

Some new approaches to the mean-field limit of integrate and fire systems

P.-E. Jabin,
joint works with D. Poyato, J. Soler, and D. Zhou

Introduction

Many-particle or multi-agent systems are used in a widespread range of applications

- Plasmas: Particles are ions or electrons.
- Astrophysics: Particles are dark matter particles, galaxies or galaxy clusters...
- Fluids: Point vortices, suspensions...
- Bio-mechanics: Medical aerosols in the respiratory tract, suspensions in the blood...
- Bio-Sciences: Collective behaviors of animals, swarming or flocking, but also dynamics of micro-organisms, chemotaxis, cell migration, **neural networks**...
- Social Sciences: Opinion dynamics, consensus formation...
- Economics: Mean-field games...

Biological neurons



Figure: Credits: CNRS Bordeaux, France; 2D reconstruction of rat hippocampus, marked for cytoskeleton protein.

Our goals

We want to derive some simple, continuous model to describe the **behavior** of networks of biological neurons **at large scales**, as

- we typically have a large number of neurons, up to **$8.6 \cdot 10^{10}$ neurons** in the human brain for example;
- each neuron has a complex map of interactions with other neurons, with an average of **6,000 – 7,000 synaptic connections**.

Some sort of **mean-field approximation** can be expected.

But classical mean-field tools cannot be applied because each neuron has a distinct map of connections.

Integrate and fire for a single neuron

One of the simplest model for a single neuron is a straightforward **spike dynamics** of the type originally proposed by Lapicque in 1907. Between spikes, the membrane potential is assumed to increase with a straightforward ODE (or SDE)

$$\frac{dV}{dt} = b(V(t)).$$

The neuron has a certain probability to spike

$$\mathbb{P}(\text{spike between } t \text{ and } t + dt) \approx \nu(V(t)) dt.$$

In that case the potential of the neuron is reset at $V = 0$ at the same time.

There exists many, many variants of such models, for example with stochastic terms, more complex dynamics, more intricate spiking rules...

Integrate and fire for a single neuron

One of the simplest model for a single neuron is a straightforward **spike dynamics** of the type originally proposed by Lapicque in 1907. Between spikes, the membrane potential is assumed to increase with a straightforward ODE (or SDE)

$$dV_t = b(V(t)) + \sigma dW_t.$$

The neuron has a certain probability to spike

$$\mathbb{P}(\text{spike between } t \text{ and } t + dt) \approx \nu(V(t)) dt.$$

In that case the potential of the neuron is reset at $V = 0$ at the same time.

There exists many, many variants of such models, for example with stochastic terms, more complex dynamics, more intricate spiking rules...

How to connect the dynamics of each neurons

For some of the simplest **integrate and fire** model, a neuron $\#i$ only interact with another neuron $\#j$ when the neuron $\#i$ fires:

For all j connected to i , V_j jumps by αw_{ji} ,

at the **exact time when i fires**, and where w_{ji} are the **synaptic connections**.

- The w_{ji} can be positive or negative with corresponding excitation or inhibition between neurons.
- The w_{ji} may all be completely different for every pair of neurons.
- We are not even taking **synaptic plasticity** into account. In fact the w_{ji} should evolve in time.
- We are seemingly very **far from the fully symmetric case of identical interacting charged particles...**

Modeling is a complex issue

The previous discussion obviously does not do justice to the complexity of the modeling of neuron dynamics.

The dynamics of each neuron involve complex biochemical processes, with temporal and spatial dependencies. Their reduction to a set of simple rules for multi-agent systems has consequences... For example while different Hodgkin-Huxley type of models may capture better the individual dynamics of neurons, they seem to have more difficulty with the interactions between neurons (spikes being not as well defined). See for example Gerstner-Kistler, Izhikevich...

However recent advances in medical imaging have also provided more accurate information. For example the full connectome of the *Drosophila* was recently obtained.

Scaling in biological neuron networks

- Mammalian brains contain between $10^8 - 10^{11}$ neurons (86 · 10⁹ actually in humans). Of course models typically applied only to sub-domains... $\rightarrow N \gg 1$.
- In the human brain, each neuron has on average 7000 synaptic connections to other neurons
 $\rightarrow \sup_{i,j} w_{ij} \sim 10^{-3} - 10^{-4} \ll 1$.

\rightarrow The map of synaptic connections is only sparsely connected but each neuron still interact with a large number of other neurons. The large N makes it interesting, for both analytical and numerical studies, to derive simplified models at macroscopic scales. To summarize, we make the following assumption extending the classical mean-field scaling

$$\sup_i \sum_j |w_{ij}| + \sup_j \sum_i |w_{ji}| \leq C, \quad \max_{i,j} |w_{ij}| \rightarrow 0 \quad \text{as } N \rightarrow \infty.$$

Biological neurons




Figure: Part of Figure 2 in Hulse et al 21'. Recent progress in experimental biology that makes detailed connectome for large neuron networks available.

A mean-field limit for neural networks?

The advantage of mean-field approximations for the analysis or simulations of large neuronal networks has long been recognized: See Omurtag-Knight-Sirovich, Renart-Brunel-Wang, Mattia-Del Giudice,...

Neural field models for example have proved enduringly popular for this reason, see for example Amari, Freeman, Griffith, Da Silva et al, Nunez, Wilson-Cowan in the 70s...

But neural field models also face strong criticism as neurons that are spatially close may still have completely different synaptic connections. And synaptic weights play a critical role for the individual dynamics of each neuron in the network.

This leaves the **question of some sort of continuum limit fully open.**

A novel statistical notion

- We propose a novel notion of the statistical distribution in the system, directly from the joint law and connection maps.
- No additional assumption is needed for the w_{ij} outside of the natural scaling.
- Those empirical observables are indexed by trees.
- For the trivial tree with only the root vertex, the observable is the classical 1-particle distribution.
- The family could be extended to include other well-known statistical objection, such as the 2-particle distribution.
- At the limit, observables provide a self-contained macroscopic description of the system: convergence of observables initially implies convergence of observables at any later time.

A statistical point of view

Instead of looking at trajectories, we consider the **joint law** f_N at time t .

For the integrate and fire models, we have

$f_N(t, v_1, \dots, v_N)$ = probability density of the membrane potentials $V_1(t), \dots, V_N(t)$ at time t .

It contains most of the **statistical information** on the system but not all the information: **Correlations in time are lost** and it may be difficult to reconstruct trajectories of the system.

A new notion of observables for integrate and fire models

Denote the marginal of the joint law f_N of the i_1, \dots, i_n -th neurons

$$\begin{aligned} & f_N^{i_1, \dots, i_n}(t, v_1, \dots, v_n) \\ &= \left(\int_{\mathbb{R}^{N-n}} f_N(t, z_1, \dots, z_N) \prod_{j \neq i_1, \dots, i_n} dz_j \right) \Big|_{z_{i_k} = v_k, k=1, \dots, n}. \end{aligned}$$

We then define the family of observables indexed by trees T

$$\begin{aligned} & \tau_N(T, w, f_N)(t, v_1, \dots, v_{|T|}) \\ &= \frac{1}{N} \sum_{i_1, \dots, i_{|T|}=1}^N \prod_{(j,k) \in T} w_{i_j, i_k} f_N^{i_1, \dots, i_{|T|}}(v_1, \dots, v_{|T|}), \end{aligned}$$

where the tree T represents a possible map of interactions starting from neuron i_1 .

Marginals in the simplified case

In the case of symmetrical systems with indistinguishable agents, the observables reduce to the classical marginals that have an immediate physical interpretation. For our simplified setting

$$f_k(t, x_1, \dots, x_k) = \text{Law at time } t \text{ of } X_1, \dots, X_k.$$

For example f_1 is the 1-particle distribution, while f_2 contains information about correlations between particles.

The various marginals are nested in a natural hierarchy

$$\begin{aligned} f_k(t, x_1, \dots, x_k) \\ = \int_{\mathbb{R}^d} f_{k+1}(t, x_1, \dots, x_{k+1}) dx_{k+1}. \end{aligned}$$

Our new observables extend this physical interpretation to non-identical agents.

A tree indexed hierarchy for integrate and fire models

In the setting of integrate and fire models, we do not have a closed hierarchy but we are still able to derive an approximate hierarchy that is now indexed by trees where we denote $v = (v_1, \dots, v_{|T|})$ and $z = (z_1, \dots, z_{|T|})$

$$\begin{aligned} & \partial_t \tau_N(T)(t, v_1, \dots, v_{|T|}) \\ &= \sum_{n=1}^{|T|} \left\{ \left[-\partial_{v_n} (b(v_n) \tau_N(T)(t, \cdot)) + \frac{\sigma^2}{2} \partial_{v_n}^2 \tau_N(T)(t, \cdot) \right. \right. \\ & \quad \left. \left. - \nu(v_n) \tau_N(T)(t, \cdot) + \delta_0(v_n) \left(\int_{\mathbb{R}} \nu(z_n) \tau_N(T)(t, z) dz_n \right) \Big|_{\forall k \neq n, z_k = v_k} \right] \right. \\ & \quad \left. - \partial_{z_n} \left[\int_{\mathbb{R}} \nu(v_{|T|+1}) \tau_N(T+n)(t, v, v_{|T|+1}) dv_{|T|+1} \right] \right\} + \text{remainders.} \end{aligned}$$

A tree indexed hierarchy for integrate and fire models

$$\begin{aligned}
 & \partial_t \tau_N(T)(t, v_1, \dots, v_{|T|}) \\
 &= \sum_{n=1}^{|T|} \left\{ \left[-\partial_{v_n} (b(v_n) \tau_N(T)(t, \cdot)) + \frac{\sigma^2}{2} \partial_{v_n}^2 \tau_N(T)(t, \cdot) \right. \right. \\
 & \quad \left. \left. - \nu(v_n) \tau_N(T)(t, \cdot) + \delta_0(v_n) \left(\int_{\mathbb{R}} \nu(z_n) \tau_N(T)(t, z) dz_n \right) \Big|_{\forall k \neq n, z_k = v_k} \right] \right. \\
 & \quad \left. - \partial_{z_n} \left[\int_{\mathbb{R}} \nu(v_{|T|+1}) \tau_N(T+n)(t, v, v_{|T|+1}) dv_{|T|+1} \right] \right\} + \text{remainders.}
 \end{aligned}$$

As before the equation on $\tau_N(T)$ depends on all observables at the next level: All trees $T+n$ where we add a leaf on vertex $\#n$.

The main steps in the proof

- We want to obtain direct estimates on the observables $\tau_N(T)$ for integrate and fire, and the marginals f_k for charged particles.
- This will also allow to derive macroscopic descriptions by passing to the limit in the hierarchies.
- However a **major difficulty** is that each equation on the hierarchy depends on **observables at the next level**.
- At the technical level, this is resolved through the introduction of new commutator estimates that are compatible with weak convergence and are uniform in the dimension.

Mean-field limit for integrate and fire models

We recall that our mean-field scaling is

$$\sup_i \sum_j |w_{ij}| + \sup_j \sum_i |w_{ji}| \leq C, \quad \max_{i,j} |w_{ij}| \rightarrow 0 \quad \text{as } N \rightarrow \infty.$$

Then there exists a notion of **limiting synaptic “kernel”** $w(\xi, \zeta)$ in a complex measure-valued space $L_\xi^\infty \mathcal{M}_\zeta \cap L_\zeta^\infty \mathcal{M}_\xi$ such that the limit of the observables can be obtained through a mean-field equation on an extended density $f(t, \nu, \xi)$

$$\begin{aligned} \partial_t f + \partial_\nu ((b(\nu) + S(t, \xi) f) &= \frac{\sigma^2}{2} \partial_\nu^2 f - \nu(\nu) f, \\ f(t, \nu = 0, \xi) &= \frac{1}{b(0) + S(t, \xi)} \int_0^\infty \nu(z) f(t, z, \xi) dz, \quad (1) \\ S(t, \xi) &= \int_0^1 \int_0^\infty \nu(z) f(t, z, \zeta) dz w(\xi, d\zeta). \end{aligned}$$

A theorem for integrate and fire

Define the limiting observables through

$$\tau_\infty(T, w, f)(t, z) := \int_{[0, 1]^{|T|}} \prod_{(l,m) \in T} w(\xi_l, \xi_m) f^{\otimes |T|}(t, z_1, \xi_1, \dots, z_{|T|}, \xi_{|T|}) d\xi_1 \dots d\xi_{|T|}.$$

Theorem

Assume b, ν are bounded and smooth, that f is a smooth, and fast decaying solution to (1) for $w \in L_\xi^\infty \mathcal{M}_\zeta \cap L_\zeta^\infty \mathcal{M}_\xi$, and that

$\tau_N(T)(t=0) \rightarrow \tau_\infty(T)(t=0)$, weak- $*$ in measure, for all trees T .

Then for any $t > 0$,

$\tau_N(T)(t) \rightarrow \tau_\infty(T)(t)$, weak- $*$ in measure, for all trees T .

In particular we have convergence of the 1-particle distribution.

Conclusions

- It allows the rigorous derivation of a sort of neural field model but the limiting synaptic kernel cannot be interpreted in spatial terms.
- Novel, straightforward quantitative estimates with minimal assumptions on the interaction kernel.
- Fits with the expected scaling of molecular chaos where $f_k = f^{\otimes k}$ but does not strictly require independence.
- Many open questions remain, including whether it is possible to derive some **notion of mean trajectory** or **how to incorporate learning** in the model.