Machine Learning Methods for Neural Data Analysis

Lecture 8: Summary Statistics and GLMs

Scott Linderman

STATS 220/320 (NBIO220, CS339N). Winter 2023.

Agenda

Unit II: Encoding and Decoding

- You’ve got spikes. Now what?
- Retinal ganglion cells
- Encoding RGC responses with generalized linear models (GLMs)

You’ve got spikes. Now what?

Plot your data!

- **Spike train**: neuron x time array of spike counts in each time bin (e.g. 10ms).
- **Empirical firing rate**: smooth the spike train (e.g. with a Gaussian kernel).
- **Sanity checks:**
  - Are the spike trains plausible (e.g. 1-50 spikes/sec)?
  - Do the firing rates look similar in the beginning, middle, and end of the recording?

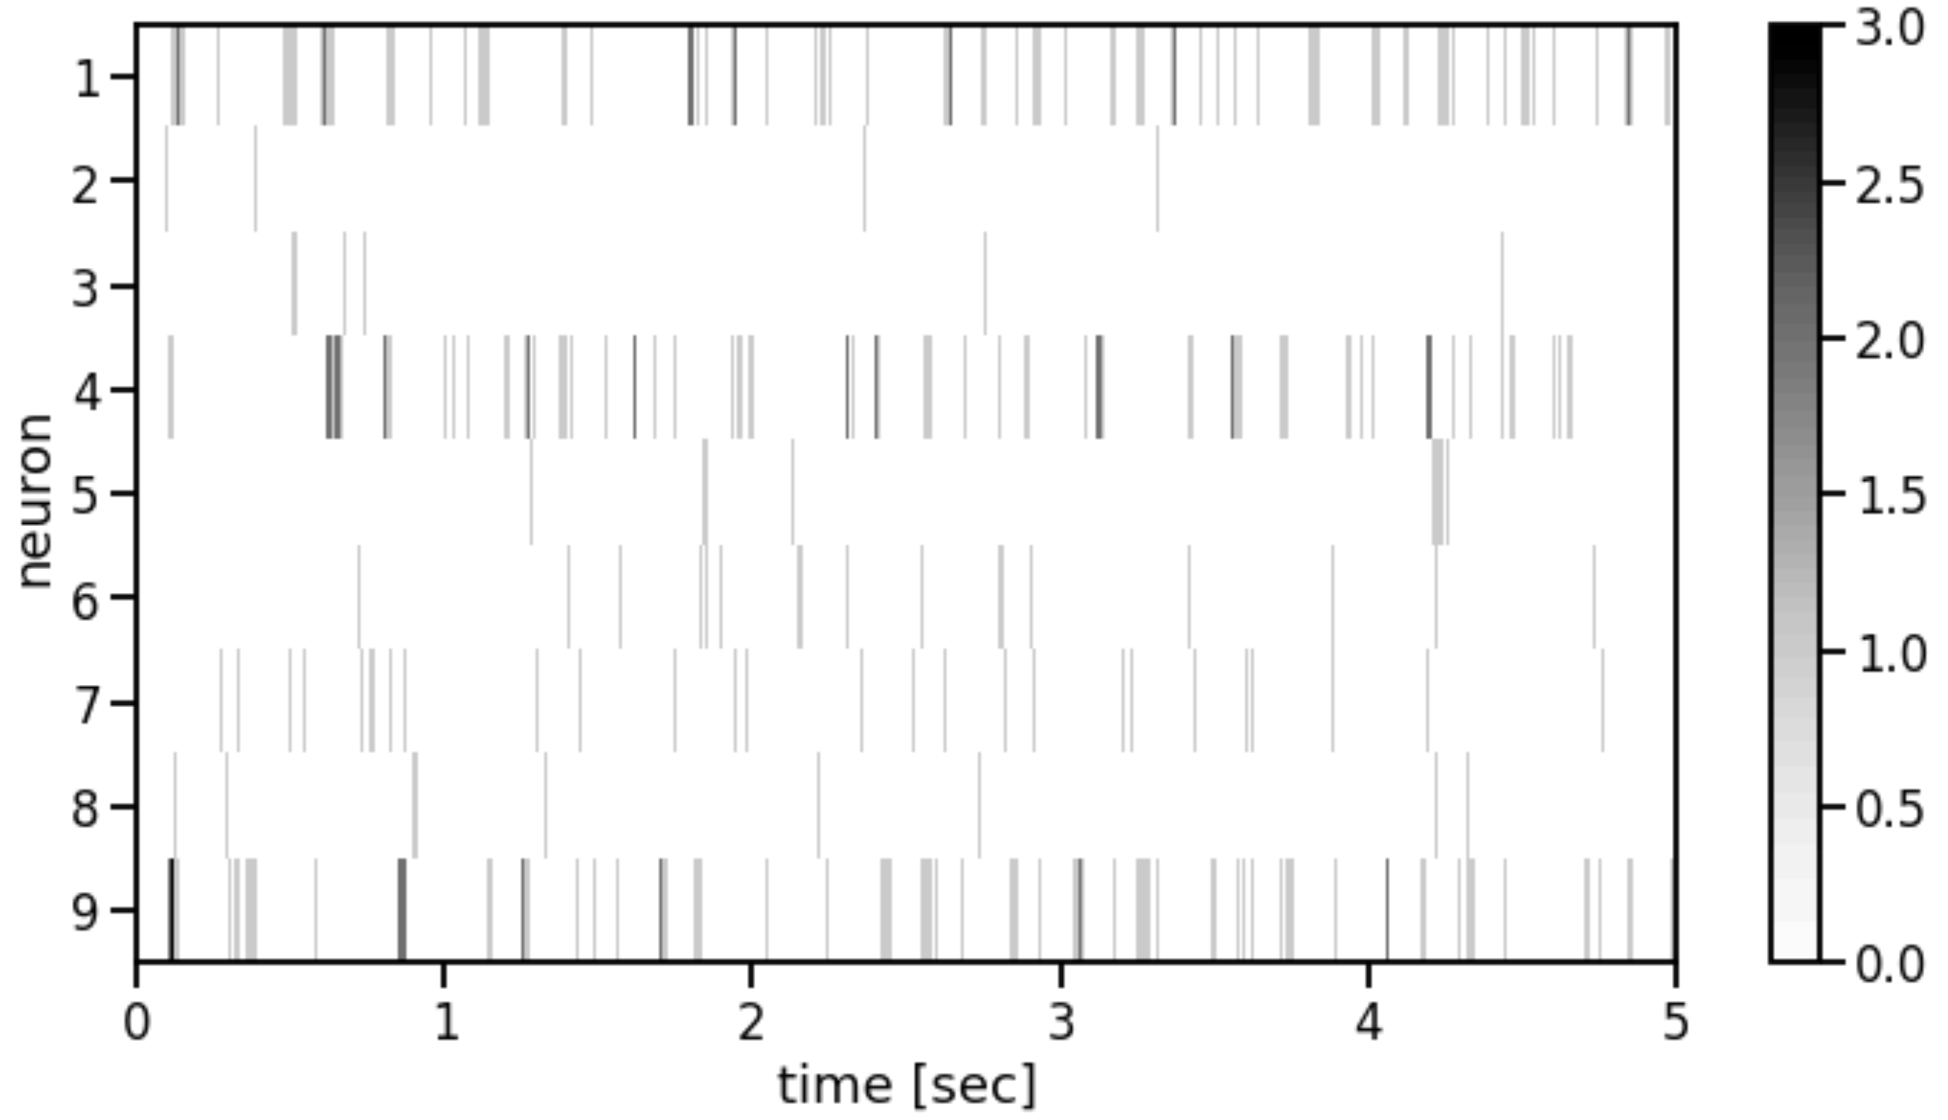

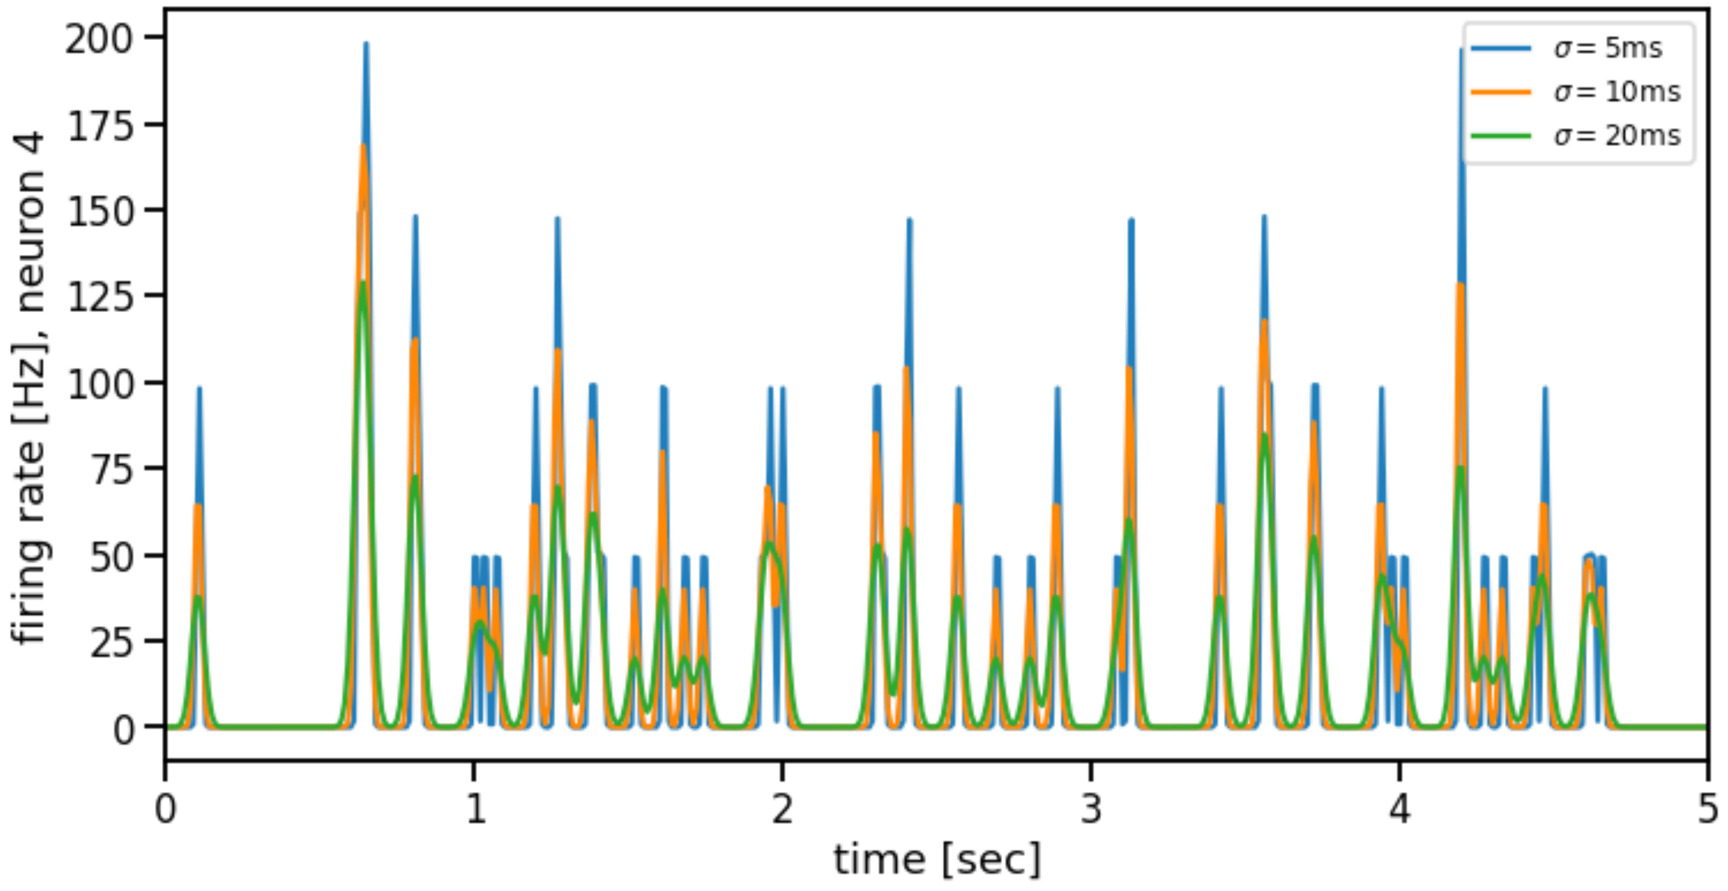

You've got spikes. Now what?

Spike triggered averages

- **Spike triggered average (STA):** What did the stimulus look like preceding (or surrounding) each spike?
- STA $\equiv$ conditional distribution of stimulus (x_{t-d}) given response ($y_t = 1$).

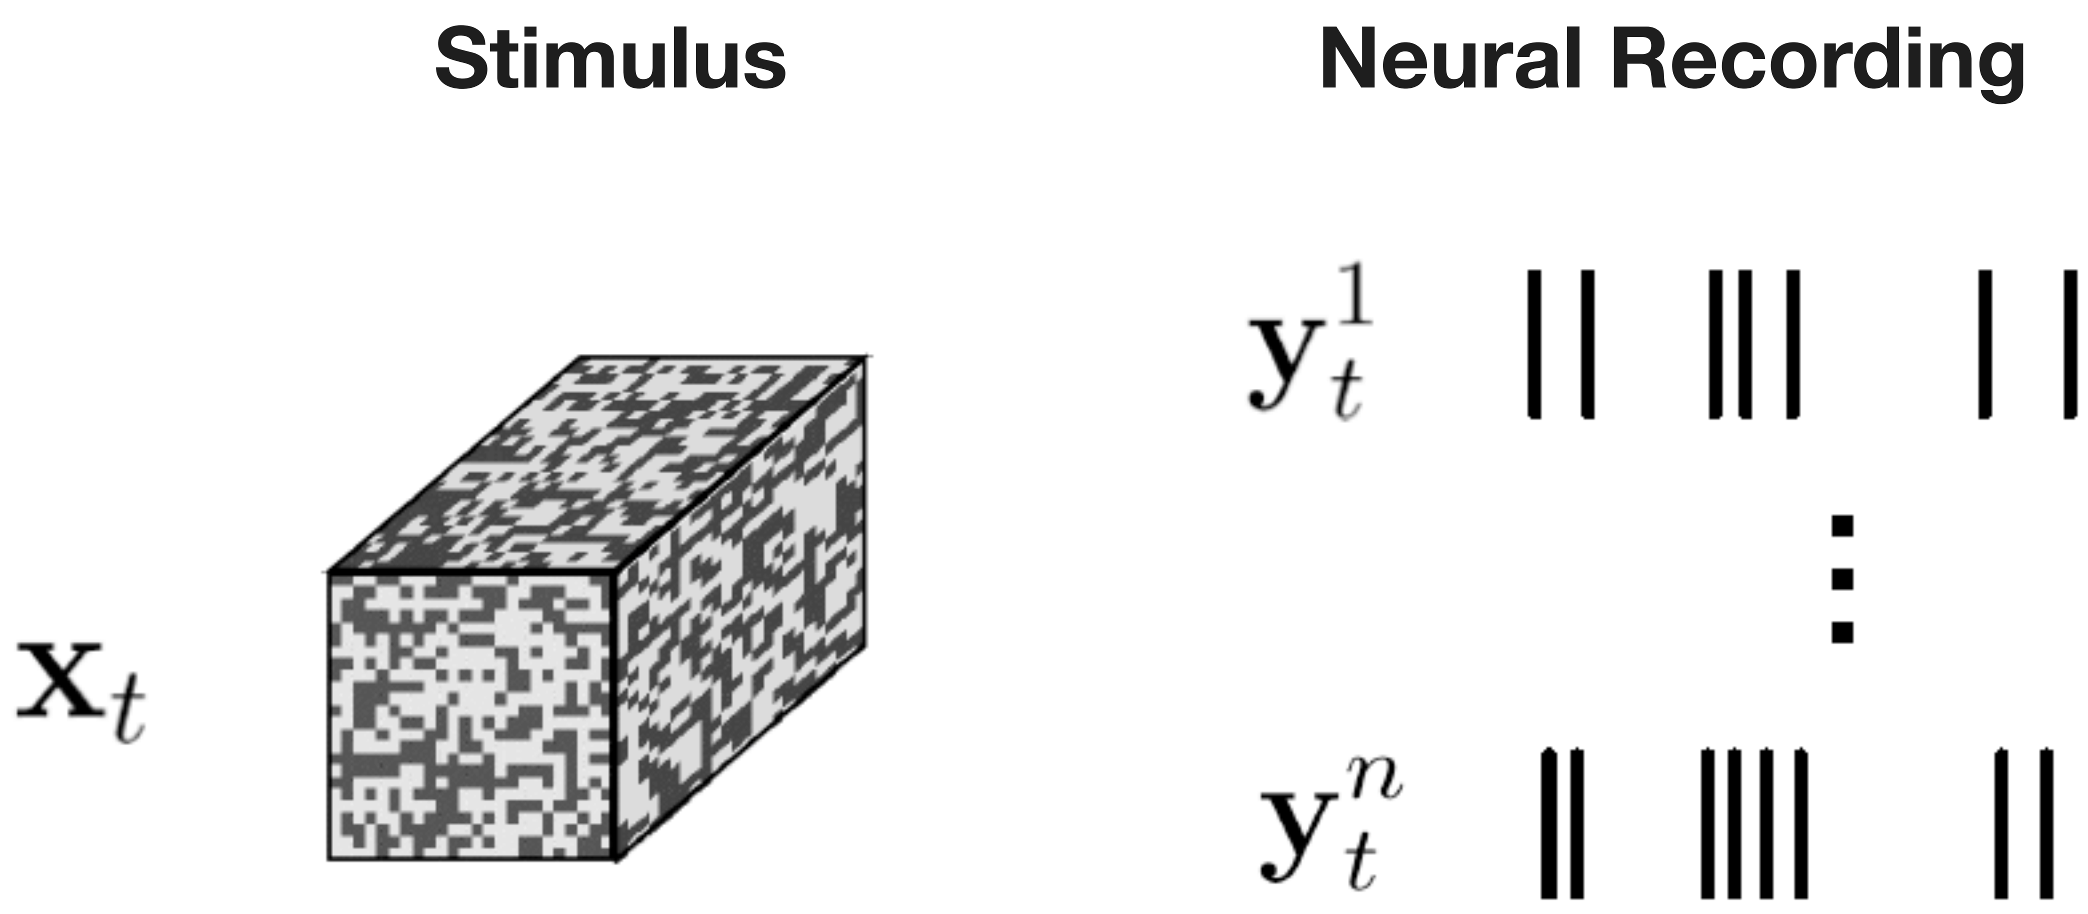

You've got spikes. Now what?

Spike triggered averages

- **Spike triggered average (STA):** What did the stimulus look like preceding (or surrounding) each spike?
- STA $\equiv$ conditional distribution of stimulus (x_{t-d}) given response ($y_t = 1$).
- **Receptive field**: portion of stimulus to which neuron responds.

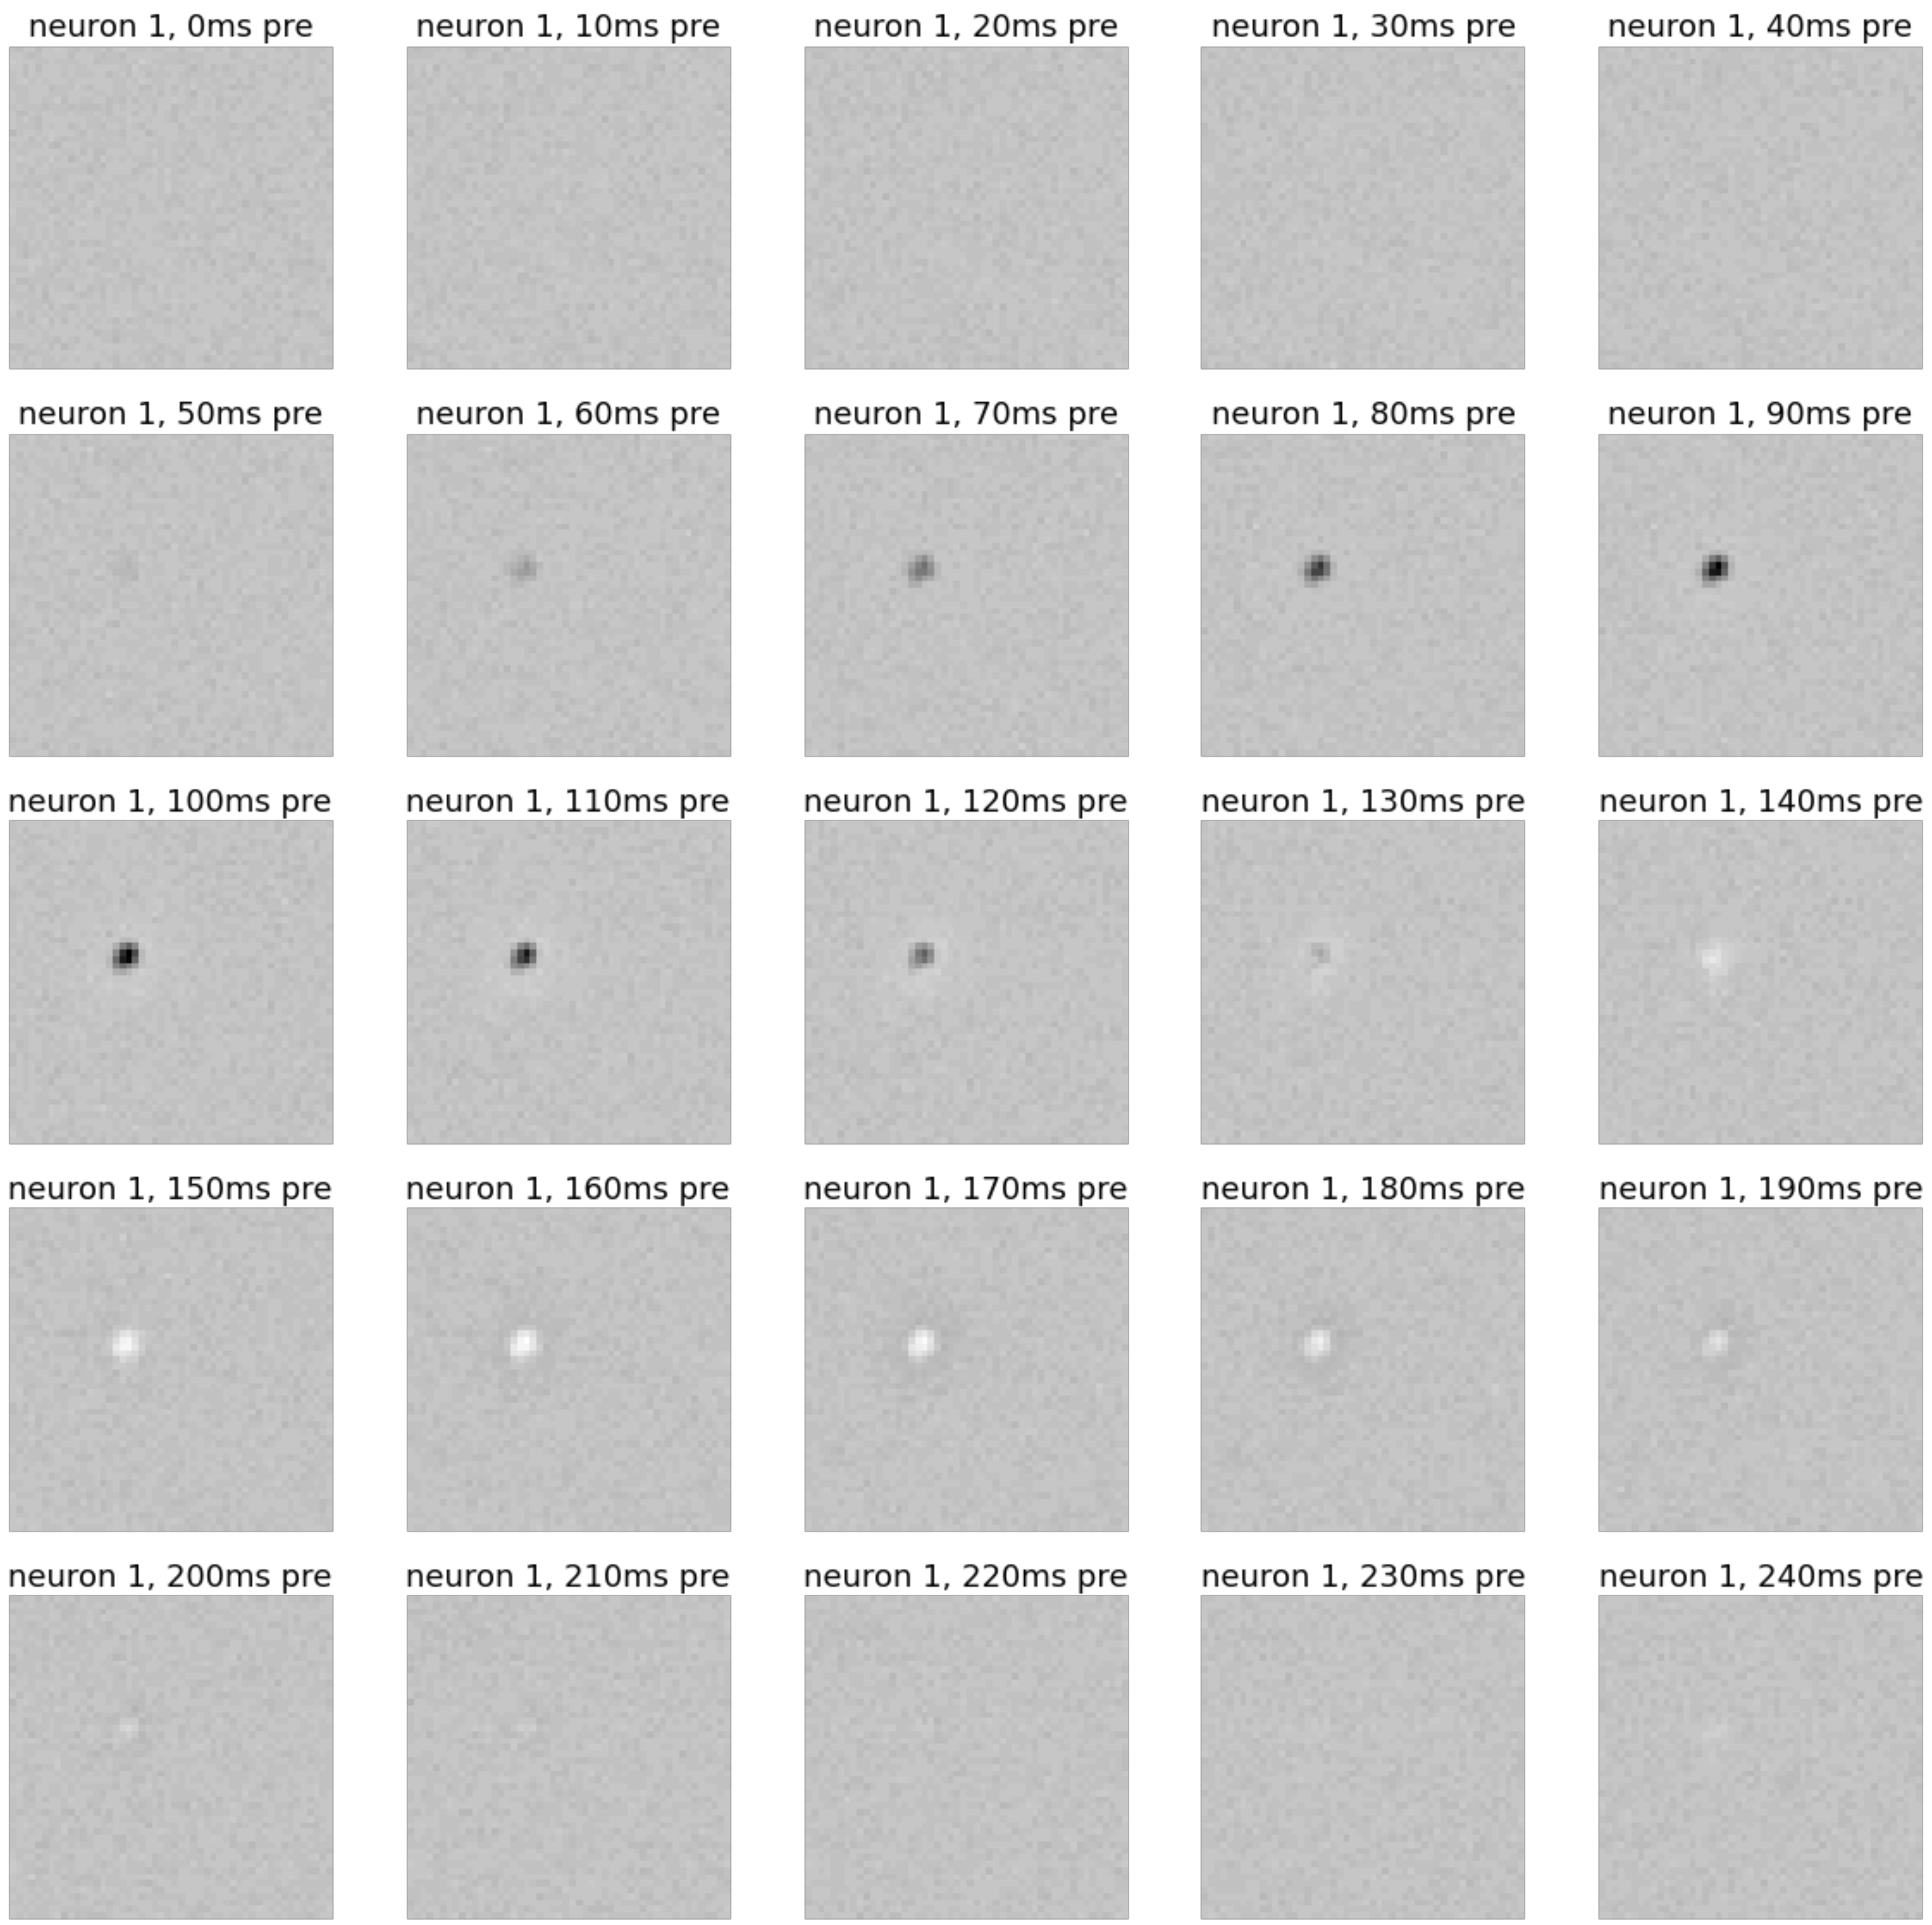

You’ve got spikes. Now what?

Peri-stimulus time histogram (PSTH)

- **Peri-stimulus time histogram (PSTH):** What did the response look like following (or surrounding) each stimulus presentation?

- PSTH $\equiv$ conditional distribution of spike train (y_{t+d}) given stimulus ($x_t = k$).

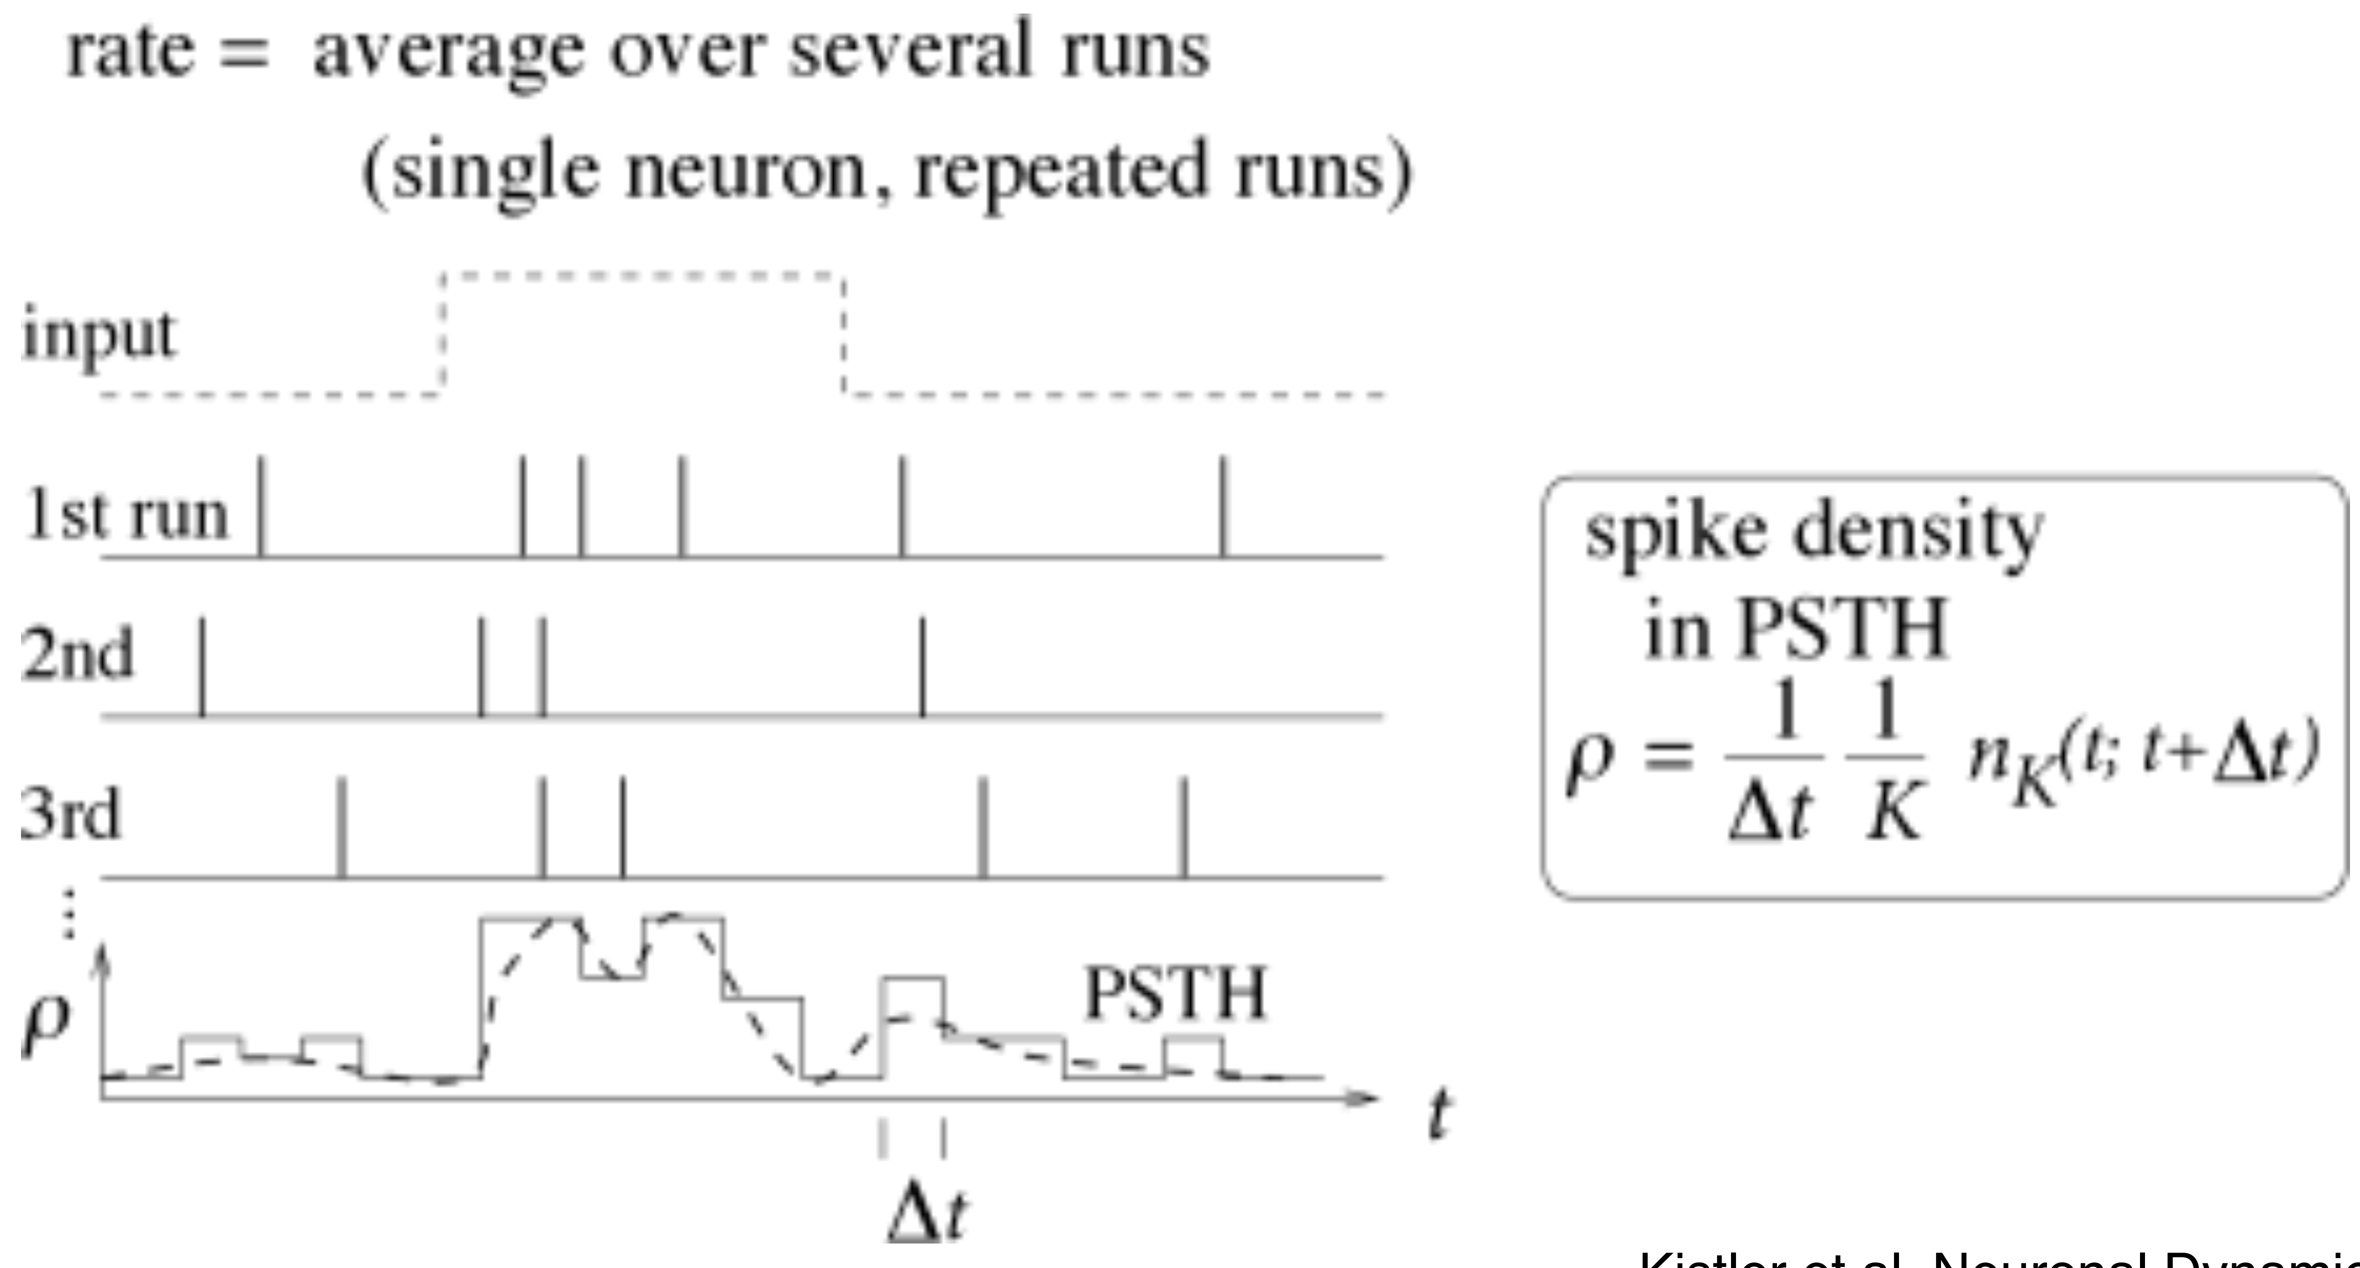

Kistler et al. Neuronal Dynamics

You’ve got spikes. Now what?

Cross-correlation function

- **Cross-correlation function (CCF):** what is the correlation between neuron n and neuron m as a function of delay of d?

$$\mathbb{E}\left[\frac{(y_{m,t}-\mu_m)}{\sigma_m}\frac{(y_{n,t+d}-\mu_n)}{\sigma_n}\right]$$

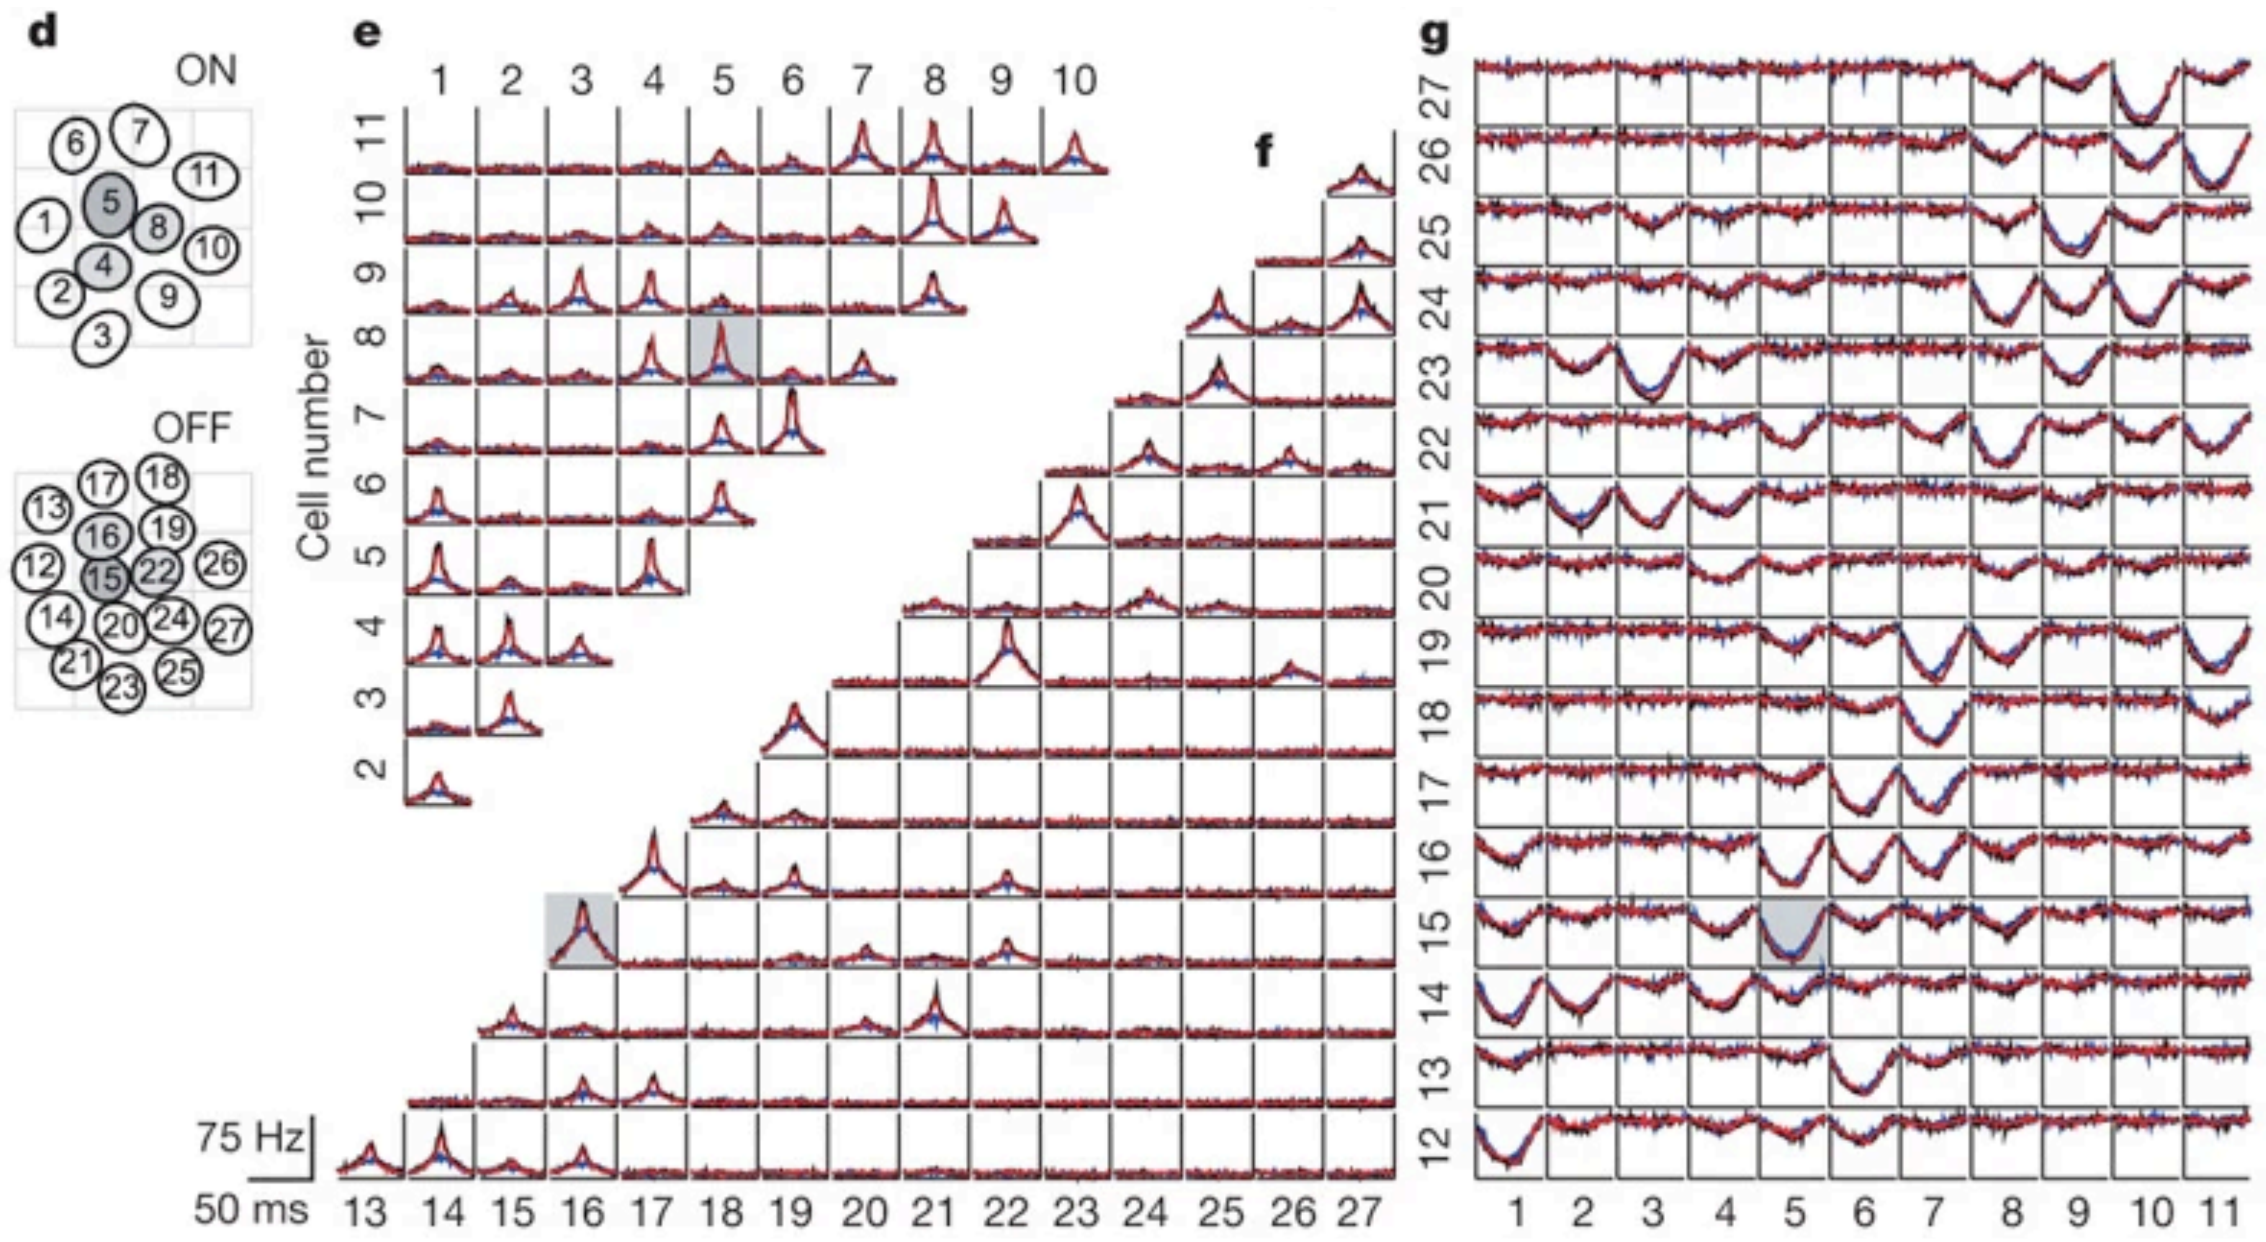

Pillow et al (*Nature*, 2008)

You’ve got spikes. Now what?

- A good model should **recapitulate these statistics** of the data.

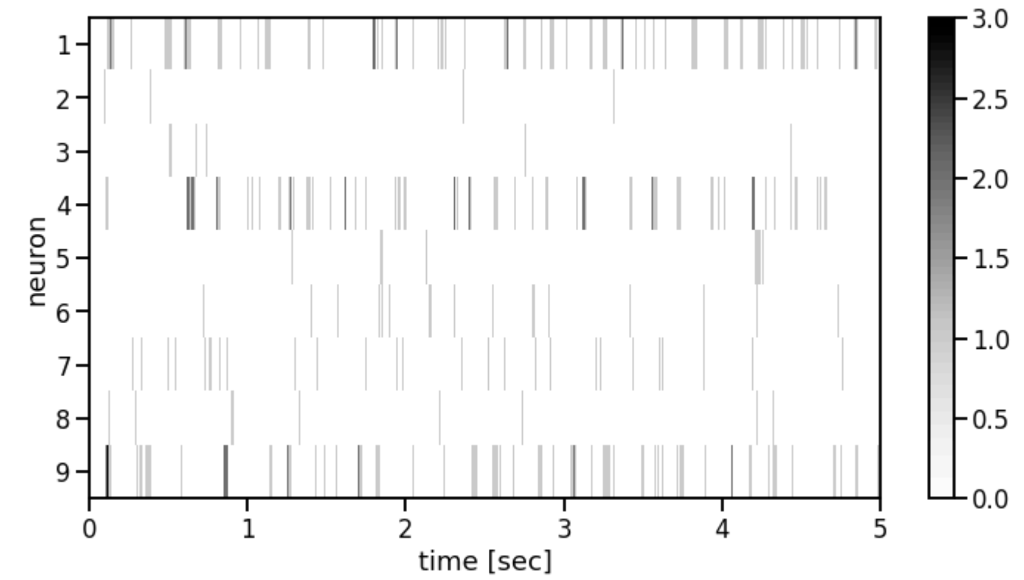

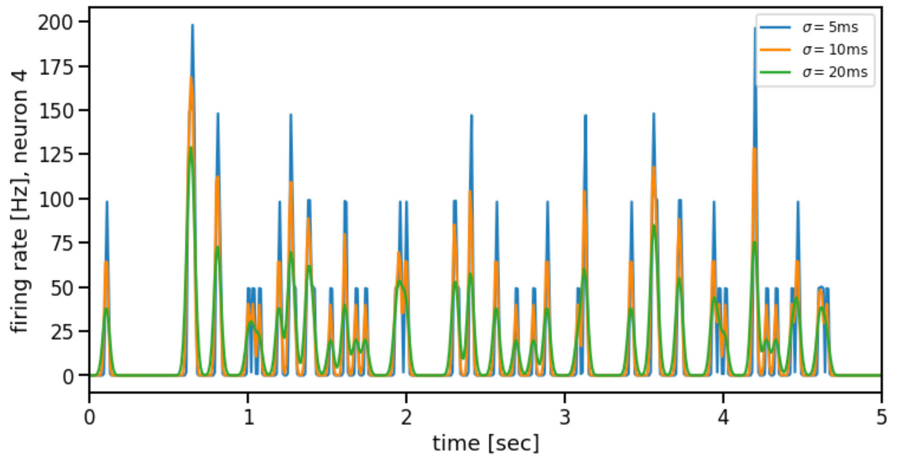

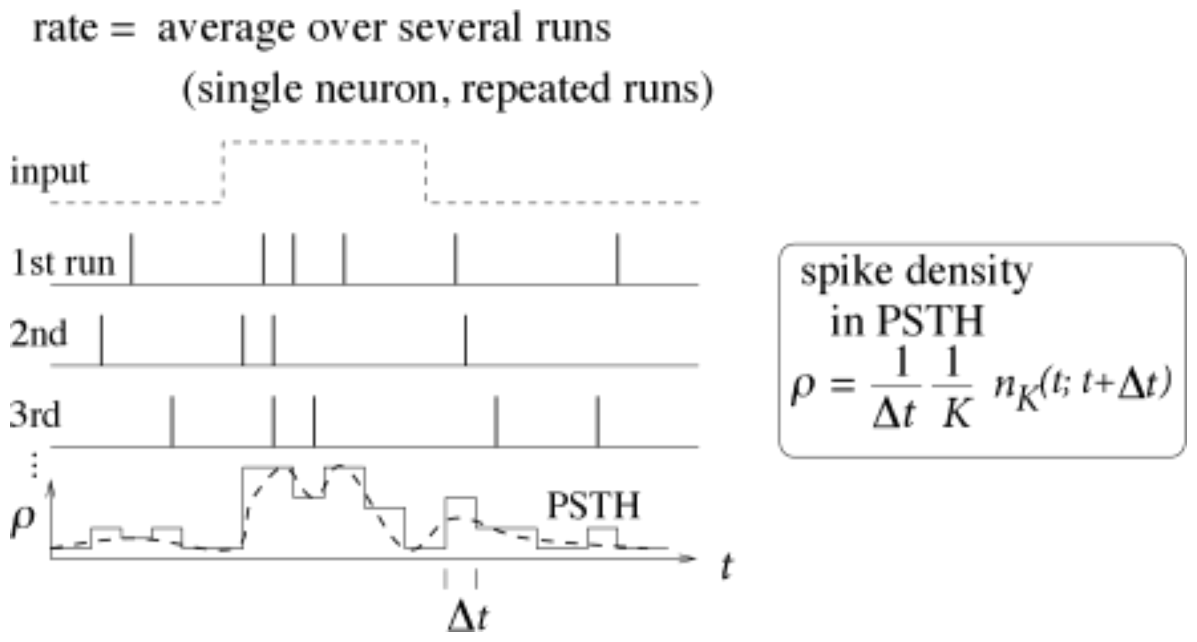

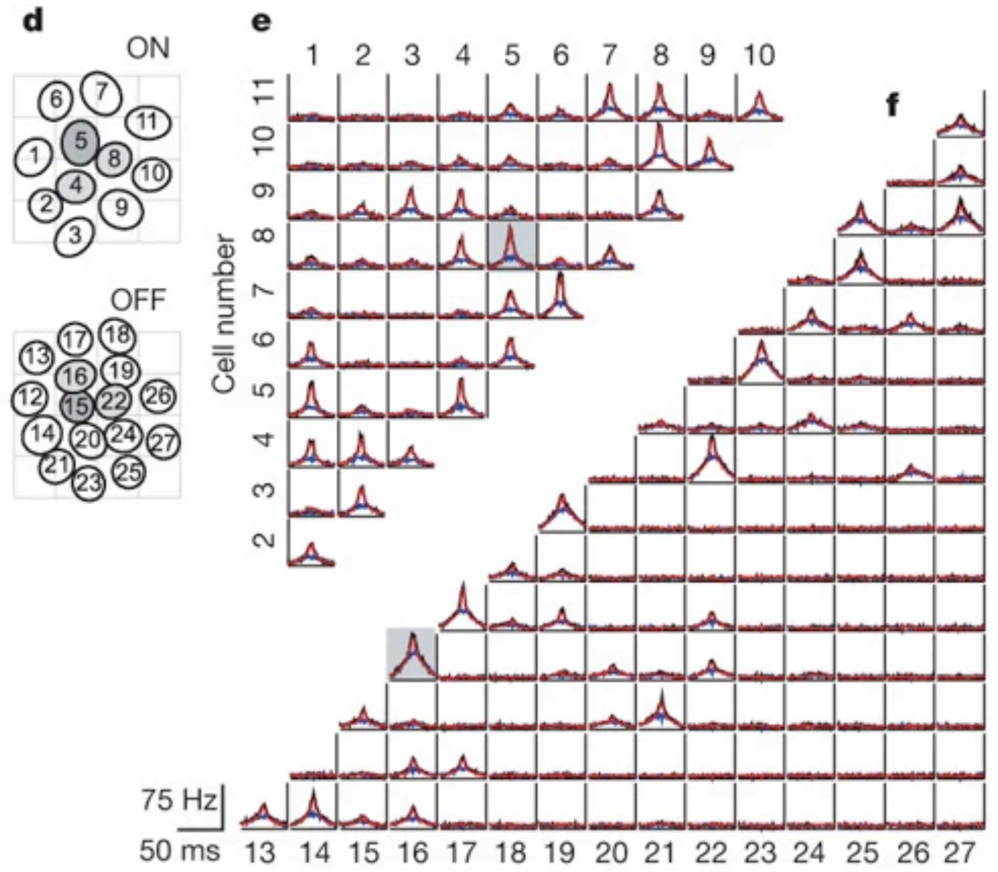

Retinal circuits

Retinal circuits

Basic architecture

- Incoming light stimulates **photoreceptors (rods and cones)** at the back of the eye.
- Rods and cones trigger an intermediate layer of **bipolar cells** and **amacrine cells.**
- Activity in these intermediate cells is pooled by **retinal ganglion cells (RGCs).**
- RGCs send action potentials down the **optic nerve** to the rest of the brain.
- The optic nerve innervates the **lateral geniculate nucleus (LGN)** of the thalamus and **primary visual cortex (V1)**

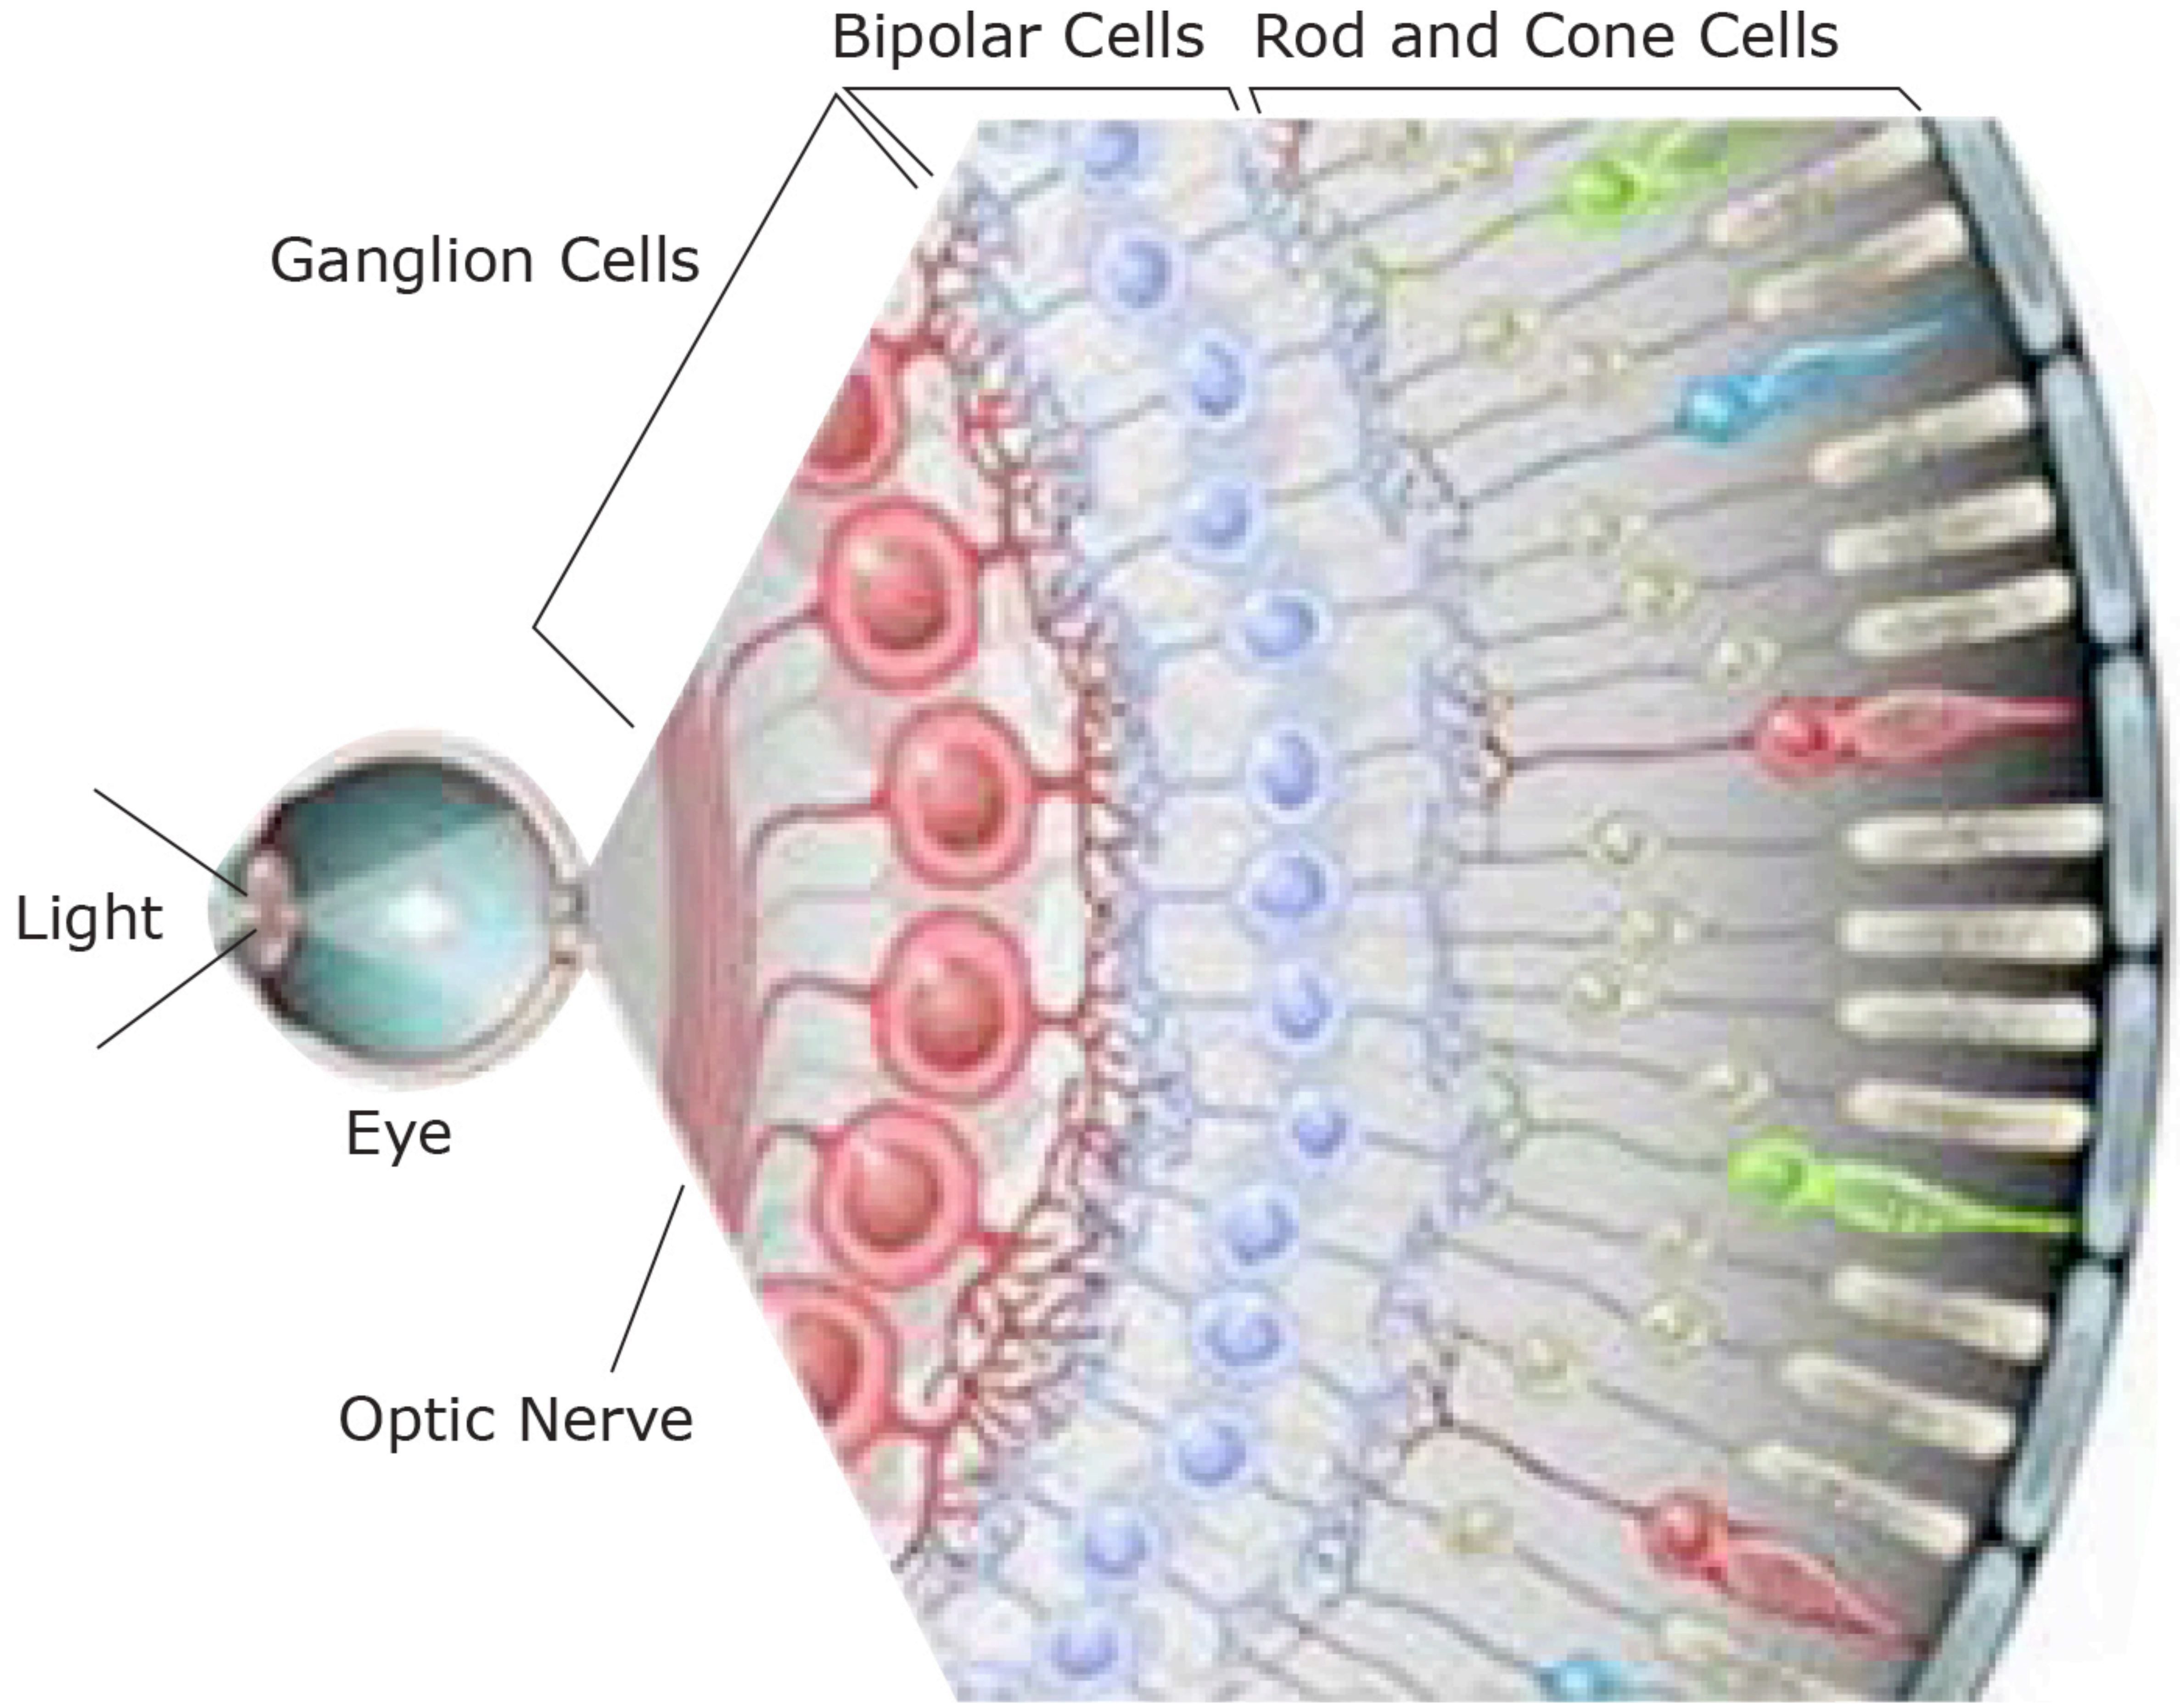

http://visionmagazineonline.co.za/2018/04/01/why-retinal-ganglion-cells-are-important-in-glaucoma/

Retinal circuits

Types of RGCs

- RGCs have been subdivided into dozens of types based on their morphology and their response properties.
- To first approximation, two main types: **ON and OFF cells**.
    - **ON cells** fire action potentials in response to **increased light intensity** at the center of their receptive field.
    - **OFF cells** fire action potentials in response to **decreased light intensity** at the center of their receptive field.
- Lots of **heterogeneity**; e.g. direction selective cells, transient and sustained responses, local edge detectors…

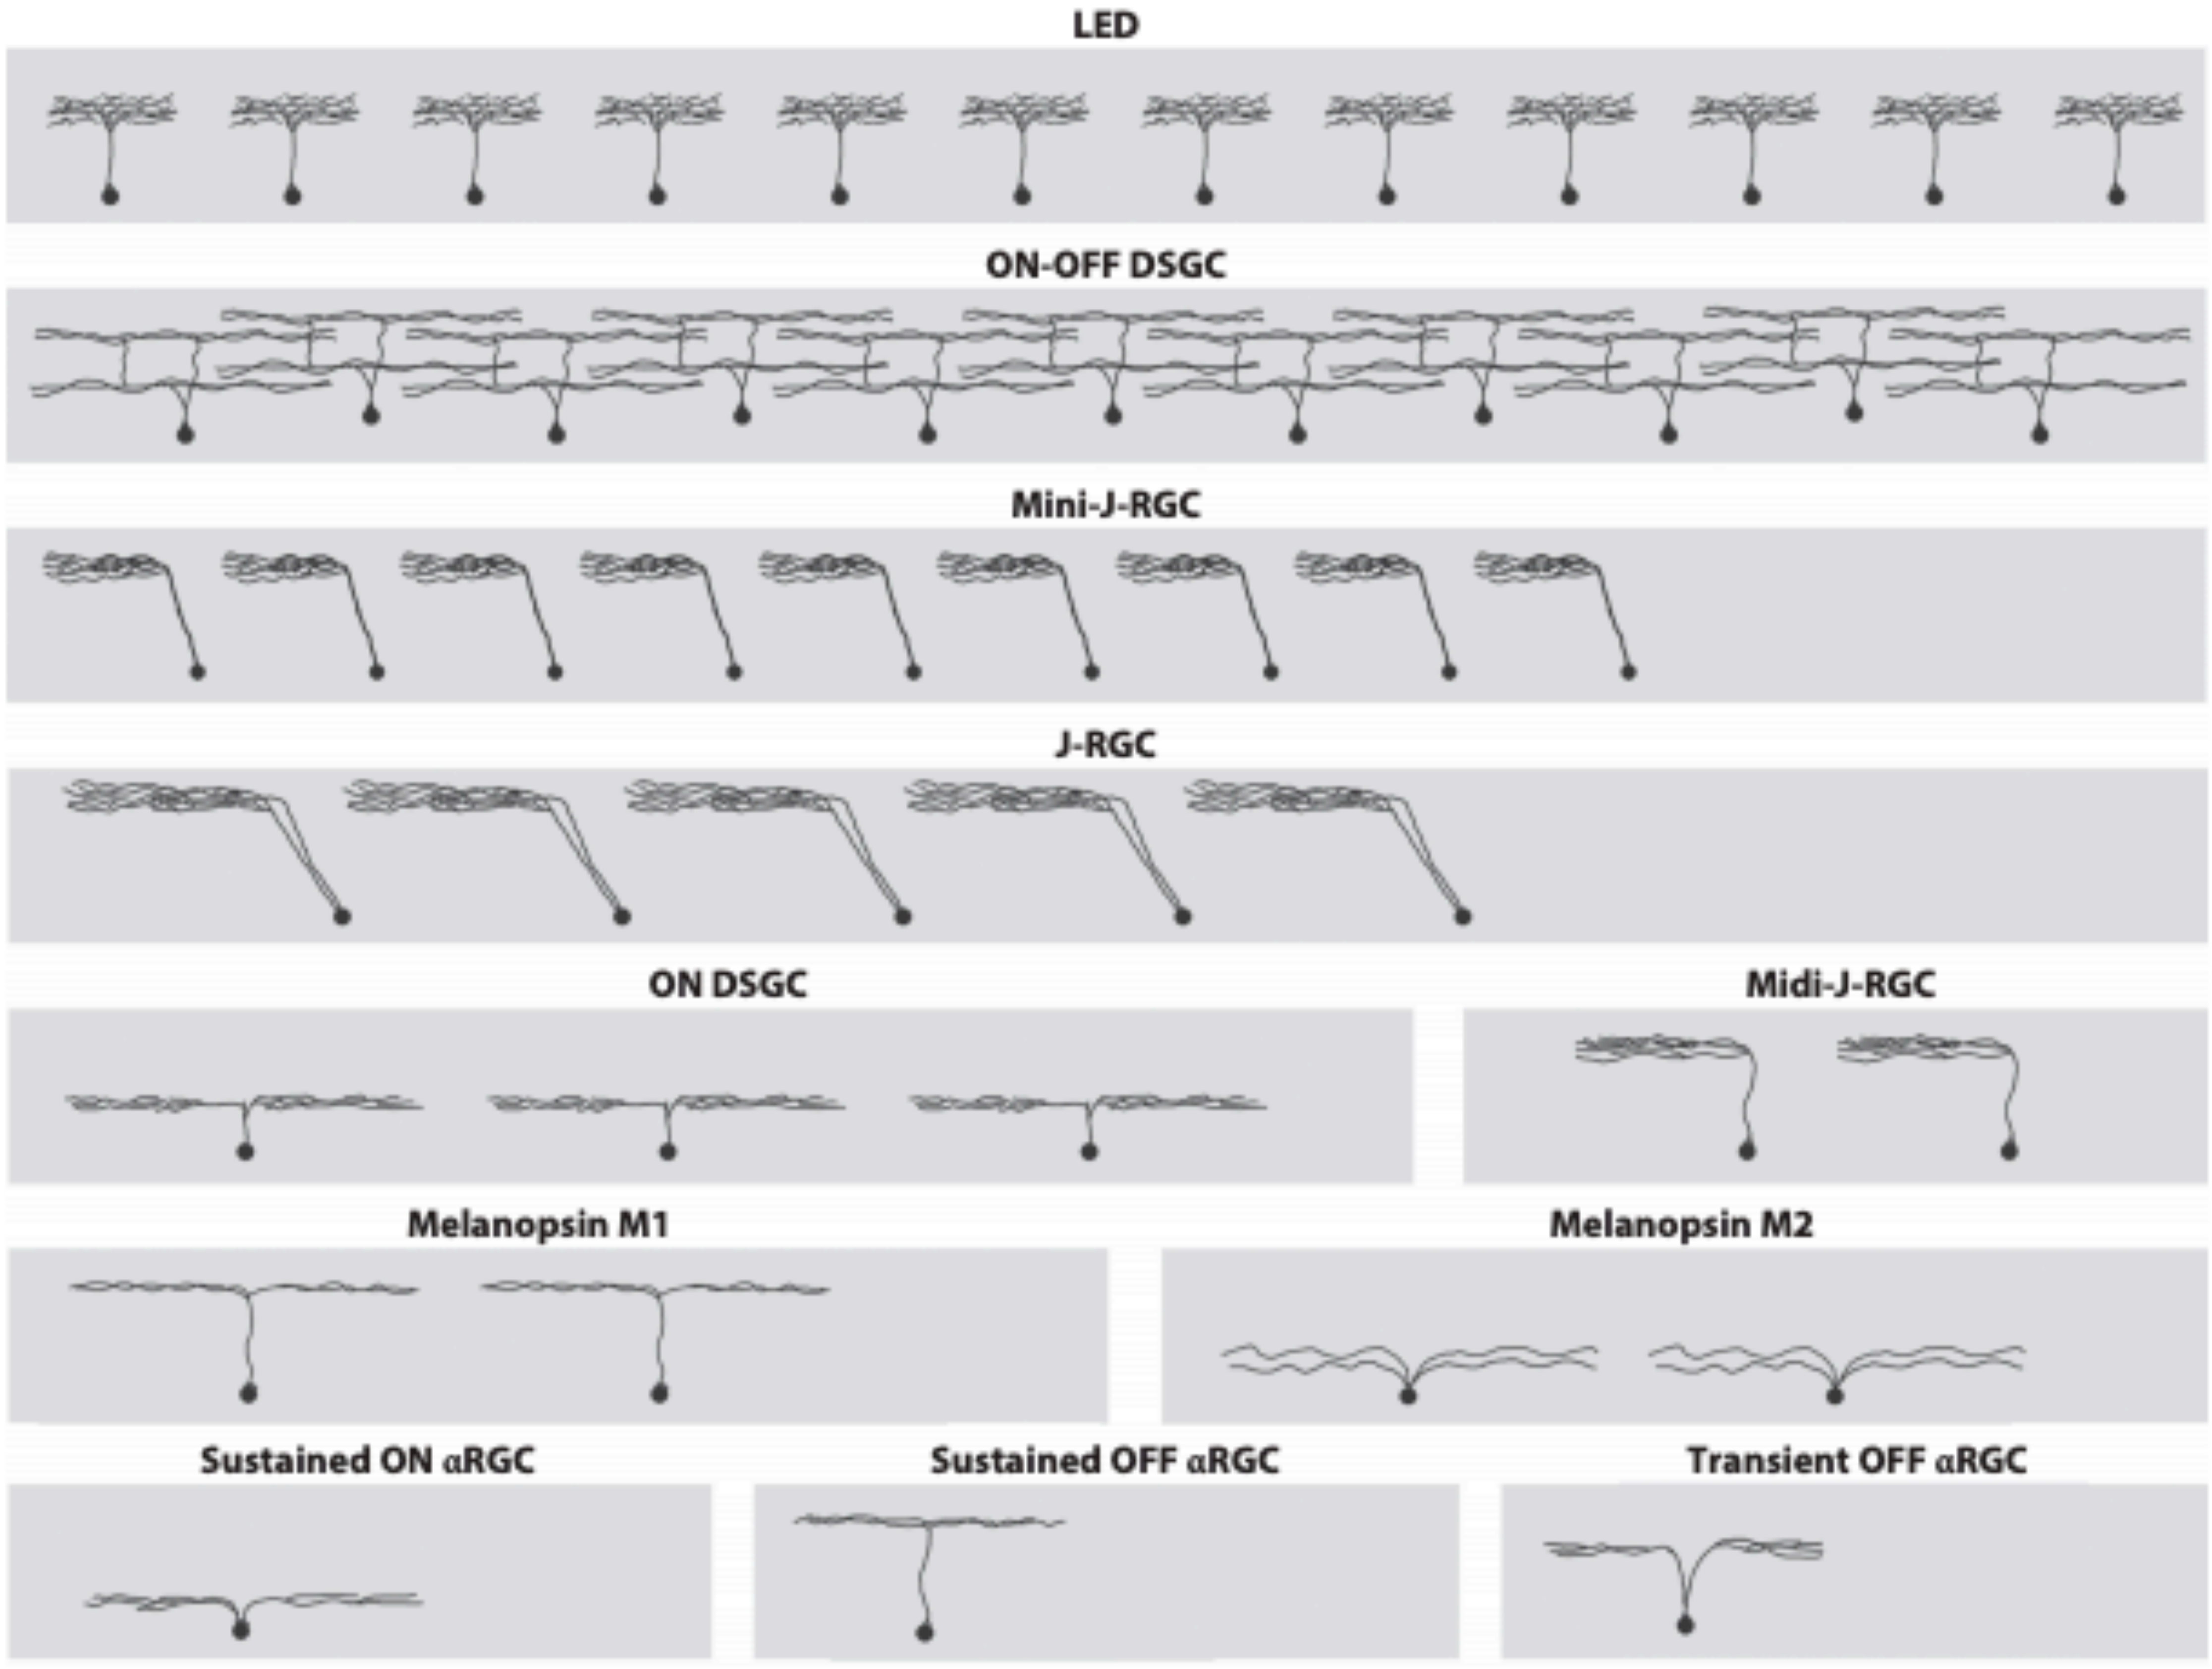

Sanes and Masland (*Ann. Rev. Neuro.*, 2014)

Retinal circuits

- **Key question**: *How are visual stimuli encoded in the output of these retinal circuits?*

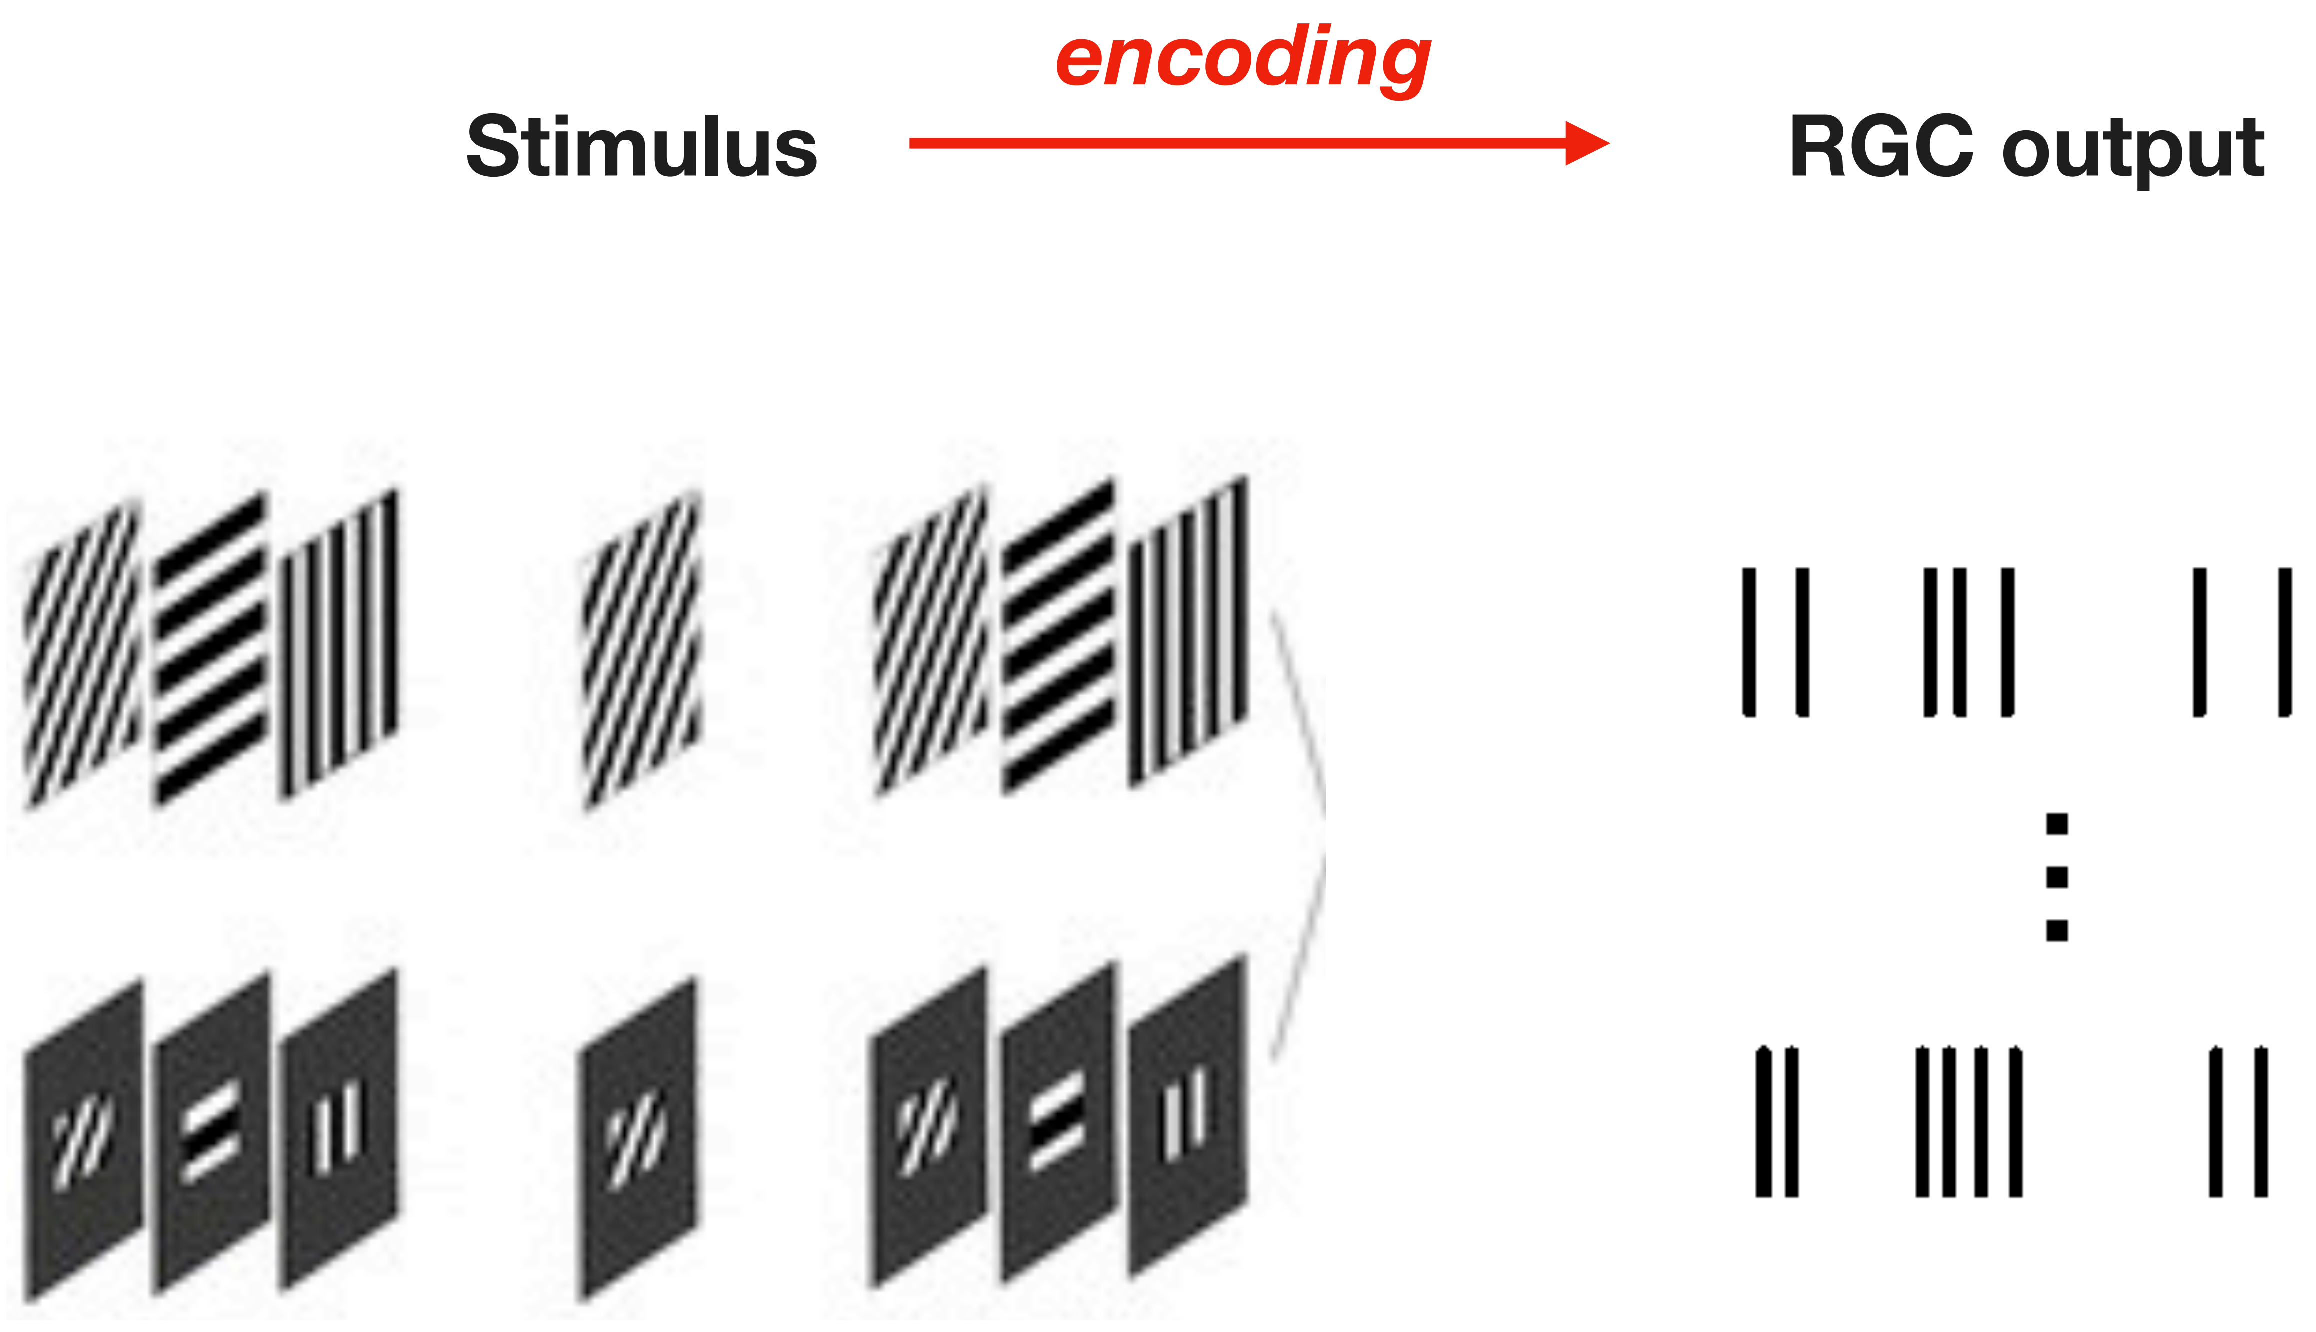

Generalized linear models of RGC responses

Encoding models of RGC responses

Basic linear-nonlinear-Poisson (LNP) model

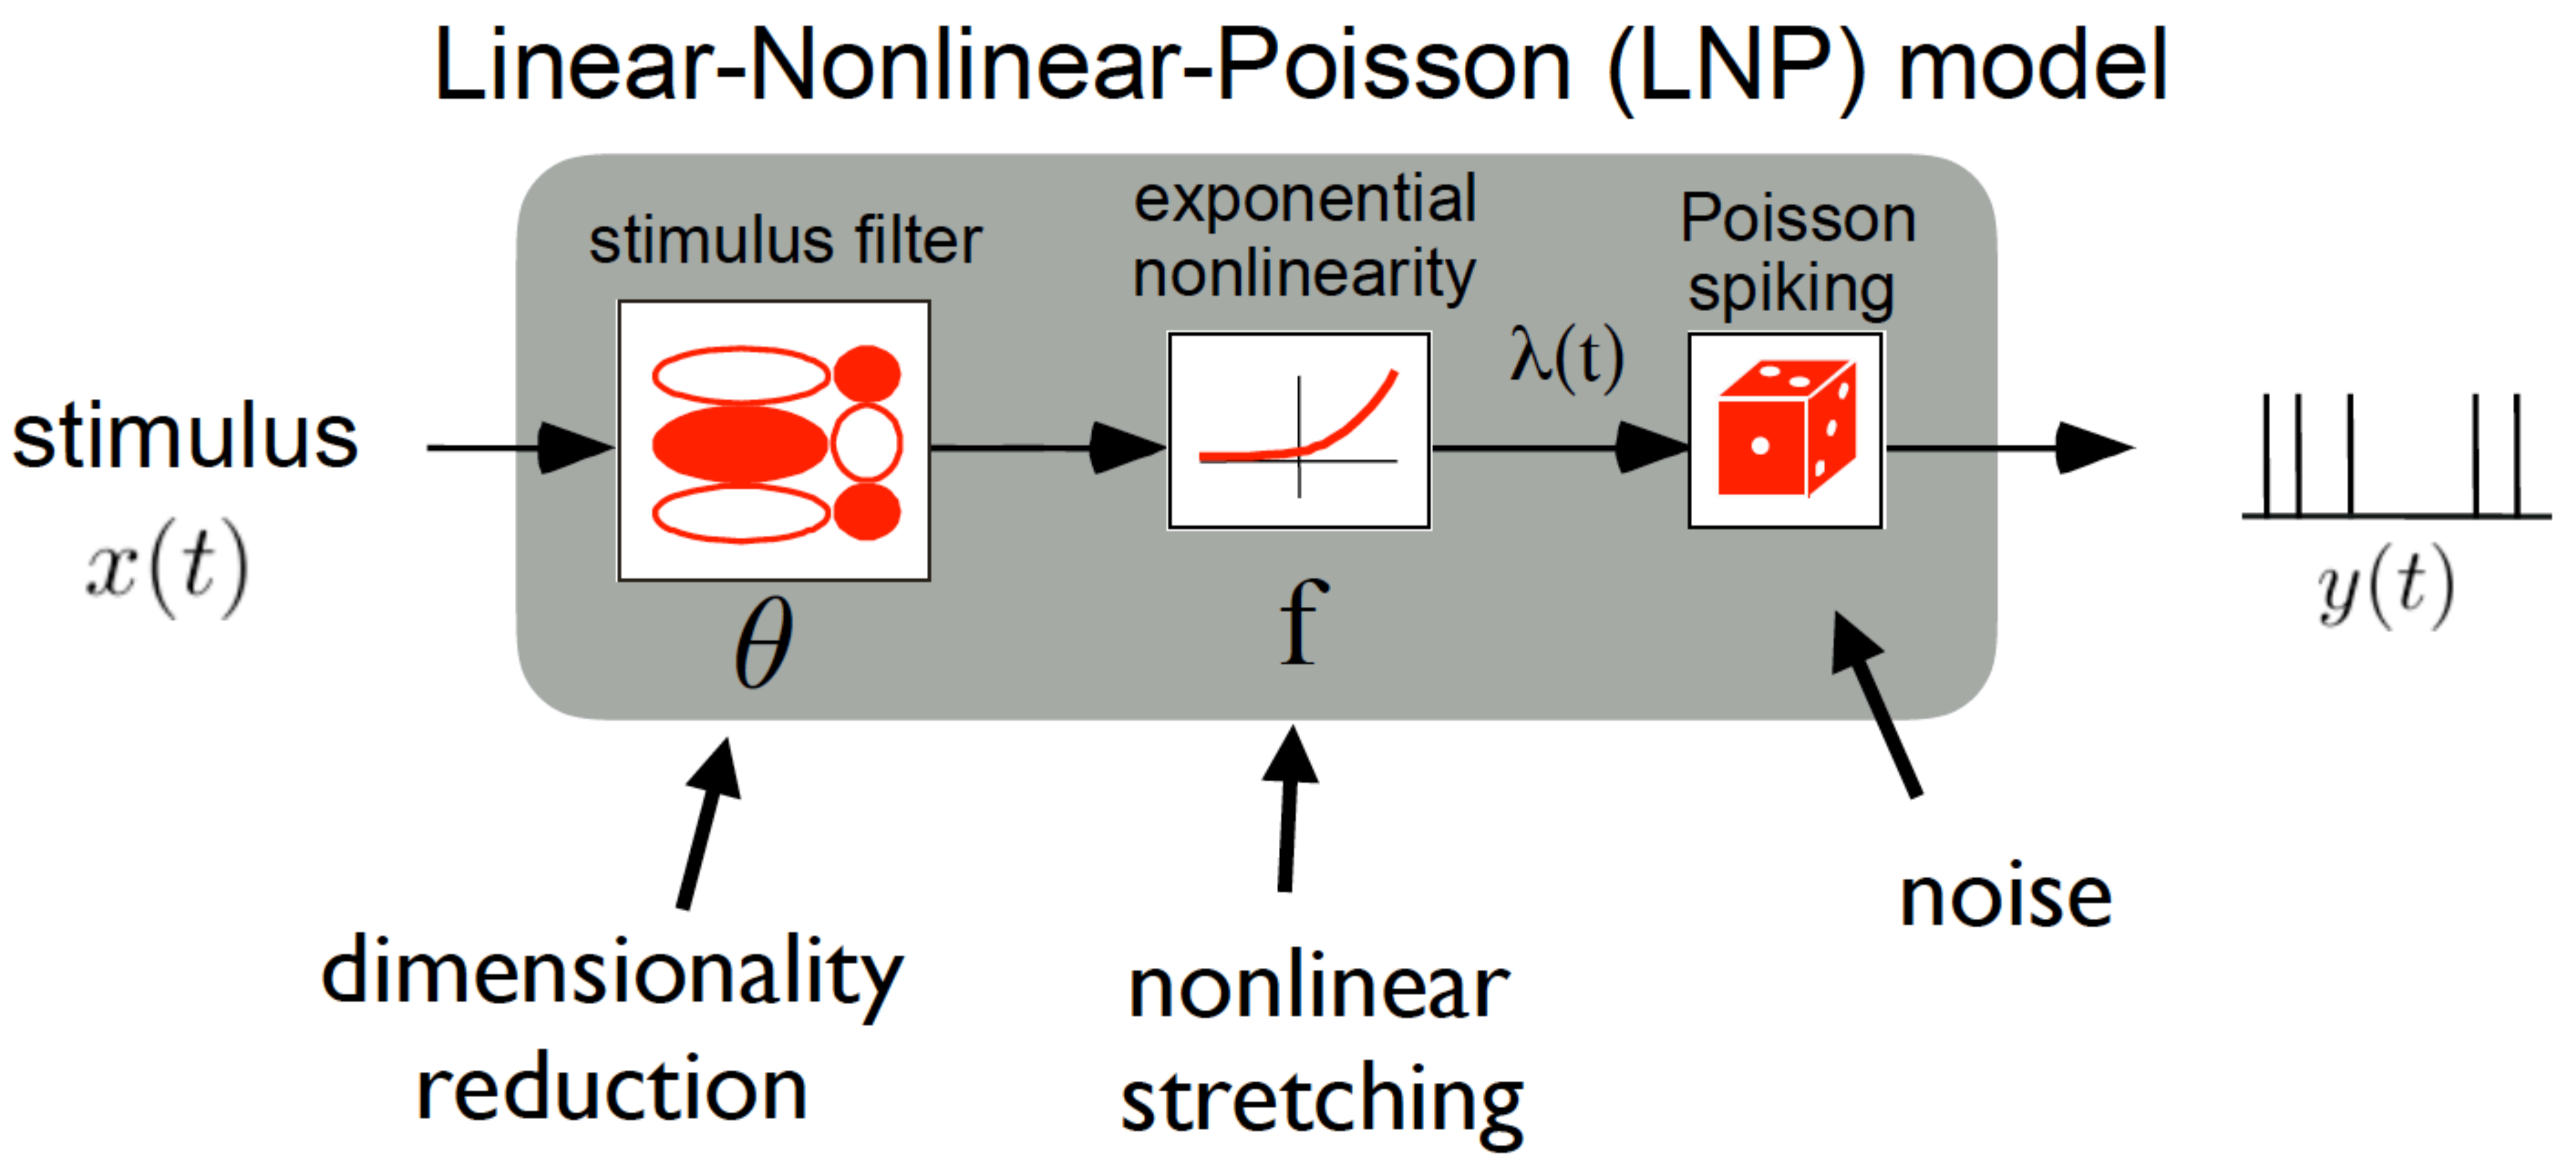

spike rate $\lambda = f(\vec{k} \cdot \vec{x})$

spike count $y \sim \text{Poiss}(\lambda)$

In statistics, we call this a generalized linear model (GLM).

Slides from Jonathan Pillow

Encoding models of RGC responses

First things first: Linear models

- Let $Y \in \mathbb{N}$ denote an integer-valued random variable; e.g. a spike count.
- Let $X \in \mathbb{R}^p$ be a p-dimensional feature vector.
- **Linear regression** estimates the conditional expectation $\mu(X) \triangleq \mathbb{E}[Y \mid X]$ via a linear function $\hat{\mu} \triangleq \beta^\top X$, where $\beta \in \mathbb{R}^p$ is a vector of regression weights.
- **Question: What are some shortcomings of this linear regression model?**

Encoding models of RGC responses

Linear models

- Let $Y \in \mathbb{N}$ denote an integer-valued random variable; e.g. a spike count.
- Let $X \in \mathbb{R}^p$ be a p-dimensional feature vector.
- **Linear regression** estimates the conditional expectation $\mu(X) \triangleq \mathbb{E}[Y \mid X]$ via a linear function $\hat{\mu} \triangleq \beta^\top X$, where $\beta \in \mathbb{R}^p$ is a vector of regression weights.
- **Question: What are some shortcomings of this linear regression model?**
  - **Negative means**
  - **Symmetric and homoskedastic noise (same variance for each input)**

Encoding models of RGC responses

Generalized linear models

- **Generalized linear models** address these shortcomings with a simple tweak:
- Let $\hat{\eta}(X) \triangleq \beta^\top X$ be a **linear predictor** defined by parameter β, which we will estimate.
- Map the predictor through a monotonic, continuous, non-linear **mean function** $g(\,\cdot\,) : \mathbb{R} \to \mathscr{M}$, where $\mathscr{M}$ is the space of conditional expectations of Y.
  - E.g. If $Y \in \mathbb{N}$ is a non-negative integer its expectations lie in $\mathscr{M} = \mathbb{R}_+$, so we might take $g(a) = e^a$.
- The inverse of the mean function, $g^{-1} : \mathscr{M} \to \mathbb{R}$, is called the **link function.**
- Finally, plug the conditional expectation into a **conditional distribution** of Y given X.
  - E.g. $Y \mid X \sim \mathrm{Po}(g(\eta(X))$
- Generally, we assume the conditional distribution is a member of the **exponential family**.

Encoding models of RGC responses

Generalized linear models

- Let $X \in \mathbb{R}^{T \times P_H \times P_W}$ denote a stimulus movie and $Y \in \mathbb{N}^{N \times T}$ denote the resulting spike train.
- Define a **Poisson GLM** to predict (encode) neural responses given the past D frames of stimulus.

Encoding models

Generalized linear models

Stimulus **Filter Weights** **Firing Rates** **Neural Recording**

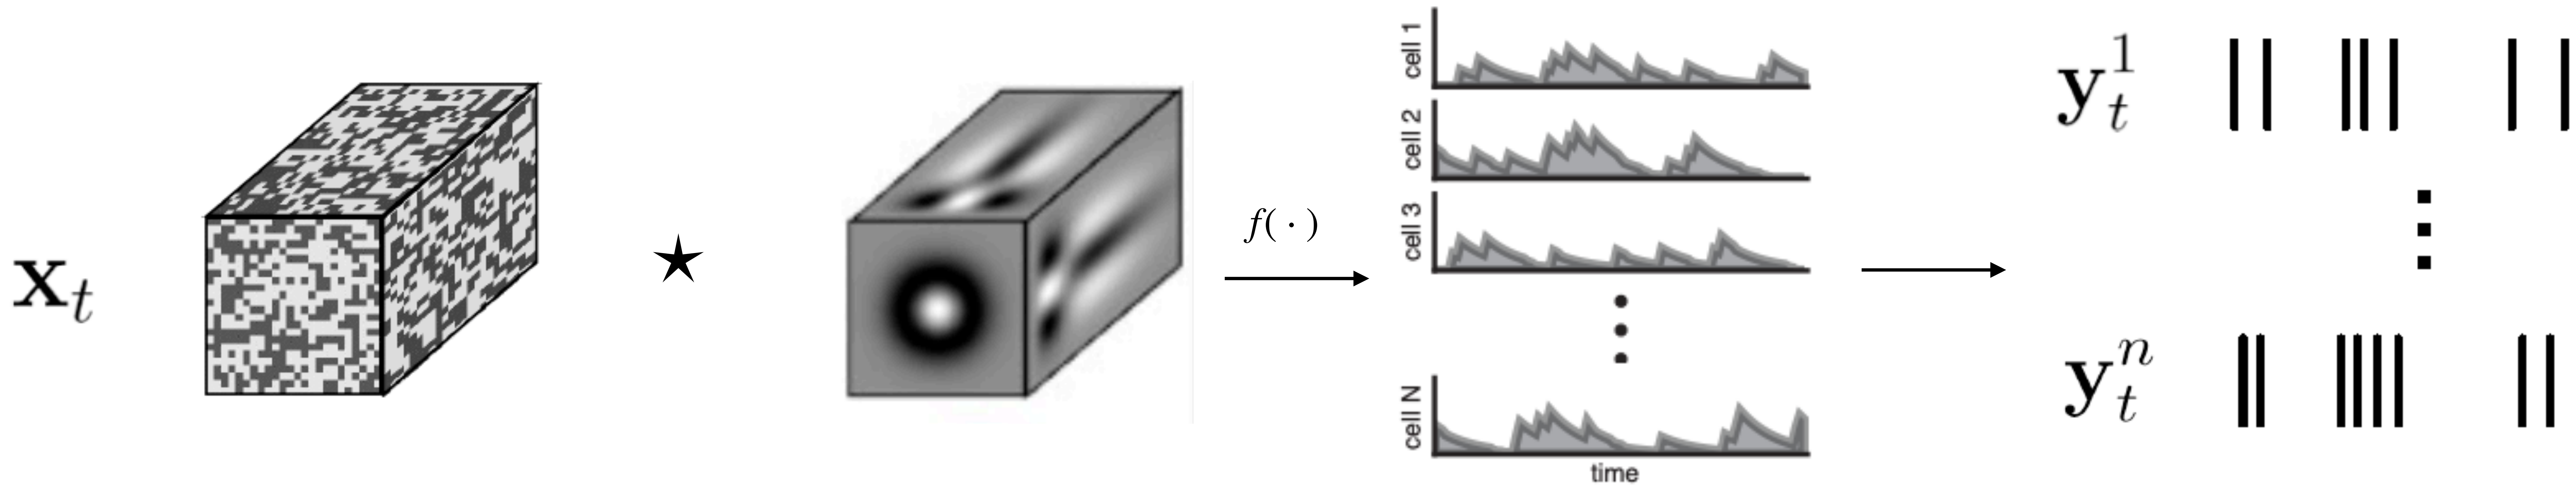

W_n

$$\lambda_{nt} = g([X \star W_n]_t)$$

Slides from Jonathan Pillow

Encoding models of RGC responses

Adding coupling between neurons

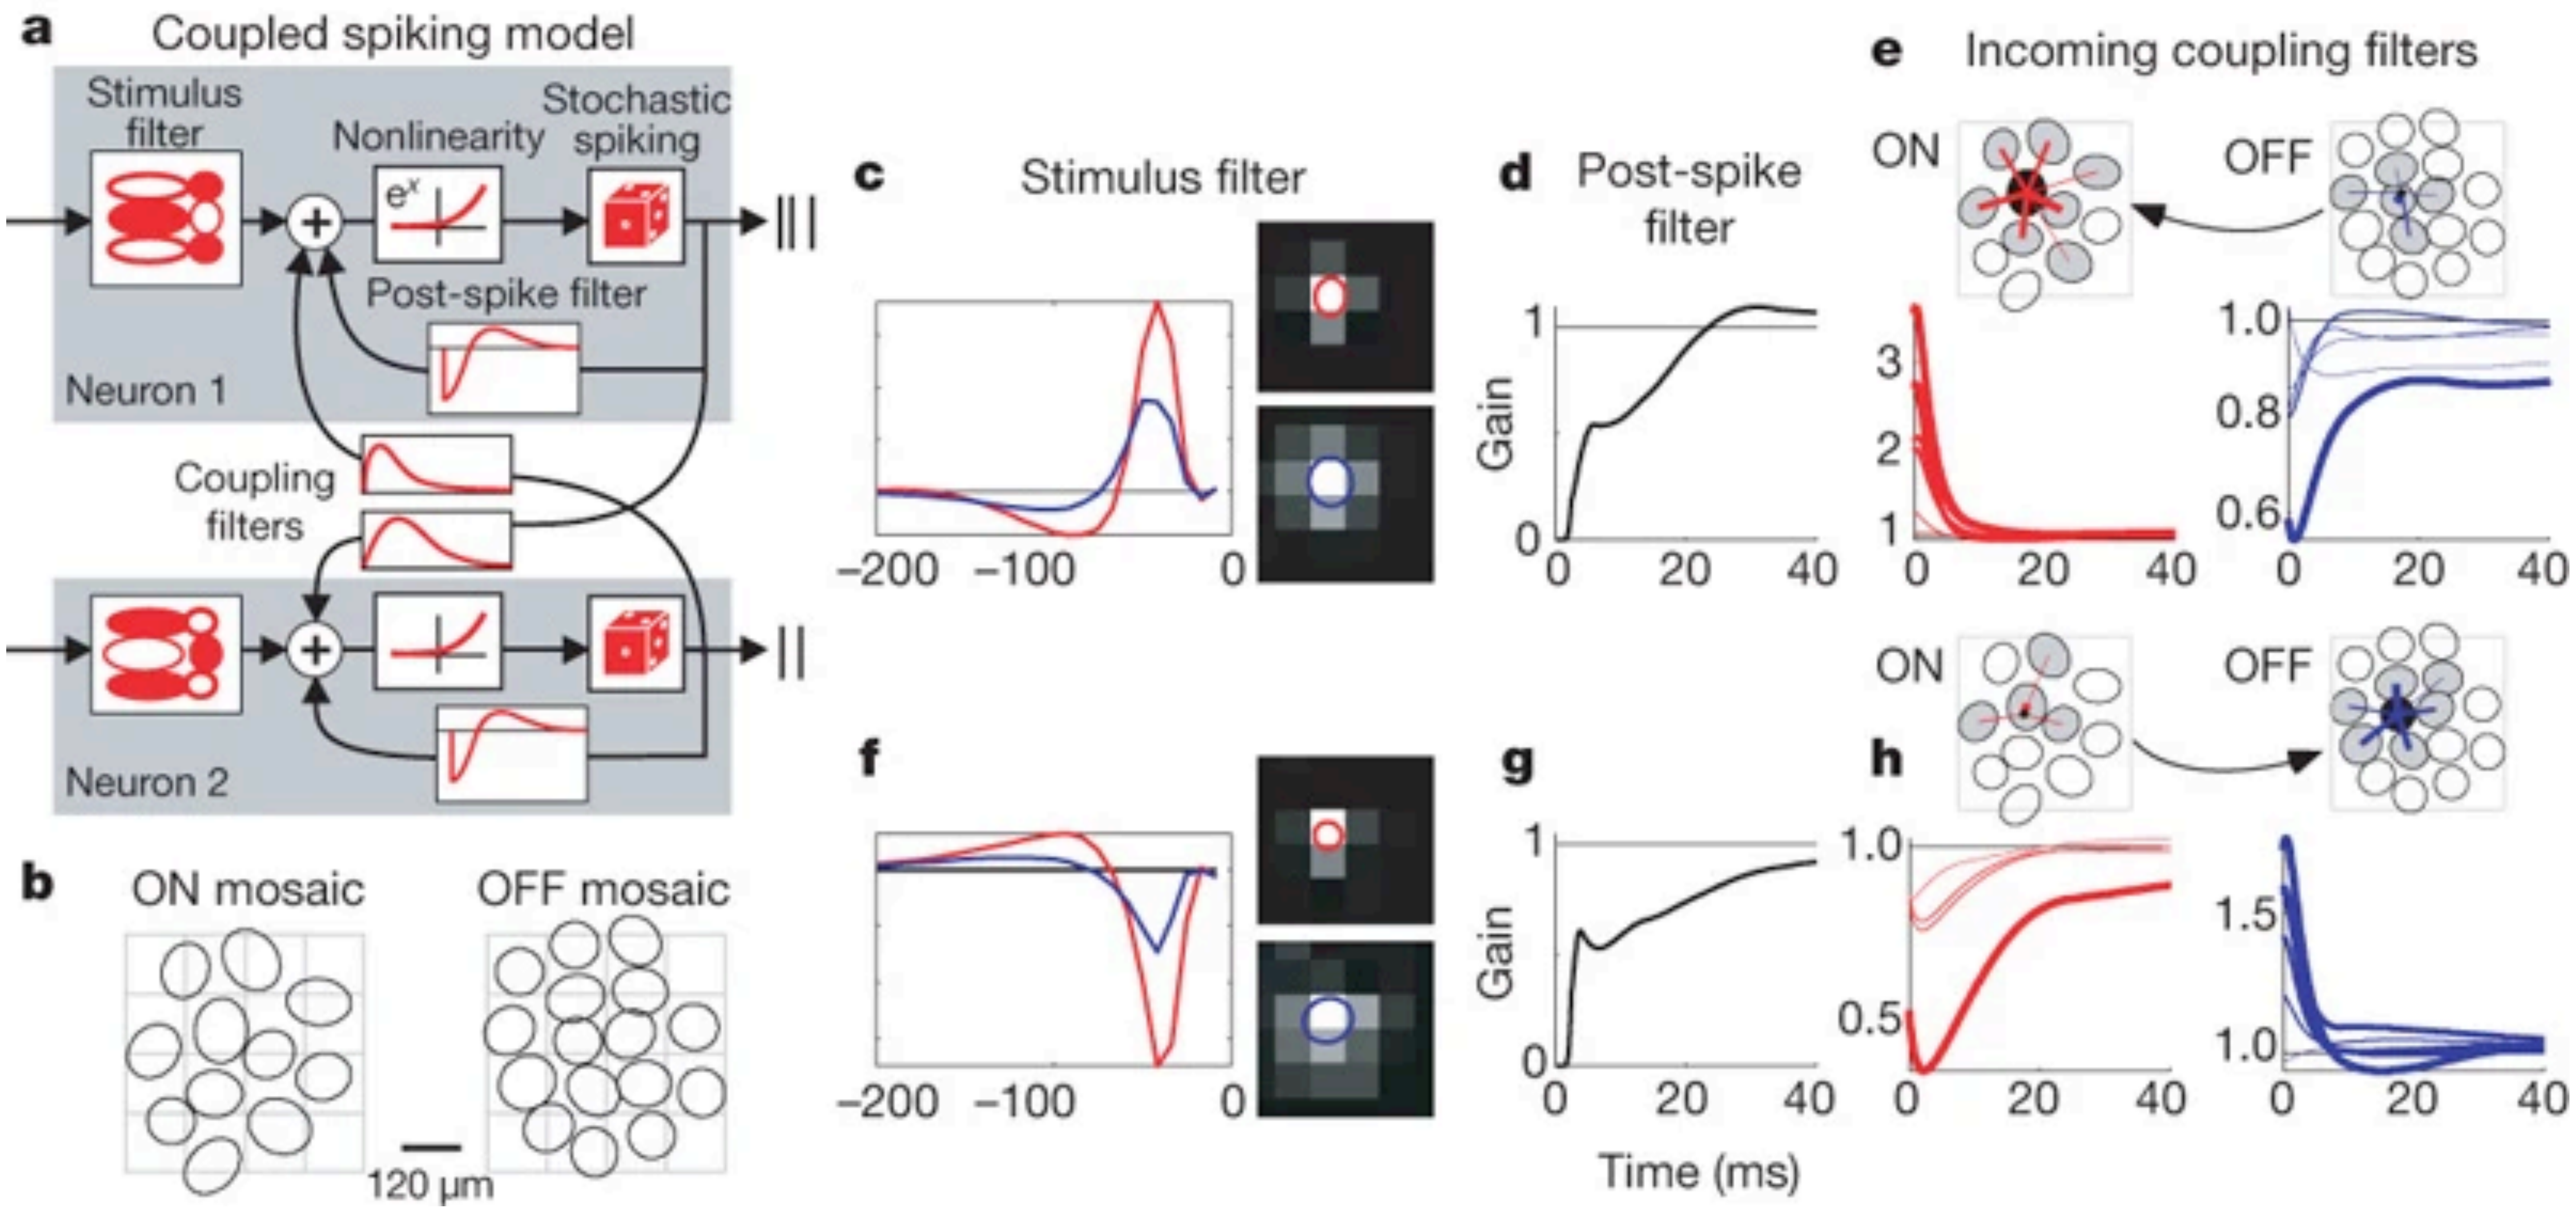

Encoding models of RGC responses

Separately modeling the coupling sparsity and weights

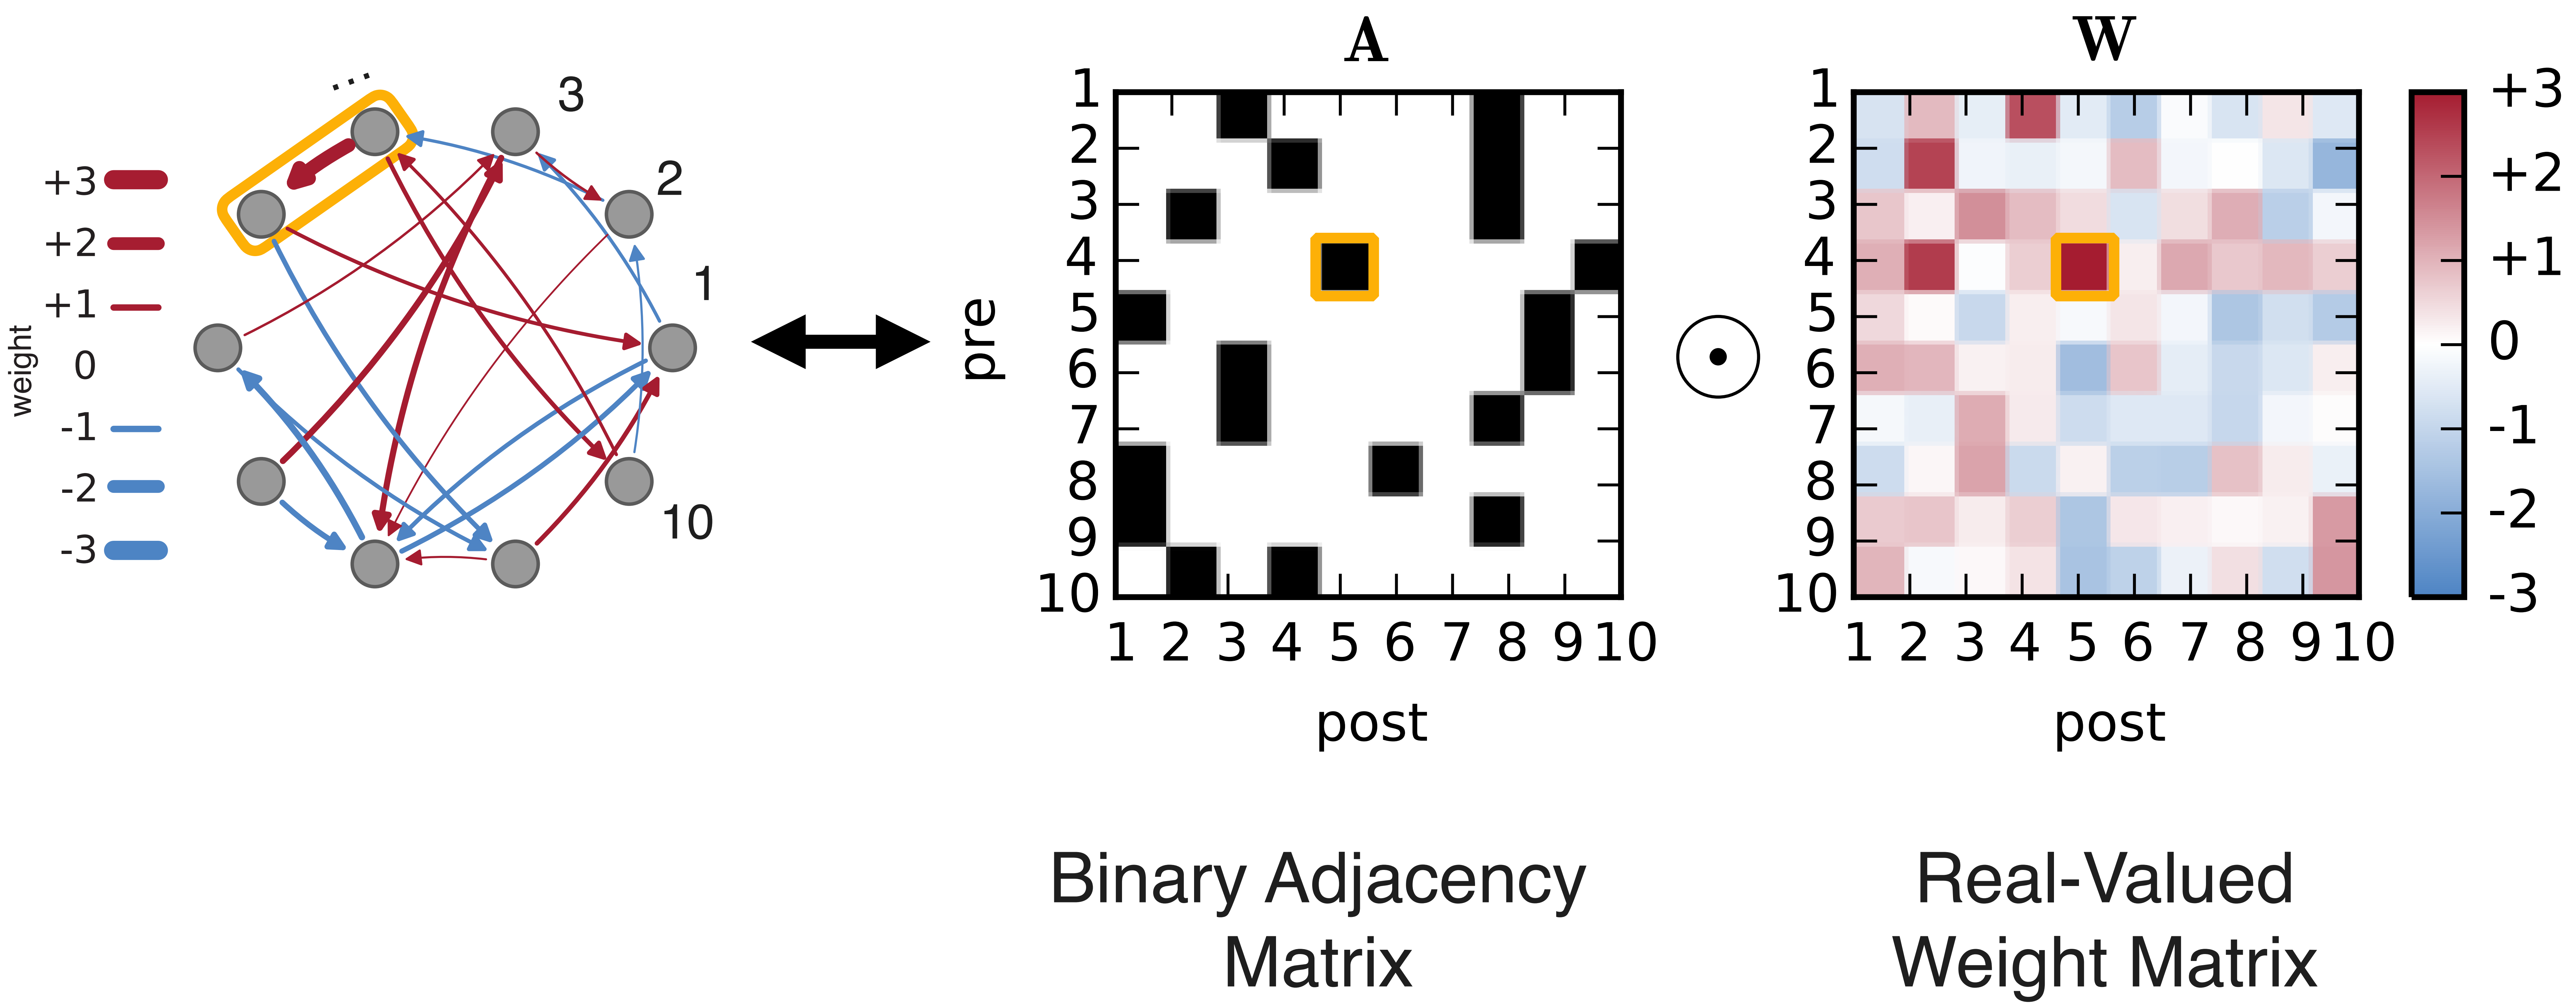

Encoding models of RGC responses

Latent variable models for networks

	Independent Edge Model	Stochastic Block Model (SBM)	Latent Distance Model
Types or Features Z	None	Type { ■ , ● }	Location $\mathbb{R}^D$
Adjacency Models A	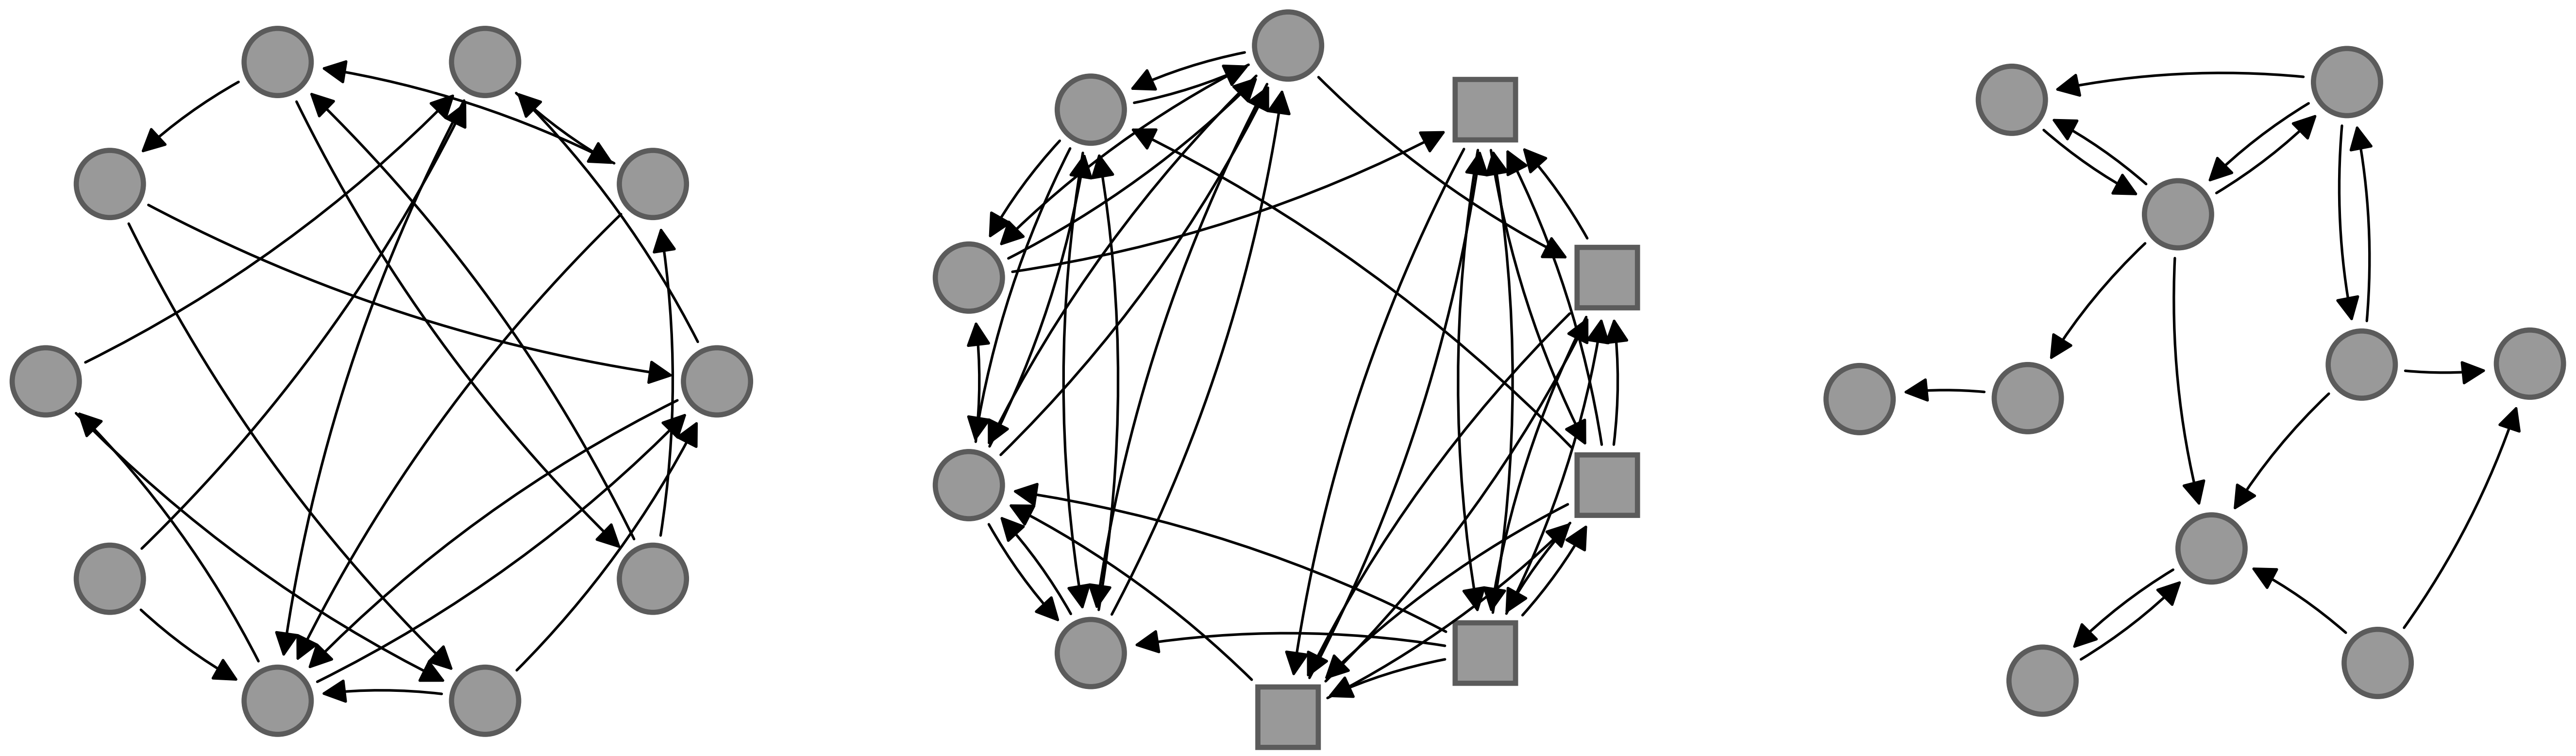		
Weight Models W	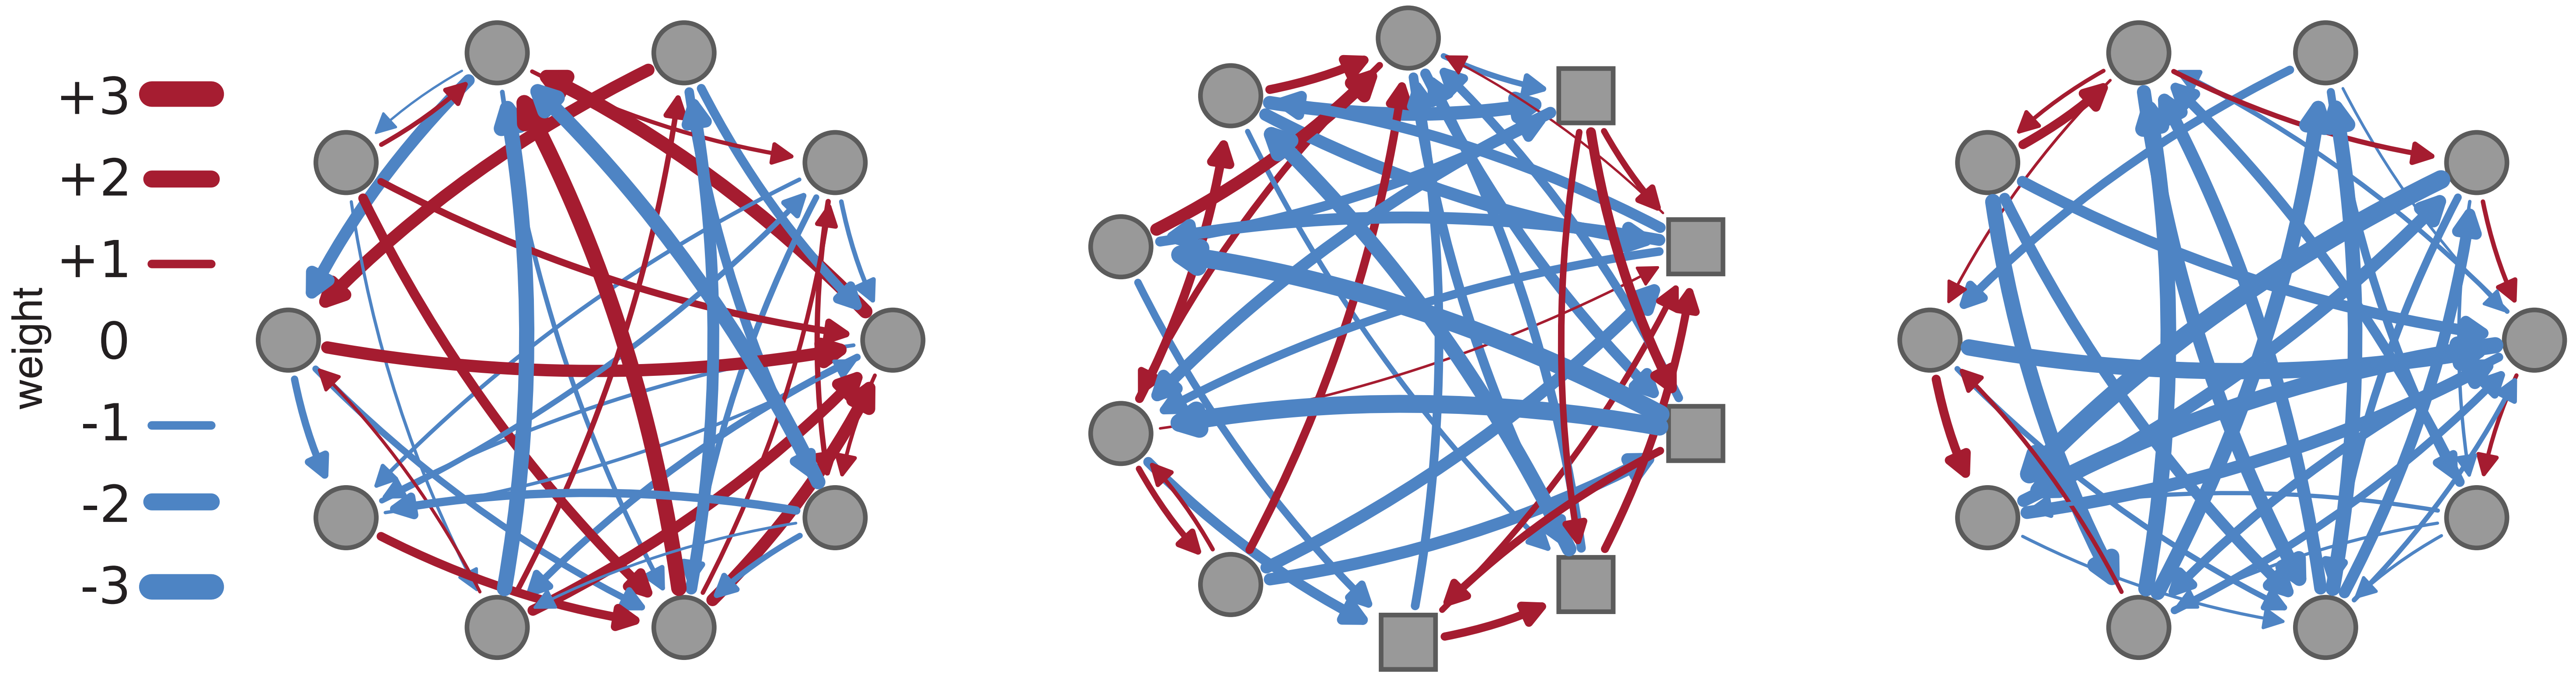		

Hoff (*NeurIPS*, 2007)

Example: a synthetic retina

Latent variables: types & locations. **Adjacency:** distance-dependent. **Weights:** type-dependent.

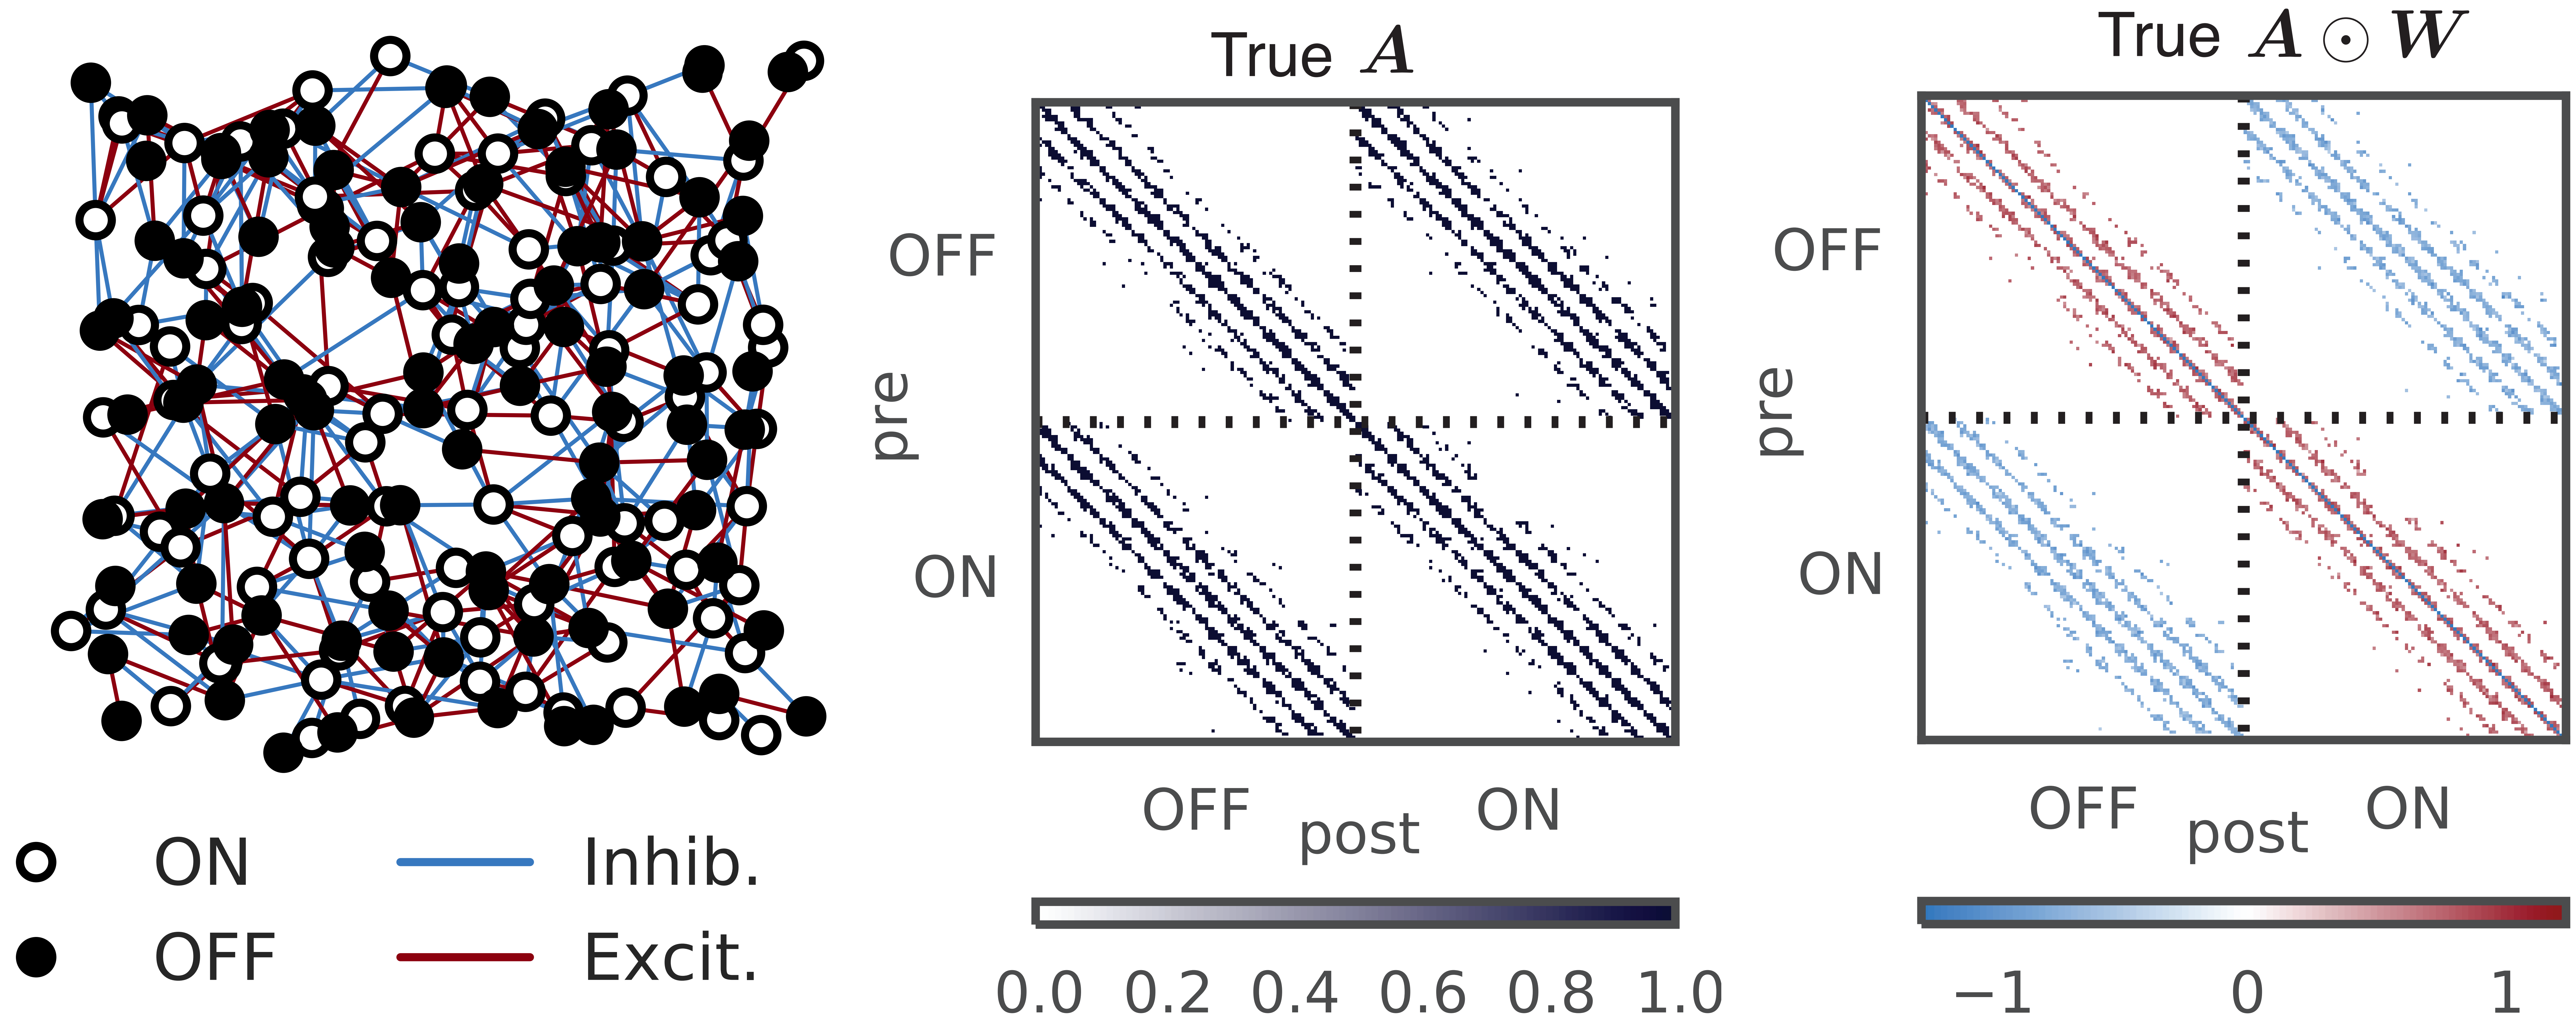

Linderman, Adams, Pillow (*NeurIPS*, 2016

Application to real primate retina data

Inferring locations

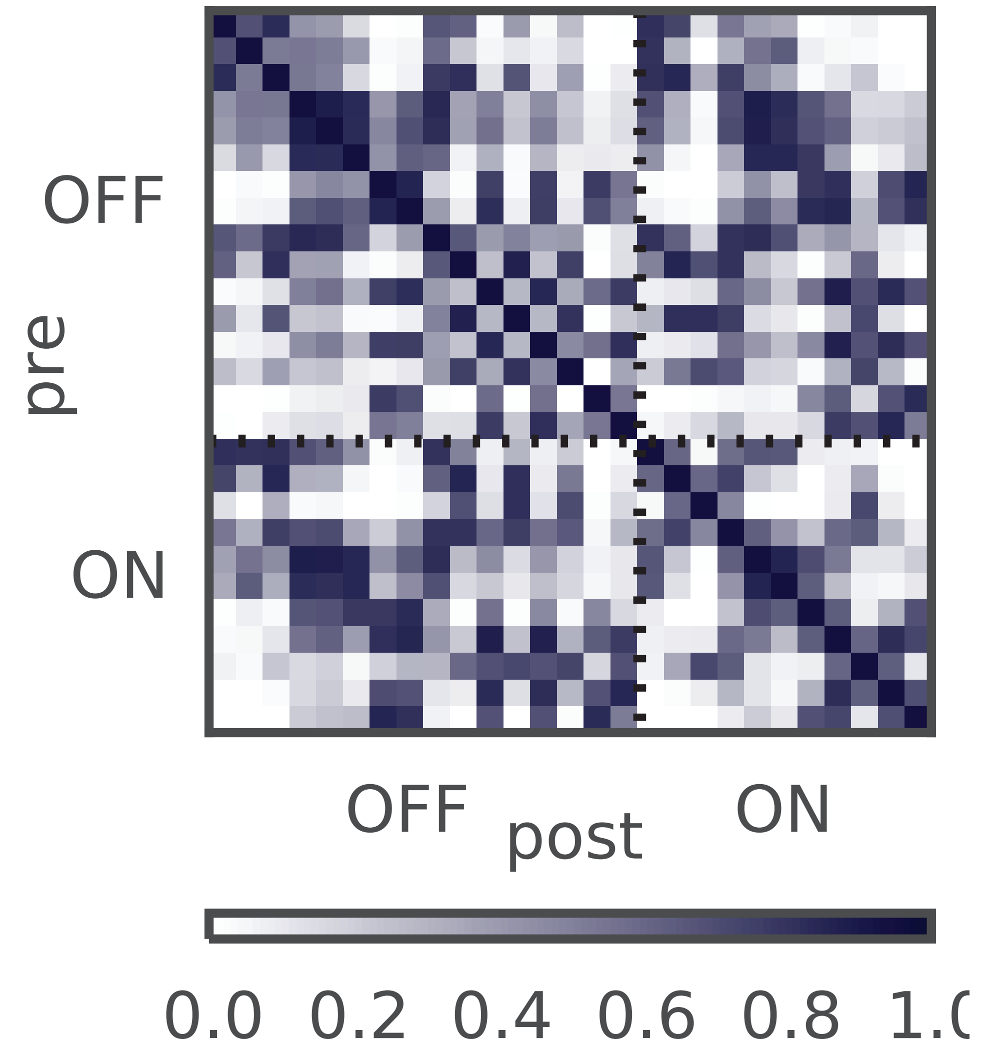

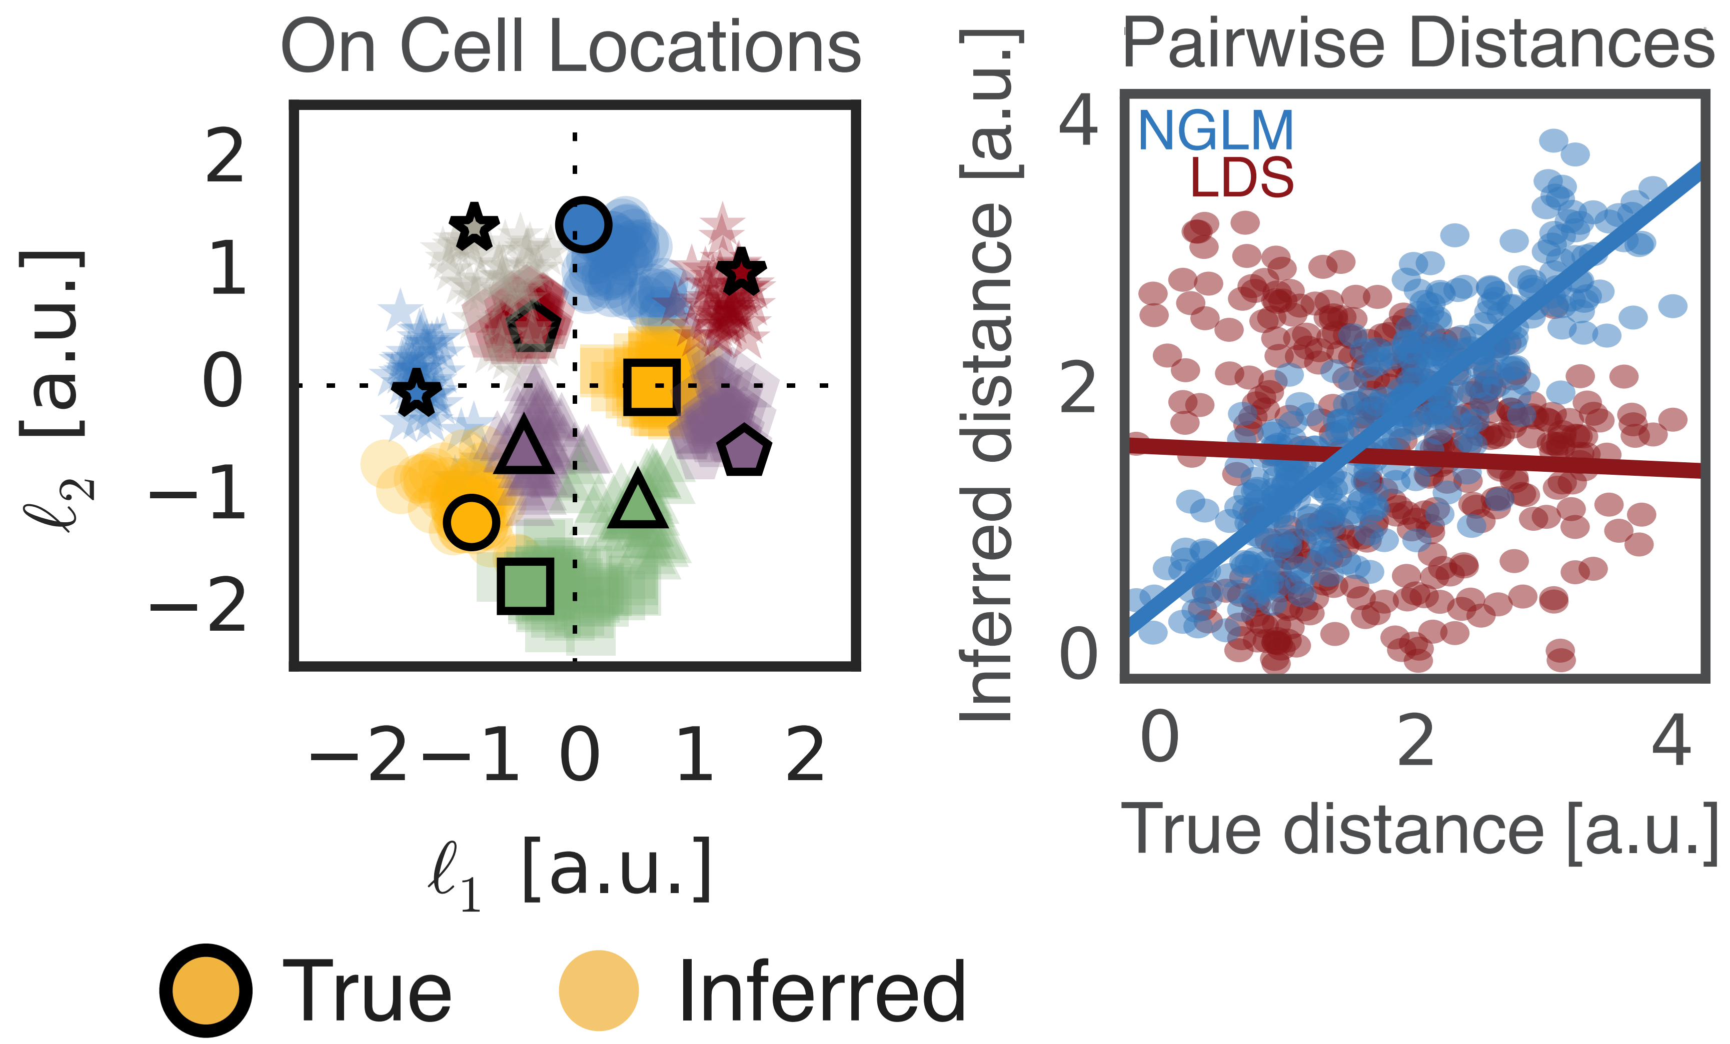

Inferring cell types

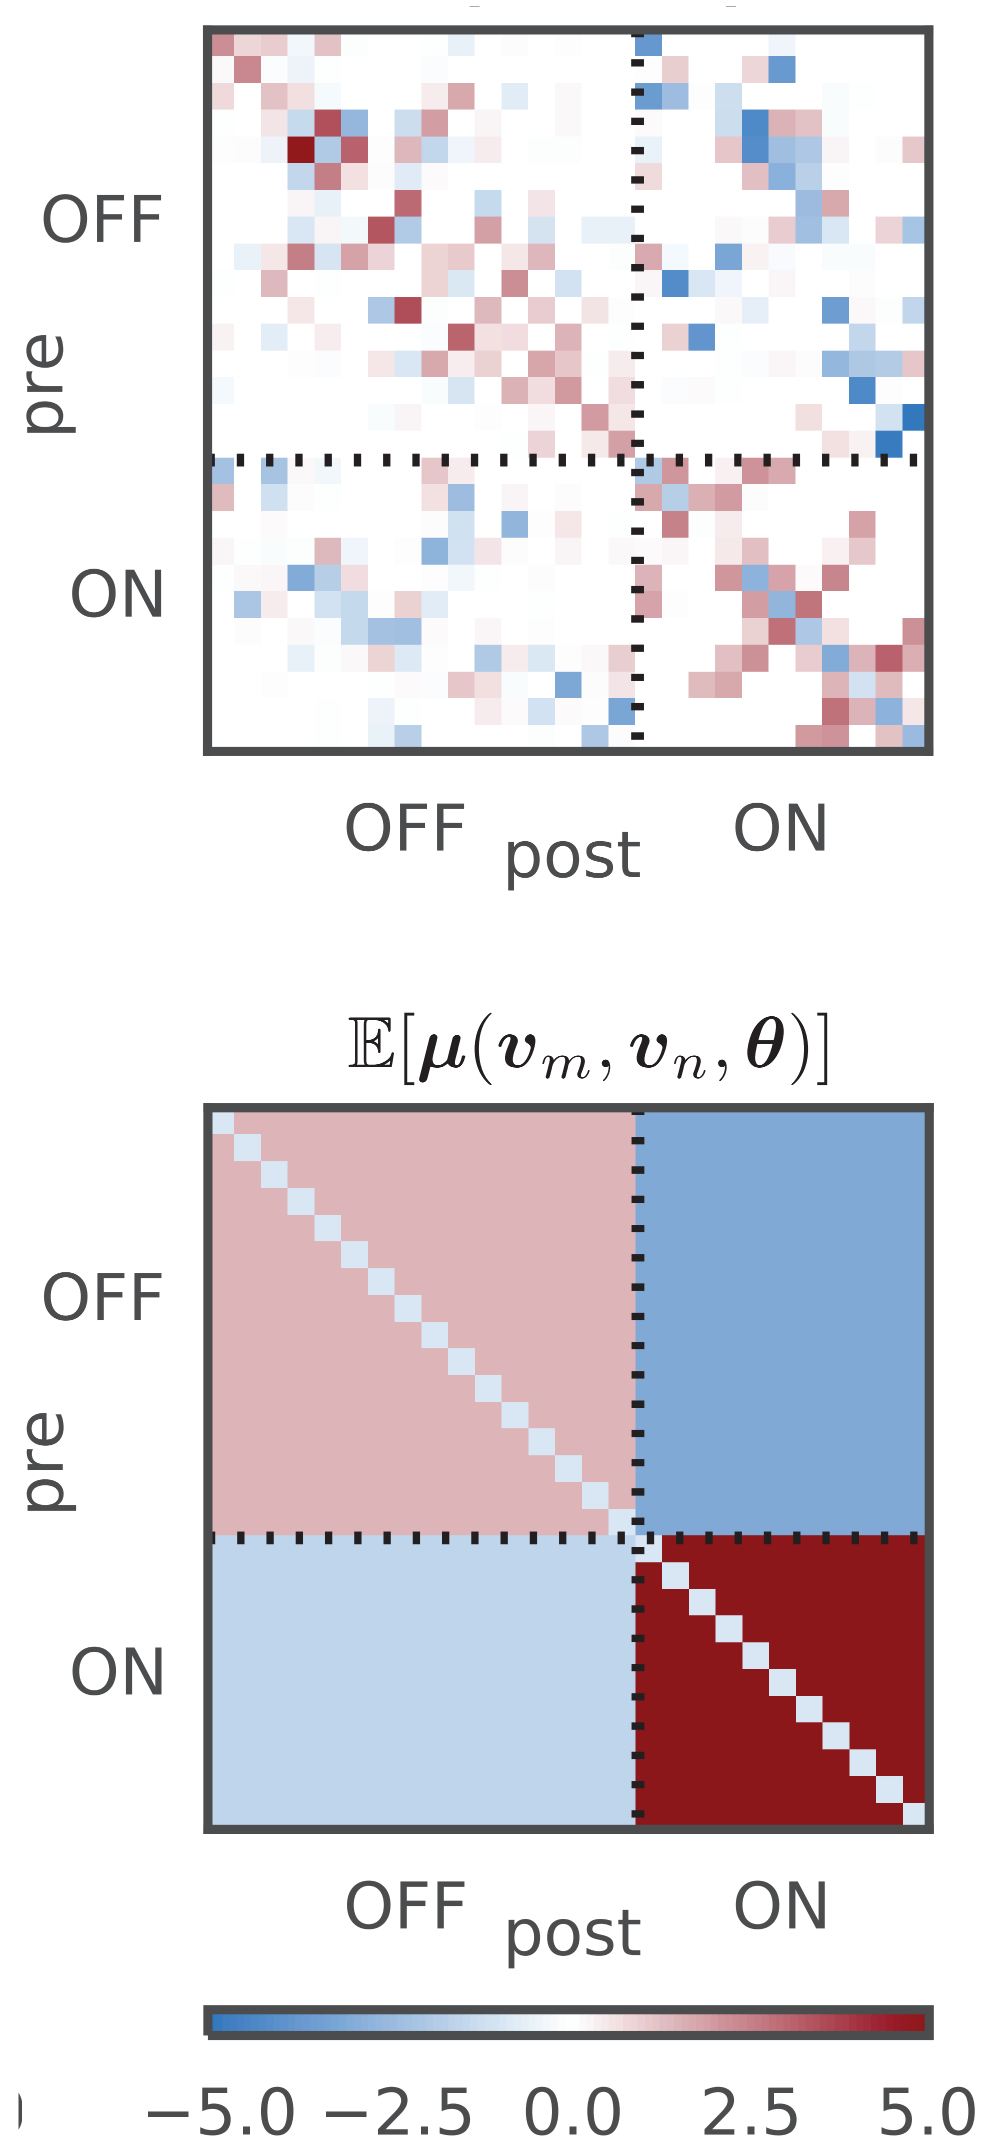

Model comparison

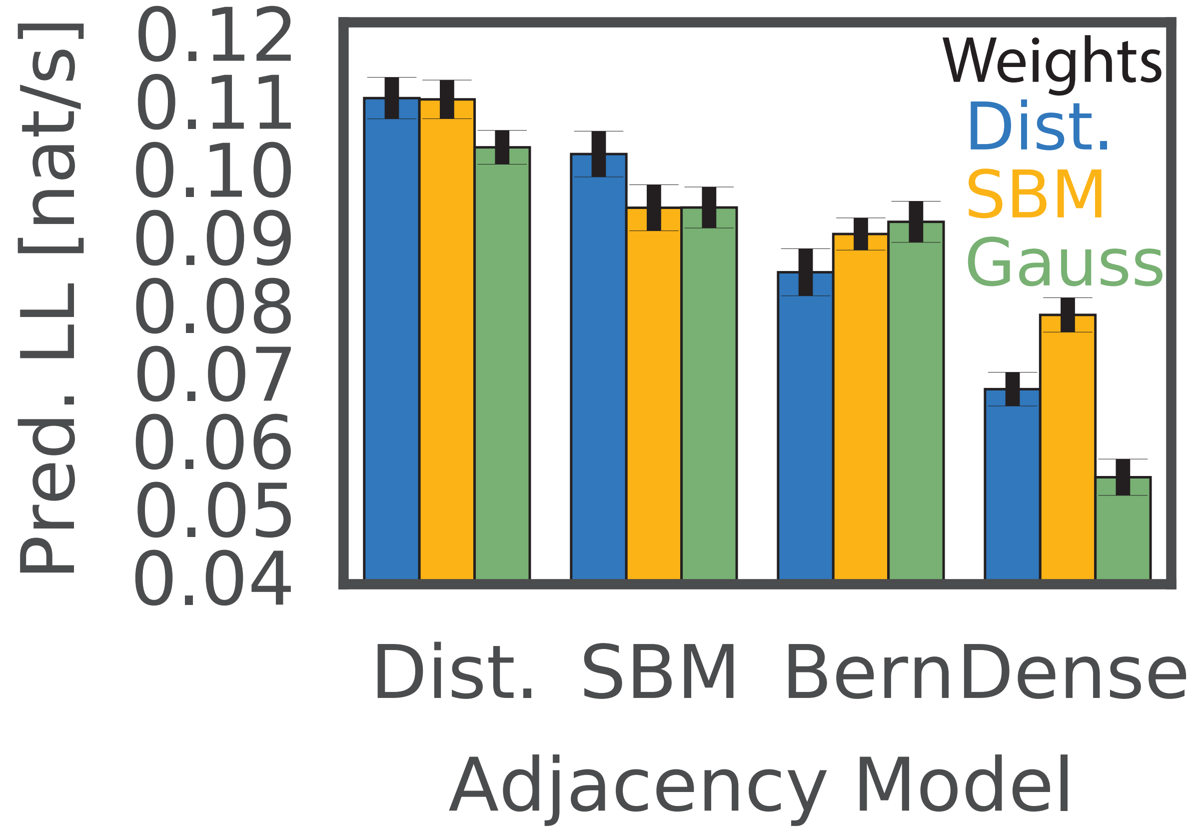

Linderman, Adams, Pillow (*NeurIPS*, 2016

Encoding models of RGC responses

Going deeper

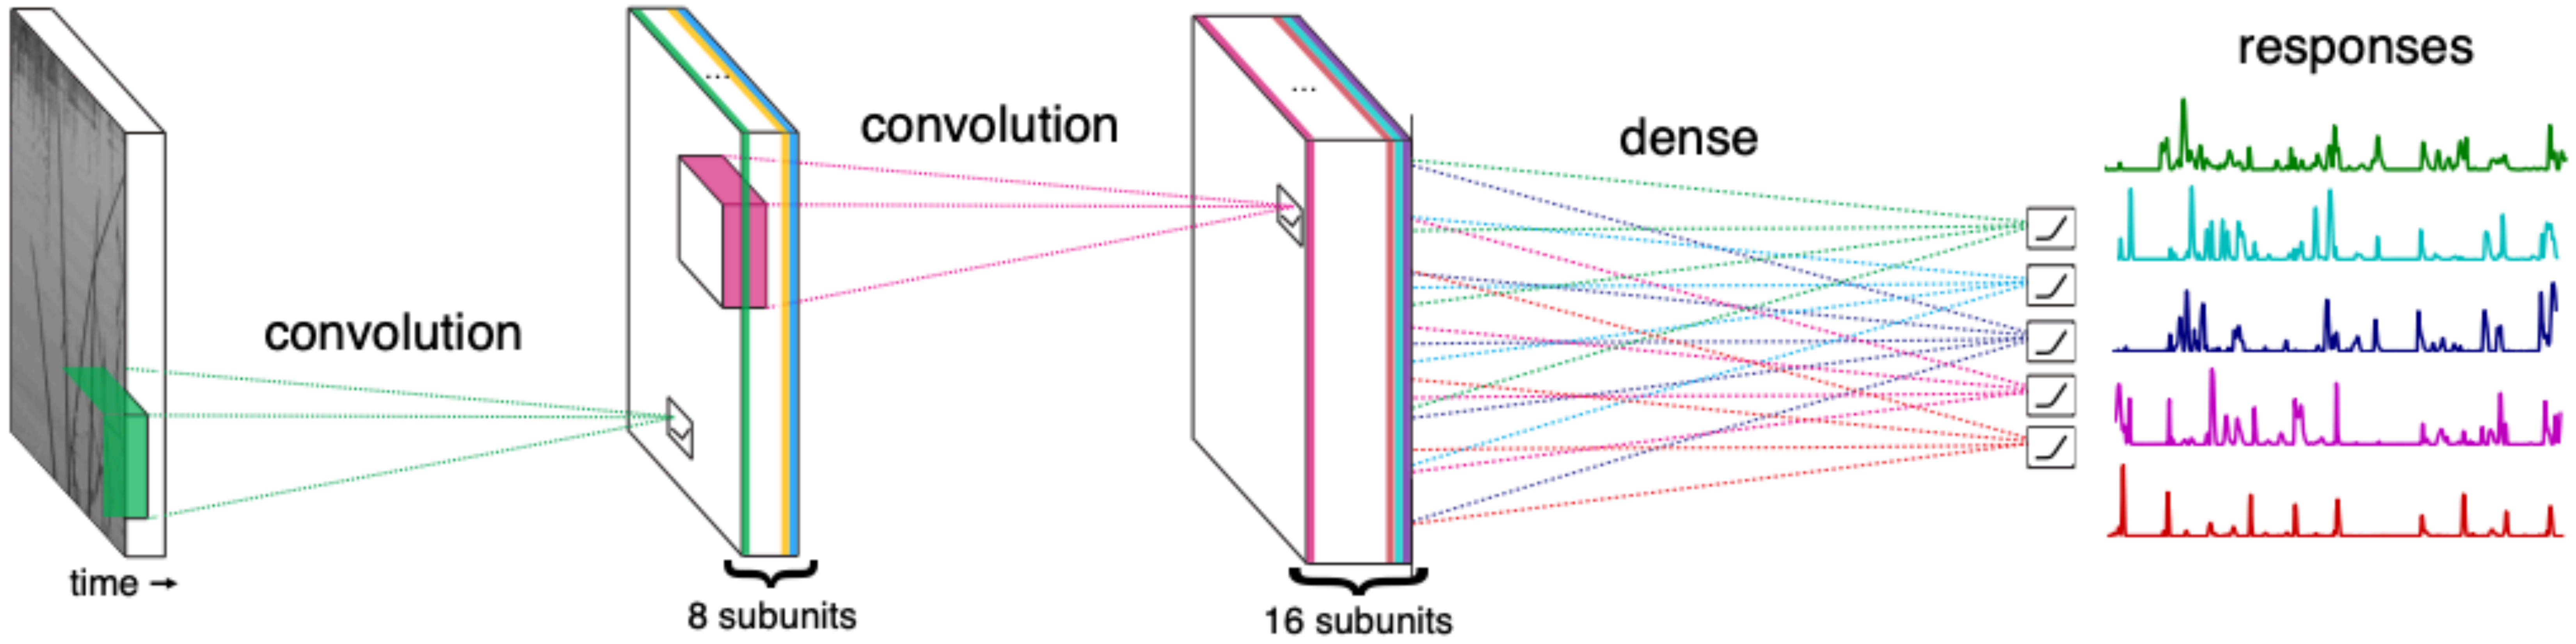

Figure 1: A schematic of the model architecture. The stimulus was convolved with 8 learned spatiotemporal filters whose activations were rectified. The second convolutional layer then projected the activity of these subunits through spatial filters onto 16 subunit types, whose activity was linearly combined and passed through a final soft rectifying nonlinearity to yield the predicted response.

McIntosh et al (*NeurIPS*, 2016)

Conclusion

- **Encoding models** predict the conditional distribution of neural responses to sensory stimuli.
  - Note, however, that we could have done the same thing by conditioning on (lagged) motor outputs instead.
- **Generalized linear models** are an effective means of modeling these conditional distributions.
- **Deep neural networks**, e.g. CNNs, essentially add multiple nonlinear layers to obtain features for estimating the conditional mean, rather than assuming its linear in the stimulus.

Further reading

- Dayan, Peter, and Laurence F. Abbott. Theoretical neuroscience: computational and mathematical modeling of neural systems. Computational Neuroscience Series, 2001.
- Pillow, Jonathan W., Jonathon Shlens, Liam Paninski, Alexander Sher, Alan M. Litke, E. J. Chichilnisky, and Eero P. Simoncelli. 2008. “Spatio-Temporal Correlations and Visual Signalling in a Complete Neuronal Population.” Nature 454 (7207): 995–99.
- McIntosh, L. T., Maheswaranathan, N., Nayebi, A., Ganguli, S., & Baccus, S. A. (2016). Deep learning models of the retinal response to natural scenes. Advances in neural information processing systems, 29, 1369.
- Linderman, Scott W., Ryan P. Adams, and Jonathan W. Pillow. "Bayesian latent structure discovery from multi-neuron recordings." Proceedings of the 30th International Conference on Neural Information Processing Systems. 2016.
- Pillow, Jonathan. *Cosyne Tutorial,* 2018. http://pillowlab.princeton.edu/pubs/pillow_TutorialSlides_Cosyne2018.pdf