Introduction to Neural Computation

Prof. Michale Fee
MIT BCS 9.40 — 2018

Lecture 9 — Receptive fields

Spatial receptive fields

LIGHT

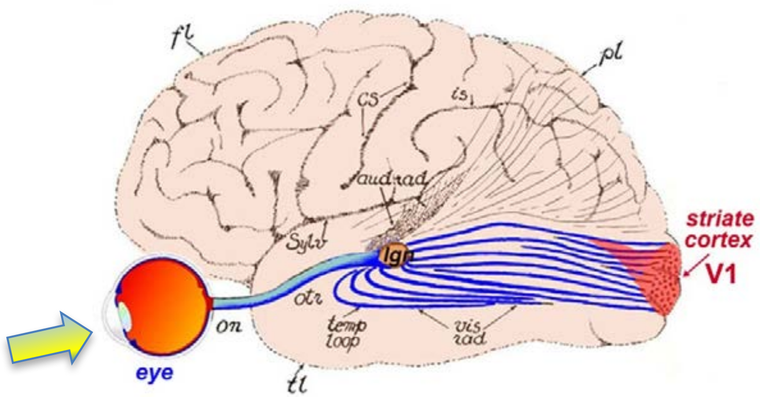

Video of visual neurons of a cat from Hubel & Wiesel's experiments.

Ali Moeeny. "Hubel & Wiesel – LGN Neuron." April 23, 2011. YouTube.

Spatial receptive fields

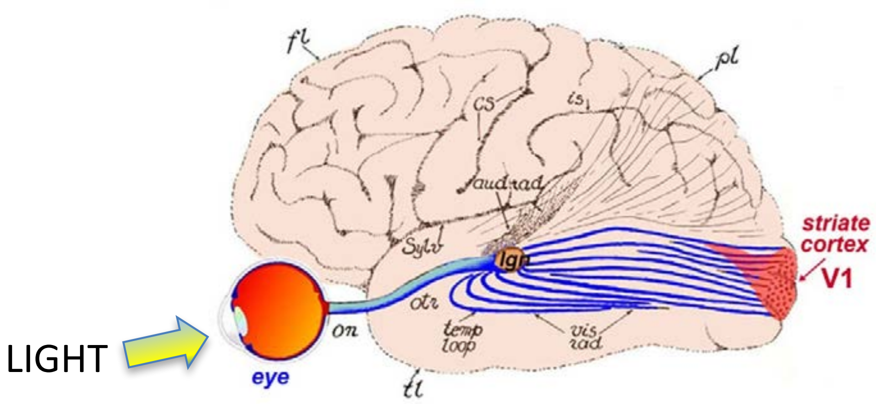

Video of visual neurons of a cat – simple and complex cells - from Hubel & Wiesel's experiments.

Ali Moeeny. "[Hubel & Wiesel – Cortical Neuron – V1]." April 23, 2011. YouTube.

Figure (annotated) from [Introduction to Visual Prostheses] on Webvision. License CC BY-NC.

Learning objectives for Lecture 9

- To be able to mathematically describe a neural response as a linear filter followed by a nonlinear function.
  - A correlation of a spatial receptive field with the stimulus
  - A convolution of a temporal receptive field with the stimulus

- To understand the concept of a Spatio-temporal Receptive Field (STRF) and the concept of ‘separability’

- To understand the idea of a Spike Triggered Average and how to use it to compute a Spatio-temporal Receptive Field and a Spectro-temporal Receptive Field (STRF).

Learning objectives for Lecture 9

- To be able to mathematically describe a neural response as a linear filter followed by a nonlinear function.
  - A correlation of a spatial receptive field with the stimulus
  - A convolution of a temporal receptive field with the stimulus

- To understand the concept of a Spatio-temporal Receptive Field (STRF) and the concept of ‘separability’

- To understand the idea of a Spike Triggered Average and how to use it to compute a Spatio-temporal Receptive Field and a Spectro-temporal Receptive Field (STRF).

Spatial receptive fields

- How do we represent receptive fields mathematically?
- At the simplest level, we think of the receptive field (RF) as the region of visual space that causes the neuron to spike.
- But a visual neuron doesn't respond to any stimulus within this RF. It responds selectively to certain 'features' in the stimulus.
- We can think of a neuron as having a filter (G) that passes certain features in both space and time.

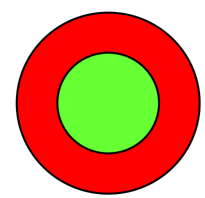

- The better the stimulus 'overlaps' with the filter, the more the neuron will spike.

Spatial receptive fields

- How do we represent receptive fields mathematically?

Start by describing the spatial part of this filter.

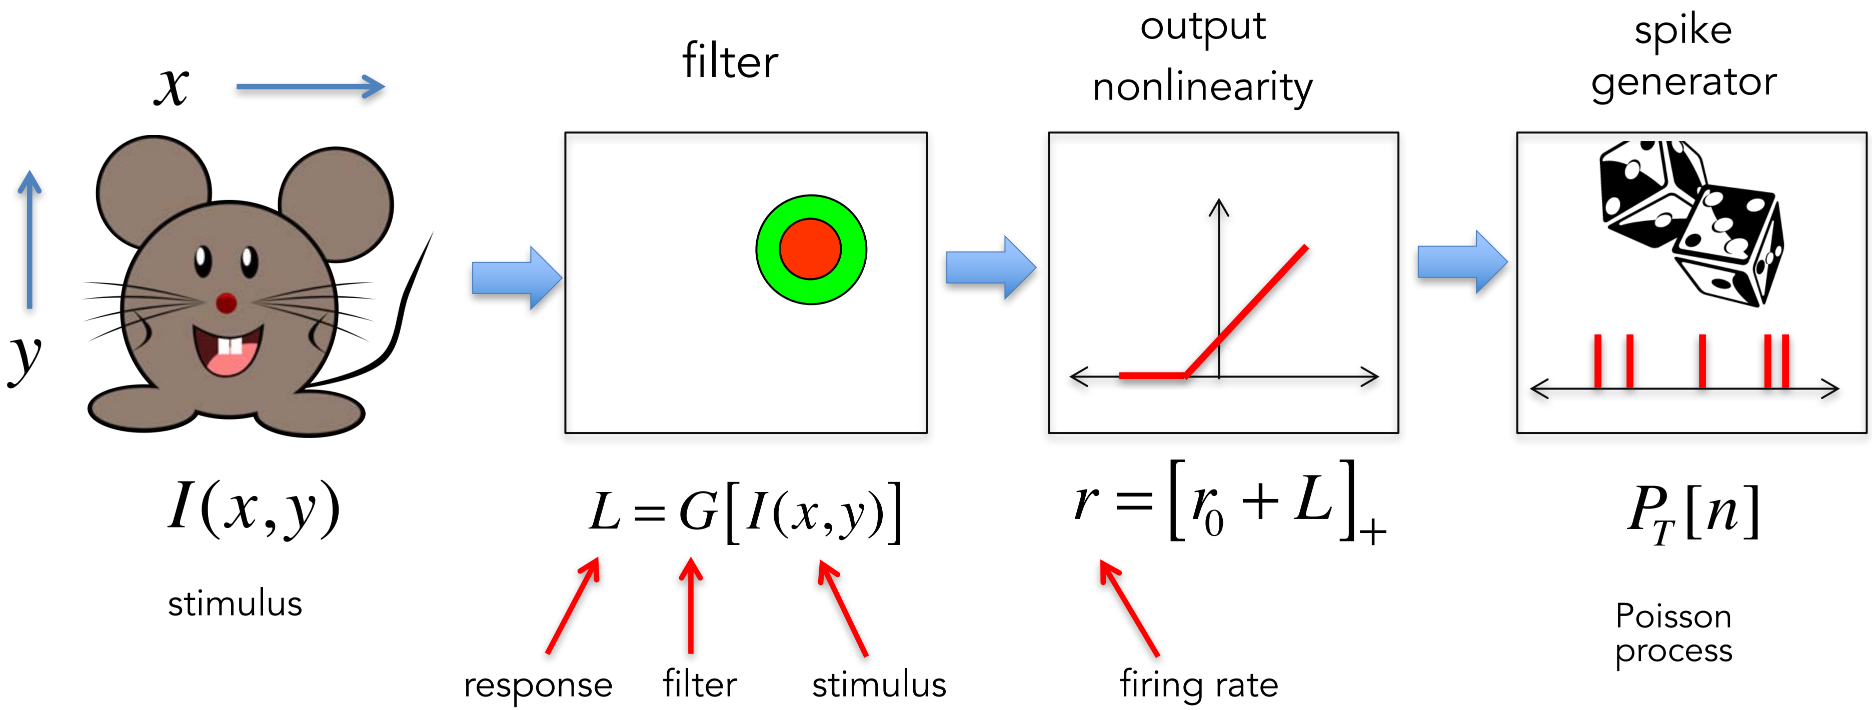

Spatial receptive fields

- How do we represent receptive fields mathematically?

We are going to consider the simplest case in which the response of a neuron is given by a linear filter acting on the stimulus.

$$r = r_0 + \iint G(x,y) I(x,y)\, dx\, dy$$

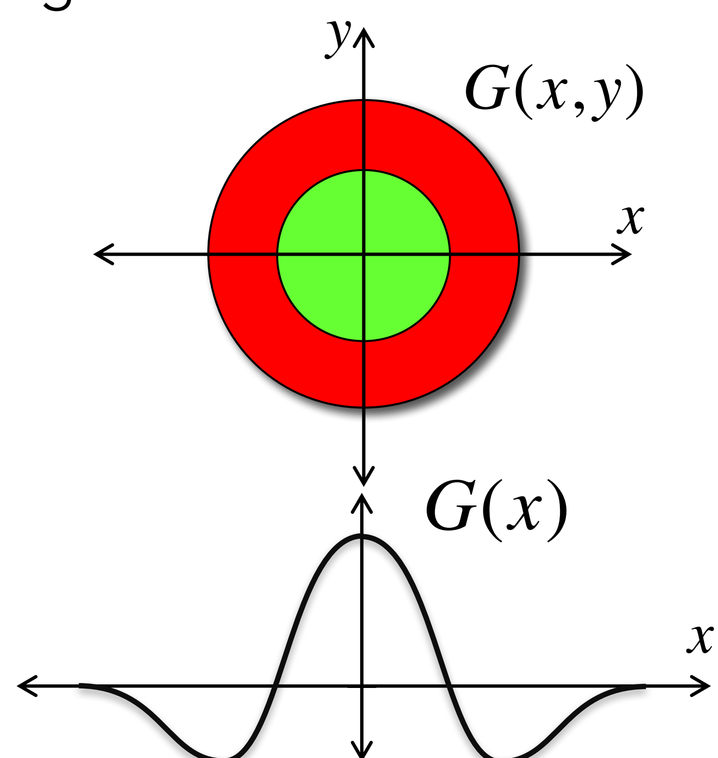

Let's look at this in one dimension

$$r = r_0 + \int G(x) I(x)\, dx$$

Like a correlation $\sum_i G_i I_i$

Spatial receptive fields

- How do we represent receptive fields mathematically?

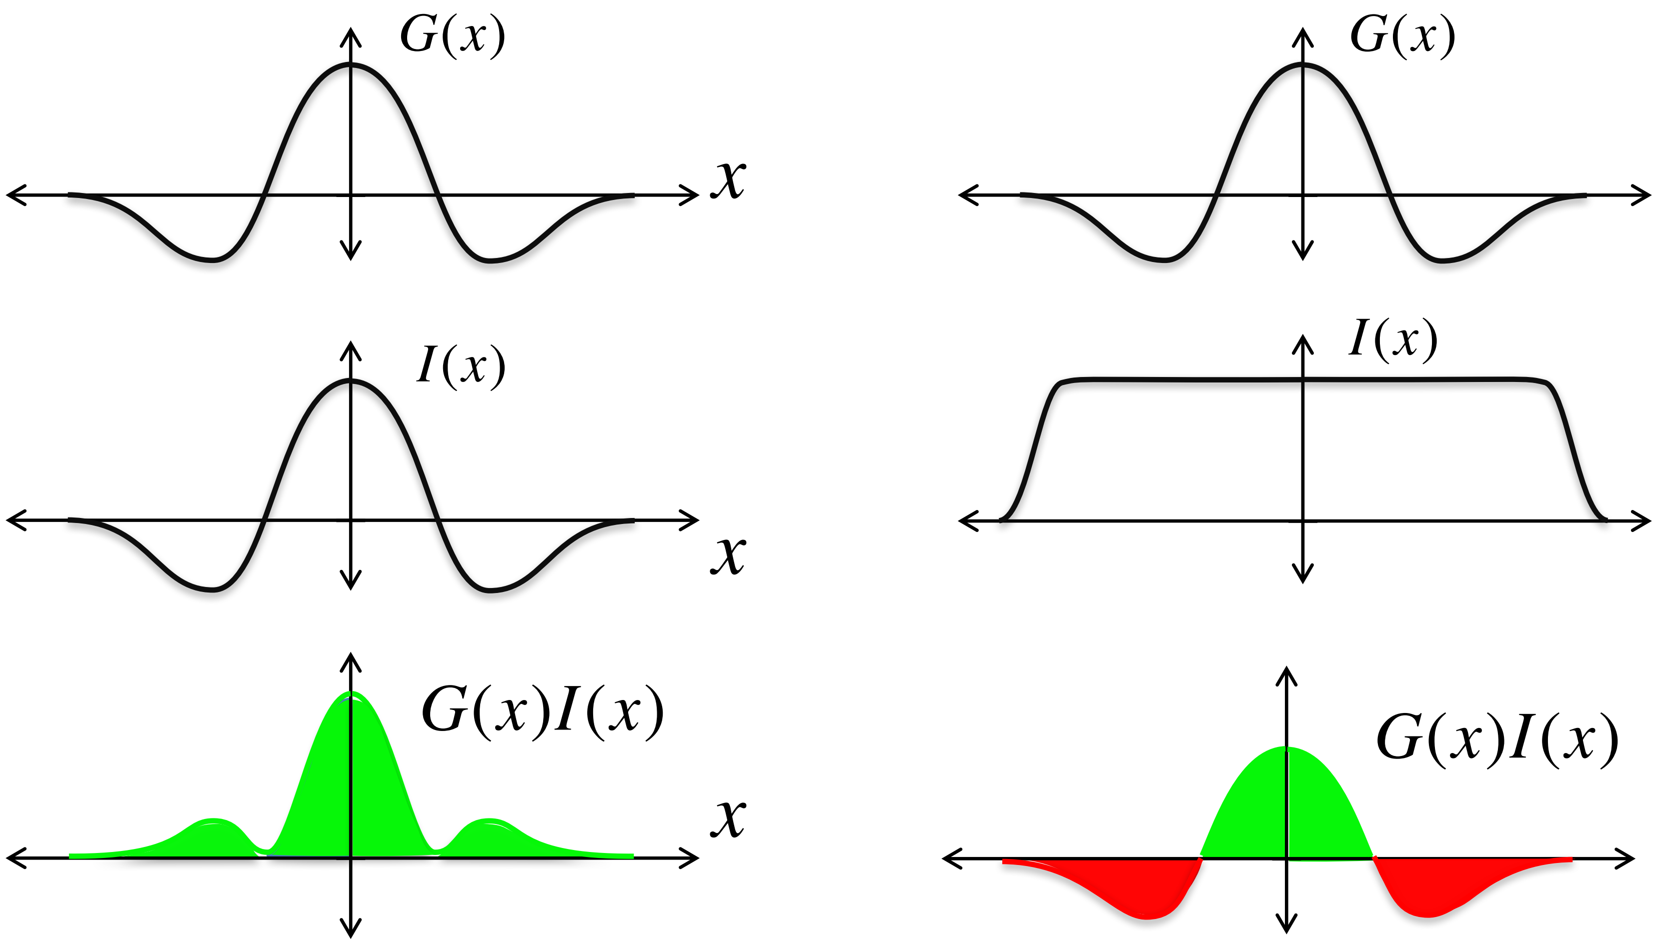

$\int G(x)I(x)dx \quad big$ $\qquad\qquad$ $\int G(x)I(x)dx \quad small$

Linearity

- Response varies linearly with overlap

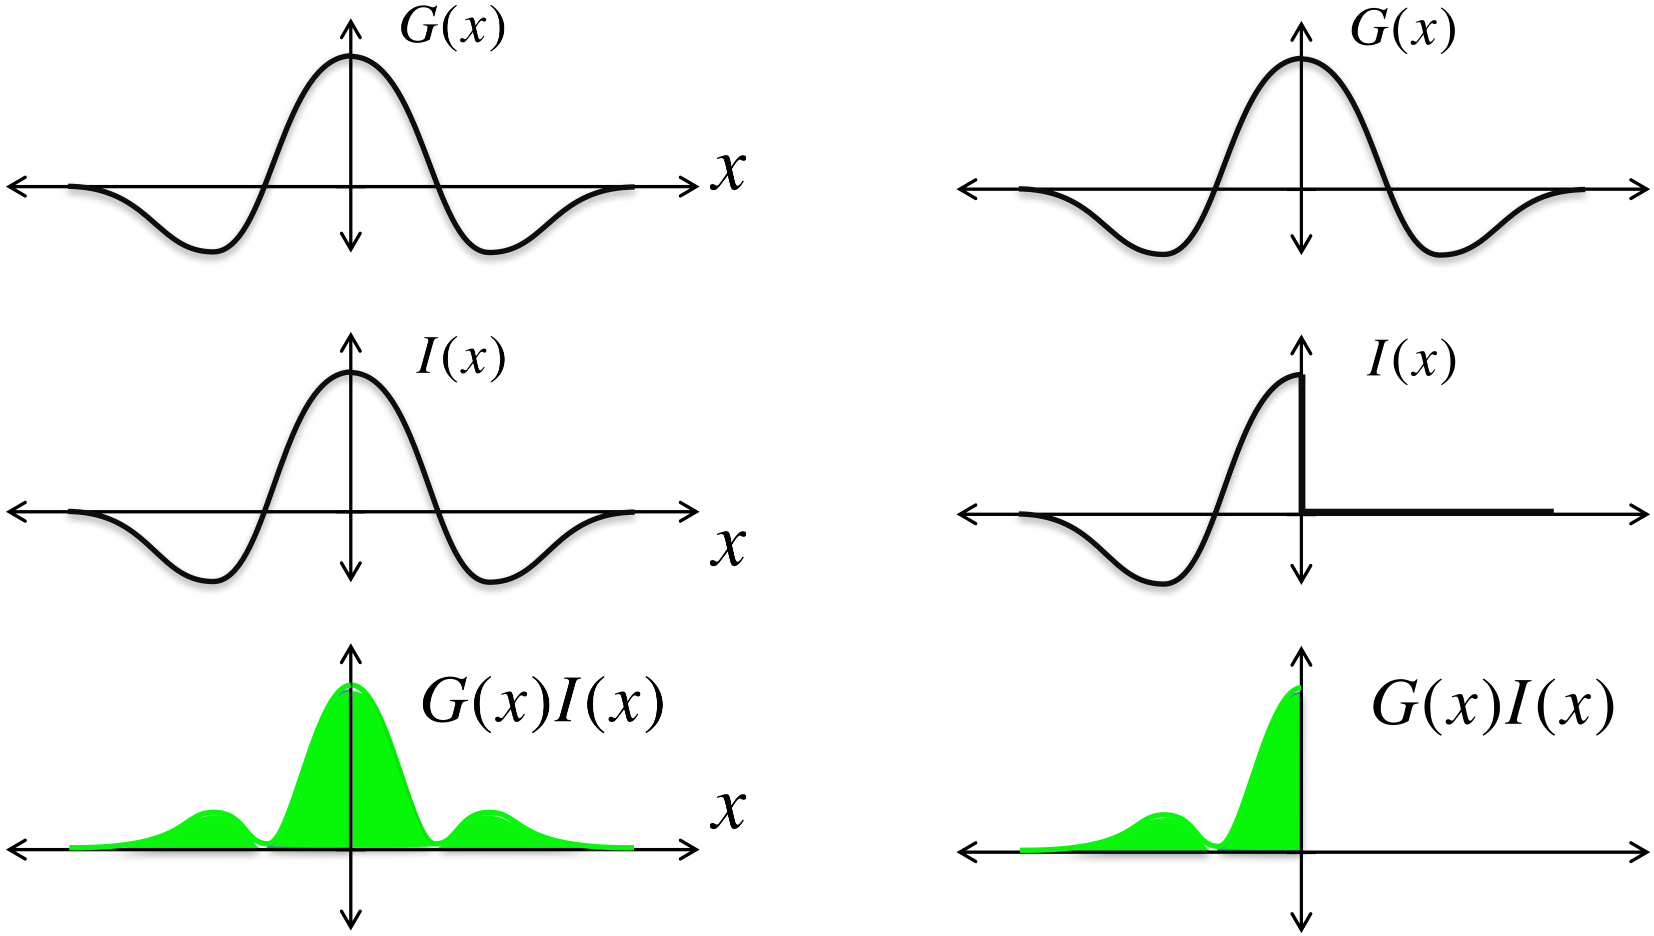

$\int G(x)I(x)dx \quad big$

$\int G(x)I(x)dx \quad half\ as\ big$

Temporal receptive fields

- We can also think of the response of a neuron as some function of the temporal variations in the stimulus.

$$r(t) = r_0 + D\left[\mathbb{S}(t)\right]$$

Time-dependent firing rate

Spontaneous firing rate

Filter

Stimulus

Temporal receptive fields

- We can think of ‘overlap’ in the time domain! That there is a particular ‘temporal profile’ of a stimulus that makes a neuron spike.

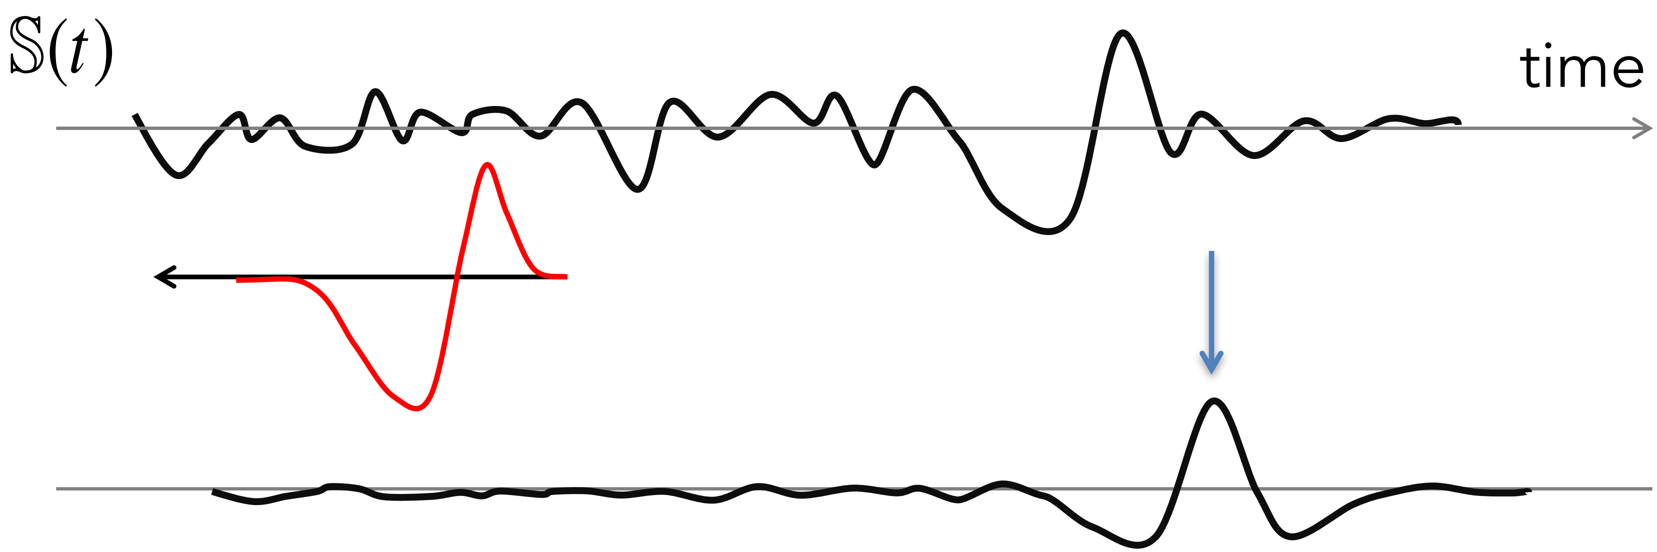

Does this look familiar?

Temporal receptive fields

Convolution!! $r(t) = r_0 + \int_{-\infty}^{\infty} D(\tau)\mathbb{S}(t-\tau)\,d\tau$

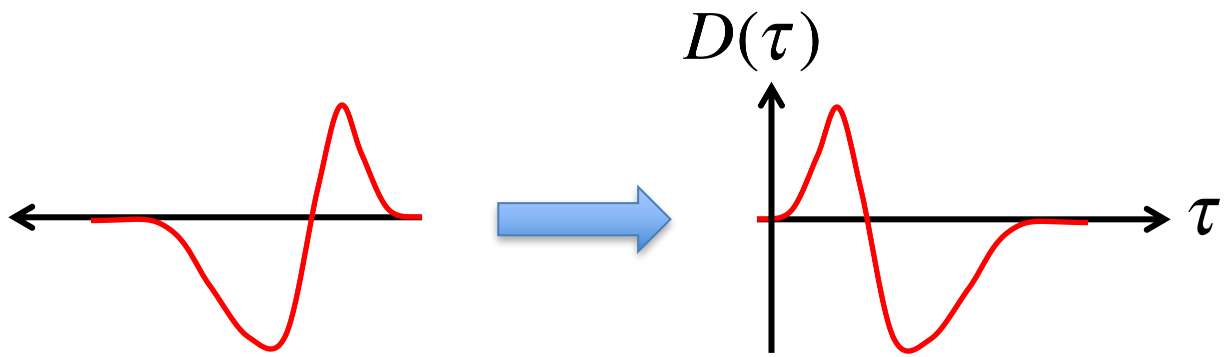

Linear temporal response kernel. (Or ‘temporal kernel’)

It is linear in the sense that if we make the stimulus partial, or weaker, the response changes linearly.

Learning objectives for Lecture 9

- To be able to mathematically describe a neural response as a linear filter followed by a nonlinear function.
  - A correlation of a spatial receptive field with the stimulus
  - A convolution of a temporal receptive field with the stimulus
- To understand the concept of a Spatio-temporal Receptive Field (STRF) and the concept of ‘separability’
- To understand the idea of a Spike Triggered Average and how to use it to compute a Spatio-temporal Receptive Field and a Spectro-temporal Receptive Field (STRF).

Spatio-temporal receptive fields

- Now we are going to put the temporal receptive field and the spatial receptive field together in a single object.

- This is called the spatio-temporal receptive field (STRF).

- Let's imagine a stimulus that is a function of space and time, like the light falling on a retina: $I(x, y, t)$

- But now we are going to simplify things by considering only one spatial dimension: $I(x, t)$

$$r(t) = r_0 + \int_{-\infty}^{\infty} dx\, d\tau\, D(x, \tau) I(x, t-\tau)$$

Spatio-temporal receptive fields

Here we are doing a correlation and a convolution at the same time! Correlation in the integral over space and a convolution in the integral over time!

$$r(t) = r_0 + \iint dx\, d\tau\, D(x,\tau) I(x,t-\tau)$$

correlation

$$r(t) = r_0 + \int_{-\infty}^{\infty} d\tau \int_{-\infty}^{\infty} dx\, D(x,\tau) I(x,t-\tau)$$

convolution

Separability

- If a receptive field is separable in space and time, then we can decompose it into a spatial receptive field and a temporal receptive field.

Separable

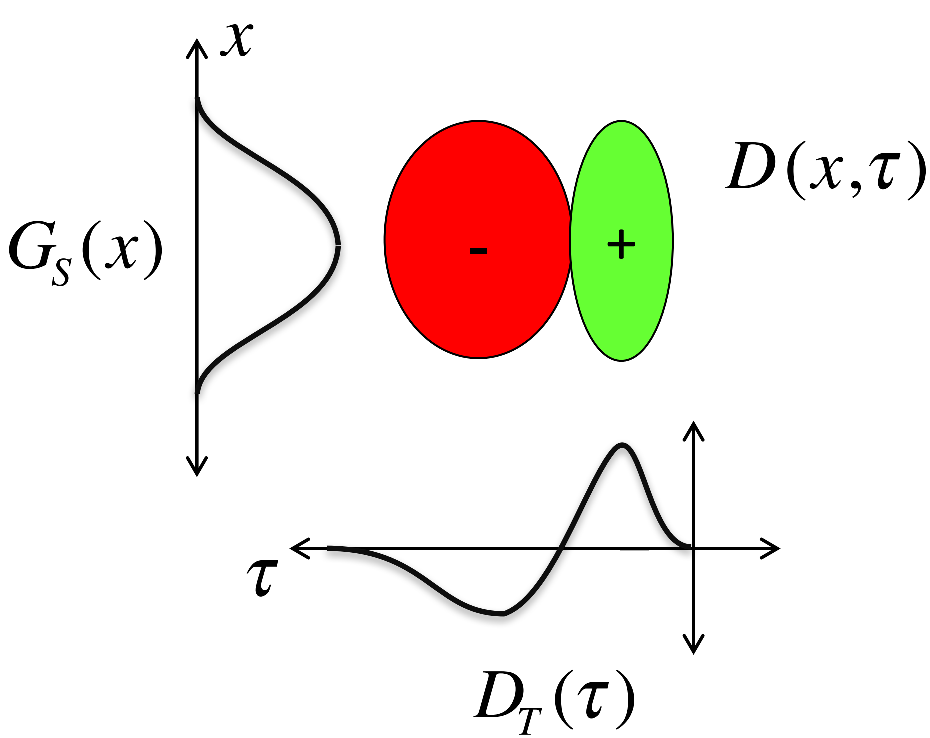

Inseparable

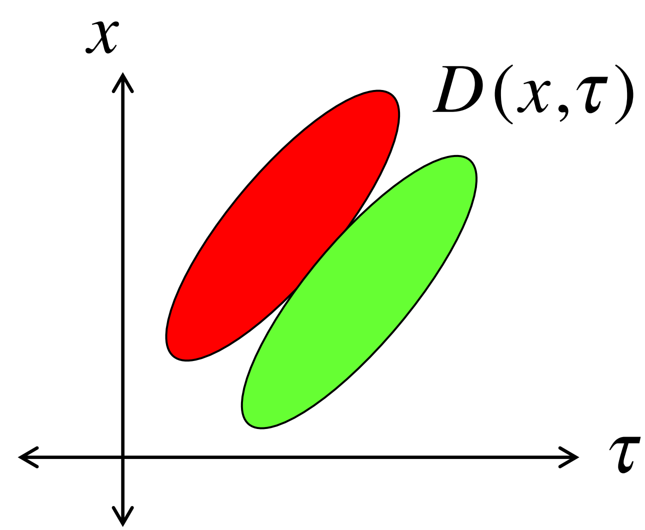

Separability

- If a receptive field is separable in space and time, then we can decompose it into a spatial receptive field and a temporal receptive field:

$$r(t) = r_0 + \int_{-\infty}^{\infty} dx \int_{-\infty}^{\infty} d\tau D(x,\tau) I(x, t-\tau)$$

$$D(x,\tau) = G_S(x) D_T(\tau)$$

$$\mathbb{S}(t) = \int_{-\infty}^{\infty} dx\, G(x) I(x,t)$$

Correlation

where

$$r(t) = r_0 + \int_{-\infty}^{\infty} d\tau\, D_T(\tau) \mathbb{S}(t-\tau)$$

Convolution

Representing stimulus and receptive fields in space and time

Suppose our stimulus is a bar of light extending from x=2 to x=4 and that is turned on for times from t=1 to t=6

We represent it on a space-time plot as follows:

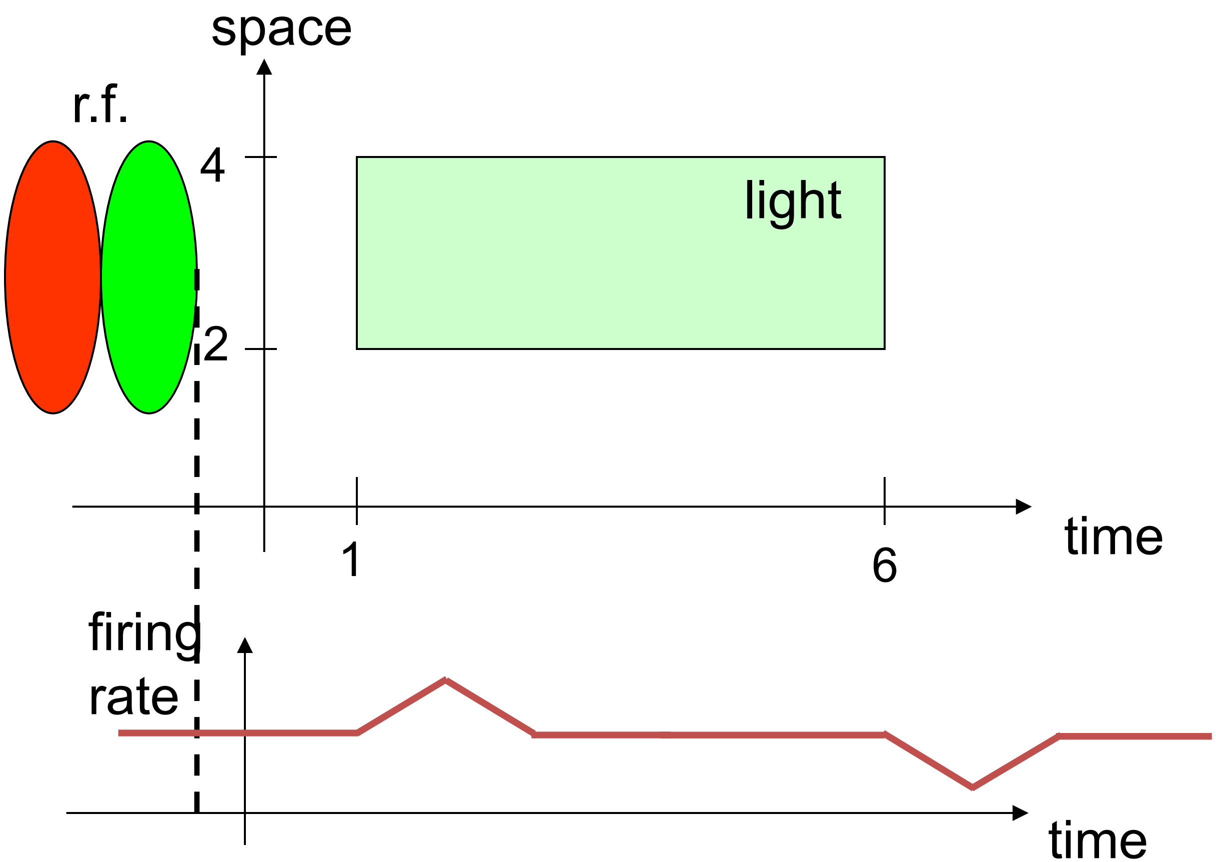

Representing stimulus and receptive fields in space and time

Suppose we now consider a receptive field D(t,x) with spatial as well as temporal structure (but 'space-time separable': D(t,x) = D(t)D(x))

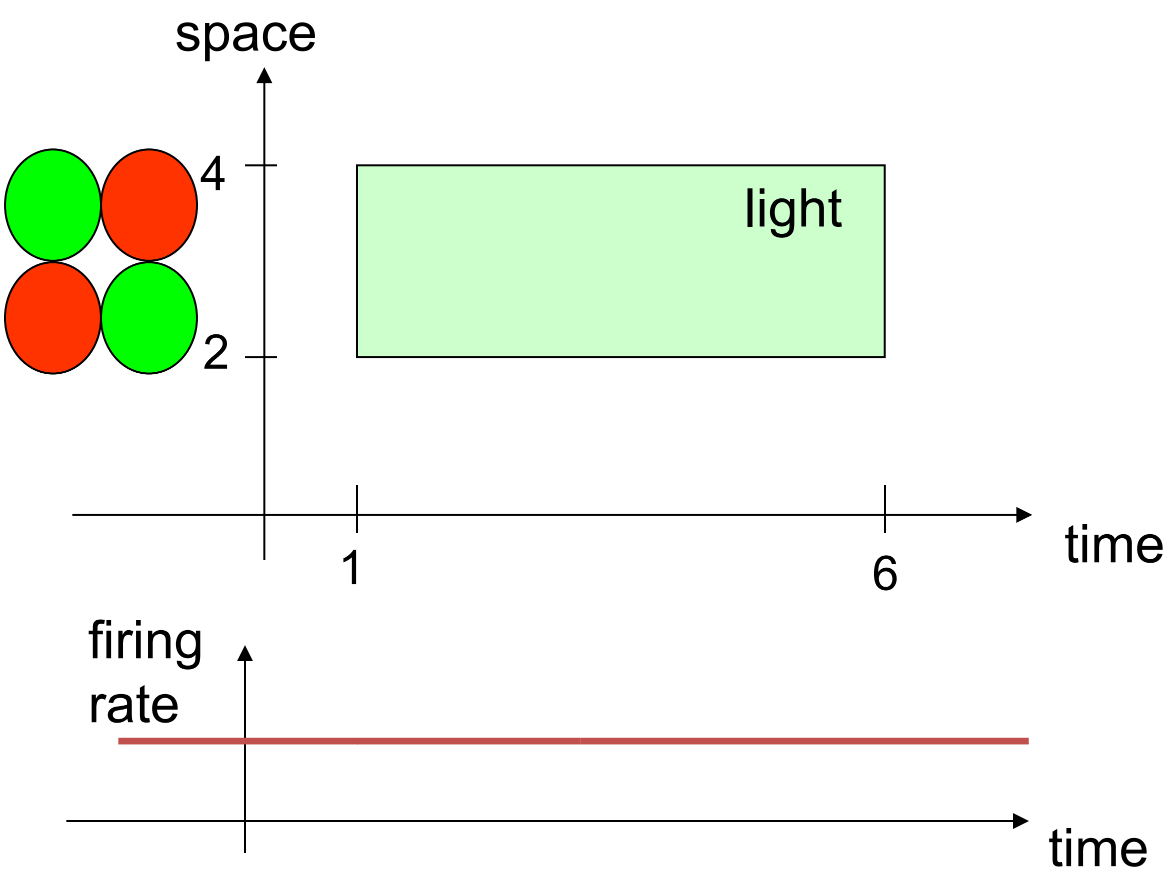

Response to a moving bar of light

Now suppose our stimulus is moving:

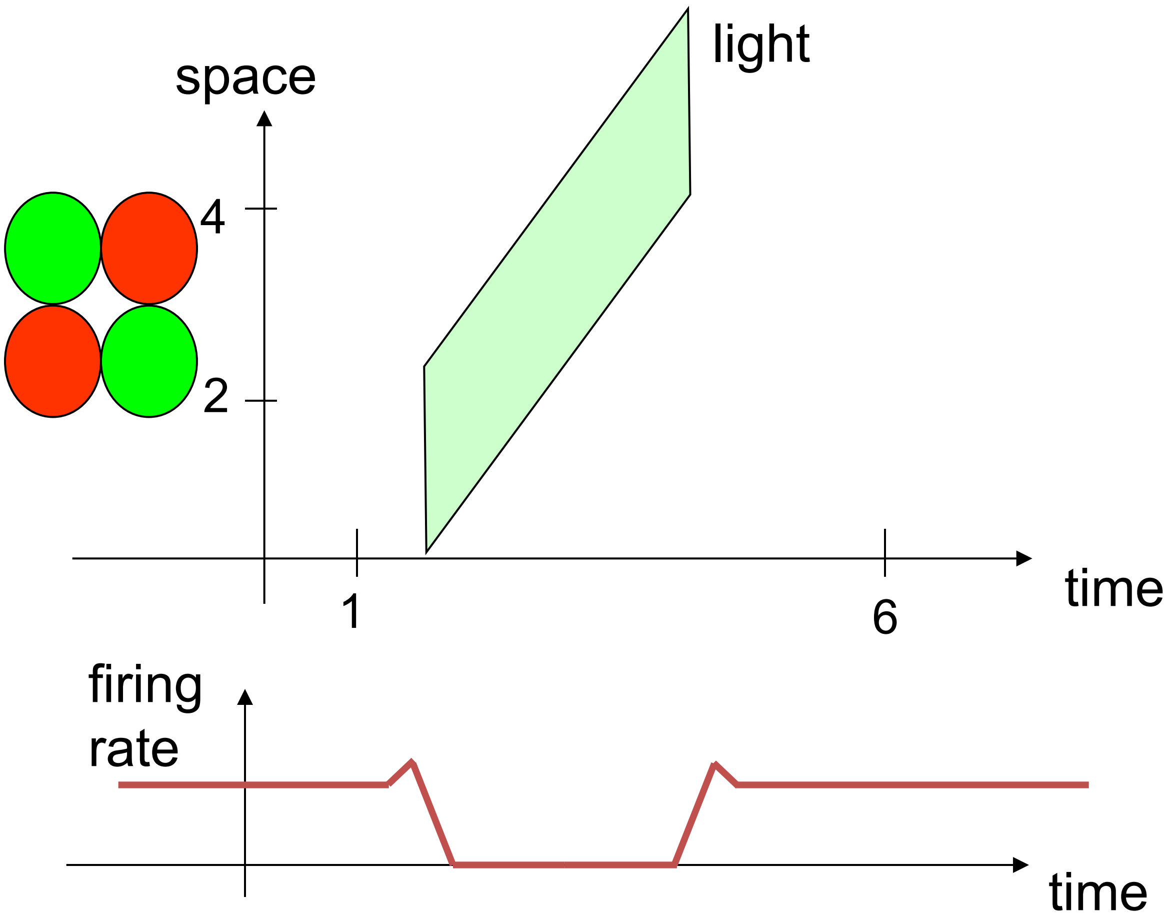

Response to a moving bar of light

Now suppose our stimulus is moving in opposite direction:

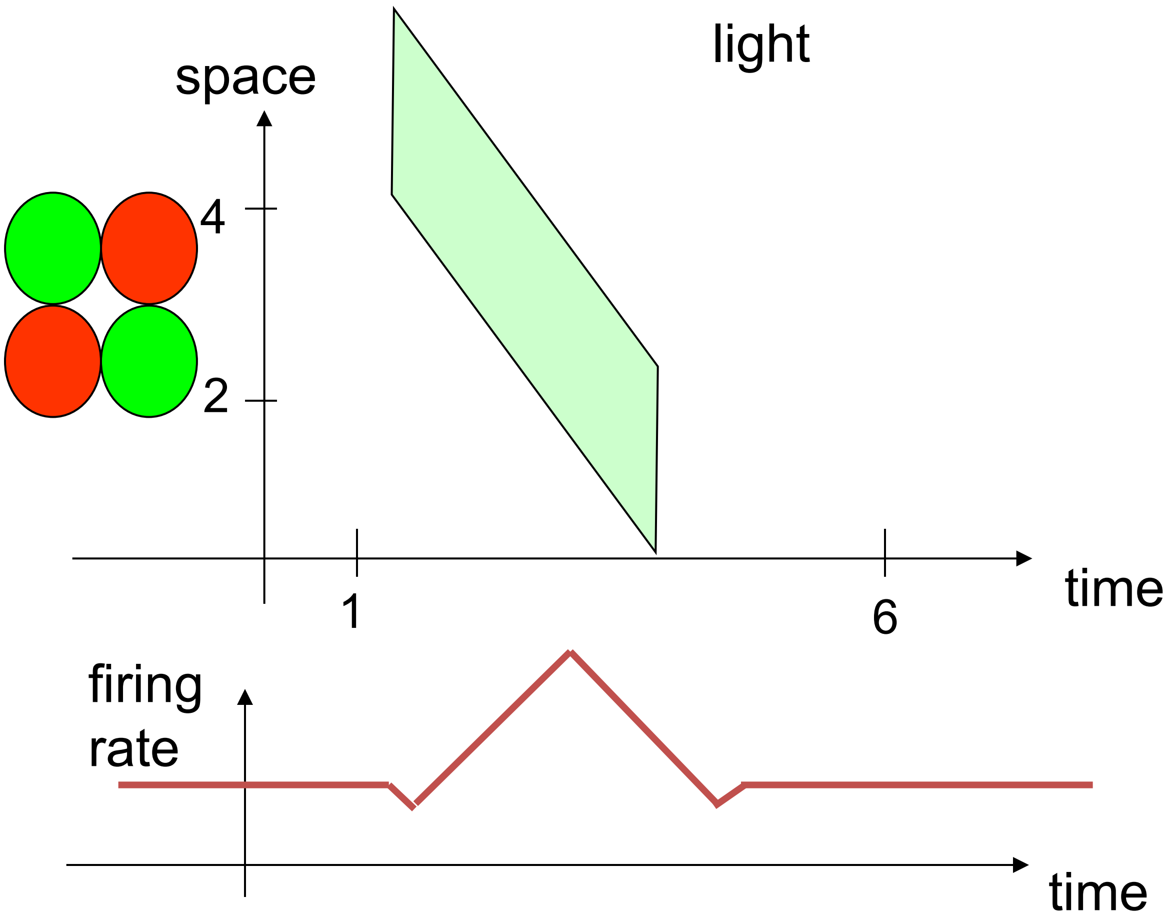

Learning objectives for Lecture 9

- To be able to mathematically describe a neural response as a linear filter followed by a nonlinear function.
  - A correlation of a spatial receptive field with the stimulus
  - A convolution of a temporal receptive field with the stimulus

- To understand the concept of a Spatio-temporal Receptive Field (STRF) and the concept of 'separability'

- To understand the idea of a Spike Triggered Average and how to use it to compute a Spatio-temporal Receptive Field and a Spectro-temporal Receptive Field (STRF).

Spike-Triggered Average

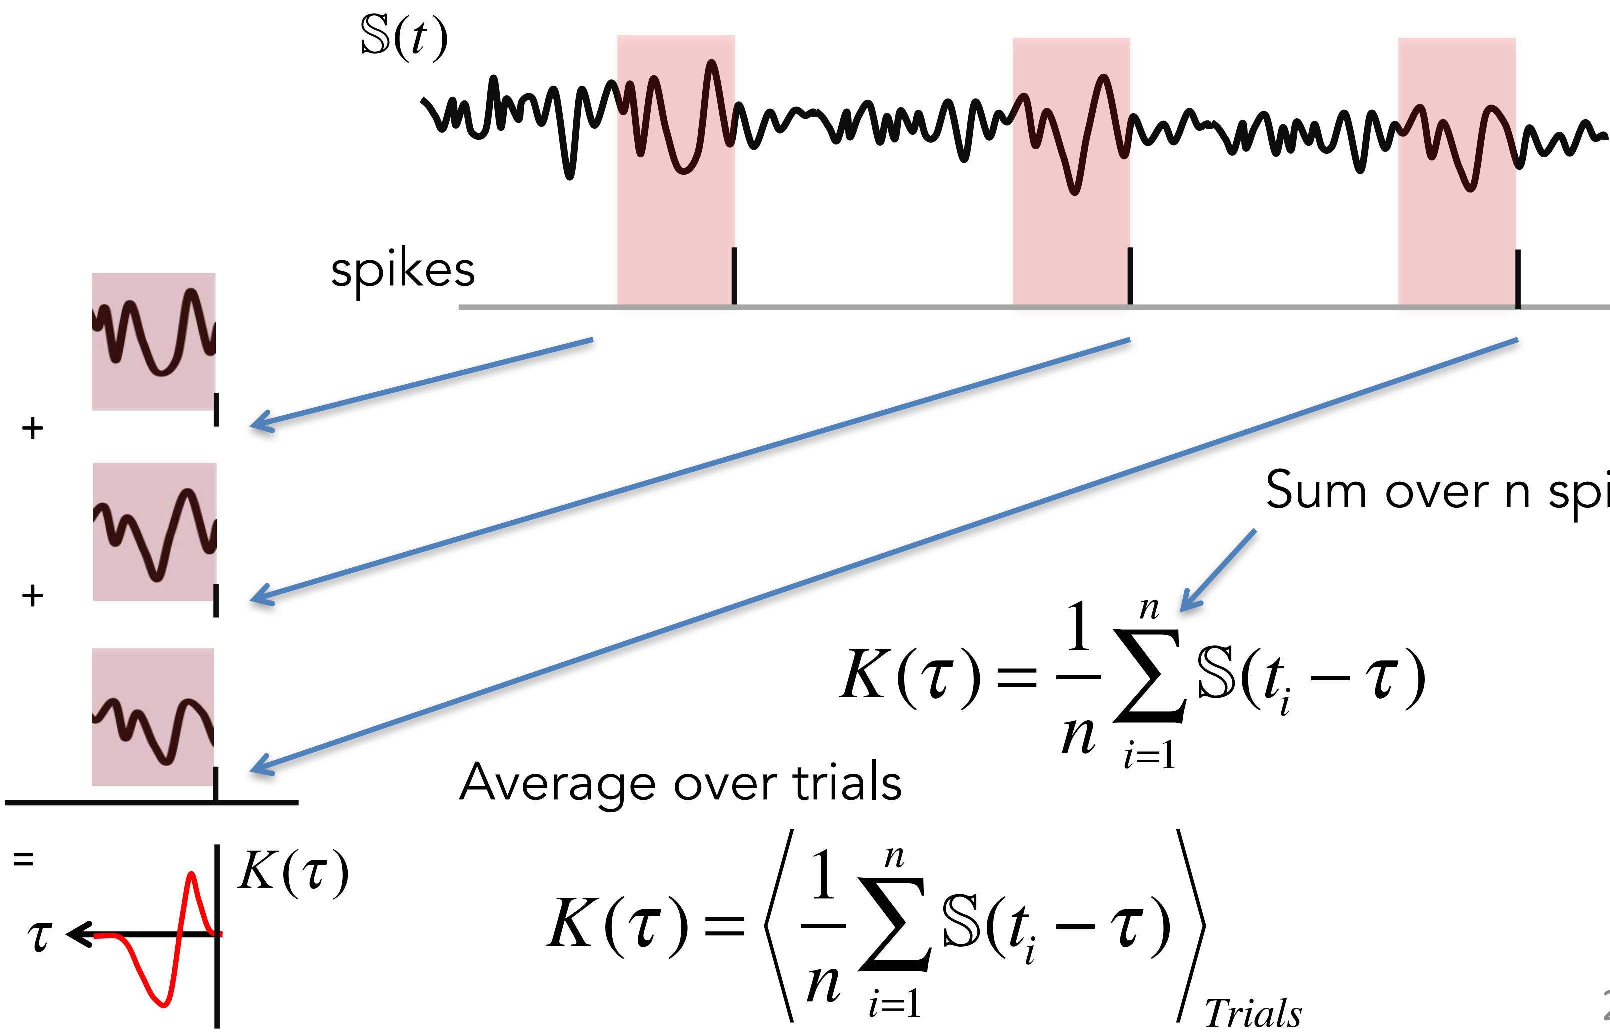

Sum over n spikes

$$K(\tau) = \frac{1}{n}\sum_{i=1}^{n}\mathbb{S}(t_i - \tau)$$

Average over trials

$$K(\tau) = \left\langle \frac{1}{n}\sum_{i=1}^{n}\mathbb{S}(t_i - \tau) \right\rangle_{Trials}$$

Measuring STRFs in the retina

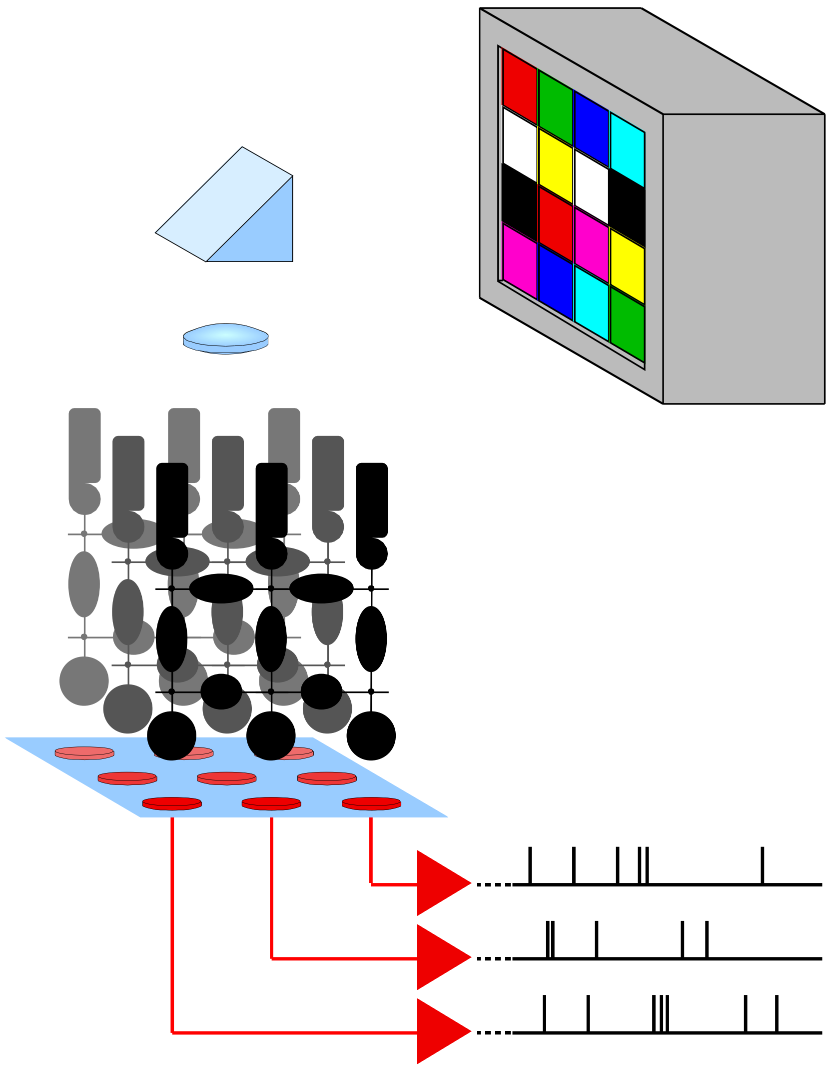

Marcus Meister

Simultaneous Recording from Retinal Ganglion Cells

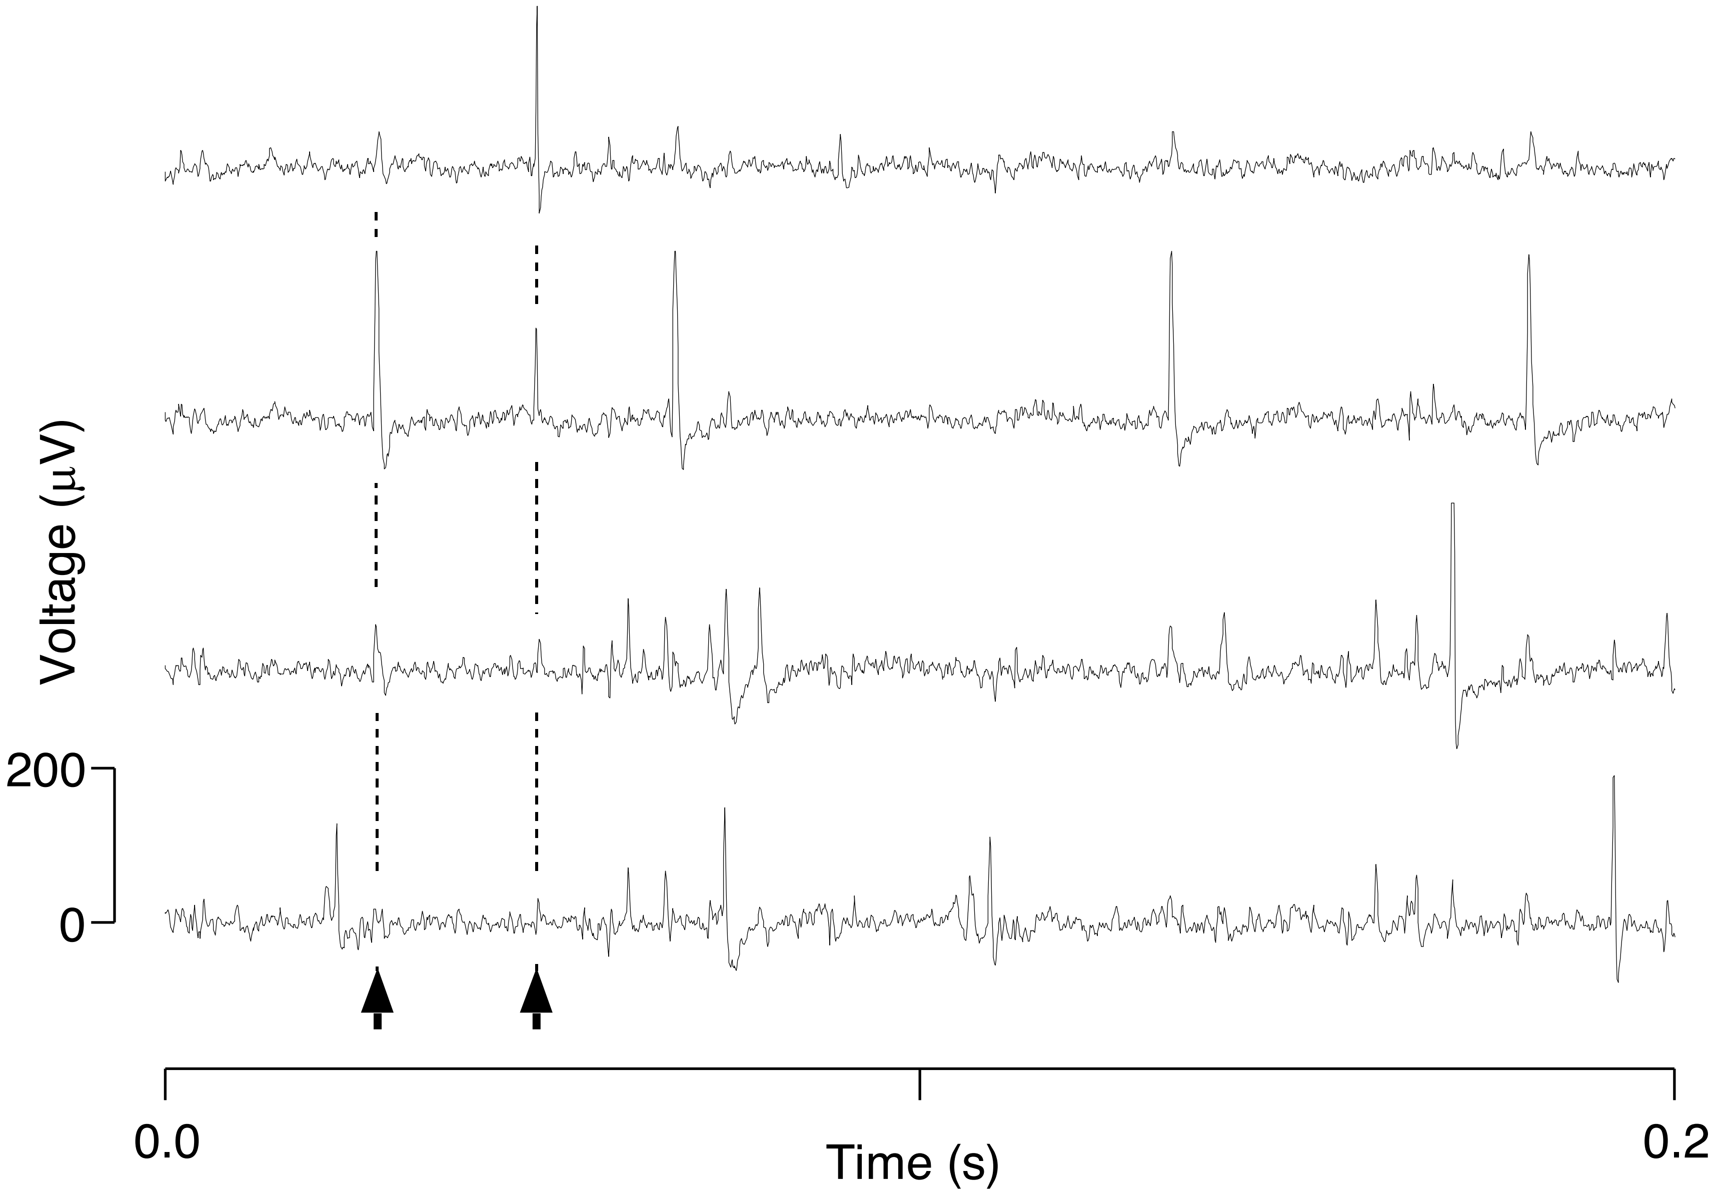

Rabbit ganglion cells responding to a natural movie

"Trees swaying in the breeze"

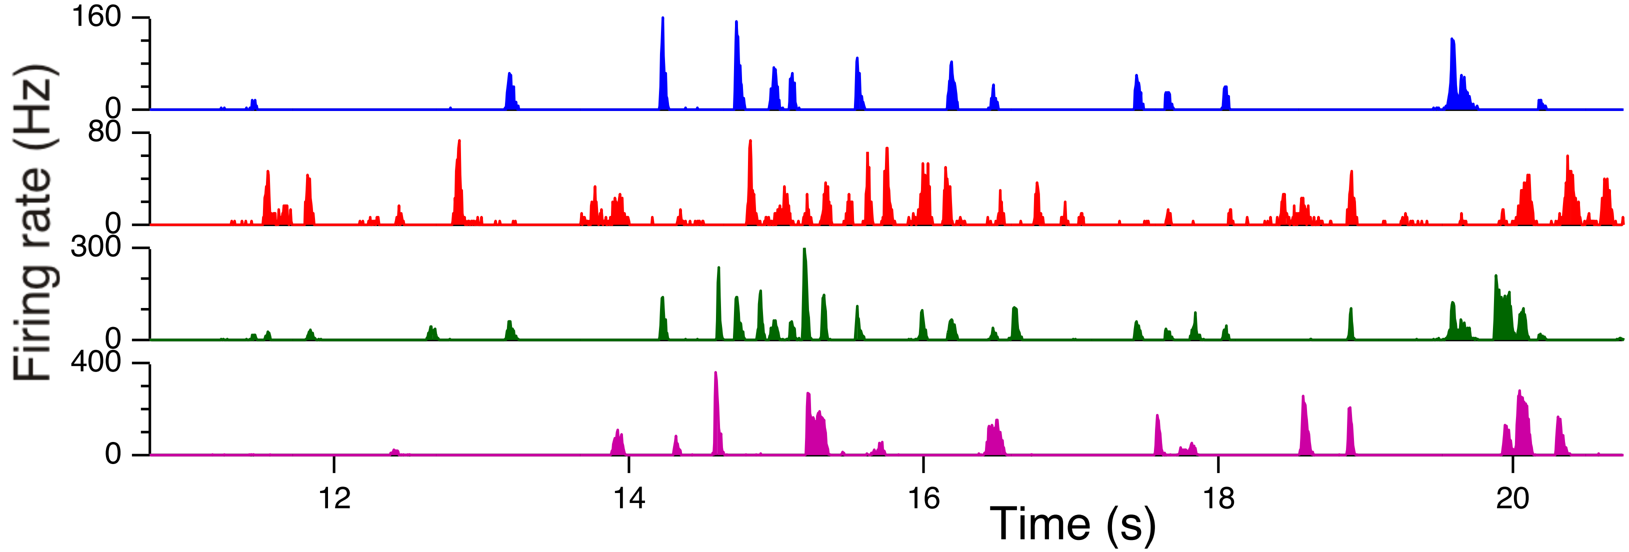

Marcus Meister

Measuring STRFs in the retina

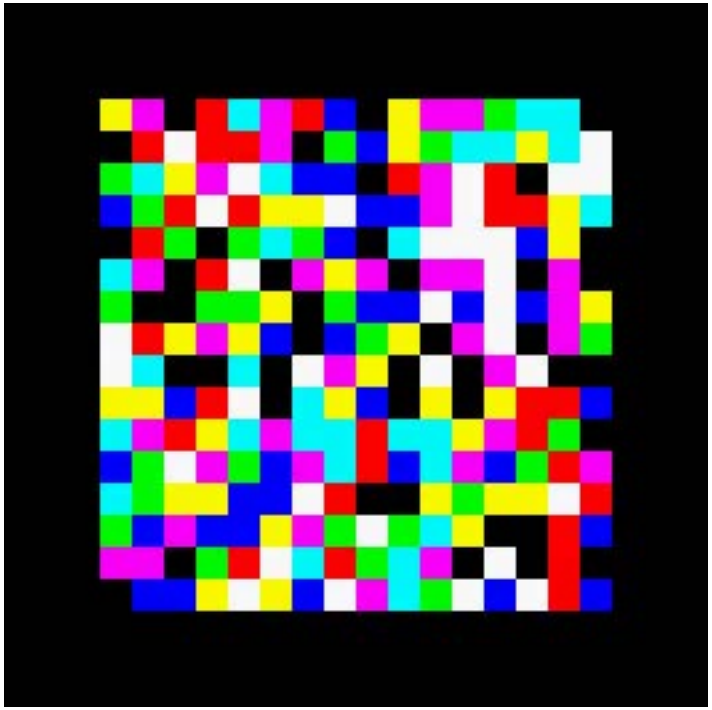

Random flicker stimulus

Reverse-Correlation to a Random Flicker Stimulus

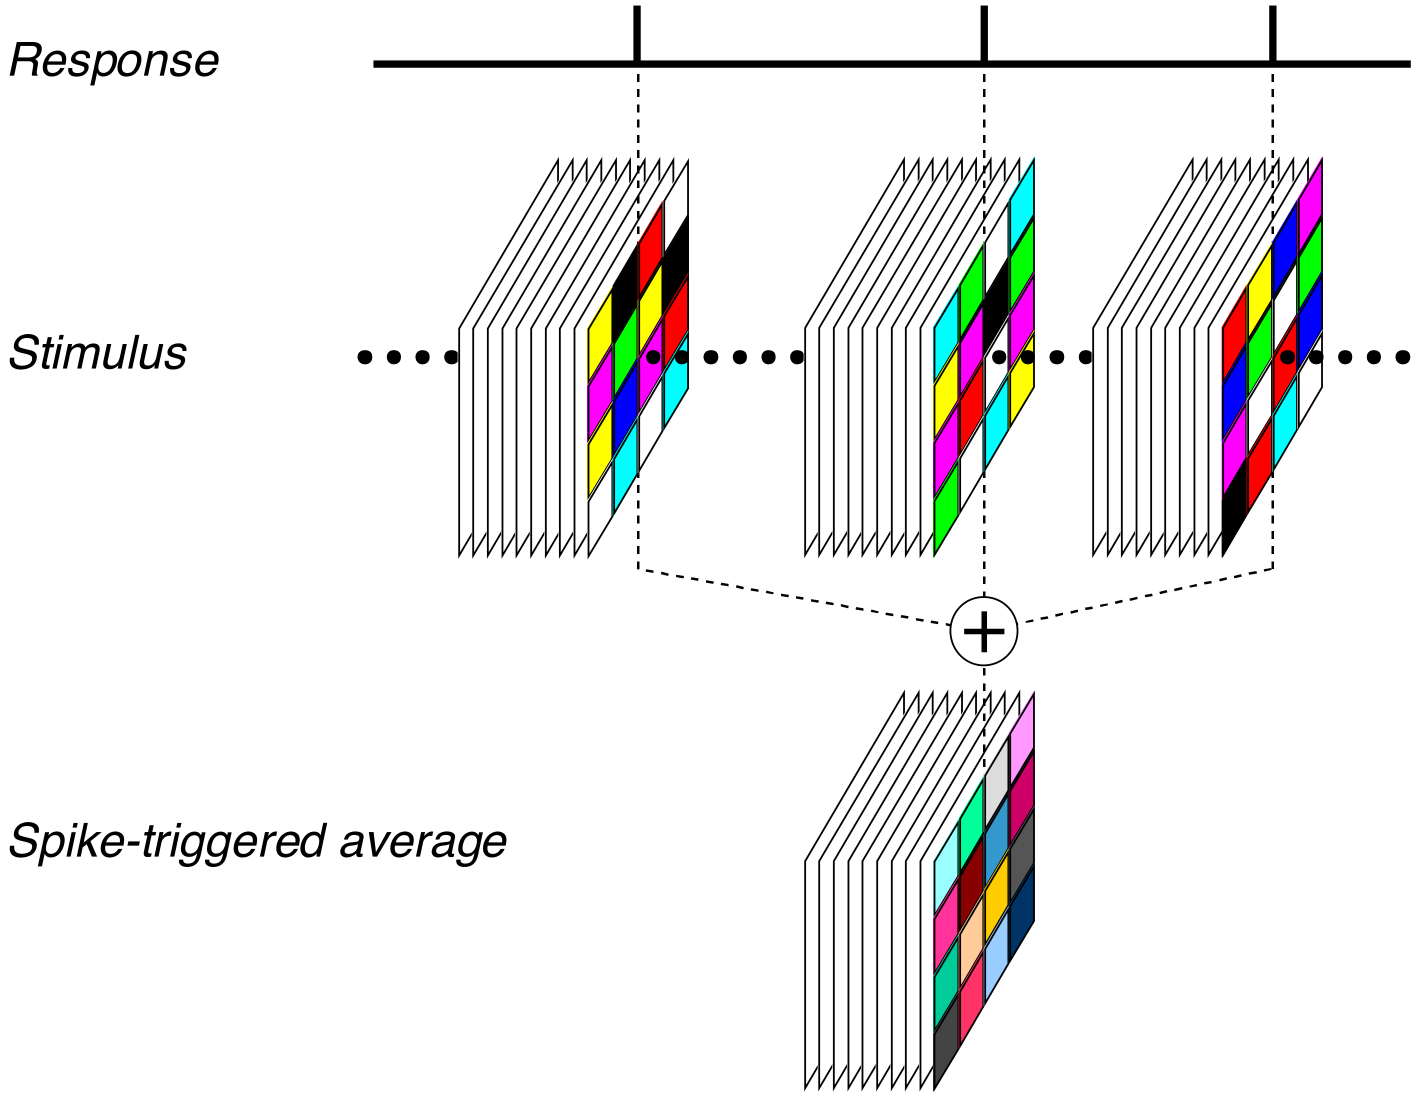

Marcus Meister

Spatio-Temporal Receptive Fields (STRF)

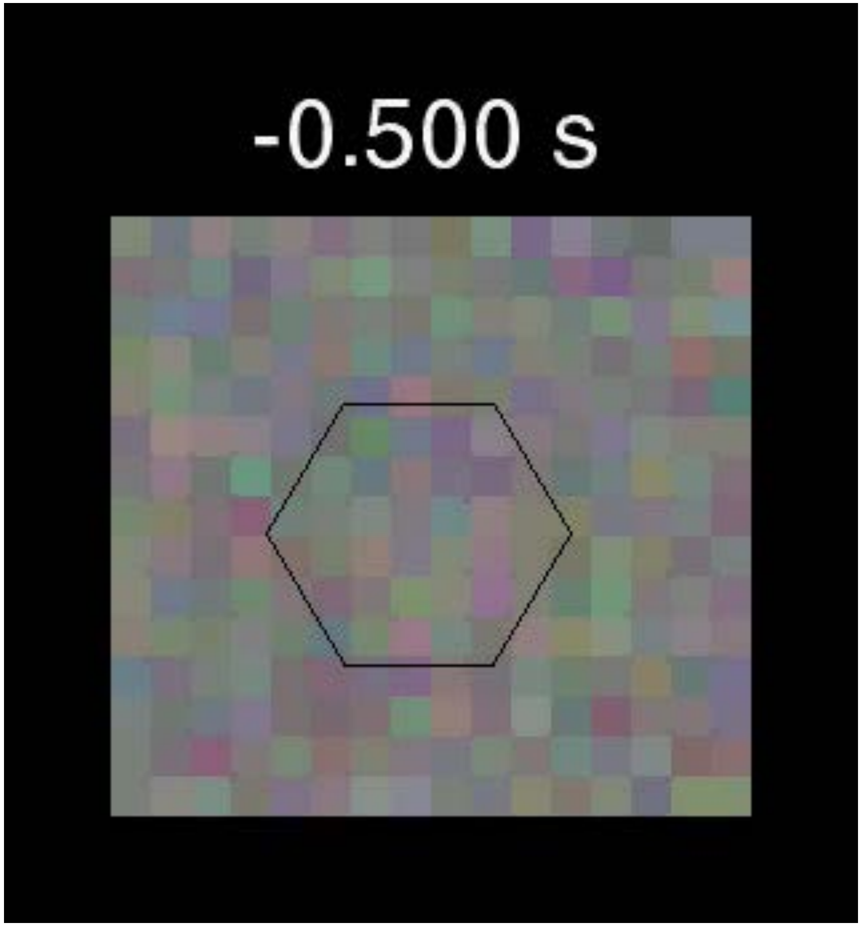

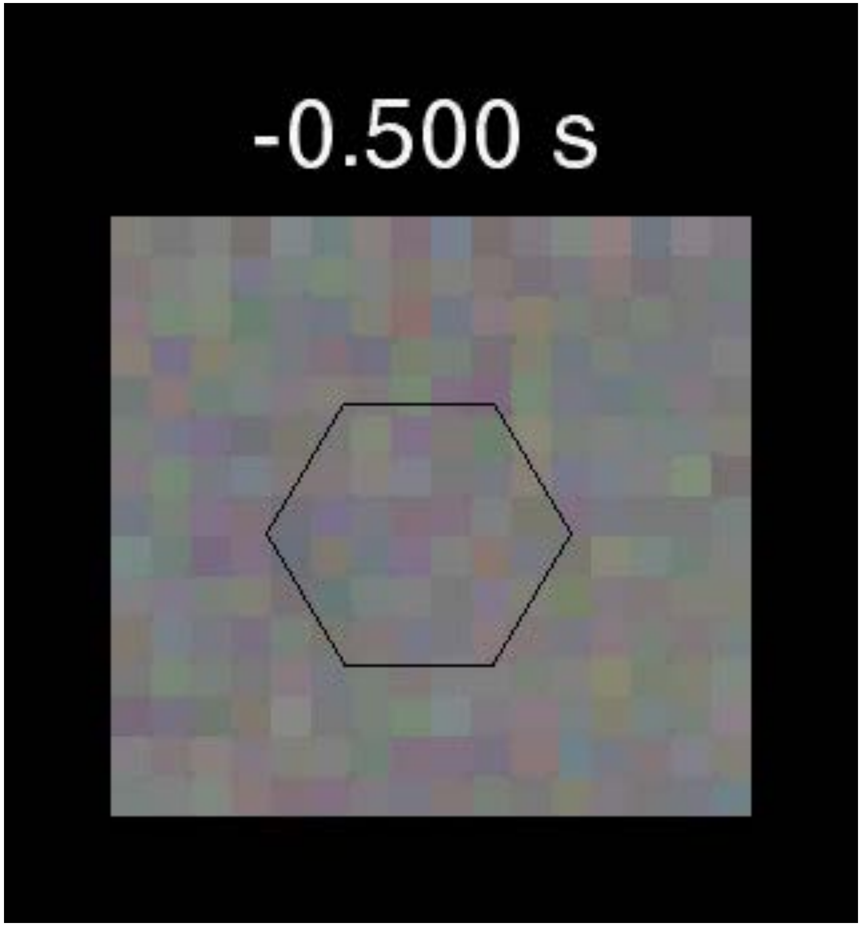

See Lecture 9 video to view the above clips.

Marcus Meister

Mean Effective Stimulus for an ON cell

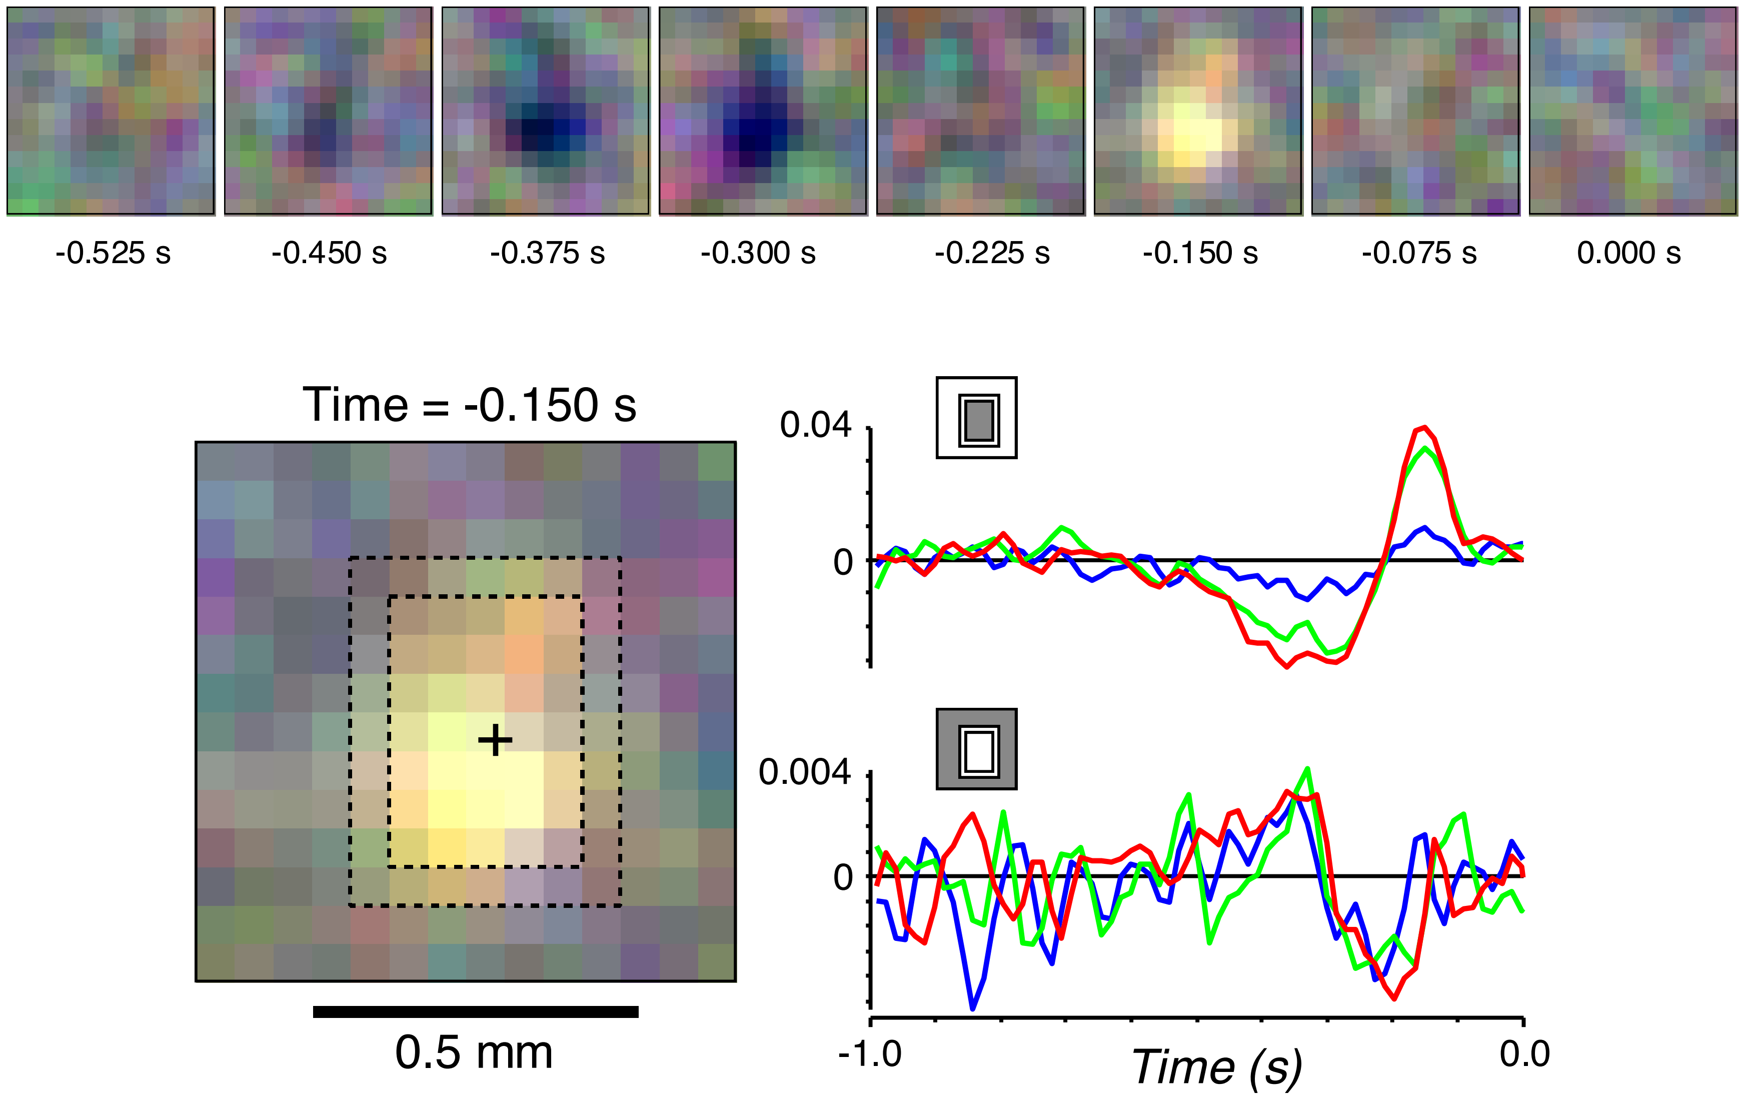

Spectro-temporal receptive fields

We can use this same approach to describe the responses of neurons in the auditory system.

We start by representing sounds in a spectral representation.

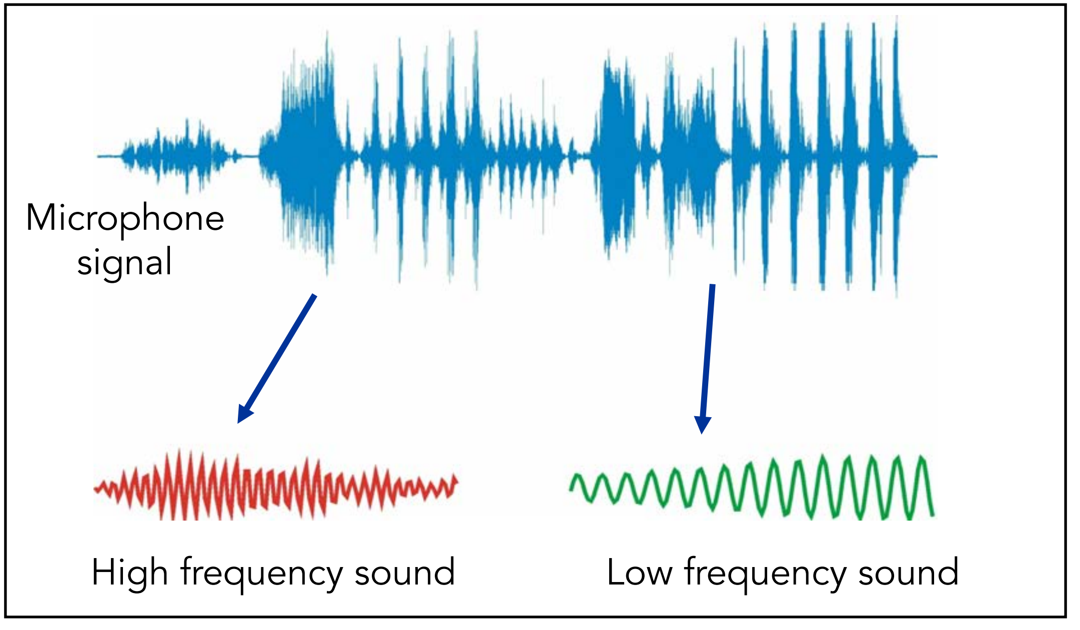

Spectrogram

A spectrogram shows how much power there is in a sound at different frequencies and at different times.

$$S(f,t)$$

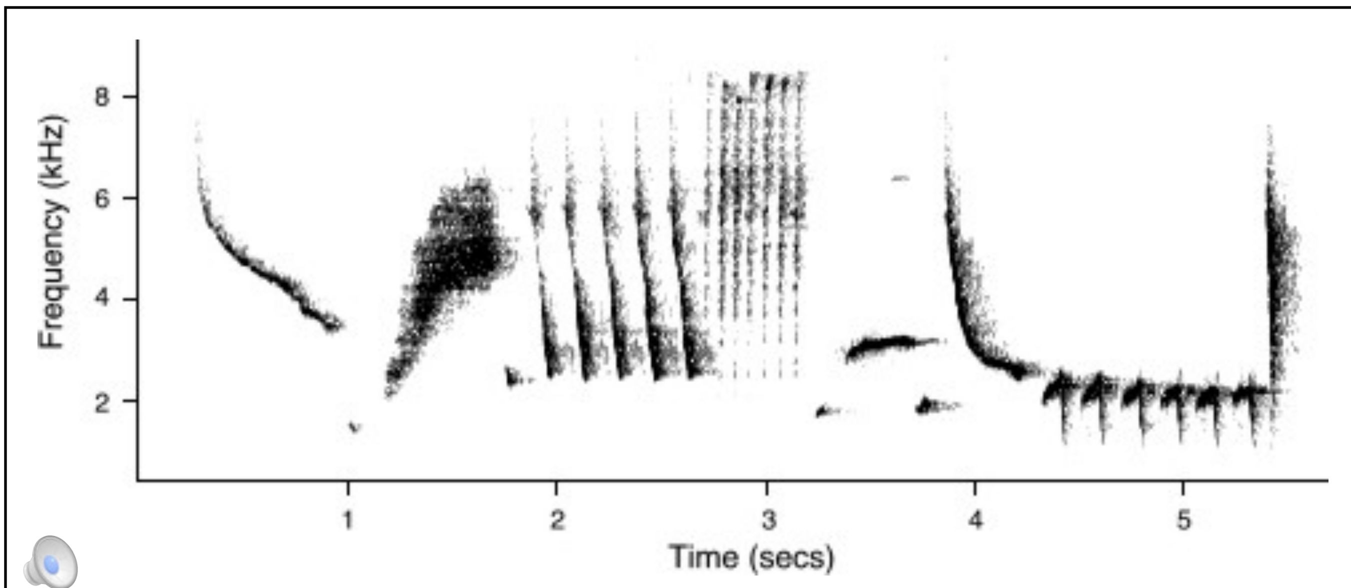

Spectro-temporal receptive fields

Spectro-temporal receptive fields from A1 in monkey.

Spectro-temporal receptive fields

Spectro-temporal receptive fields from A1 in monkey.

deCharms, Blake, Merzenich, Science, 1998

Learning objectives for Lecture 9

- To be able to mathematically describe a neural response as a linear filter followed by a nonlinear function.
  - A correlation of a spatial receptive field with the stimulus
  - A convolution of a temporal receptive field with the stimulus

- To understand the concept of a Spatio-temporal Receptive Field (STRF) and the concept of ‘separability’

- To understand the idea of a Spike Triggered Average and how to use it to compute a Spatio-temporal Receptive Field and a Spectro-temporal Receptive Field (STRF).

Extra slides on nonlinear receptive fields

Non-linearities

Imagine a neuron with these responses to the following stimuli.

Stimulus 1

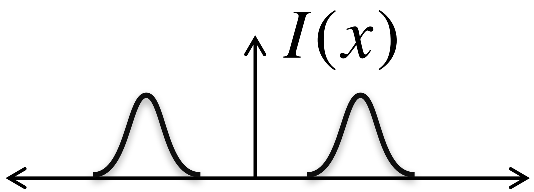

Strong response

Stimulus 2

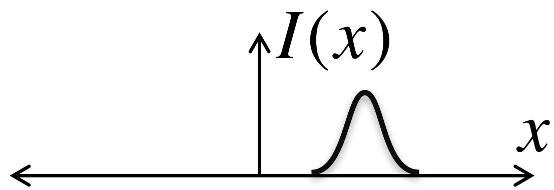

No response

Stimulus 3

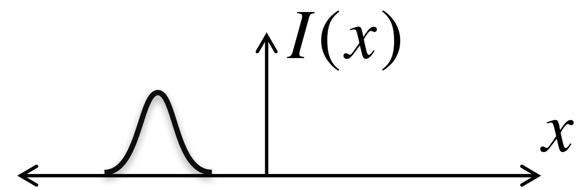

No response

Is this neuron linear?

$$r = r_0 + \int G_1(x) I(x)\, dx$$

If this response was captured by a linear kernel, then, the response to Stimulus 2 and 3 would be half as large as to Stimulus 1. Thus…

$$G_1(x) = 0$$

Non-linearities

$$r = r_0 + \int dx_1\, dx_2\, G_2(x_1, x_2)\, \boxed{I(x_1)\, I(x_2)}$$

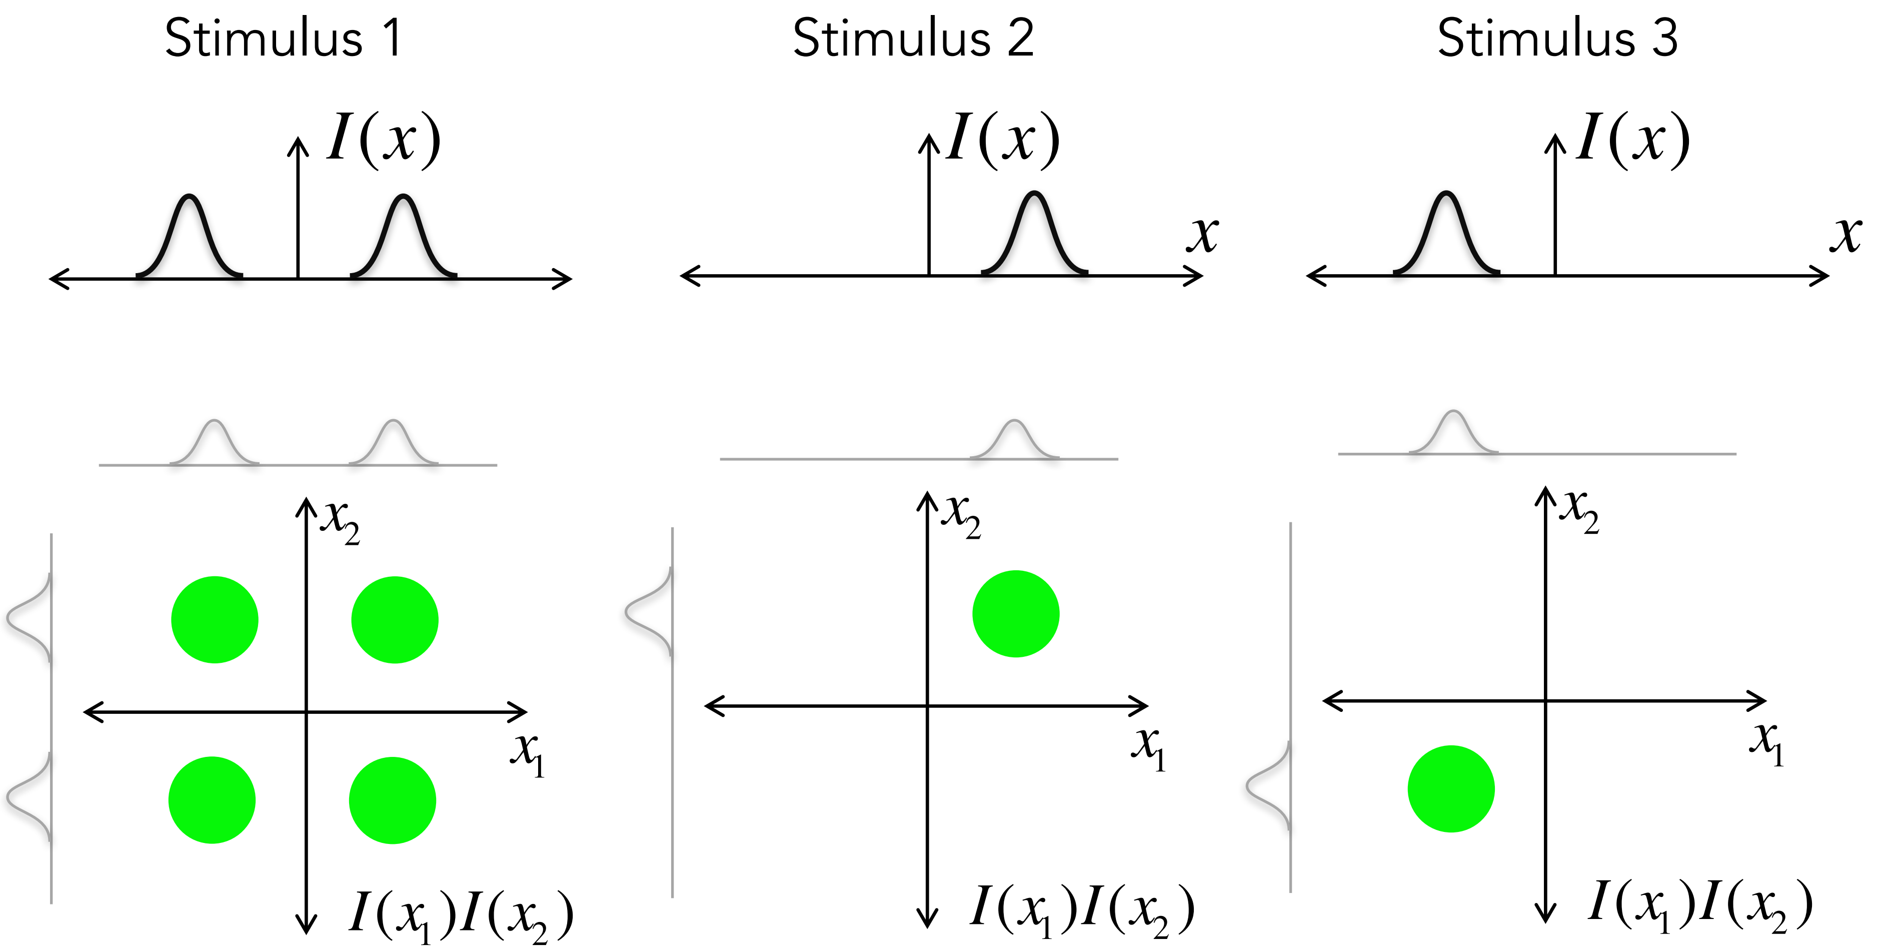

Non-linearities

$$r = r_0 + \int dx_1\, dx_2\, G_2(x_1,x_2)\, I(x_1)\, I(x_2)$$

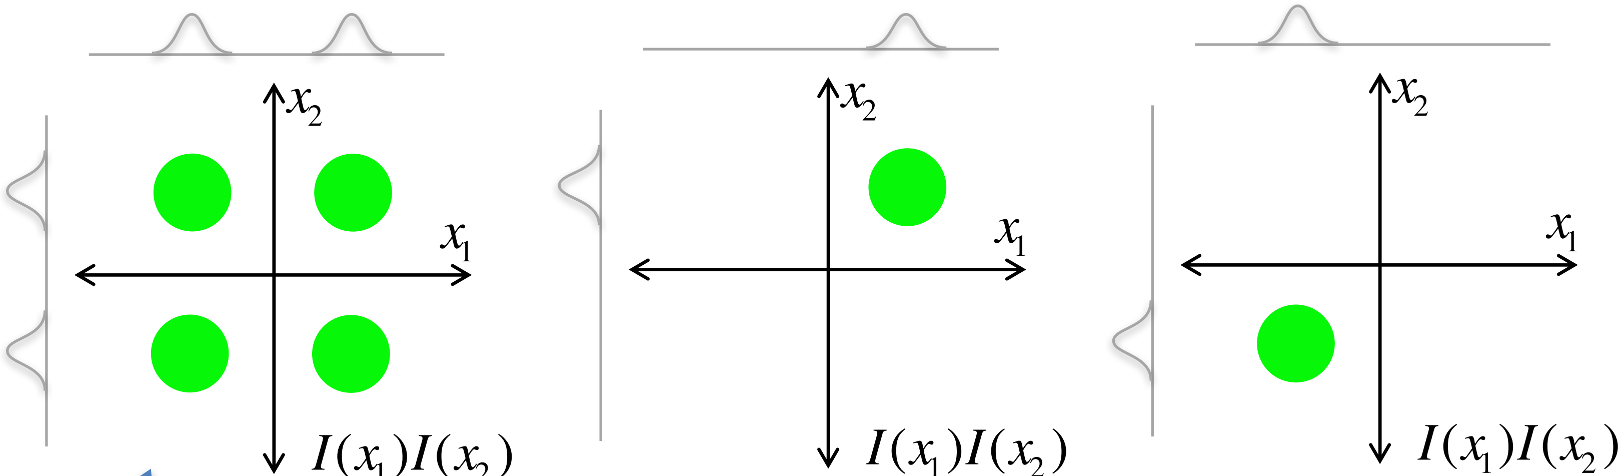

Stimulus 1 is the only one that has overlap with this nonlinear kernel.

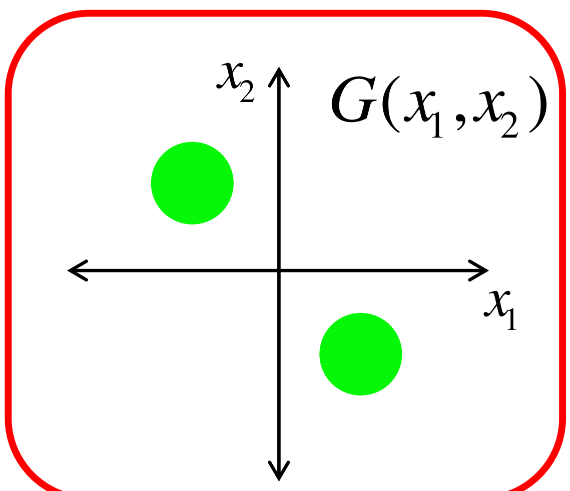

This kernel implements an AND operation!

Non-linearities

The Weiner-Volterra expansion is like a Taylor-series expansion for functions:

$$
\begin{aligned}
r = r_0 &+ \int G_1(x) I(x)\, dx \\
&+ \int dx_1\, dx_2\, G_2(x_1, x_2) I(x_1) I(x_2) \\
&+ \iiint dx_1\, dx_2\, dx_3\, G_3(x_1, x_2, x_3) I(x_1) I(x_2) I(x_3) \\
&+ \quad \ldots
\end{aligned}
$$

Spike-Triggered Average

- One can show that the spike triggered average is just the cross correlation of firing rate and the stimulus

$$K(\tau) = \frac{1}{n}\sum_{i=1}^{n} \mathbb{S}(t_i - \tau)$$

reverse correlation

$$K(\tau) = \int_{-\infty}^{\infty} r(t)\mathbb{S}(t - \tau)dt$$