Nonparametric Analysis of Simultaneously Recorded Spike Trains Considered as a Realization of a Multivariate Point Process.

Christophe Pouzat

Mathématiques Appliquées à Paris 5 (MAP5)

Université Paris-Descartes et CNRS UMR 8145

christophe.pouzat@parisdescartes.fr

CANUM: 4 April 2014

Where are we ?

Data's origin

Viewed "from the outside", neurons generate brief electrical pulses: the action potentials

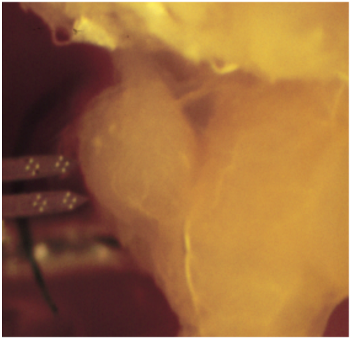
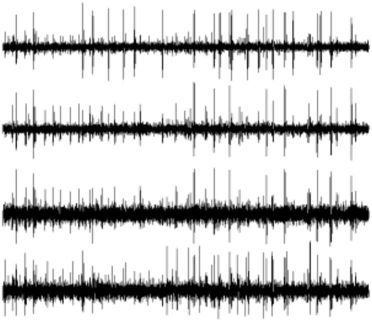

Left, the brain of an insect with the recording probe on which 16 electrodes (the bright spots) have been etched. Each probe's branch has a 80 μm width. Right, 1 sec of data from 4 electrodes. The spikes are the action potentials.

Spike trains

After a "rather heavy" pre-processing called spike sorting, the raster plot representing the spike trains can be built:

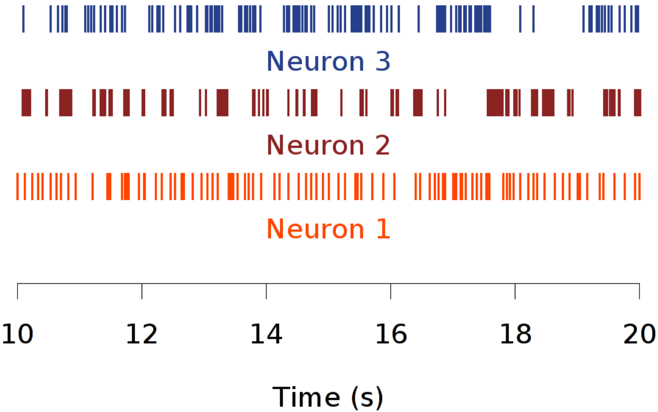

Modeling spike trains: Why and How?

- A key working hypothesis in Neurosciences states that the spikes' occurrence times, as opposed to their waveform are the only information carriers between brain region (Adrian and Zotterman, 1926).
- This hypothesis legitimates and leads to the study of spike trains *per se*.
- It also encourages the development of models whose goal is to predict the probability of occurrence of a spike at a given time, without necessarily considering the biophysical spike generation mechanisms.
- In the sequel we will identify spike trains with point process / counting process realizations.

A tough case (1)

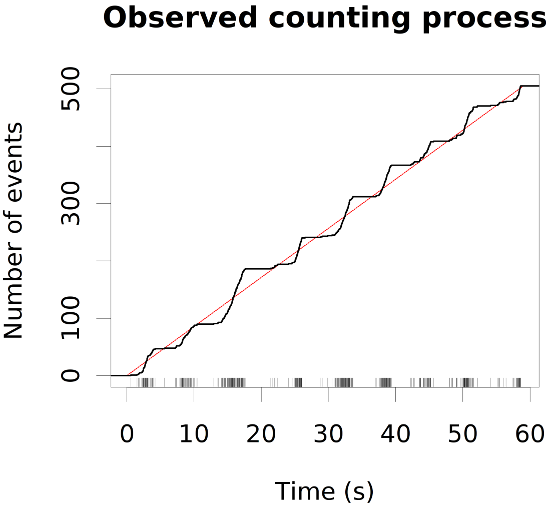

The expected counting process of a homogeneous Poisson process—with the same mean frequency—is shown in dashed red.

A tough case (2)

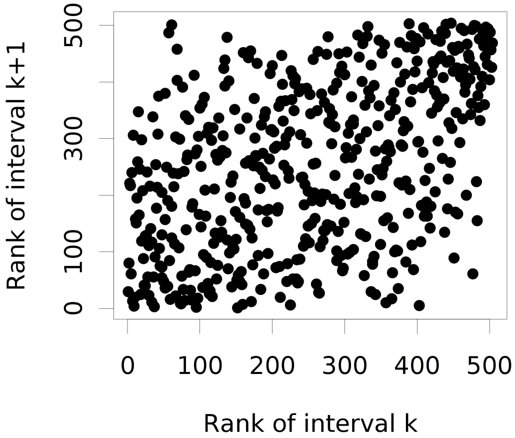

A renewal process is inadequate here: the rank of successive inter spike intervals are correlated.

Where are we ?

Model constraints

Our model should give room for:

- The elapsed time since the last spike of the neuron (enough for homogeneous renewal processes).
- Variables related to the discharge history—like the duration of the last inter spike interval.
- The elapsed time since the last spike of a "functionally coupled" neuron.
- The elapsed time since the beginning of a applied stimulation.

Filtration, history and conditional intensity

- Probabilists working on processes use the filtration or history: a family of increasing sigma algebras, $(\mathscr{F}_t)_{0 \le t \le \infty}$, such that all the information related to the process at time t can be represented by an element of $\mathscr{F}_t$.
- The conditional intensity of a counting process $N(t)$ is then defined by:

$$\lambda(t \mid \mathscr{F}_t) \equiv \lim_{h \downarrow 0} \frac{\mathrm{Prob}\{N(t+h) - N(t) = 1 \mid \mathscr{F}_t\}}{h} .$$

- λ constitutes an exhaustive description of process / spike train.

Two problems

As soon as we adopt a conditional intensity based formalism, we must:

- Find an estimator $\hat{\lambda}$ of λ.
- Find goodness of fit tests.

Where are we ?

What to do with λ: A summary

We start by associating to λ, the integrated intensity:

$$\Lambda = \int_0^t \lambda(u \mid \mathscr{F}_u) du,$$

it then easy—but a bit too long for such a brief talk—to show that:

- If our model is correct ($\hat{\lambda} \approx \lambda$), the density of successive spikes after time transformation:

  $$\{t_1, \ldots, t_n\} \rightarrow \{\Lambda(t_1) = \Lambda_1, \ldots, \Lambda(t_n) = \Lambda_n\}$$

  is exponential with parameter 1.
- Stated differently, the point process $\{\Lambda_1, \ldots, \Lambda_n\}$ is a homogeneous Poisson process with parameter 1.

The next slides illustrate this result.

Time transformation illustration (1)

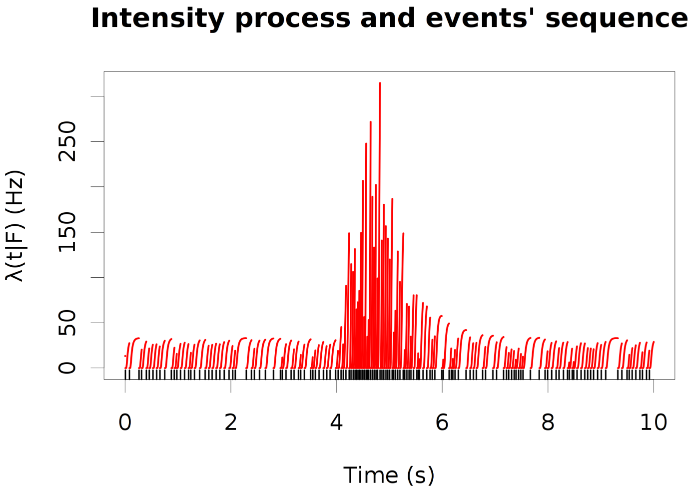

Time transformation illustration (2)

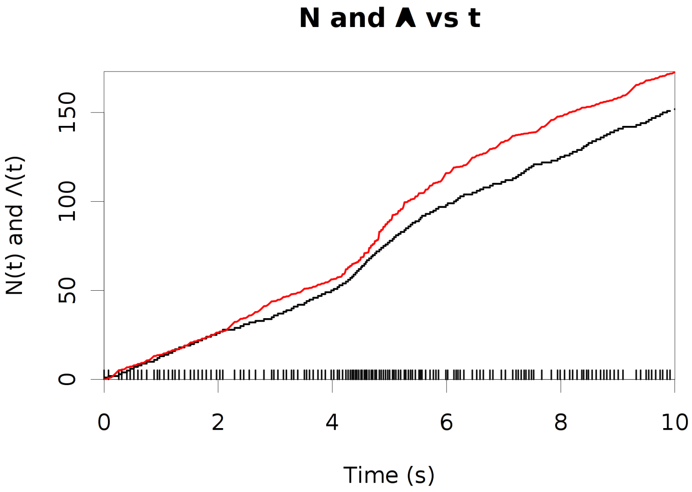

Time transformation illustration (3)

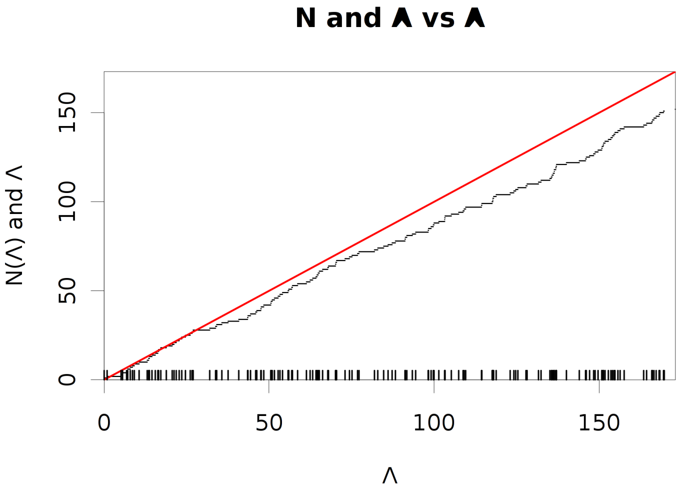

Ogata's tests

- If, for a good model, the transformed sequence of spike times, $\{\hat{\Lambda}_1, \ldots, \hat{\Lambda}_n\}$, is the realization of a homogeneous Poisson process with rate 1, we should test $\{\hat{\Lambda}_1, \ldots, \hat{\Lambda}_n\}$ against such a process.
- This is what Yosihiko Ogata proposed in 1988 (Statistical models for earthquake occurrences and residual analysis for point processes, Journal of the American Statistical Association, 83: 9-27).
- But an observation suggest nevertheless that another type of test could also be used…

A Brownian motion?

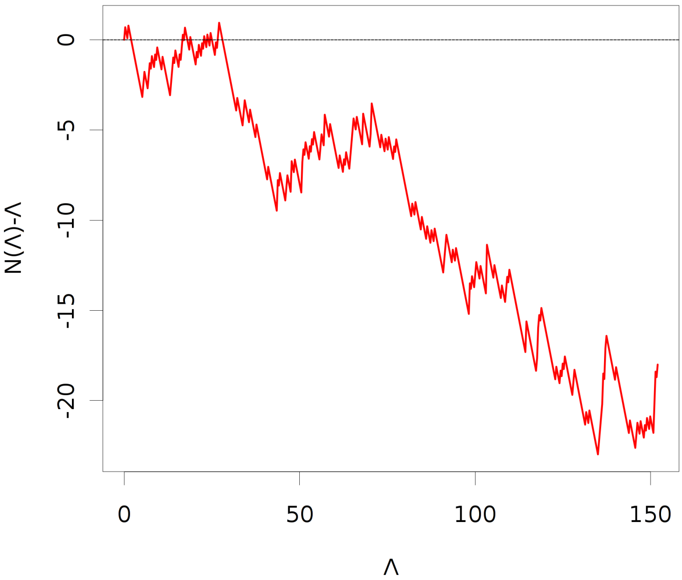

Where are we ?

Donsker's theorem and minimal area region

- The intuition of the convergence—of a properly normalized version—of the process $N(\Lambda) - \Lambda$ towards a Brownian motion is correct.
- This is an easy consequence of Donsker's theorem as Vilmos Prokaj explained to me on the R mailing and as Olivier Faugeras and Jonathan Touboul explained to me directly.
- It is moreover possible to find regions of minimal area having a given probability to contain the whole trajectory of a canonical Brownian motion (Kendall, Marin et Robert, 2007; Loader et Deely, 1987).
- We get thereby a new goodness of fit test.

Minimal area region at 95%

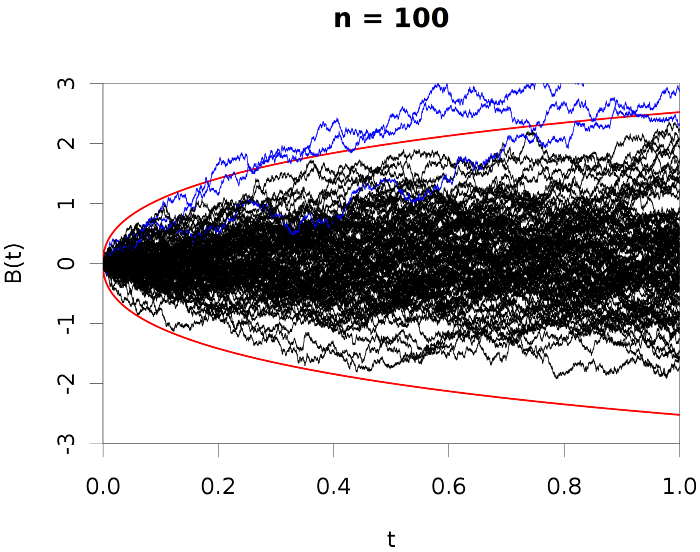

Where are we ?

Back to our "tough" case (1)

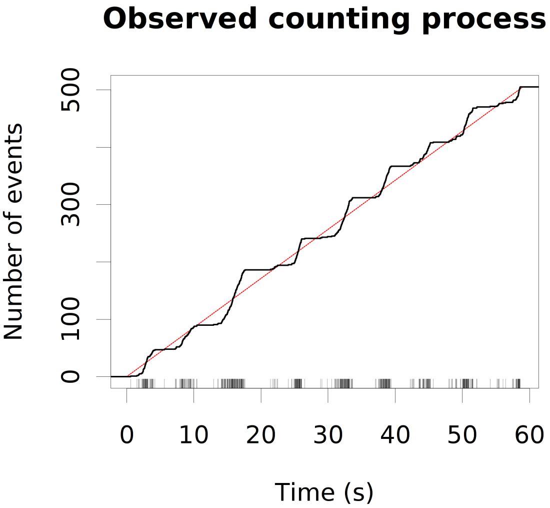

Back to our "tough" case (2)

Our former exploratory analysis leads to a minimal the following model:

$$\lambda(t|\mathscr{F}_t) = f(t - t_d, t_d - t_{ad}),$$

where t_d is the time of the last spike and t_{ad} is the time of the one-before-the-last spike.
This is known in the point process literature as a Wold process.

David Brillinger's approach

- We follow D. Brillinger (1988) who starts by binning the time axis into bins of length h, where h is small enough to observe at most one spike per bin.
- We are therefore brought back to a binomial regression problem.
- The binned data are then considered as an observation from a collection of Bernoulli random variables $\{Y_1, \ldots, Y_k\}$ with parameters: $f(t - t_d, t_d - t_{ad})\, h$.
- We estimate in fact:

$$\log\left(\frac{f(t - t_d, t_d - t_{ad})\, h}{1 - f(t - t_d, t_d - t_{ad})\, h}\right) = \eta(t - t_d, t_d - t_{ad})\,.$$

The binned data

	event	time	neuron	lN.1	i1
14604	0	58.412	1	0.012	0.016
14605	1	58.416	1	0.016	0.016
14606	0	58.420	1	0.004	0.016
14607	1	58.424	1	0.008	0.016
14608	0	58.428	1	0.004	0.008
14609	0	58.432	1	0.008	0.008
14610	1	58.436	1	0.012	0.008
14611	0	58.440	1	0.004	0.012

event is the binned spike train; time is the time at the center of the bin; neuron is the neuron to which event "belongs"; lN.1 is $t - t_d$; i1 is $t_d - t_{ad}$. Here, *h* was set to 4 ms.

Smoothing splines

- Since cellular biophysics does not provide much guidance on how to build η we have chosen to use smoothing splines (Wahba, 1990; Green and Silverman, 1994; Eubank, 1999; Gu, 2002).
- Computations are performed with `gss` an `R` package developed by Chong Gu.
- $\eta(t-t_d, t_d-t_{ad})$ is decomposed in a unique way in:

  $$\eta(t-t_d, t_d-t_{ad}) = \eta_\emptyset + \eta_1(t-t_d) + \eta_2(t_d-t_{ad}) + \eta_{1,2}(t-t_d, t_d-t_{ad}),$$

  where the variables: $t-t_d$ and t_d-t_{ad} have been linearly transformed such their domains are both the $[0,1]$ interval.
- The decomposition is made unique by imposing conditions like: $\int_0^1 \eta_i = 0$.

Where are we ?

A remark on the tests

- Ogata's tests, like the proposed "Brownian motion test", assume that the Λ use to transform the time is independent of the data.
- But our $\hat{\Lambda}$ depends strongly on the data.
- We therefore split our data sets in two parts, fit on one part and test on the other.
- Our test level is then slightly lower than the nominal level (as explained by Reynaud-Bouret et al, 2014) since our $\hat{\Lambda}$ is slightly different (at best) from Λ.

Fit early / test late

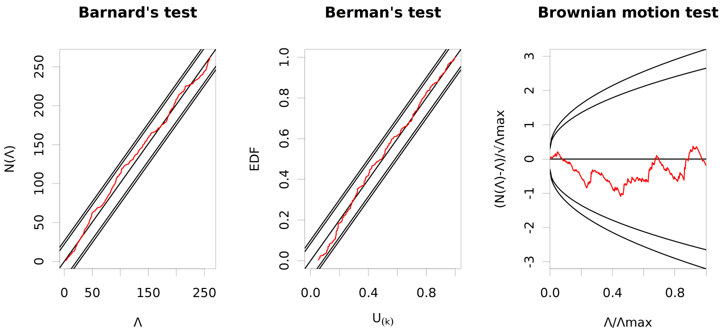

The model is:

$\eta(t-t_d, t_d-t_{ad}) = \eta_{\emptyset} + \eta_1(t-t_d) + \eta_2(t_d-t_{ad}) + \eta_{1,2}(t-t_d, t_d-t_{ad})$.

Fit late / test early

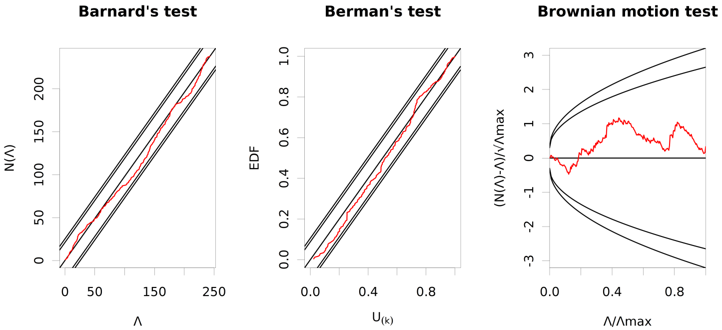

Data and $\hat{\lambda}$

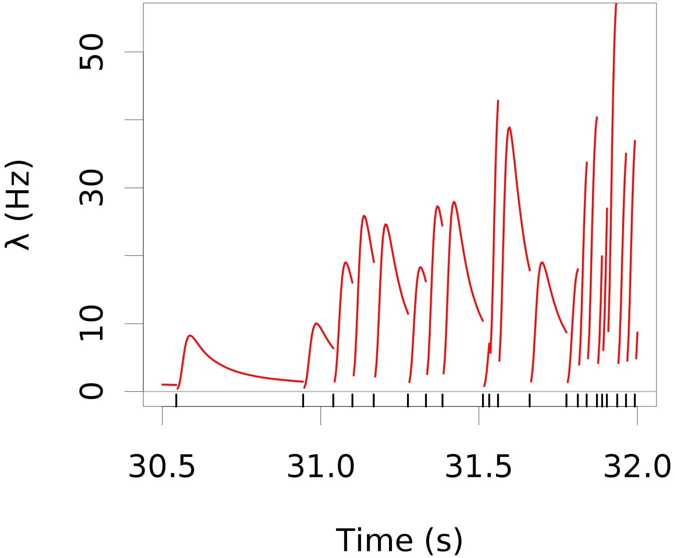

Conclusions

- We can now routinely estimate the conditional intensity of our spike trains.
- We can include interactions between neurons as well as stimulations' response in our models.
- We systematically pass much tougher tests than our competitors.
- The difficult question of model choice has not been touched upon here but we have a solution—even if computationally expensive.
- You can try all that out with the `STAR` package available on `CRAN` (a `Python` version is in development).

Thank you!

I want to thank:

- Antoine Chaffiol who recorded the data and did the spike sorting.
- Chong Gu who developed `gss`: my main collaborator on this project.
- Vilmos Prokaj, Olivier Faugeras and Jonathan Touboul who pointed out Donsker's theorem to me.
- Clément Léna and Yann-Suhan Senova for testing `STAR` on their data.
- Lyle Graham and Angelo Iollo for inviting me.
- You guys for listening.