Standard model of population coding

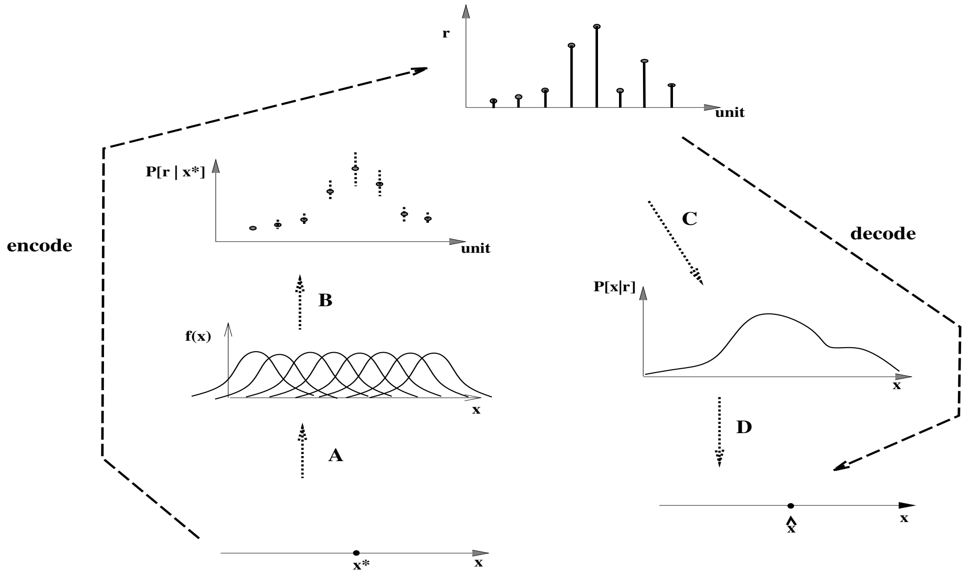

tuning curves $f_i(\mathbf{x})$ describe expected firing rate of cell i as function of stimulus value $\mathbf{x}$

encoding: value encoded in activities $\mathbf{r}$

$$
\begin{aligned}
r_i &= f_i(\mathbf{x}) + \epsilon_i \\
P[\mathbf{r}|\mathbf{x}] &\propto \prod_i e^{-f_i(\mathbf{x})} (f_i(\mathbf{x}))^{r_i}
\end{aligned}
$$

decoding: what does activities $\mathbf{r}$ imply about $\mathbf{x}$?

$$
\begin{aligned}
P[\mathbf{x}|\mathbf{r}] &= \frac{P[\mathbf{r}|\mathbf{x}] P[\mathbf{x}]}{P[\mathbf{r}]} \\
\mathcal{I}[\mathbf{x}, \mathbf{r}] &= \mathcal{H}[\mathbf{x}] - \mathcal{H}[\mathbf{x}|\mathbf{r}]
\end{aligned}
$$

Open issues in population coding

- spiking model: how do population codes arise from afferent connections?

- noise correlations: existence and effect

- relevant stimulus variables
  1. higher areas – what are underlying dimensions?
  2. joint coding – which variables coded together, and why?

Extending standard model: Representing uncertainty

standard model limited, only considers:

1. uncertainty due to noisy activities
2. information about single value of relevant variable

main claim: ***populations represent more than single value***

(A). uncertainty: not only numerous noise forms, but inherent in computation

- non-deterministic environments
- ignorance about relevant properties
- ill-posed: non-invertible projection to relatively low-dim sensor

(B). multiple values

stimulus

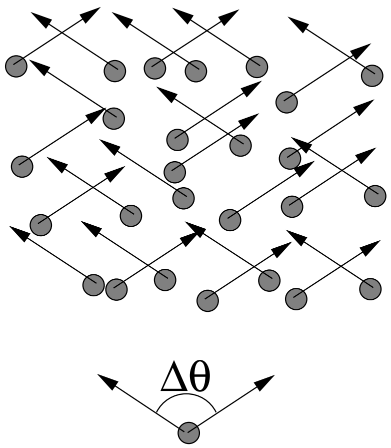

MT response

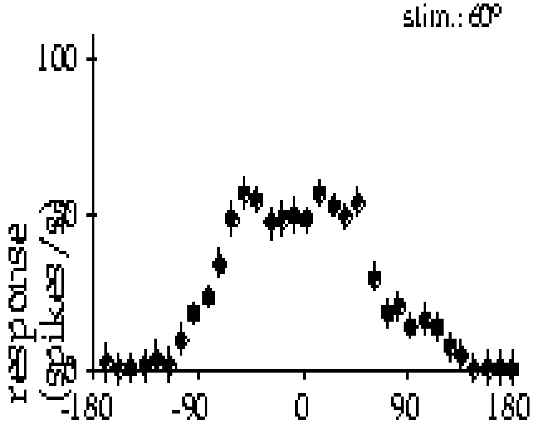

subject's response

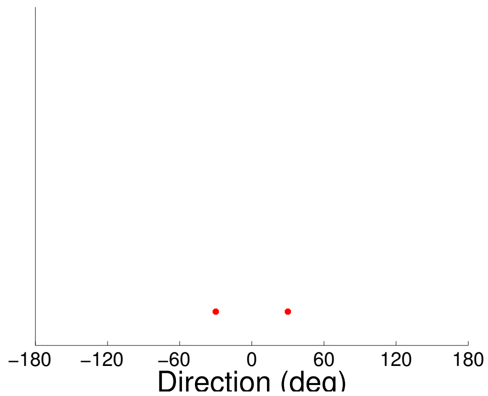

standard model inadequate: if tuning functions are Gaussian ($f_i(\mathbf{x}) \propto \exp(-(\mathbf{x} - \mathbf{x}_i)^2/2\sigma^2)$), evenly and densely packed ($|\mathbf{x}_i - \mathbf{x}_{i=-1}| << \sigma$), and noise is Poisson:

$$P[\mathbf{x}|\mathbf{r}] = \mathcal{N}\left[\frac{\sum_i r_i \mathbf{x}_i}{\sum_i r_i}, \frac{\sigma^2}{\sum_i r_i}\right]$$

corollary of main claim:

populations must propagate probabilities of underlying quantities in order to do proper inference

Psychophysical and/or neurobiological evidence?

- no experiment (yet) shows how degree of uncertainty is reflected in population responses
- psychophysical studies demonstrate perception interpretable as statistical inference problem (e.g., speed perception varies with contrast, cue combination with varying reliabilities)
- neurobiological studies show that activities of neurons representing alternative hypotheses directly reflect hypothesis likelihood (e.g., 2-AFC experiments in LIP, superior colliculus)

Neural representation of uncertainty

How can activity $\mathbf{a} = (a_1, .., a_n)$ of n neurons encode $P[\mathbf{x}]$?

1. direct codes: neuron associated with discrete proposition (motion *left*)
   $\Rightarrow$ activity strength directly, or stochastically reflects certainty

   $$a_i \sim P[\mathbf{x}_i]$$

   (e.g., Gold & Shadlen, 02; Lee & Mumford, 03; Rao, 04)

2. likelihood codes: neuron represents data likelihood given hypothesized value

   $$a_i \sim \log P[\mathbf{d}|\mathbf{x}_i]$$

   Weiss & Fleet (02) showed that motion energy filters can be interpreted as representing a form of probability distribution

Distributional Population Codes

3. convolution codes: neuron reflects true uncertainty in quantity represented according to its tuning function

$$a_i \sim \int_{\mathbf{x}} d\mathbf{x} f_i(\mathbf{x}) P[\mathbf{x}]$$

matches regular population code when encoded distribution contains single value

decoding methods:

- basis functions

$$\log \mathcal{P}[\mathbf{x}|\{a_i\}] \sim \sum_i a_i f_i(\mathbf{x})$$

- approximate ML in distributions over θ

$$\log \mathcal{P}[\mathbf{x}|\{a_i\}] \sim \sum_i a_i \log [f_i(\mathbf{x}) * \mathcal{P}(\mathbf{x}|\mathbf{a})] - \alpha g [\mathcal{P}(\mathbf{x}|\mathbf{a})]$$

Testing DPC on Transparent Motion Data

MT response ⇒

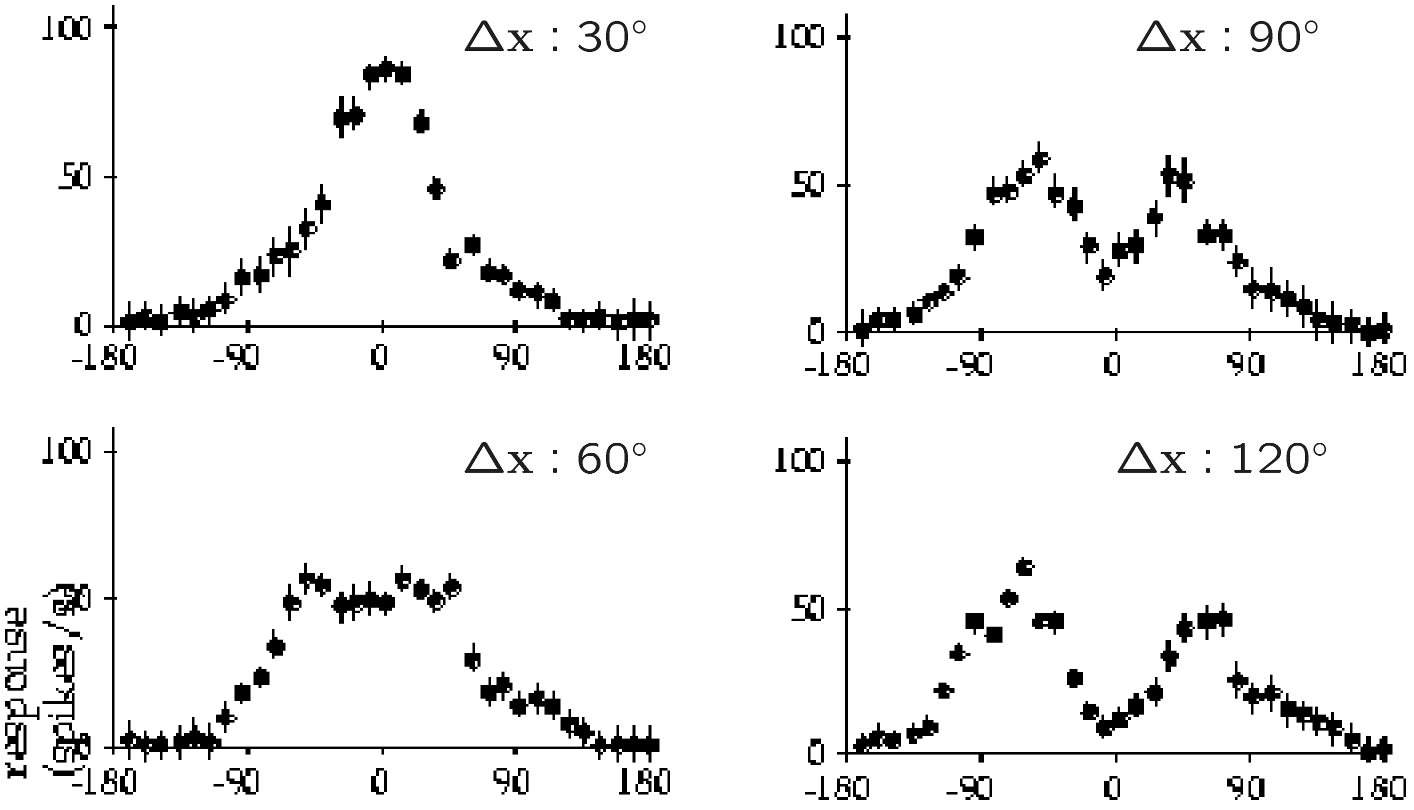

Bimodal response patterns: DPC matches generating distribution

Unimodal patterns fit well, until $\Delta x \sim 10^{\circ}$ (captures subjects' uncertainty)

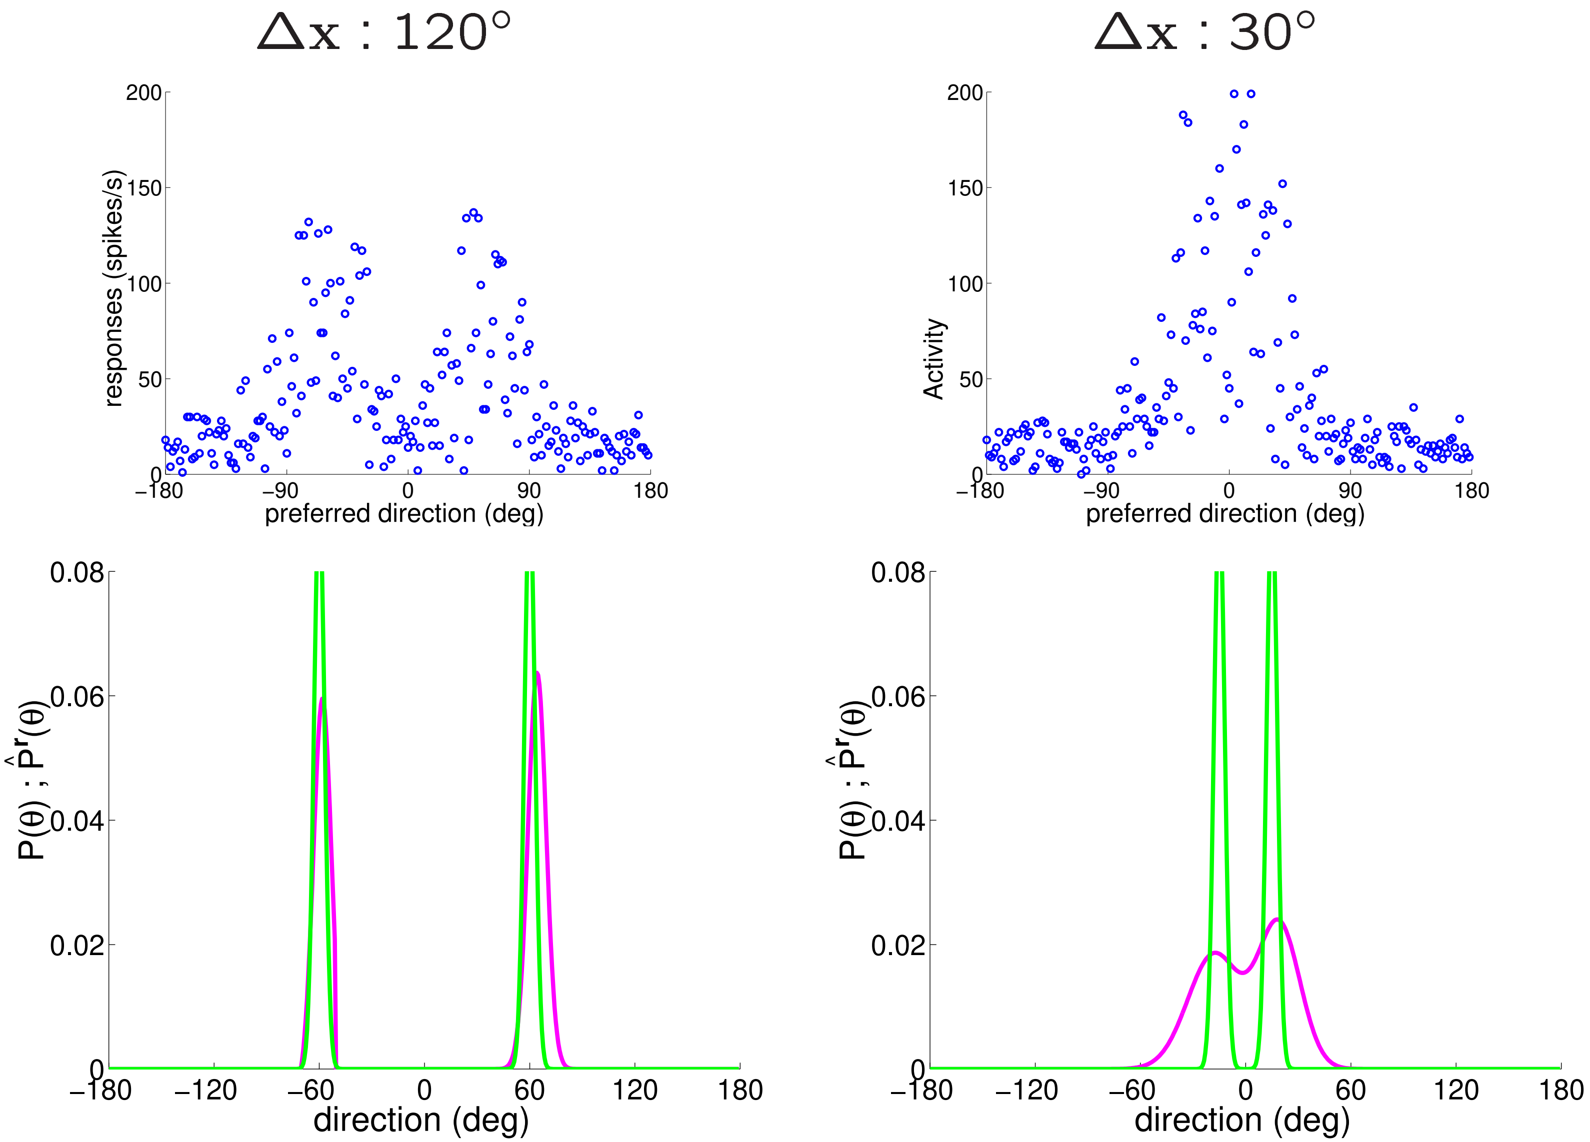

Δx : 120°
Δx : 30°
responses (spikes/s)
Activity
preferred direction (deg)
preferred direction (deg)
$P(\theta)\ ;\ \hat{P}^{\mathbf{r}}(\theta)$
$P(\theta)\ ;\ \hat{P}^{\mathbf{r}}(\theta)$
direction (deg)
direction (deg)

Current Directions I: Decoding motion energy filters

How much information is contained in motion sensitive neurons about $P[V]$?

1. random dot kinematogram stimuli, varying in:
   - spatial distribution of dots
   - $P[V]$
   - number of dots N
   - contrast
2. simulated cell responses based on motion energy (ME) filters

Initial study considering how much information ME filters contain about direction, where each $P[V]$ a Dirac: $\theta = \bar{\theta} \pm \delta\theta$; only noise due to dot distribution

Next: Is ME response to Gaussian $P[V]$ consistent with gain or convolution encoding predictions?

Eventually: general $P[V]$ – compare decoding schemes (basis functions, ML)

Current Directions II: Spike-based, dynamic DPC

How can population codes represent temporally changing distributions over variables, using spikes?

Starting point is basis function DPC decoder:

$$\log \mathcal{P}[\mathbf{x}|\{a_i\}] \sim \sum_i a_i f_i(\mathbf{x})$$

Extend to include time (Hinton & Brown, 00):

$$\log \mathcal{P}_t[\mathbf{x}|\{a_i(\tau)\}] \sim \sum_i f_i(\mathbf{x}) \sum_{\tau=0}^{\infty} a_i(\tau) g(t-\tau)$$

where $g(b) = [b \geq 0]e^{-b}$

if $\{a_i\}$ is binary vector representing spike times, then slow spiking neurons can rapidly communicate arbitrary distributions over real values

difficult problem: learning encoding scheme (which cell spikes when) in causal manner

approach: given $\{f_i(\mathbf{x})\}, g(b)$, learn encoding model for some known dynamic distribution: $a_i(t) = h(\mathcal{P}_t[\mathbf{x}], \mathbf{a}(t-1))$