Istituto Italiano di Tecnologia iit

Neural Coding

Stefano Panzeri

IIT, Robotics, Brain and Cognitive Sciences Dept
Laboratory of Brain Signals Analysis
stefano.panzeri@iit.it

Lecture 3

Building response models of single neurons and of neural populations

Content

- Building models of neural responses
- What drives a neuron to fire a spike?
- Spike triggered stimulus averages
- Linear models of the stimulus-response relationship
- Example of linear models and spike-triggered averages of visual cells
- LNP models
- Population models

Stimulus-response models

Stimulus-response models estimates the value of the time-dependent firing rate $r_{est}(t)$ of a neuron in response to any time-dependent ("dynamic") stimulus $s(t)$.

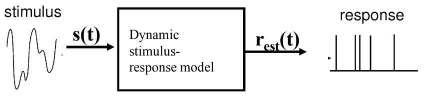

4

Stimulus-response models – why they are useful

Generalization of experimental results. The stimulus-response model allows computation of neuronal response to arbitrary type of stimuli, even stimuli never presented during the experiment. Thus one can use the stimulus-response model to predict how the neuron would respond to any complex natural stimulus.

Insights into the mechanisms generating neural representations. It gives important insights into what feature of neural population codes are important to permit the brain to form reliable sensory representations of noisy elements

Partial Replacement of use of animals

5

Ingredients needed to build a stimulus response model

Stimulus Selectivity – model of which features of the stimulus make the neuron fire

Noise – model of neural response variability at fixed stimulus

Dependence of neural response on past activity of the same neurons

Dependence of neural response on current and past activity of other neurons

6

Ingredient 1 : Stimulus selectivity

Single spikes are metabolically expensive and are not emitted without a good reason.

So, ideally we would like to know the cause and meaning of each spike

What stimulus features (or changes in stimulus features) make a neuron fire a spike?

7

Experimental technique

We present many stimuli and we record the neuronal response.

We study which stimuli are best in causing a spike.

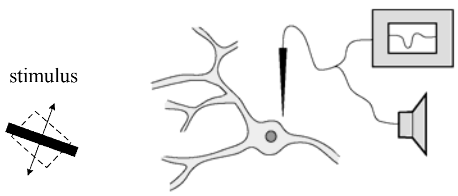

8

Single pixel stimulus: light intensity s(t)

9

What stimuli are best to determine what makes a neuron fire?

A very good type of dynamic stimulus is the “white noise stimulus”: it allows the presentation of an unbiased and wide selection of stimulus combinations.

The white noise stimulus? A “white noise stimulus” is a dynamic stimulus whose value at any time t is not related to the value of the stimulus at any previous or future time.

A white noise stimulus is essentially a “random stimulus” without any structure.

10

What are the stimulus features that evoke a spike?

What makes a neuron fire is a certain pattern of stimulation that is applied shortly before the spike.

We have to look back at what was the stimulus shortly before each spike was emitted, and check whether there were some features of the stimulus that always (or often) precede the spike.

11

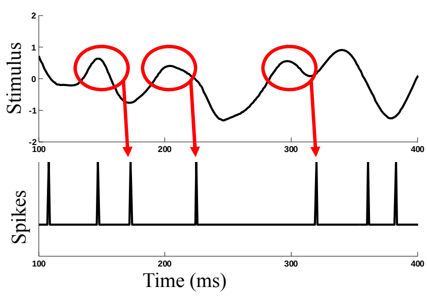

12

Spike triggered average

The spike-triggered average stimulus, C(τ), is the average value of the stimulus at a time interval τ before a spike is fired.

$$C(\tau) = \left\langle \frac{1}{n} \sum_{i=1}^{n} s(t_i - \tau) \right\rangle$$

13

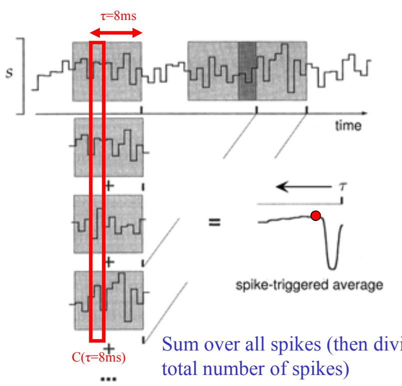

Sum over all spikes (then divided by the total number of spikes)

14

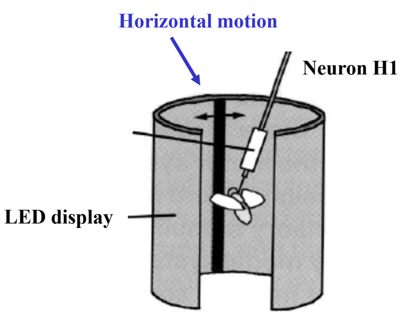

Neurophysiological experiment

15

H1 neuron of a fly (motion sensitive visual neuron).
Spike triggered stimulus average

Spikes are triggered by changes in velocity (i.e. accelerations) happening approximately 25 ms before the spike.

Dynamic stimulus-response models and stimulus history

The firing rate r(t) of a neuron at any given time t depends on the behaviour of a stimulus over a period of time starting a few hundreds of ms before t (short neuronal time constant) and ending a few ms before t (latency).

Stimuli at different times may have a different impact of the current firing rate of the neuron.

How to model the impact on the neuron firing rate of stimulus values at different times in the past?

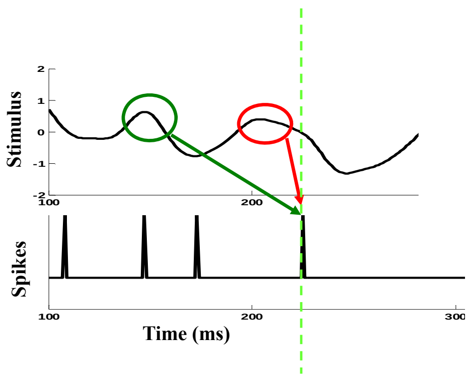

LINEAR dynamic stimulus-response models

The estimated firing rate $r_{est}(t)$ at any given time is obtained by summing over the past-time values taken by the stimulus.

To take into account that stimuli at different times may have a different impacts, the sum over stimuli happening at different past times τ is weighted by a factor D(τ).

The higher is D(τ), the higher is the impact on current neuronal firing of a stimulus happening τ ms ago.

19

LINEAR dynamic stimulus-response models

The estimated firing rate $r_{est}(t)$ at any given time is a weighted **sum** of the values taken by the stimulus at earlier times.

$$r_{est}(t) = r_0 + L(t)$$

$$L(t) = \sum_{\tau} D(\tau)s(t-\tau)$$ **(Linearly filtered stimulus)**

r_0 = **background firing rate** (response that occurs even when there is no stimulus, i.e. s(t)=0
D(τ) = **temporal filter**. It is a weighting factor that determines how strongly, and with what sign, the value of the stimulus at time t-τ affects the firing rate at time t.

20

How to determine the best choice of the temporal filter D

$$r_{est}(t) = r_0 + \sum_{\tau} D(\tau)s(t-\tau)$$

D(τ) = **temporal filter**. How to choose this?

Chose as D(τ) the filter that gives an estimated model firing rate $r_{est}(t)$ which is as close as possible (in the least-squares sense) to the true neuronal r(t) for all possible dynamic stimuli.

21

The best possible linear model

Theorem: If the stimulus is of the "white noise" type, the best possible linear model (the one that minimizes the error between real and predicted responses) is the one that uses as temporal filter the spike triggered stimulus average.

Thus, the spike triggered stimulus average (measured with white noise) is the optimal temporal filter for building a stimulus-response model.

The spike-triggered stimulus average is useful not only to characterize the neuronal responses (the "average" stimulus causing a spike) but also to build dynamic models of stimulus-response.

22

Recipe to build the optimal linear dynamic stimulus-response models

1) Compute the spike-triggered average C(τ) using a white noise stimulus
2) Replace the temporal filter D(τ) in the linear model with the spike-triggered average (which is the optimal filter)
3) Compute the expected time-dependent firing rate according to the linear equation as follows:

$$r_{est}(t) = r_0 + L(t)$$

$$L(t) = \sum_{\tau} C(\tau) s(t-\tau)$$

(Linearly filtered stimulus)

23

Linear decoder of the stimulus

Linearly decoded stimulus is obtained by summing the spike trigger averages at the times of spikes

$$s_{rec}(t) = \sum_{i} s(t_i - \tau)$$

24

Stimulus-response models of visual cells

Dynamic stimulus-response models of visual cells predict the value of the time-dependent firing rate r(t) of a neuron in response to a dynamic visual stimulus s(t) (i.e. movies!!!).

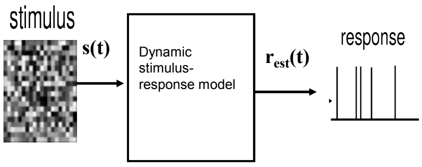

Dynamic visual stimuli (movies)

Another example of dynamic stimulus is a movie displayed on a computer screen.

A movie is described by the light intensity of each pixel (x,y coordinates) at a certain instant of time t:

Movie Stimulus = s(x,y,t)

The stimulus at each time is an image (a frame of the video)

The spike-triggered stimulus average is the average image that preceded the spike by a certain amount of time τ

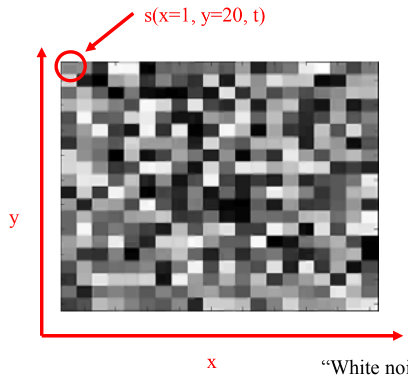

"White noise" stimulus

Spike triggered average with a movie and linear dynamic models of visual neurons.

The spike-triggered average when using a movie stimulus is the average image that precedes the spikes by a certain amount of time τ.
Being an image, the spike-triggered average is a function of the pixel coordinates x and y at any given time τ before the spike.

$$C(x, y, \tau)$$

The linear dynamic model of a visual cell is obtained as before using the spike-triggered image as temporal filter.

28

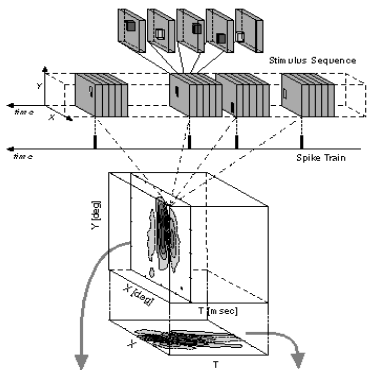

29

Separable spatio-temporal receptive fields

$$C(x, y, \tau) = C_S(x, y) C_T(\tau)$$

$$C_S(x, y) =$$

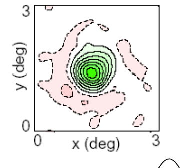

$$C_T(\tau) =$$

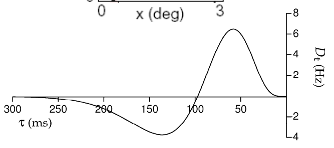

30

LINEAR dynamic stimulus-response models of visual cells

The estimated firing rate $r_{est}(t)$ at any given time is obtained by **summing** over the past values taken by the movies with the spike triggered stimulus acting as a weighting function.
We assume that contributions from different spatial locations sum linearly (spatial summation principle).

$$r_{est}(t) = r_0 + L(t)$$

$$L(t) = \sum_{\tau,x,y} C(x,y,\tau)\, s(x,y,t-\tau)$$

31

Generate the single trial noisy neuronal response

Dynamic stimulus-response models estimates the value of the time-dependent firing rate $\mathbf{r_{est}(t)}$ of a neuron in response to a dynamic stimulus $\mathbf{s(t)}$.
To generate a spike response with the estimated firing rate $r_{est}(t)$, we can use a Poisson spike generator .

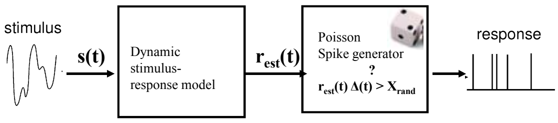

32

The Poisson spike generator

A simple procedure for generating spikes in a computer program is based on the fact that the estimated probability of firing a spike during a short interval of duration Δt is $r(s;t)\Delta t$.

The program progress through time in small steps of size Δt (e.g. Δt = 1ms) and generates, at each time step, a random number x_{rand} chosen uniformly in the range between 0 and 1.

If $r(s;t)\Delta t > x_{rand}$, then a spike is generated.

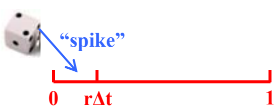

33

Good and bad feature of Poisson processes

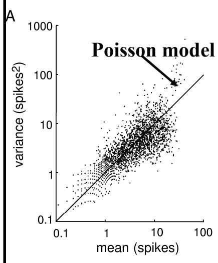

The Poisson spike generation model predicts that the variance of spike counts is equal to the average spike count itself:
Var(n) = <n>

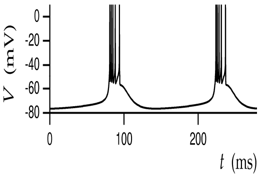

The Poisson spike generation model cannot account for temporal relationships between spikes (e.g. Bursts)

34

Including simple nonlinearities

Neurons are non-linear devices, so the linear approximation is only an approximation, albeit a useful one.
The linear prediction has two obvious problems:

There is nothing to prevent the estimated firing rate to become negative (impossible, because firing rates can only be >=0)

The predicted rate does not saturate, but instead increase without bound as the magnitude of the stimulus increases. Thus, in the linear model, firing rates can be infinite (which is impossible!)

35

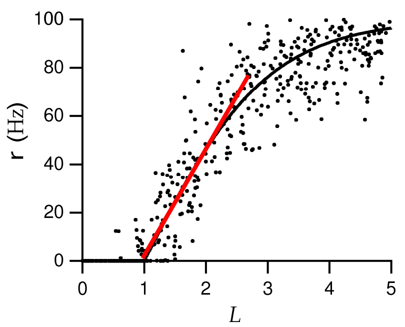

36

STATIC NON-LINEARITY

The estimated firing rate $r_{est}(t)$ is expressed as the background rate r_0 plus a non-linear function of the linearly filtered stimulus.

$$r_{est}(t) = r_0 + F(L(t))$$

$$L(t) = \sum_{\tau} D(\tau)s(t-\tau)$$ **(Linearly filtered stimulus)**

The function F is called static non-linearity to stress that it is the same type of distortion of the linear filter at all times.
If F is a bounded function, it will prevent rest to become negative or unrealistically large.

37

Graphical procedure to determine the shape of the static non-linearity function F from the data

$$r_{est}(t) = r_0 + F(L(t))$$

$$L(t) = \sum_{\tau} C(\tau)s(t-\tau)$$ **(Linearly filtered stimulus)**

1. Compute the optimal linearly filtered stimulus L(t) (using the linear equation above).
2. Plot all datapoints (L(t),r(t)) at various times and for various stimuli, where r(t) is the real neuronal firing rate.
3. From the scatterplot of the datapoints above, find a function F(L) that fits the data well.

38

Summary of Linear Non-linear Poisson (LNP) models of neural responses to stimuli

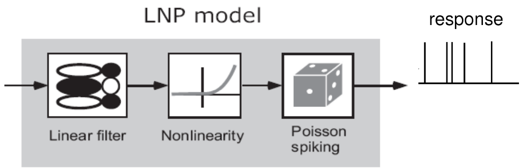

Paninski (2003)

Introducing dependencies on other neurons or on previous spiking history

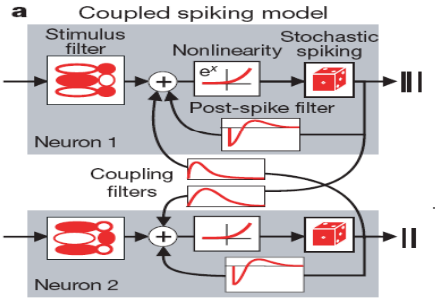

Pillow et al (2008) Nature

Coupled models

$$r^i_{est}(t) = r_0 + F\left(L(t)\right)$$

$$L(t) = \underbrace{\sum_{\tau} C(\tau)s(t-\tau)}_{linearly-filtered-stimulus} +$$

$$+\underbrace{\sum_{\tau} h(\tau)n^i(t-\tau)}_{post-spike-filters} + \underbrace{\sum_{j,\tau} m^{j\to i}(\tau)n^j(t-\tau)}_{interaction-terms}$$

41

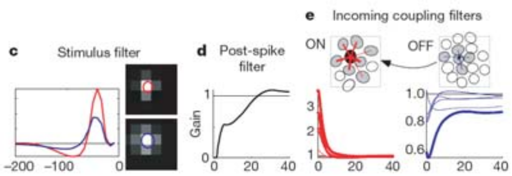

Pillow et al (2008) Nature

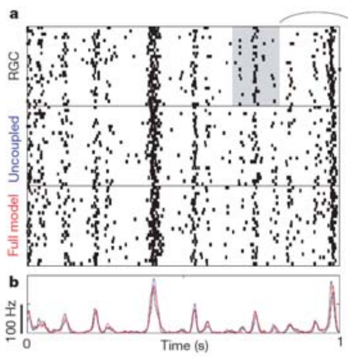

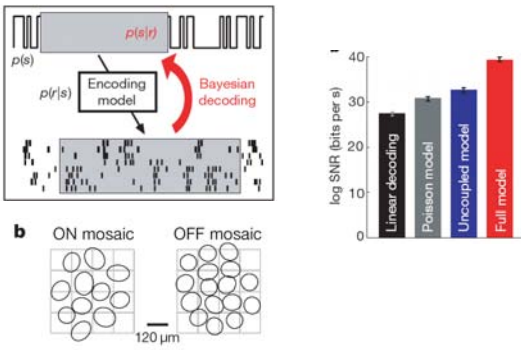

Pillow et al (2008) Nature

Summary

Building simple but satisfactory mathematical models of neural responses is possible

These models can be used to learn how populations of neurons work together to represent complex stimuli

45