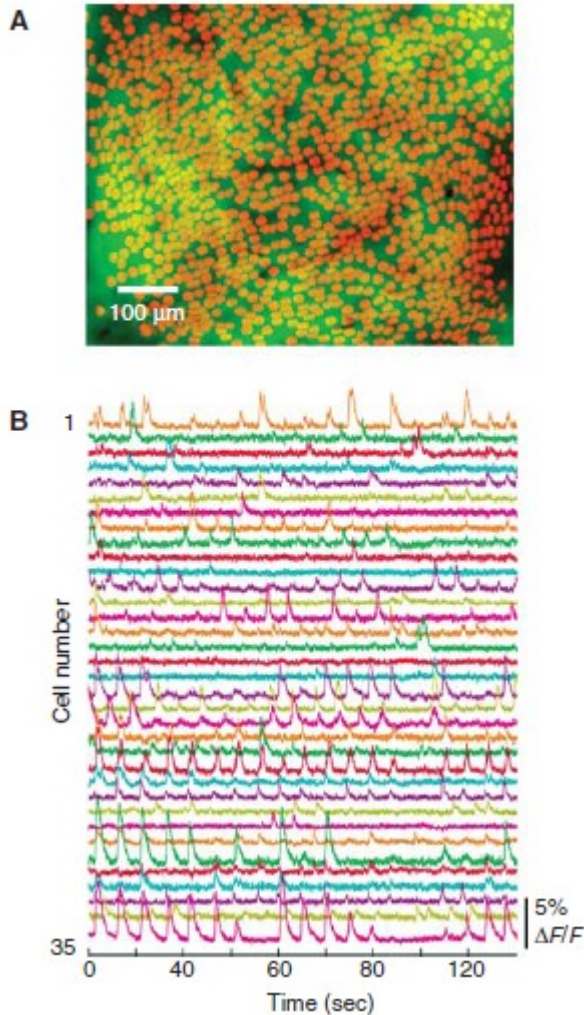


Tutorial on Meanfield Methods for Networks of Spiking Neurons

ICMNS Tutorials
May 29th, 2016

Alex Roxin
Computational Neuroscience Group
Centre de Recerca Matemàtica
Bellaterra, Spain

Behavior is driven by the collective spiking of large networks of neurons



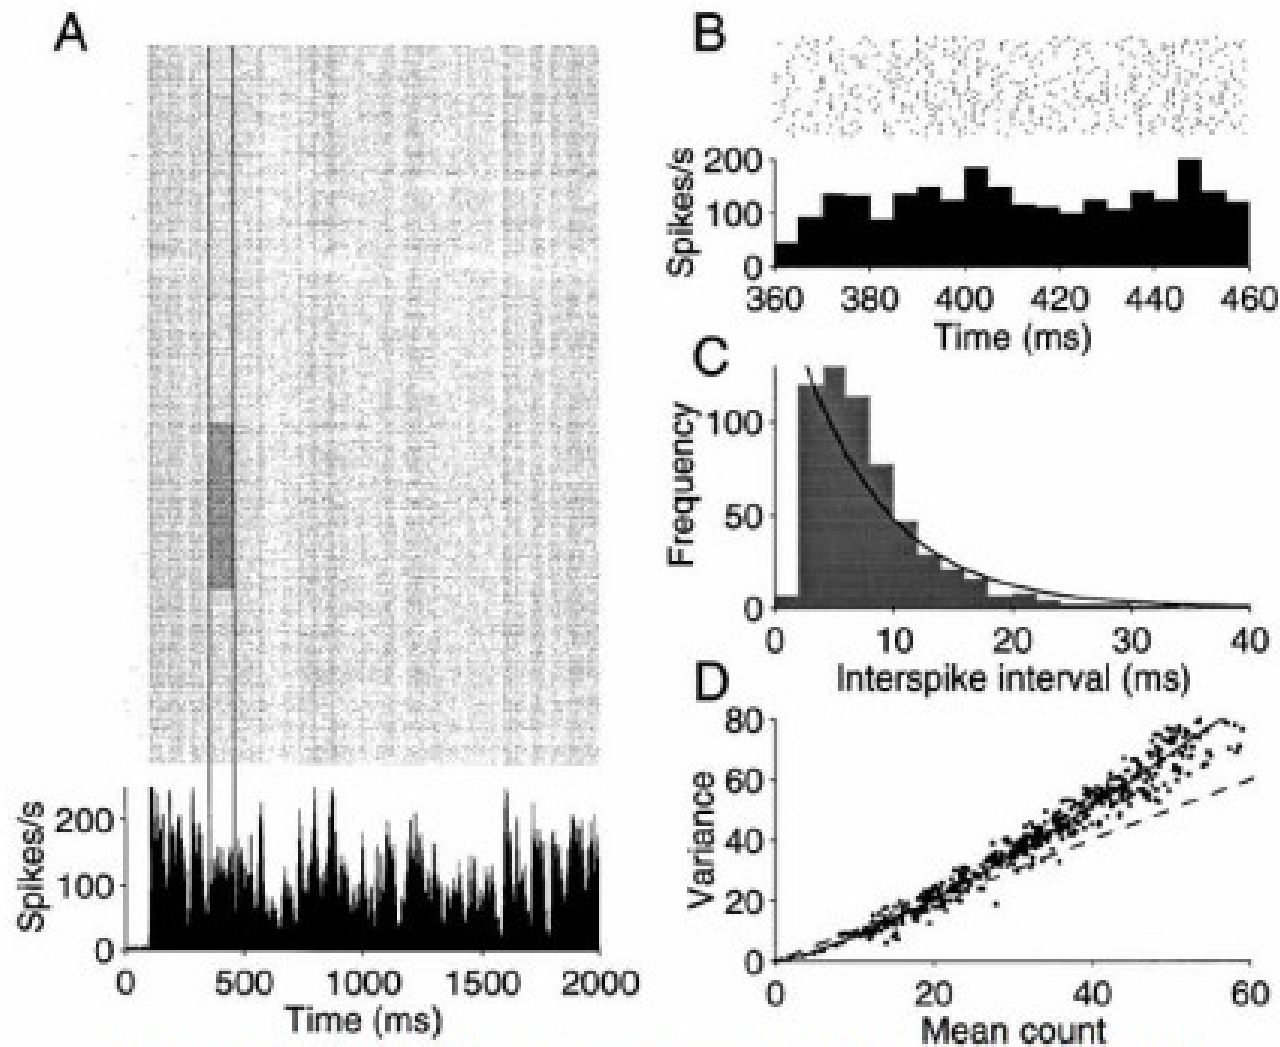
Large Scale Recording Techniques

- EEG
- MEG
- LFP
- ECoG
- fMRI
- Voltage-sensitive dyes
- Calcium imaging
- Multiple unit recordings

Jercog et al. CSHPB 2016

Single cell spiking is highly irregular


Spiking dynamics of a single cell: extracellular



Shadlen and Newsome, J. Neurosci. 1998

To understand how spiking activity drives behavior we need a network theory for irregularly spiking neurons. I will discuss the theory developed in Amit and Brunel, Network (1997), Brunel and Hakim, Neural Comp. (1999), Brunel, J. Comp. Neurosci. (2000) and other papers. As a resource I will use the review chapter entitled “Mean-field theory of irregularly spiking neuronal populations and working memory in recurrent cortical networks” by Renart, Brunel and Wang.

First we consider a single neuron receiving many inputs



Ce excitatory inputs

Leaky Integrate-and-Fire Neuron

$$\tau_m \dot{V} = -V + I(t)$$

threshold B.C. $V(t_{\text{spike}}^-) = V_t$ $V(t_{\text{spike}}^+) = V_r$

refractory time τ_{ref}

Instantaneous Post-synaptic Currents

$$I(t) = J_e \sum_{i=1}^{C_e} \sum_j \delta(t - t_j^i) - J_i \sum_{i=1}^{C_i} \sum_j \delta(t - t_j^i)$$

We assume the inputs are independent Poisson processes with rates ν_e and ν_i .

Consider the evolution of the probability distribution for the voltage $\rho(V, t)$.

Because the stochastic process is Markov, the probability distribution can be written

$$\rho(V, t + \Delta t | V_0, t_0) = \int_{-\infty}^{\infty} dV' \rho(V, t + \Delta t | V', t) \rho(V', t | V_0, t_0).$$

Please go to the notes, section **A**, to find the derivation of the evolution equation for the probability distribution. In the limit of a large number of very weak inputs the equation becomes the following Fokker-Planck equation.

$$\frac{\partial \rho}{\partial t} = \frac{\sigma^2}{2\tau_m} \frac{\partial^2 \rho}{\partial V^2} + \frac{\partial}{\partial V} \left[\left(\frac{V - \mu}{\tau_m} \right) \rho \right]$$

$$\mu = \tau_m \left(J_e C_e \nu_e - J_i C_i \nu_i \right)$$

$$\sigma^2 = \tau_m \left(J_e^2 C_e \nu_e + J_i^2 C_i \nu_i \right)$$

+ appropriate boundary conditions (next page)

To derive the boundary conditions first rewrite the Fokker-Planck equation as a continuity equation.

$$\frac{\partial \rho}{\partial t} + \frac{\partial F}{\partial V} = 0,$$

This equation says that changes in the probability density in a given voltage interval are due to a flux from that interval. The flux is the following

$$F = -\frac{\sigma^2}{2\tau_m} \frac{\partial \rho}{\partial V} - \left[\left(\frac{V - \mu}{\tau_m} \right) \rho \right].$$

The probability density is zero for any voltage above the threshold. This means it must approach zero continuously from below. Why? In the diffusion approximation all the “kicks” are vanishingly small; if we had finite kicks the density could be non-zero at threshold.

Finally we have the boundary conditions:

$$\rho(V_t, t) = 0 \quad \text{The density goes to zero at threshold.}$$

$$\nu(t) = F(V_t, t) \quad \text{The flux at threshold gives the firing rate.}$$

$$\lim_{\epsilon \rightarrow 0} F(V_r + \epsilon, t) - F(V_r - \epsilon, t) = \nu(t) \quad \text{The flux exiting at the threshold is reinjected at the rest potential.}$$

$$\int_{-\infty}^{V_t} dV \rho(V, t) = 1. \quad \text{The integral must be one for a probability density.}$$




Please refer to the notes, Section **B**, for a derivation of the solution of the steady-state probability density and firing rate.

Steady-state probability density for LIF

$$\rho_0(V) = \frac{2\tau\nu_0}{\sigma} e^{-\frac{(V-\mu)}{\sigma}} \int_{\frac{V-\mu}{\sigma}}^{\frac{V_t-\mu}{\sigma}} dy e^{y^2} H\left(y - \frac{(V_r - \mu)}{\sigma}\right)$$

Steady-state firing rate for LIF

$$\nu_0 = \frac{1}{\tau\sqrt{\pi}} \left[\int_{\frac{V_r-\mu}{\sigma}}^{\frac{V_t-\mu}{\sigma}} dy e^{y^2} \operatorname{erfc}(-y) \right]^{-1}$$






Steady-state probability density for LIF


$$\rho_0(V) = \frac{2\tau\nu_0}{\sigma} e^{-\frac{(V-\mu)}{\sigma}} \int_{\frac{V_r-\mu}{\sigma}}^{\frac{V_t-\mu}{\sigma}} dy e^{y^2} H\left(y - \frac{(V_r-\mu)}{\sigma}\right)$$

Steady-state firing rate for LIF

$$\nu_0 = \frac{1}{\tau\sqrt{\pi}} \left[\int_{\frac{V_r-\mu}{\sigma}}^{\frac{V_t-\mu}{\sigma}} dy e^{y^2} \operatorname{erfc}(-y) \right]^{-1}$$



We see that at low firing rates the spiking is highly irregular while at high rates it is nearly periodic. Neurons *in-vivo* are highly irregular; this is sometimes quantified by the coefficient of variation of the interspike interval (CV) which is 1 for a Poisson Process. The CV can be calculated for the LIF neuron, see section **C** of the notes for the derivation.



Theory for recurrent networks

So far we have calculated the firing rate of a single LIF neuron bombarded by Poisson Inputs. To extend this to a recurrent network we only need to make the input and output firing rates self-consistent. When might this work? Which assumptions may break down?

We will consider two kinds of networks: 1 – all-to-all connected networks and 2 - sparsely connected networks.

To make things consistent yet easy, from now on I will discuss networks of excitatory and inhibitory cells (E-I networks) in equal numbers, for which individual E and I cells are identical and synaptic weights are constant and $J_{ie} = J_{ii} = J_i$ and $J_{ee} = J_{ie} = J_e$.

Also, for the sparsely connected networks I will assume the sparseness is the same for all connectivities. When this is the case, then the mean and variance of the input to any cell are the same and can be written

$$\begin{aligned}\mu &= \mu_{ext} + \tau \left(J_e C_e \nu_e - J_i C_i \nu_i \right) \\ \sigma^2 &= \sigma_{ext}^2 + \tau \left(J_e^2 C_e \nu_e + J_i^2 C_i \nu_i \right)\end{aligned}$$

where the rates are those of the neurons in the network.

All-to-all connected networks

Here all neurons receive identical recurrent input (the meanfield). A proper thermodynamic limit (large system size) is only possible if the synaptic weights scale as $1/N$. The number of connections $C = N$. This means that

$$\mu = \mu_{ext} + \tau \left(J_e \nu_e - J_i \nu_i \right)$$

$$\sigma^2 = \sigma_{ext}^2 + \tau \frac{1}{N} \left(J_e^2 \nu_e + J_i^2 \nu_i \right)$$


so the only fluctuations in the input are external in the large-system limit.

The firing rates are found self-consistently through the steady-state Eq.5 from the notes


$$\nu_e = \phi(\mu, \sigma)$$

$$\nu_i = \phi(\mu, \sigma)$$

The steady-state theory for the all-to-all network works quite well even for relatively small system sizes. Even for $N = 10$ the deviation from theory is only about 10% for the average rate. Please see notes, section **D** for the parameter values I used.



$$N = 1000, \sigma = 5, \mu_{\text{ext},I} = 14.7$$



Sparse networks

Here either: 1 – all neurons receive the same number of connections C , which are chosen randomly, or 2 – connections are made with probability C/N . In the second case the number of connections is a binomial variable with mean C , an added source of fluctuations. I will not treat the effect of this quenched variability here.

If we want to keep fluctuations of order one in the thermodynamics limit, we can take a fixed connection probability ($p = C/N$) and scale the synaptic weights as $1/\sqrt{N}$. This leads to


$$\mu = \tau \sqrt{pN} \left(J_{ext} \nu_{ext} + J_e \nu_e - J_i \nu_i \right)$$
$$\sigma^2 = \sigma_{ext}^2 + \tau p \left(J_e^2 \nu_e + J_i^2 \nu_i \right)$$


Where the strength of external inputs scales the same way with N . Clearly the mean input explodes in this limit, meaning the activity is epileptic or silent...unless the rates exactly cancel. This is known as the *balance* condition and this is a *balanced network*, see (van Vreeswijk and Sompolinsky Neural Computation 1998) for a thorough discussion of this type of scaling.

Sparse networks


Alternatively, we can just study finite networks with some given number of connections per neuron. In this case finite size effects may need to be taken into account. Please see Section **E** of the notes for a description of the network I use.

$$C = pN = 200$$






Finite-size effects are due to coherent fluctuations in the populations firing rate, i.e. they are shared fluctuations in the inputs to all cells.

$$\sum_{j=1}^N \sum_k J w_{ij} \delta(t - t_j^k) = Jp \left(N\nu(t) + \delta\nu(t) \right) + \text{independent fluctuations}$$


We assume neurons are Poisson processes, therefore there are fluctuations (the variance equals the mean).

$$= JC\nu(t) + \sigma\eta(t)$$


$$\text{where } \langle \eta(t)\eta(t') \rangle = \delta(t - t')$$

$$\sigma = JC\sqrt{\frac{\nu(t)}{N}}$$



This is why coherent fluctuations vanish as N goes to infinity.

In a sparse network, quenched variability gives rise to “noisy” dynamics, even in the absence of noisy forcing. The model I simulated below is completely deterministic, yet the theory assuming neurons are Poisson processes works pretty well!

$$\mu_{\text{ext},I} = 24 \text{ mV}, J_e = 0.2, J_i = -1.0, p = 0.1, N_e = 8,000, N_i = 2,000$$



This is what the dynamics looks like when the mean excitatory input (which is just a constant drive) is 24 mV (the middle dot from the last slide). Does this look Poisson to you?



Linear Stability

Up until now we have just calculated the stationary state in networks of LIF neurons. However, E-I networks can generate oscillations and bistable regimes, i.e. the stationary state can become unstable. We can calculate when these instabilities will occur (as a function of some network parameter such as external drive) using standard linear stability techniques.

$$\frac{\partial \rho}{\partial t} = \frac{\sigma^2}{2\tau_m} \frac{\partial^2 \rho}{\partial V^2} + \frac{\partial}{\partial V} \left[\left(\frac{V - \mu}{\tau_m} \right) \rho \right]$$

$$\rho = \rho_0 + \delta \rho e^{\lambda t}$$

$$\mu = \mu_0 + \delta \mu e^{\lambda t}$$

$$\nu = \nu_0 + \delta \nu e^{\lambda t}$$

eigenvalue of perturbation

steady-state solution

small-amplitude perturbation

Please see notes, sections **F** and **G** for a discussion of how to calculate the linear stability.

The linear stability calculation leads to a rather complicated eigenvalue problem. It is actually much easier nowadays to simply solve for the stationary solution and linear stability numerically directly from the Fokker-Planck equation, thanks to the work of Magnus Richardson (Richardson, PRE 2007). This is how I made all of the theoretical curves here BTW. Please see section **G** of the notes for details on the numerical method. The numerical method works for any kind of integrate-and-fire neuron, not just leaky!

There is much more to be said about meanfield methods for spiking networks. We have only just scratched the surface, but we have only had one hour. Here's a list of topics, some of which are well understood and some less so.

- Non-instantaneous synapses.
- Quenched variability: firing rate distributions.
- Response to time-varying inputs.
- Oscillatory states.
- Working memory states.
- Microscopic state versus macroscopic state.
- Reduction of spiking networks to firing rate models.

Just to conclude on an interesting note, one may ask “What happens in recurrent networks of LIF neurons when the diffusion approximation breaks down”? I'm not sure it's entirely understood yet. Here are two attempts to do so.

ARTICLES

nature
neuroscience

PHYSICAL REVIEW E **92**, 040901(R) (2015)

Slow fluctuations in recurrent networks of spiking neurons

Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons

Srdjan Ostojic

Stefan Wieland,^{1,2} Davide Bernardi,^{1,2} Tilo Schwalger,³ and Benjamin Lindner^{1,2}

¹Bernstein Center for Computational Neuroscience Berlin, Philippstrasse 13, Haus 2, 10115 Berlin, Germany

²Physics Department of Humboldt University Berlin, Newtonstrasse 15, 12489 Berlin, Germany

³School of Computer and Communication Sciences and School of Life Sciences, Brain Mind Institute, École Polytechnique Fédérale de Lausanne, Station 15, 1015 Lausanne EPFL, Switzerland

