Two or three things I know about mean field methods in neuroscience

Olivier Faugeras

NeuroMathComp Laboratory - INRIA Sophia/ENS Paris/UNSA LJAD

CIRM Workshop on
Mean-field methods and multiscale analysis of neuronal populations

Joint work with

- Javier Baladron
- Bruno Cessac
- Diego Fasoli
- Jonathan Touboul

The neuronal activity in V1: from Ecker et al., Science 2010

- Recording neurons in V1



B



A



The neuronal activity in V1: from Ecker et al., Science 2010

- shows that their activity is HIGHLY **decorrelated** for synthetic



The neuronal activity in V1: from Ecker et al., Science 2010



- and natural images

The neuronal activity in V1: from Ecker et al., Science 2010



- as opposed to the current consensus.

The neuronal activity in V1: from Ecker et al., Science 2010

- Is this a network effect?
- Is this related to the stochastic nature of neuronal computation?

Spin glasses

- N spins x^i in interaction in a potential U (keeps spin values bounded).
- Weights of interaction: J_{ij}. Assume $J_{ij} \neq J_{ji}$.
- Weights are i.i.d. $\mathcal{N}(0,1)$.

Single spin dynamics:

$$\begin{cases} \dot{x}^i & = & -\nabla U(x^i) + \frac{\beta}{\sqrt{N}} \sum_j J_{ij} x^j + \xi^i \\ \text{Law of } x_0 & = & \mu_0^{\otimes N} \end{cases}$$

- Limit, when $N \to \infty$, of the dynamics?

Which limit?

- Let $P_\beta^N(\mathbf{J})$ be the law of the solution to the N spin equations.
- If we **anneal** it by taking the expectation over the weights: $Q_\beta^N = \mathbb{E}\left[P_\beta^N(\mathbf{J})\right]$ we can obtain two theorems.

Theorem (Ben Arous-Guionnet)

The law of the empirical measure $\hat{\mu}^N = \frac{1}{N}\sum_{i=1}^N \delta_{x^i}$ *under* Q_β^N *converges to* δ_Q.

Theorem (Ben Arous-Guionnet)

Q is the law of the solution to the following nonlinear stochastic differential equation:

$$\begin{cases} dx_t & = & -\nabla U(x_t)dt + dB_t \\ dB_t & = & dW_t + \beta^2 \left(\int_0^t f(t,s)\, dB_s\right)\, dt \\ \textit{Law of } x_0 & = & \mu_0 \end{cases}$$

Which limit?

W is a Q-Brownian motion, the function f is given by

$$f(t,s) = \mathbb{E}\left[G_s^Q G_t^Q \frac{\exp\left\{-\beta^2 \int_0^t \left(G_u^Q\right)^2 \, du\right\}}{\mathbb{E}\left[\exp\left\{-\beta^2 \int_0^t \left(G_u^Q\right)^2 \, du\right\}\right]} \right],$$

and G_t^Q is a centered Gaussian process, independent of Q, and with the same covariance:

$$\mathbb{E}\left[G_s^Q G_t^Q\right] = \int x_s x_t \, dQ(x)$$

The results of Sompolinsky and Zippelius

- S.-Z. studied the same spin-glass equation, up to minor details.

Theorem (Sompolinsky-Zippelius)

The annealed mean-field equation in the thermodynamic limit is

$$\begin{cases} dx_t & = & -\nabla U(x_t)dt + \Phi_t^x dt \\ \textit{Law of } x_0 & = & \mu_0 \end{cases}$$

Φ_t^x *is a Gaussian process with zero mean and whose autocorrelation* C *writes*

$$C(t,s) \equiv \mathbb{E}\left[\Phi_s^x \Phi_t^x\right] = \delta(t-s) + \beta^2 \int x_s x_t \, dQ(x) =$$

$$\delta(t-s) + \beta^2 \mathbb{E}\left[G_s^Q G_t^Q\right]$$

Are these two results the same?

Proposition (Faugeras)

If the function f in the Ben Arous-Guionnet theorem is continuous in $[0, T]^2$, the stochastic differential equation

$$\begin{cases} dB_t &= dW_t + \beta^2 \left(\int_0^t f(t,s)\, dB_s \right) dt \\ B_0 &= 0 \end{cases}$$

has a unique solution defined for all $t \in [0, T]$ by

$$dB_t = dW_t + \left(\int_0^t \Gamma(t,s)\, dW_s \right) dt,$$

where

$$\Gamma(t,s) = \sum_{i=1}^{\infty} g_i(t,s),$$

and

$$g_{n+1}(t,s) = \beta^2 \int_s^t f(t,\tau) g_n(\tau,s)\, d\tau \quad n \geq 1, \quad g_1 = \beta^2 f$$

Are these two results the same?

- Rewrite the Ben Arous-Guionnet mean-field equation as

$$\begin{cases} dx_t & = & -\nabla U(x_t)dt + \Psi_t^x dt \\ \text{Law of } x_0 & = & \mu_0 \end{cases},$$

where

$$\Psi_t^x = \frac{dW_t}{dt} + \int_0^t \Gamma(t,u)\, dW_u$$

- The process $\frac{dW_t}{dt}$ is interpreted as "Gaussian white noise".
- Ψ^x is a Gaussian process with zero mean and autocorrelation

$$\mathbb{E}\left[\Psi_t^x \Psi_s^x\right] = \delta(t-s) + \int_0^{t\wedge s} \Gamma(t,u)\Gamma(s,u)\, du$$

- Since in general

$$\int_0^{t\wedge s} \Gamma(t,u)\Gamma(s,u)\, du \neq \beta^2 \mathbb{E}\left[G_s^Q G_t^Q\right],$$

the two results may be contradictory!

From spin glasses to firing rate neurons

- In 1988, Sompolinsky, Crisanti and Sommers generalized the spin glass equations to firing rate neurons:

$$\begin{cases} \dot{x}^i & = & -\nabla U(x^i) + \frac{\beta}{\sqrt{N}} \sum_j J_{ij} S(x^j) + \xi^i \\ \text{Law of } x_0 & = & \mu_0^{\otimes N} \end{cases}$$

- S is the "usual" sigmoid function.
- They proposed the following mean-field equation:

$$\begin{cases} dx_t & = & -\nabla U(x_t)dt + \Phi_t^x dt + I dt \\ \text{Law of } x_0 & = & \mu_0 \end{cases}$$

Φ_t^x is a zero mean Gaussian process whose autocorrelation C writes

$$C(t,s) \equiv \mathbb{E}\left[\Phi_s^x \Phi_t^x\right] = \delta(t-s) + \beta^2 \int S(x_s) S(x_t)\, dQ(x) =$$

$$\delta(t-s) + \beta^2 \mathbb{E}\left[S(G_s^Q) S(G_t^Q)\right]$$

From spin glasses to firing rate neurons

- In 2009, Faugeras, Touboul and Cessac generalized the S.-C.-S. equation to the case of several populations.
- The weights are i.i.d. and $J_{ij} \approx \mathcal{N}\left(\frac{J_{\alpha\beta}}{N_\beta}, \frac{\sigma_{\alpha\beta}}{\sqrt{N_\beta}}\right)$
- They proposed an annealed mean-field equation inspired from that of S.-C.-S. and proved the equation had a unique solution in finite time.
- The solution is Gaussian, but non-Markov.
- The mean satisfies a first-order differential equation.
- The covariance function satisfies an integral equation.
- Both equations are nonlinearly coupled.
- Studying the solution turned out to be a formidable task (see part of Geoffroy Hermann's PhD thesis)
- From the discussion on spin glasses one may wonder whether this equation is the "correct" one.

From spin glasses to firing rate neurons

The mean process is coupled through the variance $C(t,t)$

$$\frac{d\mu(t)}{dt} = -\frac{\mu(t)}{\tau} + \bar{J}\int_{\mathbb{R}} S\left(x\sqrt{C(t,t)} + \mu(t)\right)\, Dx + I(t)$$

From spin glasses to firing rate neurons

The mean process is **coupled** through the variance $C(t,t)$

$$\frac{d\mu(t)}{dt} = -\frac{\mu(t)}{\tau} + \bar{J}\int_{\mathbb{R}} S\left(x\sqrt{C(t,t)} + \mu(t)\right)\, Dx + I(t)$$

From spin glasses to firing rate neurons

To the non-Markovian covariance

$$C(t,s) = e^{-(t+s)/\tau}\Big[C(0,0) + \frac{\sigma_{\text{ext}}^2}{2}\left(e^{2(t\wedge s)/\tau} - 1\right) + \bar{J}\int_0^t\int_0^s e^{(u+v)/\tau}\Delta(u,v)dudv\Big]$$

From spin glasses to firing rate neurons

To the **non-Markovian** covariance

$$C(t,s) = e^{-(t+s)/\tau}\Big[C(0,0) + \frac{\sigma_{\text{ext}}^2}{2}\left(e^{2(t\wedge s)/\tau} - 1\right) + \bar{J}\int_0^t\int_0^s e^{(u+v)/\tau}\Delta(u,v)\,du\,dv\Big]$$

From spin glasses to firing rate neurons

To the **non-Markovian** covariance

$$C(t,s) = e^{-(t+s)/\tau}\Big[C(0,0) + \frac{\sigma_{\text{ext}}^2}{2}\left(e^{2(t\wedge s)/\tau} - 1\right) + \bar{J}\int_0^t\int_0^s e^{(u+v)/\tau}\Delta(u,v)dudv\Big]$$

$$\Delta(t,s) = \mathbb{E}\left[S(x_t)S(x_s)\right]$$

Questions, open problems

- The neuron models are deceptively simple: can we do better?
- The synaptic connection models are deceptively simple: can we improve them?
- The completely connected graph model is too restrictive: can we develop a theory for different graphs?
- The i.i.d. model for the synaptic weights is not compatible with biological evidence: can we include correlations?

Networks of continuous spiking neurons

- Hodgkin-Huxley model or one of its 2D reductions.

Networks of continuous spiking neurons

- Hodgkin-Huxley model or one of its 2D reductions.
- Chemical and electric noisy synapses (conductance-based models)

Networks of continuous spiking neurons

- Hodgkin-Huxley model or one of its 2D reductions.
- Chemical and electric noisy synapses (conductance-based models)
- Synaptic weights are **dynamically** changing over time.

Fitzhugh Nagumo model

Stochastic Differential Equation (SDE):

$$\begin{cases} dV_t = \left(V_t - \frac{V_t^3}{3} - w_t + I^{\mathrm{ext}}(t)\right) dt + \sigma_{\mathrm{ext}}\, dW_t \\ dw_t = a\left(b\, V_t - w_t\right) dt \end{cases}$$

Takes into account external current noise.

Synapses

Synaptic current from the jth to the ith neuron:

$$I_{ij}^{\mathrm{syn}} = g_{ij}(t)(V^i - V_{\mathrm{rev}}^i)$$

Chemical conductance:

$$g_{ij}(t) = J_{ij}(t) y^j(t)$$

The function y denotes the fraction of open channels:

$$\dot{y}^j(t) = a_r^j S_j(V^j)(1 - y^j(t)) - a_d^j y^j(t),$$

The function S:

$$S(V^j) = \frac{T_{\max}}{1 + e^{-\lambda(V^j - V_T)}}$$

Synapses

Taking noise into account:

$$dy_t^j = \left(a_r S(V^j)(1 - y^j(t)) - a_d y^j(t)\right)\, dt + \sigma(V^j, y^j)\, dW_t^{j,y}$$

Keeping y^j between 0 and 1:

$$\sigma(V^j, y^j) = \sqrt{a_r S(V^j)(1 - y^j) + a_d y^j}\,\chi(y^j)$$

Synaptic weights

The synaptic weights are affected by dynamical random variations:

$$J_{ij}(t) = \frac{\bar{J}}{N} + \frac{\sigma}{N}\xi^i(t) = \frac{\bar{J}}{N} + \frac{\sigma}{N}\frac{dB_t^i}{dt},$$

Advantage : simplicity

Synaptic weights

The synaptic weights are affected by dynamical random variations:

$$J_{ij}(t) = \frac{\bar{J}}{N} + \frac{\sigma}{N}\xi^i(t) = \frac{\bar{J}}{N} + \frac{\sigma}{N}\frac{dB_t^i}{dt},$$

Advantage : simplicity

Disadvantage : an increase of the noise level increases the probability that the sign of $J_{ij}(t)$ is different from that of $\bar{J}$.
It can be fixed (technical)

Putting everything together

Each neuron is represented by a state vector of dimension 3:

$$\begin{cases} dV_t^i &= \left(V_t^i - \frac{(V_t^i)^3}{3} - w_t^i + I(t)\right) dt + \\ & \quad \left(\frac{1}{N}\sum_j \bar{J}(V_t^i - V_{\mathrm{rev}}) y_t^j\right) dt + \\ & \quad \frac{1}{N}\left(\sum_j \sigma(V_t^i - V_{\mathrm{rev}}) y_t^j\right) dB_t^i + \\ & \quad \sigma_{\mathrm{ext}}\, dW_t^i \\ dw_t^i &= a\left(b\, V_t^i - w_t^i\right) dt \\ dy_t^i &= \left(a_r S(V_t^i)(1 - y_t^i) - a_d y_t^i\right) dt + \sigma(V_t^i, y_t^i) dW_t^{i,y} \end{cases}$$

Putting everything together

The full dynamics of the neuron i can be described compactly by

$$dX_t^i = f(t, X_t^i)\, dt + g(t, X_t^i) \begin{bmatrix} dW_t^i \\ dW_t^{i,y} \end{bmatrix} + \frac{1}{N} \sum_j b(X_t^i, X_t^j)\, dt + \frac{1}{N} \sum_j \beta(X_t^i, X_t^j) dB_t^{i,j}$$

This very general equation applies to all continuous spiking neuron models.

Putting everything together

Questions:

- What happens when $N \to \infty$?
- Can we "summarize" the mean network activity with a few equations?
- What is ∞?

Answer to points 1 and 2

In the limit $N \to \infty$ (thermodynamic limit), given independent initial conditions

1. All neurons become independent (propagation of chaos) .

Answer to points 1 and 2

In the limit $N \to \infty$ (thermodynamic limit), given independent initial conditions

1. All neurons become independent (propagation of chaos) .
2. All neurons have asymptotically the same probability distribution

Answer to points 1 and 2

True under the following assumptions:

(H1). **Locally Lipschitz dynamics:** The functions f and g are uniformly **locally** Lipschitz-continuous with respect to the second variable.

(H2). **Lipschitz interactions:** The functions b and β are Lipschitz-continuous

(H3). **Monotonous growth of the dynamics:** We assume that f and g satisfy the following monotonous condition:

$$x^T f(t,x) + \frac{1}{2}\|g(t,x)\|^2 \leq K\,(1 + \|x\|^2)$$

Answer to points 1 and 2

- The equation:

$$d\bar{X}_t = f(t, \bar{X}_t)\, dt + \mathbb{E}_{\bar{Z}}[b(\bar{X}_t, \bar{Z}_t)]\, dt + g(t, \bar{X}_t)\, dW_t + \mathbb{E}_{\bar{Z}}[\beta(\bar{X}_t, \bar{Z}_t)]\, dB_t$$

$\bar{Z}$ is a process independent of $\bar{X}$ that has the same law, and $\mathbb{E}_{\bar{Z}}$ denotes the expectation under the law of $\bar{Z}$.
Note that we had to "guess" the equation

Answer to points 1 and 2

- Note $p(t, z)$ the PDF of $\bar{X}$.
- Rewrite the equations as:

$$d\bar{X}_t = f(t, \bar{X}_t)\, dt + \left(\int_{\mathbb{R}^d} b(\bar{X}_t, z) p(t, z)\, dz \right)\, dt + \\ g(t, \bar{X}_t)\, dW_t + \left(\int_{\mathbb{R}^d} \beta(\bar{X}_t, z) p(t, z) \right)\, dB_t$$

Comments

- The equation is “unsurprising”

Comments

- The equation is "unsurprising"
- It is complicated: "non-local" and the populations are independent but functionnally coupled (McKean-Vlasov equation)

Comments

- The equation is "unsurprising"
- It is complicated: "non-local" and the populations are independent but functionnally coupled (McKean-Vlasov equation)
- A "non-local" Fokker-Planck equations can be written:

$$\partial_t p(t,x) = \\ -\operatorname{div}\left(\left(f(t,x) + \int b(x,y)\, p(t,y)\, dy\right) p(t,x)\right) \\ + \frac{1}{2}\sum_{i,j=1}^{d} \frac{\partial^2}{\partial x_i \partial x_j}\left(D_{ij}(x)\, p(t,x)\right),$$

Comments

- The equation is "unsurprising"
- It is complicated: "non-local" and the populations are independent but functionnally coupled (McKean-Vlasov equation)
- A "non-local" Fokker-Planck equations can be written:

$$\partial_t p(t,x) = - \operatorname{div}\left(\left(f(t,x) + \int b(x,y)\, p(t,y)\, dy\right) p(t,x)\right) + \frac{1}{2}\sum_{i,j=1}^{d} \frac{\partial^2}{\partial x_i \partial x_j}\left(D_{ij}(x)\, p(t,x)\right),$$

$$D(x) = \mathbb{E}_Z\left[\beta(x,Z)\right] \mathbb{E}_Z^T\left[\beta(x,Z)\right]$$

Main result I

Theorem (Well-posedness)

Under the previous assumptions, there exists a unique solution to the mean-field equation on $[0, T]$ *for any* $T > 0$.

Proof: Use a fixed point argument.

Main result I

- Express the solution of the mean field equation as the solution of $x = F(x)$
- x is a measure and F a mapping from the set of measures into itself

Main result I

- Express the solution of the mean field equation as the solution of $x = F(x)$
- x is a measure and F a mapping from the set of measures into itself
- The proof involves Martingale inequalities and Gronwall lemma

Main result I

- Express the solution of the mean field equation as the solution of $x = F(x)$
- x is a measure and F a mapping from the set of measures into itself
- The proof involves Martingale inequalities and Gronwall lemma

Main result II

Theorem

Under the previous assumptions the following holds true:

- **Convergence:** *For each neuron i, the law of the multidimensional process $X^{i,N}$ converges towards the law of the solution of the mean-field equation, namely $\bar{X}$.*
- **Propagation of chaos:** *For any $k \in \mathbb{N}^*$, and any k-uplet $(i_1, \ldots, i_k)$, the law $p(t, z_1, \cdots, z_k)$ of the process $(X_t^{i_1,N}, \ldots, X_t^{i_k,N}, t \leq T)$ converges towards $p(t, z_1) \otimes \ldots \otimes p(t, z_k)$, i.e. the asymptotic processes have the law of the solution of the mean-field equations and are all independent.*

Proof: Use Sznitman coupling argument (1989), first proposed by Dobrushin (1970)

Numerical validation of the propagation of chaos

- The propagation of chaos effect appears for small populations: compatible with biology.



Numerical validation of the propagation of chaos

- The propagation of chaos effect appears for small populations: compatible with biology.





Numerical validation of the propagation of chaos

- The propagation of chaos effect appears for small populations: compatible with biology.







Numerical validation of the propagation of chaos

- It goes beyond decorrelation



Numerical validation of the propagation of chaos

- It goes beyond decorrelation





Numerical validation of the propagation of chaos

- It goes beyond decorrelation







Connection with information theory

- The propagation of chaos property implies that the cortex behaves optimally in terms of information coding.



from Ecker et al., Science 2010

Connection with information theory

- The propagation of chaos property implies that the cortex behaves optimally in terms of information coding.
- It looks as if neurons were coding probability laws



from Ecker et al., Science 2010

Finite size effects



Finite size effects



Finite size effects

Here is a **movie**

Discussion and caveats

Assumptions

- Realistic model of neurons, synapses.

Comments

- Fine

Discussion and caveats

Assumptions

- Realistic model of neurons, synapses.
- Somewhat realistic models of synaptic weights (not including plasticity)

Comments

- Fine
- Fine, but plasticity is (very) important

Discussion and caveats

Assumptions

- Realistic model of neurons, synapses.
- Somewhat realistic models of synaptic weights (not including plasticity)
- Fully connected network

Comments

- Fine
- Fine, but plasticity is (very) important
- The results (probably) hold if the number of neighbours is $o(N)$, e.g., $\log N$ (Ben Arous and Zeitouni 1999)

Discussion and caveats

Assumptions

- Realistic model of neurons, synapses.
- Somewhat realistic models of synaptic weights (not including plasticity)
- Fully connected network
- Independence of noise sources

Comments

- Fine
- Fine, but plasticity is (very) important
- The results (probably) hold if the number of neighbours is $o(N)$, e.g., $\log N$ (Ben Arous and Zeitouni 1999)
- Should look at correlated noise sources (probably very hard)

Discussion and caveats

Assumptions

- Realistic model of neurons, synapses.
- Somewhat realistic models of synaptic weights (not including plasticity)
- Fully connected network
- Independence of noise sources
- Accurate description for infinitely large populations.

Comments

- Fine
- Fine, but plasticity is (very) important
- The results (probably) hold if the number of neighbours is $o(N)$, e.g., $\log N$ (Ben Arous and Zeitouni 1999)
- Should look at correlated noise sources (probably very hard)
- Finite size effects seem to be small and can be characterized

Comparison between the static and dynamic randomness approaches

Static randomness

- Mean field equation is part of the large deviations technique

Dynamic randomness

- Mean field equation must be guessed

Comparison between the static and dynamic randomness approaches

Static randomness

- Mean field equation is part of the large deviations technique
- Mean field process is non-Markov

Dynamic randomness

- Mean field equation must be guessed
- Markov property is preserved

Comparison between the static and dynamic randomness approaches

Static randomness

- Mean field equation is part of the large deviations technique
- Mean field process is non-Markov

Dynamic randomness

- Mean field equation must be guessed
- Markov property is preserved

- Generalization to correlated noise/synaptic weights is possible

Comparison between the static and dynamic randomness approaches

Static randomness

- Mean field equation is part of the large deviations technique
- Mean field process is non-Markov

Dynamic randomness

- Mean field equation must be guessed
- Markov property is preserved

- Generalization to correlated noise/synaptic weights is possible
- Generalization to other types of connectivity is possible