
Population Coding

Peter Latham / Maneesh Sahani

Gatsby Computational Neuroscience Unit
University College London

Term 1, Autumn 2013

Common approach

The most common sort of questions asked of population codes:

- given assumed encoding functions, how well can we (or downstream areas) decode the encoded stimulus value?
- what encoding schemes would be optimal, in the sense of allowing decoders to estimate stimulus values as well as possible.

Before considering populations, we need to formulate some ideas about rate coding in the context of single cells.

Population codes

- High dimensionality (cells \times stimulus \times time).
 - usually limited to simple rate codes.
 - even prosthetic work assumes instantaneous (lagged) coding
- Limited empirical data
 - can record 10s - 100s of neurons.
 - population size more like 10^4 - 10^6 .
 - *theoretical* inferences, based on single-cell and aggregate (fMRI, LFP, optical) measurements .

Rate coding

In the rate coding context, we imagine that the firing rate of a cell r represents a single (possibly multidimensional) stimulus value s at any one time:

$$r = f(s).$$

Even if s and r are embedded in time-series we assume:

1. that coding is instantaneous (with a fixed lag),
2. that r (and therefore s) is constant over a short time Δ .

The actual number of spikes n produced in Δ is then taken to be distributed around $r\Delta$, often according to a Poisson distribution.

Tuning curves

The function $f(s)$ is known as a tuning curve.

Commonly assumed forms:

- Gaussian $r_0 + r_{\max} \exp \left[-\frac{1}{2\sigma^2} (x - x_{\text{pref}})^2 \right]$
- Cosine $r_0 + r_{\max} \cos(\theta - \theta_{\text{pref}})$
- Wrapped Gaussian $r_0 + r_{\max} \sum_n \exp \left[-\frac{1}{2\sigma^2} (\theta - \theta_{\text{pref}} - 2\pi n)^2 \right]$
- von Mises ("circular Gaussian") $r_0 + r_{\max} \exp [\kappa \cos(\theta - \theta_{\text{pref}})]$

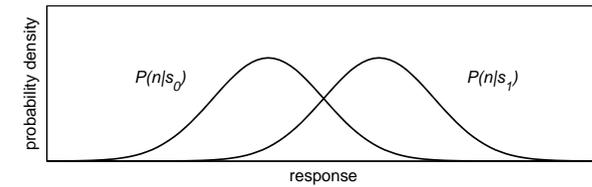
Measuring the performance of rate codes: Discrete choice

Suppose we want to make a binary choice based on firing rate:

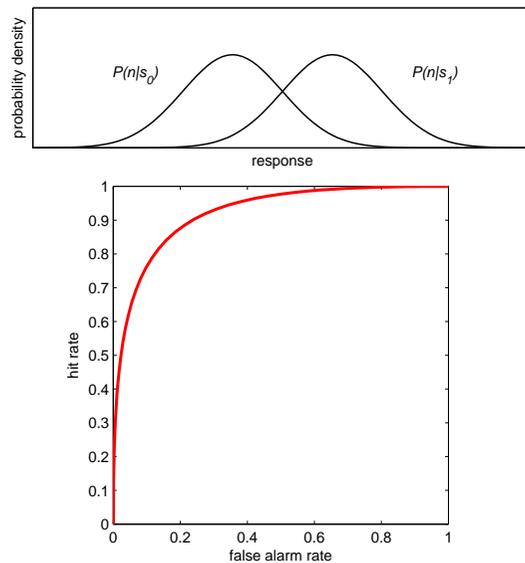
- present / absent (signal detection)
- up / down
- horizontal / vertical

Call one potential stimulus s_0 , the other s_1 .

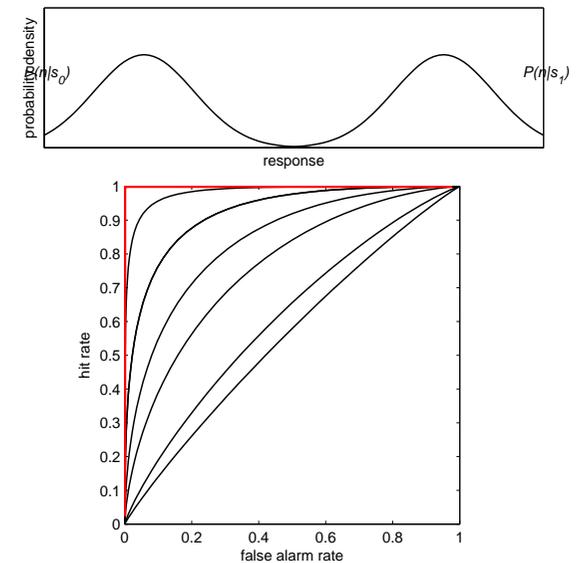
$P(n|s)$:



ROC curves



ROC curves



Summary measures

- area under the ROC curve
 - given $n_1 \sim P(n|s_1)$ and $n_0 \sim P(n|s_0)$, this equals $P(n_1 > n_0)$
- discriminability d'
 - for equal variance Gaussians $d' = \frac{\mu_1 - \mu_0}{\sigma}$.
 - for any threshold $d' = \Phi^{-1}(1 - \text{FA}) - \Phi^{-1}(1 - \text{HR})$ where Φ is a standard normal cdf.
 - definition unclear for non-Gaussian distributions.

Continuous estimation

Useful to consider a limit given $N \rightarrow \infty$ measurements n_i all generated by the same stimulus s^* .

The posterior over s is

$$\log P(s|\{n_i\}) = \sum_i \log P(n_i|s) + \log P(s) - \log Z(\{n_i\})$$

Taking $N \rightarrow \infty$ we have

$$\frac{1}{N} \log P(s|\{n_i\}) \rightarrow \langle \log P(n|s) \rangle_{n|s^*} + 0 - \log Z(s^*)$$

and so

$$\begin{aligned} P(s|\{n_i\}) &\rightarrow e^{N \langle \log P(n|s) \rangle_{n|s^*}} / Z \\ &= e^{-N \text{KL}[P(n|s^*)||P(n|s)]} / Z \end{aligned}$$

(Note: Z is being redefined as we go, but never depends on s)

Continuous estimation

Now consider a (one dimensional) stimulus that takes on continuous values (e.g. angle).

- contrast
- orientation
- motion direction
- movement speed

Suppose a neuron fires n spikes in response to stimulus s according to some distribution

$$P(n|f(s)\Delta)$$

Given an observation of n , how well can we estimate s ?

Continuous estimation

Now, Taylor expand the KL divergence in s around s^* :

$$\begin{aligned} \text{KL}[P(n|s^*)||P(n|s)] &= -\langle \log P(n|s) \rangle_{n|s^*} + \langle \log P(n|s^*) \rangle_{n|s^*} \\ &= -\langle \log P(n|s^*) \rangle_{n|s^*} - (s - s^*) \left\langle \frac{d \log P(n|s)}{ds} \Big|_{s^*} \right\rangle_{s^*} - \frac{1}{2} (s - s^*)^2 \left\langle \frac{d^2 \log P(n|s)}{ds^2} \Big|_{s^*} \right\rangle_{s^*} + \dots \\ &\quad + \langle \log P(n|s^*) \rangle_{n|s^*} \\ &= -\frac{1}{2} (s - s^*)^2 \left\langle \frac{d^2 \log P(n|s)}{ds^2} \Big|_{s^*} \right\rangle_{s^*} + \dots \\ &= \frac{1}{2} (s - s^*)^2 J(s^*) + \dots \end{aligned}$$

So in asymptopia, the posterior $\rightarrow \mathcal{N}(s^*, 1/J(s^*))$.

$J(s^*)$ is called the *Fisher Information*.

$$J(s^*) = -\left\langle \frac{d^2 \log P(n|s)}{ds^2} \Big|_{s^*} \right\rangle_{s^*} = \left\langle \left(\frac{d \log P(n|s)}{ds} \Big|_{s^*} \right)^2 \right\rangle_{s^*}$$

(You will show that these are identical in the homework.)

Cramér-Rao bound

The Fisher Information is important even outside the large data limit due to a deeper result that is due to Cramér and Rao.

This states that for any N , any unbiased estimator $\hat{s}(\{n_i\})$ of s will have the property that

$$\langle (\hat{s}(\{n_i\}) - s^*)^2 \rangle_{n_i | s^*} \geq \frac{1}{J(s^*)}.$$

Thus, Fisher Information gives a lower bound on the variance of any unbiased estimator. This is called the **Cramér-Rao bound**.

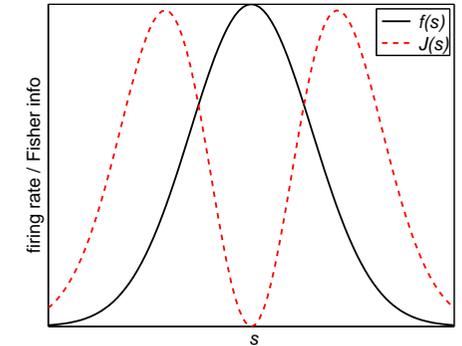
[For estimators with bias $b(s^*) = \langle \hat{s}(\{n_i\}) - s^* \rangle$ the bound is $\langle (\hat{s}(\{n_i\}) - s^*)^2 \rangle_{n_i | s^*} \geq \frac{(1+b'(s^*))^2}{J(s^*)} + b^2(s^*)$]

The Fisher Information will be our primary tool to quantify the performance of a population code.

Fisher Info and tuning curves

$$n = r\Delta + \text{noise}; \quad r = f(s) \Rightarrow$$

$$\begin{aligned} J(s^*) &= \left\langle \left(\frac{d}{ds} \Big|_{s^*} \log P(n|s) \right)^2 \right\rangle_{s^*} \\ &= \left\langle \left(\frac{d}{dr\Delta} \Big|_{f(s^*)} \log P(n|r\Delta) \Delta f'(s^*) \right)^2 \right\rangle_{s^*} \\ &= J_{\text{noise}}(r\Delta) \Delta^2 f'(s^*)^2 \end{aligned}$$



Fisher info for Poisson neurons

For Poisson neurons

$$P(n|r\Delta) = \frac{e^{-r\Delta}}{(r\Delta)^n} n!$$

so

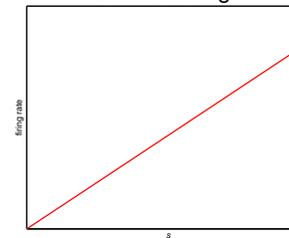
$$\begin{aligned} J_{\text{noise}}[r\Delta] &= \left\langle \left(\frac{d}{dr\Delta} \Big|_{r^*\Delta} \log P(n|r\Delta) \right)^2 \right\rangle_{s^*} \\ &= \left\langle \left(\frac{d}{dr\Delta} \Big|_{r^*\Delta} -r\Delta + n \log r\Delta - \log n! \right)^2 \right\rangle_{s^*} \\ &= \left\langle \left(-1 + n/r^*\Delta \right)^2 \right\rangle_{s^*} \\ &= \left\langle \frac{(n - r^*\Delta)^2}{(r^*\Delta)^2} \right\rangle_{s^*} \\ &= \frac{r^*\Delta}{(r^*\Delta)^2} = \frac{1}{r^*\Delta} \quad [\text{not surprising! } \widehat{r^*\Delta} = n \text{ and } \mathcal{V}ar[n] = r^*\Delta] \end{aligned}$$

and, referred back to the stimulus value:

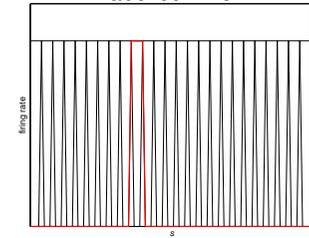
$$J[s^*] = f'(s^*)^2 \Delta / f(s^*)$$

Coding a continuous variable

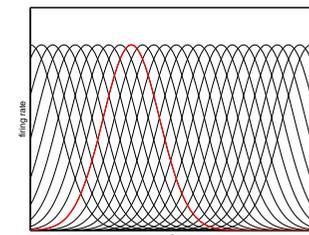
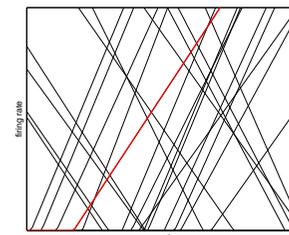
Scalar coding



Labelled Line



Distributed encoding



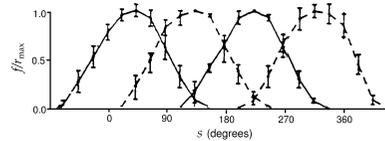
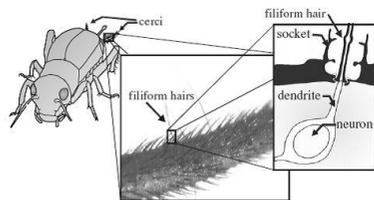
Coding a continuous variable

All of these schemes have been found in biological systems.

Issues:

1. redundancy and robustness (not scalar)
2. efficiency/resolution (not labelled line)
3. local computation (not scalar or scalar distributed)
4. multiple values (not scalar)

Cricket cercal system



$$r_a(s) = r_a^{\max} [\cos(\theta - \theta_a)]_+ = r_a^{\max} [\mathbf{c}_a^T \mathbf{v}]_+$$

$$\begin{aligned} \mathbf{c}_1^T \mathbf{c}_2 &= 0 \\ \mathbf{c}_3 &= -\mathbf{c}_1 \\ \mathbf{c}_4 &= -\mathbf{c}_2 \end{aligned}$$

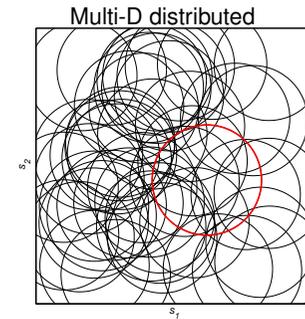
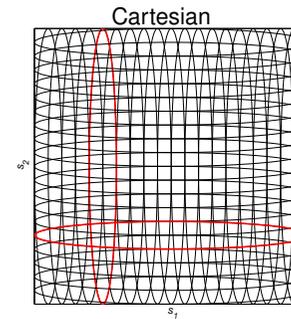
So, writing $\tilde{r}_a = r_a / r_a^{\max}$:

$$\begin{pmatrix} \tilde{r}_1 - \tilde{r}_3 \\ \tilde{r}_2 - \tilde{r}_4 \end{pmatrix} = \begin{pmatrix} \mathbf{c}_1^T \\ \mathbf{c}_2^T \end{pmatrix} \mathbf{v}$$

$$\mathbf{v} = (\mathbf{c}_1 \mathbf{c}_2) \begin{pmatrix} \tilde{r}_1 - \tilde{r}_3 \\ \tilde{r}_2 - \tilde{r}_4 \end{pmatrix} = \tilde{r}_1 \mathbf{c}_1 - \tilde{r}_3 \mathbf{c}_3 + \tilde{r}_2 \mathbf{c}_2 - \tilde{r}_4 \mathbf{c}_4 = \sum_a \tilde{r}_a \mathbf{c}_a$$

This is called **population vector decoding**.

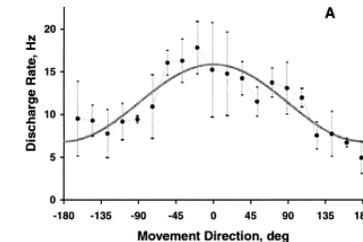
Coding in multiple dimensions



- efficient
- problems with multiple values

- represent multiple values
- may require more neurons

Motor cortex (simplified)



Cosine tuning, randomly distributed preferred directions.

In general, population vector decoding works for

- cosine tuning
- cartesian or dense (*tight*) directions

But:

- is it optimal?
- does it generalise? (Gaussian tuning curves)
- how accurate is it?

Bayesian decoding

Take $n_a \sim \text{Poisson}[f_a(s)\Delta]$, independently for different cells.

Then

$$P(\mathbf{n}|s) = \prod_a \frac{e^{-f_a(s)\Delta} (f_a(s)\Delta)^{n_a}}{n_a!}$$

and

$$\log P(\mathbf{n}|s) = -\sum_a f_a(s)\Delta + n_a \log(f_a(s)\Delta) - \log n_a! + \log P(s)$$

Assume $\sum_a f_a(s)$ is independent of s for a homogeneous population, and prior is flat.

$$\begin{aligned} \frac{d}{ds} \log P(\mathbf{n}|s) &= \frac{d}{ds} \sum_a n_a \log(f_a(s)\Delta) \\ &= \sum_a \frac{n_a}{f_a(s)\Delta} f'_a(s)\Delta \end{aligned}$$

Population Fisher Info

Fisher Informations for independent random variates add:

$$\begin{aligned} J_{\mathbf{n}}(s) &= \left\langle -\frac{d^2}{ds^2} \log P(\mathbf{n}|s) \right\rangle \\ &= \left\langle -\frac{d^2}{ds^2} \sum_a \log P(n_a|s) \right\rangle \\ &= \sum_a \left\langle -\frac{d^2}{ds^2} \log P(n_a|s) \right\rangle = \sum_a J_{n_a}(s). \\ &= \Delta \sum_a \frac{f'_a(s)^2}{f_a(s)} \quad [\text{for Poisson cells}] \end{aligned}$$

Bayesian decoding

Now, consider $f_a(s) = e^{-(s-s_a)^2/2\sigma^2}$, so $f'_a(s) = -(s-s_a)/\sigma^2 e^{-(s-s_a)^2/2\sigma^2}$ and set the derivative to 0:

$$\begin{aligned} \sum_a n_a (s-s_a)/\sigma^2 &= 0 \\ \hat{s}_{\text{MAP}} &= \frac{\sum_a n_a s_a}{\sum_a n_a} \end{aligned}$$

So the MAP estimate is a population average of preferred directions. Not exactly a population vector.

Optimal tuning properties

A considerable amount of work has been done in recent years on finding optimal properties of tuning curves for rate-based population codes. Here, we reproduce one such argument (from Zhang and Sejnowski, 1999).

Consider a population of cells that codes the value of a D dimensional stimulus, \mathbf{s} . Let the a th cell emit r spikes in an interval τ with probability distribution that is conditionally independent of the other cells (given \mathbf{s}) and has the form

$$P_a(r | \mathbf{s}, \tau) = S(r, f^a(\mathbf{s}), \tau).$$

The tuning curve of the a th cell, $f^a(\mathbf{s})$, has the form

$$f^a(\mathbf{s}) = F \cdot \phi((\xi^a)^2); \quad (\xi^a)^2 = \sum_i^D (\xi_i^a)^2; \quad \xi_i^a = \frac{s_i - c_i^a}{\sigma},$$

where F is a maximal rate and the function ϕ is monotonically decreasing. The parameters \mathbf{c}^a and σ give the centre of the a th tuning curve and the (common) width.

Optimal tuning properties

Now, the (ij) th term in the FI matrix for the a th cell is (by definition)

$$J_{ij}^a(\mathbf{s}) = \mathbb{E} \left[\frac{\partial}{\partial s_i} \log P^a(r | \mathbf{s}, \tau) \frac{\partial}{\partial s_j} \log P^a(r | \mathbf{s}, \tau) \right]$$

Applying the chain rule repeatedly, we find that

$$\begin{aligned} \frac{\partial}{\partial s_i} \log P^a(r | \mathbf{s}, \tau) &= \frac{1}{S(r, f^a(\mathbf{s}), \tau)} \frac{\partial}{\partial s_i} S(r, f^a(\mathbf{s}), \tau) \\ &= \frac{S^{(2)}(r, f^a(\mathbf{s}), \tau)}{S(r, f^a(\mathbf{s}), \tau)} \frac{\partial}{\partial s_i} f^a(\mathbf{s}) \end{aligned}$$

(where $S^{(2)}$ indicates differentiation with respect to the second argument)

$$\begin{aligned} &= \frac{S^{(2)}(r, f^a(\mathbf{s}), \tau)}{S(r, f^a(\mathbf{s}), \tau)} F \phi'((\xi^a)^2) \frac{\partial}{\partial s_i} \sum_i^D (\xi_i^a)^2 \\ &= \frac{S^{(2)}(r, f^a(\mathbf{s}), \tau)}{S(r, f^a(\mathbf{s}), \tau)} F \phi'((\xi^a)^2) \frac{2(s_i - c_i^a)}{(\sigma_i^a)^2} \end{aligned}$$

Optimal tuning properties

We assumed neurons were independent \Rightarrow Fisher information adds. Approximate by integral over the tuning curve centres, assuming uniform density η of neurons.

$$\begin{aligned} J_{ij}(\mathbf{s}) &= \sum J_{ij}^a(\mathbf{s}) \\ &\approx \int_{-\infty}^{+\infty} dc_1^a \cdots \int_{-\infty}^{+\infty} dc_D^a \eta J_{ij}^a(\mathbf{s}) \\ &= \int_{-\infty}^{+\infty} dc_1^a \cdots \int_{-\infty}^{+\infty} dc_D^a \eta A_\phi((\xi^a)^2, F, \tau) \frac{(s_i - c_i^a)(s_j - c_j^a)}{\sigma^4} \end{aligned}$$

Change variables: $c_i^a \rightarrow \xi_i^a$

$$\begin{aligned} &= \int_{-\infty}^{+\infty} \sigma d\xi_1^a \cdots \int_{-\infty}^{+\infty} \sigma d\xi_D^a \eta A_\phi((\xi^a)^2, F, \tau) \frac{\xi_i^a \xi_j^a}{\sigma^2} \\ &= \frac{\sigma^D}{\sigma^2} \eta \int_{-\infty}^{+\infty} d\xi_1^a \cdots \int_{-\infty}^{+\infty} d\xi_D^a A_\phi((\xi^a)^2, F, \tau) \xi_i^a \xi_j^a \end{aligned}$$

Now, if $i \neq j$, integral is odd in both ξ_i^a and ξ_j^a , and thus vanishes. If $i = j$, then the integral has some value $D \cdot K_\phi(F, \tau, D)$, independent of σ . Thus,

$$J_{ii} = \sigma^{D-2} \eta D K_\phi(F, \tau, D)$$

and the total Fisher information is proportional to σ^{D-2} .

Optimal tuning properties

So,

$$\begin{aligned} J_{ij}^a(\mathbf{s}) &= \mathbb{E} \left[\left(\frac{S^{(2)}(r, f^a(\mathbf{s}), \tau)}{S(r, f^a(\mathbf{s}), \tau)} \right)^2 \right] 4F^2 (\phi'((\xi^a)^2))^2 \frac{(s_i - c_i^a)(s_j - c_j^a)}{\sigma^4} \\ &= A_\phi((\xi^a)^2, F, \tau) \frac{(s_i - c_i^a)(s_j - c_j^a)}{\sigma^4} \end{aligned}$$

where the function A_ϕ does not depend explicitly on σ .

Optimal tuning properties

Thus optimal tuning width depends on the stimulus dimension.

- $D = 1$
 $\Rightarrow \sigma \rightarrow 0$ (although a lower limit is encountered when the tuning width falls below the inter-cell spacing)
- $D = 2$
 $\Rightarrow J$ independent of σ .
- $D > 2$
 $\Rightarrow \sigma \rightarrow \infty$ (actual limit set by valid stimuli).

More ...

- Correlated noise
- Extended s (feature maps etc.)
- Uncertainty