Neural Encoding

Matthias Hennig
based on material by Mark van Rossum

School of Informatics, University of Edinburgh

January 2019

From stimulus to behaviour

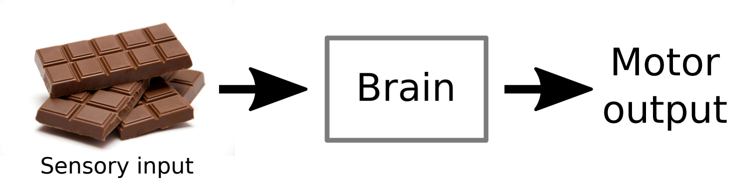

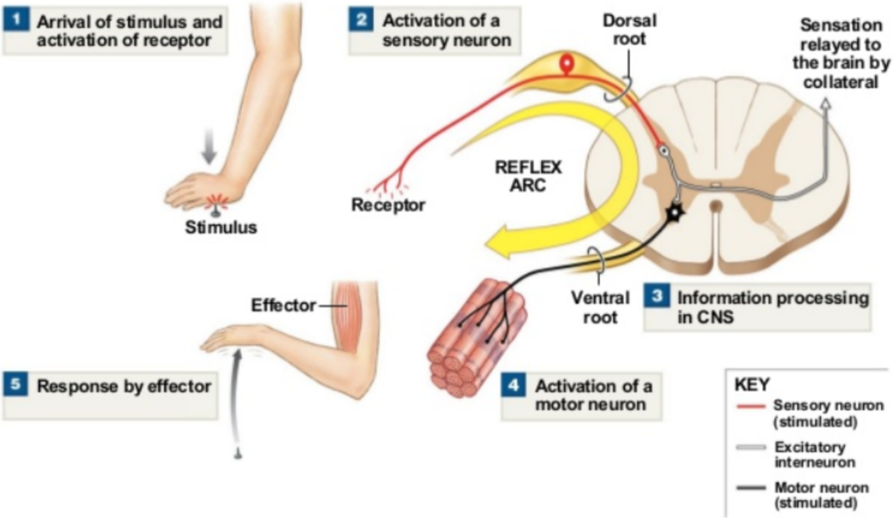
1 Arrival of stimulus and activation of receptor
Stimulus
2 Activation of a sensory neuron
Receptor
Dorsal root
Sensation relayed to the brain by collateral
REFLEX ARC
Ventral root
3 Information processing in CNS
Effector
5 Response by effector
4 Activation of a motor neuron
KEY
Sensory neuron (stimulated)
Excitatory interneuron
Motor neuron (stimulated)

The brain as a computer

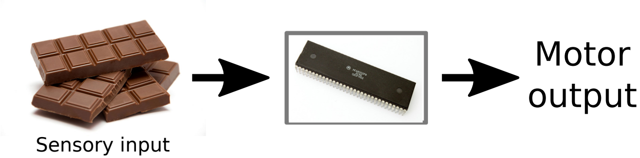

- Information processing to extract features and generate outputs
- Statistical inference

- Physical implementation irrelevant, possible to replicate *in silico*?

The neural code

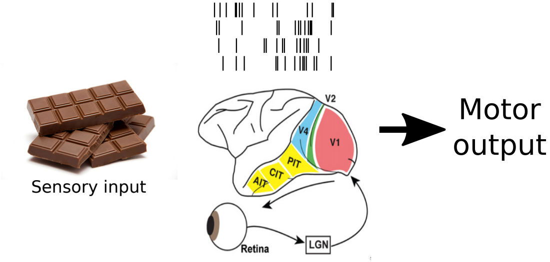

- Encoding: Prediction of neural response to a given stimulus: $P(R|S)$
- Decoding:
  - Given response, what was the stimulus: $P(S|R)$
  - Prosthetics: given firing pattern, what will be the motor output: $P(M|R)$

Understanding the *neural code* is like building a dictionary.

- Translate from outside world (sensory stimulus or motor action) to internal neural representation
- Translate from neural representation to outside world
- Like in real dictionaries, there are both one-to-many and many-to-one entries in the dictionary

Encoding: Stimulus-response relation

Predict response R to stimulus S. Black box approach.
This is a supervised learning problem, but:

- Stimulus S can be synaptic input or sensory stimulus.
- Responses are noisy and unreliable: Use probabilities.
- Typically many input (and sometimes output) dimensions
- Reponses are non-linear[1]
  - Assume non-linearity is weak. Make series expansion?
  - Or, impose a parametric non-linear model with few parameters
- Need to assume causality and stationarity (system remains the same). This excludes adaptation!

[1] Linear means: $r(\alpha s_1 + \beta s_2) = \alpha r(s_1) + \beta r(s_2)$ for all α, β.

Response: Spikes and rates

Response consists of spikes. Spikes are (largely) stochastic. Compute rates by trial-to-trial average and hope that system is stationary and noise is really noise.

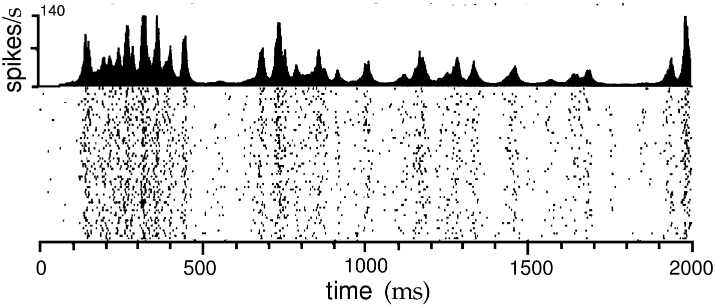

Often, we try to predict *R*, rather than predict the spikes.

Paradigm: Early Visual Pathways

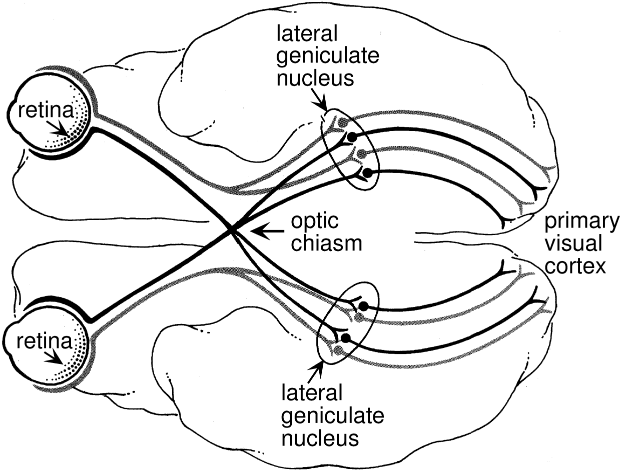

[Figure: Dayan and Abbott, 2001, after Nicholls et al, 1992]

Retinal/LGN cell response types

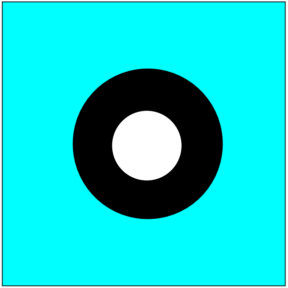

On-centre off-surround

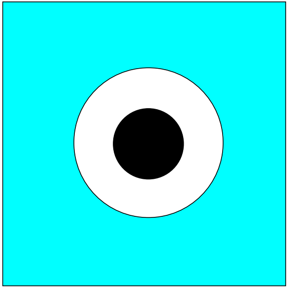

Off-centre on-surround

Mach bands

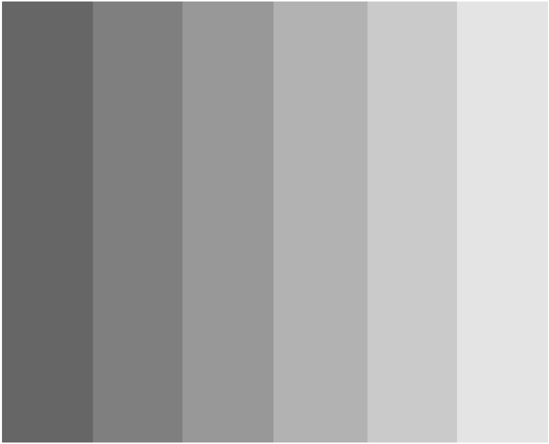

V1 cell response types (Hubel & Wiesel)

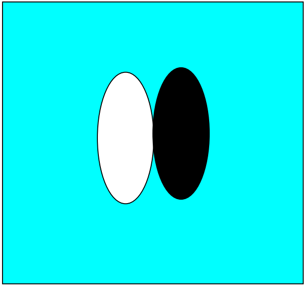

Odd

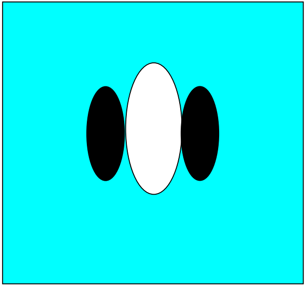

Even

- Simple cells, modelled by Gabor functions
- Also complex cells, and spatio-temporal receptive fields
- Higher areas
- Other pathways (e.g. auditory)

Not all cells are so simple...

Intermediate sensory areas (eg. IT) have face selective neurons. In the limbic system, neurons appear even more specialised [Quiroga et al., 2005].

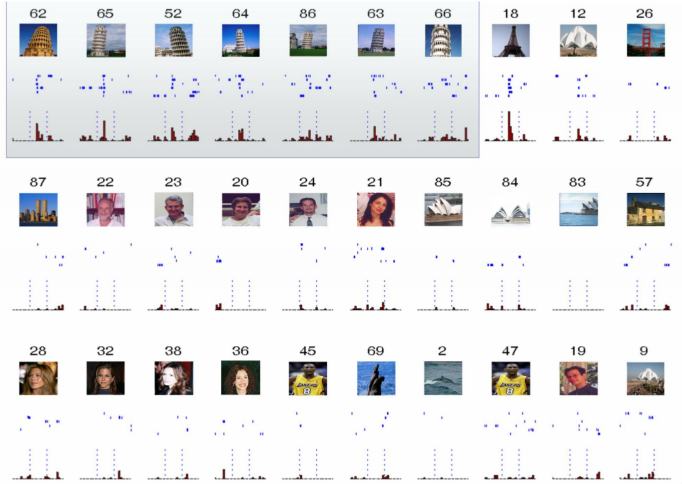

In higher areas the receptive field (RF) is not purely sensory. Example: pre-frontal cells that are task dependent [Wallis et al., 2001]

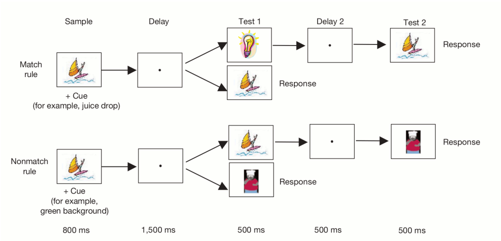

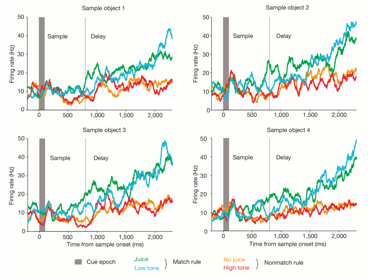

Model complexity

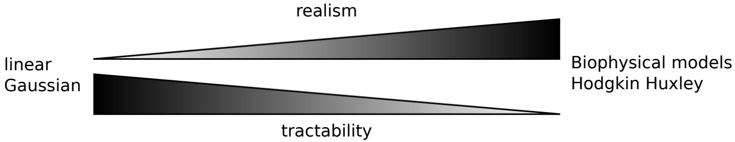

To study neural encoding, we need a model. There is an inevitable trade-off between realism and complexity.
Simple models: normative theories
Detailed models: how implemented in the brain

From stimulus to response

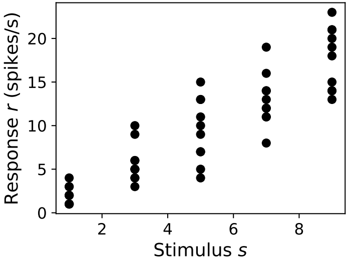

What is the correct $P(R|S,\theta)$, where θ is a model parameter?
Strategy: Maximise the likelihood $P(R|S,\theta)$

General linear model (GLM)

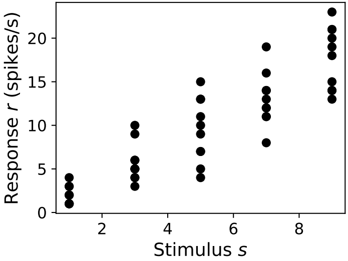

We assume a Poisson model. For N trials, we write the likelihood

$$
\begin{aligned}
P(R|S,\theta) &= \prod_{i=1}^{N} P(r_i|s_i,\theta) \\
&= \prod_{i=1}^{N} \frac{1}{r_i!}(\theta s_i)^{r_i} e^{-\theta s_i}
\end{aligned}
$$

Model likelihood

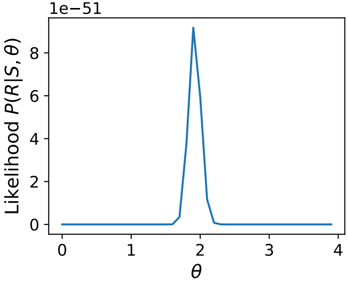

$$P(R|S,\theta) = \prod_{i=1}^{N} P(r_i|s_i,\theta)$$

$$= \prod_{i=1}^{N} \frac{1}{r_i!}(\theta s_i)^{r_i} e^{-\theta s_i}$$

has a maximum close to 2.

log-likelihood

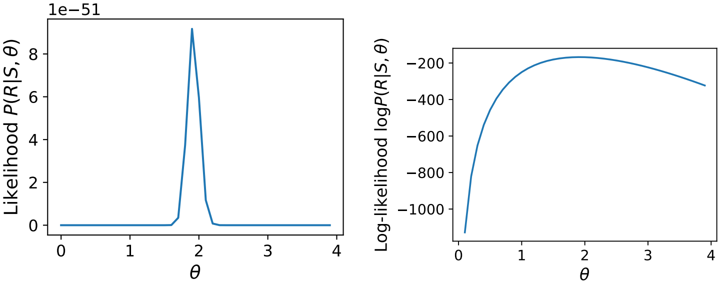

In practice, we use the logarithm

$$\log P(R|S,\theta) = \log \prod_{i=1}^{N} P(r_i|s_i,\theta)$$
$$= \sum_{i}^{N} r_i \log\theta - \theta s_i + C$$

Terms in C does not depend on θ, so can be ignored.

log-likelihood

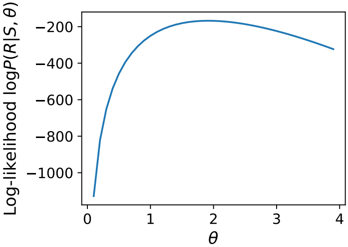

To find the maximum, differentiate:

$$\frac{\partial \log P(R|S,\theta)}{\partial \theta}$$

log-likelihood

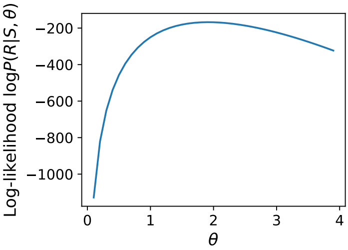

Find the maximum:

$$\log P(R|S,\theta) = \sum_i^N r_i \log\theta - \theta s_i + C$$

$$\frac{\partial \log P(R|S,\theta)}{\partial\theta} = \sum_i \frac{r_i}{\theta} - \sum_i s_i$$

log-likelihood

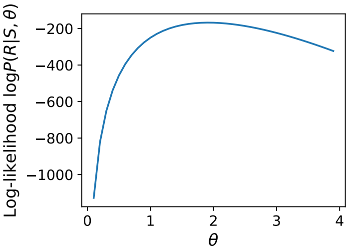

Find the maximum:

$$\frac{\partial \log P(R|S,\theta)}{\partial\theta} = \sum_i \frac{r_i}{\theta} - \sum_i s_i$$

$$\hat{\theta} = \frac{\sum r_i}{\sum s_i}$$

In this example I obtain $\hat{\theta} = 1.92$, close to the true value $\theta = 2$.

Remarks

- The predicted rate can be <0.
- In biology, unlike physics, there is no obvious small parameter that justifies neglecting higher orders. Rectification requires infinite orders, for instance. Check the accuracy of the approximation post hoc.

Averaging and ergodicity

- $\langle r \rangle$ formally means an average over many realizations over the random variables of the system (both stimuli and internal state). This definition is good to remember when conceptual problems occur.
- An ergodic system visits all realizations if one waits long enough. That means one can measure from a system long enough, true averages can be obtained.
- This however requires stationarity, internal states are not allowed to change.

A more realistic response

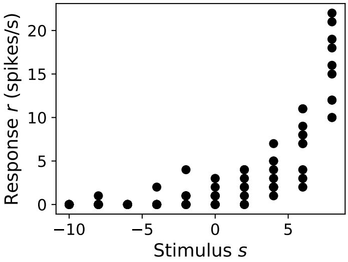

A more realistic response

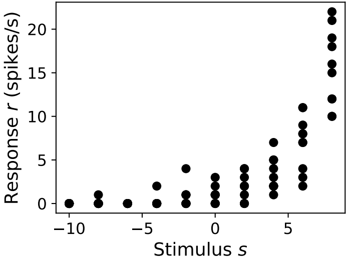

This requires a non-linear transformation $r(s) \sim Poisson(f(\theta s))$.

Neural responses depend on the stimulus history

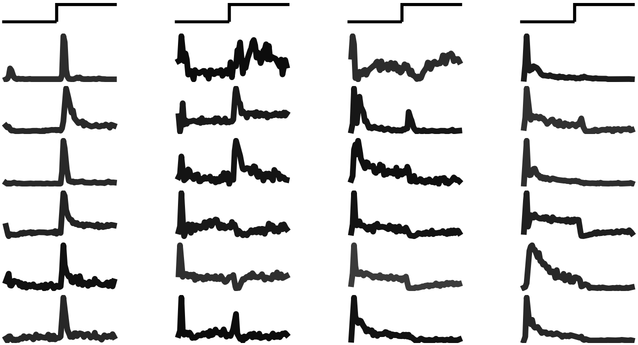

Introducing a linear temporal kernel $k(t)$ with
$r(t) = Poisson(f(\int dt' s(t')k(t-t'))\ )$.

Poisson Generalised Linear Model (also GLM!)

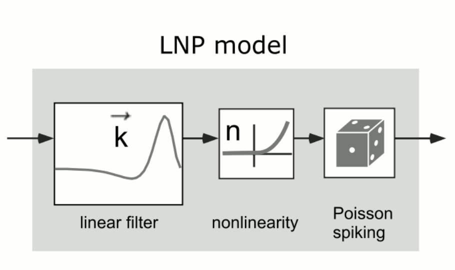

[Pillow et al., 2005]

$r(t) = Poisson(f(\int dt' s(t') k(t - t')\ )$

- Linear: spatial and temporal filter kernel **k**
- Non-linear function giving output spike probability: rectification, saturation
- Poisson spikes $p_{spike}(t) = \lambda(t)$ (noisy)

Fitting a linear model

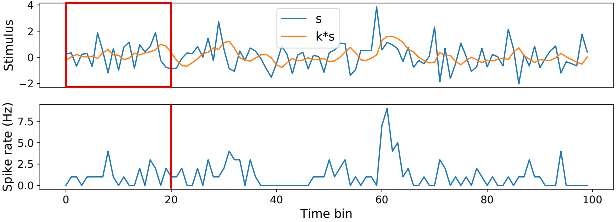

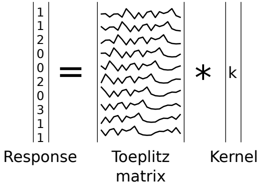

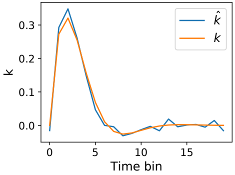

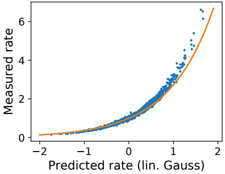

$r(t) = Gaussian(\int dt' s(t')k(t - t'))$

This has closed form MLE: $\hat{k} = (S^T S)^{-1} S^T R$

Data comes from model with exponential nonlinearity. The model recovers the kernel well, but cannot predict the rates.

Spike triggered average (STA)

Spike times t_i, $r(t) = \sum \delta(t - t_i)$
$g_1(\tau) = \frac{1}{\sigma^2} \langle r(t) s(t - \tau) \rangle = \frac{1}{\sigma^2} \sum_{t_i} s(t_i - \tau)$

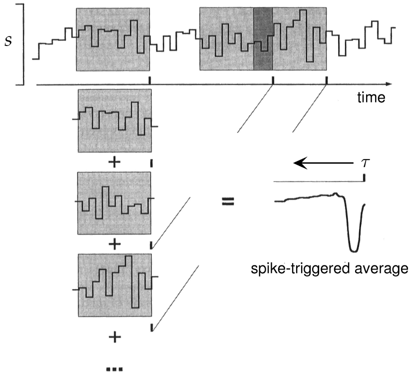

Linear models for spiking neurons

Application on H1 neuron [Rieke et al., 1996]. Prediction (solid), and actual firing rate (dashed). Prediction captures the slow modulations, but not faster structure. This is often the case.

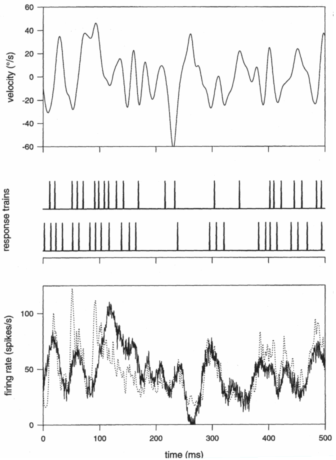

Fitting a non-linear model

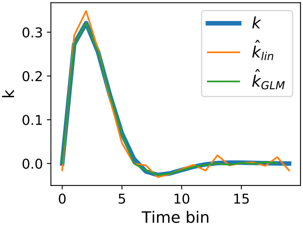

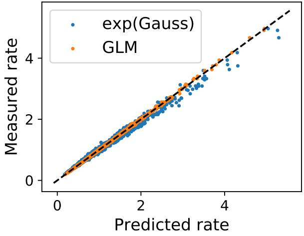

Poisson GLM log-likelihood has no closed form MLE:

$$\log P(R|S,\theta) = \sum_i r_i \log f(k * s_i) - \sum_i f(k * s_i)$$

Use numerical minimisation of the neg. log-likelihood (scipy.optimize.fmin or fminsearch in Matlab) This recovers the kernel and rates correctly.

Fitting non-linear models

Poisson GLM log-likelihood:

$$\log P(R|S,\theta) = \sum_i r_i \log f(k * s_i) - \sum_i f(k * s_i)$$

Bernoulli GLM log-likelihood:

$$\log P(R|S,\theta) = \sum_i r_i \log f(k * s_i) + \sum_i (1 - r_i) \log(1 - f(k * s_i))$$

For $f(x) = 1/(1 + \exp(-x))$, this is logistic regression.
When f is convex (*log*(f) is concave) in parameters, e.g. $f(x) = [x]_+$, or $f(x) = \exp(x)$, then $\log \mathcal{L}$ is concave, hence a global maximum exists.

Regularization

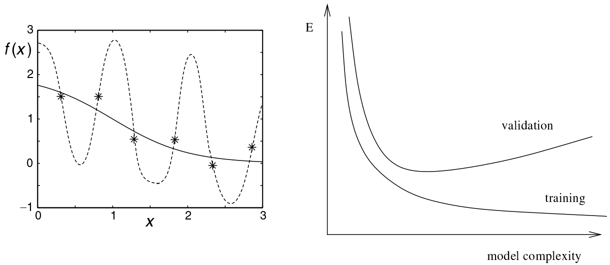

Figure: Over-fitting: Left: The stars are the data points. Although the dashed line might fit the data better, it is over-fitted. It is likely to perform worse on new data. Instead the solid line appears a more reasonable model. Right: When you over-fit, the error on the training data decreases, but the error on new data increases. Ideally *both* errors are minimal.

Regularization

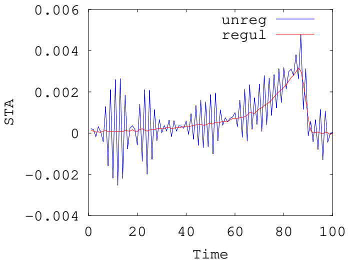

- Fits with many parameters/short data typically require regularization to prevent over-fitting
- Regularization: punish fluctuations (smooth prior, ridge regression)
- $\hat{k} = (S^T S + \lambda I)^{-1} S^T \mathbf{r}$
- Regulariser λ has to be set by hand

Poisson GLM results

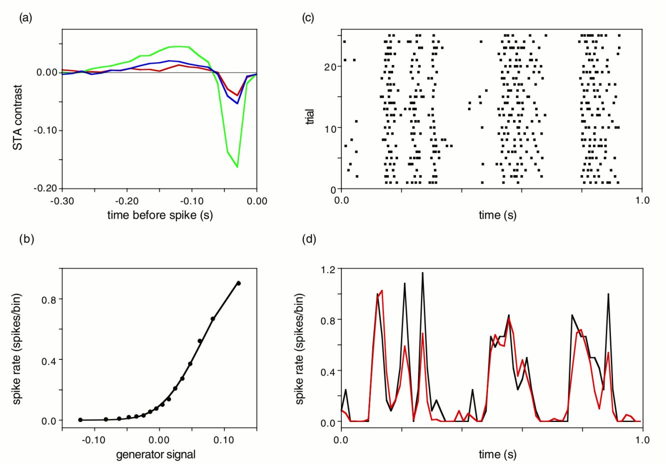

[Chichilnisky, 2001]

Colors are the kernels for the different RGB channels

Spatio-temporal kernels

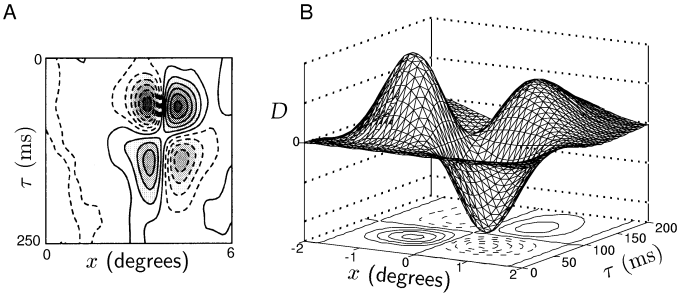

[Dayan and Abbott, 2002]

- Kernel can also be in spatio-temporal domain.
- This V1 kernel does not respond to static stimulus, but will respond to a moving grating ([Dayan and Abbott, 2002]§2.4 for more motion detectors)

Integrate and fire model

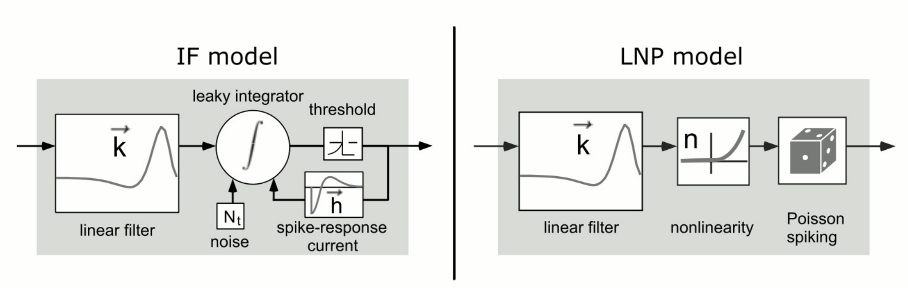

[Pillow et al., 2005]

- Parameters are the **k** and **h** kernels
- **h** can include reset and refractoriness
- For standard I&F: $h(t) = \frac{1}{R}(V_T - V_{reset})\delta(t)$

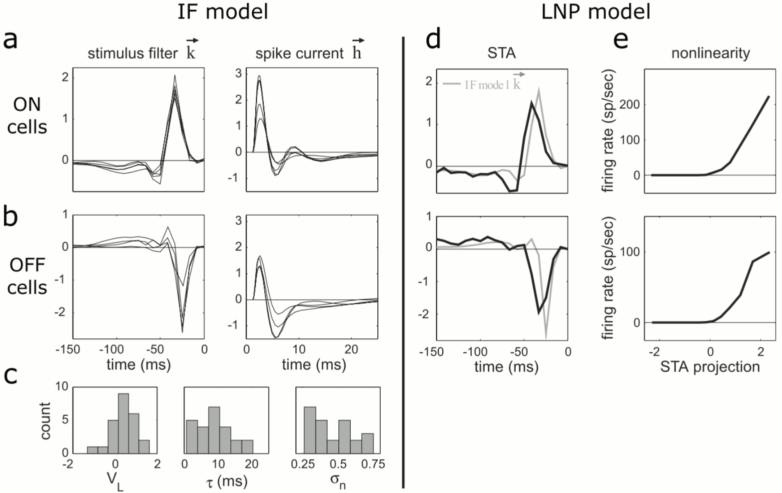

Figure 2. Parameters obtained from fits to RGC data for the IF model (***a*–*c***) and the LNP model (***d***, ***e***). ***a***, Filters $\vec{k}$ and spike–response currents $\vec{h}$ obtained for five ON cells in one retina. ***b***, Corresponding filters for four OFF cells. ***c***, Histograms of model scalar parameters and for all 24 cells in three retinas. ***d***, Comparison of linear filters for the IF model (gray) and LNP model (black) for one ON cell (top) and one OFF cell (bottom). ***e***, Measured LNP point nonlinearities for converting filter output to instantaneous spike rate.

[Pillow et al., 2005] Fig 2

[Pillow et al., 2005] Fig 3

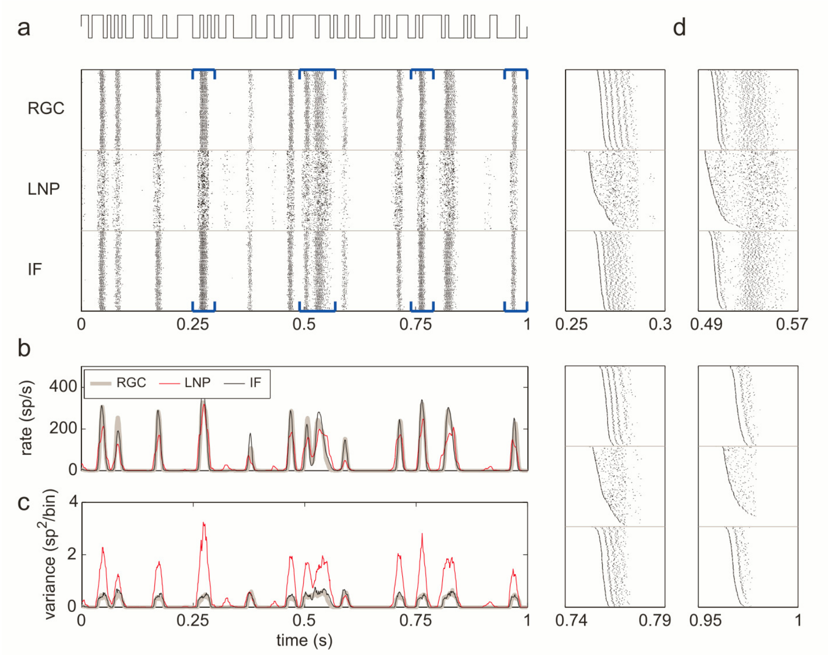

Poisson GLM with spike feedback

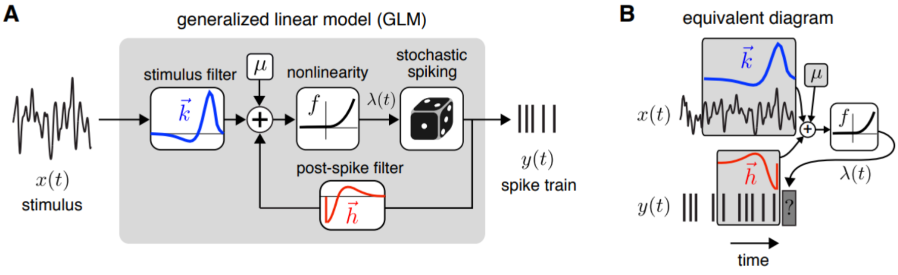

[Weber and Pillow, 2017]

Spike feedback allows modelling neuron types

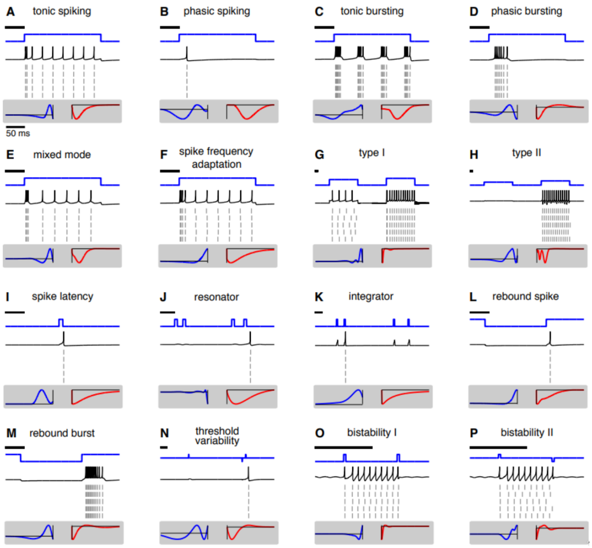

[Weber and Pillow, 2017]

Even more complicated models

A retina + ganglion cell model with multiple adaptation stages [van Hateren et al., 2002]

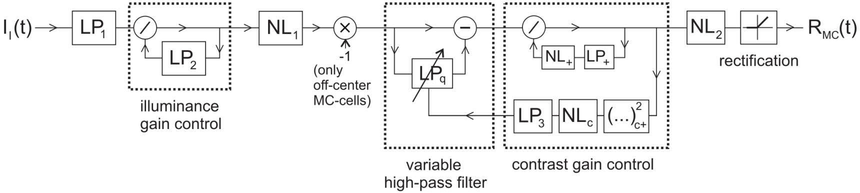

But how to fit the parameters?

Network models

Generalization to networks.

- Unlikely to have data from all neurons
- Predict of cross-neuron spike patterns and correlations
- Correlations are important for decoding (coming lectures)
- Estimate 'functional coupling', $O(N \times N)$ parameters
- Uses small set of basis functions for kernels

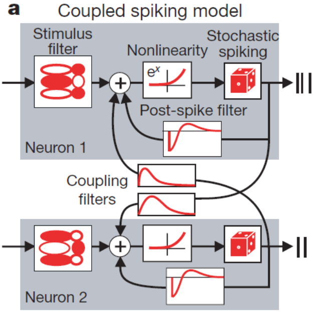

[Pillow et al., 2008]

Network models

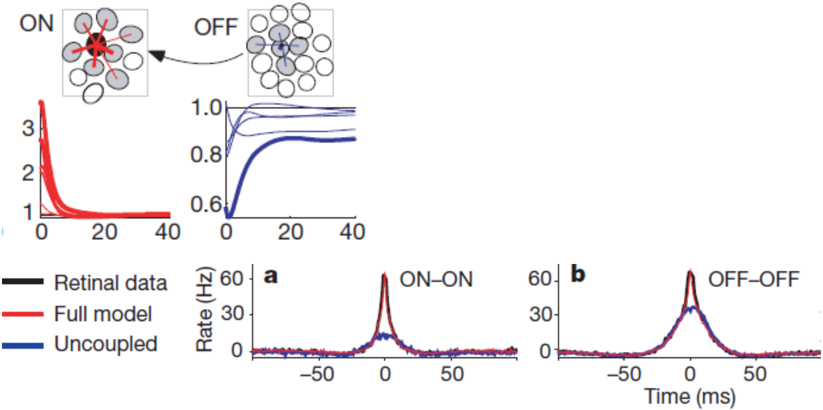

Note uncoupled case still correlations due to RF overlap, but less sharp. [Pillow et al., 2008]
Unclear however if the IF model would perform better here than the Poisson GLM.

Summary

Predicting neural responses
In order of decreasing generality

- Linear models: simple, exact inference, but miss essential aspects of neural physiology
- Note higher orders may be captured by Wiener kernels, see Dayan & Abbott, chapter 2. Require more data to fit.
- Poisson GLM model: fewer parameters, spiking output, but lacks precise spike timing
- More neurally inspired models (I&F, GLM with spike feedback): good spike timing, but hard to fit, require careful regularization
- Biophyical models: in principle very precise, but in practice unwieldy

References I

Chichilnisky, E. J. (2001).
A simple white noise analysis of neuronal light responses.
Network, 12:199–203.

Dayan, P. and Abbott, L. F. (2002).
Theoretical Neuroscience.
MIT press, Cambridge, MA.

Pillow, J. W., Paninski, L., Uzzell, V. J., Simoncelli, E. P., and Chichilnisky, E. J. (2005).
Prediction and Decoding of Retinal Ganglion Cell Responses with a Probabilistic Spiking Model.
J Neurosci, 23:11003–11013.

Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E. J., and Simoncelli, E. P. (2008).
Spatio-temporal correlations and visual signalling in a complete neuronal population.
Nature, 454(7207):995–999.

Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., and Fried, I. (2005).
Invariant visual representation by single neurons in the human brain.
Nature, 435(7045):1102–1107.

Rieke, F., Warland, D., de Ruyter van Steveninck, R., and Bialek, W. (1996).
Spikes: Exploring the neural code.
MIT Press, Cambridge.

van Hateren, J. H., Rüttiger, L., Sun, H., and Lee, B. B. (2002).
Processing of Natural Temporal Stimuli by Macaque Retinal Ganglion Cells.
J Neurosci, 22:9945–9960.

References II

Wallis, J., Anderson, K. C., and Miller, E. K. (2001).
Single neurons in prefrontal cortex encode abstract rules.
Nature, 441:953–957.

Weber, A. I. and Pillow, J. W. (2017).
Capturing the dynamical repertoire of single neurons with generalized linear models.
Neural computation, 29(12):3260–3289.