Decoding an arbitrary continuous stimulus

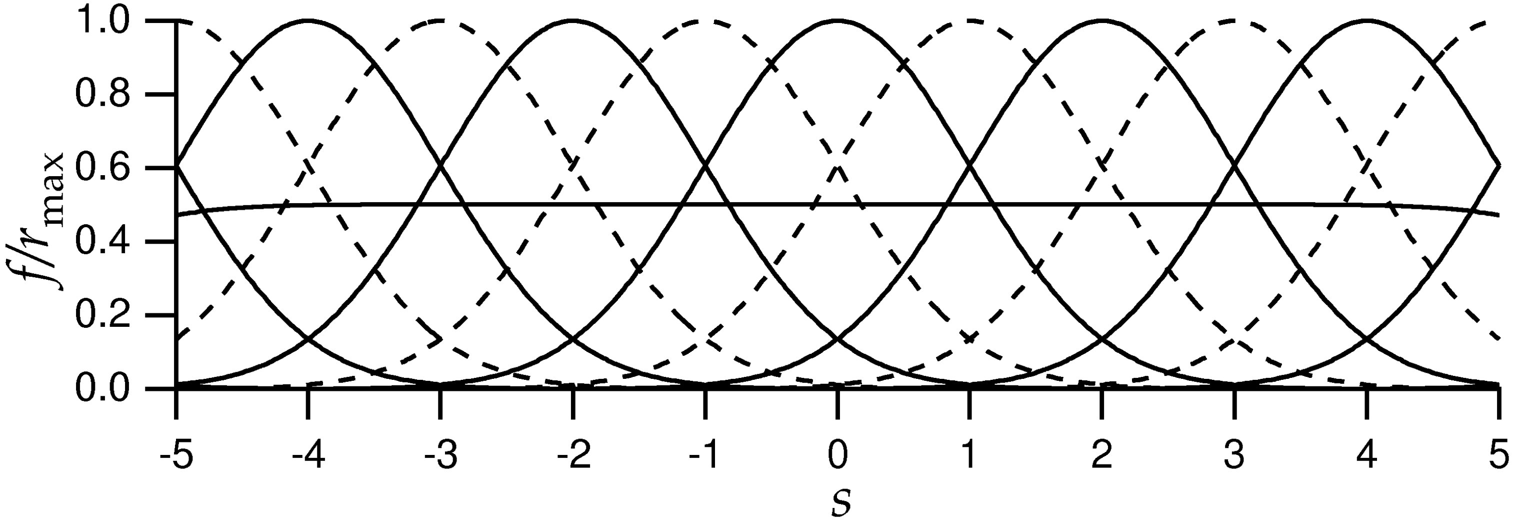

E.g. Gaussian tuning curves

$$f_a(s) = r_{\max} \exp\left(-\frac{1}{2}\left[\frac{(s - s_a)}{\sigma_a}\right]^2\right)$$

$$\sum_{a=1}^{N} f_a(s) \ \text{const.}$$

Decoding an arbitrary continuous stimulus

Many neurons "voting" for an outcome.

- assume independence
- assume Poisson firing

Noise model: Poisson distribution

$$P_T[k] = (\lambda T)^k \exp(-\lambda T)/k!$$

Need to know full P[**r**|s]

Assume Poisson: $$P[r_a|s] = \frac{(f_a(s)T)^{r_aT}}{(r_aT)!}\exp(-f_a(s)T)$$

Assume independent: $$P[\mathbf{r}|s] = \prod_{a=1}^{N} \frac{(f_a(s)T)^{r_aT}}{(r_aT)!}\exp(-f_a(s)T)$$

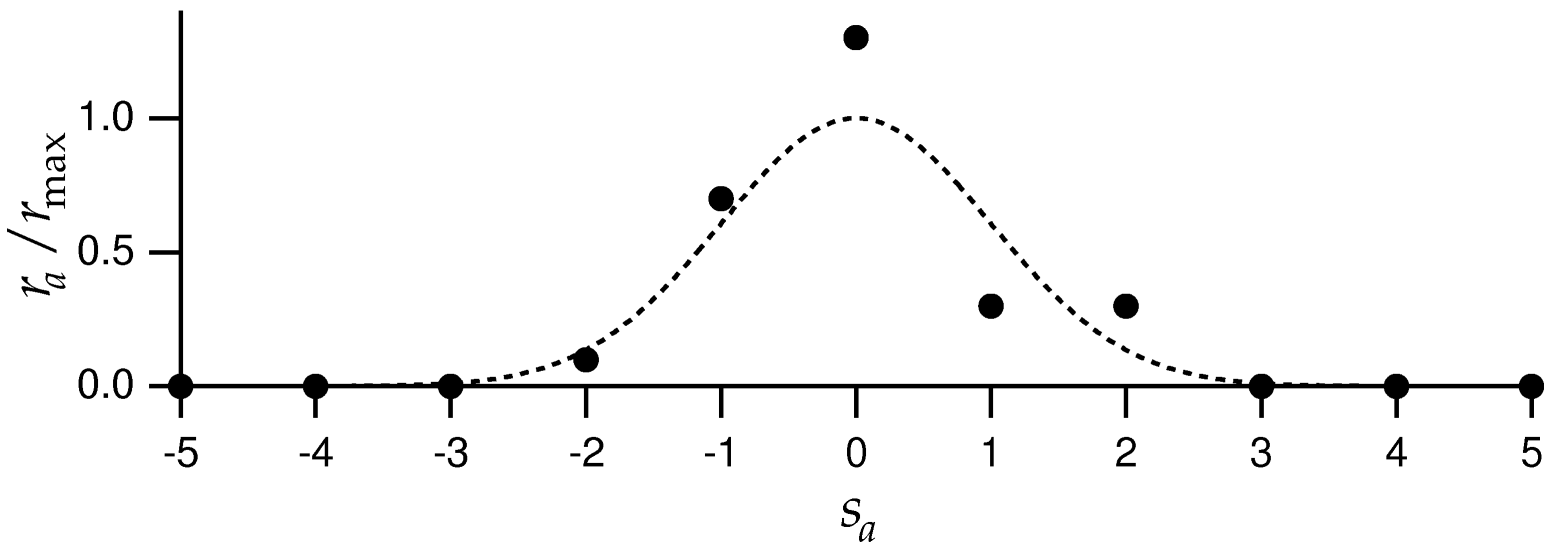

Population response of 11 cells with Gaussian tuning curves

ML

Apply ML: maximize ln P[**r**|s] with respect to s

$$\ln P[\mathbf{r}|s] = T\sum_{a=1}^{N} r_a \ln(f_a(s)) + \ldots$$

Set derivative to zero, use sum = constant

$$\sum_{a=1}^{N} r_a \frac{f'(s^*)}{f(s^*)} = 0$$

From Gaussianity of tuning curves,

$$s^* = \frac{\sum r_a s_a/\sigma_a^2}{\sum r_a/\sigma_a^2}$$

If all σ same

$$s^* = \frac{\sum r_a s_a}{\sum r_a}$$

MAP

Apply MAP: maximise ln p[s|**r**] with respect to s

$$\ln p[s|\mathbf{r}] = \ln P[\mathbf{r}|s] + \ln p[s] - \ln P[\mathbf{r}]$$

$$\ln p[s|\mathbf{r}] = T \sum_{a=1}^{N} r_a \ln(f_a(s)) + \ln p[s] + \ldots$$

Set derivative to zero, use sum = constant

$$\sum_{a=1}^{N} r_a \frac{f'(s^*)}{f(s^*)} + \frac{p'[s]}{p[s]} = 0$$

From Gaussianity of tuning curves,

$$s^* = \frac{T \sum r_a s_a / \sigma_a^2 + s_{\text{prior}} / \sigma_{\text{prior}}^2}{T \sum r_a / \sigma_a^2 + 1/\sigma_{\text{prior}}^2}$$

Given this data:

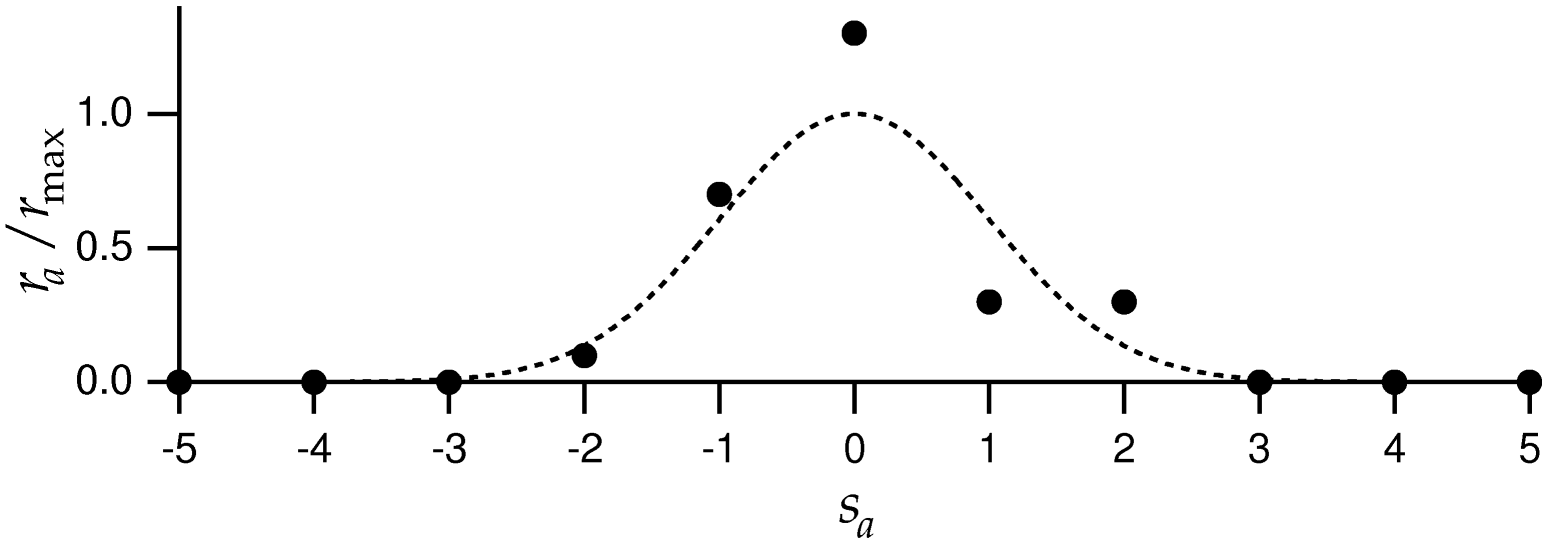

MAP:

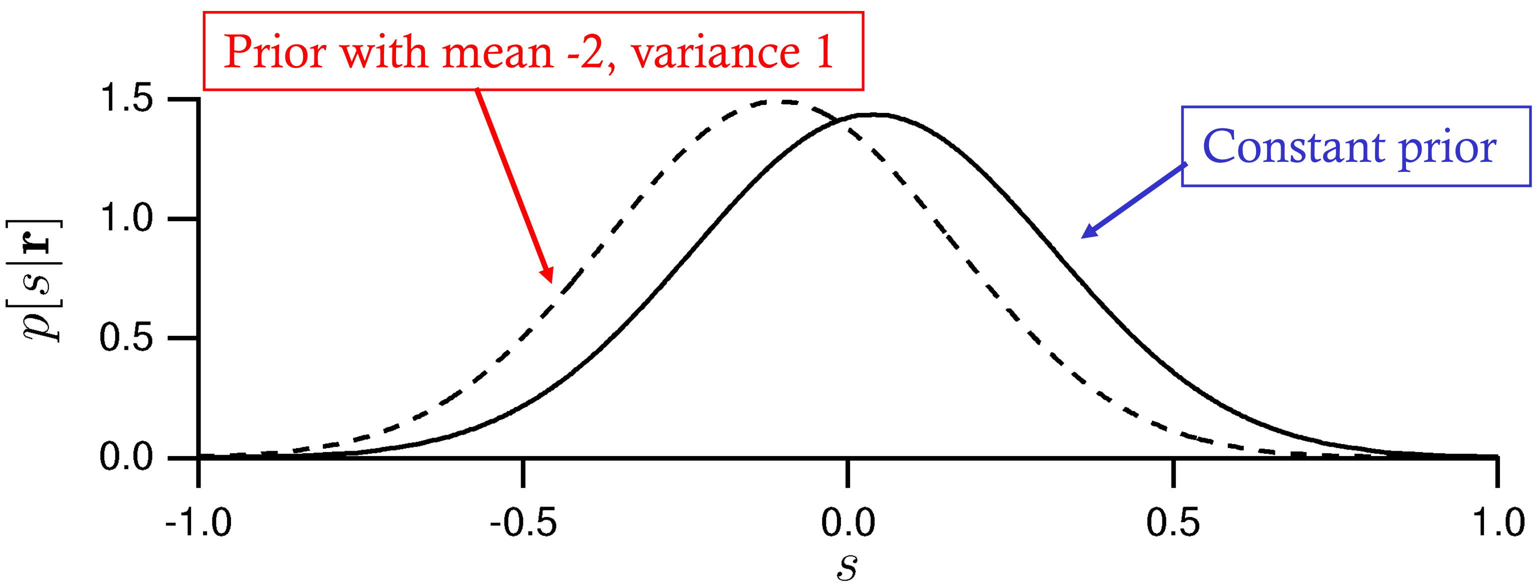

How good is our estimate?

For stimulus s, have estimated s_{est}

Bias: $$b_{\text{est}}(s) = \langle s_{\text{est}} - s \rangle$$

Variance: $$\sigma^2_{\text{est}}(s) = \langle (s_{\text{est}} - \langle s_{\text{est}} \rangle)^2 \rangle$$

Mean square error:

$$\langle (s_{\text{est}} - s)^2 \rangle = \langle (s_{\text{est}} - \langle s_{\text{est}} \rangle + b_{\text{est}}(s))^2 \rangle = \sigma^2_{\text{est}}(s) + b^2_{\text{est}}(s).$$

Cramer-Rao bound: $$\sigma^2_{\text{est}} \geq \frac{(1 + b'_{\text{est}})^2}{I_{\text{F}}(s)}$$

Fisher information

Fisher information

$$I_{\mathrm{F}}(s) = \left\langle -\frac{\partial^2 \ln p[\mathbf{r}|s]}{\partial^2 s} \right\rangle = \int d\mathbf{r}\, p[\mathbf{r}|s] \left(-\frac{\partial^2 \ln p[\mathbf{r}|s]}{\partial s^2} \right)$$

Alternatively:

$$I_{\mathrm{F}}(s) = \left\langle \left(\frac{\partial \ln p[\mathbf{r}|s]}{\partial s} \right)^2 \right\rangle = \int d\mathbf{r}\, p[\mathbf{r}|s] \left(\frac{\partial \ln p[\mathbf{r}|s]}{\partial s} \right)^2$$

Quantifies local stimulus discriminability

Entropy and Shannon information

Entropy and Shannon information

For a random variable *X* with distribution *p(x),* the **entropy** is

$$H[X] = - \Sigma_x \, p(x) \log_2 p(x)$$

Information is defined as

$$I[X] = - \log_2 p(x)$$

Mutual Information between X and Y is defined as

$$MI[X,Y] = H[X] - E_y[H[X|Y=y]]$$
$$= H[Y] - E_x[H[Y|X=x]]$$

Information in single spikes

How much information does a single spike convey about the stimulus?

Key idea: the information that a spike gives about the stimulus is the reduction in entropy between the distribution of spike times not knowing the stimulus, and the distribution of times knowing the stimulus.

The response to an (arbitrary) stimulus sequence ***s*** is r(t).

Without knowing that the stimulus was ***s***, the probability of observing a spike in a given bin is proportional to $\bar{r}$, the mean rate, and the size of the bin.

Consider a bin Δt small enough that it can only contain a single spike. Then in the bin at time t,

$$\begin{aligned} P(r=1) &= \bar{r}\Delta t, \\ P(r=0) &= 1-\bar{r}\Delta t, \\ P(r=1|s) &= r(t)\Delta t, \\ P(r=0|s) &= 1-r(t)\Delta t. \end{aligned}$$

Information in single spikes

Now compute the entropy difference: $p = \bar{r}\Delta t$, $p(t) = r(t)\Delta t$

$$I(r,s) = -p\log p - (1-p)\log(1-p) +$$ ← prior

$$+\frac{1}{T}\int_0^T dt\,\left[p(t)\log p(t) + (1-p(t))\log(1-p(t))\right].$$ ← conditional

Note substitution of a time average for an average over the *r* ensemble.

Assuming $p \ll 1$ $\log(1-p) \sim$ $-p$ and using $\frac{1}{T}\int_0^T dt\,p(t) \rightarrow p$

$$I(r,s) = \frac{1}{T}\int_0^T dt\,\Delta t\,r(t)\log\frac{r(t)}{\bar{r}} + Var(p(t))/2ln2 + O(p^3).$$

In terms of information per spike (divide by $\bar{r}\Delta t$):

$$I(r,s) = \frac{1}{T}\int_0^T dt\,\frac{r(t)}{\bar{r}}\log\frac{r(t)}{\bar{r}}$$

Using information to evaluate neural models

We can use the information about the stimulus to evaluate our reduced dimensionality models.

Evaluating models using information

Mutual information is a measure of the reduction of uncertainty about one quantity that is achieved by observing another.

Uncertainty is quantified by the entropy of a probability distribution,

$$\sum p(x) \log_2 p(x).$$

We can compute the information in the spike train directly, without direct reference to the stimulus (Brenner et al., Neural Comp., 2000)

This sets an upper bound on the performance of the model.

Repeat a stimulus of length T many times and compute the time-varying rate r(t), which is the probability of spiking given the stimulus.

Evaluating models using information

Information in timing of 1 spike: $$I_{\text{one spike}} = \frac{1}{T}\int_0^T dt\, \frac{r(t)}{\bar{r}} \log_2\left[\frac{r(t)}{\bar{r}}\right]$$

$$\boxed{\frac{r(t)}{\bar{r}} = \frac{P(\text{spike at } t|\mathbf{s})}{P(\text{spike at } t)}}$$

By definition

Evaluating models using information

Given: $$I_{\text{one spike}} = \frac{1}{T}\int_0^T dt\, \frac{r(t)}{\bar{r}} \log_2\left[\frac{r(t)}{\bar{r}}\right]$$

$$\frac{r(t)}{\bar{r}} = \boxed{\frac{P(\text{spike at } t|\mathbf{s})}{P(\text{spike at } t)} = \frac{P(\mathbf{s}|\text{spike at } t)}{P(\mathbf{s})}}$$

By definition Bayes' rule

Evaluating models using information

Given: $$I_{\text{one spike}} = \frac{1}{T}\int_0^T dt\,\frac{r(t)}{\bar{r}}\log_2\left[\frac{r(t)}{\bar{r}}\right]$$

$$\frac{r(t)}{\bar{r}} = \frac{P(\text{spike at } t|\mathbf{s})}{P(\text{spike at } t)} = \boxed{\frac{P(\mathbf{s}|\text{spike at } t)}{P(\mathbf{s})} \rightarrow \frac{P(s_1, s_2, s_3, \ldots|\text{spike at } t)}{P(s_1, s_2, s_3, \ldots)}}$$

By definition Bayes' rule Dimensionality reduction

Evaluating models using information

Given: $$I_{\mathrm{one\,spike}} = \frac{1}{T}\int_0^T dt\,\frac{r(t)}{\bar{r}}\log_2\left[\frac{r(t)}{\bar{r}}\right]$$

$$\frac{r(t)}{\bar{r}} = \frac{P(\mathrm{spike\,at}\,t|\mathbf{s})}{P(\mathrm{spike\,at}\,t)} = \boxed{\frac{P(\mathbf{s}|\mathrm{spike\,at}\,t)}{P(\mathbf{s})} \rightarrow \frac{P(s_1, s_2, s_3, \ldots|\mathrm{spike\ \ at}\,t)}{P(s_1, s_2, s_3, \ldots)}}$$

By definition Bayes' rule Dimensionality reduction

So the information in the K-dimensional model is evaluated using the distribution of projections:

$$I^K_{\mathrm{one\,spike}} = \int d^K s\,P(s_1, \ldots, s_K|\mathrm{spike\,at}\,t)\log_2\left[\frac{P(s_1, \ldots, s_K|\mathrm{spike\,at}\,t)}{P(s_1, \ldots, s_K)}\right].$$

Using information to evaluate neural models

Here we used information to evaluate reduced models of the Hodgkin-Huxley neuron.

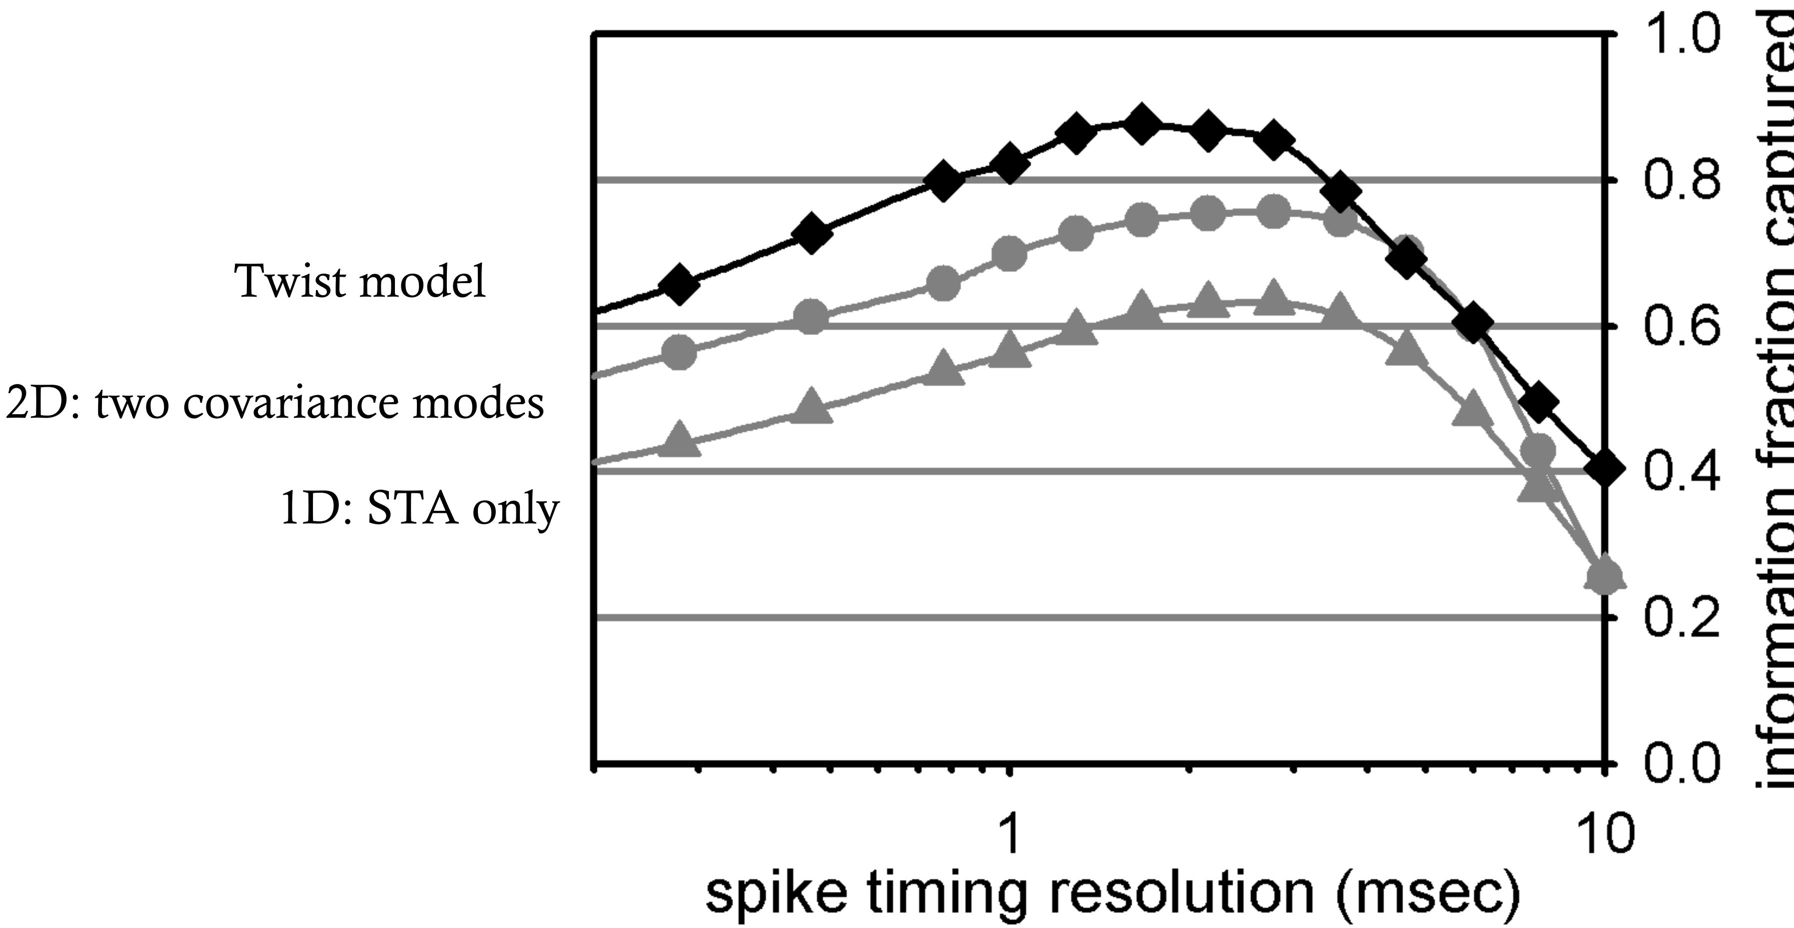

Information in 1D

The STA is the single most informative dimension.

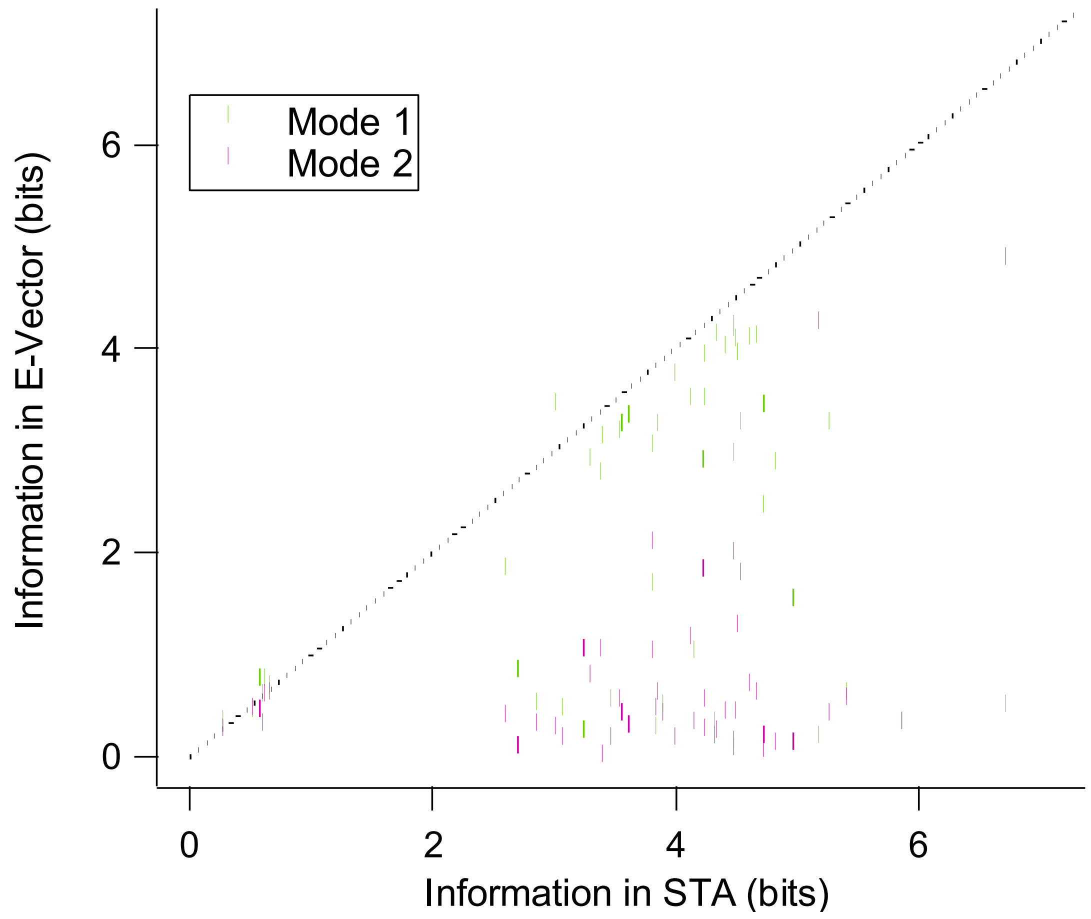

Information in 1D

•The information is related to the eigenvalue of the corresponding eigenmode

•Negative eigenmodes are much more informative

•Information in STA and leading negative eigenmodes up to 90% of the total

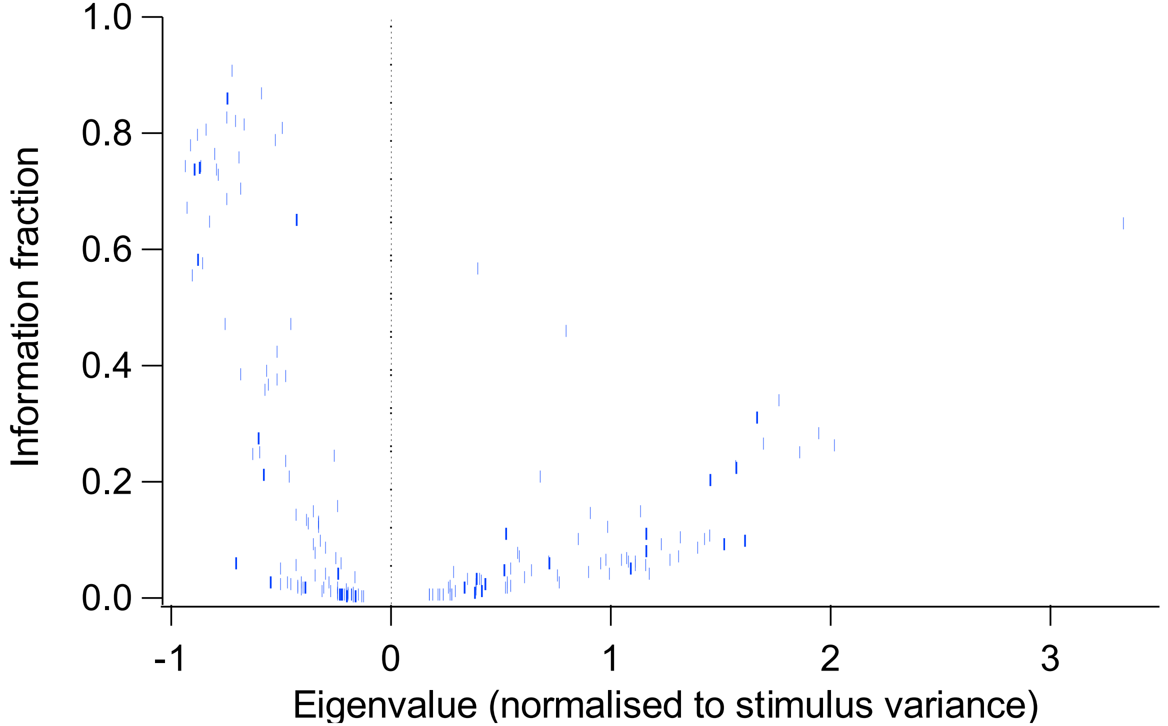

Information in 2D

- We recover significantly more information from a 2-dimensional description

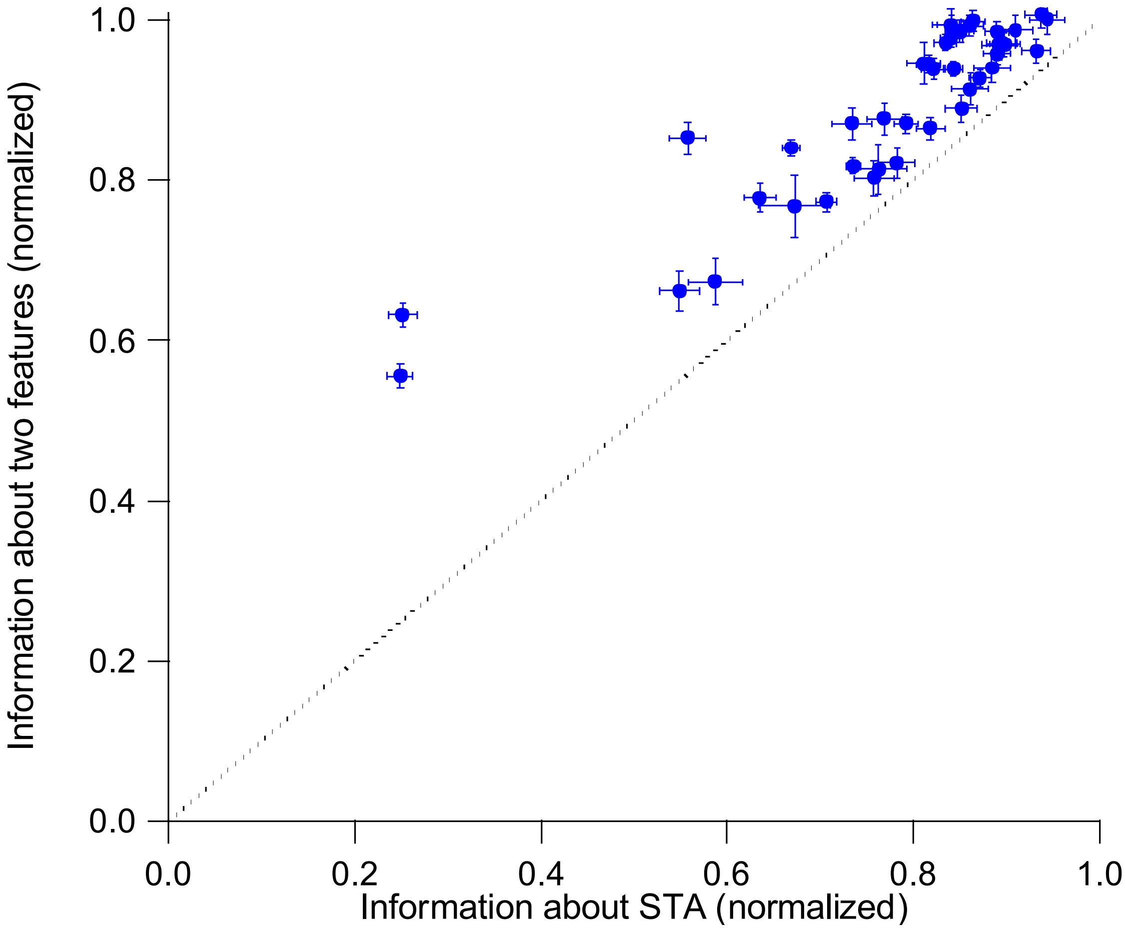

Calculating information in spike trains

How can one compute the entropy and information of spike trains?

Entropy:

Discretize the spike train into binary words *w* with letter size Δt, length T. This takes into account correlations between spikes on timescales TΔt.

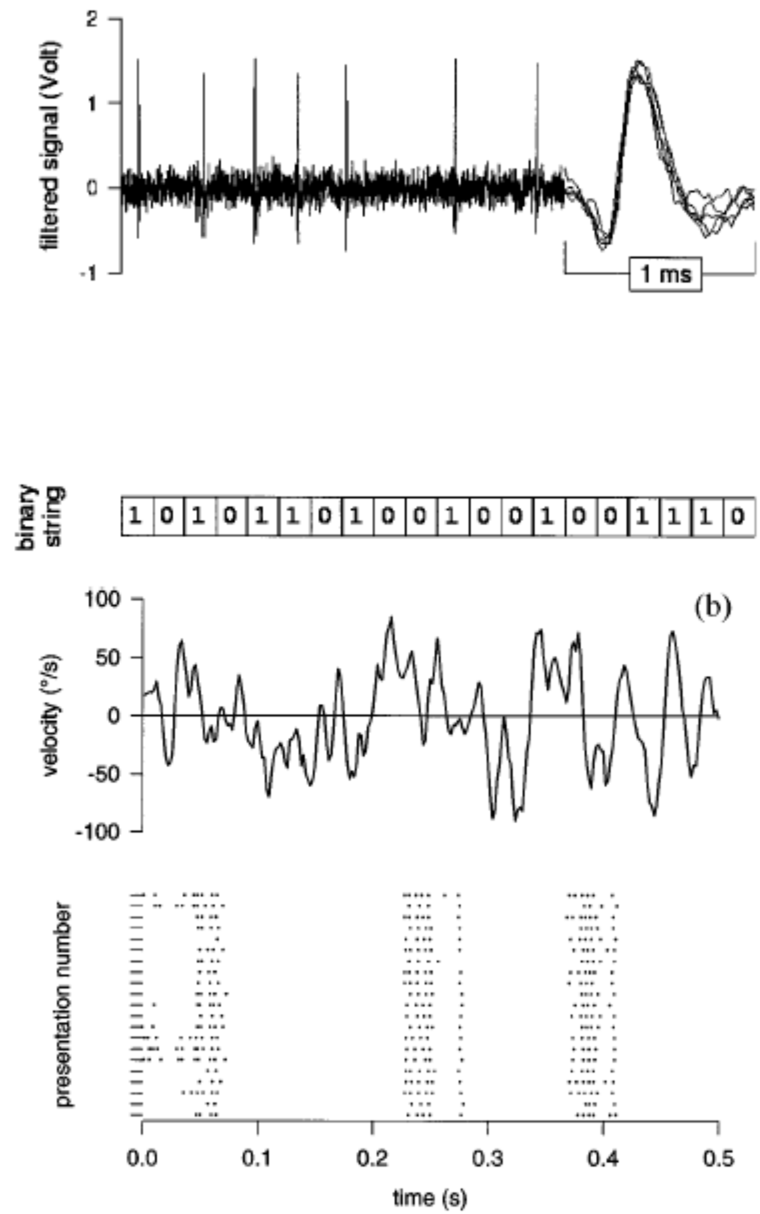

Strong et al., 1997; Panzeri et al.

Calculating information in spike trains

Information : difference between the variability driven by stimuli and that due to noise.

Take a stimulus sequence *s* and repeat many times.

For each time in the repeated stimulus, get a set of words P(*w*|*s*(t)).

Average over *s* → average over time:

$$H_{noise} = \langle H[P(w|s_i)] \rangle_i.$$

Choose length of repeated sequence long enough to sample the noise entropy adequately.

Finally, do as a function of word length T and extrapolate to infinite T.

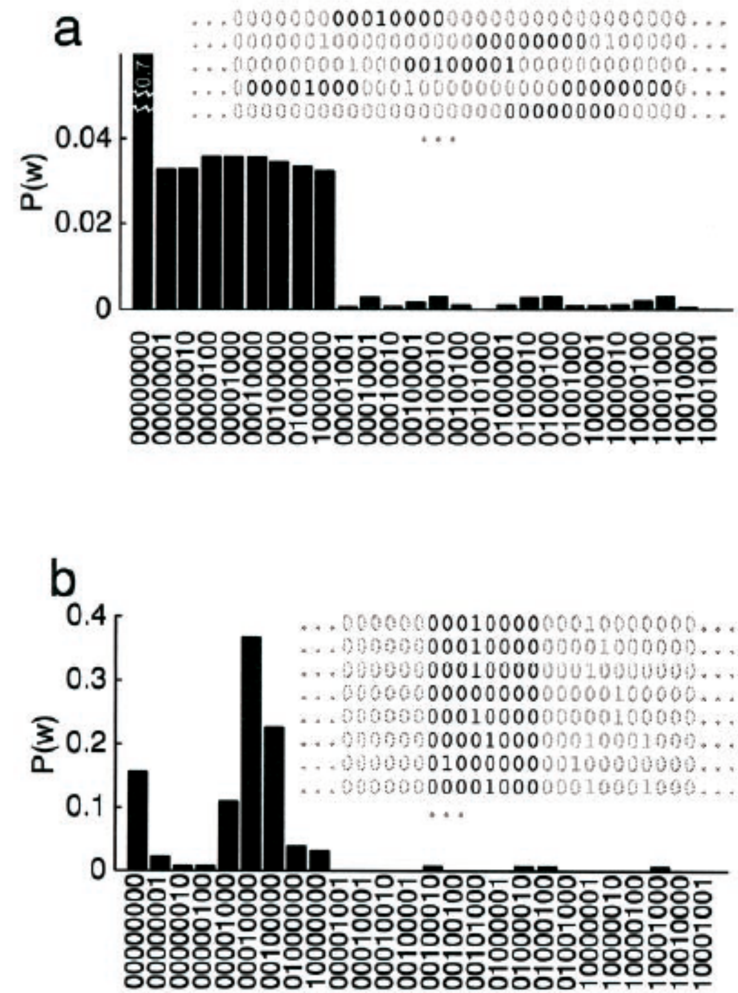

Reinagel and Reid, '00

Calculating information in spike trains

Fly H1:
obtain information rate of
~80 bits/sec or 1-2 bits/spike.

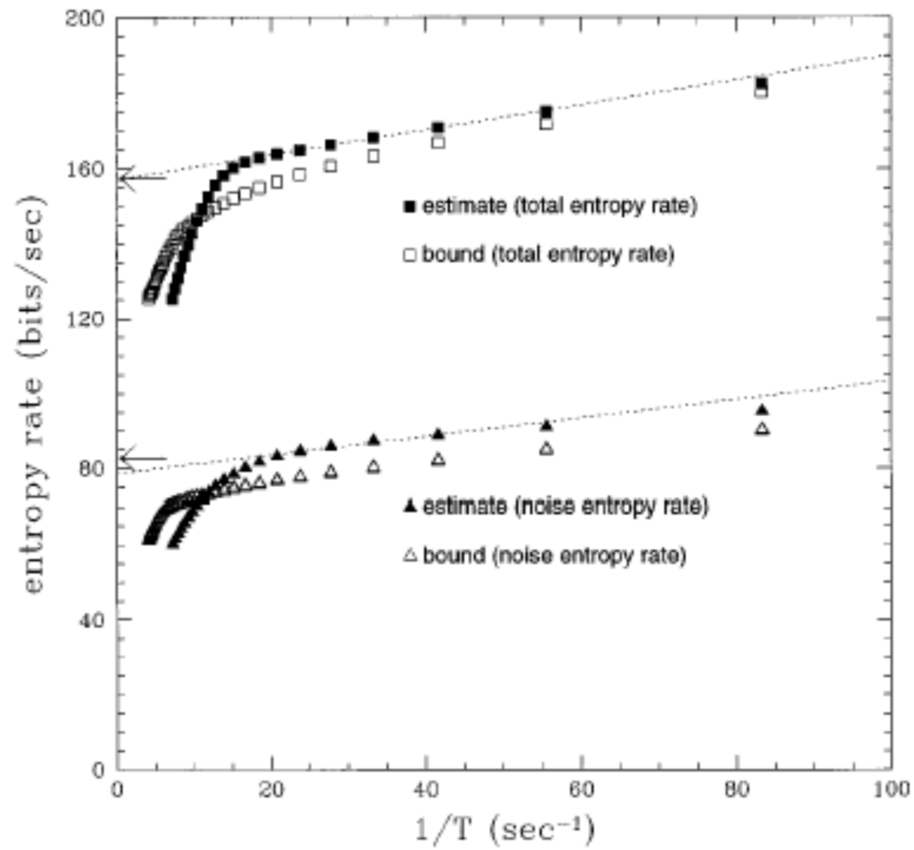

Calculating information in the LGN

Another example: temporal coding in the LGN (Reinagel and Reid '00)

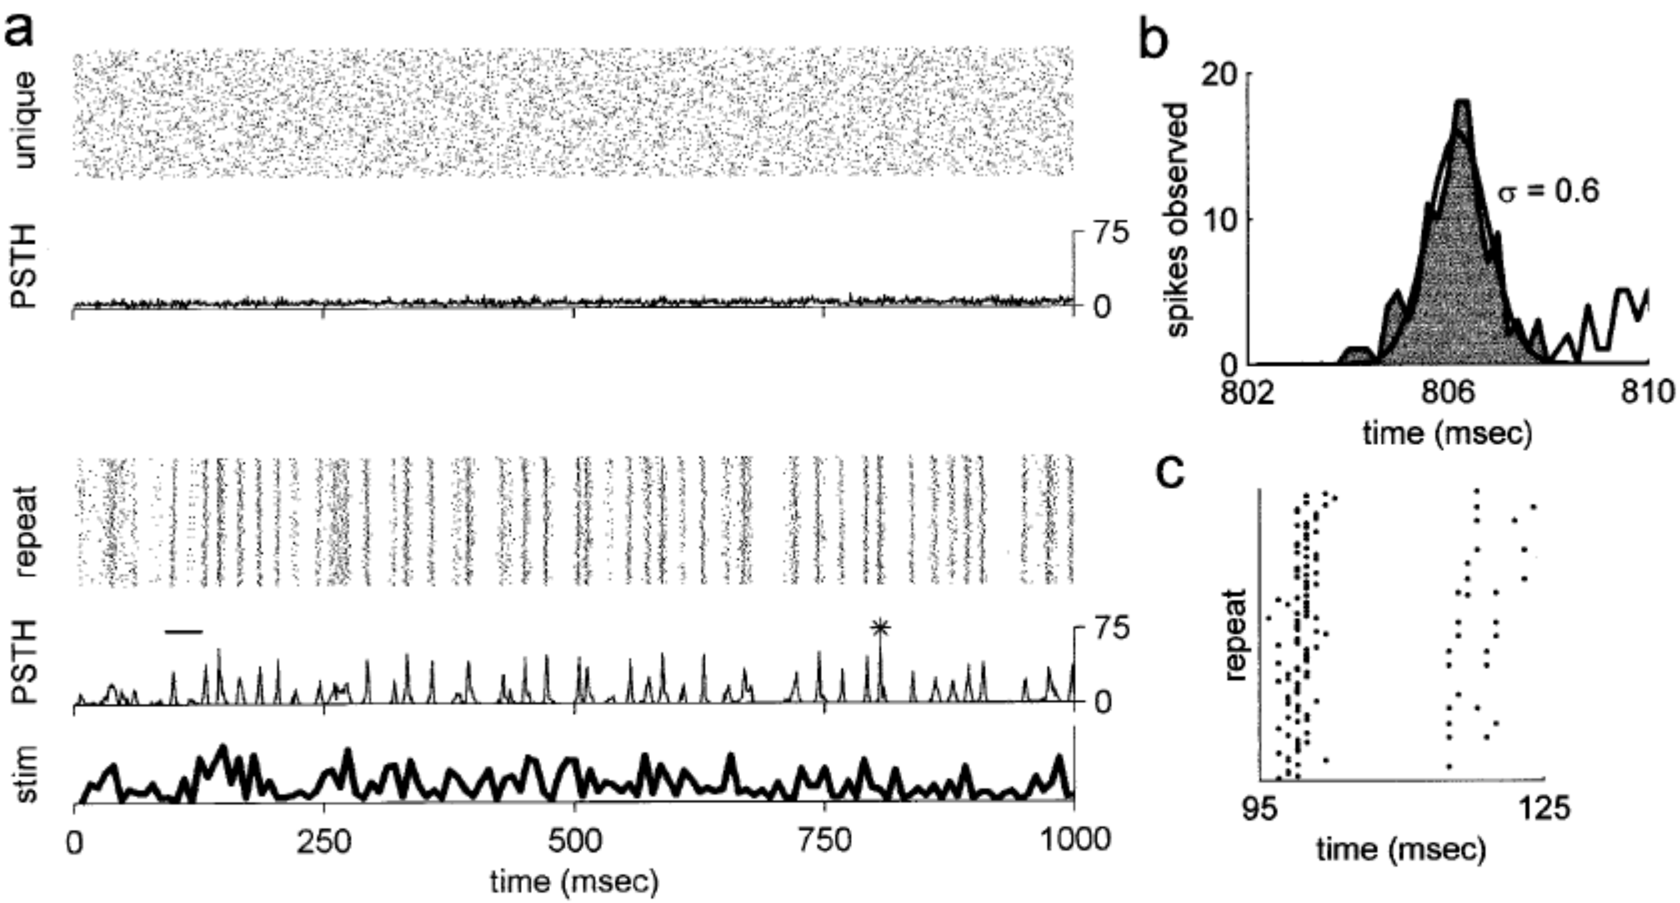

Calculating information in the LGN

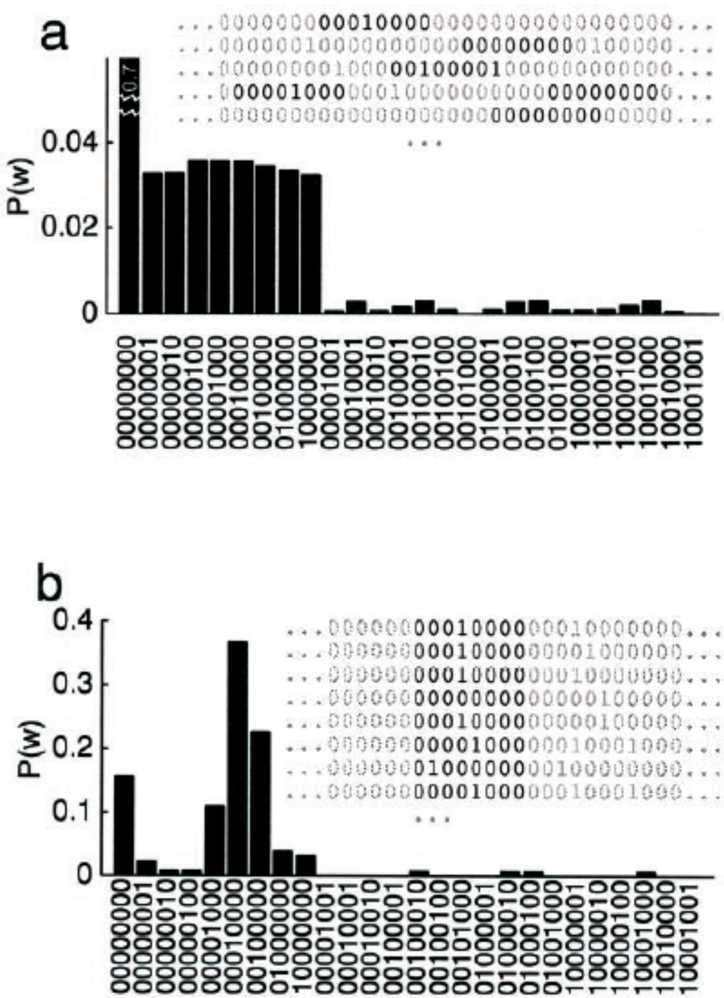

Apply the same procedure:
collect word distributions
for a random, then repeated stimulus.

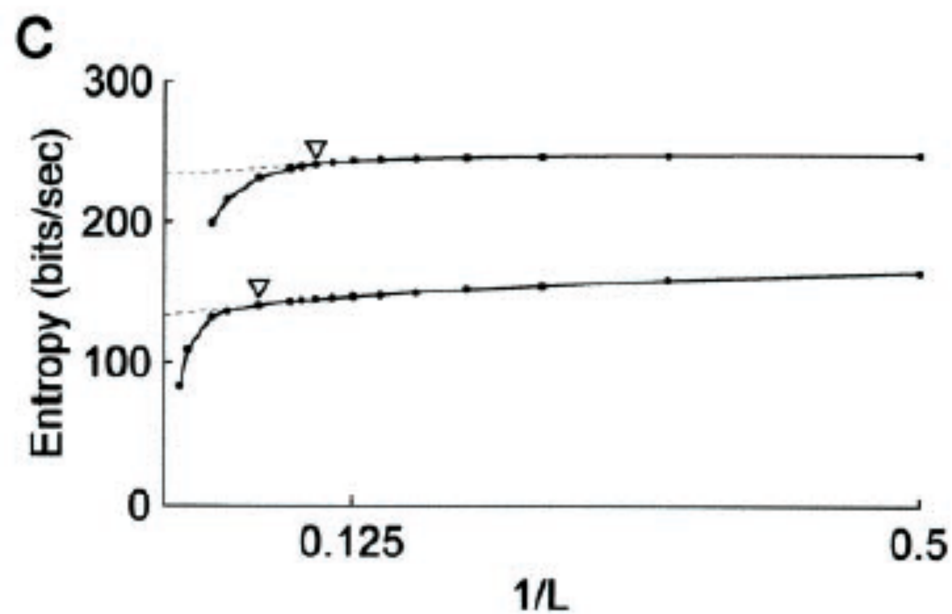

Information in the LGN

Use this to quantify how
precise the code is,
and over what timescales
correlations are important.