Decoding

How well can we learn what the stimulus is by looking at the neural responses?

Two approaches:

- devise explicit algorithms for extracting a stimulus estimate
- directly quantify the relationship between stimulus and response using information theory

Reading minds: the LGN

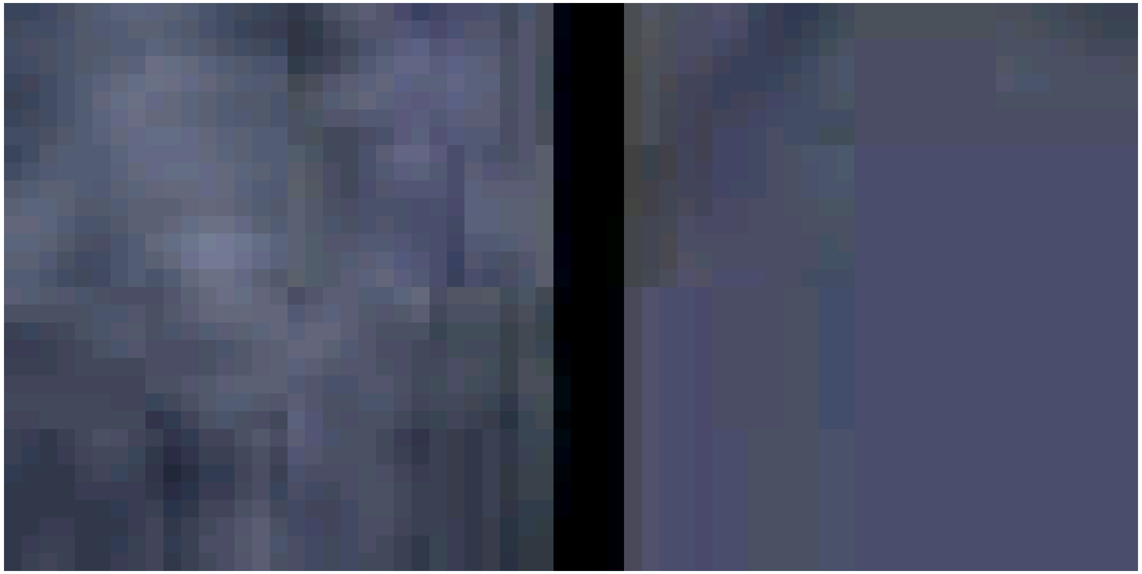

Yang Dan, UC Berkeley

Two-alternative tasks

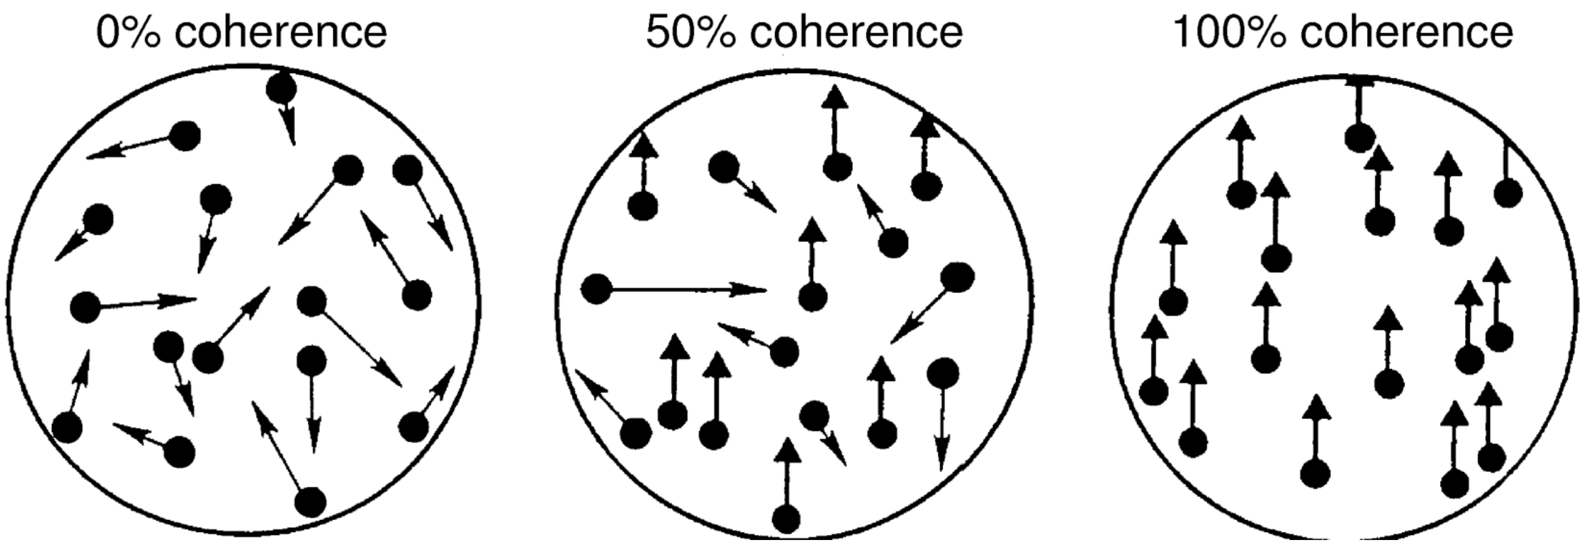

Britten et al. ‘92: behavioral monkey data + neural responses

Behavioral performance

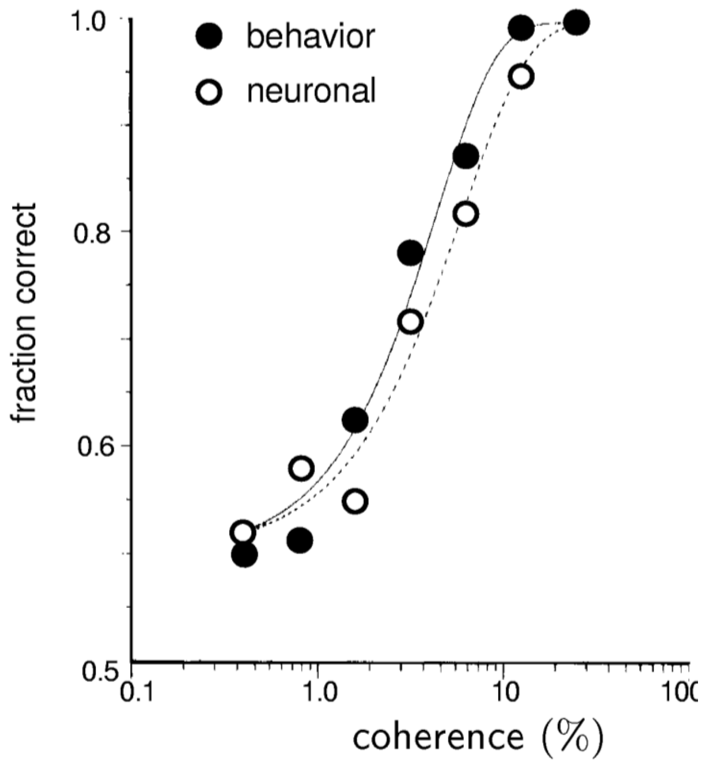

Predictable from neural activity?

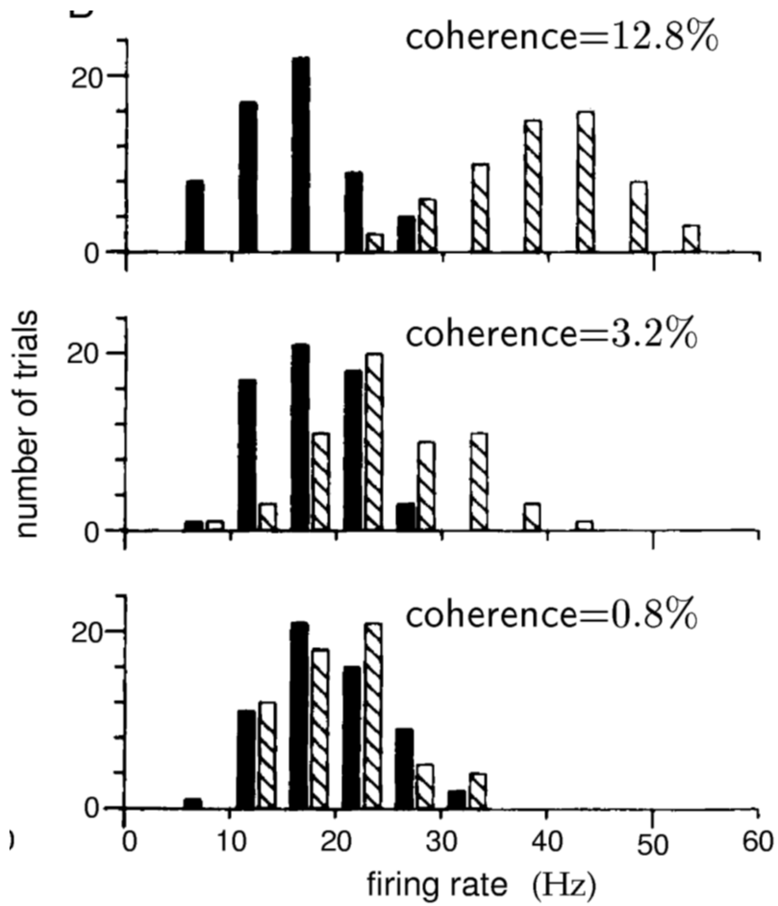

Discriminability: $d' = (\langle r \rangle_+ - \langle r \rangle_-) / \sigma_r$

Signal detection theory

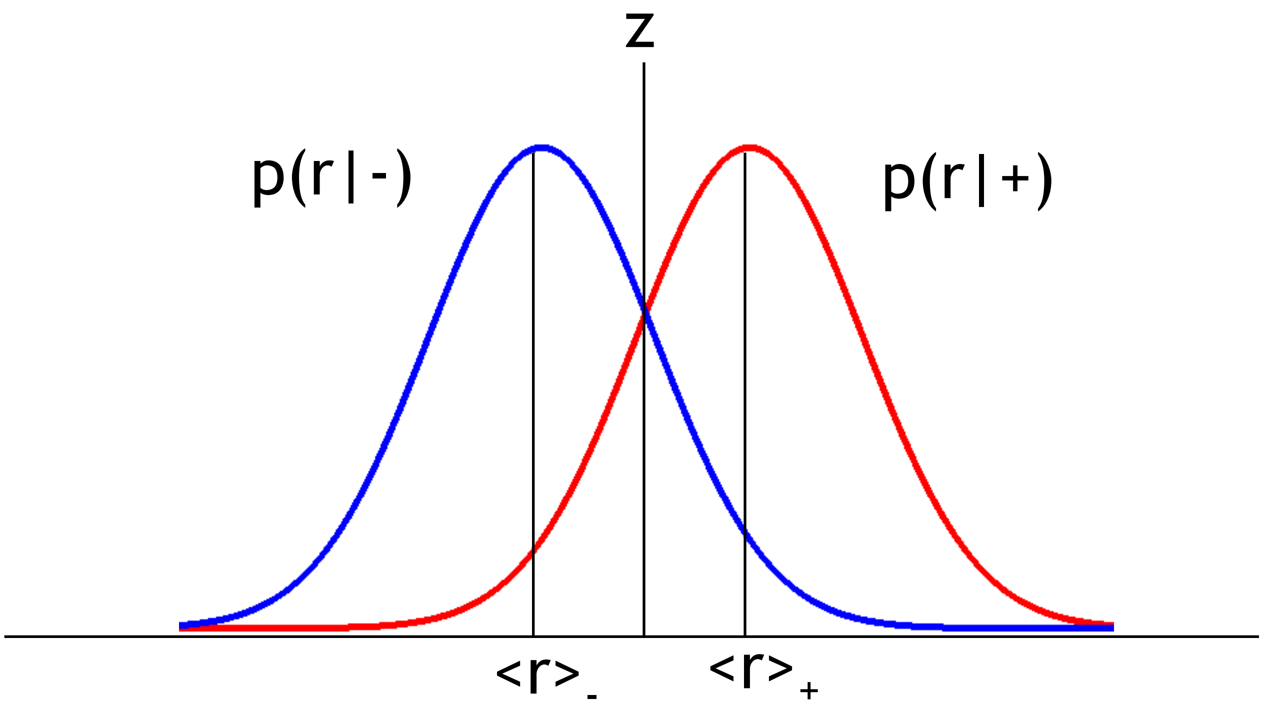

Decoding corresponds to comparing test, r, to threshold.

$\alpha(z) = P[\, r \geq z|-]$ false alarm rate, "size"

$\beta(z) = P[\, r \geq z|+]$ hit rate, "power"

Find z by maximizing $P[\text{correct}] = p(+)\,\beta(z) + p(-)(1 - \alpha(z))$

ROC curves

summarize performance of test for different thresholds *z*

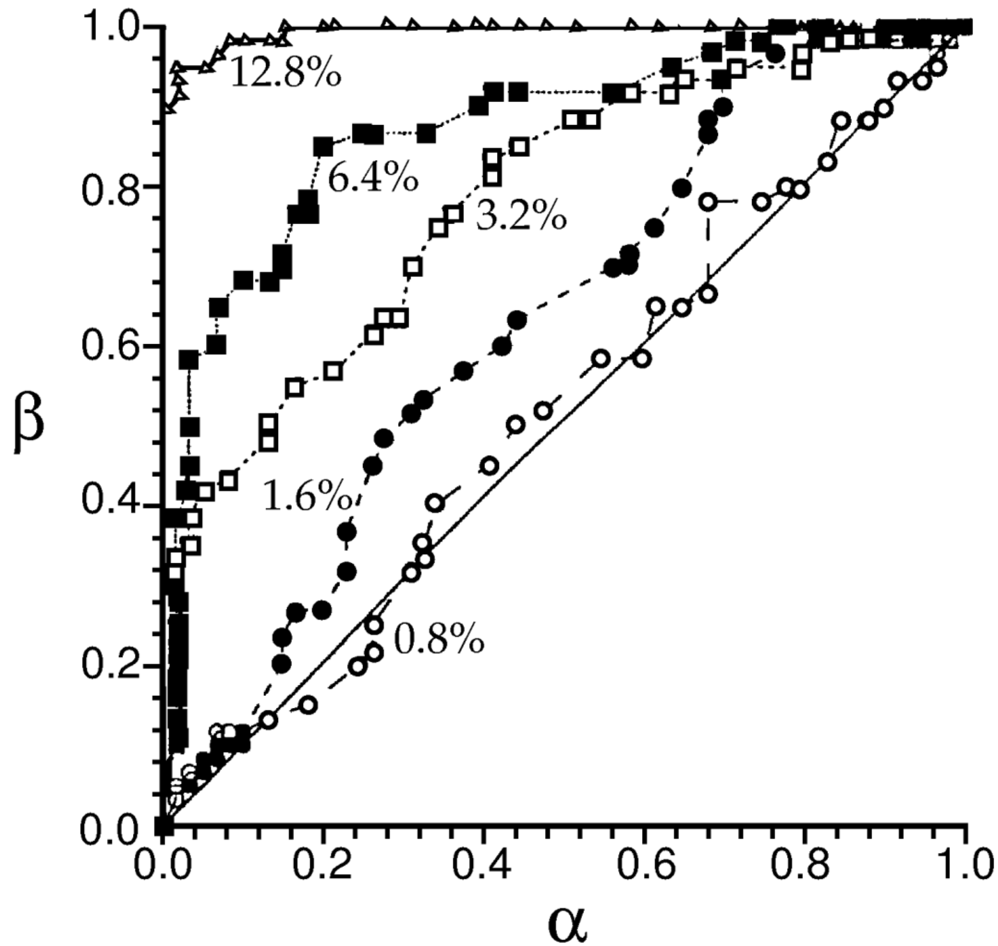

Want $\beta \rightarrow 1$, $\alpha \rightarrow 0$.

ROC curves

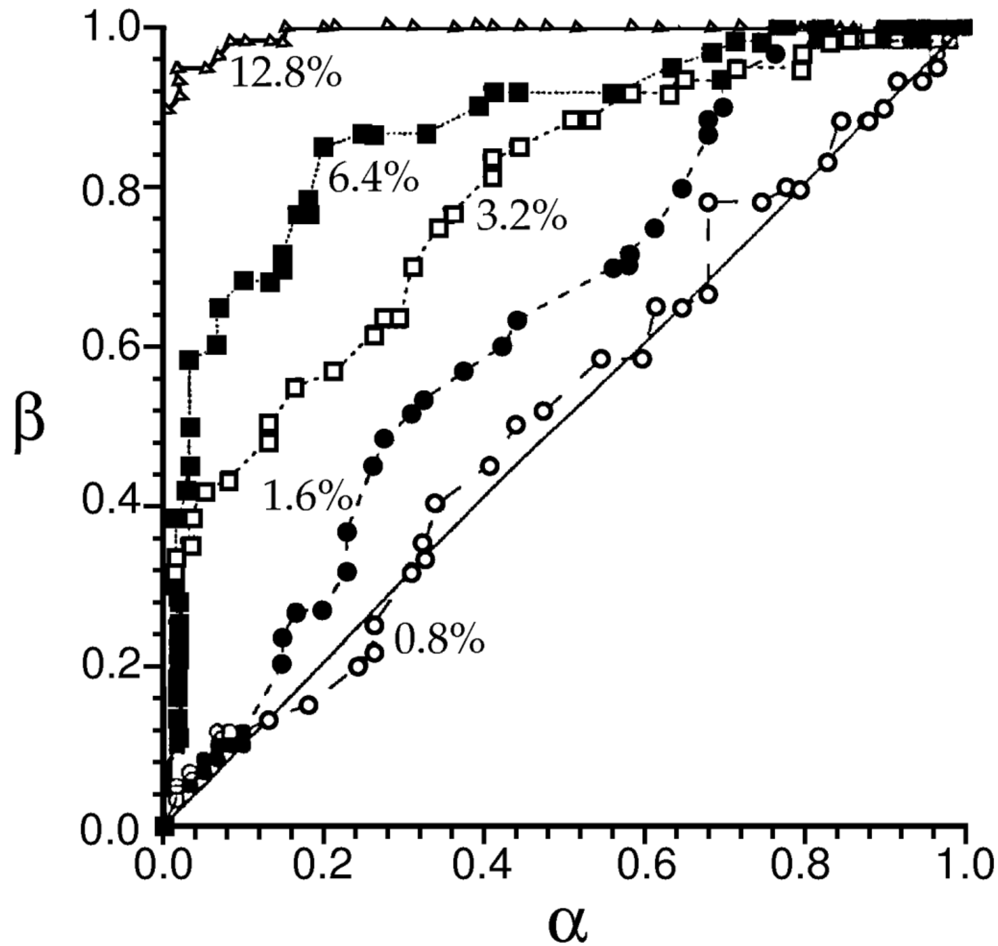

The area under the ROC curve corresponds to P[correct] for a two-alternative forced choice task

Emergence of the logistic function

If p[r|+] and p[r|-] are both Gaussian,
P[correct] = ½ erfc(-d’/2).

To interpret results as two-alternative forced choice, need simultaneous responses from + neuron and from – neuron.
Get “- neuron” responses from same neuron in response to – stimulus.

Ideal observer: performs as area under ROC curve.

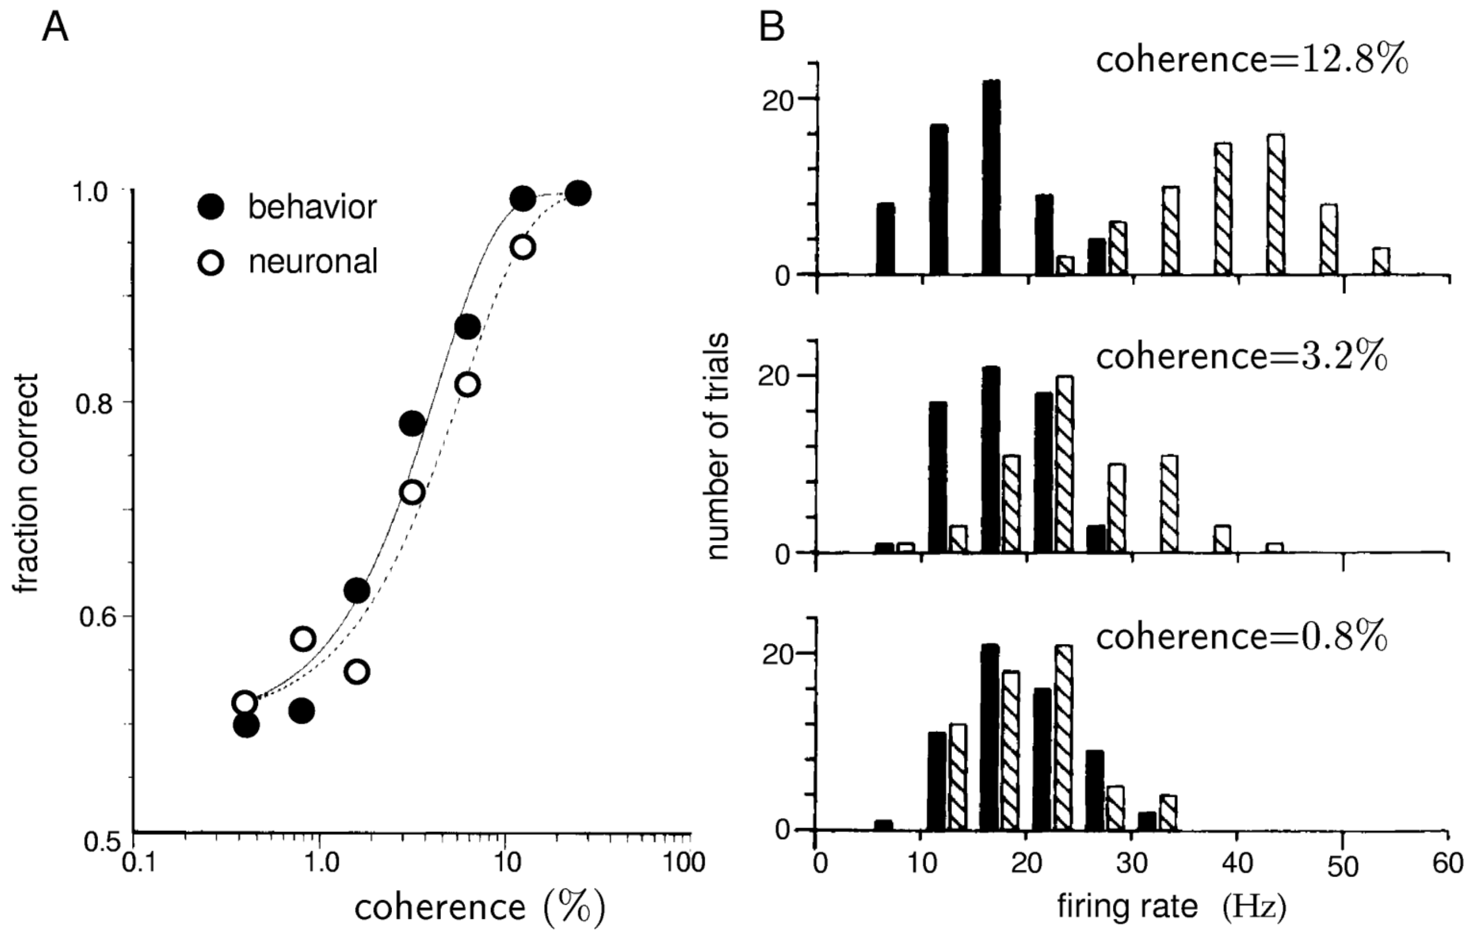

Close correspondence between neural and behaviour..

Why so many neurons? Correlations limit performance.

Is there a better test to use than r?

The optimal test function is the *likelihood ratio*,

$$l(r) = p[r|+] \,/\, p[r|-].$$

(Neyman-Pearson lemma)

Note that

$$l(z) = (d\beta/dz) \;/\; (d\alpha/dz) \;=\; d\beta/d\alpha$$

i.e. slope of ROC curve

Likelihood as loss minimization

Penalty for incorrect answer: L_+, L_-
Observe r;

Expected loss $\quad Loss_+ = L_+ P[-|r]$
$\qquad\qquad\qquad\quad Loss_- = L_- P[+|r]$

Cut your losses: answer + when $Loss_+ < Loss_-$

i.e. $L_+ P[-|r] > L_- P[+|r]$.

Using Bayes', $\quad P[+|r] = p[r|+]P[+]/p(r)$;
$\qquad\qquad\qquad\quad P[-|r] = p[r|-]P[-]/p(r)$;

$\rightarrow l(r) = p[r|+]/p[r|-] > L_+ P[-] / L_- P[+]$.

The role of d’

Again, if p[r|-] and p[r|+] are Gaussian,
and P[+] and P[-] are equal,

$$P[+|r] = 1/\ [1 + \exp(-d'\ (r - \langle r \rangle)/\ \sigma)].$$

+

→ d’ is the slope of the sigmoidal fitted to P[+|r]

Likelihood and tuning curves

For small stimulus differences s and s + δs

$$\frac{p[r|s+\delta s]}{p[r|s]} \sim \frac{p[r|s] + \delta s \partial p[r|s]/\partial s}{p[r|s]}$$

$$= 1 + \delta s \frac{\partial \ln p[r|s]}{\partial s}.$$

comparing $Z(r) = \frac{\partial \ln p[r|s]}{\partial s}$

to threshold $(z-1)/\delta s$

Population coding

- Population code formulation
- Methods for decoding:
  - population vector
  - Bayesian inference
  - maximum a posteriori
  - maximum likelihood
- Fisher information

Cricket cercal cells coding wind velocity

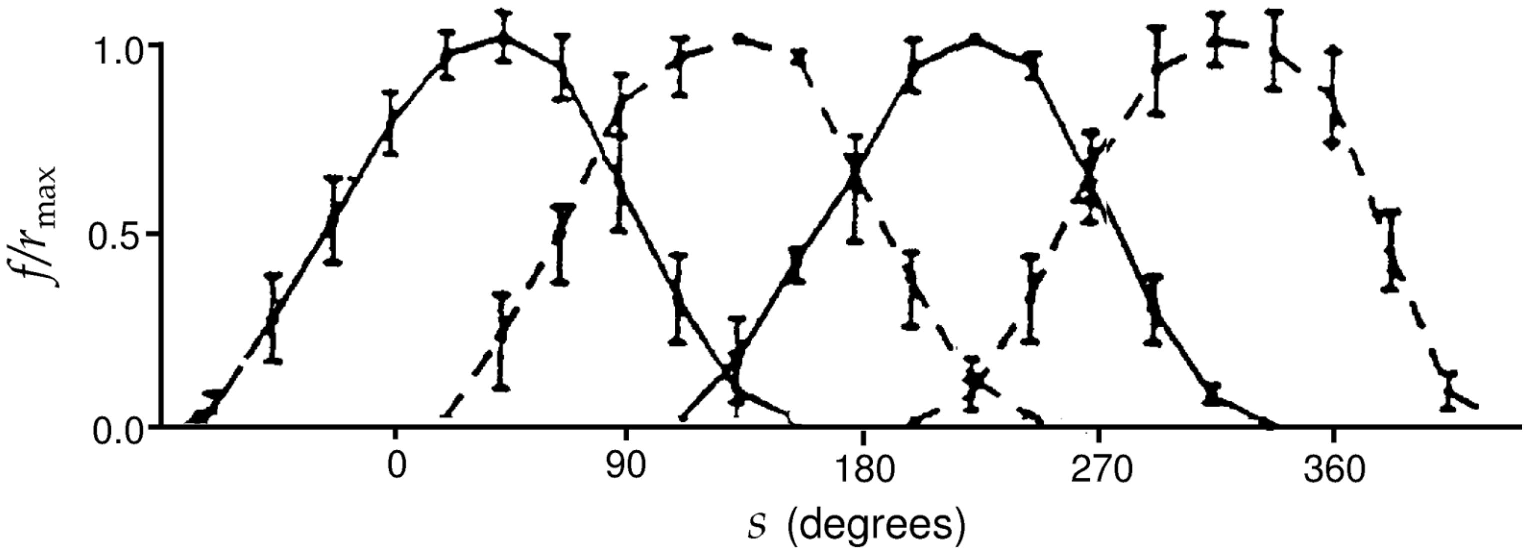

$$\left(\frac{f(s)}{r_{\max}}\right)_a = [\cos(s - s_a)]_+$$

$$\left(\frac{f(s)}{r_{\max}}\right)_a = [\vec{v} \cdot \vec{c}_a]_+$$

Population vector

$$\vec{v}_{\text{pop}} = \sum_{a=1}^{4} \left(\frac{r}{r_{\max}} \right)_a \vec{c}_a$$

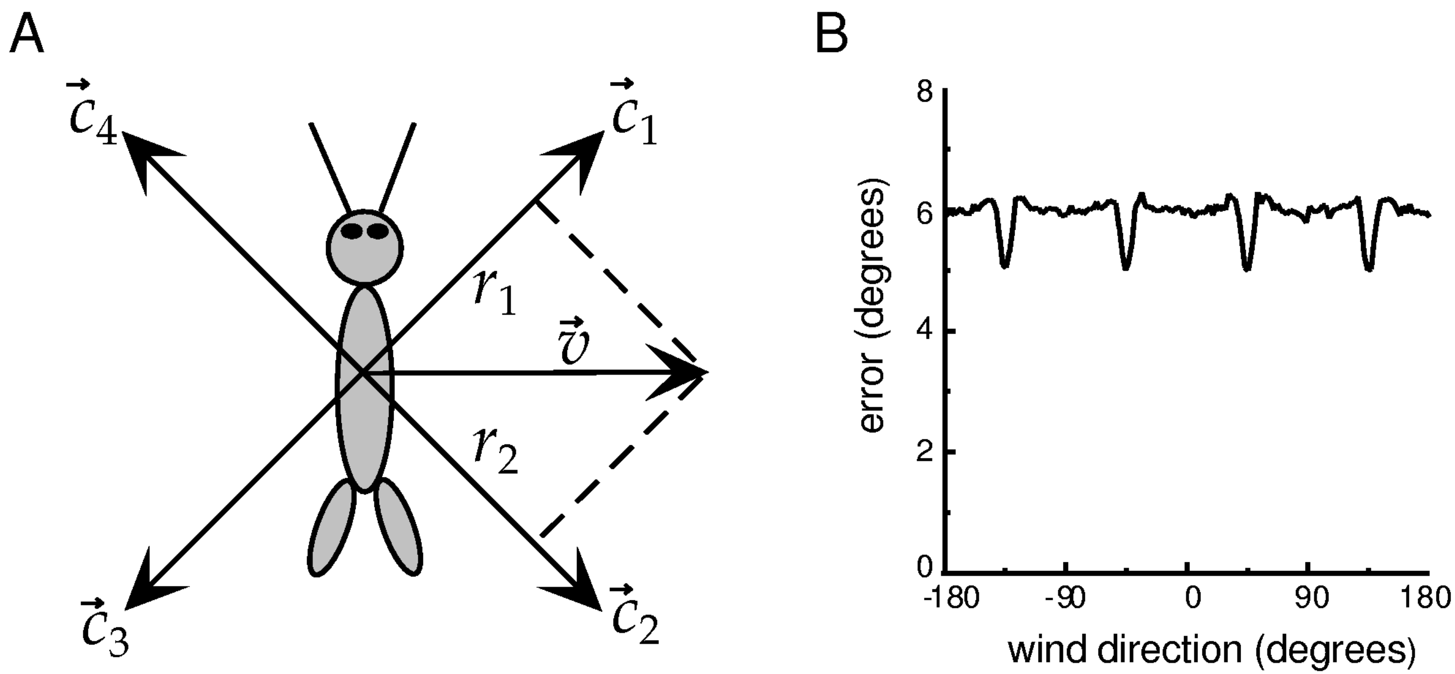

RMS error in estimate

Theunissen & Miller, 1991

Population coding in M1

Cosine tuning:

$$\left(\frac{\langle r \rangle - r_0}{r_{\max}}\right)_a = \left(\frac{f(s) - r_0}{r_{\max}}\right)_a = \vec{v} \cdot \vec{c}_a$$

Pop. vector:

$$\vec{v}_{\text{pop}} = \sum_{a=1}^{N} \left(\frac{r - r_0}{r_{\max}}\right) \vec{c}_a$$

For sufficiently large N,

$$\langle \vec{v}_{\text{pop}} \rangle = \sum_{a=1}^{N} (\vec{v} \cdot \vec{c}_a) \vec{c}_a$$

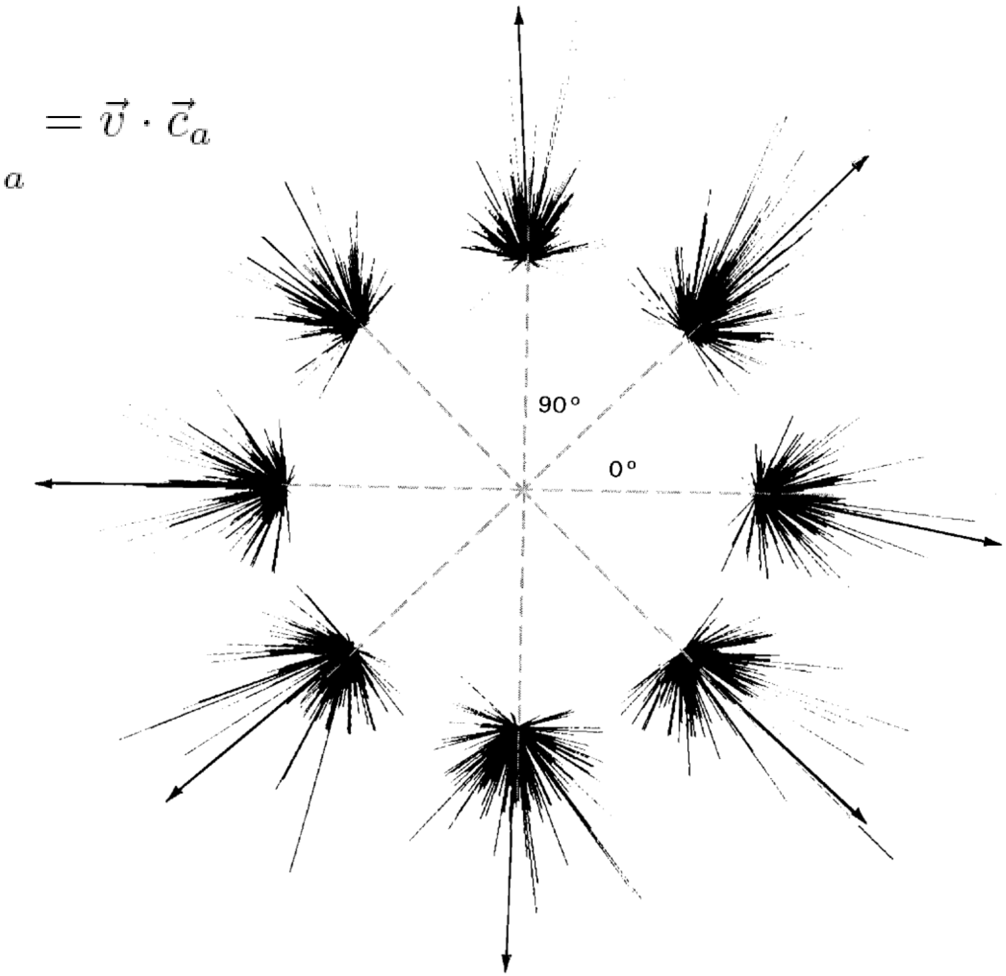

is parallel to the direction of arm movement

The population vector is neither general nor optimal.

"Optimal": Bayesian inference and MAP

Bayesian inference

By Bayes' law,

$$p[s|\mathbf{r}] = \frac{p[\mathbf{r}|s]p[s]}{p[\mathbf{r}]}$$

Introduce a cost function, $L(s,s_{Bayes})$; minimise mean cost.

$$\int ds\, L(s, s_{\text{bayes}})p[s|\mathbf{r}]$$

For least squares, $L(s,s_{Bayes}) = (s - s_{Bayes})^2$;
solution is the conditional mean.

$$s_{\text{bayes}} = \int ds\, p[s|\mathbf{r}]s$$

MAP and ML

MAP: s* which maximizes p[s|r]

ML: s* which maximizes p[r|s]

Difference is the role of the prior: differ by factor p[s]/p[r]

For cercal data:

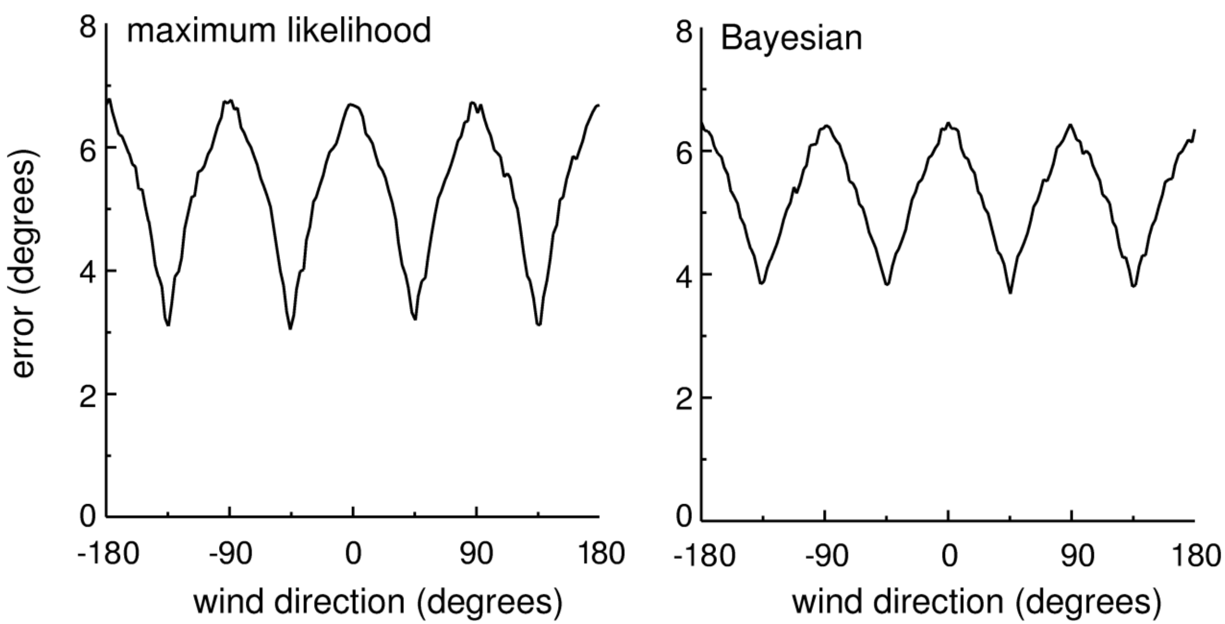

Decoding an arbitrary continuous stimulus

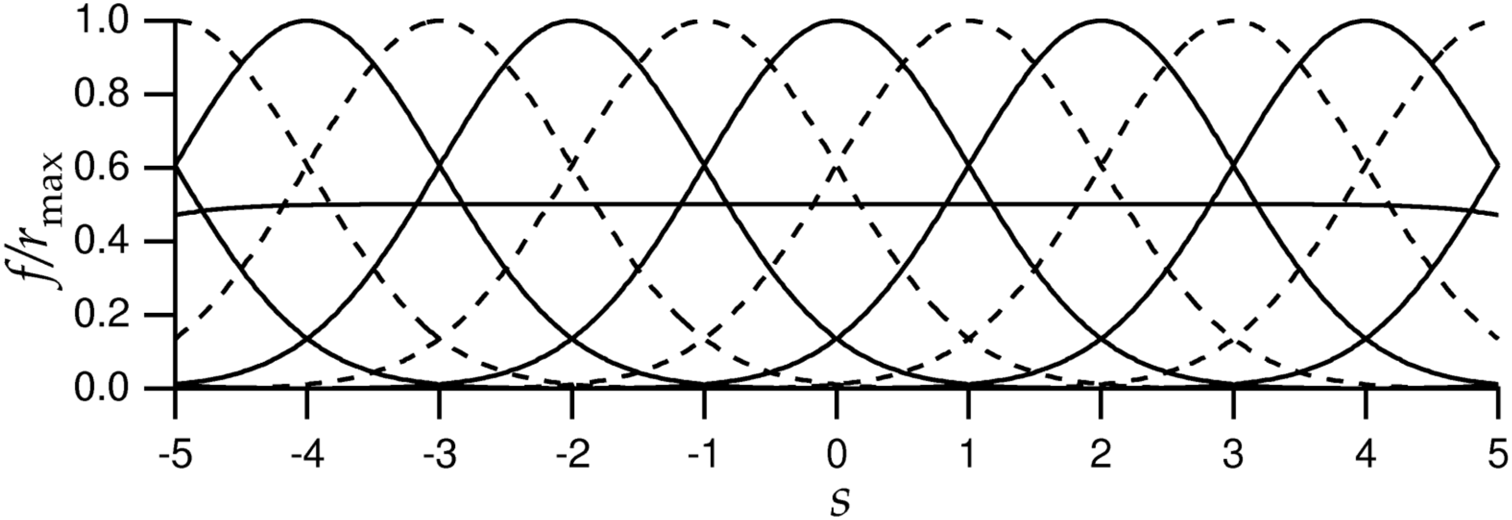

E.g. Gaussian tuning curves

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

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

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:

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

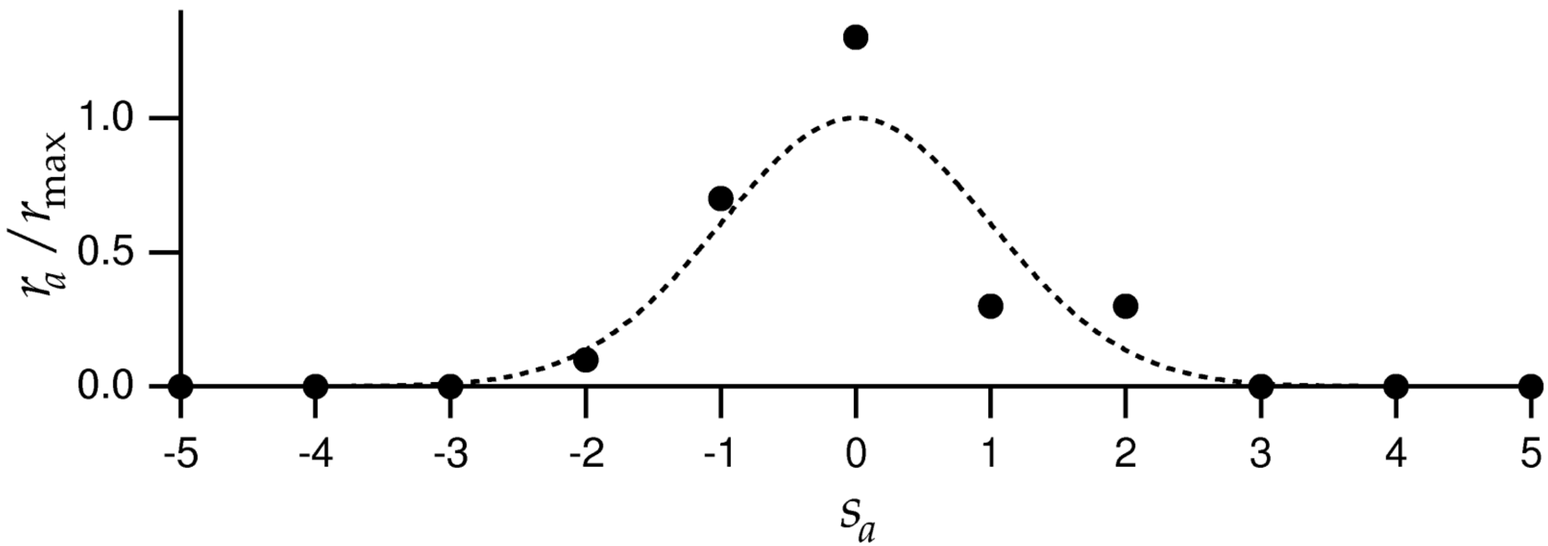

Population response of 11 cells with Gaussian tuning curves

Apply ML: maximise 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

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

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:

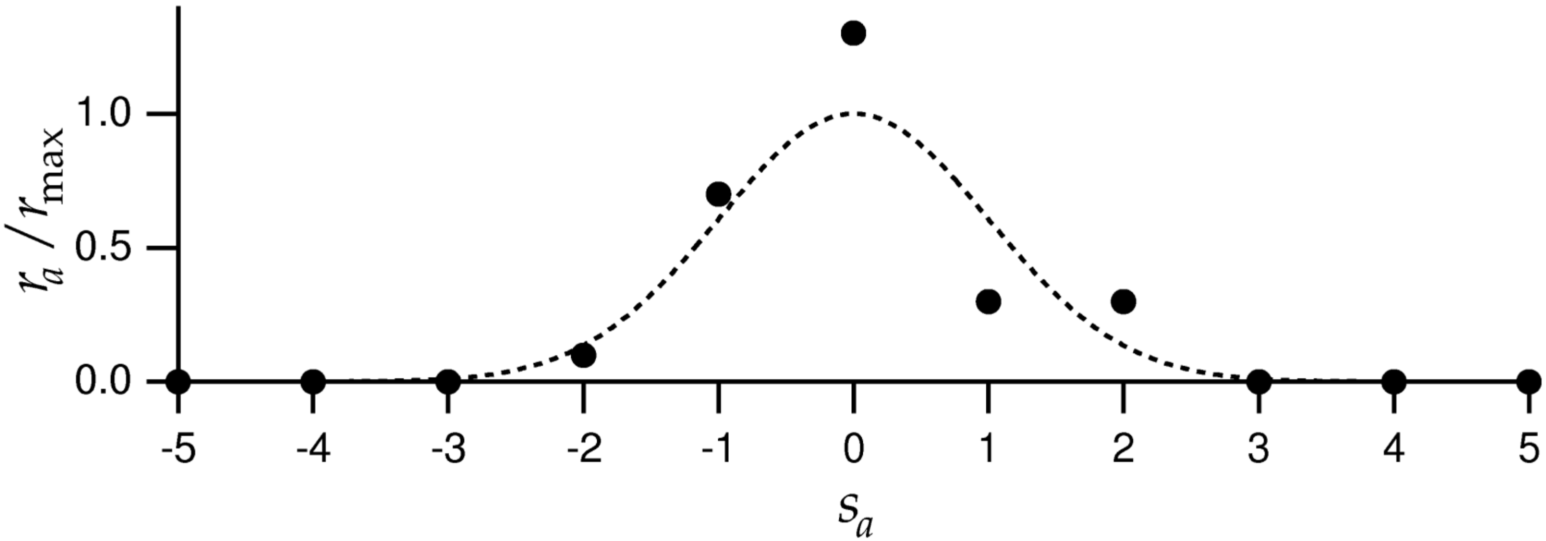

MAP:

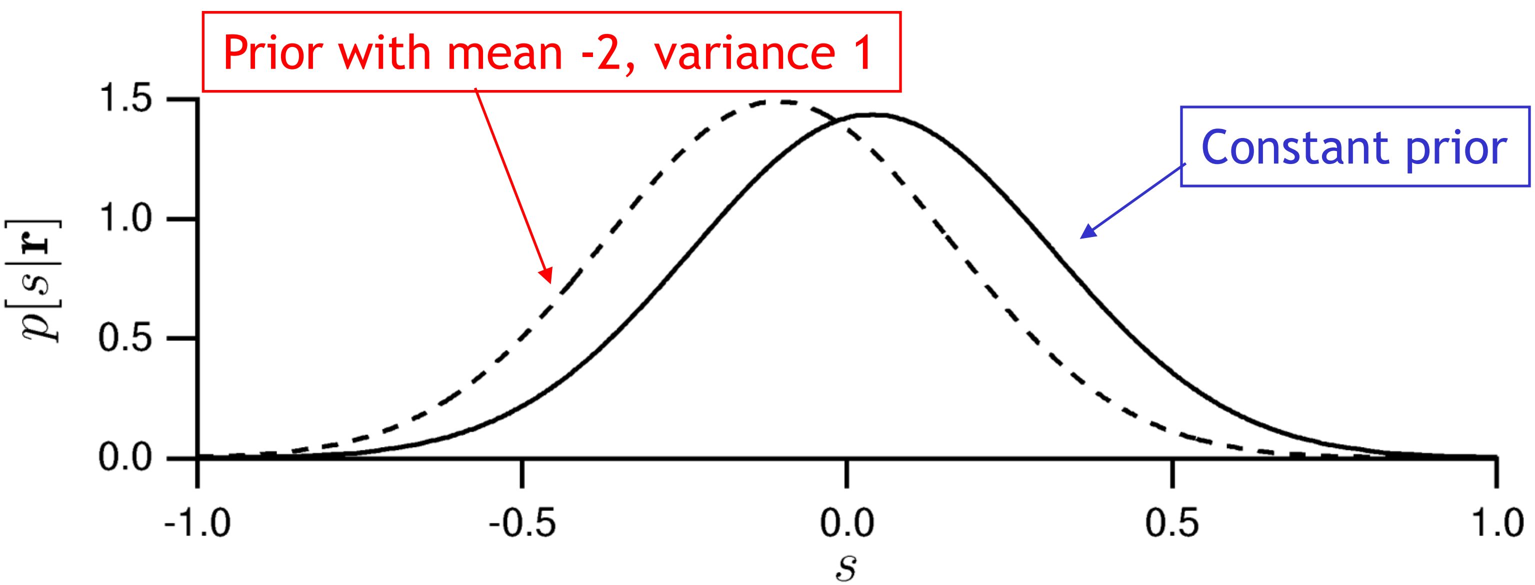

How good is our estimate?

For stimulus s, have estimated s_{est}

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

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

Mean square error:

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

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

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:

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

For the Gaussian tuning curves w/Poisson statistics:

$$I_{\mathrm{F}}(s) = \left\langle \left(\frac{d^2 \ln P[\mathbf{r}|s]}{ds^2} \right) \right\rangle = T \sum_{a=1}^{N} \langle r_a \rangle \left(\left(\frac{f_a'(s)}{f_a(s)} \right)^2 - \frac{f_a''(s)}{f_a(s)} \right)$$

Fisher information for Gaussian tuning curves

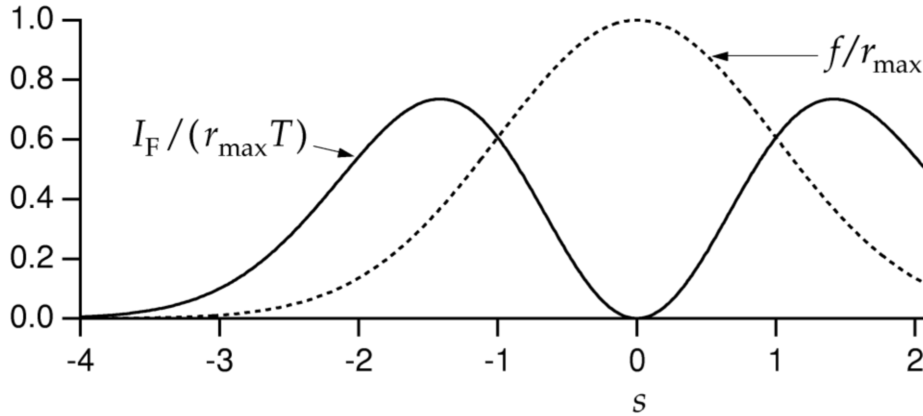

Quantifies local stimulus discriminability

Do narrow or broad tuning curves produce better encodings?

$$I_{\mathrm{F}} = T \sum_{a=1}^{N} \frac{r_{\mathrm{max}}(s - s_a)^2}{\sigma_r^4} \exp\left(-\frac{1}{2}\left(\frac{s - s_a}{\sigma_r}\right)^2\right)$$

Approximate: $I_{\mathrm{F}} \sim \frac{\sqrt{2\pi}\rho_s \sigma_r r_{\mathrm{max}} T}{\sigma_r^2}.$

Thus, $I_{\mathrm{F}} \sim 1/\sigma_r$ → Narrow tuning curves are better

But not in higher dimensions!

$$I_{\mathrm{F}} \sim (2\pi)^{D/2} D \rho_s \sigma_r^{D-2} r_{\mathrm{max}} T$$

Fisher information and discrimination

Recall d' = mean difference/standard deviation

Can also decode and discriminate using decoded values.

Trying to discriminate s and s+Δs:

Difference in estimate is Δs (unbiased)

variance in estimate is $1/I_F(s)$.

→

$$d' = \Delta s \sqrt{I_{\mathrm{F}}(s)}$$