How do you determine the stimuli encoded by a neuron?

- Get lucky/genious (lec 1 – Hubel + Wiesel, Hollywood)
- Have really simple sensor (lec 1 - wind directon)
- Try everything
  - But … there as many 16x16 black/white images than atoms in the universe
- Try random things, see what 'lights up' the neuron, and generalize!
  … Reverse engineering the brain via spike triggered averages

Multivariate Statistics

- Probability: $P(x)$
  - Say: probability of x
  - Mean: what are the chances of event x happening?
  - Example: when you roll a d6, what is the probability of landing a 5?

$$P(roll = 5) = \frac{1}{6}$$

- Conditional Probability: $P(x|y)$
  - Say: probability of x given y
  - Mean: given the knowledge of y having happened, how probable is x?
  - Example: what is the probability of landing a 5 given the roll was over 3?

$$P(roll = 5|roll > 3) = \frac{1}{3}$$

$$P(roll = 5|sky = blue)$$

$$P(roll = 5|roll > 3, isOdd(roll))$$

- Bayes Inversion
  - Conditional probabilities can be 'inverted':

$$P(x|y) = \frac{P(y|x)P(x)}{P(y)}$$

Intuition

- Consider a simple example, a 'color detector' neuron

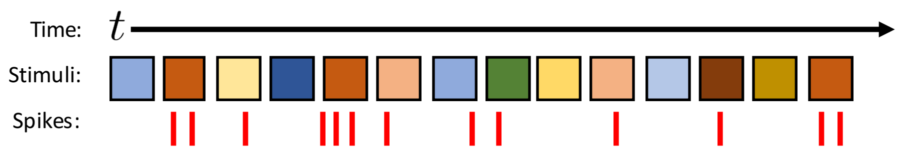

Intuition: Orange detector?

Goal: $P(\text{spike}_t|\text{stim}_t)$

Intuition

- Consider a simple example, a 'color detector' neuron

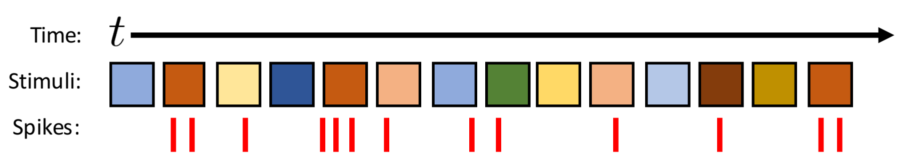

Intuition: Orange detector?

Goal: $P(\text{spike}_t|\text{stim}_t) = \frac{P(\text{stim}_t|\text{spike}_t)P(\text{spike}_t)}{P(\text{stim}_t)}$

$\sim P(\text{stim}_t|\text{spike}_t)$

Intuition

- Consider a simple example, a 'color detector' neuron

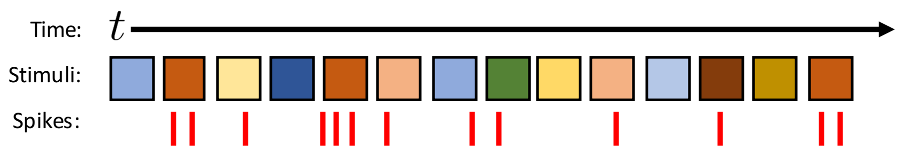

Intuition: Orange detector?

Goal: $P(\text{spike}_t|\text{stim}_t) = \dfrac{P(\text{stim}_t|\text{spike}_t)P(\text{spike}_t)}{P(\text{stim}_t)}$

$\sim P(\text{stim}_t|\text{spike}_t)$

More complicated situation, 2D Grayscale (independent, uniform)

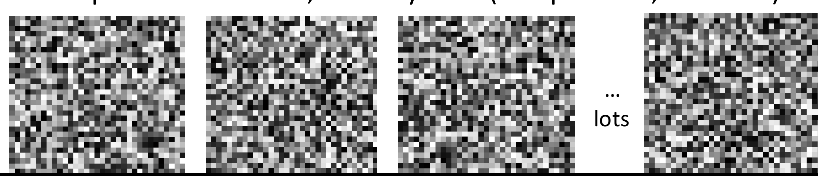

Spike triggered average (STA)

- Assume stimulus is zero-mean and completely random (independent)
  - If a pixel 'drives' a neuron, it will likely be present in stimuli evoking spikes
    This will result in a bias of that pixel in all stimuli that evoked a spike
  - If a pixel is irrelevant to neuron's response, it may/may not be in spiking stimuli
    Since pixel values are independent and zero mean, average value is 0

- Take the expected value of each pixel across spike-triggered ensemble
  - Spike triggered ensemble: the set of all stimuli that evoked a spike

$$\text{STA}(x,y) = \frac{1}{\#\text{ of spikes}} \sum_{t=0}^{T} \text{stim}_t(x,y) \mid \text{spike}_t$$

- Note the conceptual similarity to the probability $P(\text{stim}_t|\text{spike}_t)$
  - The STA then gives us an idea about neural activity $P(\text{spike}_t|\text{stim}_t)$

Spike triggered average (STA)

- Let's do an example for a 1-D, temporal stimulus

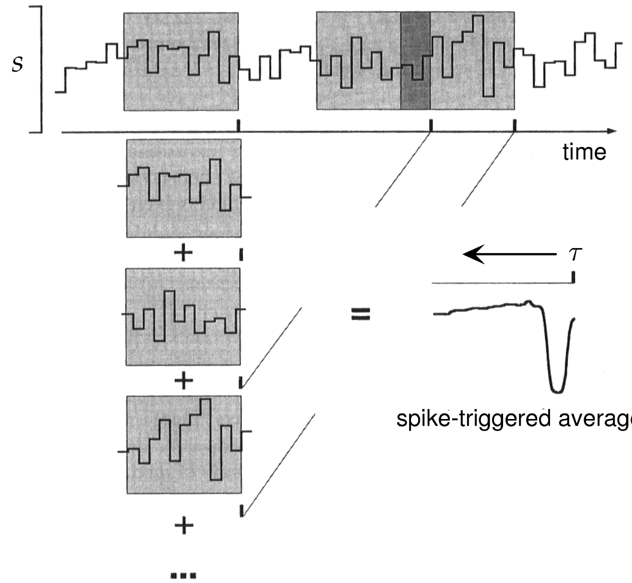

Spike triggered average (STA)

- Let's do an example for a 1-D, temporal stimulus

generate_spiketrain_from_linear_filter.m

```
T=100 * 10^3; %total duration of spike train, in milliseconds
deltat=1;   %in ms

time_list=deltat*(1:length(stim_list));   %list of times

spike_train  %list of 0/1 spike/or not each timestep
stim_list  %list of stimulus values at each timestep
…

figure;
subplot(211)
plot(time_list,stim_list);
title('stimulus','FontSize',18)
subplot(212)
stem(time_list,spike_train,'.')
xlabel('time  (ms)','FontSize',15)
title('spike  raster plot','FontSize',15)
```

Spike triggered average (STA)

- Let's do an example for a 1-D, temporal stimulus

```
Your turn!

First run generate_spiketrain_from_linear_filter.m
to make the vectors spike_train and stim_list

Then write a code that computes STA for this stimulus
and spike train.

Discuss its form, and what it means intuitively for what
stimuli drive the neuron to fire.

Note, you might need to increase T to get an
interpretable result!
```

Predicting responses to new stimuli

Spike triggered average (STA)

- Idea of the optimal filter to predict neural firing:
- Take a (brand new) stimulus stim(x,y)

- Compute "dot product"

$$L = \sum_{x,y} \mathrm{stim}(x,y) \times \mathrm{STA}(x,y)$$

- L can give the best (linear) estimate of p(spike | stim(x,y)) … for this NEW stimulus: i.e., that's the firing rate! (See Ch. 2 for conditions)

Spike triggered average (STA)

- Idea of the optimal filter to predict neural firing:
- Take a (brand new) stimulus stim(x,y)

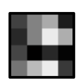

- Compute "dot product"

$$L = \sum_{x,y} \mathrm{stim}(x,y) \times \mathrm{STA}(x,y)$$

- Use L as (linear) estimate of p(spike | stim(x,y)) … for this NEW stimulus: i.e., that's the firing rate!

Literally, as in:
```
p=L*deltat
spike=round(rand + (p-1/2))
```

Spike triggered average (STA)

Extension to temporal stimuli:

$$\mathrm{STA}(x,t,\tau) = \text{average stim preceding spike by } \tau$$

$$L(t) = \sum_{x,y,\tau} \mathrm{stim}(x,y,t-\tau) \times \mathrm{STA}(x,y,\tau)$$

(2.24, Abbott and Dayan)

```
p=L(t)*deltat
spike(t)=round(rand  + (p-1/2))
```

Spike triggered average (STA)

- Idea of the optimal filter to predict neural firing:
- Take a (brand new) stimulus stim(x,y)
- Compute "dot product"

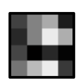

$$L = \sum_{x,y} \mathrm{stim}(x,y) \times \mathrm{STA}(x,y)$$

- Use L as (linear) estimate of p(spike | stim(x,y)) … for this NEW stimulus: i.e., that's the firing rate!

Makes INTUITIVE sense … similarity to the "average" stimulus that created a spike.

When can we show that this makes MATHEMATICAL sense? L can give the best (linear) estimate of p(spike | stim(x,y))

(See Ch. 2 for conditions)

Using linear filter to predict responses

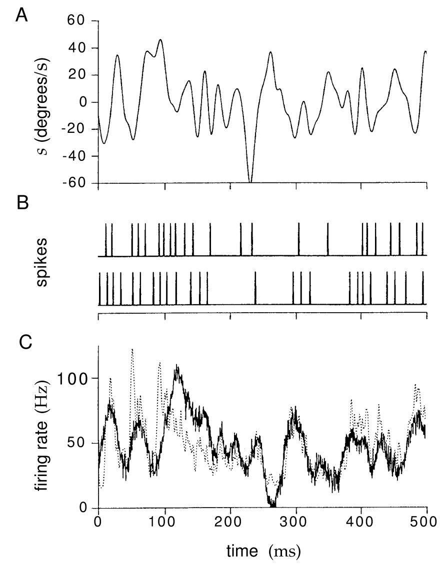

Figure 2.1 Prediction of the firing rate for an H1 neuron responding to a moving visual image. (A) The velocity of the image used to stimulate the neuron. (B) Two of

Linear-Nonlinear-Poisson Model (LNP)

- Incorporating nonlinearities into $P(\text{spike}_t|\text{stim}_t)$
- LN Cascade

Stimulus

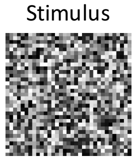

dot (inner) product

STA

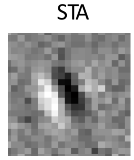

$F\left(L(t)\right) = max(0, L(t))$

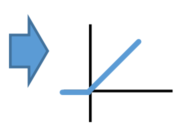

$= r(t)$

Spike rate at t

After determine STA, nonlinearity (F) may be fit based on samples of L and samples of r

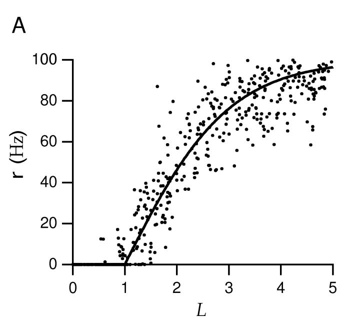

STAs used to Model "Receptive Fields"

- Gaussian
  - Gain alpha
    mean mu
    variance sigma squared

$$\alpha e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

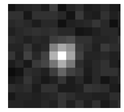

Rod/Ganglion receptive field

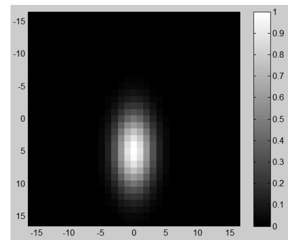

- 2D Gaussian
  Product of Gaussians in each dimension

```
%make 2D arrays of X and Y positions
[DX, DY] = meshgrid(staData.X, staData.Y);

%inline function definition
g = @(D,mu,sigma,alpha) alpha*exp(-(D-mu).^2./(2*sigma^2));

%make 2d RF
rf = g(DX, 0, 2, 1).*g(DY, 5, 4, 1);

%display RF
imagesc(staData.X, staData.Y, rf);
```

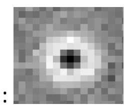

Ganglion/LGN receptive field

Difference of Gaussians:

```
rf = g(DX, 0, 2, 1.5).*g(DY, 0, 2, 1.5) -
     g(DX, 0, 1, 2).*g(DY, 0, 1, 2);
```

- Gabor
  Product of sinusoid and a Gaussian

$$\alpha \sin(2\pi\omega + \theta)\, e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

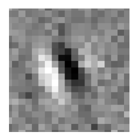

V1 simple cell receptive field

STA's used to recover V1 RFs

- Abhishek De's V1 Cells
  - Courtesy Horwitz Lab
  - Brains are noisy

Cell 1

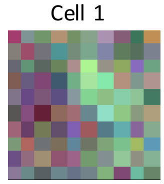

Cell 2

Cell 3

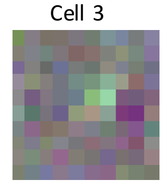

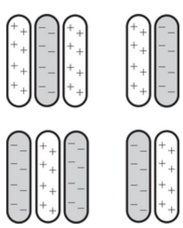

FIGURE 21.1 Typical simple cell receptive fields described by Hubel and Wiesel with even symmetric and odd symmetric spatial profile. Excitatory regions are marked by pluses and inhibitory regions by minuses.

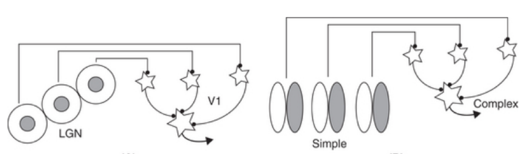

FIGURE 21.13 **A.** Hubel and Wiesel model describing

- The 'complexity' of computations between a stimulus and the neuron's spike rate effect the ability of STA to estimate the RF

Linear-Nonlinear-Poisson Model (LNP)

- We have a RF, how do you model spikes? $P(\text{spike}_t | \text{stim}_t)$
- LN Cascade

Stimulus

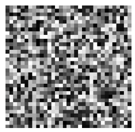

dot (inner) product

RF

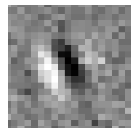

$f(r_t) = \max(0, r_t)$

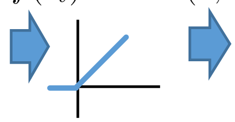

$\lambda_t = f(r_t)$

Lambda of t:
likelihood of spiking at t

$\lambda \frac{T}{\Delta t}$

Limits of STA

- Spike-triggered averages seem like magic
  Why haven't we solved the brain and vision?
- Lets look at some data recorded from V4

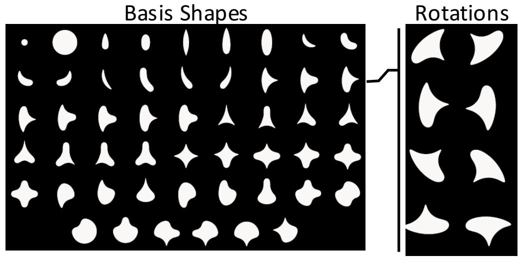

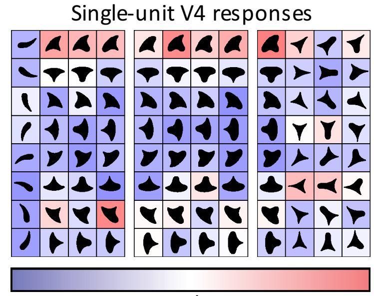

- Would these stimuli drive the cell?

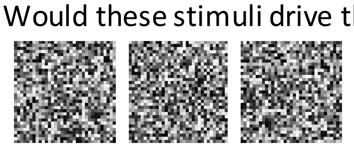

- How long would it take before you randomly sampled a shape?
- STA only guaranteed to work in the limit of infinite stimuli (not practical for experimentation)

Decoding neurons probabilistically

- STA requires uncorrelated stimuli
  - Good for Retina, LGN, V1
- "GLM" and point process methods provide important allied approaches
  - [Gerstner, Paninski et al, Book "Neuronal Dynamics"]
- Deeper regions of ventral cortex respond to complex structure and form
- Maximally Informative Dimensions
  - Analyzing Neural Responses to Natural Signals: Maximally Informative Dimensions. Tatyana Sharpee, Nicole C. Rust, and William Bialek, Neural Computation 2004 16:2, 223-250
  - Given a model of stimuli to spike output, maximize the difference between:

$$P(\text{stim}_t|\text{spike}_t) \quad P(\text{stim}_t|\text{nospike}_t)$$

  - Agnostic to stimuli and computation, hard to fit

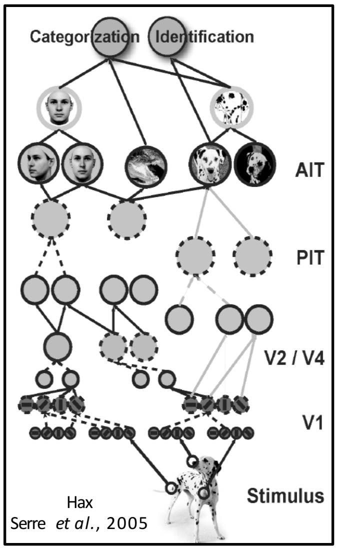

Hax
Serre *et al.*, 2005