C. R. Acad. Sci. Paris, Ser. I 352 (2014) 847–852

Contents lists available at ScienceDirect

C. R. Acad. Sci. Paris, Ser. I

www.sciencedirect.com

Dynamical systems/Probability theory

Asymptotic description of stochastic neural networks. II. Characterization of the limit law

Description asymptotique de réseaux de neurones stochastiques. II. Caractérisation de la loi limite

Olivier Faugeras, James Maclaurin

Inria Sophia-Antipolis Méditerranée, NeuroMathComp Group, France

A R T I C L E   I N F O

Article history:
Received 28 April 2014
Accepted after revision 27 August 2014
Available online 20 September 2014

Presented by the Editorial Board

A B S T R A C T

We continue the development, started in [8], of the asymptotic description of certain stochastic neural networks. We use the Large Deviation Principle (LDP) and the good rate function H announced there to prove that H has a unique minimum μ_e, a stationary measure on the set of trajectories $\mathcal{T}^{\mathbb{Z}}$. We characterize this measure by its two marginals, at time 0, and from time 1 to T. The second marginal is a stationary Gaussian measure. With an eye on applications, we show that its mean and covariance operator can be inductively computed. Finally, we use the LDP to establish various convergence results, averaged, and quenched.

R É S U M É

Nous prolongeons le développement, commencé en [8], de la description asymptotique de certains réseaux de neurones stochastiques. Nous utilisons le principe de grandes déviations (PGD) et la bonne fonction de taux H que nous y annoncions pour démontrer l'existence d'un unique minimimum, μ_e, de H, une mesure stationnaire sur l'ensemble $\mathcal{T}^{\mathbb{Z}}$ des trajectoires. Nous caractérisons cette mesure par ses deux maginales, à l'instant 0, et du temps 1 au temps T. La seconde marginale est une mesure gaussienne stationnaire. Avec un oeil sur les applications, nous montrons comment calculer de manière inductive sa moyenne et son opérateur de covariance. Nous montrons aussi comment utiliser le PGD pour établir des résultats de convergence en moyenne et presque sûrement.

Version française abrégée

Après avoir rappelé dans la section 2 les notations et le modèle de réseaux de neurones utilisés dans [8], nous montrons dans la proposition 3.1 et le théorème 3.2 que la bonne fonction de taux H du PGD annoncé dans cette publication admet

E-mail addresses: olivier.faugeras@inria.fr (O. Faugeras), james.maclaurin@inria.fr (J. Maclaurin).

http://dx.doi.org/10.1016/j.crma.2014.08.017

un minimum unique. Le théorème 3.2 fournit une méthode constructive de calcul effectif de la loi de ce minimum. Nous montrons enfin, dans la section 4, l'intérêt de ce minimum, qui apparaît comme la limite faible quand $n \to \infty$ (le nombre de neurones tend vers l'infini) de la loi Q^{V_n} du réseau moyennée par rapport aux poids synaptiques ; c'est un résultat en moyenne. Nous montrons aussi dans le corollaire 4.3 un résultat de convergence faible presque sûrement par rapport aux poids synaptiques, résultat intéressant d'un point de vue pratique, puisqu'il évite de prendre la moyenne par rapport à tous les réseaux. Le théorème 4.4 donne un résultat de convergence presque sûre de la mesure empirique vers le minimum de H.

1. Introduction

In [8] we started our asymptotic analysis of very large networks of neurons with correlated synaptic weights. We showed that the image Π^n of the averaged law Q^{V_n} through the empirical measure satisfied a large deviation principle with good rate function H. In the same article, we provided an analytical expression of this rate, which is a function of the spectral representation of certain Gaussian processes. In the next section, we recall some definitions given in [8].

2. Mathematical framework

We start by recalling the model of [8]. For some positive integer $n > 0$, we let $V_n = \{j \in \mathbb{Z} : |j| \le n\}$, and $|V_n| = 2n + 1$. The finite-size neural network below is indexed by points in V_n. We work in discrete time, over times $t \in \{0, 1, \ldots, T\}$, for some positive integer T. The state variable for each neuron is in $\mathbb{R}$, and the path space is $\mathcal{T} = \mathbb{R}^{T+1}$. We equip $\mathcal{T}$ with the Euclidean topology, $\mathcal{T}^{\mathbb{Z}}$ with the cylindrical topology, and denote the Borelian σ-algebra generated by this topology by $\mathcal{B}(\mathcal{T}^{\mathbb{Z}})$.

The equation describing the time variation of the membrane potential U^j of the jth neuron writes:

$$U_t^j = \gamma U_{t-1}^j + \sum_{i \in V_n} J_{ji}^n f\left(U_{t-1}^i\right) + \theta^j + B_{t-1}^j, \qquad U_0^j = u_0^j, \quad j \in V_n,\ t = 1, \ldots, T \tag{1}$$

$f : \mathbb{R} \to]0, 1[$ is a monotonically increasing Lipschitz continuous bijection. γ is in $[0, 1)$ and determines the time scale of the intrinsic dynamics of the neurons. The B_t^js are i.i.d. Gaussian random variables distributed as $\mathcal{N}_1(0, \sigma^2)$.[1] They represent the fluctuations of the neurons' membrane potentials. The θ^js are i.i.d. as $\mathcal{N}_1(\bar{\theta}, \theta^2)$. They are independent of the B_t^is and represent the current injected in the neurons. The u_0^js are i.i.d. random variables each governed by the law μ_I.

The J_{ij}^ns are the synaptic weights. J_{ij}^n represents the strength with which the 'presynaptic' neuron j influences the 'postsynaptic' neuron i. They arise from a stationary Gaussian random field specified by its mean and covariance function

$$\mathbb{E}\left[J_{ij}^n\right] = \frac{\bar{J}}{|V_n|}, \qquad \operatorname{cov}\left(J_{ij}^n J_{kl}^n\right) = \frac{1}{|V_n|} \Lambda\left((k - i) \bmod V_n, (l - j) \bmod V_n\right).$$

Λ is positive definite, let $\tilde{\Lambda}$ be the corresponding (positive) Fourier transform. We make the technical assumption that the summation over both indices of the series $(\Lambda(i, j))_{i,j \in \mathbb{Z}}$ is absolutely convergent to $\Lambda^{\mathrm{sum}} > 0$. We write $\Lambda_{\min}^{\mathrm{sum}} = \inf_{n>0} \sum_{j,k \in V_n} \Lambda(j, k)$ and assume that $\Lambda_{\min}^{\mathrm{sum}} > 0$.

We note J^n the $|V_n| \times |V_n|$ matrix of the synaptic weights, $J^n = (J_{ij}^n)_{i,j \in V_n}$.

The process (Y^j) defined by

$$Y_t^j = \gamma Y_{t-1}^j + \bar{\theta} + B_{t-1}^j, \quad j \in V_n,\ t = 1, \ldots, T, \qquad Y_0^j = u_0^j$$

is stationary and independent. Writing $v = \Psi(u)$, we define:

$$\begin{cases} v_0 = \Psi_0(u) = u_0 \\ v_s = \Psi_s(u) = u_s - \gamma u_{s-1} - \bar{\theta}, \quad s = 1, \cdots, T. \end{cases} \tag{2}$$

We extend Ψ to a mapping $\mathcal{T}^{\mathbb{Z}} \to \mathcal{T}^{\mathbb{Z}}$ componentwise. We now introduce some more notation.

For some topological space Ω equipped with its Borelian σ-algebra $\mathcal{B}(\Omega)$, we denote the set of all probability measures by $\mathcal{M}(\Omega)$. We equip $\mathcal{M}(\Omega)$ with the topology of weak convergence. For some $\mu \in \mathcal{M}(\mathcal{T}^{\mathbb{Z}})$ governing a process $(X^j)_{j \in \mathbb{Z}}$, we let $\mu^{V_n} \in \mathcal{M}(\mathcal{T}^{V_n})$ denote the marginal governing $(X^j)_{j \in V_n}$. For some $\mu \in \mathcal{M}(\mathcal{T}^{\mathbb{Z}})$ governing a process $(X^j)_{j \in \mathbb{Z}}$, we let $\mu^{V_n} \in \mathcal{M}(\mathcal{T}^{V_n})$ denote the marginal governing $(X^j)_{j \in V_n}$. For some $X \in \mathcal{T}$ and $0 \le a \le b \le T$, $X_{a,b}$ denotes the $b - a + 1$-dimensional subvector of X. We let $\mu_{a,b} \in \mathcal{M}(\mathcal{T}_{a,b}^{\mathbb{Z}})$ denote the marginal governing $(X_{a,b}^j)_{j \in \mathbb{Z}}$. For some $j \in \mathbb{Z}$, let the shift operator $\mathcal{S}^j : \mathcal{T}^{\mathbb{Z}} \to \mathcal{T}^{\mathbb{Z}}$ be $S(\omega)^k = \omega^{j+k}$. We let $\mathcal{M}_{\mathcal{S}}(\mathcal{T}^{\mathbb{Z}})$ be the set of all stationary probability measures μ on $(\mathcal{T}^{\mathbb{Z}}, \mathcal{B}(\mathcal{T}^{\mathbb{Z}}))$ such that for all $j \in \mathbb{Z}$, $\mu \circ (\mathcal{S}^j)^{-1} = \mu$.

We next introduce the following definitions.

[1] We note $\mathcal{N}_p(m, \Sigma)$ the law of the p-dimensional Gaussian variable with mean m and covariance matrix Σ.

Definition 2.1. For each measure $\mu \in \mathcal{M}(\mathcal{T}^{V_n})$ or $\mathcal{M}_{\mathcal{S}}(\mathcal{T}^{\mathbb{Z}})$, we define $\underline{\mu}$ to be $\mu \circ \Psi^{-1}$.

Definition 2.2. Let $\mathcal{E}_2$ be the subset of $\mathcal{M}_{\mathcal{S}}(\mathcal{T}^{\mathbb{Z}})$ defined by

$$\mathcal{E}_2 = \{\mu \in \mathcal{M}_{\mathcal{S}}(\mathcal{T}^{\mathbb{Z}}) \mid \mathbb{E}^{\underline{\mu}_{1,T}}[\|v^0\|^2] < \infty\}.$$

Let $p_n : \mathcal{T}^{V_n} \to \mathcal{T}^{\mathbb{Z}}$ be such that $p_n(\omega)^k = \omega^{k \bmod V_n}$. Here, and throughout the paper, we take $k \bmod V_n$ to be the element $l \in V_n$ such that $l = k \bmod |V_n|$. Define the process-level empirical measure $\hat{\mu}_n : \mathcal{T}^{V_n} \to \mathcal{M}_{\mathcal{S}}(\mathcal{T}^{\mathbb{Z}})$ as

$$\hat{\mu}_n(\omega) = \frac{1}{|V_n|} \sum_{k \in V_n} \delta_{S^k p_n(\omega)}. \tag{3}$$

We note $Q^{V_n}(J^n)$ the element of $\mathcal{M}(\mathcal{T}^{V_n})$ which is the law of the solution to (1) conditioned on J^n. We let $Q^{V_n} = \mathbb{E}^J[Q^{V_n}(J^n)]$ be the law averaged with respect to the weights. Finally we introduce the image law in terms of which the principal results of this paper are formulated.

Definition 2.3. Let $\Pi^n \in \mathcal{M}(\mathcal{M}_{\mathcal{S}}(\mathcal{T}^{\mathbb{Z}}))$ be the image law of Q^{V_n} through the function $\hat{\mu}_n : \mathcal{T}^{V_n} \to \mathcal{M}_{\mathcal{S}}(\mathcal{T}^{\mathbb{Z}})$ defined by (3):

$$\Pi^n = Q^{V_n} \circ \hat{\mu}_n^{-1}$$

3. Characterization of the unique minimum of the rate function

In [7], with each measure $\nu \in \mathcal{M}(\mathcal{T}^{\mathbb{Z}})$ we associate the measure, noted Q^ν of $\mathcal{M}(\mathcal{T}^{\mathbb{Z}})$ such that $\underline{Q}^\nu = \mu_I^{\mathbb{Z}} \otimes \underline{Q}^\nu_{1,T}$, where $\underline{Q}^\nu_{1,T}$ is a Gaussian measure on $\mathcal{T}^{\mathbb{Z}}_{1,T}$ with spectral density $\sigma^2\delta(\theta) + \tilde{K}^\nu(\theta)$ and mean c^ν. The spectral density $\tilde{K}^\nu$ and mean c^ν are defined in [8]. We also define the rate function H^ν, which uses a linear approximation of the functional Γ defined in [7] and satisfies the relation $H^\mu(\mu) = H(\mu)$. We prove the following lemma in [7].

Lemma 3.1. *For $\mu, \nu \in \mathcal{M}_{\mathcal{S}}(\mathcal{T}^{\mathbb{Z}})$, $H^\nu(\mu) = 0$ if and only if $\mu = Q^\nu$.*

As stated in the following proposition, there exists a unique minimum μ_e of the rate function. We provide explicit equations for μ_e which would facilitate its numerical simulation.

Proposition 3.1. *There is a unique distribution $\mu_e \in \mathcal{M}_{\mathcal{S}}(\mathcal{T}^{\mathbb{Z}})$ which minimizes H. This distribution satisfies $H(\mu_e) = 0$, which is equivalent to $\mu_e = Q^{\mu_e}$.*

Proof. The proof, which is found in [7], is an easy consequence of the explicit method we outline to actually calculate μ_e below in Theorem 3.2. $\square$

We characterize the unique measure μ_e such that $\mu_e = Q^{\mu_e}$ in terms of its image $\underline{\mu_e}$. This characterization allows one to directly numerically calculate μ_e. Since $\underline{\mu_e}$ is Gaussian, the problem becomes that of defining the latest entries (in time) of K^{μ_e} and c^{μ_e} in terms of previous ones. Hence we characterize $\underline{\mu_e}$ recursively (in time), by providing a method of determining $\underline{\mu_e}_{0,t}$ in terms of $\underline{\mu_e}_{0,t-1}$. Let $K^{\mu_e,l}_{(t-1,s-1)}$ be the $(t-1) \times (s-1)$ submatrix of $K^{\mu_e,l}$ composed of the rows from times 1 to $t-1$ and the columns from times 1 to $s-1$. Let the measure $\underline{\mu_e}^{V_0}_{0,t} \in \mathcal{M}(\mathcal{T}_{0,t})$ be given by

$$\underline{\mu_e}^{V_0}_{0,t}(\mathrm{d}v) = \mu_I(\mathrm{d}v_0) \otimes \mathcal{N}_t\big(c^{\mu_e}_{1,t}, \sigma^2\,\mathrm{Id}_t + K^{\mu_e,0}_{(t,t)}\big)\,\mathrm{d}v_{1,t},$$

and $\underline{\mu_e}^{(0,l)}_{(0,t),(0,s)} \in \mathcal{M}(\mathcal{T}_{0,t} \times \mathcal{T}_{0,s})$ be given by

$$\begin{aligned}\underline{\mu_e}^{(0,l)}_{(0,t),(0,s)}\big(\mathrm{d}v^0_{0,t}\mathrm{d}v^l_{0,s}\big) &= \mu_I\big(\mathrm{d}v^0_0\big) \otimes \mu_I\big(\mathrm{d}v^l_0\big) \\ &\quad \otimes \mathcal{N}_{t+s}\big(\big(c^{\mu_e}_{1,t}, c^{\mu_e}_{1,s}\big), \sigma^2\,\mathrm{Id}_{t+s} + K^{\mu_e,(0,l)}_{(t,s)}\big)\,\mathrm{d}v^0_{1,t}\,\mathrm{d}v^l_{1,s},\end{aligned}$$

where

$$K^{\mu_e,(0,l)}_{(t,s)} = \begin{bmatrix} K^{\mu_e,0}_{(t,t)} & K^{\mu_e,l}_{(t,s)} \\ {}^\dagger K^{\mu_e,l}_{(t,s)} & K^{\mu_e,0}_{(s,s)} \end{bmatrix},$$

and the † sign represents the transpose of a matrix or a vector. Observe that $\{c^{\mu_e}_s, K^{\mu_e,l}_{rs} | r, s \in [1,t], \}$ are sufficient to characterize $\underline{\mu_e}^{(0,l)}_{(0,t),(0,t)}$, since a Gaussian is entirely characterized by its mean and covariance. Furthermore, $\{\underline{\mu_e}^{(0,l)}_{(0,t),(0,t)} | l \in \mathbb{Z}\}$

are sufficient to characterize $\underline{\mu_e}_{0,t}$, because a Gaussian law is uniquely characterized by its pairwise marginals. Therefore the following theorem is sufficient for the inductive calculation of $\underline{\mu_e}$.

Theorem 3.2. *We may characterize $\underline{\mu_e}$ inductively as follows. Initially $\underline{\mu_e}_0 = \mu_I^{\mathbb{Z}}$. Given that we have a complete characterization of $\{\underline{\mu_e}^{(0,l)}_{(0,t-1),(0,t-1)}, \underline{\mu_e}^{V_0}_{0,t-1} : l \in \mathbb{Z}\}$, we may characterize $\{\underline{\mu_e}^{(0,l)}_{(0,t),(0,t)}, \underline{\mu_e}^{V_0}_{0,t} : l \in \mathbb{Z}\}$ according to the following identities. For $s \in [1,t]$,*

$$c_s^{\mu_e} = \bar{J} \int_{\mathbb{R}^s} f\left(\Psi^{-1}(v)_{s-1}\right) \underline{\mu_e}^{V_0}_{0,s-1}(\mathrm{d}v). \tag{4}$$

For $1 \le r, s \le t$, $K_{rs}^{\mu_e,k} = \theta^2 \delta_k 1_T{}^{\dagger} 1_T + \sum_{l=-\infty}^{\infty} \Lambda(k,l) M_{rs}^{\mu,l}$. Here, for $p = \max(r-1, s-1)$,

$$M_{rs}^{\mu_e,0} = \int_{\mathbb{R}^{p+1}} f\left(\Psi^{-1}(v)_{r-1}\right) \times f\left(\Psi^{-1}(v)_{s-1}\right) \underline{\mu_e}^{V_0}_{0,p}(\mathrm{d}v), \tag{5}$$

and for $l \neq 0$

$$M_{rs}^{\mu_e,l} = \int_{\mathbb{R}^r \times \mathbb{R}^s} f\left(\Psi^{-1}\left(v^0\right)_{r-1}\right) \times f\left(\Psi^{-1}\left(v^l\right)_{s-1}\right) \underline{\mu_e}^{(0,l)}_{(0,r-1),(0,s-1)}\left(\mathrm{d}v^0\, \mathrm{d}v^l\right). \tag{6}$$

Remark 1. From a practical point of view the s-dimensional integral in (4) and the $\max(r,s)$-dimensional integral in (5) can be reduced by a change of variable to at most two dimensions. Similarly, the $r+s$-dimensional integral in (6) can be reduced to at most four dimensions.

Remark 2. If we make the biologically realistic assumption that the synaptic weights are not correlated beyond a certain correlation distance $d \ge 0$, $\Lambda(k,l) = 0$ if k or l does not belong to V_d it is seen that the matrixes $K^{\mu_e,k}$ are 0 as soon as $k \notin V_d$: thus in this case the asymptotic description of the network of neurons is sparse.

4. Convergence results

We use the Large Deviation Principle proved in [8,7] to establish convergence results for the measures Π^n, Q^{V_n} and $Q^{V_n}(J^n)$.

Theorem 4.1. *Π^n converges weakly to δ_{μ_e}, i.e., for all $\Phi \in \mathcal{C}_b(\mathcal{M}_{\mathcal{S}}(\mathcal{T}^{\mathbb{Z}}))$,*

$$\lim_{n\to\infty} \int_{\mathcal{T}^{V_n}} \Phi\left(\hat{\mu}_n(u)\right) Q^{V_n}(\mathrm{d}u) = \Phi(\mu_e).$$

Similarly,

$$\lim_{n\to\infty} \int_{\mathcal{T}^{V_n}} \Phi\left(\hat{\mu}_n(u)\right) Q^{V_n}\left(J^n\right)(\mathrm{d}u) = \Phi(\mu_e) \quad J \text{ almost surely}$$

Proof. The proof of the first result follows directly from the existence of an LDP for the measure Π^n, see Theorem 3.1 in [8], and is a straightforward adaptation of the one in [10, Theorem 2.5.1]. The proof of the second result uses the same method, making use of Theorem 4.2 below. $\square$

We can in fact obtain the following quenched convergence analogue of the usual lower bound inequality in the definition of a Large Deviation Principle.

Theorem 4.2. *For each closed set F of $\mathcal{M}_{\mathcal{S}}(\mathcal{T}^{\mathbb{Z}})$ and for almost all J:*

$$\overline{\lim_{n\to\infty}} \frac{1}{|V_n|} \log\left[Q^{V_n}\left(J^n\right)(\hat{\mu}_n \in F)\right] \le - \inf_{\mu \in F} H(\mu).$$

Proof. The proof is a combination of Tchebyshev's inequality and of the Borel–Cantelli lemma and is an adaptation of the one in [10, Theorem 2.5.4, Corollary 2.5.6]. $\square$

We define $\check{Q}^{V_n}(J^n) = \frac{1}{|V_n|} \sum_{j \in V_n} Q^{V_n}(J^n) \circ \mathcal{S}^{-j}$, where we recall the shift operator $\mathcal{S}$. Clearly $\check{Q}^{V_n}(J^n)$ is in $\mathcal{M}_{\mathcal{S}}(\mathcal{T}^{V_n})$. We define $\check{Q}^{V_n}$ to be the expectation of $\check{Q}^{V_n}(J^n)$, with respect to the synaptic weights J.

Corollary 4.3. *Fix m and let $n > m$. For almost every J and all $h \in \mathcal{C}_b(\mathcal{T}^{V_m})$,*

$$\lim_{n\to\infty} \int_{\mathcal{T}^{V_m}} h(u)\, \check{Q}^{V_n,V_m}\big(J^n\big)(\mathrm{d}u) = \int_{\mathcal{T}^{V_m}} h(u)\, \mu_e^{V_m}(\mathrm{d}u),$$

$$\lim_{n\to\infty} \int_{\mathcal{T}^{V_m}} h(u)\, Q^{V_n,V_m}(\mathrm{d}u) = \int_{\mathcal{T}^{V_m}} h(u)\, \mu_e^{V_m}(\mathrm{d}u).$$

That is, the V_m^{th} marginals $\check{Q}^{V_n,V_m}(J^n)$ and Q^{V_n,V_m} of respectively $\check{Q}^{V_n}(J^N)$ and $\check{Q}^{V_n}$ converge weakly to $\mu_e^{V_m}$ as $n \to \infty$.

Proof. It is sufficient to apply Theorem 4.1 in the case where Φ in $\mathcal{C}_b(\mathcal{M}_{\mathcal{S}}(\mathcal{T}^{\mathbb{Z}}))$ is defined by

$$\Phi(\mu) = \int_{\mathcal{T}^{V_m}} h\, \mathrm{d}\mu^{V_m}$$

and to use the fact that $Q^{V_n}, \check{Q}^{V_n}(J) \in \mathcal{M}_{\mathcal{S}}(\mathcal{T}^{V_n})$. $\square$

We now prove the following ergodic-type theorem. We may represent the ambient probability space by $\mathfrak{W}$, where $\omega \in \mathfrak{W}$ is such that $\omega = (J_{ij}, B_t^j, u_0^j)$, where $i, j \in \mathbb{Z}$ and $0 \le t \le T-1$, recall (1). We denote the probability measure governing ω by $\mathfrak{P}$. Let $u^{(n)}(\omega) \in \mathcal{T}^{V_n}$ be defined by (1). As an aside, we may then understand $Q^{V_n}(J^n)$ to be the conditional law of $\mathfrak{P}$ on $u^{(n)}(\omega)$, for a given J^n.

Theorem 4.4. *Fix $m > 0$ and let $h \in C_b(\mathcal{T}^{V_m})$. For $u^{(n)}(\omega) \in \mathcal{T}^{V_n}$ (where $n > m$) $\mathfrak{P}$ almost surely,*

$$\lim_{n\to\infty} \frac{1}{|V_n|} \sum_{j\in V_n} h\big(\pi^{V_m}\big(\mathcal{S}^j u^{(n)}(\omega)\big)\big) = \int_{\mathcal{T}^{V_m}} h(u)\, \mathrm{d}\mu_e^{V_m}(u), \tag{7}$$

where π^{V_m} is the projection onto V_m. Hence $\hat{\mu}_n(u^{(n)}(\omega))$ converges $\mathfrak{P}$-almost-surely to μ_e.

Proof. Our proof is an adaptation of [10]. We may suppose without loss of generality that $\int_{\mathcal{T}^{V_m}} h(u)\mathrm{d}\mu_e^{V_m}(u) = 0$. For $p > 1$ let

$$F_p = \left\{\mu \in \mathcal{M}_{\mathcal{S}}\big(\mathcal{T}^{\mathbb{Z}}\big) \,\middle|\, \left|\int_{\mathcal{T}^{V_m}} h(u)\, \mu^{V_m}(\mathrm{d}u)\right| \ge \frac{1}{p}\right\}.$$

Since $\mu_e \notin F_p$, but it is the unique zero of H, it follows that $\inf_{F_p} H = m > 0$. Thus by Theorem 3.1 in [8], there exists an n_0, such that for all $n > n_0$,

$$Q^{V_n}(\hat{\mu}_n \in F_p) \le \exp\big(-m|V_n|\big).$$

However

$$\mathfrak{P}\big(\omega \mid \hat{\mu}_n\big(u^{(n)}(\omega)\big) \in F_p\big) = Q^{V_n}\big(u|\hat{\mu}_n(u) \in F_p\big).$$

Thus

$$\sum_{n=0}^{\infty} \mathfrak{P}\big(\omega \mid \hat{\mu}_n\big(u^{(n)}(\omega)\big) \in F_p\big) < \infty.$$

We may thus conclude from the Borel–Cantelli lemma that $\mathfrak{P}$ almost surely, for every $\omega \in \mathfrak{W}$, there exists n_p such that for all $n \ge n_p$,

$$\left|\frac{1}{|V_n|} \sum_{j\in V_n} h\big(\pi^{V_m}\big(\mathcal{S}^j u^{(n)}(\omega)\big)\big)\right| \le \frac{1}{p}.$$

This yields (7) because p is arbitrary. The convergence of $\hat{\mu}_n(u^{(n)}(\omega))$ is a direct consequence of (7), since this means that each of the V_m^{th} marginals converge. $\square$

5. Possible extensions and conclusion

This result can be straightforwardly extended to the case when the noise is correlated but stationary Gaussian, that is $\mathrm{cov}(B_s^j, B_t^k)$ is some function of s, t and $(k-j)$, see [6]. It can also be easily extended to the case that the initial distribution is correlated, but the spatial correlation satisfies a certain "mixing principle", which enforces a sufficiently rapid decay [4].

The hypothesis that the synaptic weights are Gaussian is somewhat unrealistic from the biological viewpoint. In his PhD thesis [10], Moynot has obtained some preliminary results in the case of uncorrelated weights. We think that this is also a promising avenue. Moynot again, in his thesis, has extended the uncorrelated weights case, to include two populations with different (Gaussian) statistics for each population. This is also an important practical problem in neuroscience. Finally, the extension of Moynot's result to the correlated case would probably constitute a low hanging fruit.

It would be of interest to compare our LDP with other analyses of the rate of convergence of neural networks to their limits as the size asymptotes to infinity. This includes the system-size expansion of Bressloff [1], the path-integral formulation of Buice and Cowan [2] and the systematic expansion of the moments by (amongst others) [9,5,3].

In recent years there has been a lot of effort to mathematically justify neural-field models, through some sort of asymptotic analysis of finite-size neural networks. Many, if not most, of these models assume/prove some sort of thermodynamic limit, whereby if one isolates a particular population of neurons in a localized area of space, they are found to fire increasingly asynchronously as the number in the population asymptotes to infinity.[2] Indeed this was the result of Moynot and Samuelides. However, our results imply that there are system-wide correlations between the neurons, even in the asymptotic limit. The key reason why we do not have propagation of chaos is that the Radon–Nikodym derivative $\frac{dQ^N}{dP^N}$ of the average laws in [8, Proposition 2.4] cannot be tensored into N i.i.d. processes, whereas the simpler assumptions on the weight function Λ in Moynot and Samuelides allow the Radon–Nikodym derivative to be tensored. A very important implication of our result is that the mean-field behavior is insufficient to characterize the behavior of a population. Our limit process μ_e is system-wide and ergodic. Our work challenges the assumption held by some that one cannot have a 'concise' macroscopic description of a neural network without an assumption of asynchronicity at the local population level.

Acknowledgements

This work was supported by INRIA FRM, ERC-NERVI number 227747, European Union Project # FP7-269921 (BrainScales), and Mathemacs # FP7-ICT-2011.9.7.

References

[1] P. Bressloff, Stochastic neural field theory and the system-size expansion, SIAM J. Appl. Math. 70 (2009) 1488–1521.
[2] M. Buice, J. Cowan, Field-theoretic approach to fluctuation effects in neural networks, Phys. Rev. E 75 (2007).
[3] M. Buice, J. Cowan, C. Chow, Systematic fluctuation expansion for neural network activity equations, Neural Comput. 22 (2010) 377–426.
[4] T. Chiyonobu, S. Kusuoka, The large deviation principle for hypermixing processes, Probab. Theory Relat. Fields 78 (1988) 627–649.
[5] S. ElBoustani, A. Destexhe, A master equation formalism for macroscopic modeling of asynchronous irregular activity states, Neural Comput. 21 (2009) 46–100.
[6] O. Faugeras, J. MacLaurin, Large deviations of an ergodic synchronous neural network with learning, arXiv depot, INRIA Sophia Antipolis, France, http://arxiv.org/abs/1404.0732, 2014.
[7] O. Faugeras, J. Maclaurin, Asymptotic description of neural networks with correlated synaptic weights, Rapport de recherche RR-8495, INRIA, March 2014.
[8] O. Faugeras, J. Maclaurin, Asymptotic description of stochastic neural networks. I. Existence of a large deviation principle, C. R. Acad. Sci. Paris, Ser. I 352 (10) (2014) 841–846.
[9] I. Ginzburg, H. Sompolinsky, Theory of correlations in stochastic neural networks, Phys. Rev. E 50 (1994).
[10] O. Moynot, Étude mathématique de la dynamique des réseaux neuronaux aléatoires récurrents, Ph.D. thesis, Université Paul-Sabatier, Toulouse, 1999.

[2] We noted in the introduction that this is termed propagation of chaos by some.