Top-down signals in cortex & sampling

Ralf M. Haefner
Brain & Cognitive Sciences, University of Rochester (NY)

Full internal model for natural vision unknown

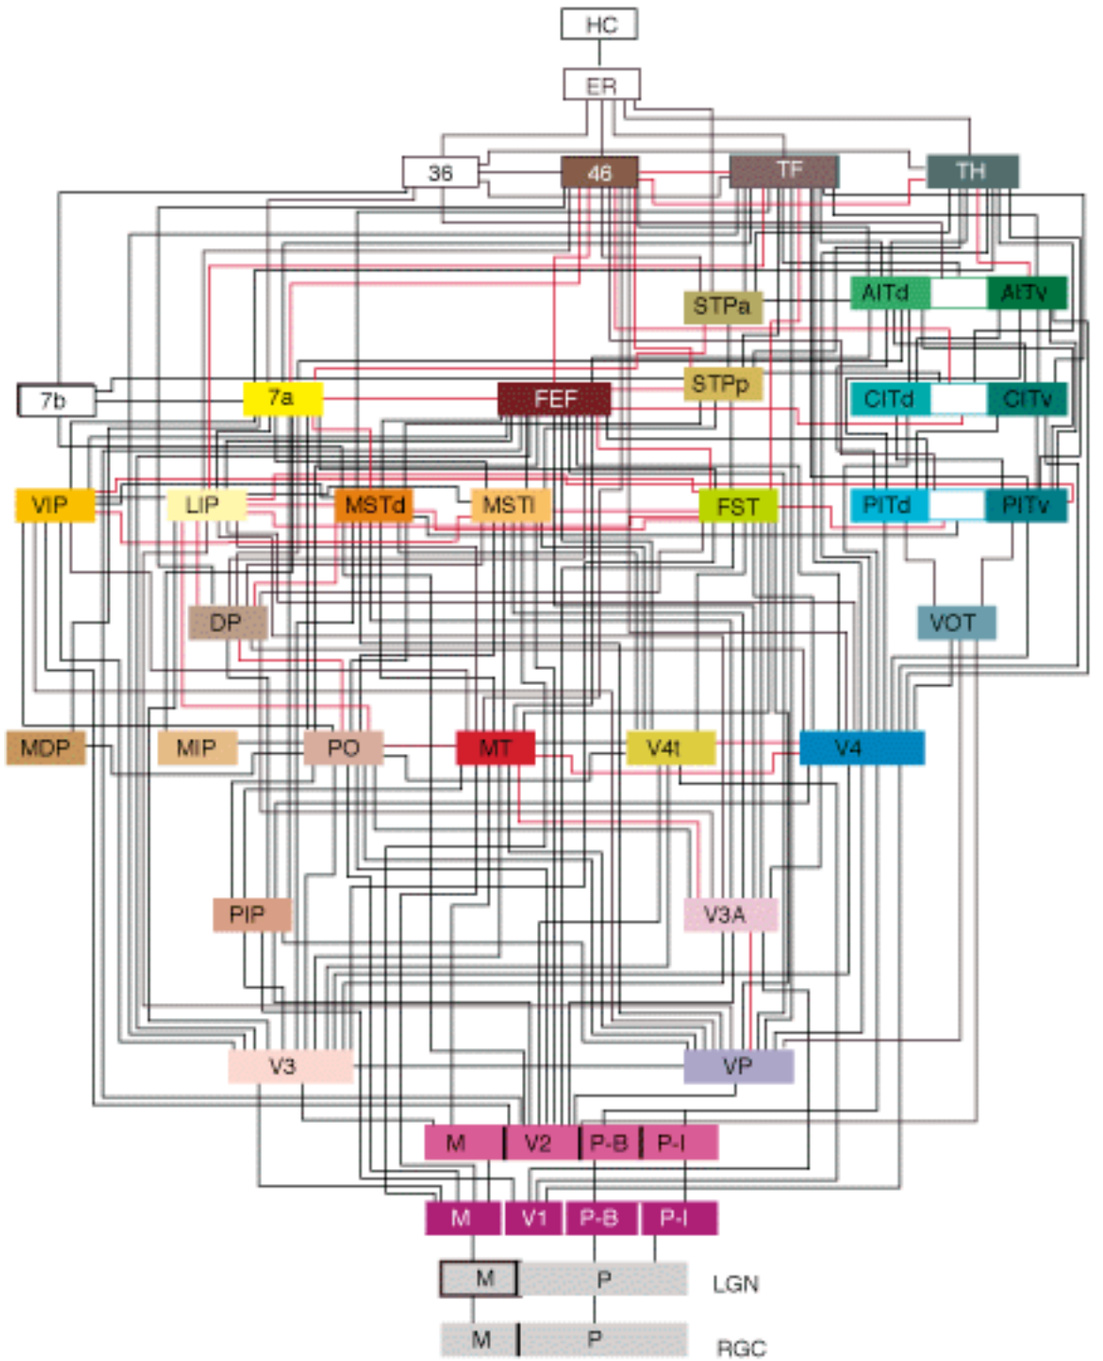

Fellemann & van Essen 1991

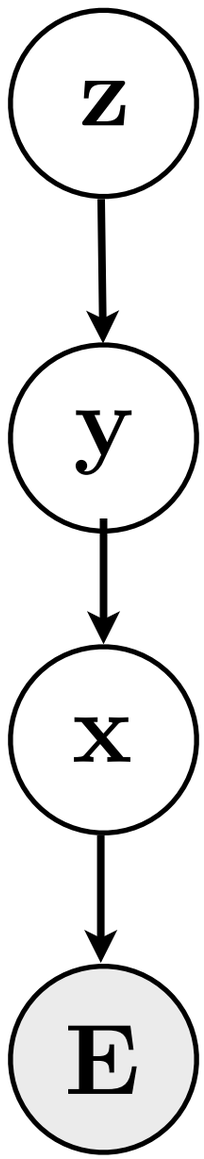

IT

V4

V1

Retina

Feed-forward & feed-back

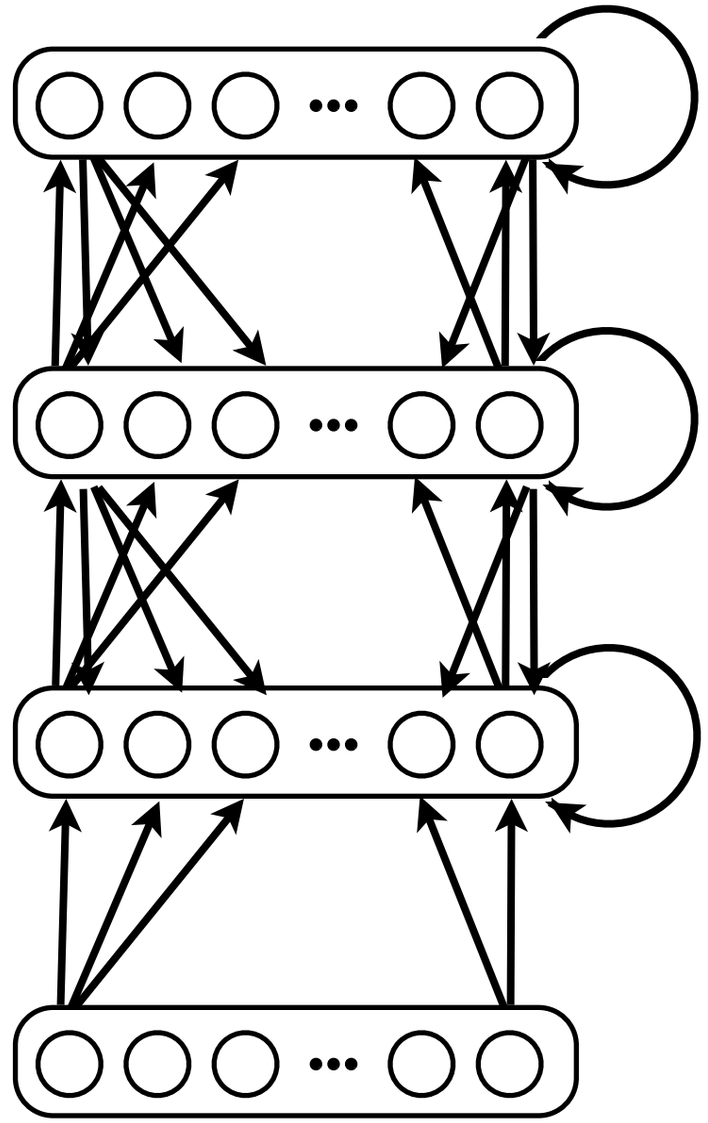

FF, FB & Recurrent connections: Mumford 1992, Lee & Mumford 2003

Full internal model for natural vision unknown

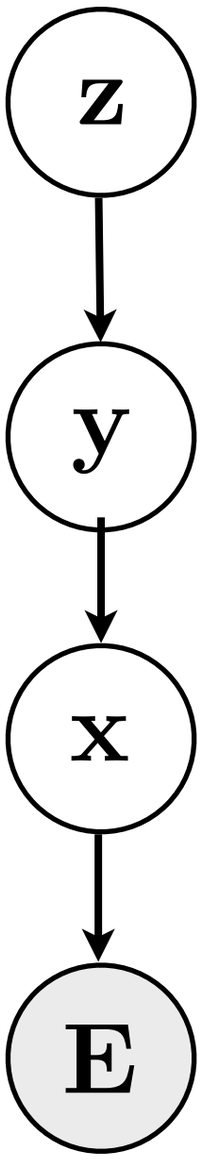

V1

Retina

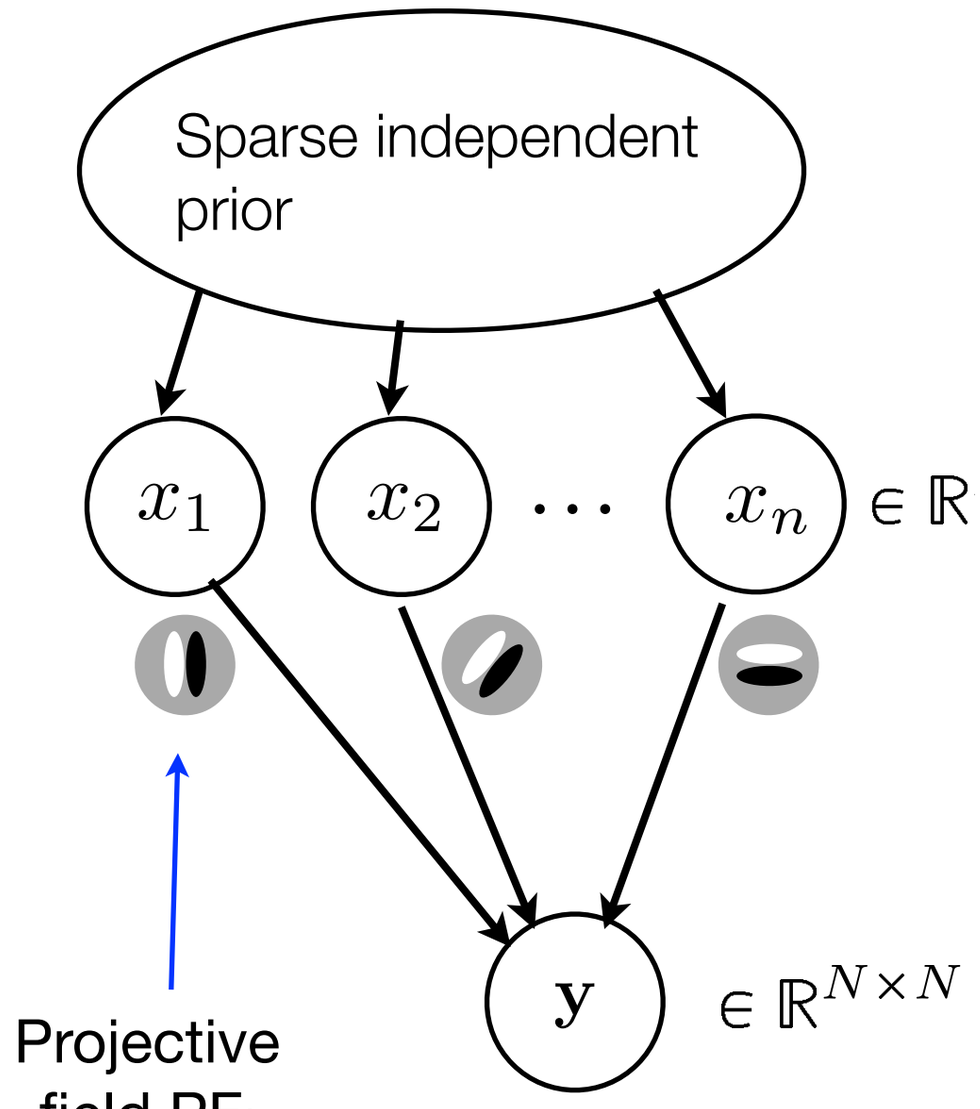

Olshausen & Field, 1996, 1997

Projective field PF_i

$$p(\mathbf{y}|\mathbf{x}) = \mathcal{N}\left(\mathbf{y} : \sum_i \mathrm{PF}_i x_i, \sigma^2\right)$$

Perturbation approach

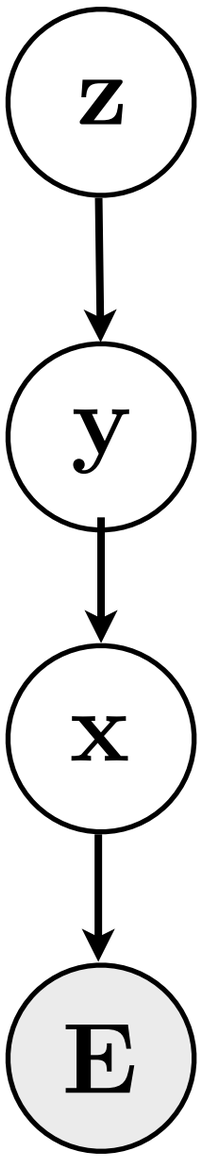

V1

Retina

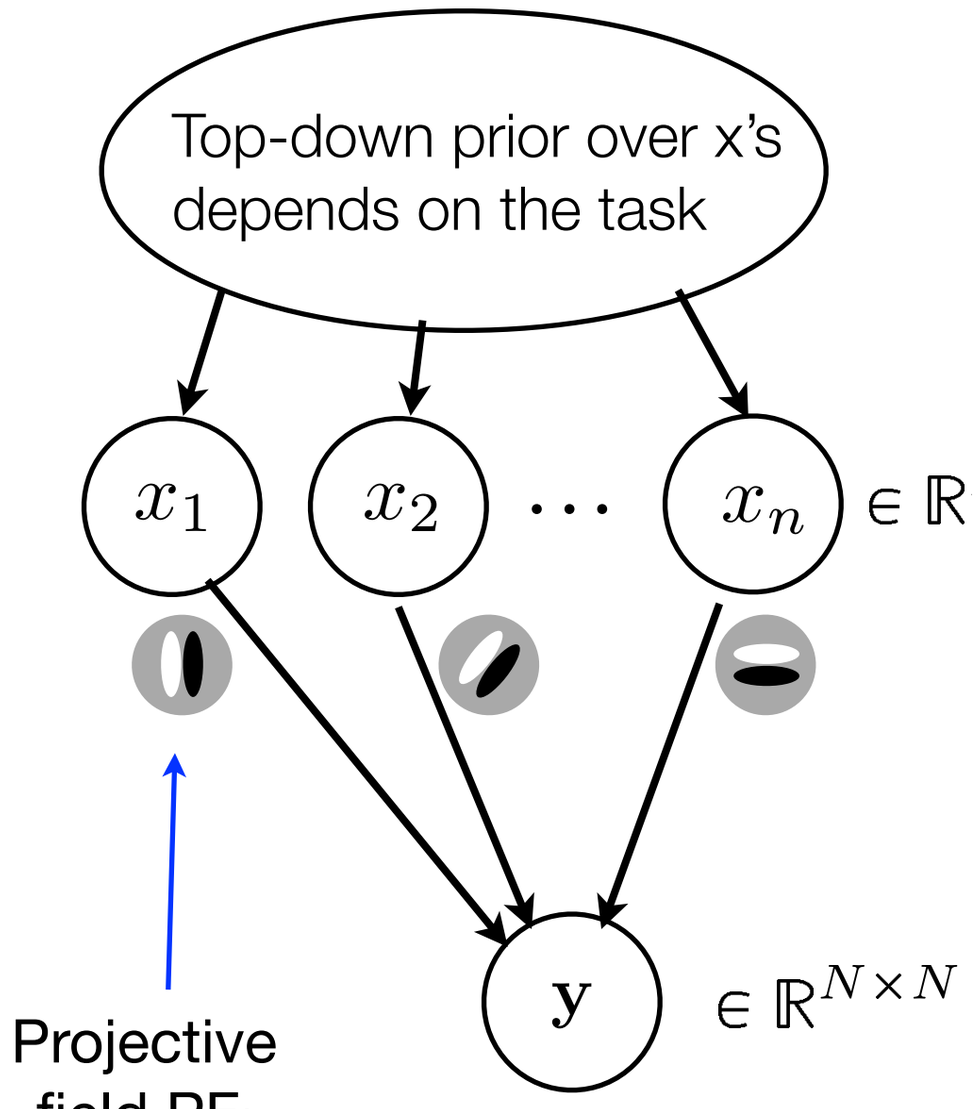

$$p(\mathbf{y}|\mathbf{x}) = \mathcal{N}\left(\mathbf{y} : \sum_i \mathrm{PF}_i x_i, \sigma^2\right)$$

Perturbation approach

Predictions for:

- Difference between before and after learning

- Difference between two different tasks

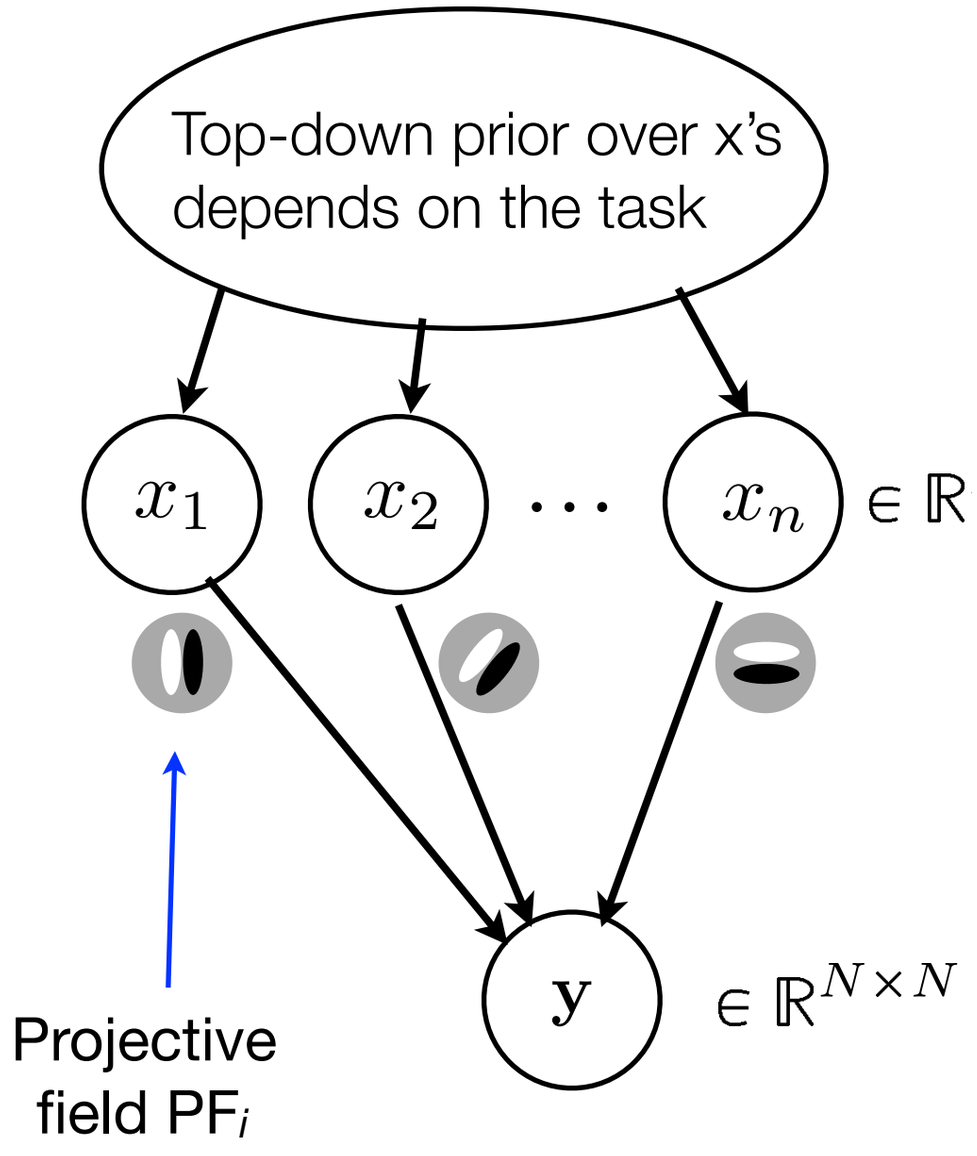

$$p(\mathbf{y}|\mathbf{x}) = \mathcal{N}\left(\mathbf{y} : \sum_i \mathrm{PF}_i x_i, \sigma^2\right)$$

Orientation discrimination task

Which of two perpendicular gratings
caused the noisy image on the screen?

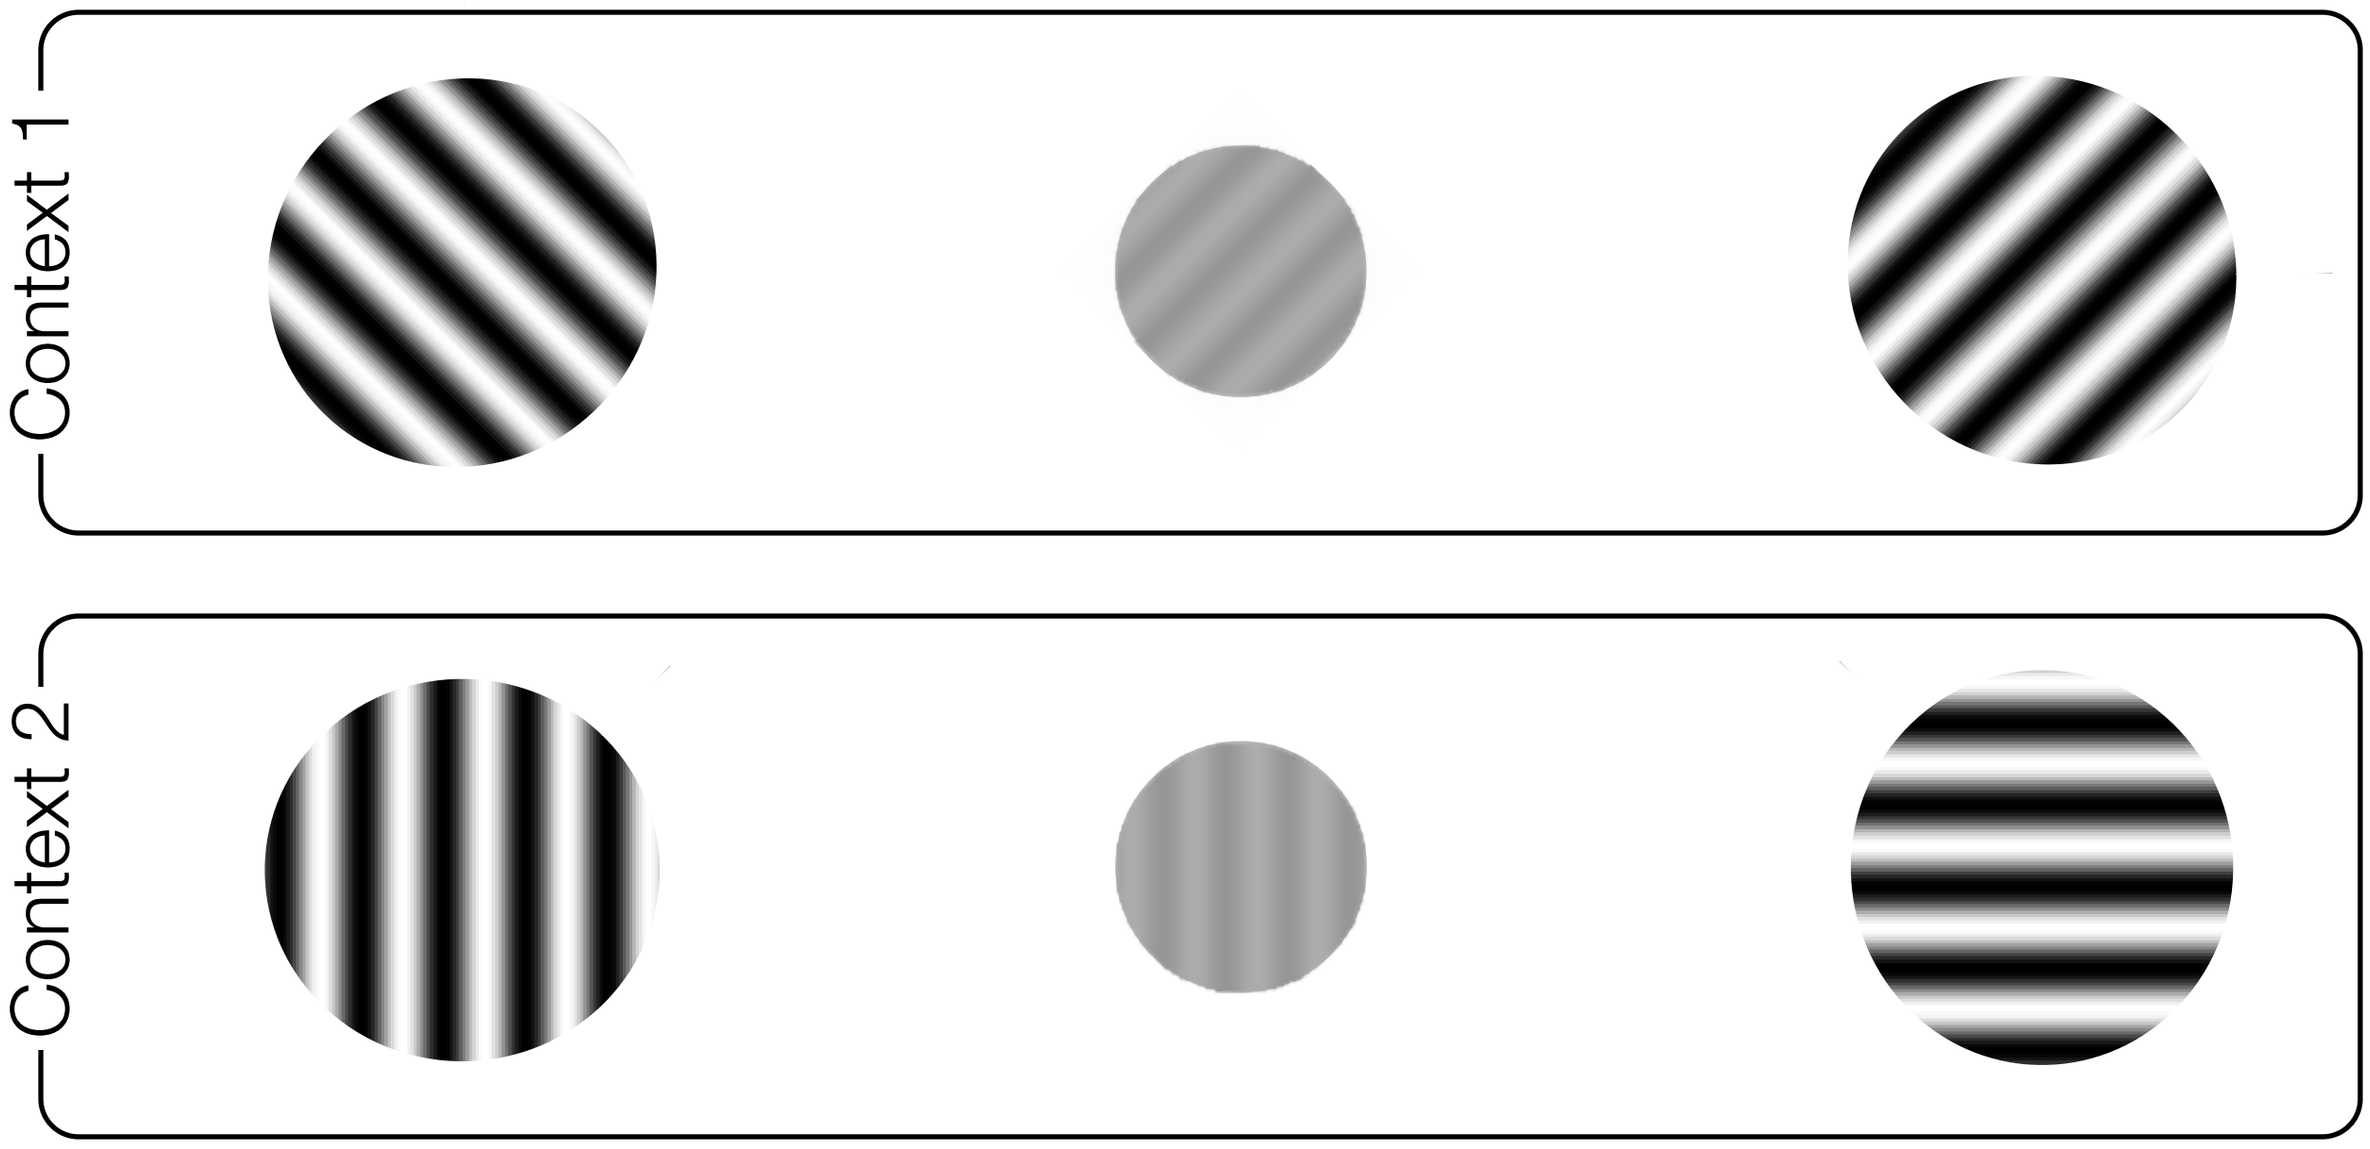

Inspired by Cohen & Newsome, Neuron 2008: Motion direction discrimination task/MT

2AFC task: generative model

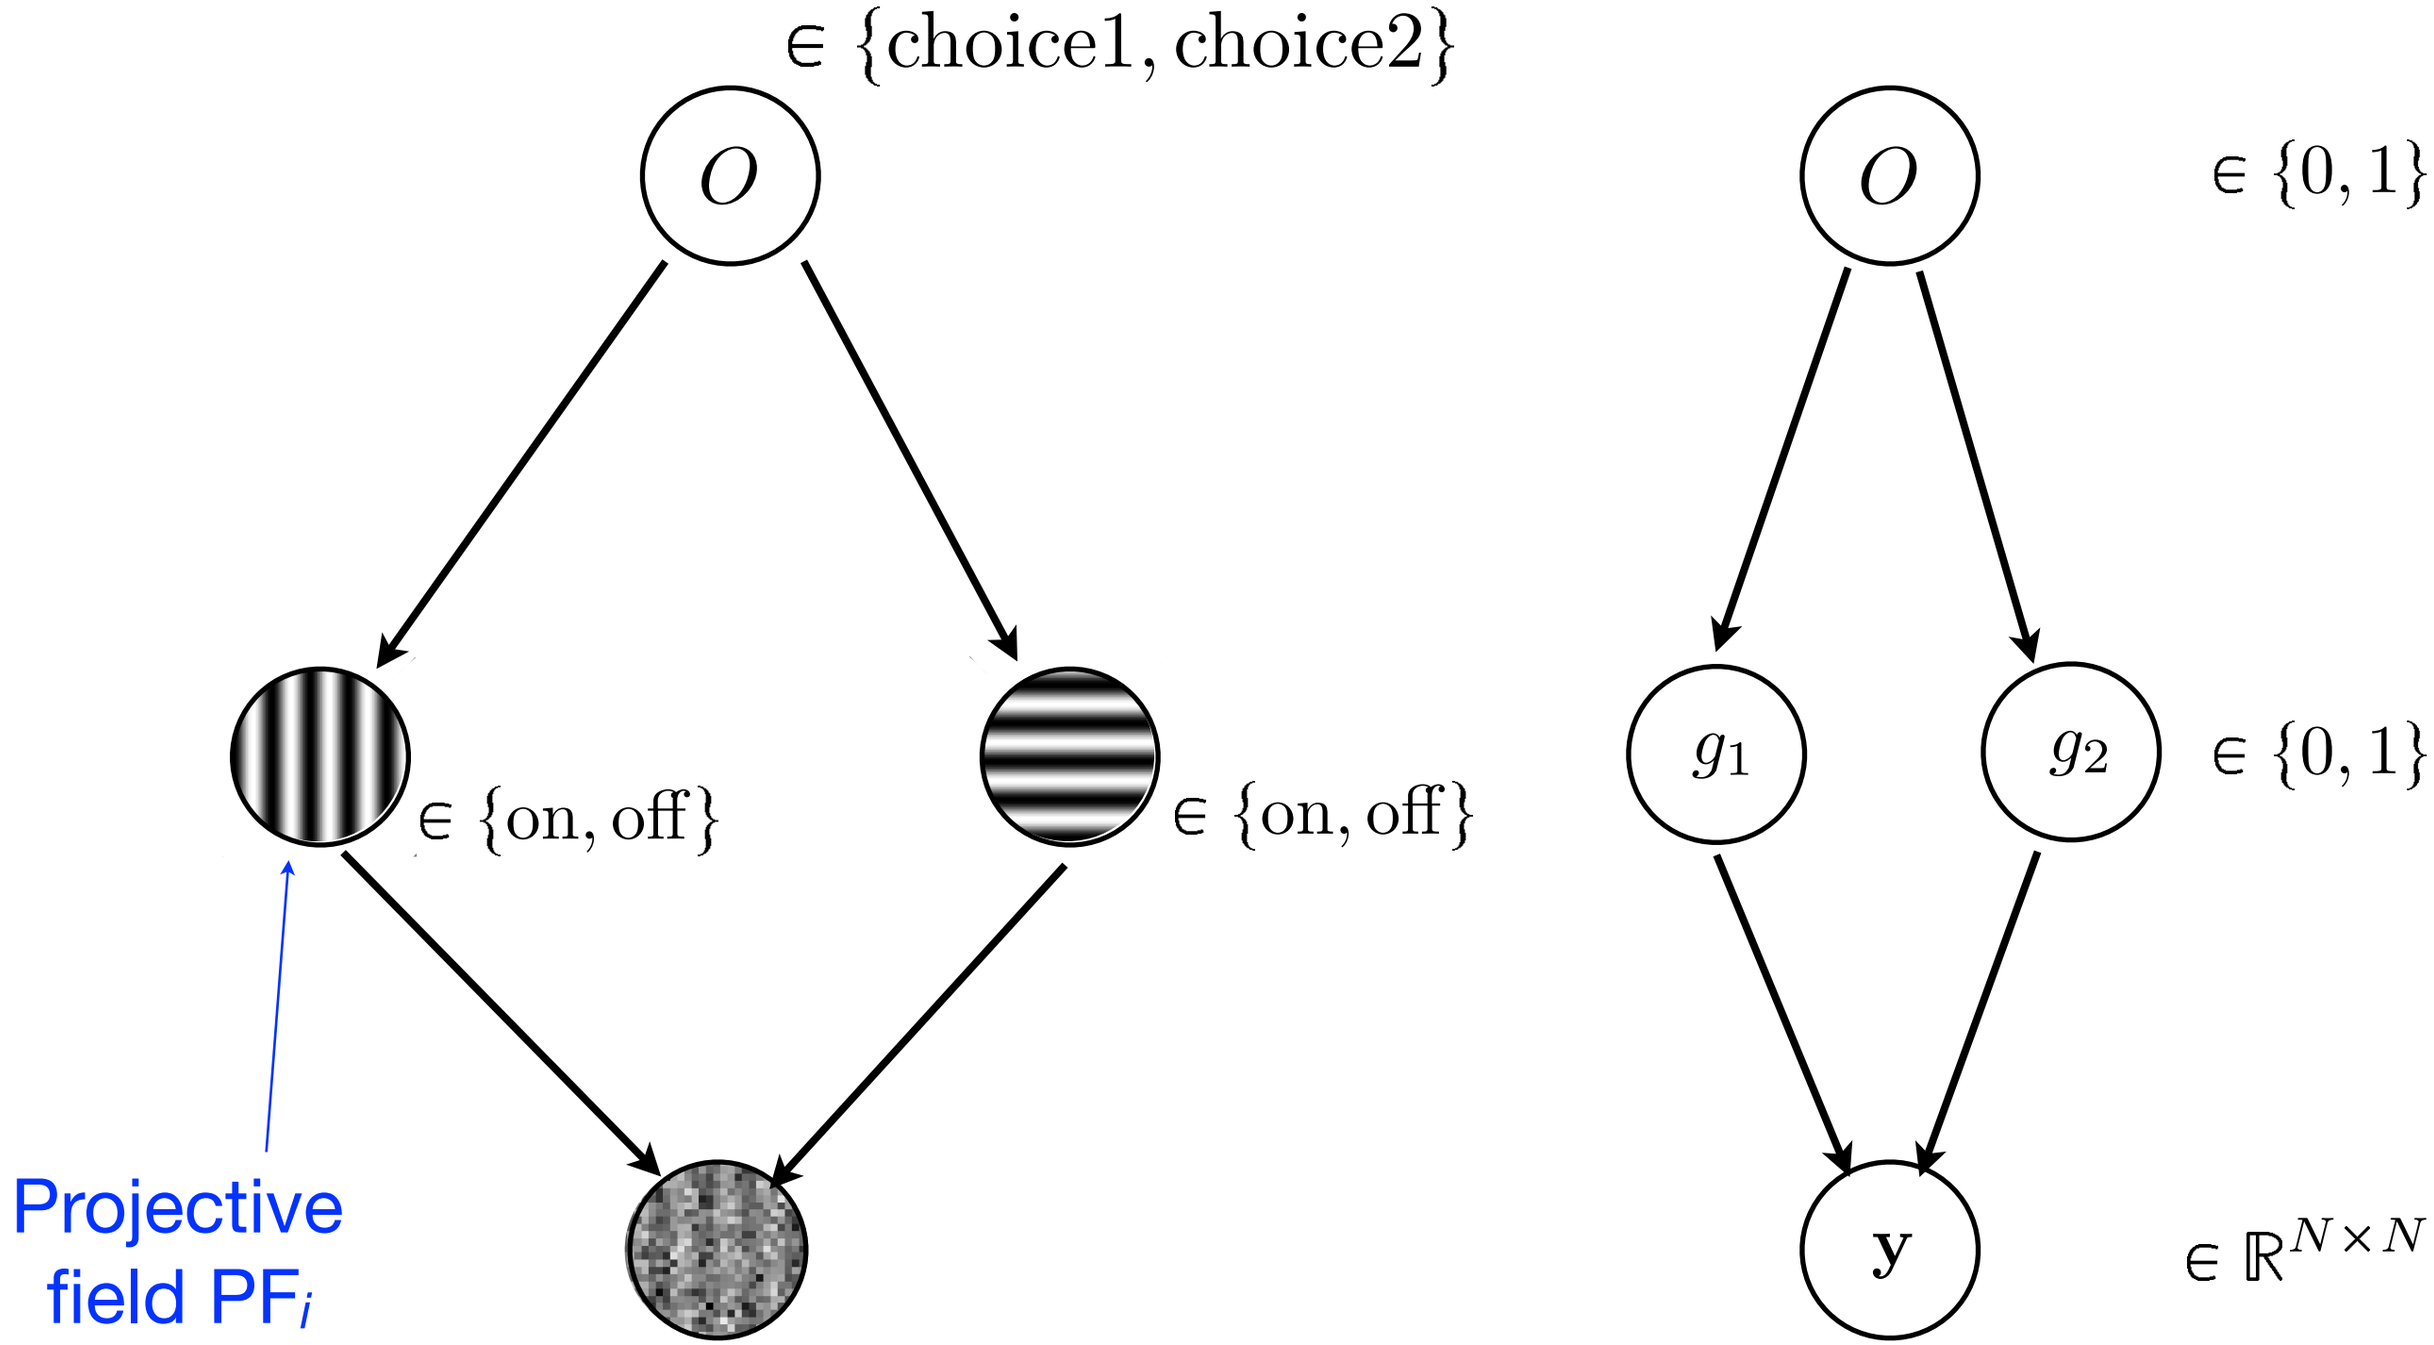

$$p(\mathbf{y}|\mathbf{g}) = \mathcal{N}\left(\mathbf{y} \,:\, \sum_i \mathrm{PF}_i g_i \,,\, \sigma_y^2\right)$$

2AFC task: generative model

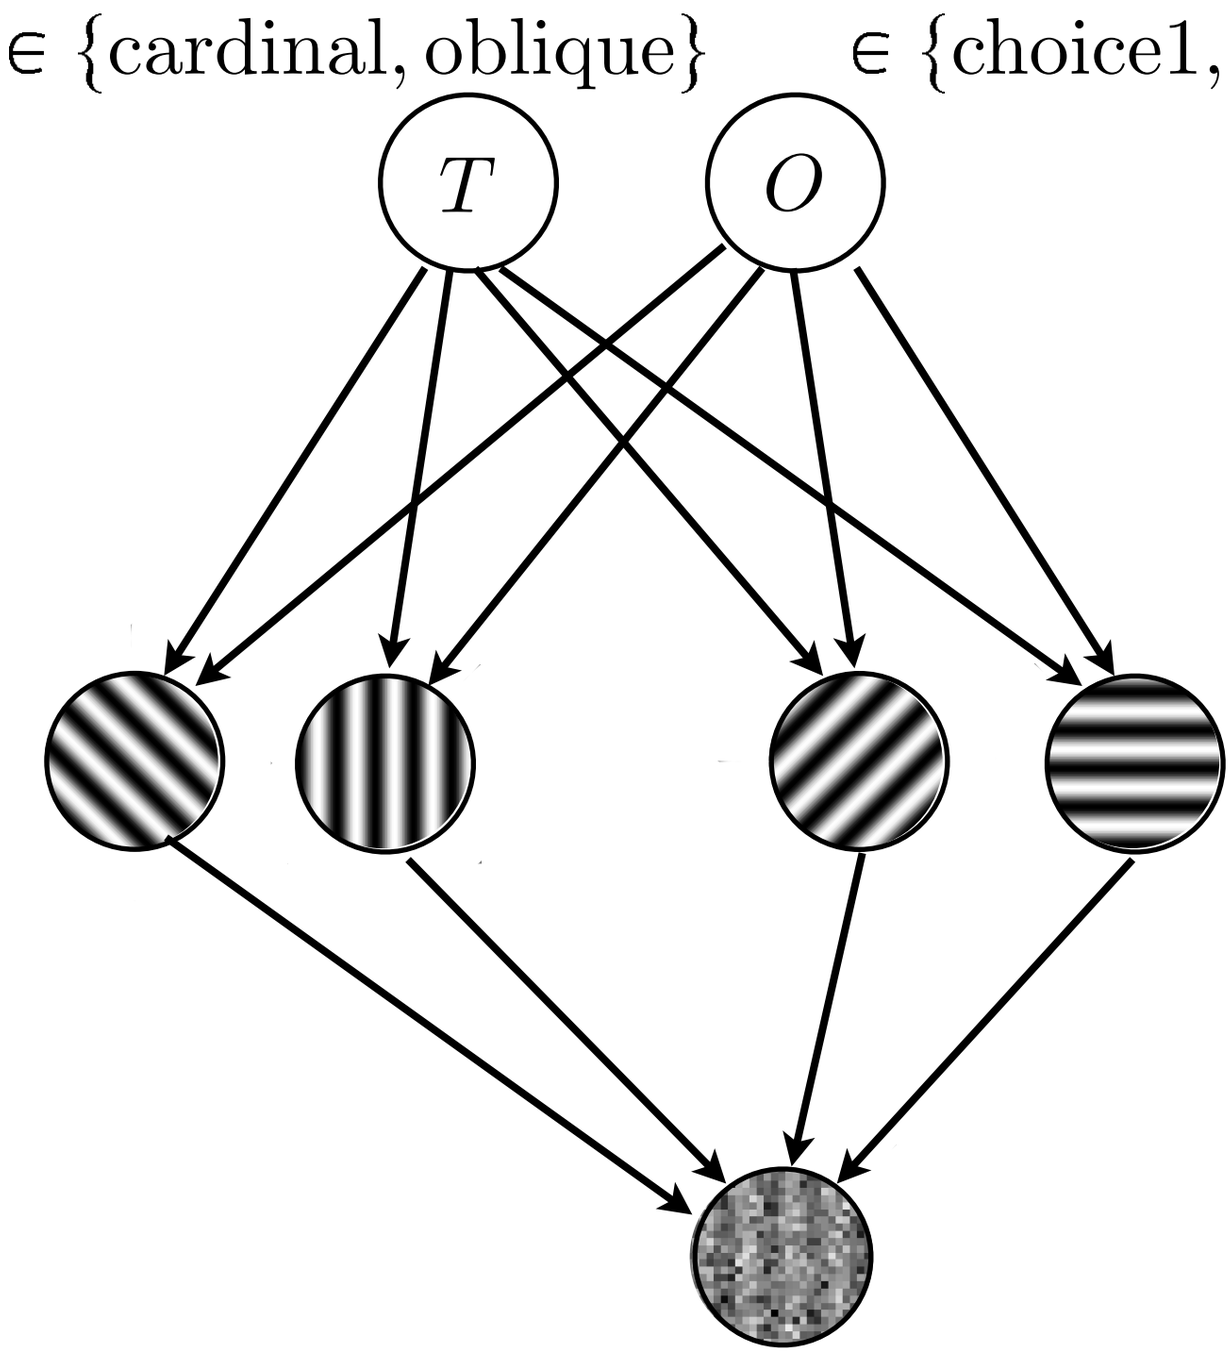

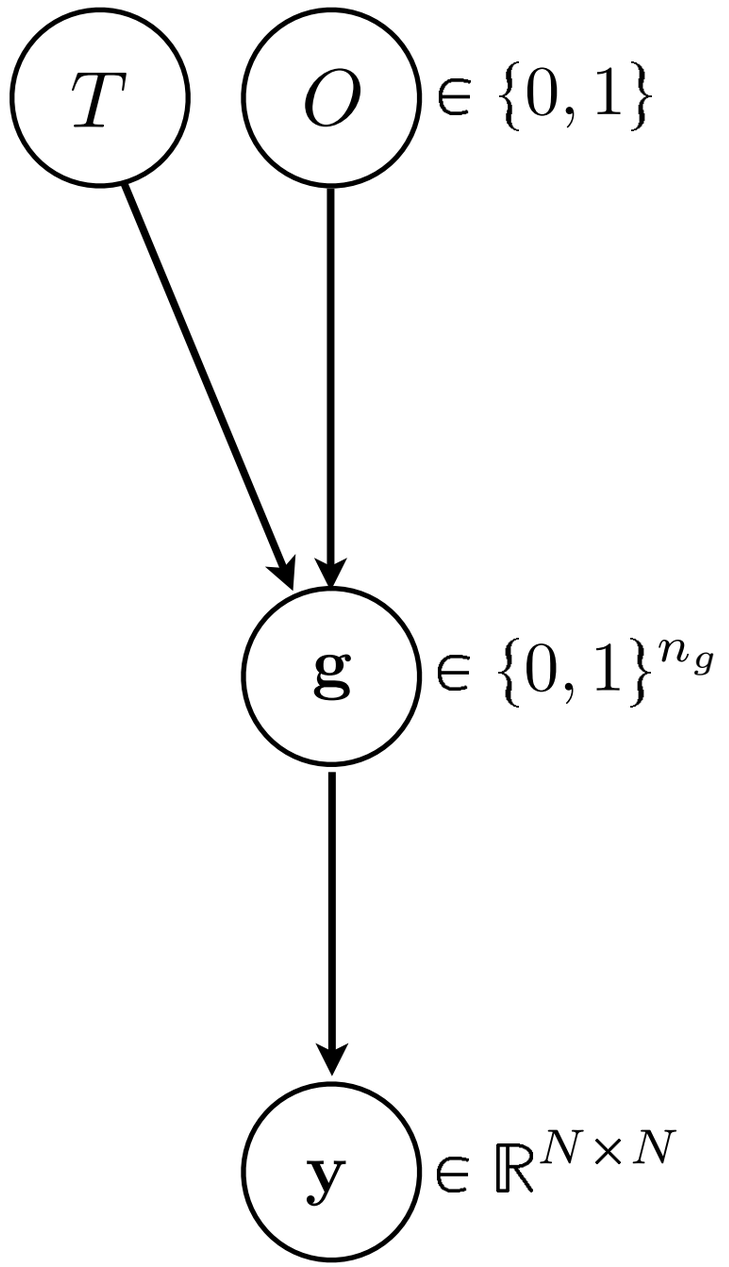

$$p(\mathbf{y}|\mathbf{g}) = \mathcal{N}\left(\mathbf{y} \,:\, \sum_i \mathrm{PF}_i g_i \,,\, \sigma_y^2\right)$$

Inference by neural sampling

$O \sim p(O|g_1, g_2)$

$g_1 \sim p(g_2|O, g_1, \mathbf{y})$

$g_2 \sim p(g_1|O, g_1, \mathbf{y})$

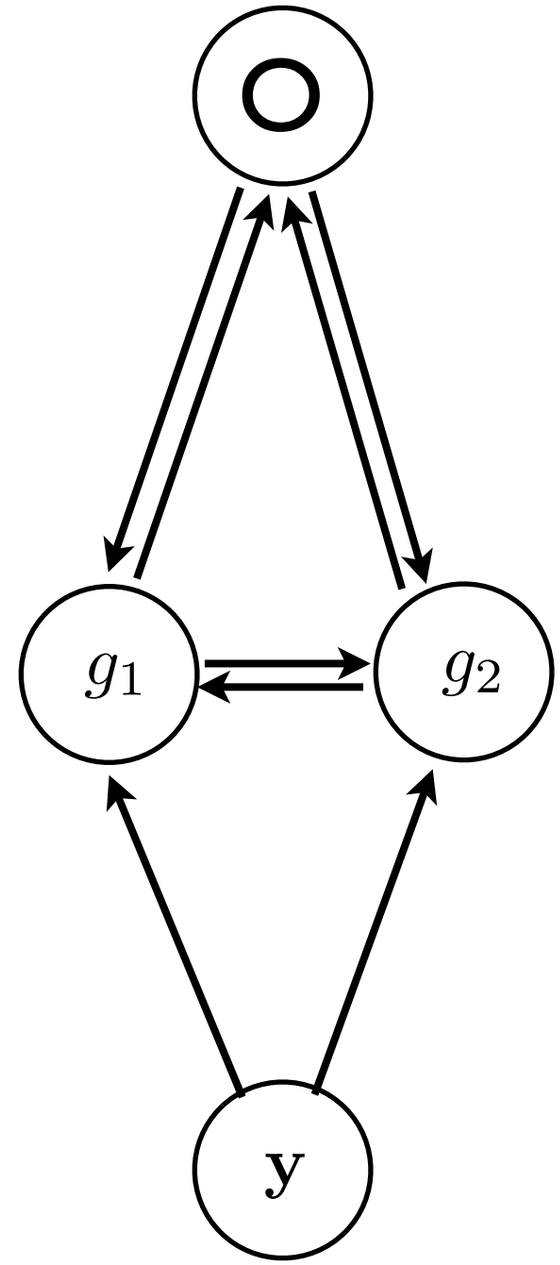

$\mathbf{y} =$

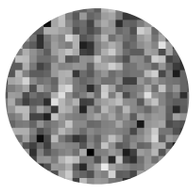

Predictions

$O \sim p(O|g_1, g_2)$

$g_1 \sim p(g_1|O, g_2, \mathbf{y})$

$g_2 \sim p(g_2|O, g_1, \mathbf{y})$

$\mathbf{y} =$ 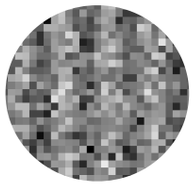

$p(O|\mathbf{y})$

$p(g_1|\mathbf{y})$

$p(g_1, g_2|\mathbf{y})$

$p(g_2|\mathbf{y})$

Noise correlations...

...and their task-dependence

Choice probabilities and psychophysical kernels

Correspondence to the visual system

LIP/PPC/
PFC...?

V2/V4/IT?

V1

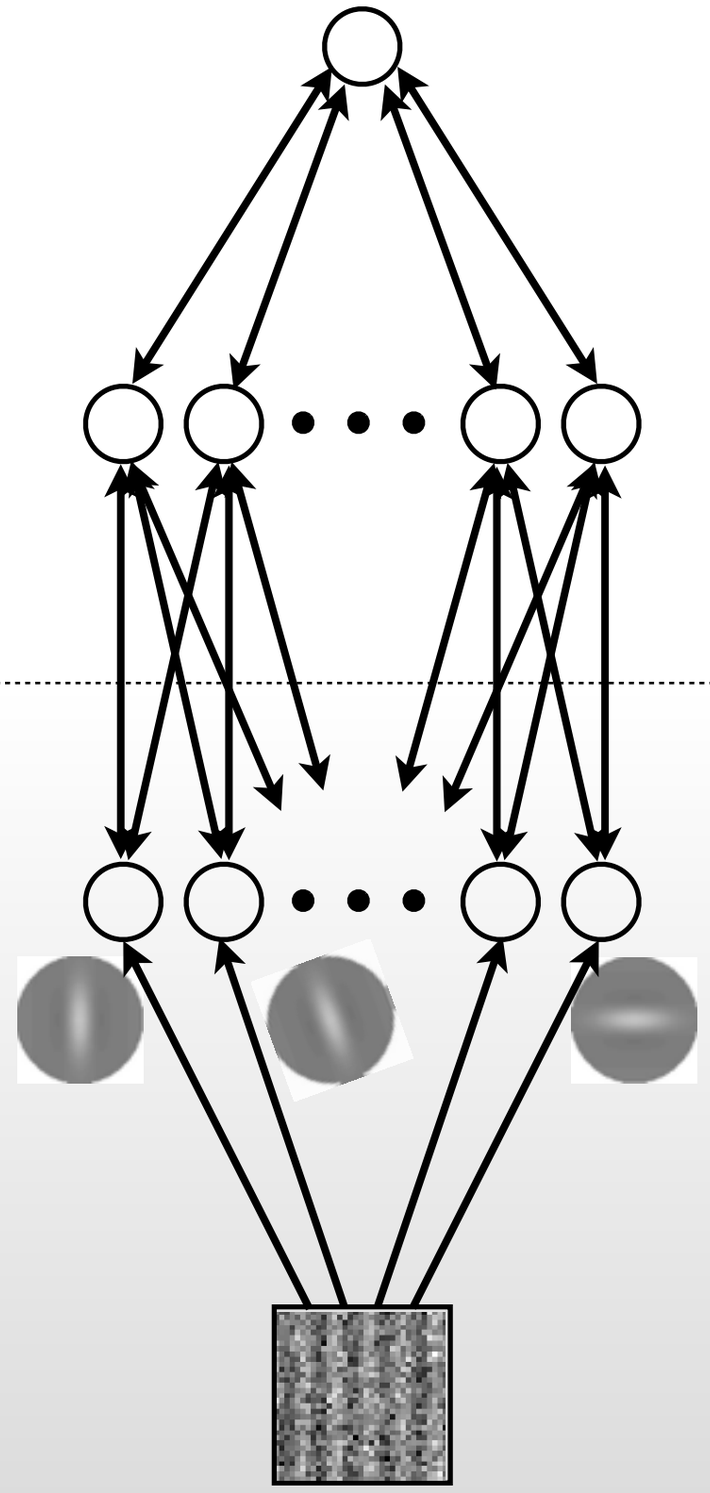

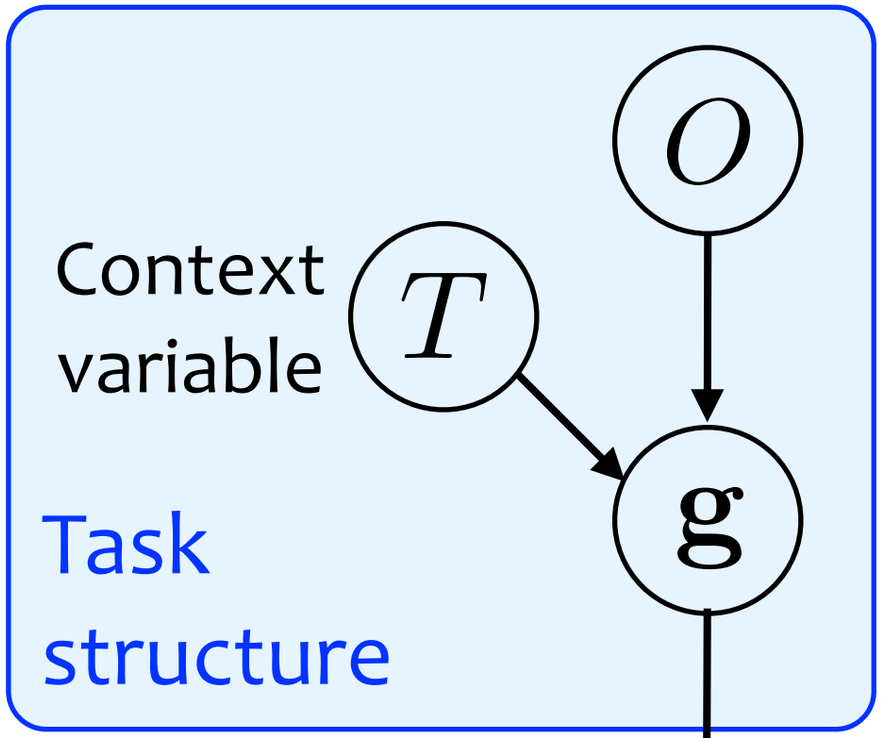

Decision
variable

Grating
variables

Neural
responses

Image
variable

Probabilistic sparse-coding model: Hoyer & Hyvarinnen 2003
Gaussian scale mixture: Schwartz & Simoncelli, 2001

V1/early vision

Dependencies and observables

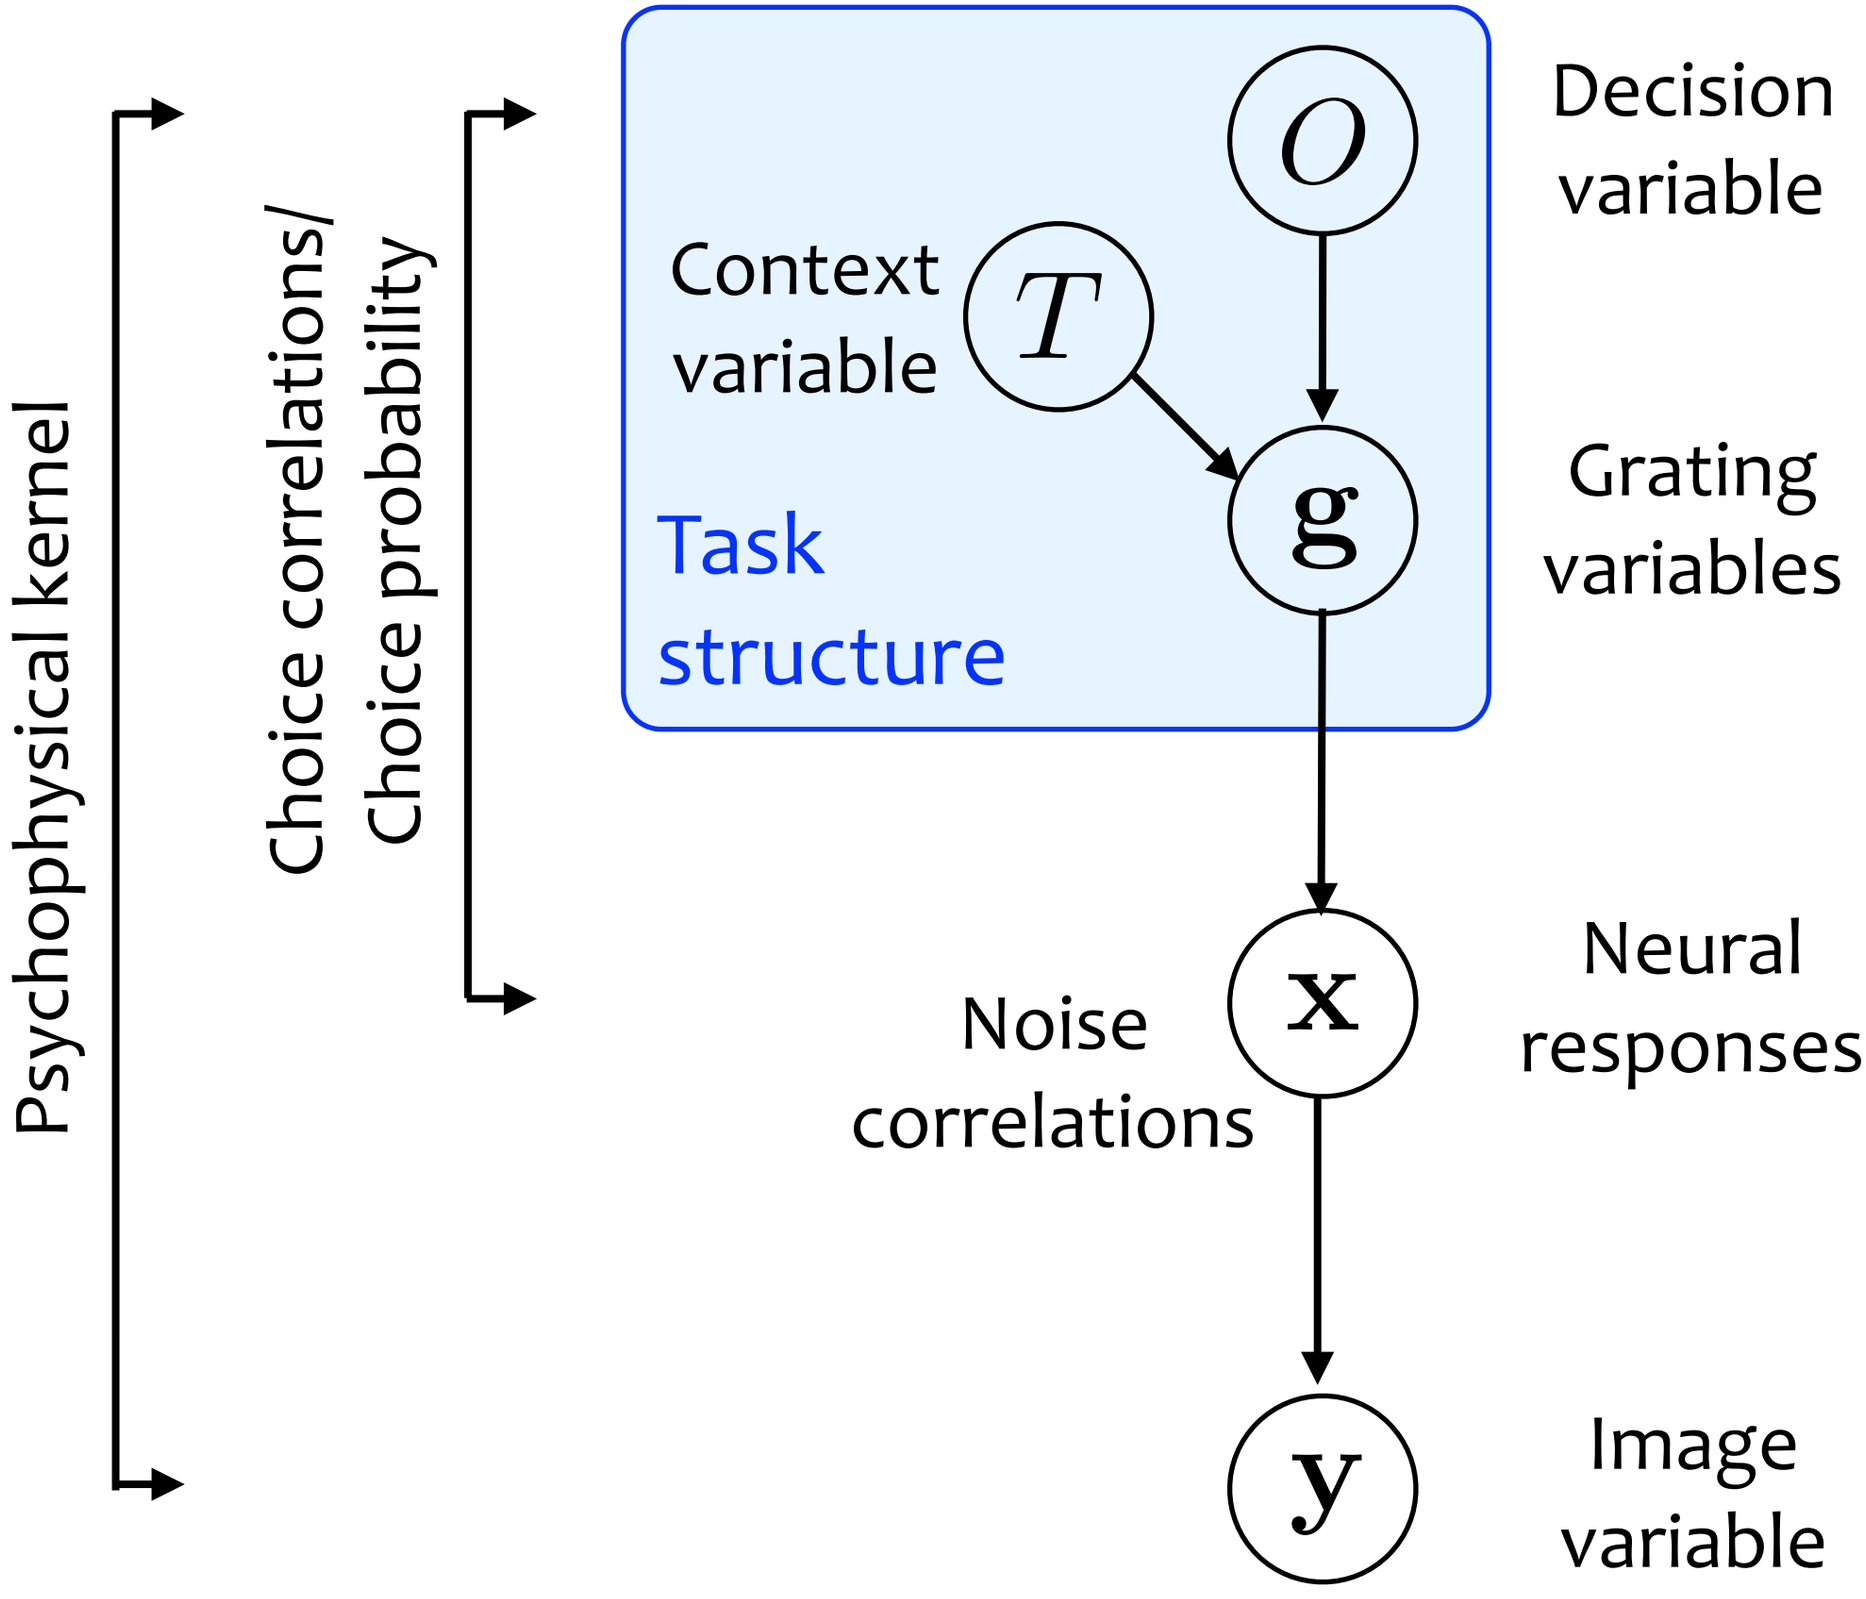

Model: noise correlations in $\mathbf{x}$

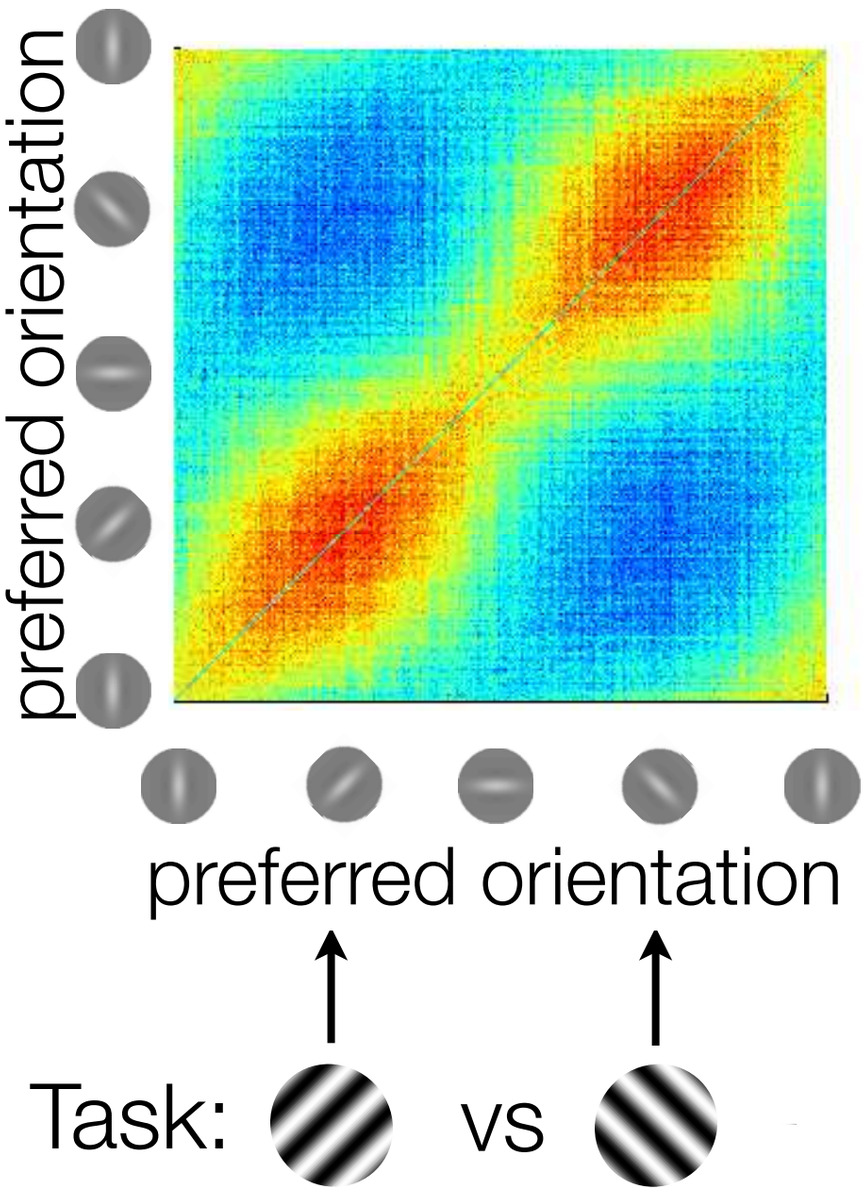

Context 1

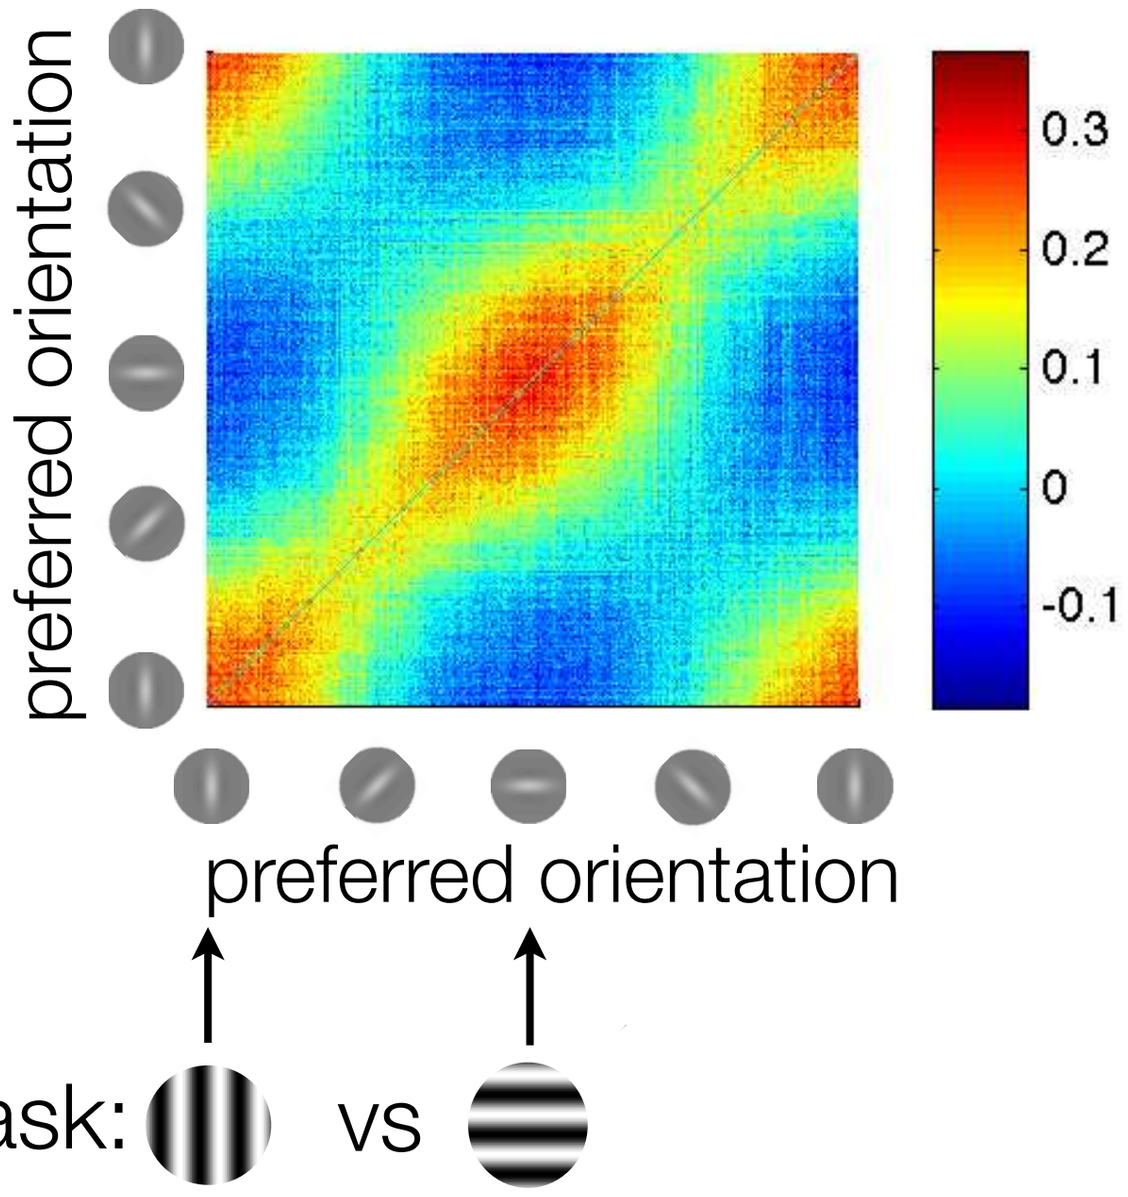

Context 2

Haefner, Berkes & Fiser, Neuron 2016

Model: noise correlations in $\mathbf{x}$

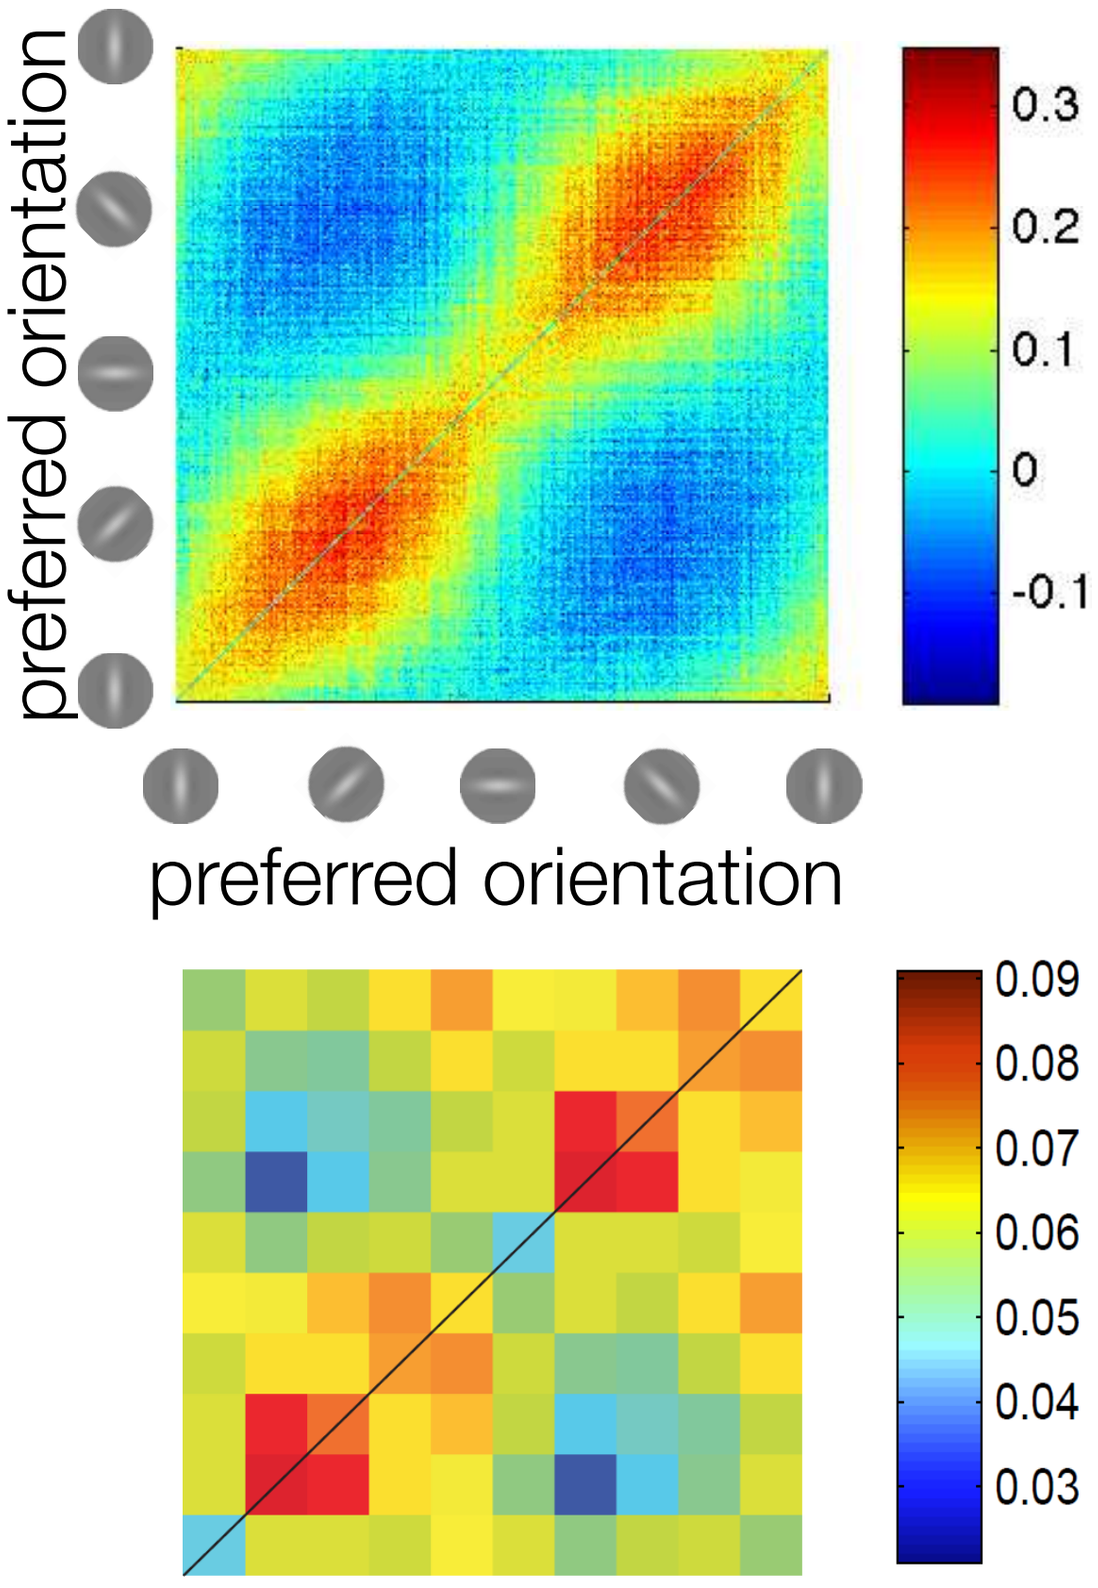

Bondy & Cumming, bioRxiv 2016

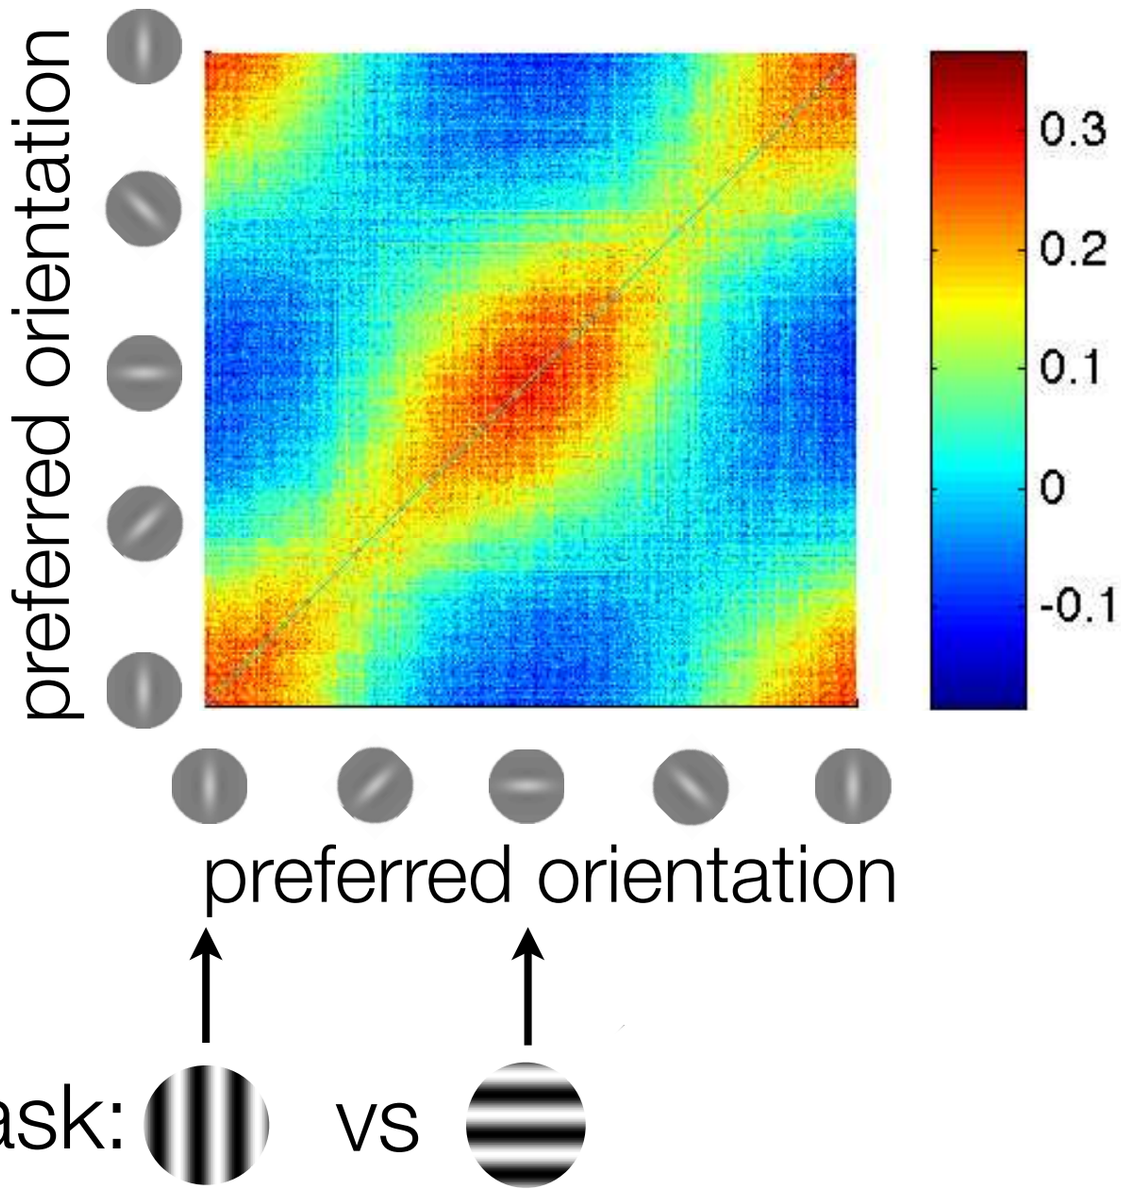

Haefner, Berkes & Fiser, Neuron 2016

Dependencies and observables

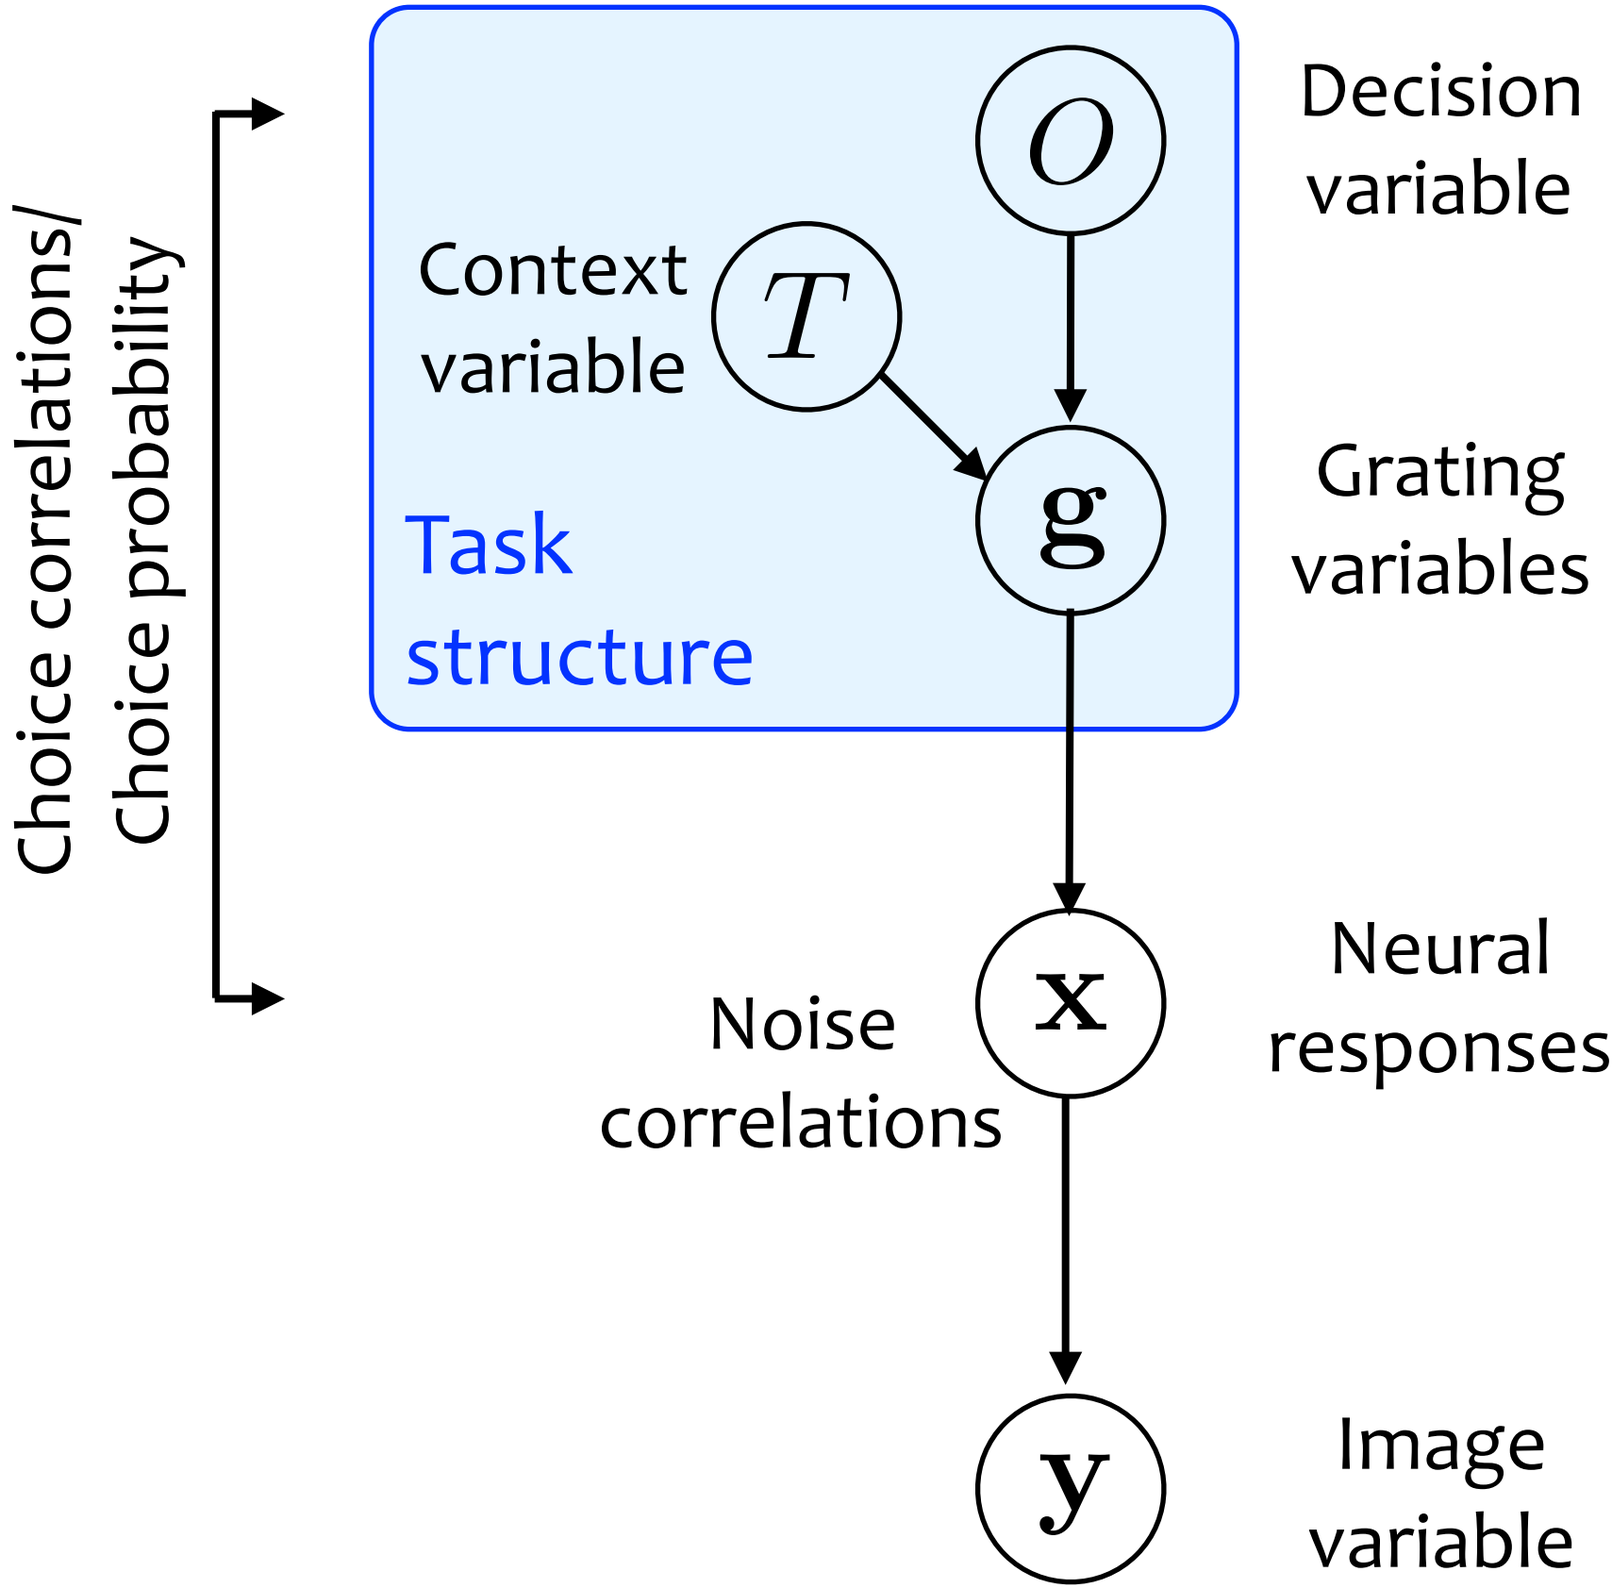

Model: noise correlations in V1

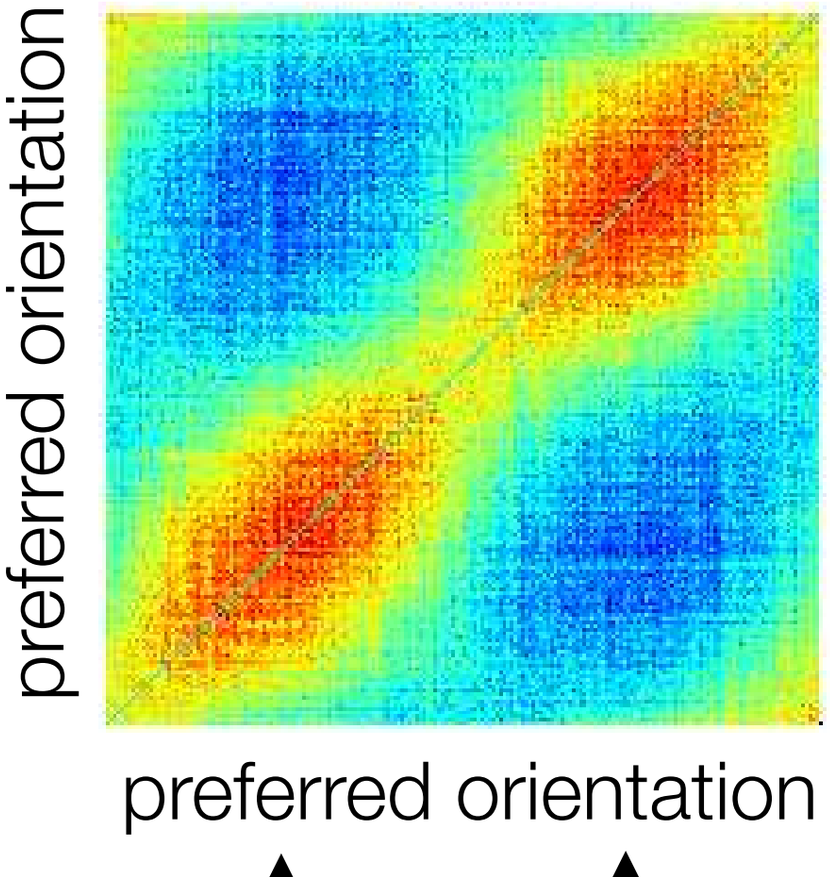

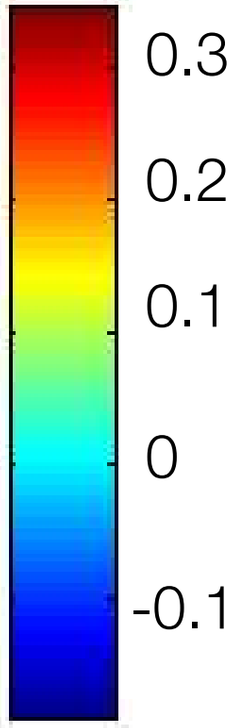

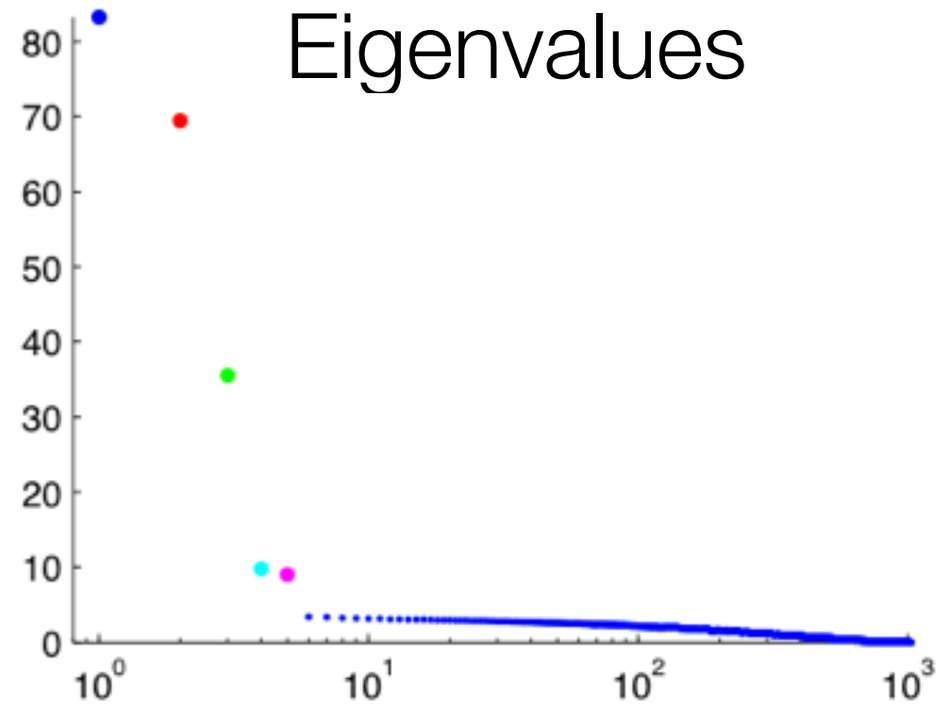

Task: 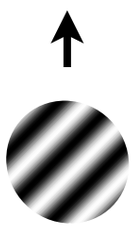 vs 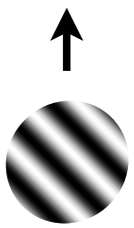

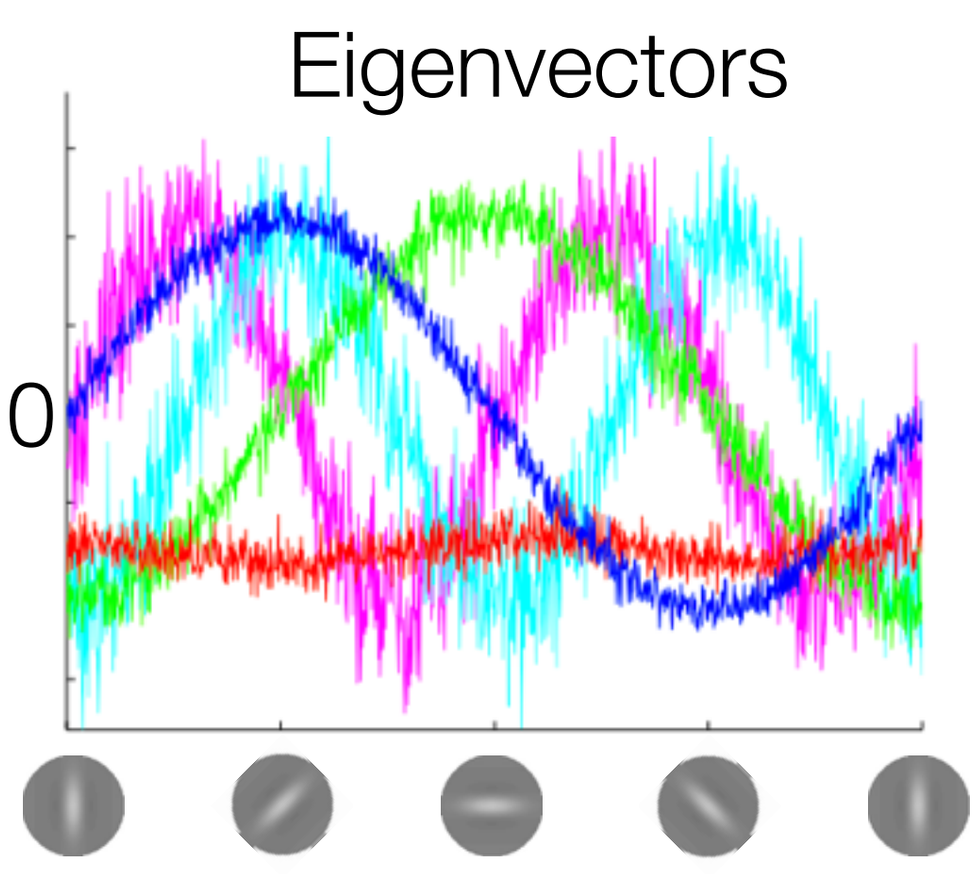

Model: noise correlations in V1

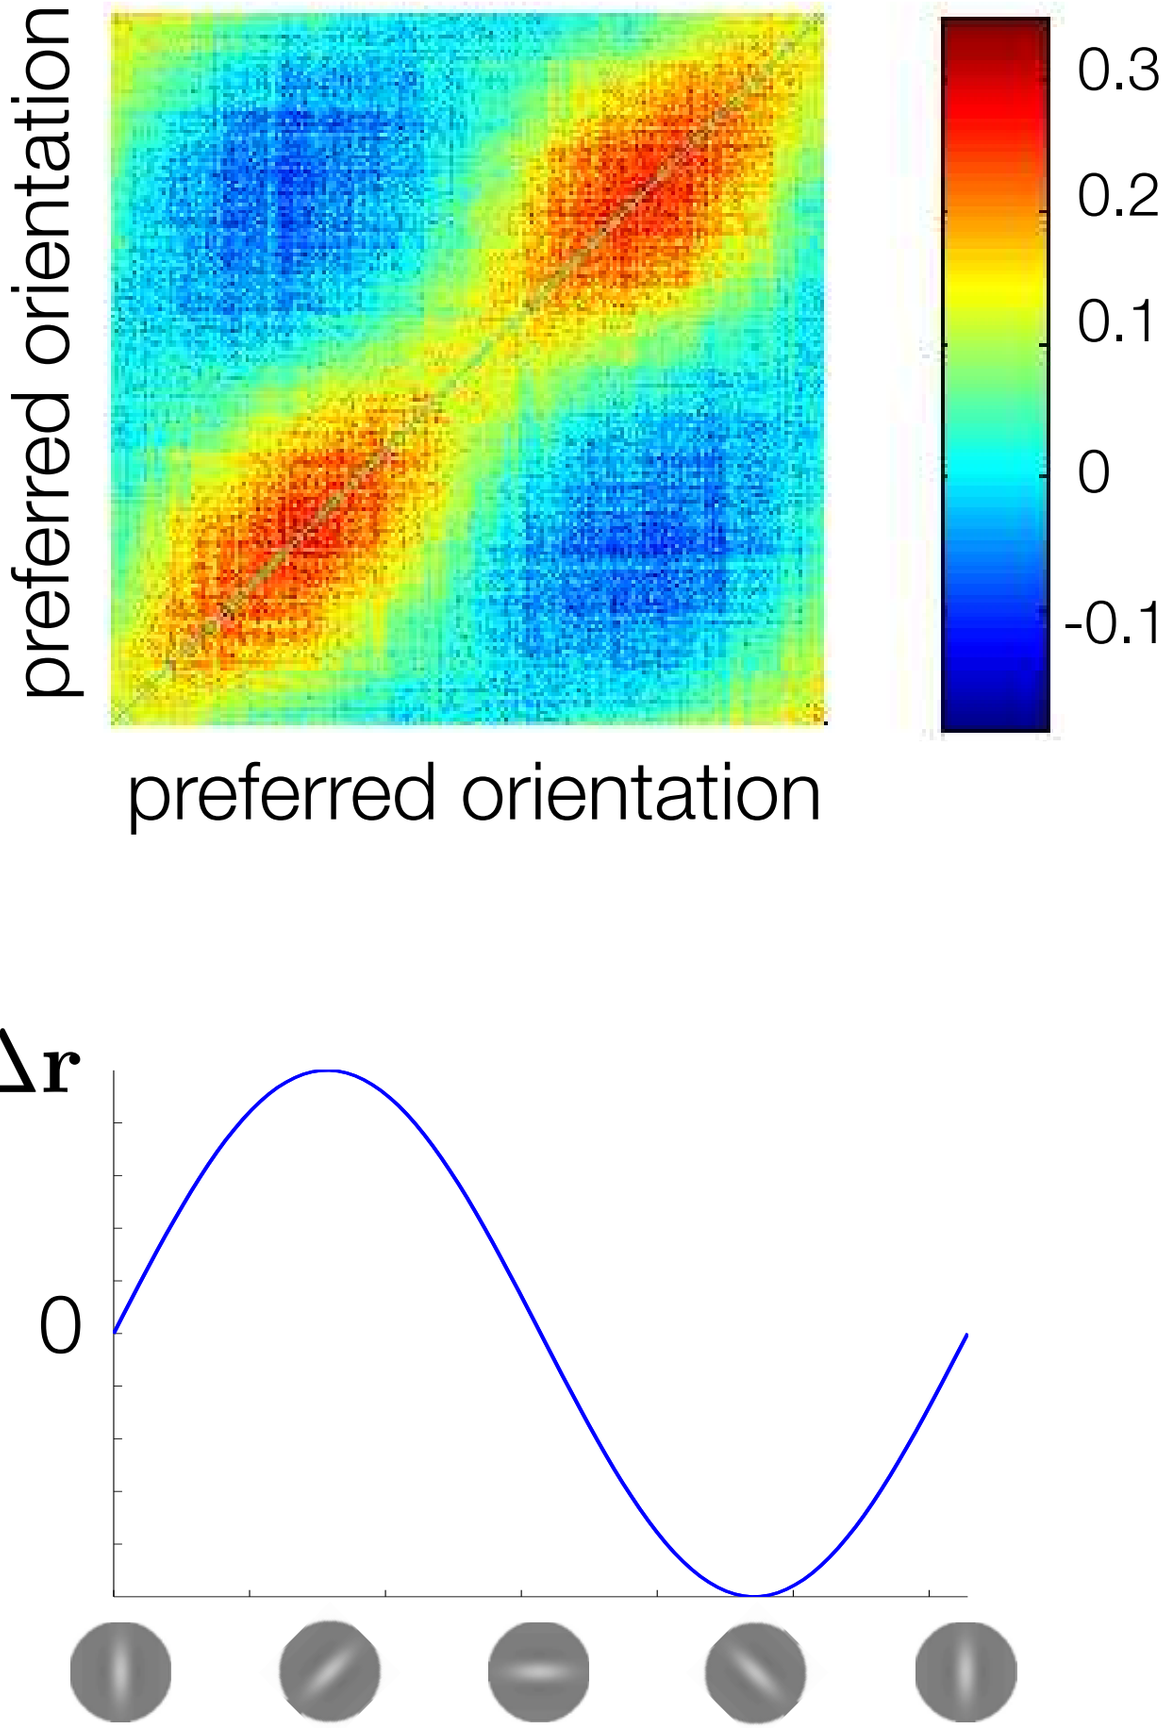

Model: noise correlations in V1

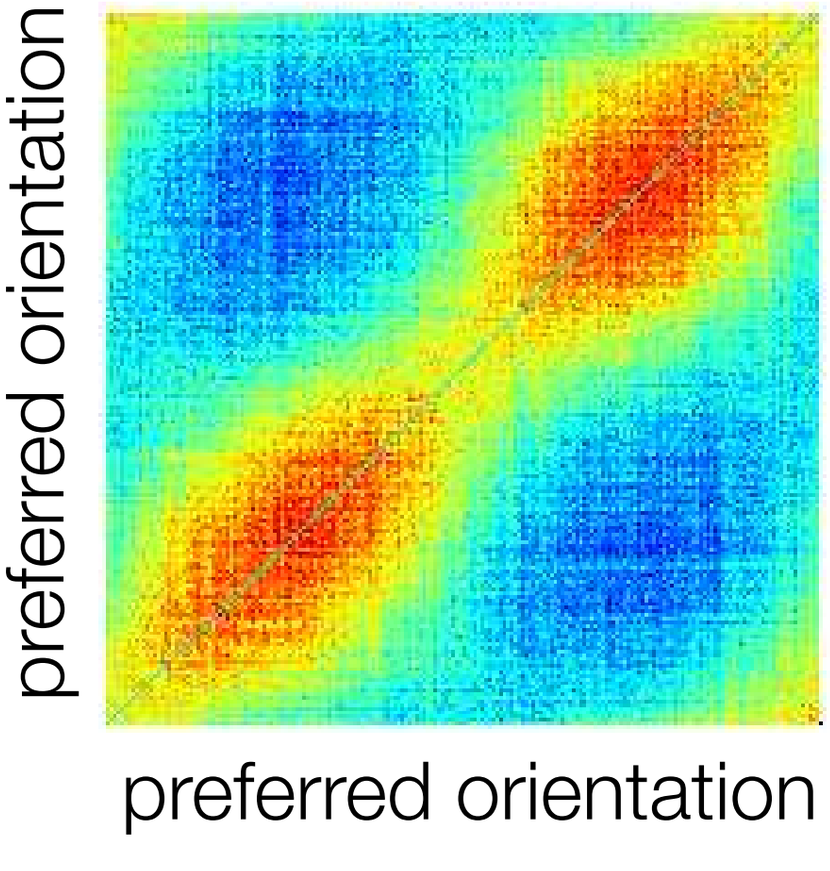

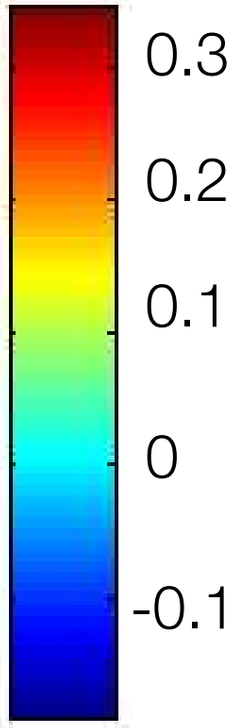

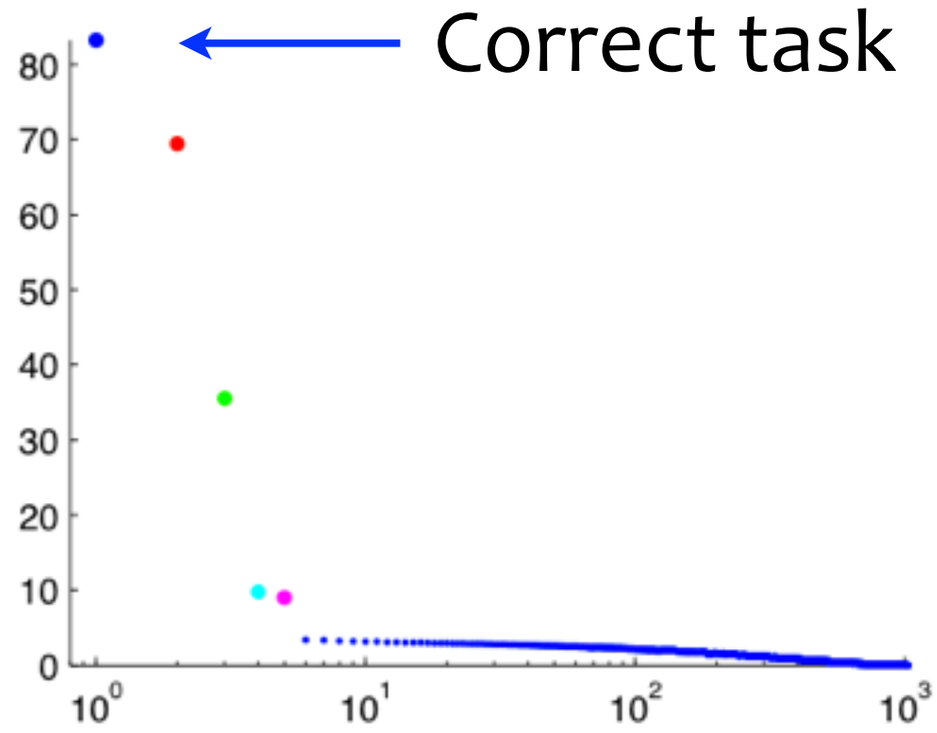

$\mathbf{f'f'^{\top}}$ – covariance

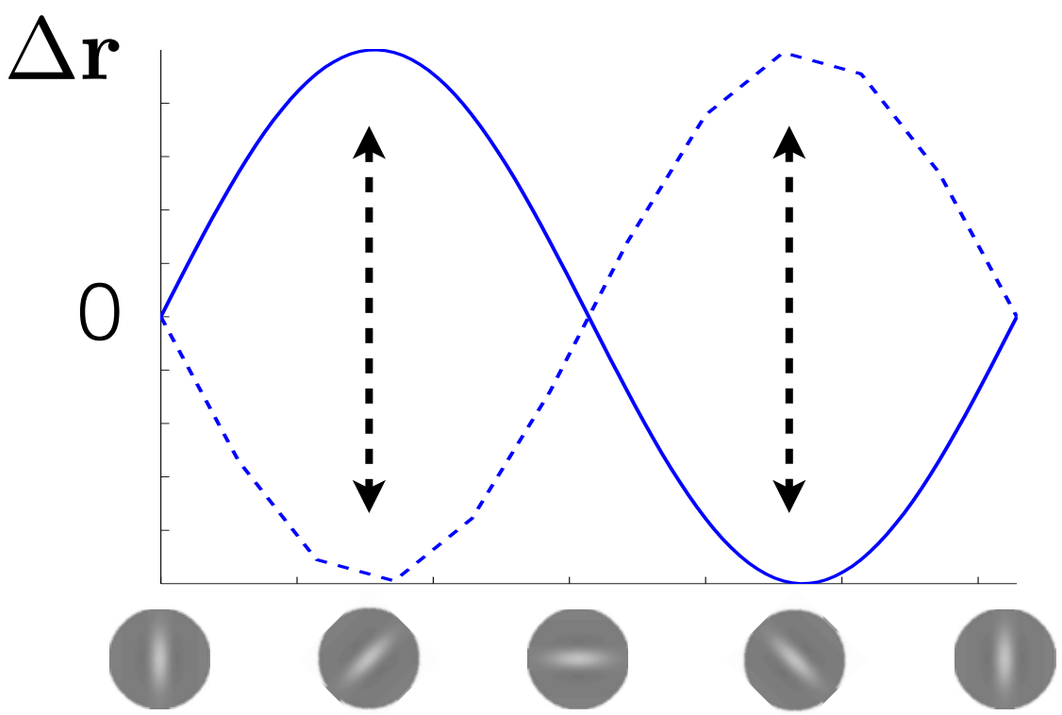

Eigenvectors

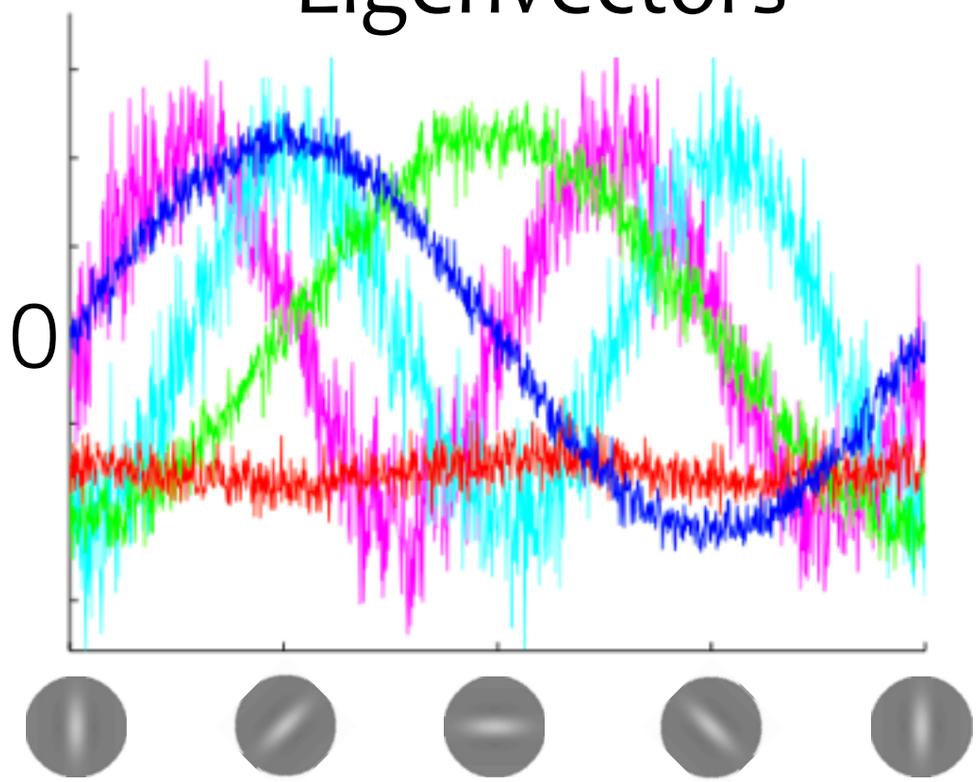

Model: noise correlations in V1

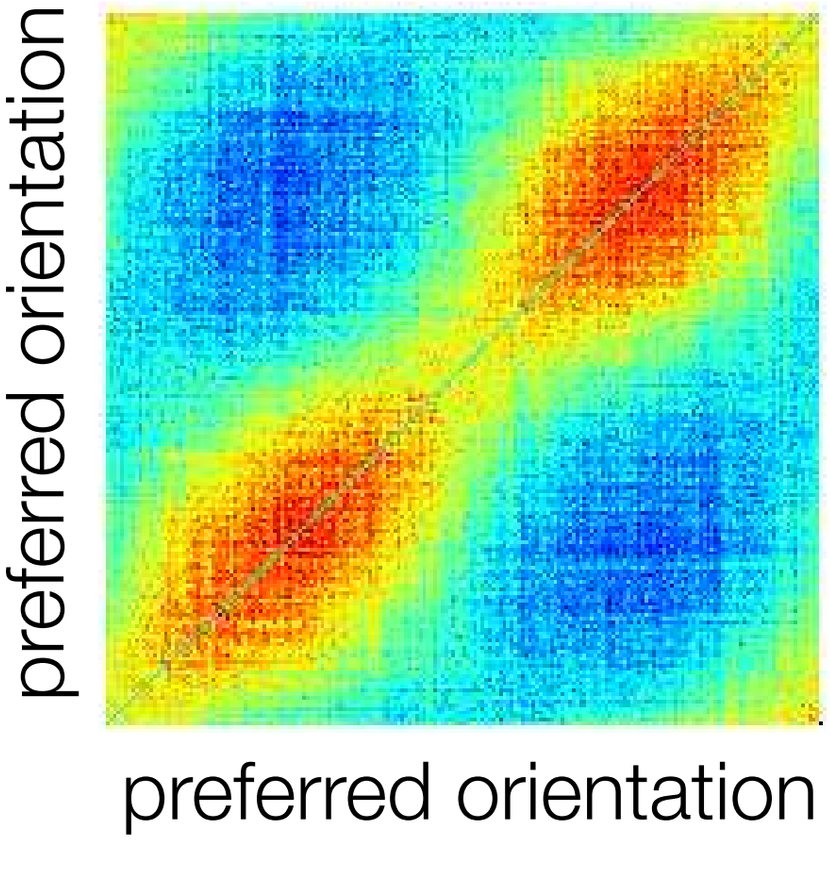

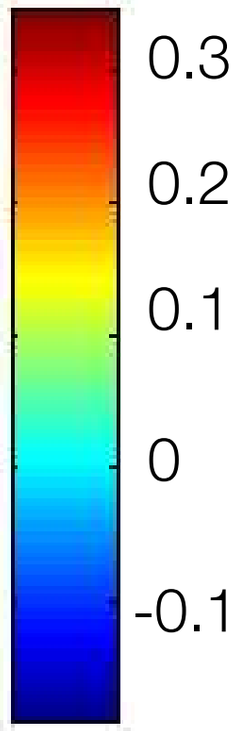

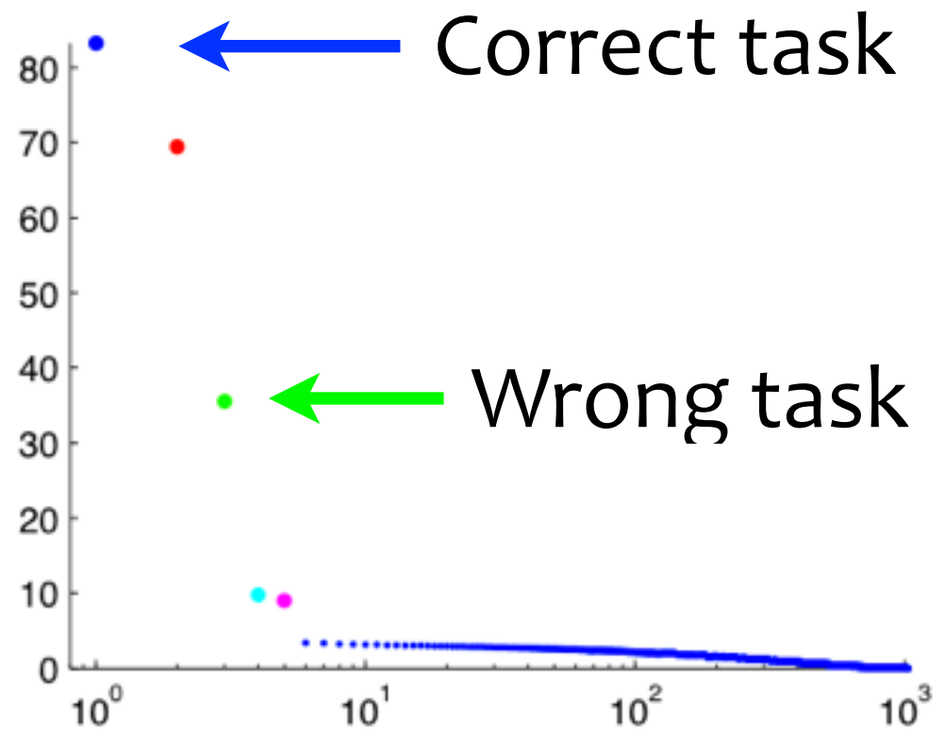

$\mathbf{f'f'^{\top}}$ – covariance

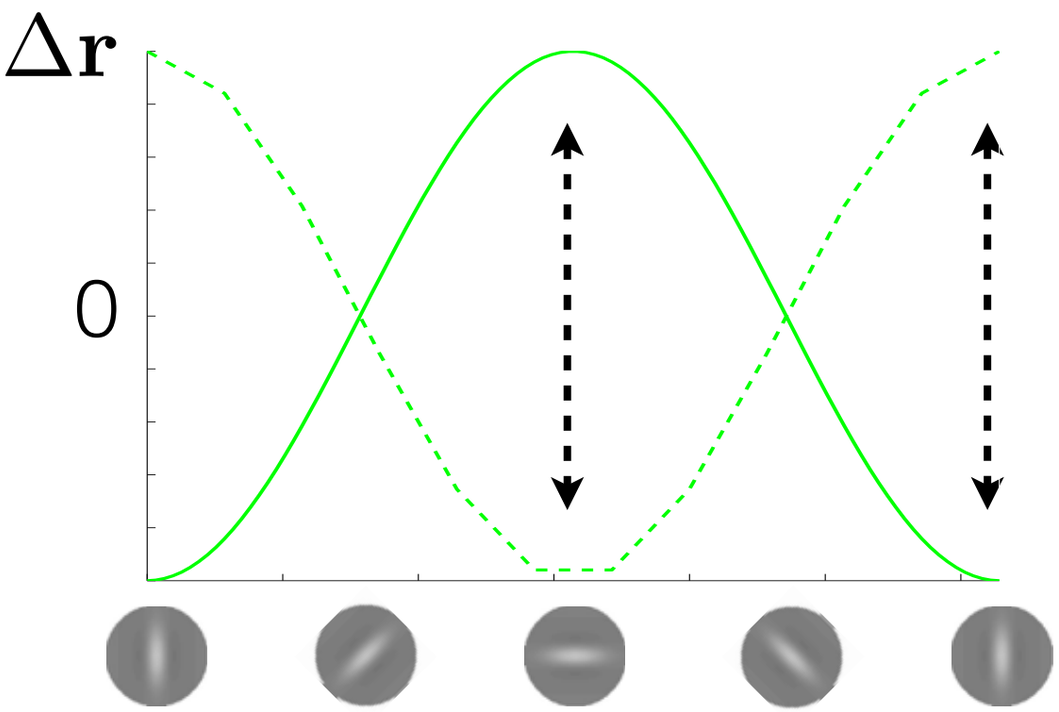

Eigenvectors

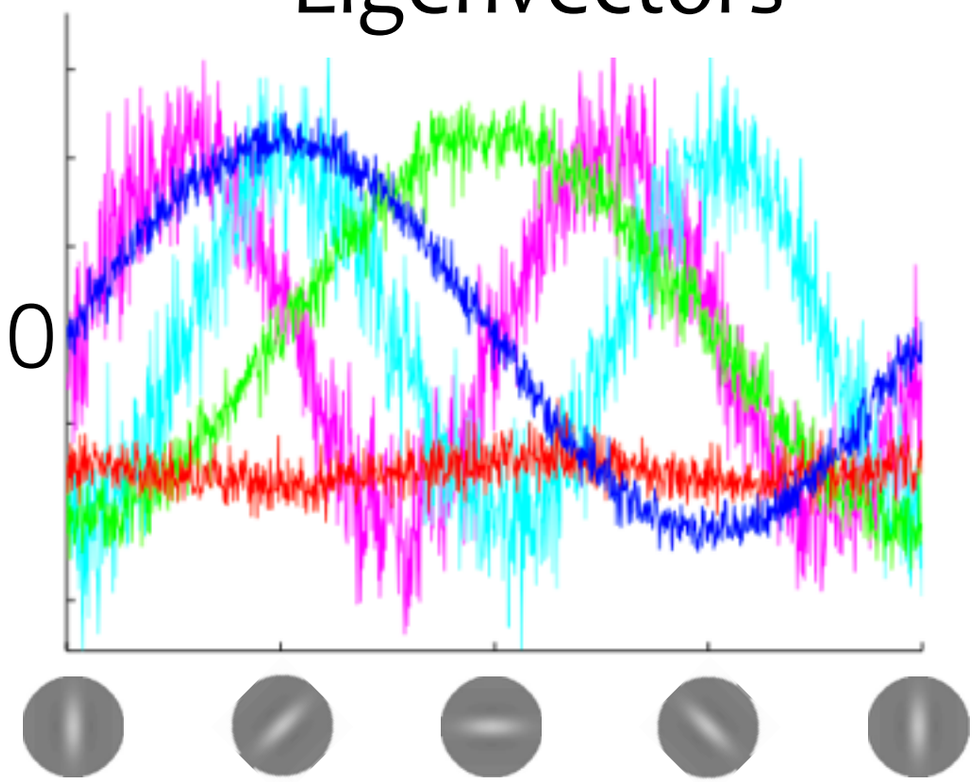

Differential correlations as a function of task

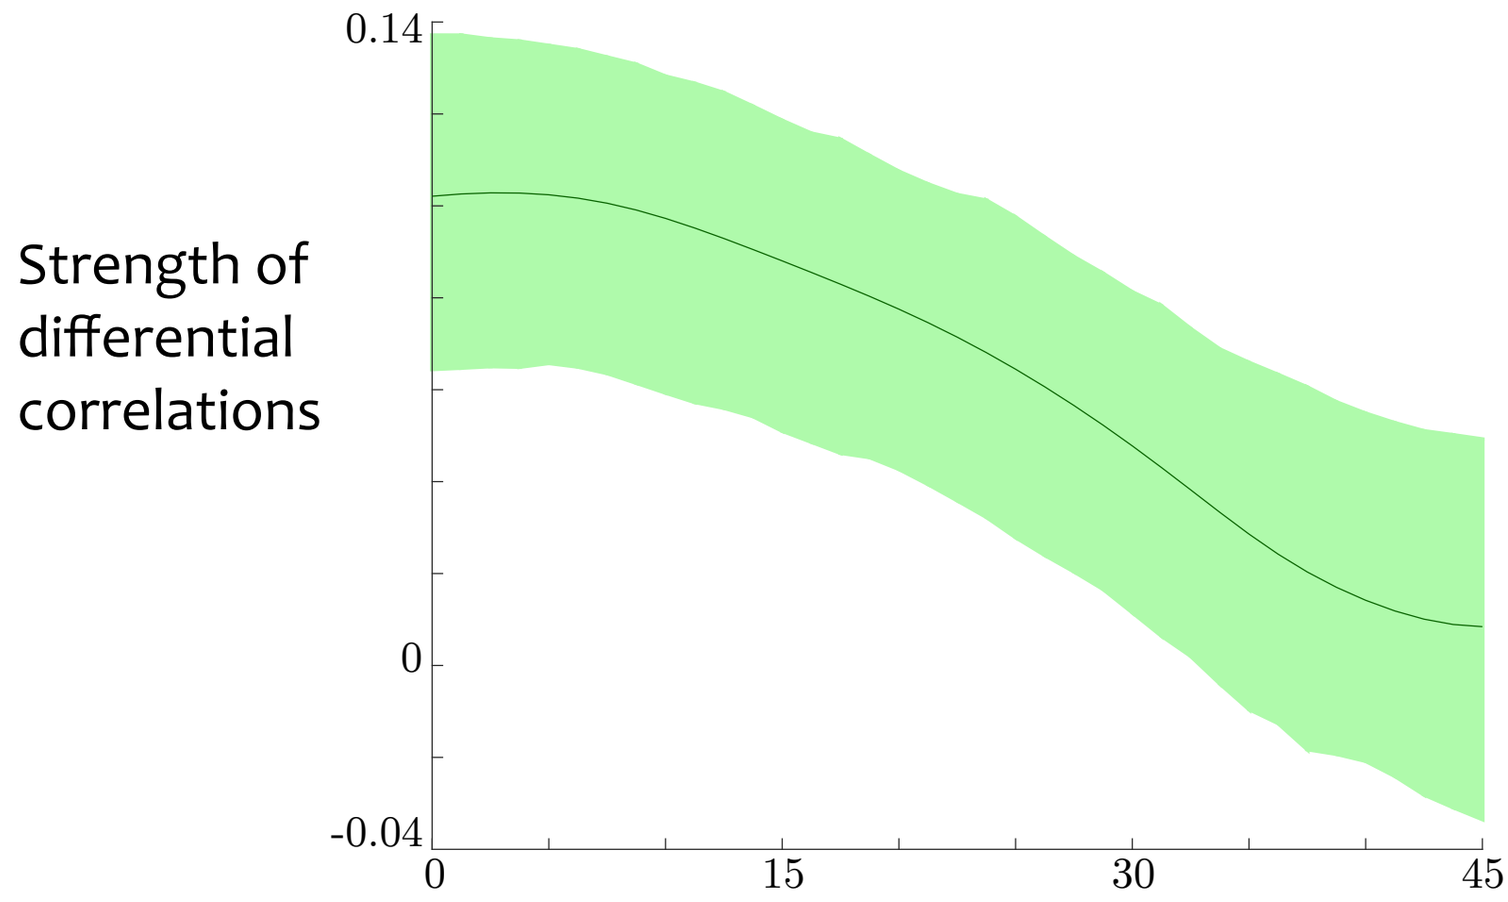

Difference between actual task performed by monkey and task for which prediction was made

Lange, Bondy, Cumming & Haefner, in prep

Learning

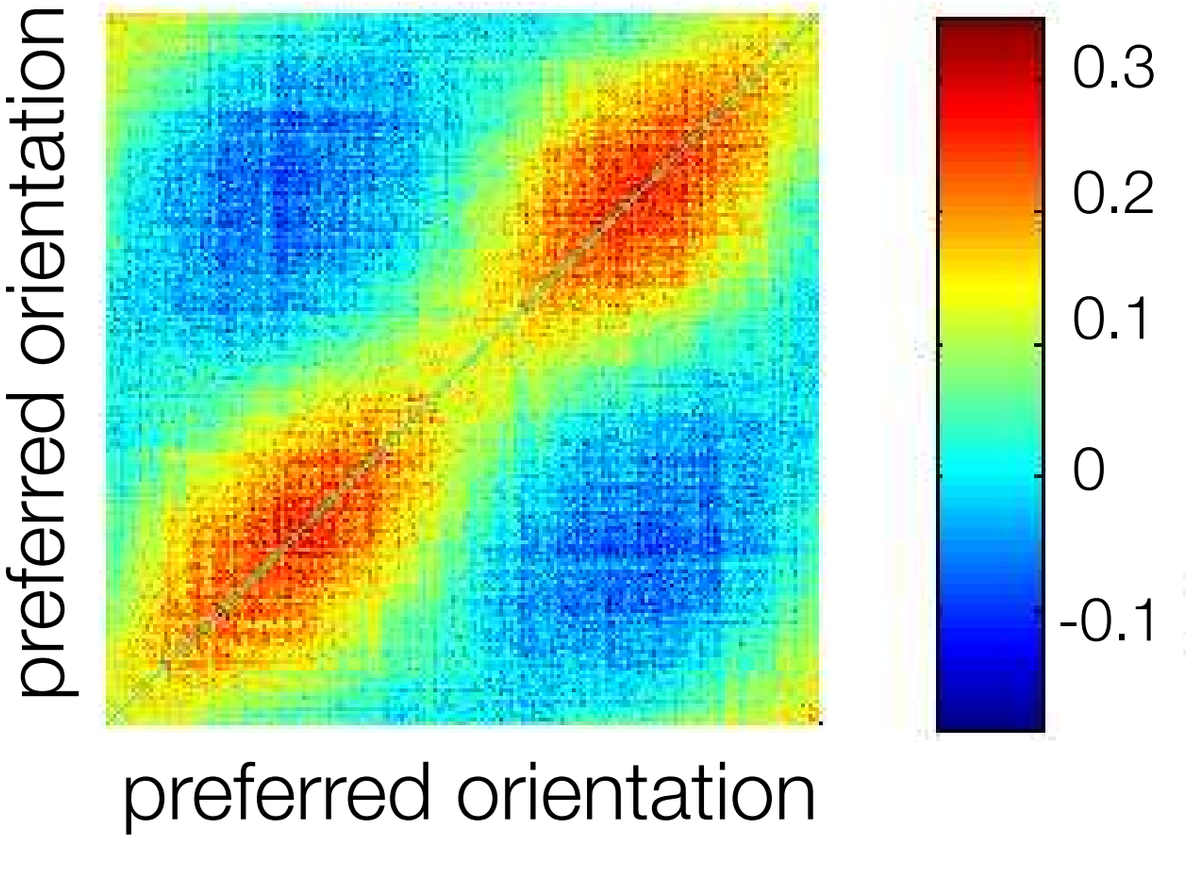

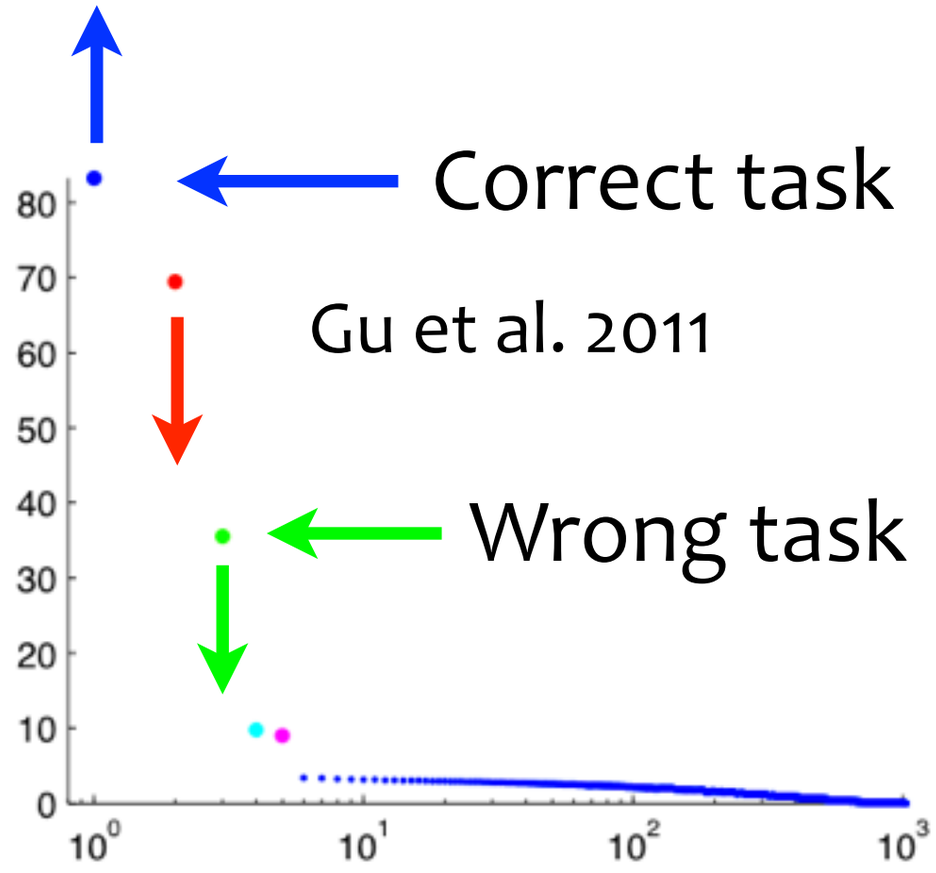

Learning:

- Increasing EV for correct task
- All other EVs decreasing

Lange & Haefner, bioRxiv 2016

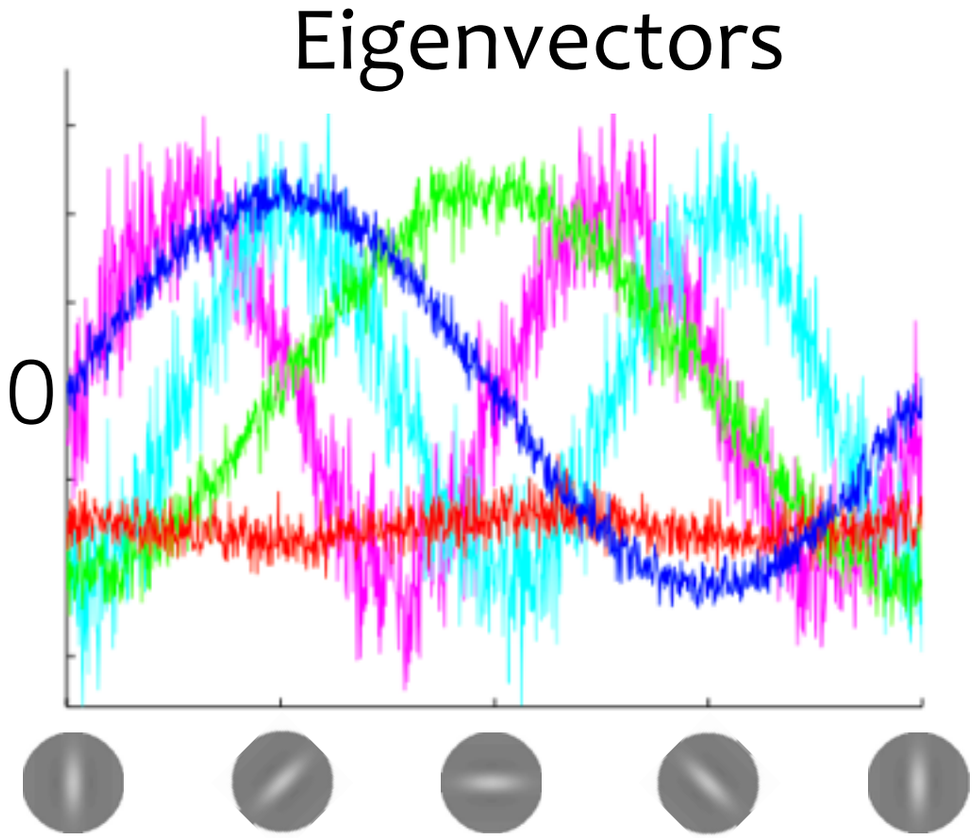

Camille Gomez-Faberge

Rick Born (Harvard)

Richard Lange

- Learning
- Inference over the task/multi-tasking
- Reversibly inactivate top-down connections to V1 (by cooling V2)

Preliminary results (1 monkey)

Cardinal task

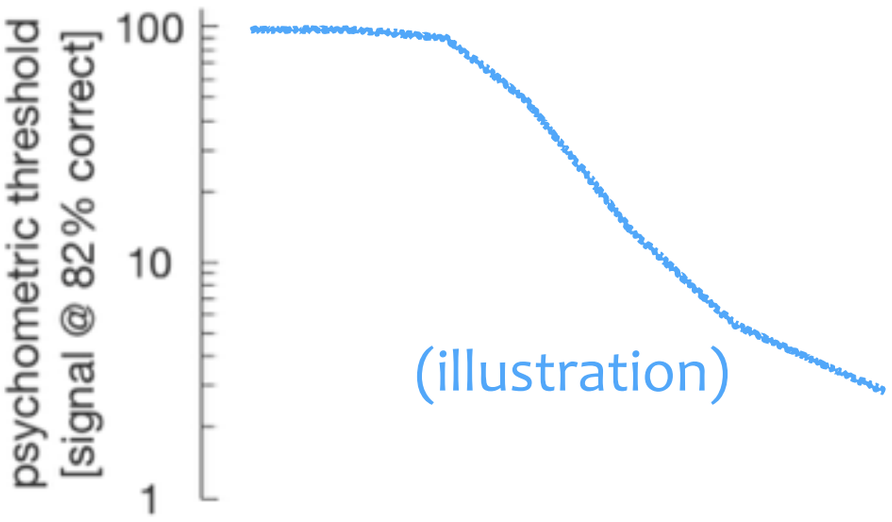

Oblique task

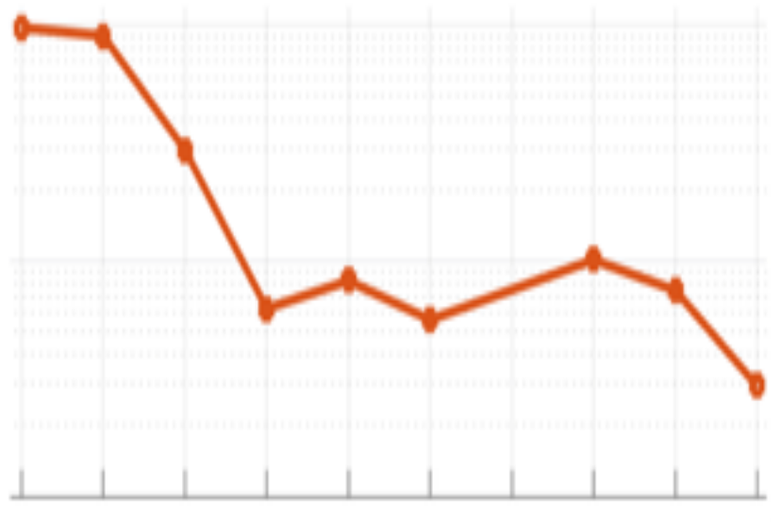

Interleaved

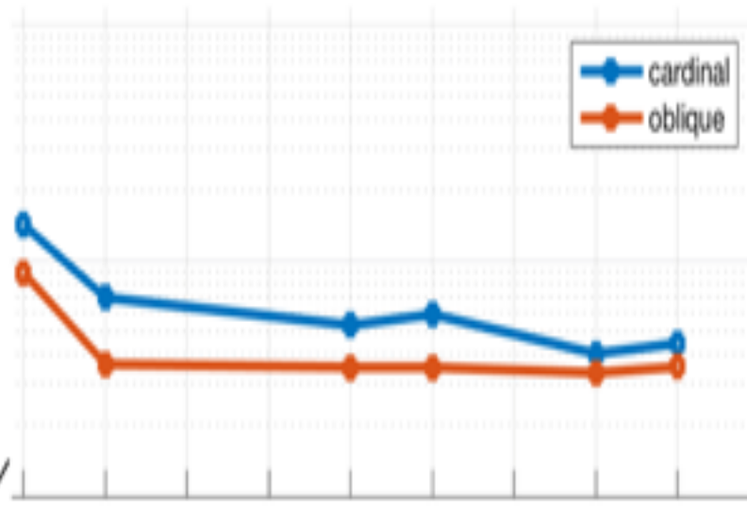

session [day]

Preliminary results! (1 monkey)

Cardinal task

Oblique task

Interleaved

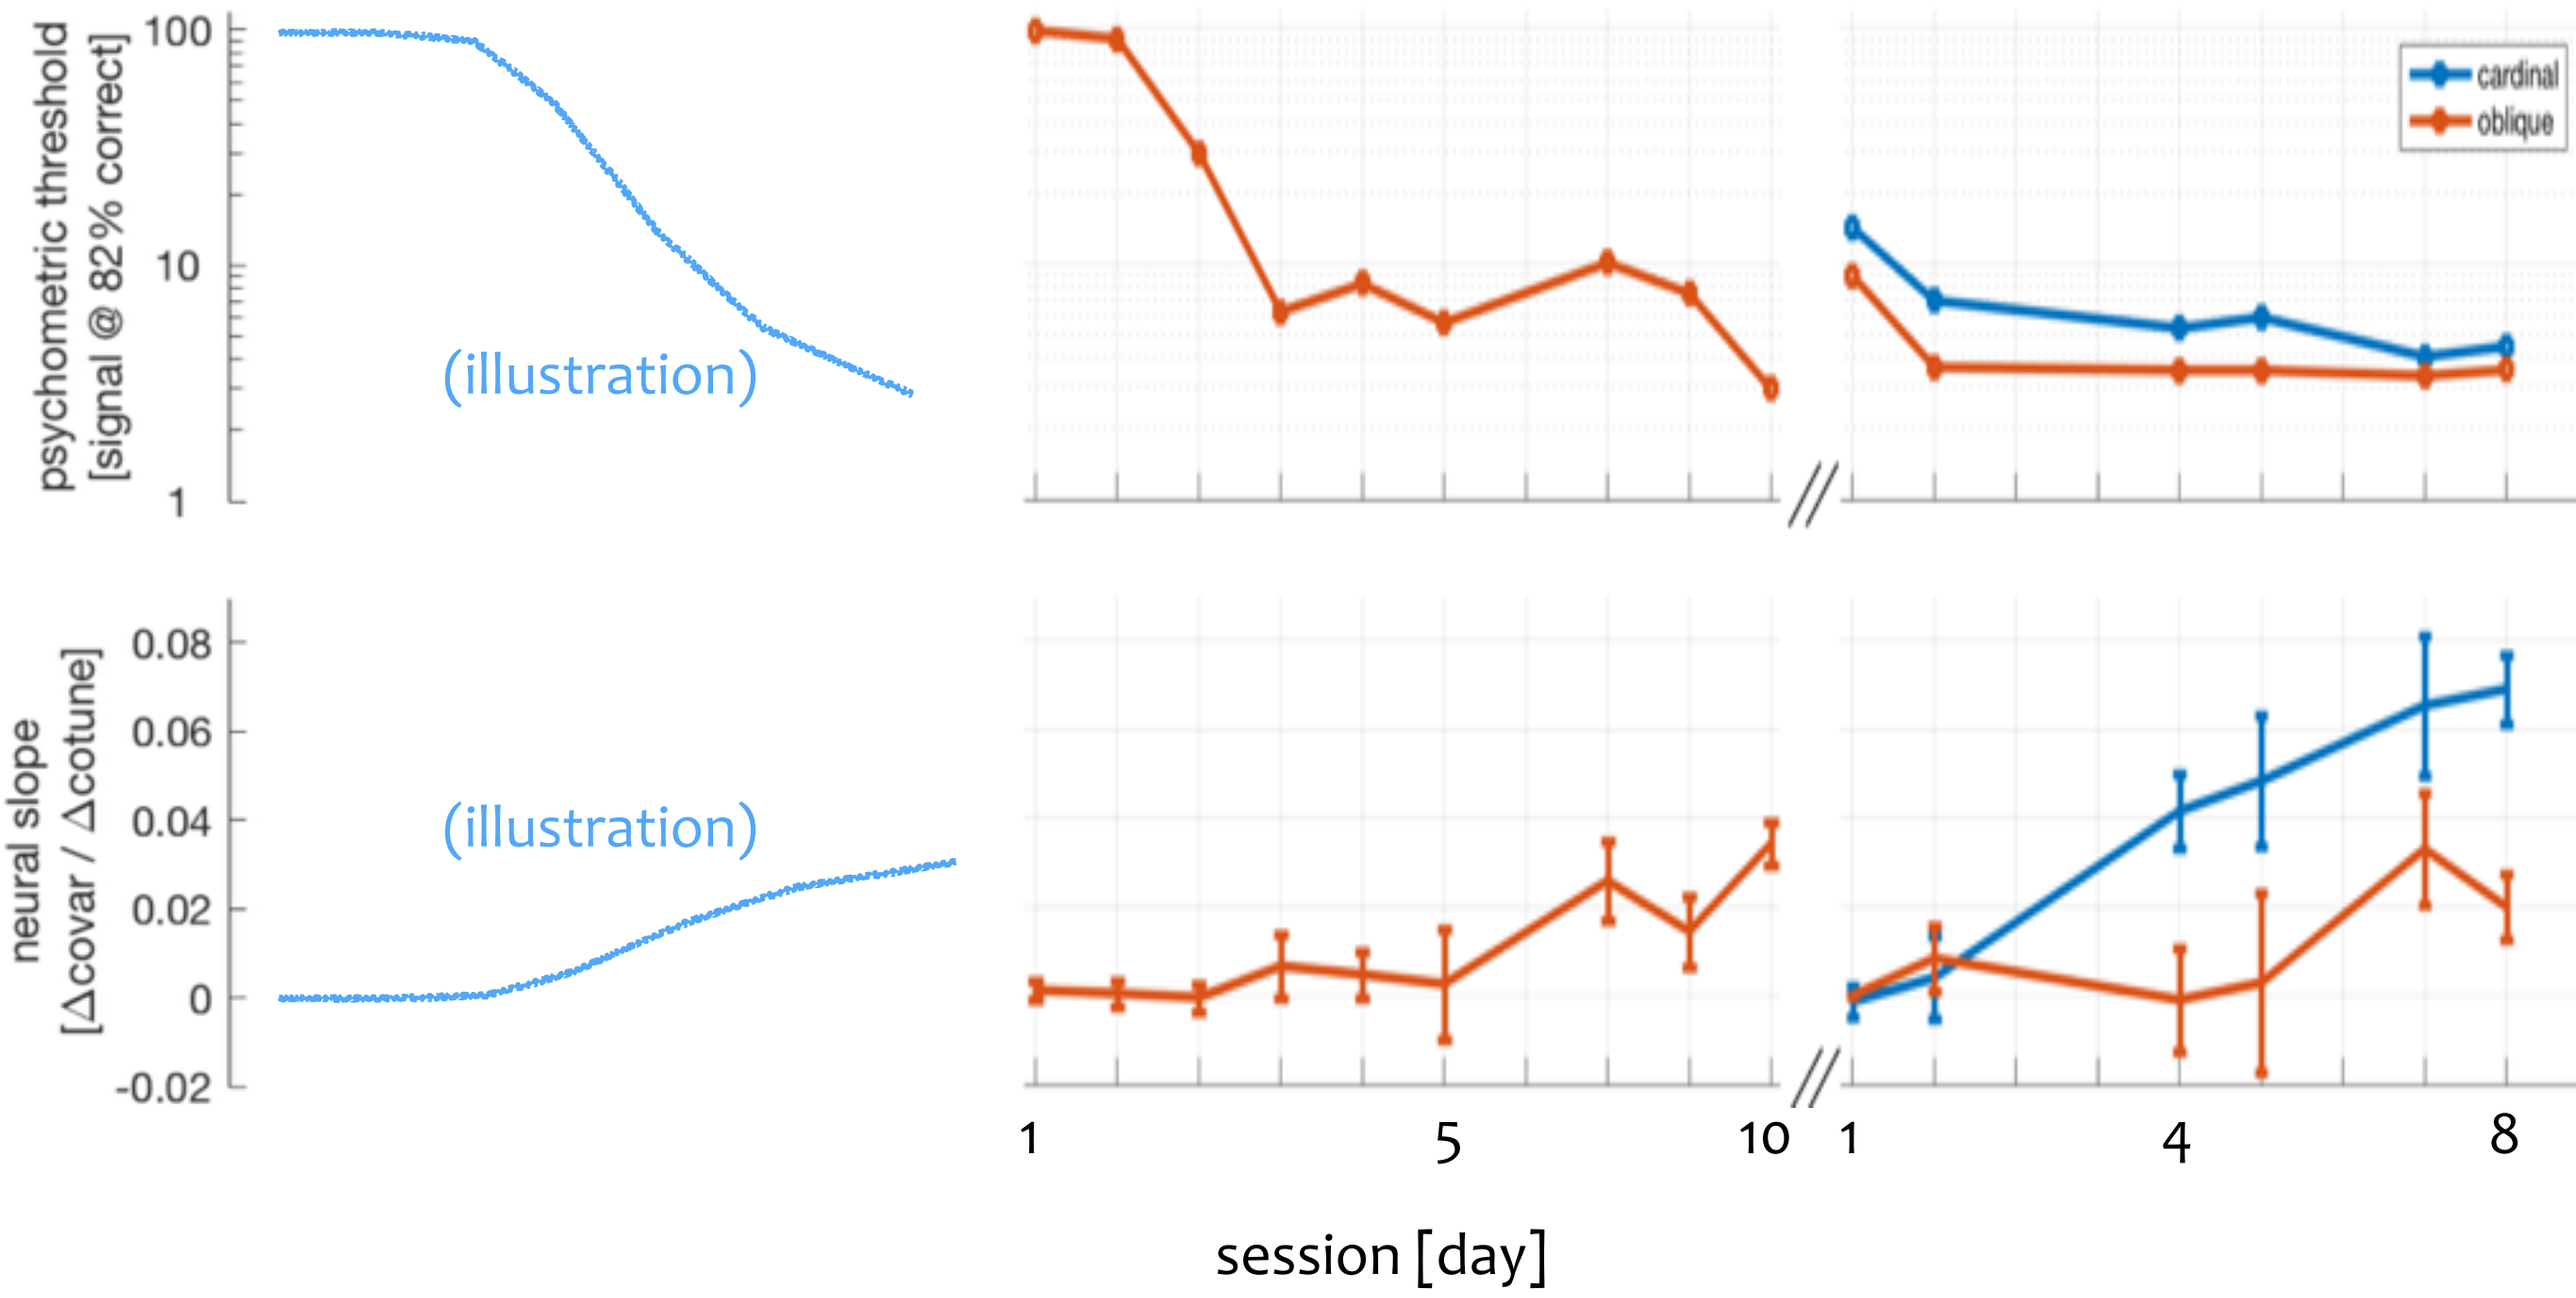

- Next steps: compare correlations & cooling data

Variability in internal states induces statistical signature in neural population responses

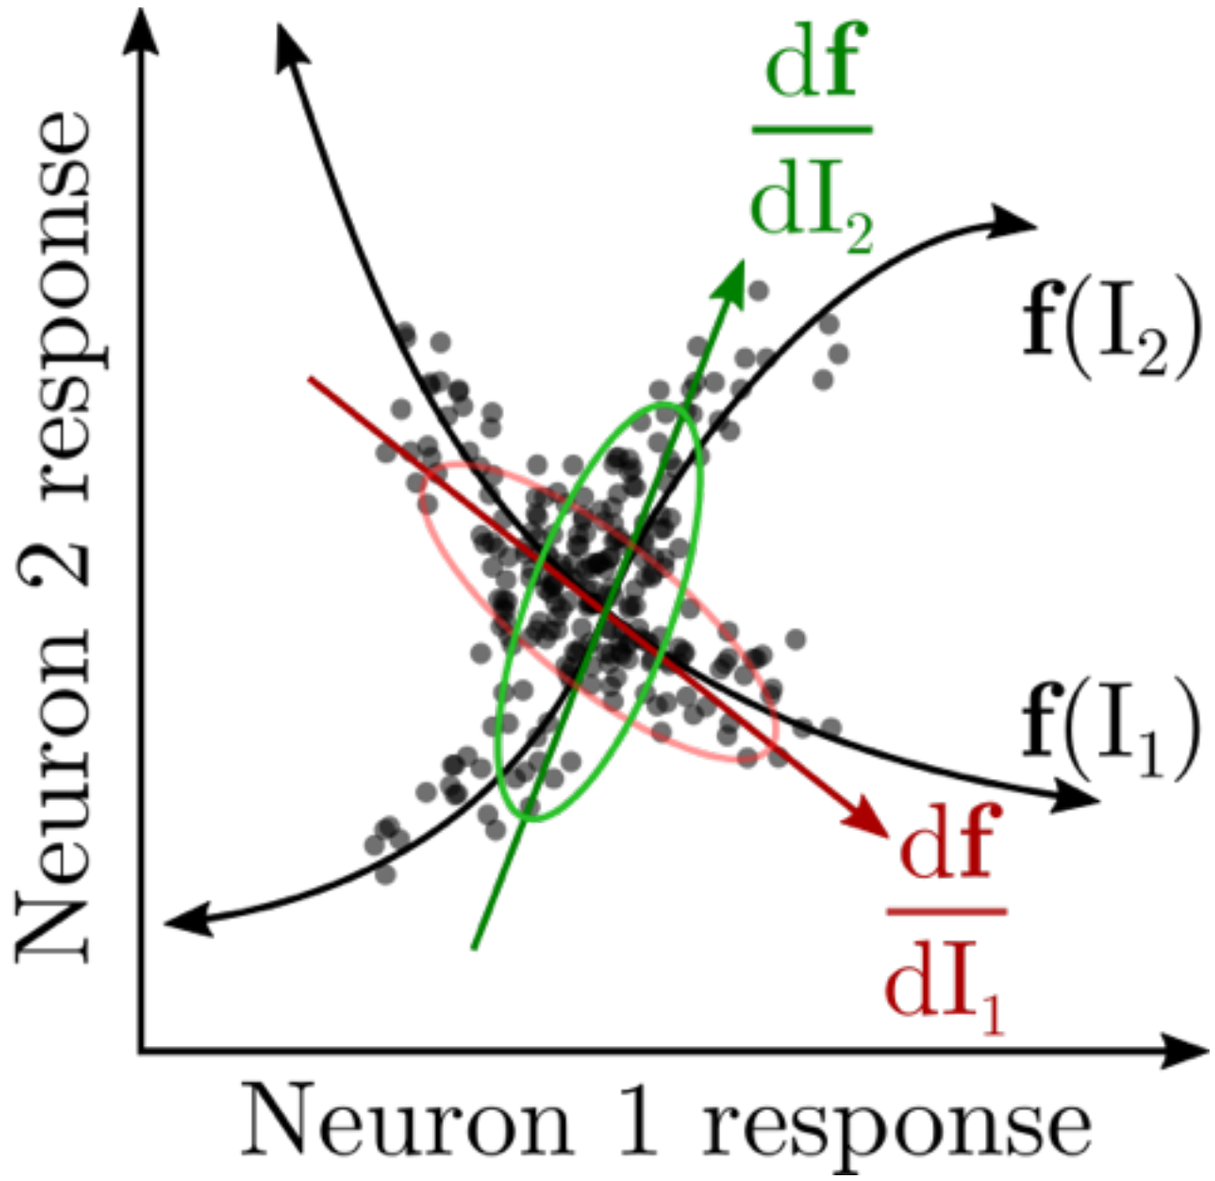

Lange & Haefner, Current Opinion in Neurobiology 2017

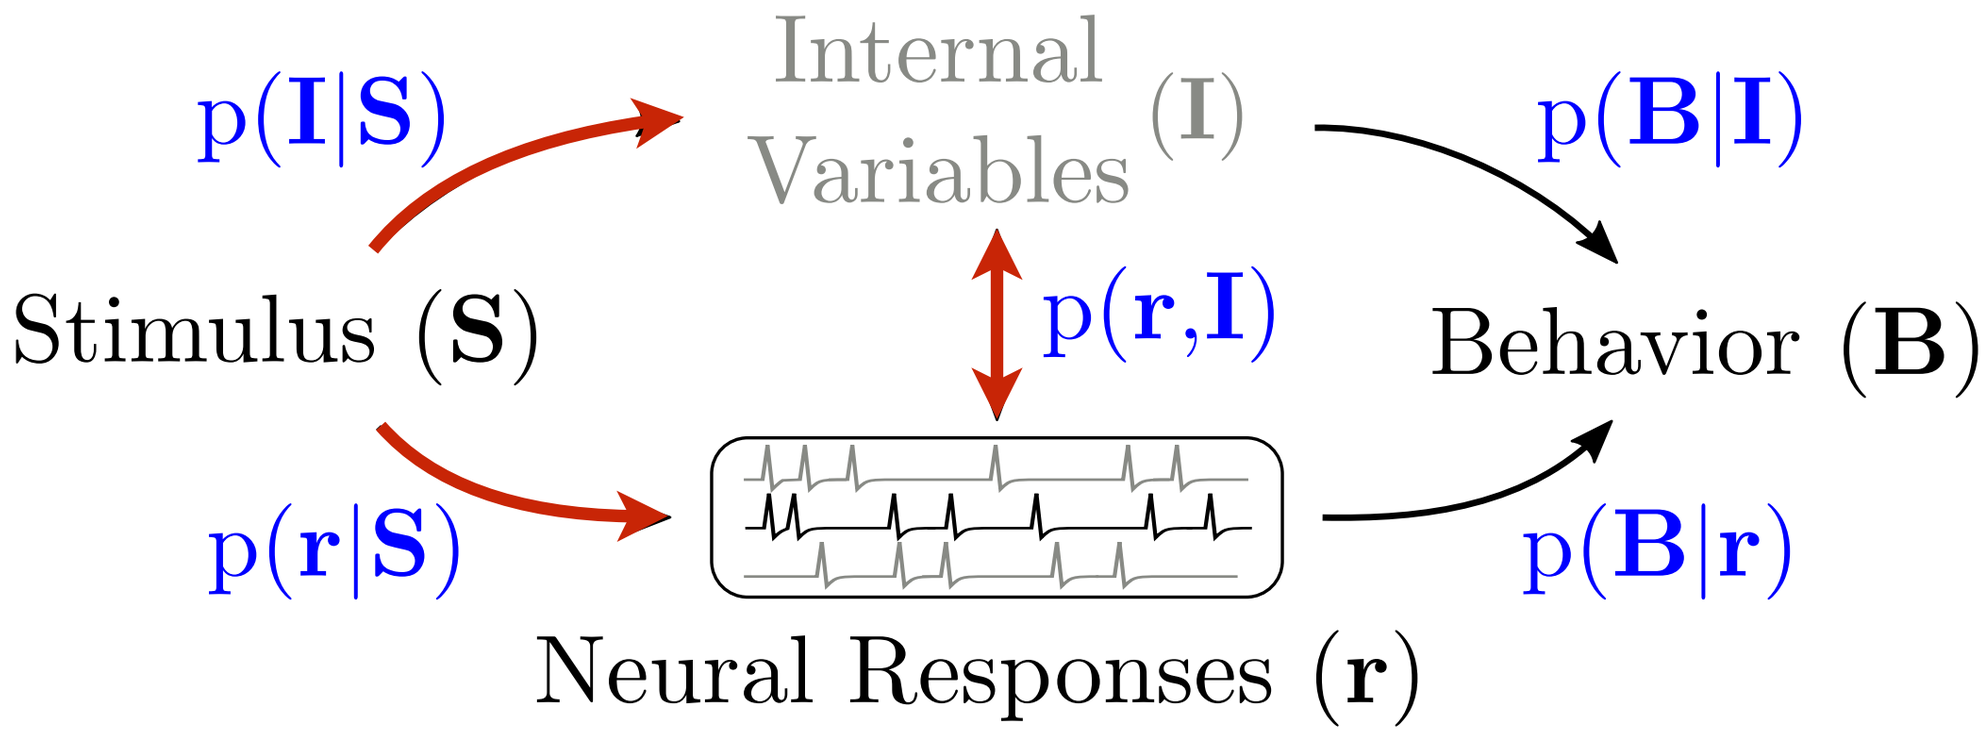

- knowledge of relationships
  stimulus–responses & responses–internal variables

▶ infer relationships: stimulus–internal variables

Lange & Haefner, Current Opinion in Neurobiology 2017

Conclusions so far

- Probabilistic inference by neural sampling makes strong predictions about the task-dependence of neural correlations

- Data so far confirm those predictions

- Population recordings let us reverse-engineer internal beliefs – e.g. track them over learning

- We can interpret them in stimulus space

- Sampling? -> time scales

Decreasing psycho-physical kernel (PK)

Confirmation bias:

- One sample at a time
- Decision based on inferred, not directly observed variables

$$p_t(O) \propto p_{t-1}(O)p(O|\mathbf{x}_t)$$

top-down

$$p_t(\mathbf{x}) \propto p(\mathbf{x}|\mathbf{y}_t)p(\mathbf{x}|O_t)$$

bottom-up

PK

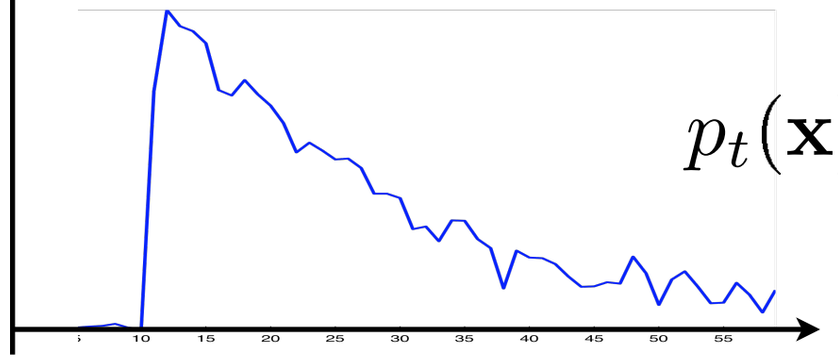

PK

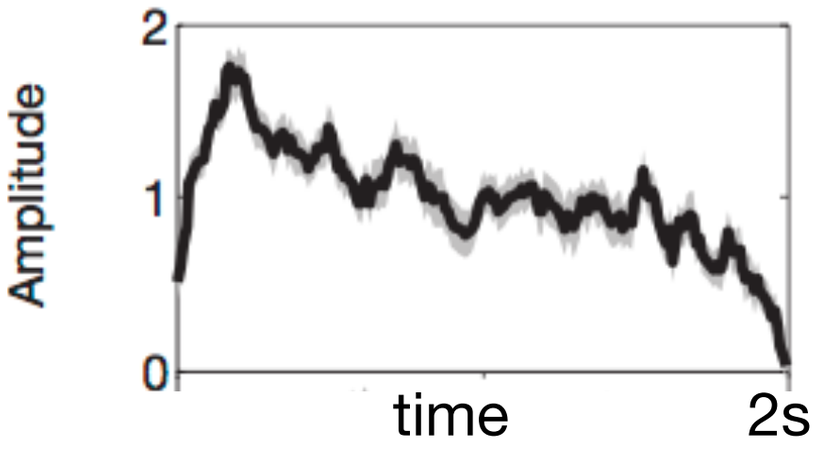

Nienborg & Cumming, Nature 2009

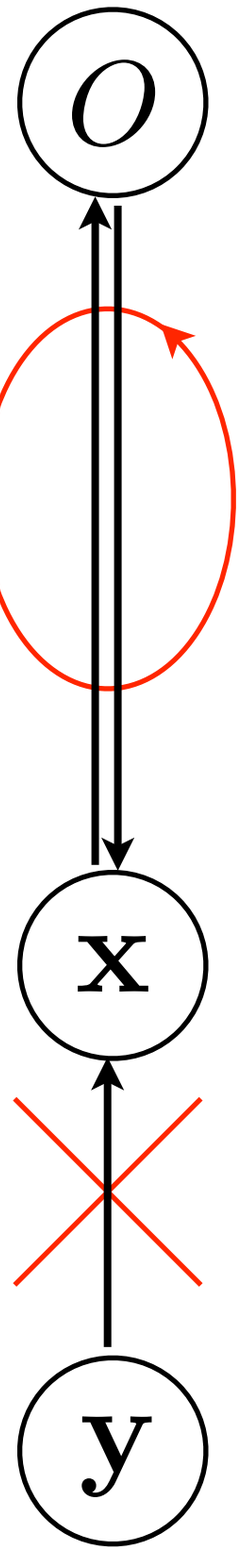

Decision variable

(Grating variables omitted for simplicity)

Gabor variables

Image variable

Why constant psycho-physical kernel?

$$p_t(O) \propto p_{t-1}(O)p(O|\mathbf{x}_t)$$

$$p_t(\mathbf{x}) \propto p(\mathbf{x}|\mathbf{y}_t)p(\mathbf{x}|O_t)$$

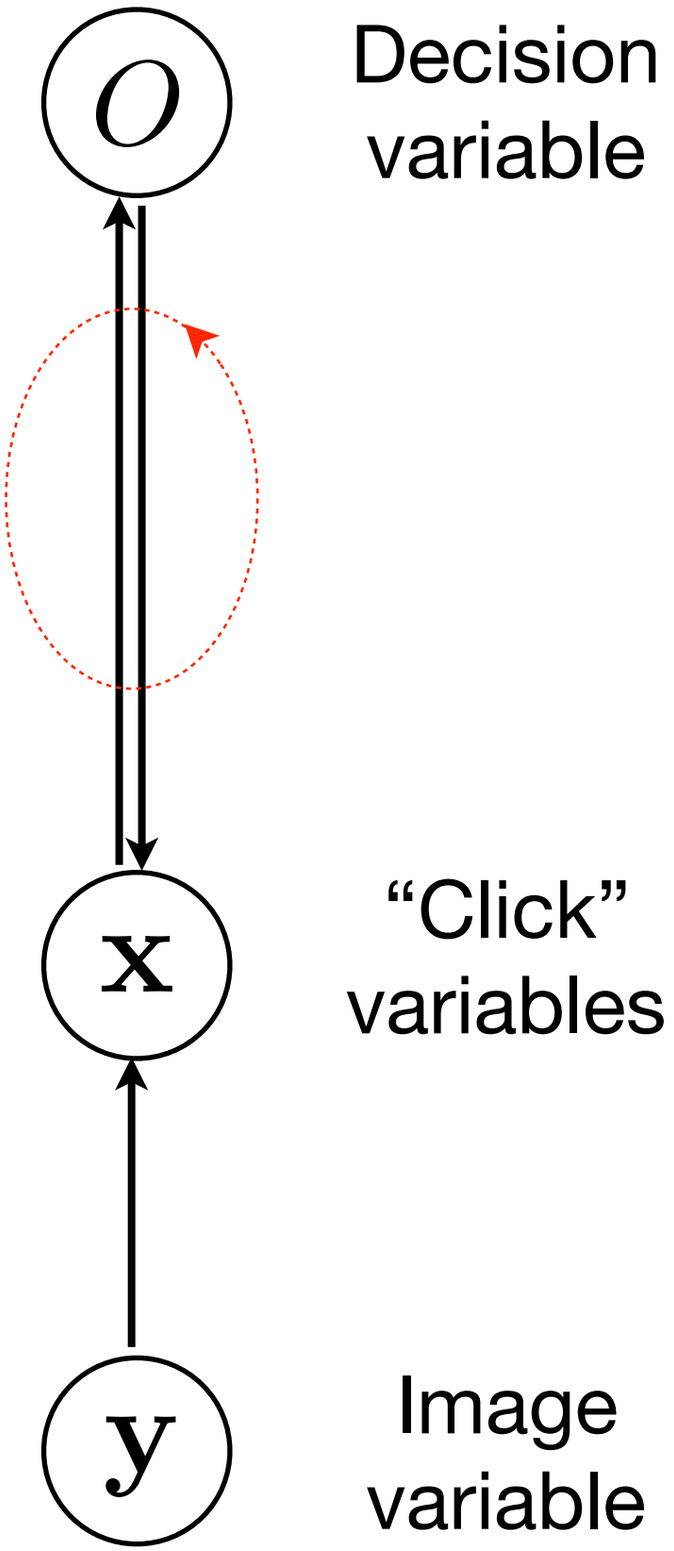

Poisson click task

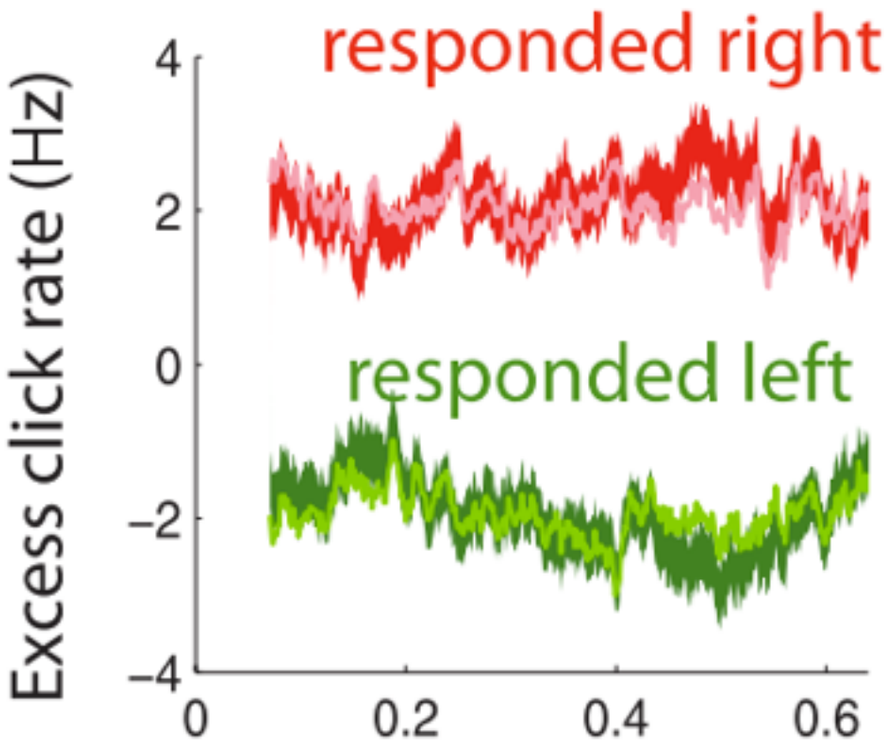

Brunton et al., Science 2013

Each click far above threshold, i.e. likelihood dominates prior.

Prediction: Soft clicks -> decreasing PK

Confirmation bias project

Richard Lange

Ankani Chattoraj

Two kinds of information

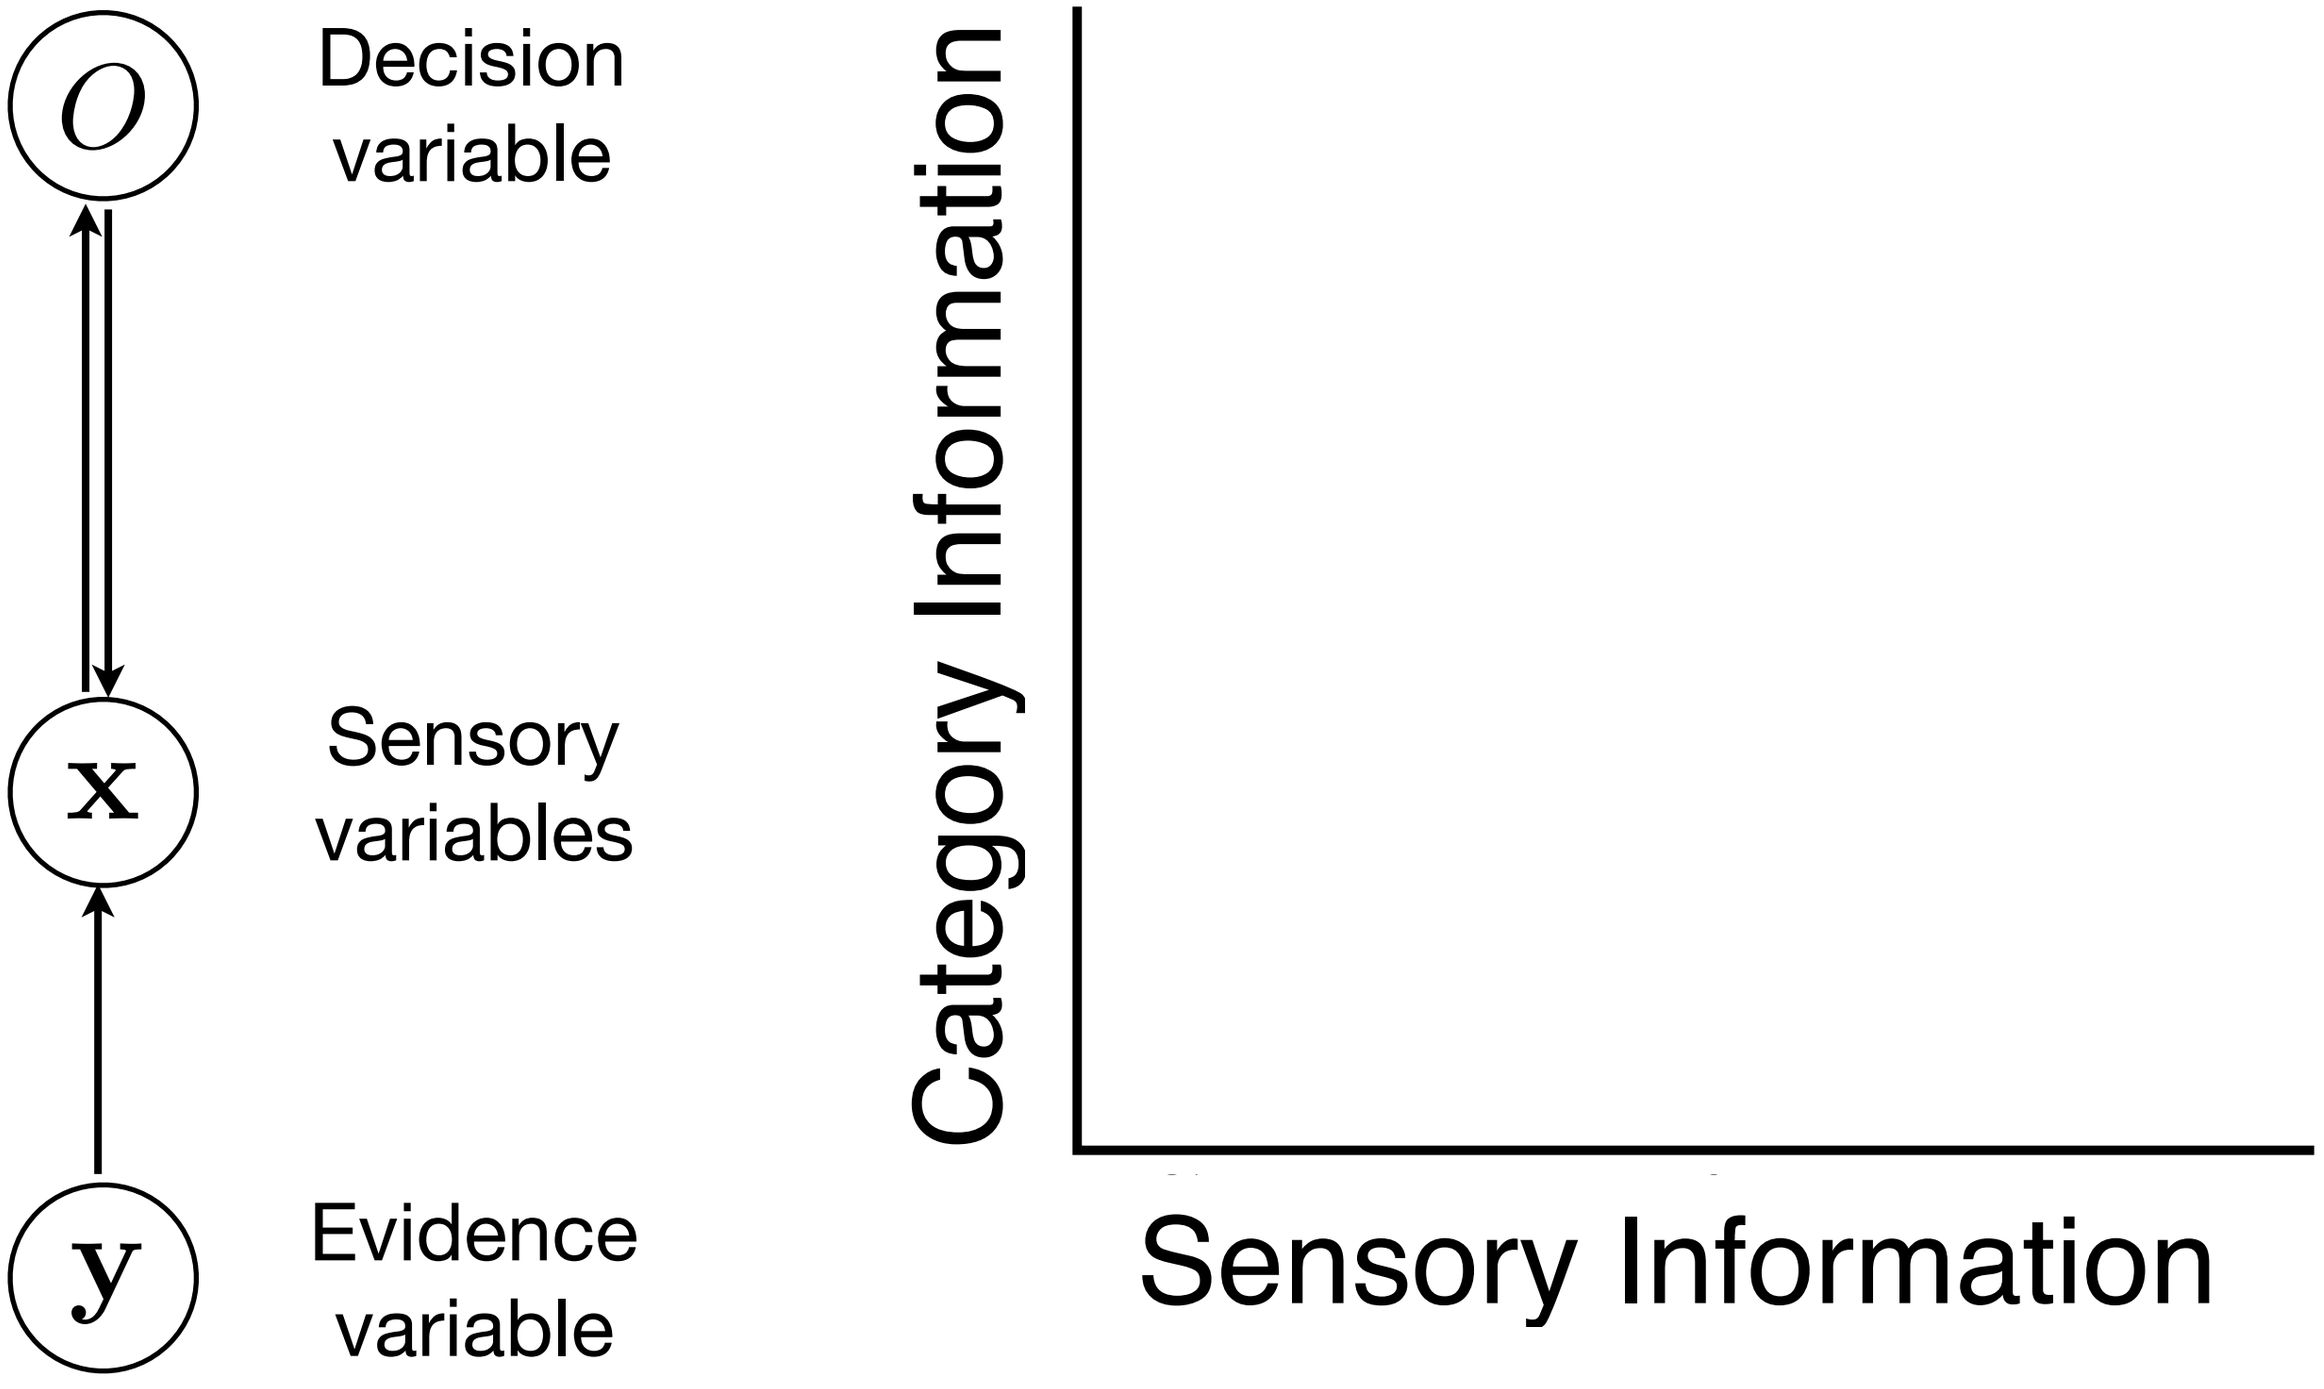

Two kinds of information

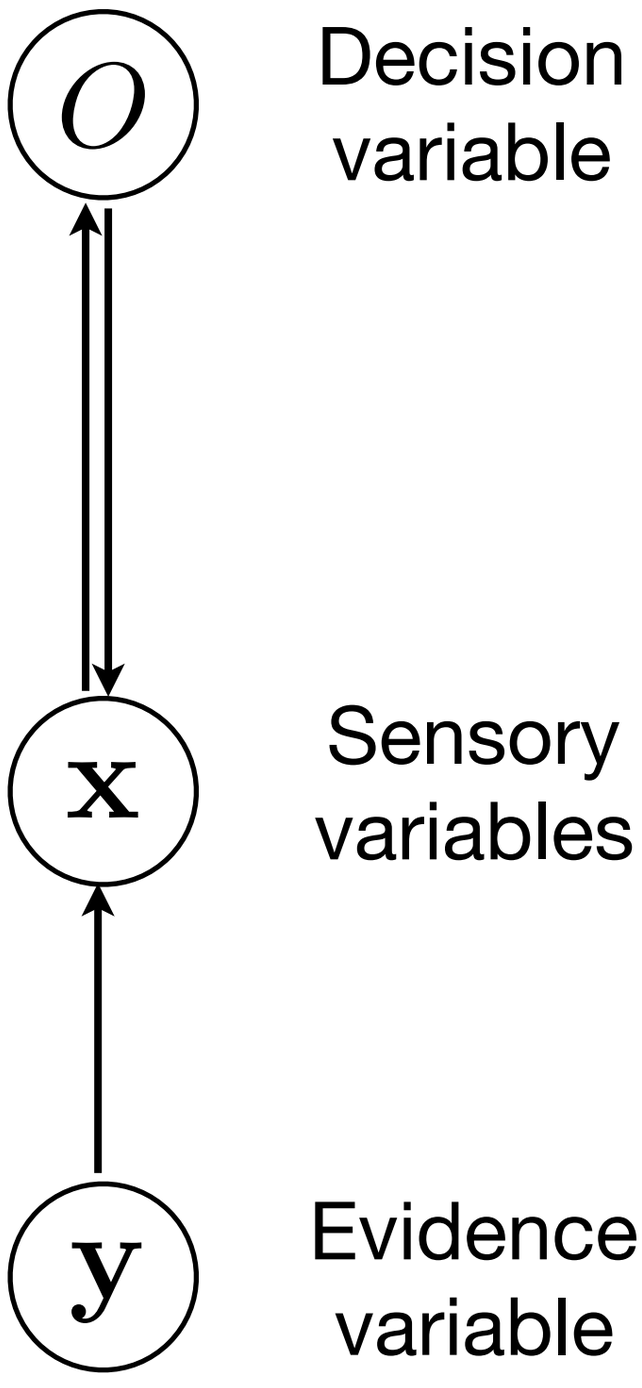

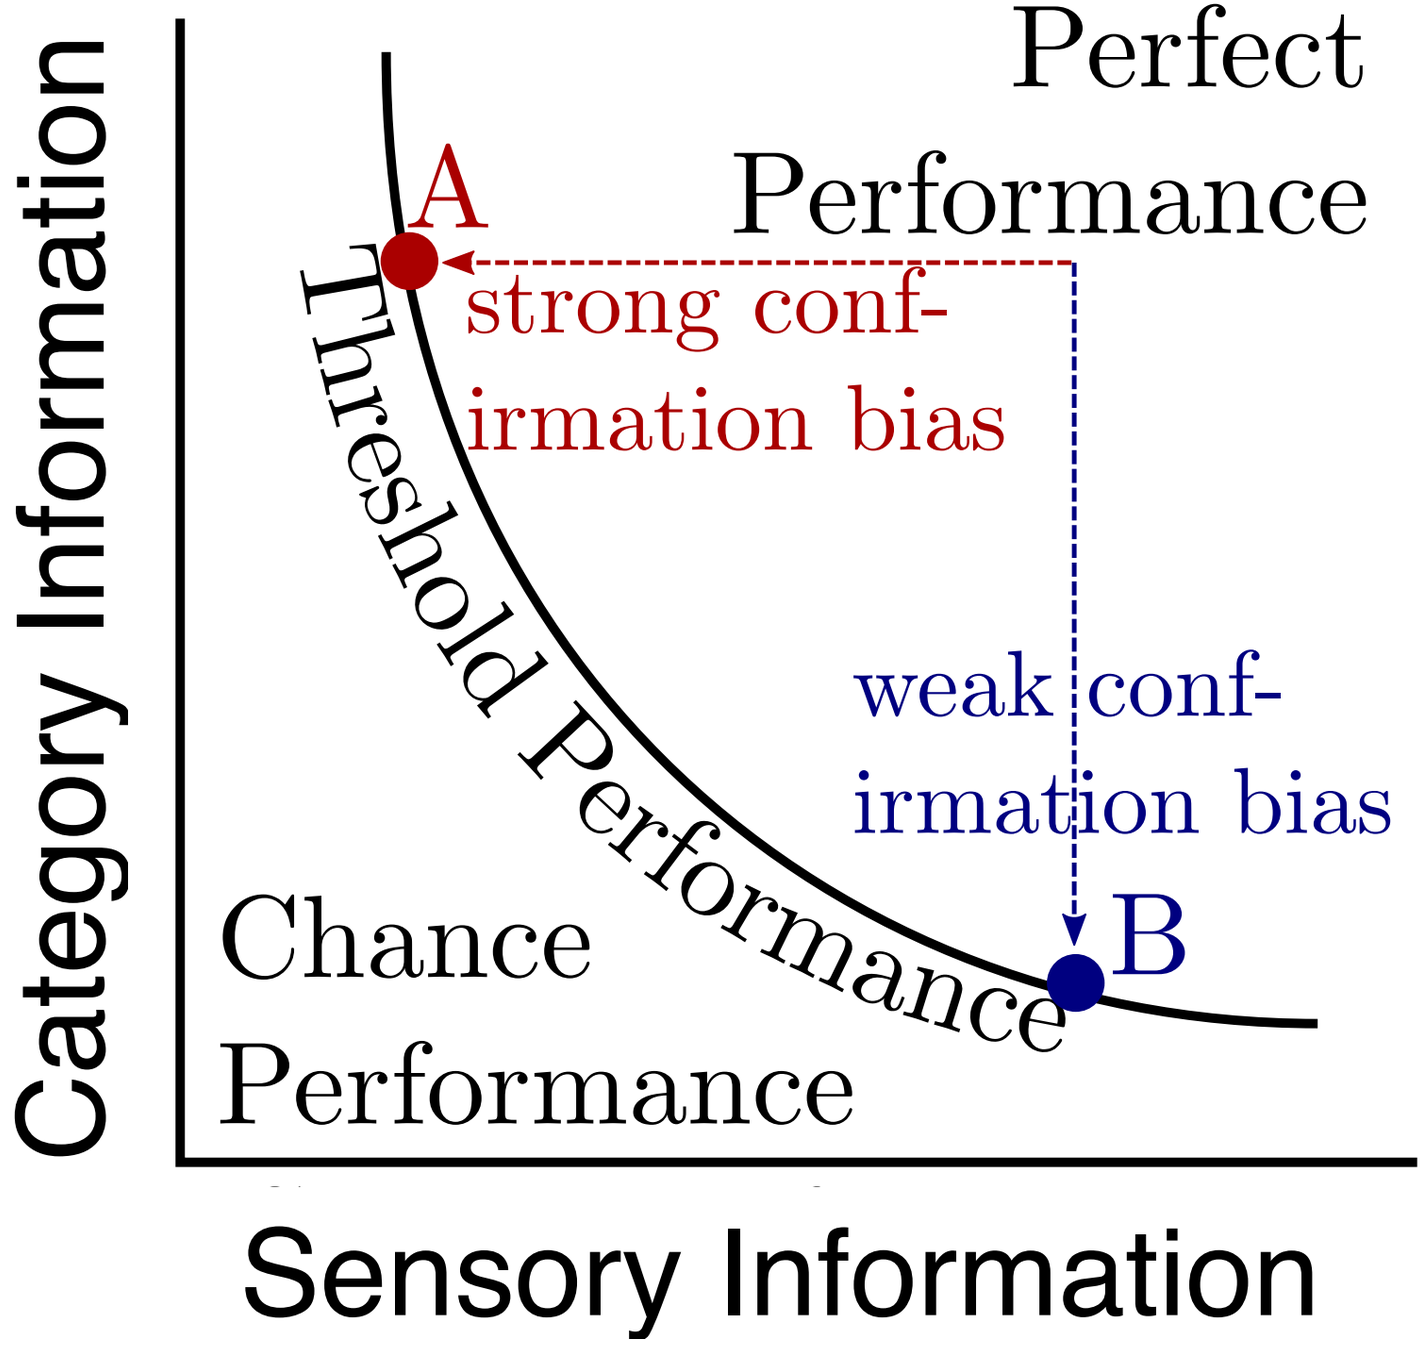

Two different tasks

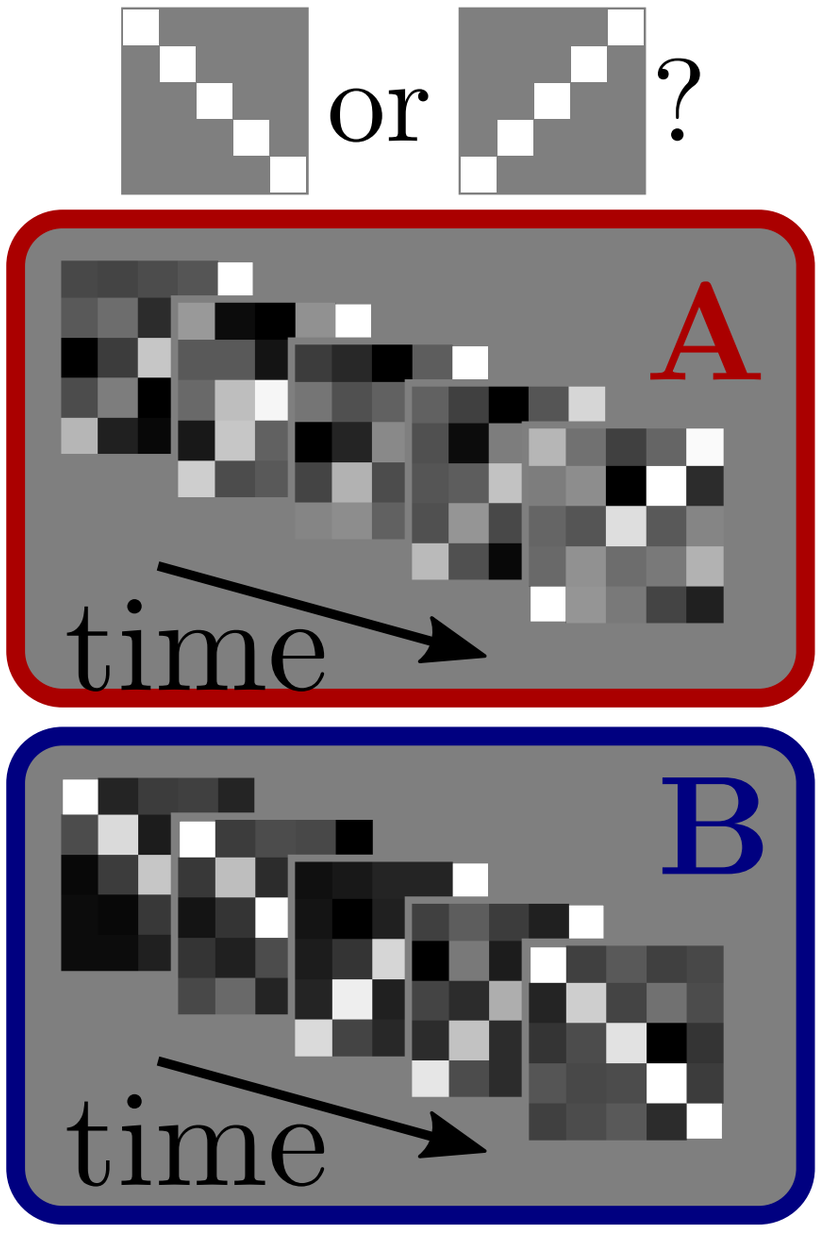

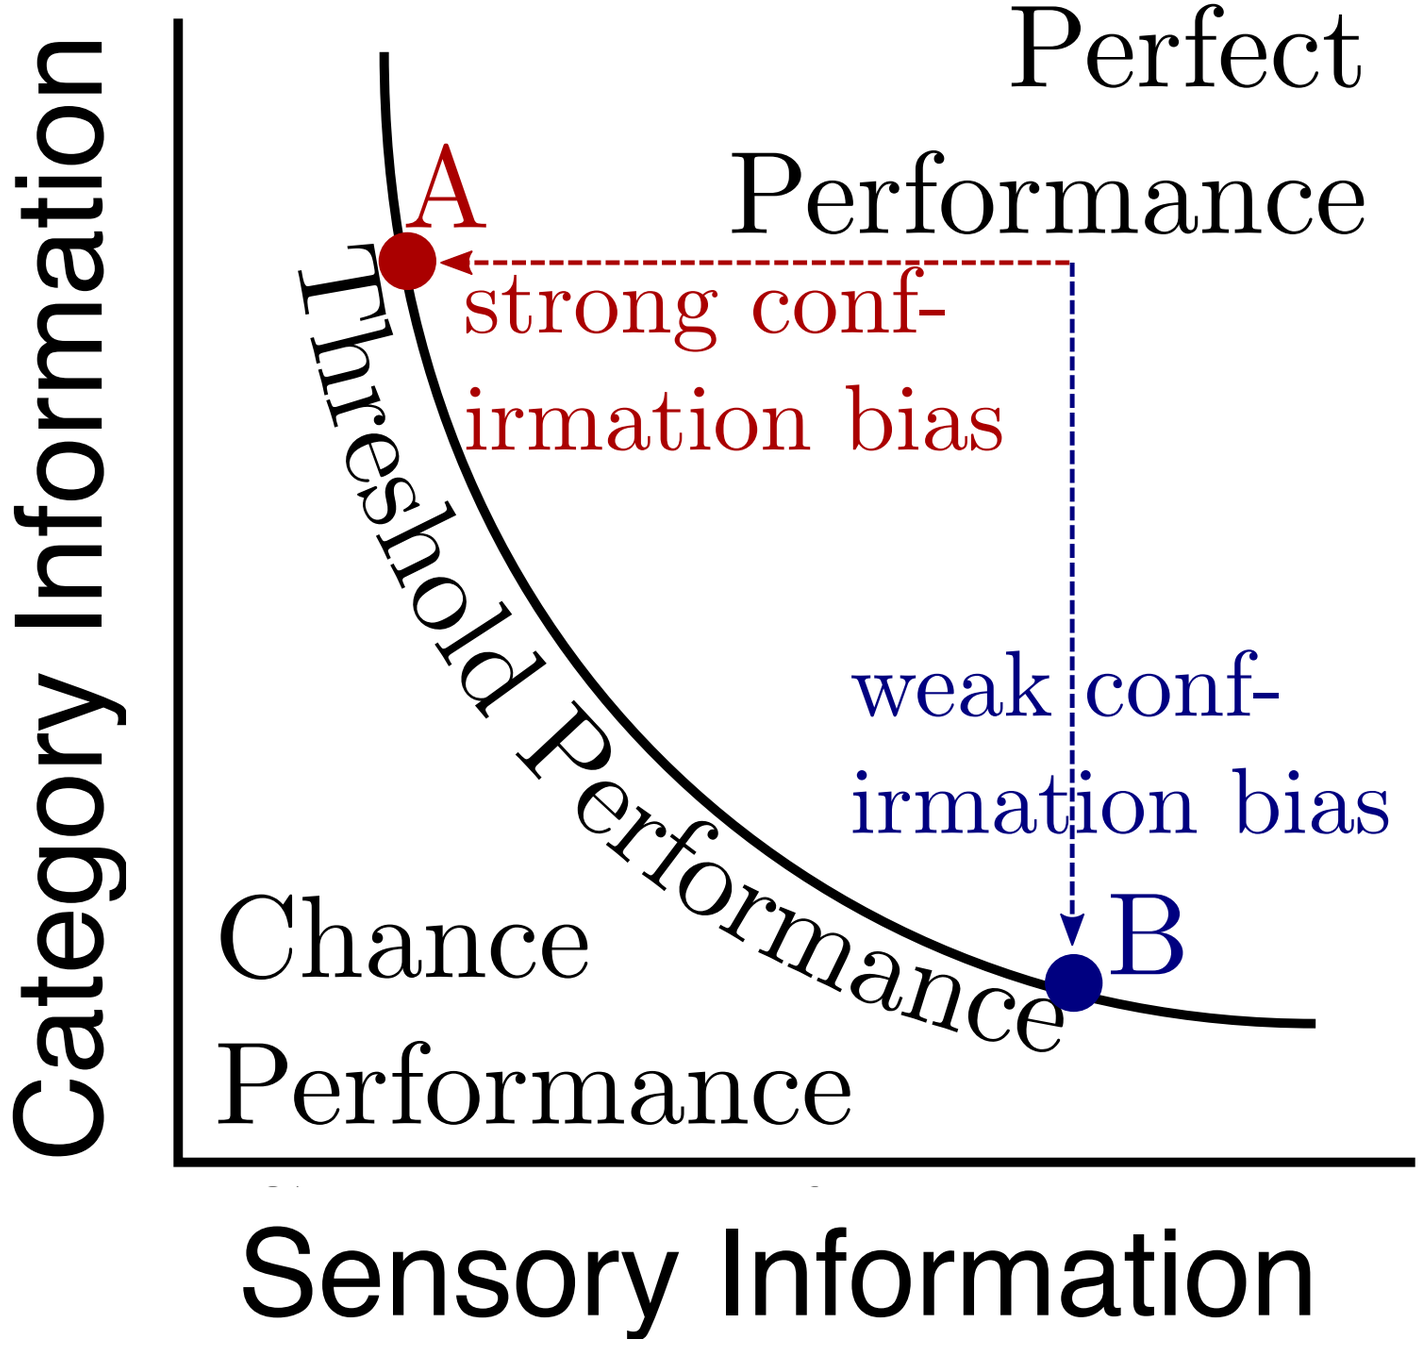

Task A: confirmation bias

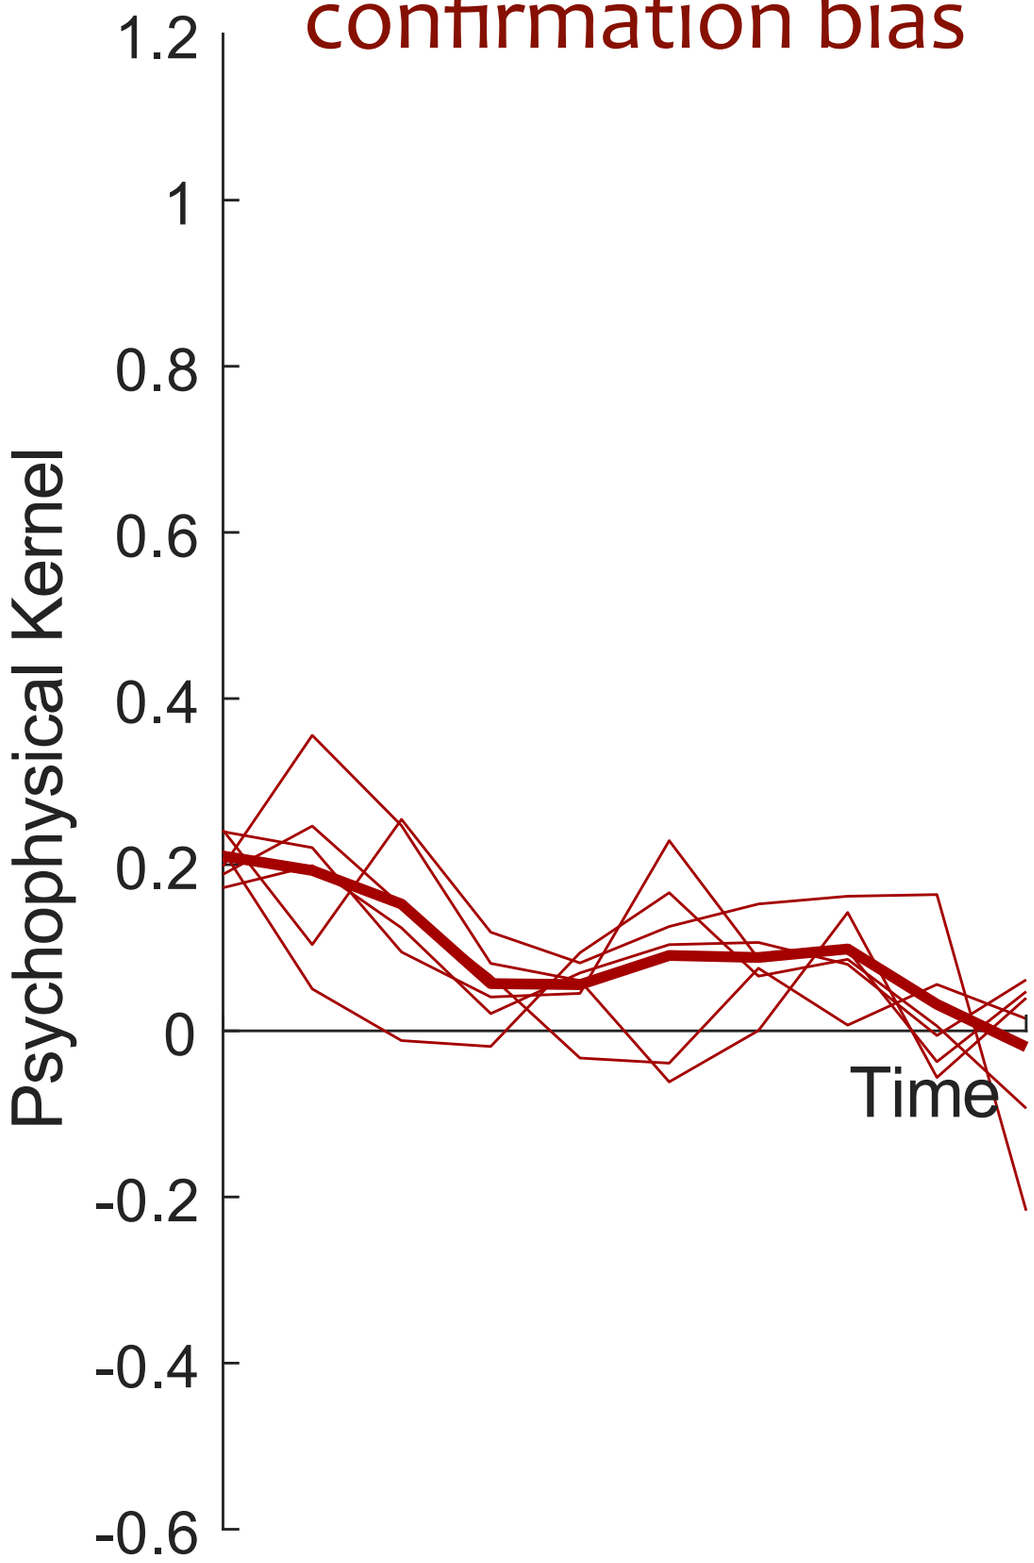

Task B: no confirmation bias

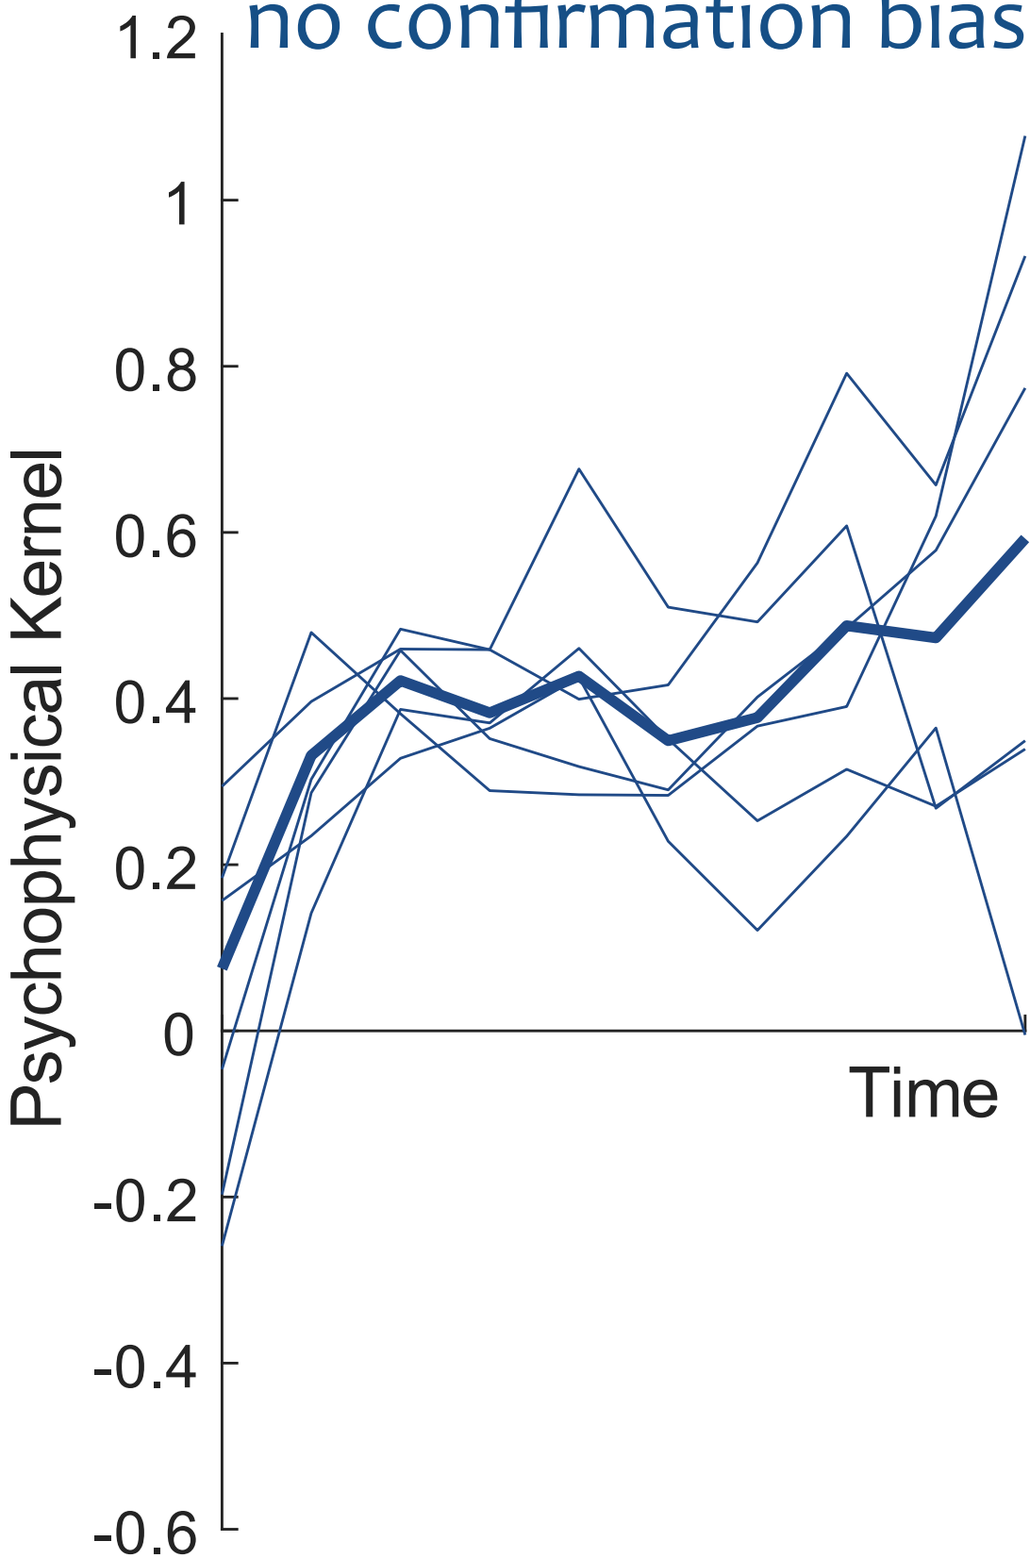

Conclusion

- Approximate nature of inference process:
  - Infer task-model learnt by the brain from population responses
  - Explain confirmation bias -> poster!

Acknowledgements

Lab

Richard Lange
Ankani Chattoraj
Sabya Shivkumar

Theoretical collaborators

Pietro Berkes
Jozsef Fiser (CEU)

Experimental collaborators

Adrian Bondy (Princeton)
Bruce Cumming (NIH)

Camille Gomez-Faberge (Harvard)
Rick Born (Harvard)

University of Rochester (NY)

Looking for postdocs!

BACKUP

Predictions

1st order

Haefner et al, NatNeuro 2013

$\Delta_{\text{choice}}\bar{r}_i$

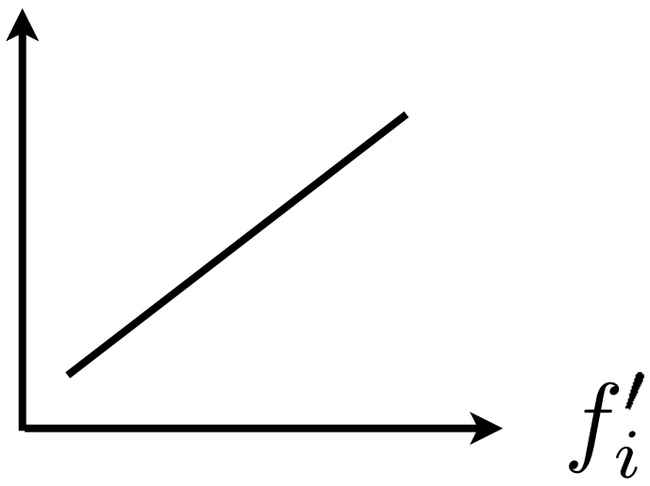

f_i'

2nd order

$\text{cov}(r_i, r_j)$

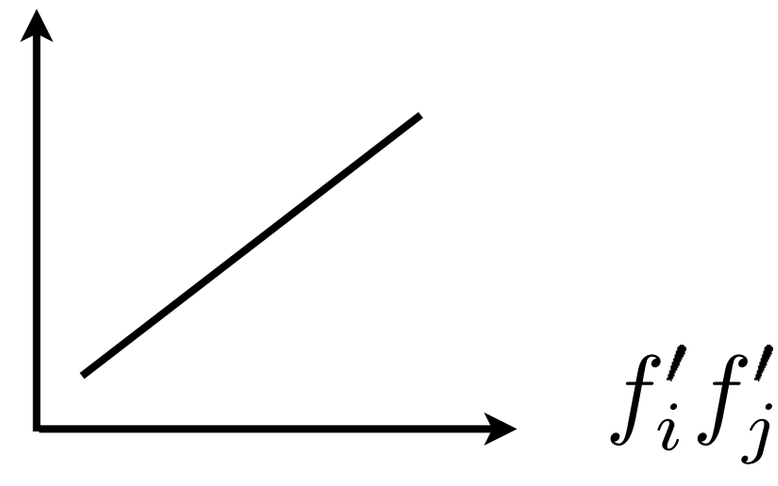

$f_i' f_j'$

Nth order

Odd *N*: choice-triggered
(like 1st order)
Even *N*: across all trials
(like 2nd order)

$(\cdots)$

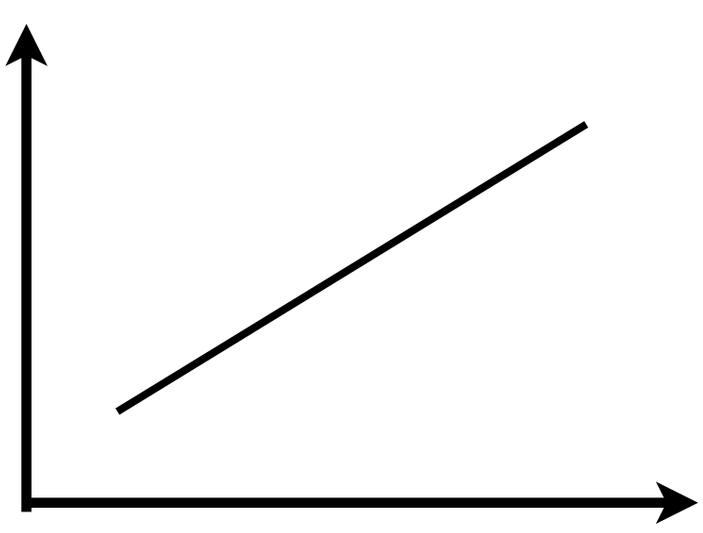

$\prod_{i=1}^{N} f_i'$

1) Brain's computation

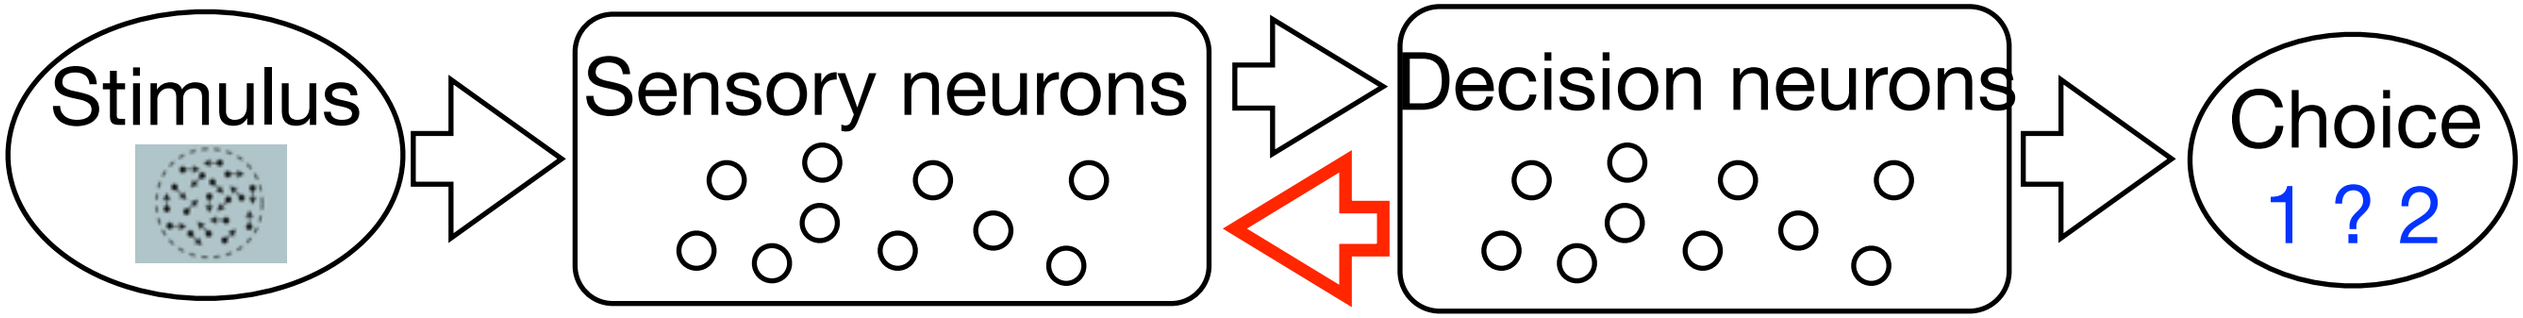

1) Brain's computation

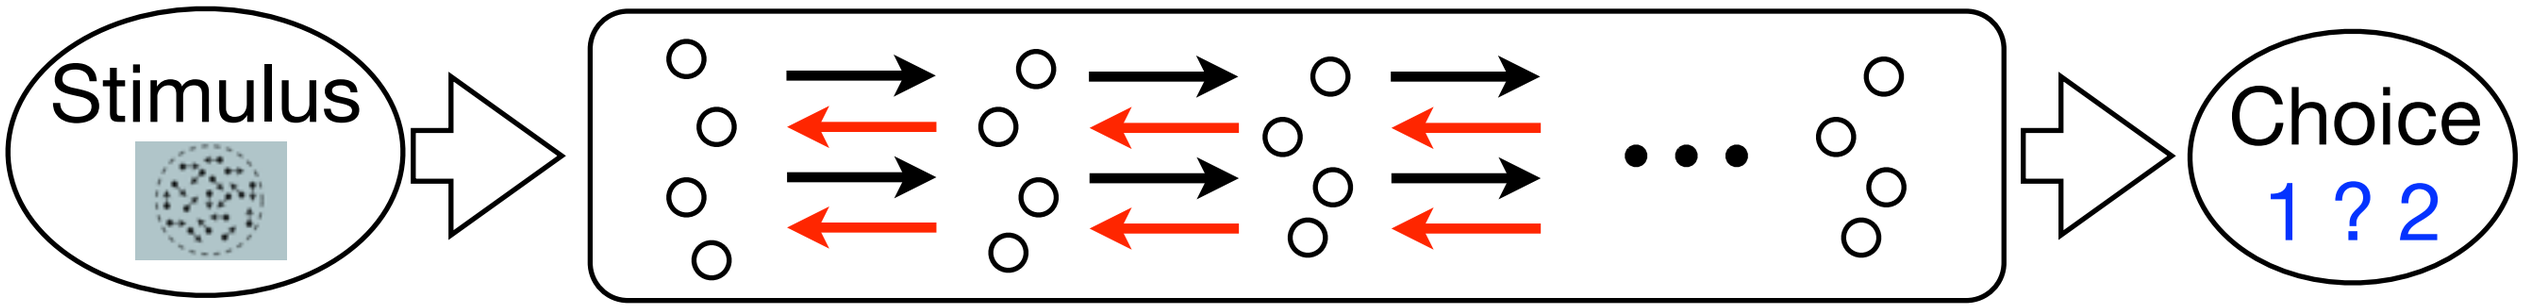

External inputs ⇨ Posterior beliefs ⇦ Prior knowledge

1) Brain's computation

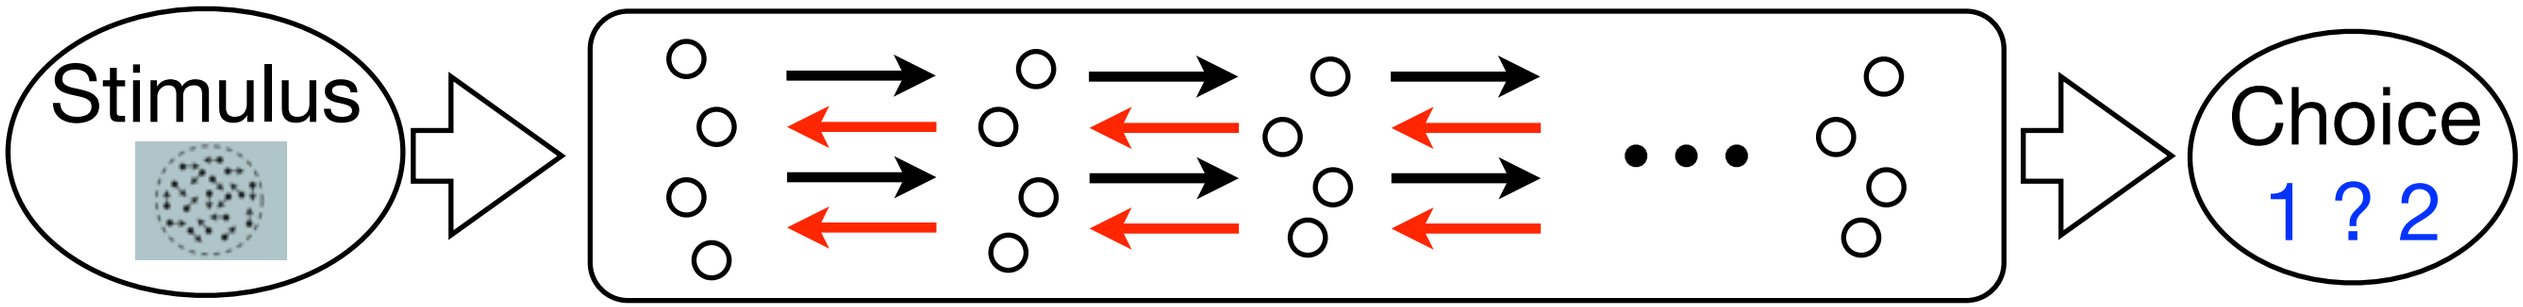

External inputs ⇨ Posterior beliefs ⇨ Prior knowledge

⇦

▶ Confirmation bias: Poster 1

1) Brain's computation

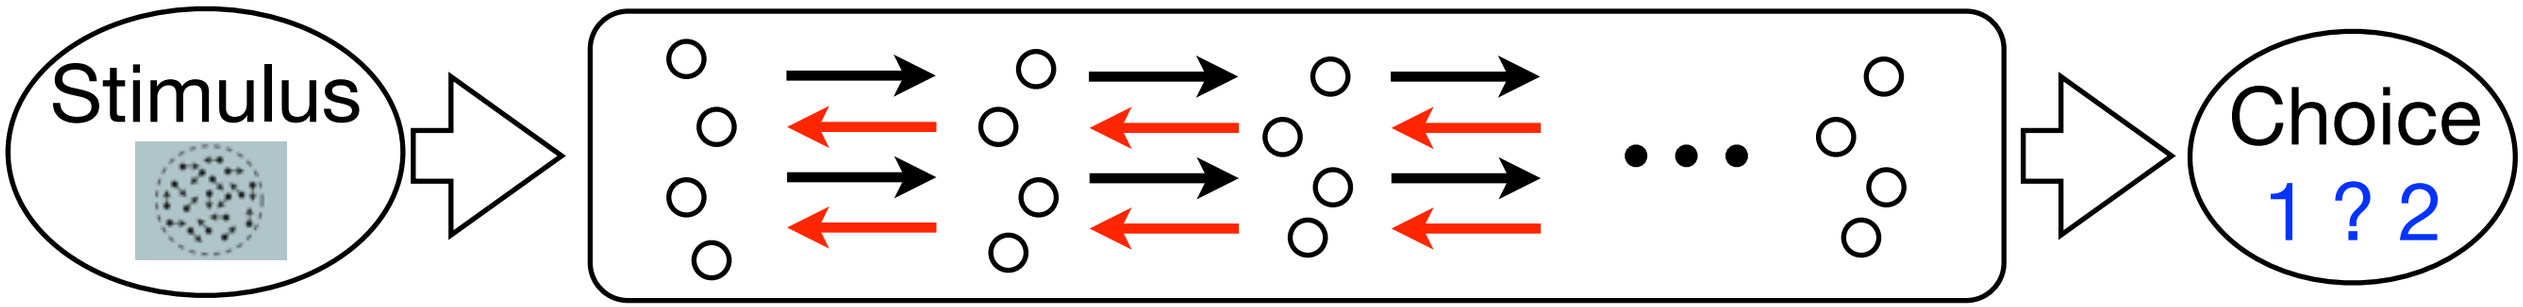

External inputs ⇨ Posterior beliefs ⇨ Prior knowledge

2) Neural data analysis

Computational Neuroscience Lab

PI: Ralf Haefner, Brain & Cognitive Sciences

Richard Lange

Ankani Chattoraj

Sabya Shivkumar

Preliminary results! (1 monkey)

Cardinal task

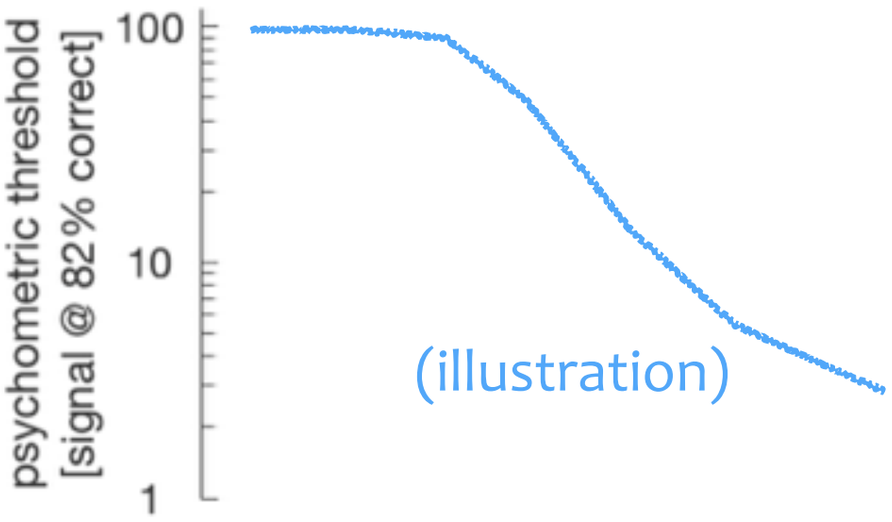

Oblique task

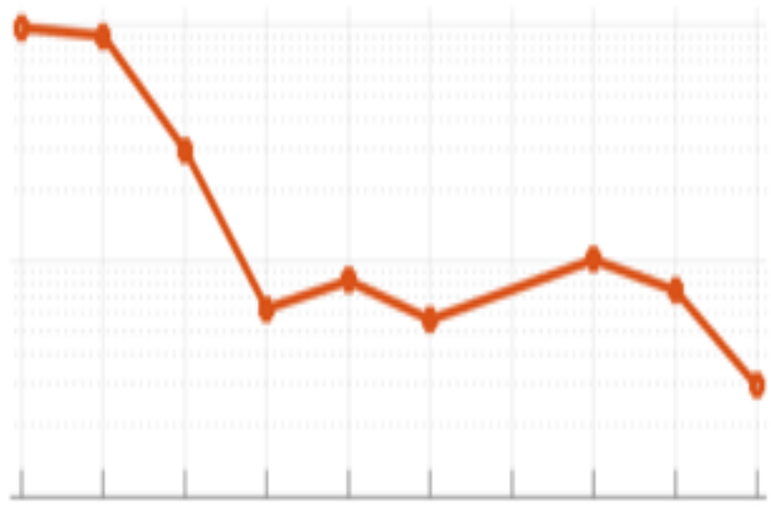

Interleaved

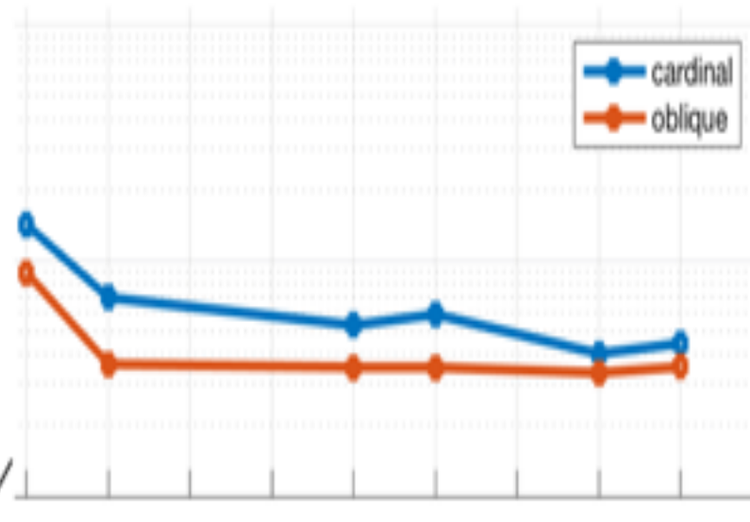

session [day]

Preliminary results! (1 monkey)

Cardinal task

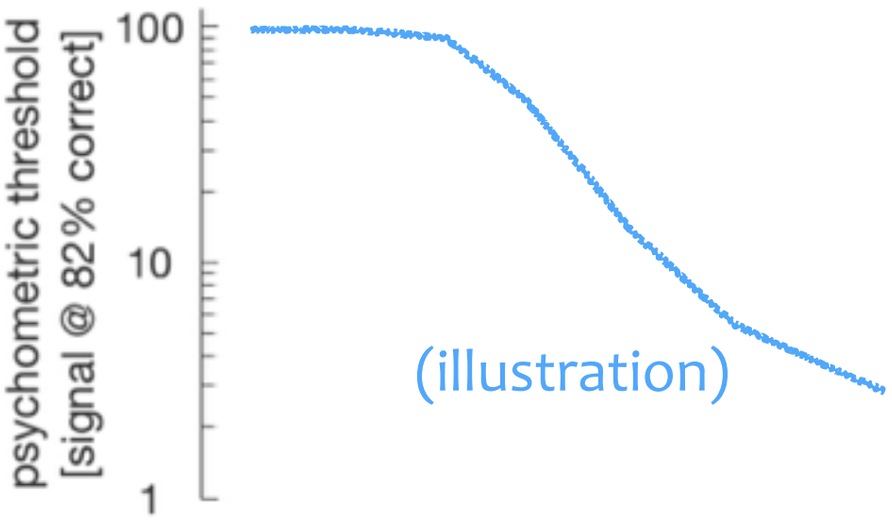

Oblique task

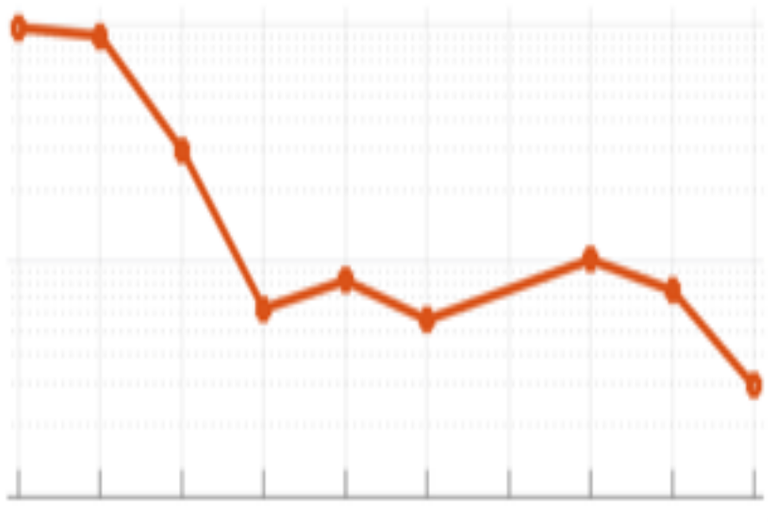

Interleaved

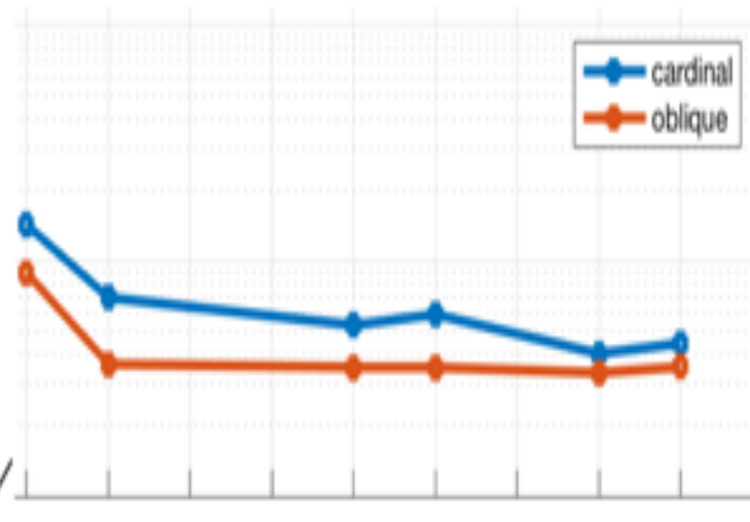

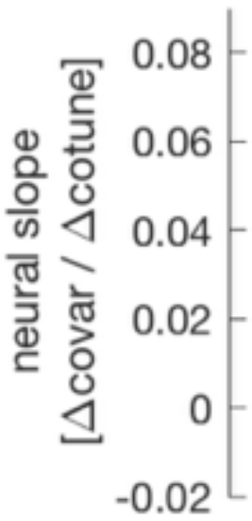

?

(not analyzed yet)

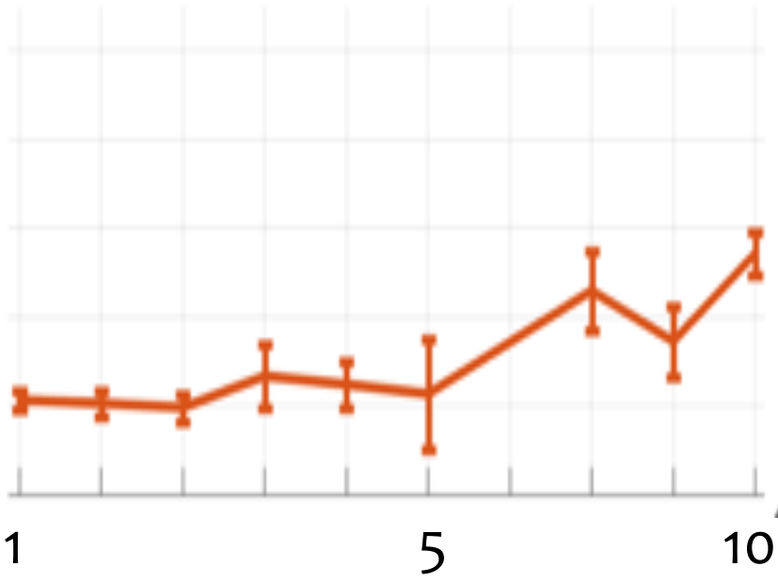

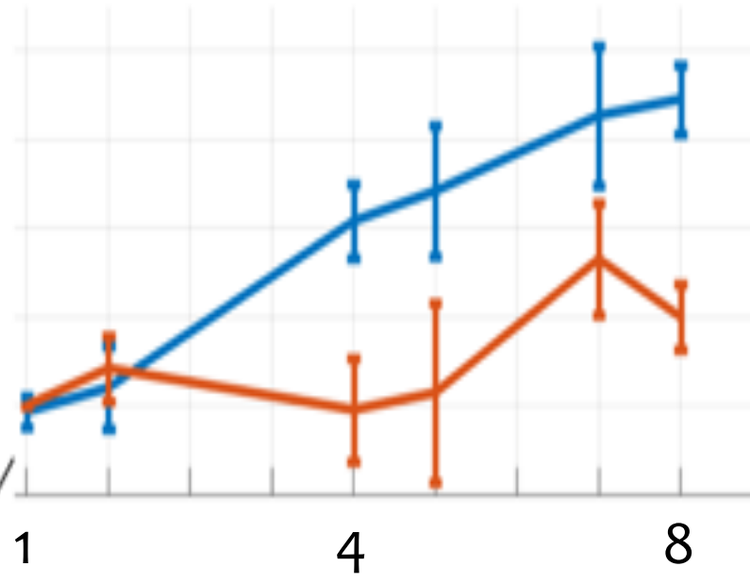

1 5 10 1 4 8

session [day]

Preliminary results! (1 monkey)

Cardinal task

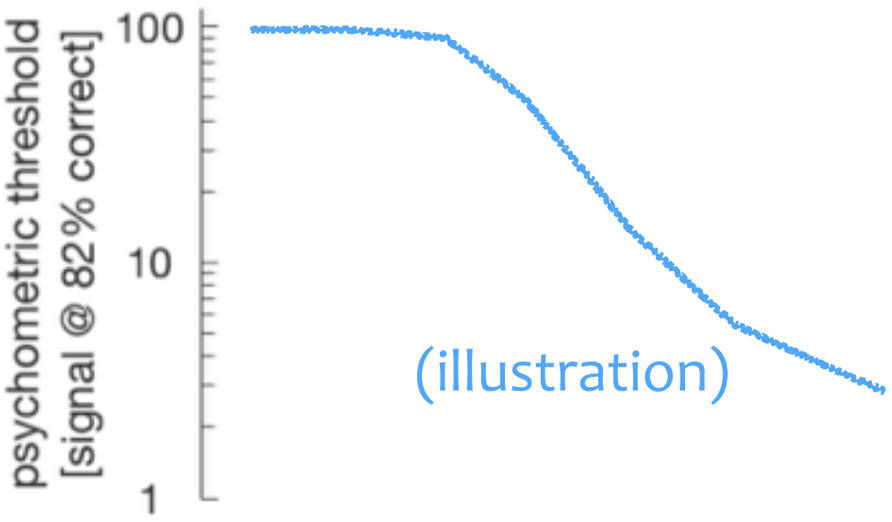

Oblique task

Interleaved

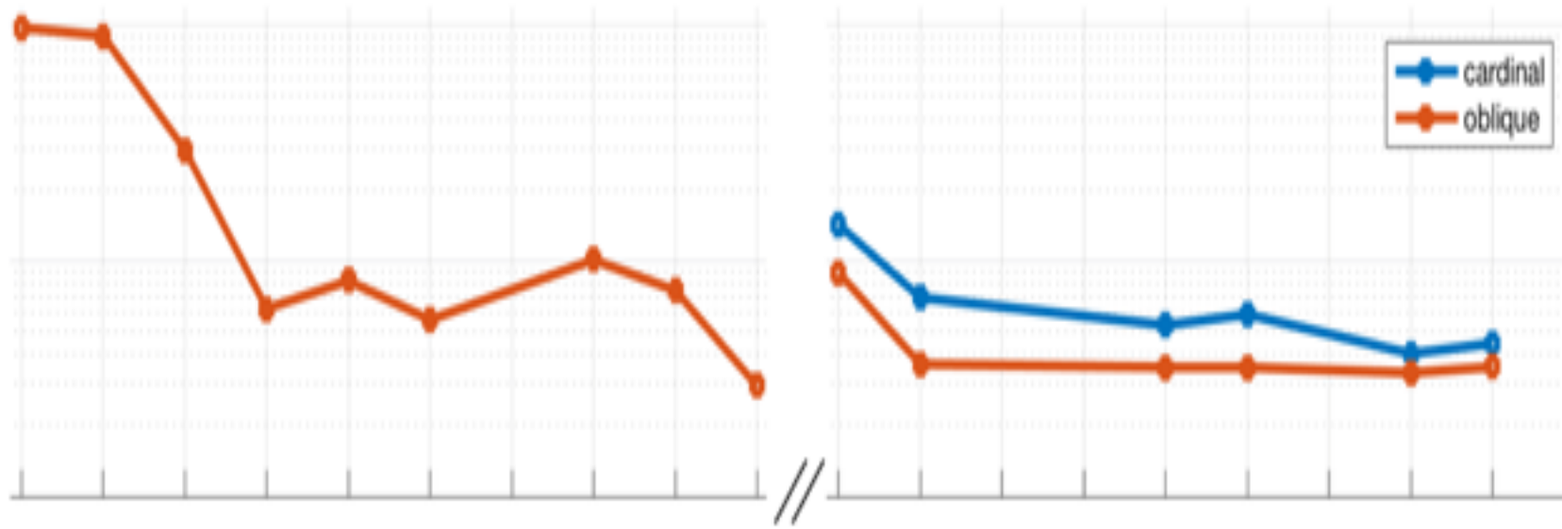

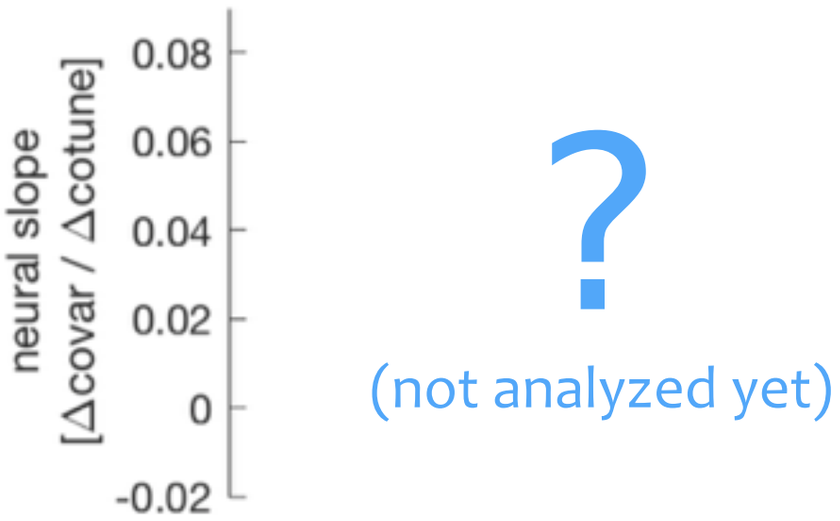

?

(not analyzed yet)

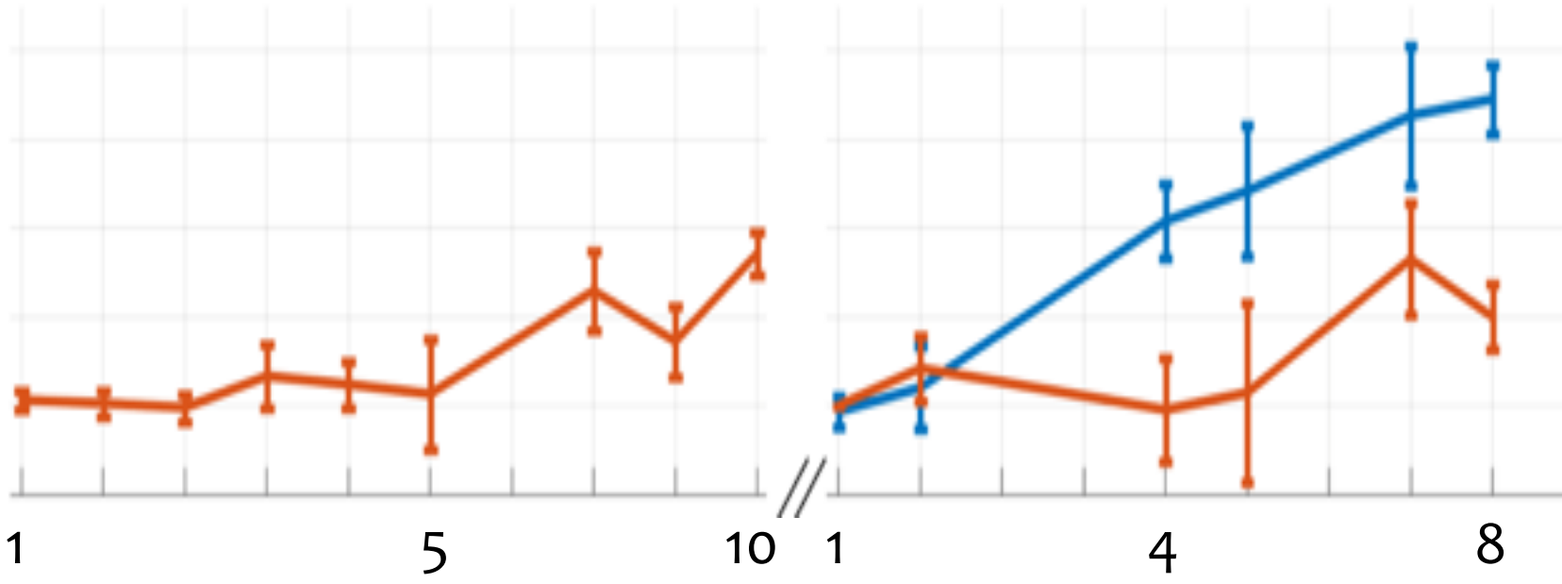

- Next steps: analyze opposite task, compare cooling data

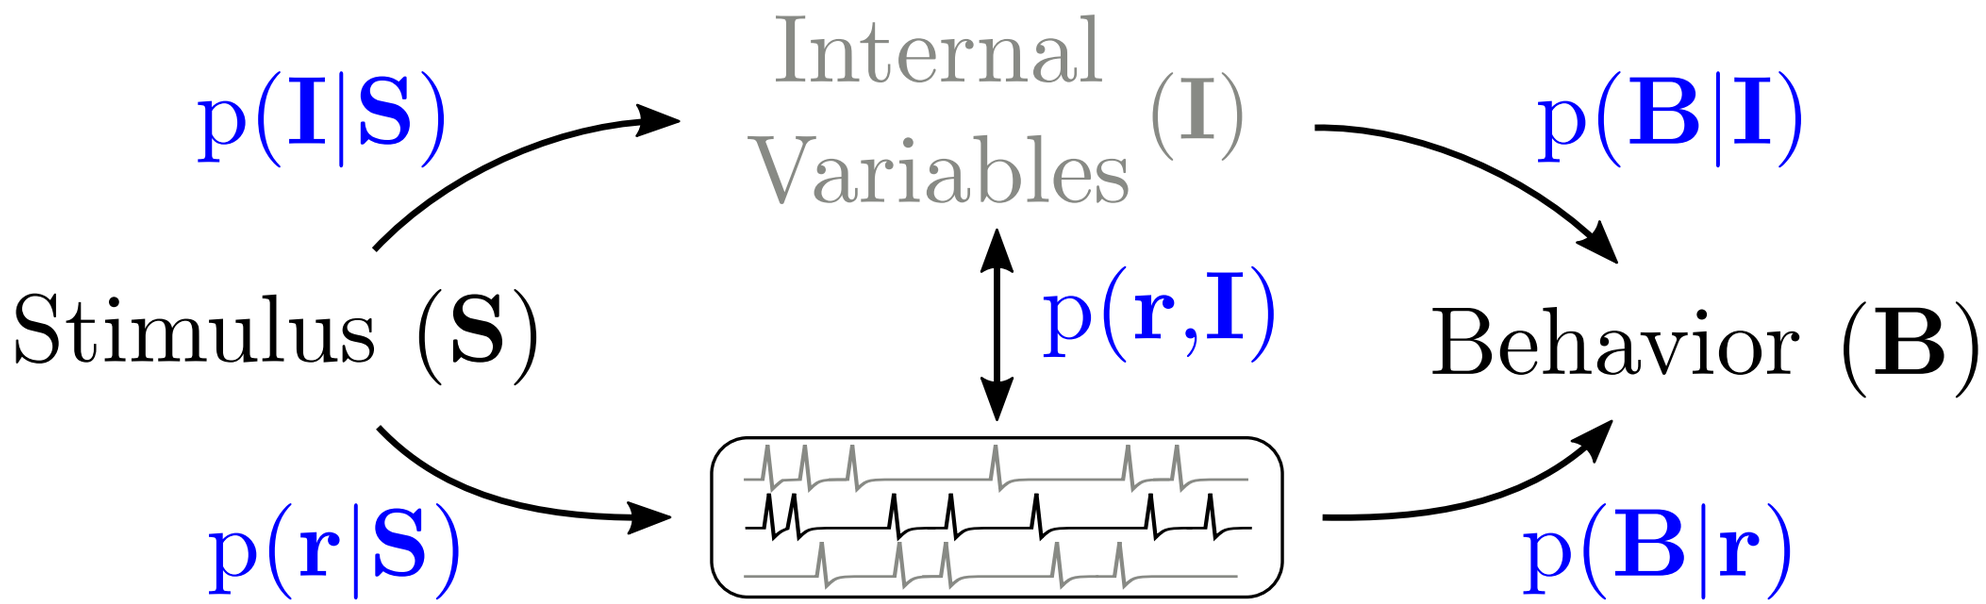
Internal Variables (I)
p(I|S)
p(B|I)
Stimulus (S)
p(r,I)
Behavior (B)
p(r|S)
p(B|r)
Neural Responses (r)

1) Brain's computation

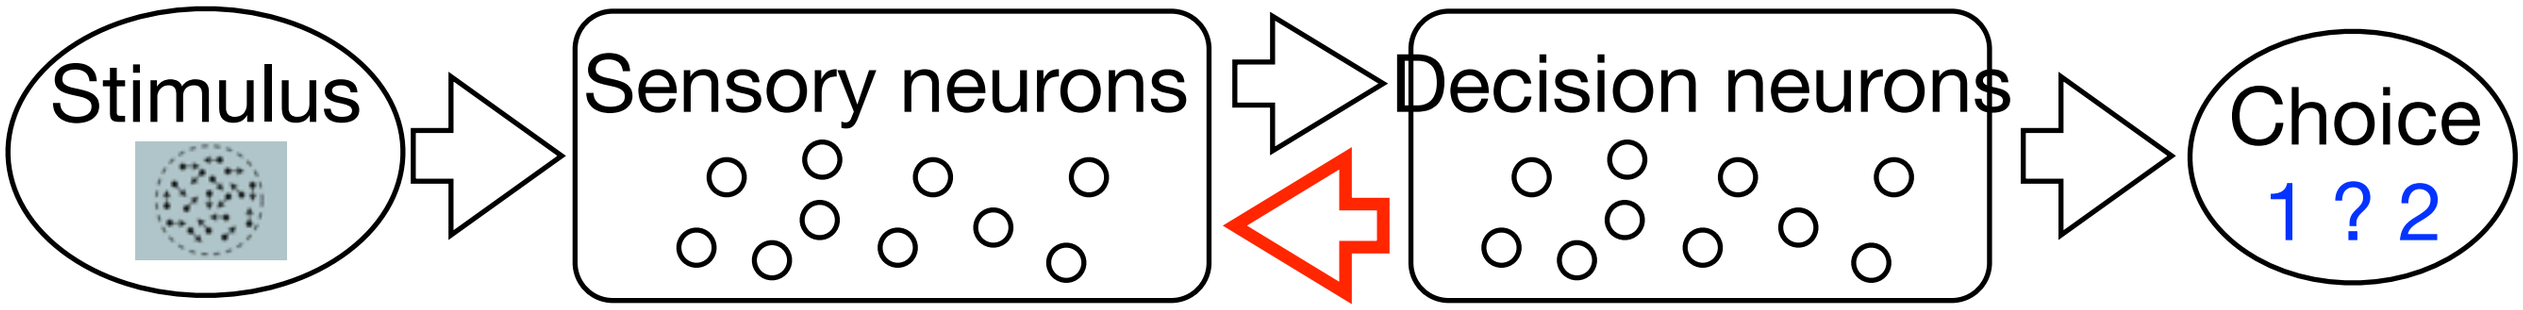

2) Neural data analysis

1) Brain's computation

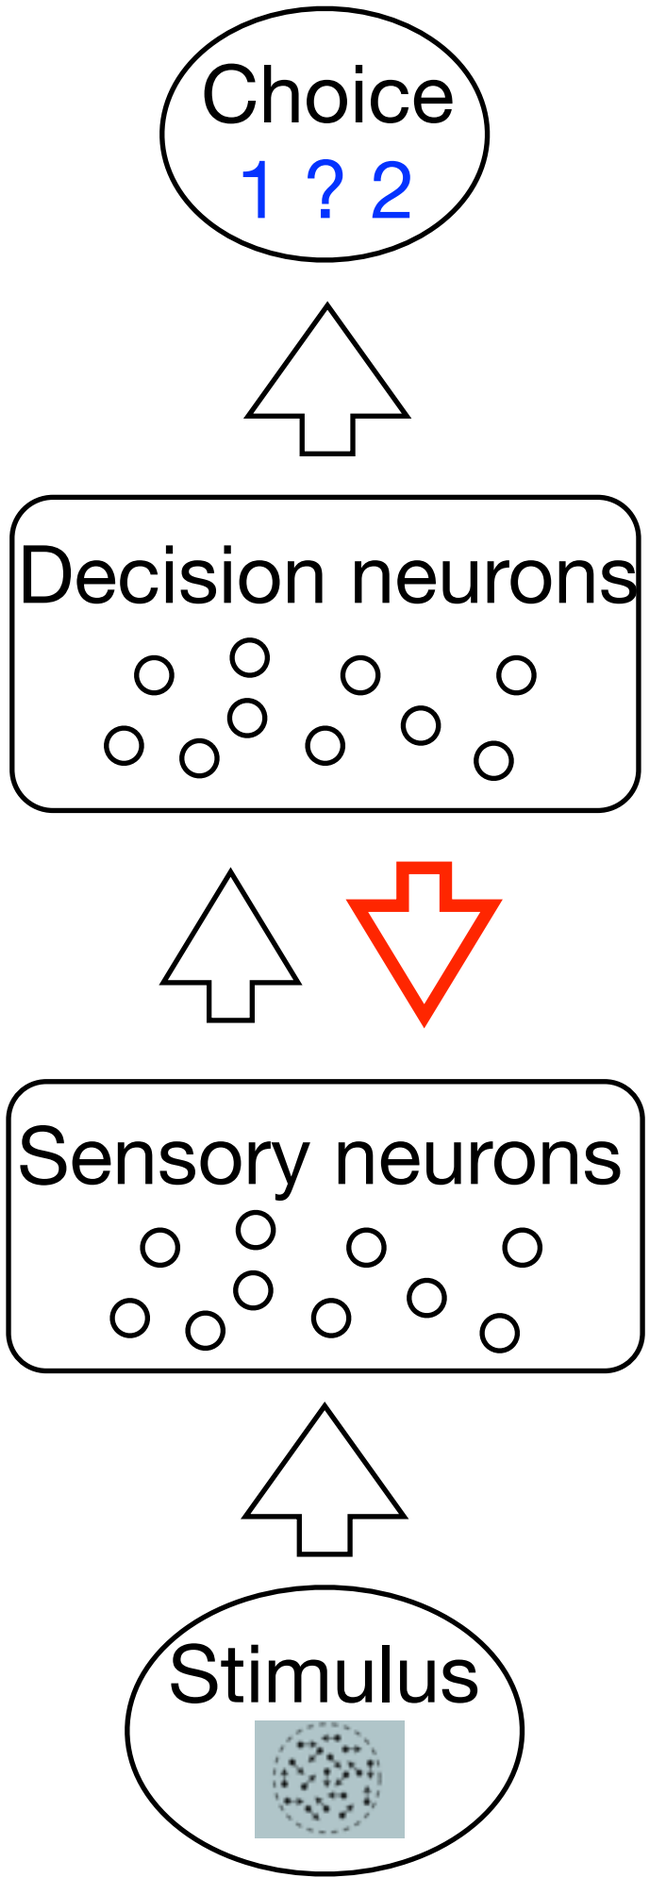

1) Brain's computation

Prior knowledge

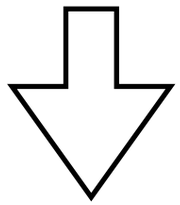

Posterior beliefs

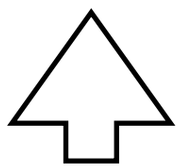

External inputs

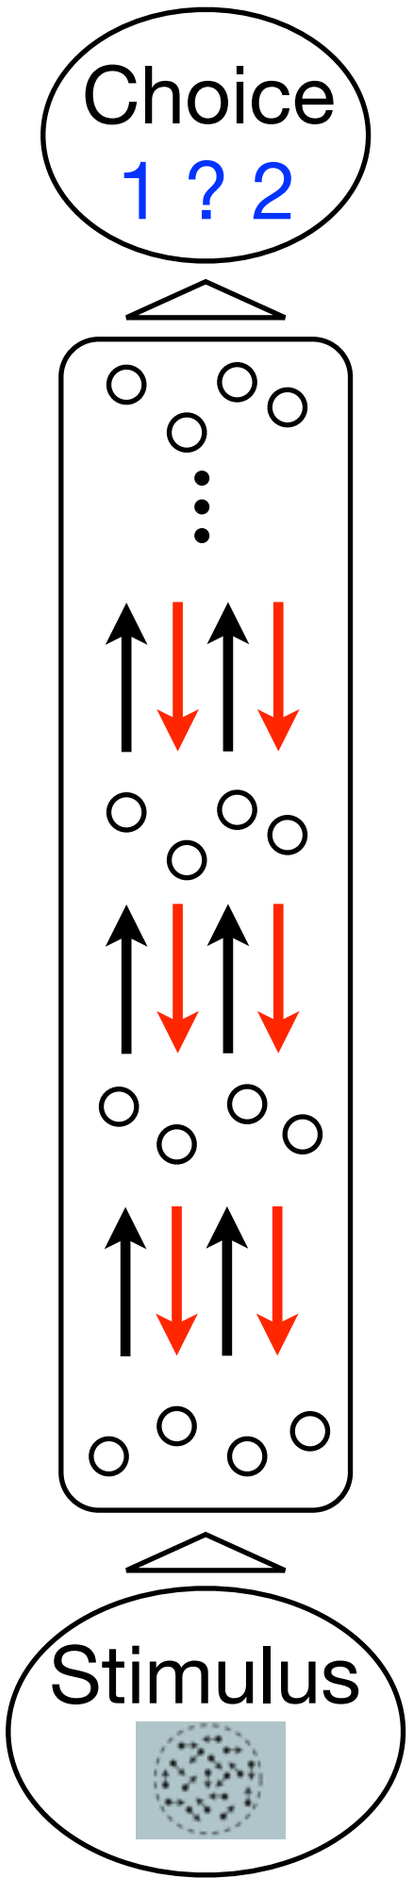

Sampling in a probabilistic model of V1

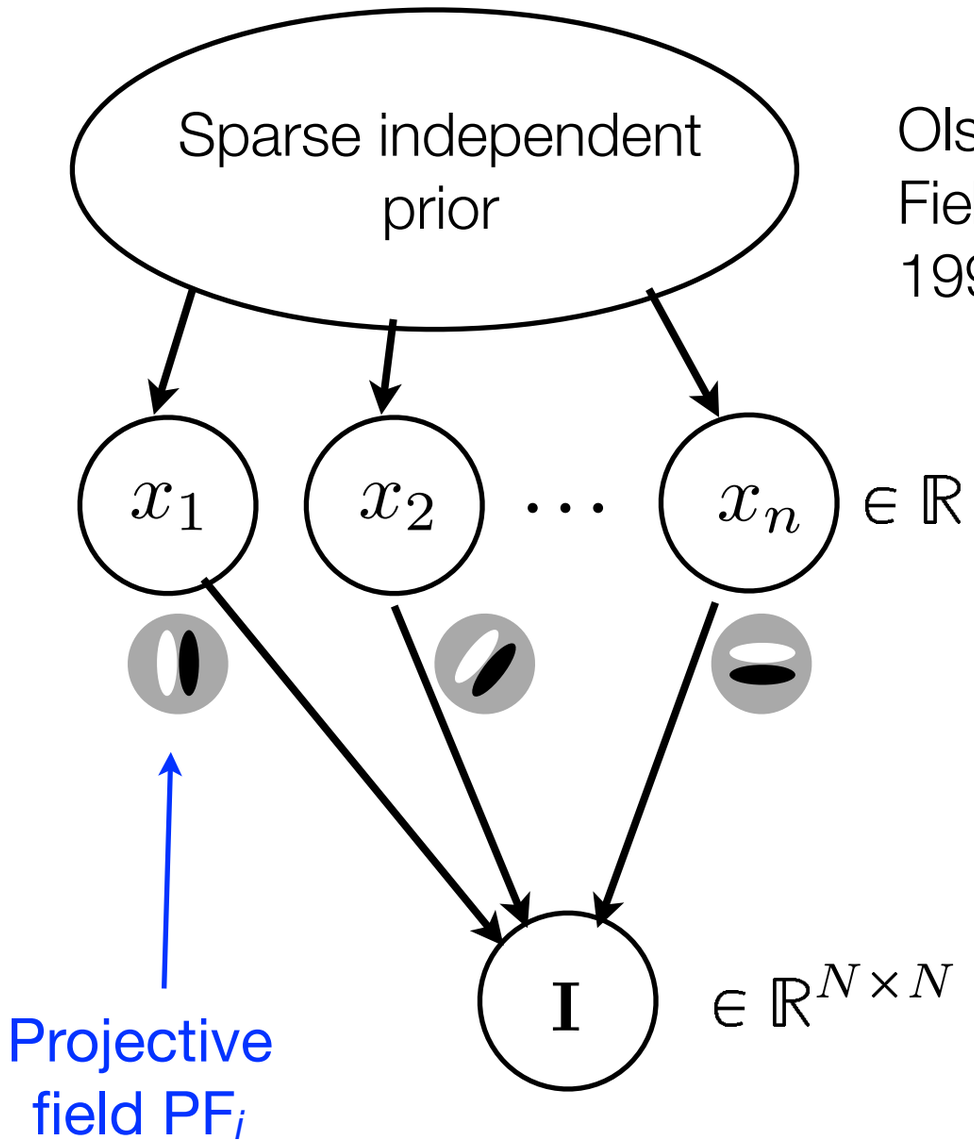

Olshausen & Field, Nature 1996

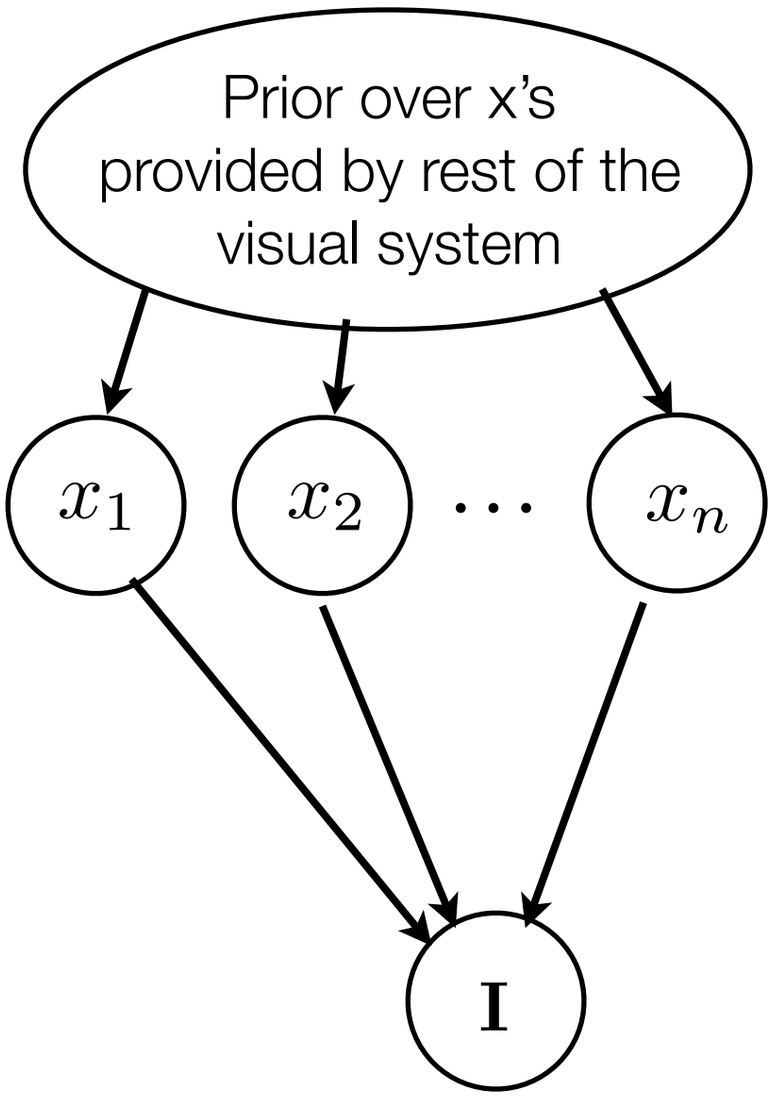

$$p(\mathbf{I}|\mathbf{x}) = \mathcal{N}\left(\mathbf{I} : \sum_i \mathrm{PF}_i x_i, \sigma^2\right)$$

Sampling: Hoyer & Hyvarinnen, NIPS 2003

Neural network implementation: Berkes et al, NIPS 2009

Perturbation approach:

Training on a psychophysical task changes internal model in predictable way

Allows predictions for neuronal behavior (e.g. correlations)

Prediction for ***changes*** in neural activity

Predictions for:

- Difference between before and after learning

- Difference between two different tasks

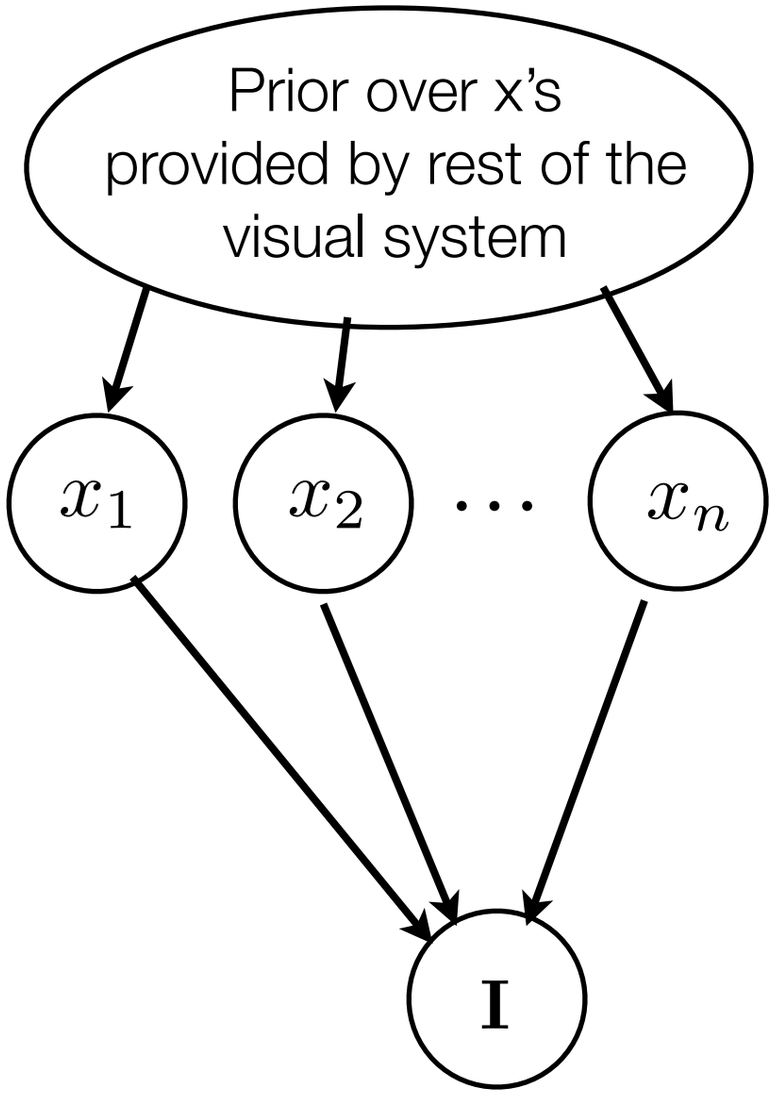

Perturbation approach:
Training on a psychophysical task changes internal model in predictable way
Allows predictions for neuronal behavior (e.g. correlations)

Orientation discrimination task

Which of two perpendicular gratings
caused the noisy image on the screen?

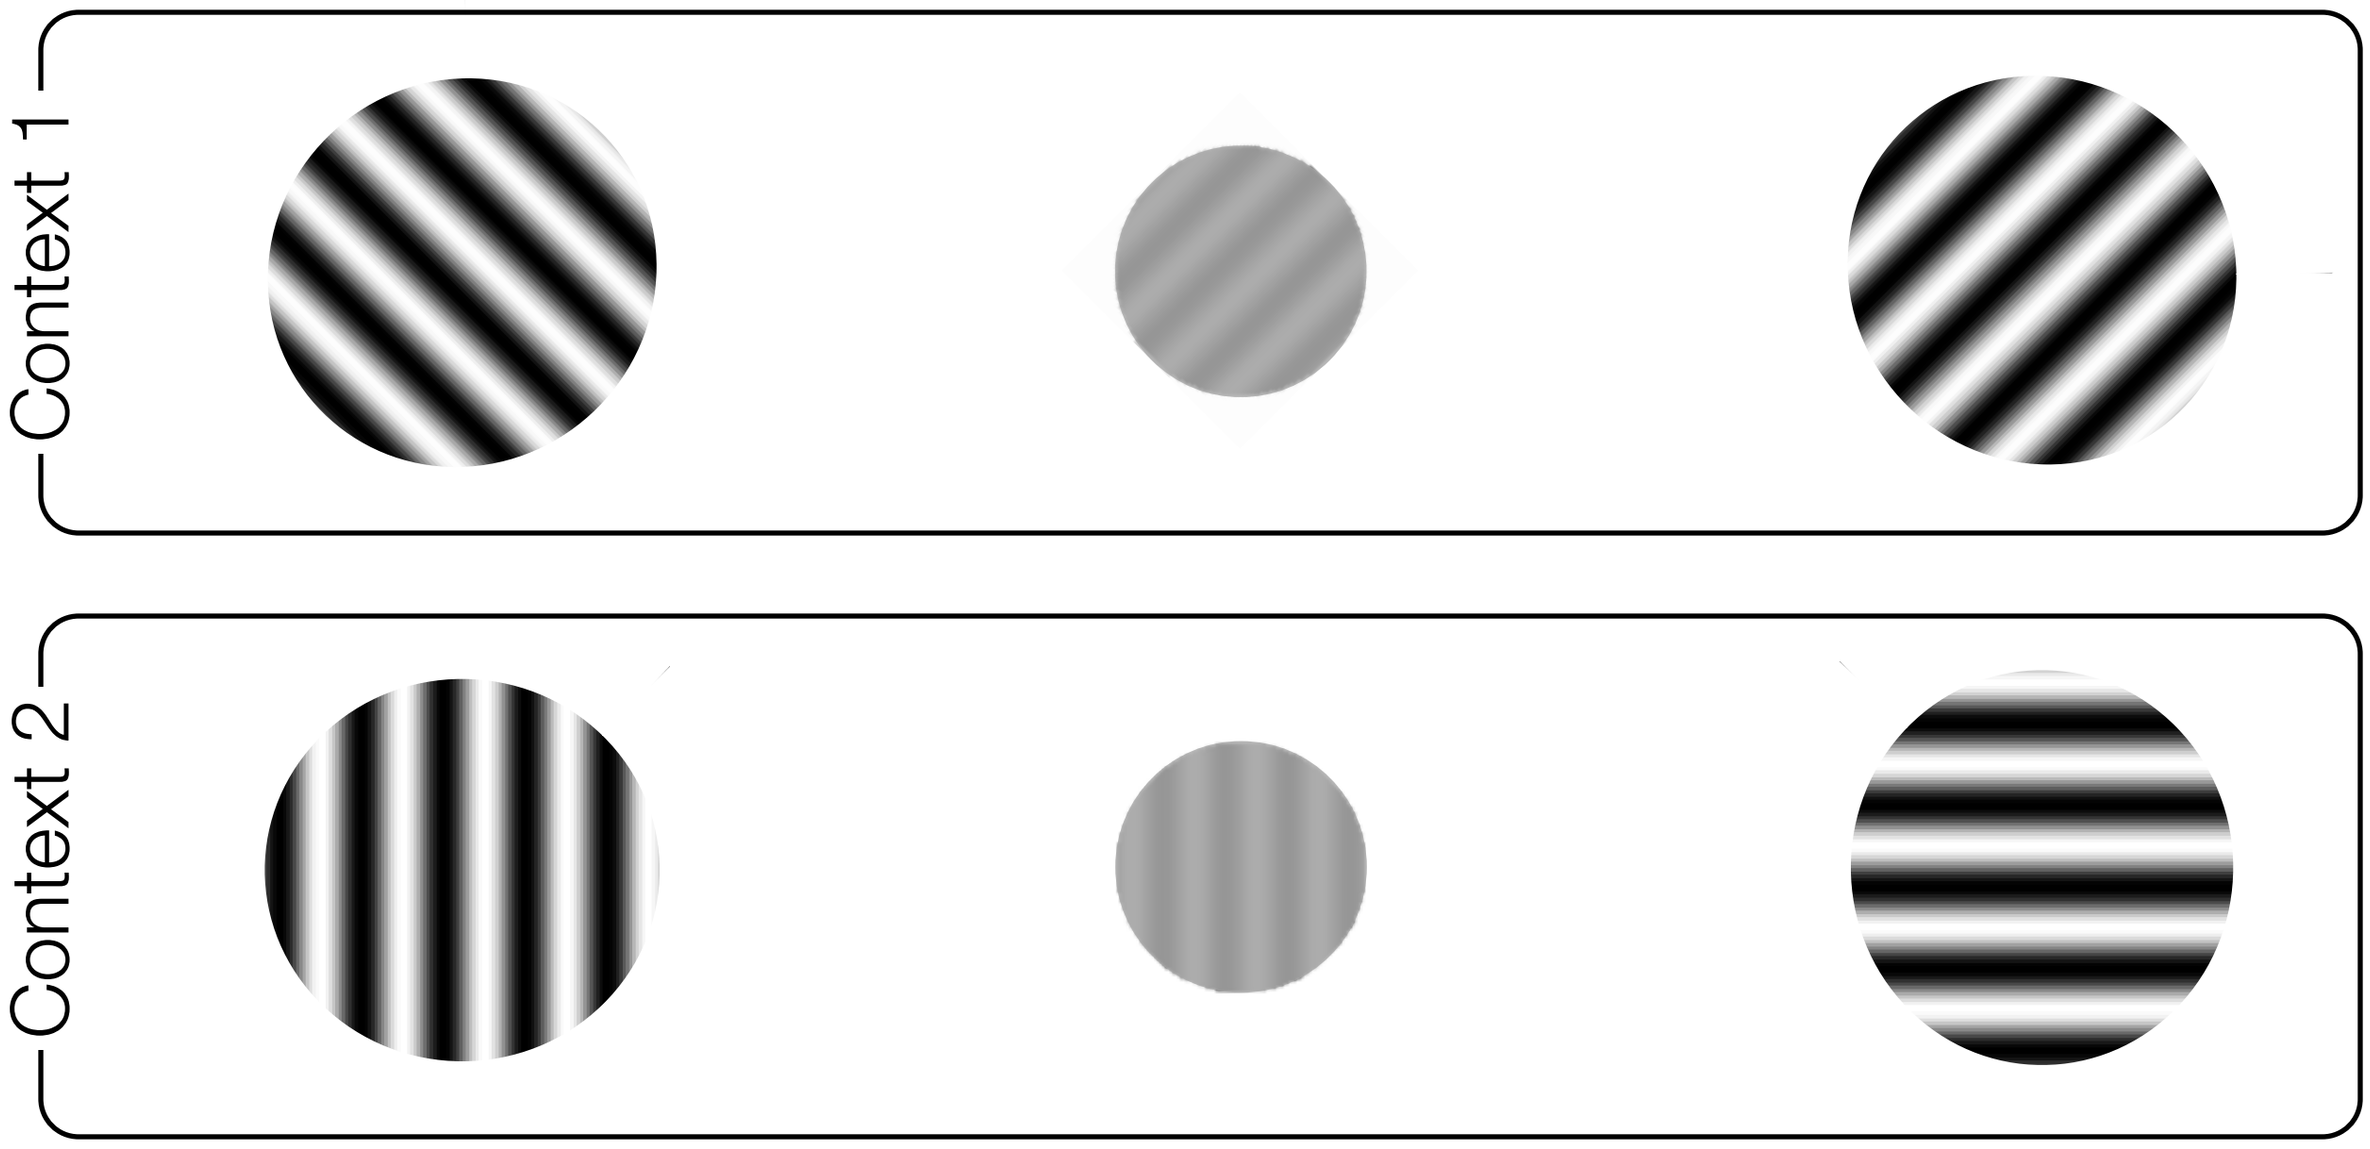

Inspired by Cohen & Newsome, Neuron 2008: Motion direction discrimination task/MT

Choice probabilities in **x**

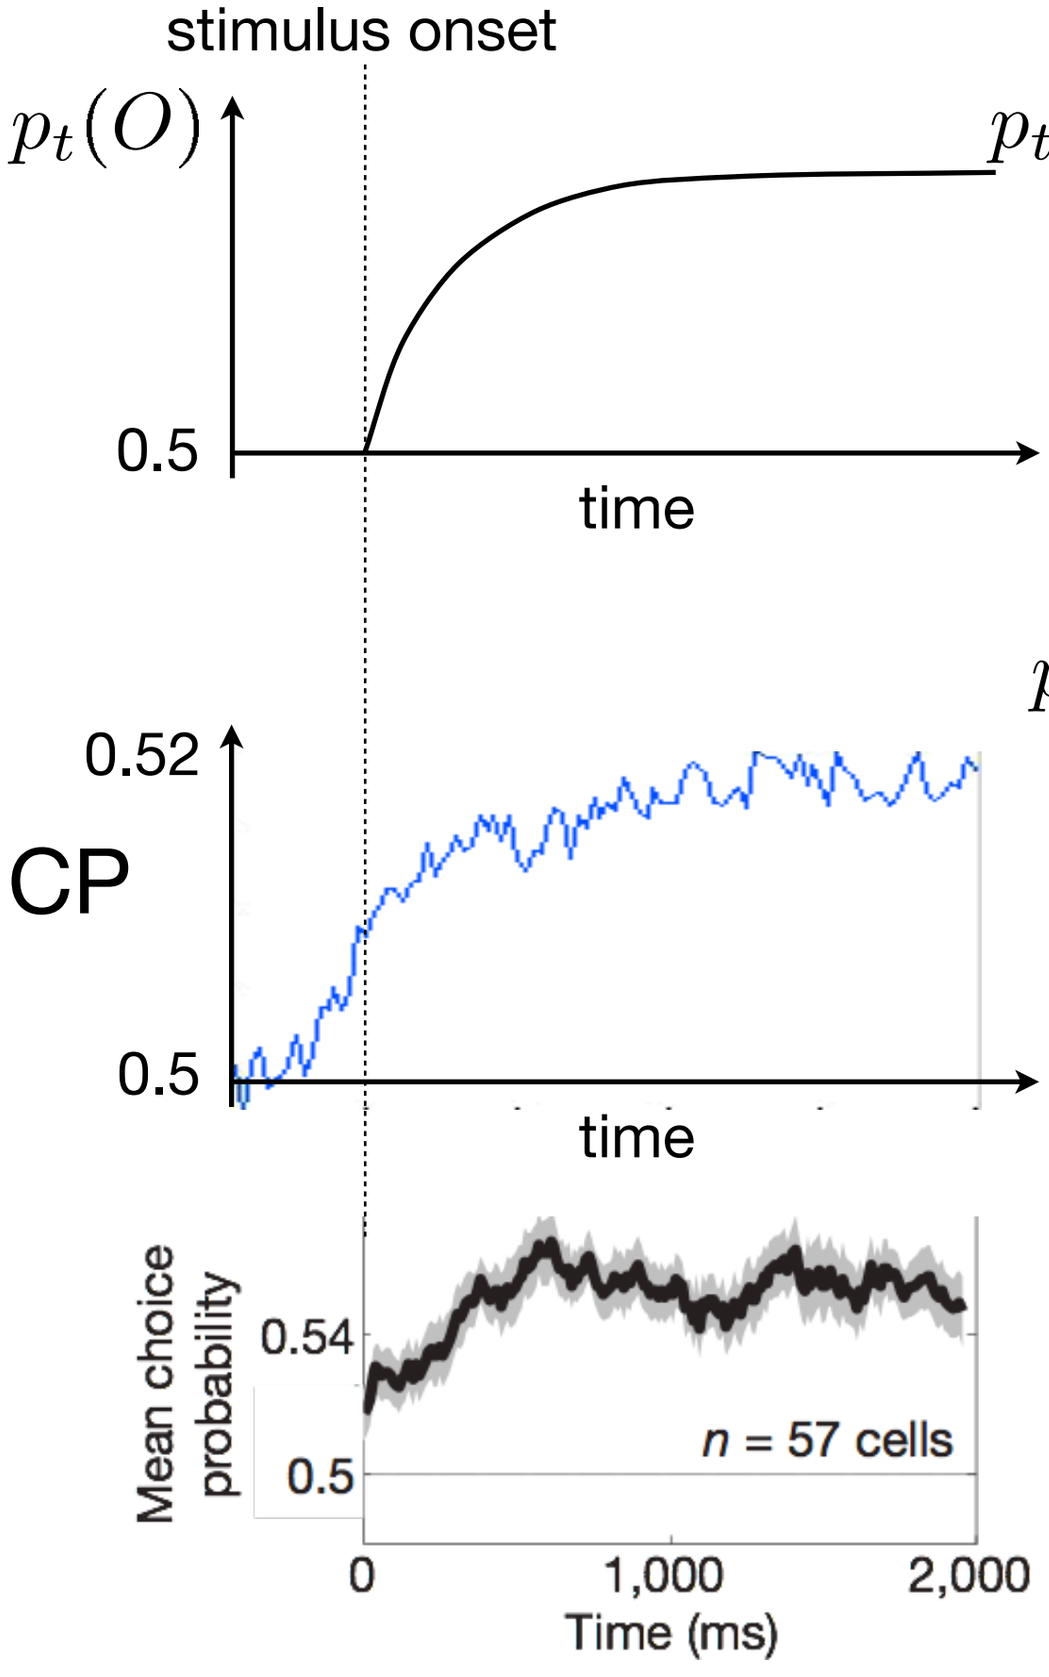

Nienborg & Cumming, Nature 2009

$p_t(O) \propto p_{t-1}(O)p(O|\mathbf{x}_t)$

top-down

$p_t(\mathbf{x}) \propto p(\mathbf{x}|\mathbf{y}_t)p(\mathbf{x}|O_t)$

bottom-up

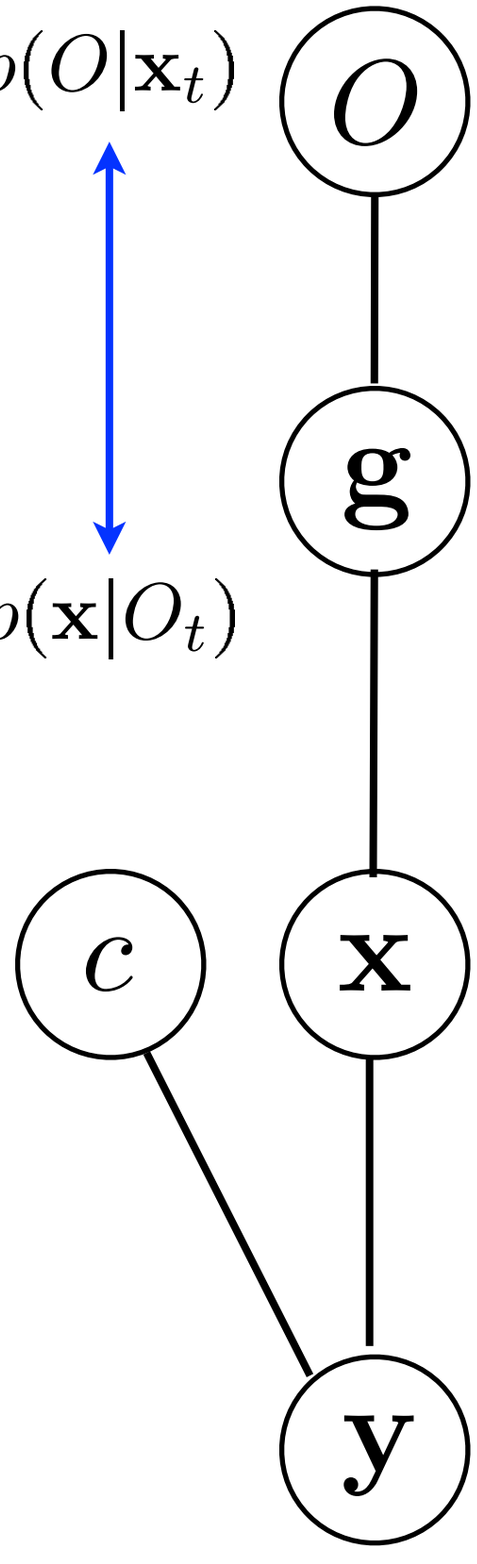

Decision variable

Grating variables

Gabor variables

Image variable

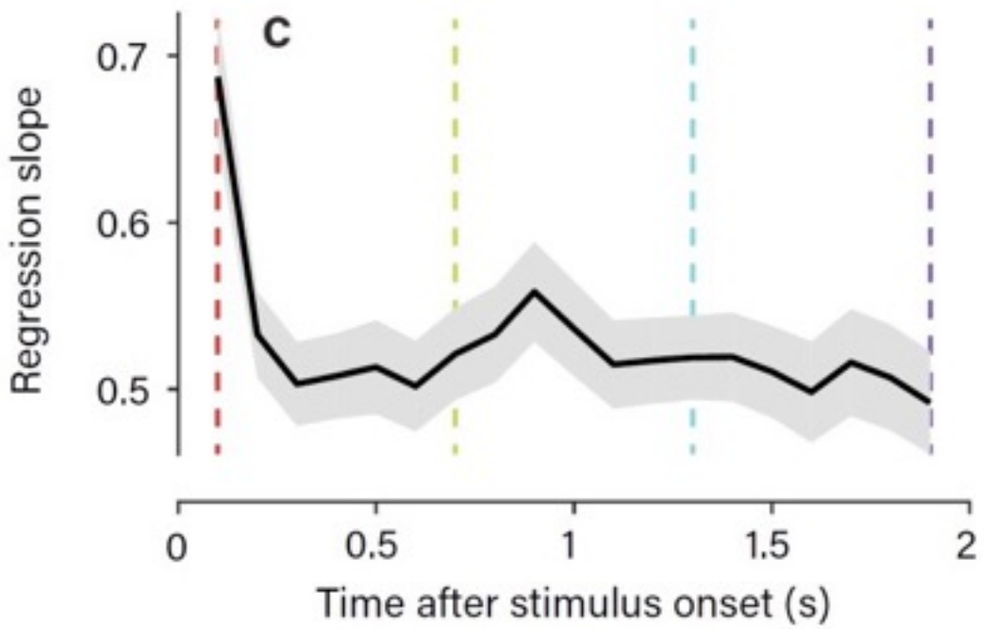
C
0.7
0.6
0.5
Regression slope
0
0.5
1
1.5
2
Time after stimulus onset (s)

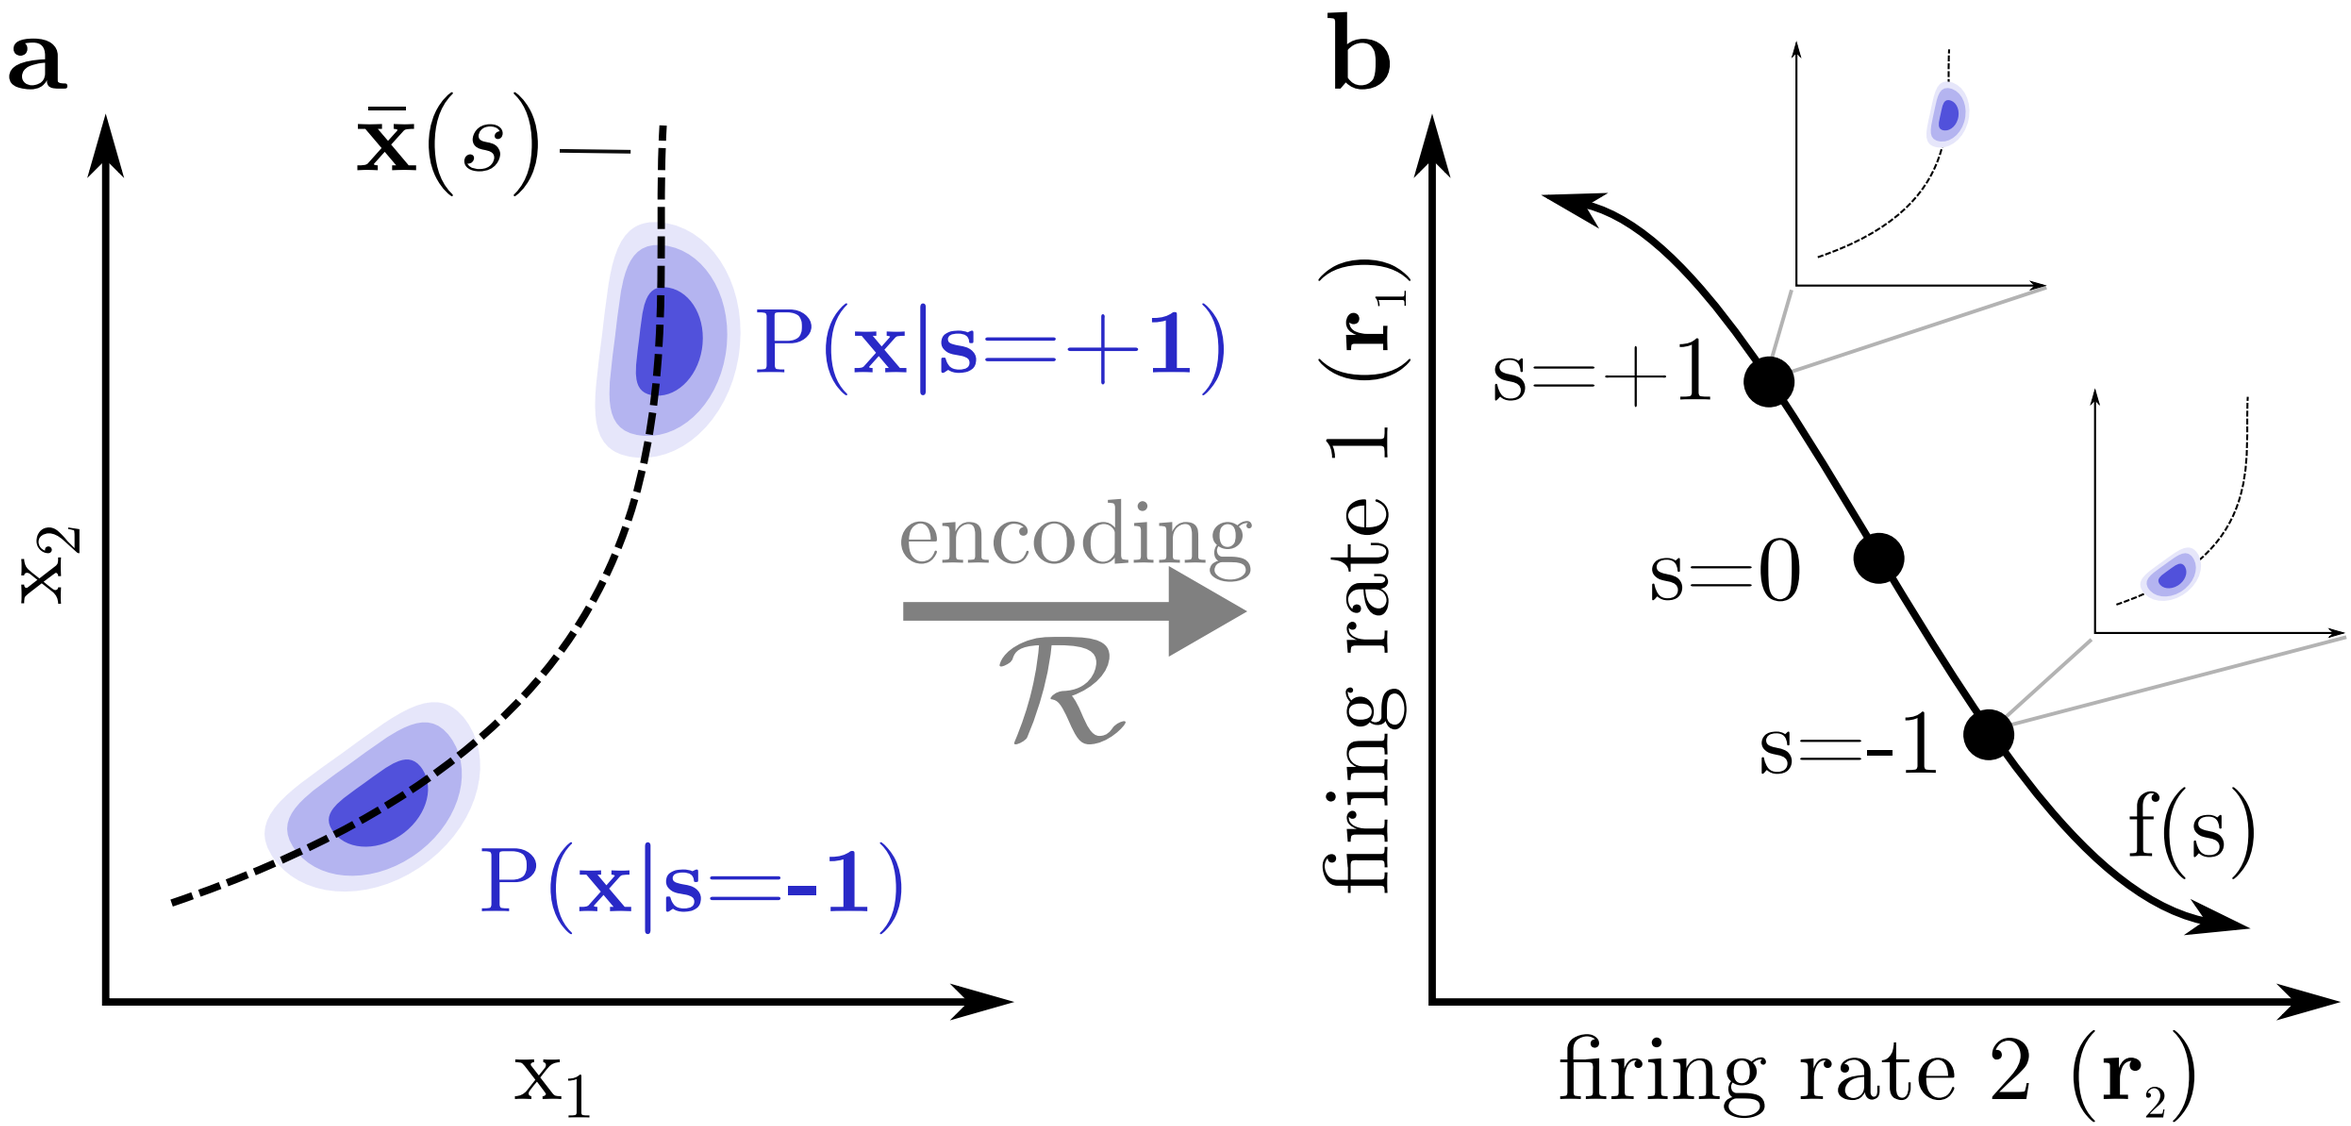
a
$\bar{\mathbf{x}}(s)$
$\mathbf{P(x|s=+1)}$
encoding
$\mathcal{R}$
$\mathbf{P(x|s=-1)}$
x_2
x_1
b
s=+1
s=0
s=-1
f(s)
firing rate 1 ($\mathbf{r}_1$)
firing rate 2 ($\mathbf{r}_2$)

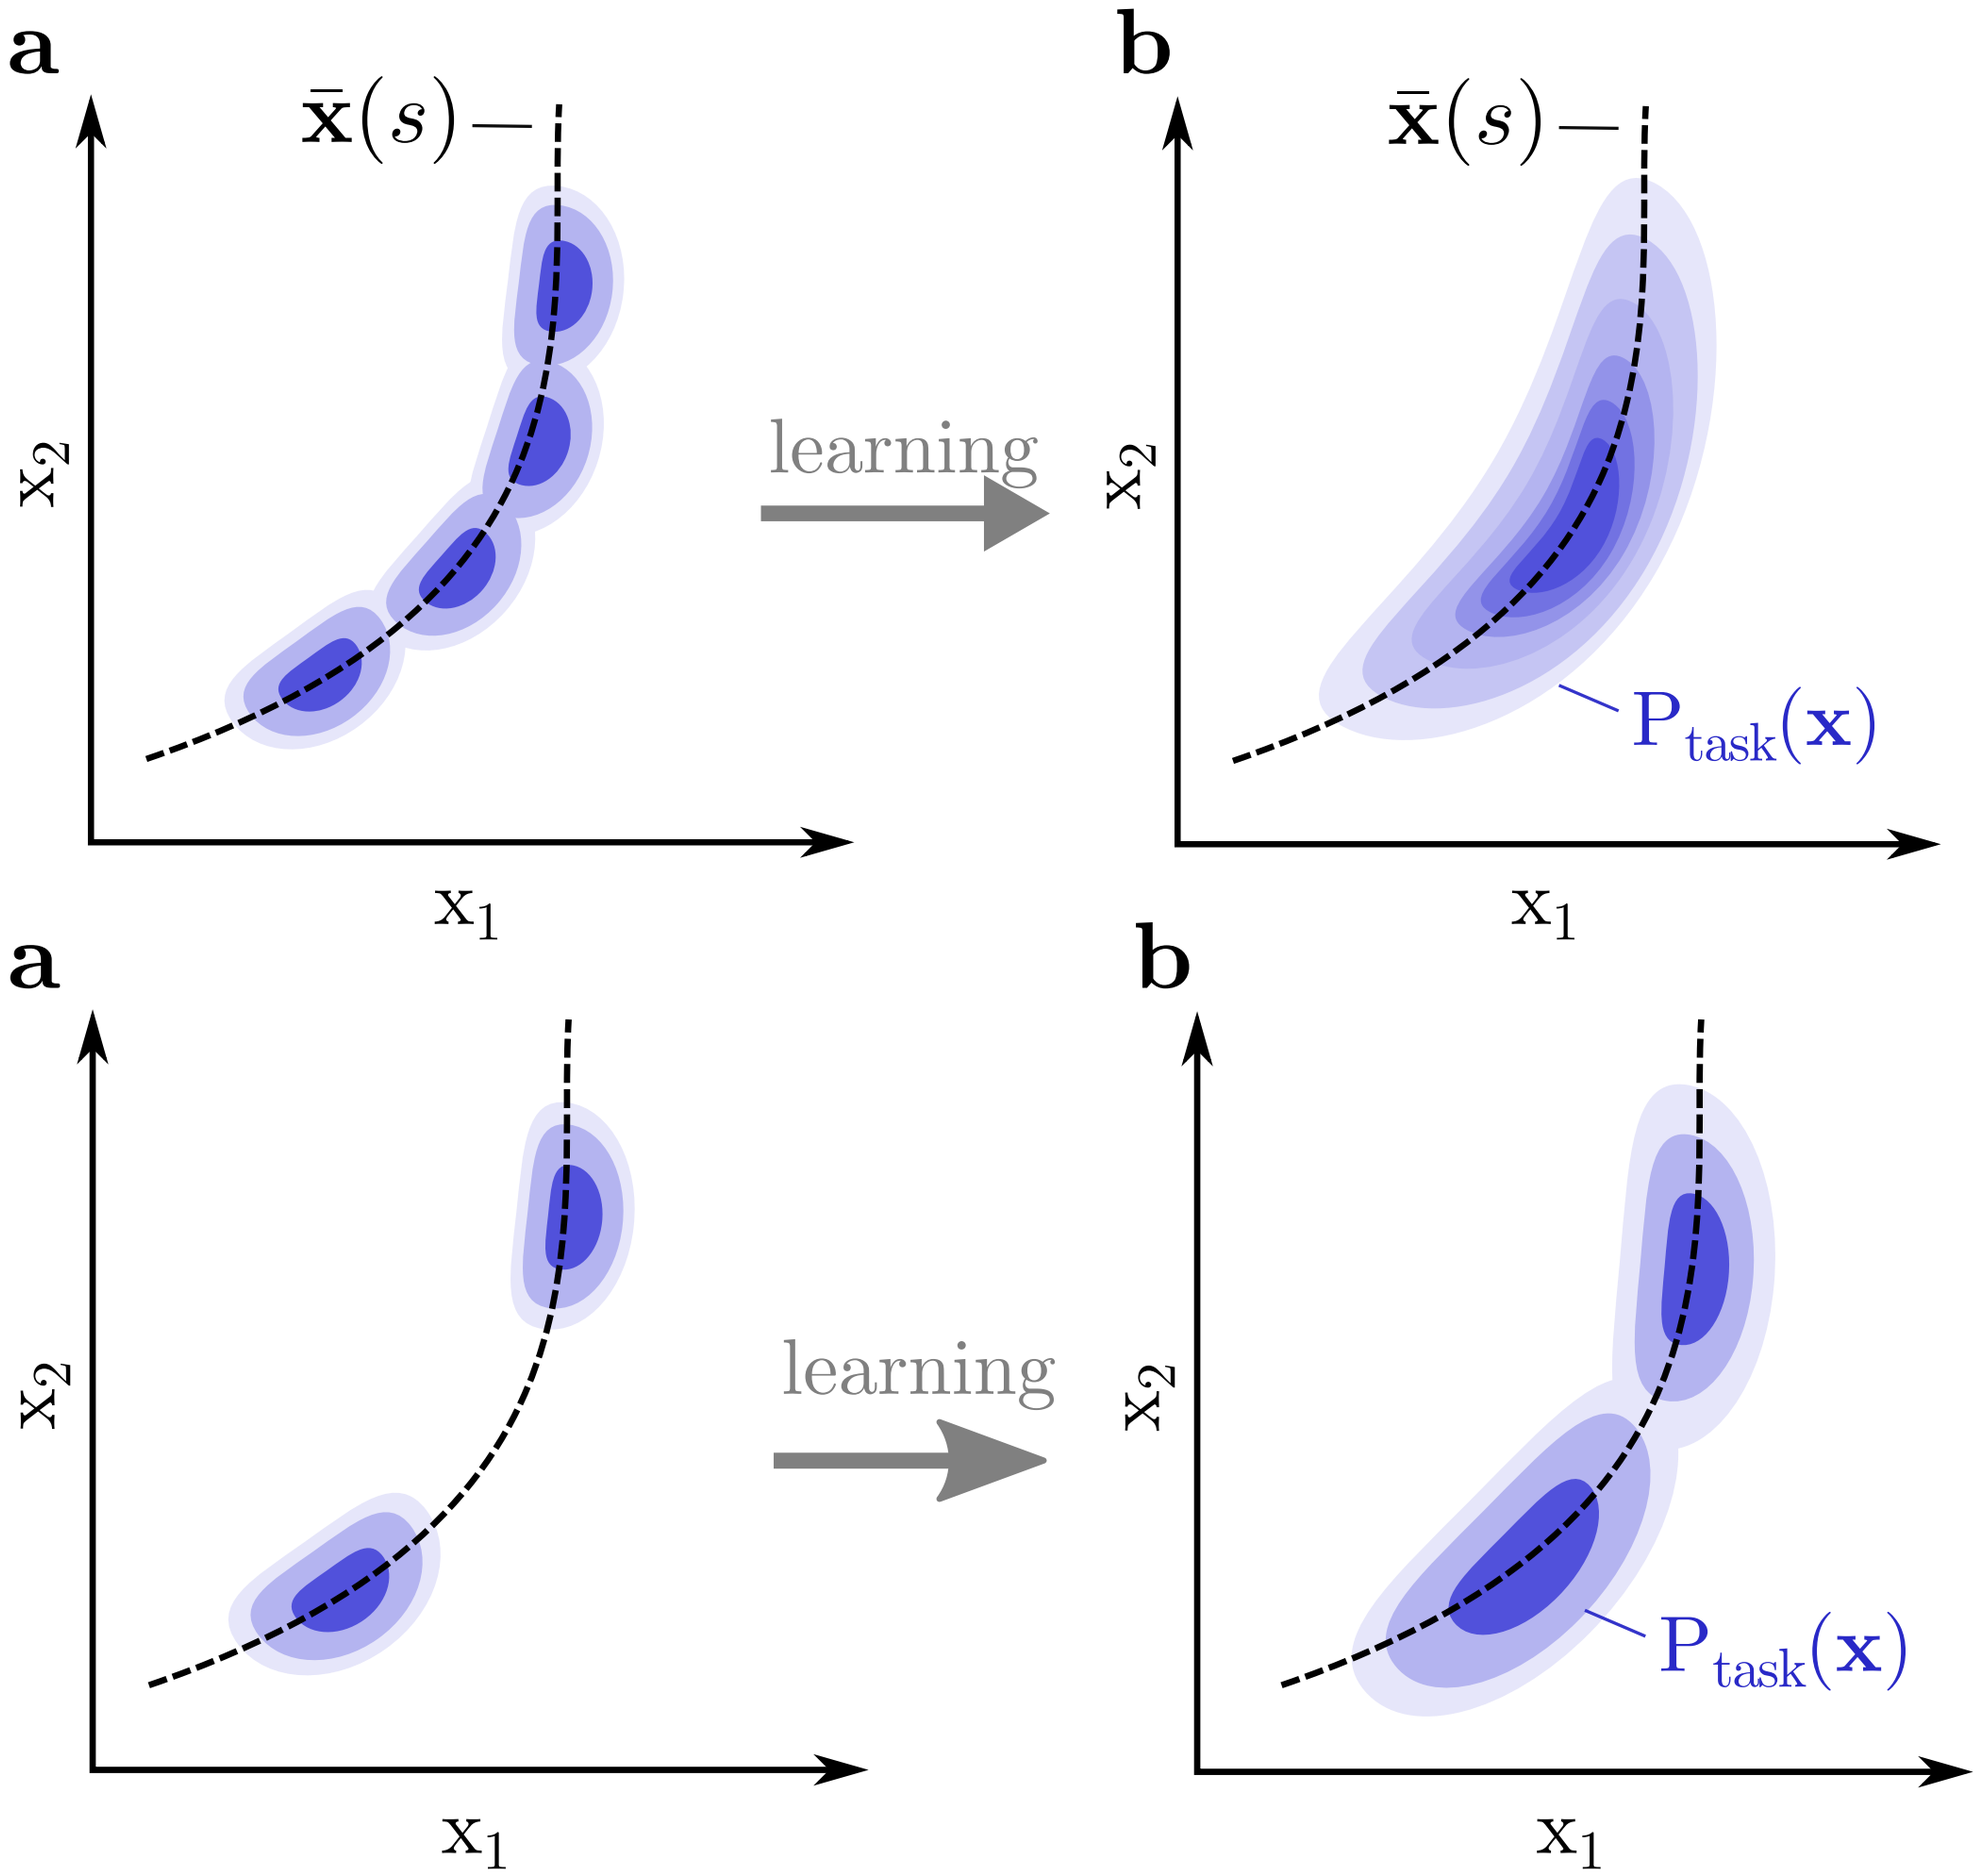
a
$\bar{\mathbf{x}}(s)$
x_2
x_1
learning
b
$\bar{\mathbf{x}}(s)$
x_2
x_1
$P_{task}(\mathbf{x})$
a
x_2
x_1
learning
b
x_2
x_1
$P_{task}(\mathbf{x})$

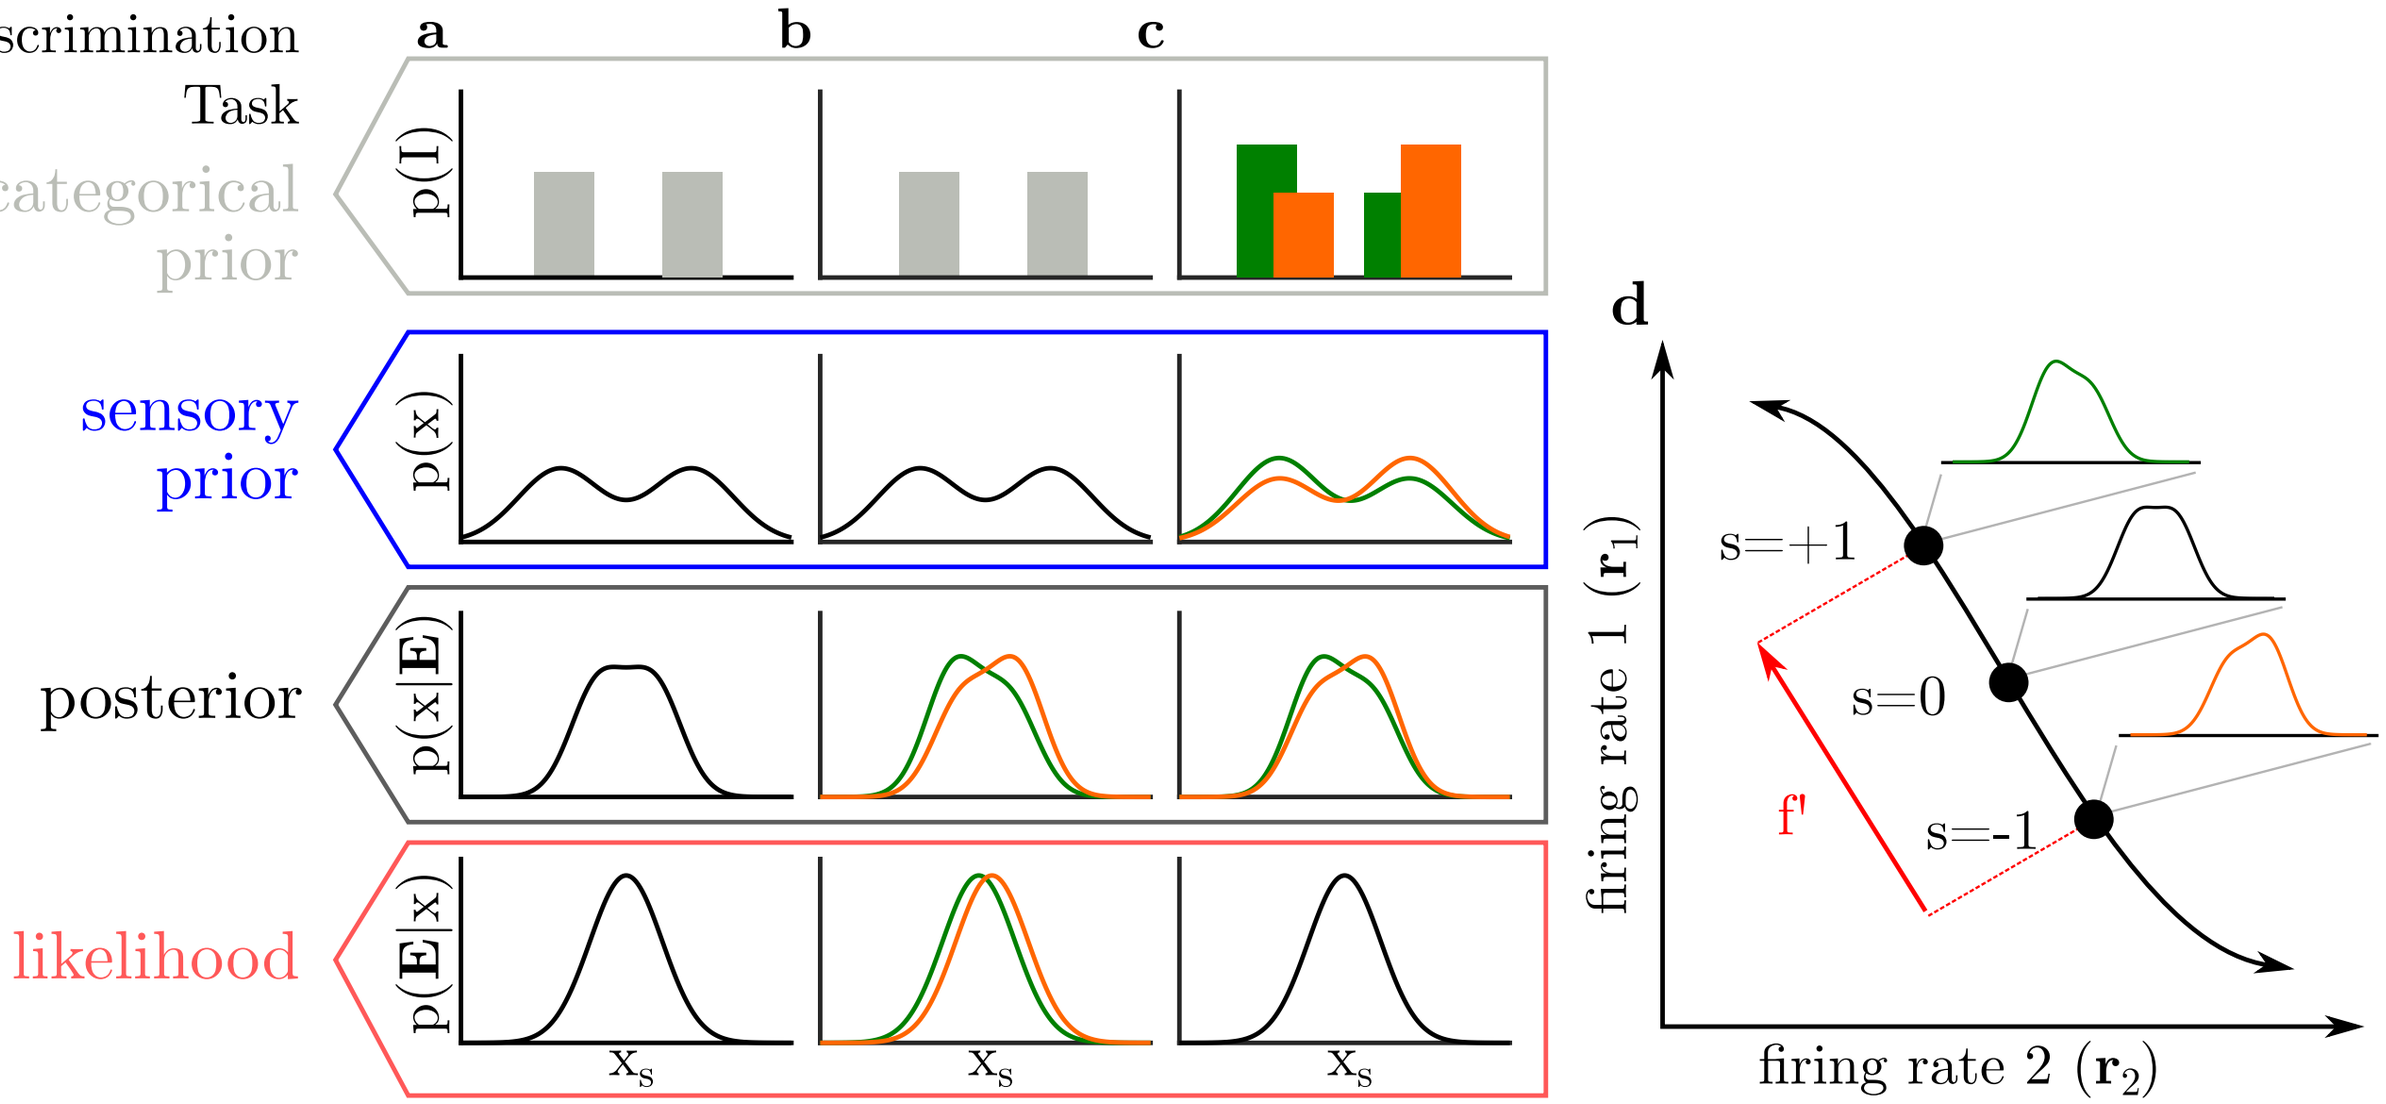

- Differential correlations (proportional to f'f')

Task

c_{ij}

r_{prior}

$d'_i d'_j$

e

r_{prior}

task mismatch

Empirical test of 2nd order prediction

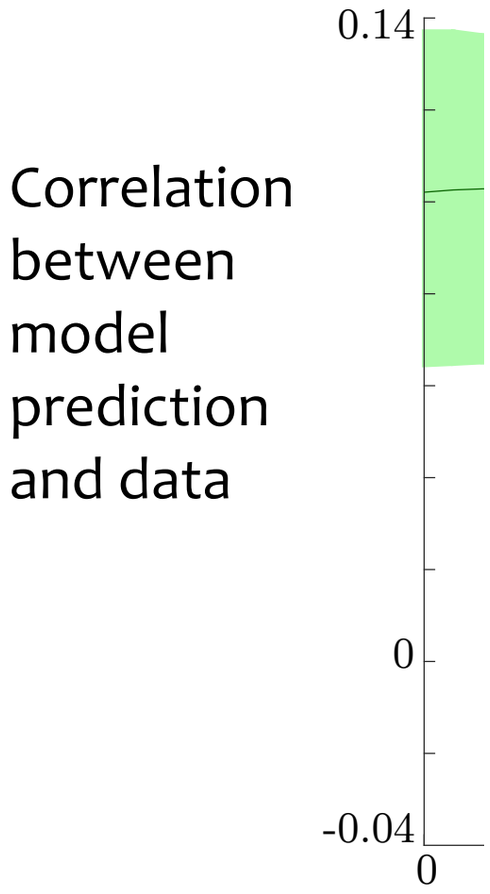

Difference between actual task performed by monkey and task for which prediction was made

a

$\bar{\mathbf{x}}(s)$

$\mathbf{P(x|s=+1)}$

$\mathbf{P(x|s=-1)}$

encoding

$\mathcal{R}$

x_2

x_1

b

firing rate 1 ($\mathbf{r}_1$)

firing rate 2 ($\mathbf{r}_2$)

s=+1

s=0

s=-1

f(s)

a

vs

vs

r_2

r_1

a

$\bar{\mathbf{x}}(s)$

$\mathbf{P(x|s=+1)}$

encoding

$\mathcal{R}$

$\mathbf{P(x|s=-1)}$

x_2

x_1

b

s=+1

s=0

s=-1

f(s)

firing rate 1 ($\mathbf{r}_1$)

firing rate 2 ($\mathbf{r}_2$)