

Correlated activity in populations

Lecture 2

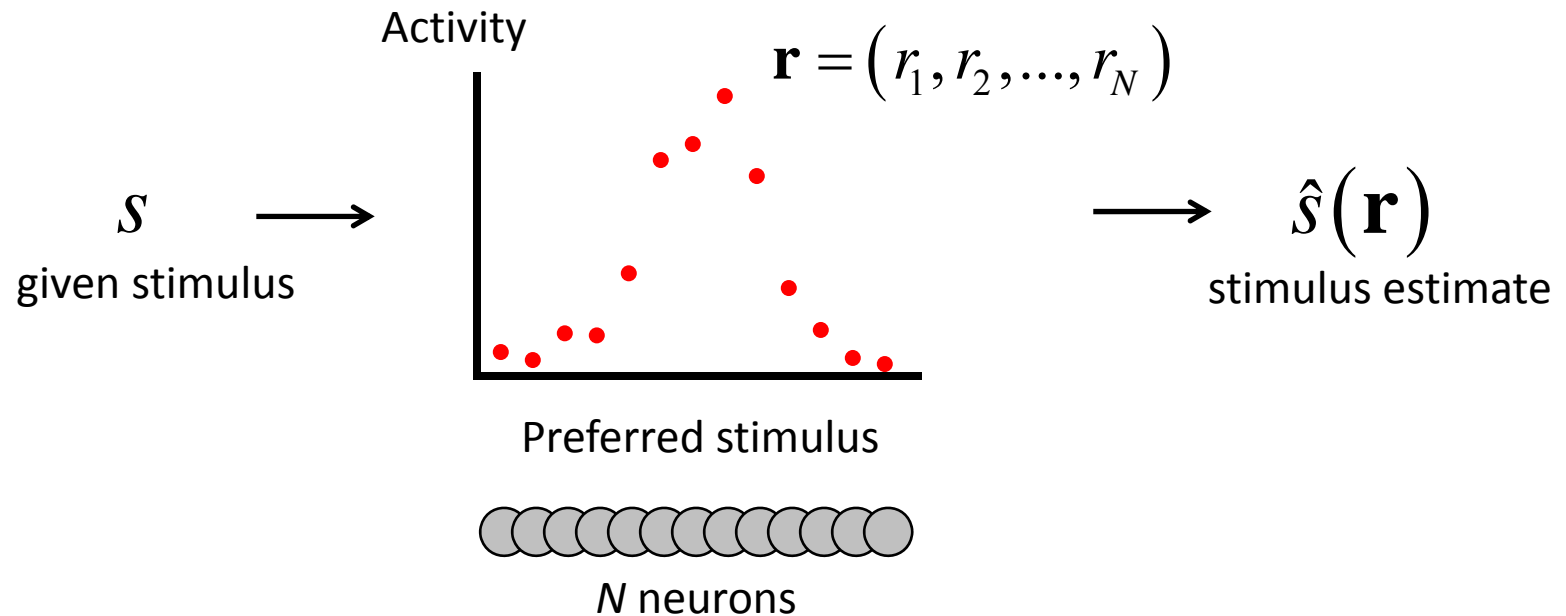
Last lecture

1. Population response (noise) distributions
2. Population decoders (winner-take-all, center-of-mass, template-matching, maximum-likelihood, Bayesian)
3. Goodness of decoders \rightarrow Fisher information
4. Decoding uncertainty (probability distributions over the stimulus)

Today

1. Quantifying the effect of correlations on information
2. Modeling correlated population activity
 - Averbeck, Latham, Pouget (2006), *Neural correlations, population coding, and computation*. Nat Rev Neurosci 7(5): 358-66.
 - Pillow et al. (2008), *Spatio-temporal correlations and visual signalling in a complete neural population*. Nature 454 (7207): 995-9.

Population codes



Quality of a population code

- How much information about a stimulus s does a population \mathbf{r} contain? \rightarrow
- How well does the best possible decoder do in decoding s from \mathbf{r} ? \rightarrow
- Fisher information

$$I(s) = - \left\langle \frac{\partial^2}{\partial s^2} \log p(\mathbf{r} | s) \right\rangle$$

response distribution of the population

Independent Poisson variability \rightarrow

$$I(s) = \sum_{i=1}^N \frac{f_i'(s)^2}{f_i(s)}$$

slopes of tuning curves

Population response distribution

- Conditional independence

$$p(\mathbf{r} | s) = \prod_{i=1}^N p(r_i | s)$$

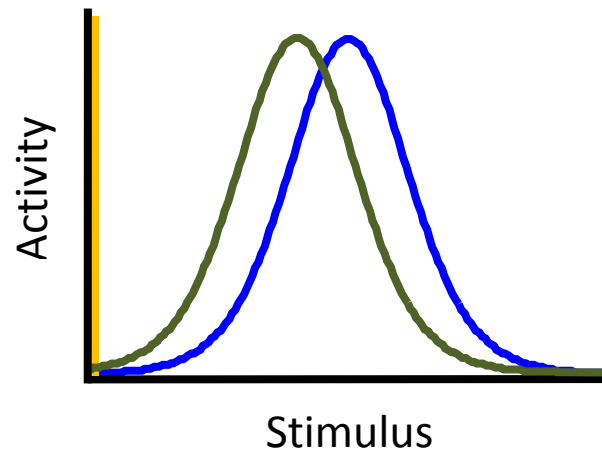
- Noise correlations

$$p(\mathbf{r} | s) \neq \prod_{i=1}^N p(r_i | s)$$

- Different from signal correlations

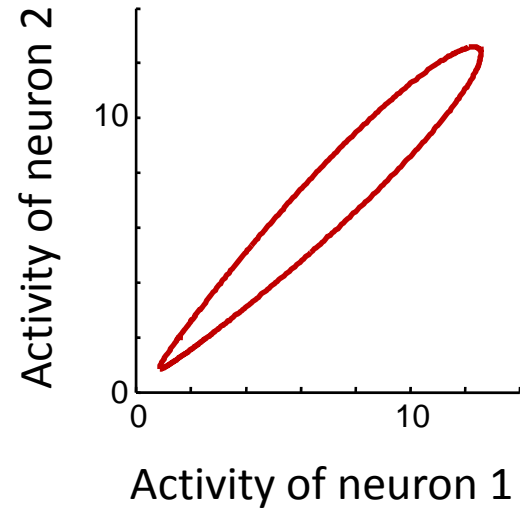
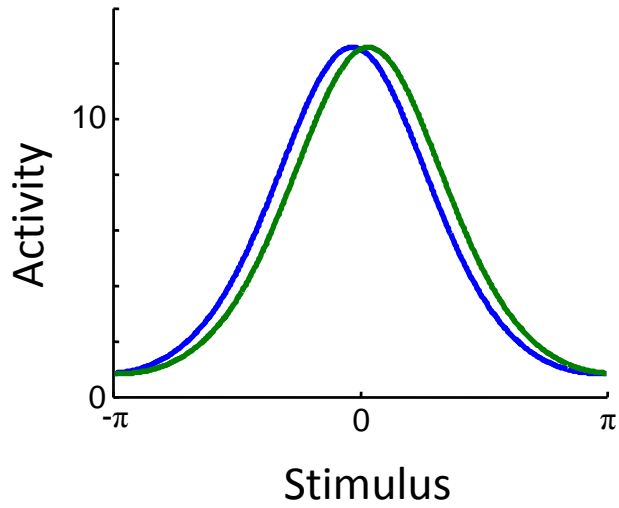
Signal correlations

Tuning curves of two neurons

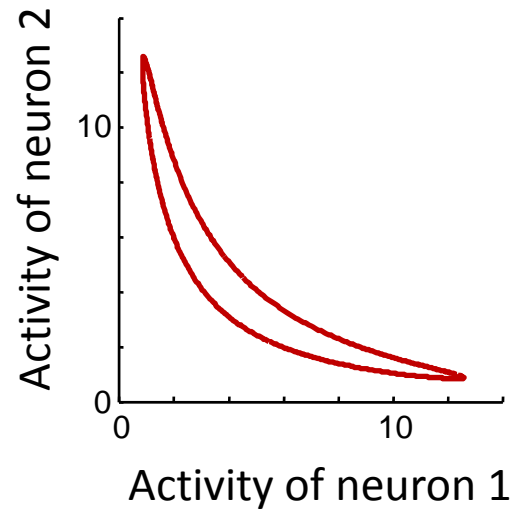



Change stimulus, ignore variability

Neurons with similar tuning




Neurons with dissimilar tuning



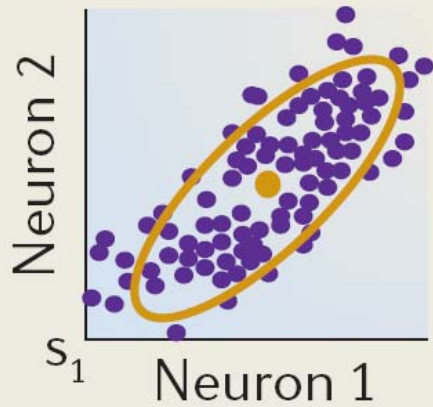
Noise correlations

Fix stimulus, examine variability across trials

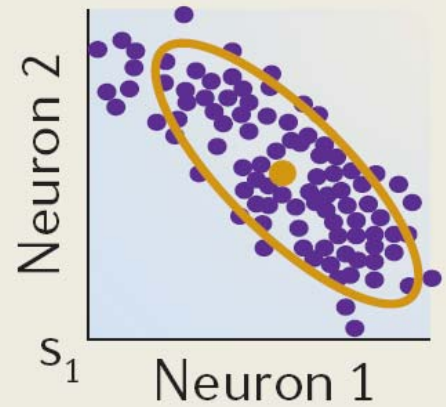


Noise correlations

+ Noise correlation



- Noise correlation




Noise correlations

c Uncorrelated



Preferred
stimulus

d Correlated



Preferred
stimulus

Noise versus signal correlations

- Noise correlations

$$p(\mathbf{r} | s) \neq \prod_{i=1}^N p(r_i | s)$$

- Signal correlations

$$p(\mathbf{r}) \neq \prod_{i=1}^N p(r_i)$$

How do noise correlations affect information?


- Can go either way!
- General conclusions about redundancy or synergy not justified
- How to quantify the impact of correlations on information? What is your “control”?

Shuffling responses

Spike counts (in response to a fixed stimulus)

Trial #	Neuron 1	Neuron 2
1	8	12
2	7	13
3	4	6
4	5	6
5	1	0
6	6	8


Positively correlated



Spike counts with trials shuffled (separately)

Neuron 1	Neuron 2
6	0
1	6
8	13
4	6
7	8
5	12

No correlation



$$\Delta I_{\text{shuffled}}$$


- Shuffling preserves variability of each neuron individually, $p(r_i|s)$
- Destroys correlations
- Measure of information in correlations:

$$\Delta I_{\text{shuffled}} = I - I_{\text{shuffled}}$$


- Information relates to discriminability

Information (I) in unshuffled responses

a $\Delta I_{\text{shuffled}} < 0$



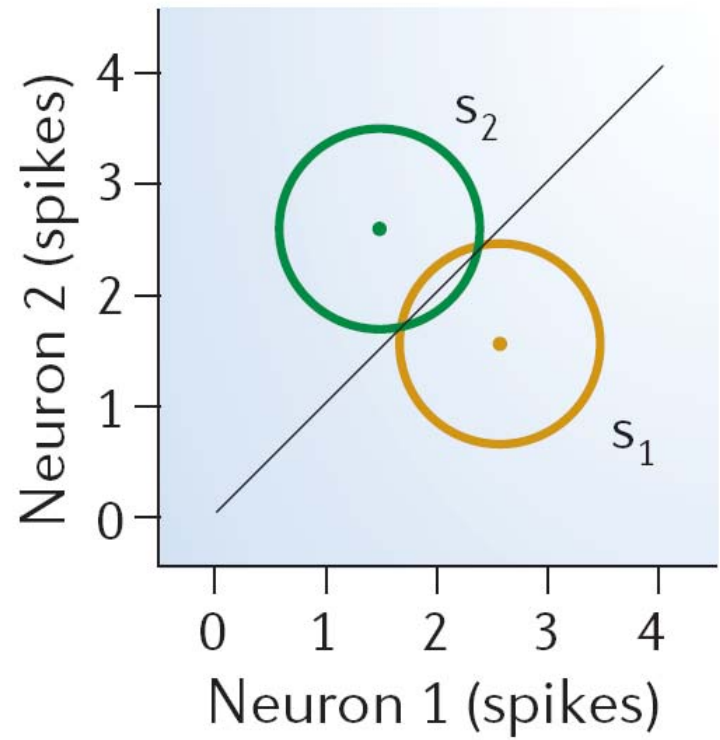

Information (I_{shuffled}) in shuffled responses



Information (I) in unshuffled responses

Information (I_{shuffled}) in shuffled responses



b $\Delta I_{\text{shuffled}} > 0$



Information (I) in unshuffled responses

Information (I_{shuffled}) in shuffled responses

c $\Delta I_{\text{shuffled}} = 0$



Interaction between signal and noise correlations

- If signal correlations are positive, positive noise correlations *decrease* information.
- If signal correlations are negative, positive noise correlations *increase* information.
- Correlations can depend on the stimulus.

In cortex

- $\Delta I_{\text{shuffled}} < 10\%$ (pairs of neurons)
 - Rat barrel cortex
 - Macaque V1, prefrontal, somatosensory cortex
- But: small effects of pair correlations can have large effects in populations

Encoding versus decoding perspective

- So far: how do correlations affect the total amount of information in a population?
(Encoding perspective: redundancy / synergy)
- Decoding: given a correlated population, how much worse would you do when ignoring the correlations?

$$\Delta I_{\text{diag}}$$



- Train a decoder on the shuffled (uncorrelated data)
- Apply the same decoder to the true, correlated data \rightarrow extract information I_{diag}

$$\Delta I_{\text{diag}} = I - I_{\text{diag}}$$

Estimate w_{diag} on shuffled responses

Apply to unshuffled responses (measures I_{diag})

a $\Delta I_{\text{diag}} = 0$

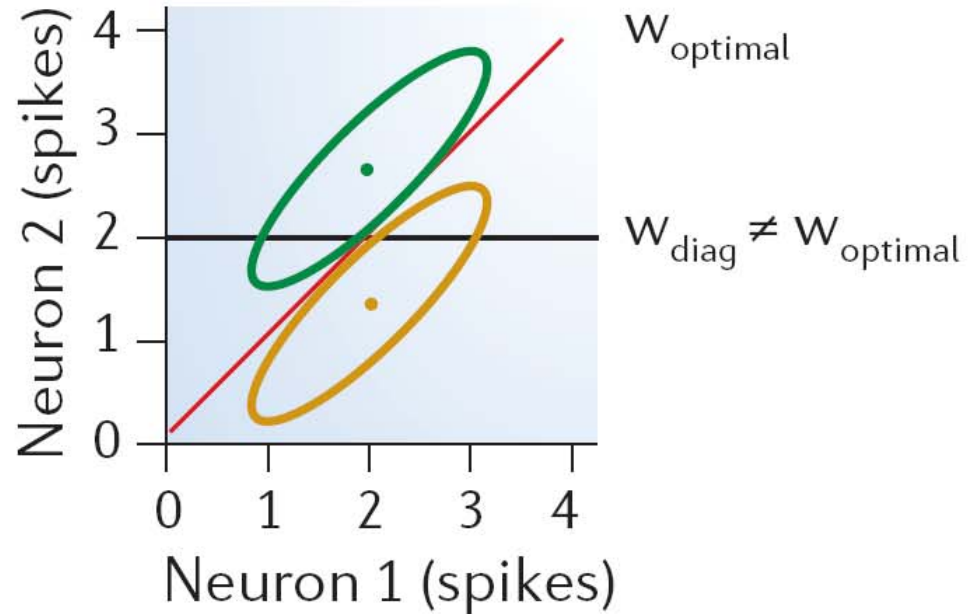
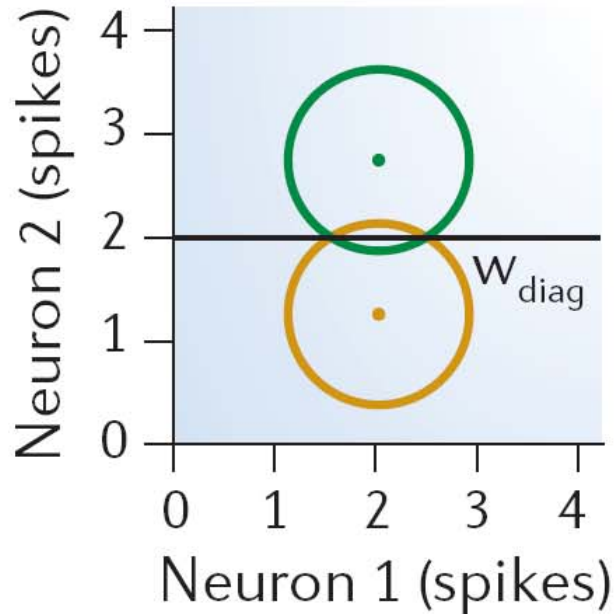


$w_{\text{diag}} = w_{\text{optimal}}$

Estimate w_{diag} on shuffled responses

Apply to unshuffled responses (measures I_{diag})

b $\Delta I_{\text{diag}} > 0$



Ignoring correlations

- I_{diag} cannot be greater than I (unlike I_{shuffled})
- Estimating correlations is data-intensive
- Trade-off between decoding performance and data needed to measure correlations
- $\Delta I_{\text{diag}} \approx 10\%$ in experiments (pairs of neurons):
 - Mouse retina
 - Rat barrel cortex
 - Macaque SMA, V1, and motor cortex

Modeling correlated populations


- Complete populations are different from pairs of neurons.
- So far, no model-based characterization of correlations

Pillow et al. (2008)

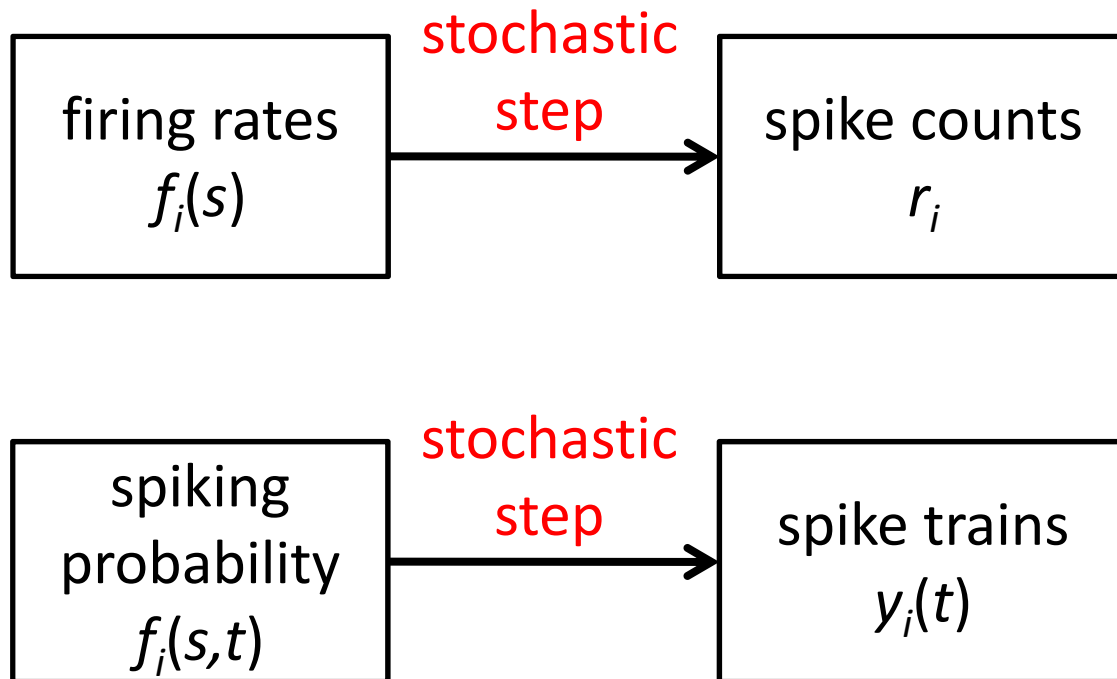
- Complete population
- Encoding model for spike *times* → allows examining temporal correlations
- Parameters can be fit to physiological data
- Stimuli: binary white noise
 - Not a single number, but a time series for every pixel i : \mathbf{x}_i

Recordings

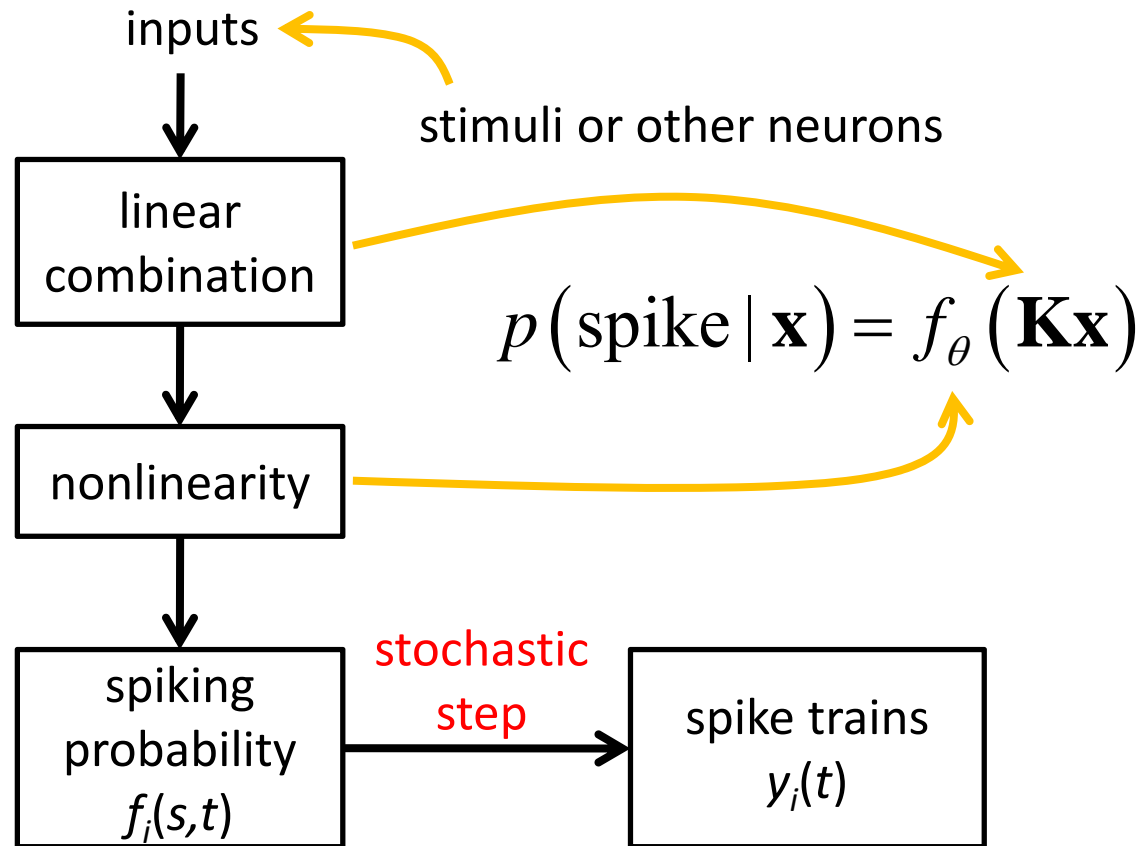
- Neural data: 27 retinal ganglion cells in vitro
 - ON and OFF cells
 - Nearly complete mosaics




Poisson neurons



LNP-neurons



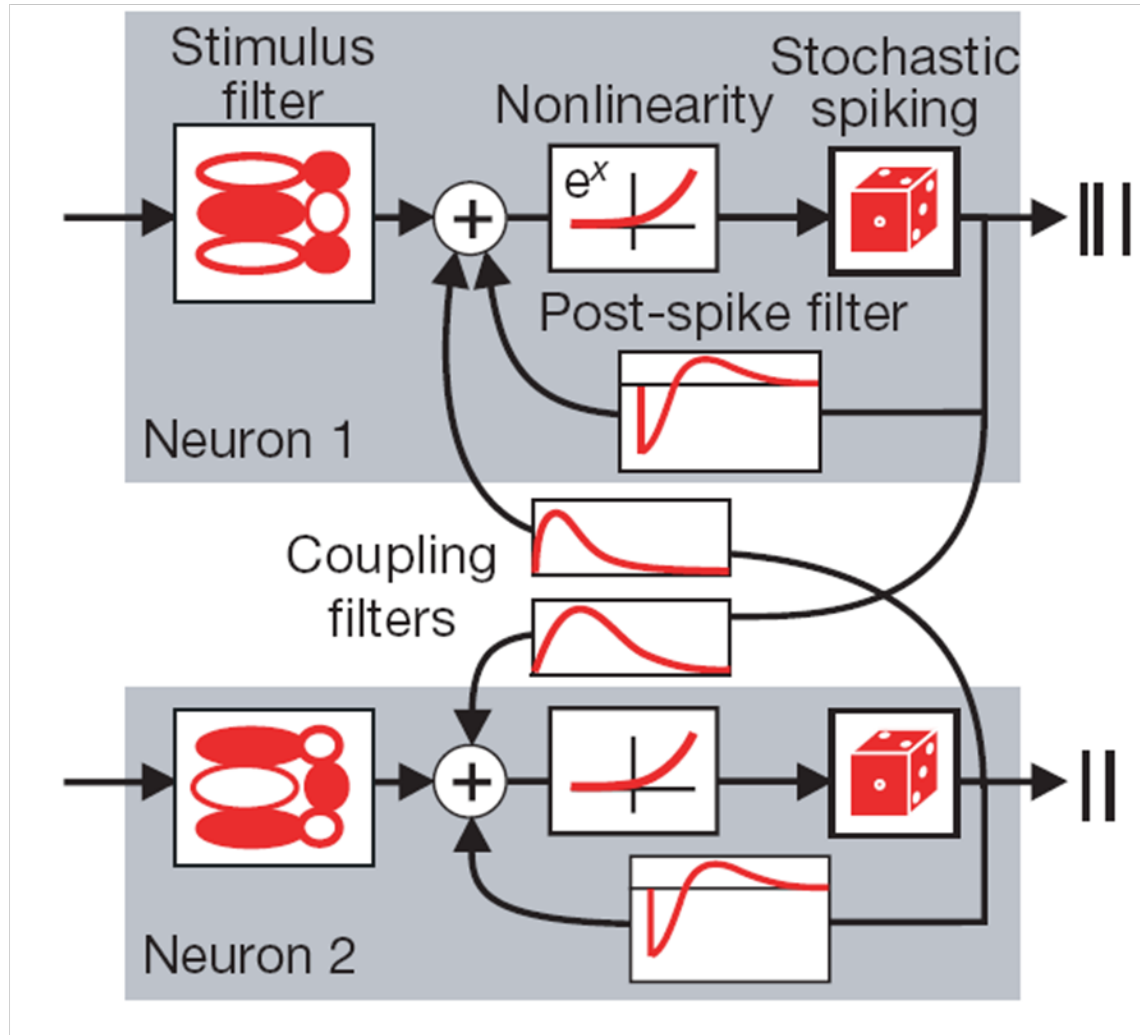
Fitting the parameters of the model to the data



Likelihood:
$$p(\text{data} \mid \text{model}) = \prod_t p(y_t \mid \text{model parameters})$$


Log likelihood:
$$\begin{aligned} \log p(\text{data} \mid \text{model}) &= \sum_t \log p(y_t \mid \text{model parameters}) \\ &= \sum_{\text{spiking } t} \log p(y_t = 1 \mid \mathbf{K}, \theta) + \sum_{\text{non-spiking } t} \log p(y_t = 0 \mid \mathbf{K}, \theta) \\ &= \sum_{\text{spiking } t} \log(f_\theta(\mathbf{K}\mathbf{x}(t)) \delta t) + \sum_{\text{non-spiking } t} \log(1 - f_\theta(\mathbf{K}\mathbf{x}(t)) \delta t) \end{aligned}$$

Coupled spiking model



“uncoupled model”:
ignoring correlations

Cross-correlation function



$$C(\tau) = \frac{\langle y_1(t) y_2(t + \tau) \rangle - \langle y_1(t) \rangle \langle y_2(t) \rangle}{\langle y_2(t) \rangle dt}$$

Other tests

- Triplet correlations
- Peri-stimulus time histograms (average single-cell responses to new stimuli)
- Predicting a single cell's spike train from the stimulus and the activity of the rest of the population
- So far: encoding perspective. What about decoding perspective?

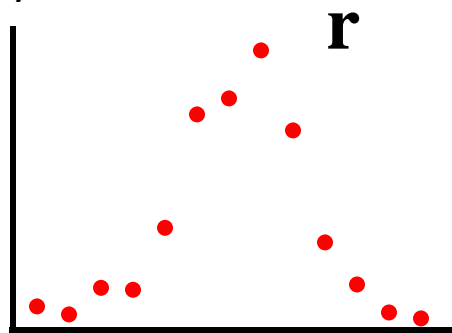
Bayesian decoding

- Bayes' rule:

$$p(s | \mathbf{r}) \propto p(\mathbf{r} | s) p(s)$$

posterior likelihood prior

Activity




Preferred stimulus

Bayes' rule




Probability




Stimulus

Bayesian decoding

single pixel
18 samples
 2^{18} stimuli



Decoder performance



Bayesian decoder under coupled model extracts about **20%** more information than under the uncoupled model.

Pairwise → just 10%, as in previous studies

Summary

- Signal versus noise correlations
- Shuffling trials is a way to study the effect of correlations on information.
- Correlations can increase or decrease information.
- If signal and noise correlations have the same sign, they tend to decrease information.
- Encoding versus decoding perspective
- Ignoring pairwise correlations reduces information by ~10%.
- Effect of pairwise correlations on entire population is not clear. Pillow et al. → 20% more information when exploiting full correlations
- LNP and coupled spiking models are convenient phenomenological models for correlated populations.

Exercises

- All in notes of lecture 1 (on class website)
- Due Saturday March 21 (end of day)
- “Bonus” exercises are optional.