)) = \sum_i \lambda_i \omega_i(0) + \sum_{ij} \lambda_{ij} \omega_i(0) \omega_j(0), \quad (3)$$

or the Gannor-Schneidman-Segev model [13], [14]

$$\mathcal{H}_{GSS}(\omega(0)) = \sum_i \lambda_i \omega_i(0) + \sum_{ij} \lambda_{ij} \omega_i(0) \omega_j(0) + \sum_{ijk} \lambda_{ijk} \omega_i(0) \omega_j(0) \omega_k(0). \quad (4)$$

(We restrict the form (4) to triplet although Gannor et al were also considering quadruplets). In these two examples the potential is a function of the spike

pattern at a given time. Here, we choose this time equal to 0, without loss of generality, since we are considering time-translation invariant potentials. More generally, the form (2) affords the consideration of spatio-temporal neurons interactions: this allows us to introduce delays, memory and causality in spike statistics estimation. A simple example is a pairwise model with delays such as:

$$\mathcal{H}_{PR}(\omega_0^D) = \sum_i \lambda_i \omega_i(D) + \sum_{s=0}^D \sum_{ij} \lambda_{ij}^s \omega_i(0) \omega_j(s), \quad (5)$$

where 'PR' stands for 'Pairwise with range R', takes into account the events where neuron i fires s time steps after a neuron j with $s = 0 \dots D$.

2.2 The Maximum Entropy Principle

Assigning equal probabilities (uniform probability distribution) to possible outcomes goes back to Laplace and Bernoulli ([15]) ("principle of insufficient reason"). Maximizing the statistical entropy *without* constraints is equivalent to this principle. In general, however, one has some knowledge about data, typically characterized by empirical average of prescribed observables (e.g. for spike trains, firing rates, probability that a fixed group of neurons fire at the same time, probability that K neurons fire at the same time [42]): this constitutes a set of constraints. The Maximum Entropy Principle (MaxEnt) is a method to obtain, from the observation of a statistical sample, a probability distribution that approaches at best the statistics of the sample, taking into account these constraints without additional assumptions [21]. Maximizing the statistical entropy given those constraints provides a distribution as far as possible from the uniform and as close as possible to the empirical distribution. For instance, considering the empirical mean and variance of the sample of a random variable as constraints results in a Gaussian distribution.

Although some attempts have been made to extend MaxEnt to non stationary data [19, 20, 22, 32] it is mostly applied in the context of stationary statistics: the average of an observable does not depend explicitly on time. We shall work with this hypothesis. In its simplest form, the MaxEnt also assumes that the sample has no memory: the probability of an outcome at time t does not depend on the past. We first discuss the MaxEnt in this context in the next section, before considering the case of processes with memory in the section 2.2.2.

2.2.1 Spatial constraints

In our case, the natural constraints are represented by the empirical probability of occurrence of characteristic spike events in the spike train, or, equivalently, by the average of specific monomials. Classical examples of constraints are the probability that a neuron fires at a given time (firing rate) or the probability that two neurons fire at the same time. For a raster ω of length T we note $\pi_\omega^{(T)}$ the empirical distribution, and $\pi_\omega^{(T)}[\mathcal{O}]$ the empirical average of the observable \mathcal{O} in the raster ω . For example, the empirical firing rate of neuron i is $\pi_\omega^{(T)}[\omega_i] = \frac{1}{T} \sum_{n=0}^{T-1} \omega_i(n)$, the empirical probability that two neurons i, j fire at the same time is $\pi_\omega^{(T)}[\omega_i \omega_j] = \frac{1}{T} \sum_{n=0}^{T-1} \omega_i(n) \omega_j(n)$ and so on. Given a set of

L monomials m_l , their empirical average, $\pi_\omega^{(T)} [m_l]$, measured in the raster ω , constitute a set of constraints shaping the sought probability distribution. We consider here monomials corresponding to events occurring *at the same time*, i.e. $m_l(\omega) \equiv m_l(\omega(0))$ postponing to section 2.2.2 the general case of events occurring at distinct times.

In this context, the MaxEnt problem is stated as follows. Find *the* distribution μ that maximizes the entropy:

$$\mathcal{S} [\mu] = - \sum_{\omega(0)} \mu [\omega(0)] \log \mu [\omega(0)], \quad (6)$$

(where the sum holds on the 2^N possible spike patterns $\omega(0)$), given the constraints:

$$\mu [m_l] = \pi_\omega^{(T)} [m_l], \quad l = 1 \dots L. \quad (7)$$

The average of monomials (noted here $\mu [m_l]$), predicted by the statistical model μ , must be equal to the average $\pi_\omega^{(T)} [m_l]$ measured in the sample. There is, additionally, the probability normalization constraint:

$$\sum_{\omega(0)} \mu [\omega(0)] = 1 \quad (8)$$

This provides a variational problem

$$\mu = \arg \max_{\nu \in \mathcal{M}} \left[\mathcal{S} [\nu] + \lambda_0 \left(\sum_{\omega(0)} \nu [\omega(0)] - 1 \right) + \sum_{l=1}^L \lambda_l \left(\nu [m_l] - \pi_\omega^{(T)} [m_l] \right) \right] \quad (9)$$

where \mathcal{M} is the set of (stationary) probabilities on spike trains.

Stated in this form the MaxEnt is a Lagrange multipliers problem. The sought probability distribution is the classical Gibbs distribution:

$$\mu_\lambda [\omega(0)] = \frac{1}{Z_\lambda} e^{\mathcal{H}_\lambda [\omega(0)]}, \quad (10)$$

where $Z_\lambda = \sum_{\omega(0)} e^{\mathcal{H}_\lambda [\omega(0)]}$ is the *partition function*, whereas $\mathcal{H}_\lambda [\omega(0)] = \sum_{l=1}^L \lambda_l m_l [\omega(0)]$. Note that the time index (here 0) does not play a role since we have assumed μ_λ to be stationary (time-translation invariant).

The value of λ_l s is fixed by the relation:

$$\mu_\lambda m_l = \frac{\partial \log Z_\lambda}{\partial \lambda_l} = \pi_\omega^{(T)} [m_l], \quad l = 1 \dots L. \quad (11)$$

Additionally, note that the matrix $\frac{\partial^2 \log Z_\lambda}{\partial \lambda_l \partial \lambda_{l'}}$ is positive. This ensures the convexity of the problem and the uniqueness of the solution of the variational problem.

Note that we do not expect in general μ_λ to be equal to the (hidden) probability shaping the observed sample. It is only the closest one satisfying the constraints (7) [3]. The notion of closeness is related to the Kullback-Leibler divergence, defined in the next section.

It is easy to check that the Gibbs distribution (10) obeys:

$$\mu_{\lambda} [\omega_{n_1}^{n_2}] = \prod_{n=n_1}^{n_2} \mu_{\lambda} [\omega(n)], \quad (12)$$

for any spike block $\mu [\omega_{n_1}^{n_2}]$. Indeed, the potential of the spike block $\omega_{n_1}^{n_2}$ is $\mathcal{H}_{\lambda} (\omega_{n_1}^{n_2}) = \sum_{n=n_1}^{n_2} \mathcal{H}_{\lambda} (\omega(n))$ whereas the partition function on spike blocks $\omega_{n_1}^{n_2}$ is $Z_{n_2-n_1} = \sum_{\omega_{n_1}^{n_2}} e^{\mathcal{H}_{\lambda}[\omega_{n_1}^{n_2}]} = Z_{\lambda}^{n_2-n_1}$. Equation (12) expresses that spiking pattern occurring at different times are *independent* under the Gibbs distribution (10). This is expected: since the constraints shaping μ_{λ} take only into account spiking events occurring at the same time, we have no information on causality between spikes generation or on memory effects. The Gibbs distributions obtained when constructing constraints only with spatial events leads to statistical models where spike patterns are renewed at each time step, without reference to the past activity.

2.2.2 Spatio-temporal constraints

On the opposite, one expects that spike trains generation involves causal interactions between neurons and memory effects. We would therefore like to construct Gibbs distributions taking into account information on spatio-temporal interactions between neurons and leading to a statistical model not assuming anymore that successive spikes patterns are independent. Although the notion of Gibbs distribution extends to process with infinite memory [11] we shall concentrate here to Gibbs distributions associated with Markov processes with finite memory depth D . That is, the probability to have a spike pattern $\omega(n)$ at time n , given the past history of spikes reads $P [\omega(n) | \omega_{n-1}^{n-D}]$. Note that those transition probabilities are assumed not to depend explicitly on time (stationarity assumption).

Such a family of transition probabilities $P [\omega(n) | \omega_{n-1}^{n-D}]$ define an homogeneous Markov chain. Provided² $P [\omega(n) | \omega_{n-1}^{n-D}] > 0$ for all ω_{n-D}^n , there is a unique probability μ , called the *invariant probability* of the Markov chain such that:

$$\mu [\omega_1^D] = \sum_{\omega_0^{D-1}} P [\omega(D) | \omega_0^{D-1}] \mu [\omega_0^{D-1}]. \quad (13)$$

In a Markov process the probability of a block $\omega_{n_1}^{n_2}$, for $n_2 - n_1 + 1 > D$, is:

$$\mu [\omega_{n_1}^{n_2}] = \prod_{n=n_1+D}^{n_2} P [\omega(n) | \omega_{n-D}^{n-1}] \mu [\omega_{n_1}^{n_1+D-1}]. \quad (14)$$

the Chapman-Kolmogorov relation [17]. Comparing with (13) one sees that, to determine the probability of $\omega_{n_1}^{n_2}$, one has to know the transition probabilities and the probability $\mu [\omega_{n_1}^{n_1+D-1}]$. When attempting to construct a Gibbs distribution obeying (14) from a set of spatio-temporal constraints one has therefore to determine simultaneously the family of transition probabilities and the invariant probability. Remark that setting:

$$\phi(\omega_0^D) = \log P [\omega(D) | \omega_0^{D-1}], \quad (15)$$

²This is a sufficient but a necessary condition. In the remaining of the paper we shall work with this assumption.

we may write (14) in the form:

$$\mu [\omega_{n_1}^{n_2} | \omega_{n_1}^{n_1+D-1}] = e^{\sum_{n=n_1+D}^{n_2} \phi(\omega_n^{n+D})}. \quad (16)$$

The probability of observing the spike pattern $\omega_{n_1}^{n_2}$ given the past $\omega_{n_1}^{n_1+D-1}$ of depth D has an exponential form, similar to (10). Actually, the invariant probability of a Markov chain is a Gibbs distribution in the following sense.

In view of (14), probabilities must be defined whatever even if $n_2 - n_1$ is arbitrary large. In this setting, the right objects are probabilities on infinite rasters [17]. Then, the *entropy rate* (or Kolmogorov-Sinai entropy) of μ is:

$$\mathcal{S}[\mu] = - \limsup_{n \rightarrow \infty} \frac{1}{n+1} \sum_{\omega_0^n} \mu[\omega_0^n] \log \mu[\omega_0^n], \quad (17)$$

where the sum holds over all possible blocks ω_0^n . This reduces to (6) when μ obeys (12).

The MaxEnt takes now the following form. We consider a set of L spatio-temporal spike events (monomials) whose empirical average value $\pi_\omega^{(T)}[m_l]$ has been computed. We only restrict to monomials with a range *at most* equal to $R = D + 1$, for some $D > 0$. This provide us a set of constraints of the form (7). To maximize the entropy rate (17) under the constraints (7) we construct a range- R potential $\mathcal{H}_\lambda = \sum_{l=1}^L \lambda_l m_l$. The generalized form of the MaxEnt states that there is a unique probability measure $\mu_\lambda \in \mathcal{M}$ such that [7]:

$$\mathcal{P}[\lambda] = \sup_{\nu \in \mathcal{M}} (\mathcal{S}[\nu] + \nu[\mathcal{H}_\lambda]) = \mathcal{S}[\mu_\lambda] + \mu_\lambda[\mathcal{H}_\lambda]. \quad (18)$$

This is the extension of the variational principle (9) to Markov chains. It selects, among all possible probability ν , a *unique* probability μ_λ which realizes the supremum. μ_λ is called *the Gibbs distribution with potential \mathcal{H}_λ* .

The quantity $\mathcal{P}[\lambda]$ is called *topological pressure* or *free energy density*. For a potential of the form (2) [36, 24]:

$$\frac{\partial \mathcal{P}[\lambda]}{\partial \lambda_l} = \mu_\lambda[m_l]. \quad (19)$$

This is the analog of (11) which allows to tune the parameters λ_l . Thus, $\mathcal{P}[\lambda]$ plays the role of $\log Z_\lambda$ in (10). Actually, it is equal to $\log Z_\lambda$ when restricting to the memory less case³. $\mathcal{P}[\lambda]$ is strictly convex⁴ which guarantees the uniqueness of μ_λ .

Note that μ_λ *has not* the form (10) for $D > 0$. Indeed a probability distribution e.g. of the form $\mu[\omega_0^{n-1}] = \frac{1}{Z_n} e^{\mathcal{H}_\lambda(\omega_0^{n-1})}$ with:

$$\mathcal{H}_\lambda(\omega_0^{n-1}) \equiv \sum_{r=0}^{n-D-1} \mathcal{H}_\lambda(\omega_r^{r+D}) = \sum_l \lambda_l \sum_{r=0}^{n-D-1} m_l(\omega_r^{r+D}), \quad (20)$$

³In statistical physics the free energy is $-kT \log Z$. The minus sign comes from the minus sign in the Hamiltonian.

⁴Thanks to the assumption $P[\omega(n) | \omega_{n-1}^{n-D}] > 0$.

the potential of the block ω_0^{n-1} , and:

$$Z_n[\boldsymbol{\lambda}] = \sum_{\omega_0^{n-1}} e^{\mathcal{H}_{\boldsymbol{\lambda}}(\omega_0^{n-1})}, \quad (21)$$

the " n -time steps" partition function would not obey the Chapman-Kolmogorov relation (14).

However, the following holds [37, 4, 16, 7].

1. There exist $A, B > 0$ such that, for any block ω_0^{n-1} :

$$A \leq \frac{\mu_{\boldsymbol{\lambda}}[\omega_0^{n-1}]}{e^{-(n-D)\mathcal{P}[\boldsymbol{\lambda}]} e^{\mathcal{H}_{\boldsymbol{\lambda}}(\omega_0^{n-1})}} \leq B. \quad (22)$$

2. We have:

$$\mathcal{P}[\boldsymbol{\lambda}] = \lim_{n \rightarrow \infty} \frac{1}{n} \log Z_n[\boldsymbol{\lambda}]. \quad (23)$$

In the spatial case, $Z_n[\boldsymbol{\lambda}] = Z^n[\boldsymbol{\lambda}]$ and $\mathcal{P}[\boldsymbol{\lambda}] = \log Z[\boldsymbol{\lambda}]$, whereas $A = B = 1$ in (22).

These relations are crucial for the developments made in the next section.

To finish this section let us introduce the Kullback-Leibler divergence $d_{KL}(\nu, \mu)$ which provides a notion of similarity between two probabilities ν, μ . We have $d_{KL}(\nu, \mu) \geq 0$ with equality if and only if $\mu = \nu$. The Kullback-Leibler divergence between an invariant probability $\nu \in \mathcal{M}$ and the Gibbs distribution $\mu_{\boldsymbol{\lambda}}$ with potential $\mathcal{H}_{\boldsymbol{\lambda}}$ is given by [7]:

$$d_{KL}(\nu, \mu_{\boldsymbol{\lambda}}) = \mathcal{P}[\boldsymbol{\lambda}] - \nu[\mathcal{H}_{\boldsymbol{\lambda}}] - \mathcal{S}[\nu]. \quad (24)$$

3 Inferring the coefficients of a potential from data

Equations (11) or (19) provide an analytical way to compute the coefficients of the Gibbs distribution from data. However, they require the computation of the partition function or of the topological pressure which becomes rapidly intractable as the number of neurons increases. Thus, researchers have attempted to find alternative methods to compute reliably and efficiently the λ_i s. An efficient method has been introduced in [10] and applied to spike trains in [5]. Although these papers are restricted to Gibbs distributions of the form (10) (models without memory) we show in this section how their method can be extended to general Gibbs distributions.

3.1 Bounding the Kullback-Leibler divergence variation

3.1.1 The spatial case

The method developed in [10] by Dudik et al is based on the so-called convex duality principle, used in mathematical optimization theory. Due the difficulty in maximizing the entropy (which is a concave function), one looks for a convex

function easier to investigate. Dudik et al showed that, for spatially constrained Maxent distributions, finding the Gibbs distribution amounts to finding the minimum of the negative log likelihood⁵:

$$L_{\pi_{\omega}^{(T)}}(\boldsymbol{\lambda}) = -\pi_{\omega}^{(T)} [\log \mu_{\boldsymbol{\lambda}}]. \quad (25)$$

Indeed, in the spatial case, the Kullback-Leibler divergence between the empirical measure $\pi_{\omega}^{(T)}$ and the Gibbs distribution at $\mu_{\boldsymbol{\lambda}}$ is:

$$d_{KL}(\pi_{\omega}^{(T)}, \mu_{\boldsymbol{\lambda}}) = \pi_{\omega}^{(T)} \left[\frac{\log \pi_{\omega}^{(T)}}{\log \mu_{\boldsymbol{\lambda}}} \right] = \pi_{\omega}^{(T)} [\log \pi_{\omega}^{(T)}] - \pi_{\omega}^{(T)} [\log \mu_{\boldsymbol{\lambda}}], \quad (26)$$

so that, from (24):

$$L_{\pi_{\omega}^{(T)}}(\boldsymbol{\lambda}) = \mathcal{P} [\boldsymbol{\lambda}] - \pi_{\omega}^{(T)} [\mathcal{H}_{\boldsymbol{\lambda}}]$$

Since \mathcal{P} is convex and $\pi_{\omega}^{(T)} [\mathcal{H}_{\boldsymbol{\lambda}}]$ linear in $\boldsymbol{\lambda}$, $L_{\pi_{\omega}^{(T)}}(\boldsymbol{\lambda})$ is convex. Its unique minimum is given by (11).

Moreover, we have:

$$L_{\pi_{\omega}^{(T)}}(\boldsymbol{\lambda}') - L_{\pi_{\omega}^{(T)}}(\boldsymbol{\lambda}) = \mathcal{P} [\boldsymbol{\lambda}'] - \mathcal{P} [\boldsymbol{\lambda}] - \pi_{\omega}^{(T)} [\Delta \mathcal{H}_{\boldsymbol{\lambda}}], \quad (27)$$

with $\Delta \mathcal{H}_{\boldsymbol{\lambda}} = \mathcal{H}_{\boldsymbol{\lambda}'} - \mathcal{H}_{\boldsymbol{\lambda}}$. From (10):

$$\begin{aligned} \frac{Z [\boldsymbol{\lambda}']}{Z [\boldsymbol{\lambda}]} &= \frac{1}{Z [\boldsymbol{\lambda}]} \sum_{\omega(0)} e^{\mathcal{H}_{\boldsymbol{\lambda}'}(\omega(0))} \\ &= \sum_{\omega(0)} e^{\Delta \mathcal{H}_{\boldsymbol{\lambda}}(\omega(0))} \mu_{\boldsymbol{\lambda}} [\omega(0)] \\ &= \mu_{\boldsymbol{\lambda}} [e^{\Delta \mathcal{H}_{\boldsymbol{\lambda}}}] \end{aligned} \quad (28)$$

and since $P[\boldsymbol{\lambda}] = \log Z[\boldsymbol{\lambda}]$ in the spatial case:

$$\mathcal{P} [\boldsymbol{\lambda}'] - \mathcal{P} [\boldsymbol{\lambda}] = \log \mu_{\boldsymbol{\lambda}} [e^{\Delta \mathcal{H}_{\boldsymbol{\lambda}}}]. \quad (29)$$

Therefore:

$$L_{\pi_{\omega}^{(T)}}(\boldsymbol{\lambda}') - L_{\pi_{\omega}^{(T)}}(\boldsymbol{\lambda}) = \log \mu_{\boldsymbol{\lambda}} [e^{\Delta \mathcal{H}_{\boldsymbol{\lambda}}}] - \pi_{\omega}^{(T)} [\Delta \mathcal{H}_{\boldsymbol{\lambda}}]. \quad (30)$$

The idea proposed by Dudik et al is then to bound this difference by an easier-to-compute convex quantity, with the same minimum as $L_{\pi_{\omega}^{(T)}}(\boldsymbol{\lambda})$, and to reach this minimum by iterations on $\boldsymbol{\lambda}$. They proposed a sequential and a parallel method. Let us summarize first the sequential method. The goal here is not to rewrite their paper [10] but to explain some crucial elements that are not directly applicable to the spatio-temporal case.

In the sequential case one updates $\boldsymbol{\lambda}$ as $\boldsymbol{\lambda}' = \boldsymbol{\lambda} + \delta \mathbf{e}_l$, for some l , where \mathbf{e}_l is the canonical basis vector in direction l , so that $\Delta \mathcal{H}_{\boldsymbol{\lambda}} = \delta m_l$, and

$$L_{\pi_{\omega}^{(T)}}(\boldsymbol{\lambda}') - L_{\pi_{\omega}^{(T)}}(\boldsymbol{\lambda}) = \log \mu_{\boldsymbol{\lambda}} [e^{\delta m_l}] - \delta \pi_{\omega}^{(T)} [m_l].$$

⁵We have adapted [10] to our notations. Moreover, in our case $\pi_{\omega}^{(T)}$ corresponds to the empirical average on a raster ω whereas π in [10] corresponds to an average over independent samples.

Using the following property:

$$e^{\delta x} \leq 1 + (e^\delta - 1)x, \quad (31)$$

for $x \in [0, 1]$ and since $m_l \in \{0, 1\}$, we have:

$$\log \mu_{\boldsymbol{\lambda}} [e^{\delta m_l}] \leq \log (1 + (e^\delta - 1)\mu_{\boldsymbol{\lambda}}[m_l]). \quad (32)$$

This bound, proposed by Dudik et al, is remarkably clever. Indeed, it replaces the computation of the average $\mu_{\boldsymbol{\lambda}} [e^{\delta m_l}]$, which is computationally hard, by the computation of $\mu_{\boldsymbol{\lambda}} [m_l]$, which is computationally easy. Finally,

$$L_{\pi_\omega^{(T)}}(\boldsymbol{\lambda}') - L_{\pi_\omega^{(T)}}(\boldsymbol{\lambda}) \leq -\delta \pi_\omega^{(T)} [m_l] + \log (1 + (e^\delta - 1)\mu_{\boldsymbol{\lambda}} [m_l]). \quad (33)$$

In the parallel case, the computation and results differ. One now updates $\boldsymbol{\lambda}$ as $\boldsymbol{\lambda}' = \boldsymbol{\lambda} + \sum_{l=1}^L \delta_l \mathbf{e}_l$. Moreover, one has to renormalize the m_l s in $m'_l = \frac{m_l}{L}$ in order that eq. (34) below holds. We have therefore $\Delta \mathcal{H}_{\boldsymbol{\lambda}} = \sum_{l=1}^L \delta_l m'_l$.

Thus,

$$L_{\pi_\omega^{(T)}}(\boldsymbol{\lambda}') - L_{\pi_\omega^{(T)}}(\boldsymbol{\lambda}) = \log \mu_{\boldsymbol{\lambda}} [e^{\sum_{l=1}^L \delta_l m'_l}] - \sum_{l=1}^L \delta_l \pi_\omega^{(T)} [m'_l].$$

Using the following property [9]:

$$e^{\sum_{l=1}^L \delta_l m'_l} \leq 1 + \sum_{l=1}^L m'_l (e^{\delta_l} - 1), \quad (34)$$

for $\delta_l \in \mathbb{R}$ and $m'_l \geq 0$, $\sum_{l=1}^L m'_l \leq 1$, we have:

$$\log \mu_{\boldsymbol{\lambda}} [e^{\sum_{l=1}^L \delta_l m'_l}] \leq \log \left(1 + \sum_{l=1}^L (e^{\delta_l} - 1) \mu_{\boldsymbol{\lambda}}[m'_l] \right).$$

Since $\log(1+x) \leq x$ for $x > -1$, Dudick et al obtain that:

$$\log \mu_{\boldsymbol{\lambda}} [e^{\sum_{l=1}^L \delta_l m'_l}] \leq \sum_{l=1}^L (e^{\delta_l} - 1) \mu_{\boldsymbol{\lambda}}[m'_l],$$

provided $\sum_{l=1}^L (e^{\delta_l} - 1) \mu_{\boldsymbol{\lambda}}[m'_l] > -1$. (This constraint has to be checked during iterations). Finally, using the definition of m'_l :

$$L_{\pi_\omega^{(T)}}(\boldsymbol{\lambda}') - L_{\pi_\omega^{(T)}}(\boldsymbol{\lambda}) \leq \frac{1}{L} \left[-\sum_{l=1}^L \delta_l \pi_\omega^{(T)} [m_l] + \sum_{l=1}^L (e^{\delta_l} - 1) \mu_{\boldsymbol{\lambda}}[m_l] \right]. \quad (35)$$

To be complete, let us mention that Dudik et al consider the case where some error ϵ_l is allowed in the estimation of the coefficient λ_l . This relaxation on the parameters alleviates the overfitting.

In this case, the bound on the right hand side in (33) (sequential case) becomes:

$$F_l(\boldsymbol{\lambda}, \delta) = -\delta \pi_\omega^{(T)} [m_l] + \log (1 + (e^\delta - 1)\mu_{\boldsymbol{\lambda}} [m_l]) + \epsilon_l (|\lambda_l + \delta| - |\lambda_l|). \quad (36)$$

whereas the right hand side in (35) becomes $\sum_{l=1}^L G_l(\boldsymbol{\lambda}, \boldsymbol{\delta})$ with:

$$G_l(\boldsymbol{\lambda}, \boldsymbol{\delta}) = \frac{1}{L} \left[-\delta_l \pi_\omega^{(T)} [m_l] + (e^{\delta_l} - 1) \mu_\lambda [m_l] \right] + \epsilon_l (|\lambda_l + \delta| - |\lambda_l|), \quad (37)$$

The minimum of these functions is easy to find and one obtains, for a given $\boldsymbol{\lambda}$ the variation $\boldsymbol{\delta}$ required to lower bound the log-likelihood variation. The authors have shown that both sequential and parallel method produce a sequence $\boldsymbol{\lambda}^{(k)}$ which converges to the minimum of $L_{\pi_\omega^{(T)}}$ as $k \rightarrow +\infty$.

3.1.2 Extension to the spatio-temporal case

We now show how to extend these computations to the spatio-temporal case, provided one replaces the log-likelihood $L_{\pi_\omega^{(T)}}$ by the Kullback-Leibler divergence (24). The main obstacle is that the Gibbs distribution does not have the form $\frac{e^{\mathcal{H}}}{Z}$. We obtain thus a convex criterion to minimize Kullback-Leibler divergence variation, hence reaching its minimum, $\pi_\omega^{(T)}$.

Replacing ν in eq. (24) by $\pi_\omega^{(T)}$, the empirical measure, one has:

$$d_{KL}(\pi_\omega^{(T)}, \mu_{\boldsymbol{\lambda}'}) - d_{KL}(\pi_\omega^{(T)}, \mu_{\boldsymbol{\lambda}}) = \mathcal{P}[\boldsymbol{\lambda}'] - \mathcal{P}[\boldsymbol{\lambda}] - \pi_\omega^{(T)}[\Delta\mathcal{H}_\lambda], \quad (38)$$

because the entropy $\mathcal{S}[\pi_\omega^{(T)}]$ cancels. This is the analog of (27). The main problem now is to compute $\mathcal{P}[\boldsymbol{\lambda}'] - \mathcal{P}[\boldsymbol{\lambda}]$.

From (22) we have:

$$\begin{aligned} & A e^{-(n-D)\mathcal{P}[\boldsymbol{\lambda}]} \sum_{\omega_0^{n-1}} e^{\mathcal{H}_\lambda(\omega_0^{n-1})} e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} \\ & \leq \sum_{\omega_0^{n-1}} \mu_\lambda[\omega_0^{n-1}] e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} \\ & \leq B e^{-(n-D)\mathcal{P}[\boldsymbol{\lambda}]} \sum_{\omega_0^{n-1}} e^{\mathcal{H}_\lambda(\omega_0^{n-1})} e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} \end{aligned}$$

so that:

$$\begin{aligned} & \lim_{n \rightarrow \infty} \frac{1}{n} \left[\log A - (n-D)\mathcal{P}[\boldsymbol{\lambda}] + \log \left(\sum_{\omega_0^{n-1}} e^{\mathcal{H}_\lambda(\omega_0^{n-1})} e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} \right) \right] \\ & \leq \lim_{n \rightarrow \infty} \frac{1}{n} \log \left(\sum_{\omega_0^{n-1}} \mu_\lambda[\omega_0^{n-1}] e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} \right) \\ & \leq \lim_{n \rightarrow \infty} \frac{1}{n} \left[\log B - (n-D)\mathcal{P}[\boldsymbol{\lambda}] + \log \left(\sum_{\omega_0^{n-1}} e^{\mathcal{H}_\lambda(\omega_0^{n-1})} e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} \right) \right]. \end{aligned} \quad (39)$$

Since $\mathcal{H}_{\boldsymbol{\lambda}'}(\omega_0^{n-1}) = \mathcal{H}_\lambda(\omega_0^{n-1}) + \Delta\mathcal{H}_\lambda(\omega_0^{n-1})$, from (23)

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \sum_{\omega_0^{n-1}} e^{\mathcal{H}_\lambda(\omega_0^{n-1})} e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} = \mathcal{P}[\boldsymbol{\lambda}'].$$

Therefore:

$$\mathcal{P}[\boldsymbol{\lambda}'] - \mathcal{P}[\boldsymbol{\lambda}] = \lim_{n \rightarrow \infty} \frac{1}{n} \log \sum_{\omega_0^{n-1}} \mu_{\boldsymbol{\lambda}}[\omega_0^{n-1}] e^{\Delta \mathcal{H}_{\boldsymbol{\lambda}}(\omega_0^{n-1})}. \quad (40)$$

This is the extension of (29) to the spatio temporal case. In the spatial case it reduces to (29) from (12). This equation is obviously numerically intractable, but it has two advantages: on one hand it allows to extend the bounds (33) (sequential case) and (35) (parallel case), and on the other hand it can be used to get a δ -power expansion of $\mathcal{P}[\boldsymbol{\lambda}'] - \mathcal{P}[\boldsymbol{\lambda}]$. This last point is used in the section 3.2.3.

To get the analogue of (33) in the sequential case where $\Delta \mathcal{H}_{\boldsymbol{\lambda}}(\omega_0^{n-1}) = \delta \sum_{r=0}^{n-D-1} m_l(\omega_r^{r+D})$, one may still apply (31) which holds provided:

$$m_l(\omega_0^{n-1}) \equiv \sum_{r=0}^{n-1-D} m_l(\omega_r^{r+D}) < 1 \quad (41)$$

So, compared to the spatial we have to replace m_l by $\frac{m_l}{n-D}$ in $\Delta \mathcal{H}_{\boldsymbol{\lambda}}(\omega_0^{n-1})$. We have therefore:

$$\begin{aligned} \sum_{\omega_0^{n-1}} \mu_{\boldsymbol{\lambda}}[\omega_0^{n-1}] e^{\Delta \mathcal{H}(\omega_0^{n-1})} &= \sum_{\omega_0^{n-1}} \mu_{\boldsymbol{\lambda}}[\omega_0^{n-1}] e^{\delta \frac{1}{n-D} m_l(\omega_0^{n-1})} \\ &\leq 1 + (e^\delta - 1) \frac{1}{n-D} \sum_{\omega_0^{n-1}} \mu_{\boldsymbol{\lambda}}[\omega_0^{n-1}] m_l(\omega_0^{n-1}). \end{aligned}$$

From the time translation invariance of $\mu_{\boldsymbol{\lambda}}$ we have:

$$\begin{aligned} \frac{1}{n-D} \sum_{\omega_0^{n-1}} \mu_{\boldsymbol{\lambda}}[\omega_0^{n-1}] m_l(\omega_0^{n-1}) &= \frac{1}{n-D} \sum_{r=0}^{n-D-1} \sum_{\omega_0^{n-1}} \mu_{\boldsymbol{\lambda}}[\omega_0^{n-1}] m_l(\omega_r^{r+D}) \\ &= \frac{1}{n-D} \sum_{r=0}^{n-D-1} \mu_{\boldsymbol{\lambda}}[m_l] \\ &= \mu_{\boldsymbol{\lambda}}[m_l] \end{aligned}$$

so that:

$$\sum_{\omega_0^{n-1}} \mu_{\boldsymbol{\lambda}}[\omega_0^{n-1}] e^{\delta \frac{1}{n-D} m_l(\omega_0^{n-1})} \leq 1 + (e^\delta - 1) \mu_{\boldsymbol{\lambda}}[m_l].$$

At first glance this bound is not really useful. Indeed, from (40) we obtain:

$$\mathcal{P}[\boldsymbol{\lambda}'] - \mathcal{P}[\boldsymbol{\lambda}] \leq \lim_{n \rightarrow \infty} \frac{1}{n} \log (1 + (e^\delta - 1) \mu_{\boldsymbol{\lambda}}[m_l]) = 0.$$

Since this holds for any δ this implies $\mathcal{P}[\boldsymbol{\lambda}'] = \mathcal{P}[\boldsymbol{\lambda}]$. The reason for this is evident. Renormalizing m_l as we did to match the condition imposed by bound (31) is equivalent to renormalizing δ by $\frac{\delta}{n-D}$. As $n \rightarrow +\infty$ this perturbation tends to 0 and $\boldsymbol{\lambda}' = \boldsymbol{\lambda}$. Therefore, the clever bound (31) would be of no interest if we were seeking exact results. However, the goal here is to propose a numerical scheme, where, obviously n is finite. We replace therefore the limit $n \rightarrow \infty$ by

a fixed n in the computation of $\mathcal{P}[\boldsymbol{\lambda}'] - \mathcal{P}[\boldsymbol{\lambda}]$. Keeping in mind that m_l must also be renormalized in $\pi_\omega^{(T)}[\Delta\mathcal{H}_\lambda]$ and using $\frac{1}{n} < \frac{1}{n-D}$ the Kullback-Leibler difference (38) obeys:

$$d_{KL}(\pi_\omega^{(T)}, \mu_{\boldsymbol{\lambda}'}) - d_{KL}(\pi_\omega^{(T)}, \mu_{\boldsymbol{\lambda}}) \leq \frac{1}{n-D} \left[-\delta\pi_\omega^{(T)}[m_l] + \log(1 + (e^\delta - 1)\mu_{\boldsymbol{\lambda}}[m_l]) \right], \quad (42)$$

the analog of (33).

In the parallel case, similar remarks holds. In order to apply the bound (34) we have to renormalize the m_l s in $m'_l = \frac{1}{L(n-D)}$. As for the spatial case we also need to check that $\sum_{l=1}^L (e^{\delta_l} - 1) \mu_{\boldsymbol{\lambda}}[m'_l] > -1$. (This constraint is not guarantee and has to be checked during iterations). One obtains finally:

$$d_{KL}(\pi_\omega^{(T)}, \mu_{\boldsymbol{\lambda}'}) - d_{KL}(\pi_\omega^{(T)}, \mu_{\boldsymbol{\lambda}}) \leq \frac{1}{L(n-D)} \left[-\sum_{l=1}^L \delta_l \pi_\omega^{(T)}[m_l] + \sum_{l=1}^L (e^{\delta_l} - 1) \mu_{\boldsymbol{\lambda}}[m_l] \right], \quad (43)$$

the analogue of (35).

Compared with the spatial case, we see therefore that n mustn't be too large to have a reasonable Kullback-Leibler divergence variation. It mustn't be too small, however, to get a good approximation of the topological pressure.

3.2 Updating the target distribution when the parameters change

When updating the parameters $\boldsymbol{\lambda}$, one has to compute again the average values $\mu_{\boldsymbol{\lambda}}[m_l]$ since the probability $\mu_{\boldsymbol{\lambda}}$ has changed. This has a huge computational cost. The exact computation (e.g. from (11, 19)) is not tractable for large N so approximate methods have to be used, like Montecarlo [31]. Again, this is also CPU time consuming especially if one recomputes it again at each iteration, but at least it is tractable.

In this spirit, Broderick et al [5] propose to generate a Montecarlo raster distributed according to $\mu_{\boldsymbol{\lambda}}$ and to use it to compute $\mu_{\boldsymbol{\lambda}'}$ when $\|\boldsymbol{\lambda}' - \boldsymbol{\lambda}\|$ is sufficiently small. We expose their method, limited to the spatial case, in the next section, and we explain why it is not applicable in the spatio-temporal case. We then propose an alternative method.

3.2.1 The spatial case

The average of m_l is obtained by the derivative of the topological pressure $\mathcal{P}[\boldsymbol{\lambda}]$. In the spatial case, where $\mathcal{P}(\boldsymbol{\lambda}) = \log Z_{\boldsymbol{\lambda}}$, we have:

$$\begin{aligned} \mu_{\boldsymbol{\lambda}'}[m_l] &= \frac{\partial \mathcal{P}(\boldsymbol{\lambda}')}{\partial \lambda'_j} \\ &= \frac{1}{Z[\boldsymbol{\lambda}']} \sum_{\omega(0)} m_l(\omega(0)) e^{\mathcal{H}_{\boldsymbol{\lambda}'}(\omega(0))} \\ &= \frac{Z[\boldsymbol{\lambda}]}{Z[\boldsymbol{\lambda}']} \sum_{\omega(0)} m_l(\omega(0)) e^{\Delta\mathcal{H}_\lambda(\omega(0))} \mu_{\boldsymbol{\lambda}}[\omega(0)] \end{aligned} \quad (44)$$

Using (28), one finally obtains:

$$\mu_{\lambda'} [m_l] = \frac{\mu_{\lambda} [m_l(\omega(0)) e^{\Delta\mathcal{H}_{\lambda}(\omega(0))}]}{\mu_{\lambda} [e^{\Delta\mathcal{H}_{\lambda}(\omega(0))}]}, \quad (45)$$

which is eq. (18) in [5]. Using this formula one is able to compute the average of m_l with respect to the new probability $\mu_{\lambda'}$ only using the old one, μ_{λ} .

3.2.2 Extension to the spatio-temporal case

We now explain why the Broderick et al method does not extend to the spatio-temporal case. The main problem is that if one tries to obtain the analogue of the equality (45) one obtains in fact an inequality:

$$\frac{A}{B} \mu_{\lambda'} [m_l] \leq \lim_{n \rightarrow \infty} \frac{1}{n} \frac{\mu_{\lambda} [m_l(\omega_0^{n-1}) e^{\Delta\mathcal{H}_{\lambda}(\omega_0^{n-1})}]}{\mu_{\lambda} [e^{\Delta\mathcal{H}_{\lambda}(\omega_0^{n-1})}]} \leq \frac{B}{A} \mu_{\lambda'} [m_l], \quad (46)$$

where A, B are the constants in (22). They are not known in general (they depend on the potential) and they are different. However, in the spatial case $A = B = 1$ whereas $\mu_{\lambda} [m_l(\omega_0^{n-1}) e^{\Delta\mathcal{H}_{\lambda}(\omega_0^{n-1})}] = \mu_{\lambda} [m_l(\omega(0)) e^{\Delta\mathcal{H}_{\lambda}(\omega(0))}]$ because the potential has range 1. Then, one recovers (45). Let us now explain how we obtain (46).

The averages of quantities are obtained by the derivative of the topological pressure (Eq. (19)). We have:

$$\mu_{\lambda'} [m_l] = \frac{\partial \mathcal{P}}{\partial \lambda'_l} = \frac{\partial \lim_{n \rightarrow \infty} \frac{1}{n} \log Z_n [\lambda']}{\partial \lambda'_l}. \quad (47)$$

Assuming that the limit and the derivative commute (see e.g. [29]), gives:

$$\begin{aligned} \mu_{\lambda'} [m_l] &= \lim_{n \rightarrow \infty} \frac{1}{n} \frac{1}{Z_n [\lambda']} \sum_{\omega_0^{n-1}} m_l(\omega_0^{n-1}) e^{\mathcal{H}_{\lambda'}(\omega_0^{n-1})} \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \frac{1}{Z_n [\lambda']} \sum_{\omega_0^{n-1}} m_l(\omega_0^{n-1}) e^{\Delta\mathcal{H}_{\lambda}(\omega_0^{n-1})} e^{\mathcal{H}_{\lambda}(\omega_0^{n-1})} \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \frac{\sum_{\omega_0^{n-1}} m_l(\omega_0^{n-1}) e^{\Delta\mathcal{H}_{\lambda}(\omega_0^{n-1})} e^{\mathcal{H}_{\lambda}(\omega_0^{n-1})}}{\sum_{\omega_0^{n-1}} e^{\Delta\mathcal{H}_{\lambda}(\omega_0^{n-1})} e^{\mathcal{H}_{\lambda}(\omega_0^{n-1})}} \end{aligned} \quad (48)$$

From (22):

$$\begin{aligned} &A e^{-(n-D)\mathcal{P}[\lambda]} \sum_{\omega_0^{n-1}} m_l(\omega_0^{n-1}) e^{\Delta\mathcal{H}_{\lambda}(\omega_0^{n-1})} e^{\mathcal{H}_{\lambda}(\omega_0^{n-1})} \\ &\leq \sum_{\omega_0^{n-1}} m_l(\omega_0^{n-1}) e^{\Delta\mathcal{H}_{\lambda}(\omega_0^{n-1})} \mu_{\lambda} [\omega_0^{n-1}] \\ &\leq B e^{-(n-D)\mathcal{P}[\lambda]} \sum_{\omega_0^{n-1}} m_l(\omega_0^{n-1}) e^{\Delta\mathcal{H}_{\lambda}(\omega_0^{n-1})} e^{\mathcal{H}_{\lambda}(\omega_0^{n-1})} \end{aligned} \quad (49)$$

and:

$$\begin{aligned}
& A e^{-(n-D)\mathcal{P}[\boldsymbol{\lambda}]} \sum_{\omega_0^{n-1}} e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} e^{\mathcal{H}_\lambda(\omega_0^{n-1})} \\
& \leq \sum_{\omega_0^{n-1}} e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} \mu_\lambda [\omega_0^{n-1}] \\
& \leq B e^{-(n-D)\mathcal{P}[\boldsymbol{\lambda}]} \sum_{\omega_0^{n-1}} e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} e^{\mathcal{H}_\lambda(\omega_0^{n-1})}.
\end{aligned}$$

Therefore:

$$\begin{aligned}
& \frac{A \sum_{\omega_0^{n-1}} m_l (\omega_0^{n-1}) e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} e^{\mathcal{H}_\lambda(\omega_0^{n-1})}}{B \sum_{\omega_0^{n-1}} e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} e^{\mathcal{H}_\lambda(\omega_0^{n-1})}} \\
& \leq \frac{\sum_{\omega_0^{n-1}} m_l (\omega_0^{n-1}) e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} \mu_\lambda [\omega_0^{n-1}]}{\sum_{\omega_0^{n-1}} e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} \mu_\lambda [\omega_0^{n-1}]} \\
& \leq \frac{B \sum_{\omega_0^{n-1}} m_l (\omega_0^{n-1}) e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} e^{\mathcal{H}_\lambda(\omega_0^{n-1})}}{A \sum_{\omega_0^{n-1}} e^{\Delta\mathcal{H}_\lambda(\omega_0^{n-1})} e^{\mathcal{H}_\lambda(\omega_0^{n-1})}}.
\end{aligned}$$

Now, from [7, 24], (48) gives (46).

3.2.3 Taylor expansion of the pressure

The idea is here to use a Taylor expansion of the topological pressure. This approach is very much in the spirit of [23], but extended here to the spatio-temporal case. Since $\boldsymbol{\lambda}' = \boldsymbol{\lambda} + \boldsymbol{\delta}$, we have:

$$\begin{aligned}
\mu_{\boldsymbol{\lambda}'} [m_l] &= \mu_\lambda [m_l] + \sum_{k=1}^K \frac{\partial \mu_\lambda [m_l]}{\partial \lambda_k} \delta_k + \frac{1}{2} \sum_{k,l=1}^K \frac{\partial^2 \mu_\lambda [m_l]}{\partial \lambda_k \partial \lambda_l} \delta_k \delta_l + \dots \\
&= \mu_\lambda [m_l] + \sum_{k=1}^K \frac{\partial^2 \mathcal{P}[\boldsymbol{\lambda}]}{\partial \lambda_j \partial \lambda_k} \delta_k + \frac{1}{2} \sum_{k,l=1}^K \frac{\partial^3 \mathcal{P}[\boldsymbol{\lambda}]}{\partial \lambda_j \partial \lambda_k \partial \lambda_l} \delta_k \delta_l + \dots \quad (50)
\end{aligned}$$

The second derivative of the pressure is given by [37, 4, 16, 7]:

$$\frac{\partial^2 \mathcal{P}[\boldsymbol{\lambda}]}{\partial \lambda_j \partial \lambda_k} = \sum_{n=-\infty}^{+\infty} C_{jk}(n) \equiv \chi_{jk}[\boldsymbol{\lambda}], \quad (51)$$

where:

$$C_{jk}(n) = \mu_\lambda [m_l m_k \circ \sigma^n] - \mu_\lambda [m_l] \mu_\lambda [m_k], \quad (52)$$

is the correlation function between m_l, m_k at time n , computed with respect to μ_λ . (51) is a version of the fluctuation-dissipation theorem in the spatio-temporal case. σ^n is the time shift applied n times. The third derivatives can be computed as well by taking the derivative (51) and using (47). This generates terms with third order correlations and so on [29]. Up to second order we have:

$$\mu_{\lambda'} [m_l] = \mu_{\lambda} [m_l] + \sum_{l=1}^L \chi_{jl} [\boldsymbol{\lambda}] \delta_l + \dots \quad (53)$$

Since the observable are monomials they only take the values 0 or 1 and the computation of χ_{ij} is straightforward, reducing to counting the occurrence of time pairs $t, t+n$ such that $m_l(t) = 1$ and $m_k(t+n) = 1$.

On practical grounds we introduce a parameter $\Delta = \|\boldsymbol{\lambda}' - \boldsymbol{\lambda}\|$ which measures the variation in the parameters after update. If Δ is small enough (smaller than some Δ_c), the terms of order 3 in the Taylor expansion are negligible, then we can use (53). Otherwise, if Δ is big, we compute a new Montecarlo estimation of μ'_{λ} (as described in [31]). We explain in section 4.2 how Δ_c was chosen in our data. Then, we use the following trick. If $\|\boldsymbol{\delta}\| > \Delta_c$ we compute the new value $\mu_{\lambda'} [m_j]$. If $\Delta_c > \|\boldsymbol{\delta}\| > \frac{\Delta_c}{10}$, we use the linear response approximation (53) of $\mu_{\lambda'}$. Finally, if $\|\boldsymbol{\delta}\| < \frac{\Delta_c}{10}$ we use $\mu_{\lambda} [m_j]$ instead of $\mu_{\lambda'} [m_j]$ in the next iteration of the method. Thus, in the case, $\|\boldsymbol{\delta}\| < \Delta_c$, we use the Gibbs distribution computed at some time step, say n , to infer the values at the next iteration. If we do that several successive time steps the distance to the original value $\boldsymbol{\lambda}_n$ of the parameters increases. So we compute the norm $\|\boldsymbol{\lambda}_n - \boldsymbol{\lambda}_{n+k}\|$ at each time step k , and we do not compute a new raster until this norm is larger than Δ_c .

3.3 The algorithms

We have two algorithms, sequential and parallel, which are very similar to Dudik et al. Especially, the convergence of their algorithms, proved in their paper, extends to our case since it only depends on the shape of the cost functions (36, 37). We describe here the algorithms coming out from the presented mathematical framework, in a sequential and parallel version. We iterate the algorithms until the distance $\epsilon = d(\mu_{\lambda}, \pi_{\omega}^{(T)})$ is smaller than some ϵ_c . We use the Hellinger distance:

$$d(\mu_{\lambda}, \pi_{\omega}^{(T)}) = \frac{1}{\sqrt{2}} \sqrt{\sum_{j=1}^L \left(\sqrt{\pi_{\omega}^{(T)}(m_l)} - \sqrt{\mu_{\lambda}(m_l)} \right)^2} \quad (54)$$

3.3.1 Sequential algorithm

Input: The features probabilities $\pi_\omega^{(T)} [m_l]$

Output: The vector of parameters λ

initialization: $\lambda_l = 0$ for every l , $\Delta = 0$

```

while  $\epsilon > \epsilon_c$  do
   $(\delta, l) = \arg \min_{j, \delta} F_j(\lambda, \delta)$ 
   $\lambda_l \leftarrow \lambda_l + \delta$ 
   $\Delta \leftarrow \Delta + |\delta|$ 
  if  $\Delta > \Delta_c$  then
    | Compute a new Gibbs sample using Montecarlo method [31]
  else
    | Compute the new features probabilities using Taylor expansion
    | (Equation 53)
  end
end

```

Algorithm 1: Sequential algorithm. δ is the learning rate by which we change the value of a parameter λ_l . ϵ is the convergence criterion (54). Δ is the parameter allowing us to decide whether we update the parameters change by computing a new Gibbs sample or by the Taylor expansion. F_j is given by eq. (36)

3.4 Parallel algorithm

Input: features probabilities $\pi_\omega^{(T)} [m_l]$

Output: parameters λ_l

initialization: $\lambda_l = 0$ for every l , $\Delta = 0$

```

while  $\epsilon > \epsilon_c$  do
  for  $i \leftarrow 1$  to  $L$  do
    |  $\delta_j = \arg \min_{\delta} G_j(\lambda, \delta)$ 
  end
   $\lambda \leftarrow \lambda + \delta$ 
   $\Delta \leftarrow \sqrt{\Delta^2 + \sum_{l=1}^L \delta_l^2}$ 
  if  $\Delta > \Delta_c$  then
    | Compute a new Gibbs sample using Montecarlo method [31]
  else
    | Compute the new features probabilities using Taylor expansion
    | (Equation 53)
  end
end

```

Algorithm 2: The parallel algorithm. G_l is given by (37).

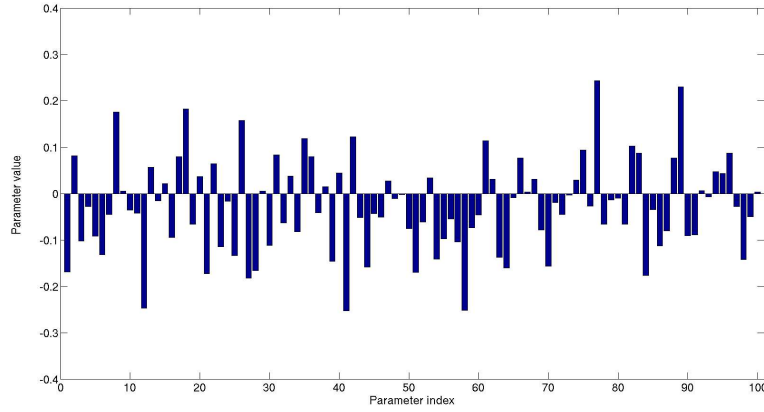
4 Results

In this section we perform several tests on our method. We first consider synthetic data generated with a known Gibbs potential and recover its parameters. This step also allows us to tune the parameter Δ_c in the algorithms. Then, we consider real data analysis where the Gibbs potential form is unknown.

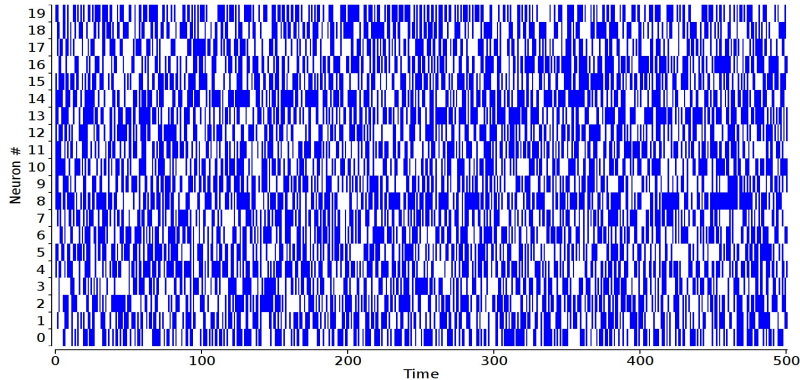
4.1 Synthetic data

Synthetic data are obtained by generating a raster distributed according to a Gibbs distribution whose potential (2) is known. We consider two families of Gibbs potentials. For each family there are $L > N$ monomials whose range belongs to $\{1, \dots, R\}$. Among them, there are N "rate monomials" $\omega_i(D)$, $i = 1 \dots N$, whose average gives the firing rate of neuron i , denoted r_i ; the $L - N$ other monomials, with degree $k > 1$, are chosen at random with a probability law $\sim e^{-k}$ which favors therefore pairwise interactions. The difference between the two families comes from the distribution of coefficients λ_l .

1. **"Dense" rasters family.** The coefficients are drawn with a Gaussian distribution with mean 0 and variance $\frac{1}{L}$ to ensure a correct scaling of the coefficients dispersion as L increases (fig. 1(a)). This produces typically a dense raster (fig. 1(b)) with strong multiple correlations.



(a) Example of coefficients distribution in the dense rasters family.

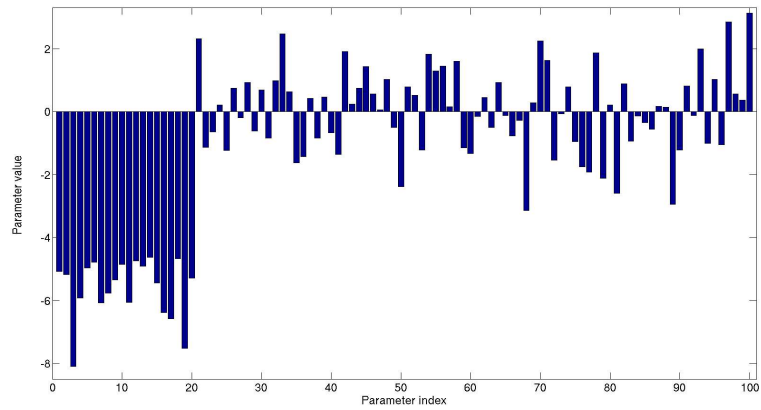


(b) Dense spike train

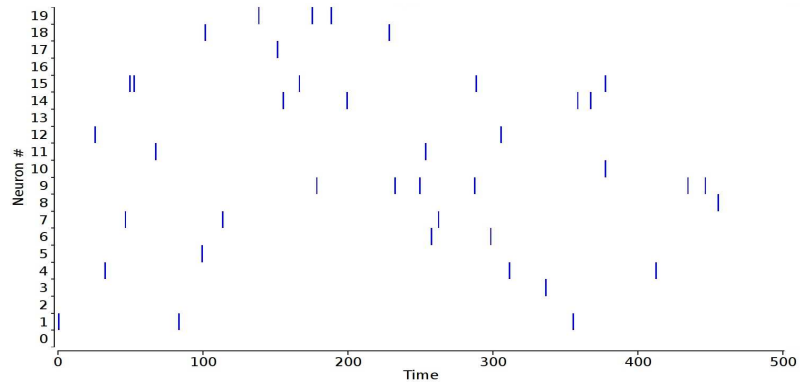
Figure 1: Dense family.

2. **"Sparse" rasters family.** The rate coefficients in the potential are very negative: the coefficient h_i of the rate monomial $\omega_i(D)$ is $h_i = \log\left(\frac{r_i}{1-r_i}\right)$

where $r_i \in [0 : 0.01]$ with a uniform probability distribution. Other coefficients are drawn with a Gaussian distribution with mean 0.8 and variance 1 (fig. 2(a)a). This produces a sparse raster (fig. 2(b)b) with strong multiple correlations.



(a) Example of coefficients distribution in the sparse rasters family.



(b) Sparse spike train

Figure 2: Sparse family.

4.2 Tuning Δ_c

For small N, R ($NR \leq 20$) it is possible to exactly compute the topological pressure using a transfer matrix technique [44]. We have therefore a way to compare the Taylor expansion (51) and the exact value.

If we perturb λ by an amount δ in the direction j , this induces a variation on $\mu_\lambda[m_j]$, $j = 1 \dots L$, given by the Taylor expansion (53). To the lowest order $\mu_{\lambda'}[m_j] = \mu_\lambda[m_j] + O^{(1)}$, so that:

$$\epsilon^{(1)} = \frac{1}{L} \sum_{j=1}^L \frac{|\mu_{\lambda'}[m_j] - \mu_\lambda[m_j]|}{|\mu_{\lambda'}[m_j]|}$$

is a measure of the relative error when considering the lowest order expansion.

In the same way, to the second order:

$$\mu_{\lambda'}[m_j] = \mu_\lambda[m_j] + \chi_{jk}[\lambda] \delta_k + O^{(2)},$$

so that:

$$\epsilon^{(2)} = \frac{1}{L} \sum_{j=1}^L \frac{|\mu_{\lambda'}[m_j] - \mu_\lambda[m_j] - \chi_{jk}[\lambda] \delta_k|}{|\mu_{\lambda'}[m_j]|},$$

is a measure of the relative error when considering the next order expansion.

In fig. 3 we show the relative errors $\epsilon^{(1)}, \epsilon^{(2)}$ (in %), as a function of δ . For each point we generate 25 potentials, with $N = 5, R = 3, L = 12$. For each of these potentials we randomly perturb the λ_j s, with a random sign, so that the norm of the perturbation $\|\delta\|$ is fixed. The linear response χ is computed from a raster of length $T = 100000$.

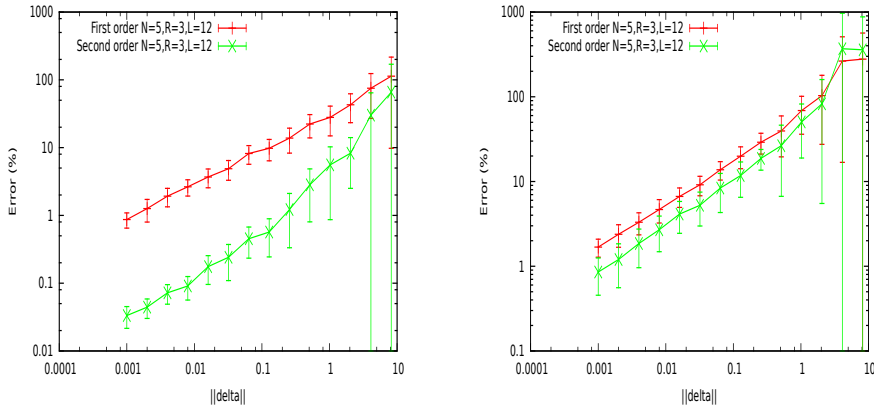


Figure 3: Error on the average $\mu_{\lambda'}[m_j]$ as a function of the perturbation amplitude δ . First order corresponds to $\epsilon^{(1)}$ and second order to $\epsilon^{(2)}$ (see text). The curves correspond to $N = 5, R = 3, L = 12$. Left: Dense case; Right: Sparse case.

These curves show a big difference between the dense and sparse case. In the dense case, the second order error is about 5% for $\Delta_c = 1$ whereas we need a $\Delta_c \sim 0.03$ to get the same 5% in the sparse case. We choose to align on the sparse case and in typical experiments we take $\Delta_c = 0.1$ corresponding to about 10% of error on the second order.

4.3 Performances on synthetic data

Here, we test the method on synthetic data where the shape of the sought potential is known: only the λ_l s have to be estimated. Experiments were designed according to the following steps:

- We start from a potential $\mathcal{H}_{\lambda^*} = \sum_{l \in \mathcal{L}} \lambda_l^* m_l$. The goal is to estimate the coefficient values λ_l^* knowing the set \mathcal{L} of monomials spanning the potential.
- We generate a synthetic spike train (ω_s) distributed according to the Gibbs distribution of \mathcal{H}_{λ^*} .
- We take a potential $\mathcal{H}_{\lambda} = \sum_{l \in \mathcal{L}} \lambda_l m_l$ with random initial coefficients λ_l . Then we fit the parameters λ_l to the synthetic spike train $\omega_s^{(T)}$.
- We evaluate the goodness of fit.

For the last step (goodness of fit) we have used two criteria. The first one simply consists of computing the L_1 error $d_1 = \frac{1}{L} \sum_{l=1}^L \left| \lambda_l^* - \lambda_l^{(est)} \right|$ where $\lambda_k^{(est)}$ is the final estimated value. d_1 is then averaged on 10 random potentials. Fig. 4 shows the committed error in the case of sparse and dense potentials. The method showed a good performance, both in dense and sparse case, for large $N \times R \sim 60$.

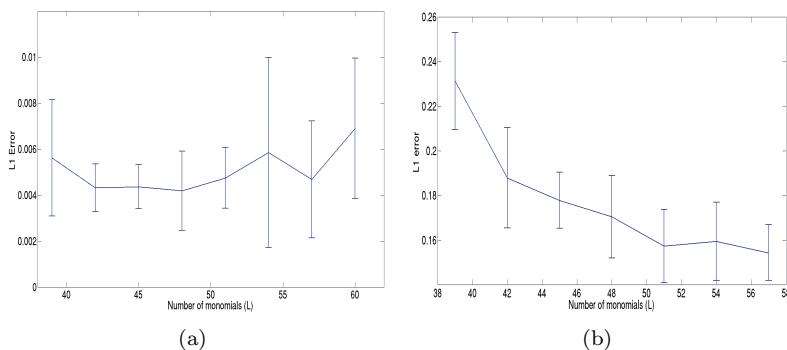


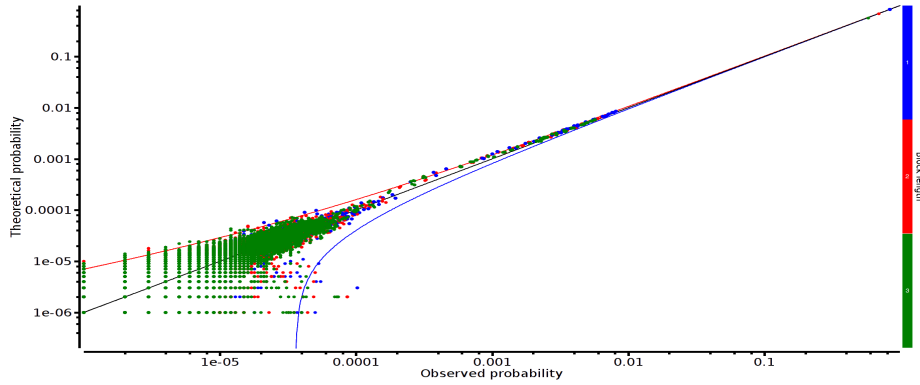
Figure 4: Max distance between the exact value of coefficients and the estimated value, averaged on the set of 10 random potentials for $NR = 60$. (a) Dense spike trains (b) Sparse spike trains.

The main advantage of this criterion is to provide an exact estimation of the error made on coefficients estimation. Its drawback is that we have to know the shape of the potential which generated the raster: this is not the case anymore for real neural networks data. We therefore used a second criterion: confidence plots. For each spike block ω_0^D appearing in the raster ω_s we draw a point in a two dimensional diagram with, on abscissa, the observed empirical probability $\pi_{\omega_s}^{(T)} [\omega_0^D]$ and, on ordinate, the predicted probability $\mu_{\lambda} [\omega_0^D]$. Ideally, all points should align on the diagonal $y = x$ (equality line). However, since the raster is finite there are finite-size fluctuations ruled by the central limit theorem.

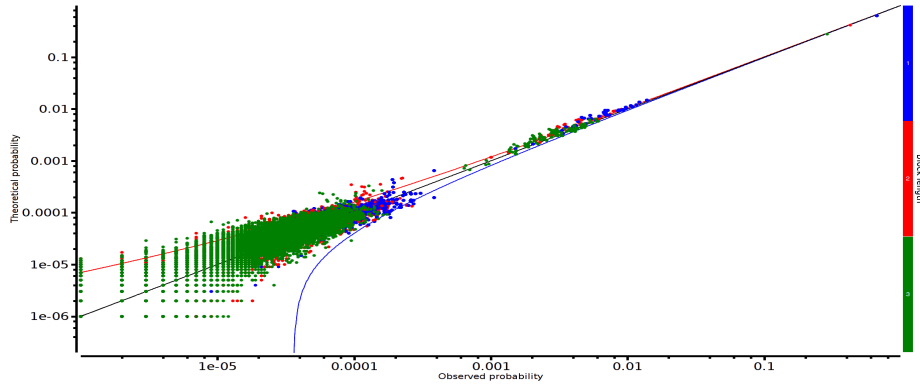
For a block ω_0^D generated by a Gibbs distribution μ_λ and having an exact probability $\mu_\lambda[\omega_0^D]$ the empirical probability $\pi_{\omega_s}^{(T)}[\omega_0^D]$ is a Gaussian random variable with mean $\mu_\lambda[\omega_0^D]$ and mean-square deviation $\sigma = \frac{\sqrt{\mu_\lambda[\omega_0^D](1-\mu_\lambda[\omega_0^D])}}{\sqrt{T}}$. The probability that $\pi_{\omega_s}^{(T)}[\omega_0^D] \in [\mu_\lambda[\omega_0^D] - 3\sigma, \mu_\lambda[\omega_0^D] + 3\sigma]$ is therefore of about 99,6%. This interval is represented by confidence lines spreading around the diagonal.

We have tested the following cases.

1. Spatial case, 40 neurons, ($NR = 40$): Ising model (3). Fig. 5(a) .
2. Spatio-temporal, 40 neurons, $R = 2$ ($NR = 80$): Pairwise model with delays (5). Fig. 5(b)



(a)



(b)

Figure 5: The confidence plots for two examples on synthetic data: (a) Ising and (b) pairwise of range 2. The 2 figures show the comparison between the observed probabilities of patterns (x-axis) and the predicted probabilities of patterns (those which are predicted by the fitted model). In order to show examples for different patterns sizes, we plotted each size with a color. For instance, the blue points correspond to patterns of size 1, red dots for patterns of size 2 and green dots for patterns of size 3.

4.4 The performance on real data

Here we show the inferring of MaxEnt distribution on real spike trains. We analyzed a set of 40 neurons (courtesy of M. J. Berry and O. Marre) with spatial and spatio-temporal constraints. Data are binned at 20 ms. In this case we don't know what is the right potential which generated data. Consequently, we cannot use the L_1 error, as in the previous section, to test the goodness of fit. We restrict therefore here to confidence plots. The goal here is to check the goodness of fit not only for spatial patterns (as done in [39, 33, 14, 13]), but also for spatio-temporal patterns. Consequently, the figures below draw 3 models comparison where each line corresponds to patterns of range 1, 2, 3.

We have analysed the data with the following potentials.

- Spatial models:
 1. Ising model (3). Fig. 6(a), 6(d), 6(g).
 2. GSS model (4). Fig. 6(b), 6(e), 6(h).
- Pairwise spatio-temporal model (5), $R = 2$. Fig. 6(c), 6(f), 6(i).

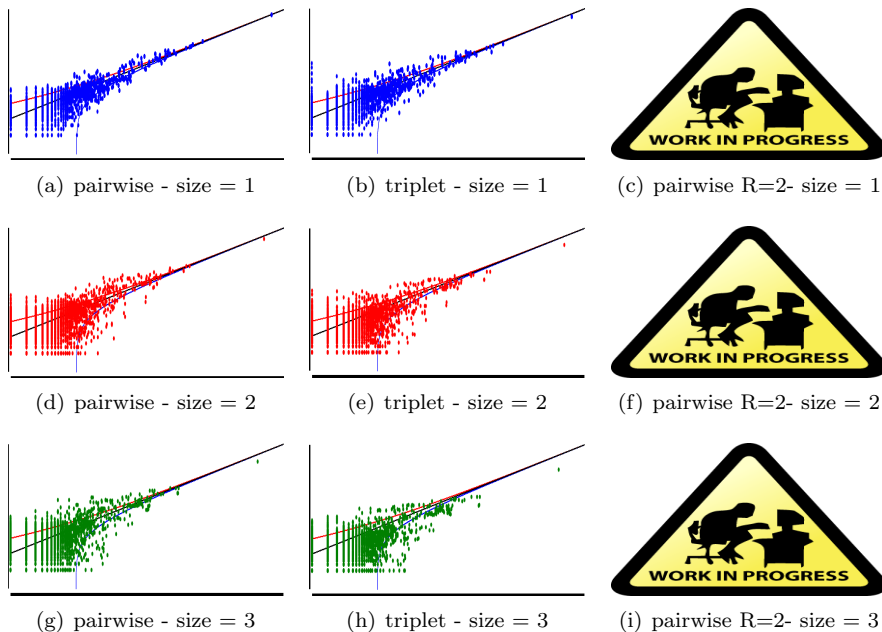


Figure 6: The comparison between observed probability (x-axis) of patterns and predicted probability (y-axis). (a), (b) and (c) corresponds to the comparison of probabilities of patterns of size 1 with respect to 3 models (respectively: pairwise instantaneous or Ising, triplets and pairwise with memory 1 time step). (d), (e) and (f) show the comparison for patterns of size 2. Finally, (g), (h) and (i) show patterns of size 3.

5 Discussion and conclusion

The method shows clearly quite better performances for synthetic data than for real data. The main reason, we believe, is that in the second case we don't know the form of the potential. As a consequence, we stick at existing canonical forms of potentials (Ising, GSS, pairwise). The main problem with this approach is that the number of parameters to estimate dramatically grows with NR . The increase is moderate for the Ising model (N rates + $\frac{N(N-1)}{2}$ symmetric pairwise couplings) but it becomes prohibitively large even for pairwise range R models. On the opposite, our analysis of synthetic data used a relatively small number of parameters to fit.

The large number of parameters has 2 drawbacks: the increasing of computation time and the increase of errors in the estimation. Let us comment on the second problem. It is not intrinsic to our method; it is neither intrinsic to MaxEnt; this is a well known problem which arises already when doing linear regression analysis. Increasing the number of parameters may eventually lead to catastrophic estimations where the addition of degree of freedom can seriously hinder the resolution.

In the case of MaxEnt the situation can be described as follows. We generate a finite raster ω_0^T from a known distribution μ_{λ^*} with a potential of the form 2. Denote $\mu_{\lambda^*}[\mathbf{m}]$ the vector with entries $\mu_{\lambda^*}[m_l]$ and $\pi_{\omega}^{(T)}[\mathbf{m}]$ the vector with entries $\pi_{\omega}^{(T)}[m_l]$. From (19) we have $\mu_{\lambda^*}[\mathbf{m}] = \nabla_{\lambda^*} \mathcal{P}$. This exact solution is obtained when the Gibbs distribution μ_{λ^*} can be exactly sampled, namely, for an infinite raster. For a finite raster, if T is large enough to apply the central limit theorem, the empirical distribution $\pi_{\omega}^{(T)}[\mathbf{m}]$ is Gaussian with mean $\mu_{\lambda}[\mathbf{m}]$ and covariance $\frac{1}{T}\chi$ given by (51). We have therefore $\pi_{\omega}^{(T)}[\mathbf{m}] = \mu_{\lambda^*}[\mathbf{m}] + \boldsymbol{\eta}$ where $\boldsymbol{\eta}$ is centered Gaussian with covariance $\frac{1}{T}\chi$. Solving (19) where the exact probability μ_{λ^*} is replaced by the empirical one $\pi_{\omega}^{(T)}$, one obtains an approximate solution of $\boldsymbol{\lambda}$, $\boldsymbol{\lambda}^*$ with $\boldsymbol{\lambda} = \boldsymbol{\lambda}^* + \boldsymbol{\epsilon}$, where $\nabla_{\boldsymbol{\lambda}} \mathcal{P} = \pi_{\omega}^{(T)}[\mathbf{m}]$. Therefore, $\nabla_{\boldsymbol{\lambda}} \mathcal{P} = \mu_{\lambda^*}[\mathbf{m}] + \boldsymbol{\eta} = \nabla_{\boldsymbol{\lambda}^* + \boldsymbol{\epsilon}} \mathcal{P} = \nabla_{\boldsymbol{\lambda}^*} \mathcal{P} + \boldsymbol{\epsilon} \chi + O(\|\boldsymbol{\epsilon}\|^2)$. Hence, $\boldsymbol{\epsilon} = \chi^{-1} \boldsymbol{\eta}$. χ is invertible since \mathcal{P} is convex.

The fluctuations of the estimated solution $\boldsymbol{\lambda}$ around the exact solution $\boldsymbol{\lambda}^*$ are therefore Gaussian, centered, with covariance $\mathbb{E}[\boldsymbol{\epsilon} \cdot \tilde{\boldsymbol{\epsilon}}] = \mathbb{E}[\chi^{-1} \cdot \boldsymbol{\eta} \cdot \tilde{\boldsymbol{\eta}} \cdot \chi^{-1}]$. Since χ is symmetric we have $\mathbb{E}[\boldsymbol{\epsilon} \cdot \tilde{\boldsymbol{\epsilon}}] = \chi^{-1} \cdot \mathbb{E}[\boldsymbol{\eta} \cdot \tilde{\boldsymbol{\eta}}] \cdot \chi^{-1} = \frac{1}{T} \chi^{-1}$. We arrive therefore at the conclusion that the fluctuations on the estimated coefficients $\boldsymbol{\lambda}$ are highly constrained by the convexity of the pressure, as expected. Mathematically, everything goes nicely since \mathcal{P} is convex. However, when we estimate the solution numerically, it may happen that \mathcal{P} is quite flat in some directions/monomials. This arises a fortiori since we are not handling the exact values but estimations. Therefore, when considering potentials of the form (2) it is expected that some terms (monomials) not only are irrelevant, but also dramatically deteriorate the estimation problem, introducing almost zero eigenvalues in χ .

At this stage, the main question is therefore: Can we have an idea of the potential shape from data before fitting the parameters? This question is not only related to the goodness of fit but, it is also a question of concept. Is it useful to represent a pairwise distribution for 40 neurons with nearly 2000 parameters? The idea would then be to filter irrelevant monomials. For that a feature selection method is useful and should complement this work. There are

many direction we can take in the favor of the features selection. For instance, selecting the features on threshold ([35, 25]), using a χ^2 method ([8]) as well as incremental feature selection algorithm ([2], [45]). Other methods based on periodic orbit sampling ([6]) and information geometry ([30, 1]) are under current investigation.

We have presented a method to fit the parameters of MaxEnt distribution with spatio-temporal constraints. In the process of exploring the dynamics of neural data, we hypothesize the model, fit it and finally judge the quality of the suggested model. Hence, this work is positioned as an important intermediate step in the neural coding using the MaxEnt framework, opening the door to analyze the dynamics of large networks being not limited to spatial and/or traditional MaxEnt models.

The implementation of parameters fitting consists on an important part in a software developed at INRIA and called EnaS (Event Neural Assembly Simulation). The software is freely available at <http://enas.gforge.inria.fr/v3/download.html>.

6 Acknowledgments

We warmly acknowledge M. Berry, O. Marre for providing us MEA recordings from the retina and G. Tkacik who provided us the references [10, 5]. This work was supported by the ERC-NERVI number 227747, KEOPS ANR-CONICYT and European Union Project # FP7-269921 (BrainScales), and Renvision # FP7-600847.

References

- [1] Shunichi Amari. Information geometry on hierarchy of probability distributions. *IEEE Transactions on Information Theory*, 47(5):1701–1711, 2001.
- [2] Adam L. Berger, Stephen A. Della Pietra, and Vincent J. Della Pietra. A maximum entropy approach to natural language processing. *Computational linguistics*, 22:39–71, 1996.
- [3] J. Besag. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, 36(2):192–236, 1974.
- [4] R. Bowen. *Equilibrium states and the ergodic theory of Anosov diffeomorphisms*, volume 470 of *Lect. Notes.in Math.* Springer-Verlag, New York, 1975.
- [5] Tamara Broderick, Miroslav Dudik, Gasper Tkacik, Robert E Schapire, and William Bialek. Faster solutions of the inverse pairwise ising problem. *arXiv preprint arXiv:0712.2437*, 2007.
- [6] Bruno Cessac and Rodrigo Cofre. Estimating maximum entropy distributions from periodic orbits in spike trains. research report RR-8329, INRIA, July 2013.

- [7] J.R. Chazottes and G. Keller. Pressure and equilibrium states in ergodic theory. *Israel Journal of Mathematics*, 131(1), 2008.
- [8] Stanley F. Chen and Ronald Rosenfeld. Efficient sampling and feature selection in whole sentence maximum entropy language models, 1999.
- [9] Michael Collins, Robert E. Schapire, and Yoram Singer. Logistic Regression, AdaBoost and Bregman Distances. *Machine Learning*, 48:253–285, 2002.
- [10] M. Dudík, S. Phillips, and R. Schapire. Performance guarantees for regularized maximum entropy density estimation. In *Proceedings of the 17th Annual Conference on Computational Learning Theory*, 2004.
- [11] Roberto Fernandez and Grégory Maillard. Chains with complete connections : General theory, uniqueness, loss of memory and mixing properties. *J. Stat. Phys.*, 118(3-4):555–588, 2005.
- [12] E. Ferrea, A Maccione, L Medrihan, T Nieu, D Ghezzi, P Baldelli, F Benfenati, and L. Berdondini. Large-scale, high-resolution electrophysiological imaging of field potentials in brain slices with microelectronic multielectrode arrays. *Frontiers in Neural Circuits.*, 6(80), 2012.
- [13] Elad Ganmor, Ronen Segev, and Elad Schneidman. The architecture of functional interaction networks in the retina. *The journal of neuroscience*, 31(8):3044–3054, 2011.
- [14] Elad Ganmor, Ronen Segev, and Elad Schneidman. Sparse low-order interaction network underlies a highly correlated and learnable neural population code. *PNAS*, 108(23):9679–9684, 2011.
- [15] U. Garibaldi and M. A. Penco. Probability theory and physics between bernoulli and laplace: The contribution of j. h. lambert (1728-1777). In *Proc. Fifth National Congress on the History of Physics*, volume 9, pages 341–346, 1985.
- [16] Hans-Otto Georgii. *Gibbs measures and phase transitions*. De Gruyter Studies in Mathematics:9. Berlin; New York, 1988.
- [17] I.I. Gikhman and A.V. Skorokhod. *The Theory of Stochastic Processes*. Springer, 1979.
- [18] Daniel N. Hill, Samar B. Mehta, and David Kleinfeld. Quality Metrics to Accompany Spike Sorting of Extracellular Signals. *The Journal of Neuroscience*, 31(24):8699–8705, 2011.
- [19] Edwin T Jaynes. Where do we stand on maximum entropy. *The maximum entropy formalism*, pages 15–118, 1978.
- [20] Edwin T Jaynes. The minimum entropy production principle. *Annual Review of Physical Chemistry*, 31(1):579–601, 1980.
- [21] E.T. Jaynes. Information theory and statistical mechanics. *Phys. Rev.*, 106:620, 1957.

- [22] ET Jaynes. Macroscopic prediction. In *Complex Systems - Operational Approaches in Neurobiology, Physics, and Computers*, pages 254–269. Springer, 1985.
- [23] H.J. Kappen and F.B. Rodriguez. Boltzmann machine learning using mean field theory and linear response correction. In *Advances in Neural Information Processing Systems*, volume 12, pages 280–286. The MIT Press, 1998.
- [24] G. Keller. *Equilibrium States in Ergodic Theory*. Cambridge University Press, 1998.
- [25] Rob Koeling. Chunking with maximum entropy models, 2000.
- [26] A. M. Litke, N. Bezayiff, E. J. Chichilnisky, W. Cunningham, W. Dabrowski, A. A. Grillo, M. Grivich, P. Grybos, P. Hottowy, S. Kachiguine, R. S. Kalmar, K. Mathieson, D. Petrusca, M. Rahman, and A. Sher. What does the eye tell the brain?: Development of a system for the large scale recording of retinal output activity. In *Nuclear Science Symposium Conference Record, 2003 IEEE*, volume 2, pages 951–955 Vol.2, October 2003.
- [27] O. Marre, D. Amodei, N. Deshmukh, K. Sadeghi, F. Soo, T.E Holy, and M.J. Berry II. Mapping a Complete Neural Population in the Retina. *The Journal of Neuroscience*, 43(32):14859–14873, October 2012.
- [28] O. Marre, S. El Boustani, Y. Frégnac, and A. Destexhe. Prediction of spatiotemporal patterns of neural activity from pairwise correlations. *Physical review letters*, 102(13), April 2009.
- [29] Volker Mayer and Mariusz Urbański. Thermodynamical formalism and multifractal analysis for meromorphic functions of finite order. *Memoirs of the American Mathematical Society*, 203(954), 2010.
- [30] Hiroyuki Nakahara and Shunichi Amari. Information-geometric decomposition in spike analysis. In *NIPS*, pages 253–260, 2001.
- [31] Hassan Nasser, Olivier Marre, and Bruno Cessac. Spatio-temporal spike train analysis for large scale networks using the maximum entropy principle and montecarlo method. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(03):P03006, 2013.
- [32] Moritz Otten and Gerhard Stock. Maximum caliber inference of nonequilibrium processes. *The Journal of Chemical Physics*, 133(3):–, 2010.
- [33] J W Pillow, J Shlens, L Paninski, A Sher, A M Litke, E J Chichilnisky, and E P Simoncelli. Spatio-temporal correlations and visual signaling in a complete neuronal population. *Nature*, 454(7206):995–999, Aug 2008.
- [34] R. Quian Quiroga, Z. Nadasdy, and Y. Ben-Shaul. Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural computation*, 16(8):1661–1687, aug 2004.

- [35] Ronald Rosenfeld, Jaime Carbonell, and Alexander Rudnicky. Adaptive statistical language modeling: A maximum entropy approach. Technical report, School of Computer Science, Carnegie Mellon University, 1994.
- [36] D. Ruelle. *Statistical Mechanics: Rigorous results*. Benjamin, New York, 1969.
- [37] D. Ruelle. *Thermodynamic formalism*. Addison-Wesley, Reading, Massachusetts, 1978.
- [38] Michael T. Schaub and Simon R. Schultz. The ising decoder: reading out the activity of large neural ensembles. *arXiv:1009.1828*, 2010.
- [39] E. Schneidman, M.J. Berry, R. Segev, and W. Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012, 2006.
- [40] Ian H. Stevenson and Konrad P. Kording. How advances in neural recording affect data analysis. *Nature Neuroscience*, 14(2):139–142, January 2011.
- [41] Aonan Tang, David Jackson, Jon Hobbs, Wei Chen, Jodi L. Smith, Hema Patel, Anita Prieto, Dumitru Petrusca, Matthew I. Grivich, Alexander Sher, Pawel Hottowy, Wladyslaw Dabrowski, Alan M. Litke, and John M. Beggs. A maximum entropy model applied to spatial and temporal correlations from cortical networks *In Vitro*. *The Journal of Neuroscience*, 28(2):505–518, January 2008.
- [42] G. Tkacik, O. Marre, T. Mora, D. Amodei, M.J. Berry 2nd, and W. Bialek. The simplest maximum entropy model for collective behavior in a neural network. *J Stat Mech*, page P03011, 2013.
- [43] G. Tkačik, Elad Schneidman, Michael J. Berry II, and William Bialek. Spin glass models for a network of real neurons. *arXiv preprint arXiv:0912.5409*, 2009.
- [44] Juan Carlos Vasquez, Olivier Marre, Adrian G Palacios, Michael J Berry, and Bruno Cessac. Gibbs distribution analysis of temporal correlation structure on multicell spike trains from retina ganglion cells. *J. Physiol. Paris*, 106(3-4):120–127, May 2012.
- [45] Yaqian Zhou and Lide Wu. A fast algorithm for feature selection in conditional maximum entropy modeling. In *in Proceedings of the EMNLP 2003*, pages 153–159, 2003.