

# the statistical structure of noise in large neural populations

neil rabinowitz


robbe goris

eero simoncelli





sept 2015

# a computational model of the brain



repeated  
stimuli



same  
brain




variable  
responses





$f_{\text{brain}}$  is stochastic




# a computational model of a single neuron



repeated  
stimuli



same  
neuron





variable  
responses






$f_{\text{neuron}}$  is stochastic

repeated  
stimuli



same  
neuron



variable  
responses



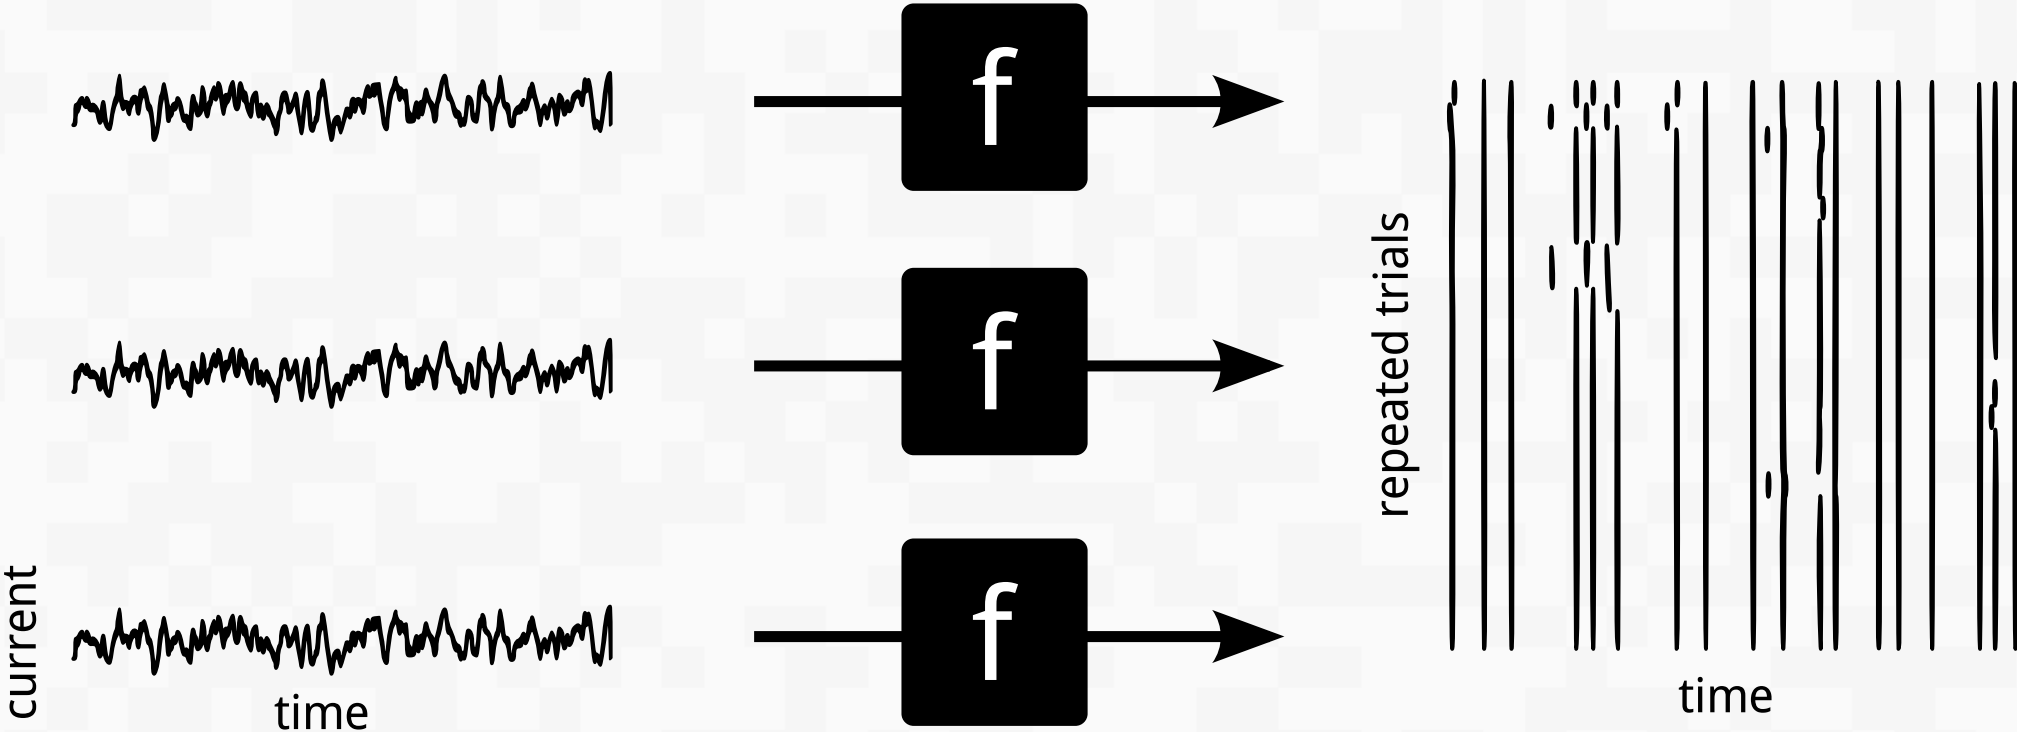
$f_{\text{neuron}}$  is stochastic

# take the neuron out of the brain...

repeated stimuli

same neuron


reliable responses



[Mainen & Sejnowski, 1995]

$f_{\text{neuron}}$  is near-deterministic


# there are all these other processes going on




attention




arousal




network state



emotion




fear



depth of  
anaesthesia






reward





neuromodulators




are these noise?



they are  
latent processes




# are these noise?



arousal




they are  
complex and structured



depth of  
anaesthesia



# the big question of this talk



what is  $p(\text{brain} \mid \text{stimuli})$  ?

# the challenges

(1) understand this density and its structure

- limited measurements (but growing)
- appropriate statistical models
- tied in to physiological processes  
(scientifically interpretable components)

(2) understand how it relates to computation

# this talk

a brief history of models of  $p(\text{brain} | \text{stimuli})$

structured latent variable models


a large neural population

feature or bug?

# single neuron model

[last ~50 years]

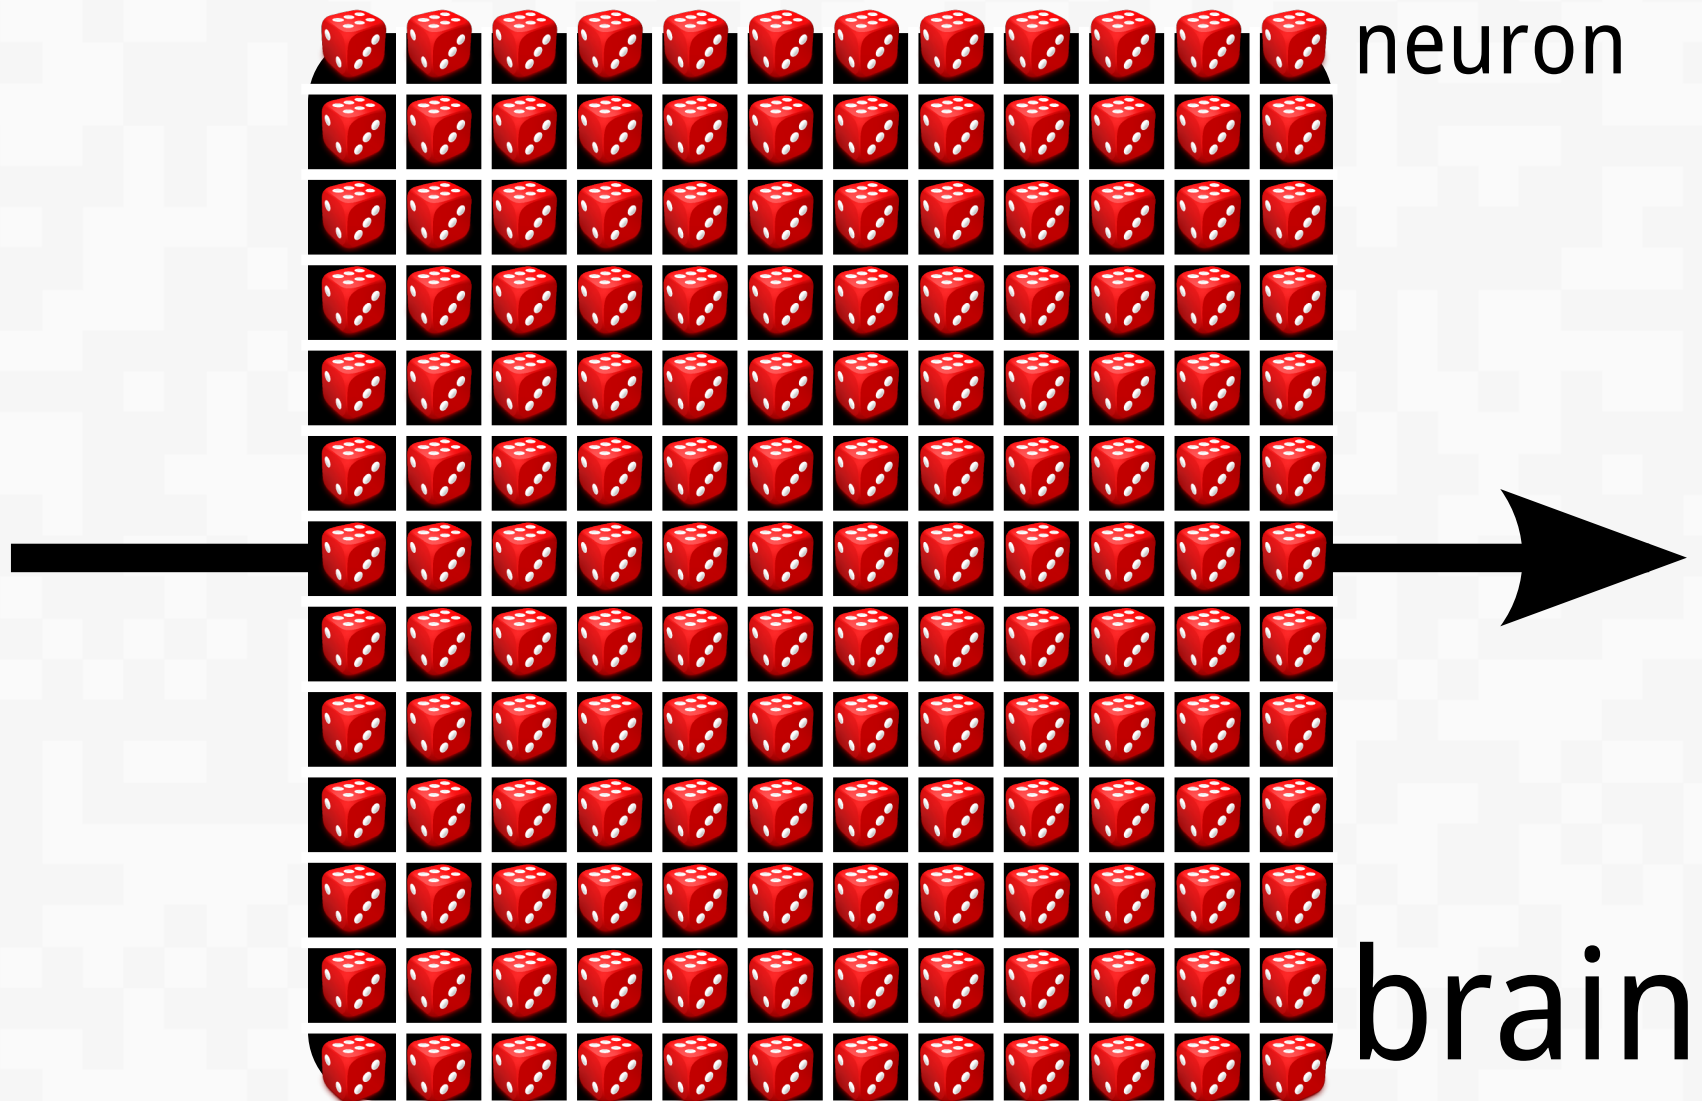
we know nothing about the latent processes



# single neuron model

[last ~50 years]


we know nothing about the latent processes



# single neuron model

[last ~50 years]

we know nothing about the latent processes




$$p(\text{brain} \mid \text{stimuli}) \approx \prod_{\text{neurons } n} \prod_{\text{time } t} p(\mathbf{r}_{n,t} \mid \text{stimulus}(t))$$



# single neuron model

[last ~50 years]

we know nothing about the latent processes




spike counts  $r$  within a window,


$$r_t \mid \text{stim} \sim \text{Poiss}( f(\text{stim}_t) \cdot \Delta t )$$

# poisson: the marginal statistics are wrong

one stimulus




many stimuli



in the style of  
[Goris et al, 2014]

# how to fix

[Goris et al, 2014]




$$g_t \sim p(g)$$

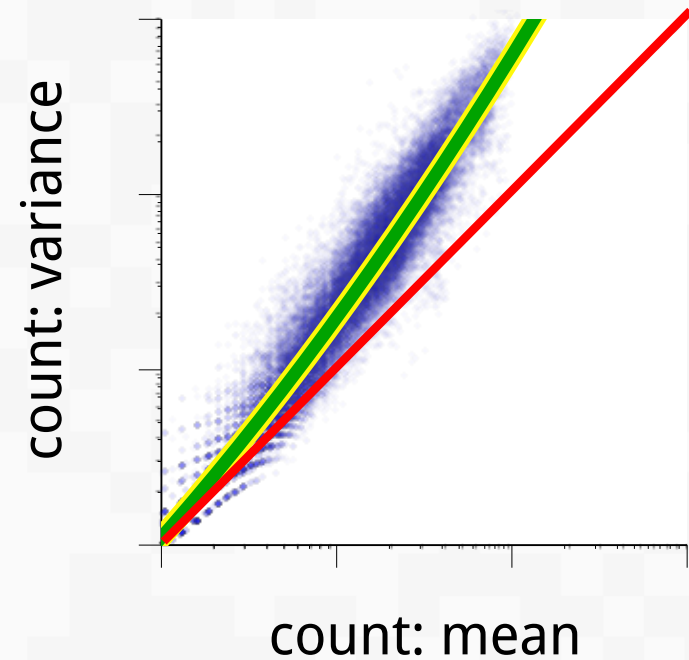
$$r_t \mid \text{stim} \sim \text{Poiss}( g_t \cdot f(\text{stim}_t) \cdot \Delta t )$$

# it gets the marginal statistics right

one stimulus




many stimuli




in the style of  
[Goris et al, 2014]

# temporal independence




over a long timespan


cell 1



cell 2



cell 3

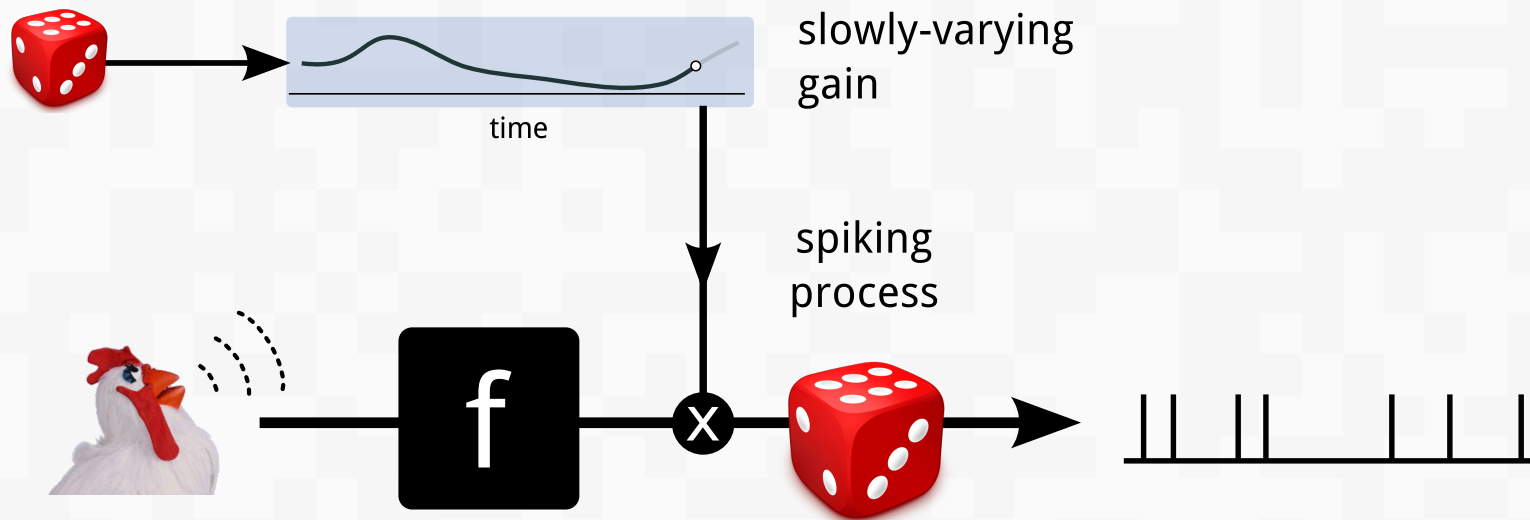


30 mins

presence of slow drifts

# slow gain fluctuations

[Paninski et al, 2010;  
Rabinowitz et al, in prep]



$$\log(\vec{g}) \sim \mathcal{GP}(\vec{0}, K(\varphi))$$

$$r_t \mid \text{stim} \sim \text{Pois}(g_t \cdot f(\text{stim}_t; \theta))$$




# results

cell 1




cell 2

cell 3

spike  
counts




slow gain  
modulators



  
30 mins



# one neuron: factorial model




we can measure each neuron  
in isolation

spike counts  $r_i$  for neuron  $i$

$$p(r_1, \dots, r_N \mid \text{stim}) \\ = \prod_i q(r_i \mid \text{stim})$$

## two neurons



what about pairwise statistics?


"noise correlations"

$r \sim 0.05$  to  $0.4$

these correlations  
can change!

[Bach & Kruger 1986; Zohary et al 1994; Gawne & Richmond 1993; Gawne et al 1996; Bair et al 2001, Kohn & Smith 2005; Gutnisky & Dragoi 2008; Smith & Kohn, 2008; Cohen & Newsome 2008; Mitchell et al 2009; Cohen & Maunsell 2009; Ecker et al 2010; Hansen et al 2012; Smith et al 2013; Smith & Sommer, 2013; Herrero et al 2013]


# whole population?



what about pairwise statistics?

$$p(r_1, \dots, r_N \mid \text{stim}) \\ = \frac{1}{Z} \prod_{i, j} \phi(r_i, r_j \mid \text{stim})$$

# problems



(1) the anatomy

(2) can't infer connectivity

(3) scales poorly

(4) how to interpret

$$p(r_1, \dots, r_N \mid \text{stim}) \\ = \frac{1}{Z} \prod_{i,j} \phi(r_i, r_j \mid \text{stim})$$

# this talk

a brief history of models of  $p(\text{brain} | \text{stimuli})$


structured latent variable models

a large neural population

feature or bug?

# hypothesis: the noise is structured

[Kulkarni & Paninski 2007; Paninski et al 2010; Macke et al 2011; Vidne et al 2012; Pfau et al 2013; Archer et al 2014; Rabinowitz et al submitted]



shared  
latent factors

a low-dimensional  
shared "modulator",  $\mathbf{m}$

$$p(r_1, \dots, r_N, \mathbf{m} \mid \text{stim})$$
$$= p(\mathbf{m}) \prod_i p(r_i \mid \mathbf{m}, \text{stim})$$

# this talk


a brief history of models of  $p(\text{brain} | \text{stimuli})$

structured latent variable models

a large neural population

feature or bug?

# a visual experiment



100-150 neurons

macaque V4

both hemispheres

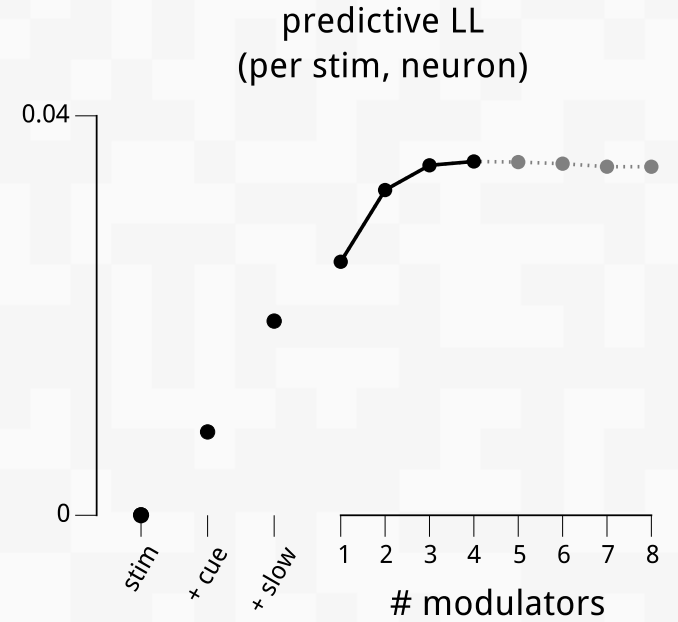
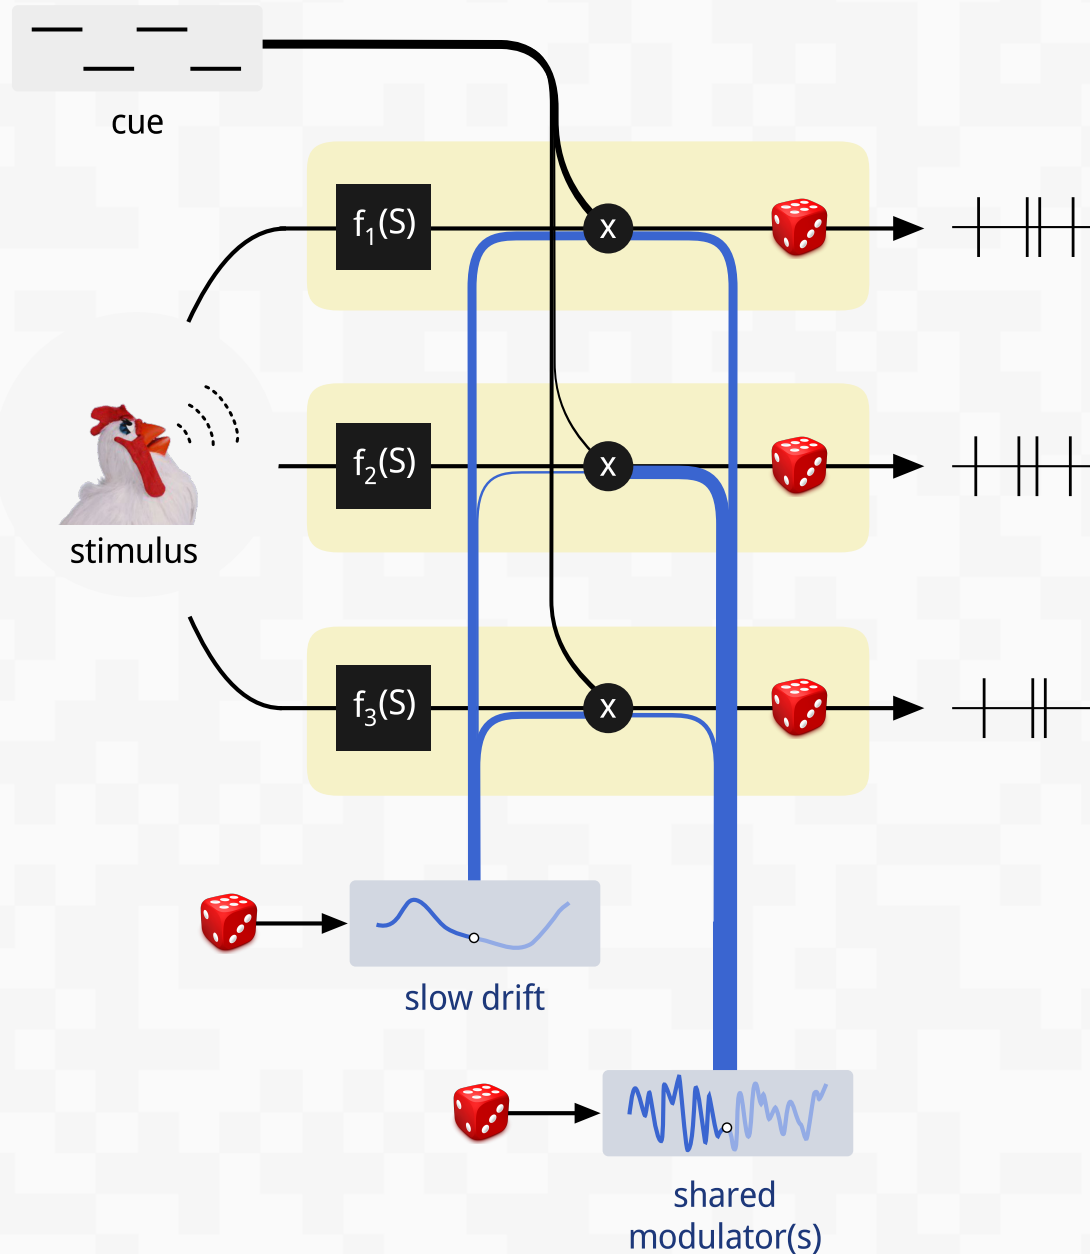
behavioural task

spatial attention

[Cohen & Maunsell, 2009]



# functional model



$$r_{t,n} \mid \text{stim} \sim$$

$$\text{Pois} [ f_n(\text{stim}_t)$$

$$\cdot \exp( u_n \cdot \text{cue}_t$$

$$+ v_n \cdot \text{drift}_t$$

$$+ \sum_{k=1}^K w_n^{(k)} \cdot m_t^{(k)} ) ]$$