Neural networks do not become asynchronous in the thermodynamic limit: there is no propagation of chaos

Olivier Faugeras
Joint work with James Maclaurin and Etienne Tanré

Inria and
University of Utah

ICMNS 2018
Antibes, France

The question

Find concise mathematical descriptions of large networks of neurons

This talk

- Fully connected networks of rate neurons
- Random synaptic weights
- Average and almost sure results

The mathematical model

- Intrinsic dynamics:

$$\mathcal{S} := \begin{cases} dV_t & = & -\alpha V_t dt + \sigma dW_t, \ 0 \leq t \leq T \\ \text{Law of } V_0 & = & \mu_0, \end{cases}$$

- There is a unique strong solution to $\mathcal{S}$ (Ornstein-Uhlenbeck process):
- Note P its law on the set $\mathcal{T} := \mathcal{C}([0, T]; \mathbb{R})$ of trajectories
- For this talk: $\alpha = 0$.

The mathematical model

- N neurons, $N = 2n + 1$; completely connected network
- Coupled dynamics

$$\mathcal{S}^N(J_n) := \begin{cases} dV_t^i & = \sum_{j \in I_n} J_n^{ij} f(V_t^j) dt + \sigma dW_t^i \ \forall i \in I_n \\ \text{Law of } V_N(0) & = \mu_0^{\otimes N} \end{cases}$$

$i \in I_n := [-n, \cdots, n]$.

- f is bounded, Lipschitz continuous (usually a sigmoid), defining the firing rate
- W^i: independent Brownians: intrinsic noise on the membrane potentials

The mathematical model

- There is a unique solution to $\mathcal{S}^N(J_n)$
- Note $P(J_n)$ its law on the set $\mathcal{T}^N$ of N-tuples of trajectories.

Modeling the synaptic weights

- J_n^{ij}: stationary Gaussian field: random synaptic weights

$$\mathbb{E}[J_n^{ij}] = \frac{\overline{J}}{N} \text{ for this talk } \bar{J} = 0$$

$$cov(J_n^{ij} J_n^{kl}) = \frac{R_{\mathcal{J}}(k-i, l-j)}{N}$$

- $R_{\mathcal{J}}(k, l)$ is a covariance function.
- Analogy with random media

Consequences

- $P(J_n)$ is a random law on $\mathcal{T}^N$
- Consider the law $P^{\otimes N}$ of N independent uncoupled neurons
- Girsanov theorem allows us to compare the law of the solution to the coupled system, $P(J_n)$, with the law of the uncoupled system, $P^{\otimes N}$:

$$\frac{dP(J_n)}{dP^{\otimes N}} = \exp\left\{ \sum_{i \in I_n} \frac{1}{\sigma} \int_0^T \left(\sum_{j \in I_n} J_n^{ij} f(V_t^j) \right) dW_t^i - \frac{1}{2\sigma^2} \int_0^T \left(\sum_{j \in I_n} J_n^{ij} f(V_t^j) \right)^2 dt \right\}$$

Uncorrelated case

- Consider the empirical measure:

$$\hat{\mu}_{N,u}(V_N) = \frac{1}{N} \sum_{i \in I_n} \delta_{V^i},$$

$V_N = (V^{-n}, \cdots, V^n)$

- It defines the mapping

$$\hat{\mu}_{N,u} : \mathcal{T}^N \to \mathcal{P}(\mathcal{T})$$

Correlated case

- Consider the empirical measure

$$\hat{\mu}_{N,c}(V_N) = \frac{1}{N} \sum_{i \in I_n} \delta_{S^i(V_{N,p})},$$

a probability measure on $\mathcal{T}^{\mathbb{Z}}$.

- $V_{N,p}$ is the periodic extension of the finite sequence of trajectories $V_N = (V^{-n}, \cdots, V^n)$. $V_{N,p} = (\cdots, V_N, V_N, \cdots)$
- S is the shift operator acting on elements of $\mathcal{T}^{\mathbb{Z}}$.
- Case $N = 3$

$$V_{3,p} = (\cdots, V^{-1}, V^0, V^1, V^{-1}, V^0, V^1, \cdots)$$
$$S^1(V_{3,p}) = (\cdots, V^1, V^{-1}, V^0, V^1, V^{-1}, V^0, \cdots)$$
$$S^2(V_{3,p}) = (\cdots, V^0, V^1, V^{-1}, V^0, V^1, V^{-1}, \cdots)$$

- It defines the mapping

$$\hat{\mu}_{N,c}(V_N) : \mathcal{T}^N \to \mathcal{P}_S(\mathcal{T}^{\mathbb{Z}})$$

1) Metric on $\mathcal{T}^{\mathbb{Z}}$

$$d_T^{\mathbb{Z}}(u,v) = \sum_{i\in\mathbb{Z}} 2^{-|i|}(\|u^i - v^i\|_T \wedge 1)$$

where

$$\|u^i - v^i\|_T = \sup_{t\in[0,T]} |u_t^i - v_t^i|$$

2) Metric on $\mathcal{P}(\mathcal{T}^{\mathbb{Z}})$
Induced by the Wasserstein-1 distance:

$$D_T(\mu,\nu) = \inf_{\xi\in C(\mu,\nu)} \int d_T^{\mathbb{Z}}(u,v)\, d\xi(u,v)$$

- We are interested in the laws of $\hat{\mu}_{N,u}$ and $\hat{\mu}_{N,c}$ under $P(J_n)$
- Define

$$Q^N = \int_\Omega P(J_n(\omega))\, d\omega,$$

  the average of $P(J_n)$ w.r.t. to the "random medium", i.e. the synaptic weights.
- We study the law of $\hat{\mu}_{N,u}$ and $\hat{\mu}_{N,c}$ under Q^N: average results.

The strategy

- Consider the law Π_u^N of $\hat{\mu}_{N,u}$ under Q^N: it is a probability measure on $\mathcal{P}(\mathcal{T})$:

$$\Pi_u^N(B) = \left(Q^N \circ (\hat{\mu}_{N,u})^{-1}\right)(B) = Q^N(\hat{\mu}_{N,u} \in B),$$

  B measurable set of $\mathcal{P}(\mathcal{T})$

- Consider the law Π_c^N of $\hat{\mu}_{N,c}$ under Q^N: it is a probability measure on $\mathcal{P}_S(\mathcal{T}^{\mathbb{Z}})$:

$$\Pi_c^N(B) = Q^N(\hat{\mu}_{N,c} \in B),$$

  B measurable set of $\mathcal{P}_S(\mathcal{T}^{\mathbb{Z}})$

The strategy

- Establish a Large Deviation Principle for the sequences of probability measures $(\Pi_u^N)_{N\geq 1}$ and $(\Pi_c^N)_{N\geq 1}$, i.e.
- Design a rate function (non-negative lower semi-continuous) H_u (resp. H_c) on $\mathcal{P}(\mathcal{T})$ (resp. $\mathcal{P}(\mathcal{T}^{\mathbb{Z}})$)
- The intuitive meaning of H is the following

$$Q^N(\hat{\mu}_N \simeq Q) \simeq e^{-NH(Q)}$$

- The measures $\hat{\mu}_N$ concentrate on the measures Q such that $H(Q) = 0$.
- If H reaches 0 at a single measure Q then Π^N converges in law toward the Dirac mass δ_Q

Minimum of H_u

By adapting the results of Ben Arous and Guionnet [BAG95] and of Moynot and Samuelides [MS02] one obtains:

Theorem

$$H_u(\mu) = I^{(2)}(\mu; P) - \Gamma_u(\mu),$$

where $I^{(2)}(\mu; P)$ *is the relative entropy of* μ *w.r.t.* P $I^{(2)}(\mu; P) = \int \log \frac{d\mu}{dP} \, d\mu$, *and* Γ_u *is defined by*

$$\frac{dQ^N}{dP^{\otimes N}} = e^{N\Gamma_u(\hat{\mu}_{N,u})}$$

H_u *achieves its minimum at a unique point* μ_u *of* $\mathcal{P}(\mathcal{T})$.

Minimum of H_u

and

Theorem

μ_u is the law of the solution to a linear non-Markovian stochastic system.

Average results

Two main results:

Theorem (1)

The law Π_u^N of the empirical measure $\hat{\mu}_{N,u}$ under Q^N converges weakly to δ_{μ_u}

This means that

$$\forall F \in C_b(\mathcal{P}(\mathcal{T}))$$

$$\lim_{N\to\infty} \int_\Omega \left(\int_{\mathcal{T}^N} F\left(\frac{1}{N} \sum_1^N \delta_{v^i} \right) P(J_n(\omega))(dv_N) \right) d\gamma(\omega) = F(\mu_u)$$

Average results

Theorem (2)

Q^N is μ_u-chaotic.

i.e. for all $m \geq 2$ and f_i, $i = 1, \ldots, m$ in $C_b(\mathcal{T})$

$$\lim_{N\to\infty} \int_{\mathcal{T}^N} f_1(v^1)\cdots f_m(v^m)\, dQ^N(v^1, \cdots, v^N) = \prod_{i=1}^{m} \int_{\mathcal{T}} f_i(v)\, d\mu_u(v)$$

"In the thermodynamic limit ($N \to \infty$) and on average, the neurons in any finite-size group become independent. One observes the propagation of chaos. The neurons become asynchronous."

Strategy in the correlated case: exponential approximation

1. Note, e.g. [Ell85], that the sequence $\Pi_0^N = P^{\otimes N} \circ (\hat{\mu}_{N,c})^{-1}$ satisfies the LDP with good rate function

$$I^{(3)}(\mu; P^{\mathbb{Z}}) = \lim_{N \to \infty} \frac{1}{N} I^{(2)}(\mu^N; P^{\otimes N})$$

   with

$$I^{(2)}(\mu^N; P^{\otimes N}) = \int \log \frac{d\mu^N}{dP^{\otimes N}} \, d\mu^N$$

2. If there existed a continuous function $\Psi : \mathcal{P}(\mathcal{T}^{\mathbb{Z}}) \to \mathcal{P}(\mathcal{T}^{\mathbb{Z}})$ such that

$$\Psi(\hat{\mu}_N(W_N)) = \hat{\mu}_N(V_N)$$

   then we would be done.

Strategy in the correlated case: exponential approximation

1. Show that there exists a sequence Ψ^m of continuous functions $\mathcal{P}_S(\mathcal{T}^{\mathbb{Z}}) \to \mathcal{P}_S(\mathcal{T}^{\mathbb{Z}})$ and a measurable map $\Psi^\infty : \mathcal{P}_S(\mathcal{T}^{\mathbb{Z}}) \to \mathcal{P}_S(\mathcal{T}^{\mathbb{Z}})$ such that for every $\alpha < \infty$

$$\limsup_{m\to\infty} \sup_{\mu: I^{(3)}(\mu)\le\alpha} D_T(\Psi^m(\mu), \Psi^\infty(\mu)) = 0$$

2. Show that the family $\Pi_0^N \circ (\Psi^m)^{-1}$ is an exponentially good approximation of the family Π_c^N,
3. and conclude from a general theorem (e.g. [DZ97]) that Π_c^N satisfies the LDP with good rate function

$$H_c(\mu) = \inf_\nu \left\{ I^{(3)}(\nu) : \mu = \Psi^\infty(\nu) \right\}$$

Definition of Ψ^m

- Note that

$$\left.\frac{dQ^N}{dP^{\otimes N}}\right|_{\mathcal{F}_t} = \exp\left(\sum_{j\in I_n}\int_0^t \theta_s^j dW_s^j - \frac{1}{2}\sum_{j\in I_n}\int_0^t \left(\theta_s^j\right)^2 ds\right)$$

where

$$\theta_t^j = \frac{1}{\sigma^2}\mathbb{E}^{\bar{\gamma}_t^{\hat{\mu}_{N,c}(V_N)}}\left[\sum_{k\in I_n} G_t^j \int_0^t G_s^k dW_s^k\right]$$

Definition of Ψ^m

$$G_t^i = \sum_{j \in I_n} J_n^{ij} f(V_t^j),\ i \in I_n$$

It can be verified that the covariance is entirely determined by the empirical measure

$$\mathbb{E}\left[G_t^i G_s^k\right] = \int_\Omega G_t^i(\omega) G_s^k(\omega)\, d\gamma(\omega) =$$

$$\sum_{l \in I_n} R_{\mathcal{J}}((k-i) \bmod I_n, l) \mathbb{E}^{\hat{\mu}_{N,c}}\left[f(v_t^0) f(v_s^l)\right]$$

Definition of Ψ^m

- We note $\gamma^{\hat{\mu}_{N,c}}$ the probability under which the Gs have the above covariance.

We also introduce

$$\Lambda_t := \frac{\exp\left\{-\frac{1}{2\sigma^2}\sum_{i\in I_n}\int_0^t \left(G_s^i\right)^2\,ds\right\}}{\mathbb{E}^{\gamma^{\hat{\mu}_{N,c}(V_N)}}\left[\exp\left\{-\frac{1}{2\sigma^2}\sum_{i\in I_n}\int_0^t \left(G_s^i\right)^2\,ds\right\}\right]},$$

and define the new probability law

$$\bar{\gamma}_t^{\hat{\mu}_{N,c}(V_N)} := \Lambda_t \cdot \gamma^{\hat{\mu}_{N,c}(V_N)}.$$

Definition of Ψ^m

- Consider the SDE

$$Z_t^j = W_t^j + \sigma^{-2} \sum_{k \in I_n} \int_0^t \mathbb{E}^{\bar{\gamma}_t^{\hat{\mu}_{N,c}(Z)}} \left[G_s^j \int_0^s G_u^k dZ_u^k \right] ds,$$

  $j \in I_n$, it follows from the above that the law of $\hat{\mu}_{N,c}(Z_N)$ is Π_c^N.
- Construct Ψ^m by time discretization (T/m) and space truncation (q_m).

Exponential approximation of Π_c^N by $\Pi_0^N \circ (\Psi^m)^{-1}$

We prove

Lemma (O.F., J. Maclaurin, E. Tanré)

For any $\delta > 0$,

$$\lim_{m\to\infty} \overline{\lim_{n\to\infty}} \frac{1}{N} \log Q^N\left(D_T\left(\Psi^m\left(\hat{\mu}_{N,c}(W_N)\right), \hat{\mu}_{N,c}(V_N)\right) > \delta \right) = -\infty.$$

Consequences

- Π_c^N satisfies an LDP with good rate function

$$H_c(\mu) = \inf_{\nu} \left\{ I^{(3)}(\nu) : \mu = \Psi^{\infty}(\nu) \right\}$$

- Π_c^N converges in law toward δ_{μ_c}, $\mu_c \in \mathcal{P}_S(\mathcal{T}^{\mathbb{Z}})$.
- μ_c is the limit law of Q^N, the averaged law of the finite size system.

Summary

Theorem (O.F., J. Maclaurin, E. Tanré)

H_c achieves its minimum at a unique point μ_c of $\mathcal{P}_S(\mathcal{T}^{\mathbb{Z}})$.

and

Theorem (O.F., J. Maclaurin, E. Tanré)

μ_c is the law of the solution to an infinite dimensional linear non-Markovian stochastic system, hence it is a Gaussian measure (in $\mathcal{P}_S(\mathcal{T}^{\mathbb{Z}})$) if the initial condition is Gaussian.

$$V_t^j = V_0^j + \sigma W_t^j + \sigma \int_0^t \theta_s^j ds$$

$$\theta_t^j = \sigma^{-3} \sum_{i \in \mathbb{Z}} \int_0^t L_{\mu_c}^{t,i-j}(t,s) dV_s^i, \; j \in \mathbb{Z}$$

$$\mathrm{Law}(V) = \mu_c$$

Definition of L_{μ_c}

The covariance operator: $\mu \in \mathcal{P}_S(\mathcal{T}^{\mathbb{Z}})$

$$K_\mu^k(t,s) = \sum_{j\in\mathbb{Z}} R_{\mathcal{J}}(k,l)\mathbb{E}^\mu\left[f(v_t^0)f(v_s^l)\right]$$

defines an operator $\bar{K}_\mu : L^1(\mathbb{Z}\times[0,T]) \to L^1(\mathbb{Z}\times[0,T])$

$$f \in L^1(\mathbb{Z}\times[0,T]) \to (\bar{K}_\mu f)_t^k = \sum_{l\in\mathbb{Z}} \int_0^T K_\mu^{k-l}(t,s) f_s^l \, ds$$

Definition of $\bar{L}_\mu$:

$$\sigma^2 \bar{L}_\mu = \mathrm{Id} - (\mathrm{Id} + \sigma^{-2}\bar{K}_\mu)^{-1}$$

The average results

$g \in C_b(\mathcal{T}^M)$, $M \geq 1$

$$\lim_{N \to \infty} \int_\Omega \left(\int_{\mathcal{T}^N} \frac{1}{N} \left(\sum_1^N g(S^i v_{N,p}) \right) P(J^N(\omega))(dv_N) \right) d\gamma(\omega) = \int_{\mathcal{T}^M} g(v)\, d\mu_c^M(v),$$

where μ_c^M is the Mth-order marginal of μ_c.

The almost sure results

- The existence of an LDP for the annealed law of the empirical measure μ_c implies that "half" the same principle applies to the quenched law.
- This implies that the law of the empirical measure converges exponentially fast to δ_{μ_c} for almost all choices of the weights and therefore

$g \in C_b(\mathcal{T}^M)$, $M \geq 1$, for almost all ω:

$$\lim_{N\to\infty} \int_{\mathcal{T}^N} \frac{1}{N} \left(\sum_1^N g(S^i v_{N,p}) \right) P(J^N(\omega))(dv_N) = \int_{\mathcal{T}^M} g(v)\, d\mu_c^M(v),$$

where μ_c^M is the Mth-order marginal of μ_c.

Numerical results: uncorrelated synaptic weights

- When

$$R_J(k,l) = R\ \delta_{kl}$$

- We are in the uncorrelated case studied by Sompolinsky et al. [SCS88]: propagation of chaos

$K^0(t,s)$

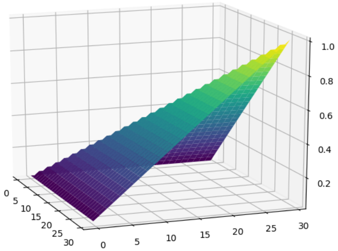

$K^1(t,s)$

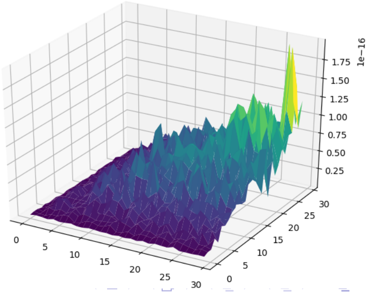

Numerical results: uncorrelated synaptic weights

Some trajectories

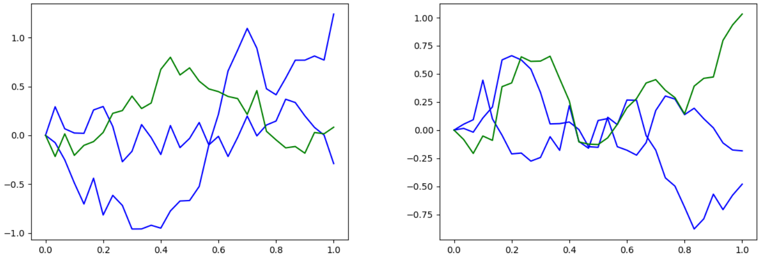

Numerical results: correlated synaptic weights

- When

$$R_J(k,l) = Q_J(k) \times Q_J(l), \quad Q_j = [1/2,\ 2.0,\ 1/2]$$

- No propagation of chaos

$$K^0(t,s)$$

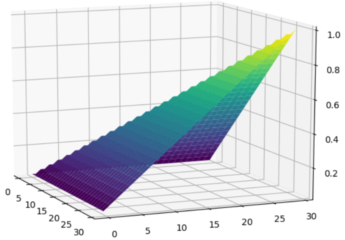

Numerical results: correlated synaptic weights

$K^1(t, s)$, scale values (0.1, 0.5, 1.0, 5.0)

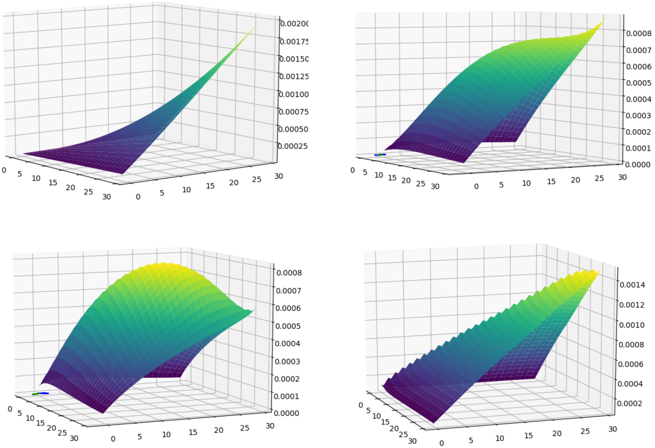

Numerical results: correlated synaptic weights

Some trajectories

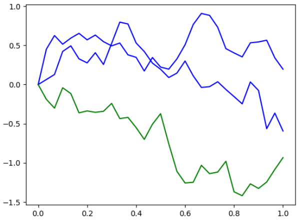

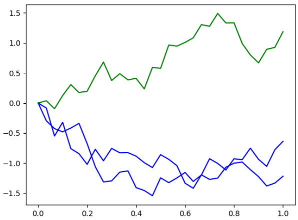

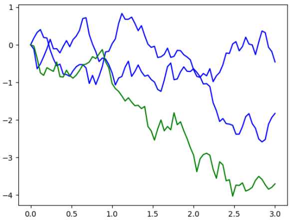

Summary

- We have started the analysis of the thermodynamic limit of completely connected networks of rate neurons in the case of correlated synaptic weights.
- In the uncorrelated case the network becomes asynchronous (propagation of chaos).
- In the correlated case there is no propagation of chaos and the neurons behaviours are completely different from those of the uncorrelated case.
- In both cases (uncorrelated and correlated synaptic weights) the thermodynamic limit is described by a Gaussian process if the initial conditions are Gaussian.

References I

- G. Ben-Arous and A. Guionnet, *Large deviations for langevin spin glass dynamics*, Probability Theory and Related Fields **102** (1995), no. 4, 455–509.
- Mireille Bossy, Olivier Faugeras, and Denis Talay, *Clarification and complement to "mean-field description and propagation of chaos in networks of Hodgkin–Huxley and FitzHugh–Nagumo neurons"*, The Journal of Mathematical Neuroscience (JMN) **5** (2015), no. 19.
- A. Dembo and O. Zeitouni, *Large deviations techniques*, Springer, 1997, 2nd Edition.
- R.S. Ellis, *Entropy, large deviations and statistical mechanics*, Springer, 1985.

References II

- Nicolas Fournier and Eva Löcherbach, *On a toy model of interacting neurons*, Ann. Inst. H. Poincaré Probab. Statist. **52** (2016), no. 4, 1844–1876.
- Eric Luçon and Wilhelm Stannat, *Mean field limit for disordered diffusions with singular interactions*, Ann. Appl. Probab. **24** (2014), no. 5, 1946–1993.
- O. Moynot and M. Samuelides, *Large deviations and mean-field theory for asymmetric random recurrent neural networks*, Probability Theory and Related Fields **123** (2002), no. 1, 41–75.

References III

H. Sompolinsky, A. Crisanti, and HJ Sommers, *Chaos in Random Neural Networks*, Physical Review Letters **61** (1988), no. 3, 259–262.

Large deviation principle: I

For all open sets $\mathcal{O}$ of $\mathcal{P}(\mathcal{T})$

$$-\inf_{\mu \in \mathcal{O}} H(\mu) \leq \liminf_{N \to \infty} \frac{1}{N} \log \Pi^N(\mathcal{O})$$

Large deviation principle: II

The sequence Π^N is exponentially tight.

Large deviation principle: III

For every compact set F of $\mathcal{P}(\mathcal{T})$

$$\limsup_{N\to\infty} \frac{1}{N} \log \Pi^N(F) \leq - \inf_{\mu \in F} H(\mu)$$

Exponential approximation

for all $\delta > 0$

$$\lim_{m\to\infty} \overline{\lim_{n\to\infty}} \frac{1}{N} \log P^{\otimes N}\left(D_T\left(\Psi_m\left(\hat{\mu}_c^N(B)\right), \hat{\mu}_c^N(Z)\right) > \delta \right) = -\infty$$