

MENTÁLIS MODELLEK BAYES-I MEGKÖZELÍTÉSSEN

A BAYESIAN APPROACH TO INTERNAL MODELS

LENGYEL MÁTÉ

Computational and Biological Learning Lab
Department of Engineering
University of Cambridge



Department of Cognitive Science
Central European University



BAYESIAN INFERENCE FOR DATA ANALYSIS

BAYESIAN INFERENCE FOR DATA ANALYSIS

data, x

BAYESIAN INFERENCE FOR DATA ANALYSIS

parameters, y



data, x

BAYESIAN INFERENCE FOR DATA ANALYSIS

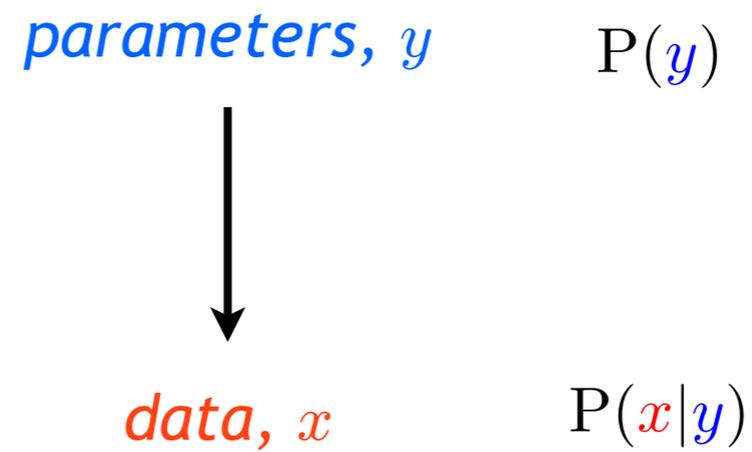
parameters, y



data, x

$$P(x|y)$$

BAYESIAN INFERENCE FOR DATA ANALYSIS



BAYESIAN INFERENCE FOR DATA ANALYSIS

parameters, y



data, x

$P(y)$

$$P(y|x) = \frac{P(x|y) P(y)}{P(x)}$$

$P(x|y)$

$$P(x) = \sum_y P(x|y) P(y)$$

BAYESIAN INFERENCE FOR NEURAL DATA ANALYSIS

parameters, y



data, x

$P(y)$

$$P(y|x) = \frac{P(x|y) P(y)}{P(x)}$$

$P(x|y)$

$$P(x) = \sum_y P(x|y) P(y)$$

BAYESIAN INFERENCE FOR NEURAL DATA ANALYSIS

network connectivity

parameters, y



data, x

$P(y)$

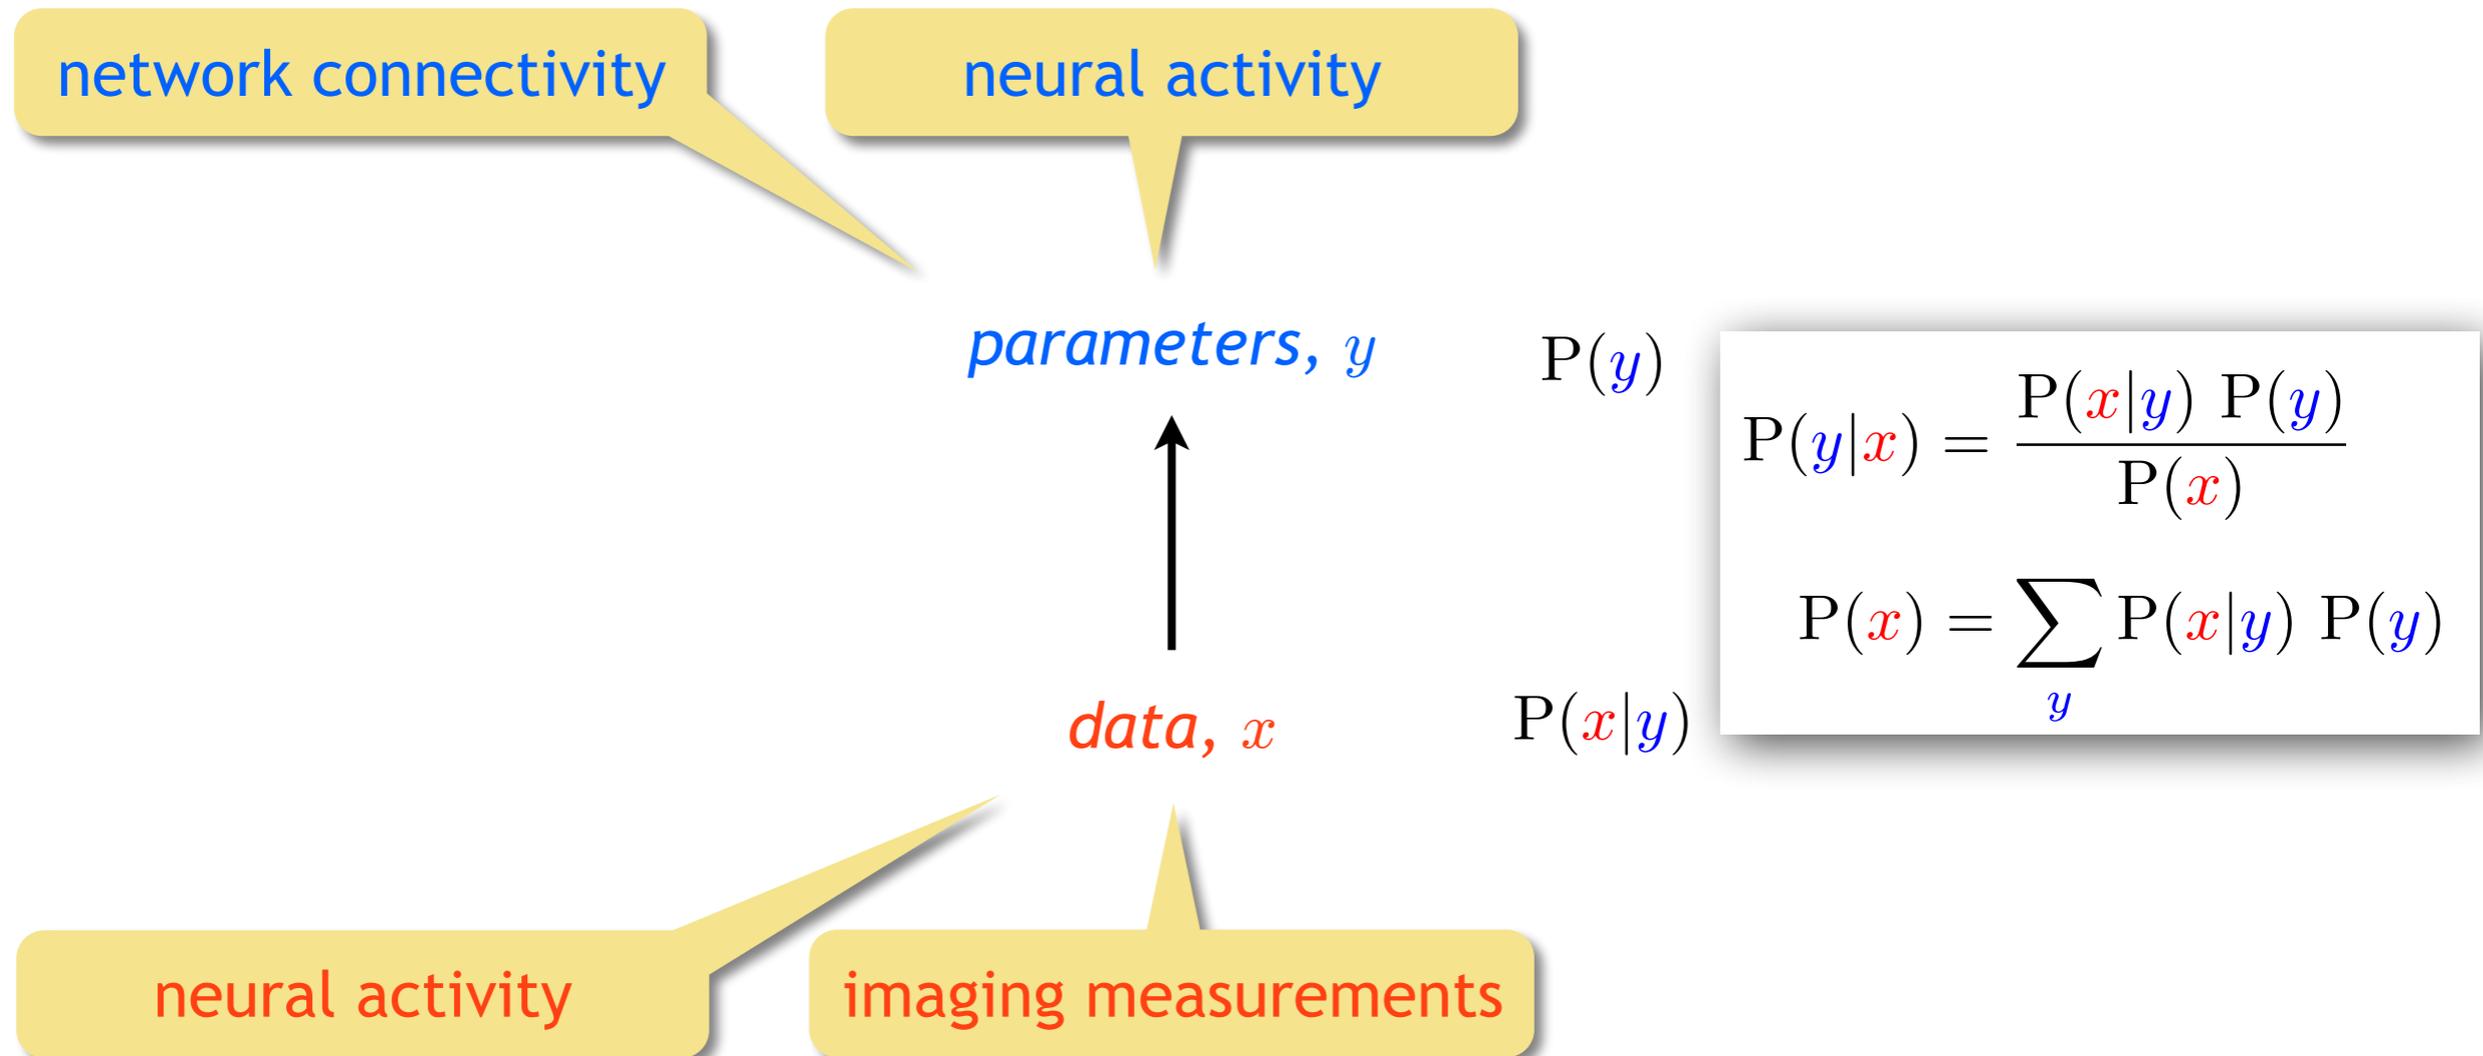
$$P(y|x) = \frac{P(x|y) P(y)}{P(x)}$$

$P(x|y)$

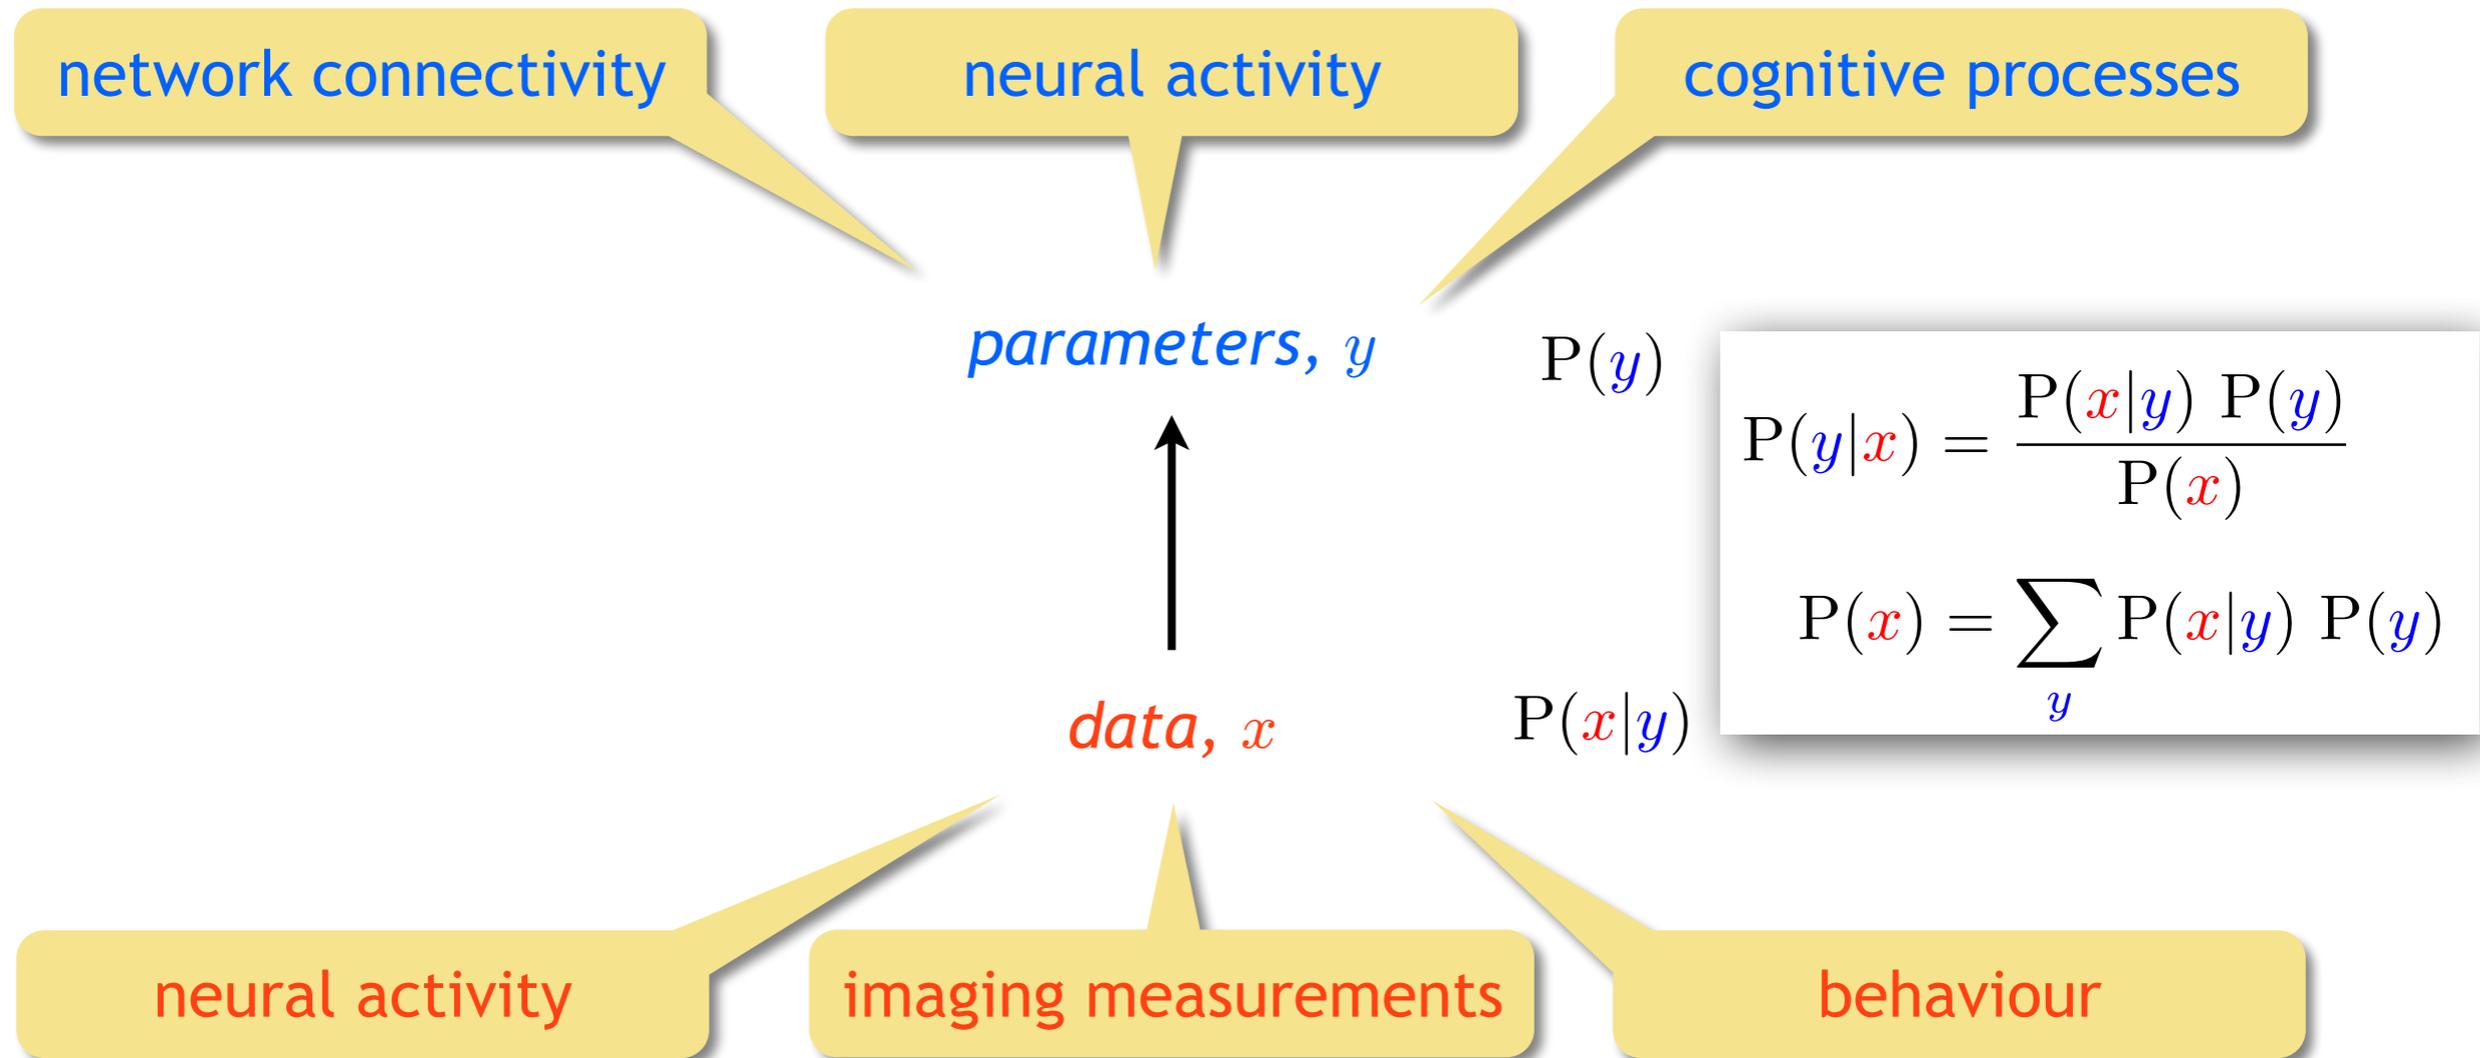
$$P(x) = \sum_y P(x|y) P(y)$$

neural activity

BAYESIAN INFERENCE FOR NEURAL DATA ANALYSIS



BAYESIAN INFERENCE FOR NEURAL DATA ANALYSIS



NEURAL COMPUTATION

How do neural circuits solve challenging computational tasks?

NEURAL COMPUTATION

How do neural circuits solve challenging computational tasks?

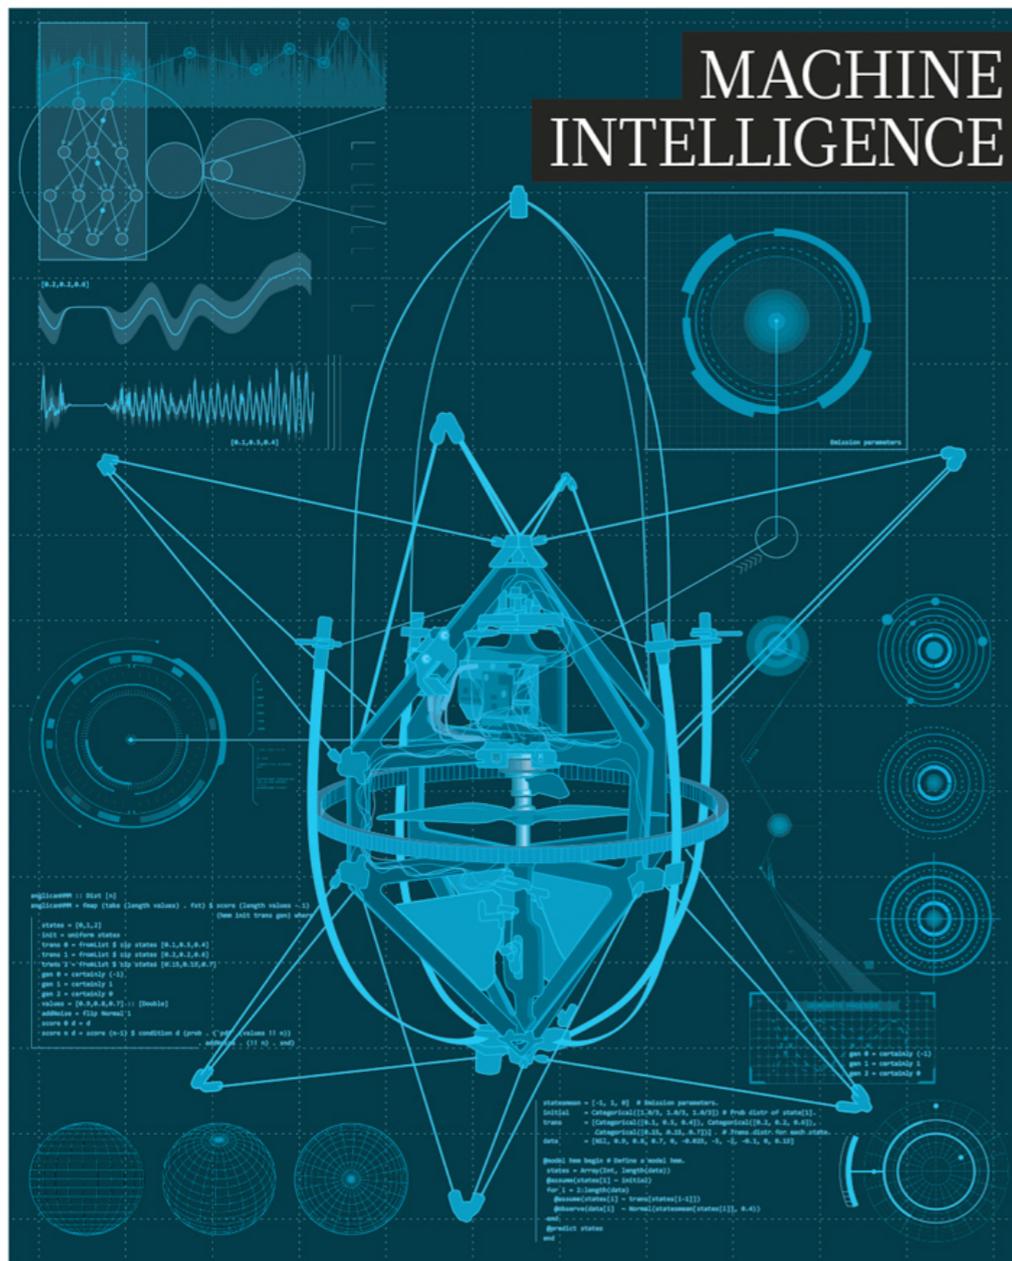
What is the computational role of elements of neural circuits?

NEURAL COMPUTATION

How do neural circuits solve challenging computational tasks?

What is the computational role of elements of neural circuits?

nature **INSIGHT**



REVIEW

doi:10.1038/nature14539

Deep learning

Yann LeCun^{1,2}, Yoshua Bengio³ & Geoffrey Hinton^{4,5}

Deep learning allows computational models that are composed of multiple processing layers to learn representations of data with multiple levels of abstraction. These methods have dramatically improved the state-of-the-art in speech recognition, visual object recognition, object detection and many other domains such as drug discovery and genomics. Deep learning discovers intricate structure in large data sets by using the backpropagation algorithm to indicate how a machine should change its internal parameters that are used to compute the representation in each layer from the representation in the previous layer. Deep convolutional nets have brought about breakthroughs in processing images, video, speech and audio, whereas recurrent nets have shone light on sequential data such as text and speech.

REVIEW

doi:10.1038/nature14541

Probabilistic machine learning and artificial intelligence

Zoubin Ghahramani¹

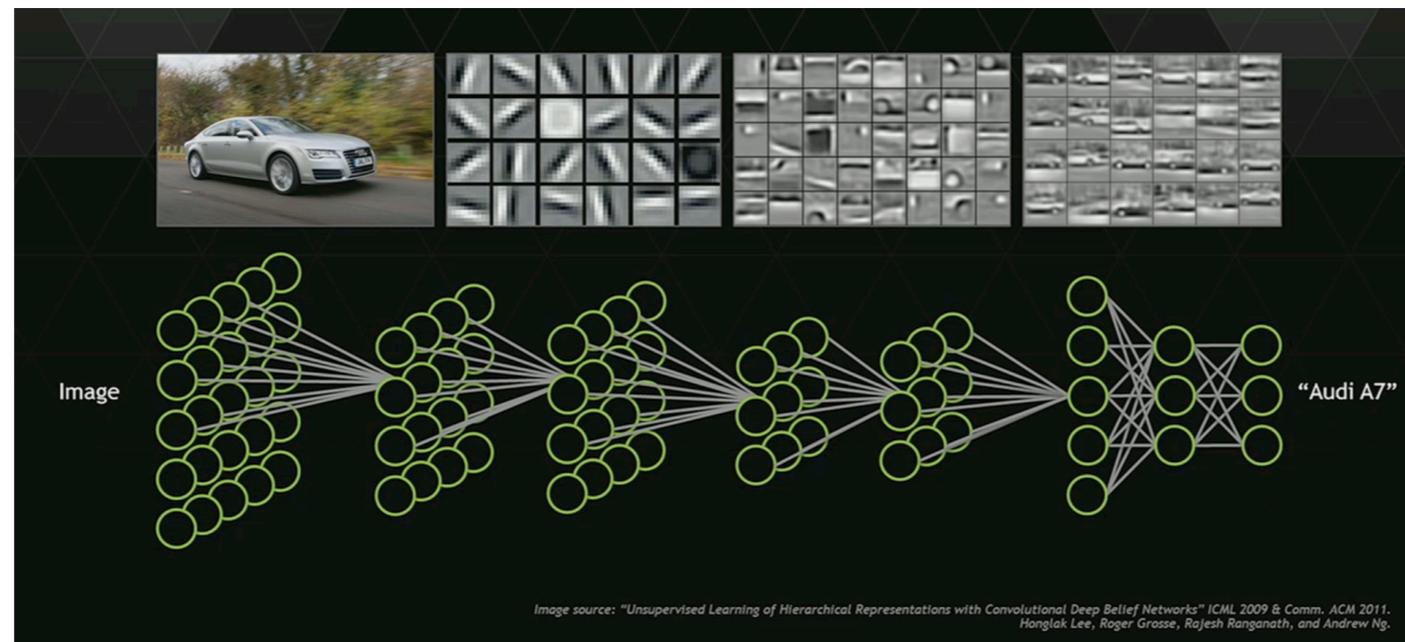
How can a machine learn from experience? Probabilistic modelling provides a framework for understanding what learning is, and has therefore emerged as one of the principal theoretical and practical approaches for designing machines that learn from data acquired through experience. The probabilistic framework, which describes how to represent and manipulate uncertainty about models and predictions, has a central role in scientific data analysis, machine learning, robotics, cognitive science and artificial intelligence. This Review provides an introduction to this framework, and discusses some of the state-of-the-art advances in the field, namely, probabilistic programming, Bayesian optimization, data compression and automatic model discovery.

NEURAL COMPUTATION

deep learning algorithms

NEURAL COMPUTATION

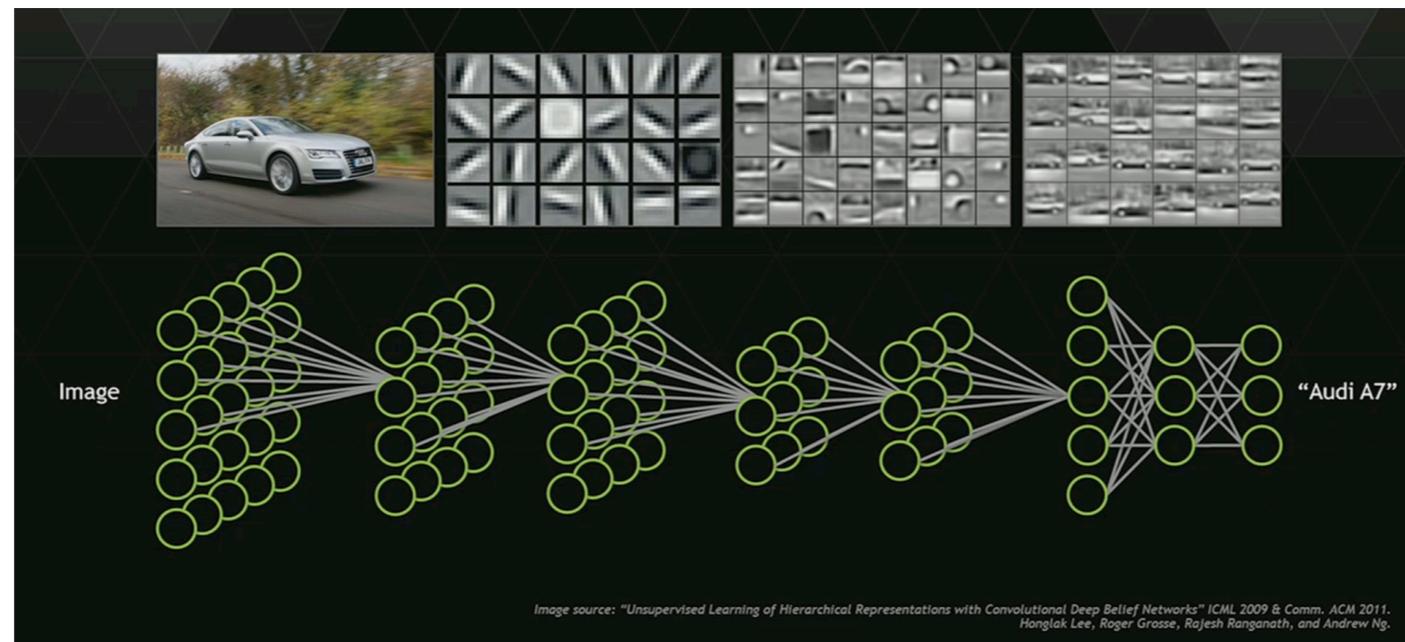
deep learning algorithms



NEURAL COMPUTATION

deep learning algorithms

- + perform well on difficult tasks
- + scale up to large data sets
/ parameter spaces
- + end-to-end training

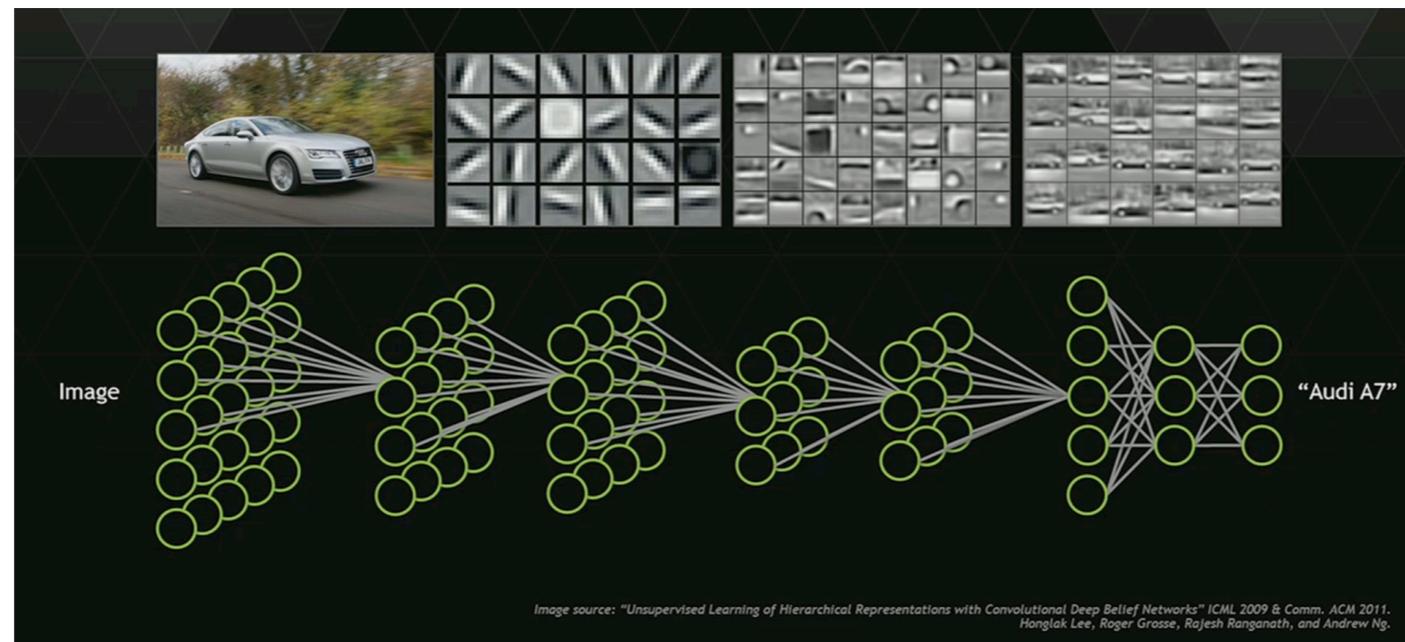


NEURAL COMPUTATION

deep learning algorithms

- + perform well on difficult tasks
- + scale up to large data sets
/ parameter spaces
- + end-to-end training

- + they are **neural** networks
- + realistic receptive fields



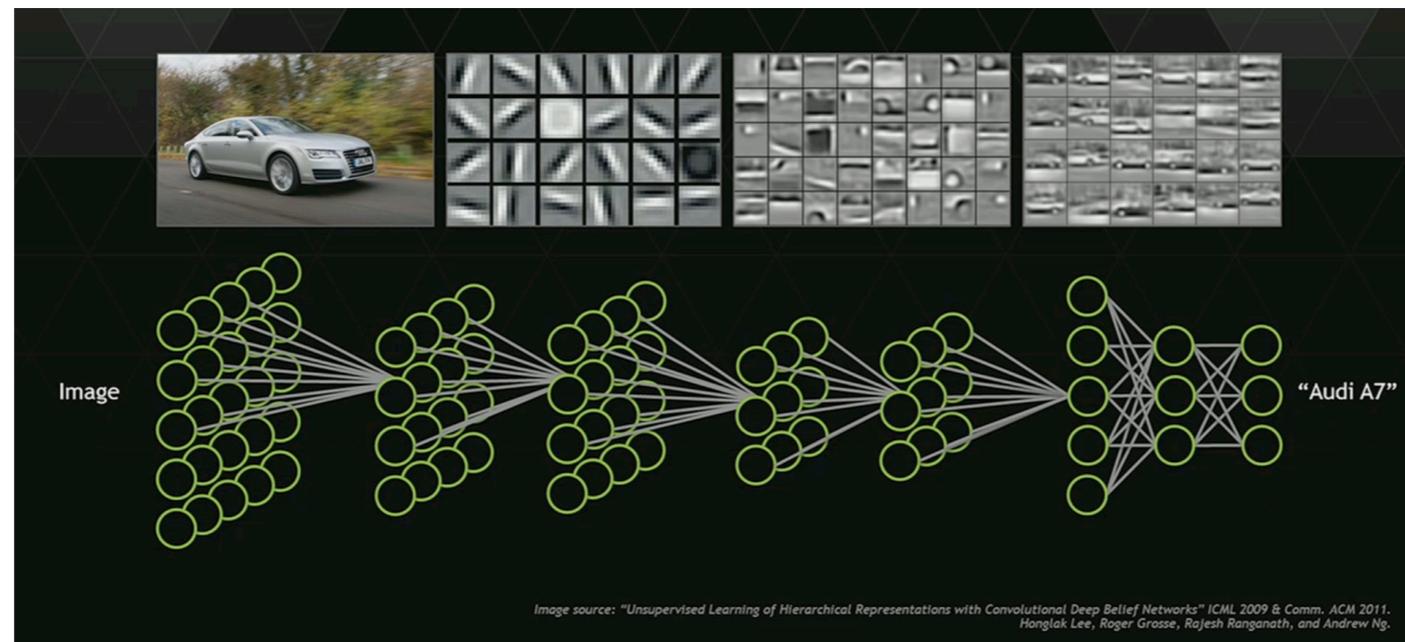
NEURAL COMPUTATION

deep learning algorithms

- + perform well on difficult tasks
- + scale up to large data sets / parameter spaces
- + end-to-end training

- + they are **neural networks**
- + realistic receptive fields

- supervised learning
- need lots of data
- susceptible to perturbations
adversarial samples
- do not represent uncertainty
or only at the final stage of computation



NEURAL COMPUTATION

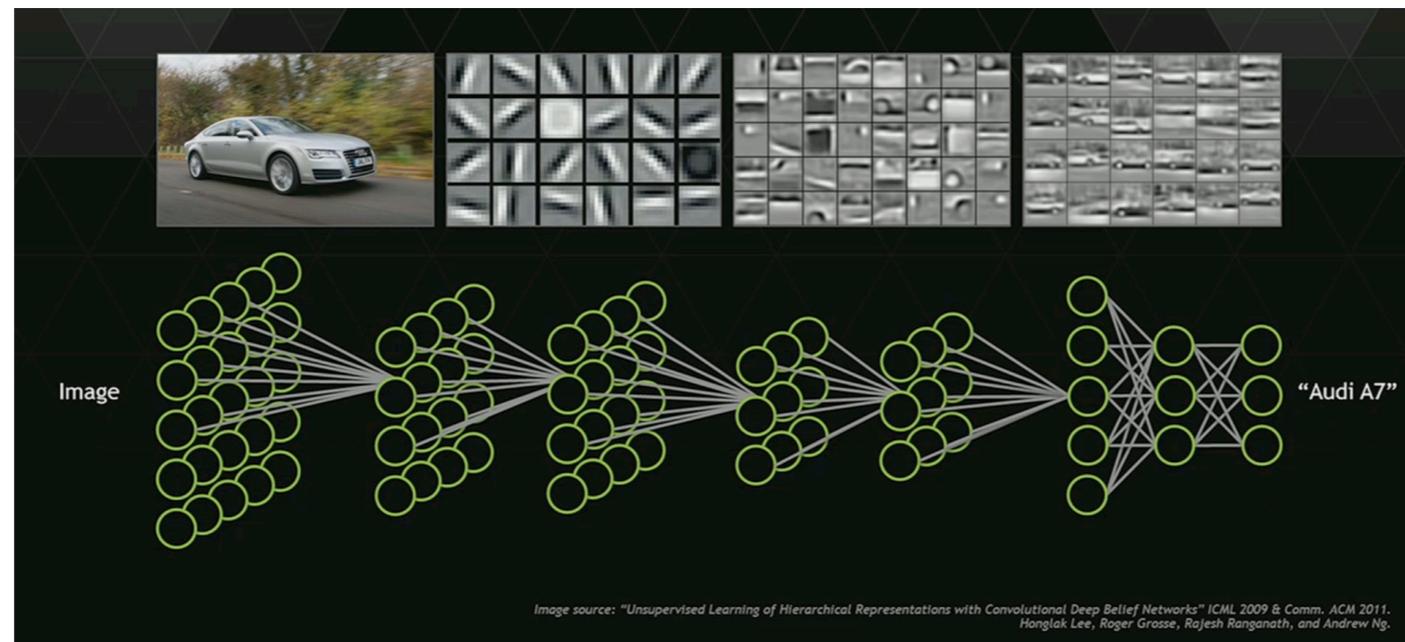
deep learning algorithms

- + perform well on difficult tasks
- + scale up to large data sets / parameter spaces
- + end-to-end training

- supervised learning
- need lots of data
- susceptible to perturbations
adversarial samples
- do not represent uncertainty
or only at the final stage of computation

- + they are **neural** networks
- + realistic receptive fields

- feed-forward
- mixed E and I
- saturating or threshold-linear units
- no dynamics, trial-average behaviour



NEURAL COMPUTATION

deep learning algorithms

- + perform well on difficult tasks
- + scale up to large data sets / parameter spaces
- + end-to-end training
- supervised learning
- need lots of data
- susceptible to perturbations
adversarial samples
- do not represent uncertainty
or only at the final stage of computation
- + they are **neural** networks
- + realistic receptive fields
- feed-forward
- mixed E and I
- saturating or threshold-linear units
- no dynamics, trial-average behaviour

probabilistic algorithms

NEURAL COMPUTATION

deep learning algorithms

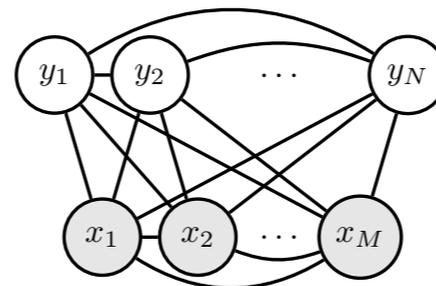
- + perform well on difficult tasks
- + scale up to large data sets / parameter spaces
- + end-to-end training

- supervised learning
- need lots of data
- susceptible to perturbations
adversarial samples
- do not represent uncertainty
or only at the final stage of computation

- + they are **neural networks**
- + realistic receptive fields

- feed-forward
- mixed E and I
- saturating or threshold-linear units
- no dynamics, trial-average behaviour

probabilistic algorithms



latent causes “Audi A7”

observations



NEURAL COMPUTATION

deep learning algorithms

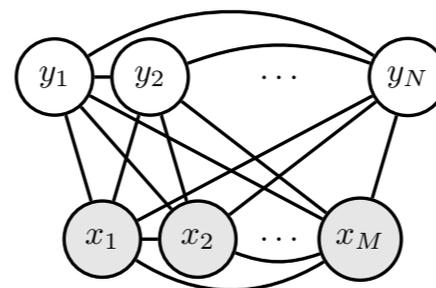
- + perform well on difficult tasks
- + scale up to large data sets / parameter spaces
- + end-to-end training

- supervised learning
- need lots of data
- susceptible to perturbations
adversarial samples
- do not represent uncertainty
or only at the final stage of computation

- + they are **neural networks**
- + realistic receptive fields

- feed-forward
- mixed E and I
- saturating or threshold-linear units
- no dynamics, trial-average behaviour

probabilistic algorithms



latent causes

“Audi A7”

observations



noise
ambiguity
partial knowledge

NEURAL COMPUTATION

deep learning algorithms

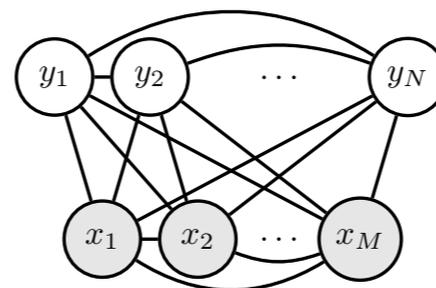
- + perform well on difficult tasks
- + scale up to large data sets / parameter spaces
- + end-to-end training

- supervised learning
- need lots of data
- susceptible to perturbations
adversarial samples
- do not represent uncertainty
or only at the final stage of computation

- + they are **neural networks**
- + realistic receptive fields

- feed-forward
- mixed E and I
- saturating or threshold-linear units
- no dynamics, trial-average behaviour

probabilistic algorithms



latent causes

“Audi A7”

$P(x, y)$

observations



noise
ambiguity
partial knowledge

NEURAL COMPUTATION

deep learning algorithms

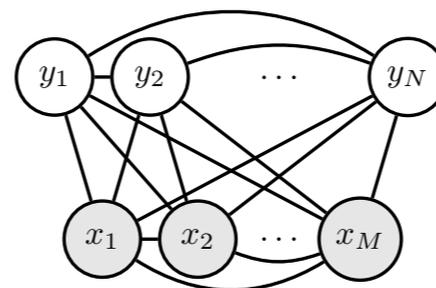
- + perform well on difficult tasks
- + scale up to large data sets / parameter spaces
- + end-to-end training

- supervised learning
- need lots of data
- susceptible to perturbations
adversarial samples
- do not represent uncertainty
or only at the final stage of computation

- + they are **neural networks**
- + realistic receptive fields

- feed-forward
- mixed E and I
- saturating or threshold-linear units
- no dynamics, trial-average behaviour

probabilistic algorithms



latent causes

“Audi A7”

$$P(y|x) \propto \frac{P(x, y)}{\int P(x, y') dy'}$$

observations



noise
ambiguity
partial knowledge

NEURAL COMPUTATION

deep learning algorithms

- + perform well on difficult tasks
- + scale up to large data sets / parameter spaces
- + end-to-end training

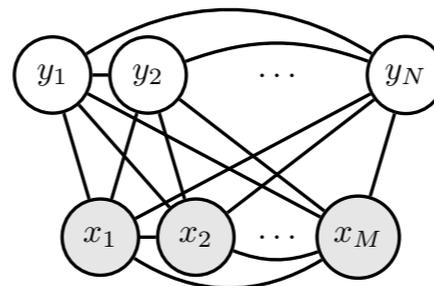
- supervised learning
- need lots of data
- susceptible to perturbations
adversarial samples
- do not represent uncertainty
or only at the final stage of computation

- + they are **neural networks**
- + realistic receptive fields

- feed-forward
- mixed E and I
- saturating or threshold-linear units
- no dynamics, trial-average behaviour

probabilistic algorithms

- + perform well on some tasks
- + unsupervised learning
- + generalise from few data points
one shot learning
- + robust to noisy and missing data
- + represent **uncertainty**
at every stage stage of computation



latent causes

“Audi A7”

observations

$$P(y|x) \propto \frac{P(x, y)}{\int P(x, y') dy'}$$



noise
ambiguity
partial knowledge

NEURAL COMPUTATION

deep learning algorithms

- + perform well on difficult tasks
- + scale up to large data sets / parameter spaces
- + end-to-end training

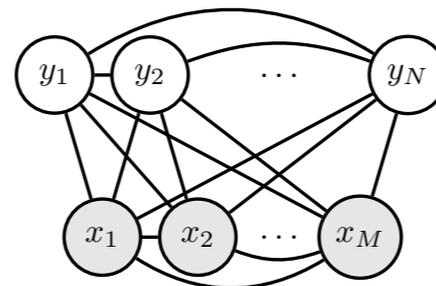
- supervised learning
- need lots of data
- susceptible to perturbations
adversarial samples
- do not represent uncertainty
or only at the final stage of computation

- + they are **neural networks**
- + realistic receptive fields

- feed-forward
- mixed E and I
- saturating or threshold-linear units
- no dynamics, trial-average behaviour

probabilistic algorithms

- + perform well on some tasks
- + unsupervised learning
- + generalise from few data points
one shot learning
- + robust to noisy and missing data
- + represent **uncertainty**
at every stage of computation



latent causes

“Audi A7”

$$P(y|x) \propto \frac{P(x, y)}{\int P(x, y') dy'}$$

observations



noise
ambiguity
partial knowledge

- fewer success stories for performance
- computationally intensive
- do not scale well
- end-to-end training difficult

NEURAL COMPUTATION

deep learning algorithms

- + perform well on difficult tasks
- + scale up to large data sets / parameter spaces
- + end-to-end training

- supervised learning
- need lots of data
- susceptible to perturbations
adversarial samples
- do not represent uncertainty
or only at the final stage of computation

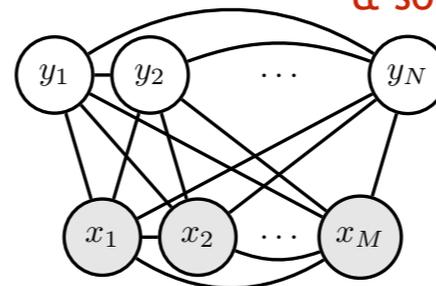
- + they are **neural networks**
- + realistic receptive fields

- feed-forward
- mixed E and I
- saturating or threshold-linear units
- no dynamics, trial-average behaviour

probabilistic algorithms

- + perform well on some tasks
- + unsupervised learning
- + generalise from few data points
one shot learning
- + robust to noisy and missing data
- + represent **uncertainty**
at every stage of computation

- + behavioural evidence
& some neural evidence



latent causes

“Audi A7”

$$P(y|x) \propto \frac{P(x, y)}{\int P(x, y') dy'}$$

observations



noise
ambiguity
partial knowledge

- fewer success stories for performance
- computationally intensive
- do not scale well
- end-to-end training difficult

NEURAL COMPUTATION

deep learning algorithms

- + perform well on difficult tasks
- + scale up to large data sets / parameter spaces
- + end-to-end training

- supervised learning
- need lots of data
- susceptible to perturbations
adversarial samples
- do not represent uncertainty
or only at the final stage of computation

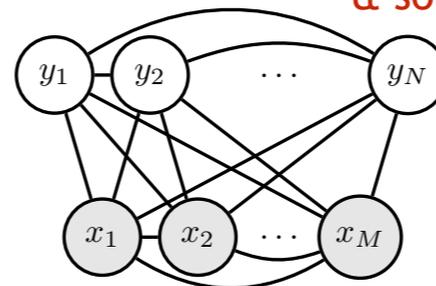
- + they are **neural networks**
- + realistic receptive fields

- feed-forward
- mixed E and I
- saturating or threshold-linear units
- no dynamics, trial-average behaviour

probabilistic algorithms

- + perform well on some tasks
- + unsupervised learning
- + generalise from few data points
one shot learning
- + robust to noisy and missing data
- + represent **uncertainty**
at every stage of computation

- + behavioural evidence
& some neural evidence



latent causes

“Audi A7”

$$P(y|x) \propto \frac{P(x, y)}{\int P(x, y') dy'}$$

observations



- not neural-like

noise
ambiguity
partial knowledge

- fewer success stories for performance
- computationally intensive
- do not scale well
- end-to-end training difficult

UNCONSCIOUS INFERENCES

UNCONSCIOUS INFERENCES

which words?

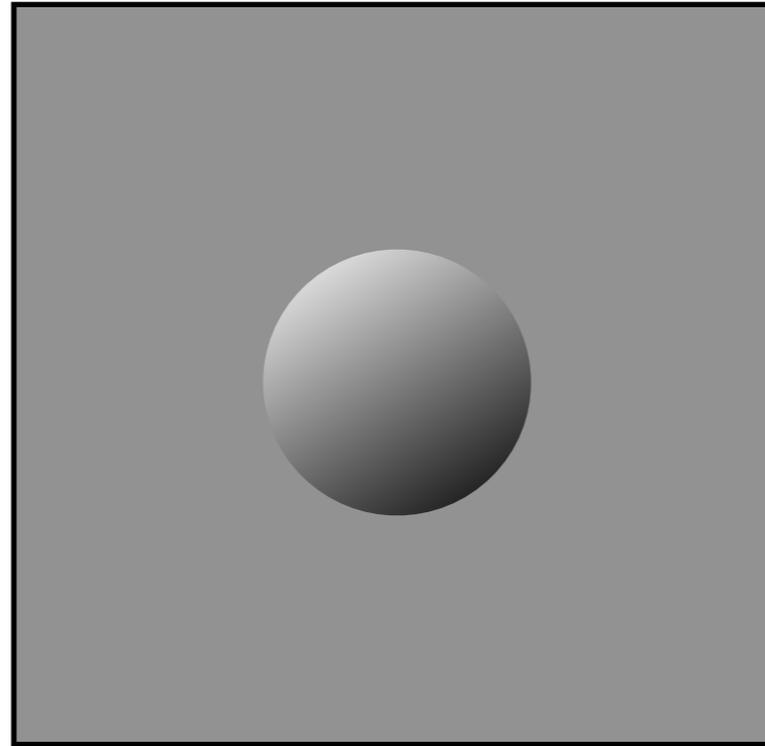
UNCONSCIOUS INFERENCES

which words?

UNCONSCIOUS INFERENCES

which words?

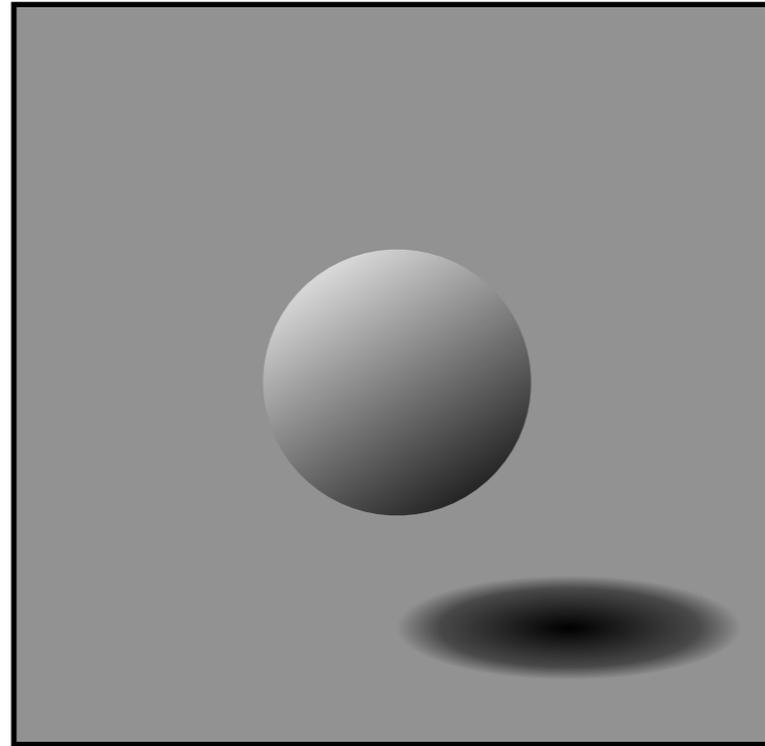
bubble or dimple?



UNCONSCIOUS INFERENCES

which words?

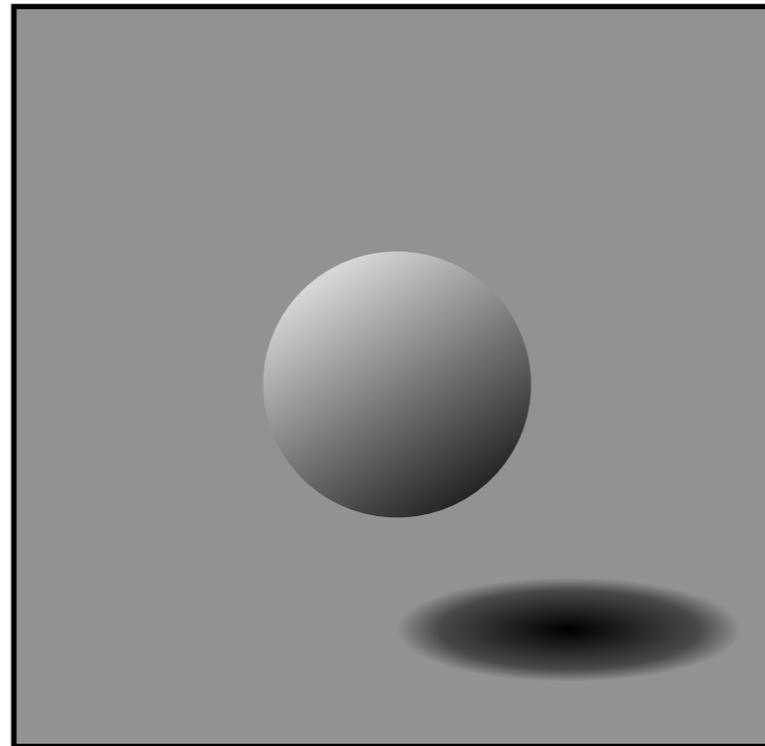
bubble or dimple?



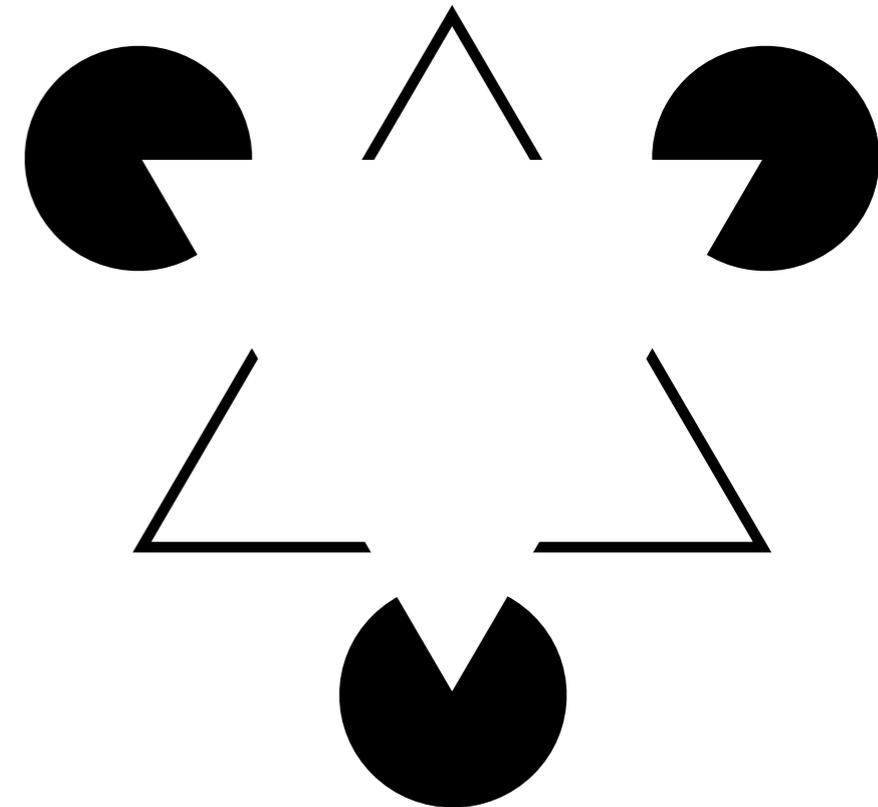
UNCONSCIOUS INFERENCES

which words?

bubble or dimple?



how many triangles?

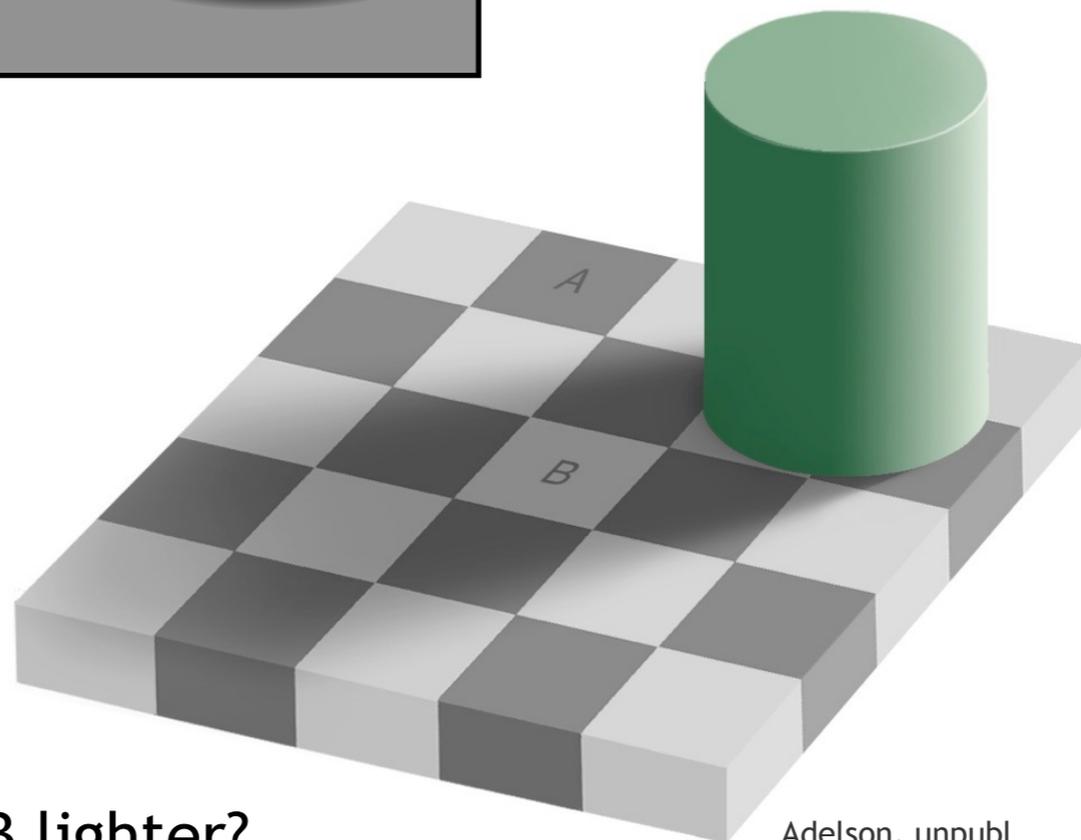
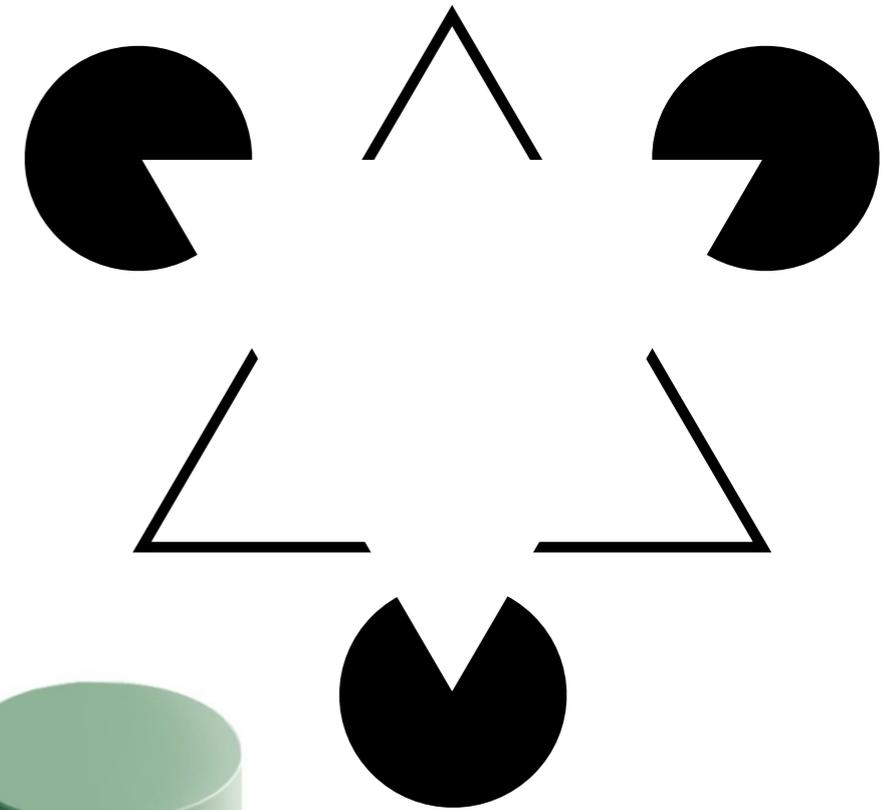
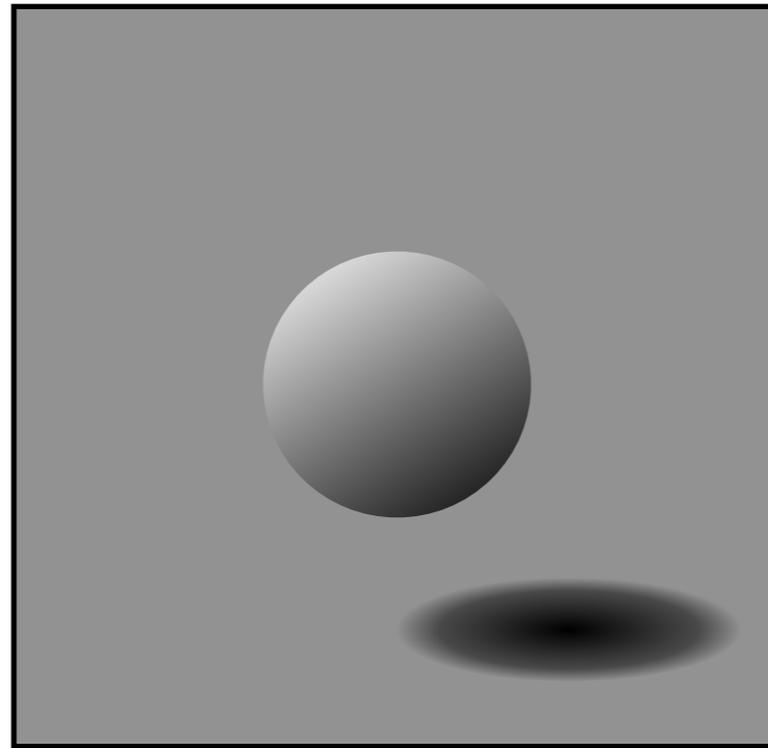


UNCONSCIOUS INFERENCES

which words?

bubble or dimple?

how many triangles?



Adelson, unpubl

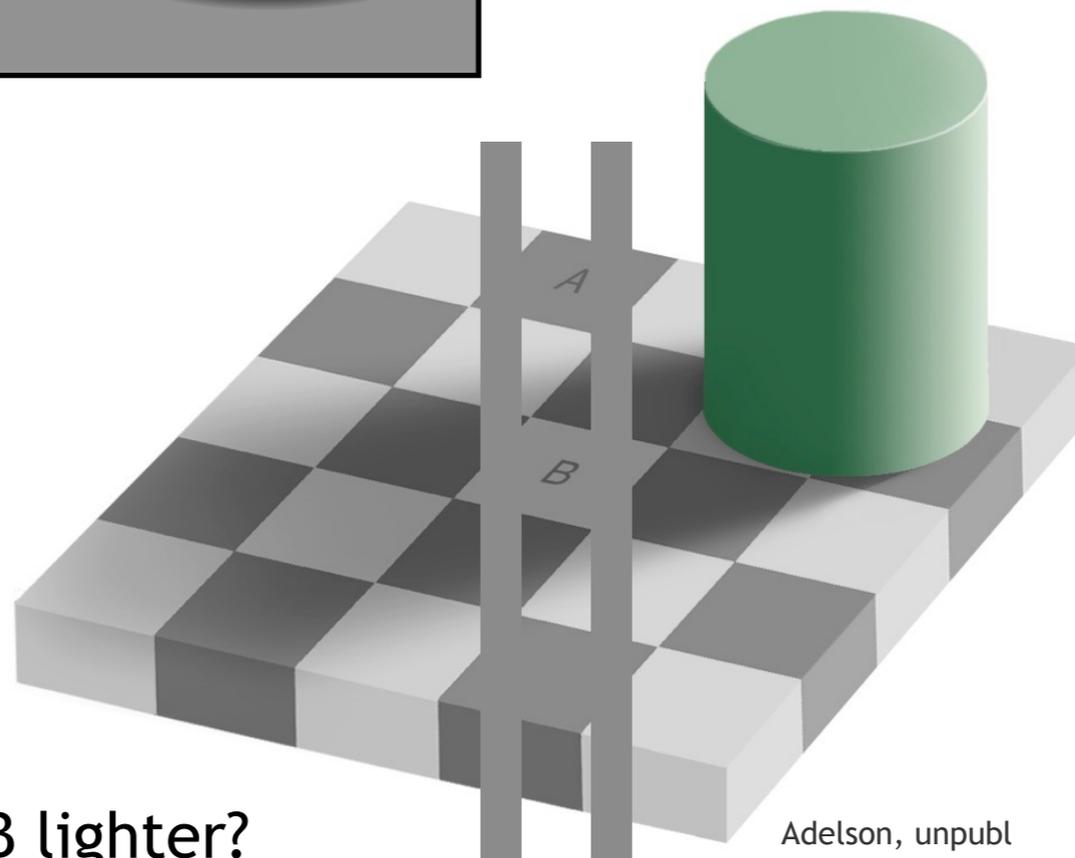
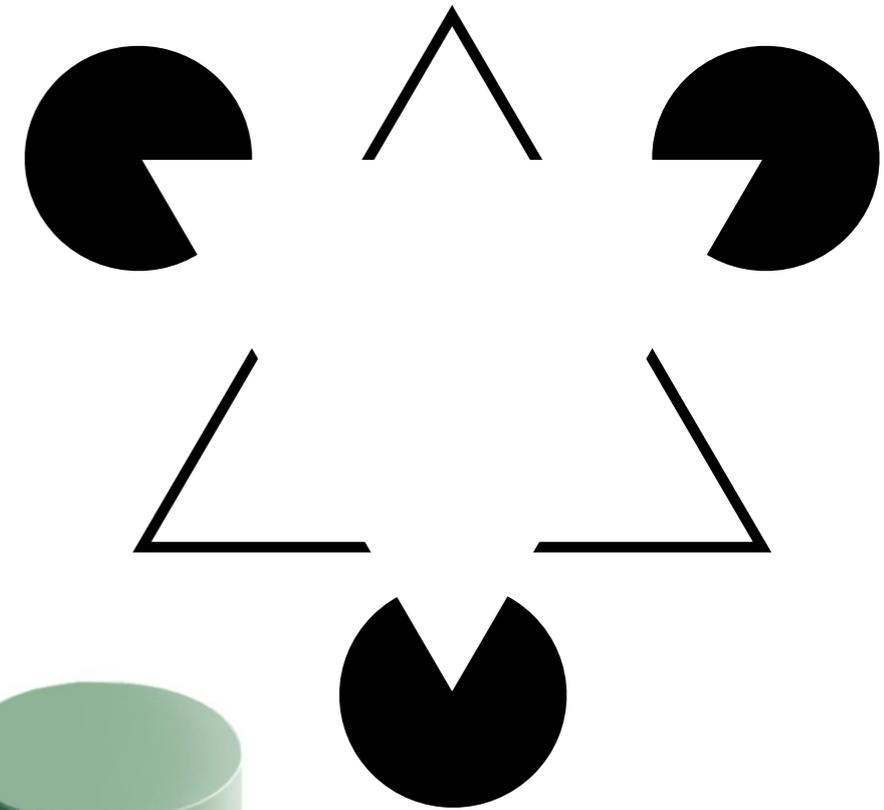
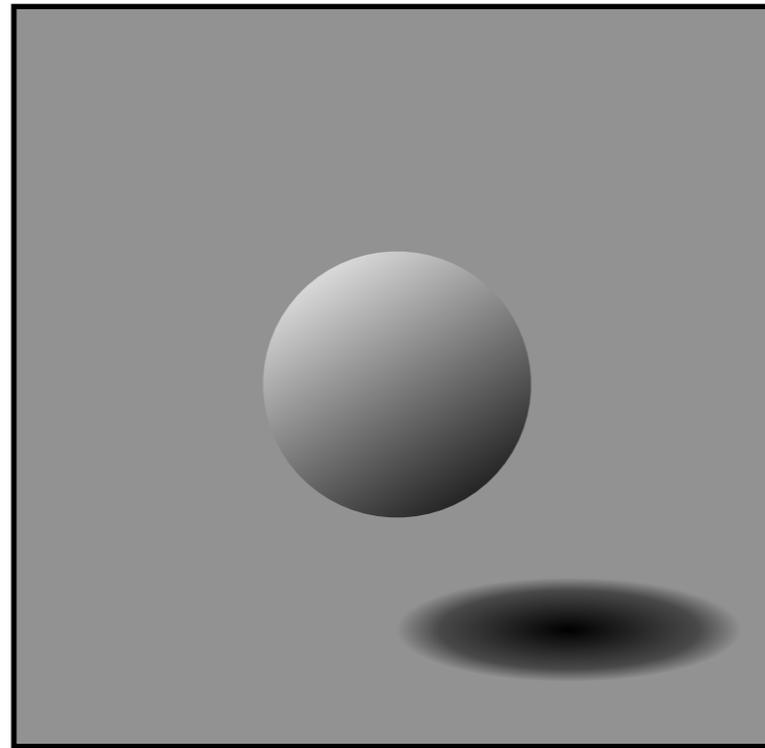
is A or B lighter?

UNCONSCIOUS INFERENCES

which words?

bubble or dimple?

how many triangles?



Adelson, unpubl

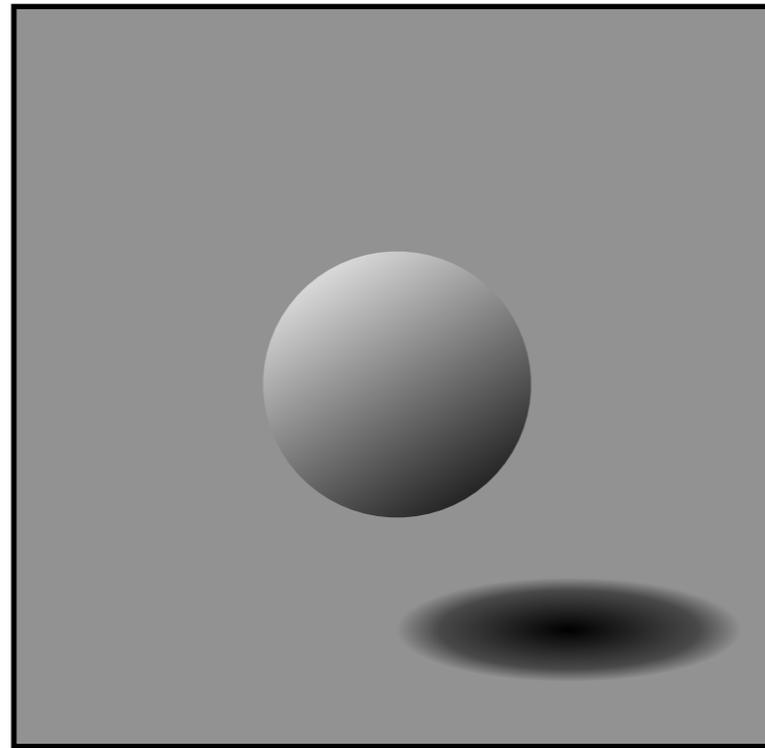
is A or B lighter?

UNCONSCIOUS INFERENCES

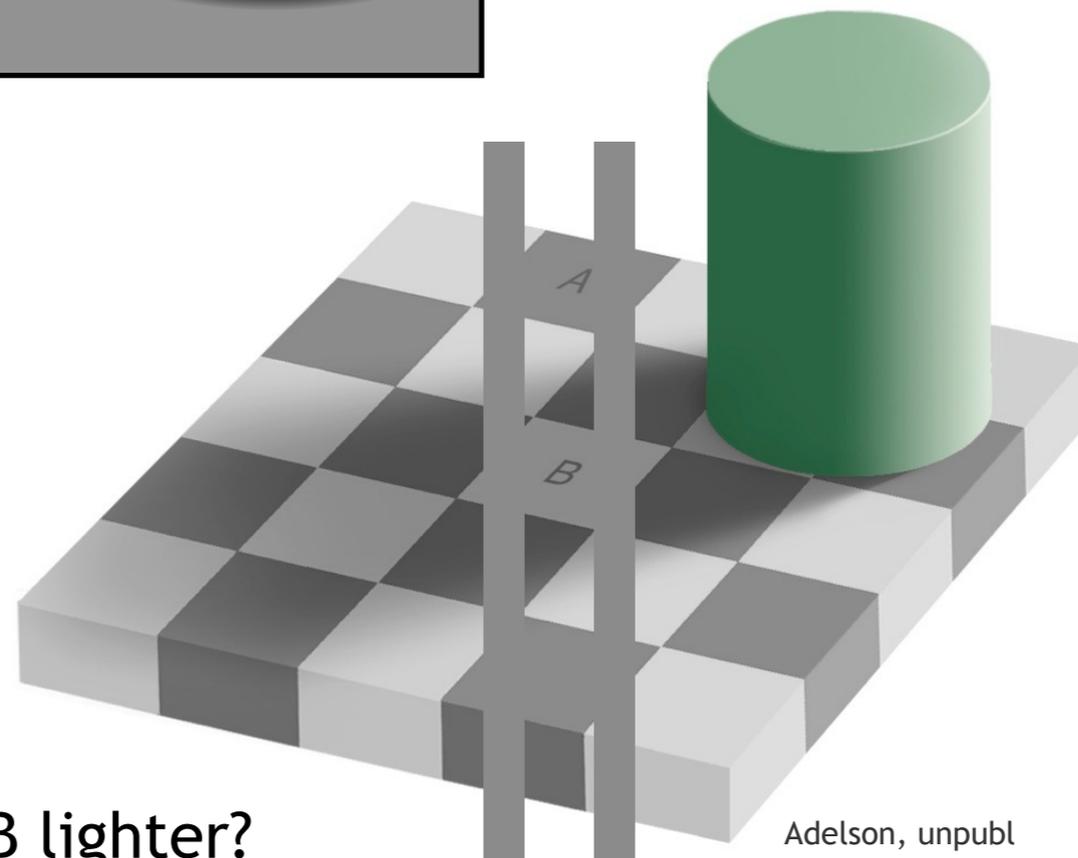
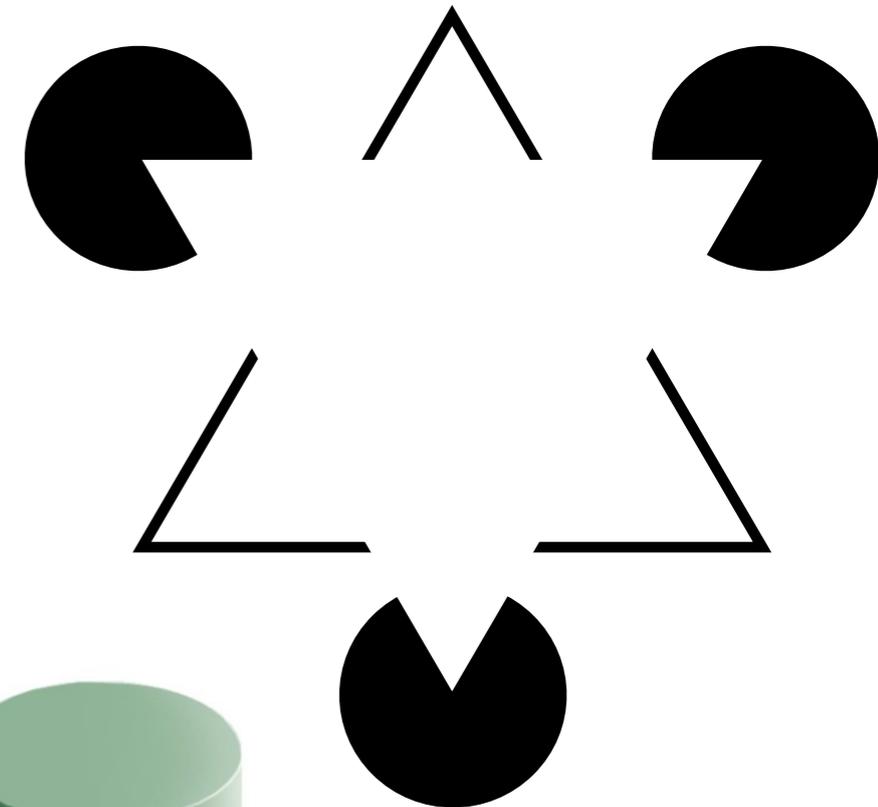
which words?

seat

bubble or dimple?



how many triangles?



is A or B lighter?

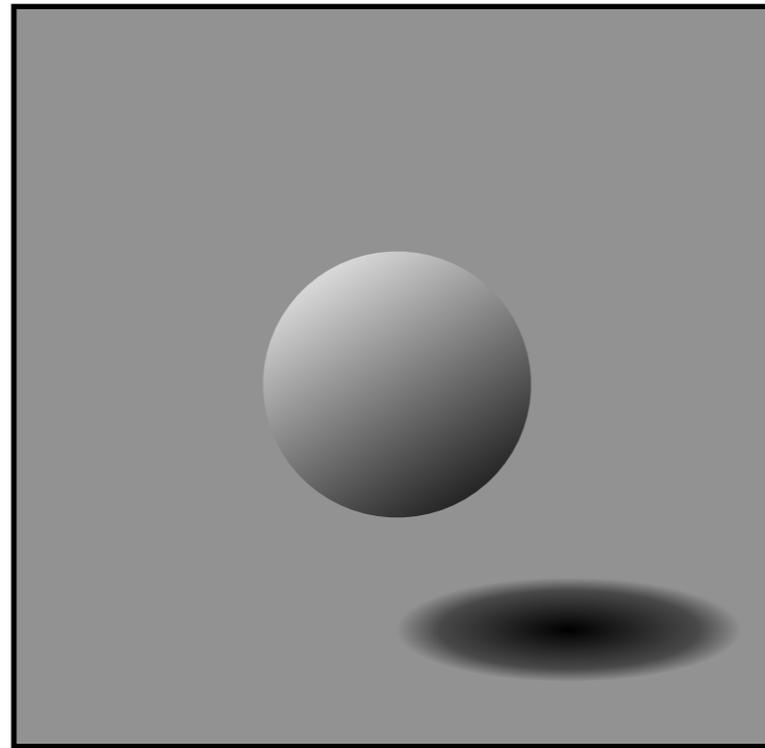
Adelson, unpubl

UNCONSCIOUS INFERENCES

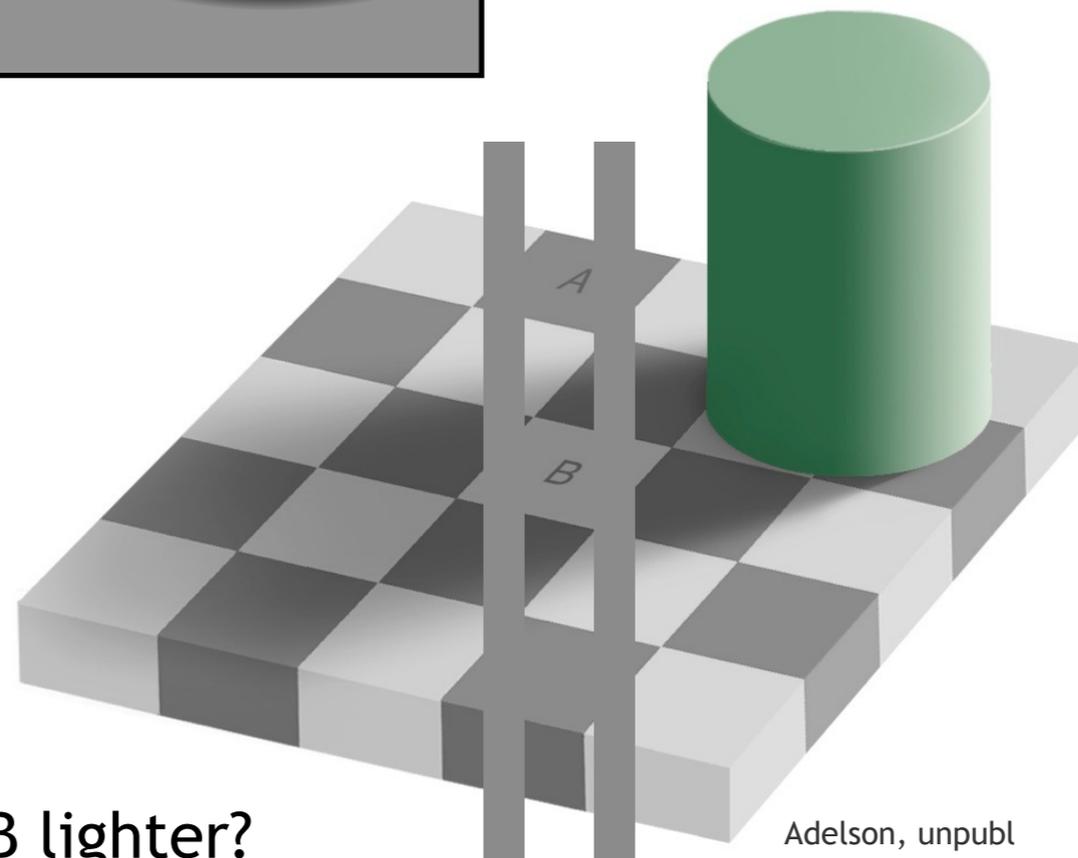
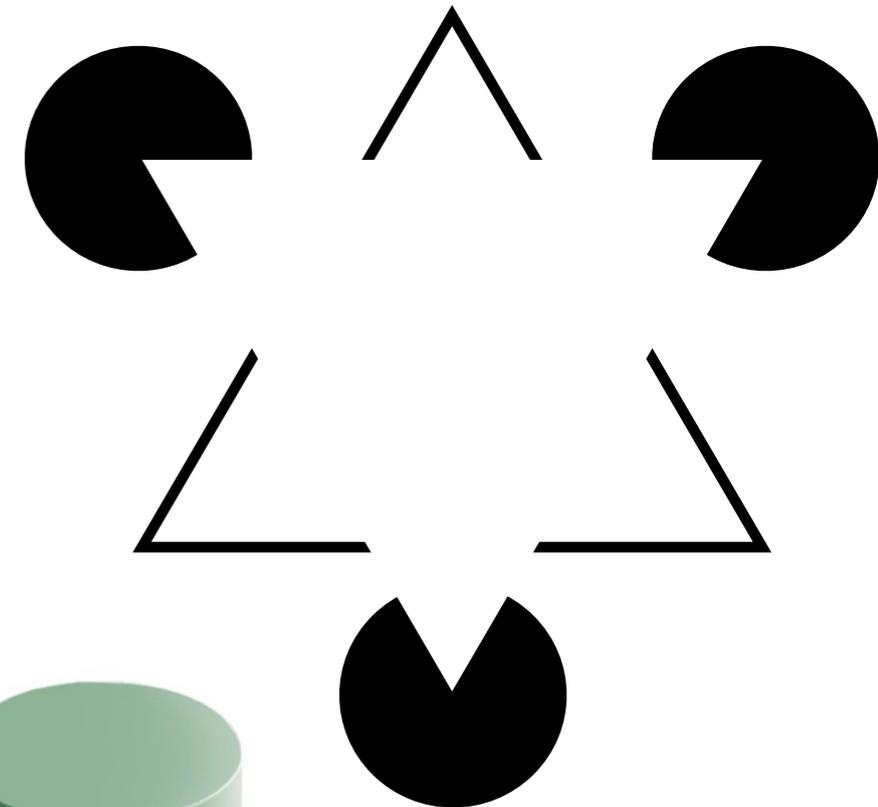
which words?

✓ seat

bubble or dimple?



how many triangles?



is A or B lighter?

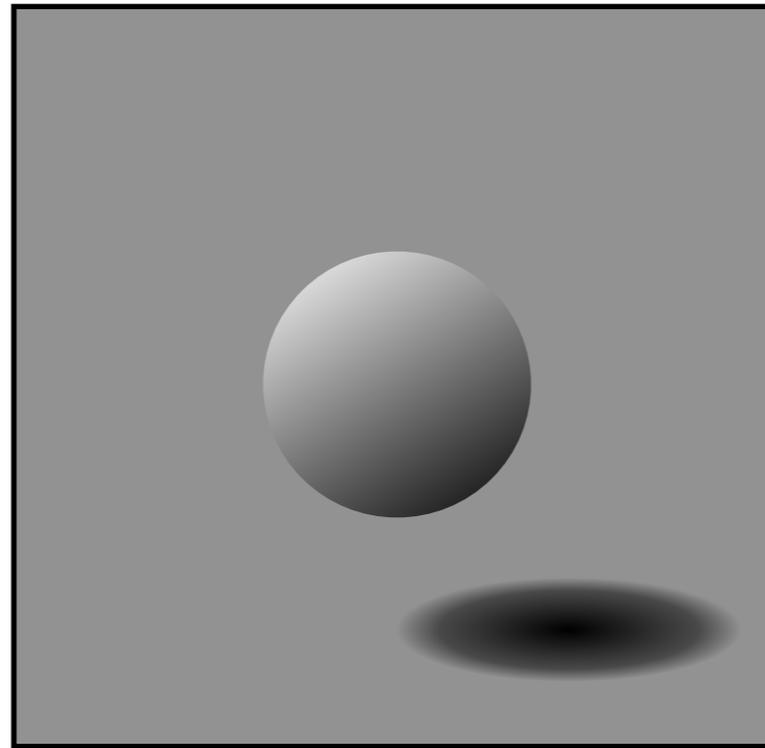
Adelson, unpubl

UNCONSCIOUS INFERENCES

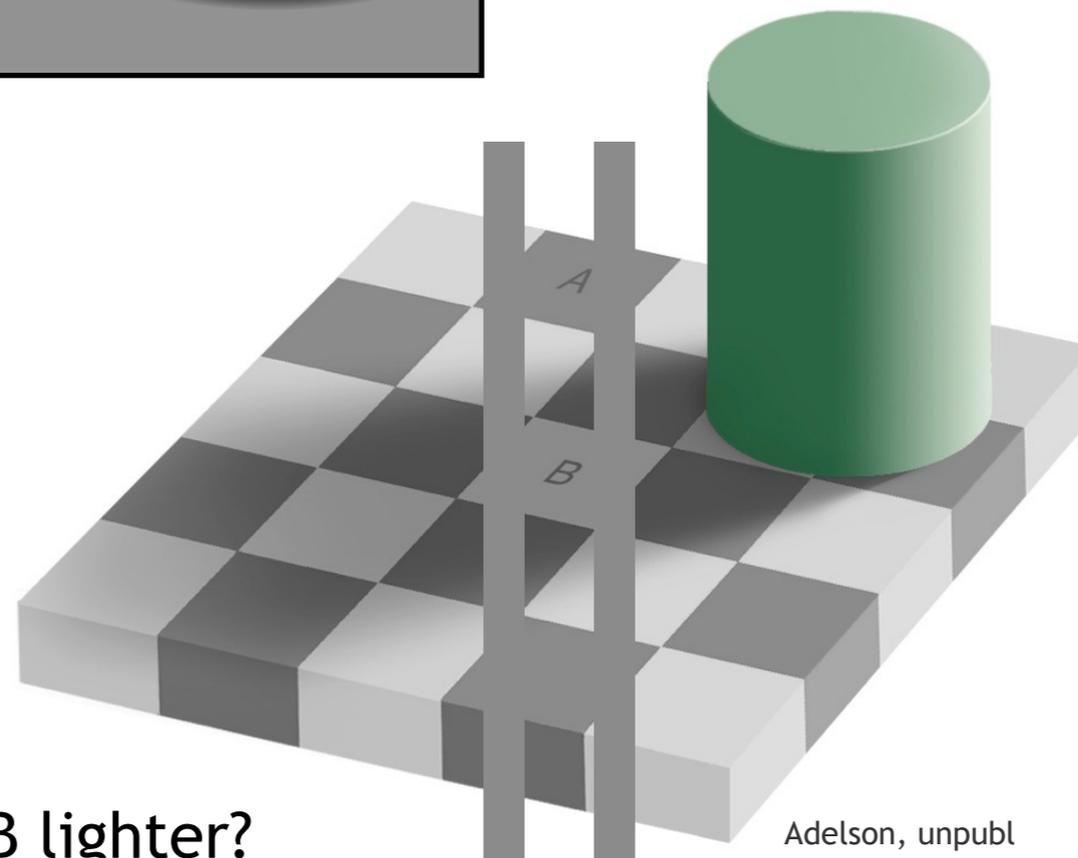
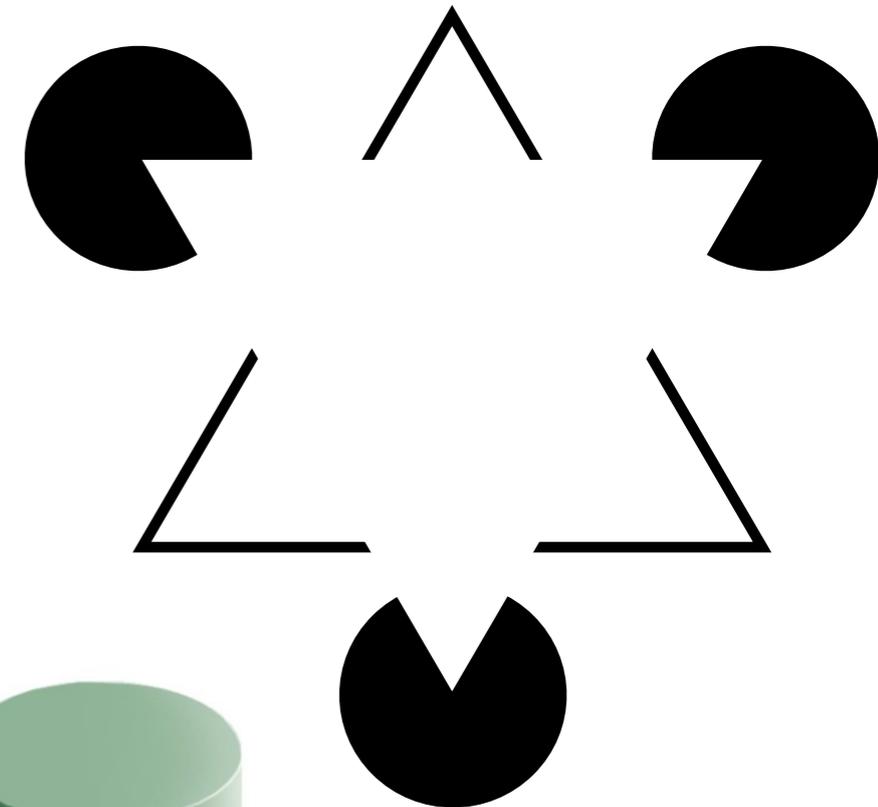
which words?

✓ seat
radio

bubble or dimple?



how many triangles?



is A or B lighter?

Adelson, unpubl

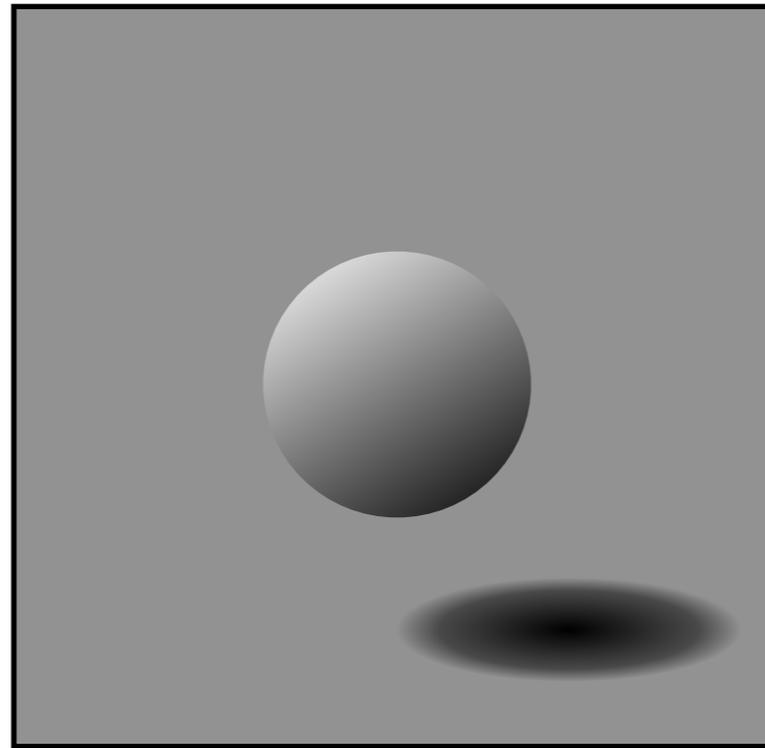
UNCONSCIOUS INFERENCES

which words?

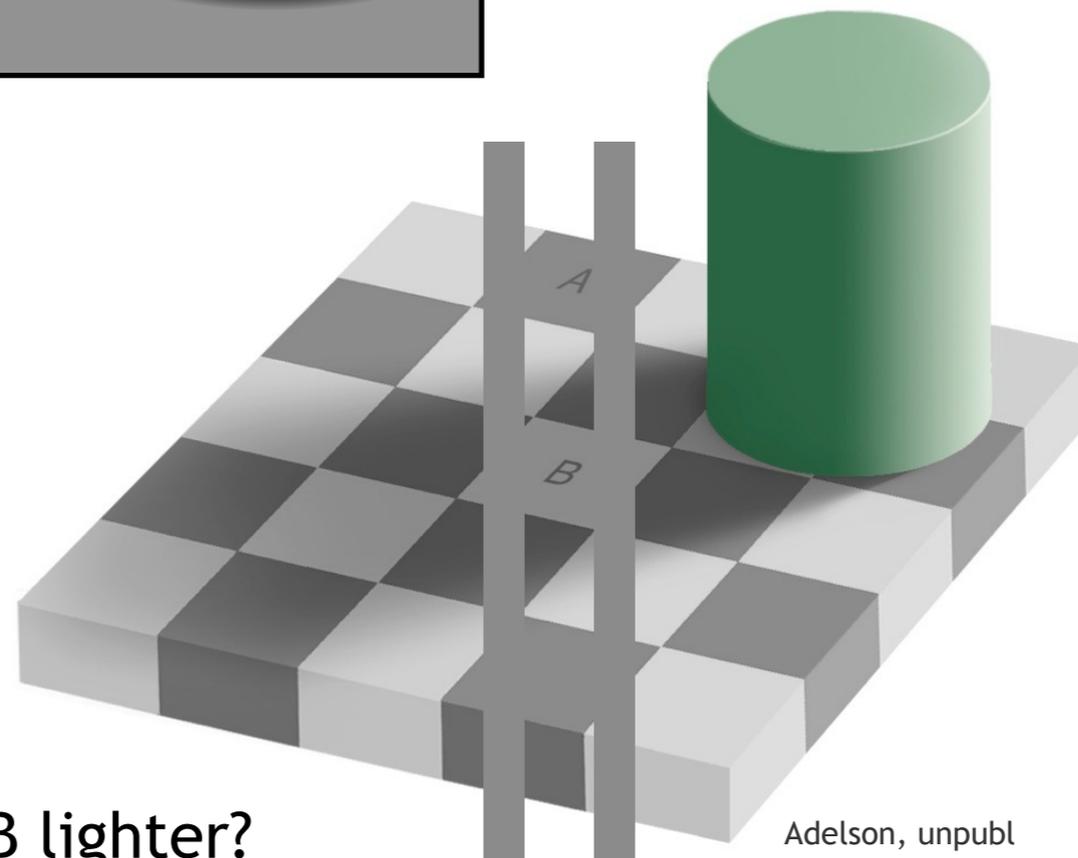
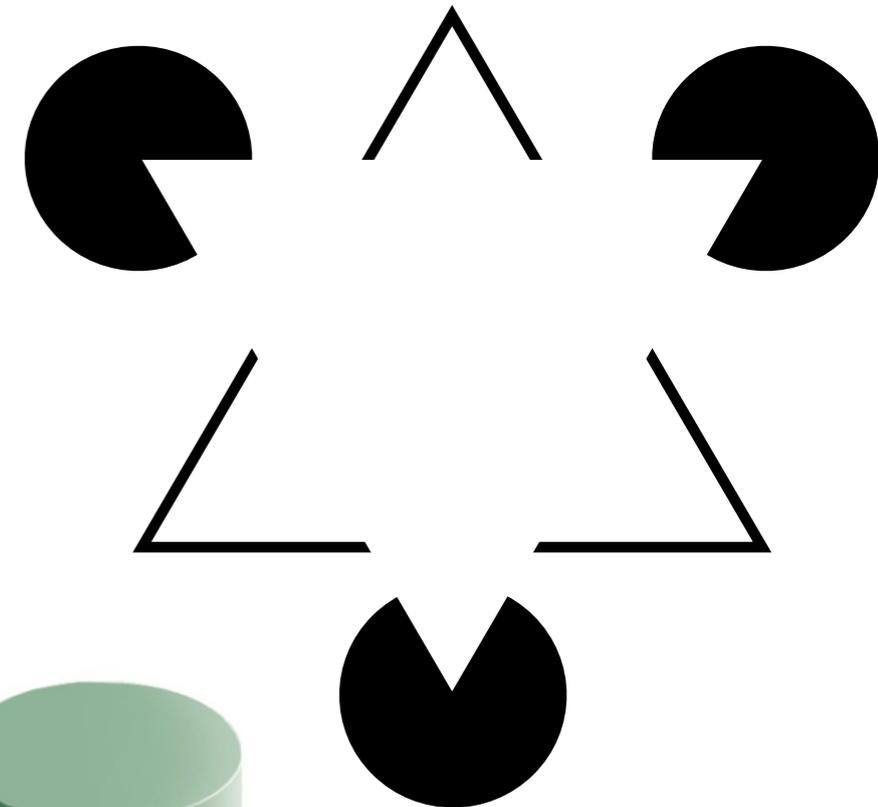
✓ seat

✗ radio

bubble or dimple?



how many triangles?



is A or B lighter?

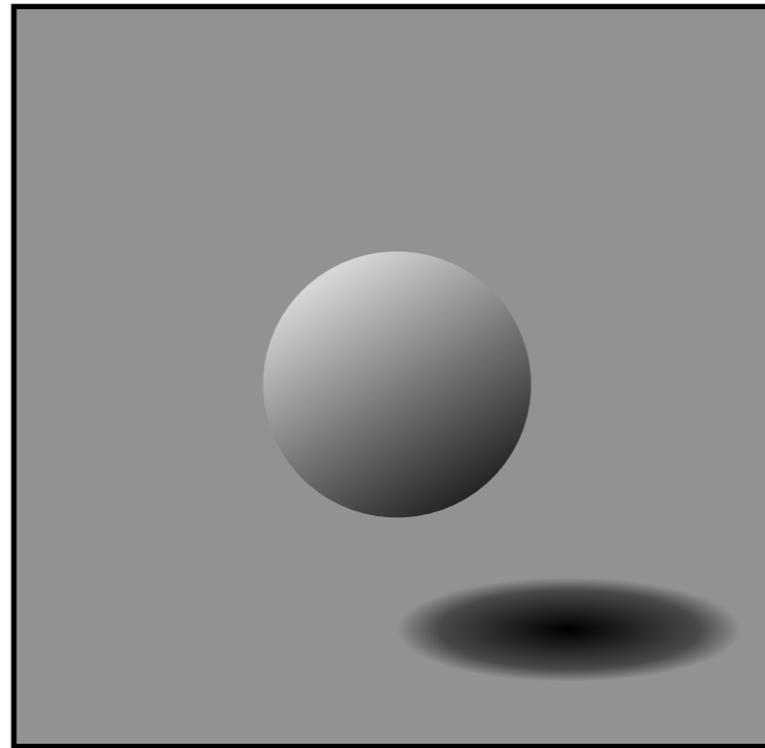
Adelson, unpubl

UNCONSCIOUS INFERENCES

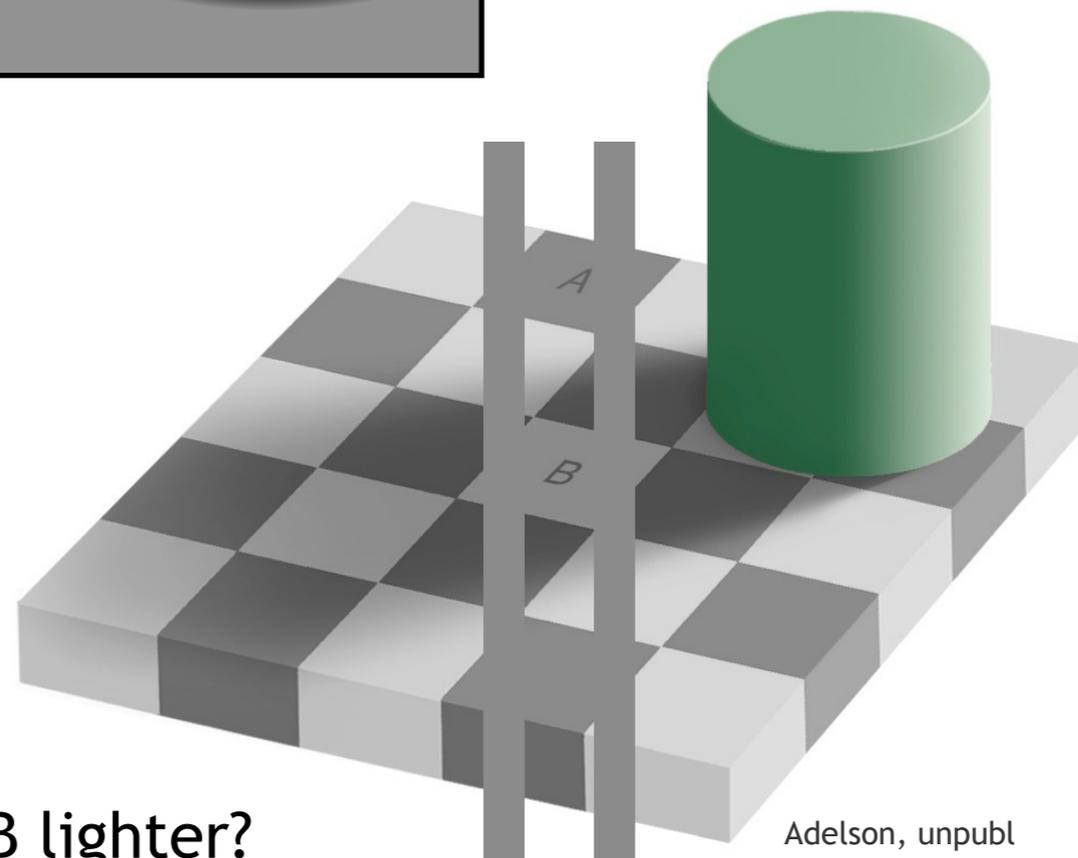
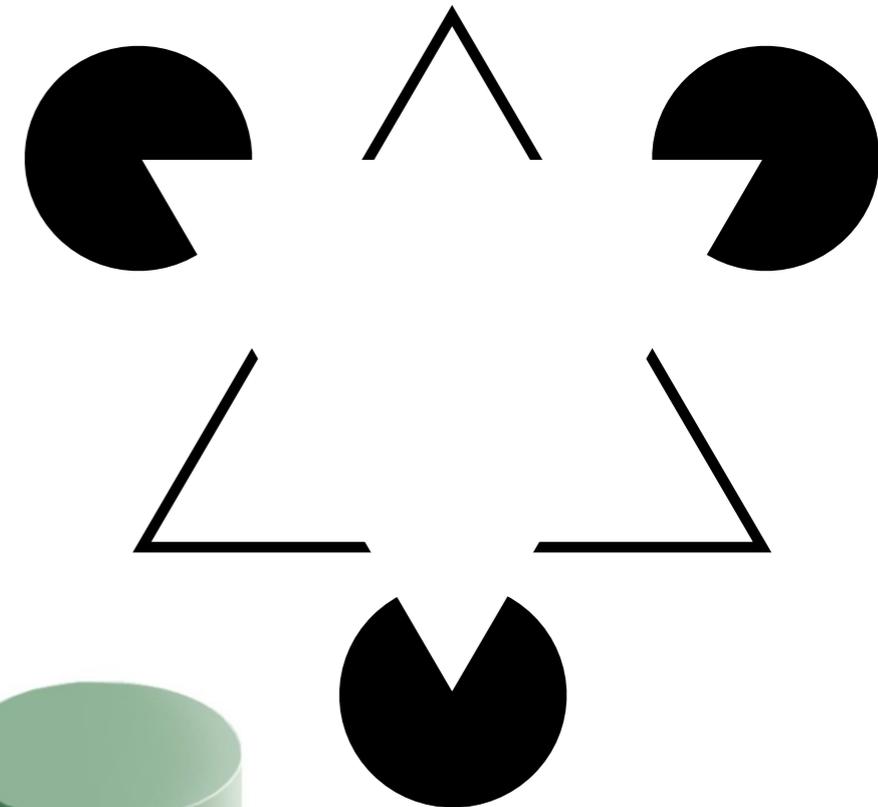
which words?

- ✓ seat
- ✗ radio
- table

bubble or dimple?



how many triangles?



is A or B lighter?

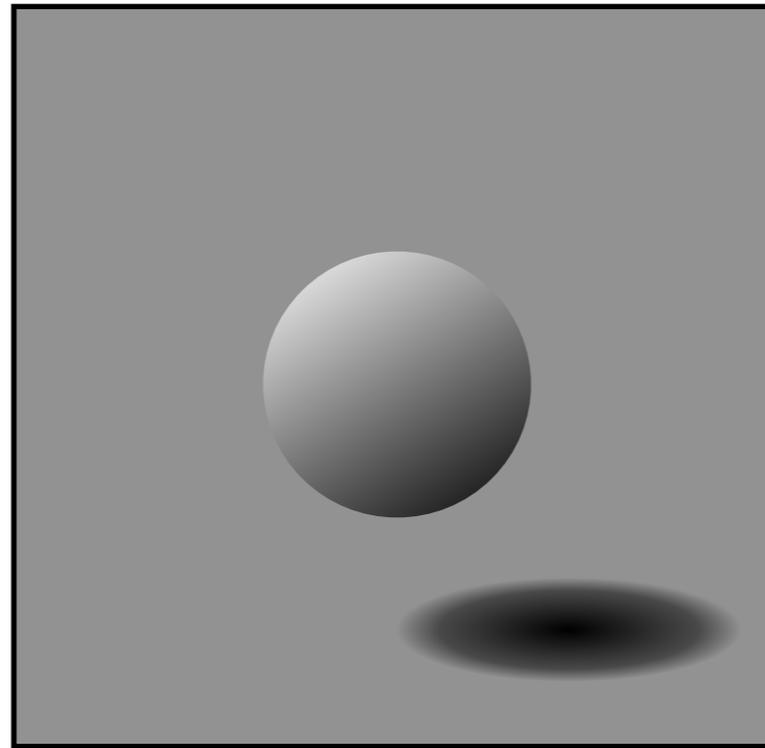
Adelson, unpubl

UNCONSCIOUS INFERENCES

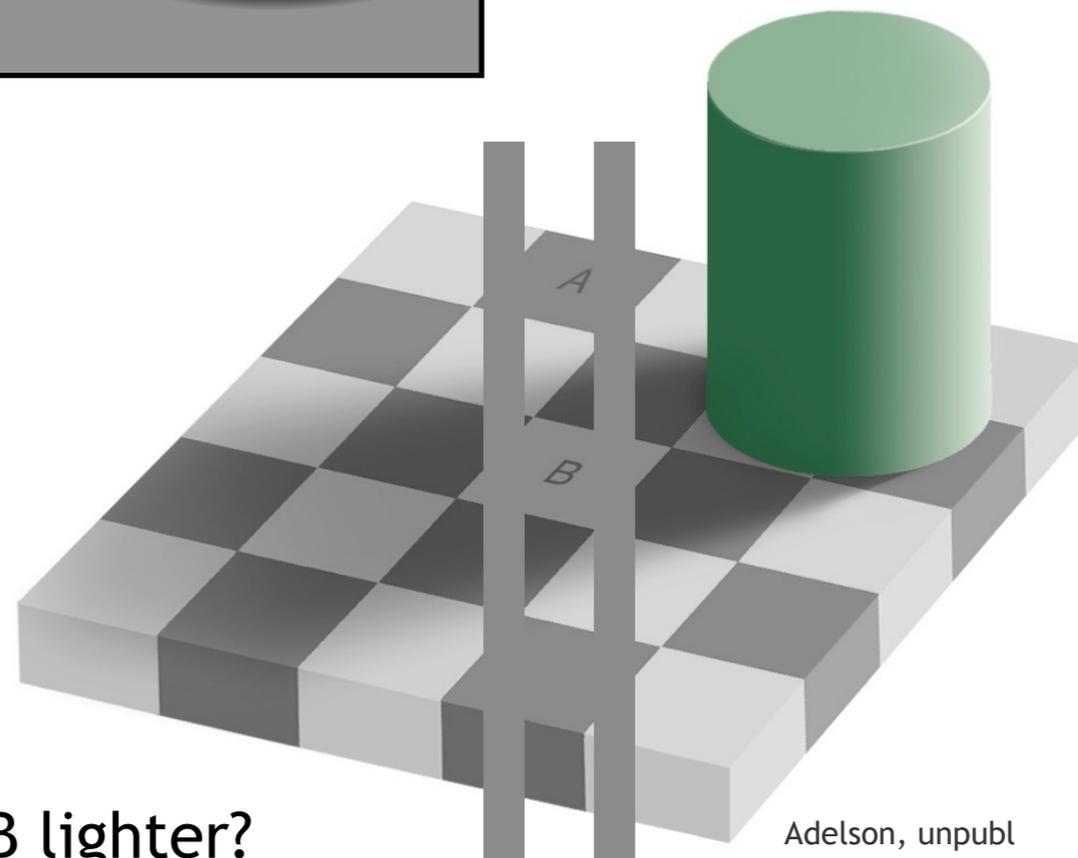
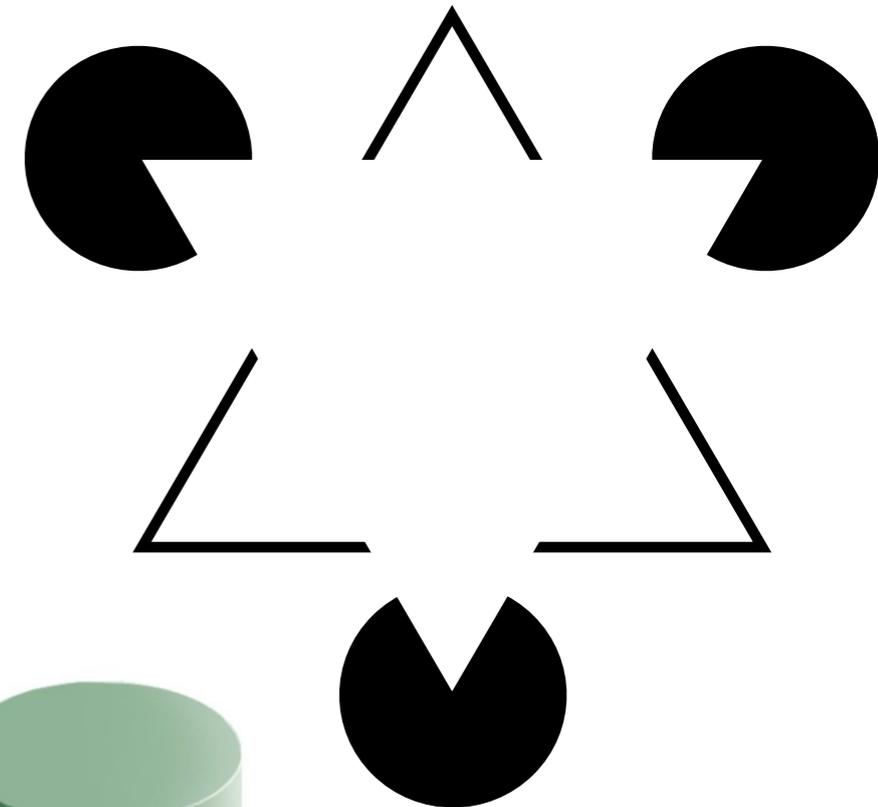
which words?

- ✓ seat
- ✗ radio
- ✓ table

bubble or dimple?



how many triangles?



is A or B lighter?

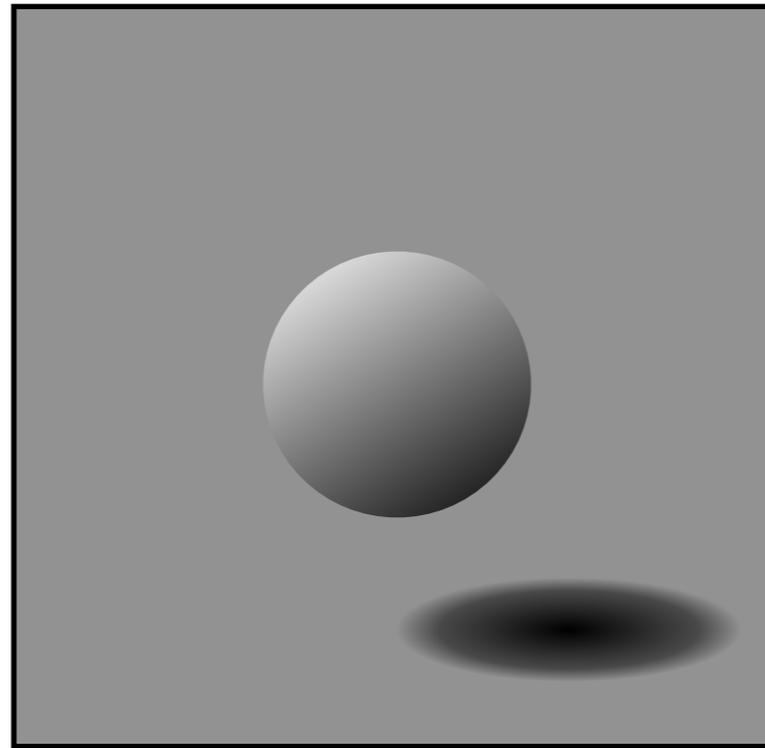
Adelson, unpubl

UNCONSCIOUS INFERENCES

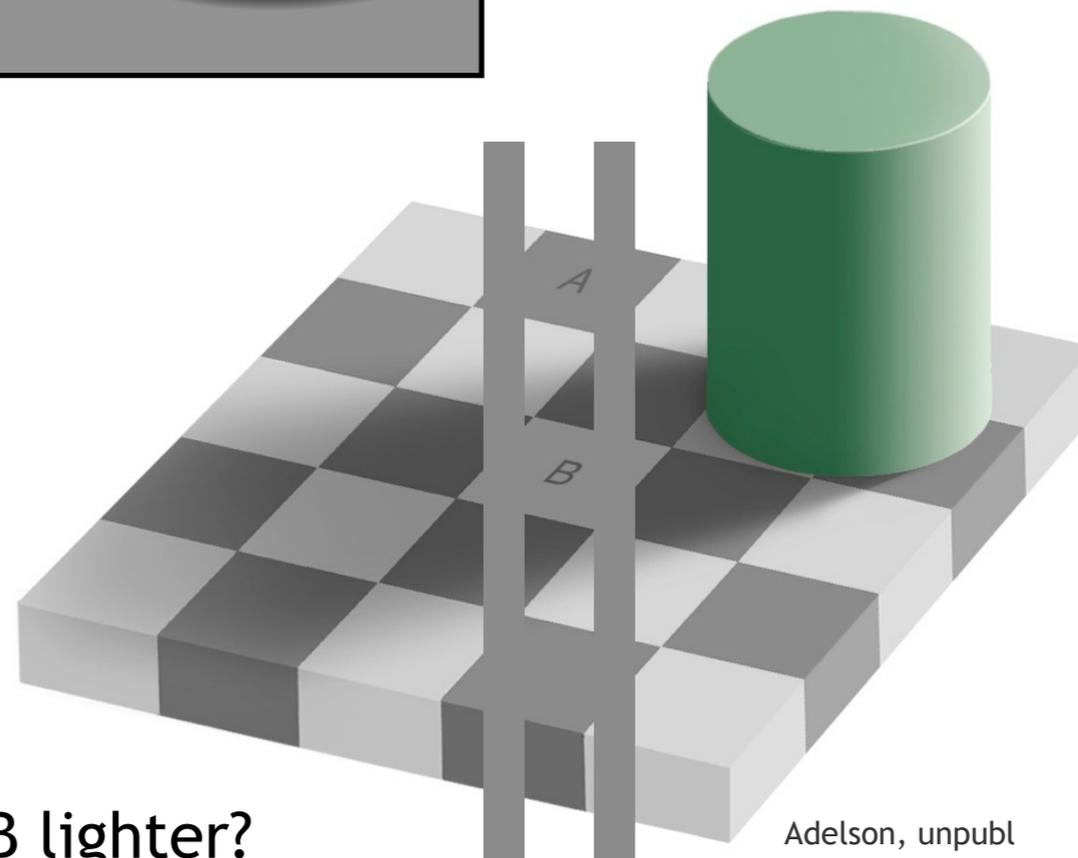
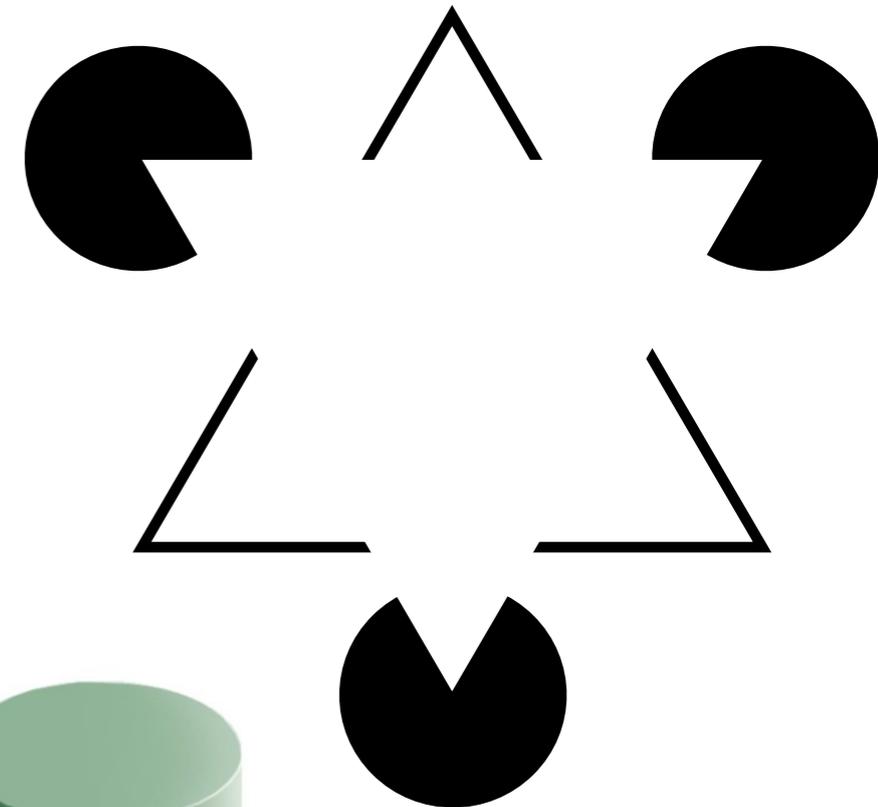
which words?

- ✓ seat
- ✗ radio
- ✓ table
- rocking

bubble or dimple?



how many triangles?



is A or B lighter?

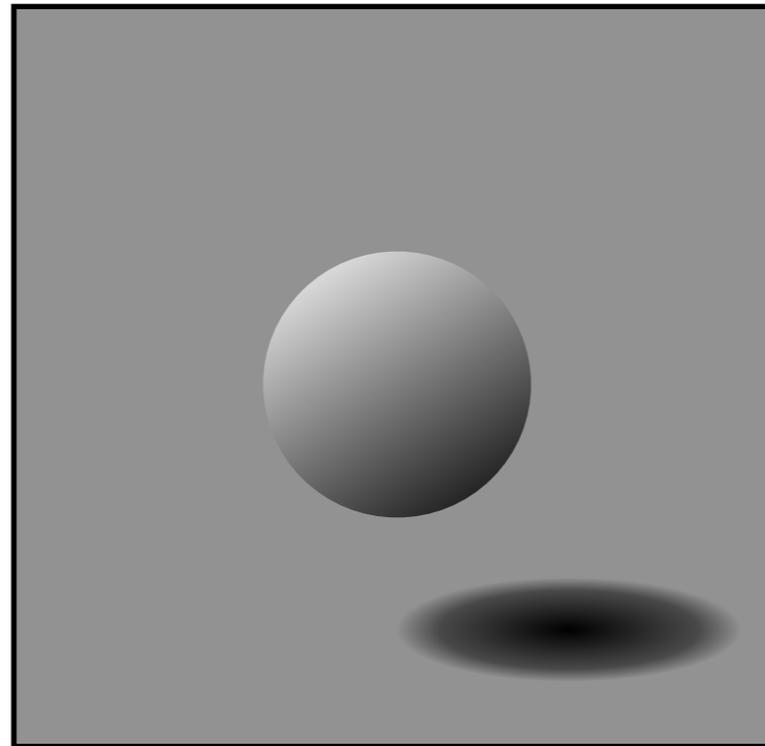
Adelson, unpubl

UNCONSCIOUS INFERENCES

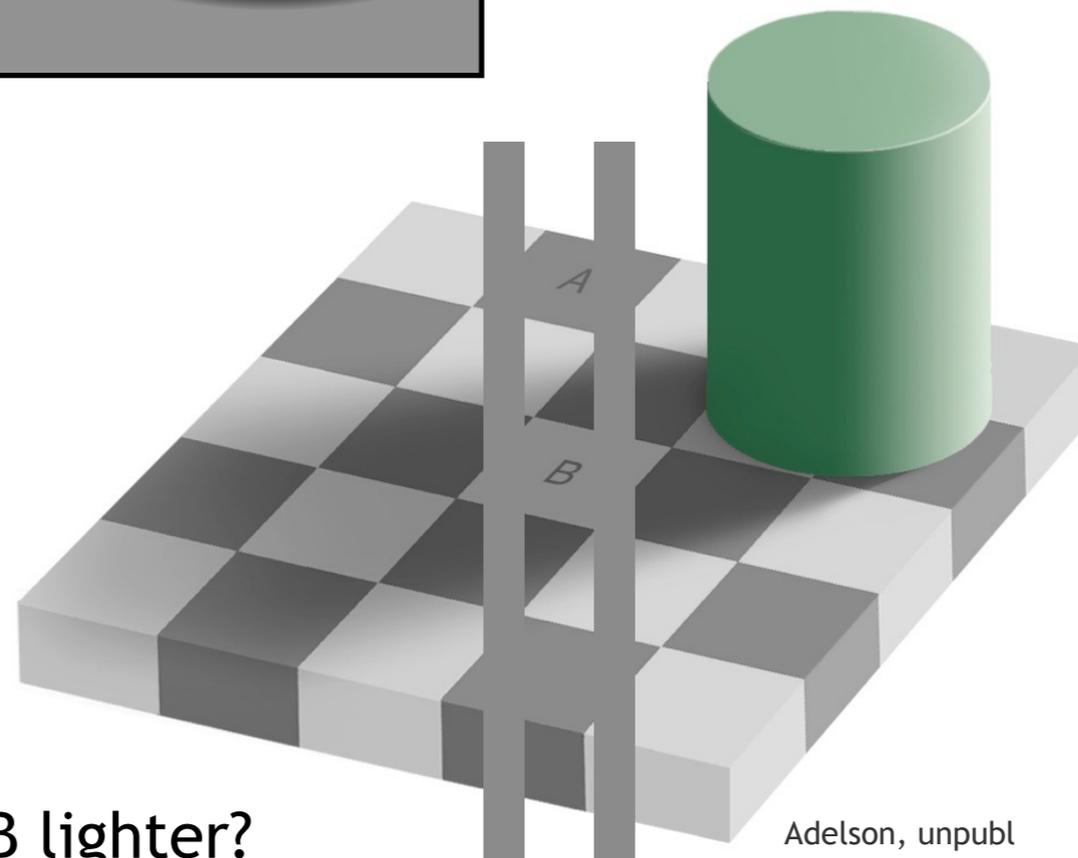
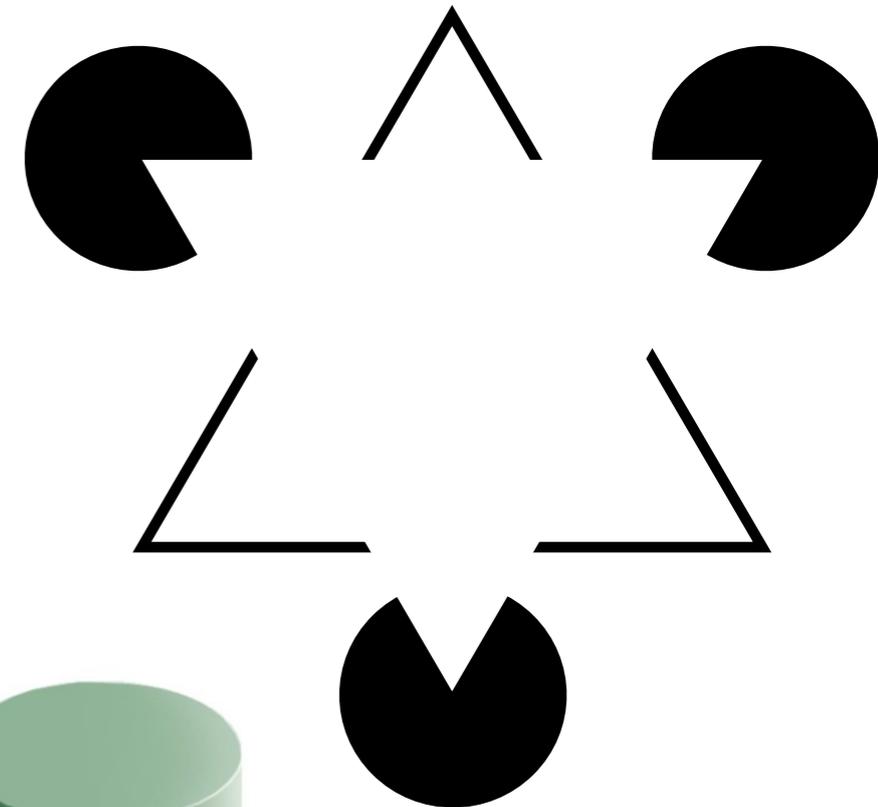
which words?

- ✓ seat
- ✗ radio
- ✓ table
- ✗ rocking

bubble or dimple?



how many triangles?



is A or B lighter?

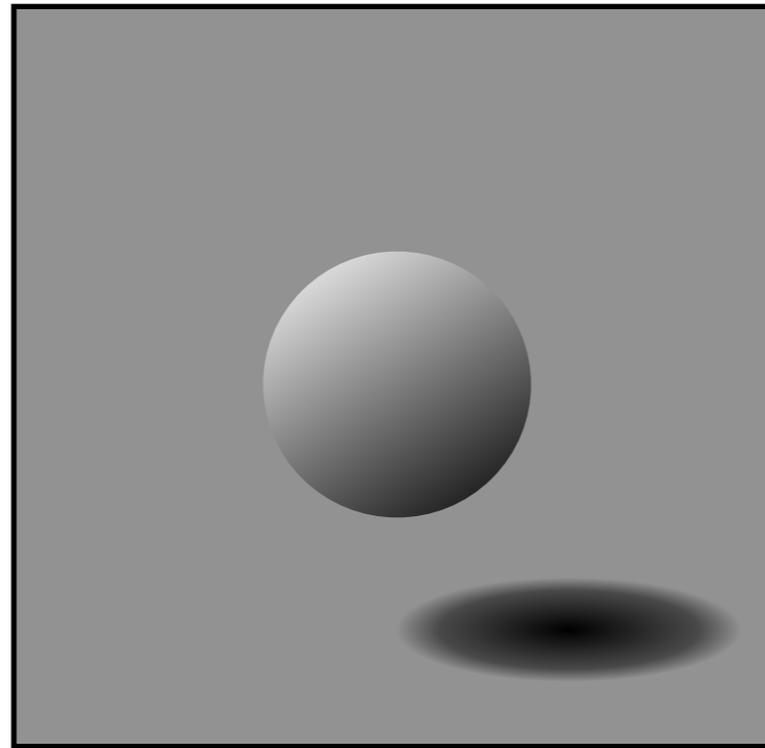
Adelson, unpubl

UNCONSCIOUS INFERENCES

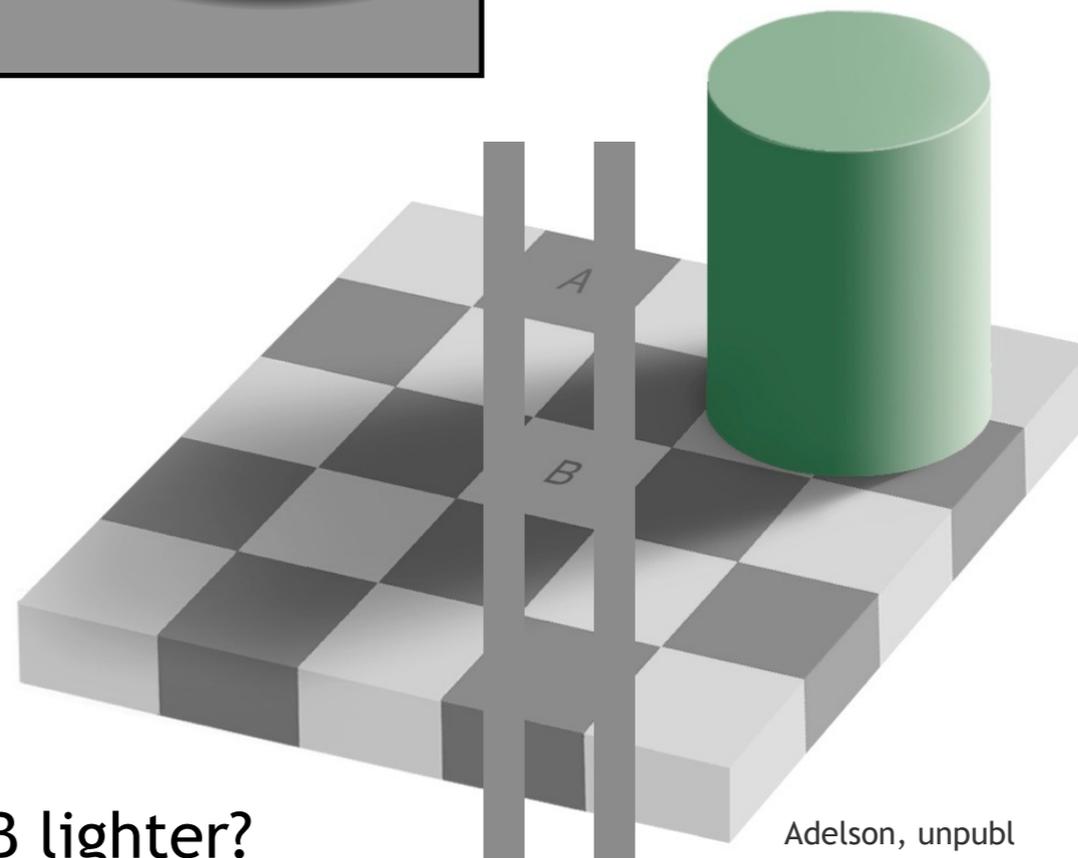
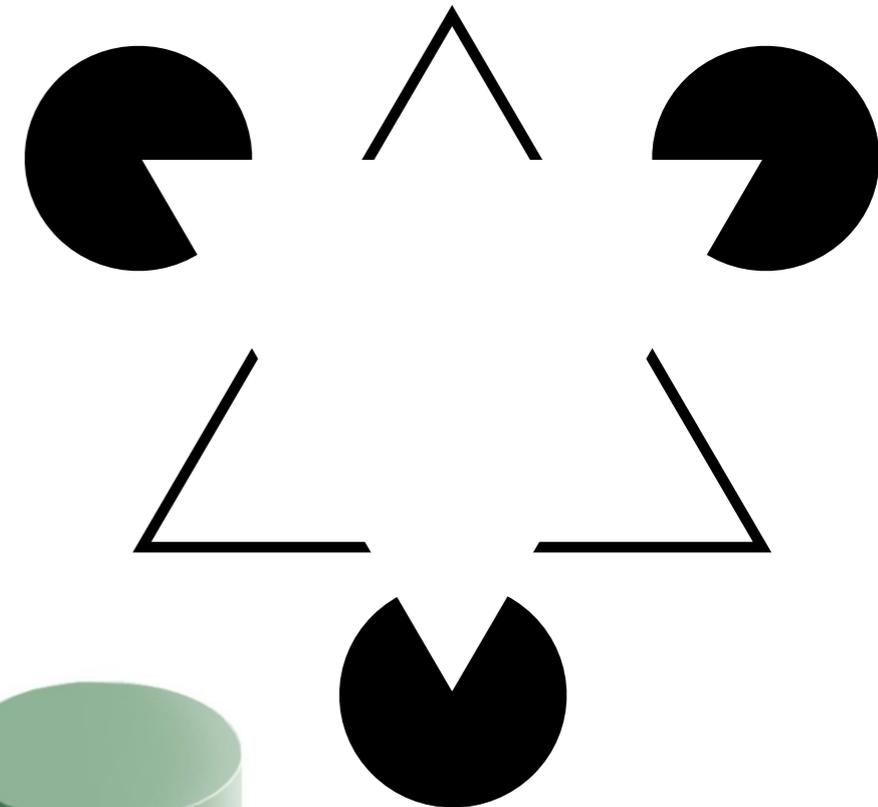
which words?

- ✓ seat
- ✗ radio
- ✓ table
- ✗ rocking bench

bubble or dimple?



how many triangles?



is A or B lighter?

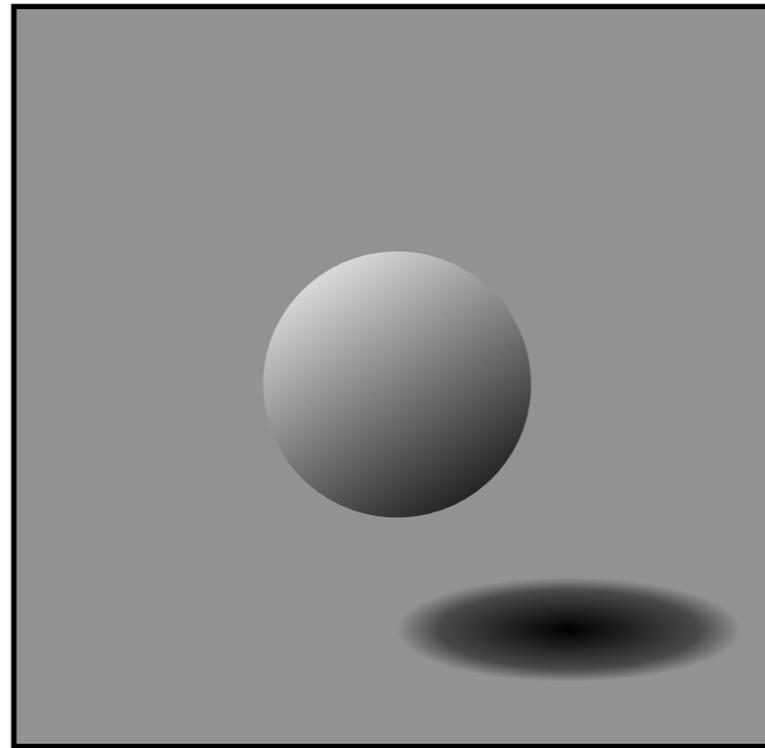
Adelson, unpubl

UNCONSCIOUS INFERENCES

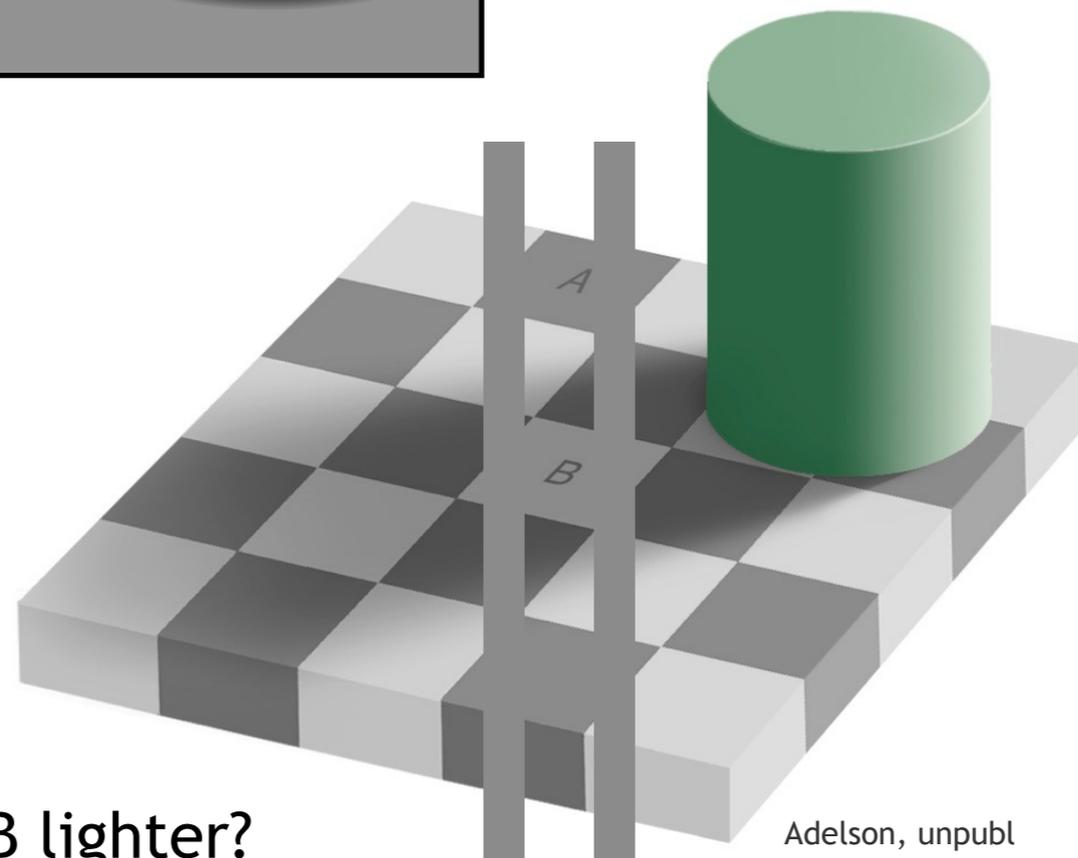
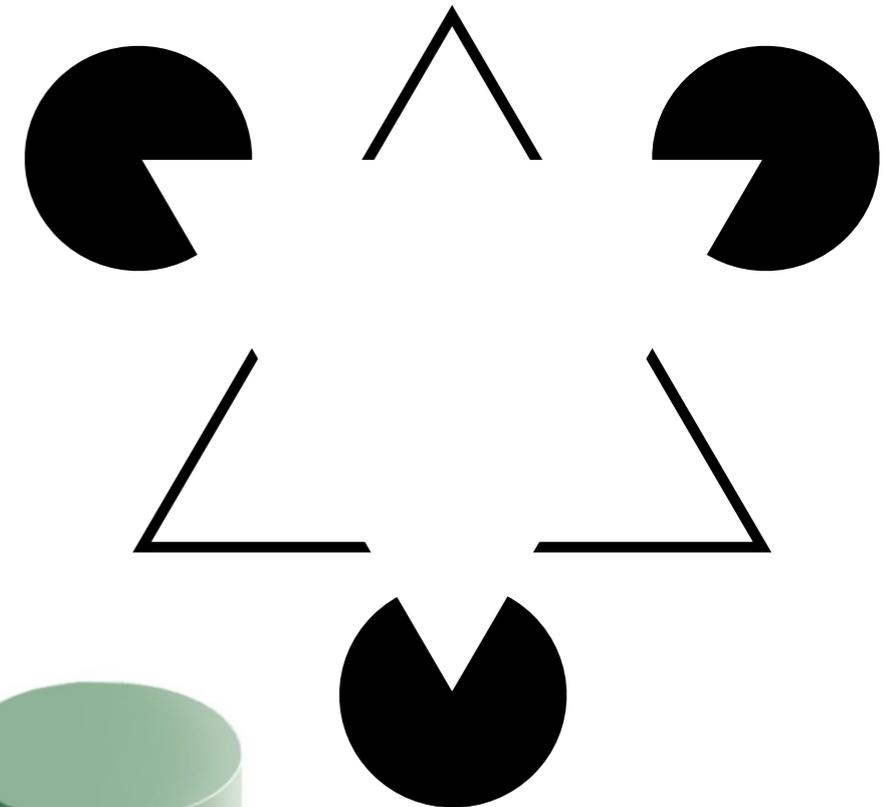
which words?

- ✓ seat
- ✗ radio
- ✓ table
- ✗ rocking
- ✗ bench

bubble or dimple?



how many triangles?



is A or B lighter?

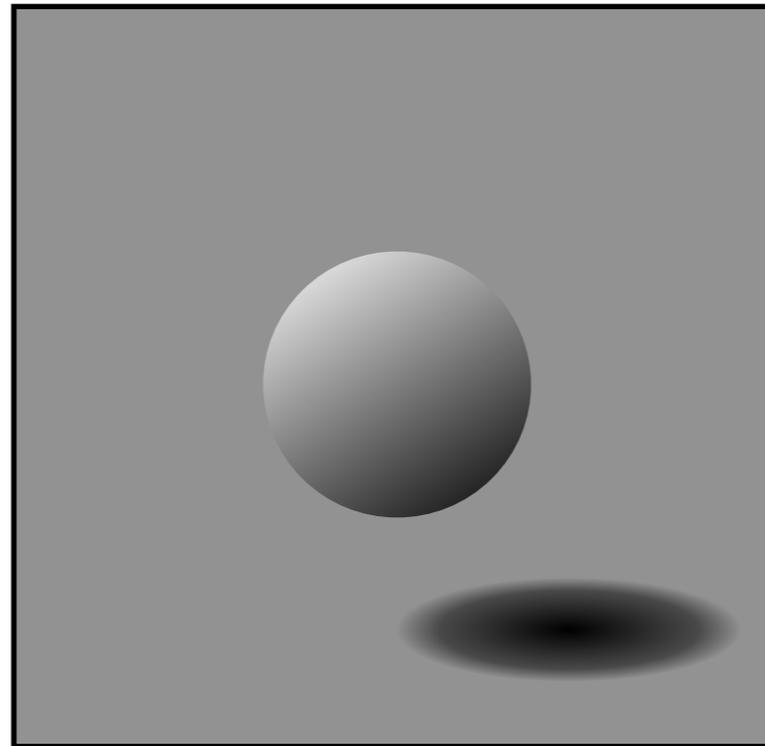
Adelson, unpubl

UNCONSCIOUS INFERENCES

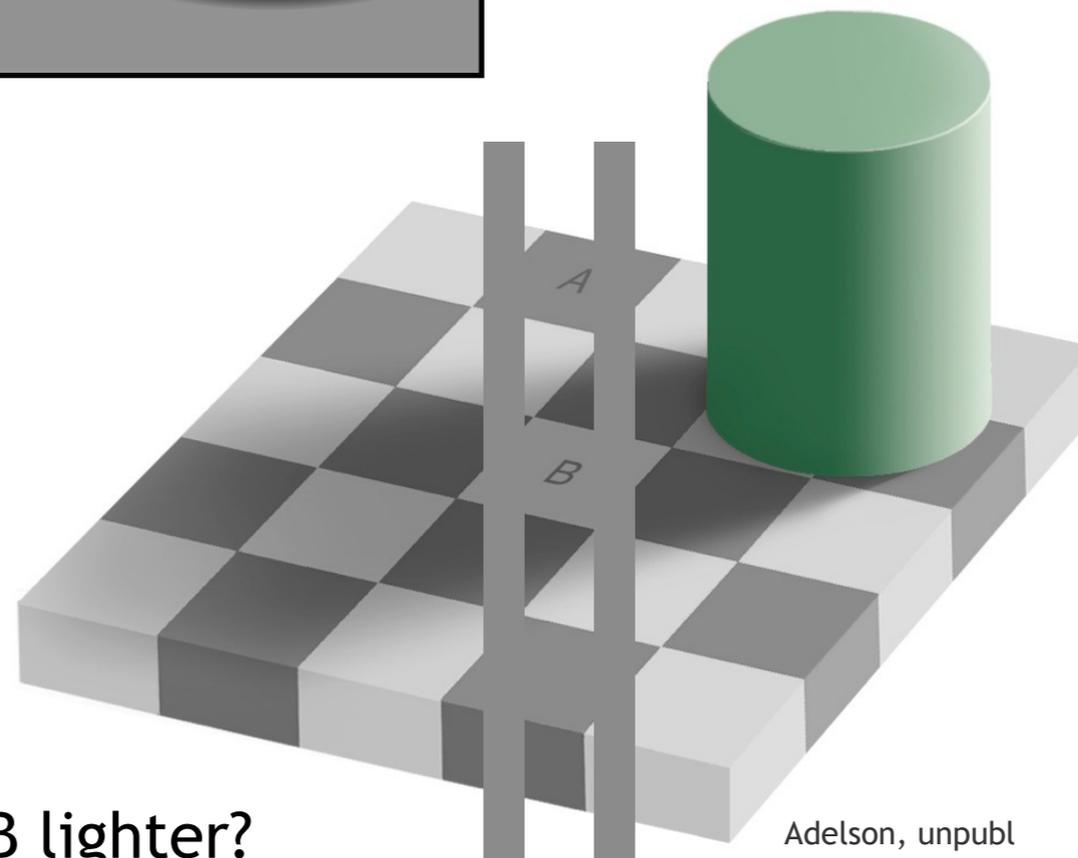
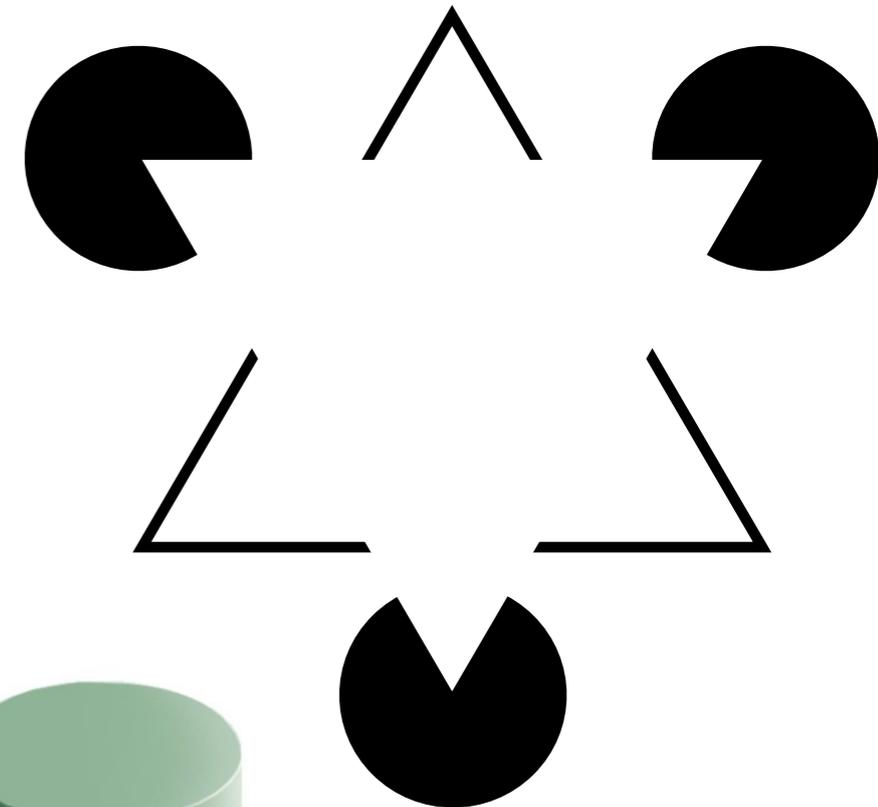
which words?

- ✓ seat
- ✗ radio
- ✓ table
- ✗ rocking
- ✗ bench
- boat

bubble or dimple?



how many triangles?



is A or B lighter?

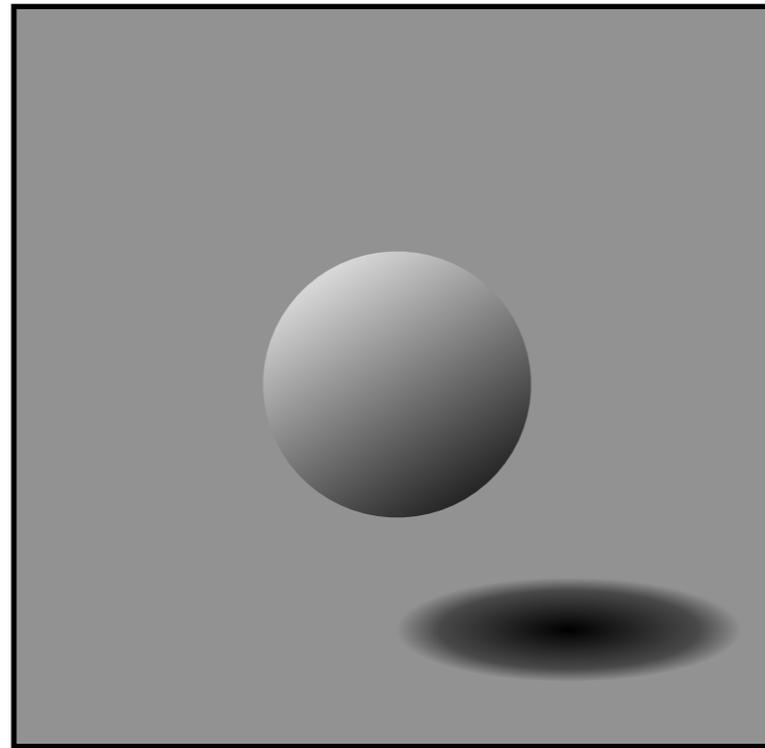
Adelson, unpubl

UNCONSCIOUS INFERENCES

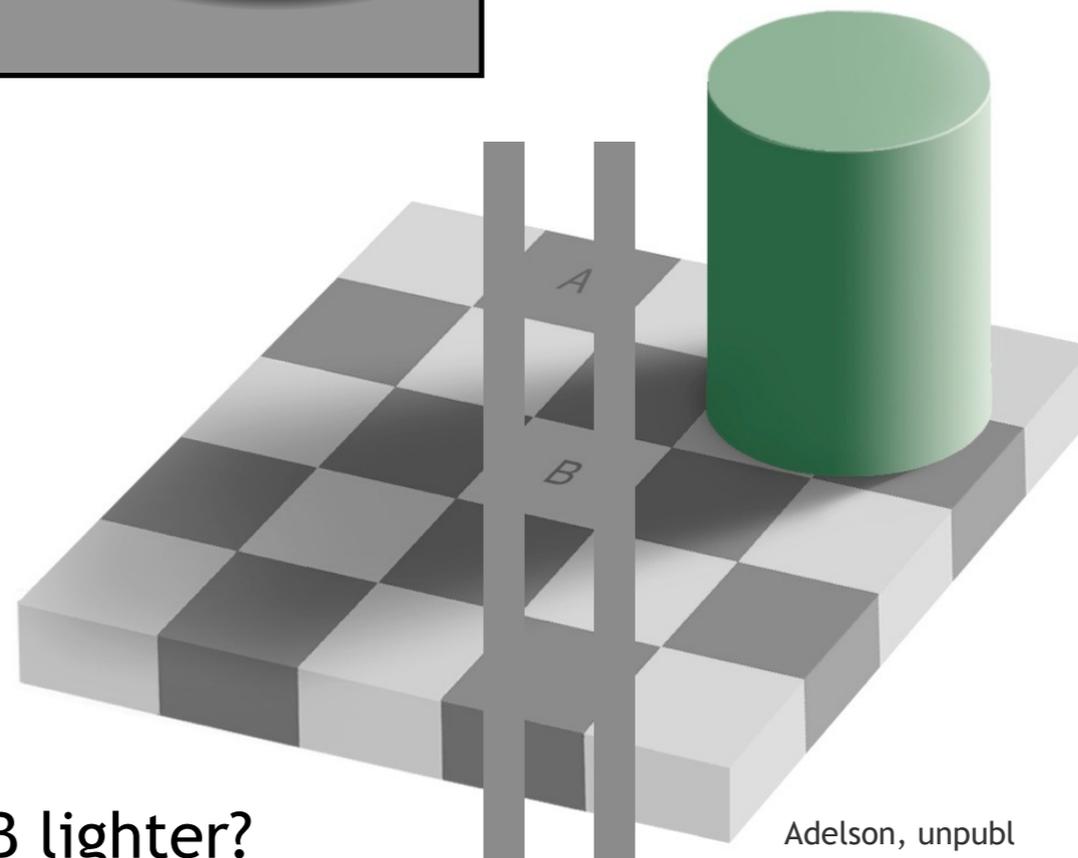
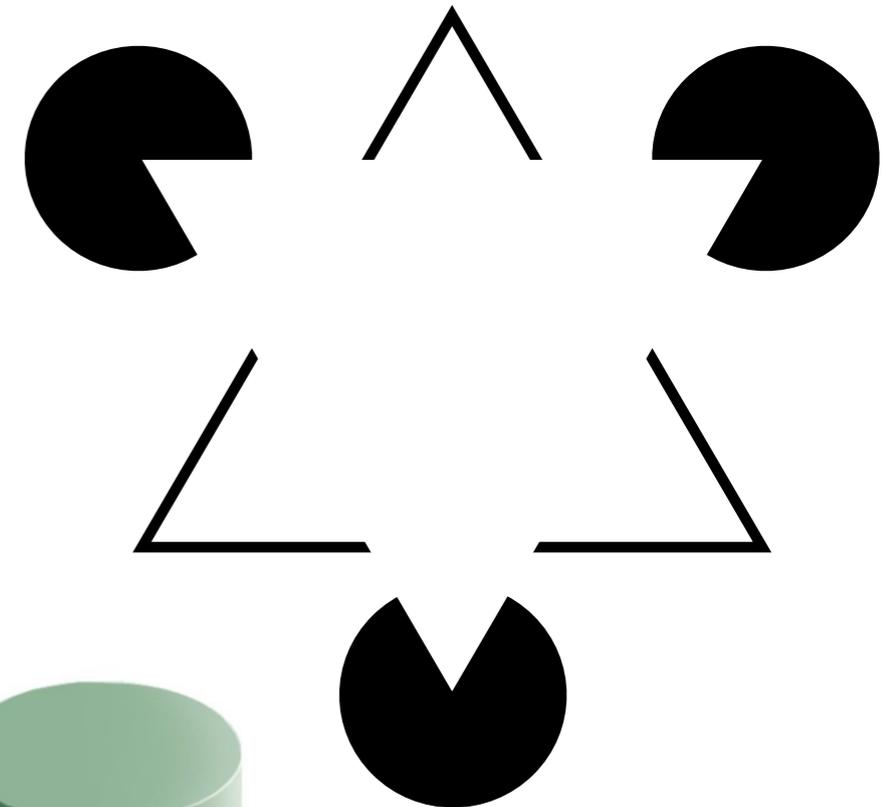
which words?

- ✓ seat
- ✗ radio
- ✓ table
- ✗ rocking
- ✗ bench
- ✗ boat
- chair

bubble or dimple?



how many triangles?



is A or B lighter?

Adelson, unpubl

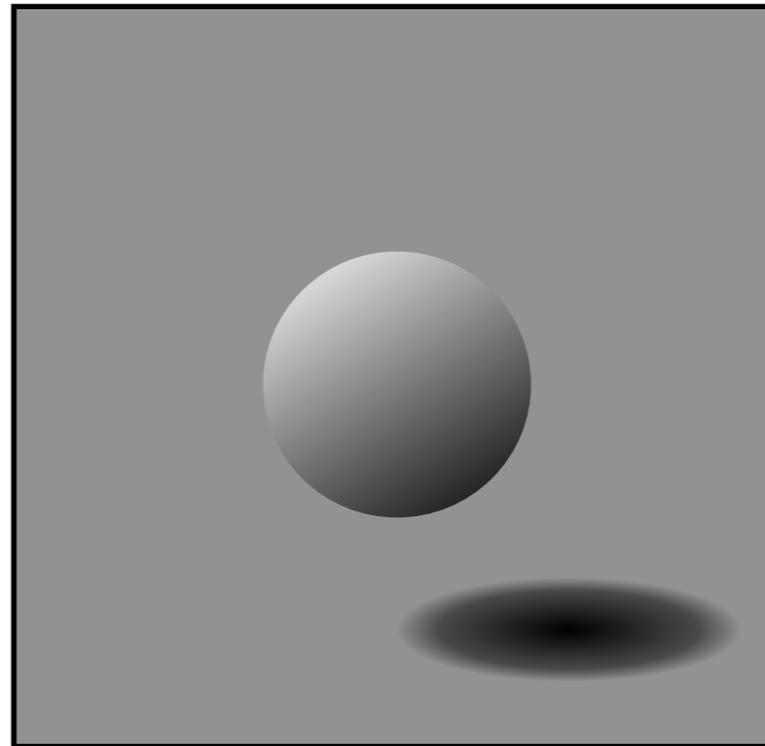
UNCONSCIOUS INFERENCES

which words?

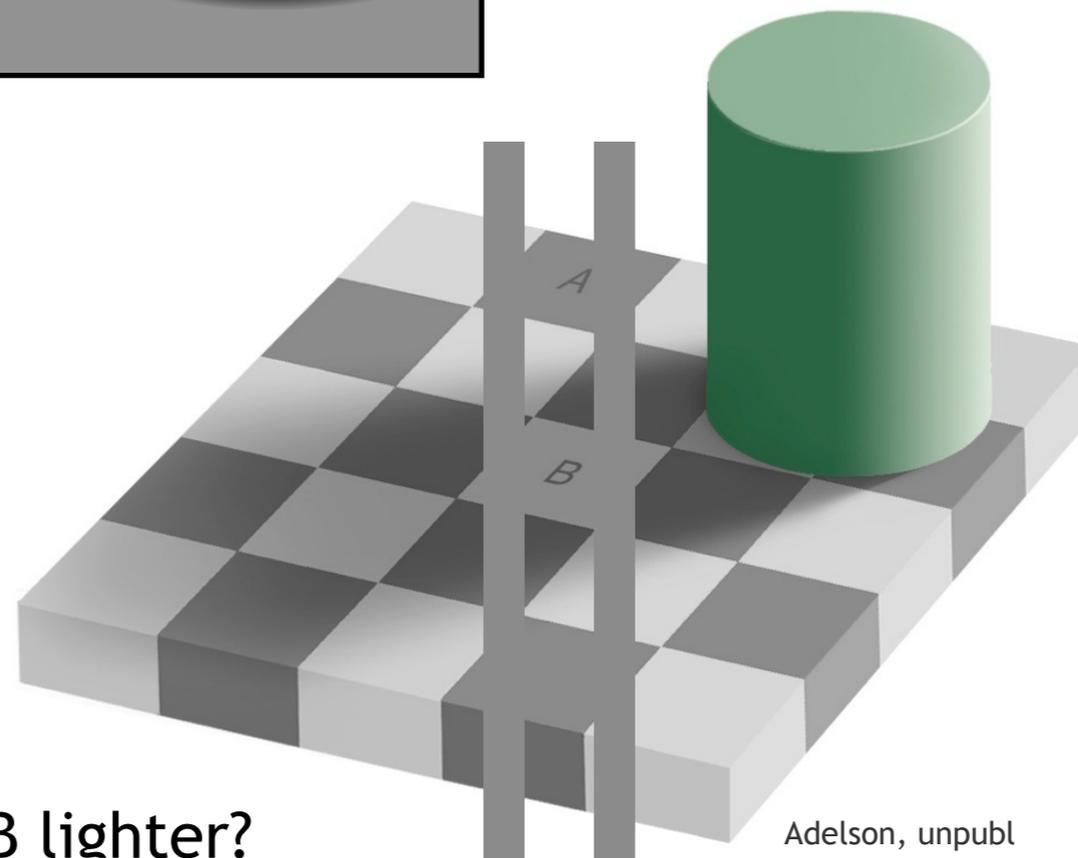
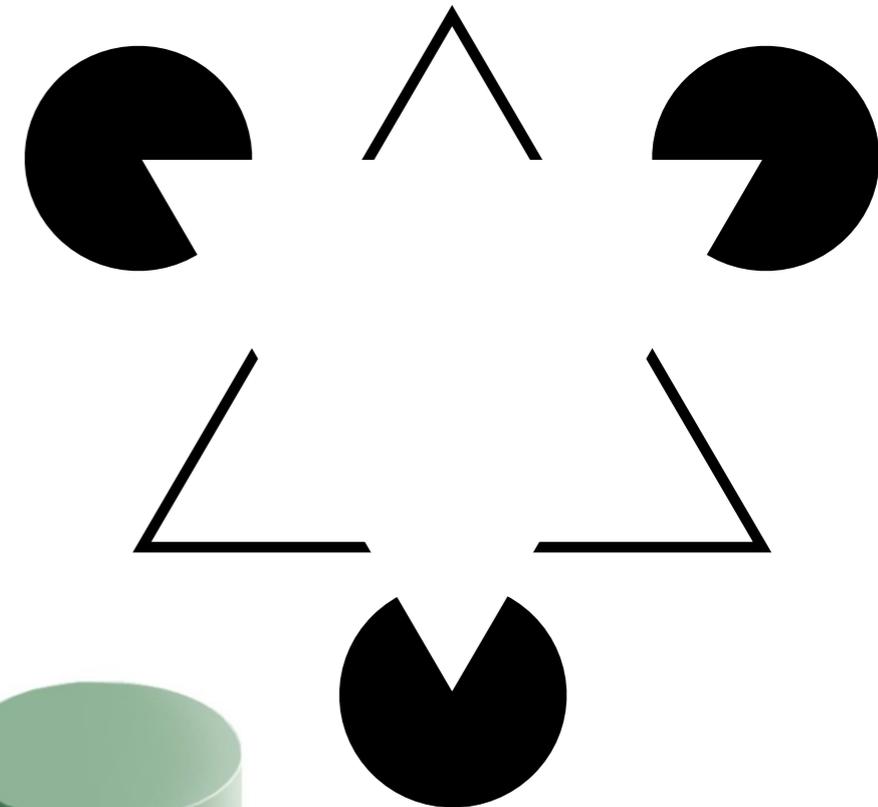
- ✓ seat
- ✗ radio
- ✓ table
- ✗ rocking
- ✗ bench
- ✗ boat
- ✗ chair

Roediger & McDermott, 1995

bubble or dimple?



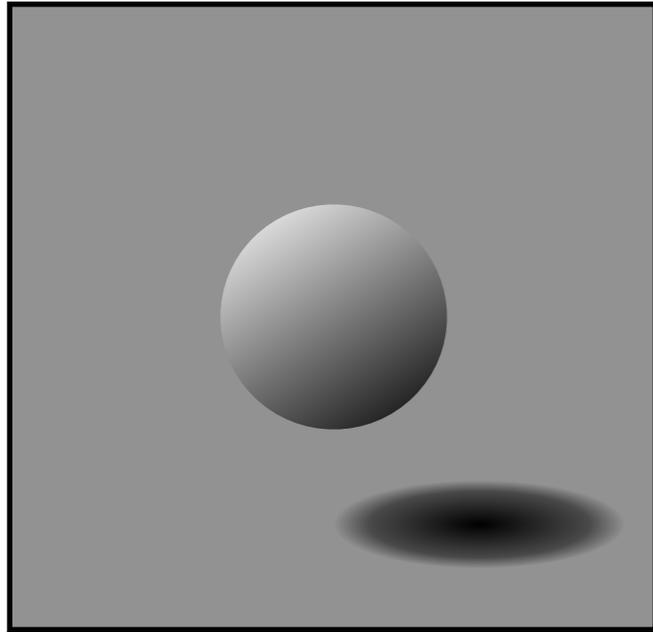
how many triangles?



is A or B lighter?

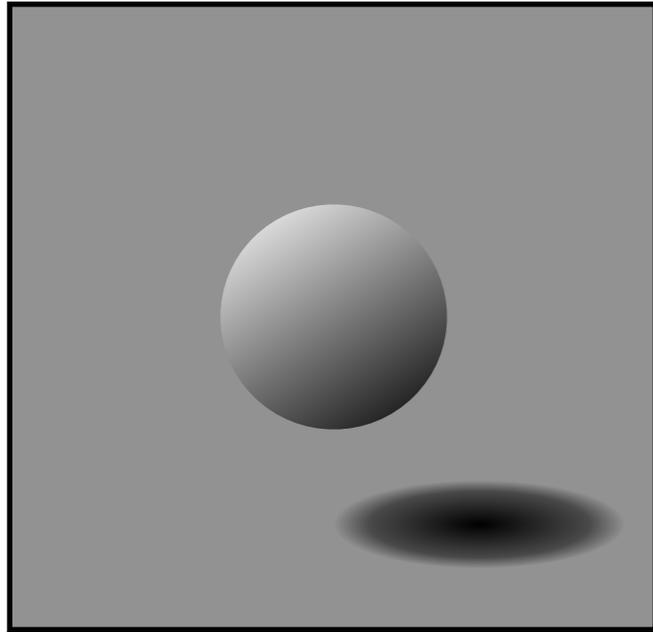
Adelson, unpubl

THE BAYESIAN BRAIN



THE BAYESIAN BRAIN

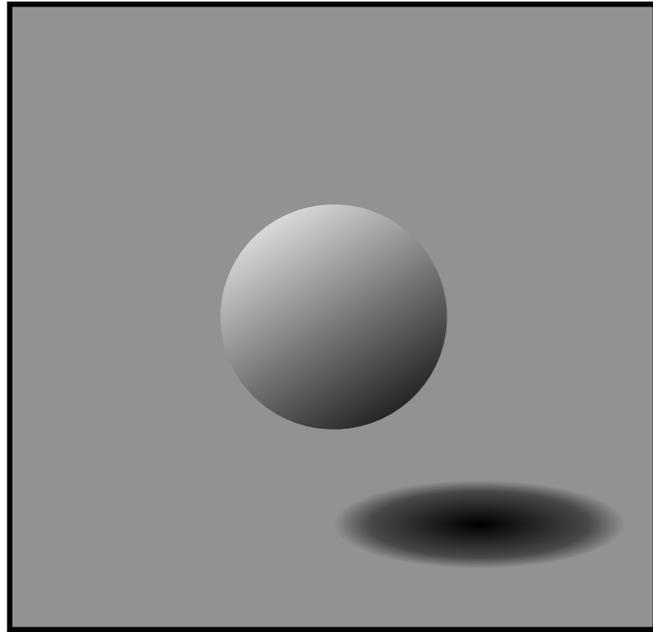
“perception is unconscious inference”



Hermann von Helmholtz
1821-1894

THE BAYESIAN BRAIN

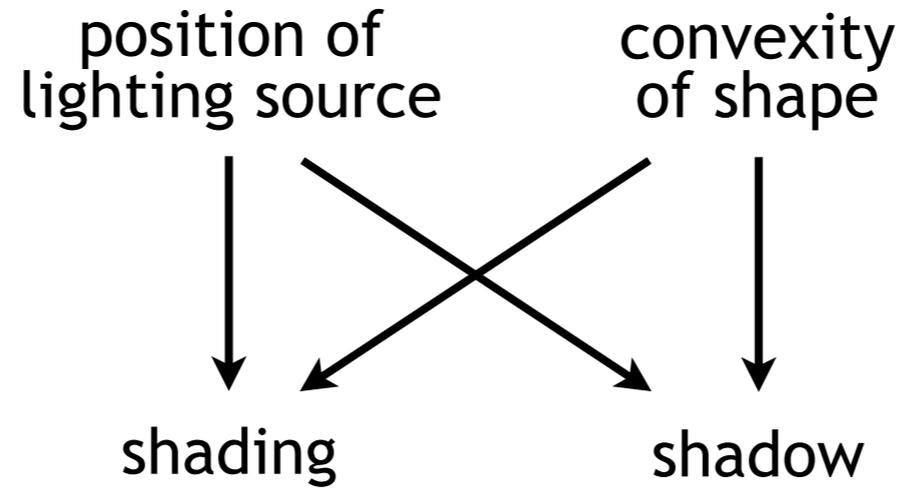
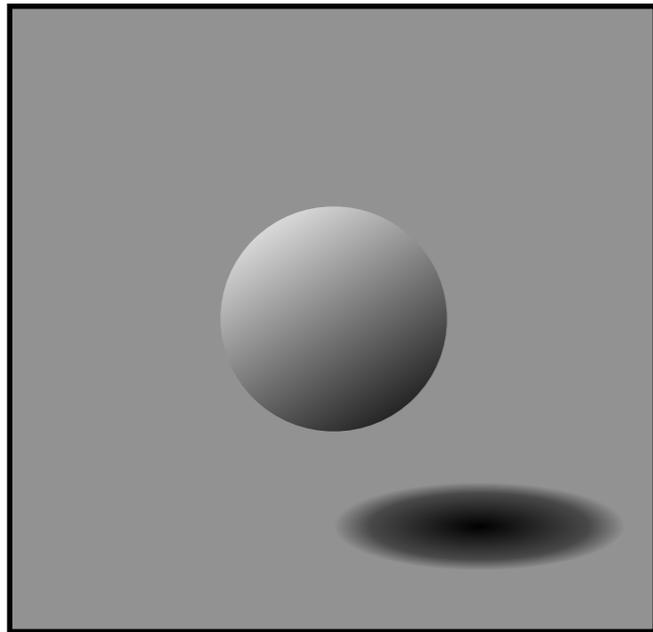
“perception is unconscious inference”
& memory & learning & ...



Hermann von Helmholtz
1821-1894

THE BAYESIAN BRAIN

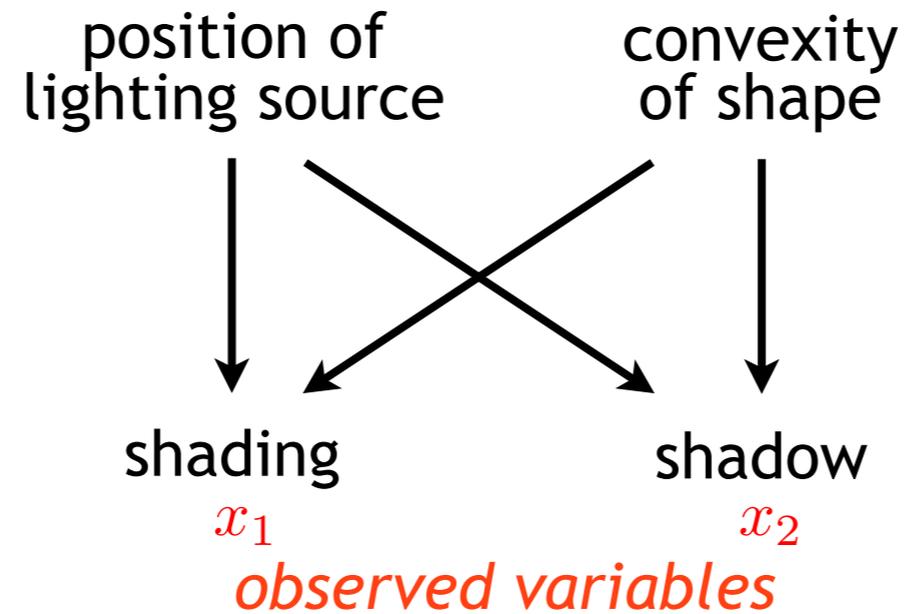
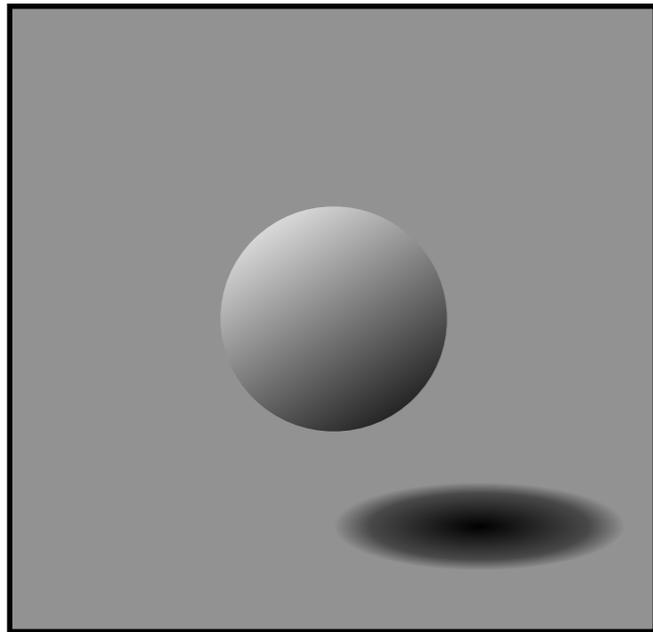
“perception is unconscious inference”
& memory & learning & ...



Hermann von Helmholtz
1821-1894

THE BAYESIAN BRAIN

“perception is unconscious inference”
& memory & learning & ...



Hermann von Helmholtz
1821-1894

THE BAYESIAN BRAIN

“perception is unconscious inference”
& memory & learning & ...

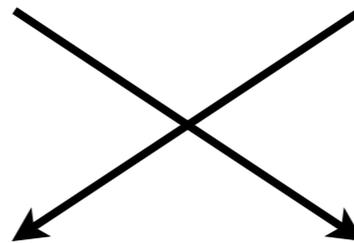
hidden variables

y_1

position of
lighting source

y_2

convexity
of shape



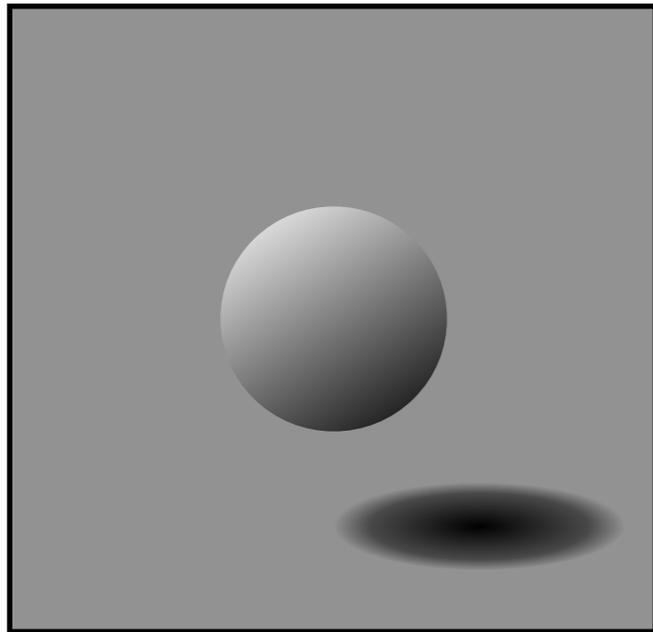
shading

shadow

x_1

x_2

observed variables



Hermann von Helmholtz
1821-1894

THE BAYESIAN BRAIN

“perception is unconscious inference”
& memory & learning & ...

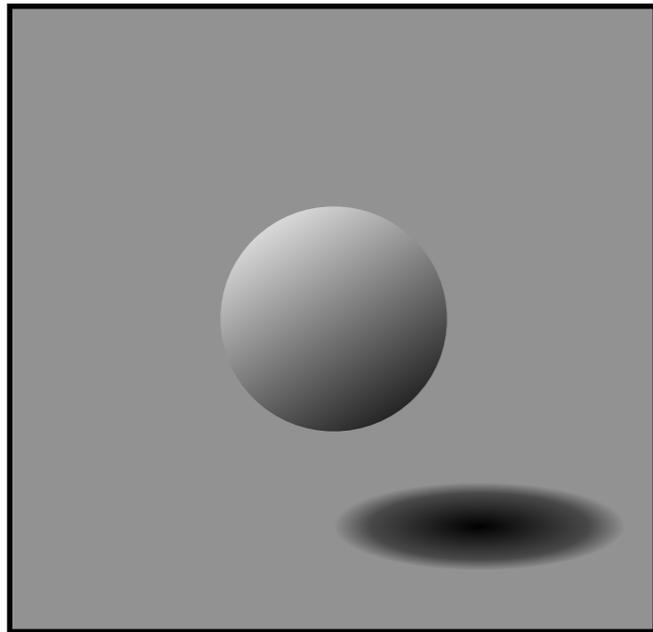
hidden variables

y_1

position of
lighting source

y_2

convexity
of shape



shading

shadow

x_1

x_2

observed variables



Hermann von Helmholtz
1821-1894

“There are **things known** and there are **things unknown**, and between are
the doors of perception”

Jim Morrison



THE BAYESIAN BRAIN

“perception is unconscious inference”
& memory & learning & ...

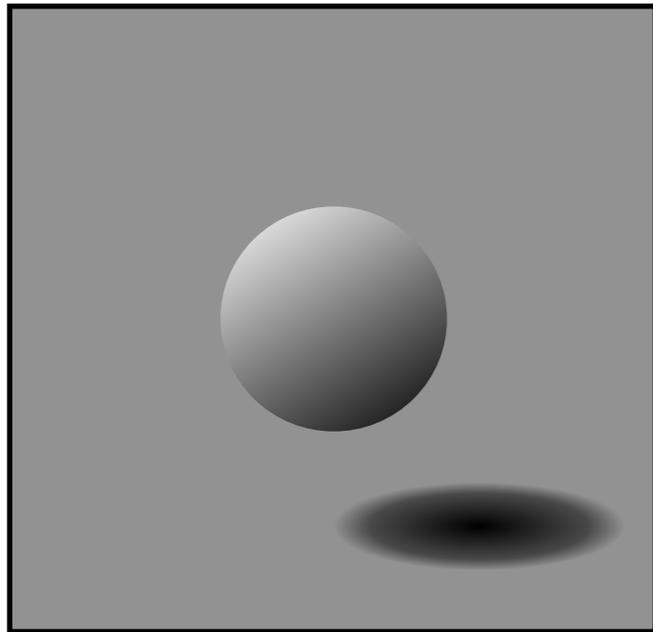
hidden variables

y_1

position of
lighting source

y_2

convexity
of shape



shading

shadow

x_1

x_2

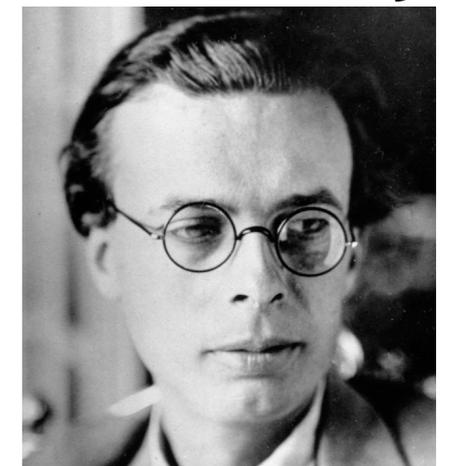
observed variables



Hermann von Helmholtz
1821-1894

“There are **things known** and there are **things unknown**, and between are
the doors of perception”

Aldous Huxley



THE BAYESIAN BRAIN

“perception is unconscious inference”
& memory & learning & ...

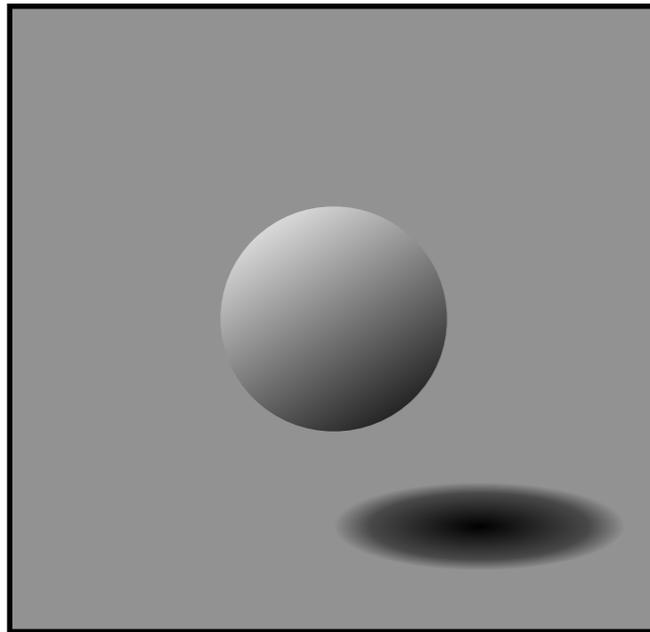
hidden variables

y_1

position of
lighting source

y_2

convexity
of shape



shading

shadow

x_1

x_2

observed variables



Hermann von Helmholtz
1821-1894

There are **things known** and there are **things unknown**, and between are
the rules of probability

THE BAYESIAN BRAIN

“perception is unconscious inference”
& memory & learning & ...

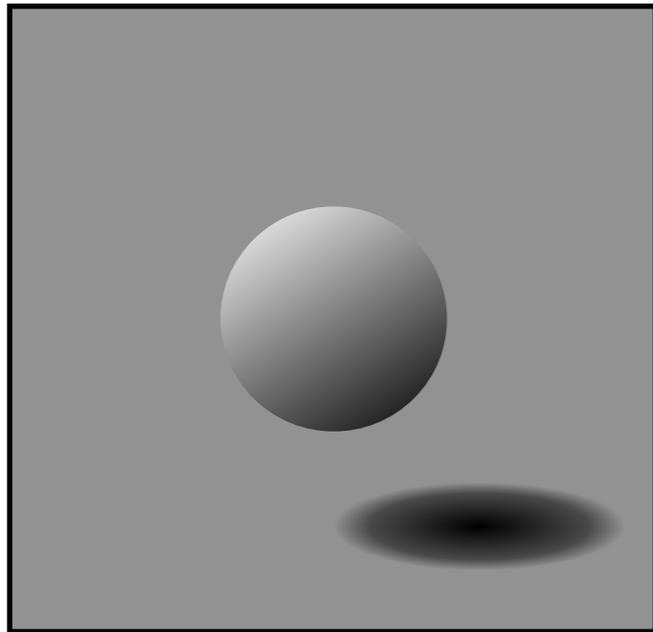
hidden variables

y_1

position of
lighting source

y_2

convexity
of shape



shading

shadow

x_1

x_2

observed variables



Hermann von Helmholtz
1821-1894

There are **things known** and there are **things unknown**, and between are
the rules of probability

product: $P(X, Y) = P(Y, X) = P(X|Y) P(Y)$

THE BAYESIAN BRAIN

“perception is unconscious inference”
& memory & learning & ...

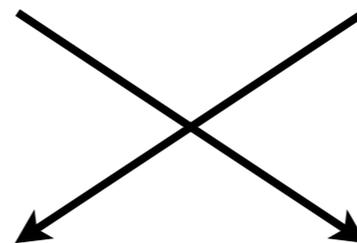
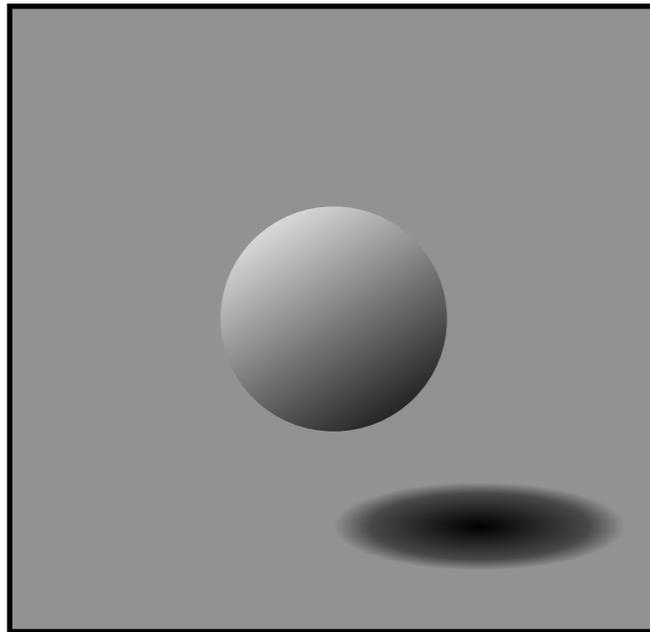
hidden variables

y_1

position of
lighting source

y_2

convexity
of shape



shading

shadow

x_1

x_2

observed variables



Hermann von Helmholtz
1821-1894

There are **things known** and there are **things unknown**, and between are
the rules of probability

product: $P(X, Y) = P(Y, X) = P(X|Y) P(Y)$

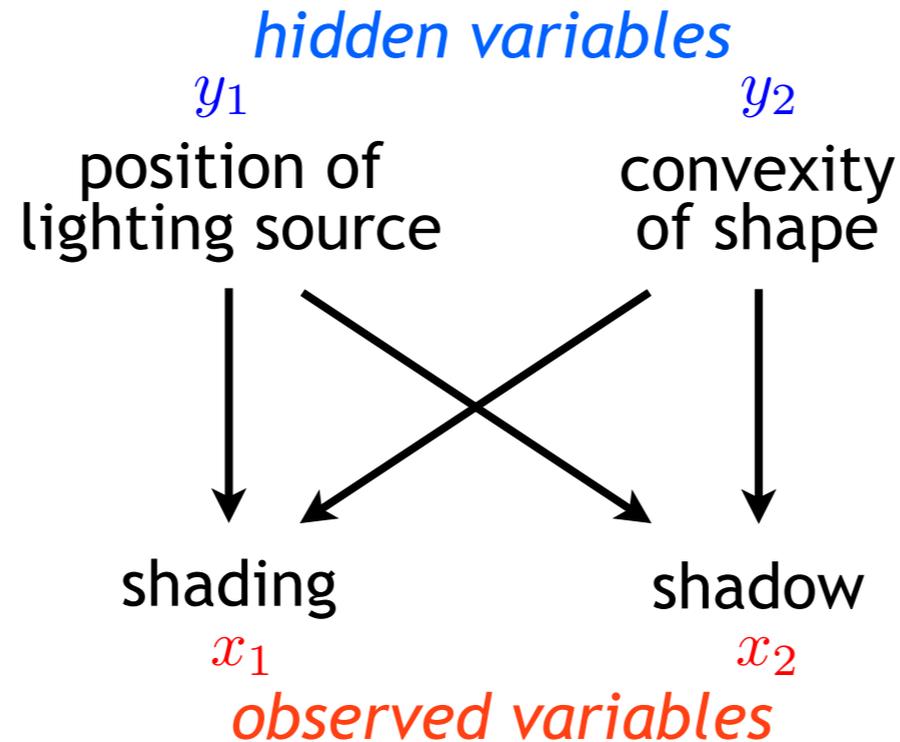
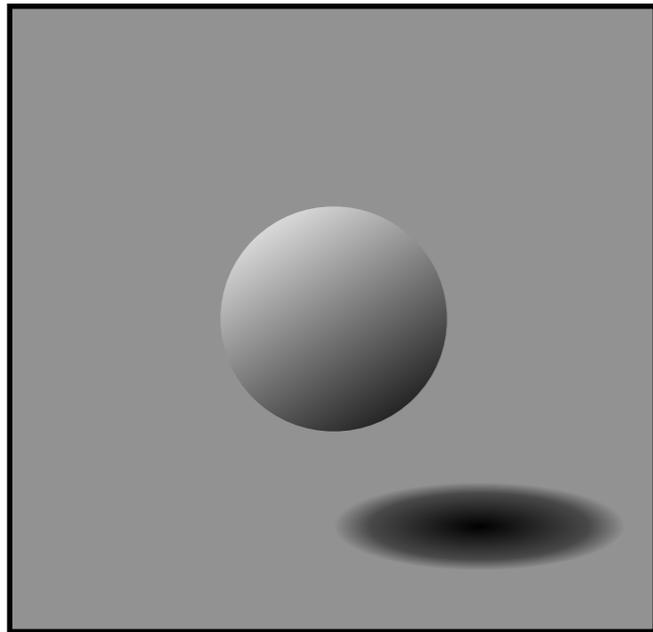
➔ Bayes' rule: $P(Y|X) = \frac{P(X|Y) P(Y)}{P(X)}$



Rev. Thomas Bayes
1702-1761

THE BAYESIAN BRAIN

“perception is unconscious inference”
& memory & learning & ...



Hermann von Helmholtz
1821-1894

There are **things known** and there are **things unknown**, and between are
the rules of probability

product: $P(X, Y) = P(Y, X) = P(X|Y) P(Y)$

➔ Bayes' rule: $P(Y|X) = \frac{P(X|Y) P(Y)}{P(X)}$

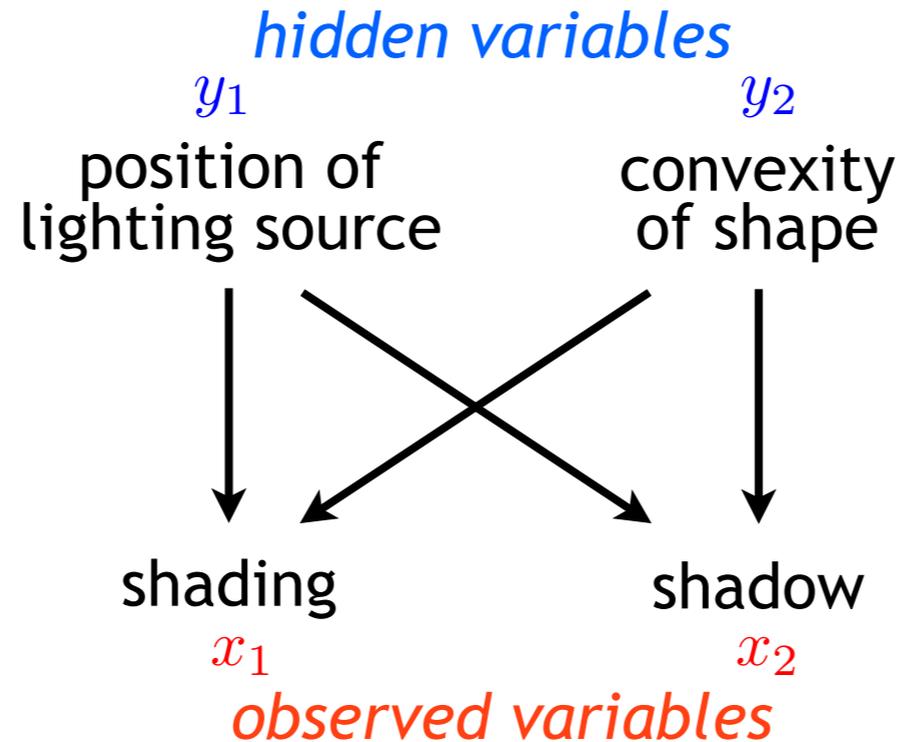
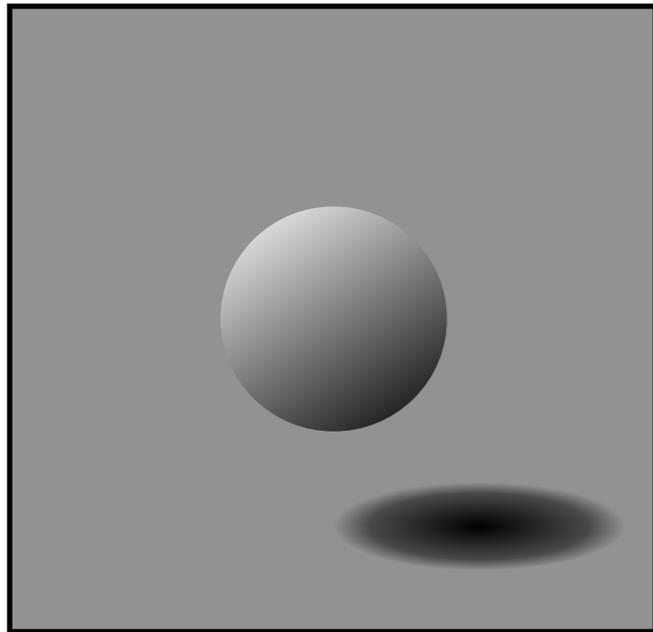
$P(X, Y) = P(X) P(Y)$ iff X and Y are independent!



Rev. Thomas Bayes
1702-1761

THE BAYESIAN BRAIN

“perception is unconscious inference”
& memory & learning & ...



Hermann von Helmholtz
1821-1894

There are **things known** and there are **things unknown**, and between are
the rules of probability

product: $P(X, Y) = P(Y, X) = P(X|Y) P(Y)$

➔ Bayes' rule: $P(Y|X) = \frac{P(X|Y) P(Y)}{P(X)}$

$P(X, Y) = P(X) P(Y)$ iff X and Y are independent!

sum:
(marginalisation)

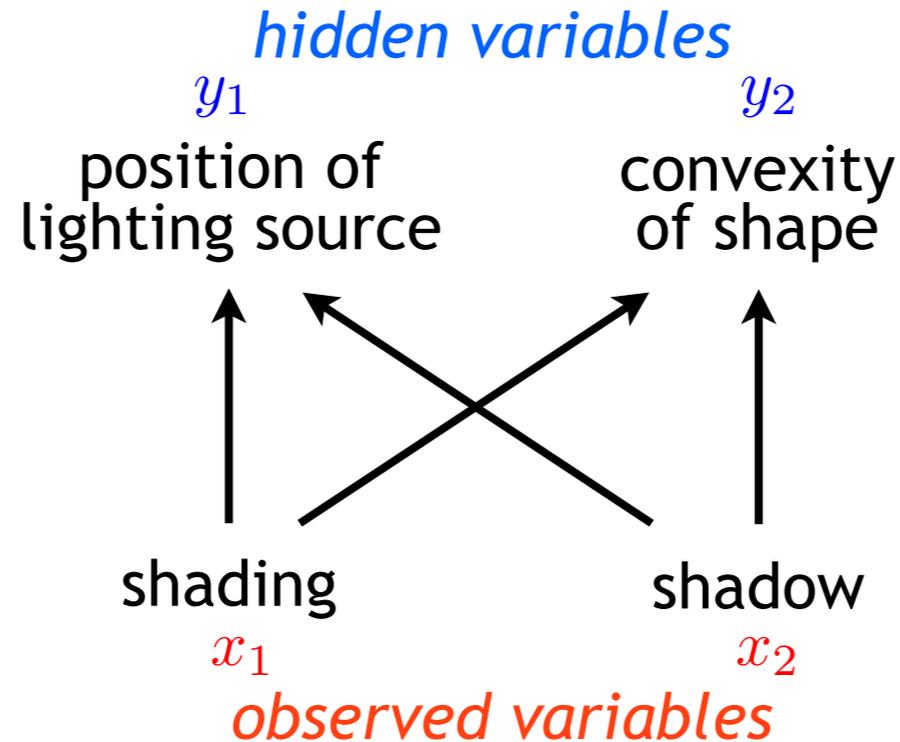
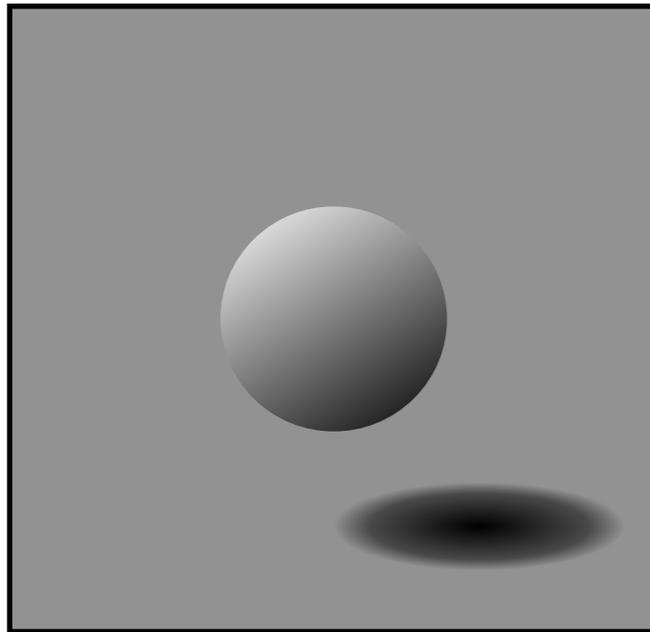
$$P(X) = \sum_Y P(X, Y)$$



Rev. Thomas Bayes
1702-1761

THE BAYESIAN BRAIN

“perception is unconscious inference”
& memory & learning & ...



Hermann von Helmholtz
1821-1894

There are **things known** and there are **things unknown**, and between are
the rules of probability

product: $P(X, Y) = P(Y, X) = P(X|Y) P(Y)$

➔ Bayes' rule: $P(Y|X) = \frac{P(X|Y) P(Y)}{P(X)}$

$P(X, Y) = P(X) P(Y)$ iff X and Y are independent!

sum:
(marginalisation)

$$P(X) = \sum_Y P(X, Y)$$



Rev. Thomas Bayes
1702-1761

BAYESIAN DECISION THEORY (and how to make point estimates)

BAYESIAN DECISION THEORY (and how to make point estimates)

state of the world

loss function

action

	y_1	y_2	y_3	...
a_1	$L(a_1, y_1)$	$L(a_1, y_2)$	$L(a_1, y_3)$	
a_2	$L(a_2, y_1)$	$L(a_2, y_2)$	$L(a_2, y_3)$	
a_3	$L(a_3, y_1)$	$L(a_3, y_2)$	$L(a_3, y_3)$	
...				

BAYESIAN DECISION THEORY (and how to make point estimates)

state of the world

loss function

action

	y_1	y_2	y_3	...
a_1	$L(a_1, y_1)$	$L(a_1, y_2)$	$L(a_1, y_3)$	
a_2	$L(a_2, y_1)$	$L(a_2, y_2)$	$L(a_2, y_3)$	
a_3	$L(a_3, y_1)$	$L(a_3, y_2)$	$L(a_3, y_3)$	
...				

action to choose:

$$a^*(x) = \operatorname{argmin}_a \sum_y L(a, y) P(y | x)$$

note: a and y need not live in the same space

BAYESIAN DECISION THEORY (and how to make point estimates)

state of the world

	y_1	y_2	y_3	...
a_1	$L(a_1, y_1)$	$L(a_1, y_2)$	$L(a_1, y_3)$	
a_2	$L(a_2, y_1)$	$L(a_2, y_2)$	$L(a_2, y_3)$	
a_3	$L(a_3, y_1)$	$L(a_3, y_2)$	$L(a_3, y_3)$	
⋮				

action

loss function

action to choose:

$$a^*(x) = \operatorname{argmin}_a \sum_y L(a, y) P(y | x)$$

note: a and y need not live in the same space

special cases when $a = \hat{y}$ and y do live in the same space

BAYESIAN DECISION THEORY (and how to make point estimates)

state of the world

	y_1	y_2	y_3	...
a_1	$L(a_1, y_1)$	$L(a_1, y_2)$	$L(a_1, y_3)$	
a_2	$L(a_2, y_1)$	$L(a_2, y_2)$	$L(a_2, y_3)$	
a_3	$L(a_3, y_1)$	$L(a_3, y_2)$	$L(a_3, y_3)$	
⋮				

action

loss function

action to choose:

$$a^*(x) = \operatorname{argmin}_a \sum_y L(a, y) P(y | x)$$

note: a and y need not live in the same space

special cases when $a = \hat{y}$ and y do live in the same space

$$L(\hat{y}, y) = (\hat{y} - y)^2$$

BAYESIAN DECISION THEORY (and how to make point estimates)

state of the world

	y_1	y_2	y_3	...
a_1	$L(a_1, y_1)$	$L(a_1, y_2)$	$L(a_1, y_3)$	
a_2	$L(a_2, y_1)$	$L(a_2, y_2)$	$L(a_2, y_3)$	
a_3	$L(a_3, y_1)$	$L(a_3, y_2)$	$L(a_3, y_3)$	
⋮				

action

loss function

action to choose:

$$a^*(x) = \operatorname{argmin}_a \sum_y L(a, y) P(y | x)$$

note: a and y need not live in the same space

special cases when $a = \hat{y}$ and y do live in the same space

$$L(\hat{y}, y) = (\hat{y} - y)^2 \quad \longrightarrow \quad \hat{y}(x) = \sum_y y P(y | x) \quad \text{posterior mean}$$

BAYESIAN DECISION THEORY (and how to make point estimates)

state of the world

	y_1	y_2	y_3	...
a_1	$L(a_1, y_1)$	$L(a_1, y_2)$	$L(a_1, y_3)$	
a_2	$L(a_2, y_1)$	$L(a_2, y_2)$	$L(a_2, y_3)$	
a_3	$L(a_3, y_1)$	$L(a_3, y_2)$	$L(a_3, y_3)$	
⋮				

action

loss function

action to choose:

$$a^*(x) = \operatorname{argmin}_a \sum_y L(a, y) P(y | x)$$

note: a and y need not live in the same space

special cases when $a = \hat{y}$ and y do live in the same space

$$L(\hat{y}, y) = (\hat{y} - y)^2 \quad \longrightarrow \quad \hat{y}(x) = \sum_y y P(y | x) \quad \text{posterior mean}$$

$$L(\hat{y}, y) = |\hat{y} - y|$$

BAYESIAN DECISION THEORY (and how to make point estimates)

state of the world

	y_1	y_2	y_3	...
a_1	$L(a_1, y_1)$	$L(a_1, y_2)$	$L(a_1, y_3)$	
a_2	$L(a_2, y_1)$	$L(a_2, y_2)$	$L(a_2, y_3)$	
a_3	$L(a_3, y_1)$	$L(a_3, y_2)$	$L(a_3, y_3)$	
⋮				

action

loss function

action to choose:

$$a^*(x) = \operatorname{argmin}_a \sum_y L(a, y) P(y | x)$$

note: a and y need not live in the same space

special cases when $a = \hat{y}$ and y do live in the same space

$$L(\hat{y}, y) = (\hat{y} - y)^2 \quad \longrightarrow \quad \hat{y}(x) = \sum_y y P(y | x) \quad \text{posterior mean}$$

$$L(\hat{y}, y) = |\hat{y} - y| \quad \longrightarrow \quad \sum_{y=-\infty}^{\hat{y}(x)} P(y | x) = \frac{1}{2} \quad \text{posterior median}$$

BAYESIAN DECISION THEORY (and how to make point estimates)

state of the world

	y_1	y_2	y_3	...
a_1	$L(a_1, y_1)$	$L(a_1, y_2)$	$L(a_1, y_3)$	
a_2	$L(a_2, y_1)$	$L(a_2, y_2)$	$L(a_2, y_3)$	
a_3	$L(a_3, y_1)$	$L(a_3, y_2)$	$L(a_3, y_3)$	
⋮				

action

loss function

action to choose:

$$a^*(x) = \operatorname{argmin}_a \sum_y L(a, y) P(y | x)$$

note: a and y need not live in the same space

special cases when $a = \hat{y}$ and y do live in the same space

$$L(\hat{y}, y) = (\hat{y} - y)^2 \quad \longrightarrow \quad \hat{y}(x) = \sum_y y P(y | x) \quad \text{posterior mean}$$

$$L(\hat{y}, y) = |\hat{y} - y| \quad \longrightarrow \quad \sum_{y=-\infty}^{\hat{y}(x)} P(y | x) = \frac{1}{2} \quad \text{posterior median}$$

$$L(\hat{y}, y) = \begin{cases} 0 & \text{if } y = \hat{y} \\ 1 & \text{otherwise} \end{cases}$$

BAYESIAN DECISION THEORY (and how to make point estimates)

state of the world

	y_1	y_2	y_3	...
a_1	$L(a_1, y_1)$	$L(a_1, y_2)$	$L(a_1, y_3)$	
a_2	$L(a_2, y_1)$	$L(a_2, y_2)$	$L(a_2, y_3)$	
a_3	$L(a_3, y_1)$	$L(a_3, y_2)$	$L(a_3, y_3)$	
⋮				

action

loss function

action to choose:

$$a^*(x) = \operatorname{argmin}_a \sum_y L(a, y) P(y | x)$$

note: a and y need not live in the same space

special cases when $a = \hat{y}$ and y do live in the same space

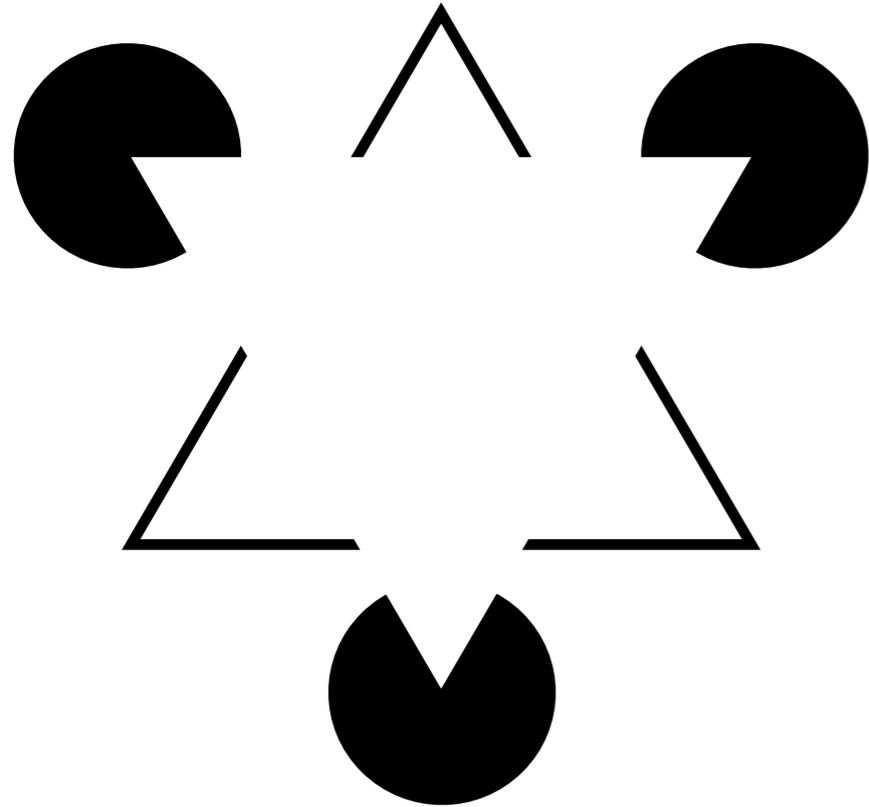
$$L(\hat{y}, y) = (\hat{y} - y)^2 \longrightarrow \hat{y}(x) = \sum_y y P(y | x) \quad \text{posterior mean}$$

$$L(\hat{y}, y) = |\hat{y} - y| \longrightarrow \sum_{y=-\infty}^{\hat{y}(x)} P(y | x) = \frac{1}{2} \quad \text{posterior median}$$

$$L(\hat{y}, y) = \begin{cases} 0 & \text{if } y = \hat{y} \\ 1 & \text{otherwise} \end{cases} \longrightarrow \hat{y}(x) = \operatorname{argmax}_y P(y | x) \quad \text{maximum a posteriori (MAP)}$$

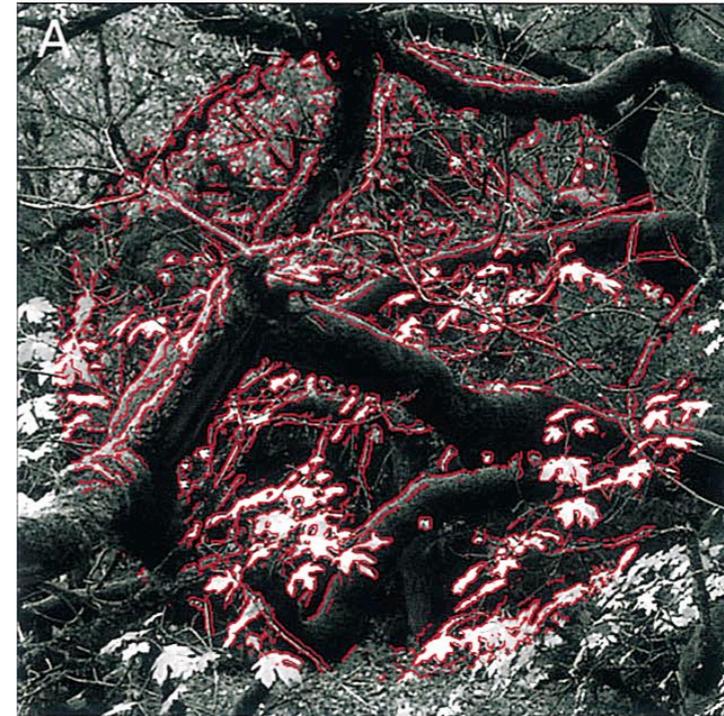
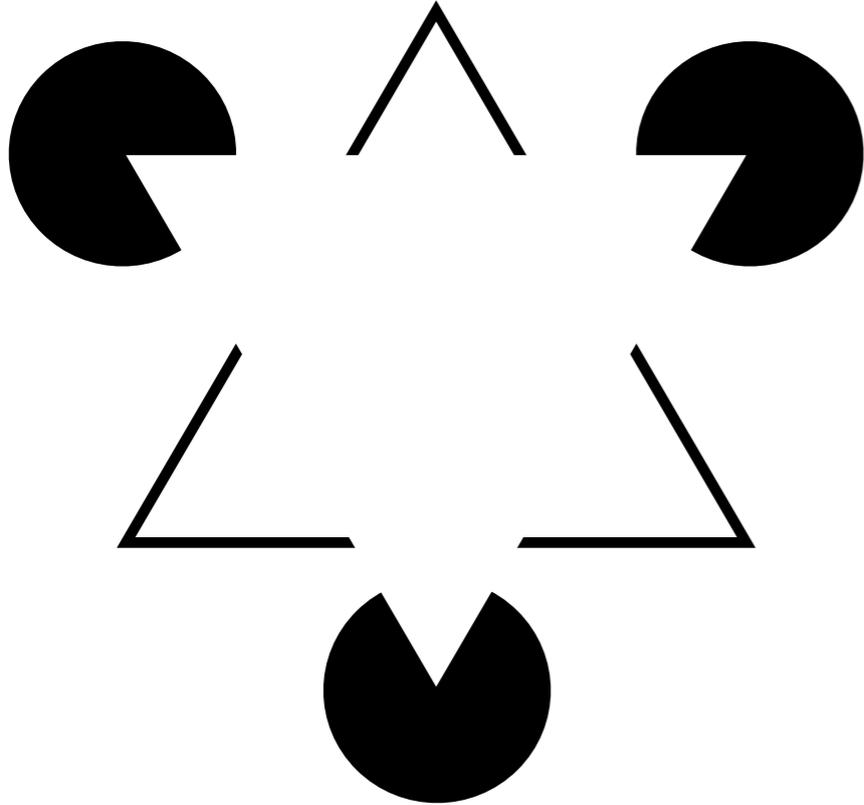
STATISTICAL ADAPTATION IN LOW LEVEL VISION

adaptation to natural image statistics?



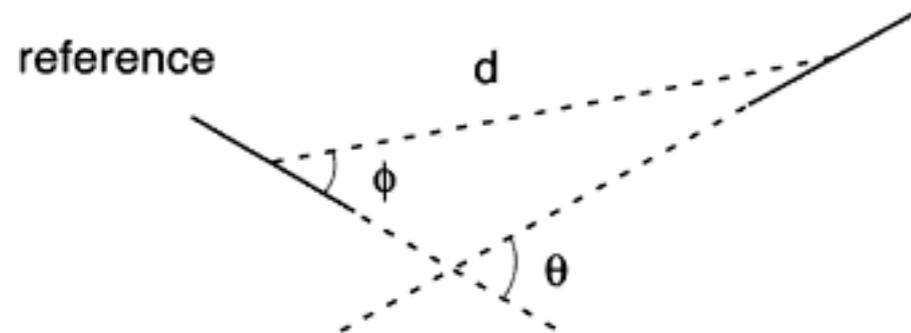
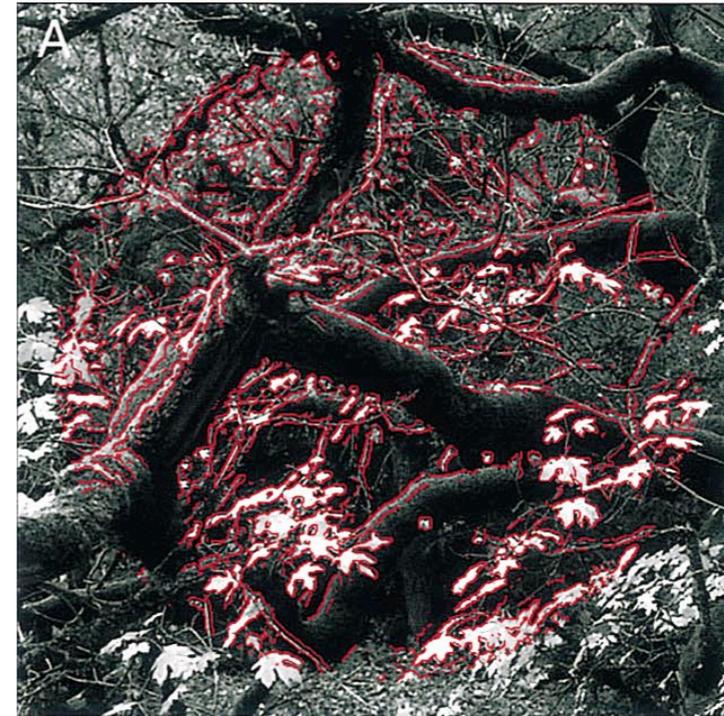
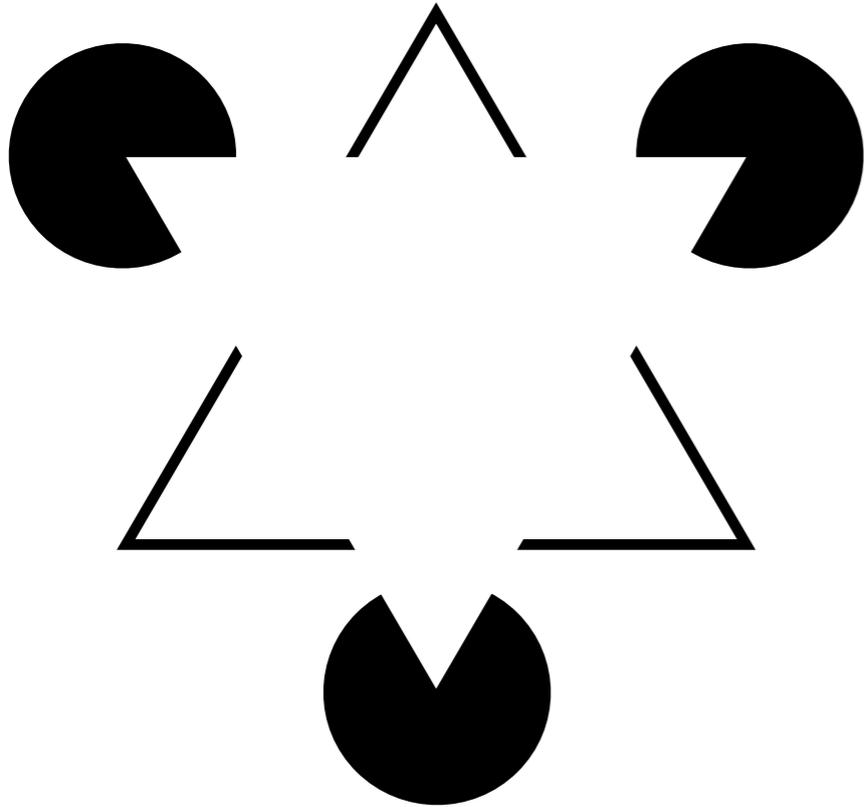
STATISTICAL ADAPTATION IN LOW LEVEL VISION

adaptation to natural image statistics?



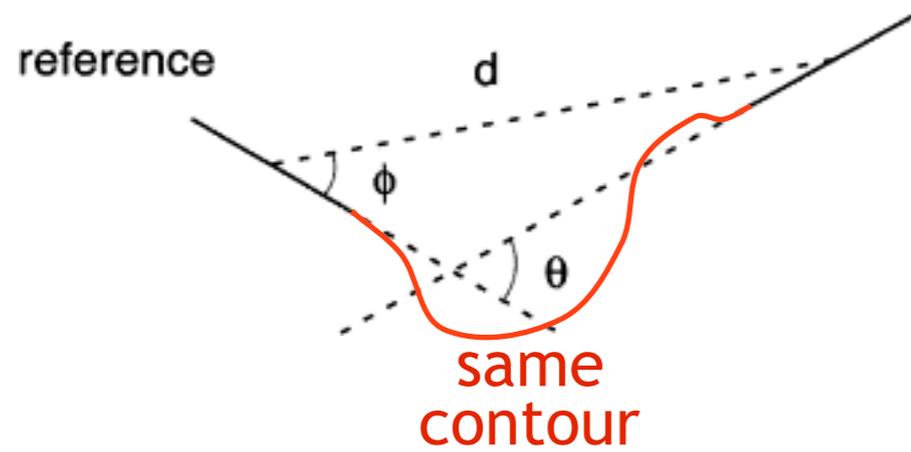
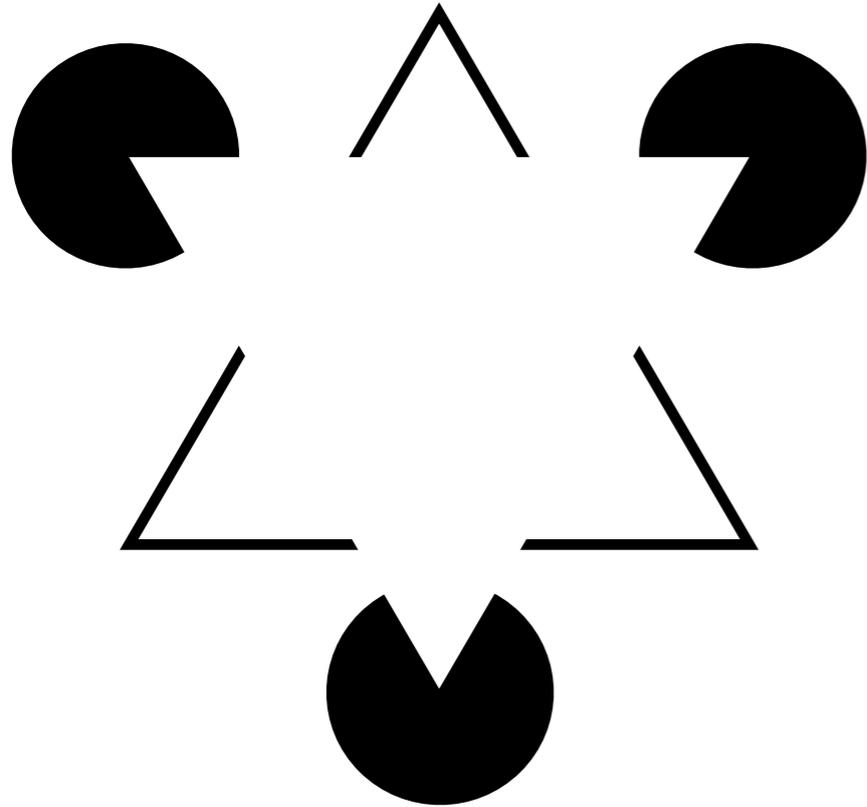
STATISTICAL ADAPTATION IN LOW LEVEL VISION

adaptation to natural image statistics?



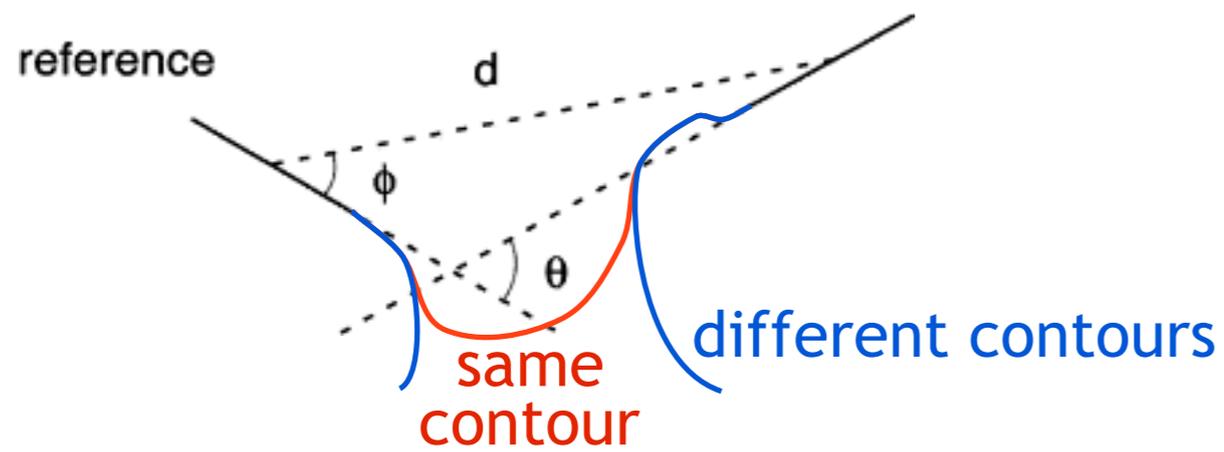
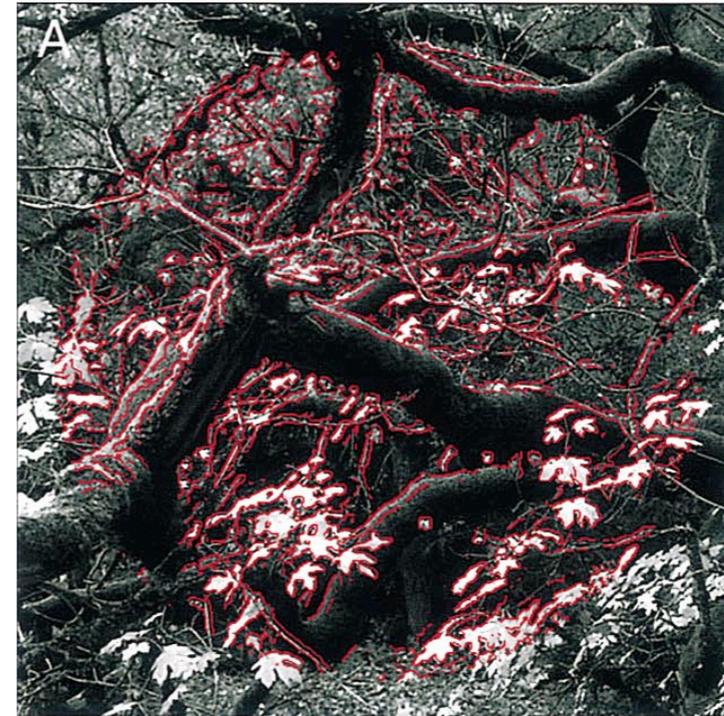
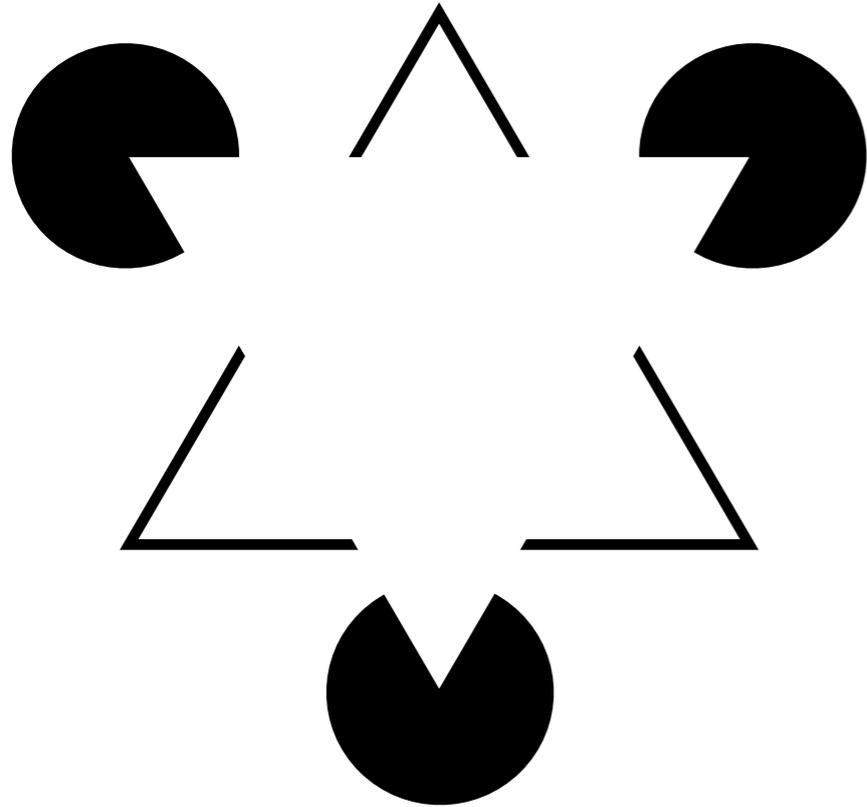
STATISTICAL ADAPTATION IN LOW LEVEL VISION

adaptation to natural image statistics?



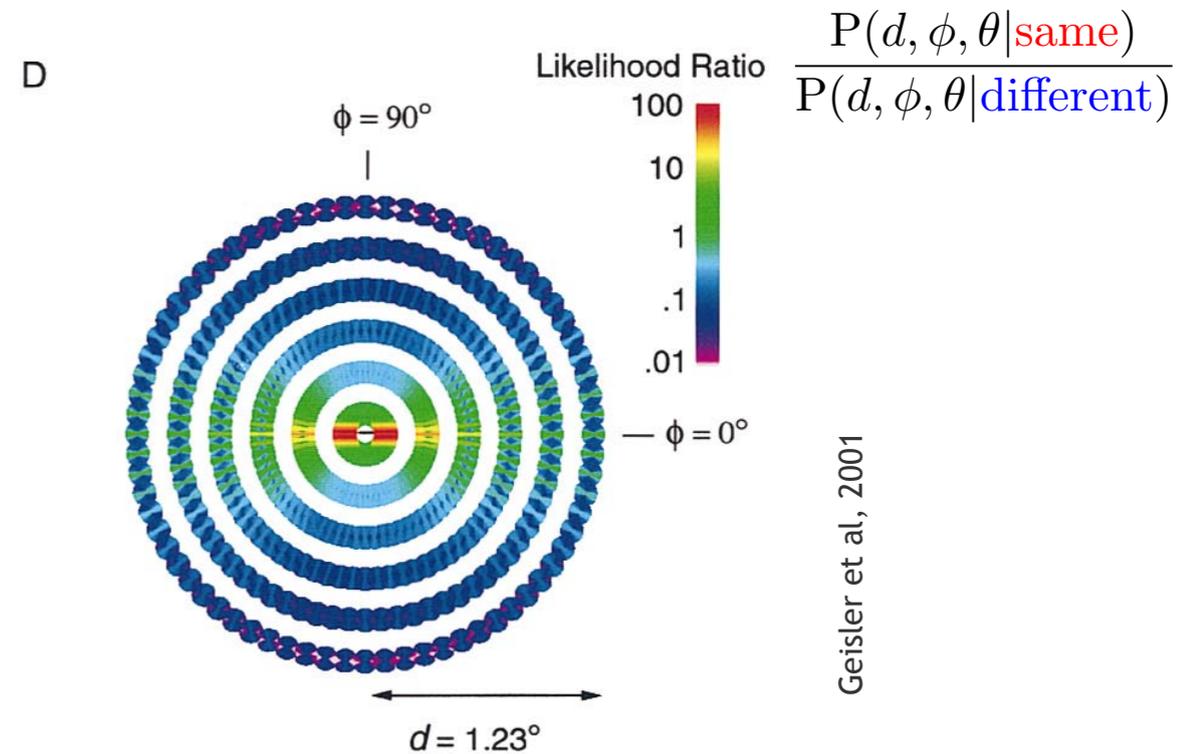
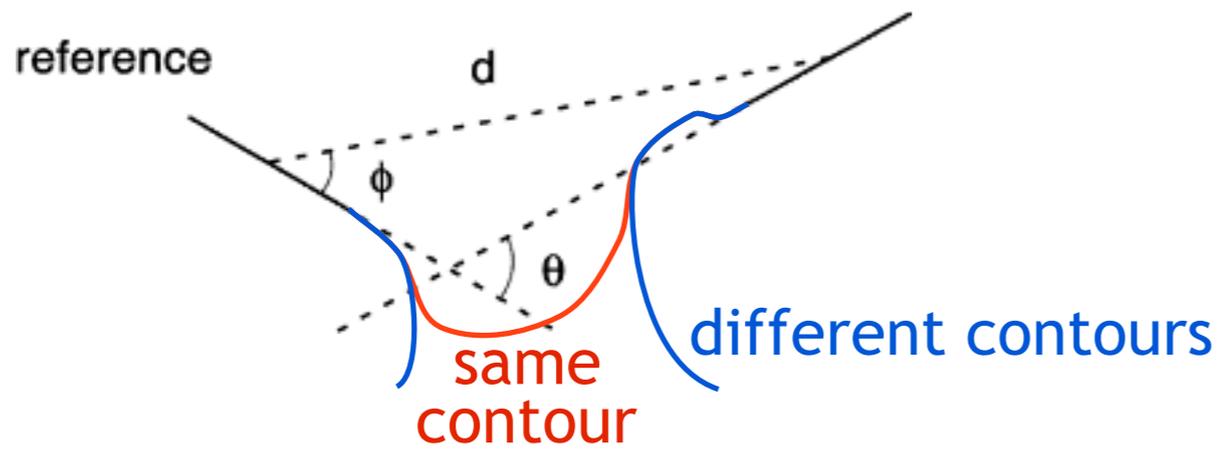
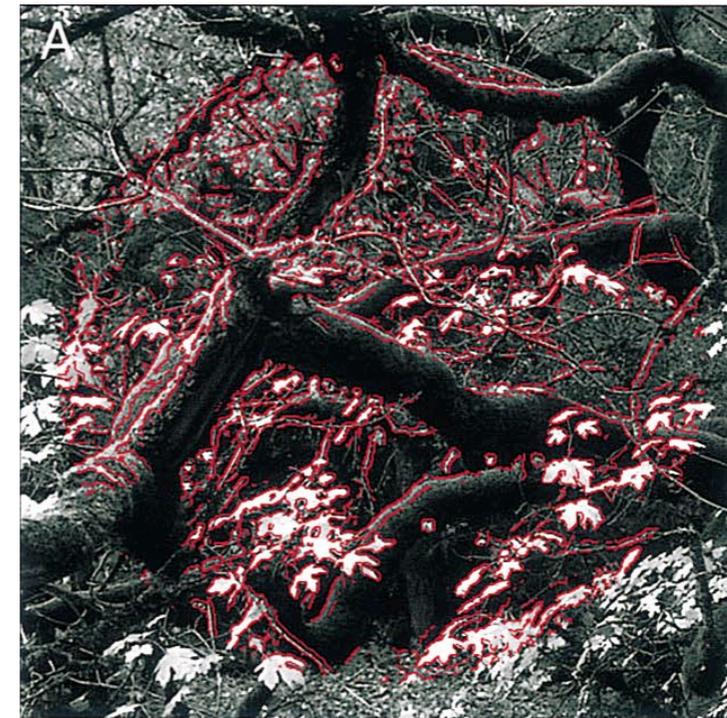
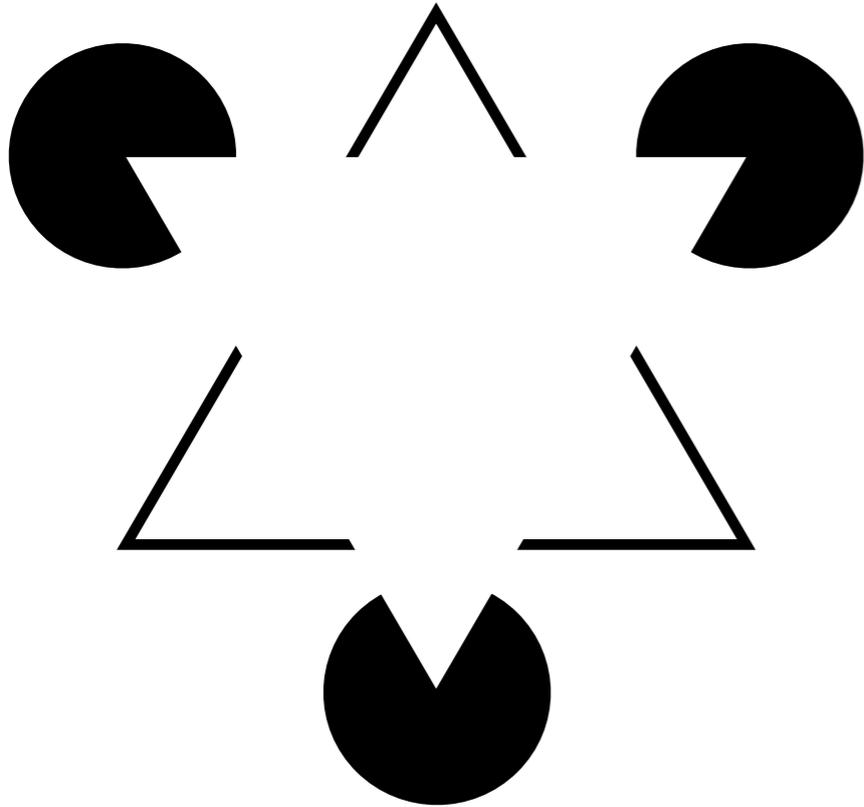
STATISTICAL ADAPTATION IN LOW LEVEL VISION

adaptation to natural image statistics?



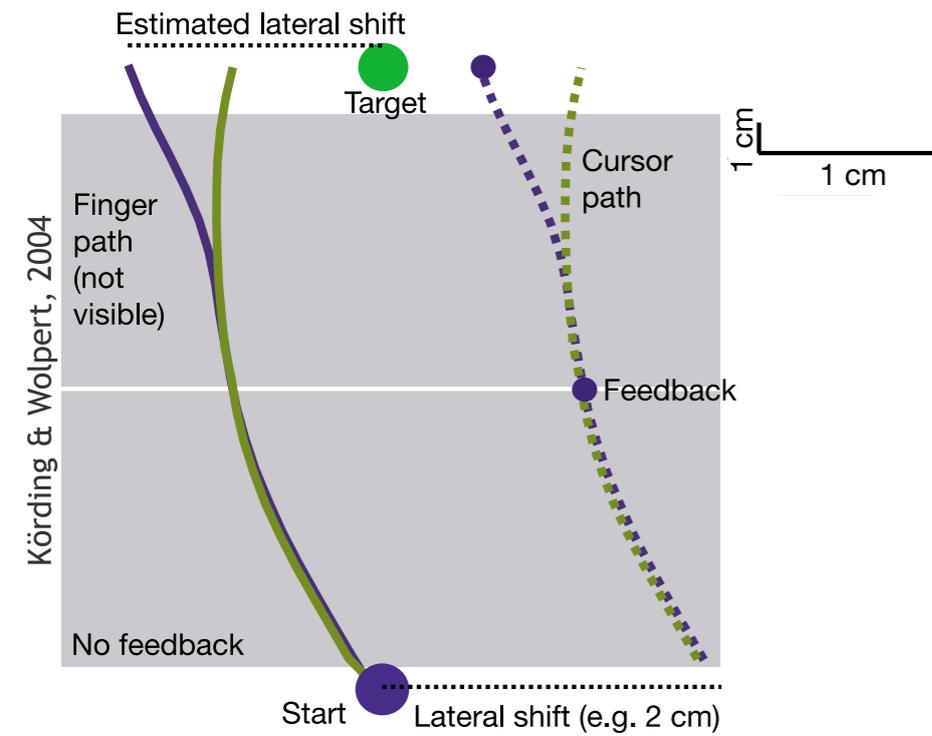
STATISTICAL ADAPTATION IN LOW LEVEL VISION

adaptation to natural image statistics?

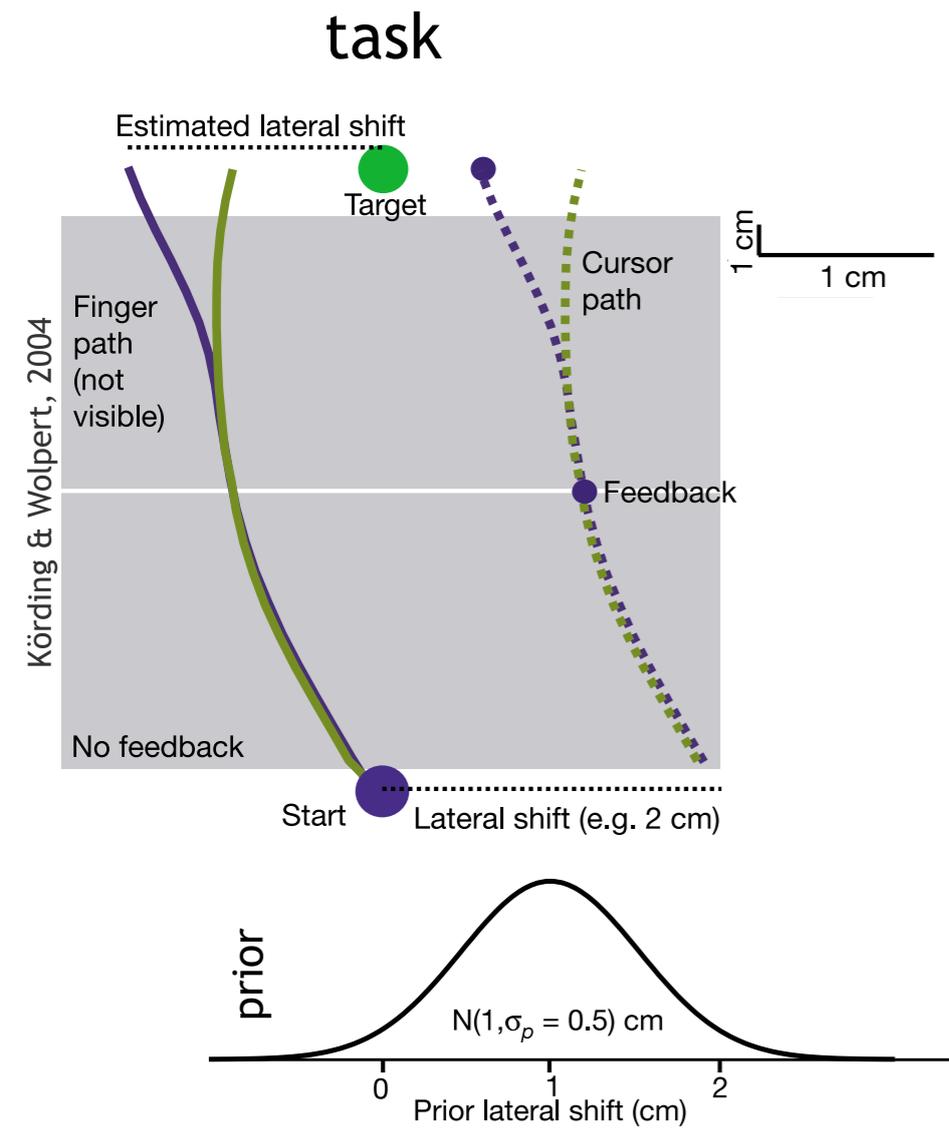


COMBINING PRIORS AND LIKELIHOODS: MOTOR CONTROL

task

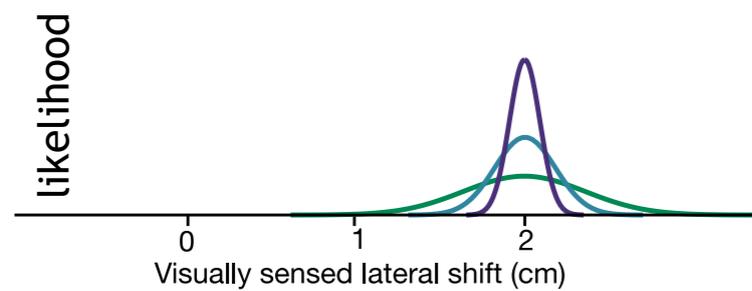
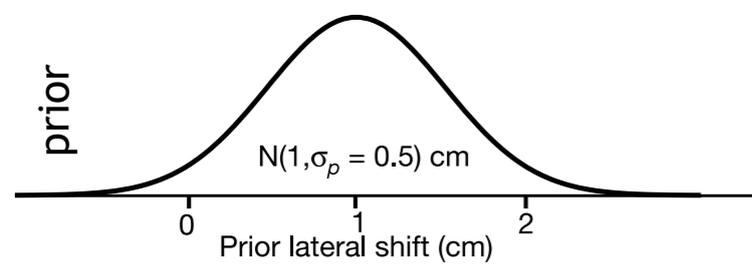
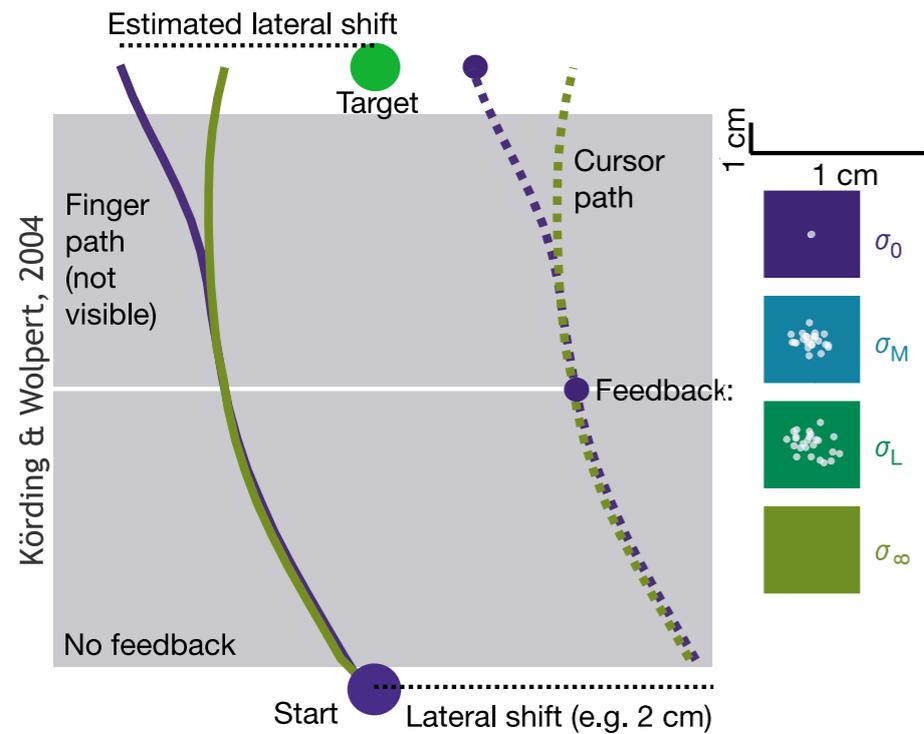


COMBINING PRIORS AND LIKELIHOODS: MOTOR CONTROL



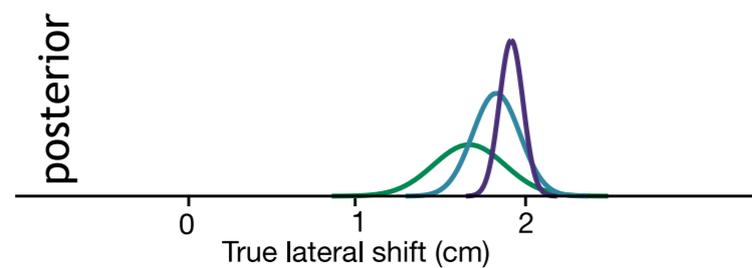
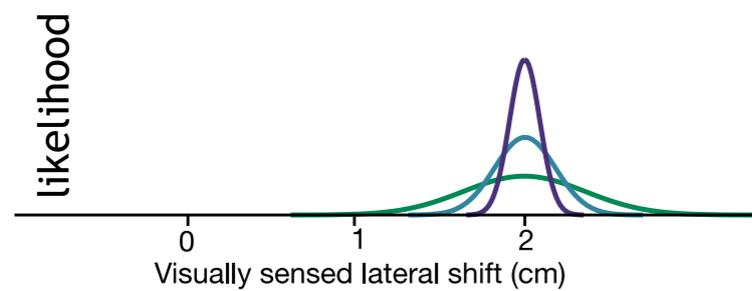
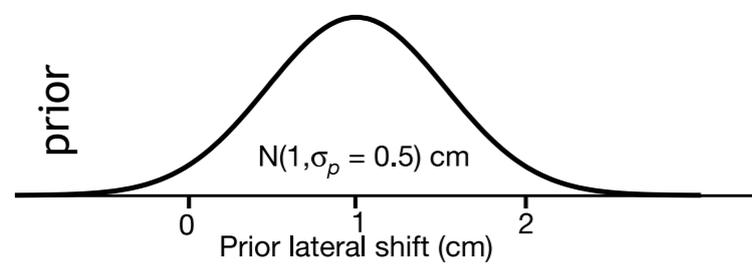
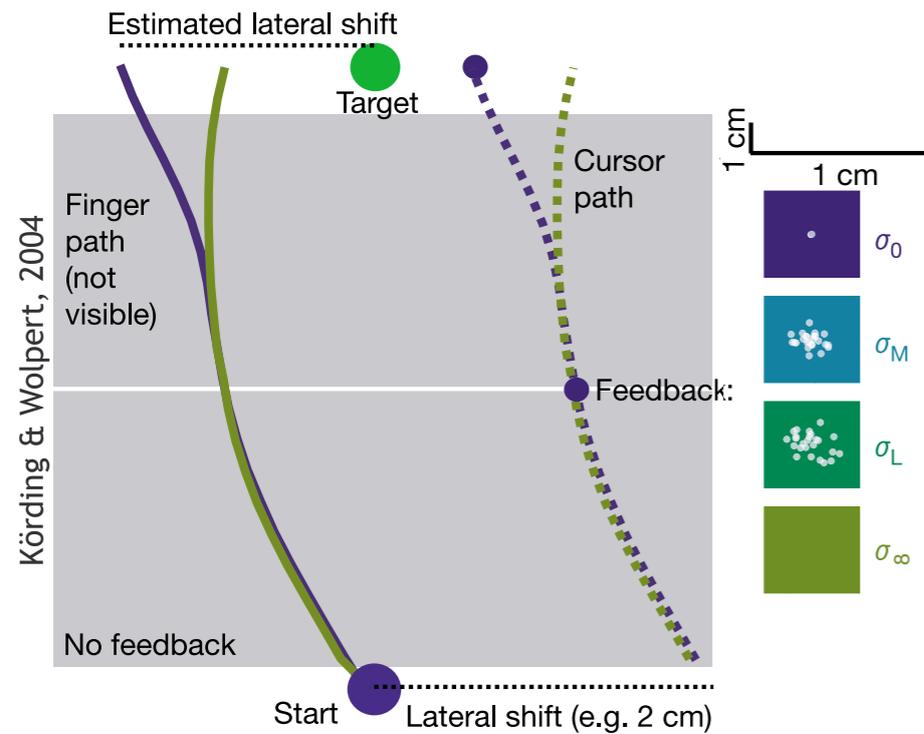
COMBINING PRIORS AND LIKELIHOODS: MOTOR CONTROL

task



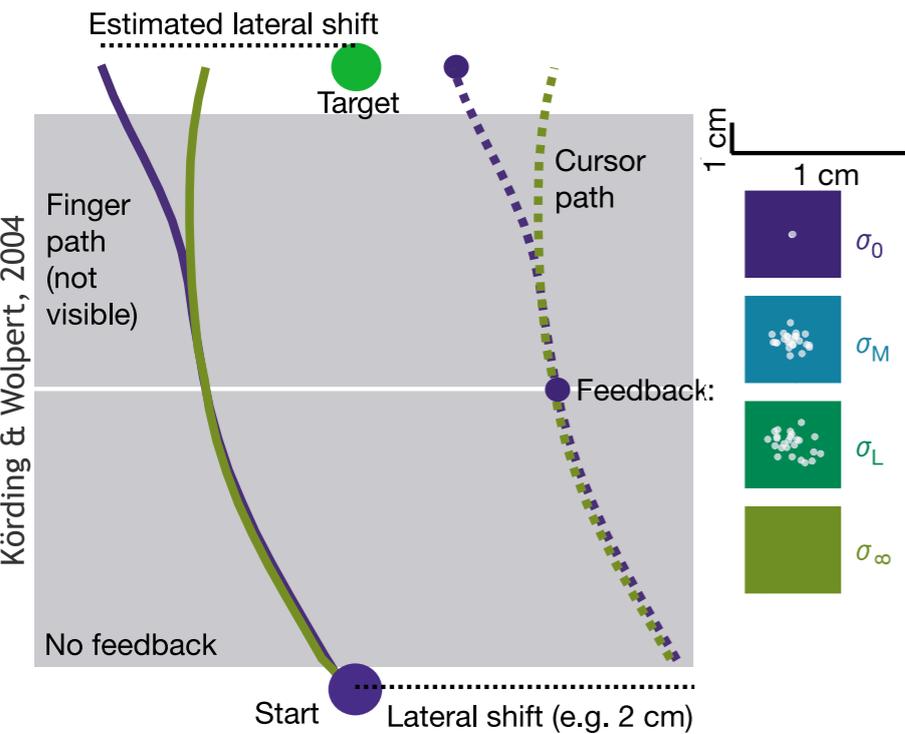
COMBINING PRIORS AND LIKELIHOODS: MOTOR CONTROL

task



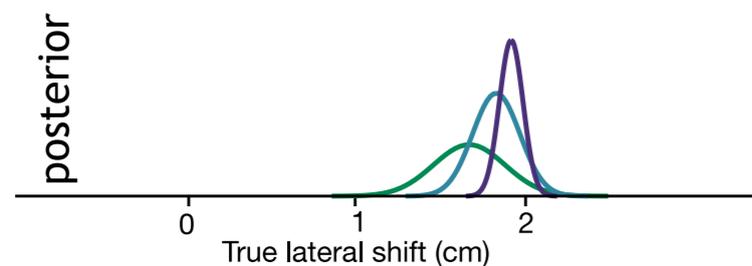
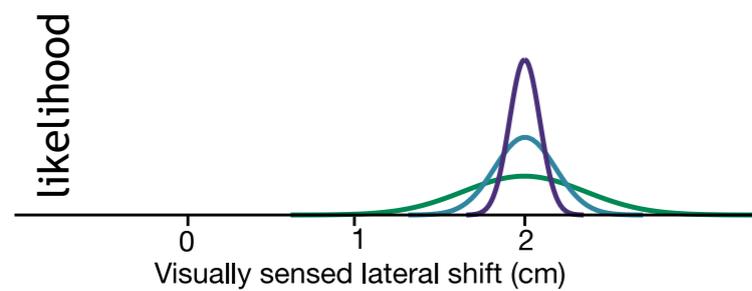
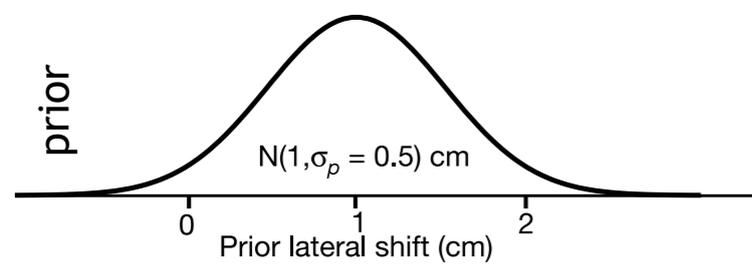
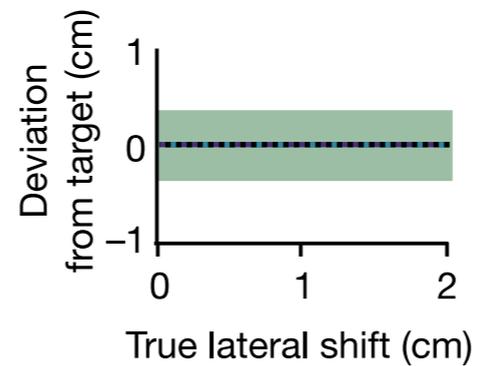
COMBINING PRIORS AND LIKELIHOODS: MOTOR CONTROL

task



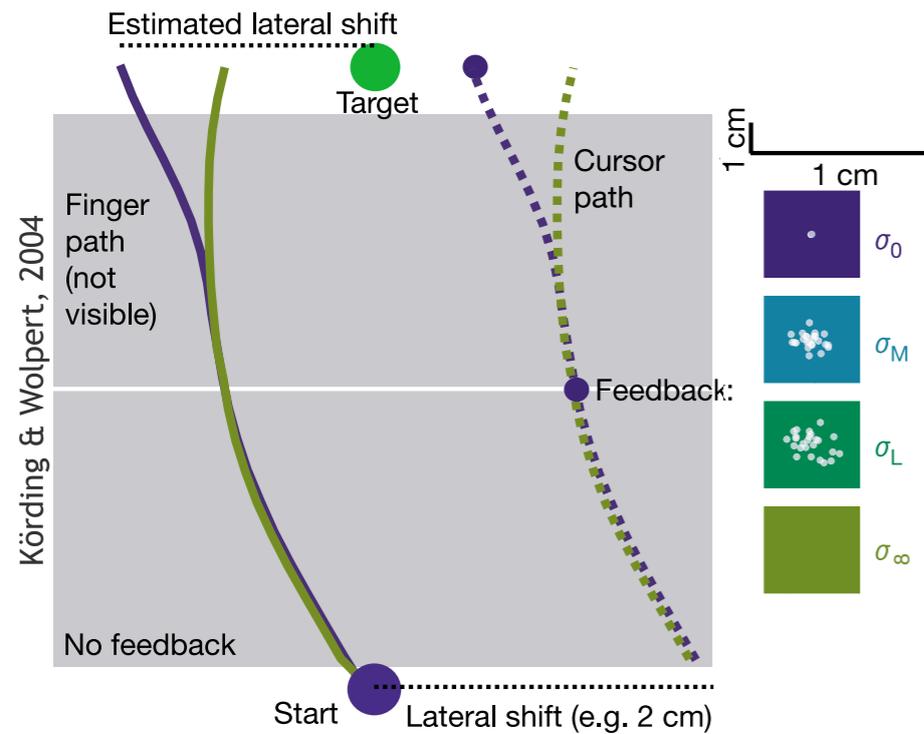
possible models

full compensation



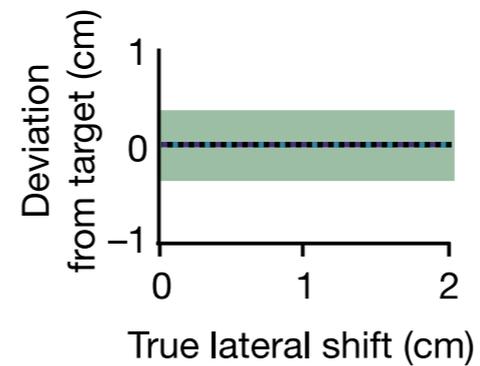
COMBINING PRIORS AND LIKELIHOODS: MOTOR CONTROL

task

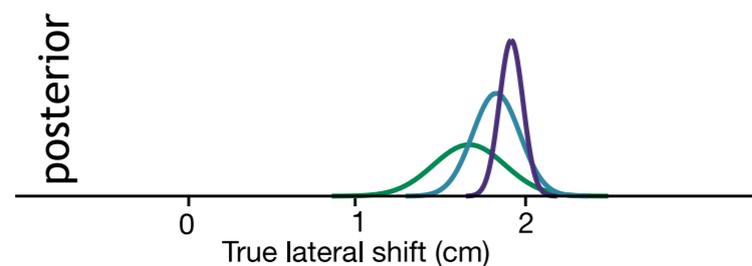
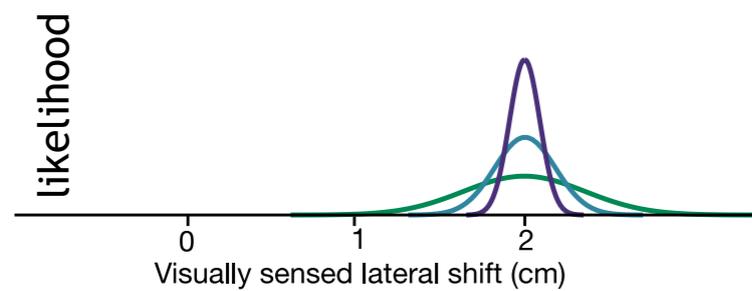
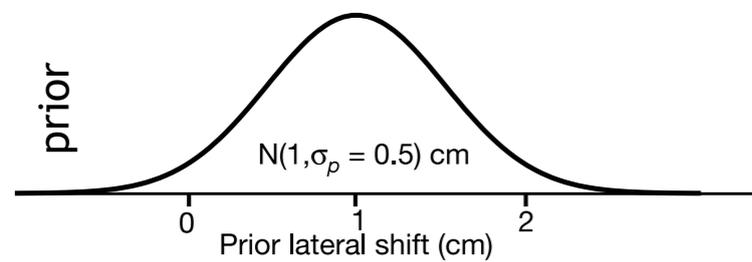
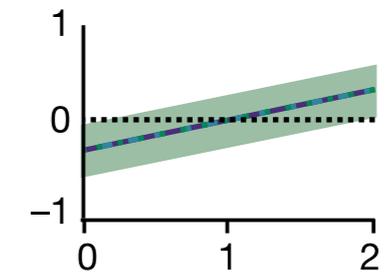


possible models

full compensation

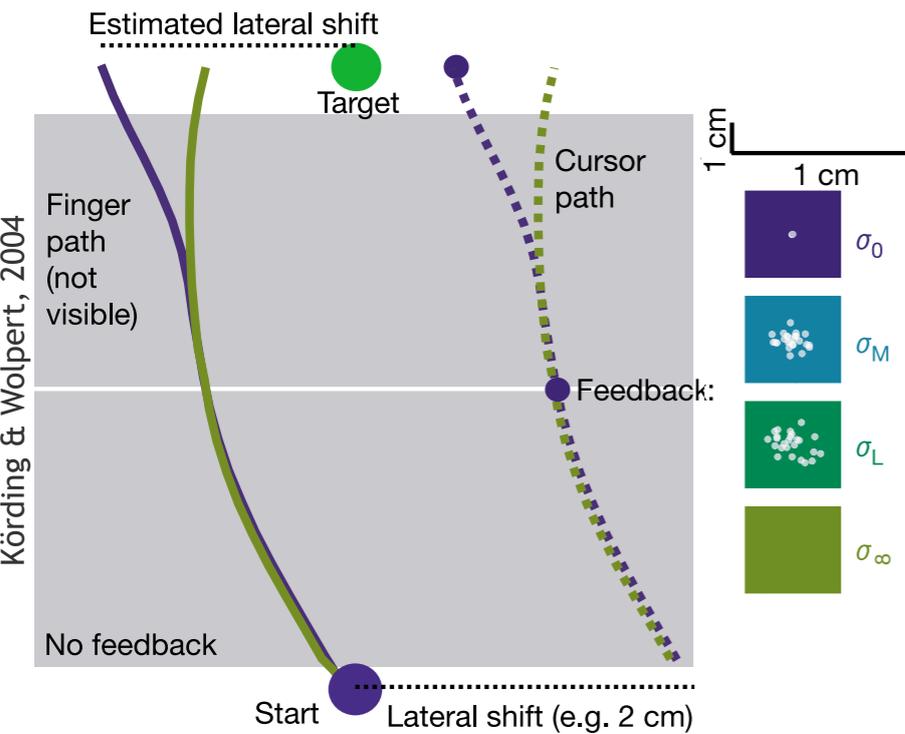


feedback-to-estimate mapping

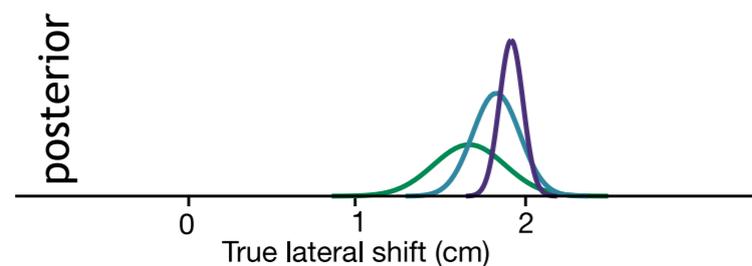
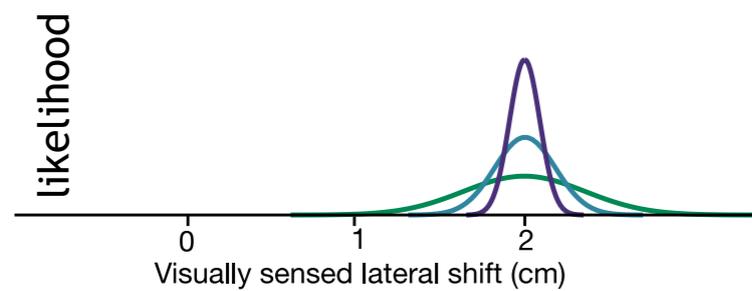
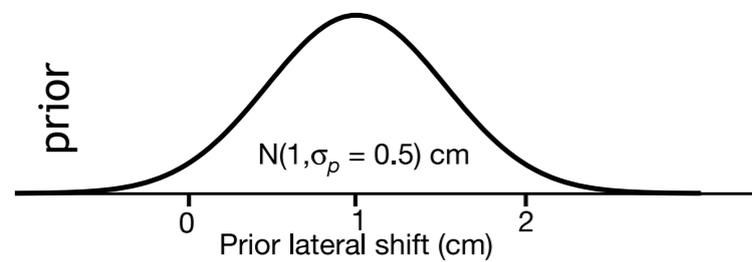
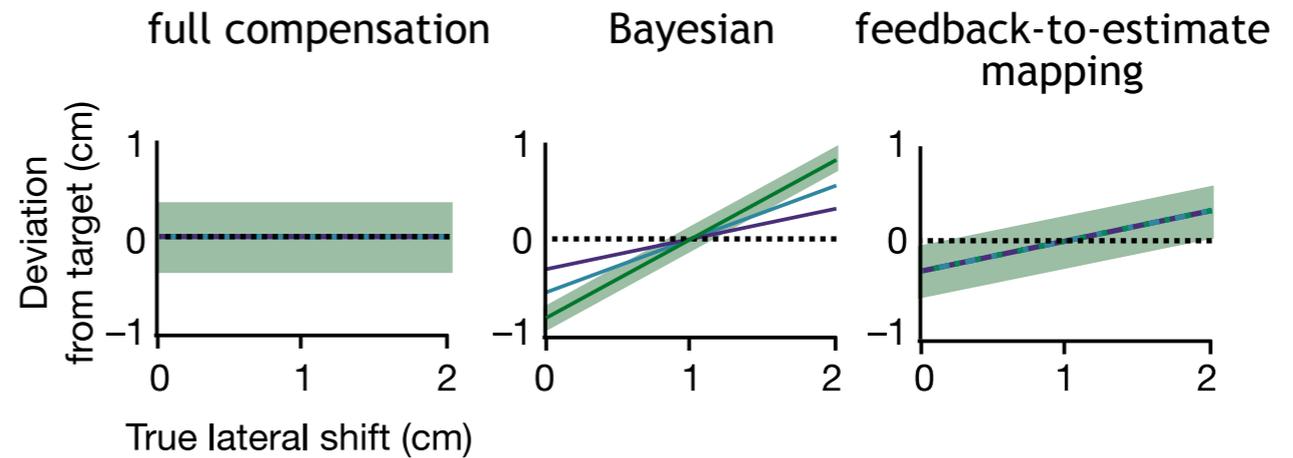


COMBINING PRIORS AND LIKELIHOODS: MOTOR CONTROL

task

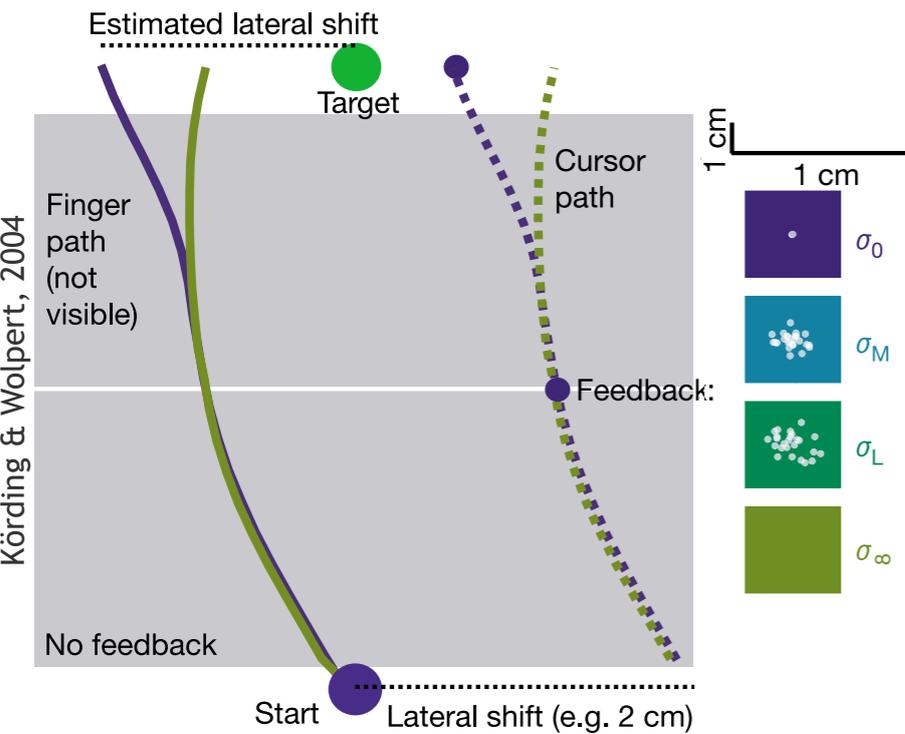


possible models

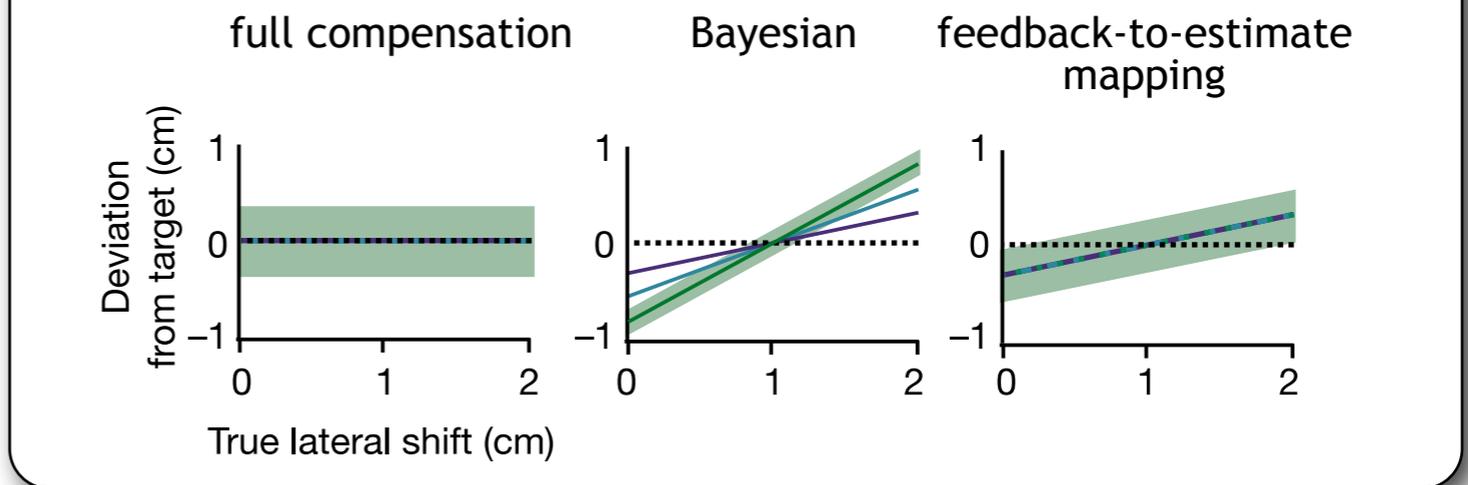


COMBINING PRIORS AND LIKELIHOODS: MOTOR CONTROL

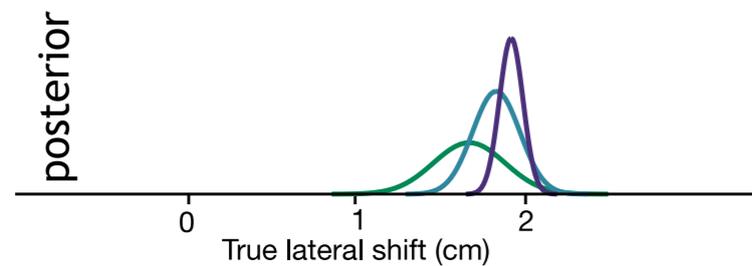
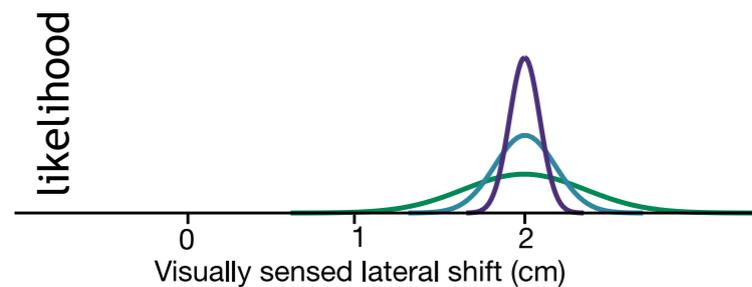
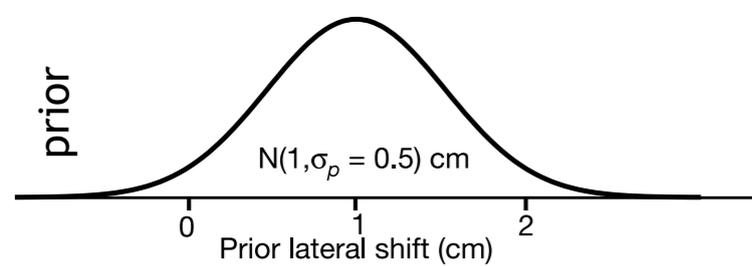
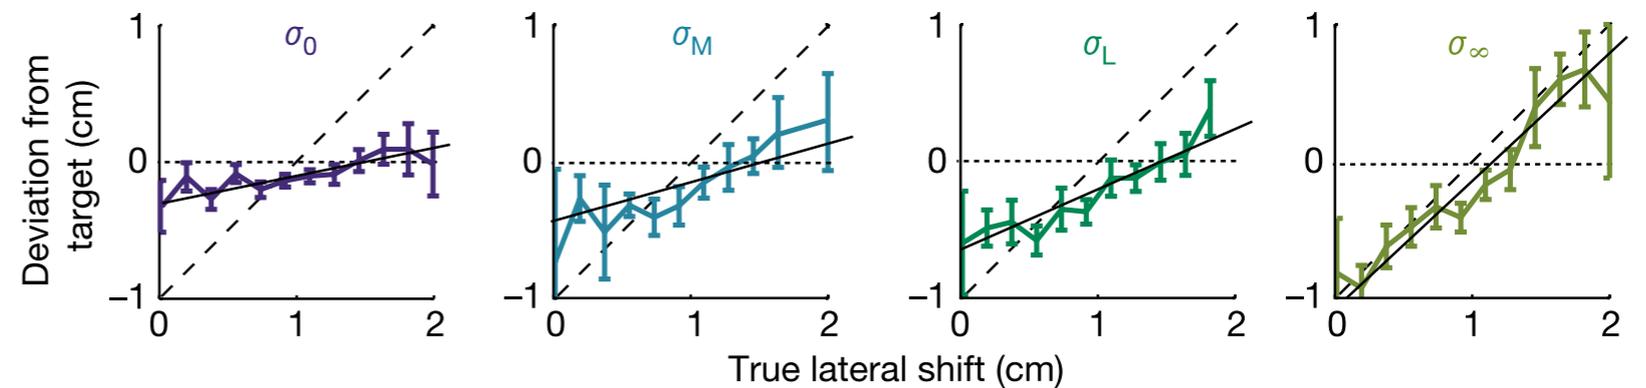
task



possible models

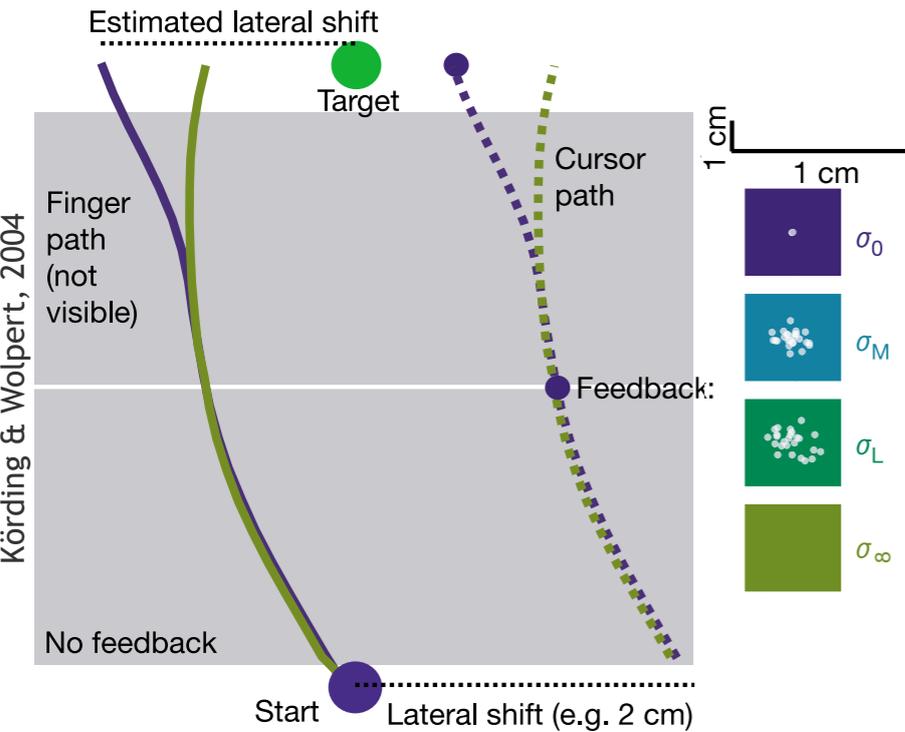


humans

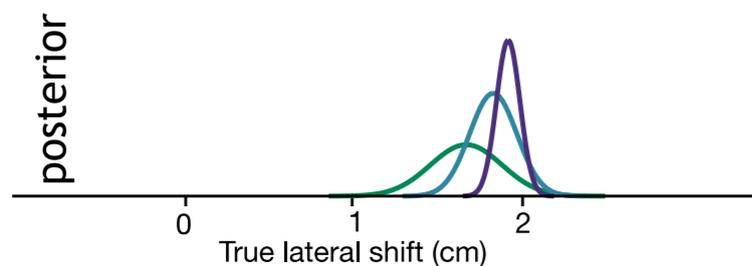
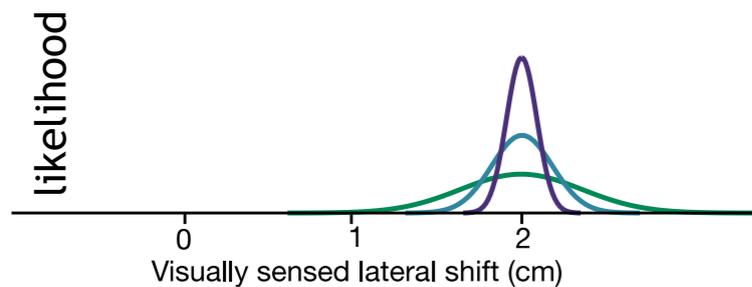
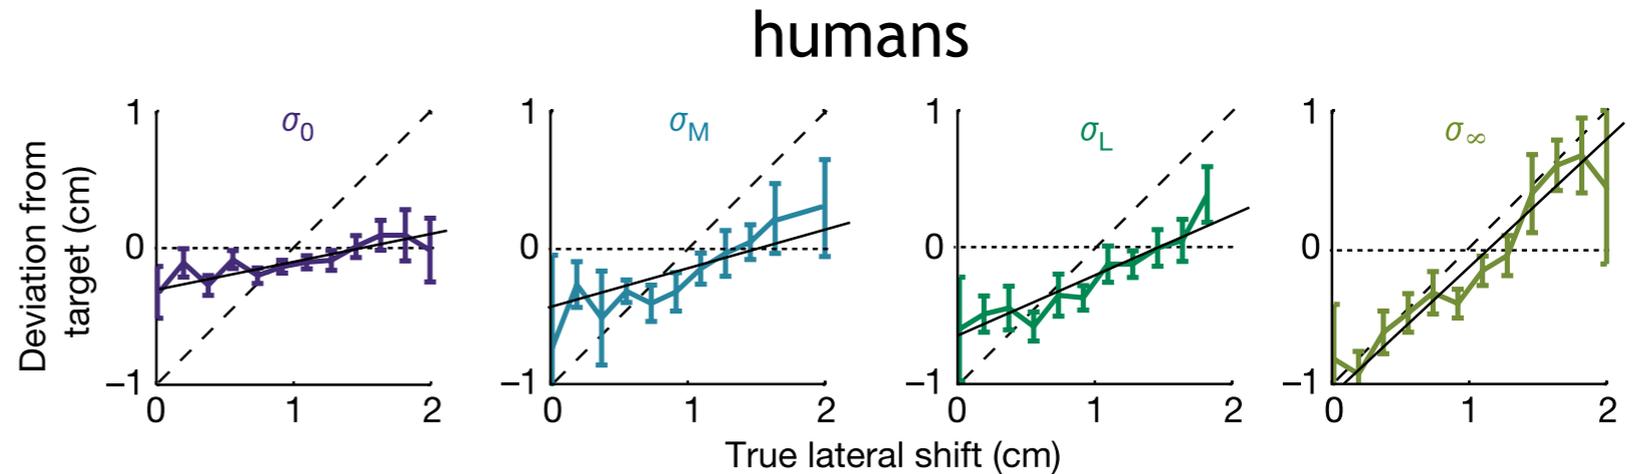
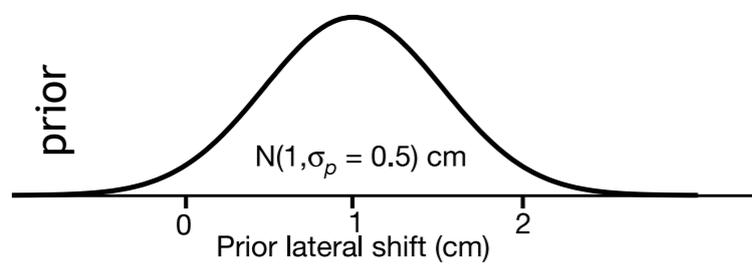
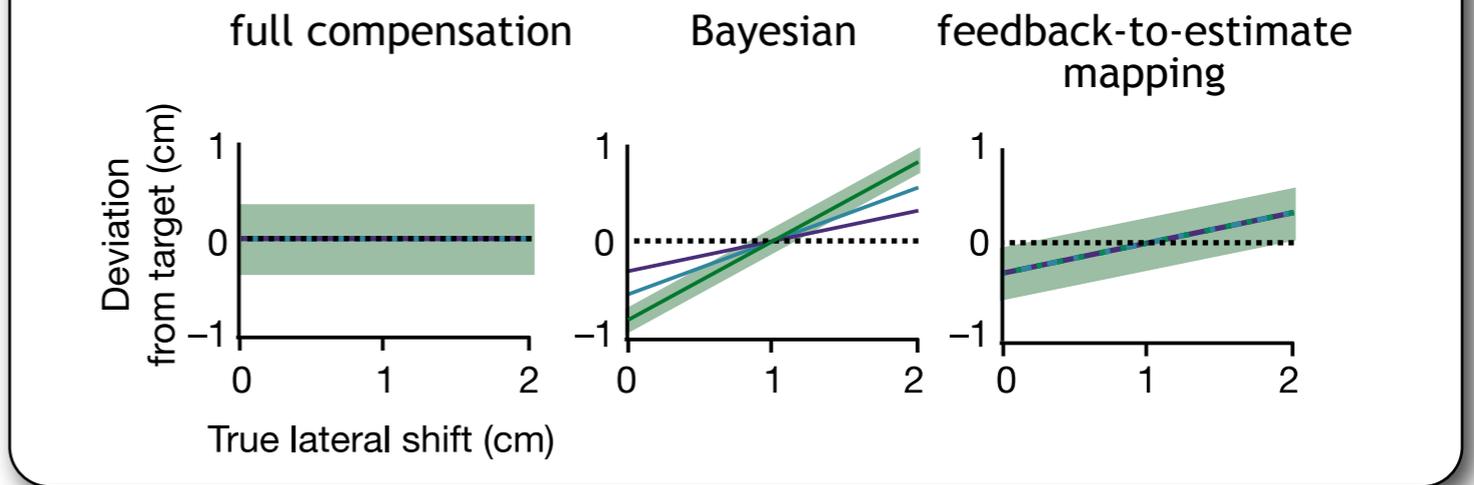


COMBINING PRIORS AND LIKELIHOODS: MOTOR CONTROL

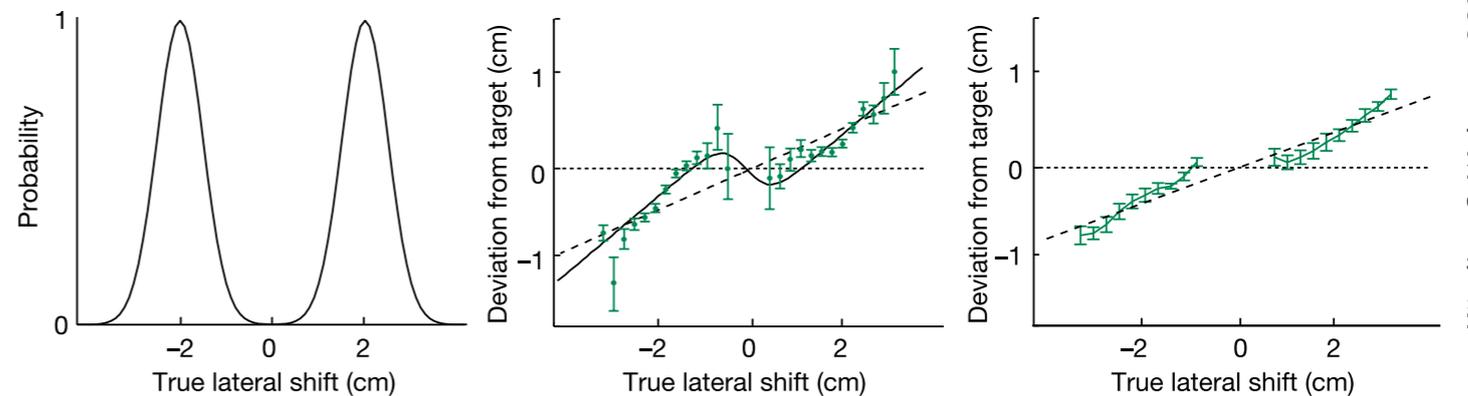
task



possible models



bimodal prior



Körding & Wolpert, 2004

COMBINING PRIORS AND LIKELIHOODS: COGNITION

You meet someone who is t years old. What will be his total life span t_{total} ?

COMBINING PRIORS AND LIKELIHOODS: COGNITION

You meet someone who is t years old. What will be his total life span t_{total} ?

$$P(t_{\text{total}}|t) \propto P(t|t_{\text{total}}) P(t_{\text{total}})$$

COMBINING PRIORS AND LIKELIHOODS: COGNITION

You meet someone who is t years old. What will be his total life span t_{total} ?

$$P(t_{\text{total}}|t) \propto P(t|t_{\text{total}}) P(t_{\text{total}})$$



the probability that
you meet someone
at the age of t
when s/he will have
a total life span of t_{total}

COMBINING PRIORS AND LIKELIHOODS: COGNITION

You meet someone who is t years old. What will be his total life span t_{total} ?

$$P(t_{\text{total}}|t) \propto P(t|t_{\text{total}}) P(t_{\text{total}})$$



the probability that
you meet someone
at the age of t
when s/he will have
a total life span of t_{total}

$$\approx \begin{cases} \frac{1}{t_{\text{total}}} & \text{if } t < t_{\text{total}} \\ 0 & \text{otherwise} \end{cases}$$

COMBINING PRIORS AND LIKELIHOODS: COGNITION

You meet someone who is t years old. What will be his total life span t_{total} ?

$$P(t_{\text{total}}|t) \propto P(t|t_{\text{total}}) P(t_{\text{total}})$$

the probability that
you meet someone
at the age of t
when s/he will have
a total life span of t_{total}

$$P(t_{\text{total}})$$

prior on life span
distribution of people

$$\approx \begin{cases} \frac{1}{t_{\text{total}}} & \text{if } t < t_{\text{total}} \\ 0 & \text{otherwise} \end{cases}$$

COMBINING PRIORS AND LIKELIHOODS: COGNITION

You meet someone who is t years old. What will be his total life span t_{total} ?

$$P(t_{\text{total}}|t) \propto P(t|t_{\text{total}}) P(t_{\text{total}})$$

the probability that you meet someone at the age of t when s/he will have a total life span of t_{total}

$P(t_{\text{total}})$
prior on life span distribution of people

+ decision theory
e.g. report the median of the posterior

$$\approx \begin{cases} \frac{1}{t_{\text{total}}} & \text{if } t < t_{\text{total}} \\ 0 & \text{otherwise} \end{cases}$$

COMBINING PRIORS AND LIKELIHOODS: COGNITION

You meet someone who is t years old. What will be his total life span t_{total} ?

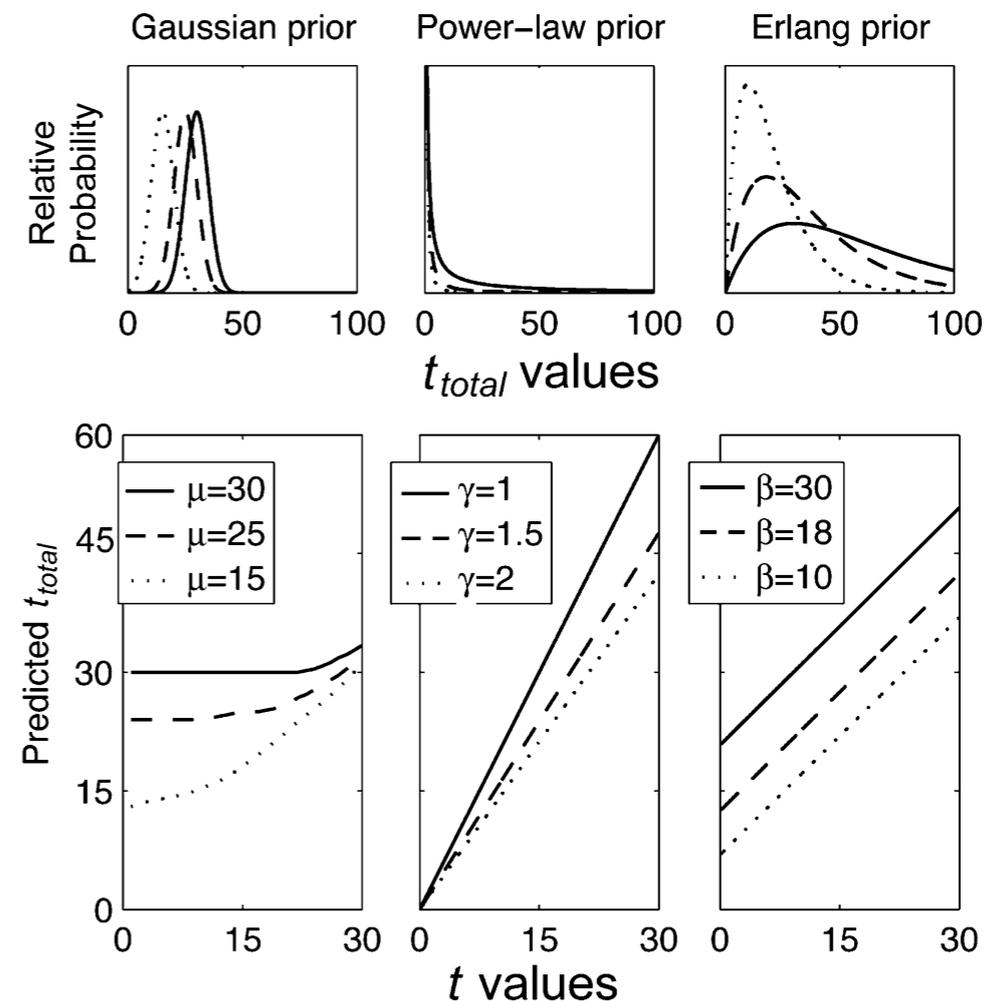
$$P(t_{total}|t) \propto P(t|t_{total}) P(t_{total})$$

the probability that you meet someone at the age of t when s/he will have a total life span of t_{total}

$P(t_{total})$
prior on life span distribution of people

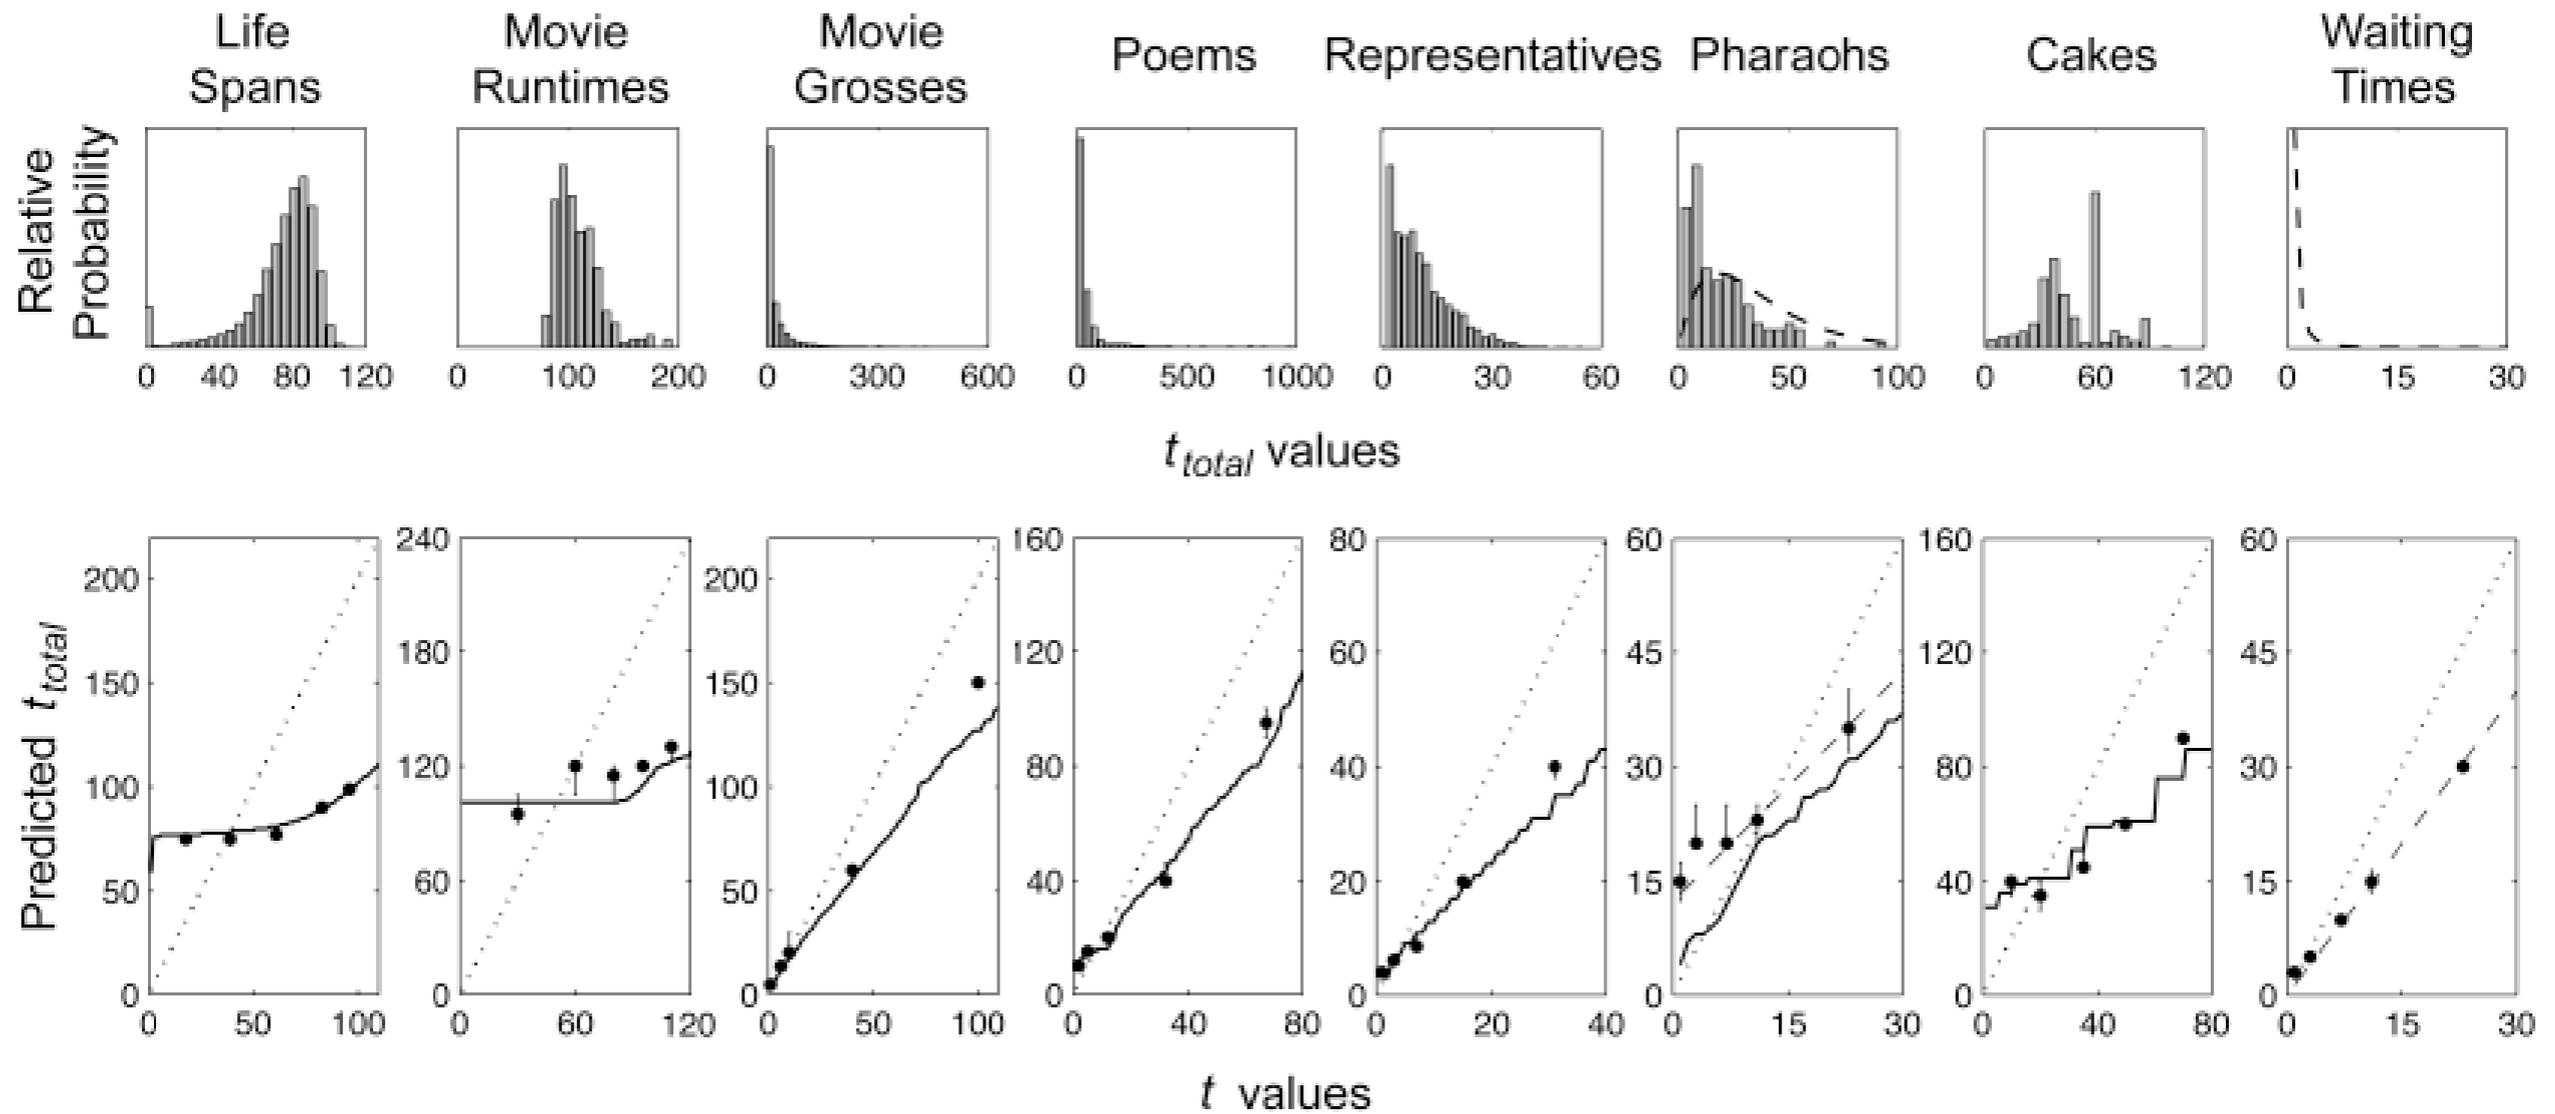
+ decision theory
e.g. report the median of the posterior

$$\approx \begin{cases} \frac{1}{t_{total}} & \text{if } t < t_{total} \\ 0 & \text{otherwise} \end{cases}$$



Griffiths & Tenenbaum, 2006

ADAPTATION TO ENVIRONMENTAL STATISTICS



Griffiths & Tenenbaum, 2006

WHAT ARE NATURAL PRIORS LIKE?

WHAT ARE NATURAL PRIORS LIKE?

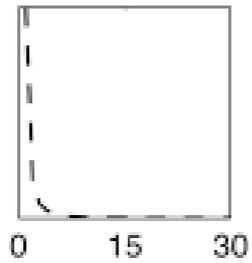
- ▶ adapted to ‘natural’ statistics

WHAT ARE NATURAL PRIORS LIKE?

► adapted to ‘natural’ statistics

box office phone
waiting times

*Griffiths &
Tenenbaum, 2006*

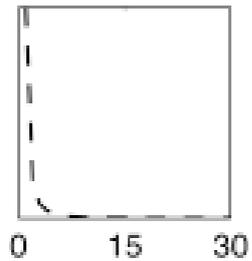


WHAT ARE NATURAL PRIORS LIKE?

- ▶ adapted to 'natural' statistics

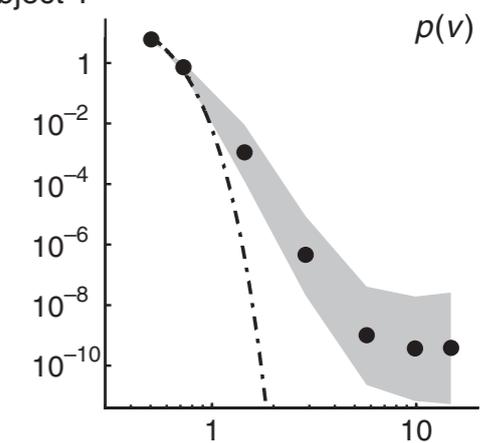
box office phone waiting times

Griffiths & Tenenbaum, 2006

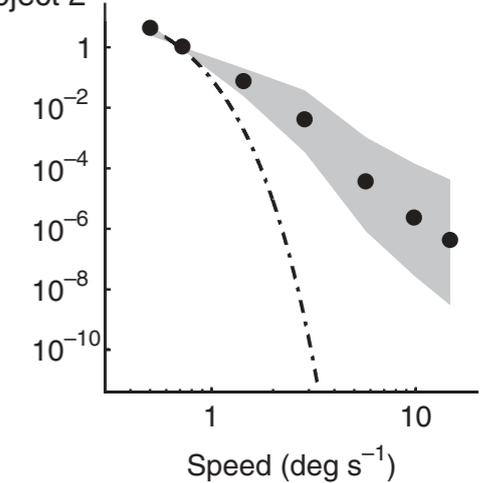


speed of visual motion

Subject 1



Subject 2



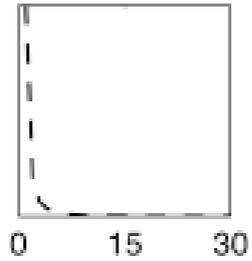
Stocker & Simoncelli, 2006

WHAT ARE NATURAL PRIORS LIKE?

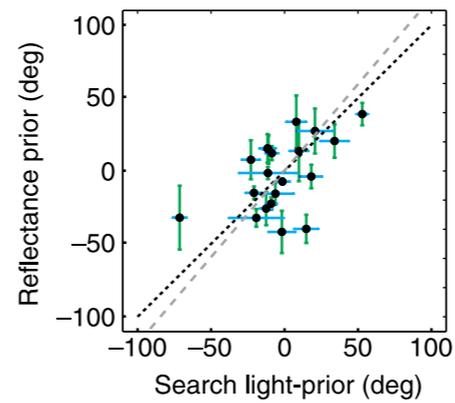
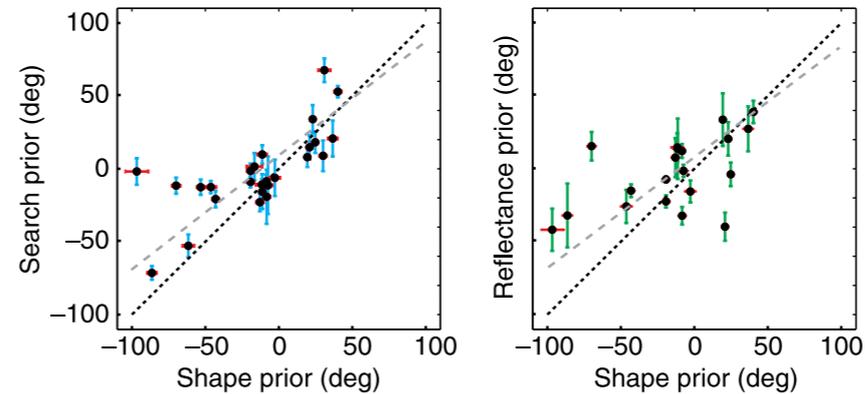
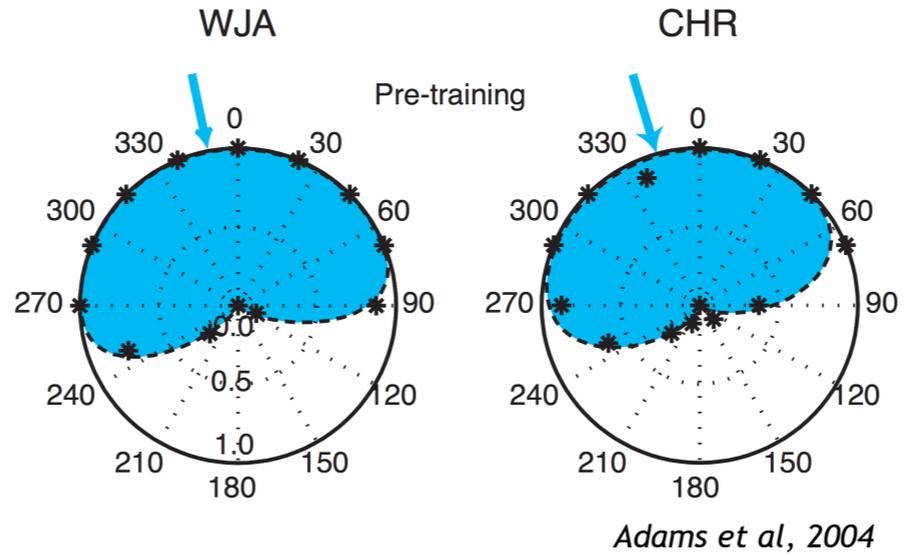
► adapted to ‘natural’ statistics

box office phone waiting times

Griffiths & Tenenbaum, 2006

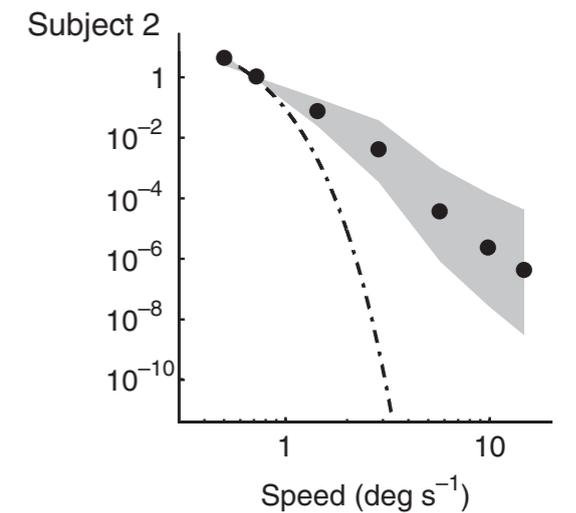
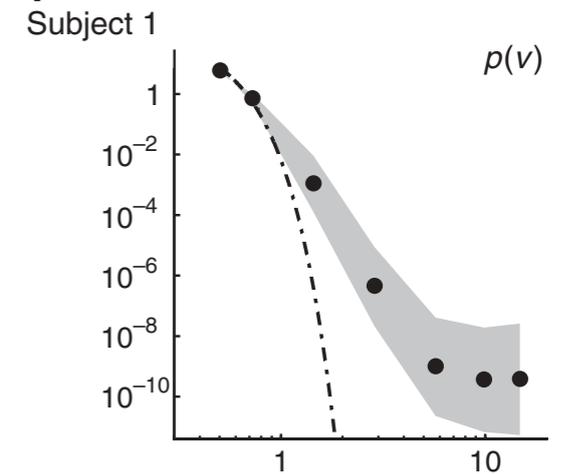


direction of light



Adams, 2007

speed of visual motion



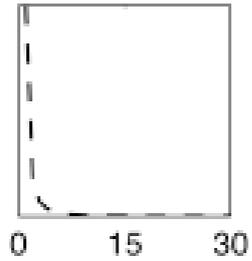
Stocker & Simoncelli, 2006

WHAT ARE NATURAL PRIORS LIKE?

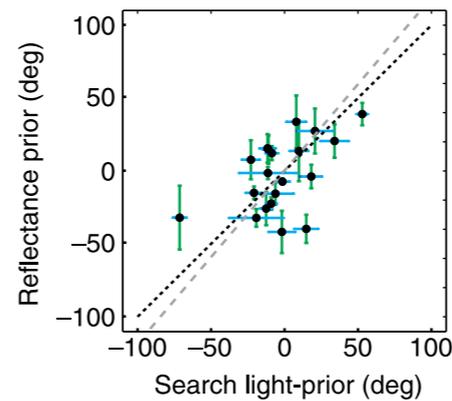
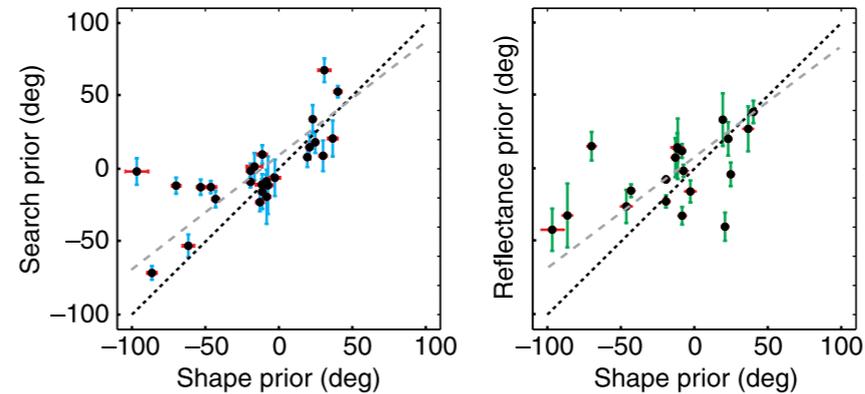
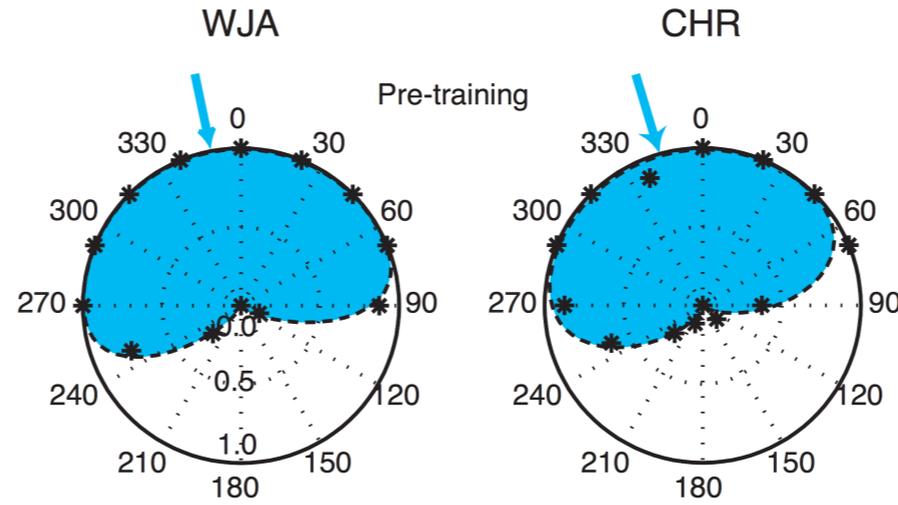
► adapted to ‘natural’ statistics

box office phone waiting times

Griffiths & Tenenbaum, 2006

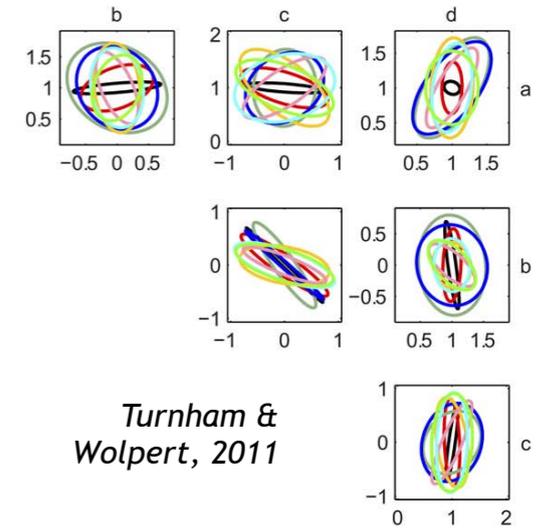


direction of light

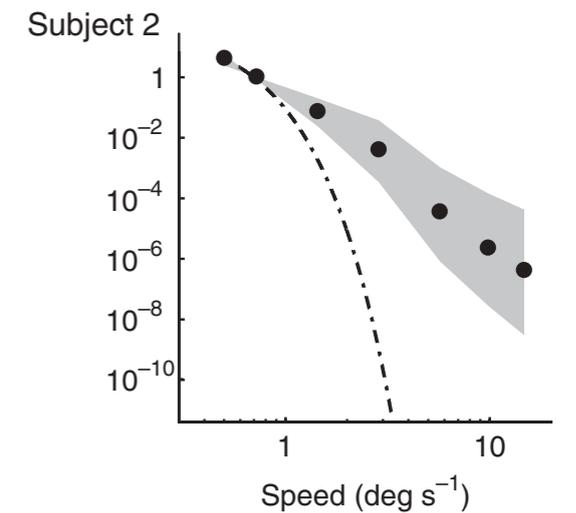
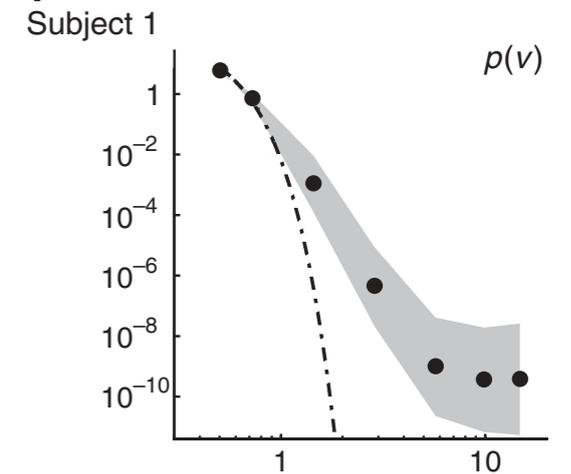


Adams, 2007

visuomotor transformations



speed of visual motion



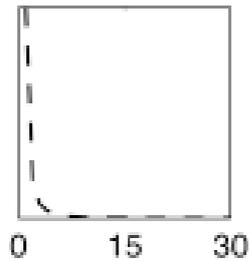
Stocker & Simoncelli, 2006

WHAT ARE NATURAL PRIORS LIKE?

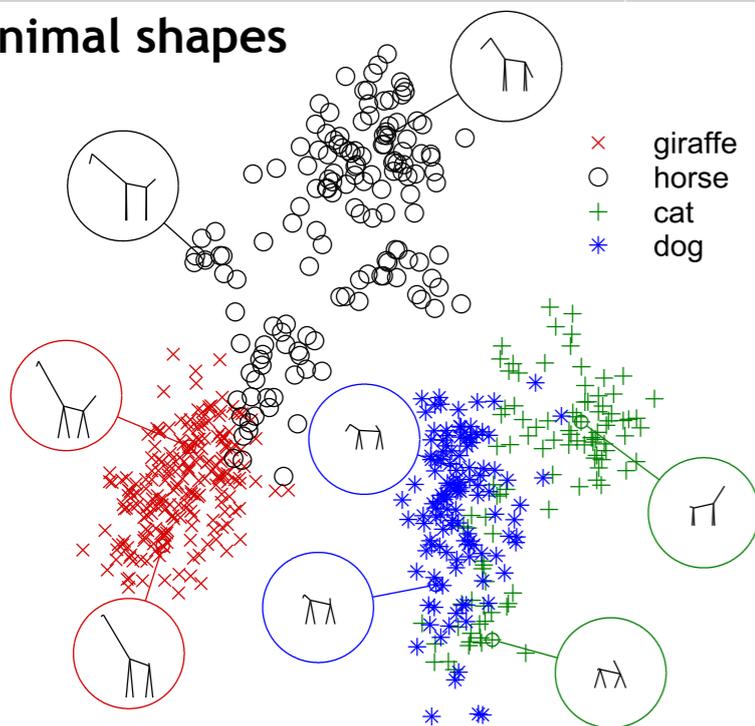
► adapted to 'natural' statistics

box office phone waiting times

Griffiths & Tenenbaum, 2006

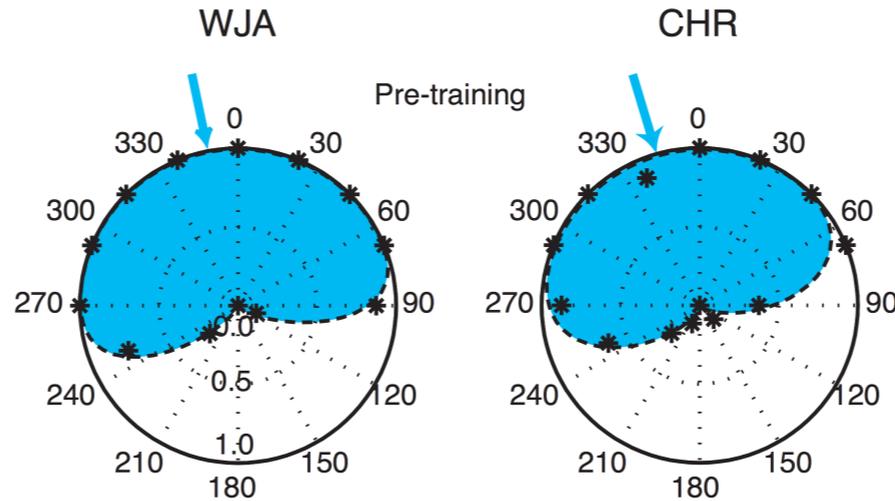


animal shapes

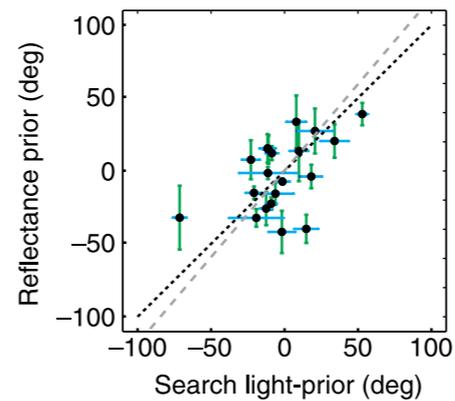
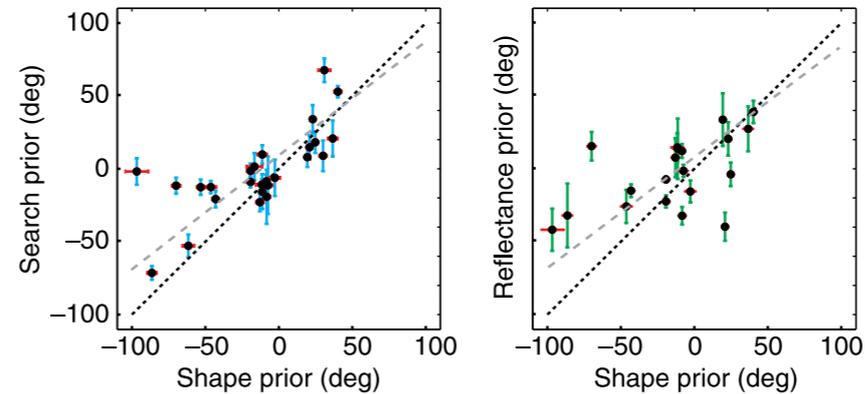


Sanborn & Griffiths, 2008

direction of light

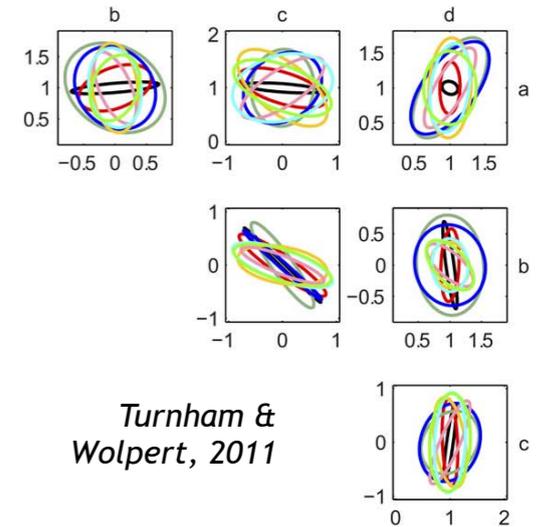


Adams et al, 2004



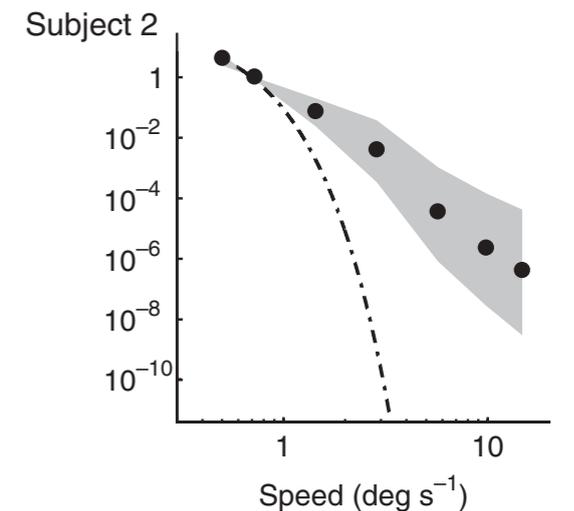
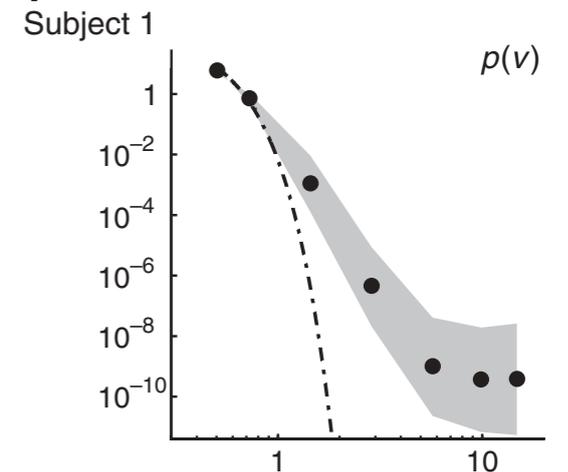
Adams, 2007

visuomotor transformations



Turnham & Wolpert, 2011

speed of visual motion



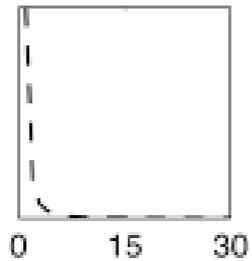
Stocker & Simoncelli, 2006

WHAT ARE NATURAL PRIORS LIKE?

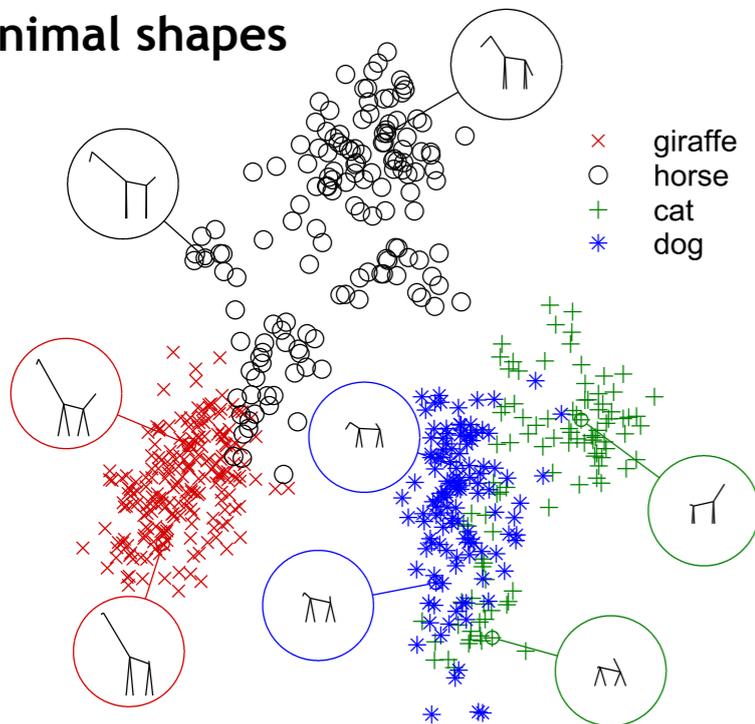
- ▶ adapted to 'natural' statistics
- ▶ pool data across subjects

box office phone waiting times

Griffiths & Tenenbaum, 2006

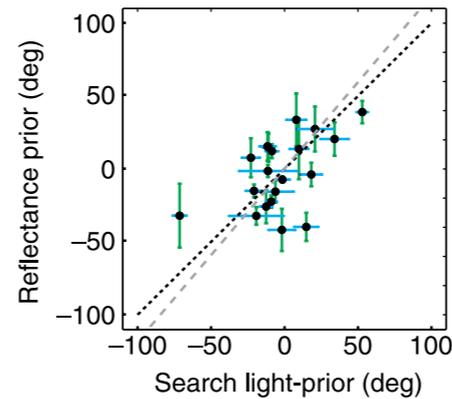
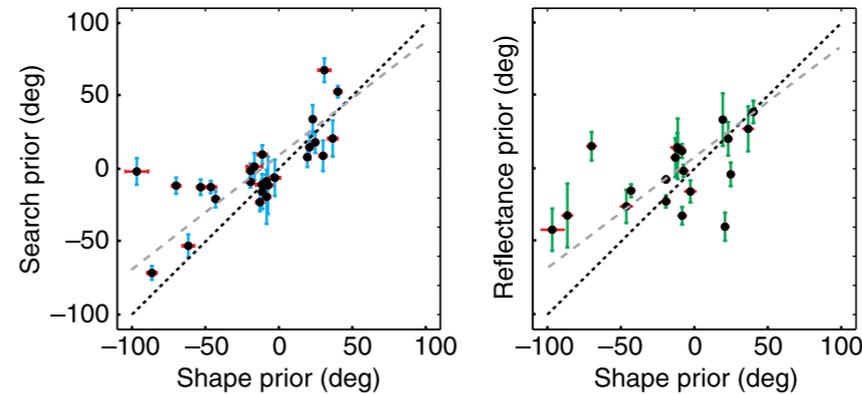
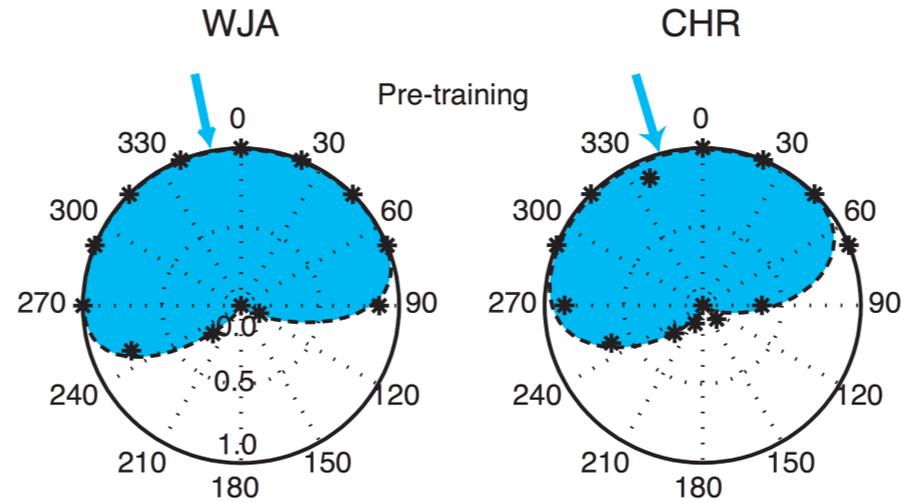


animal shapes



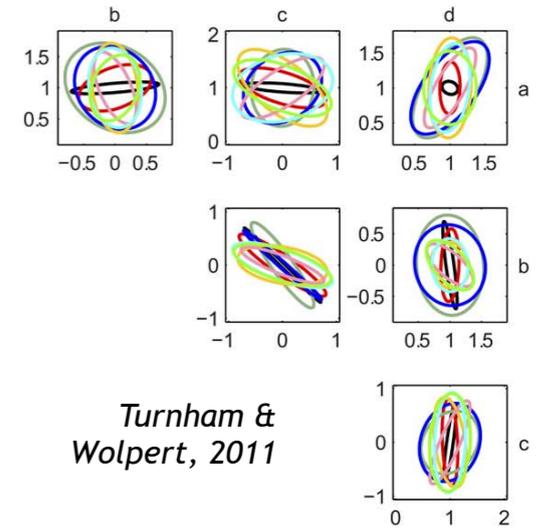
Sanborn & Griffiths, 2008

direction of light

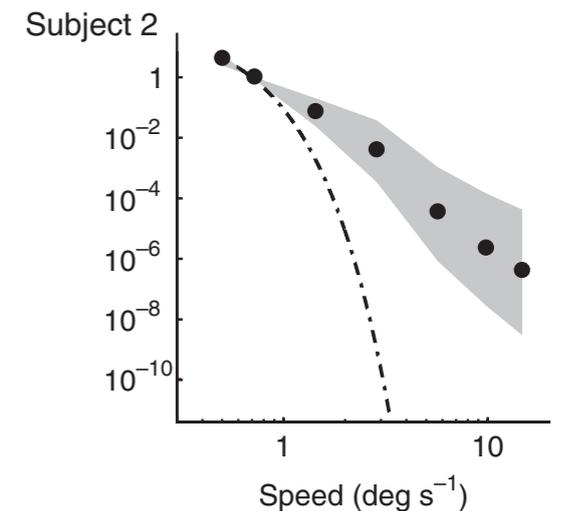
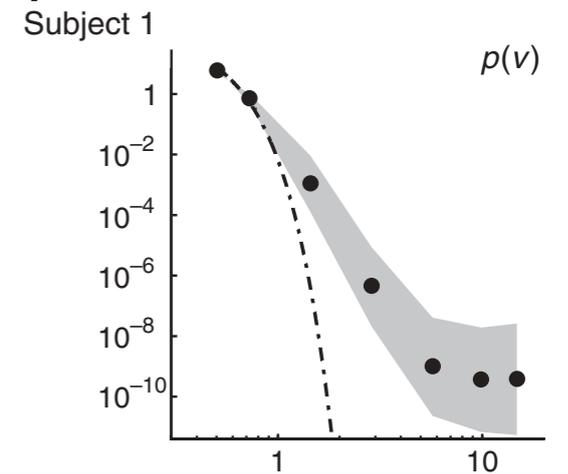


Adams, 2007

visuomotor transformations



speed of visual motion



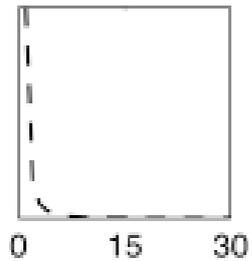
Stocker & Simoncelli, 2006

WHAT ARE NATURAL PRIORS LIKE?

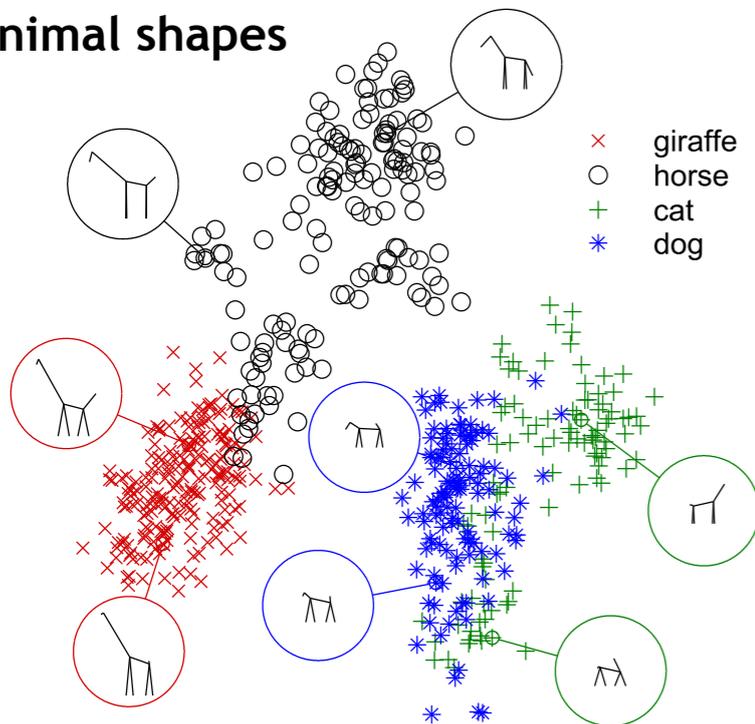
- ▶ adapted to ‘natural’ statistics
- ▶ pool data across subjects
- ▶ simple scalar measures

box office phone waiting times

Griffiths & Tenenbaum, 2006

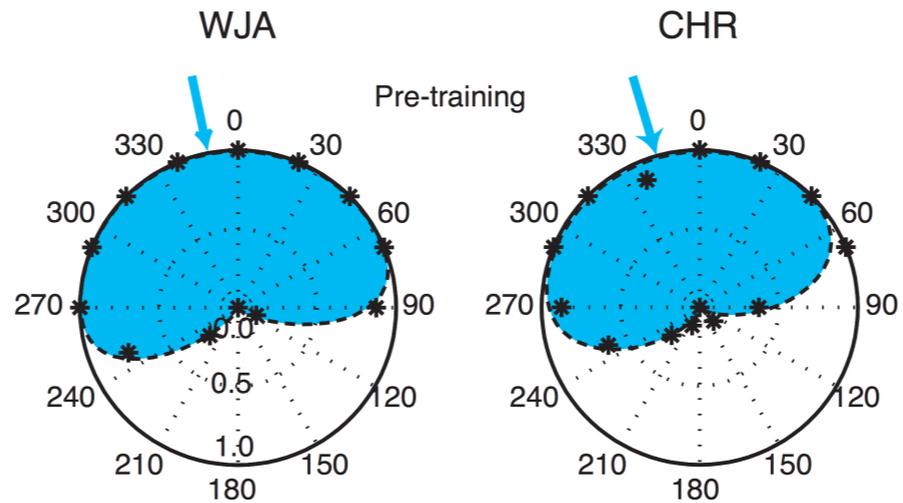


animal shapes

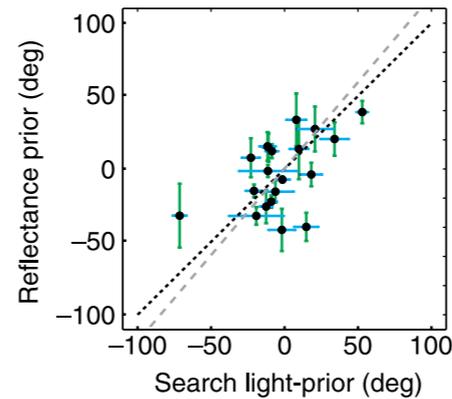
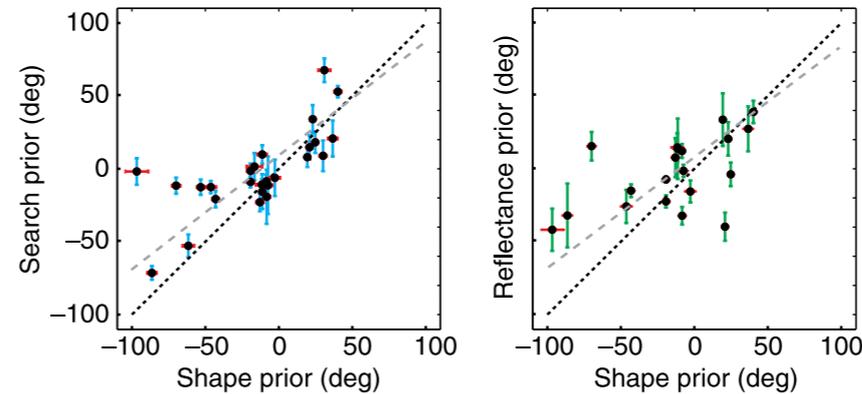


Sanborn & Griffiths, 2008

direction of light

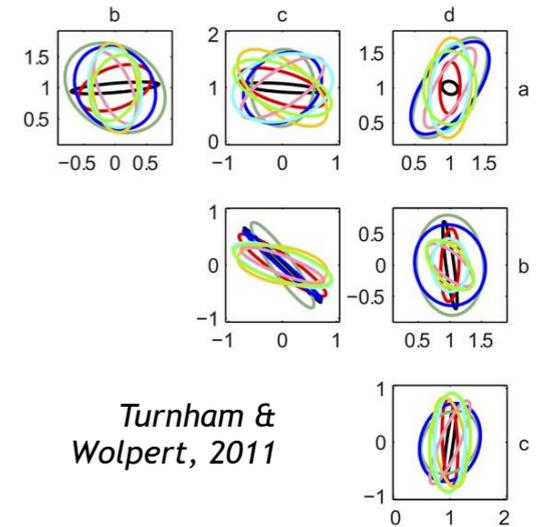


Adams et al, 2004



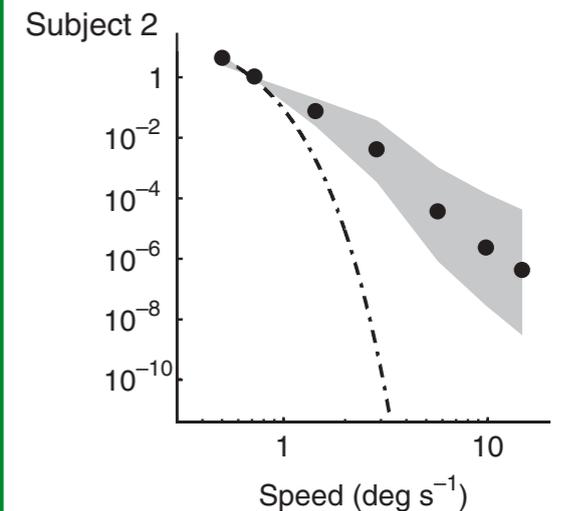
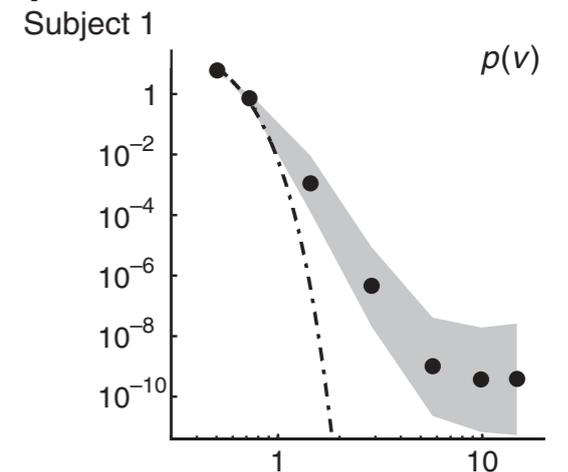
Adams, 2007

visuomotor transformations



Turnham & Wolpert, 2011

speed of visual motion



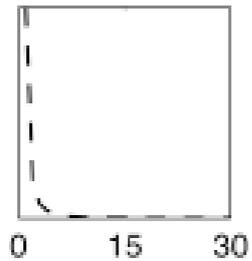
Stocker & Simoncelli, 2006

WHAT ARE NATURAL PRIORS LIKE?

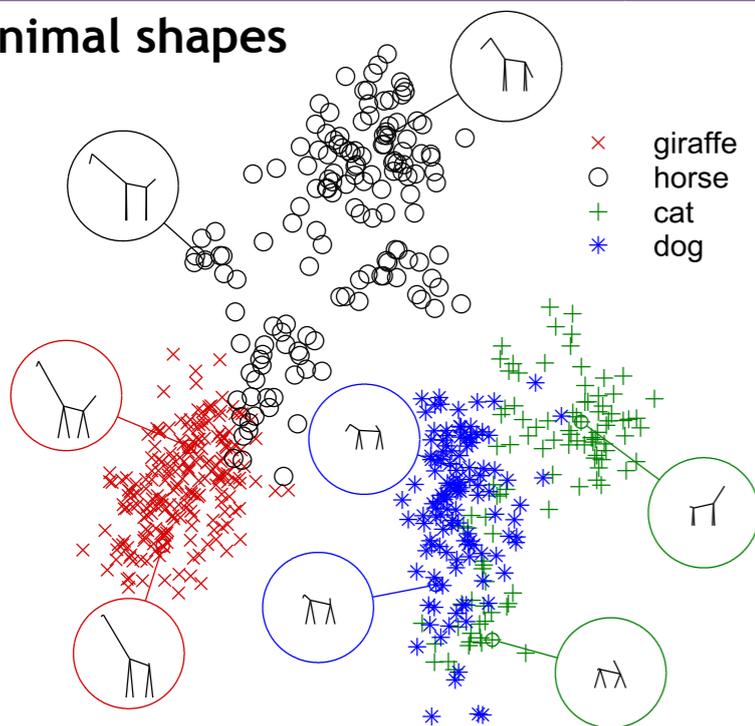
- ▶ adapted to ‘natural’ statistics
- ▶ pool data across subjects
- ▶ simple scalar measures
- ▶ simple low-level psychophysics

box office phone waiting times

Griffiths & Tenenbaum, 2006

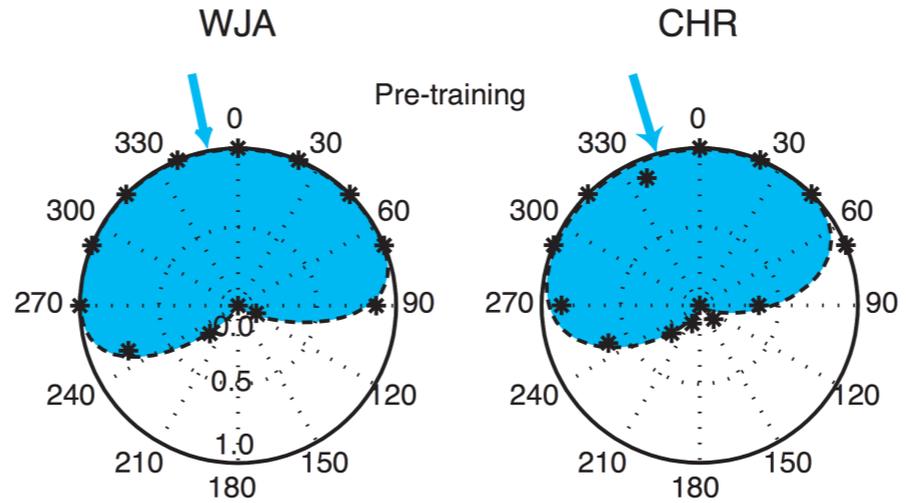


animal shapes

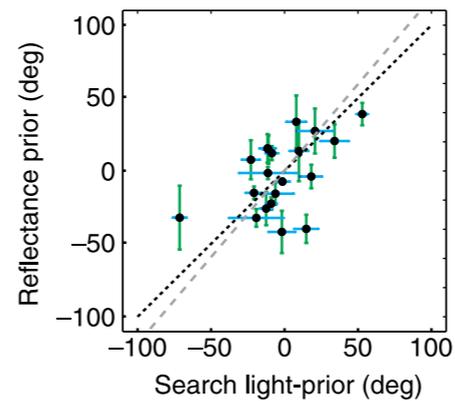
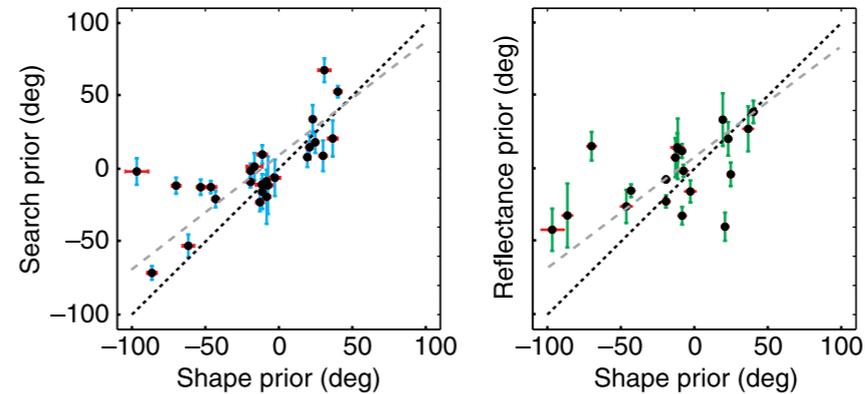


Sanborn & Griffiths, 2008

direction of light

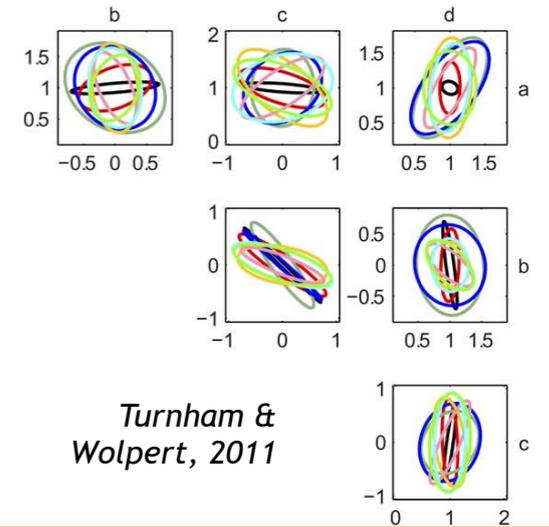


Adams et al, 2004



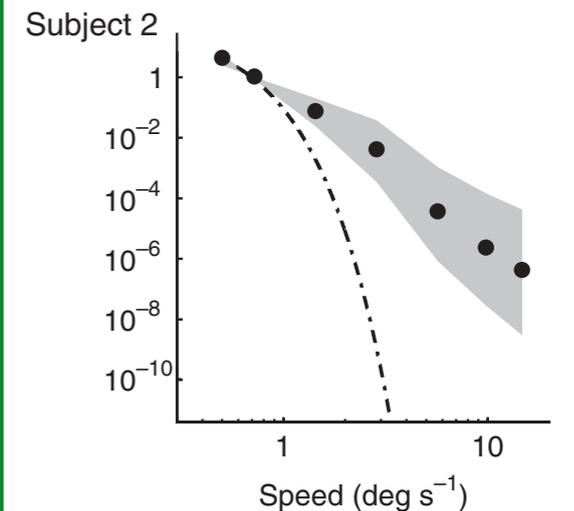
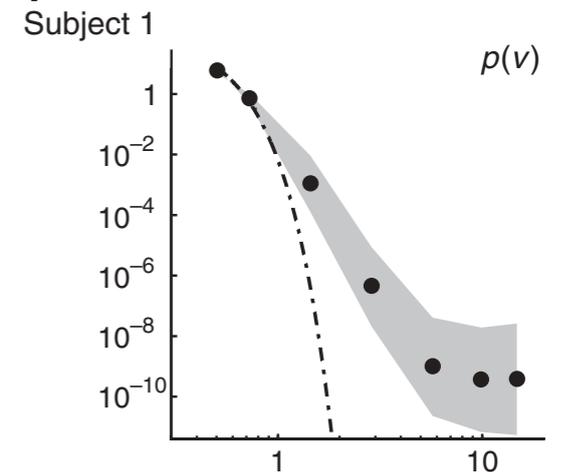
Adams, 2007

visuomotor transformations



Turnham & Wolpert, 2011

speed of visual motion



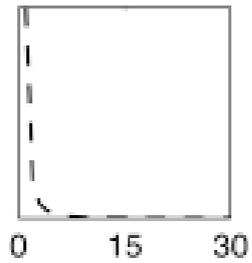
Stocker & Simoncelli, 2006

WHAT ARE NATURAL PRIORS LIKE?

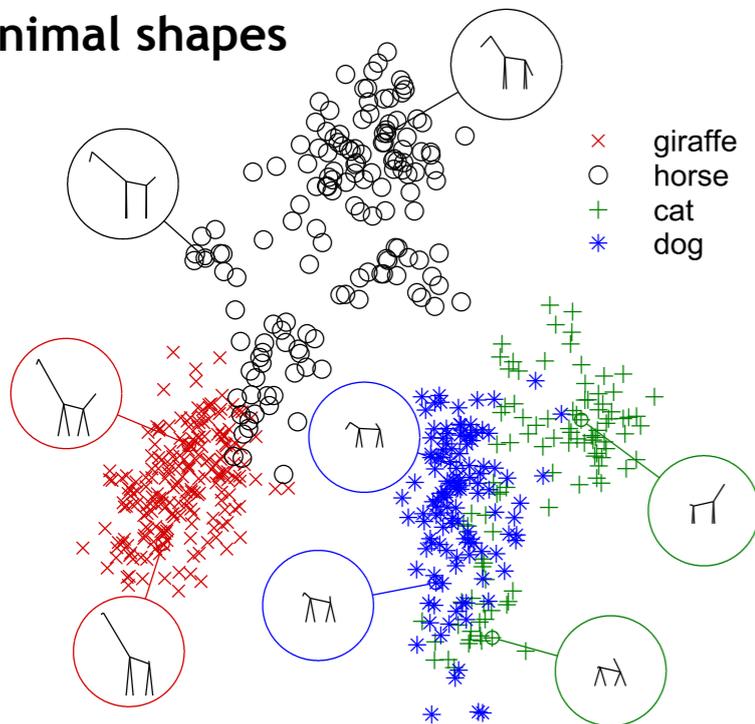
- ▶ adapted to 'natural' statistics
- ▶ pool data across subjects
- ▶ simple scalar measures
- ▶ simple low-level psychophysics
- ▶ limited to one task

box office phone waiting times

Griffiths & Tenenbaum, 2006

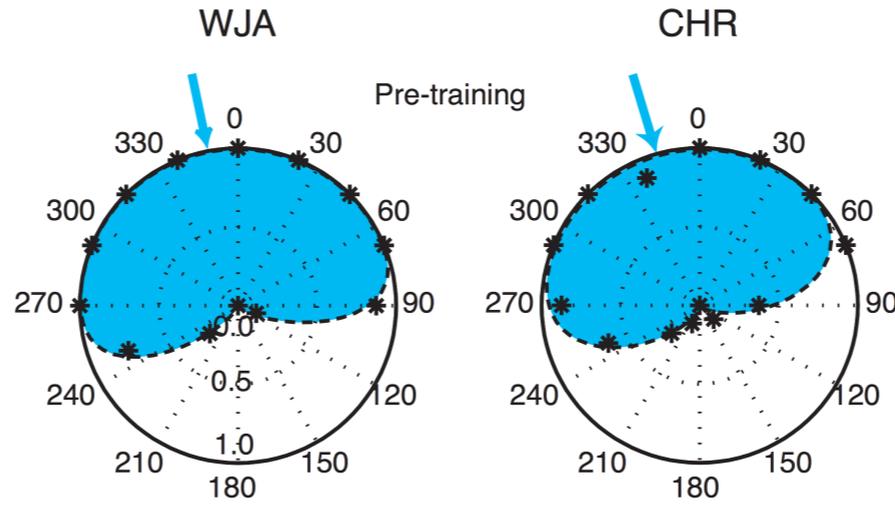


animal shapes

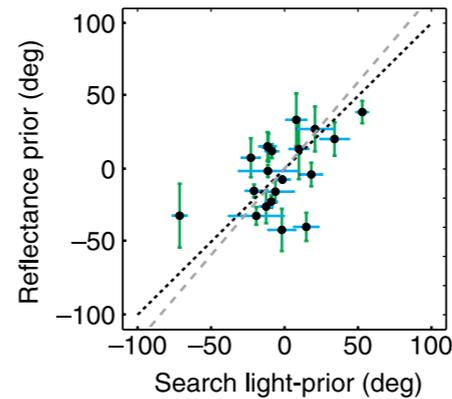
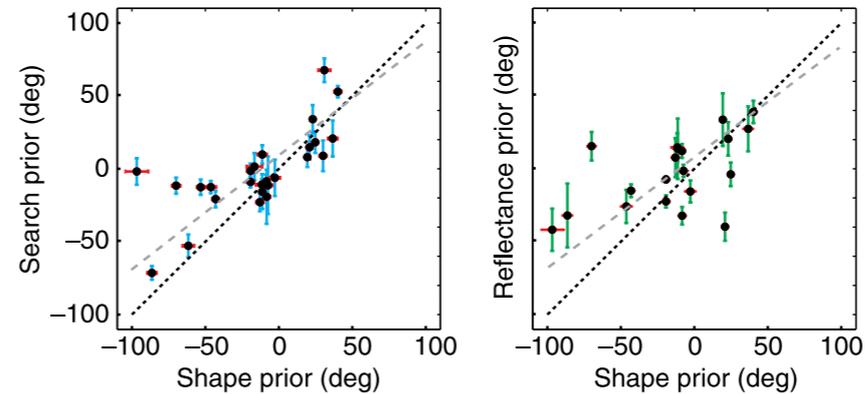


Sanborn & Griffiths, 2008

direction of light

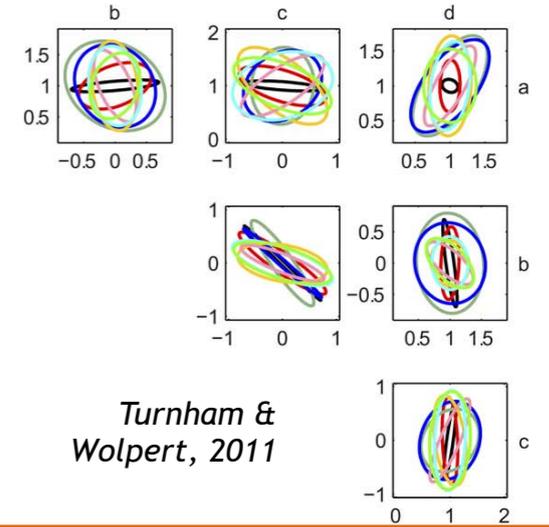


Adams et al, 2004



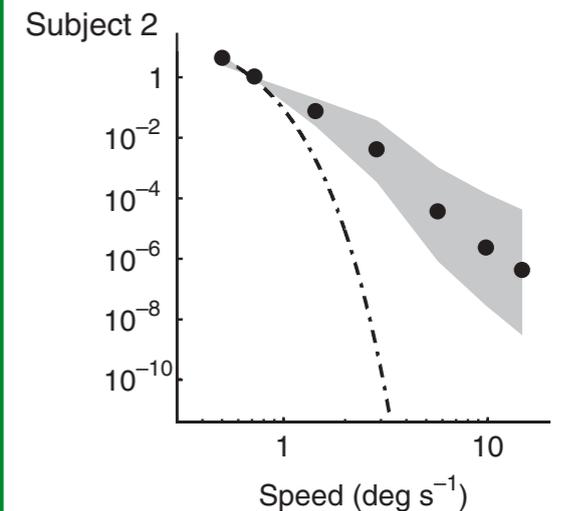
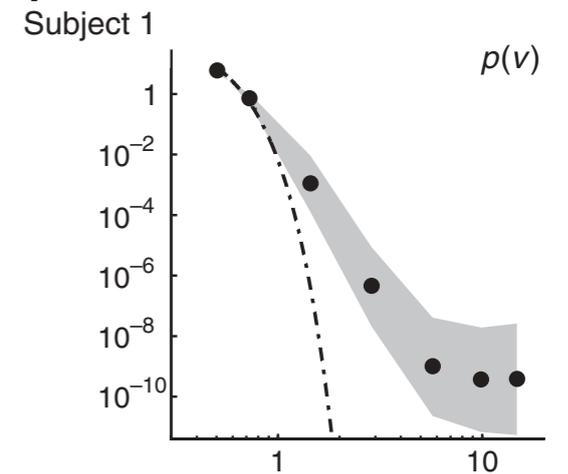
Adams, 2007

visuomotor transformations



Turnham & Wolpert, 2011

speed of visual motion



Stocker & Simoncelli, 2006

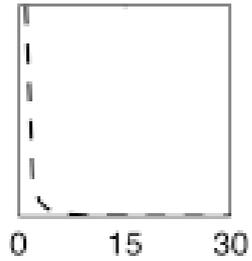
WHAT ARE NATURAL PRIORS LIKE?

- ▶ adapted to 'natural' statistics
- ▶ pool data across subjects
- ▶ simple scalar measures
- ▶ simple low-level psych
- ▶ limited to one task

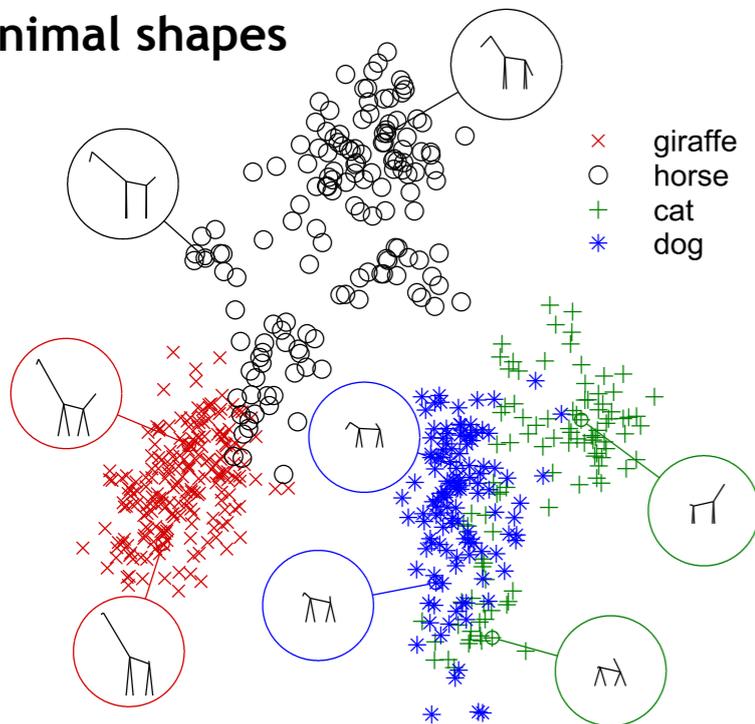
prior = bunch of free parameters → overfitting?

box office phone waiting times

Griffiths & Tenenbaum, 2006

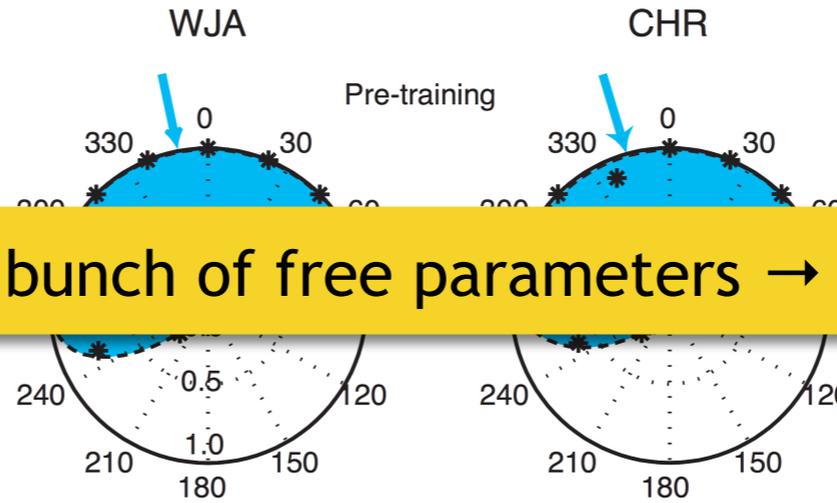


animal shapes

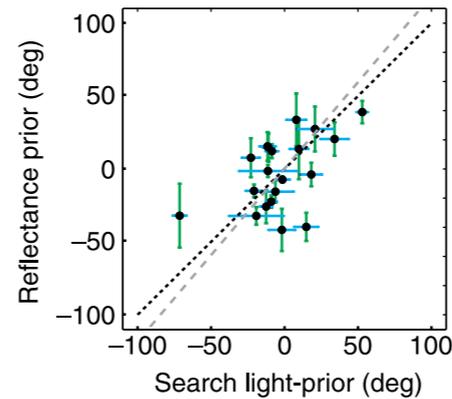
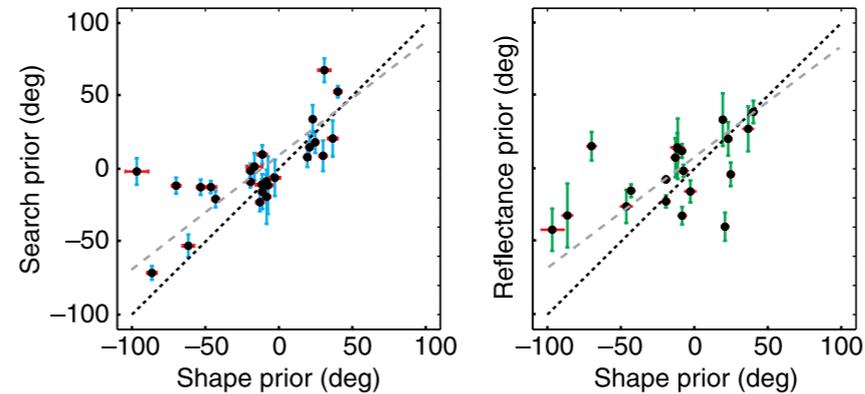


Sanborn & Griffiths, 2008

direction of light

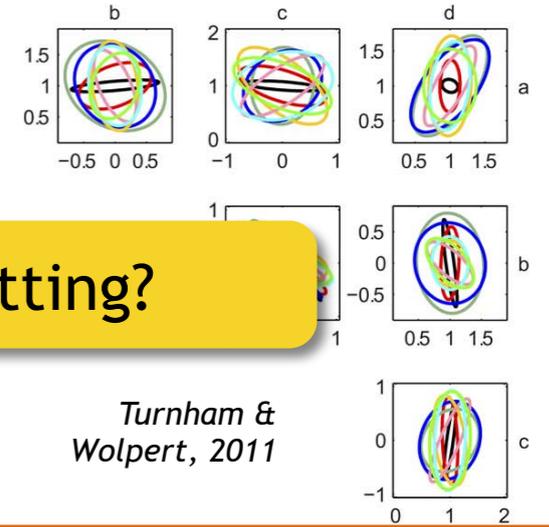


Adams et al, 2004



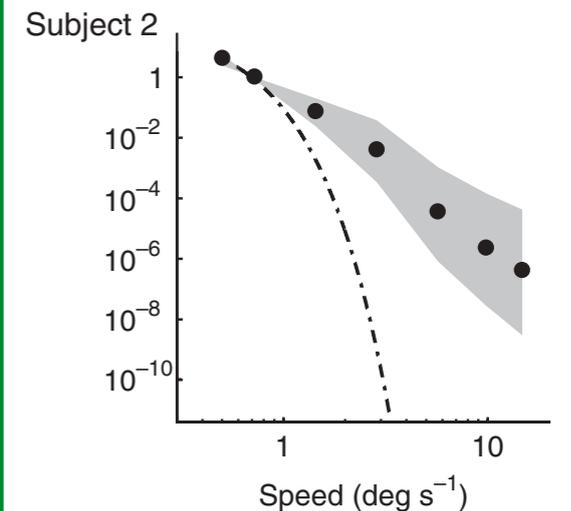
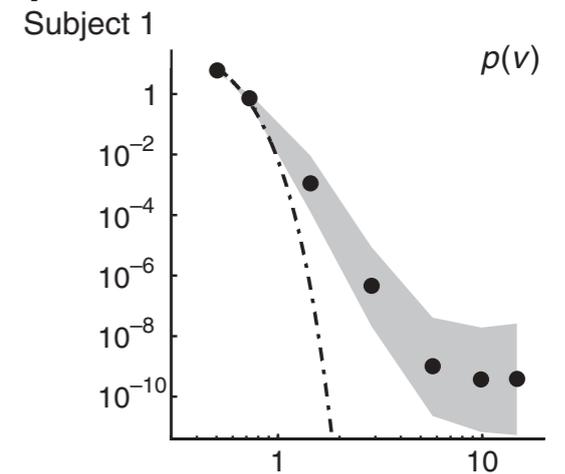
Adams, 2007

visuomotor transformations



Turnham & Wolpert, 2011

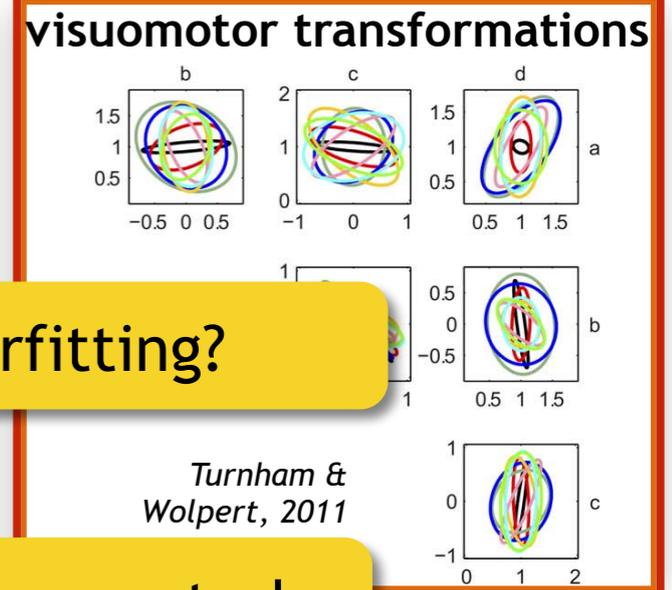
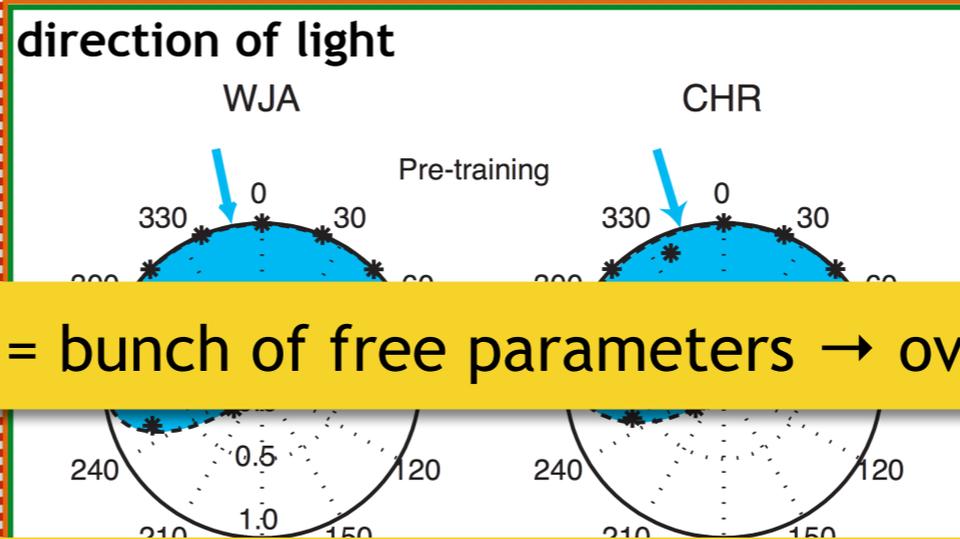
speed of visual motion



Stocker & Simoncelli, 2006

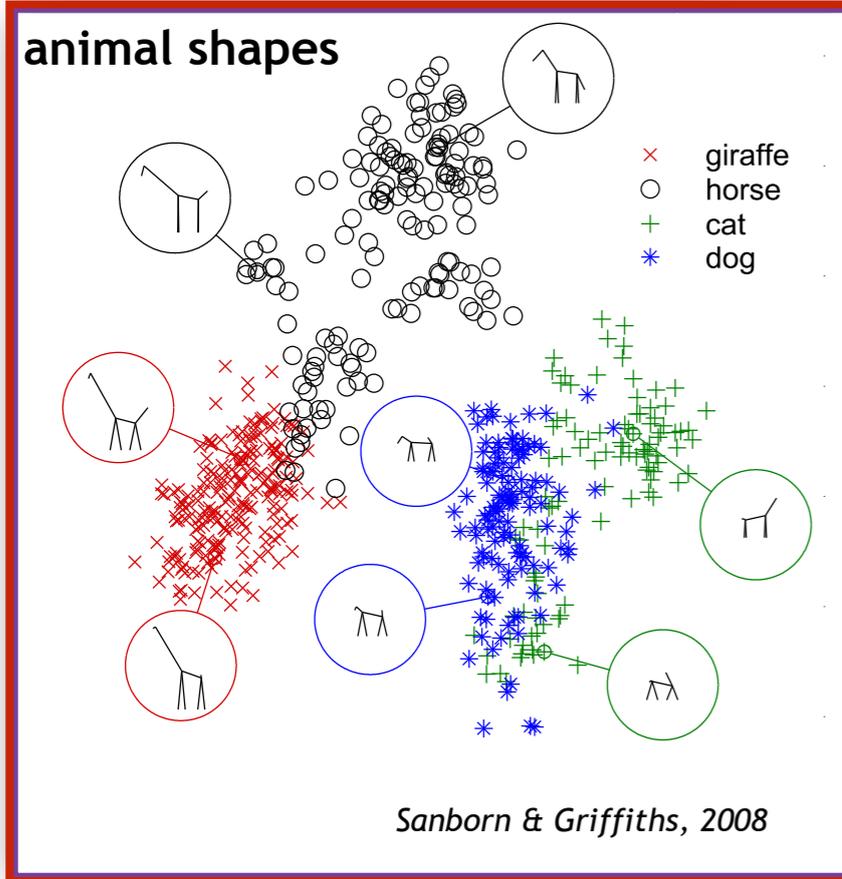
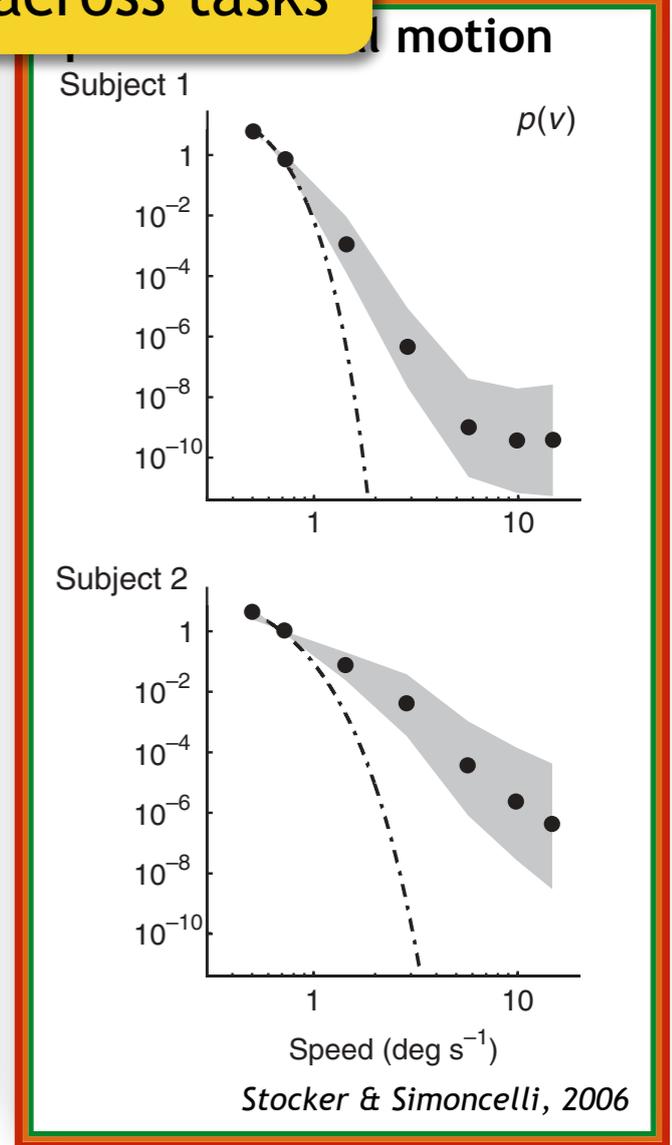
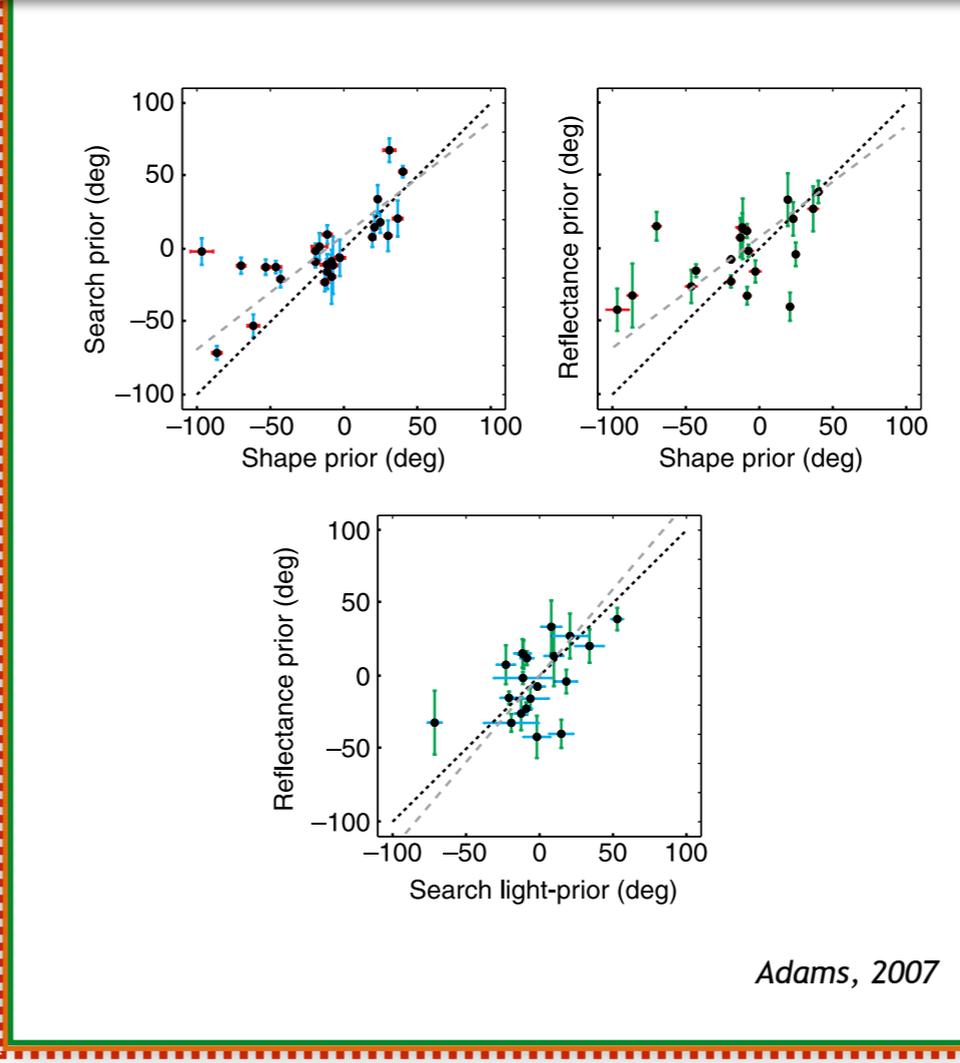
WHAT ARE NATURAL PRIORS LIKE?

- ▶ adapted to 'natural' statistics
- ▶ pool data across subjects
- ▶ simple scalar measures
- ▶ simple low-level psych
- ▶ limited to one task



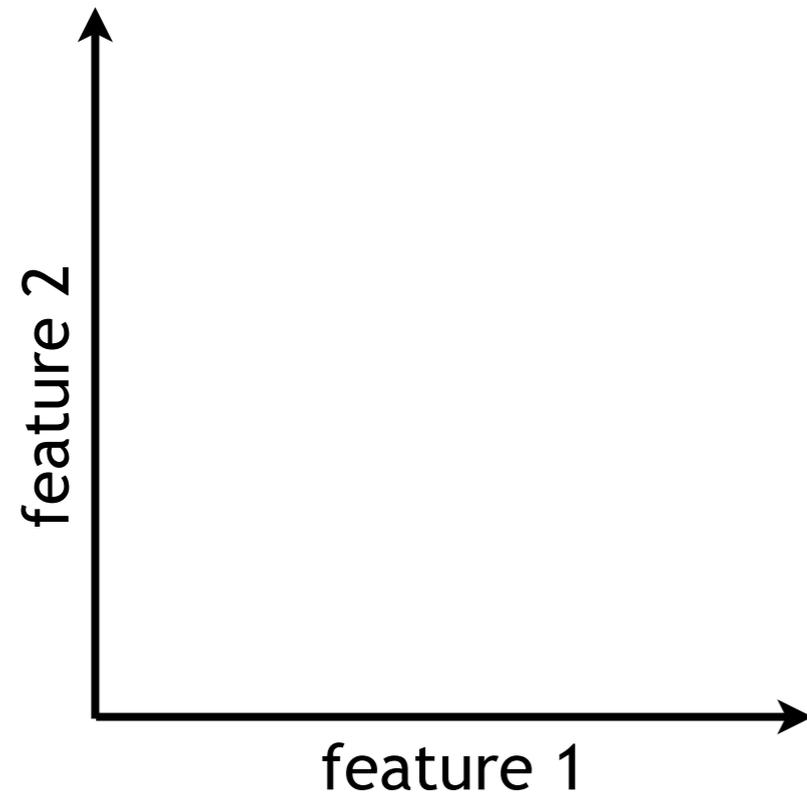
prior = bunch of free parameters → overfitting?

advantage of having a prior: generalisation across tasks



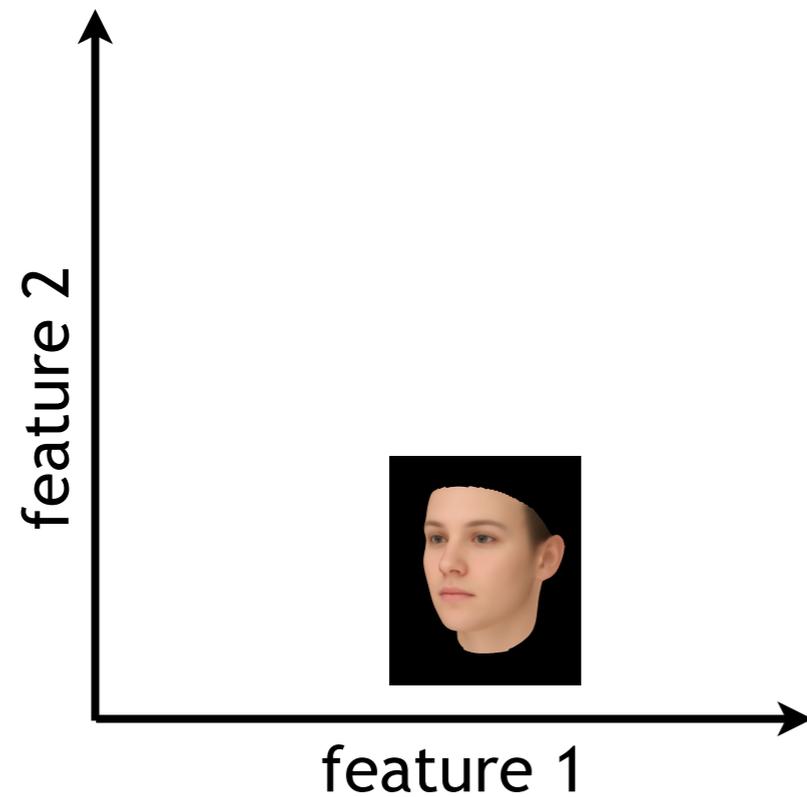
DISCRIMINATIVE VS. GENERATIVE LEARNING

DISCRIMINATIVE



DISCRIMINATIVE VS. GENERATIVE LEARNING

DISCRIMINATIVE



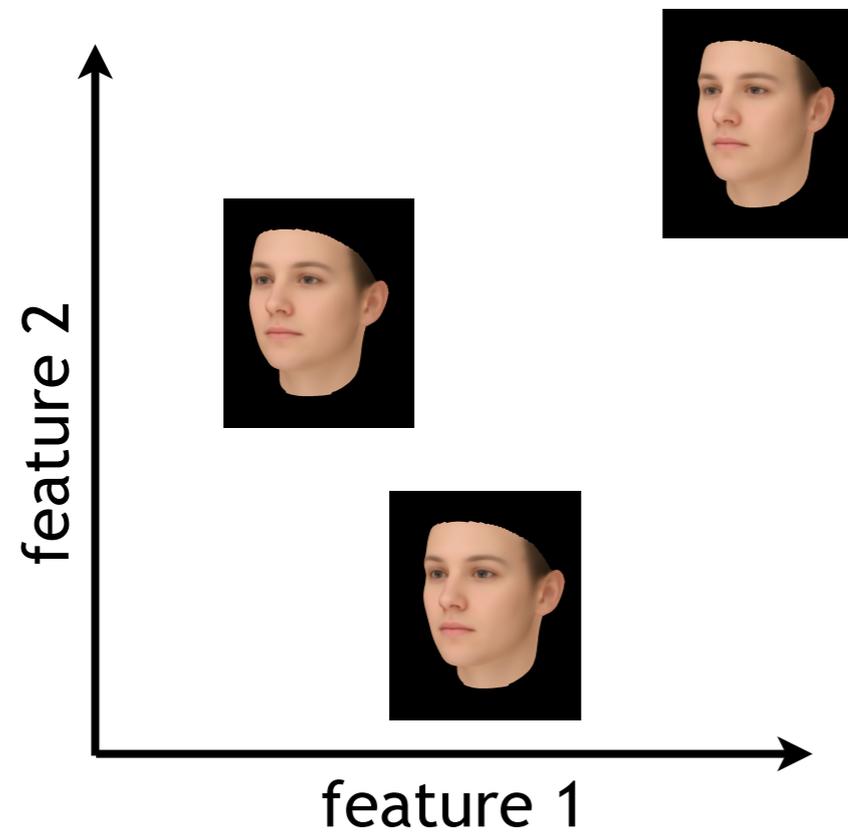
DISCRIMINATIVE VS. GENERATIVE LEARNING

DISCRIMINATIVE



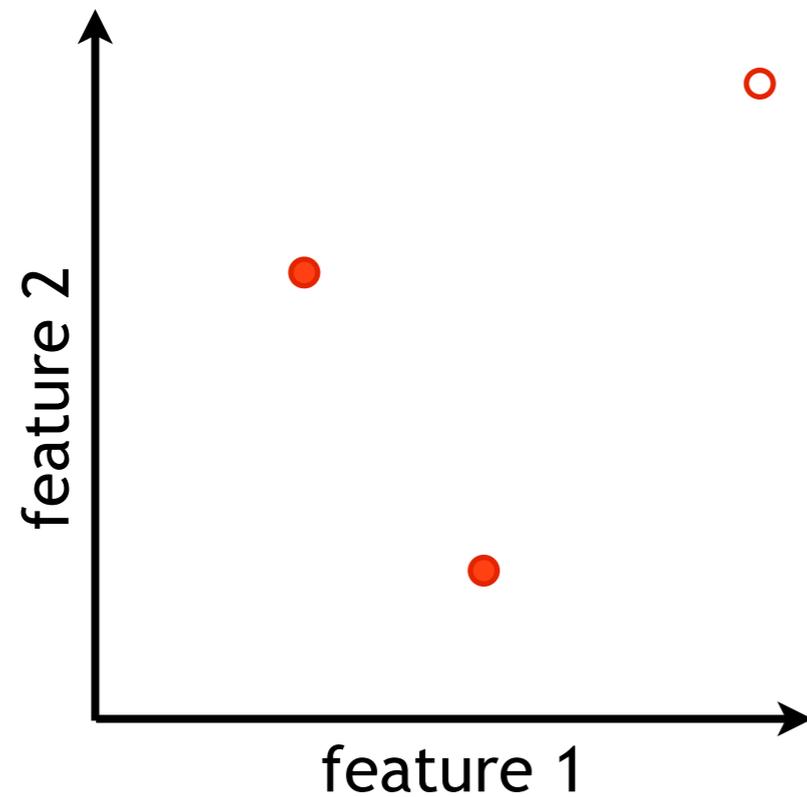
DISCRIMINATIVE VS. GENERATIVE LEARNING

DISCRIMINATIVE



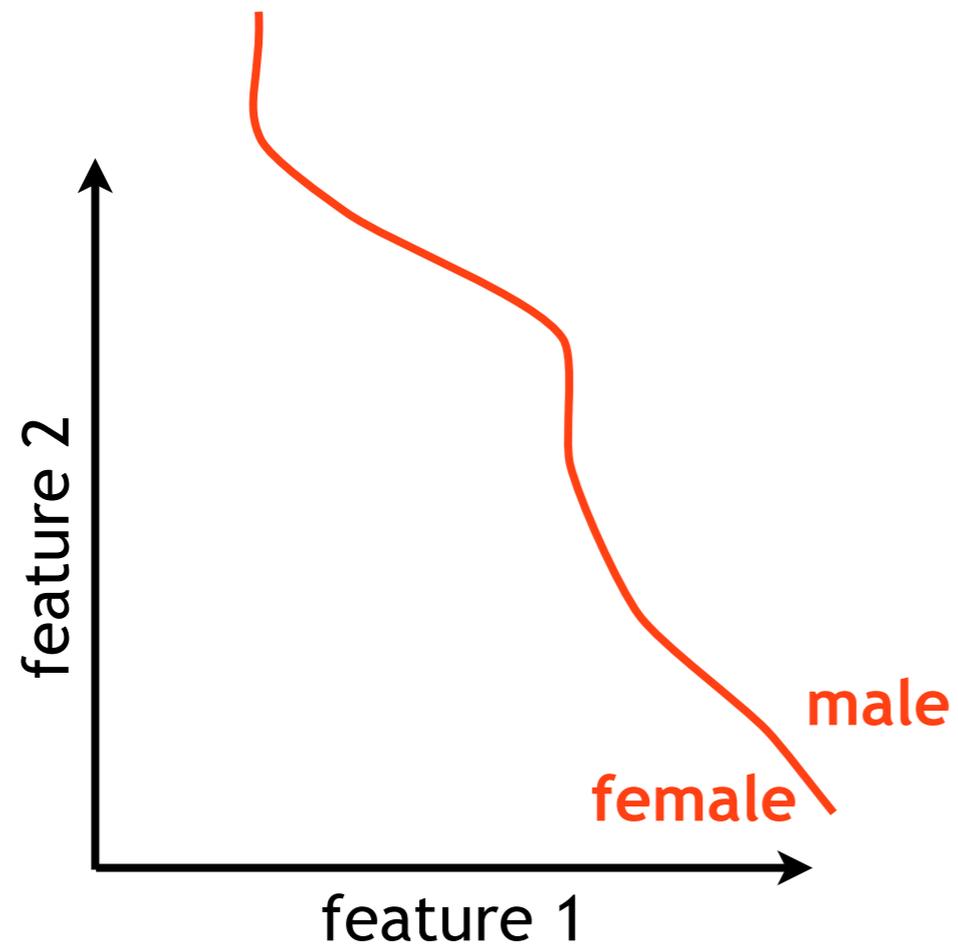
DISCRIMINATIVE VS. GENERATIVE LEARNING

DISCRIMINATIVE



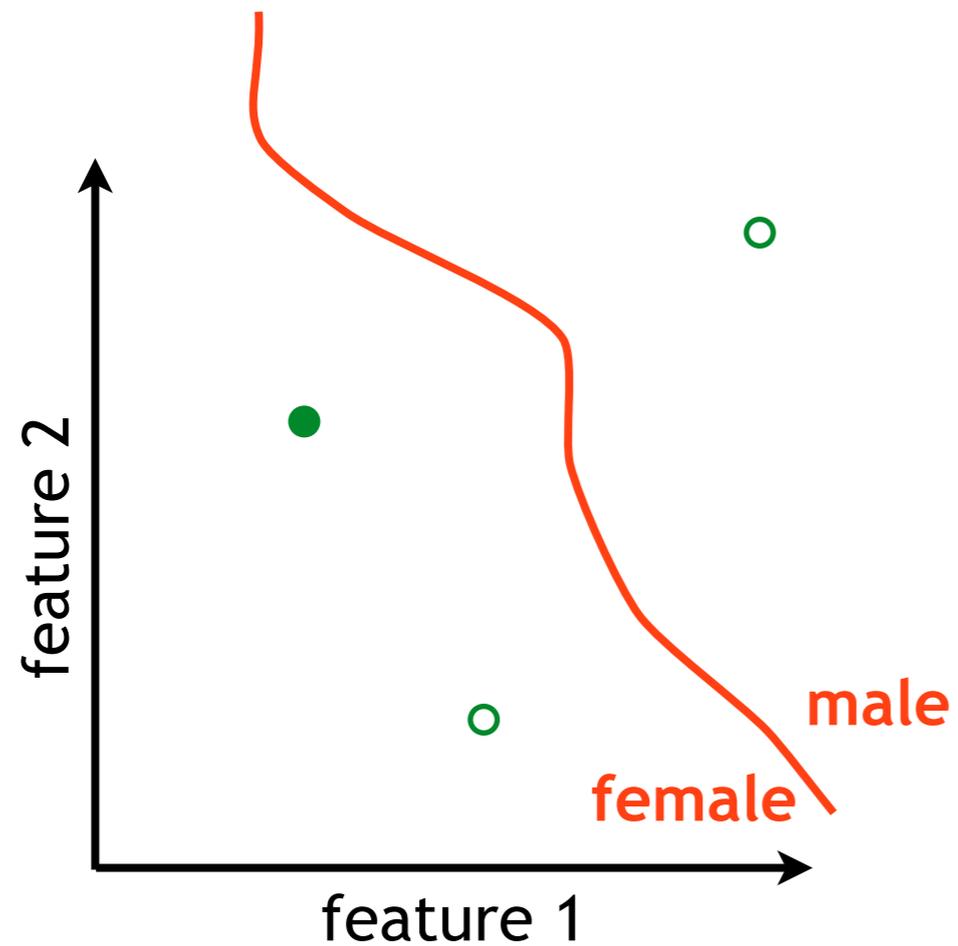
DISCRIMINATIVE VS. GENERATIVE LEARNING

DISCRIMINATIVE



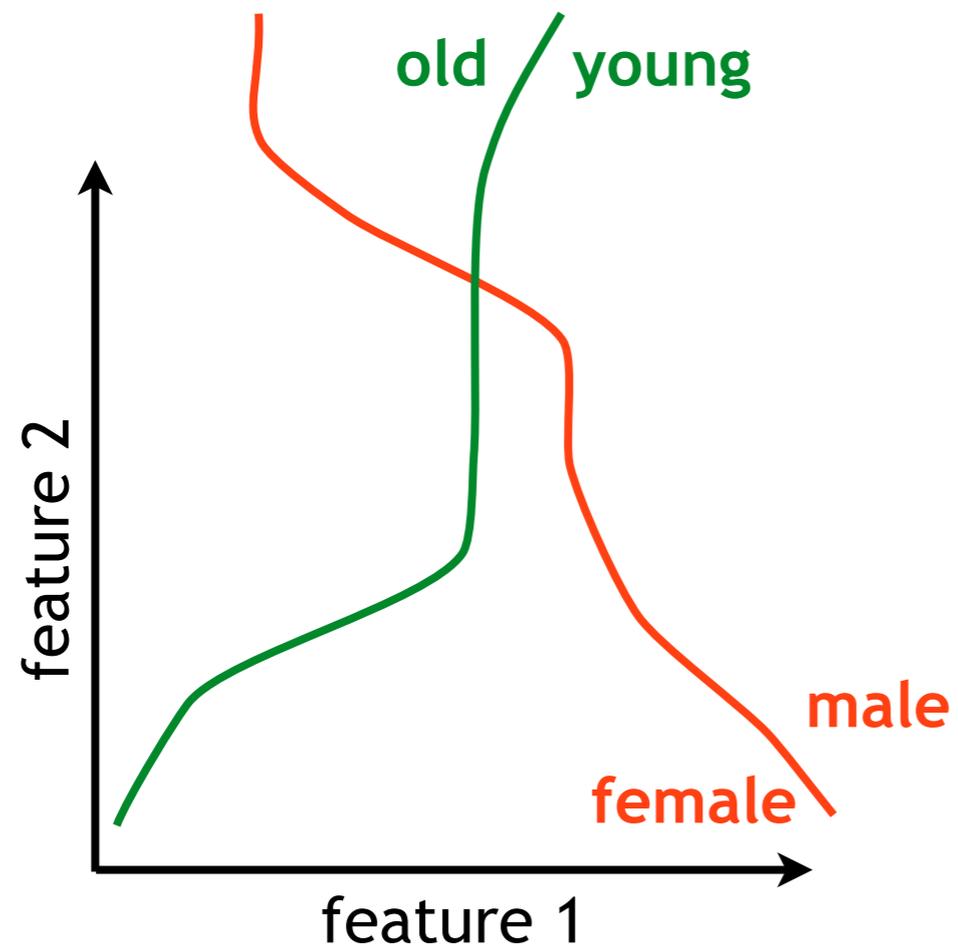
DISCRIMINATIVE VS. GENERATIVE LEARNING

DISCRIMINATIVE



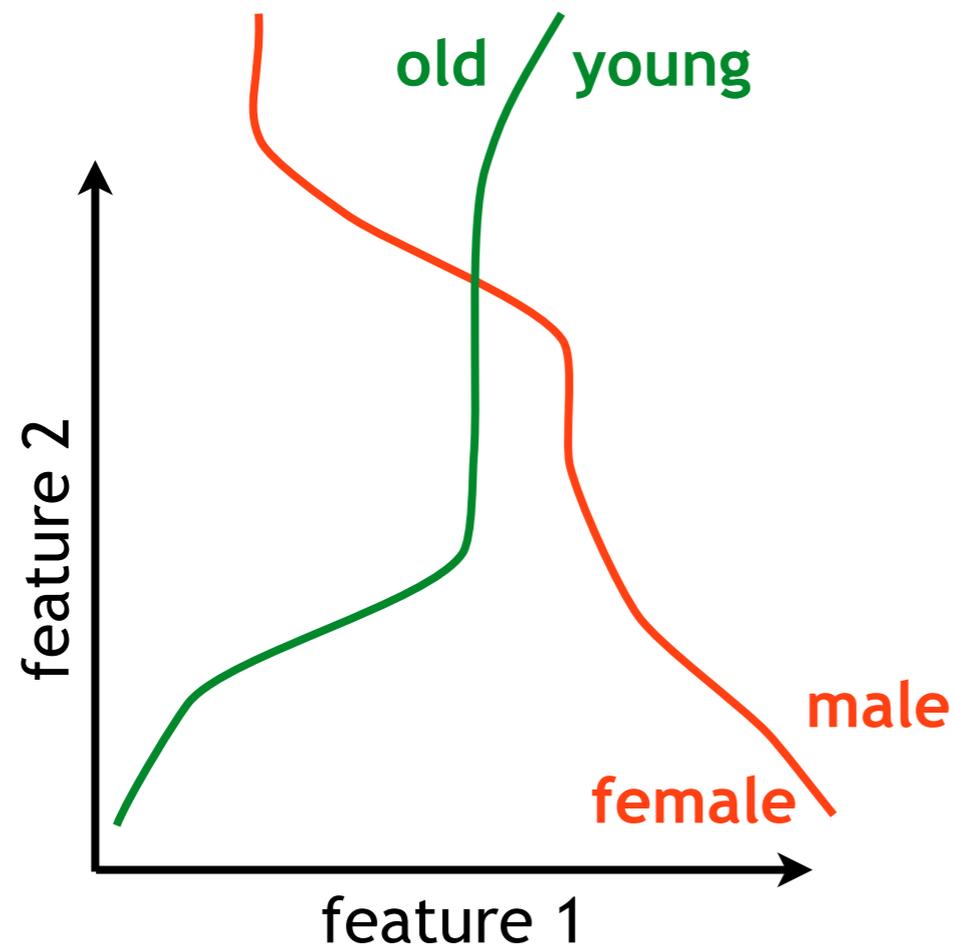
DISCRIMINATIVE VS. GENERATIVE LEARNING

DISCRIMINATIVE

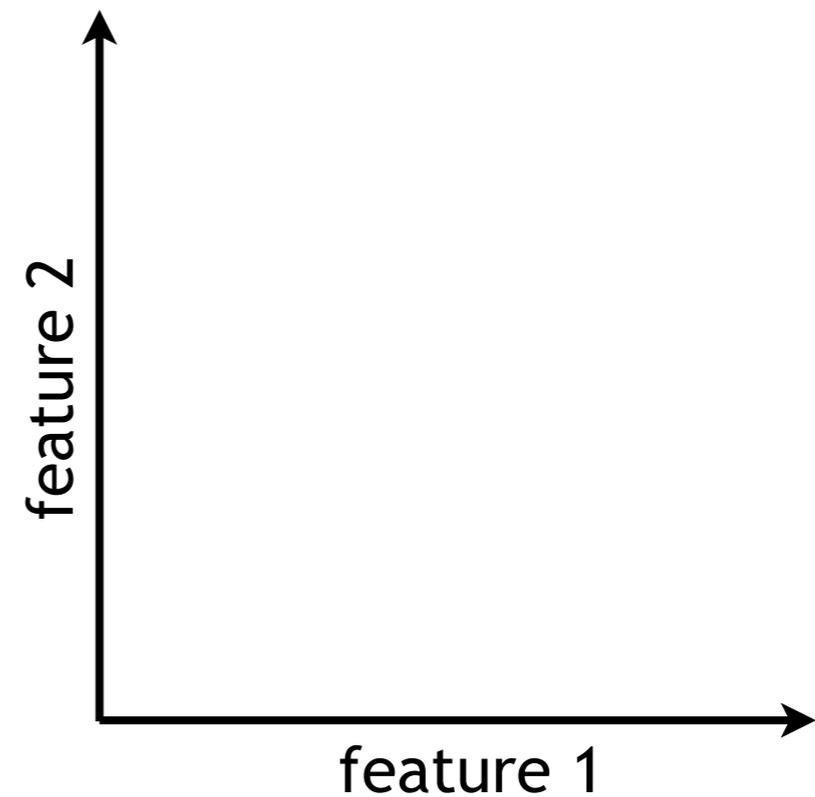


DISCRIMINATIVE VS. GENERATIVE LEARNING

DISCRIMINATIVE

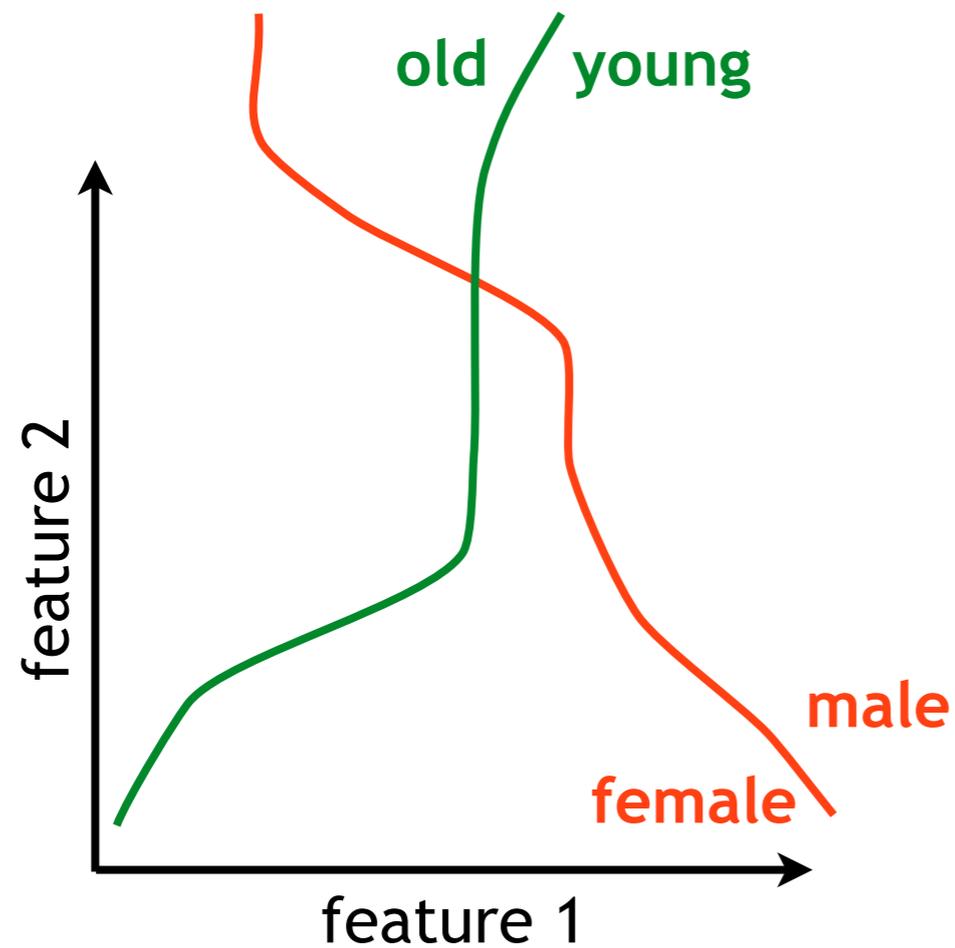


GENERATIVE

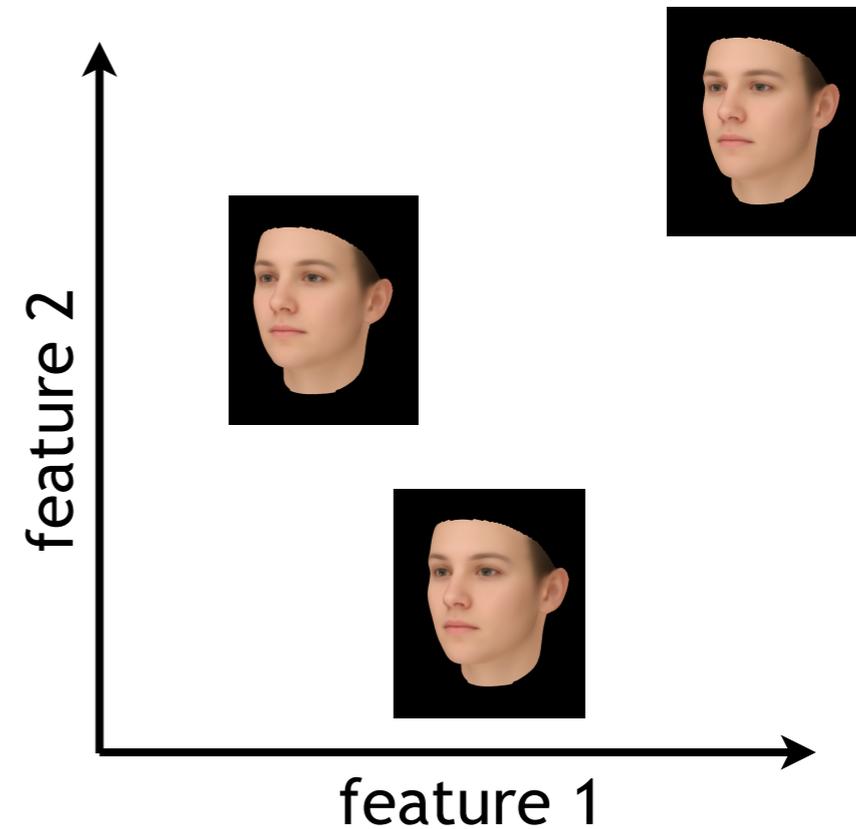


DISCRIMINATIVE VS. GENERATIVE LEARNING

DISCRIMINATIVE

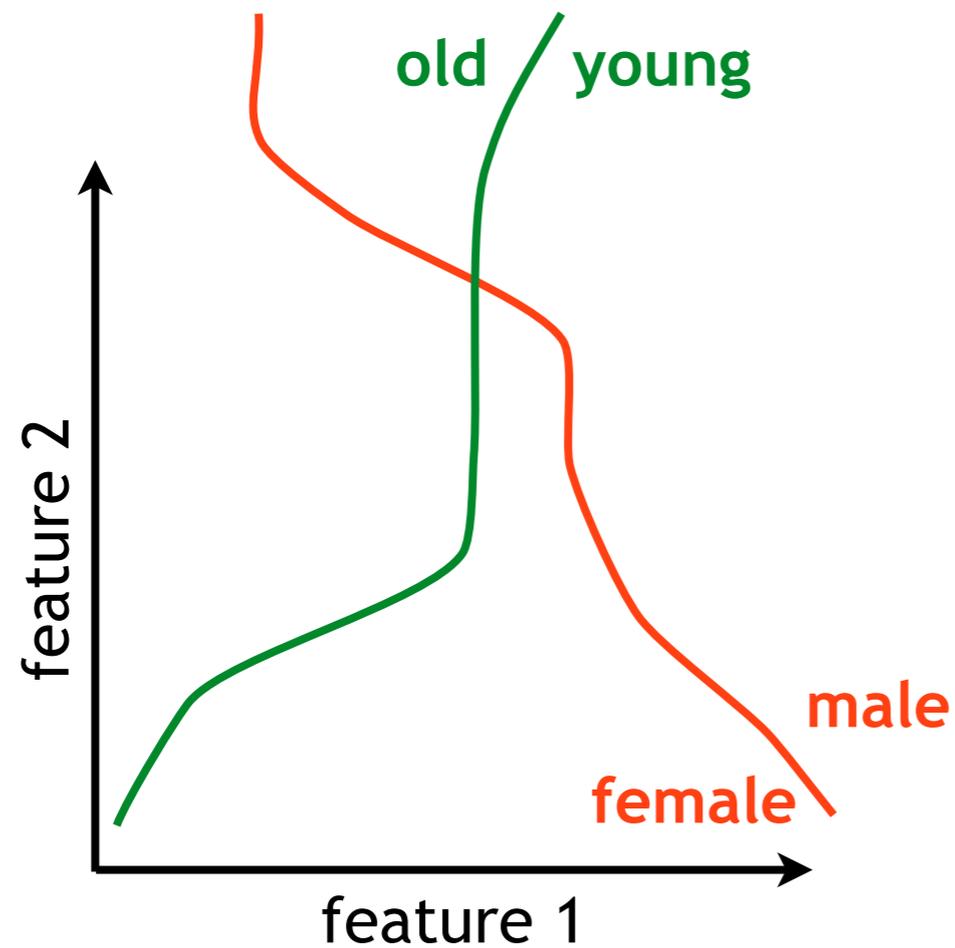


GENERATIVE

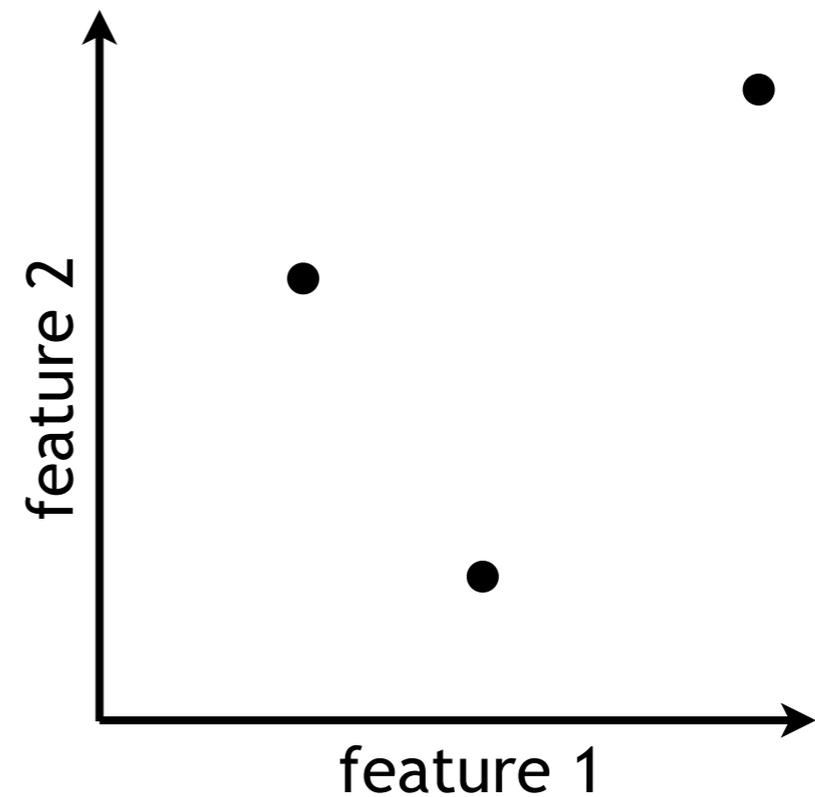


DISCRIMINATIVE VS. GENERATIVE LEARNING

DISCRIMINATIVE

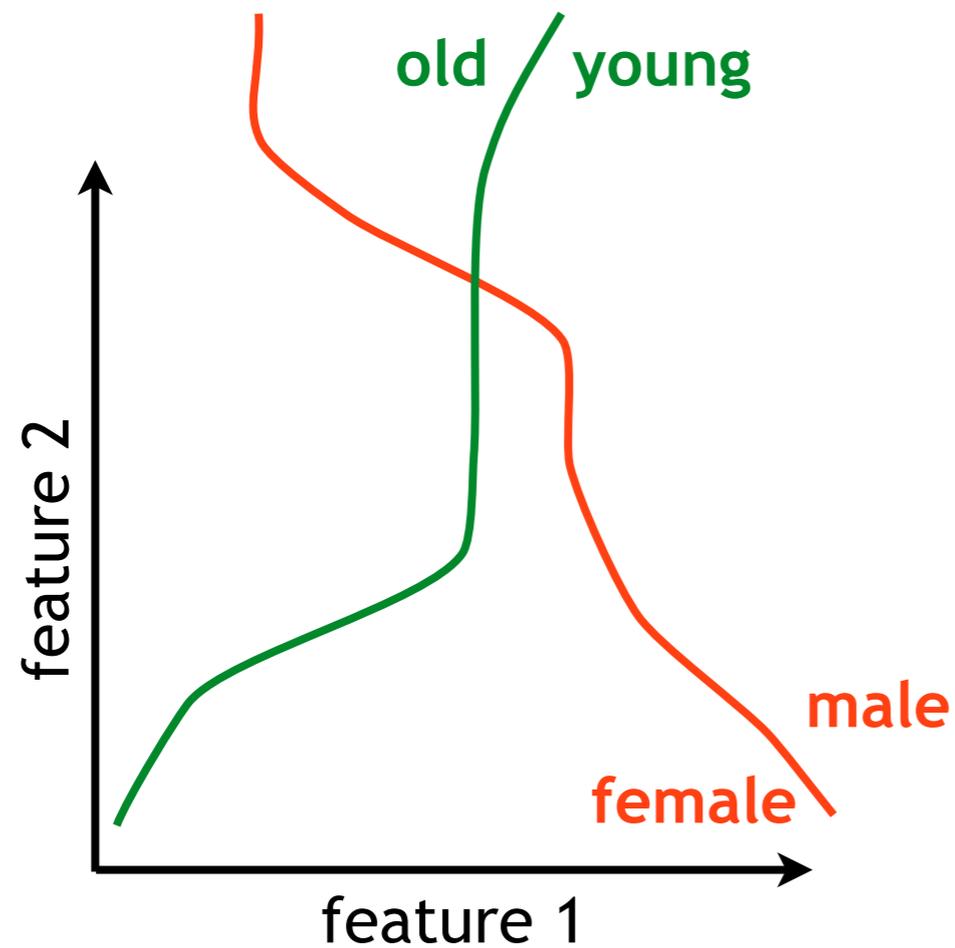


GENERATIVE

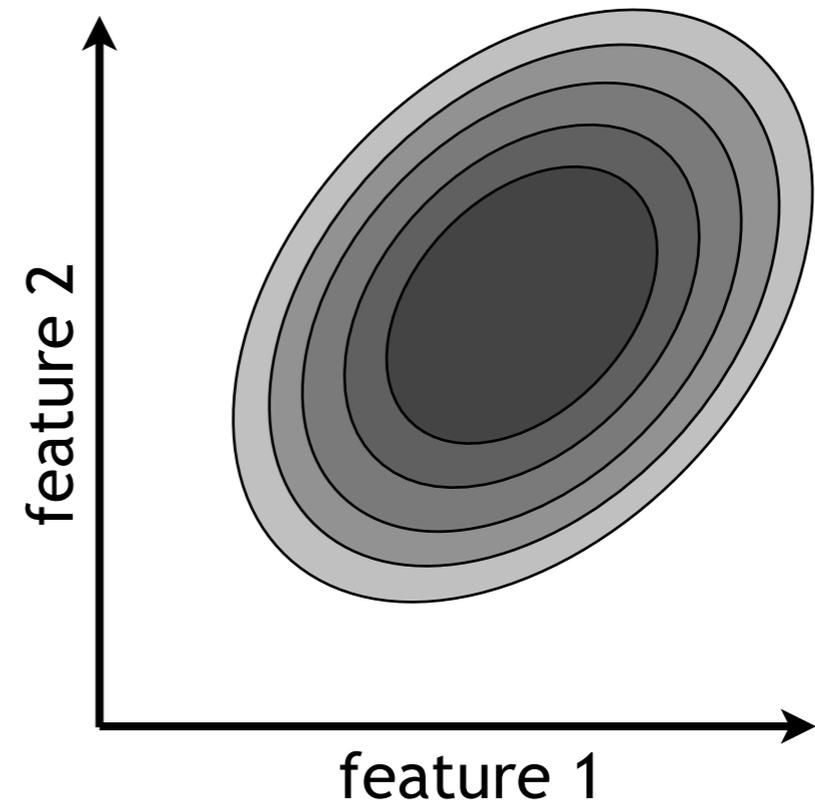


DISCRIMINATIVE VS. GENERATIVE LEARNING

DISCRIMINATIVE



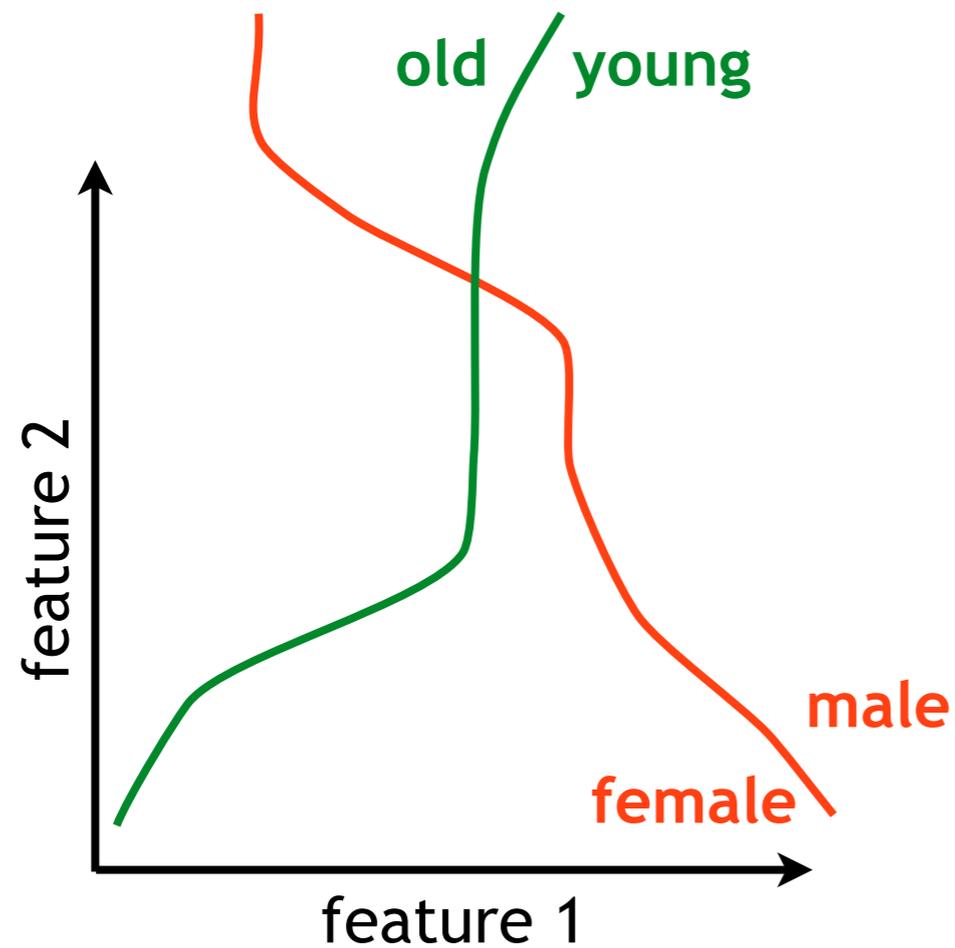
GENERATIVE



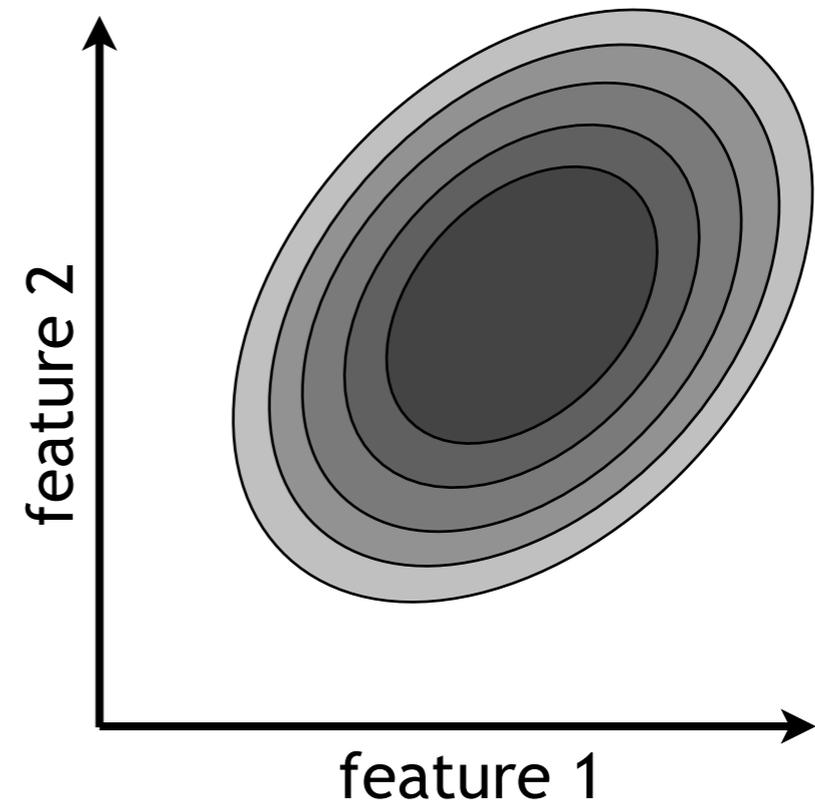
DISCRIMINATIVE VS. GENERATIVE LEARNING

DISCRIMINATIVE

task-dependent representation



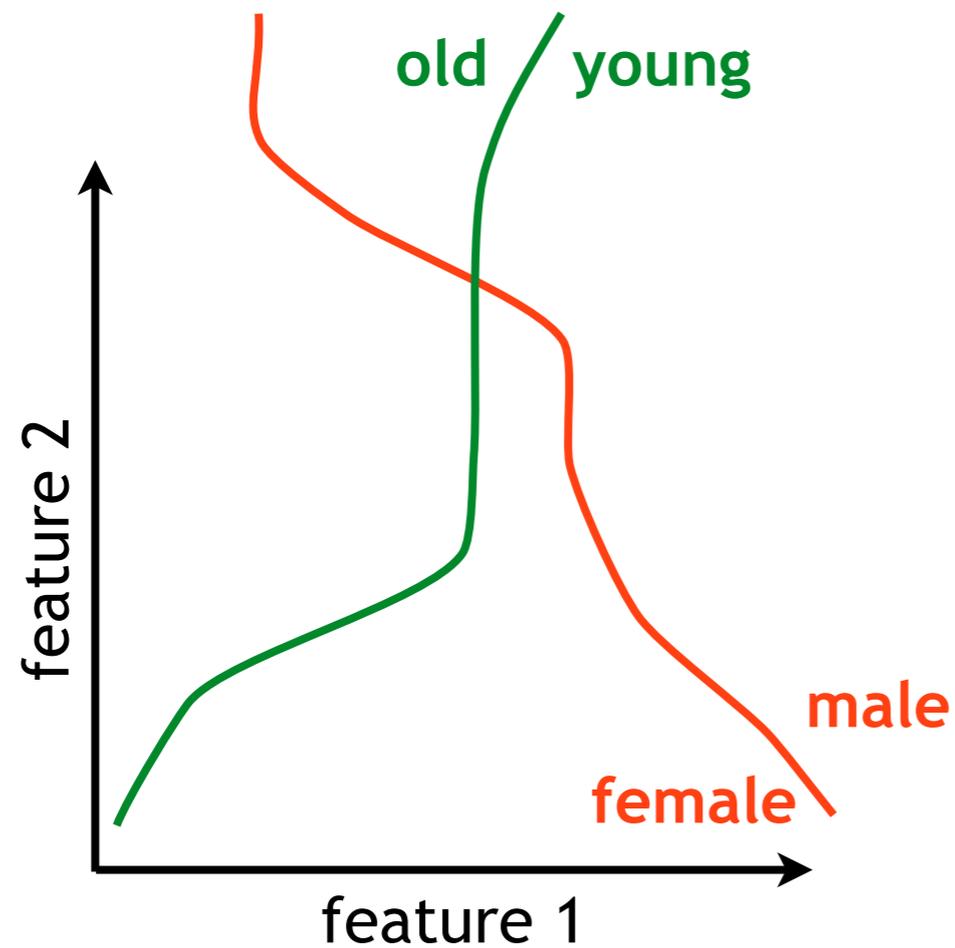
GENERATIVE



DISCRIMINATIVE VS. GENERATIVE LEARNING

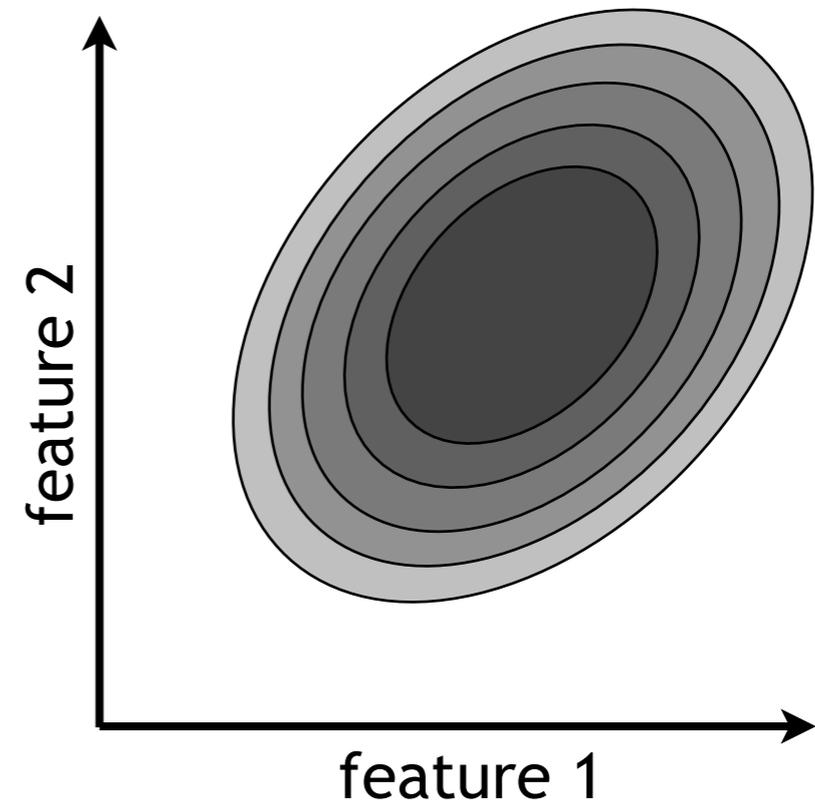
DISCRIMINATIVE

task-dependent representation



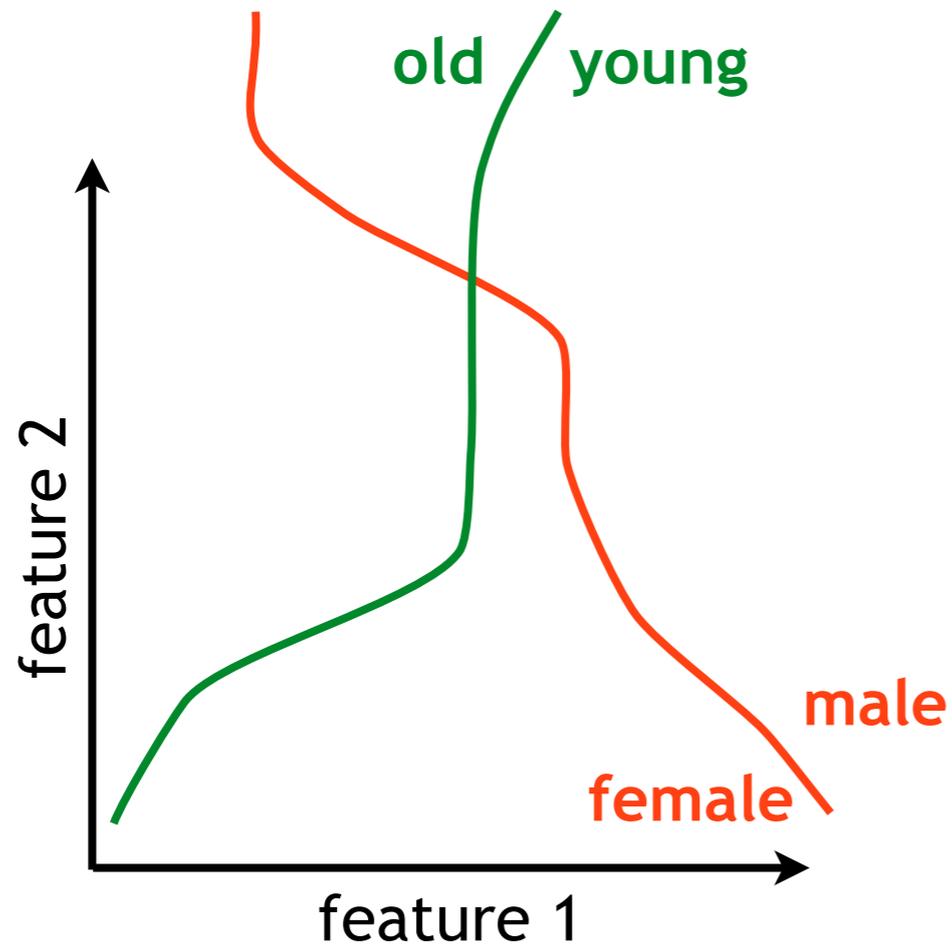
GENERATIVE

task-independent representation

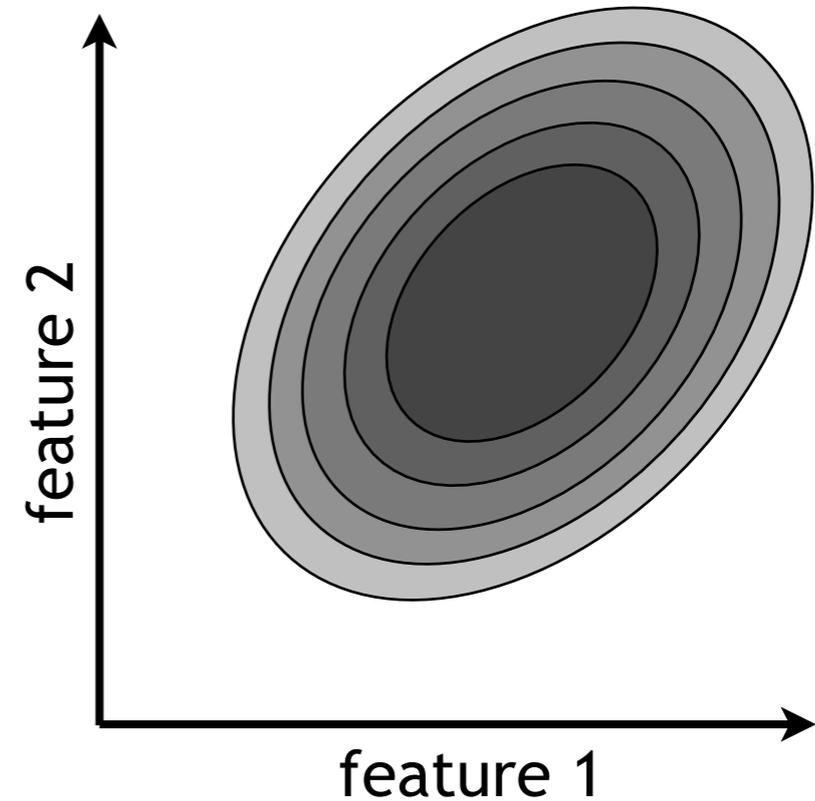


DISCRIMINATIVE VS. GENERATIVE LEARNING

DISCRIMINATIVE
task-dependent representation

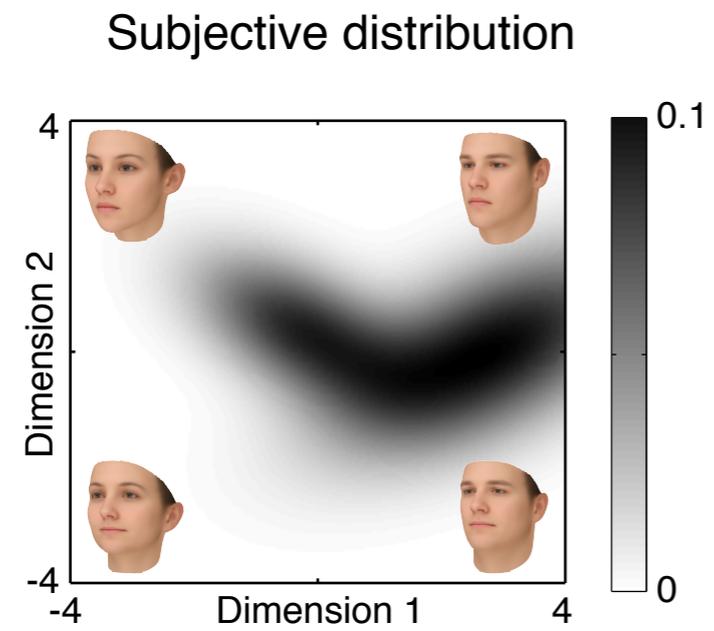


GENERATIVE
task-independent representation

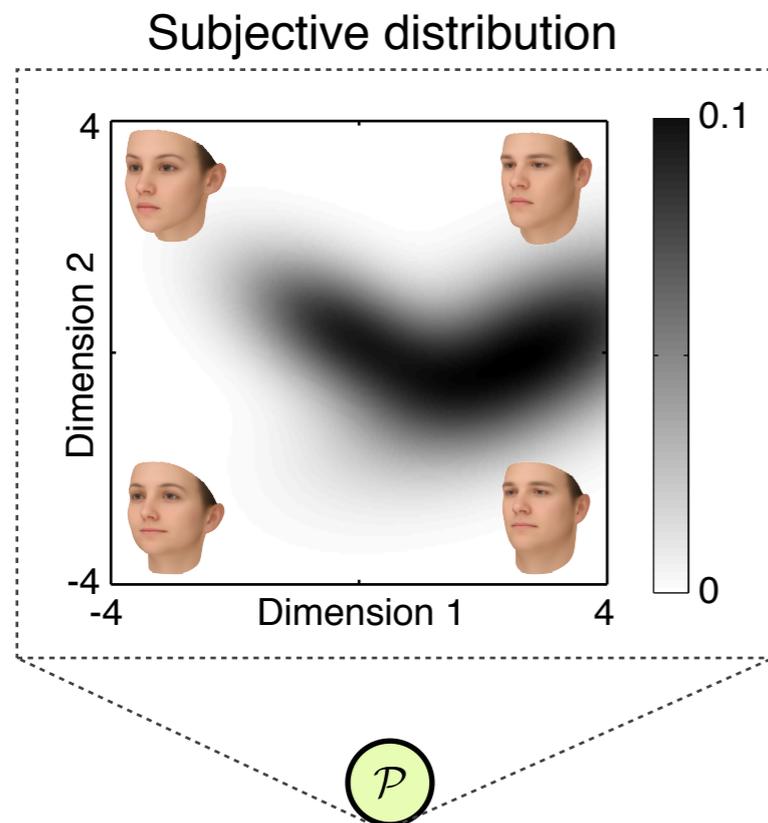


COGNITIVE TOMOGRAPHY

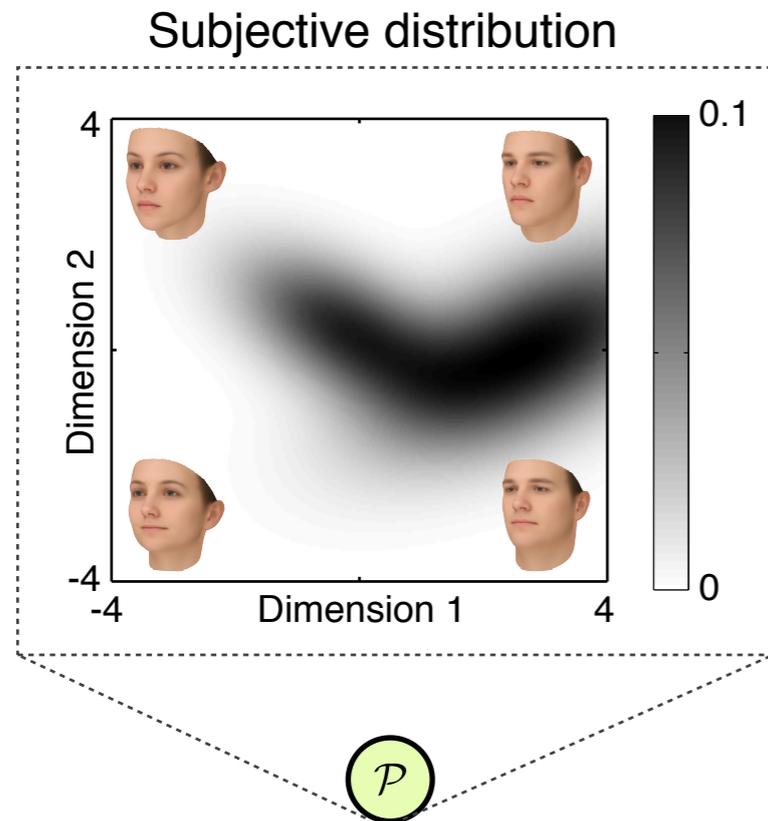
COGNITIVE TOMOGRAPHY



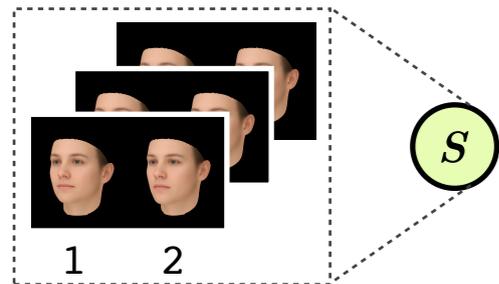
COGNITIVE TOMOGRAPHY



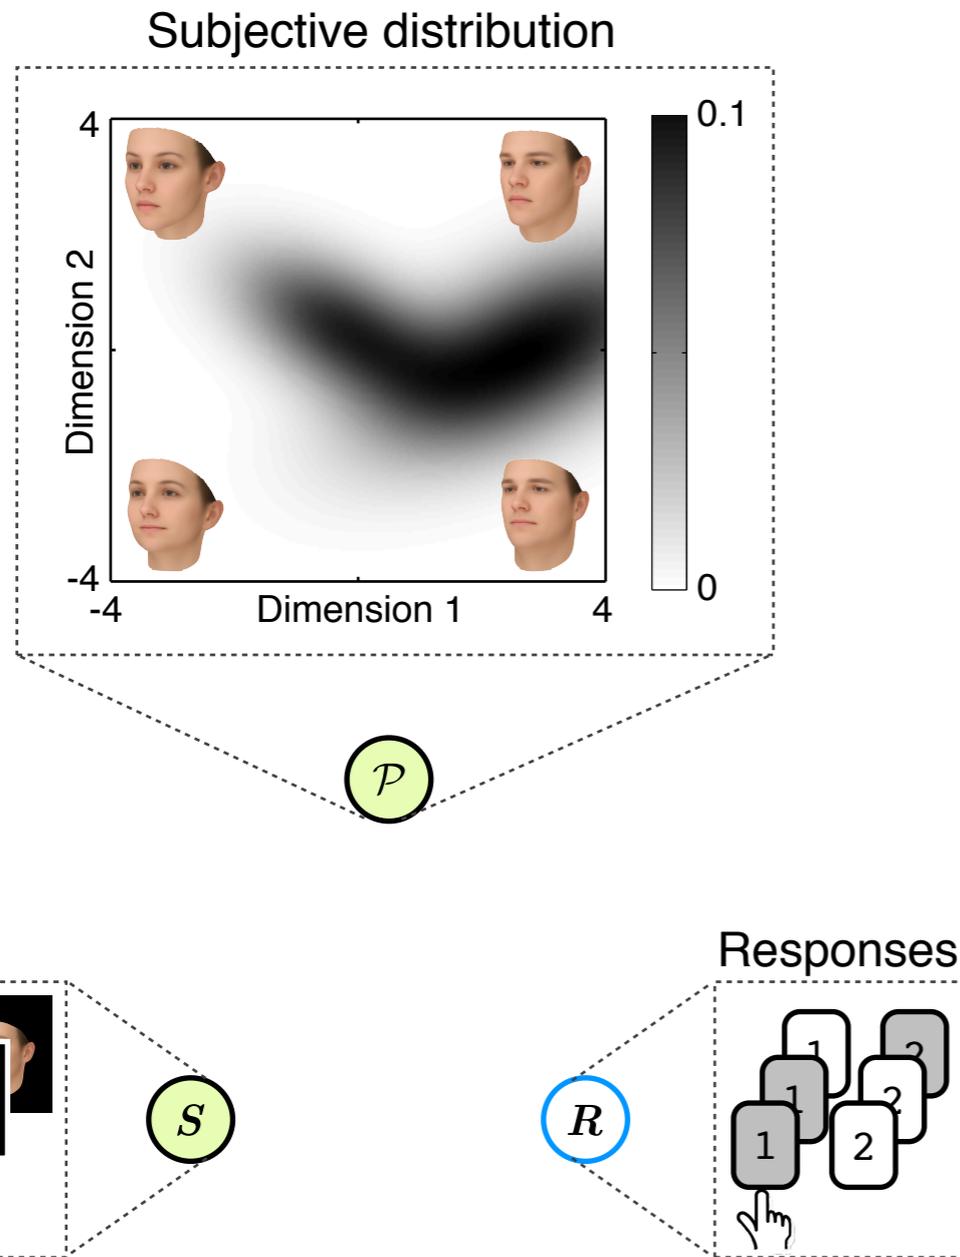
COGNITIVE TOMOGRAPHY



Stimuli



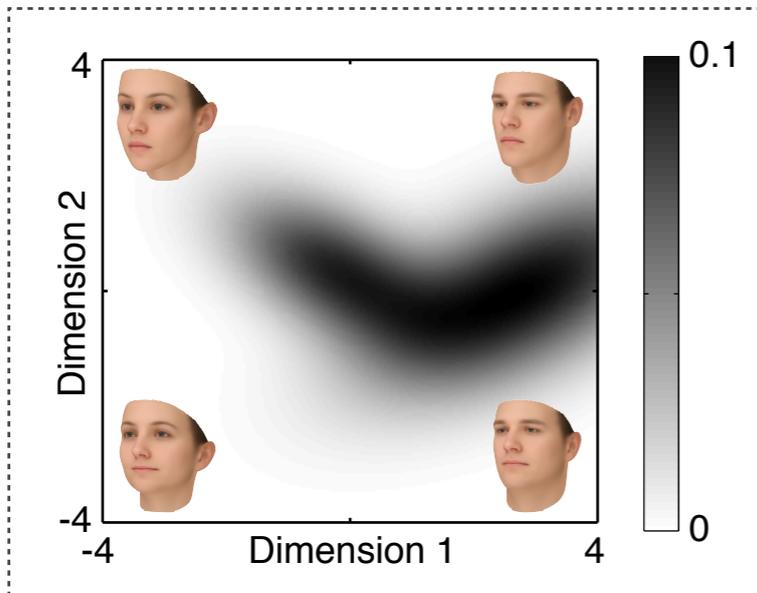
COGNITIVE TOMOGRAPHY



COGNITIVE TOMOGRAPHY

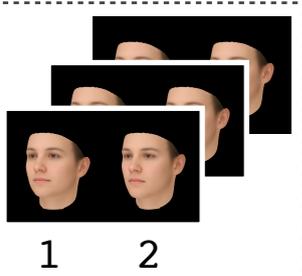
ideal observer model

Subjective distribution



\mathcal{P}

Stimuli

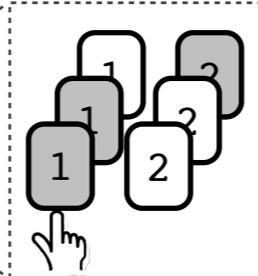


S

Hypotheses

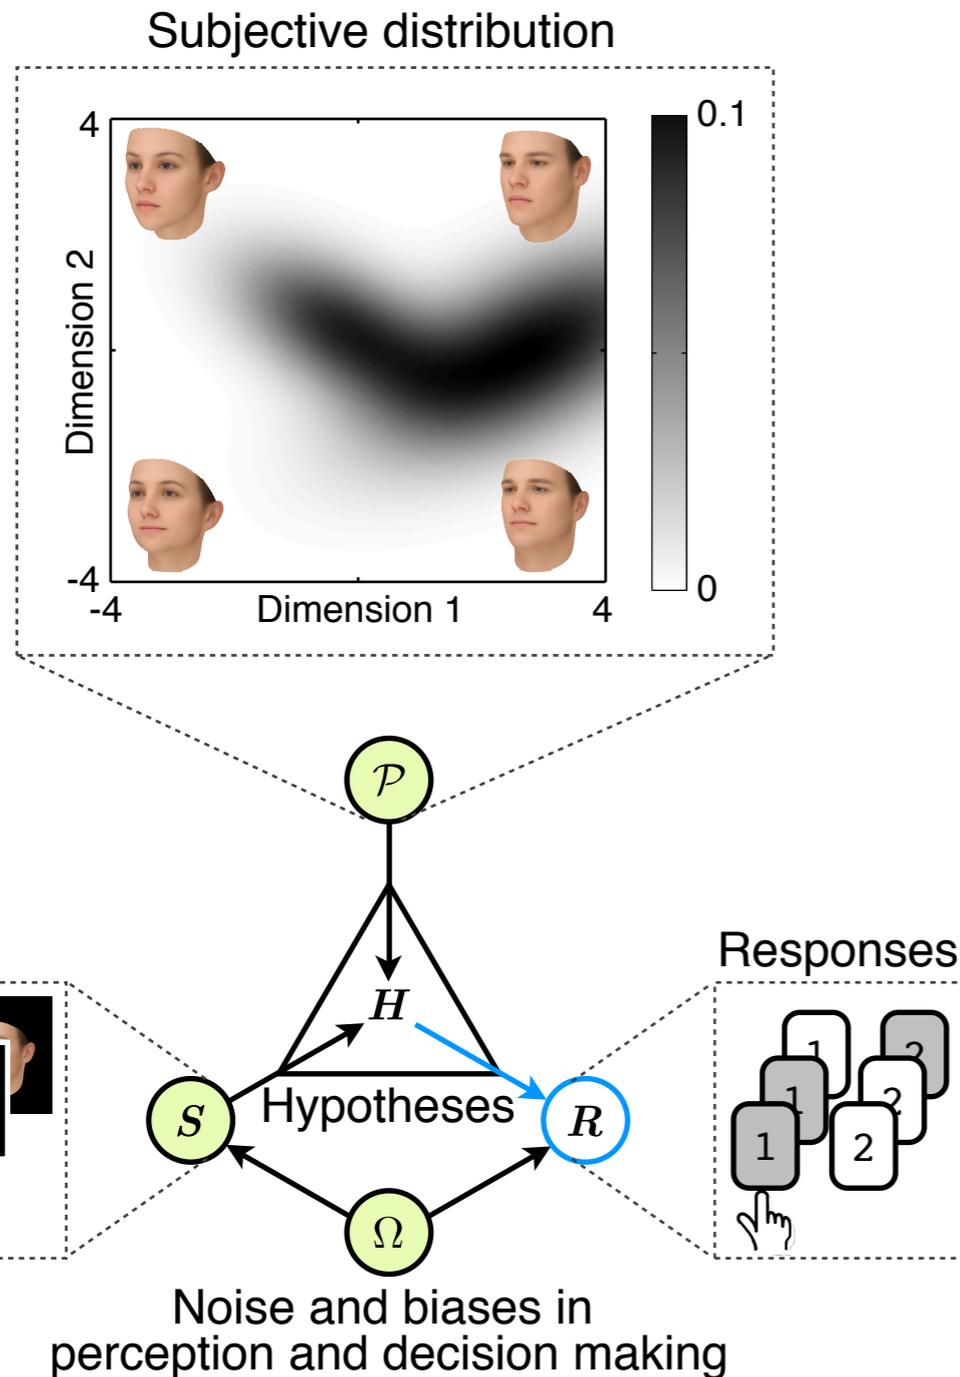
R

Responses



COGNITIVE TOMOGRAPHY

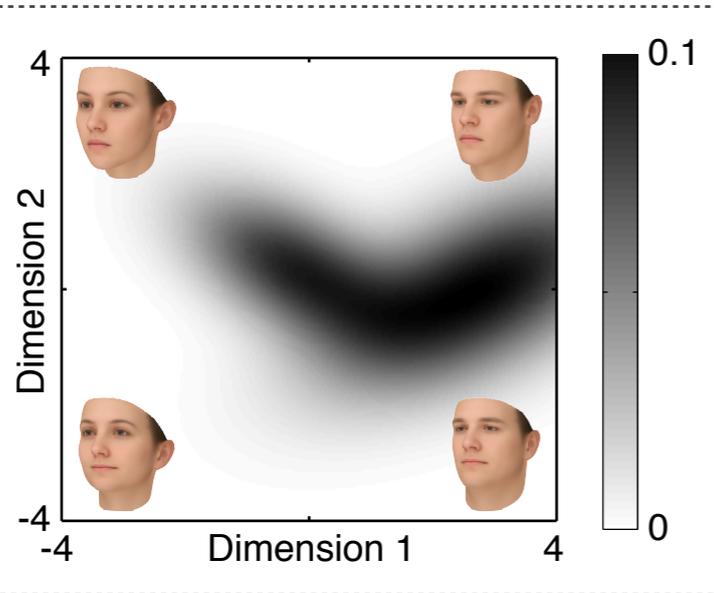
quasi-ideal observer model



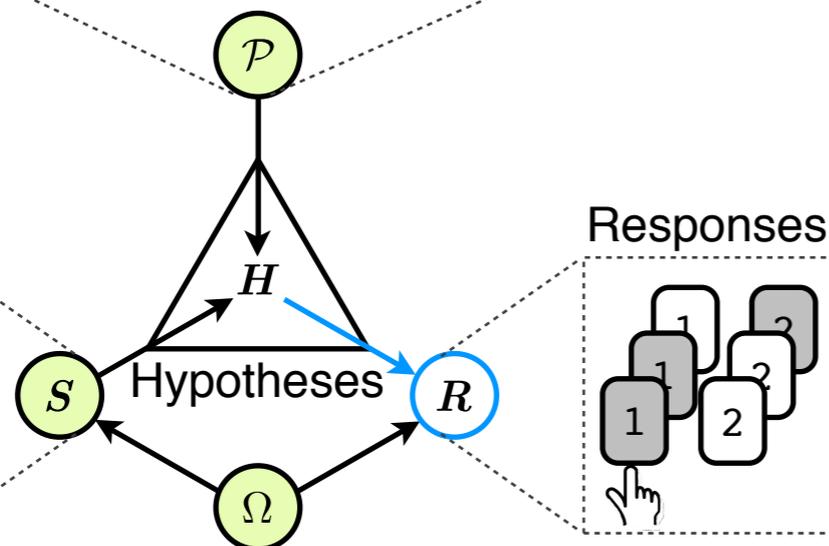
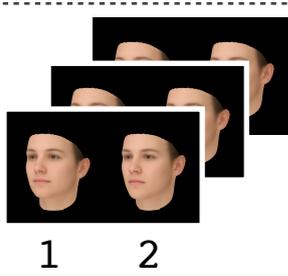
COGNITIVE TOMOGRAPHY

quasi-ideal observer model

Subjective distribution



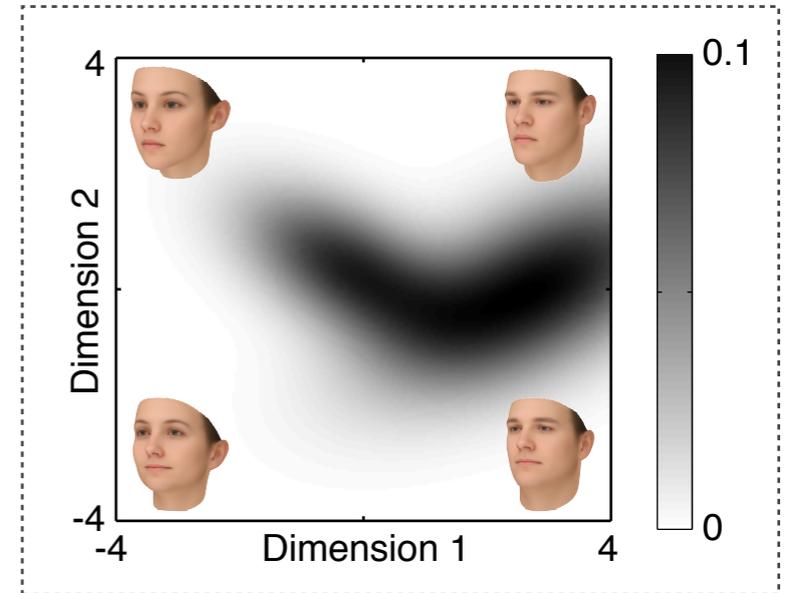
Stimuli



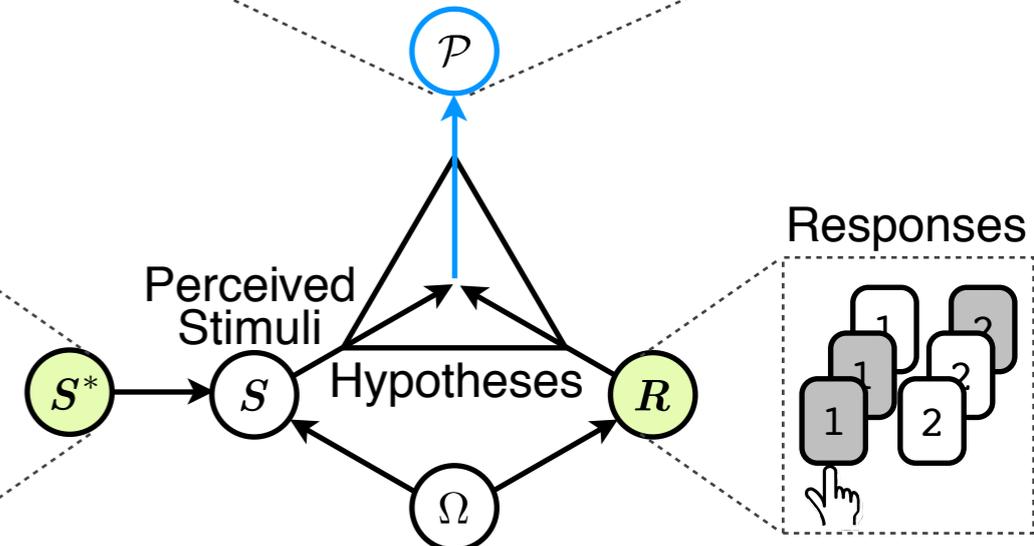
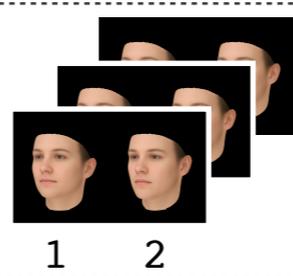
Noise and biases in perception and decision making

cognitive tomography inverting the ideal observer model

Subjective distribution



Presented Stimuli



Noise and biases in perception and decision making

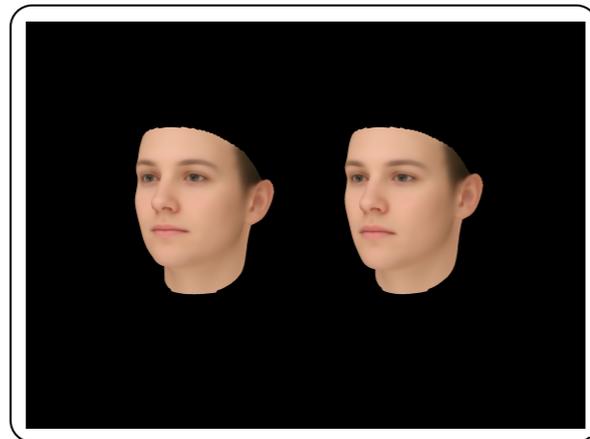
Houlsby et al, Curr Biol 2013

TWO TASKS

TWO TASKS

familiarity task

“which face looks more familiar?”



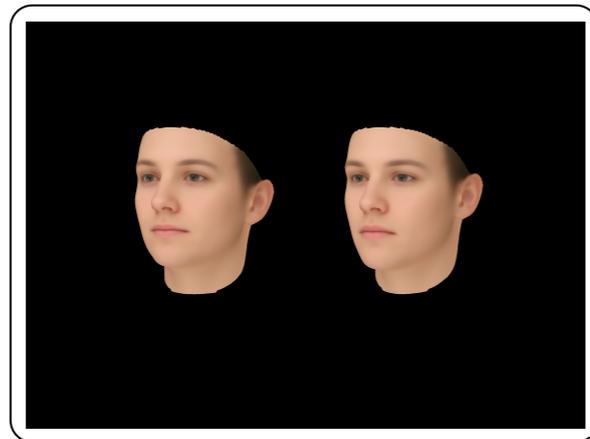
1

2

TWO TASKS

familiarity task

“which face looks more familiar?”

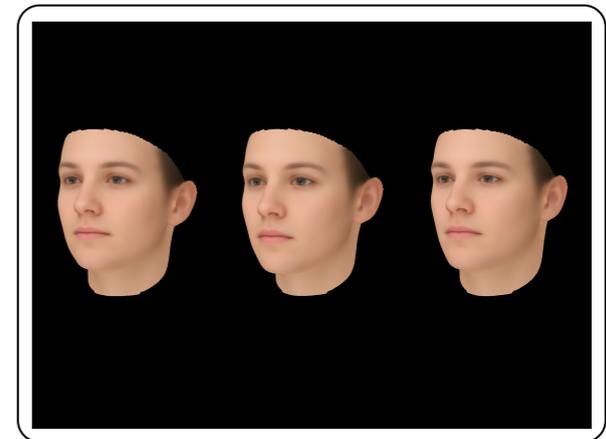


1

2

odd-one-out task

“which face is the odd-one-out?”



1

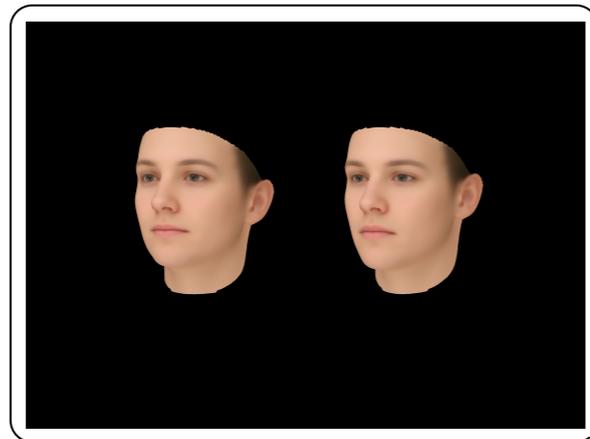
2

3

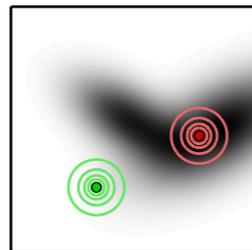
TWO TASKS

familiarity task

“which face looks more familiar?”

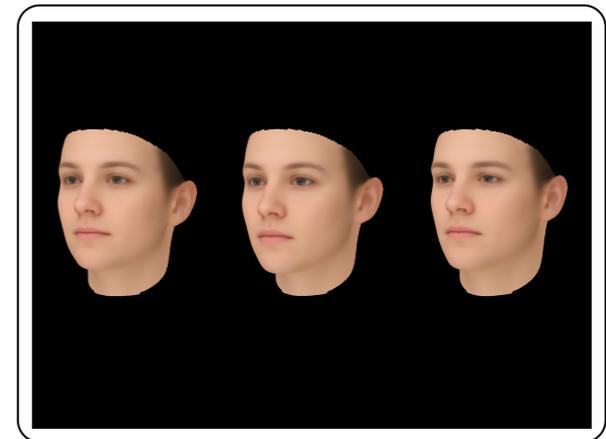


1 2



odd-one-out task

“which face is the odd-one-out?”

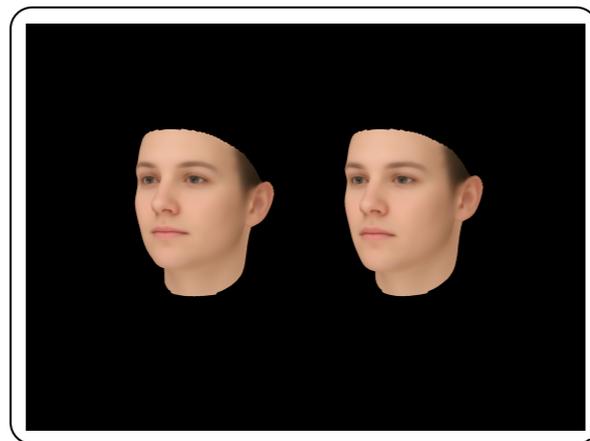


1 2 3

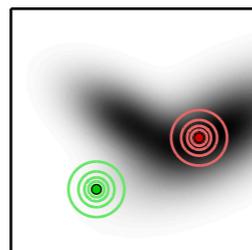
TWO TASKS

familiarity task

“which face looks more familiar?”

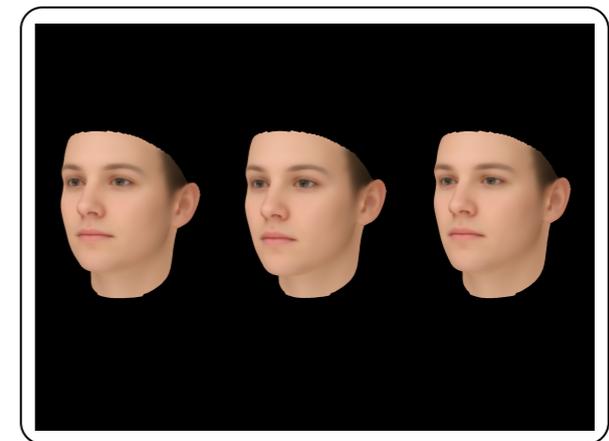


1 2

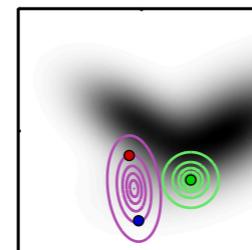


odd-one-out task

“which face is the odd-one-out?”



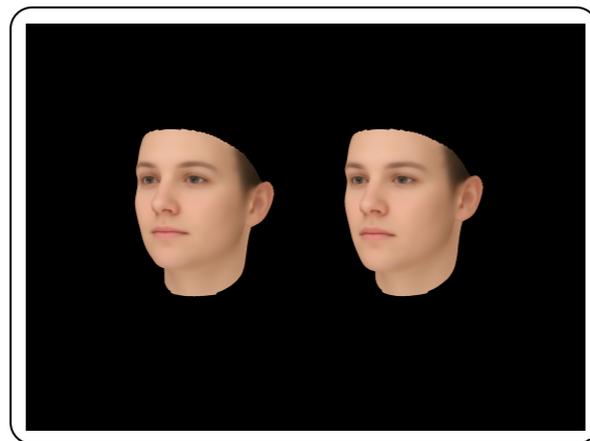
1 2 3



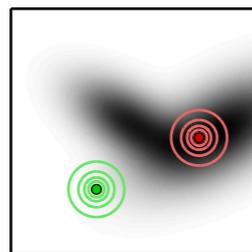
TWO TASKS

familiarity task

“which face looks more familiar?”

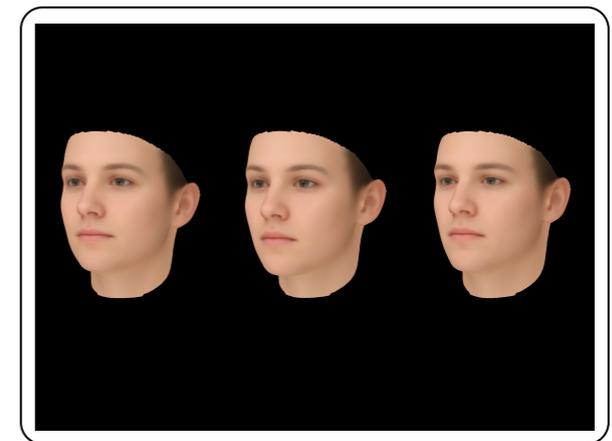


1 2

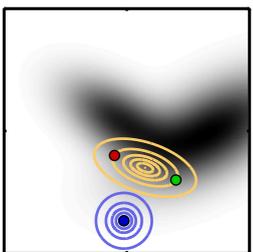
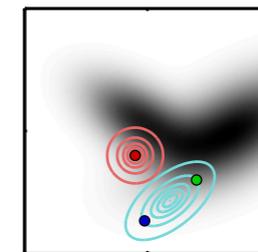
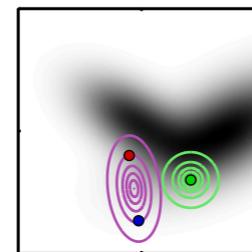


odd-one-out task

“which face is the odd-one-out?”



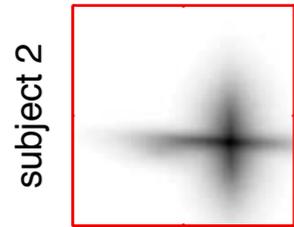
1 2 3



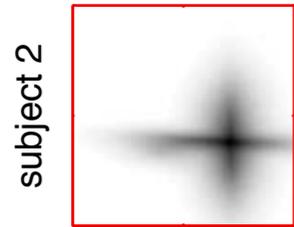
Houlsby et al, Curr Biol 2013

SUBJECTIVE DISTRIBUTIONS

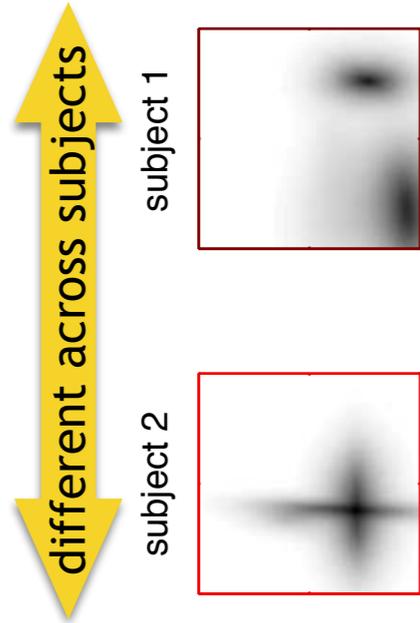
SUBJECTIVE DISTRIBUTIONS



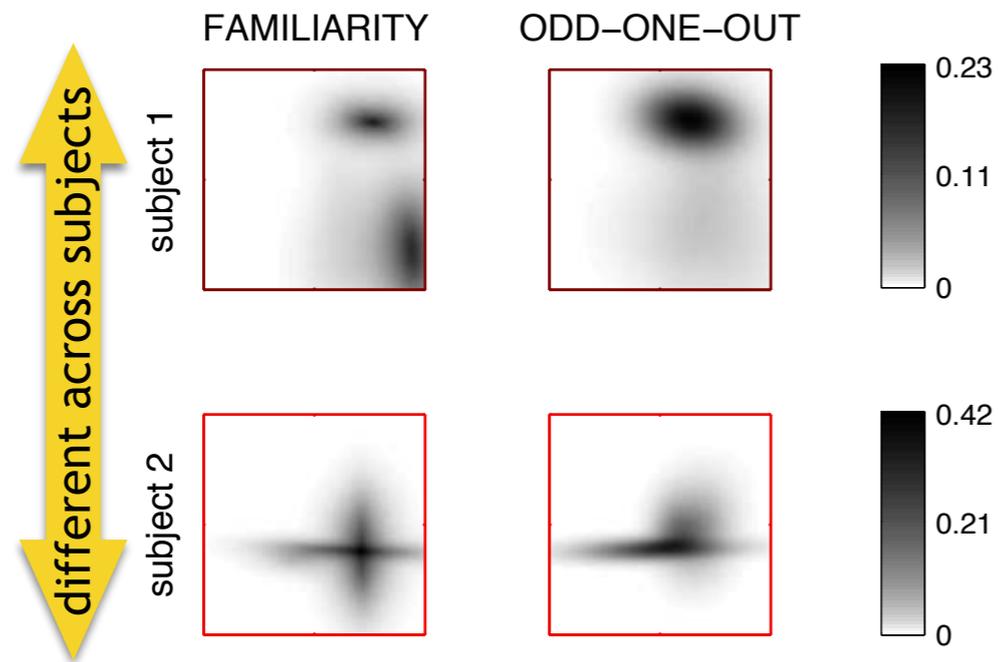
SUBJECTIVE DISTRIBUTIONS



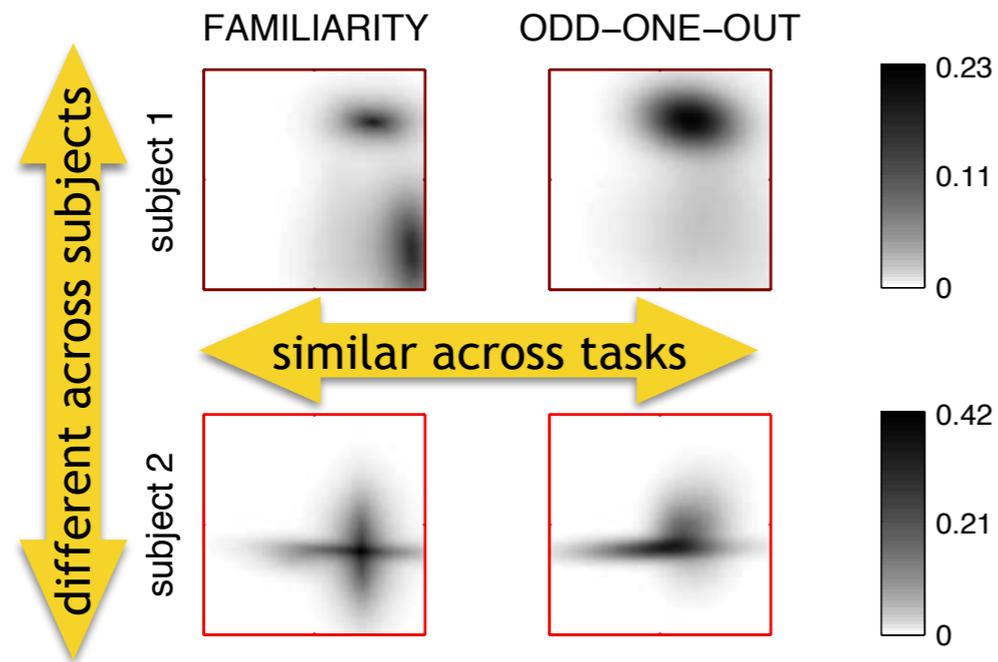
SUBJECTIVE DISTRIBUTIONS



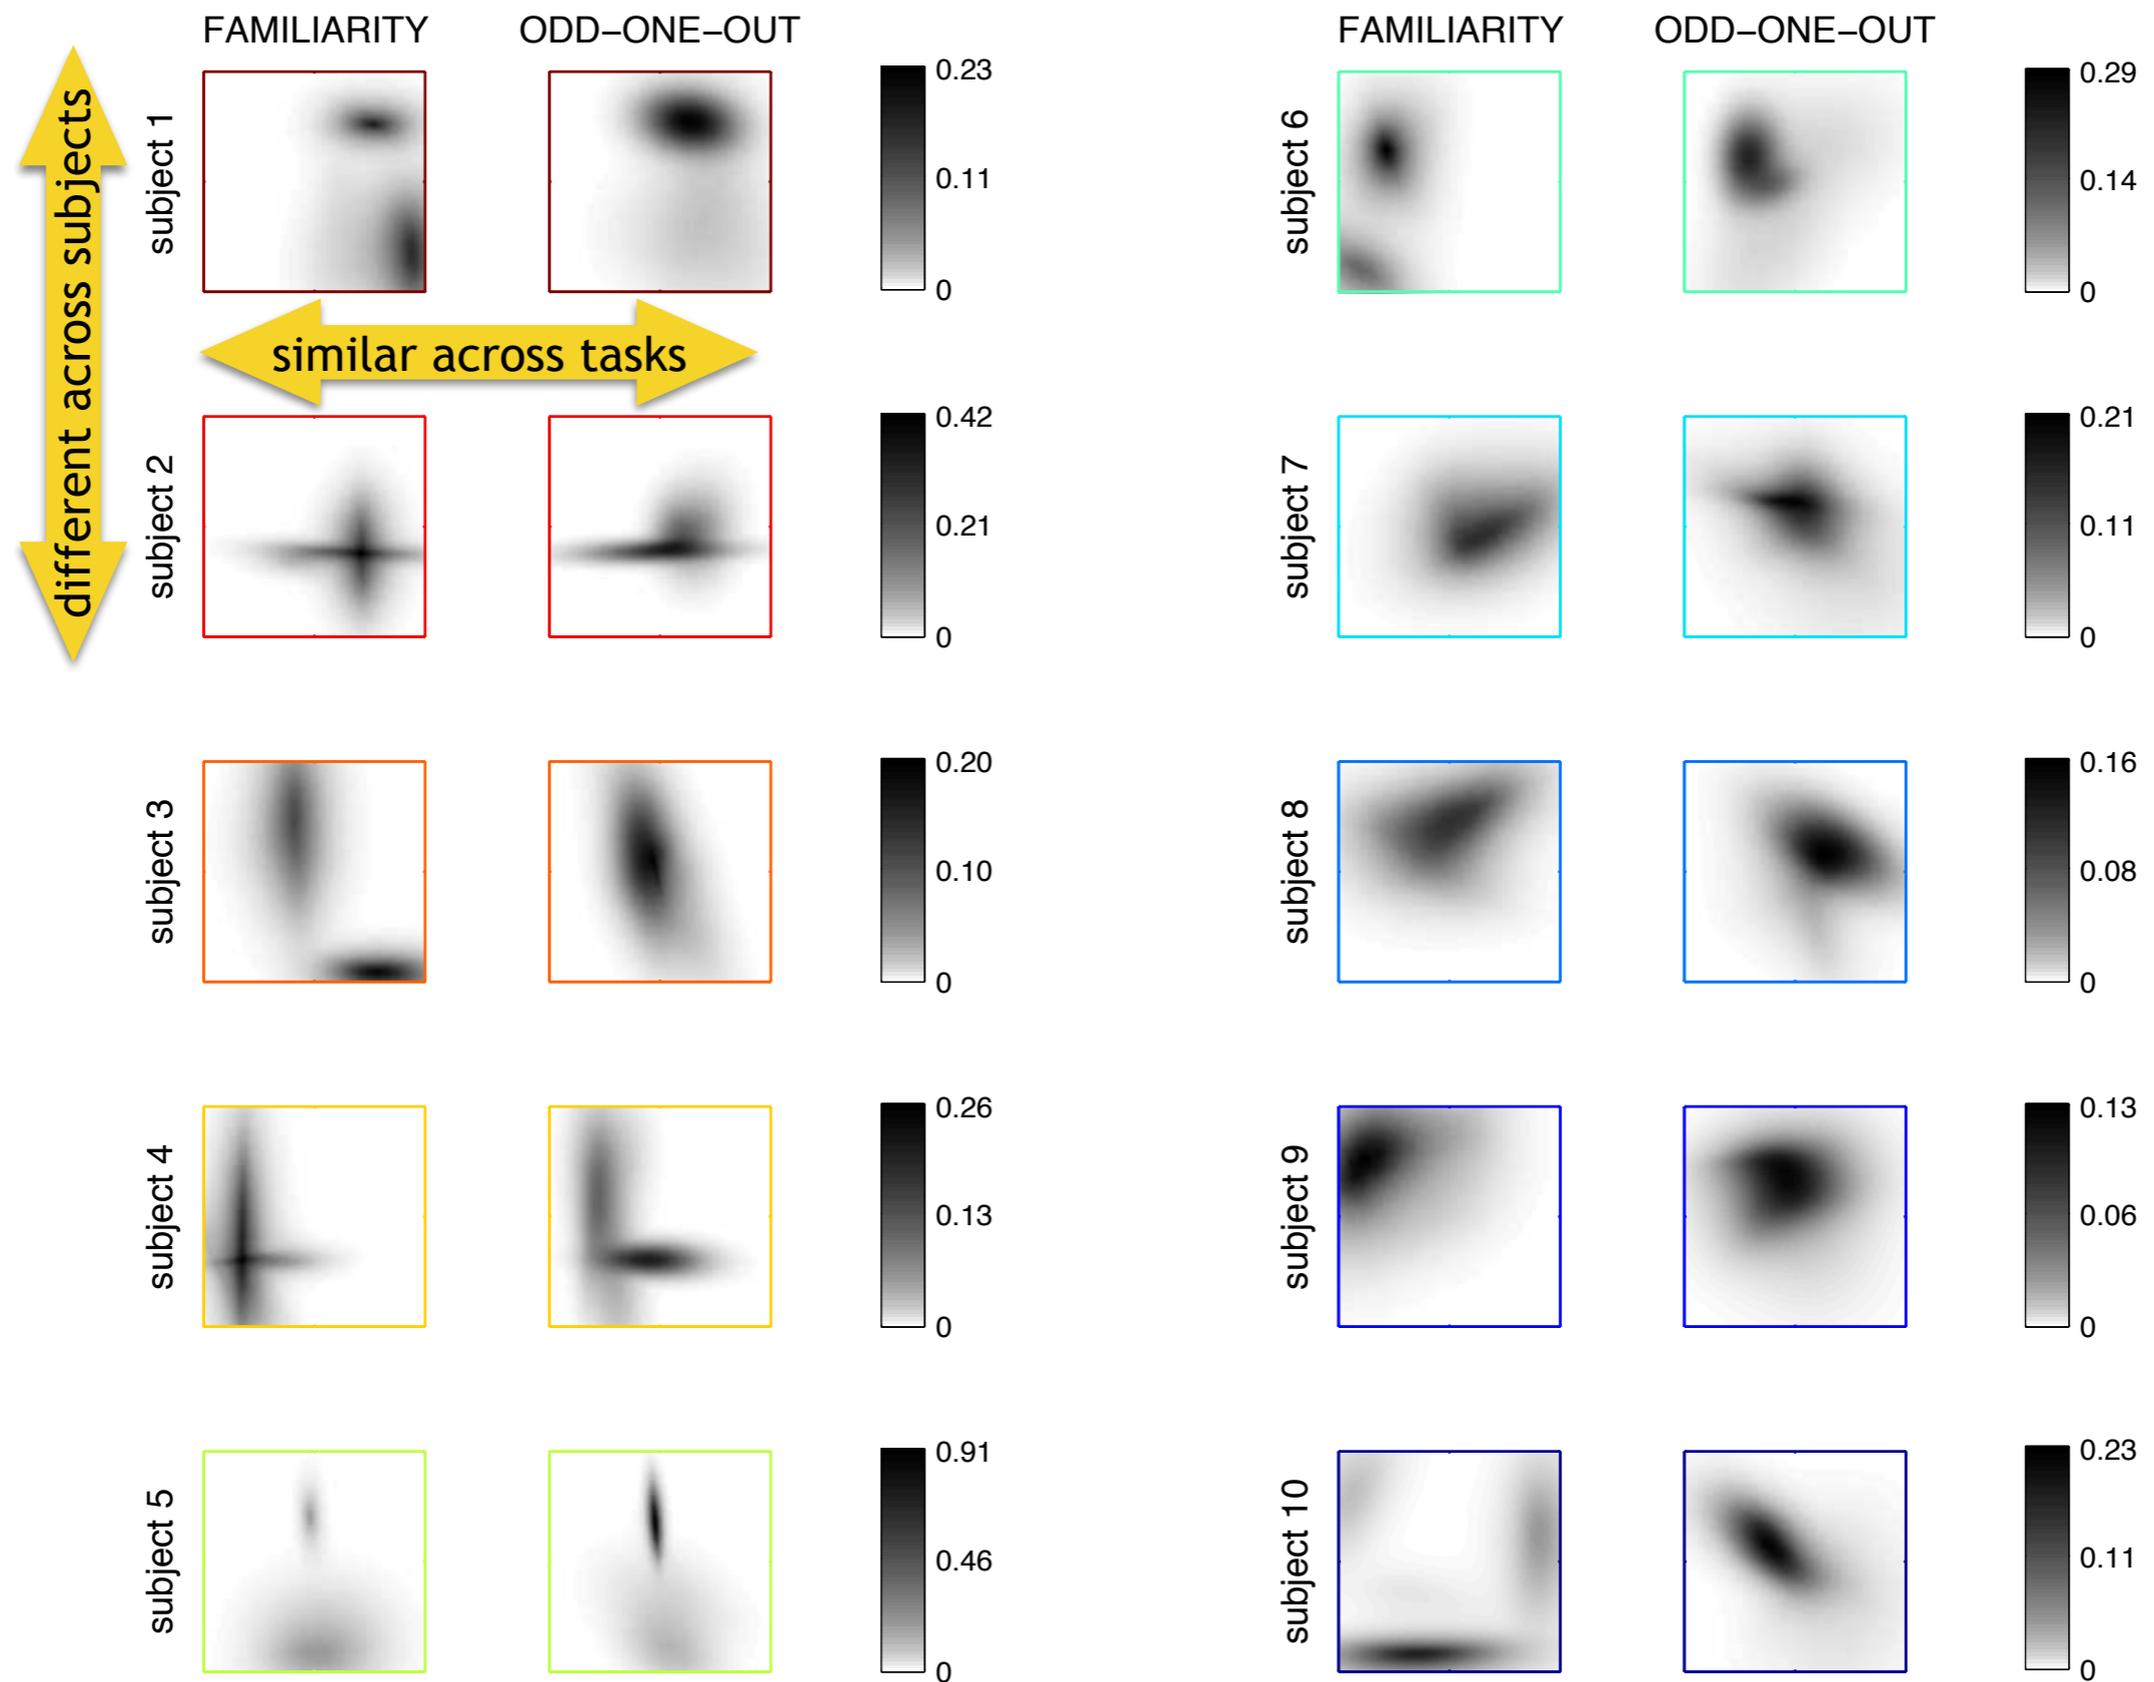
SUBJECTIVE DISTRIBUTIONS



SUBJECTIVE DISTRIBUTIONS



SUBJECTIVE DISTRIBUTIONS



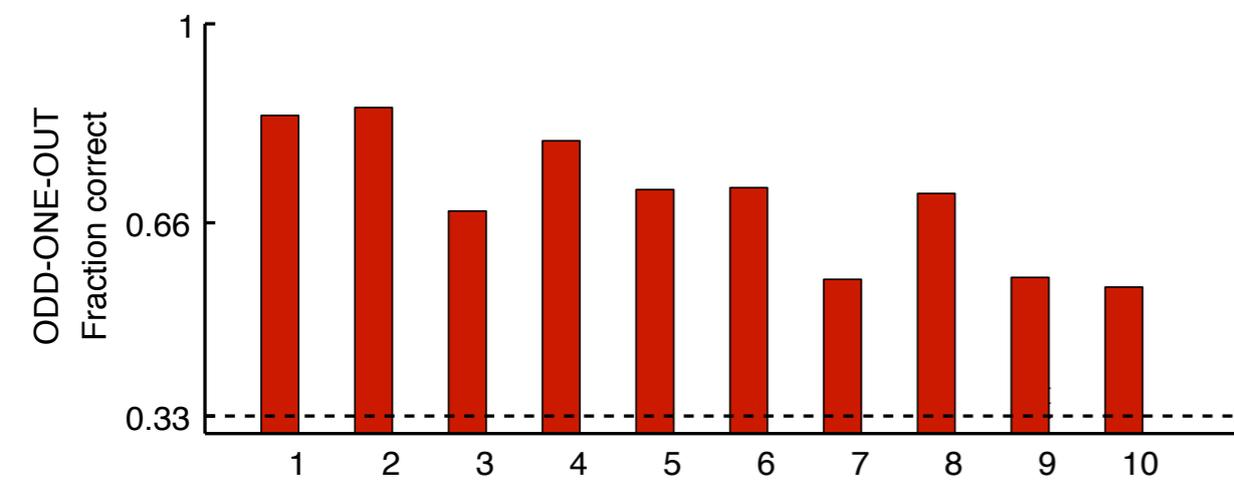
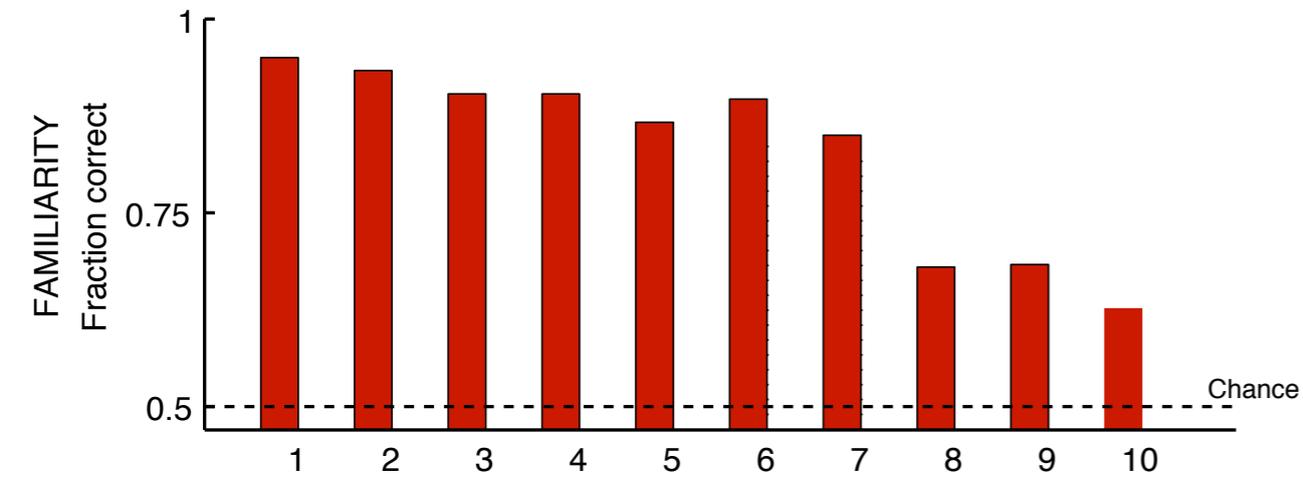
Houlsby et al, Curr Biol 2013

PREDICTING BEHAVIOUR

PREDICTING BEHAVIOUR

well above chance

Individual subjects

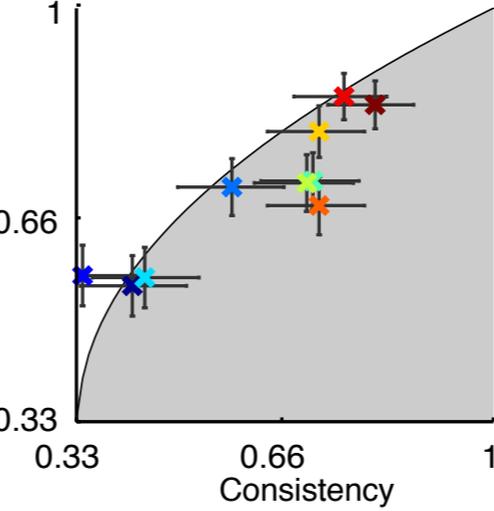
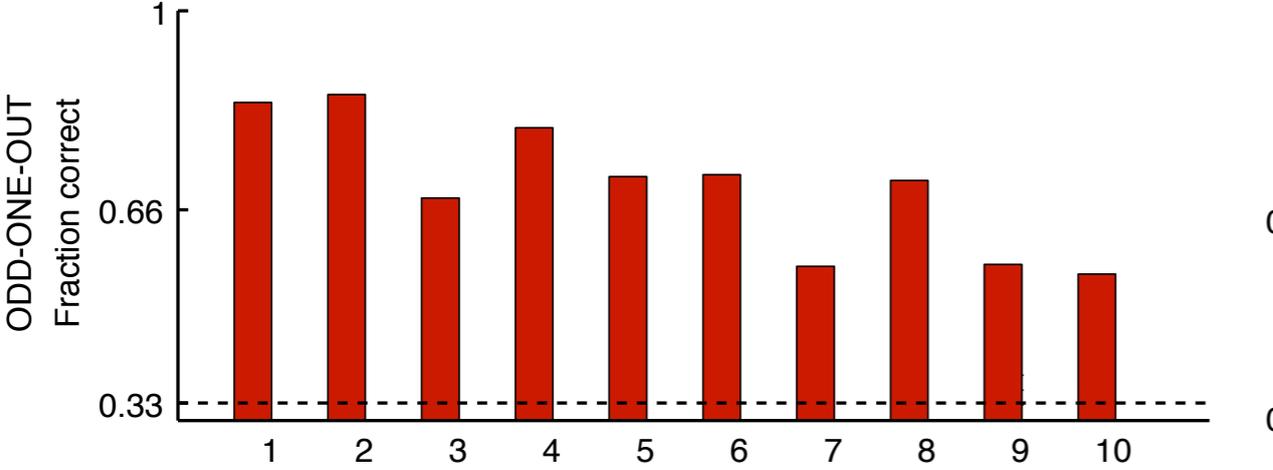
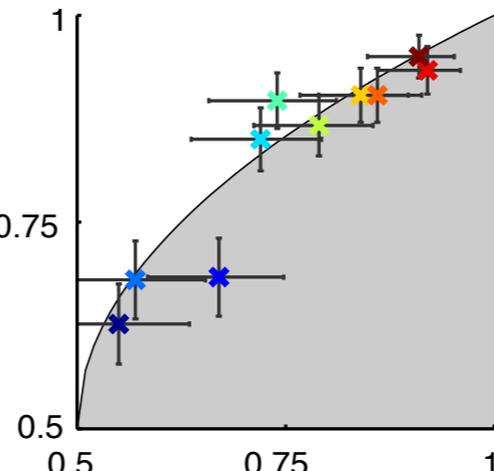
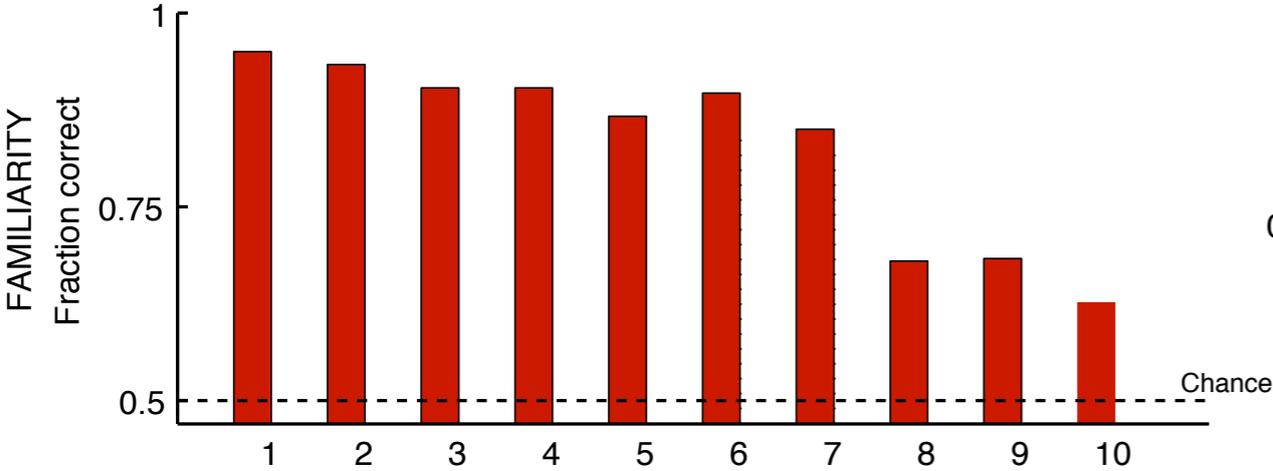


PREDICTING BEHAVIOUR

well above chance

close to theoretical upper bound

Individual subjects



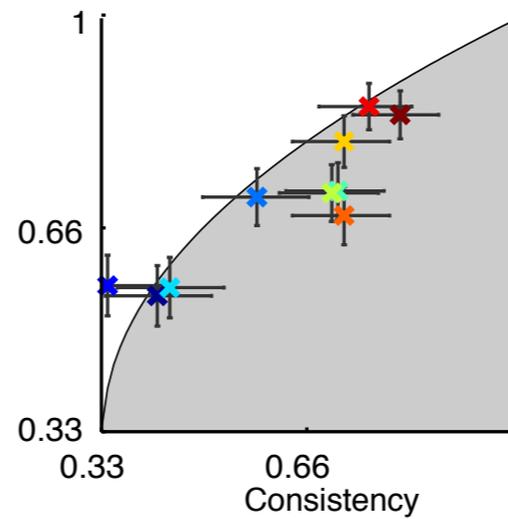
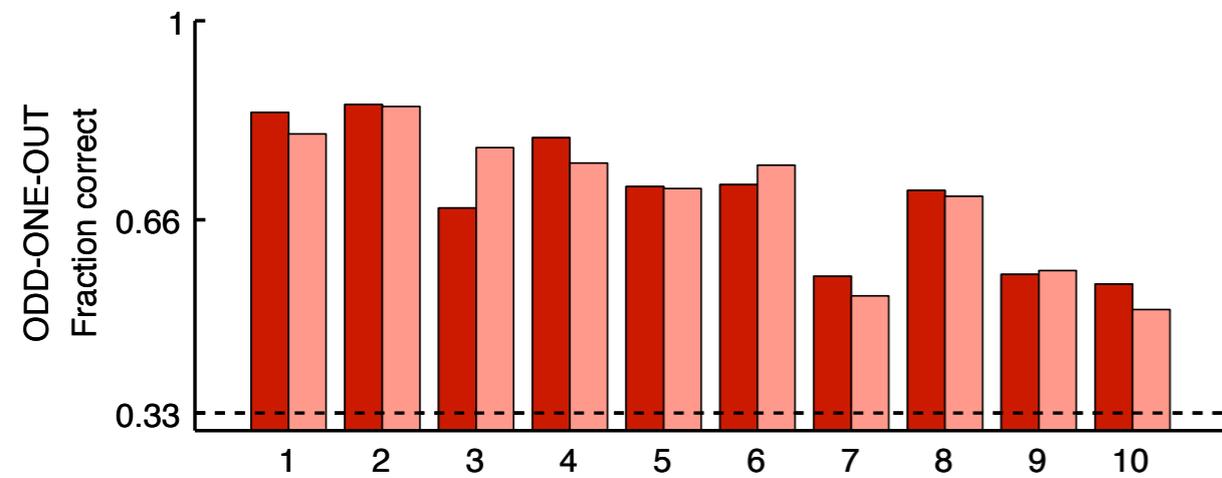
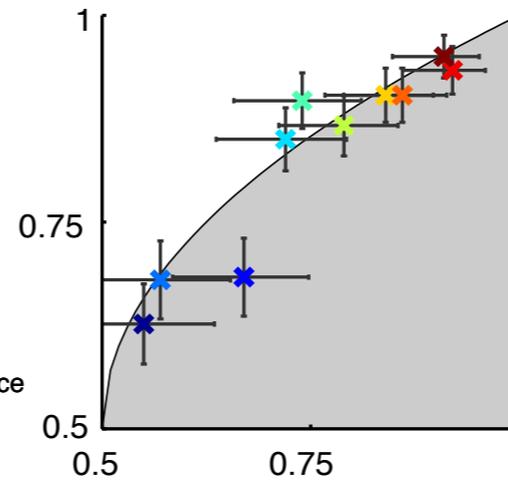
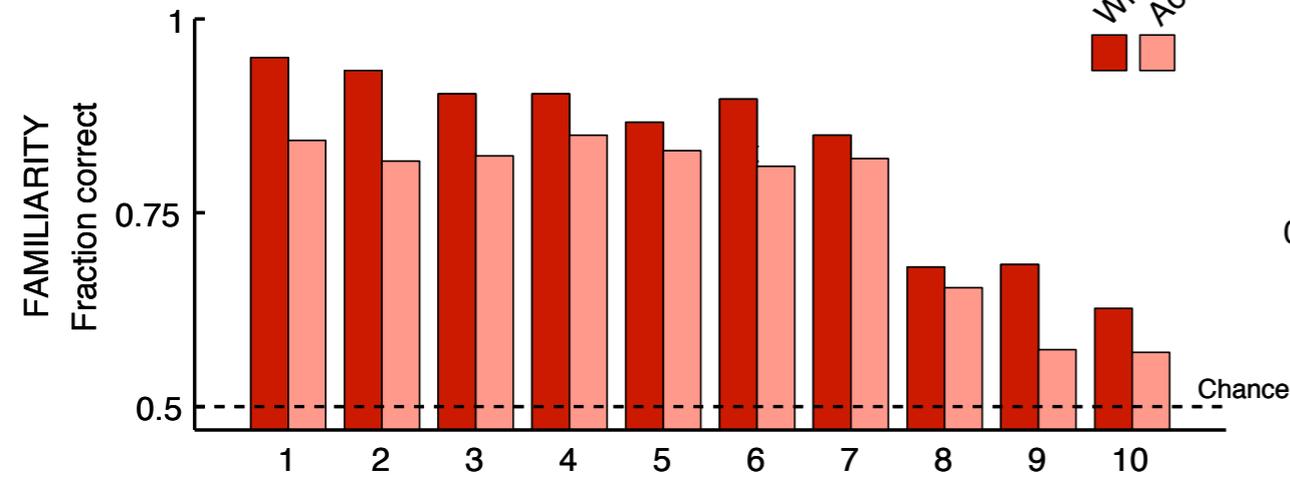
PREDICTING BEHAVIOUR

well above chance
both within and across tasks

close to theoretical
upper bound

Individual subjects

Within task
Across task



PREDICTING BEHAVIOUR

well above chance
both within and across tasks

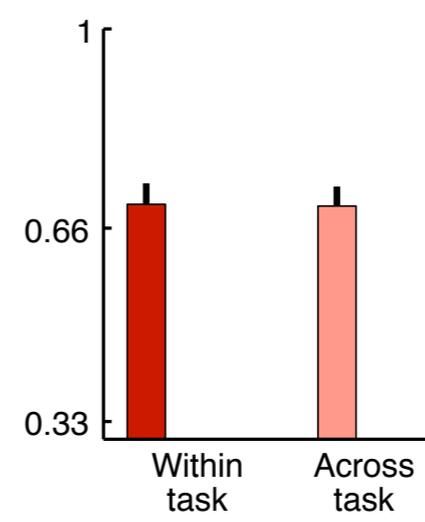
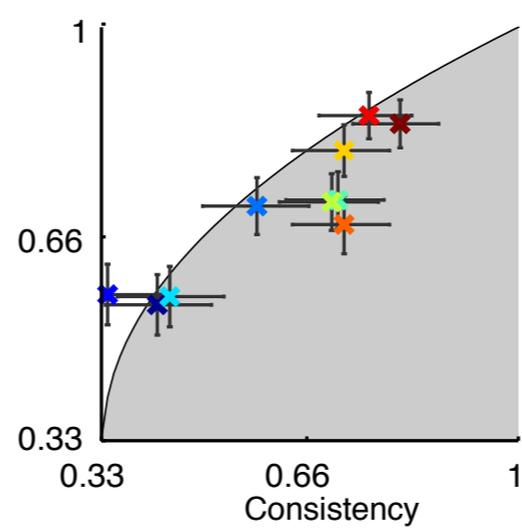
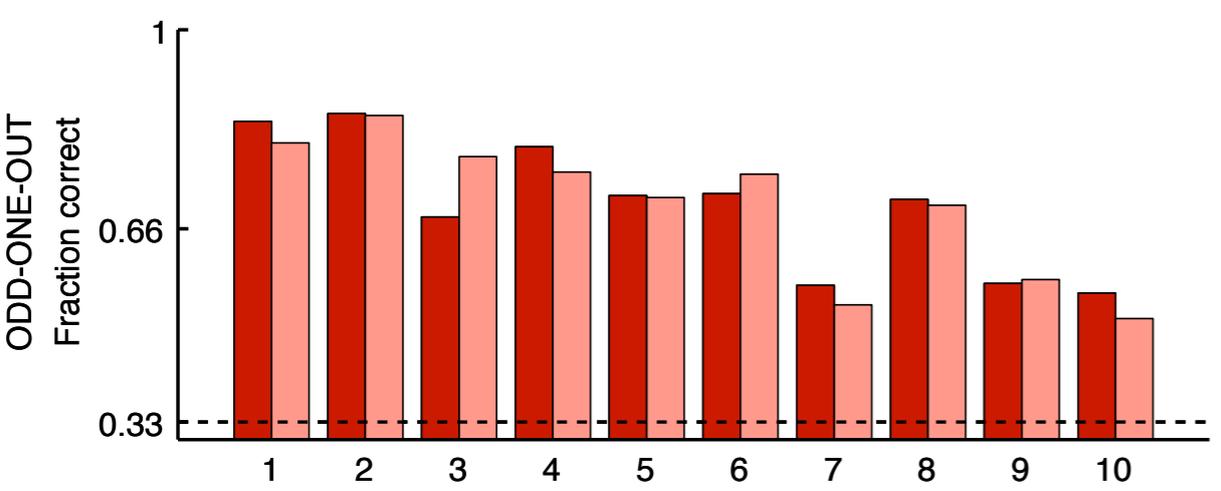
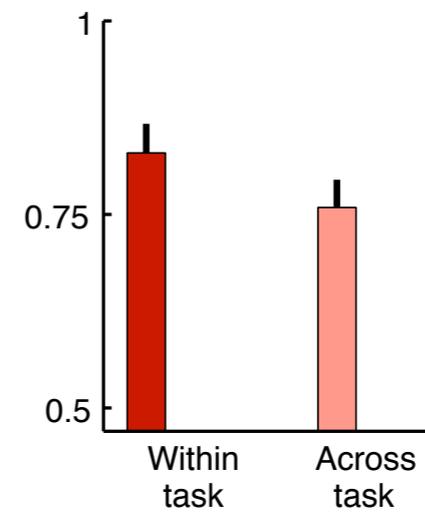
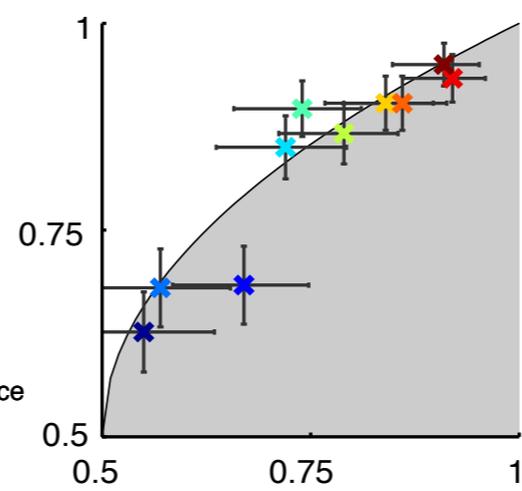
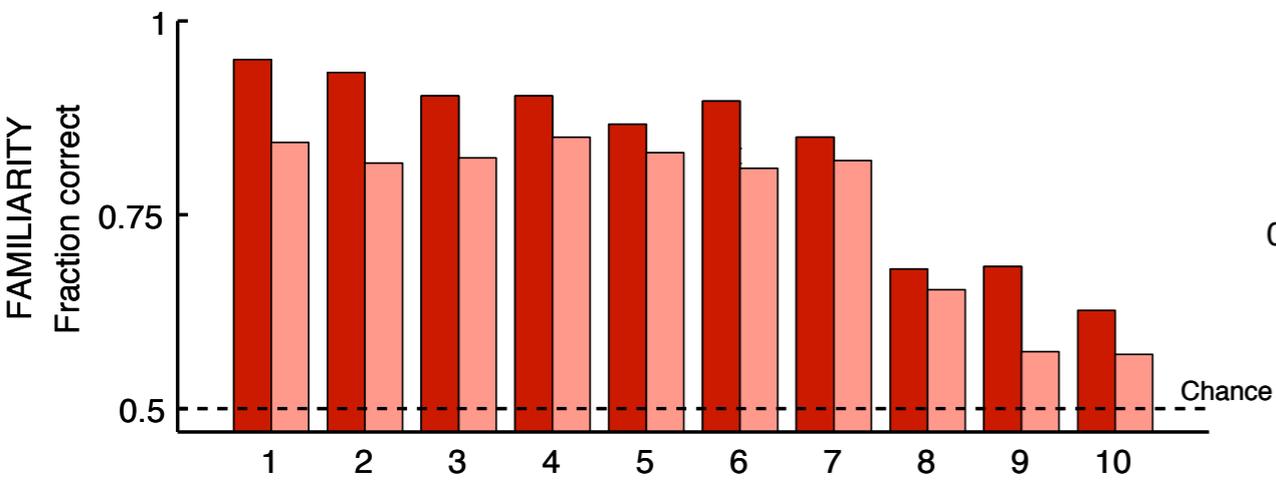
close to theoretical
upper bound

better than
controls

Individual subjects

Group averages

Within task
Across task
Cognitive tomography



Houlsby et al, Curr Biol 2013

PREDICTING BEHAVIOUR

well above chance
both within and across tasks

close to theoretical
upper bound

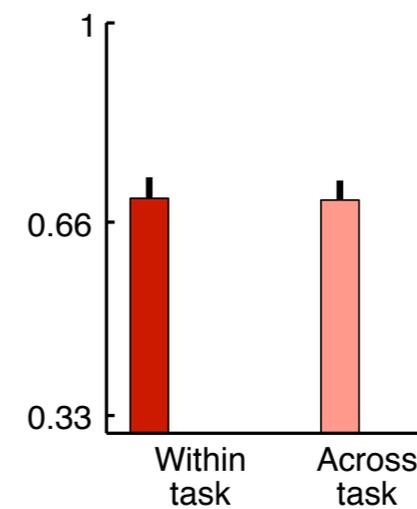
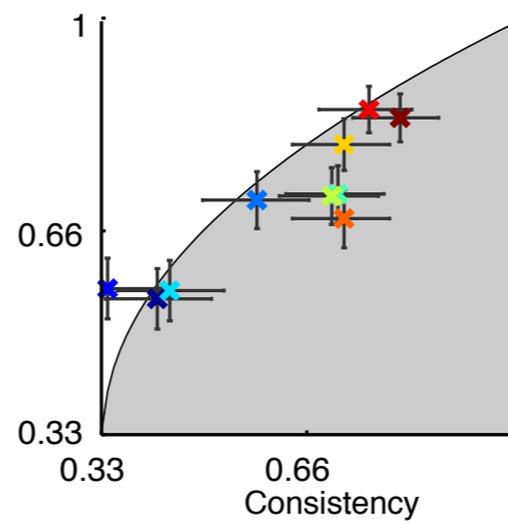
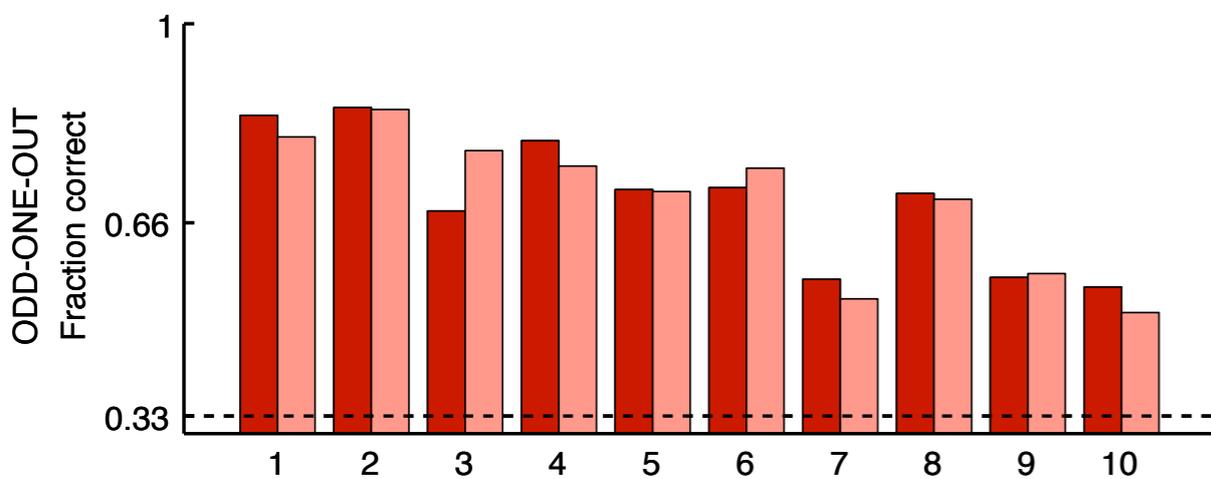
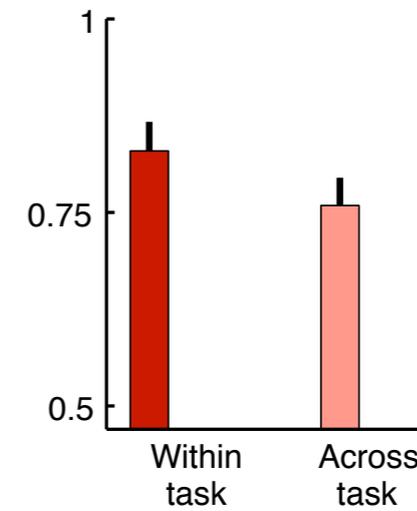
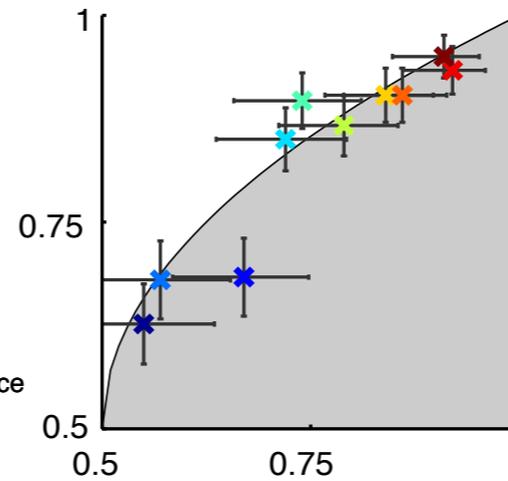
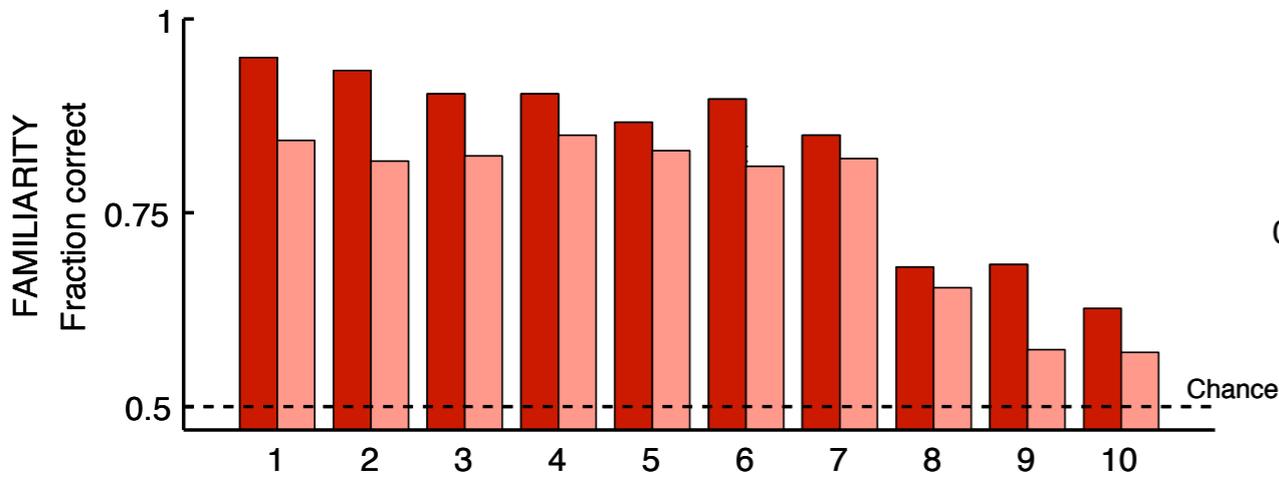
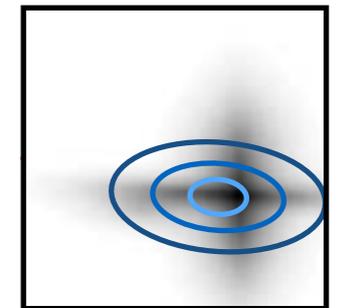
better than
controls

Individual subjects

Group averages

Within task
Across task

Cognitive tomography
Moment-matched



Houlsby et al, Curr Biol 2013

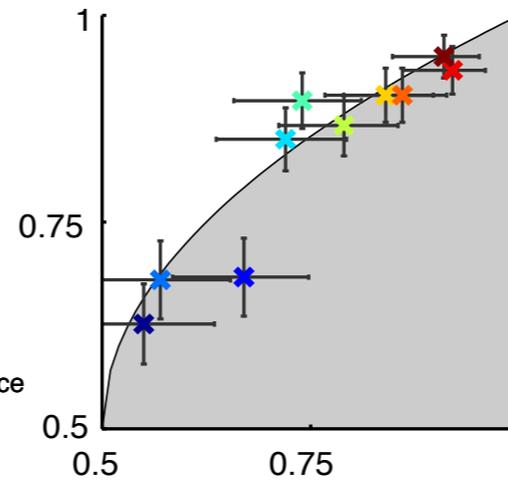
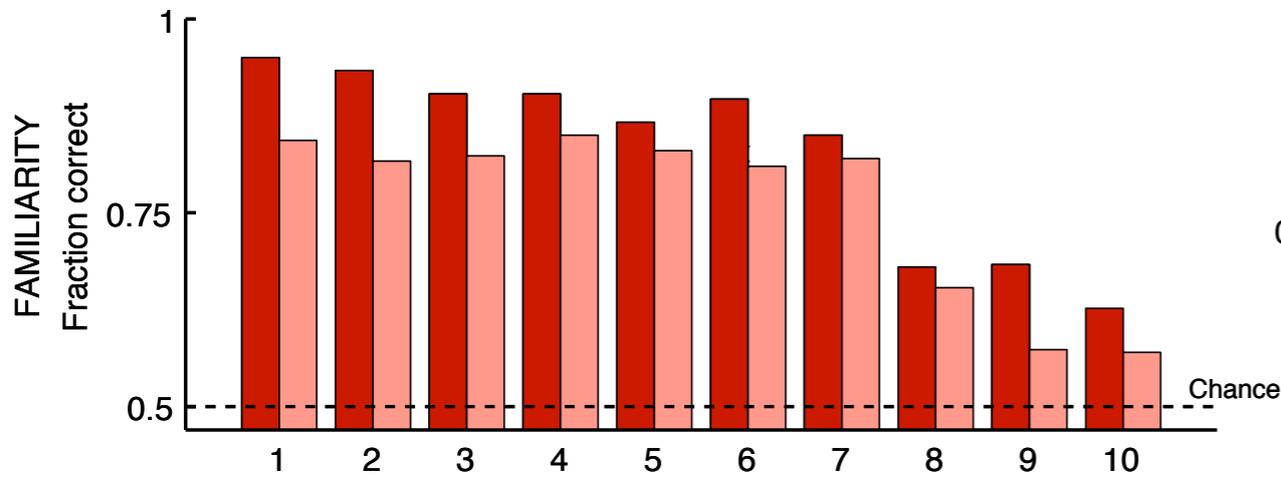
PREDICTING BEHAVIOUR

well above chance
both within and across tasks

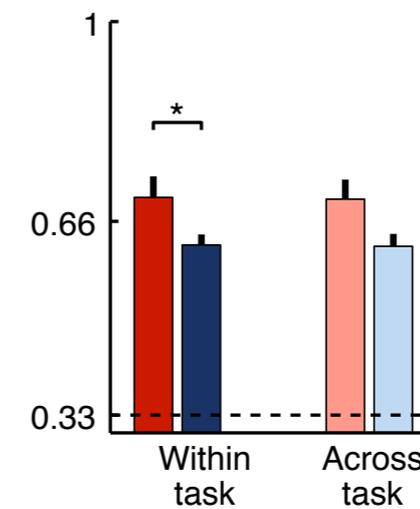
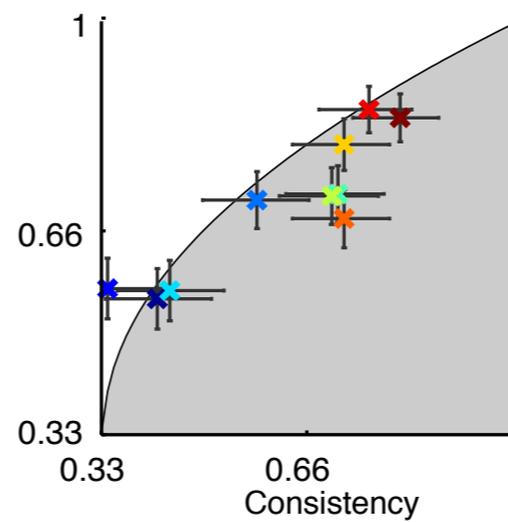
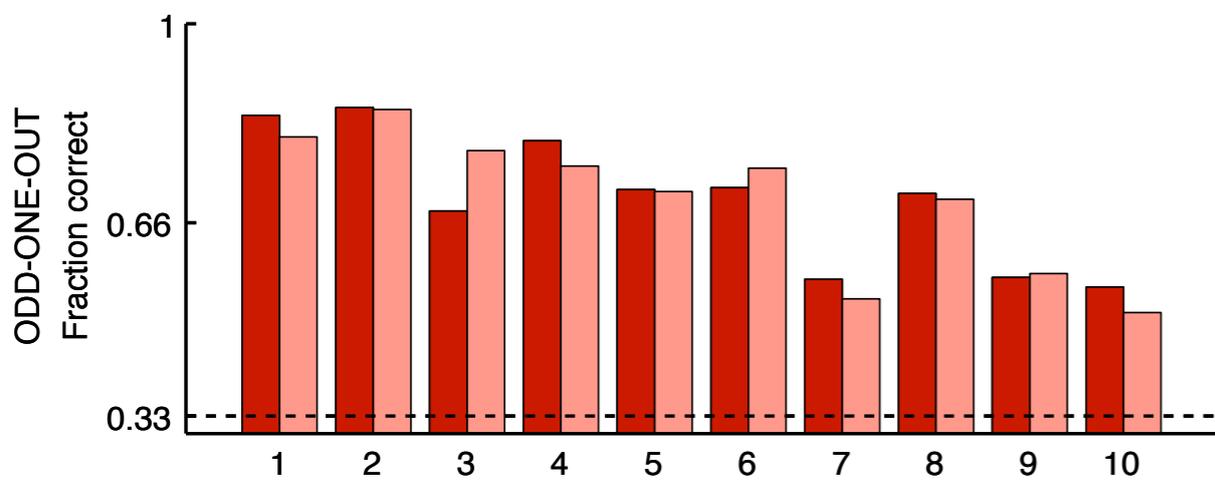
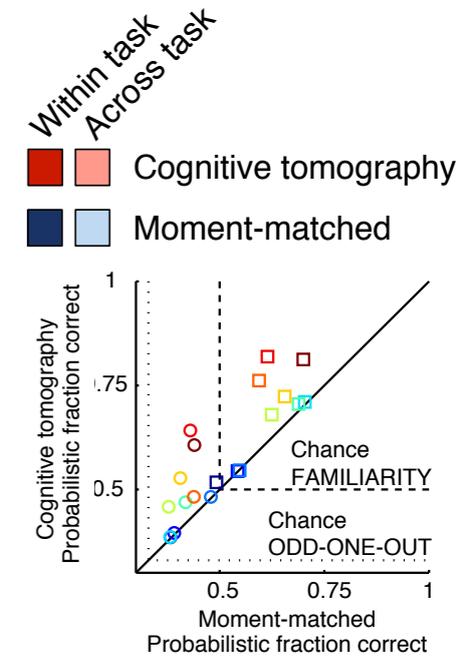
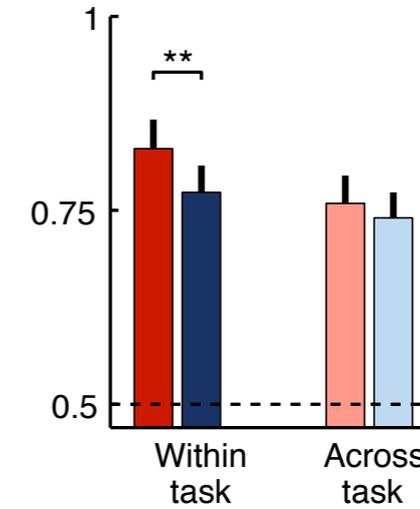
close to theoretical
upper bound

better than
controls

Individual subjects



Group averages



Houlsby et al, Curr Biol 2013

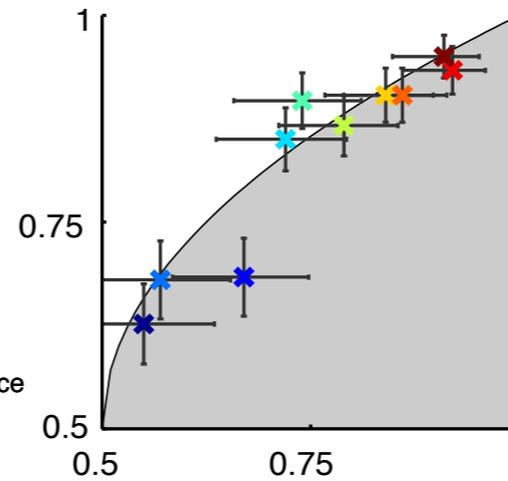
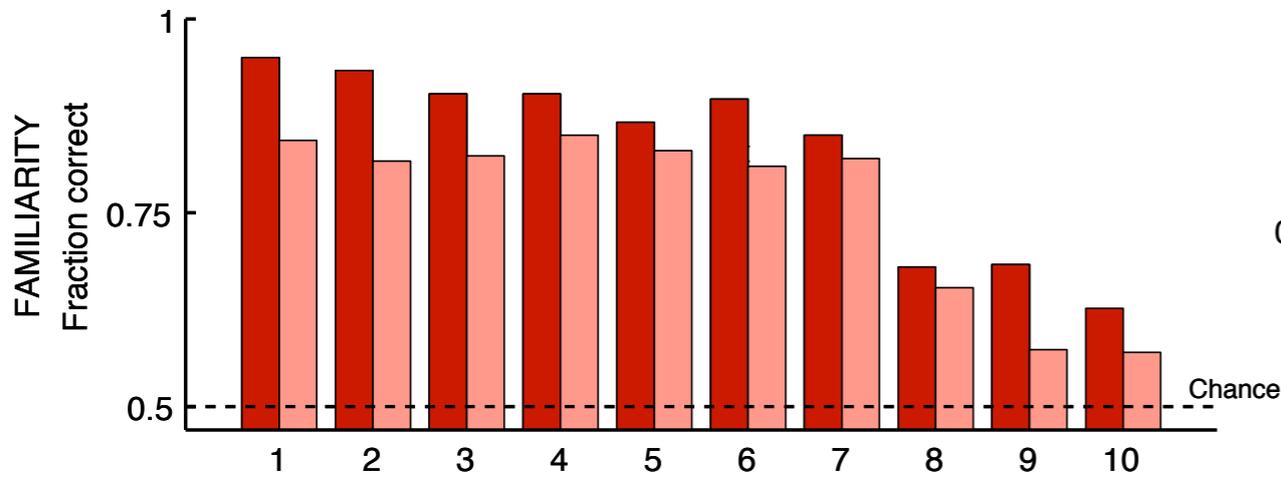
PREDICTING BEHAVIOUR

well above chance
both within and across tasks

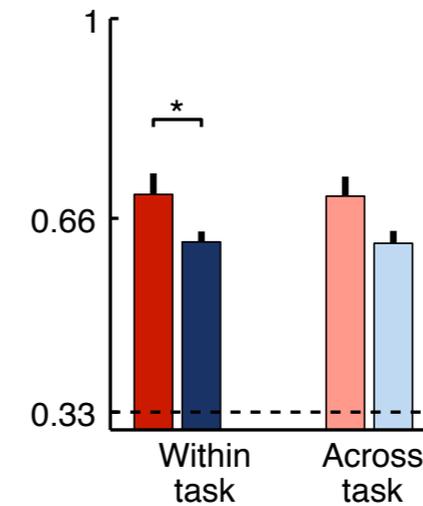
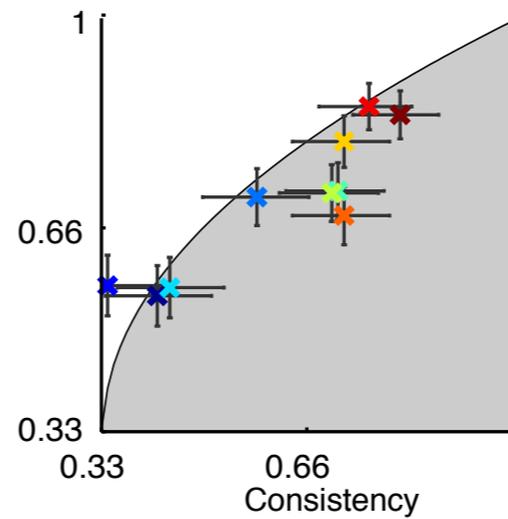
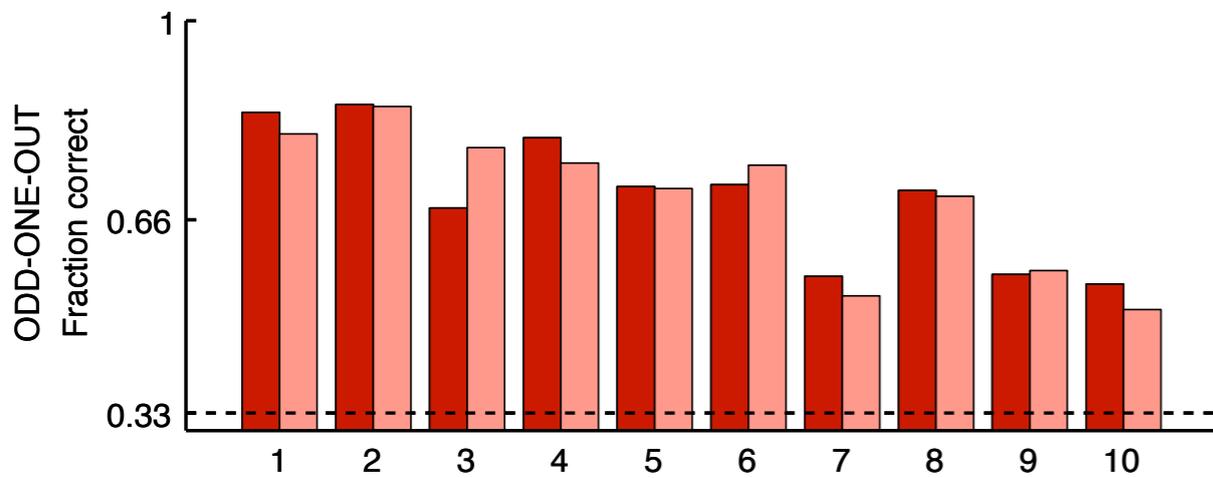
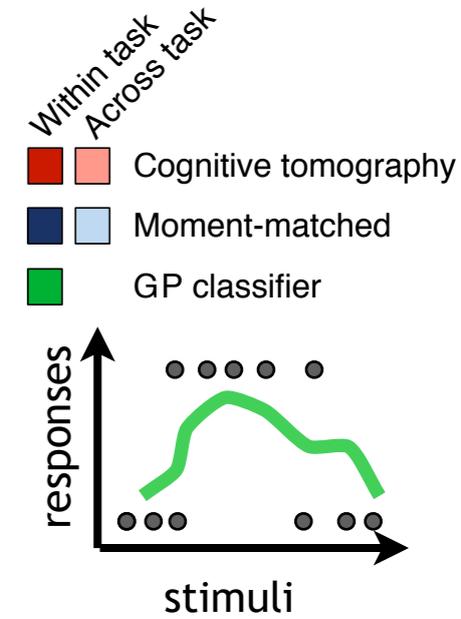
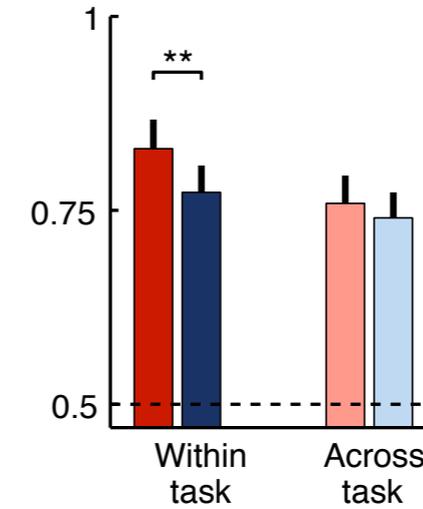
close to theoretical
upper bound

better than
controls

Individual subjects



Group averages



Houlsby et al, Curr Biol 2013

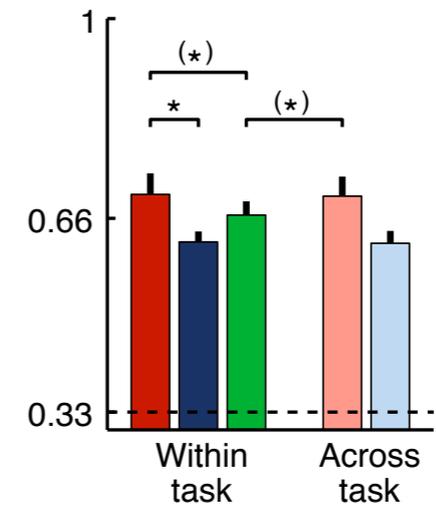
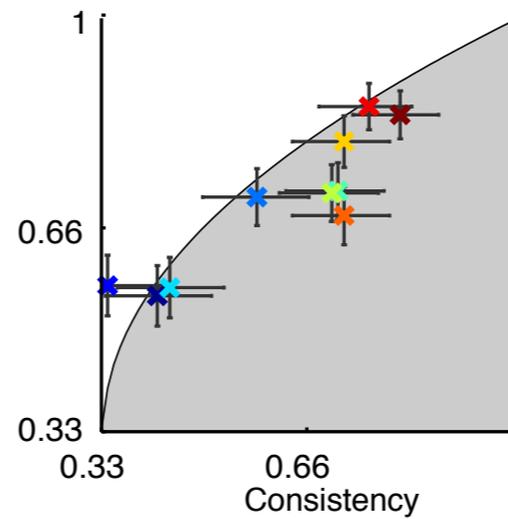
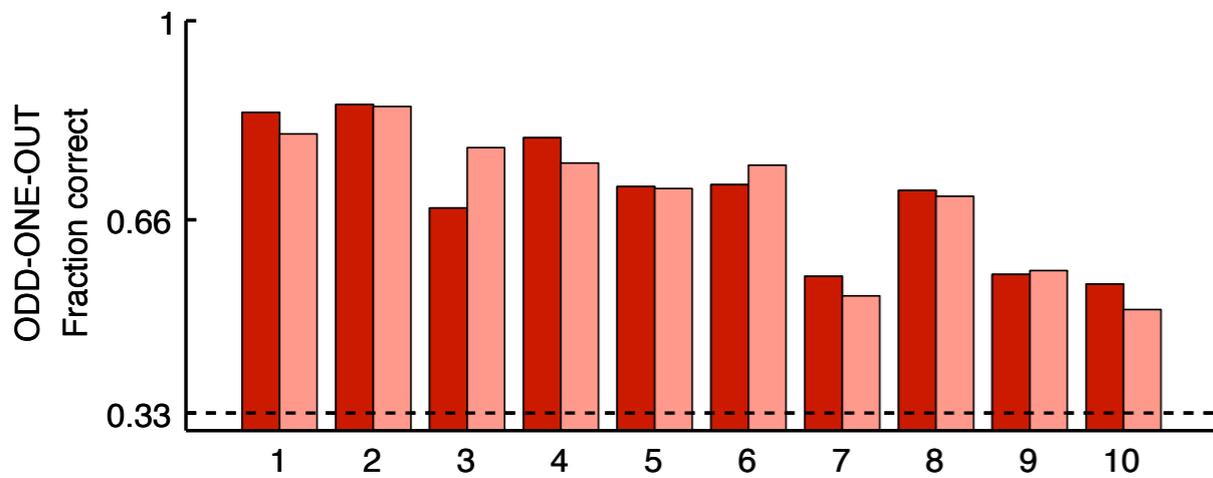
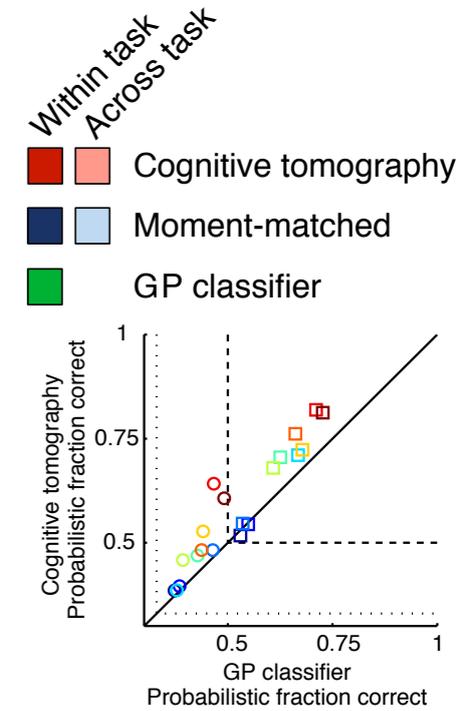
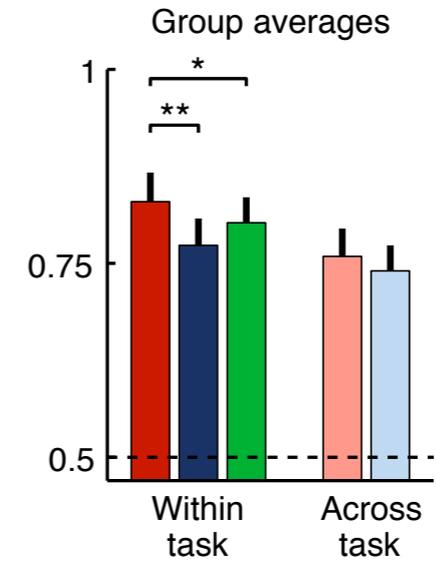
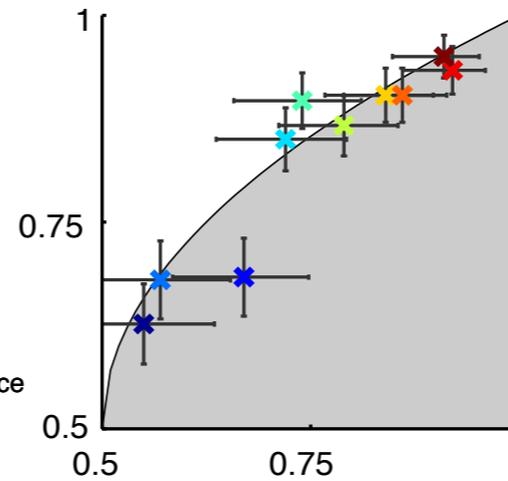
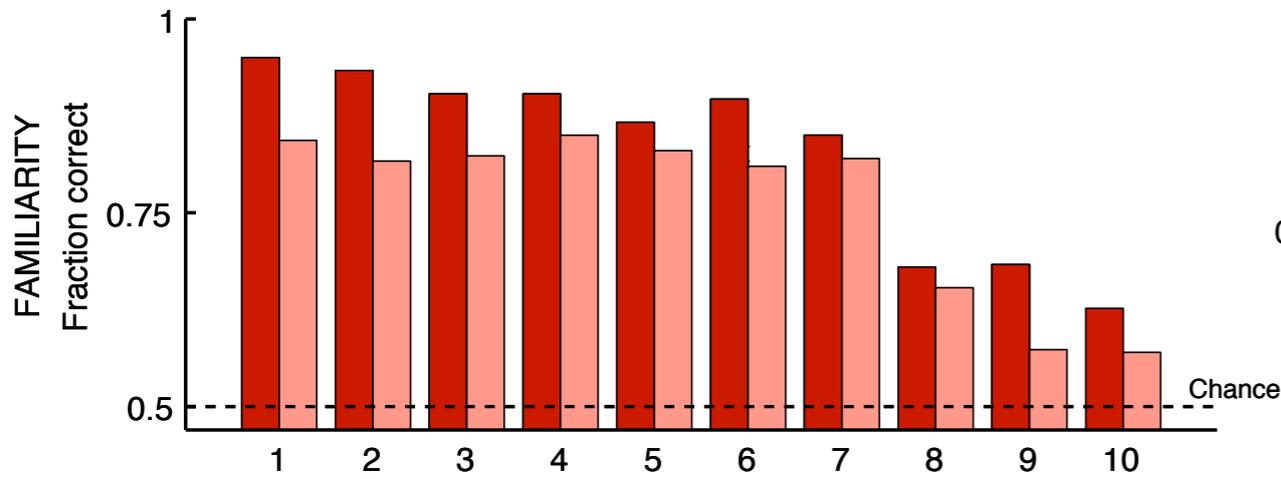
PREDICTING BEHAVIOUR

well above chance
both within and across tasks

close to theoretical
upper bound

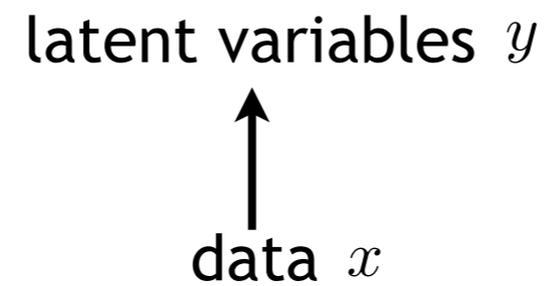
better than
controls

Individual subjects

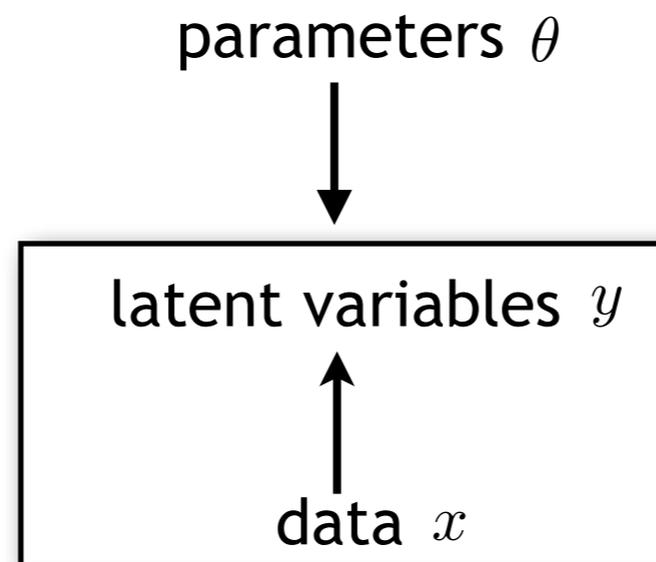


Houlsby et al, Curr Biol 2013

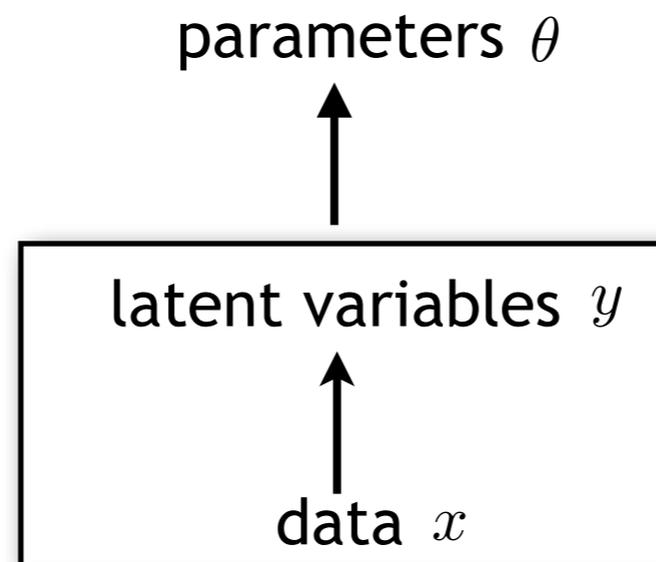
PROBABILISTIC INFERENCE AND LEARNING



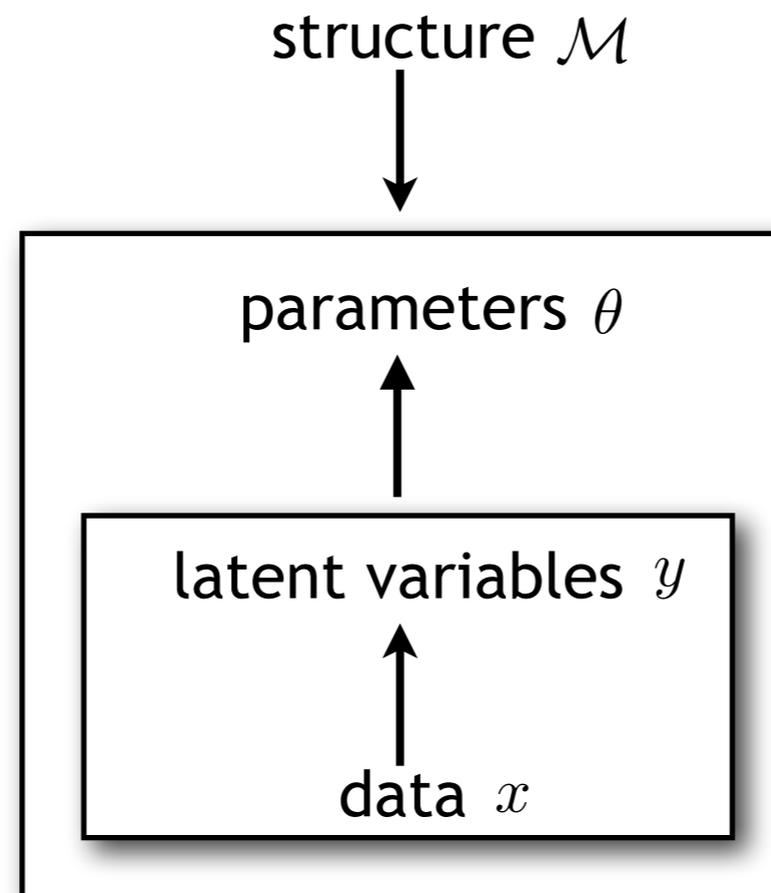
PROBABILISTIC INFERENCE AND LEARNING



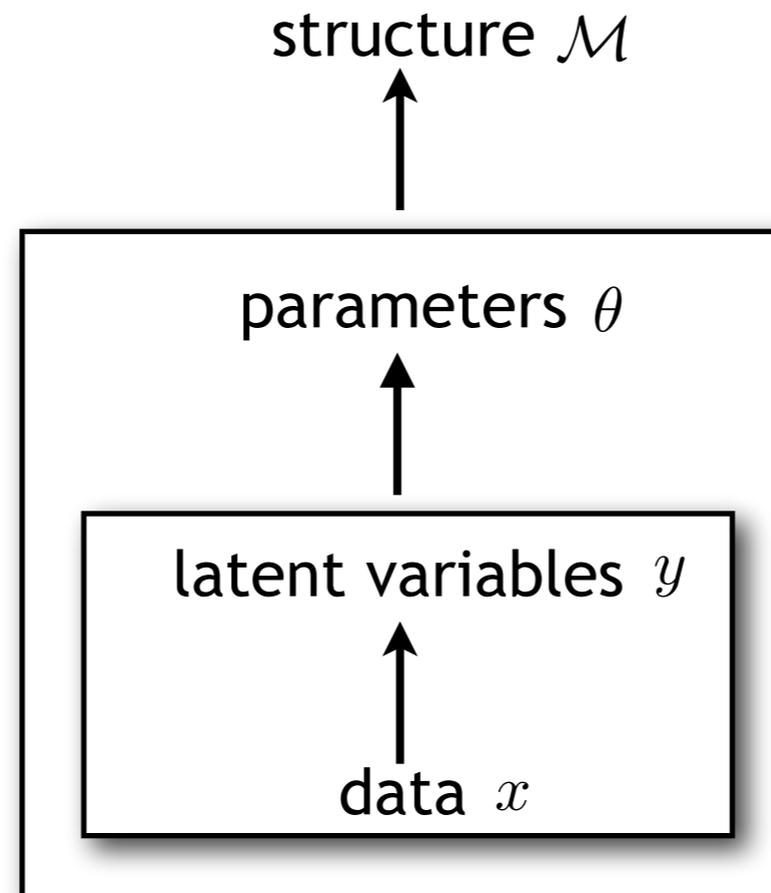
PROBABILISTIC INFERENCE AND LEARNING



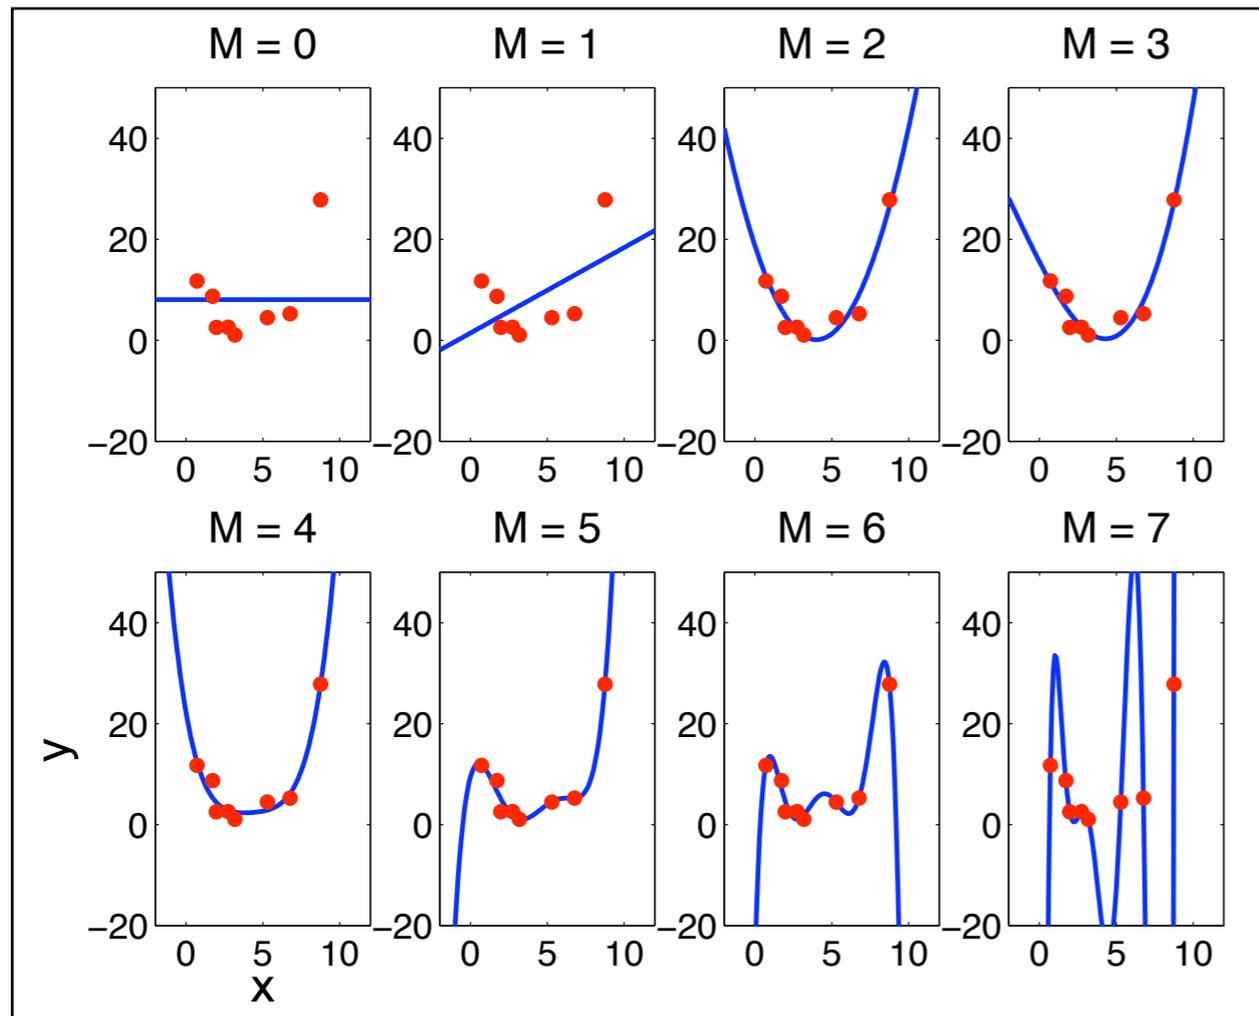
PROBABILISTIC INFERENCE AND LEARNING



PROBABILISTIC INFERENCE AND LEARNING

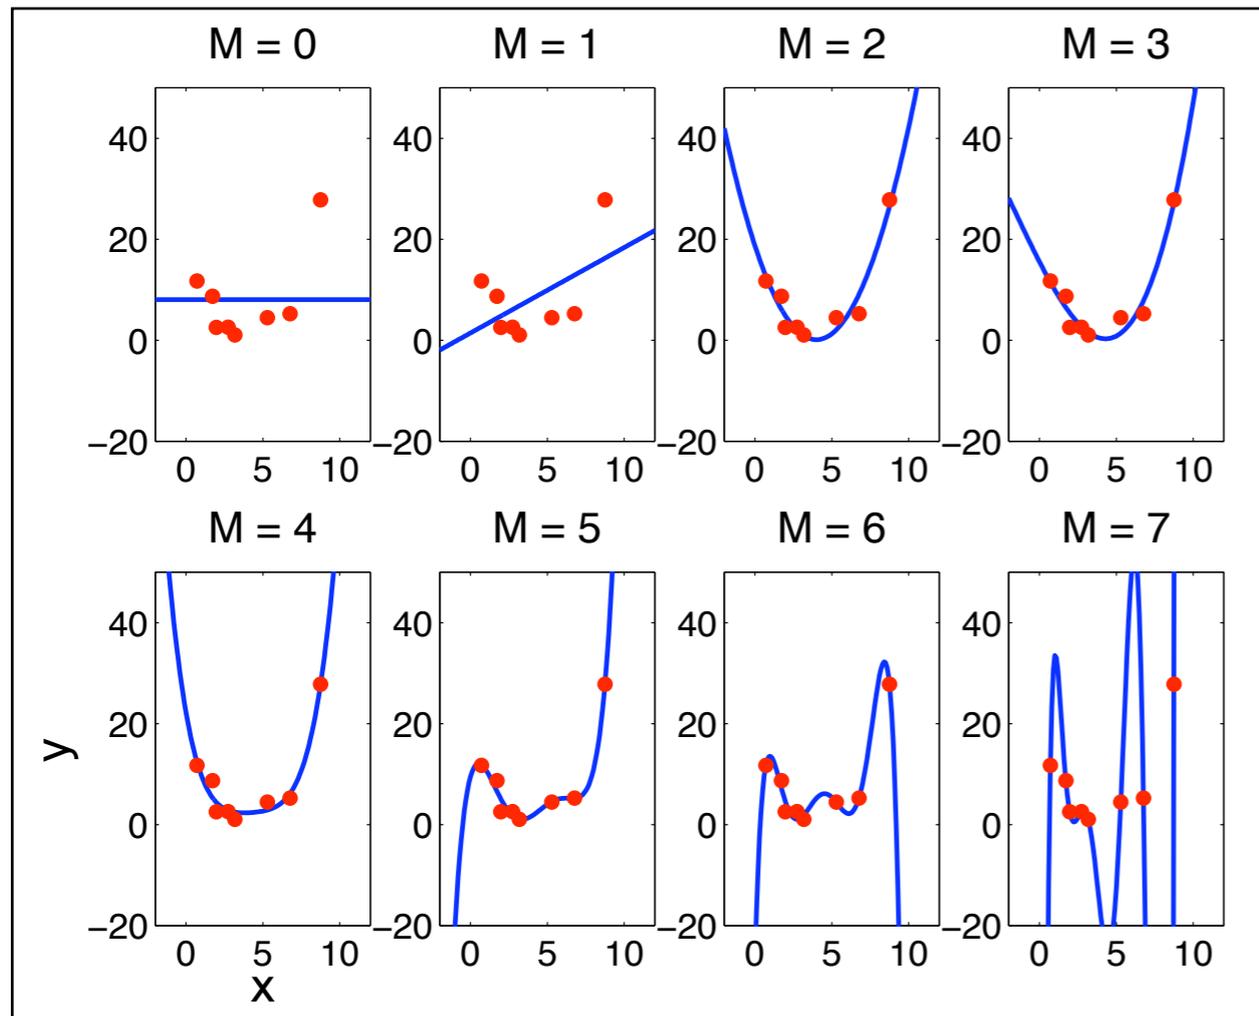


MODEL COMPARISON



courtesy of Zoubin Ghahramani

MODEL COMPARISON



courtesy of Zoubin Ghahramani

- under which model do I get the best fit?

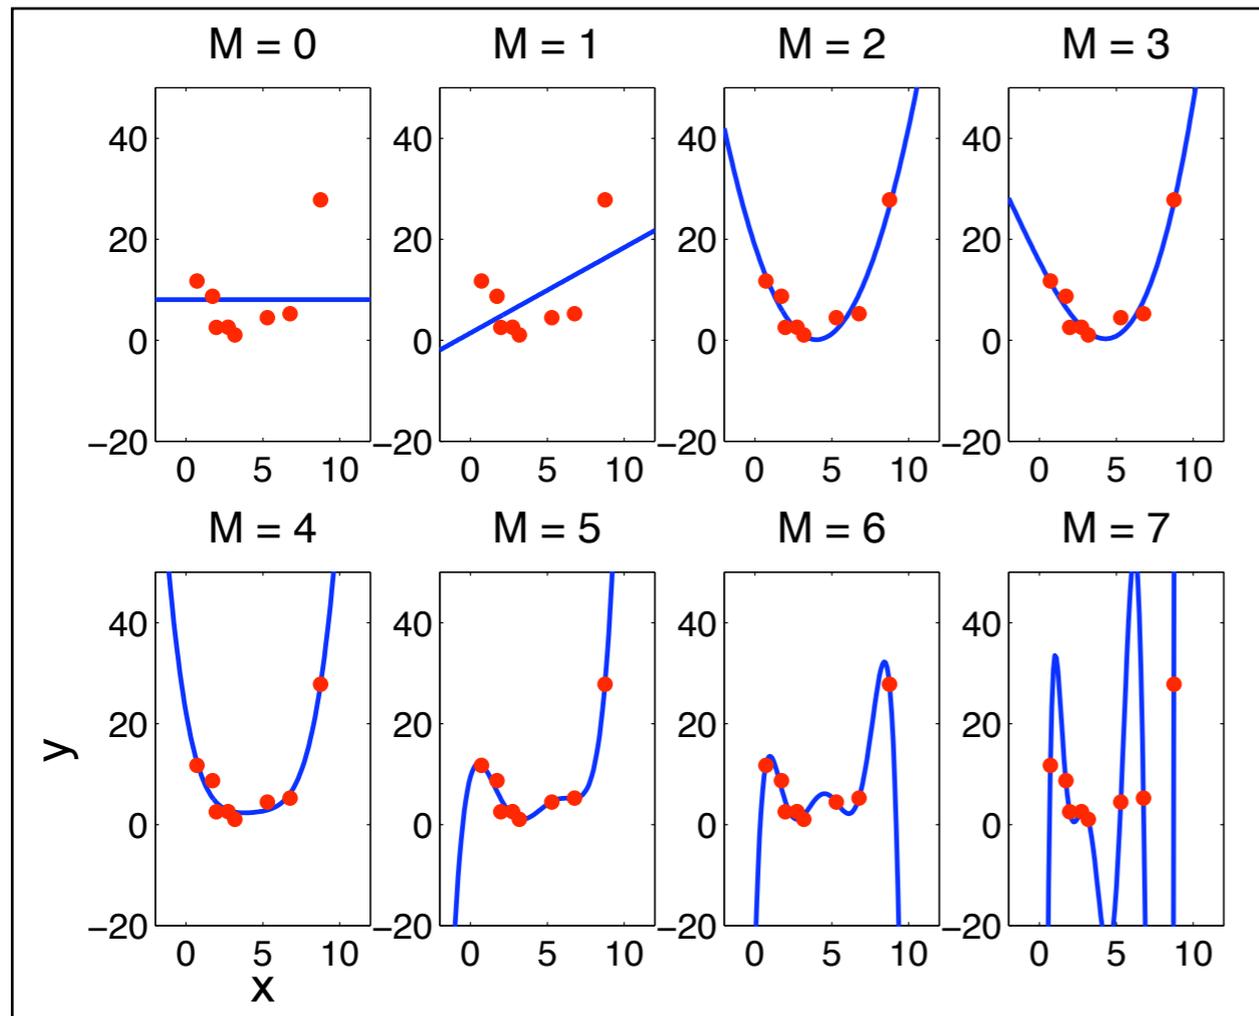
$$P(\mathcal{D} | \hat{\theta}_{ML}, \mathcal{M})$$

↑
parameters

←
model
structure

what is the likelihood of the model with *the best* parameters?

MODEL COMPARISON



courtesy of Zoubin Ghahramani

- under which model do I get the best fit?

$$P(\mathcal{D} | \hat{\theta}_{ML}, \mathcal{M})$$

↑
parameters

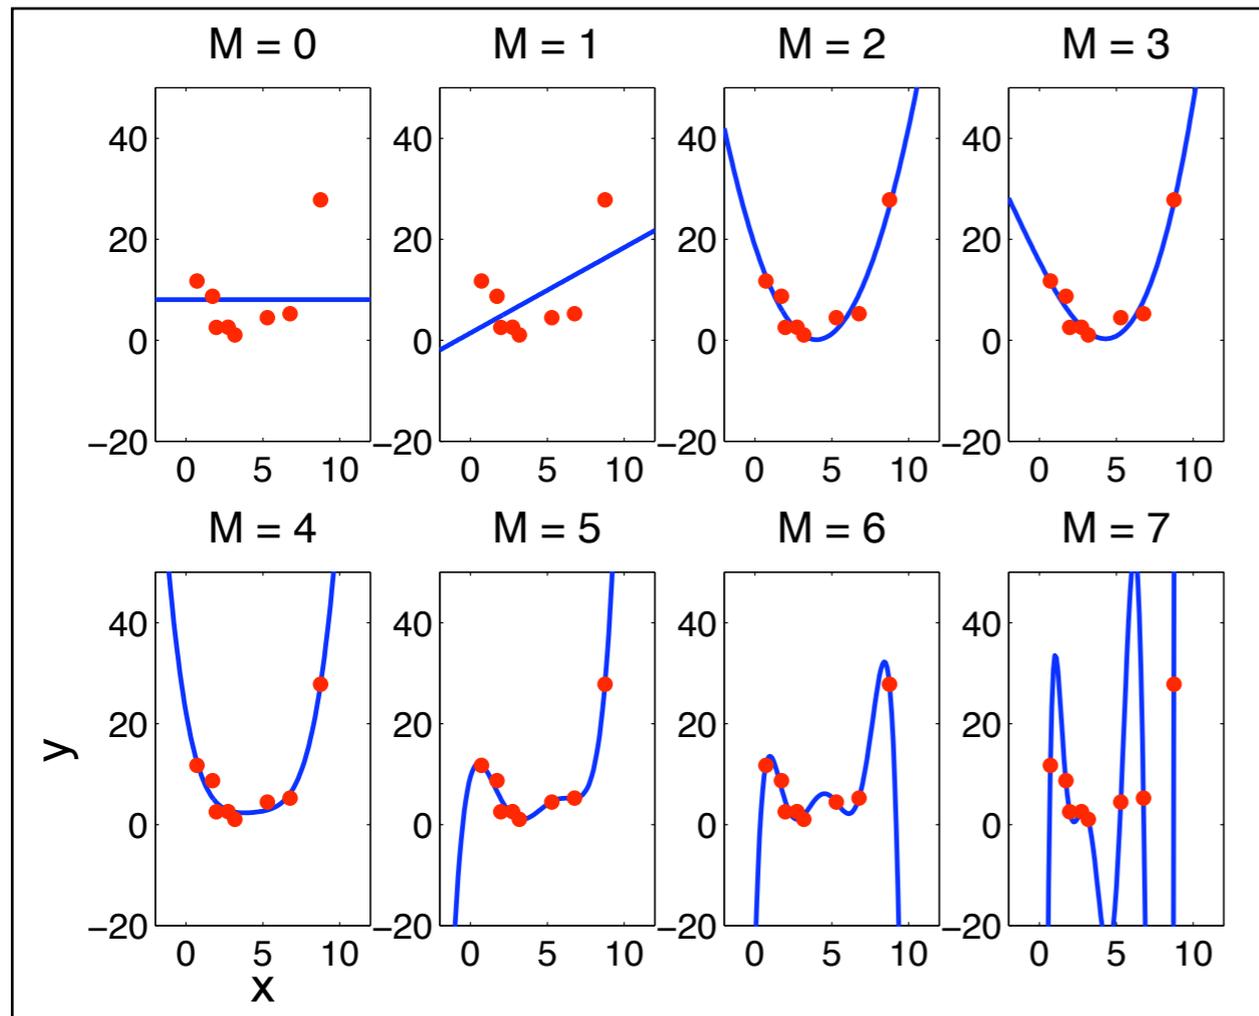
←
model structure

what is the likelihood of the model with *the best* parameters?



overfitting!

MODEL COMPARISON



courtesy of Zoubin Ghahramani

- under which model do I get the best fit?

$$P(\mathcal{D} | \hat{\theta}_{ML}, \mathcal{M})$$

↑
↑

parameters
model structure

what is the likelihood of the model with *the best* parameters?



overfitting!

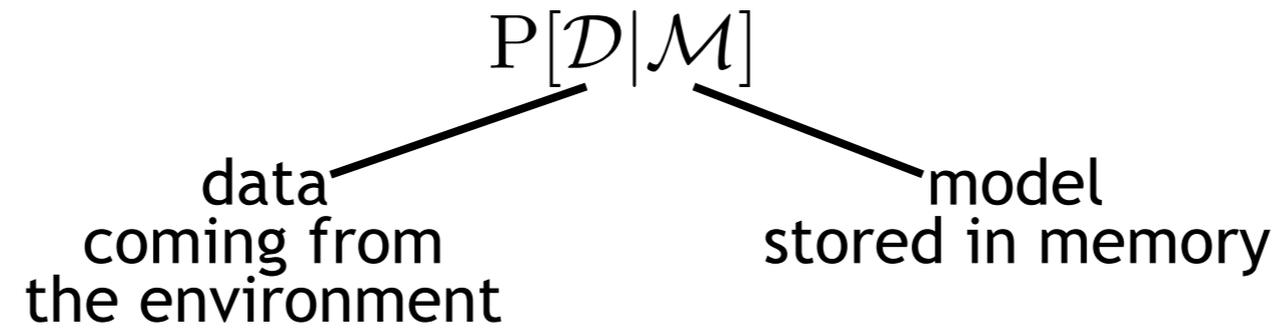
- which model has the highest likelihood?

$$P(\mathcal{D} | \mathcal{M}) = \sum_{\theta} P(\mathcal{D} | \theta, \mathcal{M}) P(\theta | \mathcal{M})$$

what is the average likelihood of the model with *randomly chosen* parameters?

BAYESIAN MODEL SELECTION

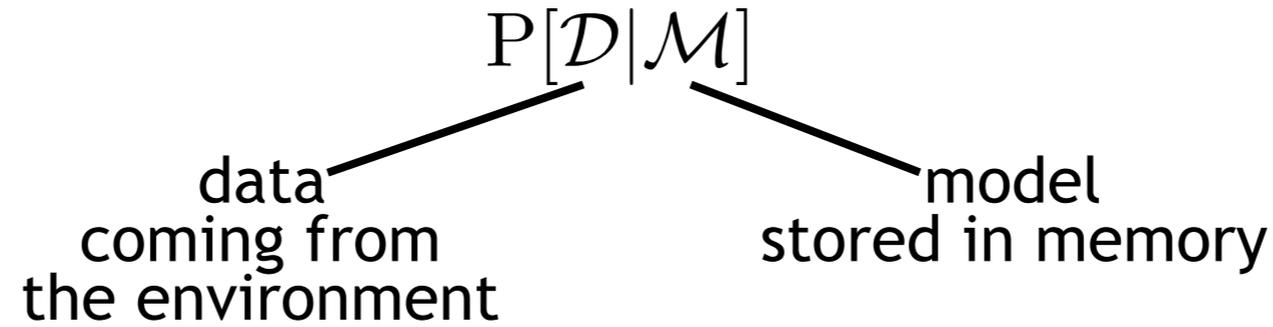
a model defines a probability distribution
over data sets



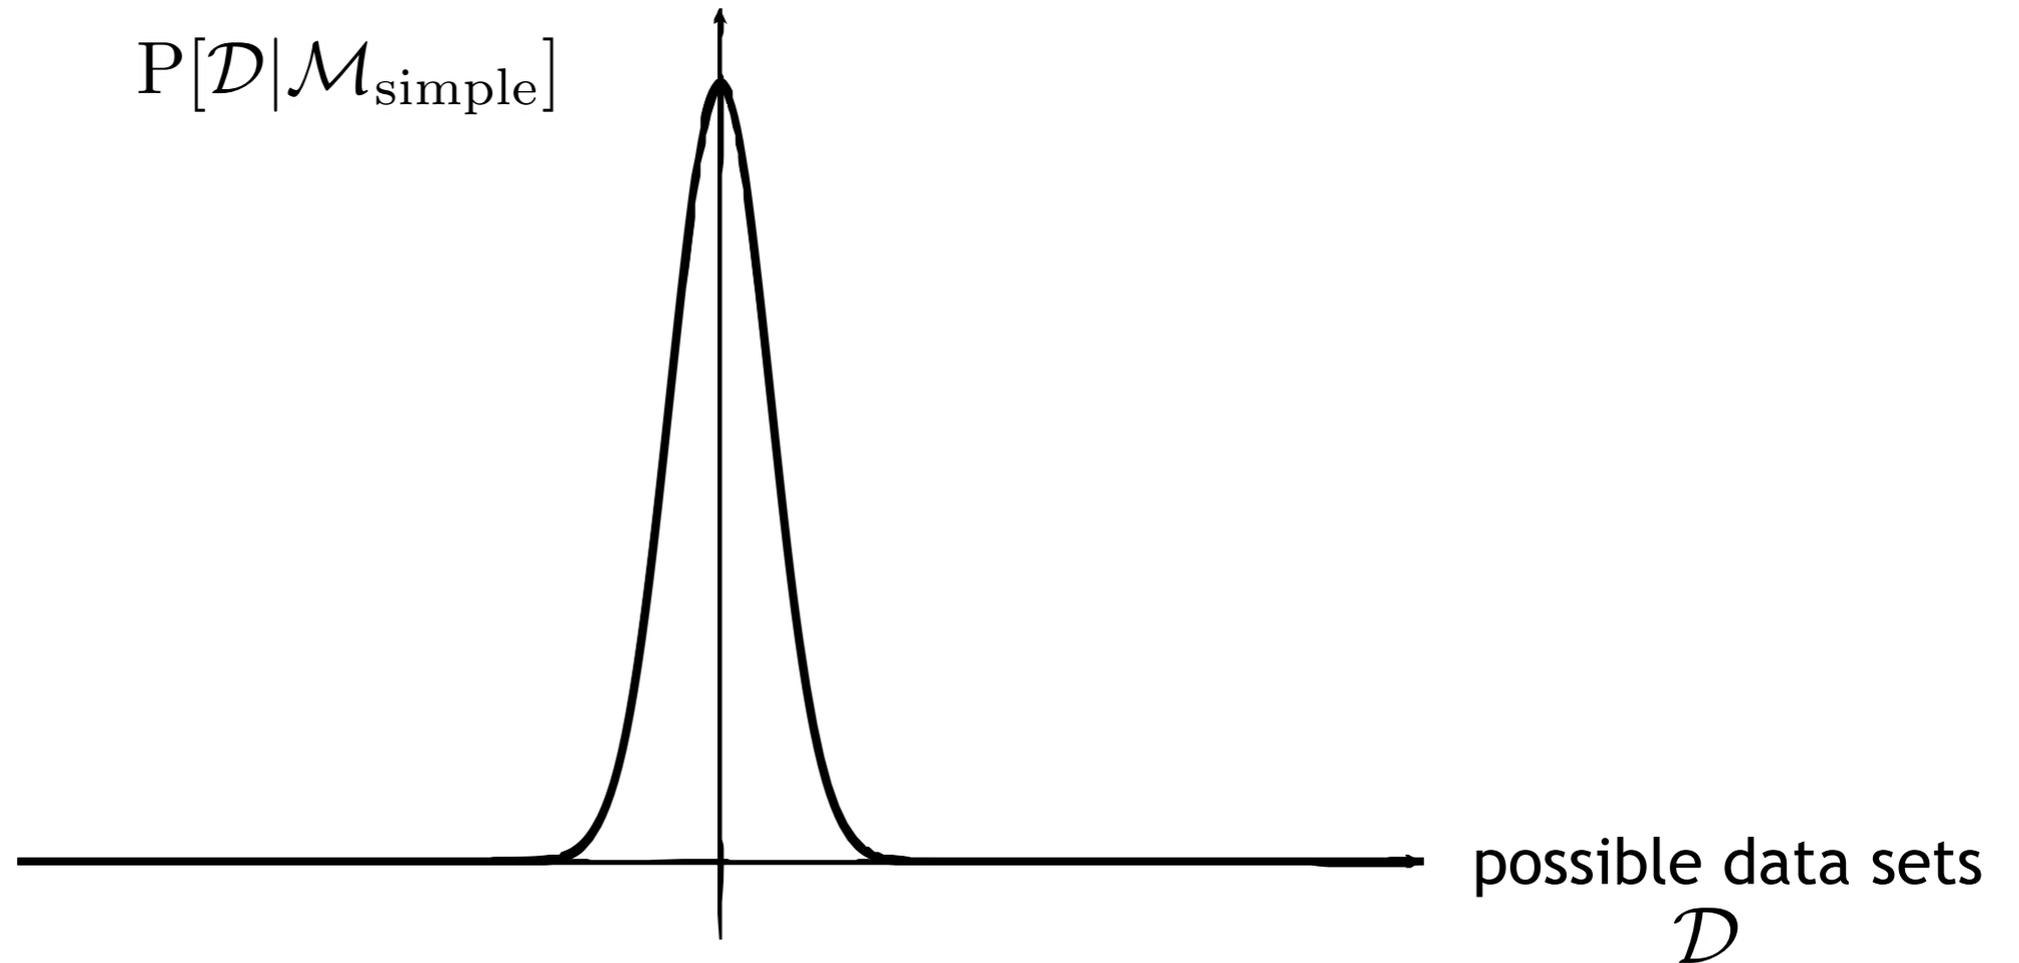
Rev. Bayes

BAYESIAN MODEL SELECTION

a model defines a probability distribution
over data sets



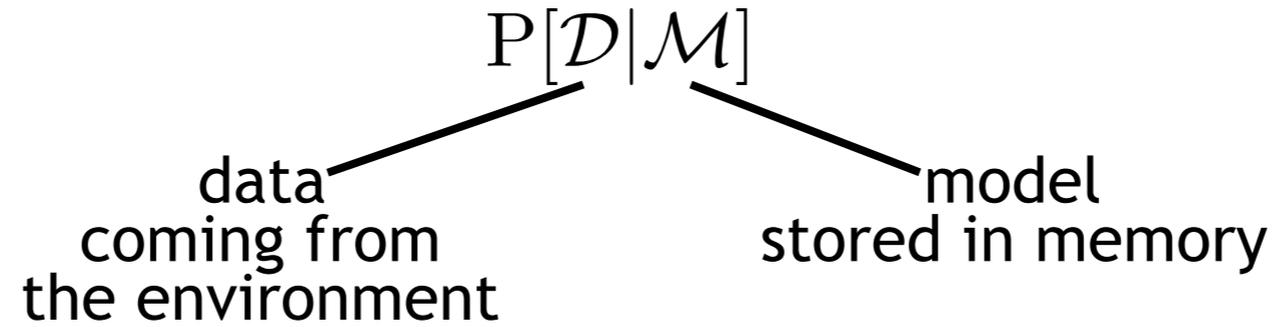
$$P[\mathcal{D}|\mathcal{M}_{\text{simple}}]$$



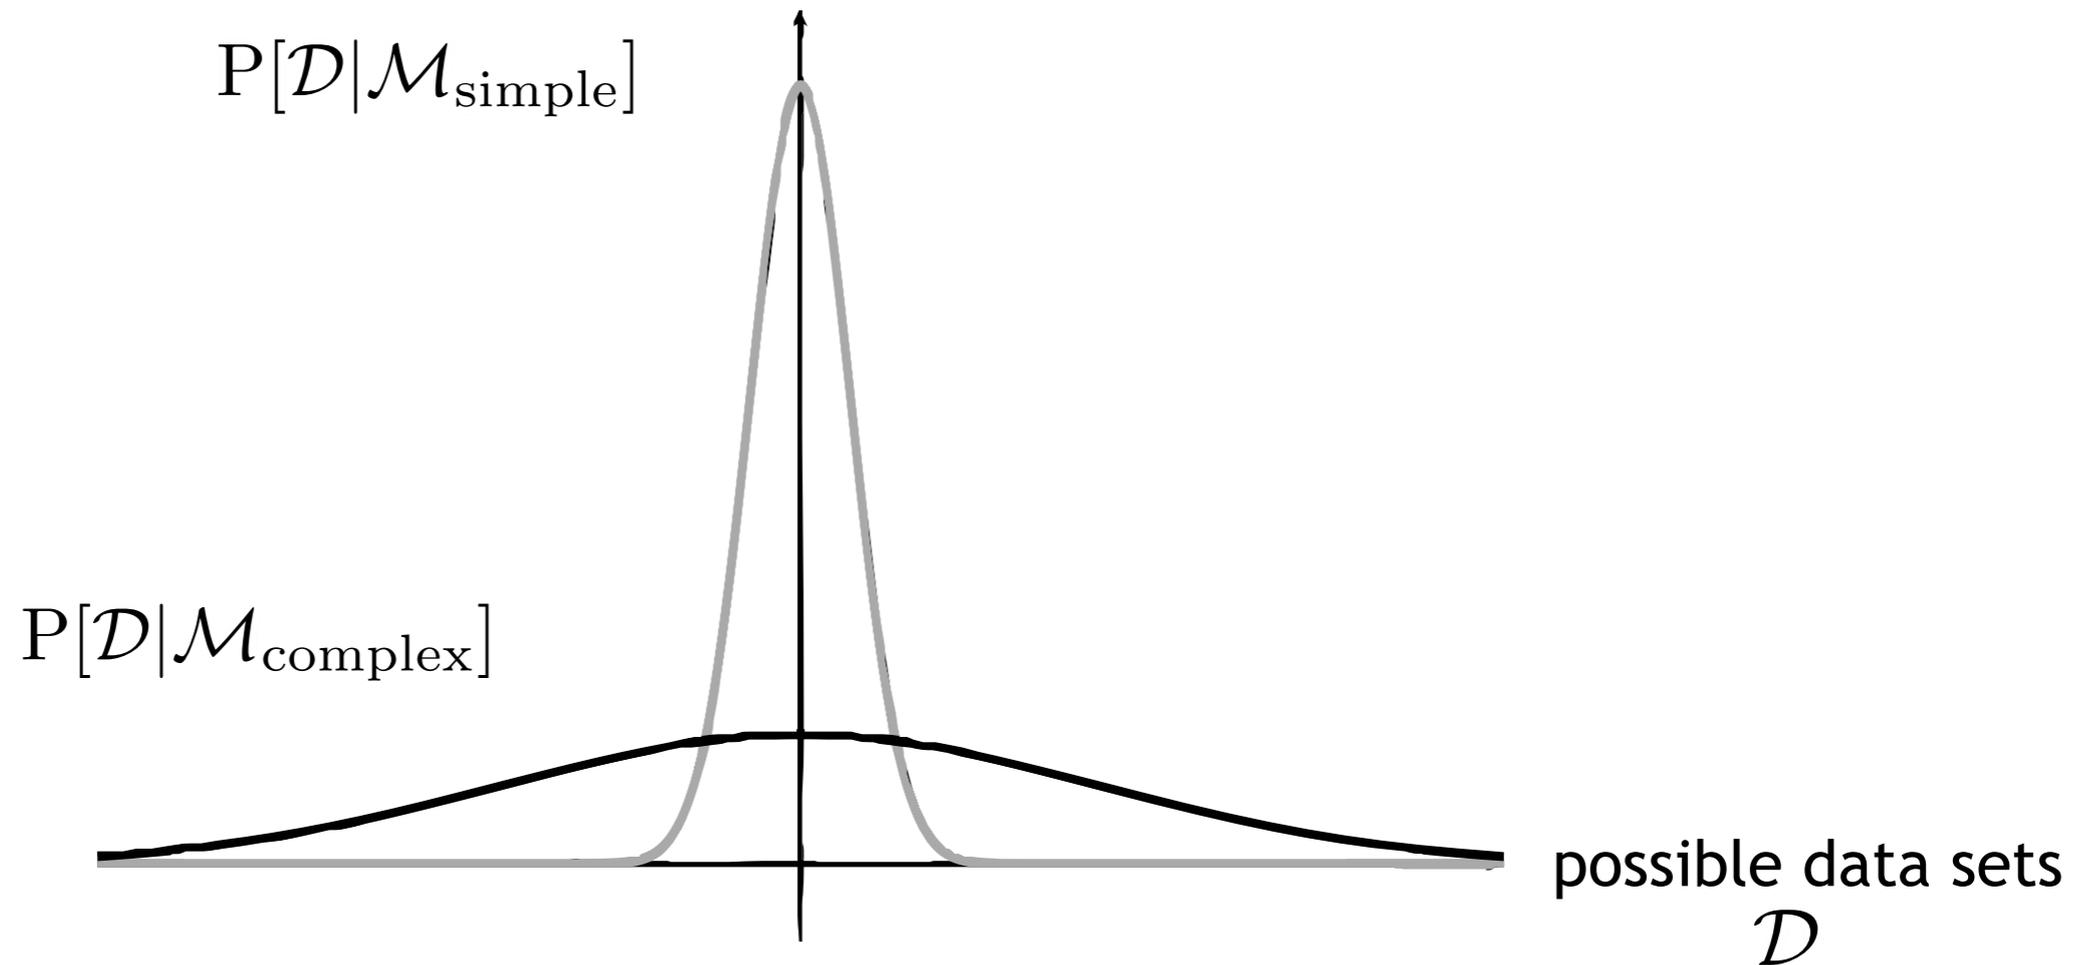
Rev. Bayes

BAYESIAN MODEL SELECTION

a model defines a probability distribution over data sets

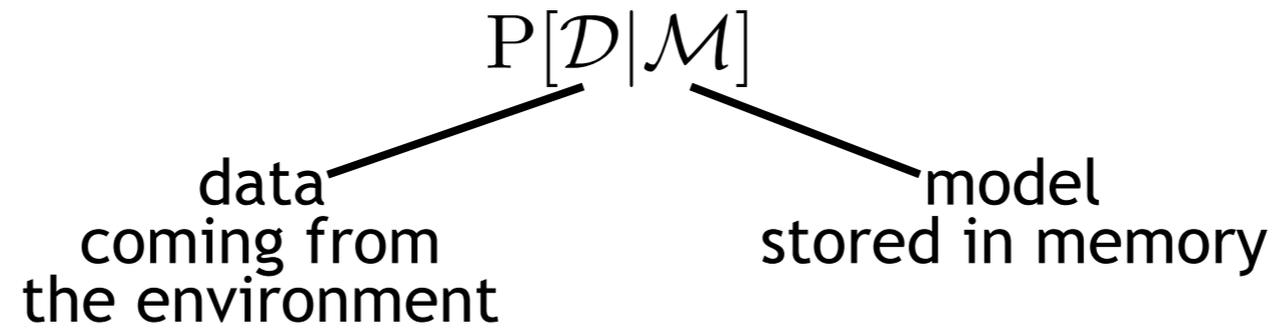


Rev. Bayes

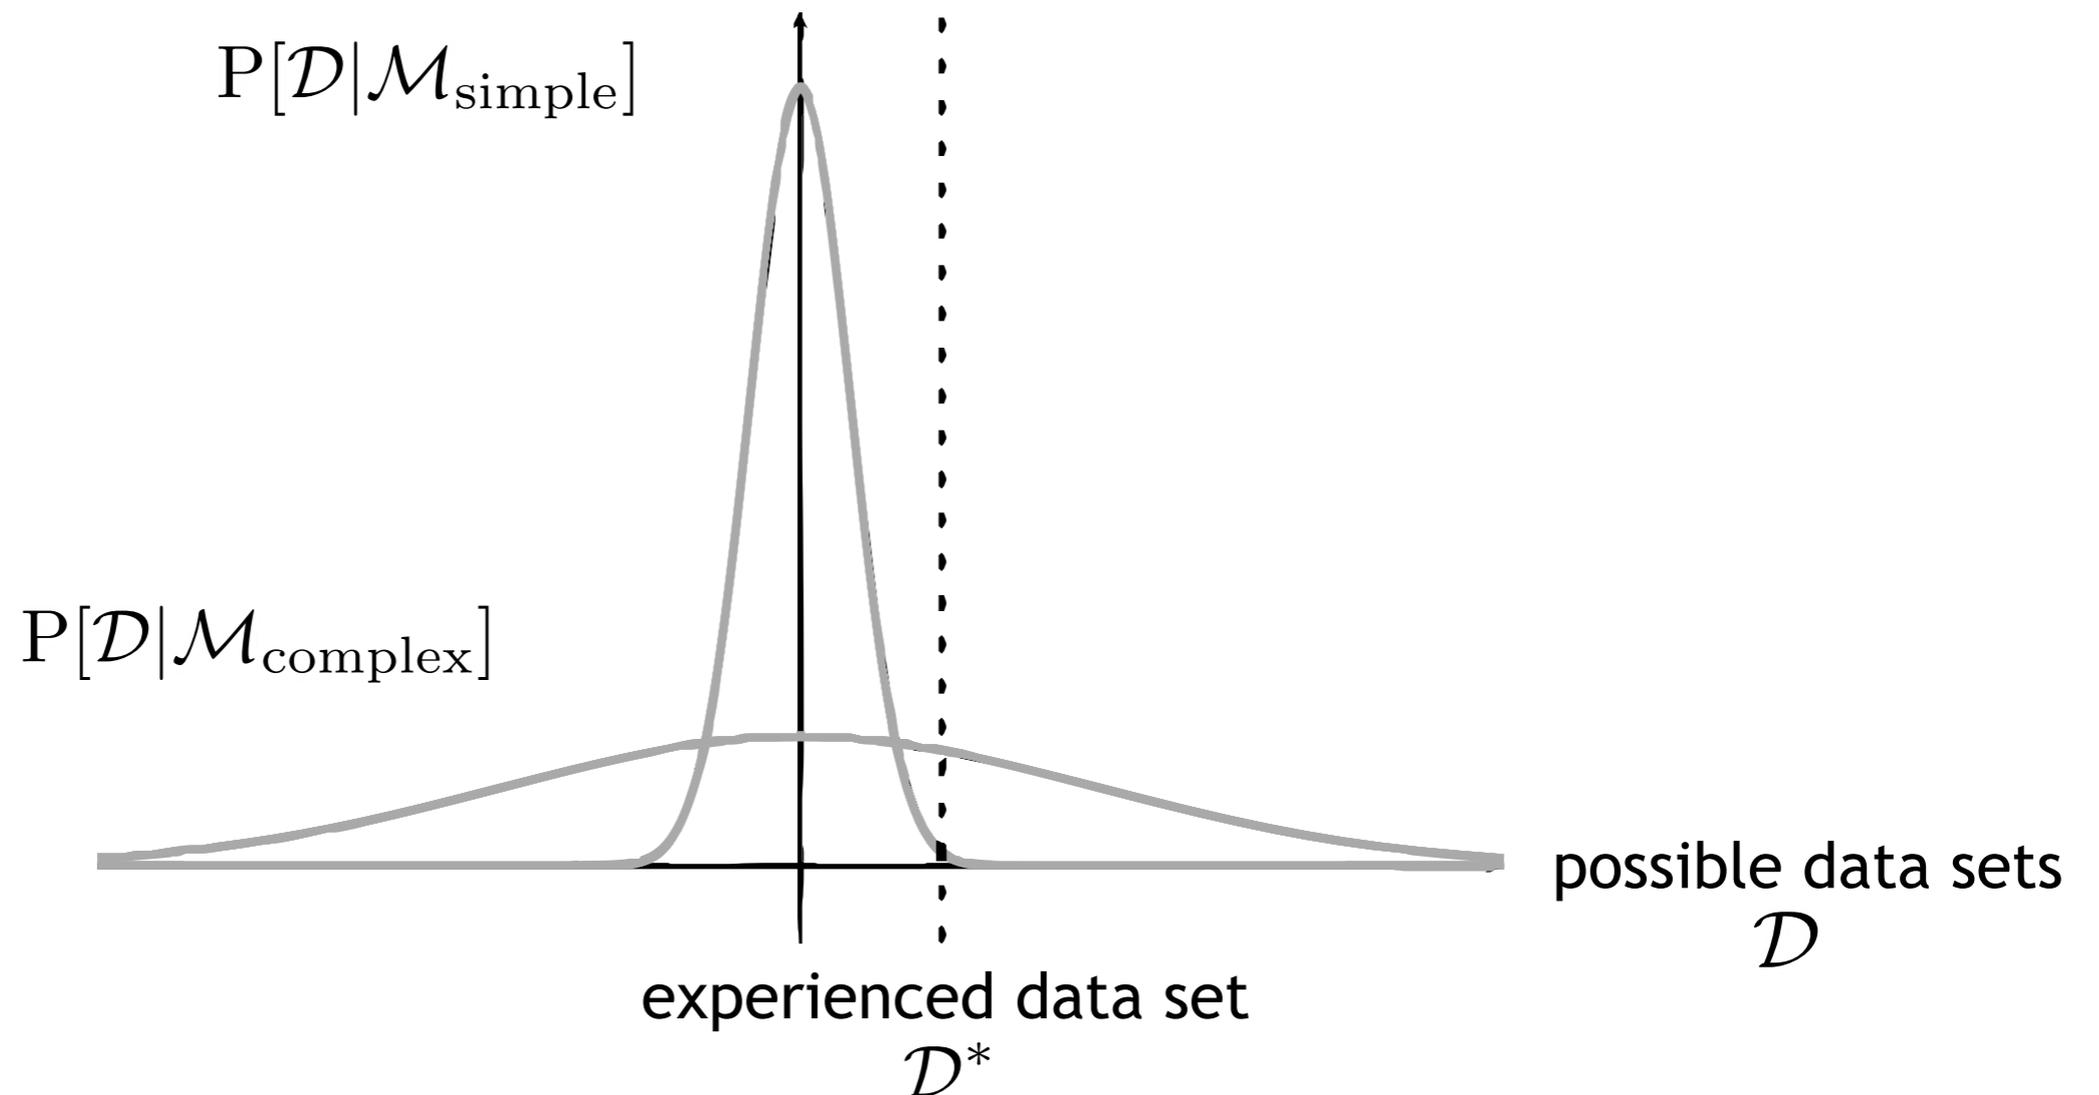


BAYESIAN MODEL SELECTION

a model defines a probability distribution over data sets

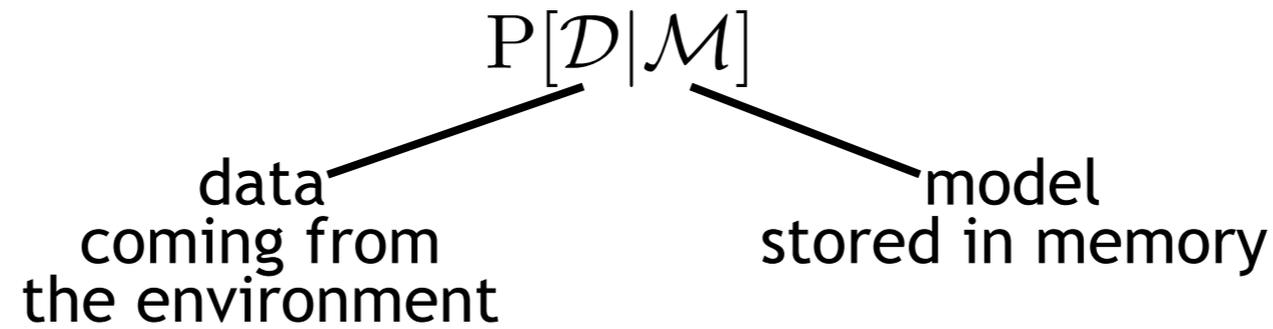


Rev. Bayes

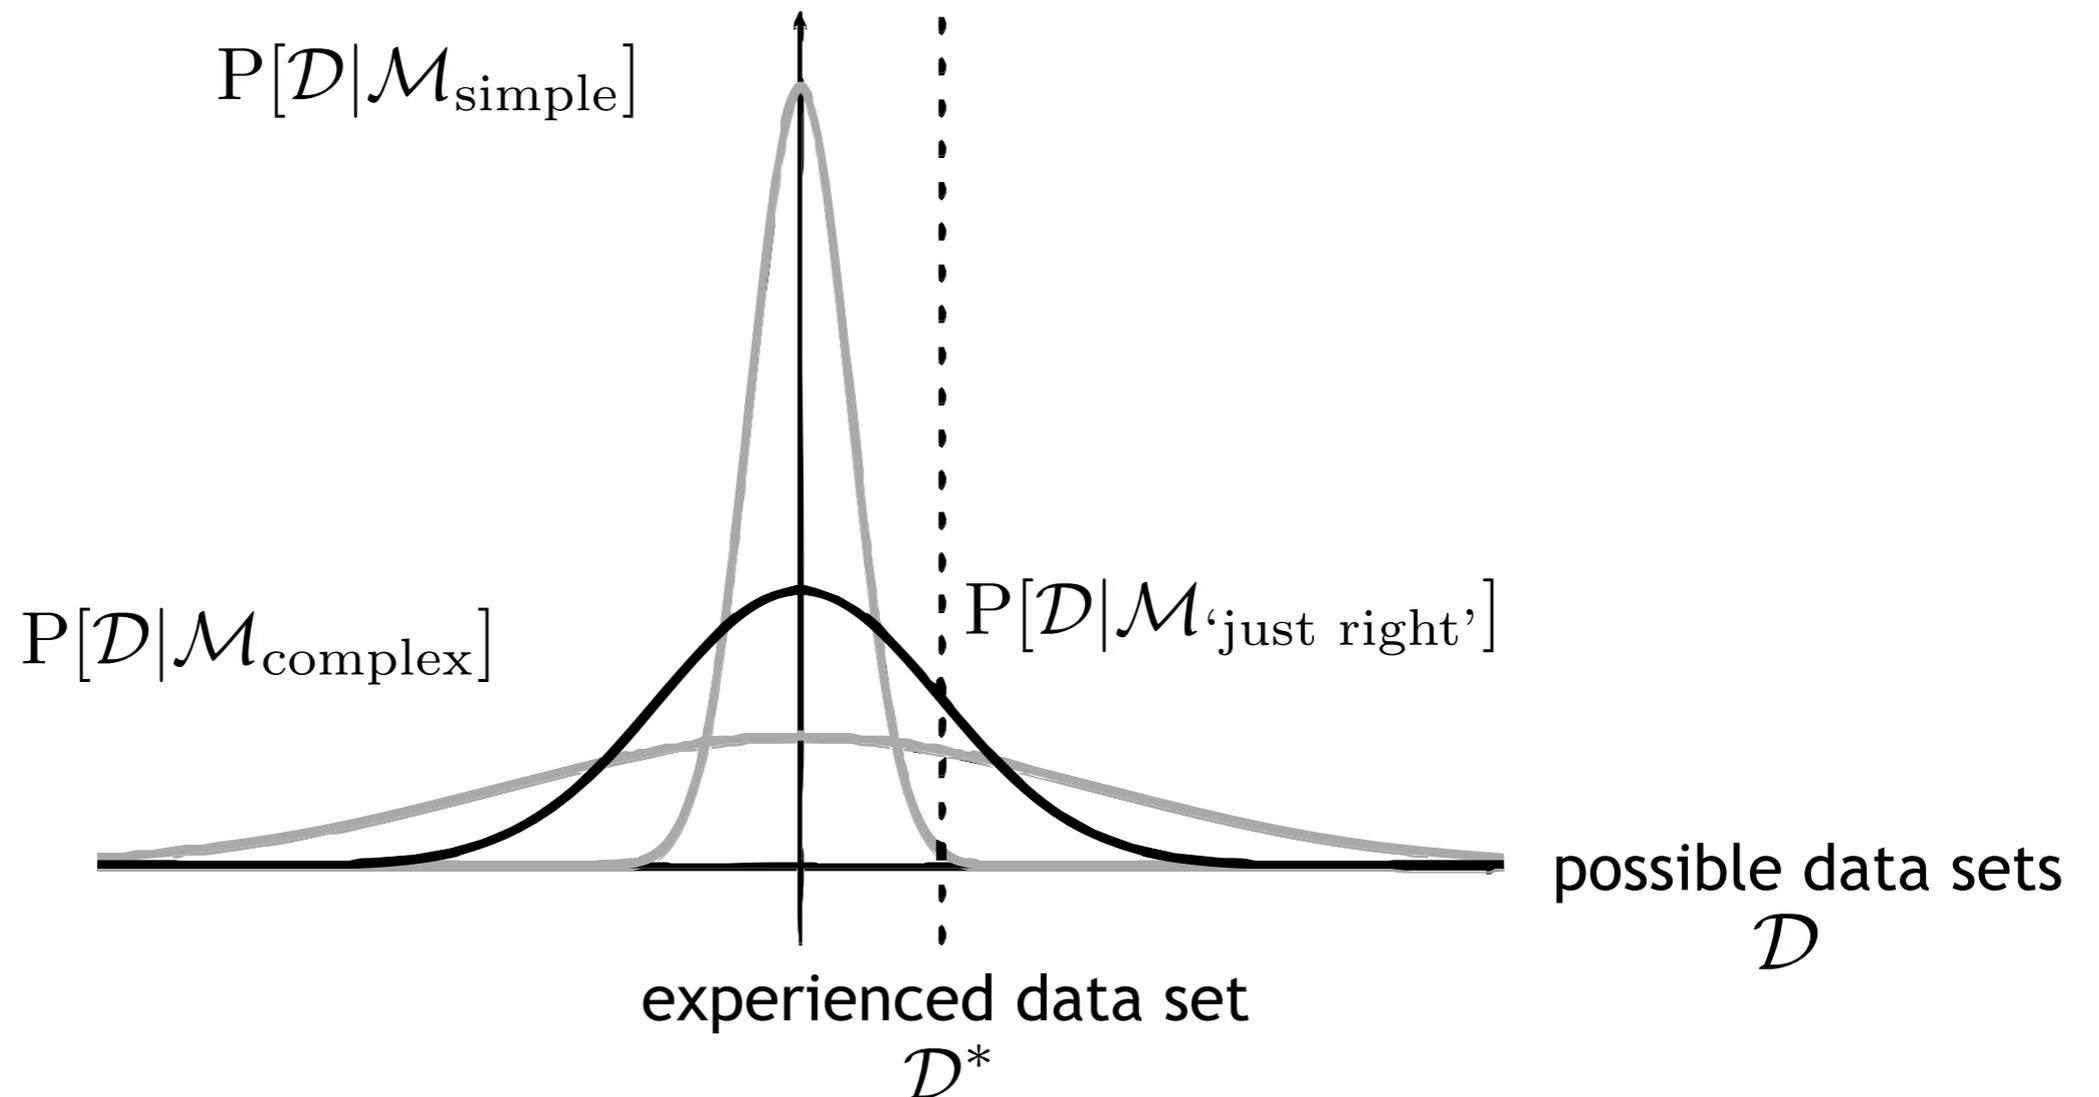


BAYESIAN MODEL SELECTION

a model defines a probability distribution over data sets



Rev. Bayes



BAYESIAN MODEL SELECTION

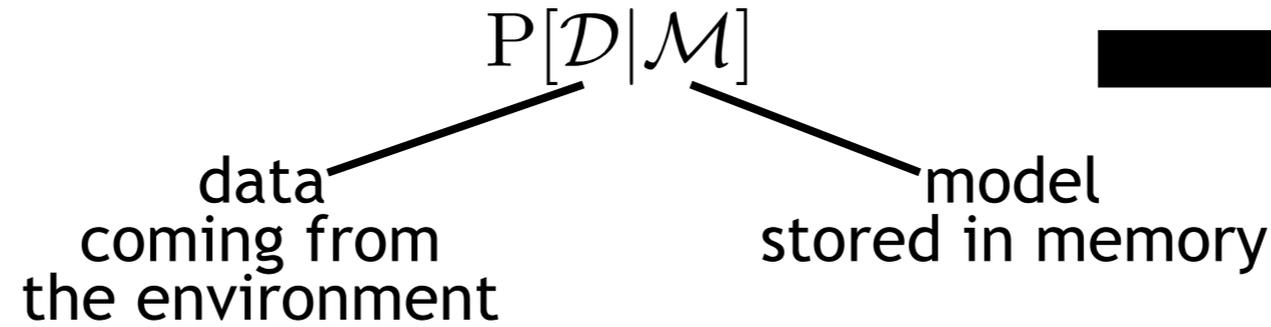
a model defines a probability distribution over data sets



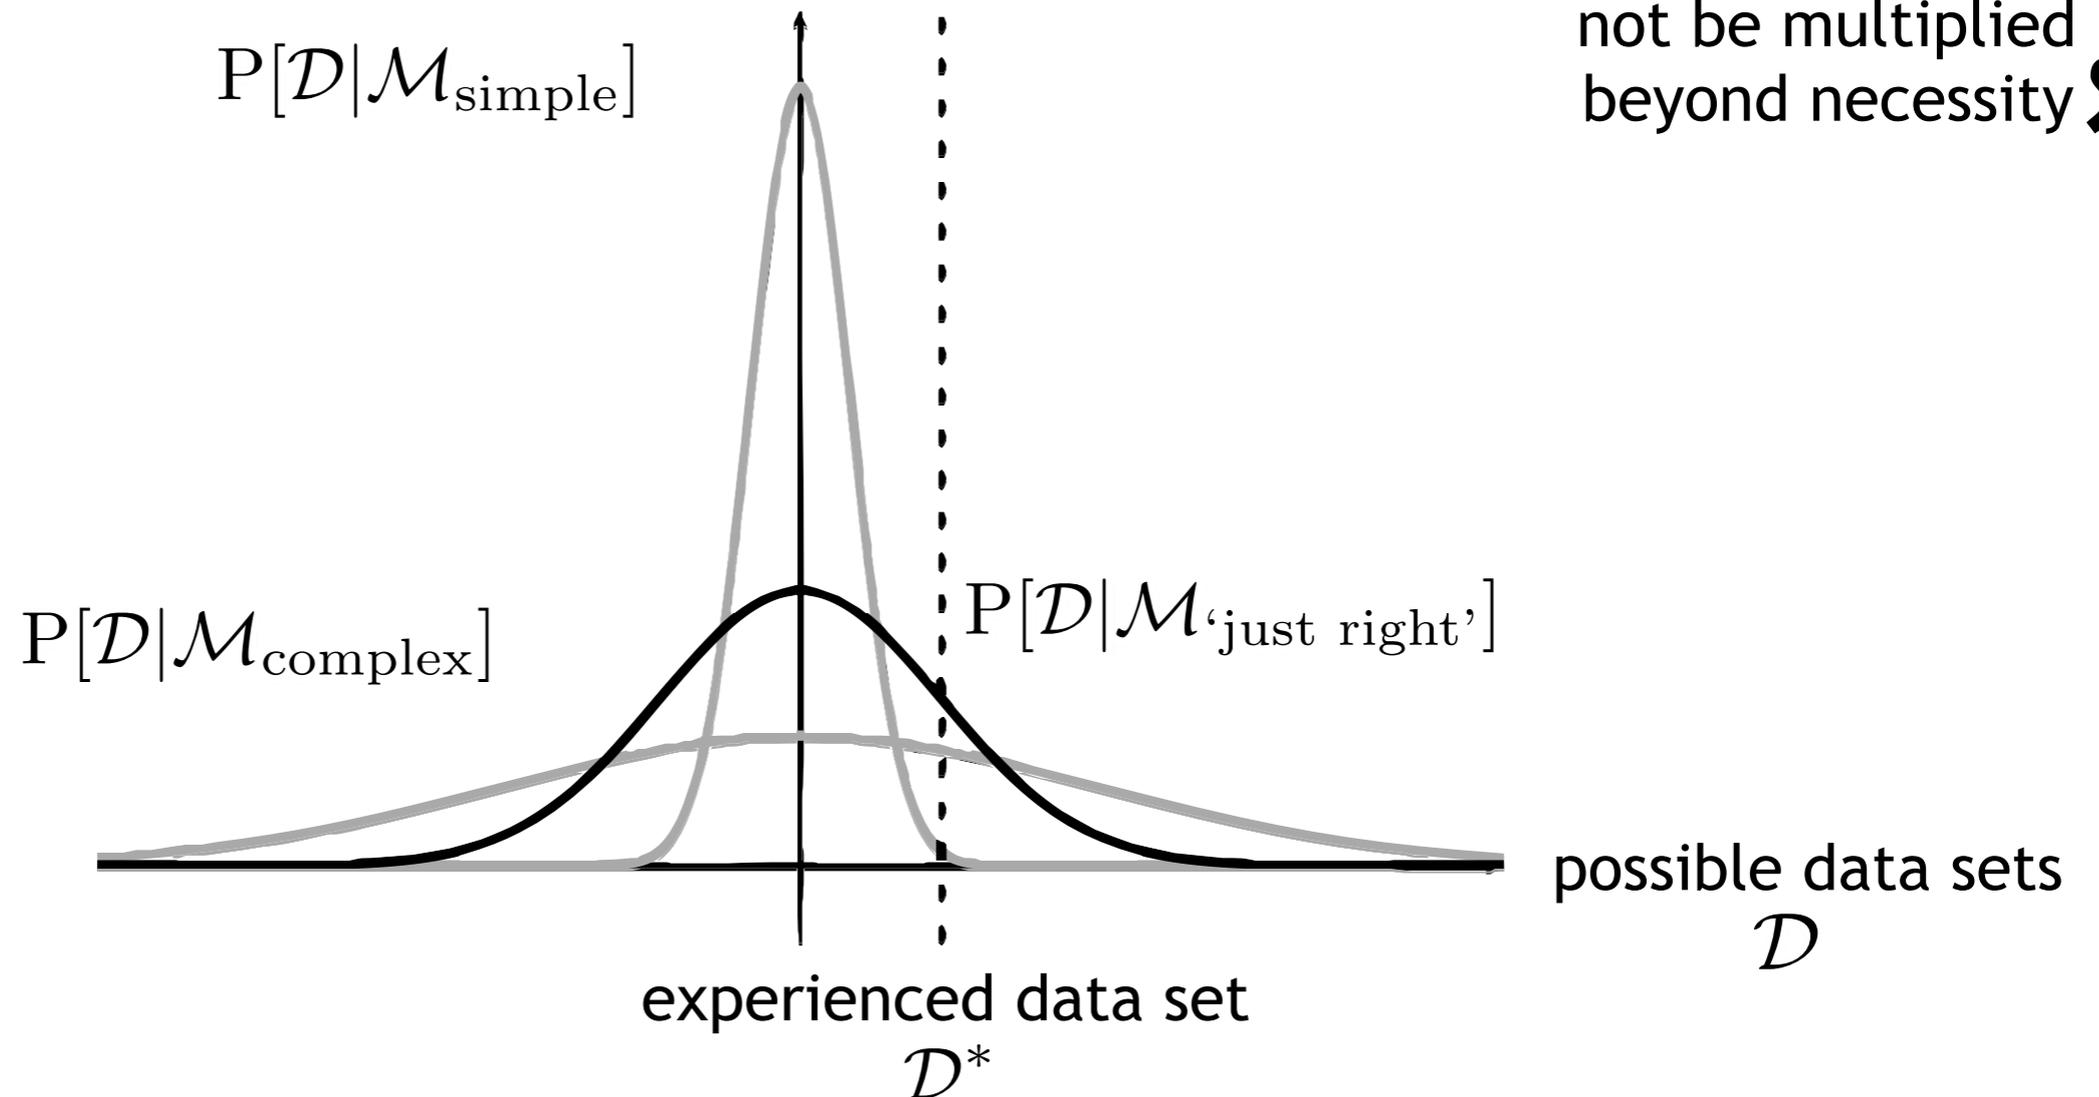
Rev. Bayes



Occam's razor



‘ entities should not be multiplied beyond necessity ’



VISUAL PATTERN LEARNING

VISUAL PATTERN LEARNING

familiarization

‘pay attention’

VISUAL PATTERN LEARNING

familiarization

‘pay attention’

VISUAL PATTERN LEARNING

familiarization

‘pay attention’

test

‘which one looks more familiar?’

VISUAL PATTERN LEARNING

familiarization

‘pay attention’

test

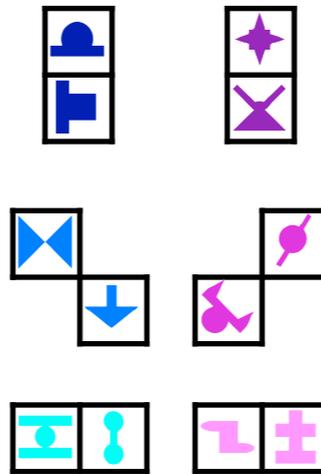
‘which one looks more familiar?’

VISUAL PATTERN LEARNING

familiarization
'pay attention'

inventory

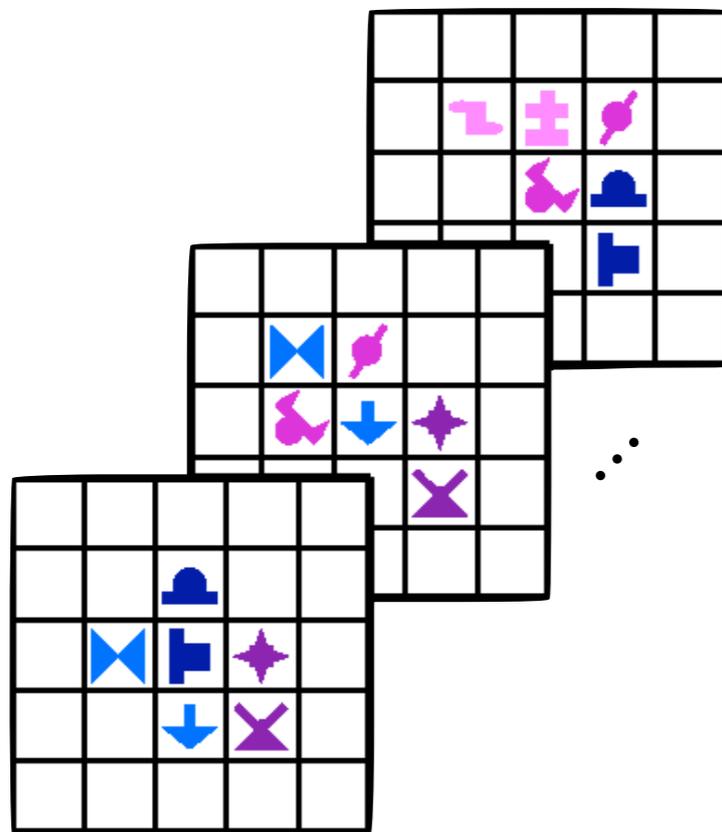
test
'which one looks more familiar?'



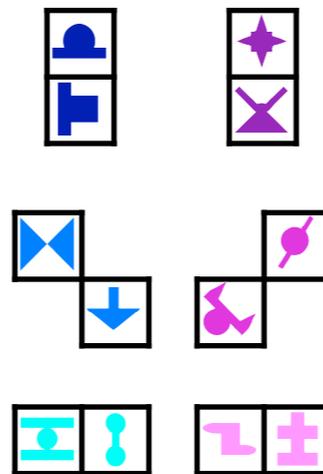
VISUAL PATTERN LEARNING

familiarization

‘pay attention’



inventory



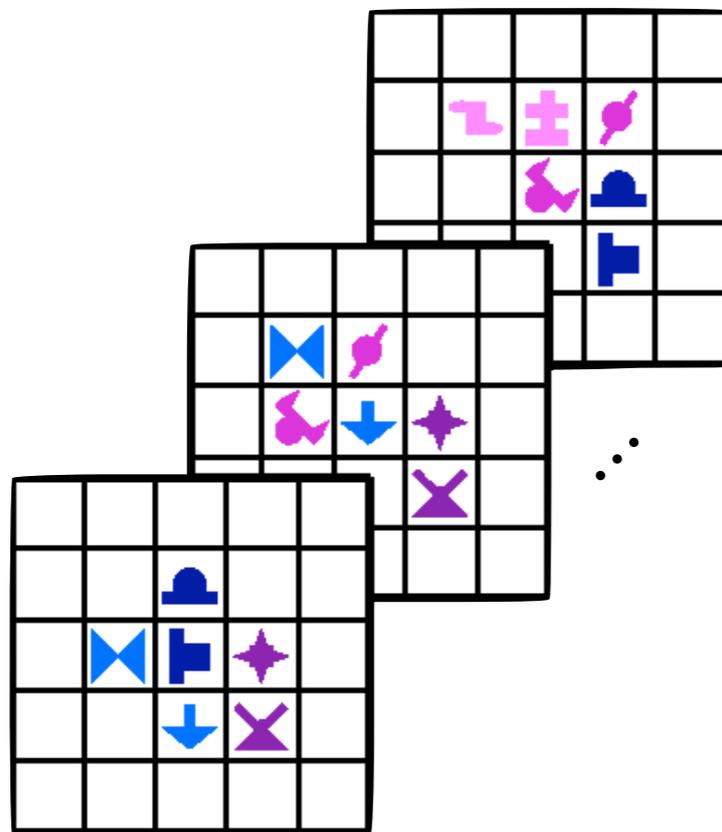
test

‘which one looks more familiar?’

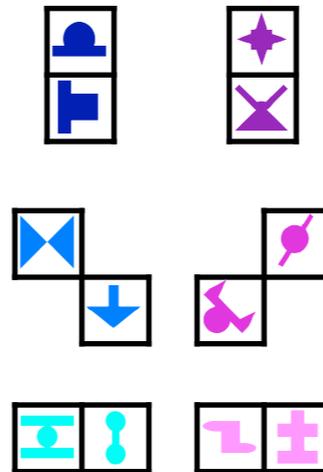
VISUAL PATTERN LEARNING

familiarization

‘pay attention’

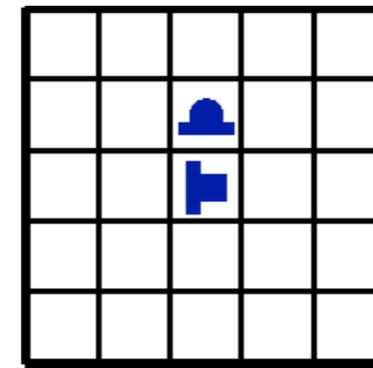


inventory

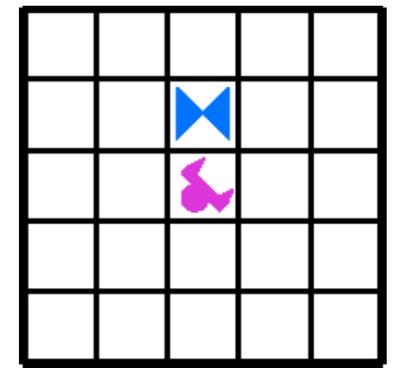


test

‘which one looks more familiar?’



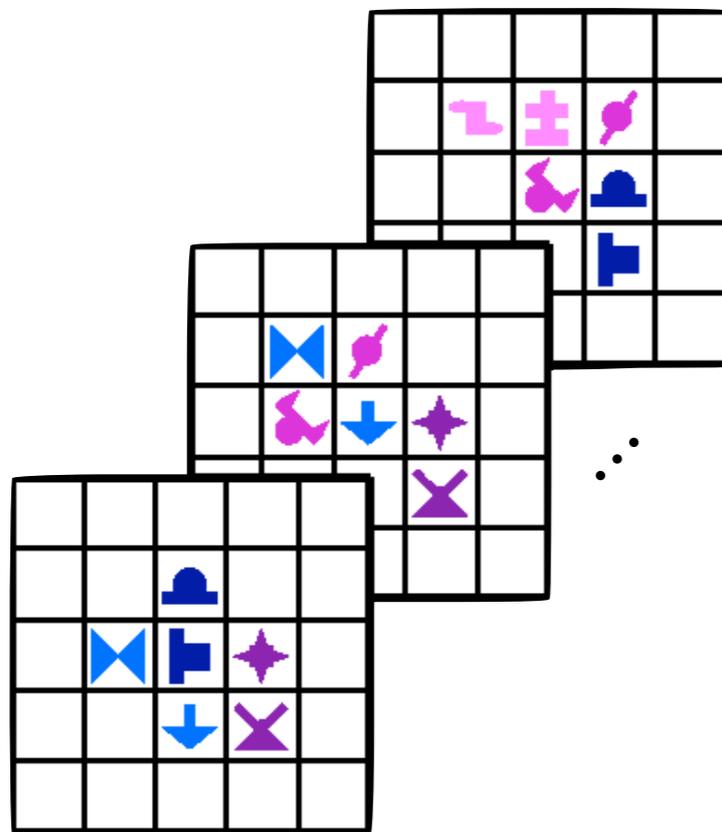
vs.



Orbán & al, PNAS 2008

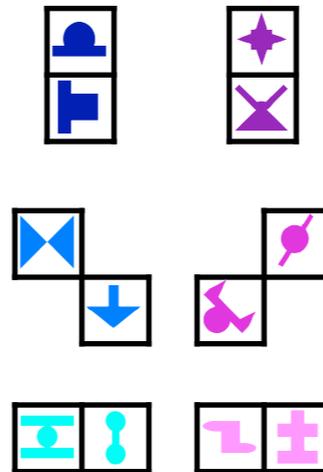
VISUAL PATTERN LEARNING

familiarization
'pay attention'



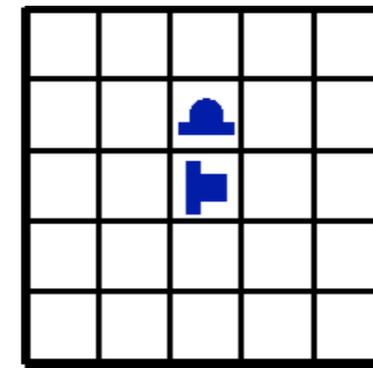
$$\mathcal{D} = \text{scene}_1, \dots, \text{scene}_n$$

inventory

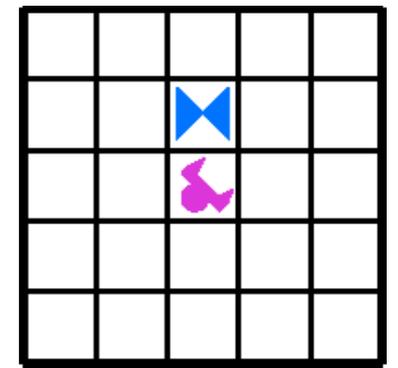


test

'which one looks more familiar?'



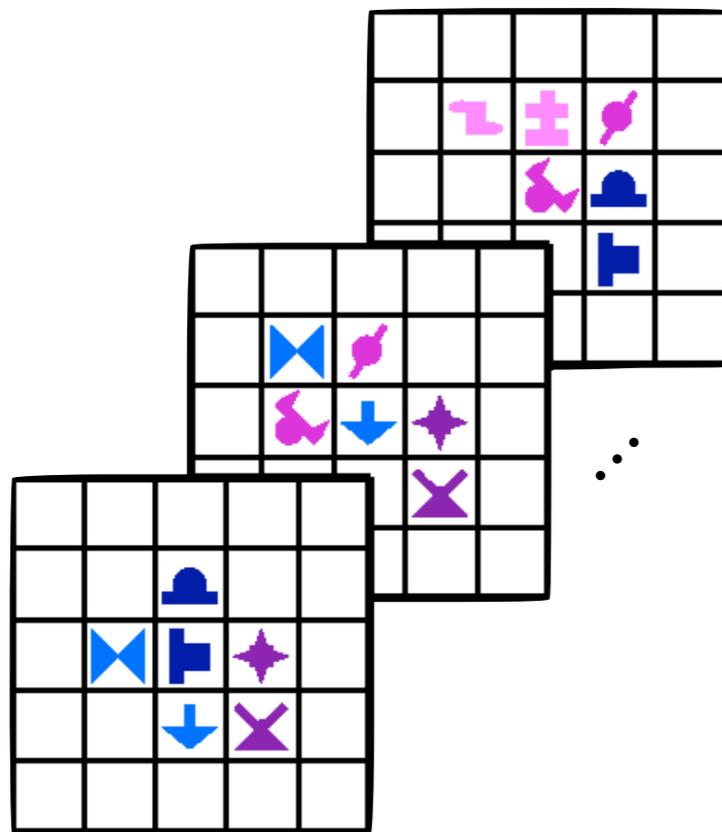
vs.



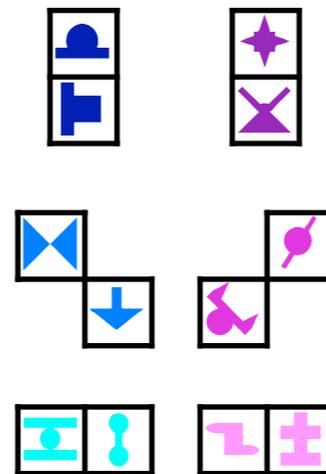
Orbán & al, PNAS 2008

VISUAL PATTERN LEARNING

familiarization
'pay attention'



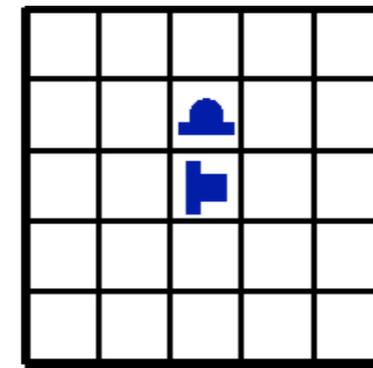
inventory



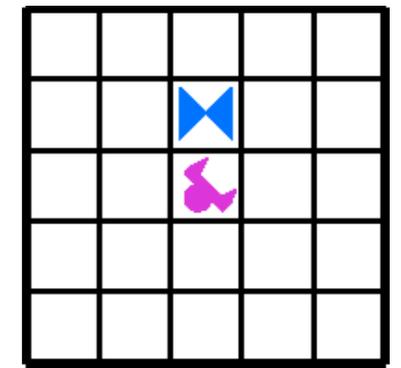
\mathcal{M}

test

'which one looks more familiar?'



vs.

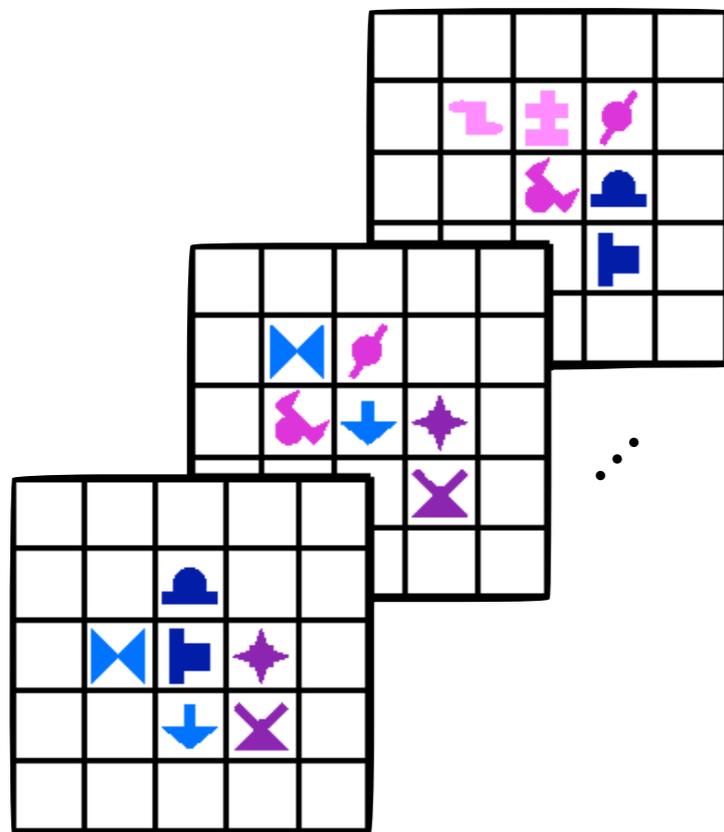


Orbán & al, PNAS 2008

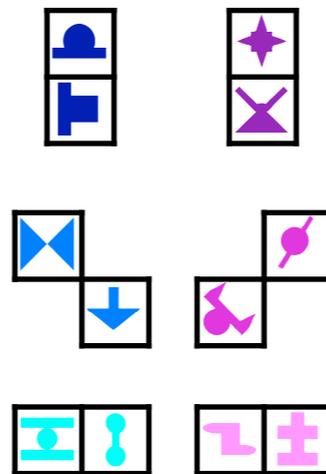
$$\mathcal{D} = \text{scene}_1, \dots, \text{scene}_n \rightarrow \hat{\mathcal{M}} = \underset{\mathcal{M}}{\operatorname{argmax}} P(\mathcal{M}|\mathcal{D})$$

VISUAL PATTERN LEARNING

familiarization
'pay attention'



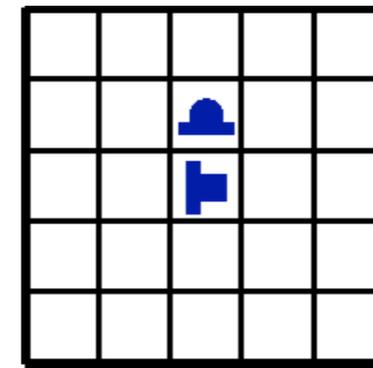
inventory



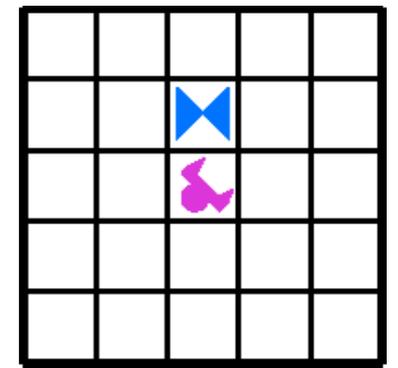
\mathcal{M}

test

'which one looks more familiar?'



vs.

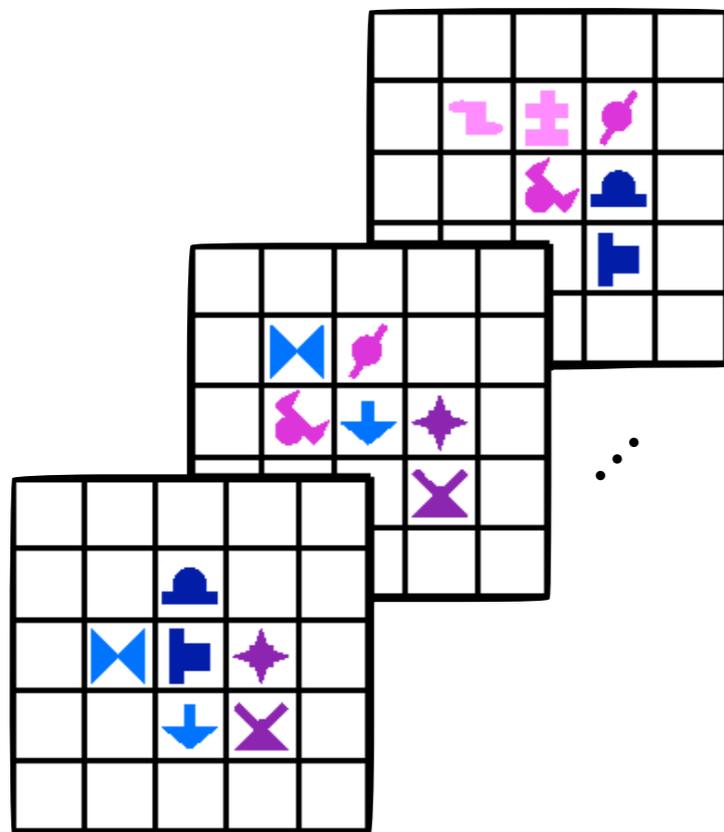


Orbán & al, PNAS 2008

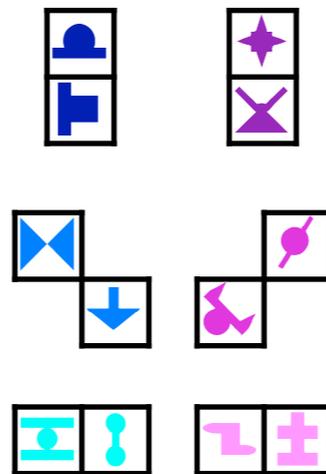
$$\mathcal{D} = \text{scene}_1, \dots, \text{scene}_n \rightarrow \hat{\mathcal{M}} = \underset{\mathcal{M}}{\operatorname{argmax}} P(\mathcal{M}|\mathcal{D}) \rightarrow P(\text{scene}_A|\hat{\mathcal{M}}) \text{ vs. } P(\text{scene}_B|\hat{\mathcal{M}})$$

VISUAL PATTERN LEARNING

familiarization
'pay attention'



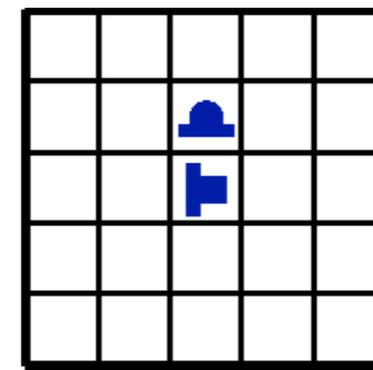
inventory



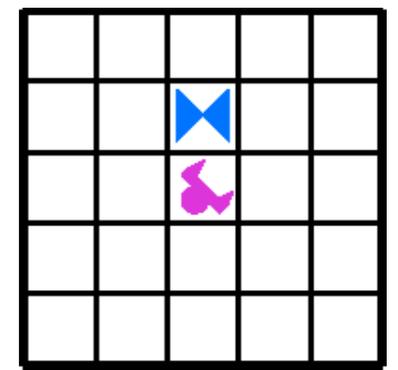
\mathcal{M}

test

'which one looks more familiar?'



vs.



Orbán & al, PNAS 2008

$$\mathcal{D} = \text{scene}_1, \dots, \text{scene}_n \rightarrow \hat{\mathcal{M}} = \underset{\mathcal{M}}{\operatorname{argmax}} P(\mathcal{M}|\mathcal{D}) \rightarrow P(\text{scene}_A|\hat{\mathcal{M}}) \text{ vs. } P(\text{scene}_B|\hat{\mathcal{M}})$$

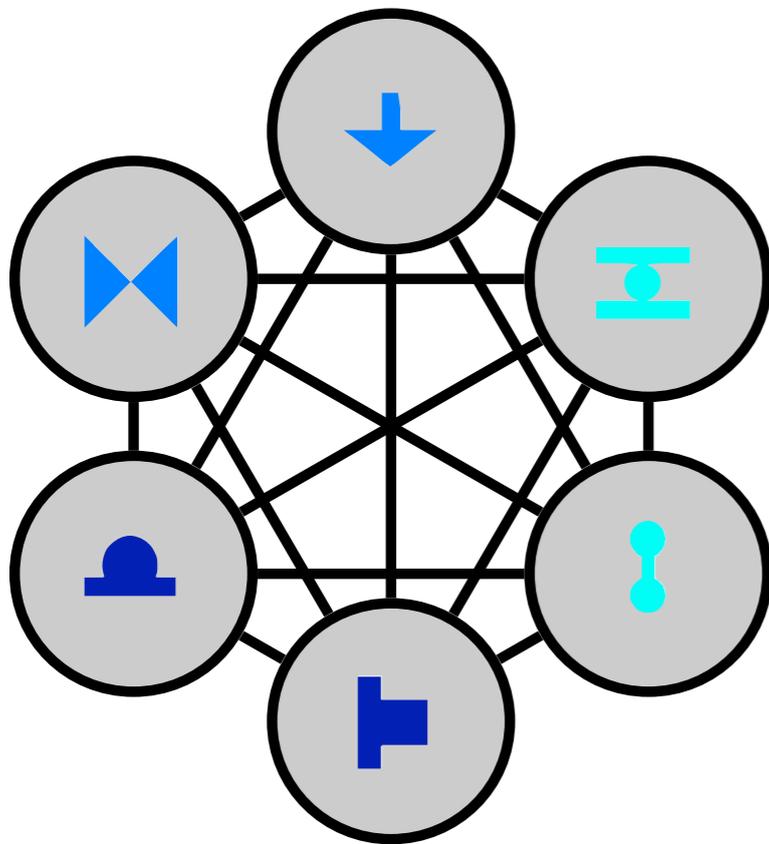
how do humans learn a statistical model of their environment?

- associative learning (fitting 2nd order max-entropy model)
- Bayesian model selection (inferring hidden causal structure)

ALTERNATIVE THEORIES

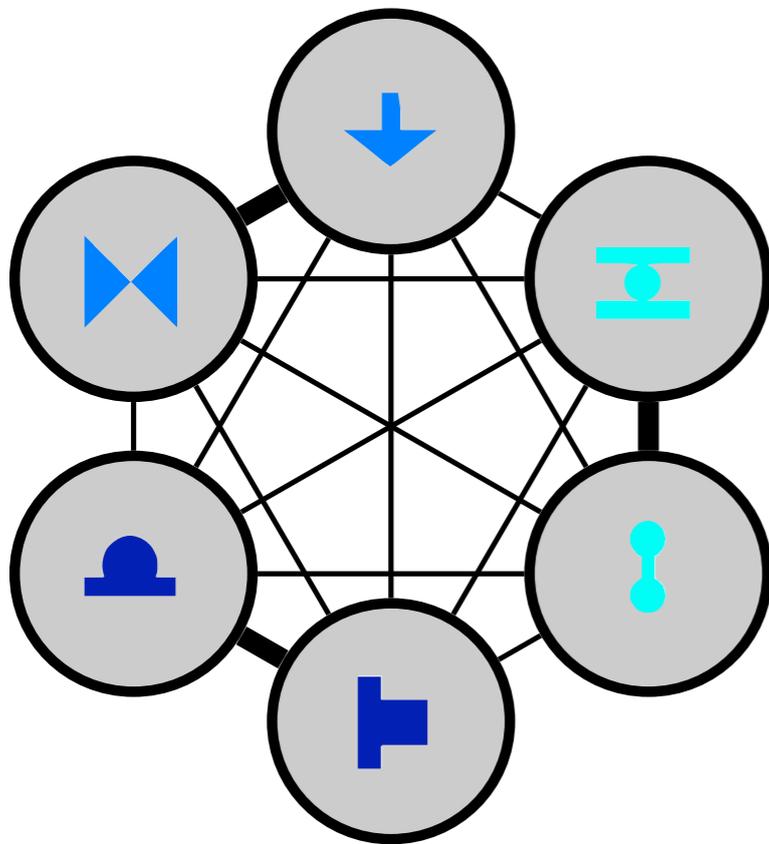
ALTERNATIVE THEORIES

associative learning



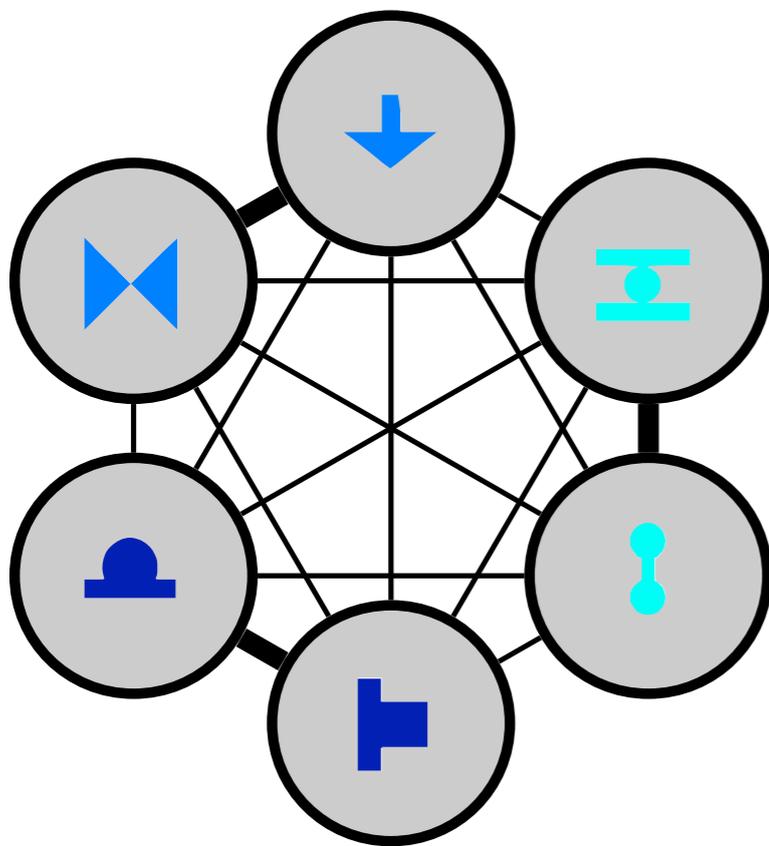
ALTERNATIVE THEORIES

associative learning

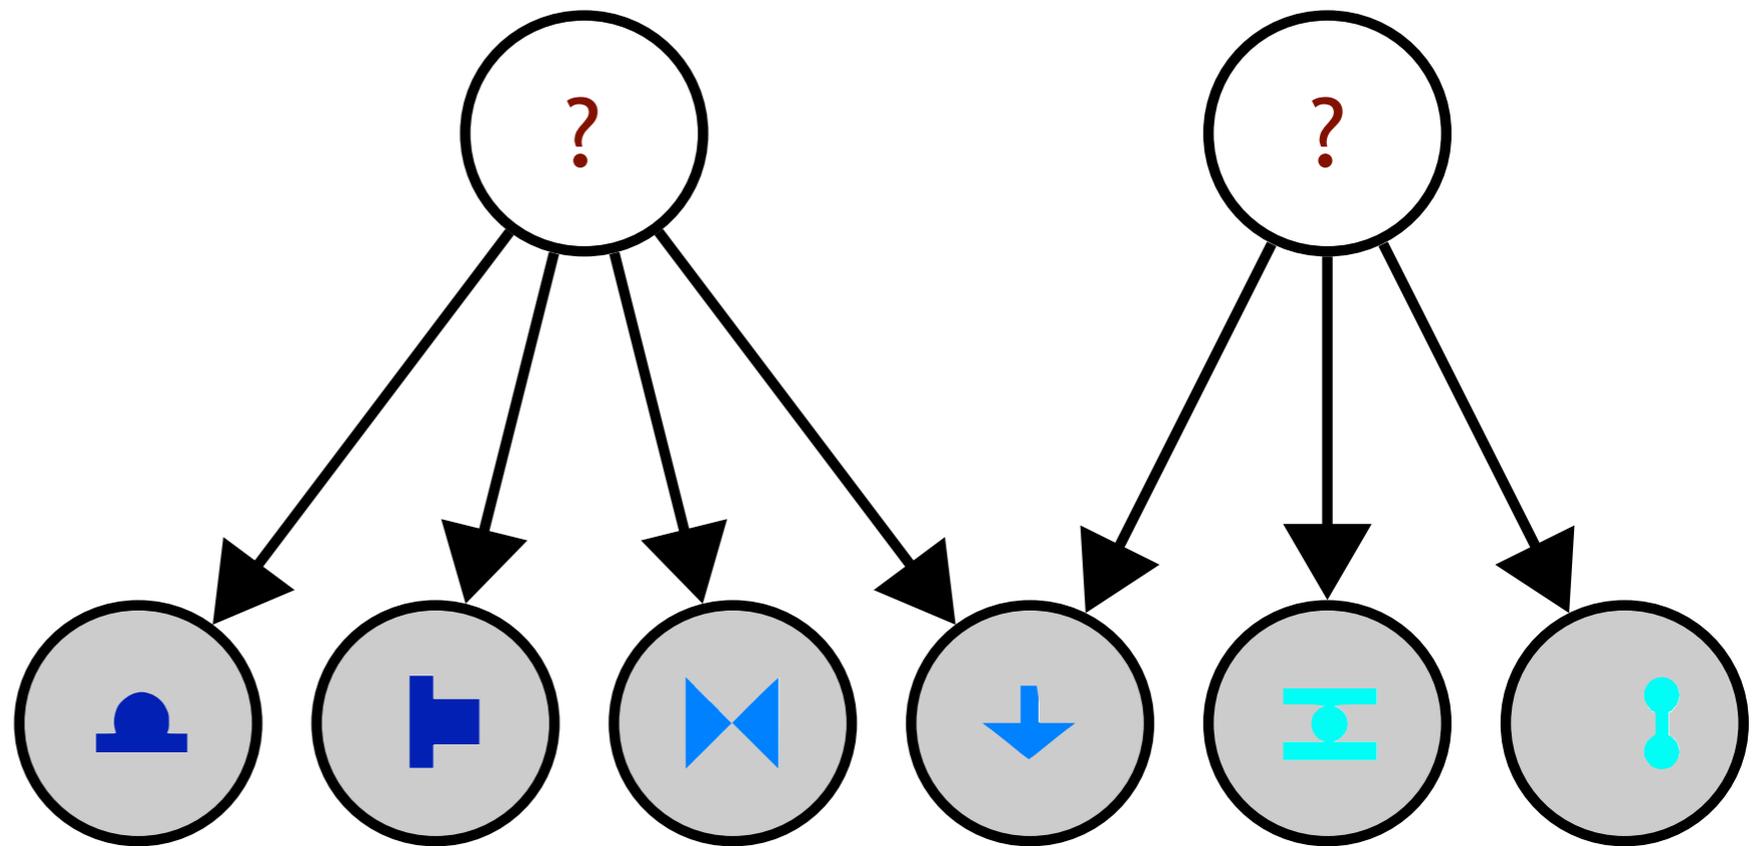


ALTERNATIVE THEORIES

associative learning

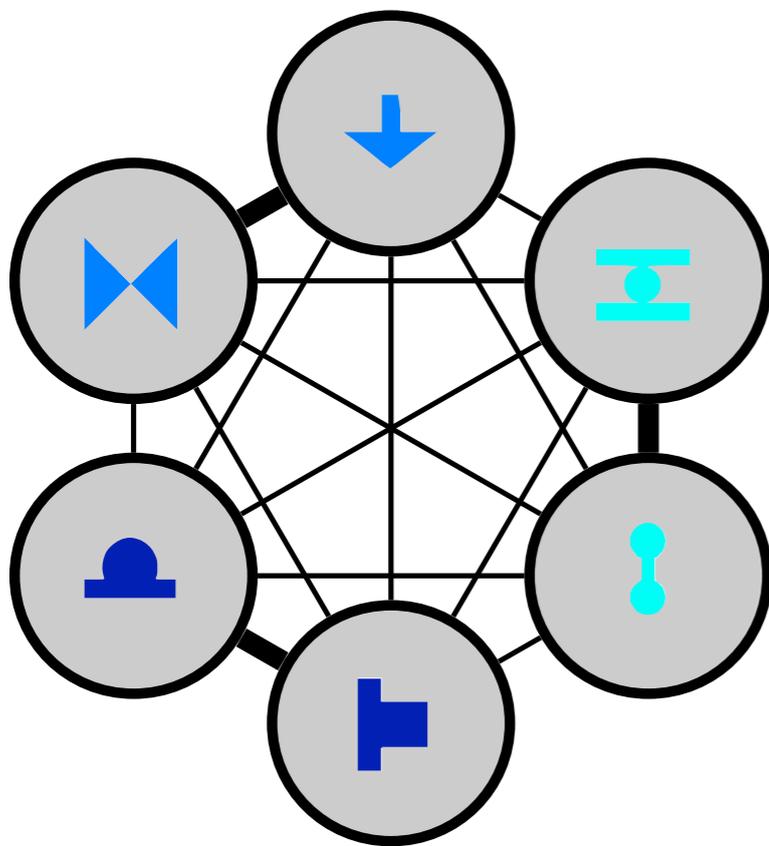


Bayesian learning

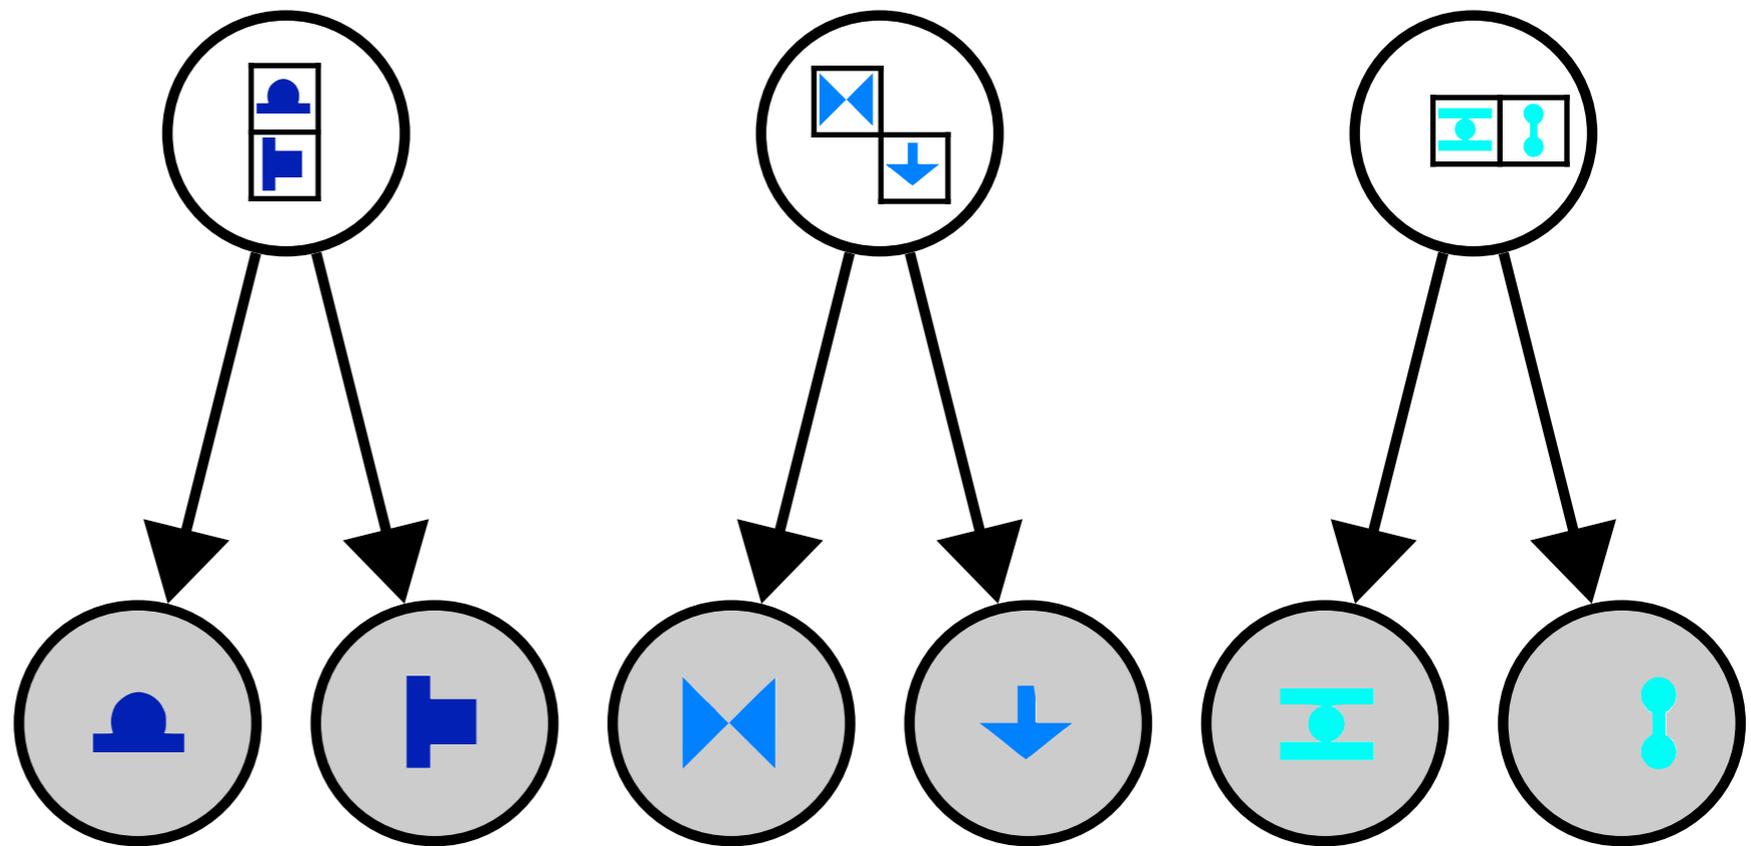


ALTERNATIVE THEORIES

associative learning



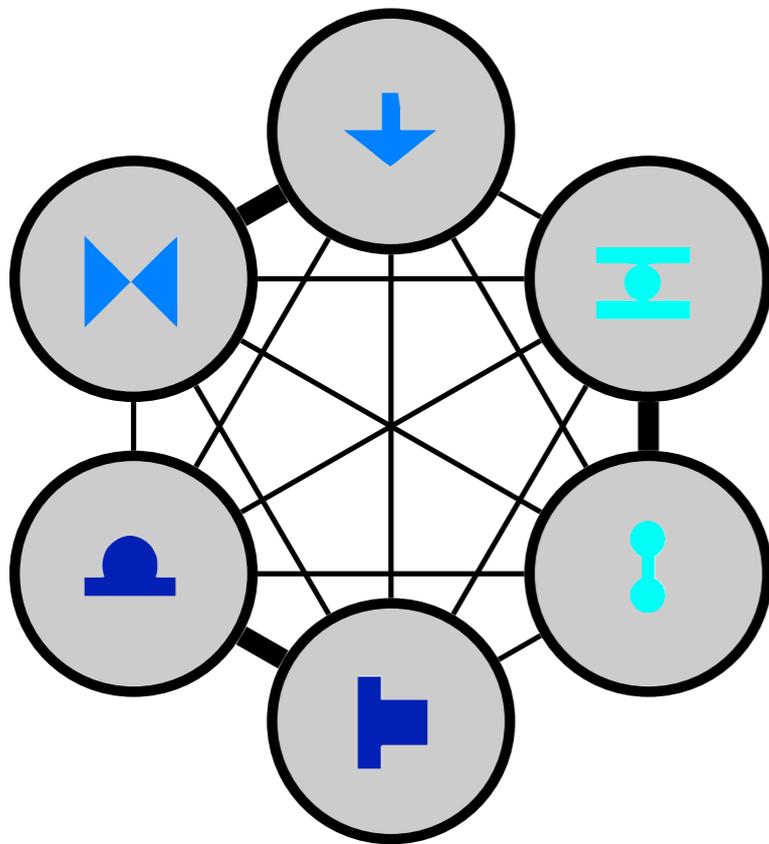
Bayesian learning



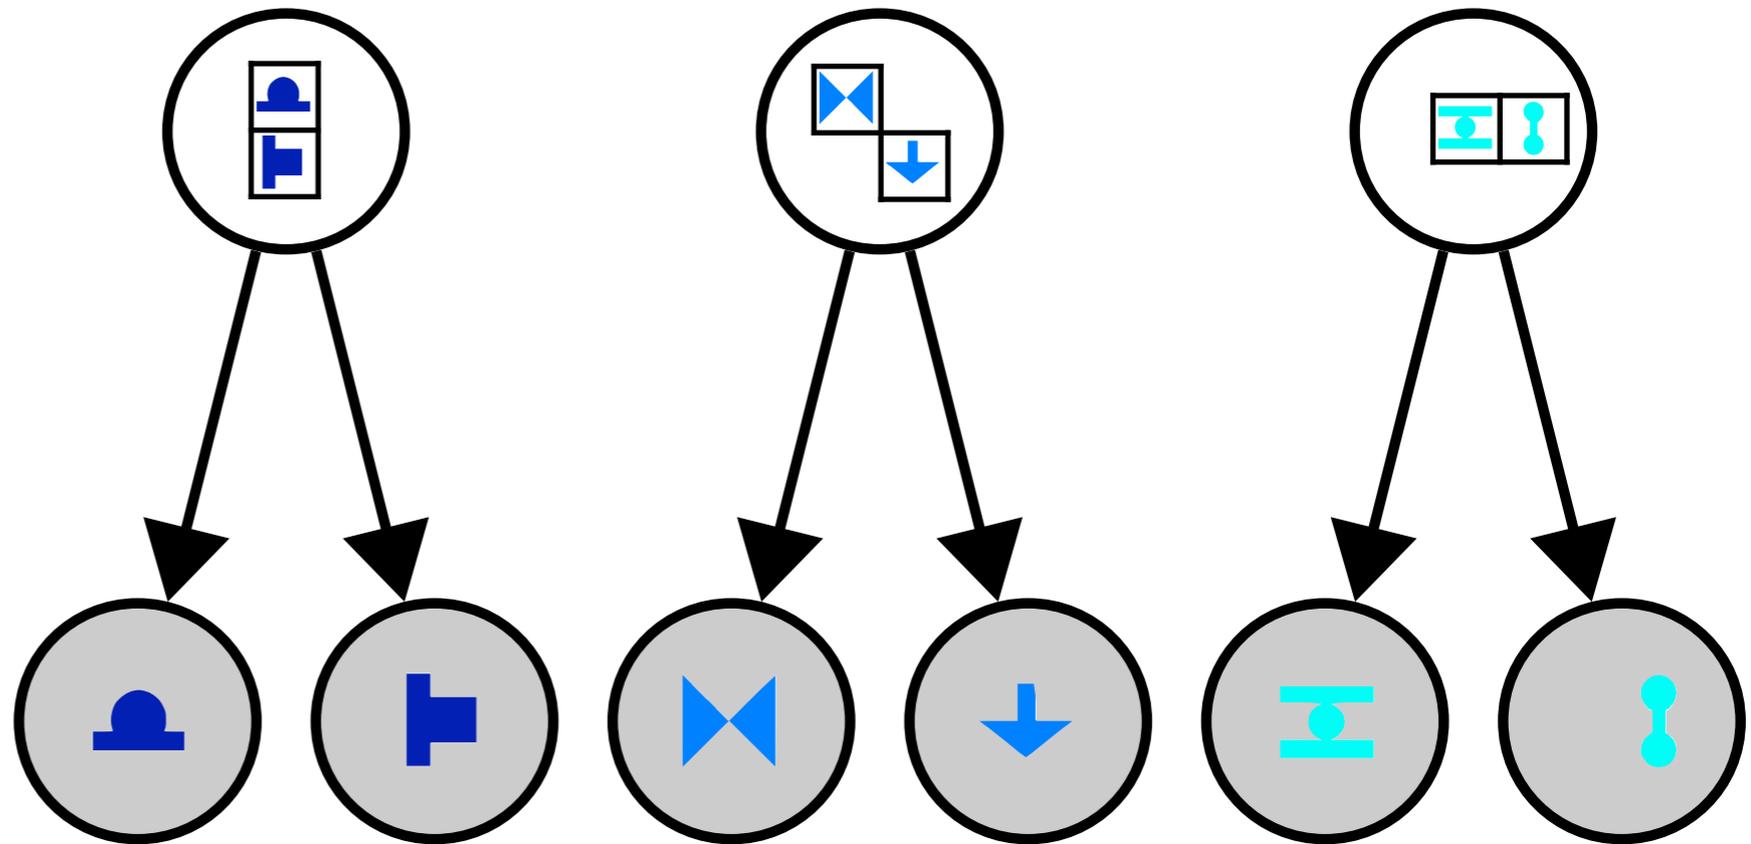
Orbán & al, PNAS 2008

ALTERNATIVE THEORIES

associative learning



Bayesian learning



Orbán & al, PNAS 2008

Boltzmann machine

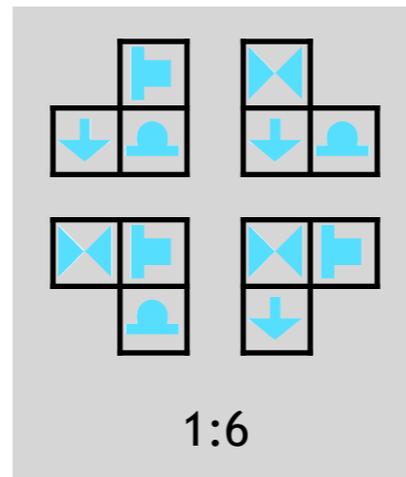
+ Gaussian Markov random field

sigmoid belief network

+ product of (conditional) Gaussian experts

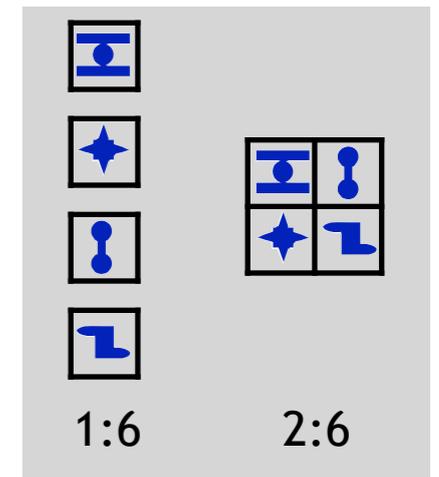
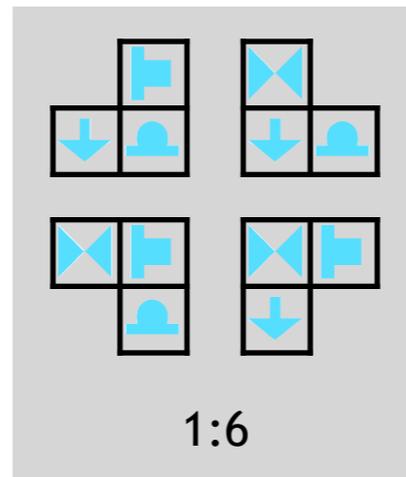
ASSOCIATIVE VS. BAYESIAN LEARNING

inventory:



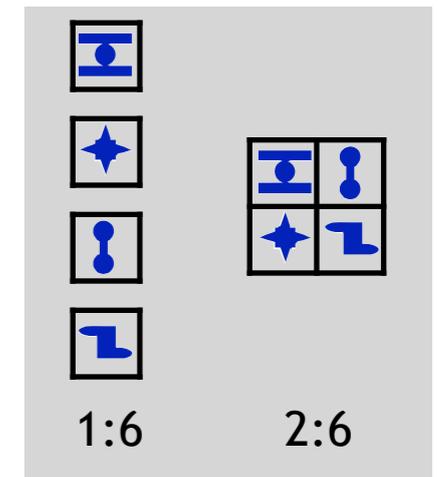
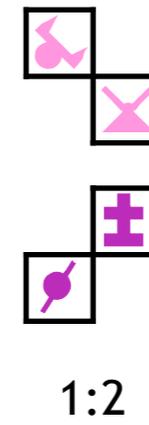
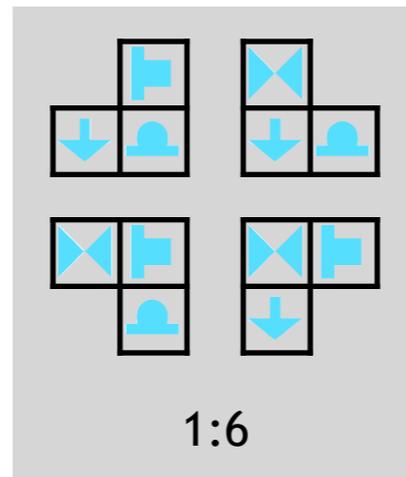
ASSOCIATIVE VS. BAYESIAN LEARNING

inventory:



ASSOCIATIVE VS. BAYESIAN LEARNING

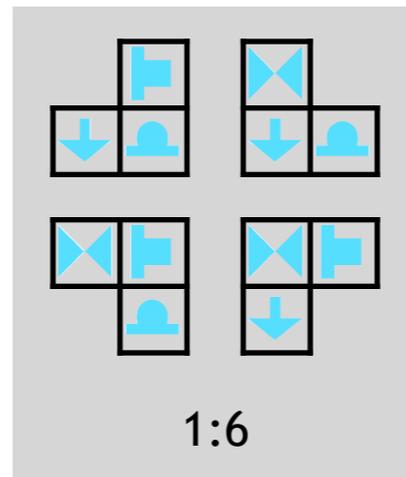
inventory:



ASSOCIATIVE VS. BAYESIAN LEARNING

inventory:

1st order statistic:
shape frequencies

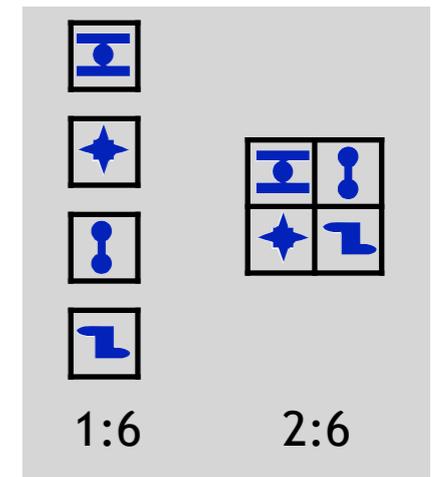


1:6

$$3 \times 1/6$$



1:2



1:6

2:6

$$1/6 + 2/6$$

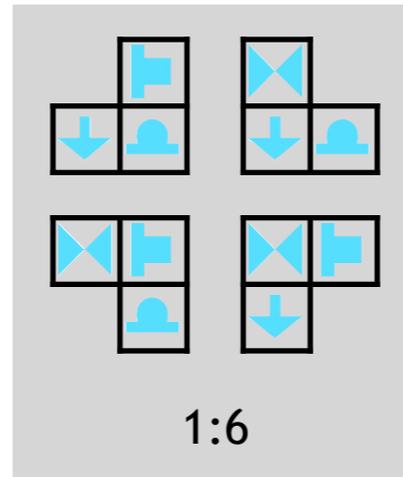
ASSOCIATIVE VS. BAYESIAN LEARNING

inventory:

1st order statistic:
shape frequencies

2nd order statistic:
pairwise correlations

both present
both absent
one present



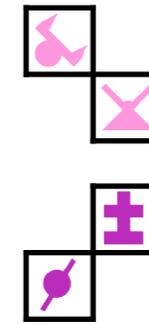
1:6

$$3 \times 1/6$$

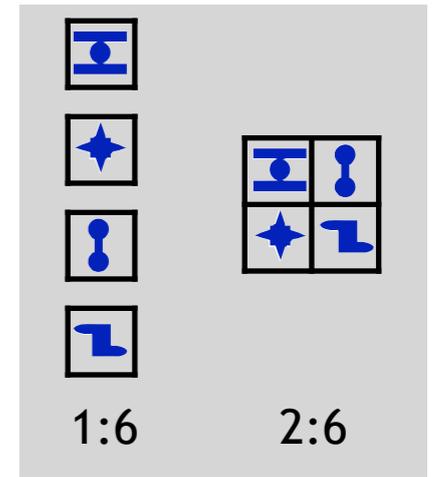
$$2 \times 1/6$$

$$2/6$$

$$2 \times 1/6$$



1:2



1:6

2:6

$$1/6 + 2/6$$

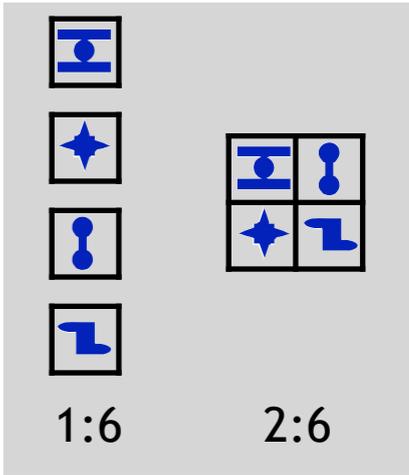
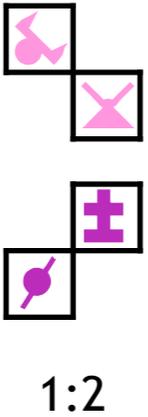
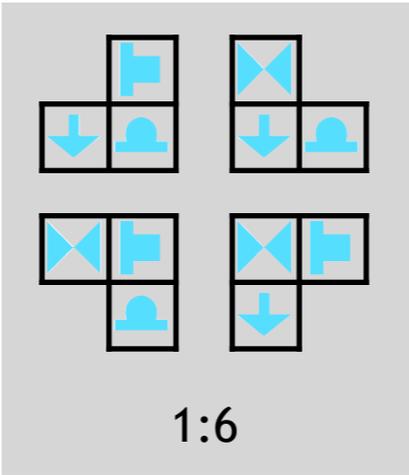
$$2/6$$

$$2 \times 1/6$$

$$2 \times 1/6$$

ASSOCIATIVE VS. BAYESIAN LEARNING

inventory:



1st order statistic:
shape frequencies

$3 \times 1/6$

$1/6 + 2/6$

2nd order statistic:
pairwise correlations

both present
both absent
one present

$2 \times 1/6$

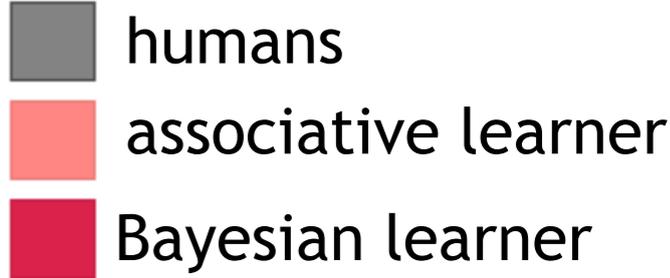
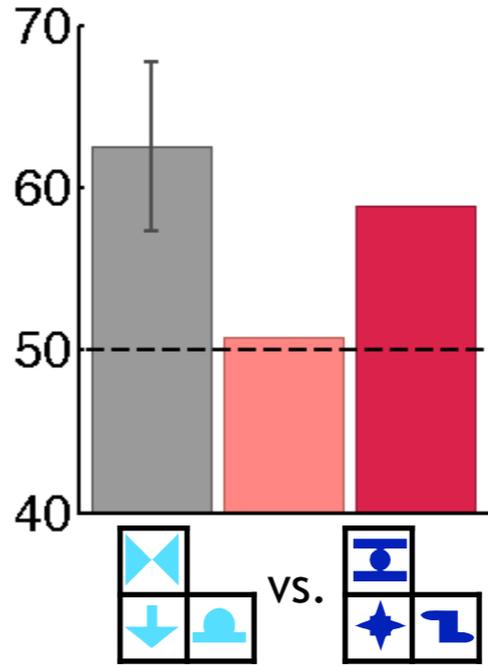
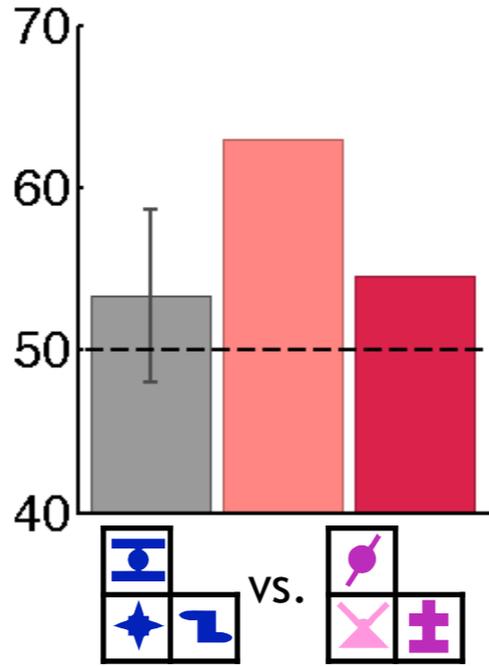
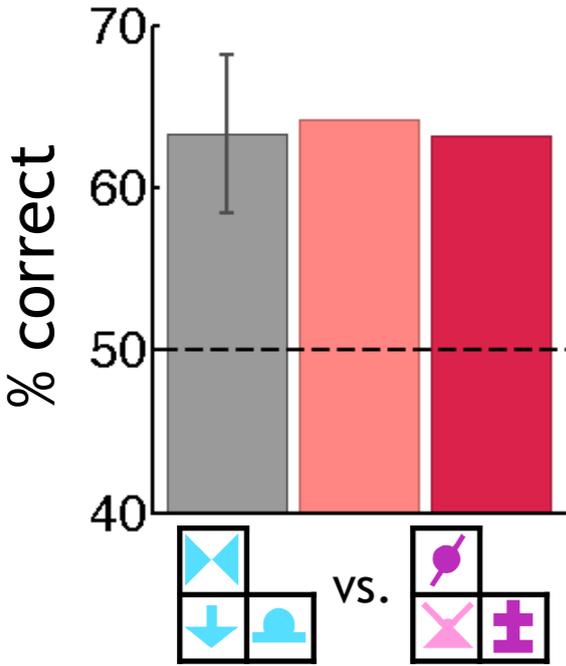
$2/6$

$2/6$

$2 \times 1/6$

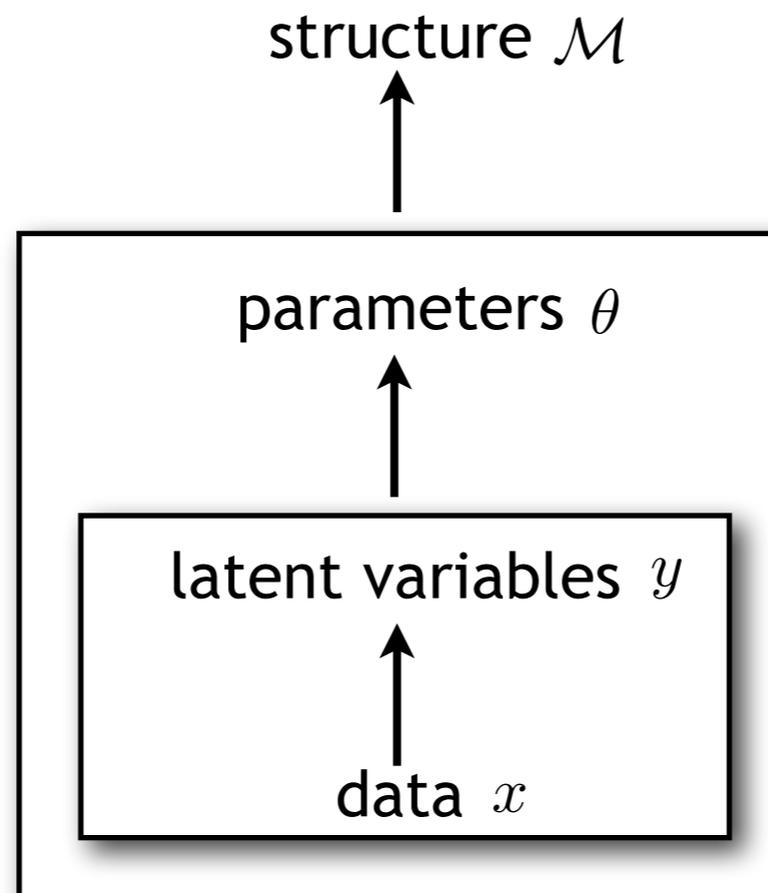
$2 \times 1/6$

test performance:

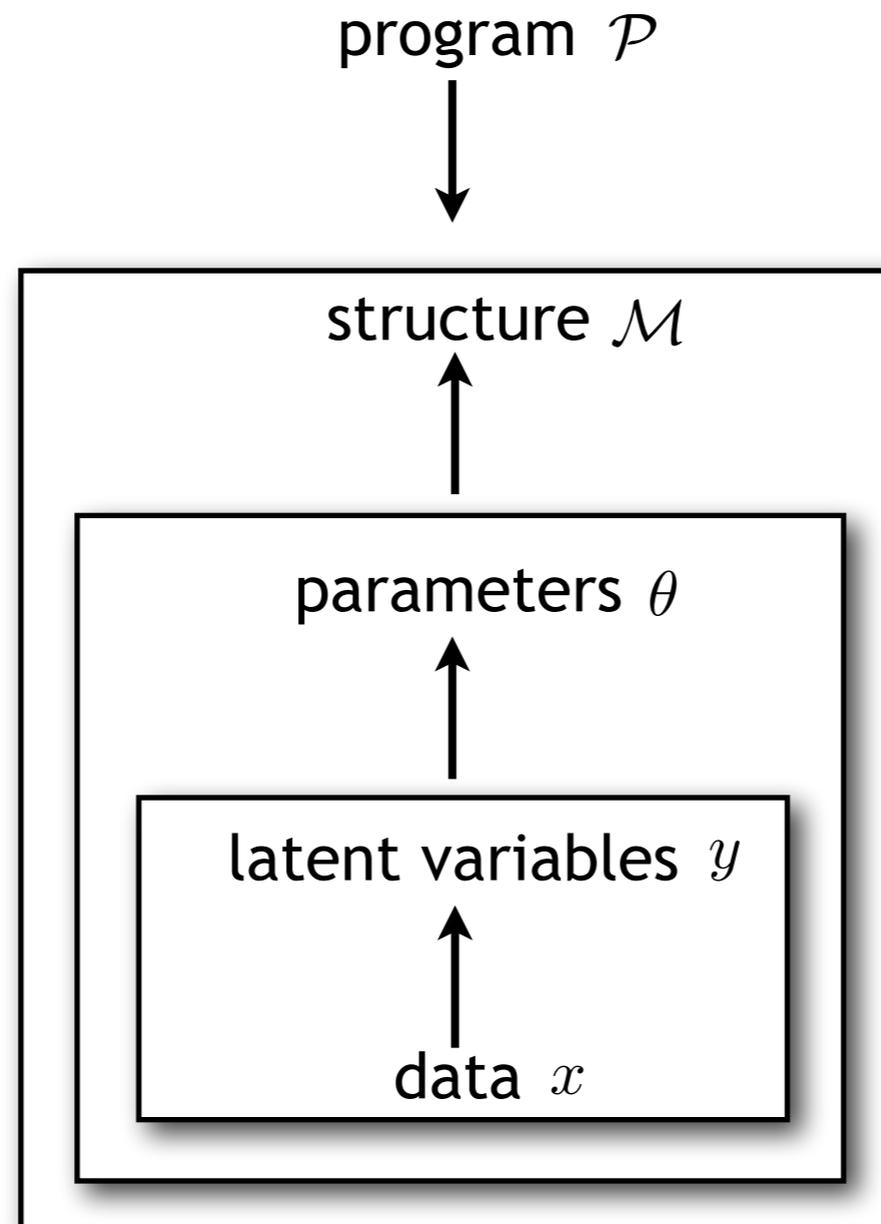


Orbán & al, PNAS 2008

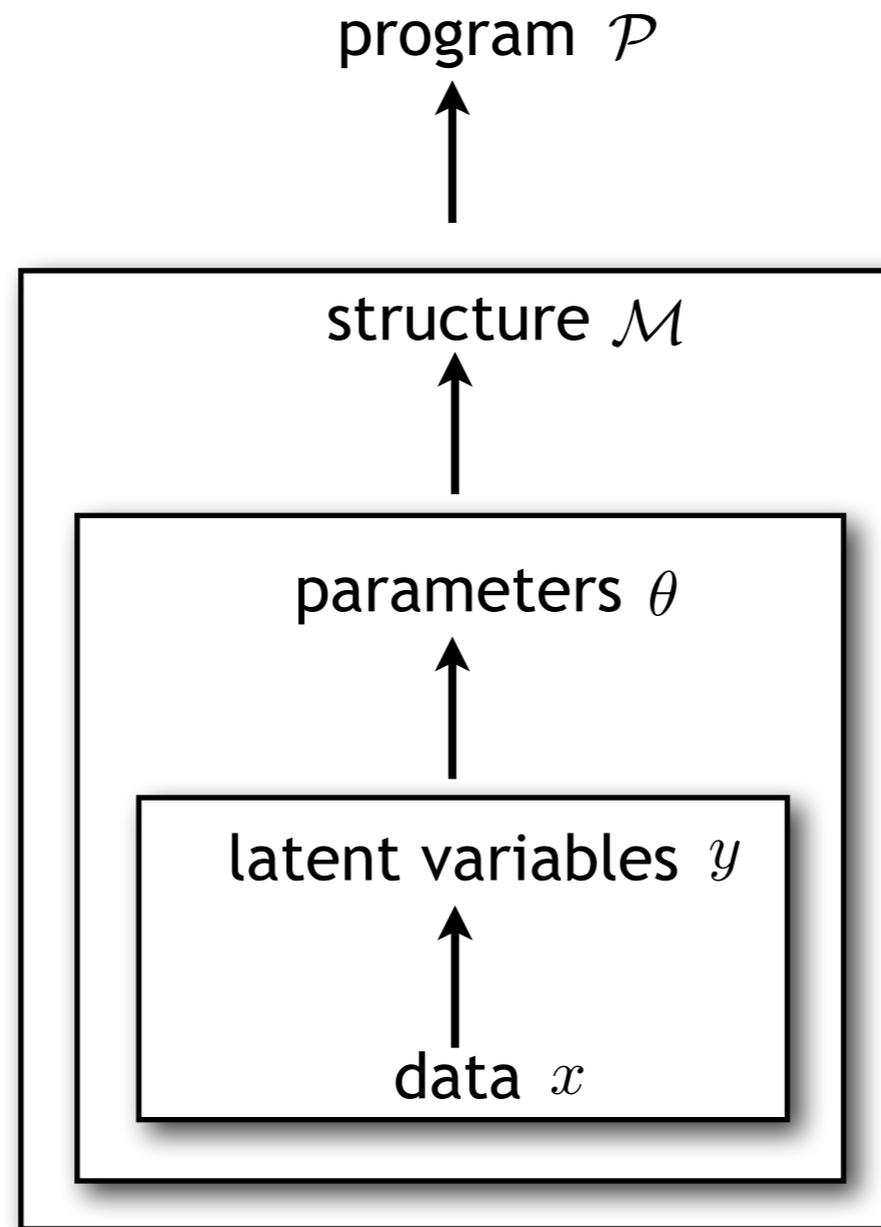
PROBABILISTIC INFERENCE AND LEARNING



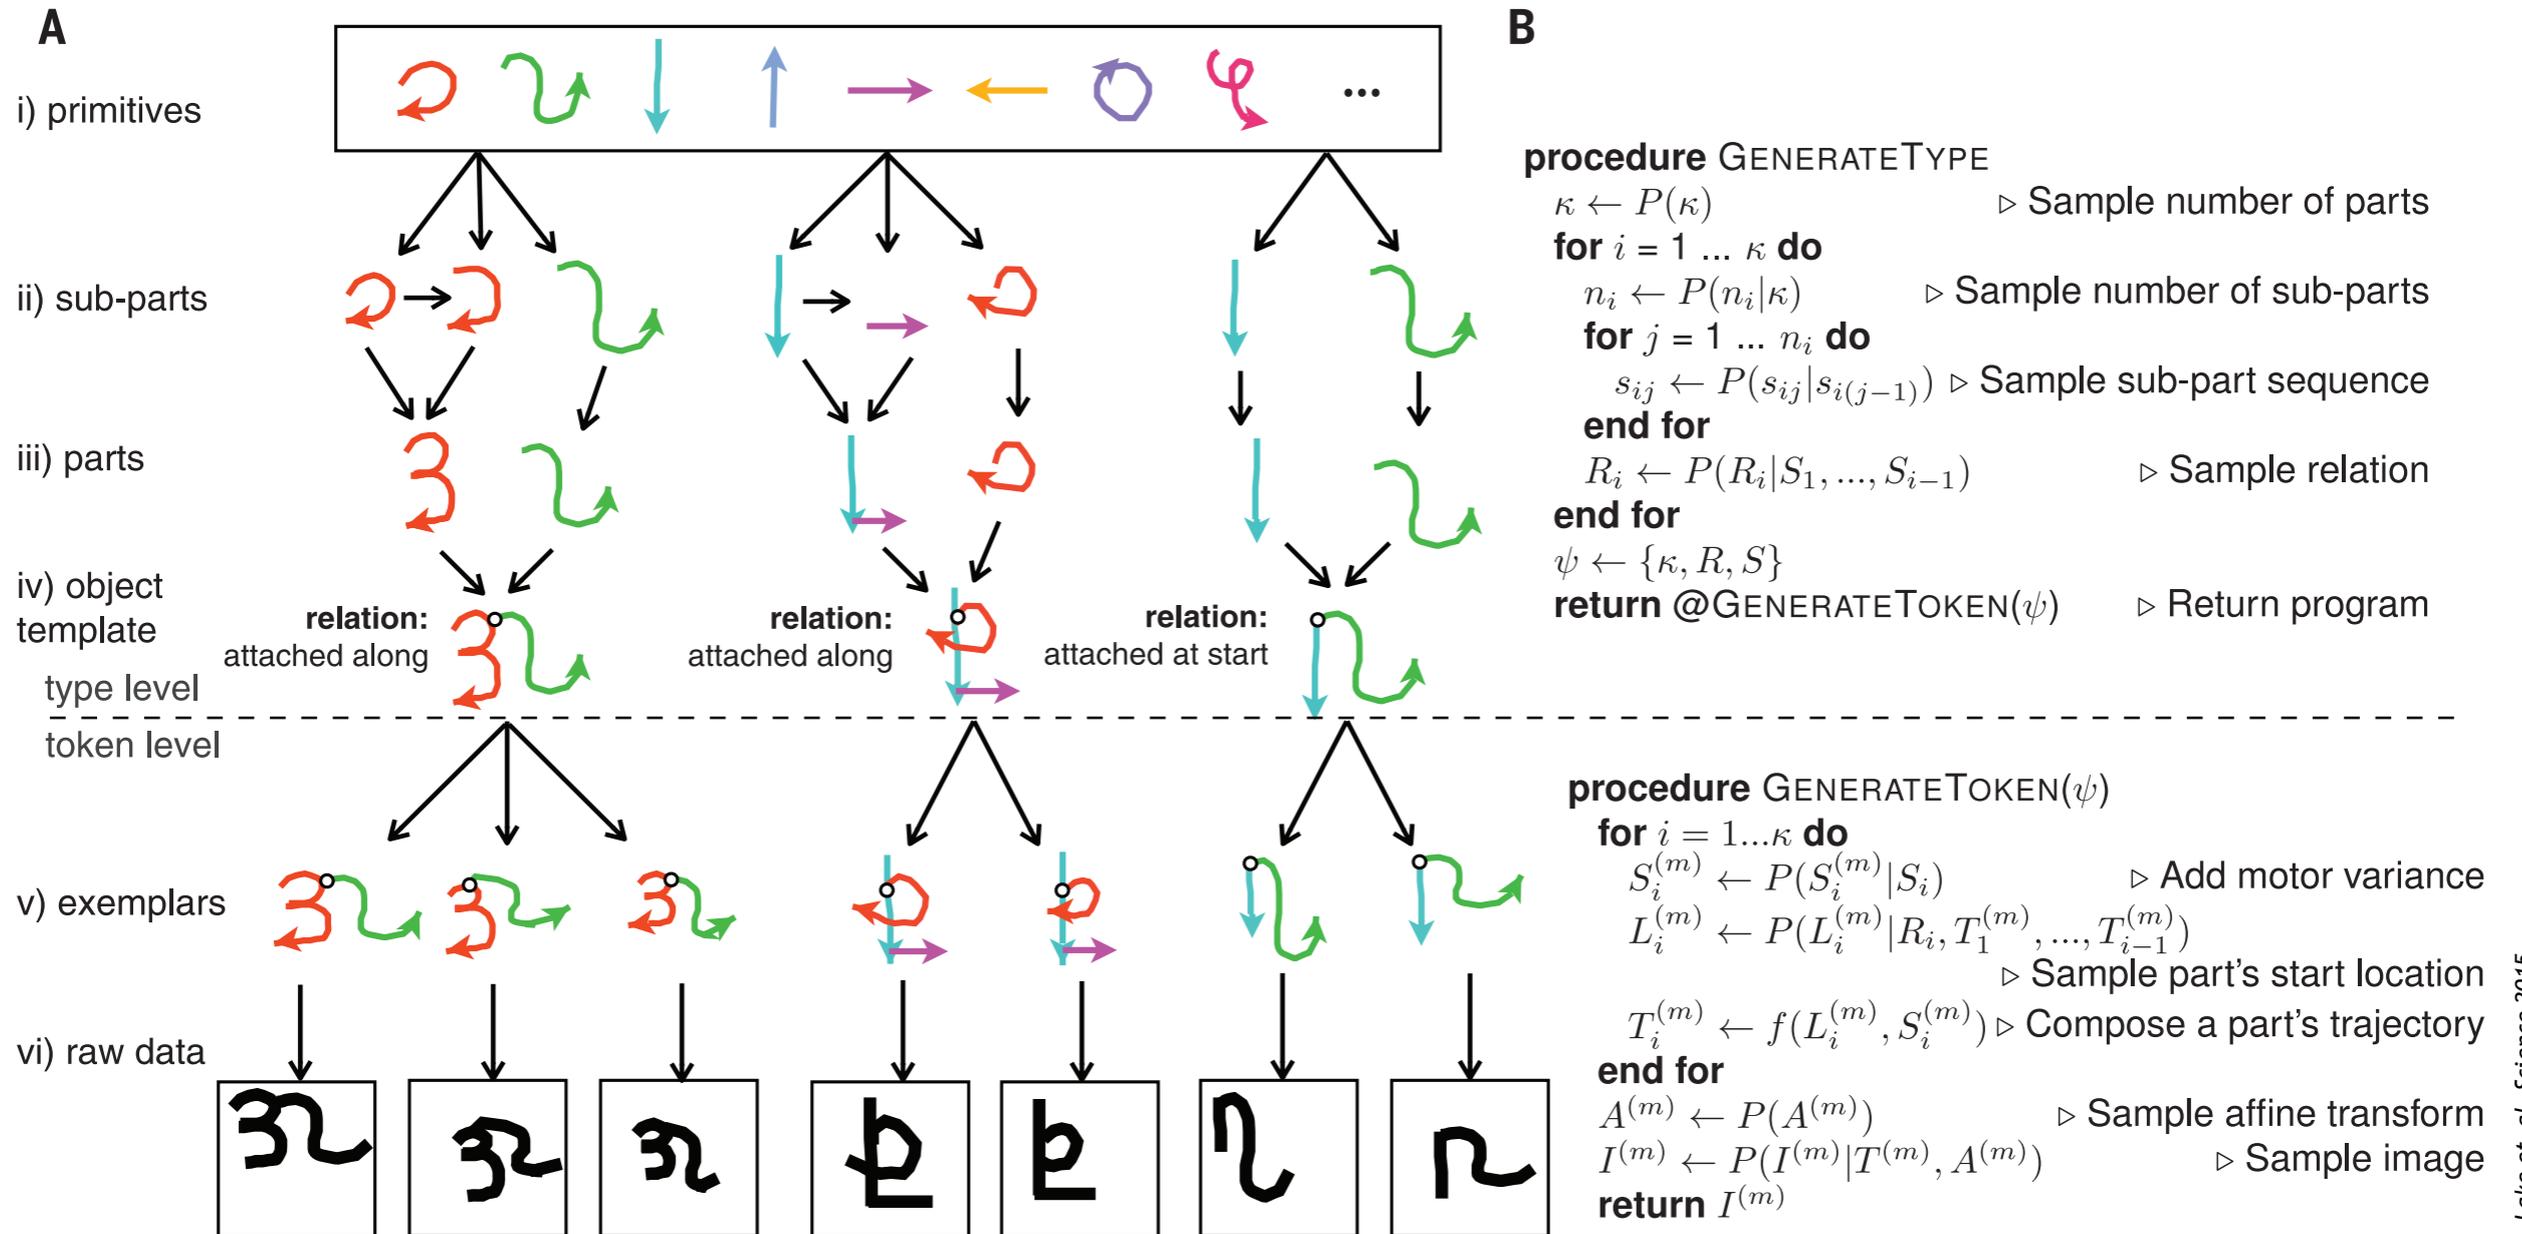
PROBABILISTIC INFERENCE AND LEARNING



PROBABILISTIC INFERENCE AND LEARNING

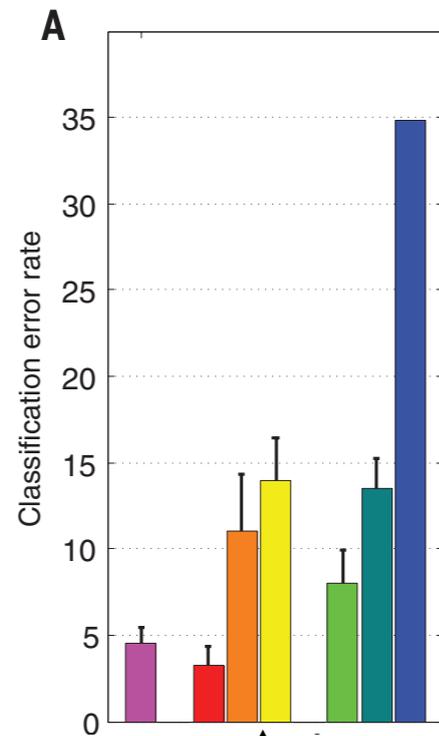
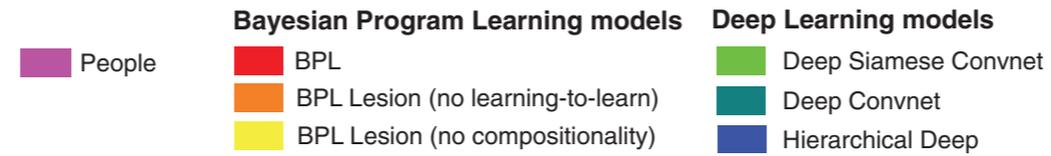


CONCEPT LEARNING BY BAYESIAN PROGRAM LEARNING



CONCEPT LEARNING BY BAYESIAN PROGRAM LEARNING

CONCEPT LEARNING BY BAYESIAN PROGRAM LEARNING



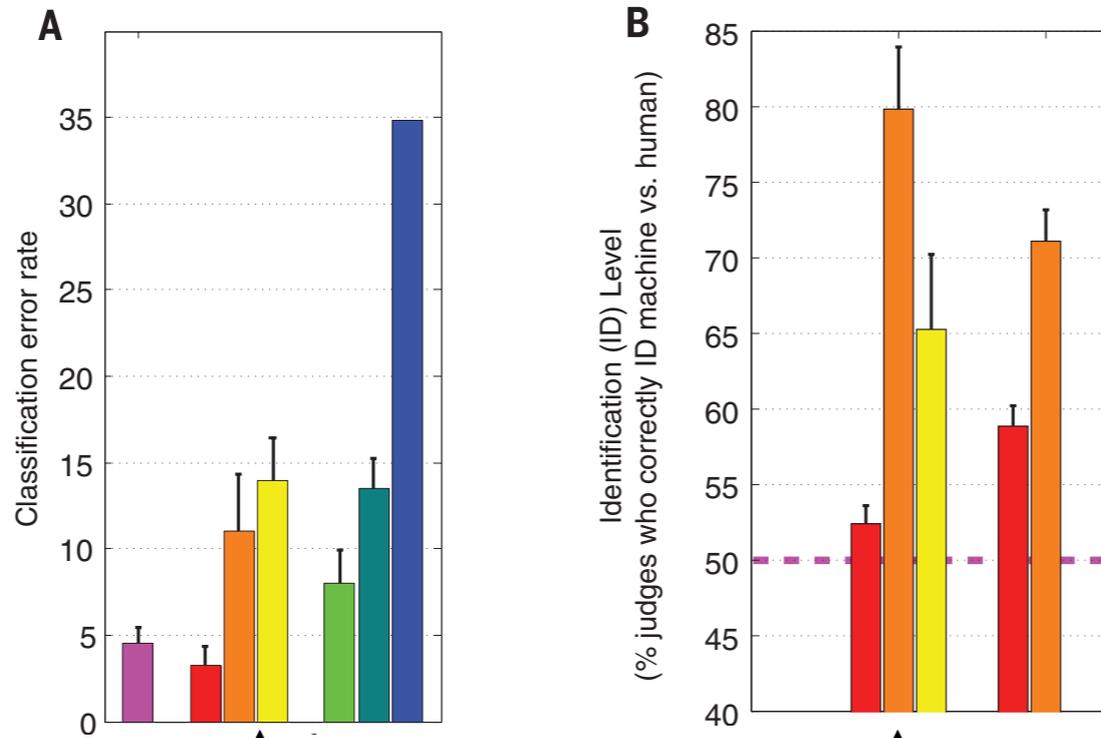
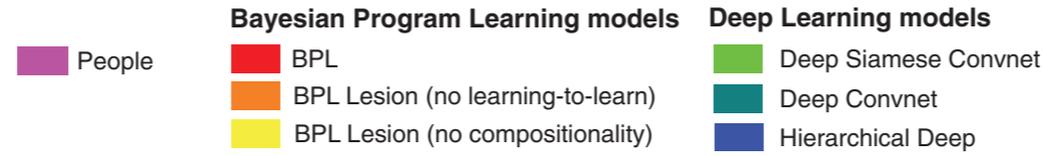
human:
copy this

ಕು

machine:
which one is the same?

ಅ	ಇ	ಉ	ಊ	ಋ
ಕ	ಖ	ಗ	ಘ	ಙ
ಚ	ಛ	ಜ	ಝ	ಞ
ಟ	ಠ	ಡ	ಢ	ನ

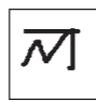
CONCEPT LEARNING BY BAYESIAN PROGRAM LEARNING



human:
copy this



human / machine:
copy this



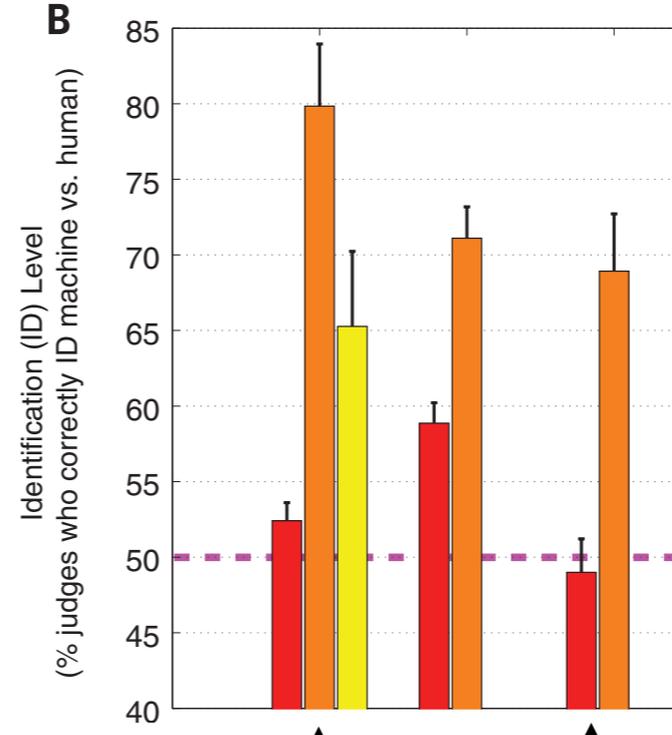
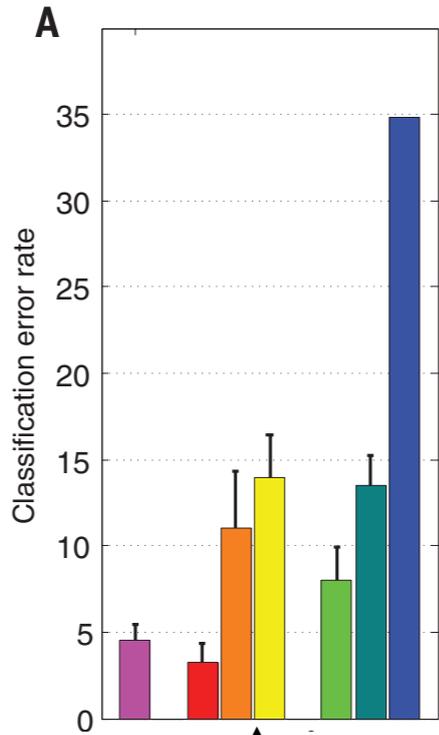
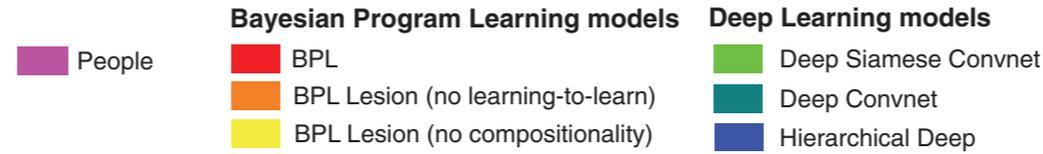
machine:
which one is the same?

అ	ఇ	ఊ	ఎ	ఔ
కా	ఖ	గా	ఙ	ఞు
శై	త	ణ	త్ర	ద
న	య	ల	ఝ	ఞ

human:
human or machine?

ఱ	ఱ	ఱ	ఱ	ఱ
ఱ	ఱ	ఱ	ఱ	ఱ
ఱ	ఱ	ఱ	ఱ	ఱ
ఱ	ఱ	ఱ	ఱ	ఱ

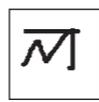
CONCEPT LEARNING BY BAYESIAN PROGRAM LEARNING



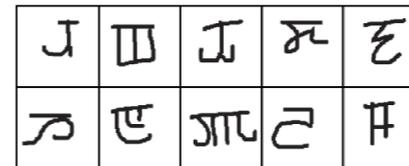
human:
copy this



human / machine:
copy this



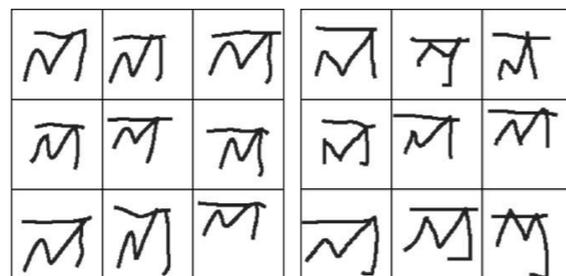
human / machine:
generate new character from alphabet



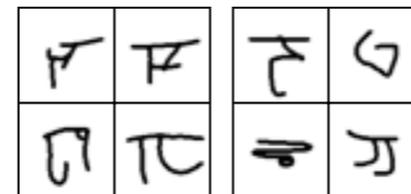
machine:
which one is the same?



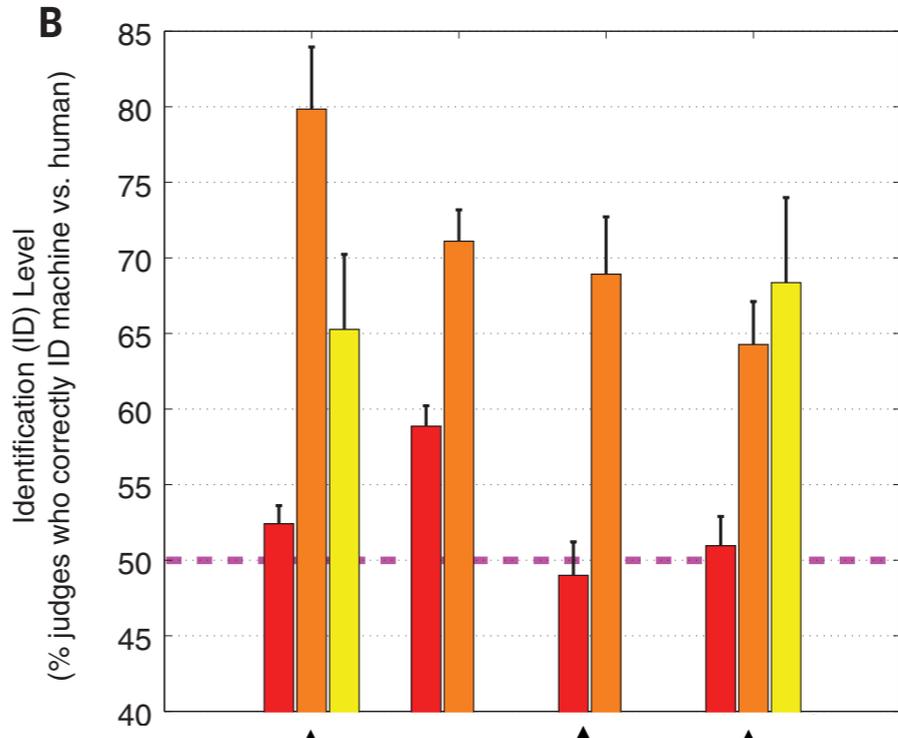
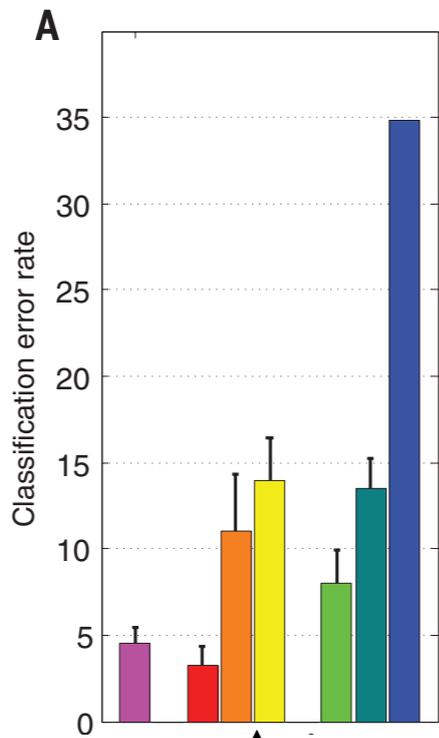
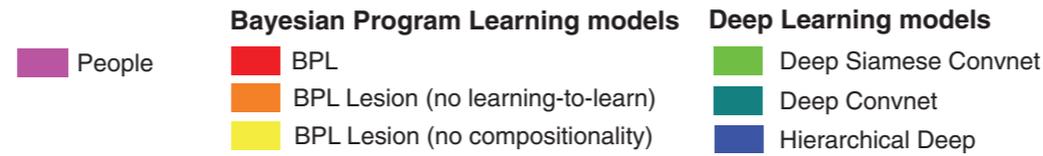
human:
human or machine?



human:
human or machine?



CONCEPT LEARNING BY BAYESIAN PROGRAM LEARNING



Lake et al., Science 2015

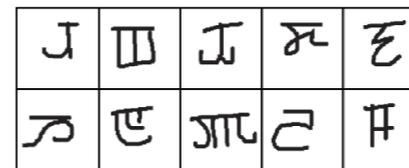
human:
copy this



human / machine:
copy this



human / machine:
generate new character from alphabet

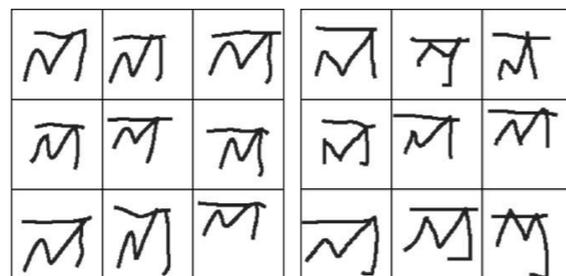


human / machine:
generate new character

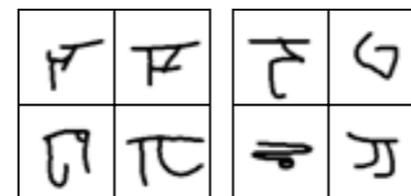
machine:
which one is the same?



human:
human or machine?



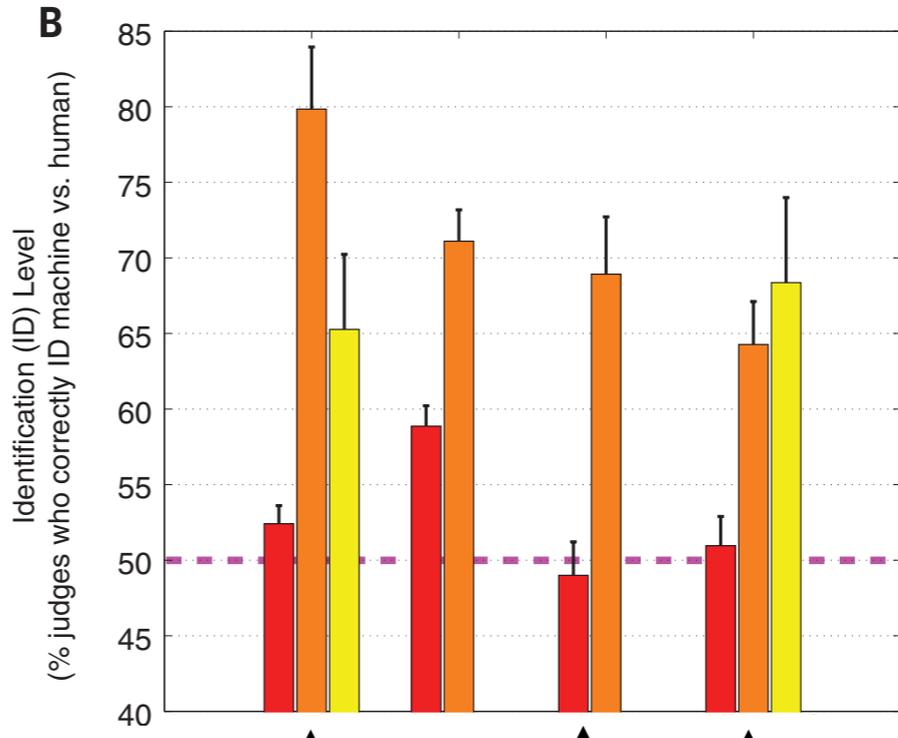
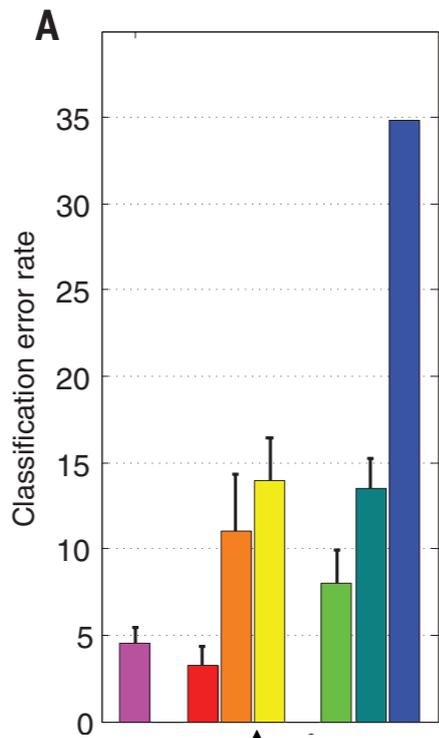
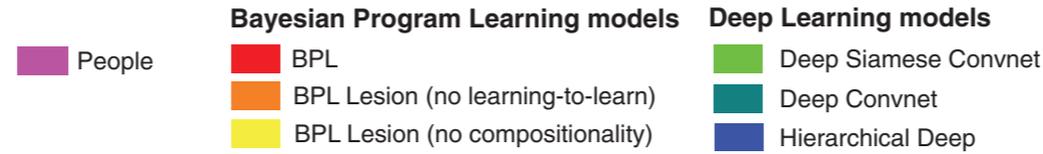
human:
human or machine?



human:
human or machine?



CONCEPT LEARNING BY BAYESIAN PROGRAM LEARNING



Lake et al., Science 2015

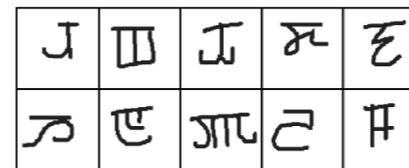
human:
copy this



human / machine:
copy this



human / machine:
generate new character from alphabet

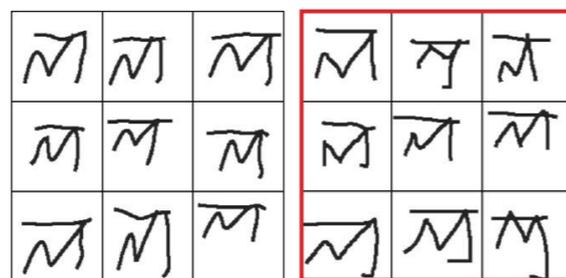


human / machine:
generate new character

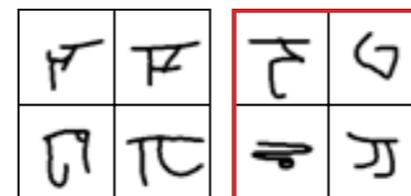
machine:
which one is the same?



human:
human or machine?



human:
human or machine?



human:
human or machine?

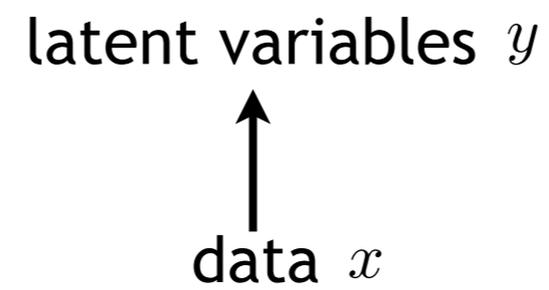


PROBABILISTIC INFERENCE AND LEARNING

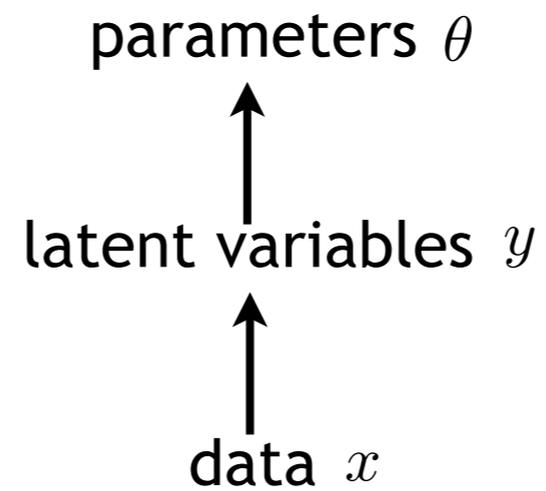
PROBABILISTIC INFERENCE AND LEARNING

data x

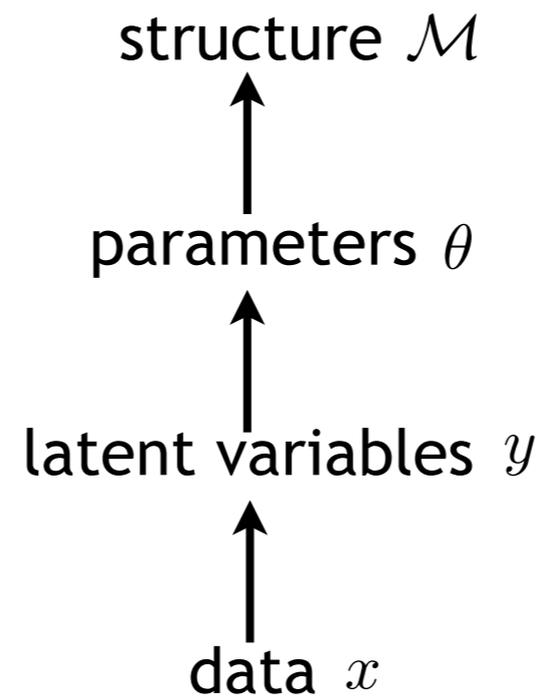
PROBABILISTIC INFERENCE AND LEARNING



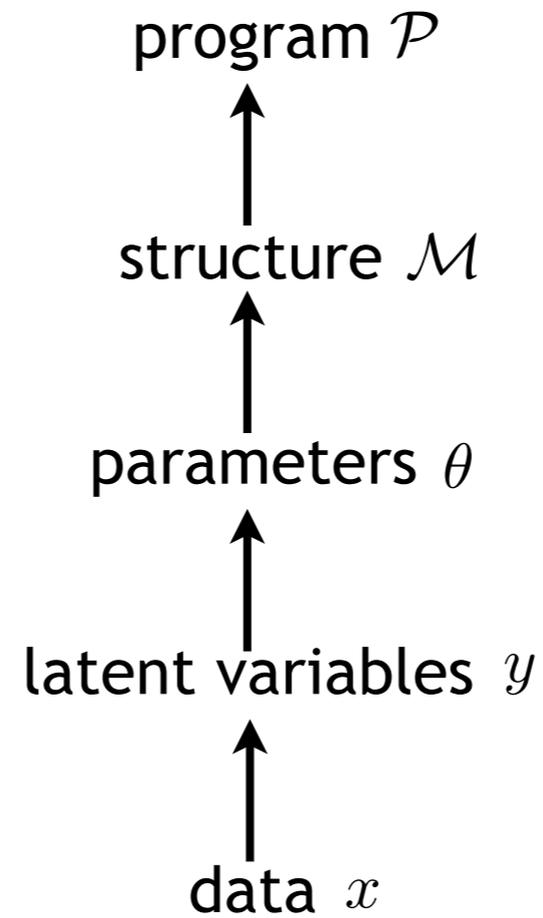
PROBABILISTIC INFERENCE AND LEARNING



PROBABILISTIC INFERENCE AND LEARNING

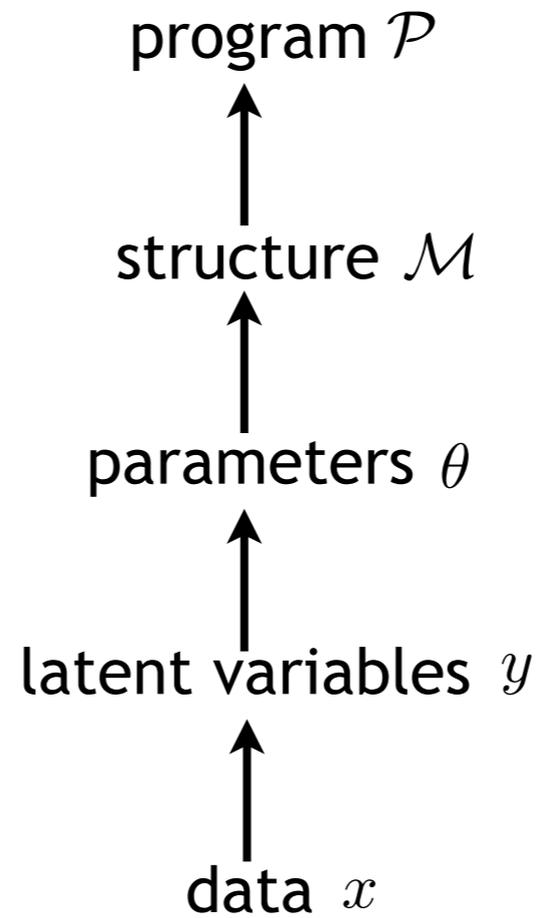


PROBABILISTIC INFERENCE AND LEARNING



PROBABILISTIC INFERENCE AND LEARNING

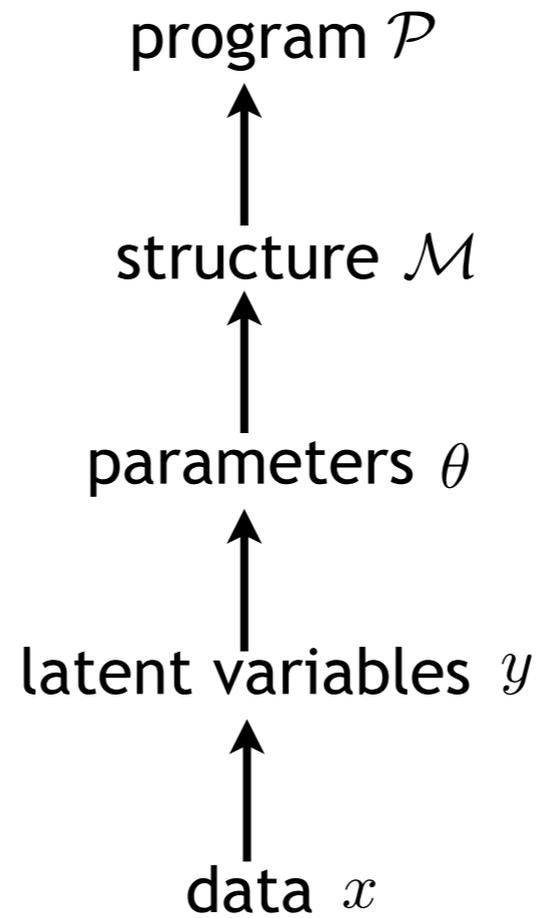
cognitive science



PROBABILISTIC INFERENCE AND LEARNING

cognitive science

theory



PROBABILISTIC INFERENCE AND LEARNING

cognitive science

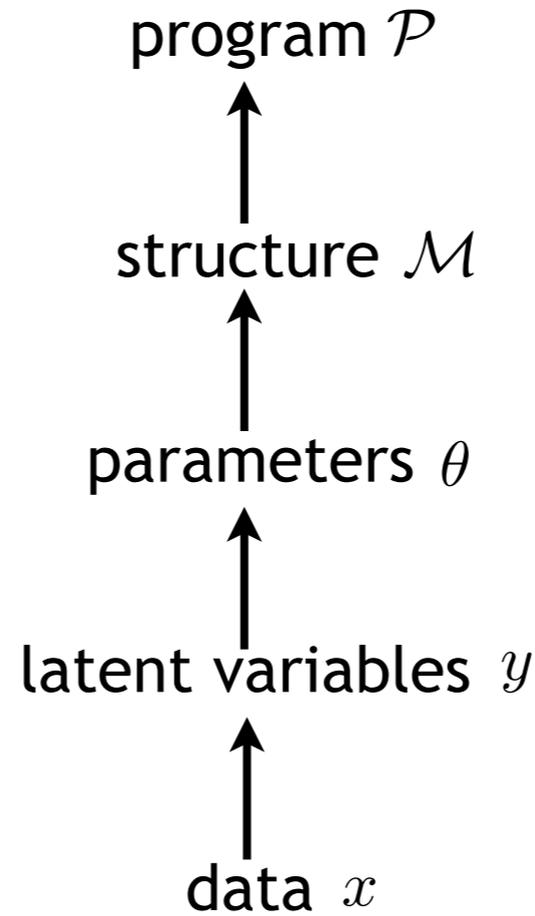
theory *experiments*

✓

✓

✓

✓



PROBABILISTIC INFERENCE AND LEARNING

cognitive science

theory *experiments*

✓

?

✓

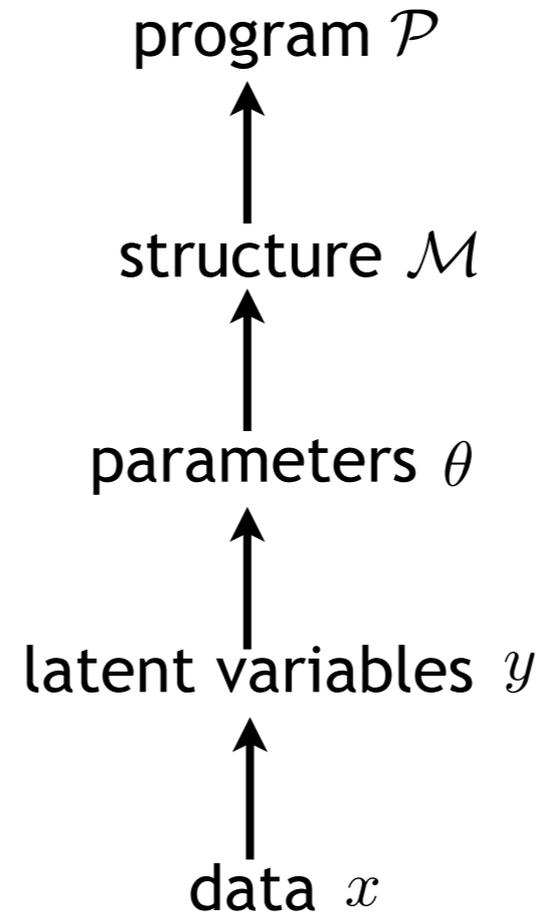
✓

✓

✓

✓

✓



PROBABILISTIC INFERENCE AND LEARNING

cognitive science

neuroscience

theory *experiments*

✓

?

✓

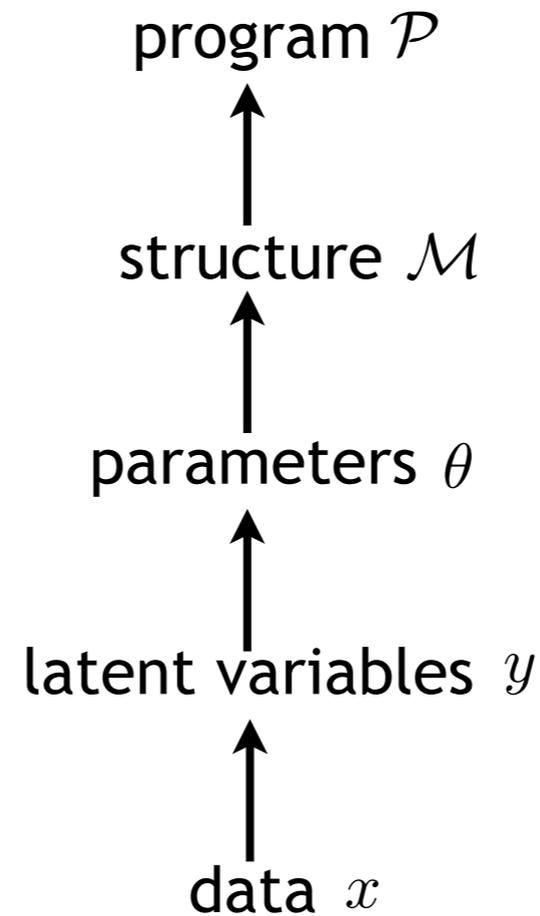
✓

✓

✓

✓

✓



PROBABILISTIC INFERENCE AND LEARNING

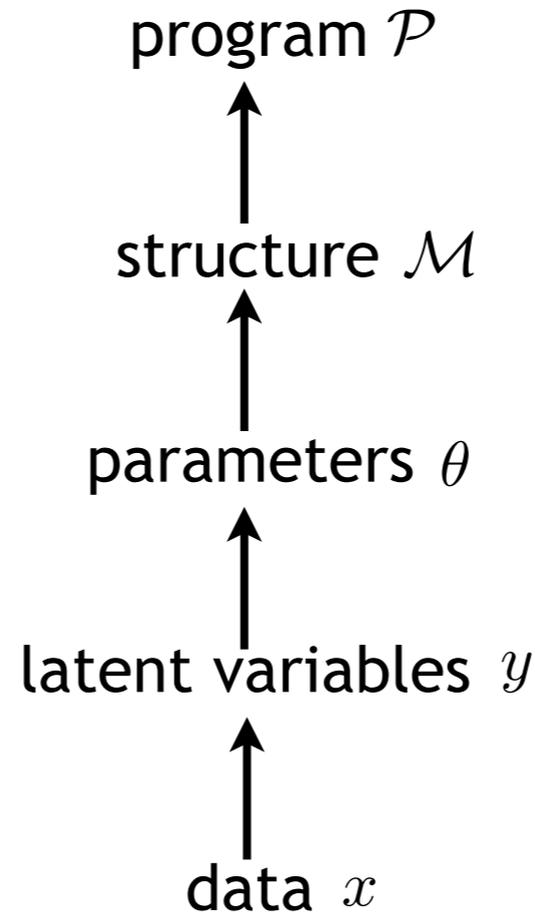
cognitive science

neuroscience

theory

experiments

theory



PROBABILISTIC INFERENCE AND LEARNING

cognitive science

neuroscience

theory

experiments

theory

✓

✓?

✓

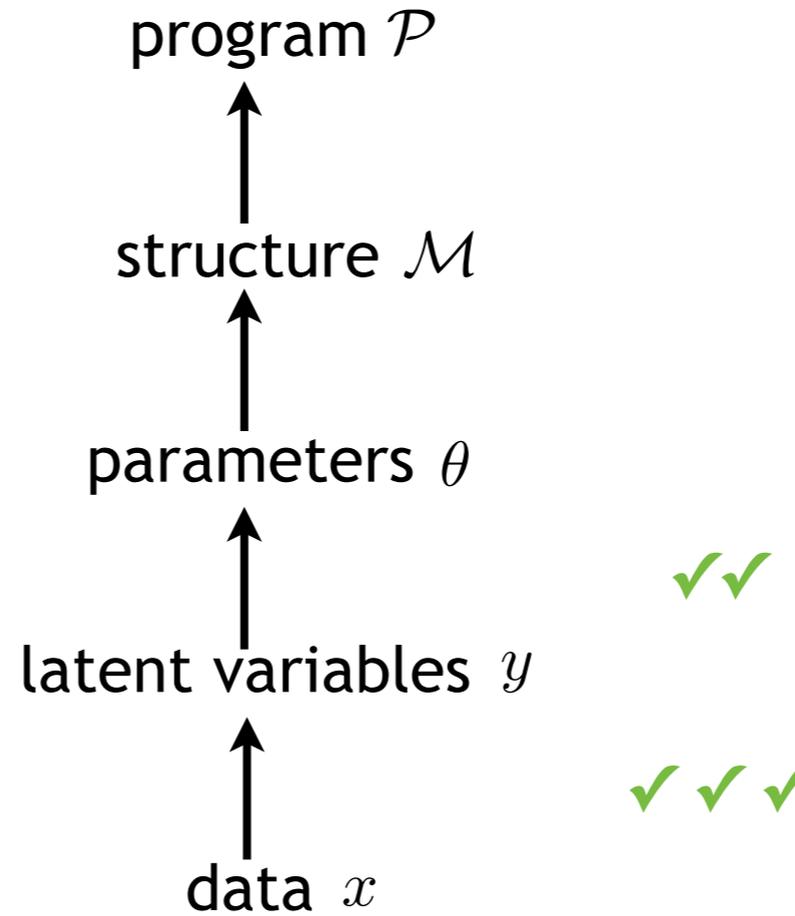
✓

✓

✓

✓

✓



PROBABILISTIC INFERENCE AND LEARNING

cognitive science

neuroscience

theory

experiments

theory

✓

?

✓

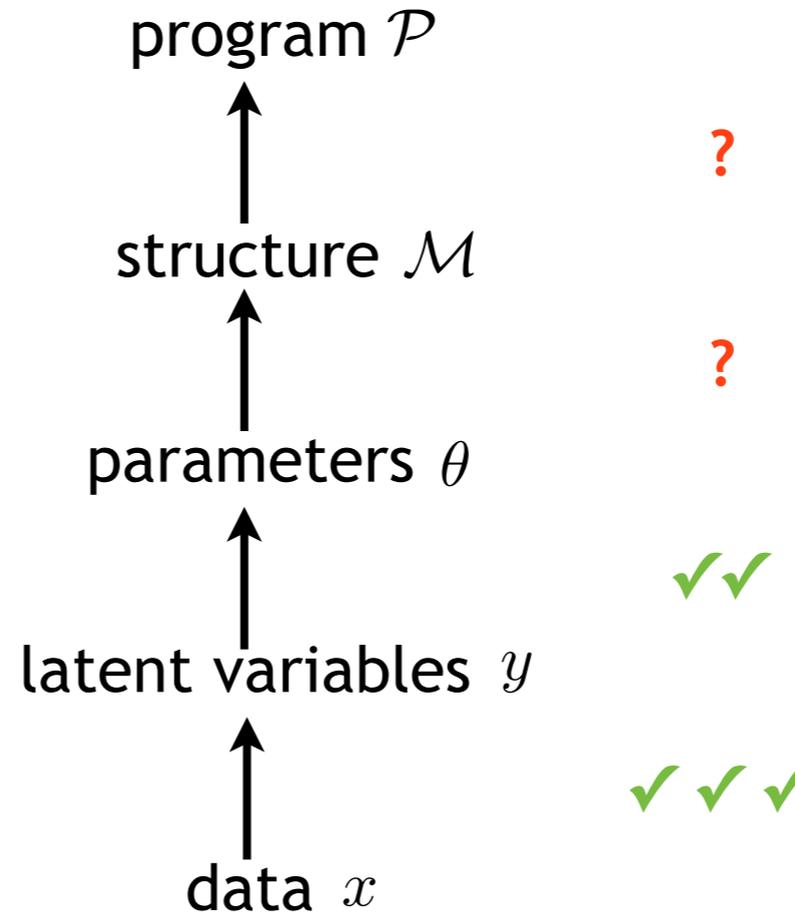
✓

✓

✓

✓

✓



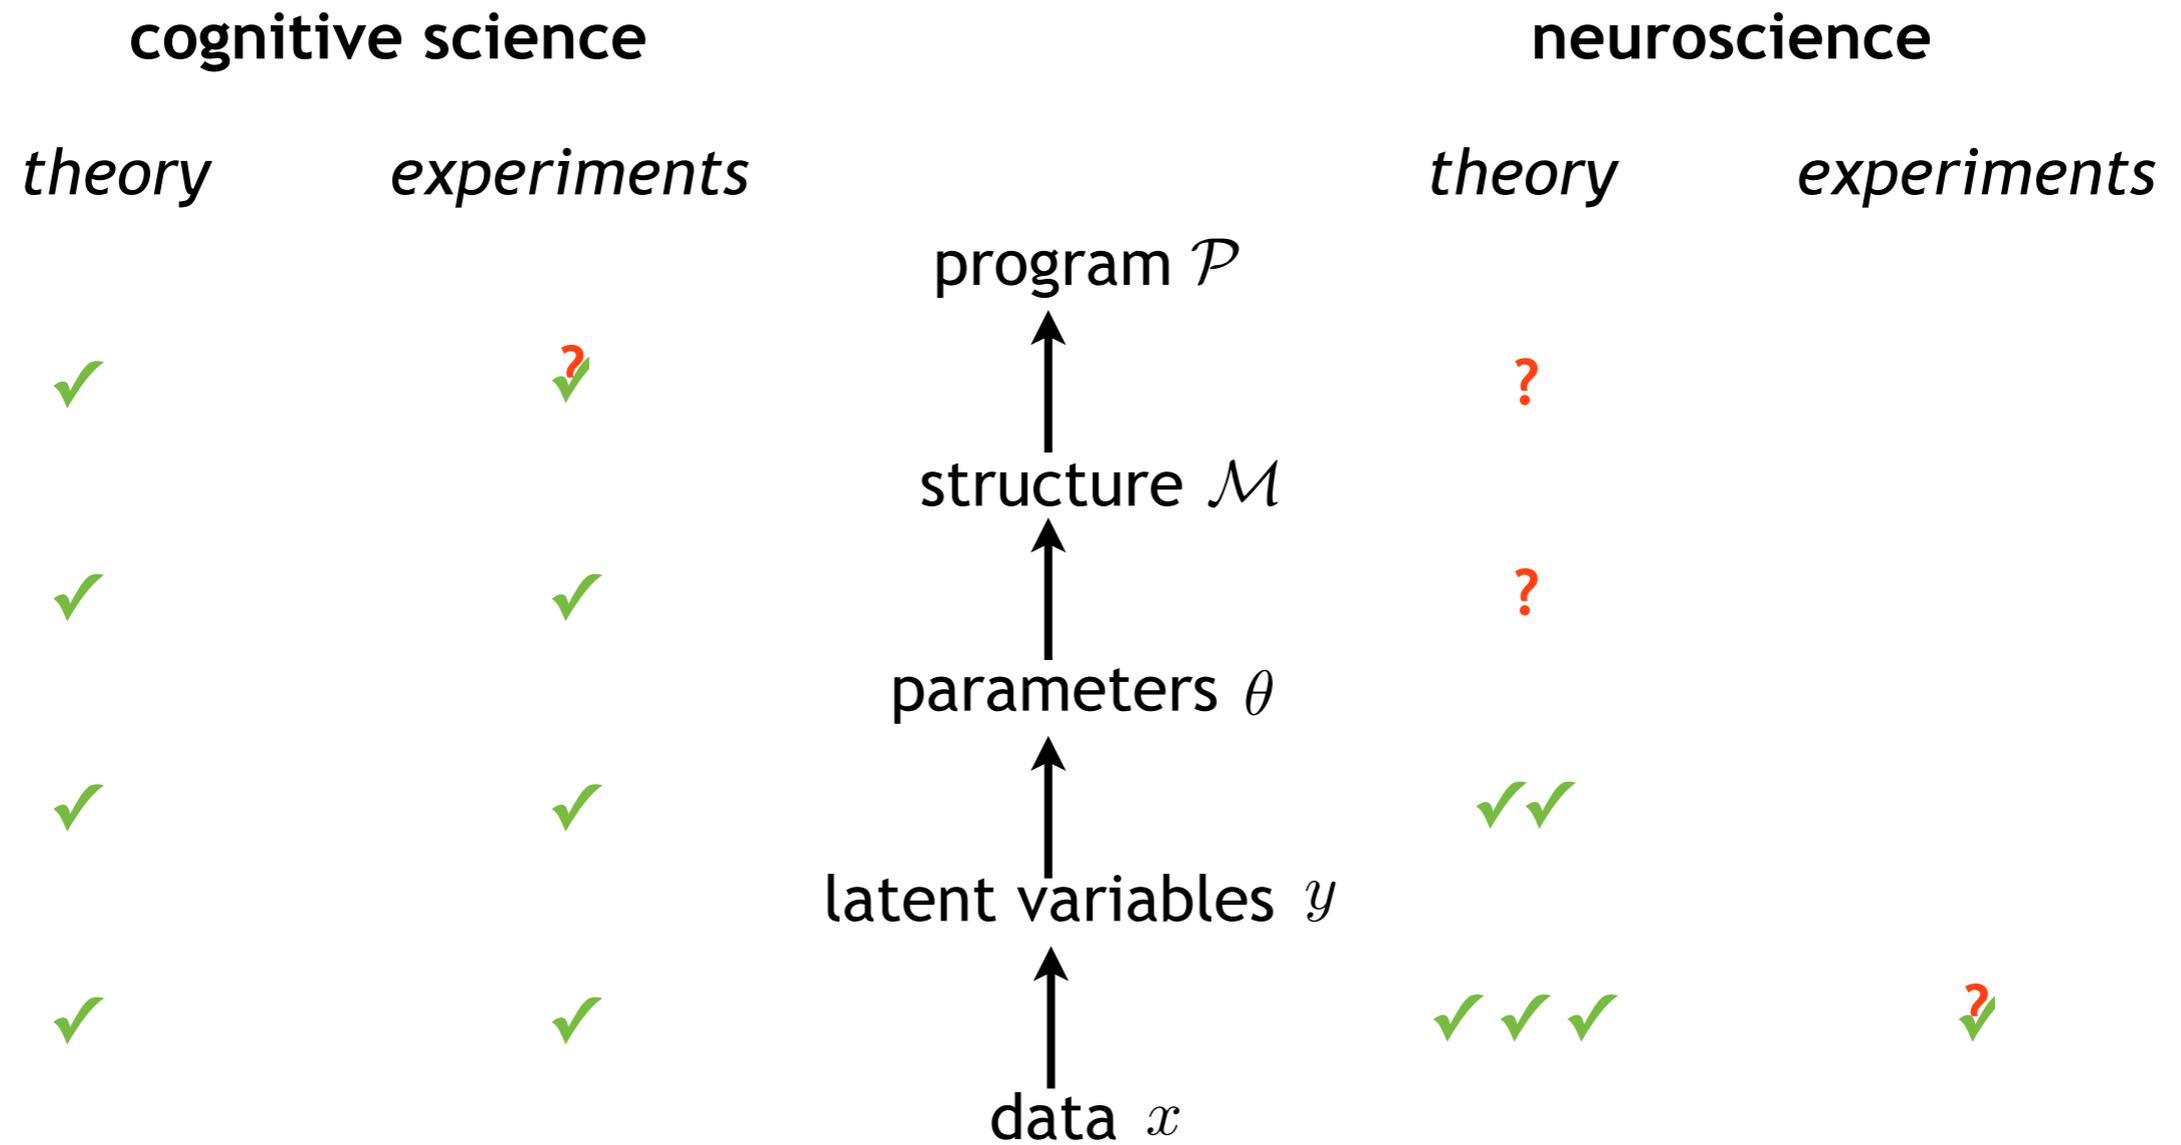
?

?

✓✓

✓✓✓

PROBABILISTIC INFERENCE AND LEARNING



PROBABILISTIC INFERENCE AND LEARNING

cognitive science

neuroscience

theory

experiments

theory

experiments

✓

?

✓

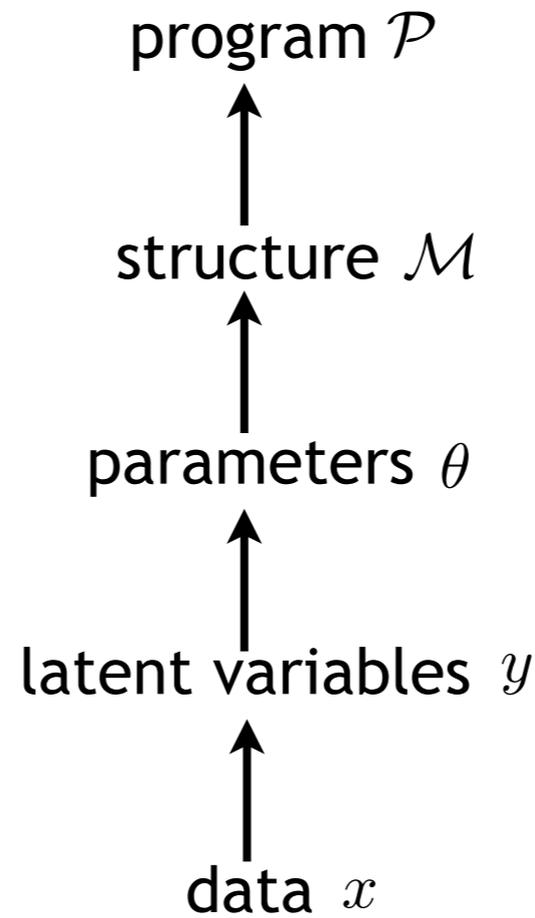
✓

✓

✓

✓

✓



?

?

✓✓

?

✓✓✓

?

PROBABILISTIC INFERENCE AND LEARNING

cognitive science

neuroscience

theory

experiments

theory

experiments

✓

✓?

✓

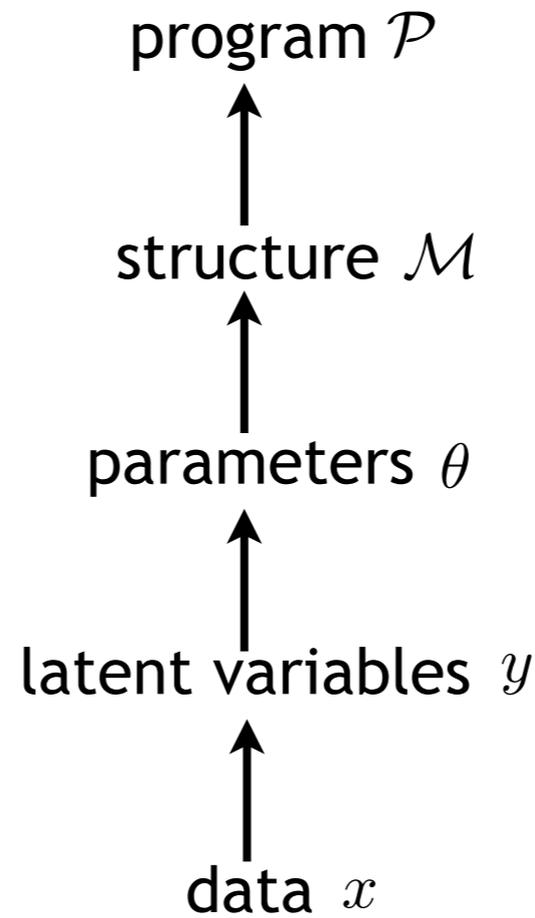
✓

✓

✓

✓

✓



?

?

?

?

✓✓

✓?

✓✓✓

✓?

BAYESIAN INFERENCE IN THE BRAIN BY MCMC SAMPLING

BAYESIAN INFERENCE IN THE BRAIN BY MCMC SAMPLING



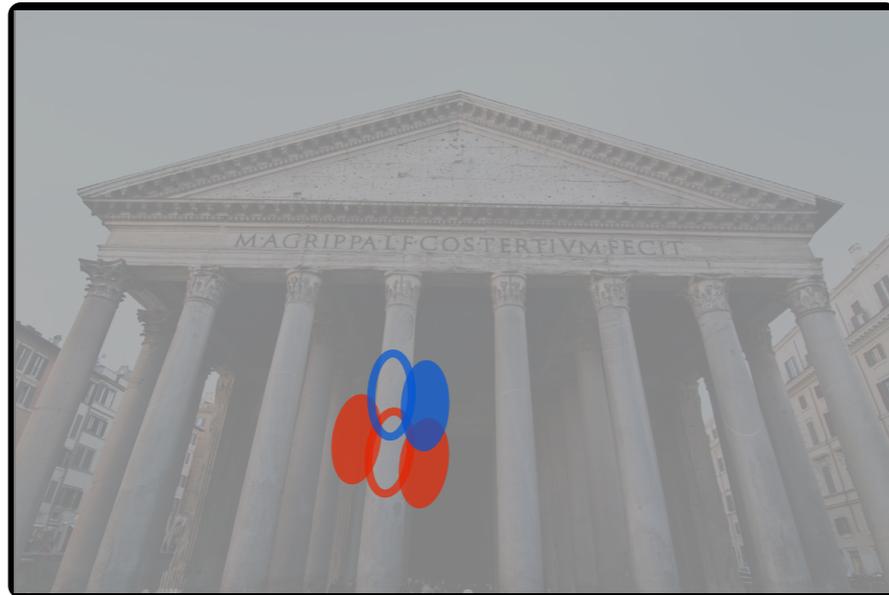
stimulus

BAYESIAN INFERENCE IN THE BRAIN BY MCMC SAMPLING



stimulus

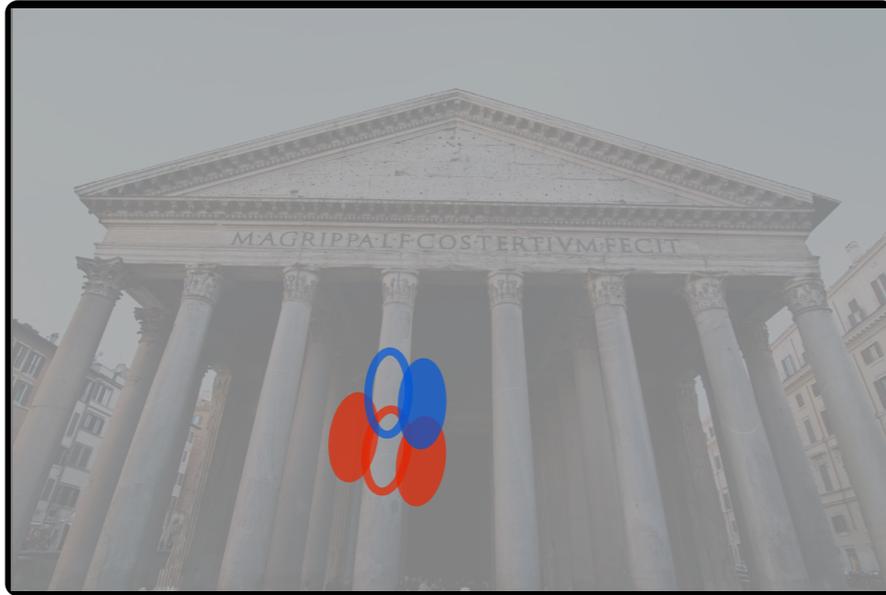
BAYESIAN INFERENCE IN THE BRAIN BY MCMC SAMPLING



stimulus

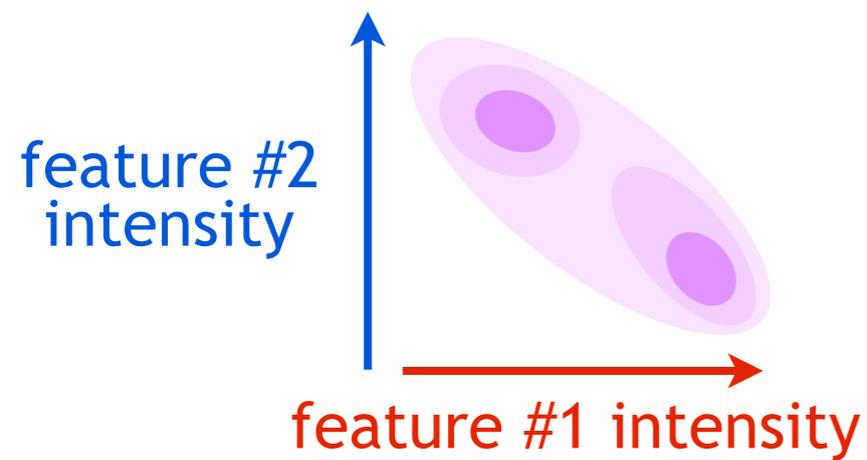
BAYESIAN INFERENCE IN THE BRAIN BY MCMC SAMPLING

ideal observer



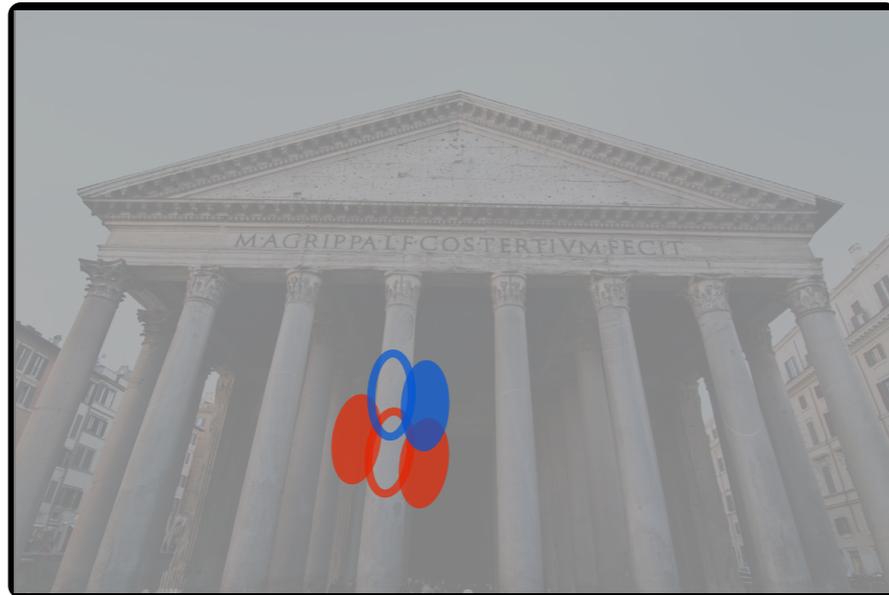
stimulus

$$P(\text{feature 1, feature 2} \mid \text{stimulus})$$



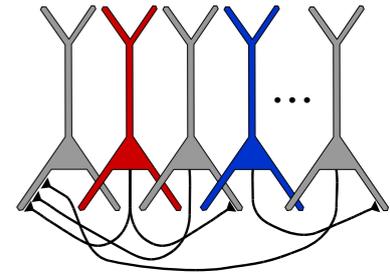
BAYESIAN INFERENCE IN THE BRAIN BY MCMC SAMPLING

ideal observer

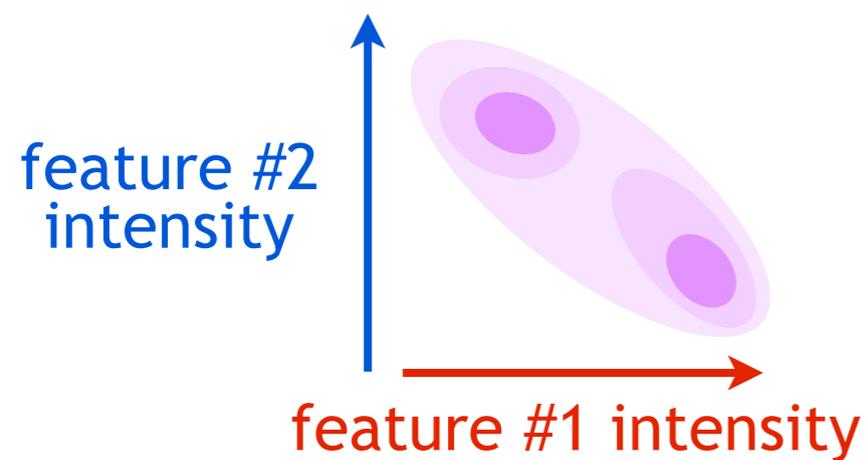


stimulus

visual cortex

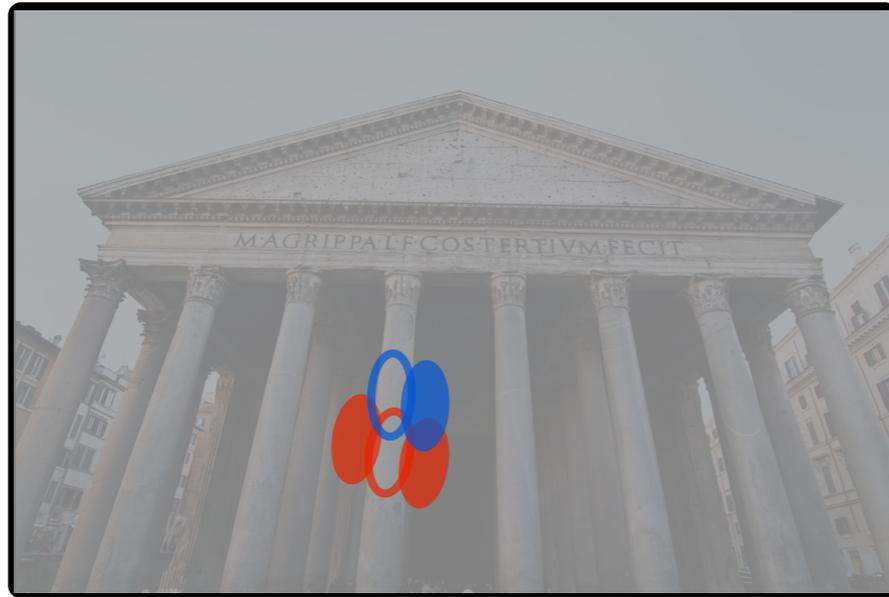


$$P(\text{feature 1, feature 2} \mid \text{stimulus})$$



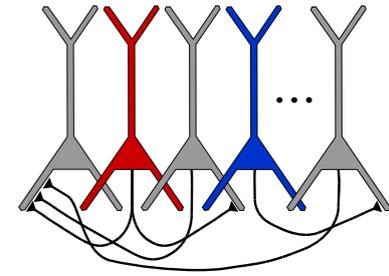
BAYESIAN INFERENCE IN THE BRAIN BY MCMC SAMPLING

ideal observer



stimulus

visual cortex

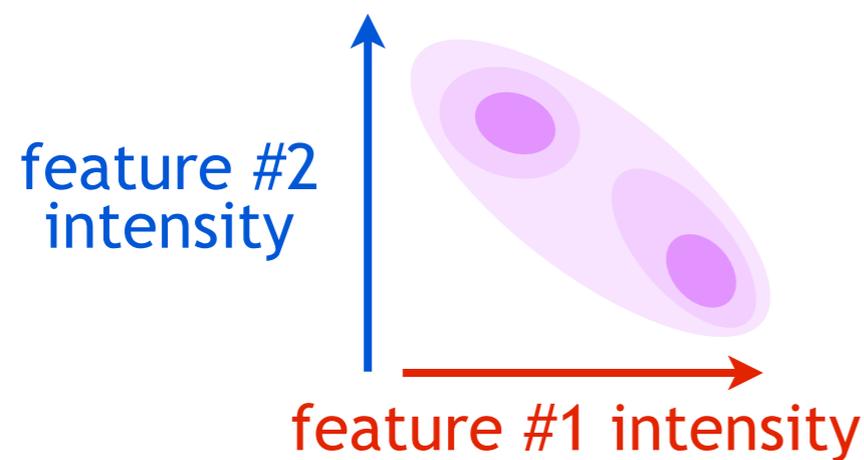


cell #1
response



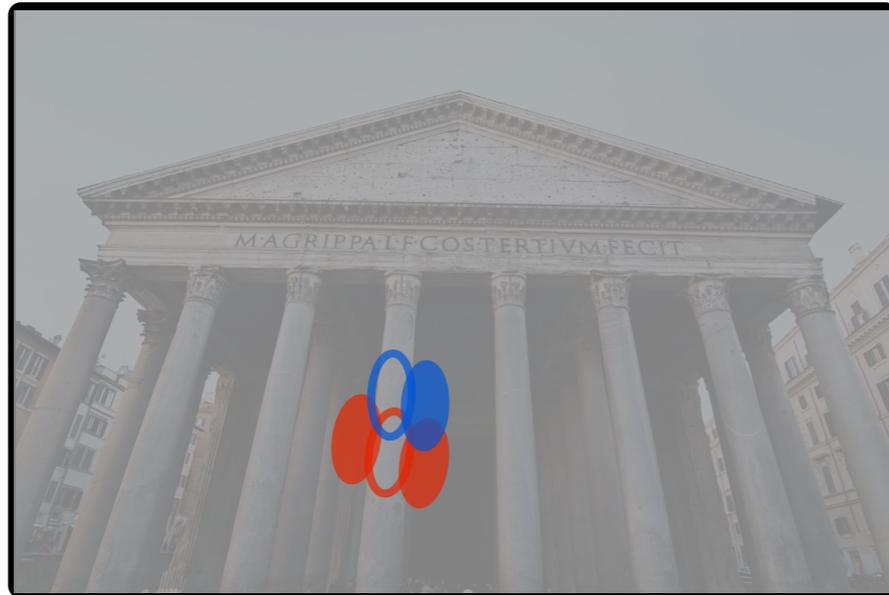
time

$$P(\text{feature 1, feature 2} \mid \text{stimulus})$$



BAYESIAN INFERENCE IN THE BRAIN BY MCMC SAMPLING

ideal observer



stimulus

visual cortex

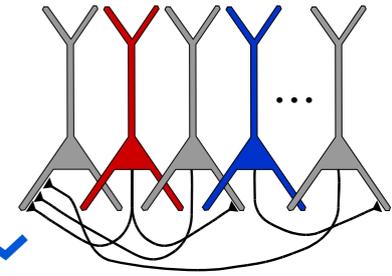
cell #2
response



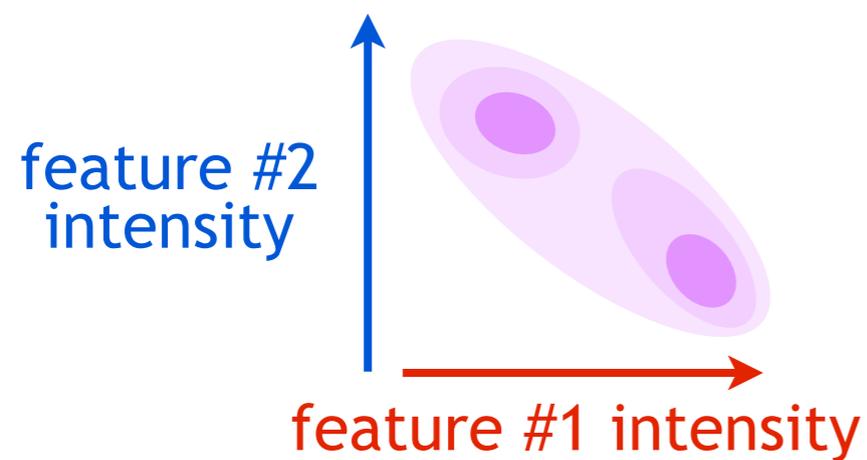
cell #1
response



time

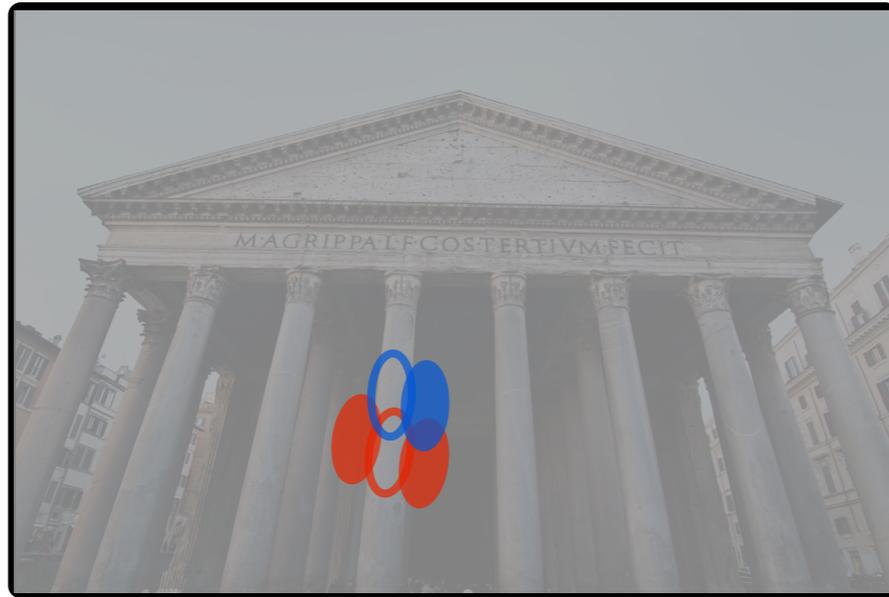


$$P(\text{feature 1, feature 2} \mid \text{stimulus})$$



BAYESIAN INFERENCE IN THE BRAIN BY MCMC SAMPLING

ideal observer



stimulus

visual cortex

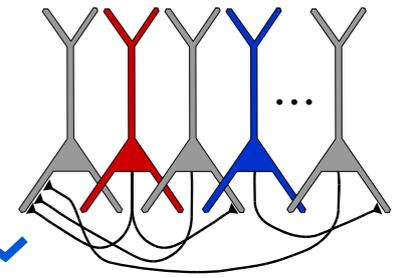
cell #2
response



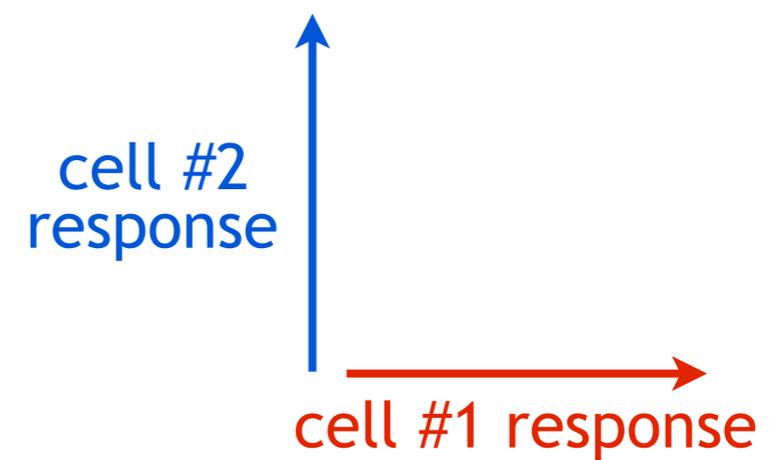
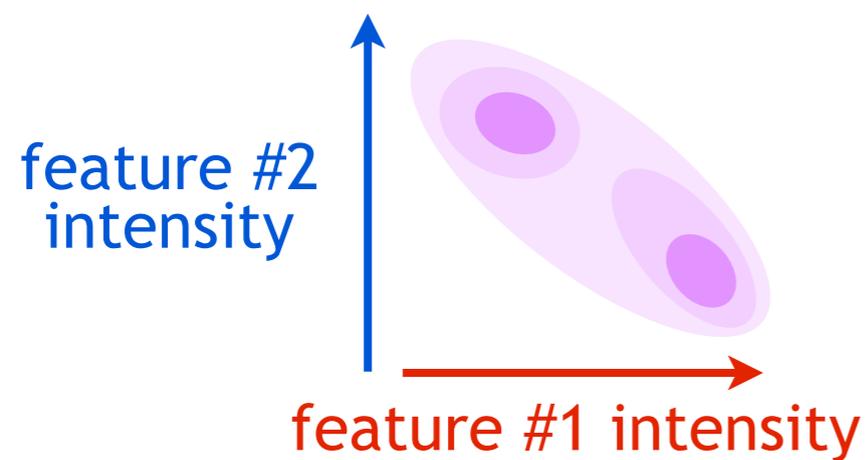
cell #1
response



time

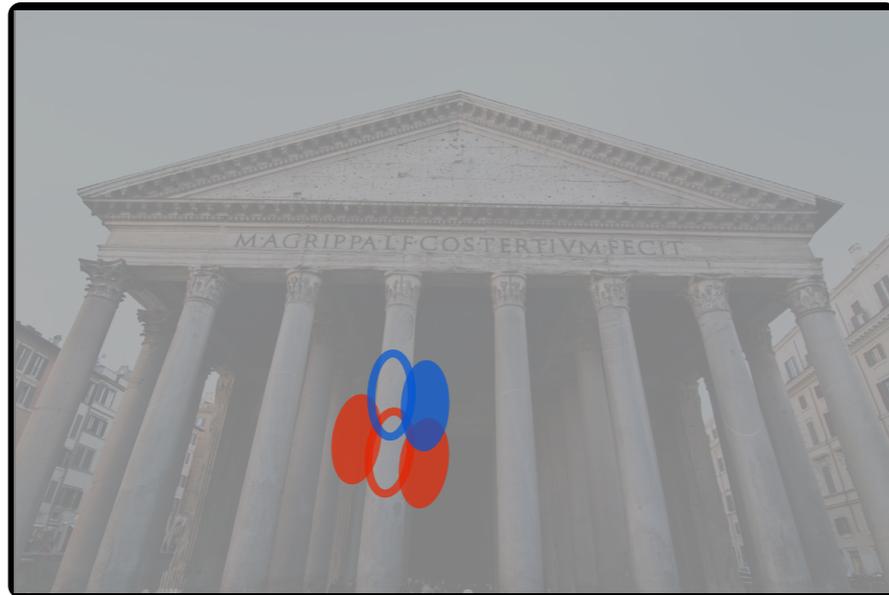


$$P(\text{feature 1, feature 2} \mid \text{stimulus})$$



BAYESIAN INFERENCE IN THE BRAIN BY MCMC SAMPLING

ideal observer



stimulus

visual cortex

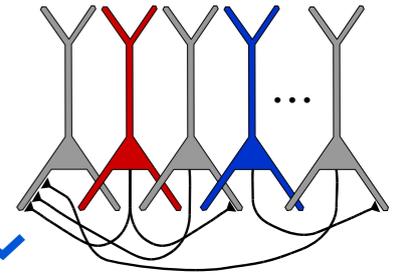
cell #2
response



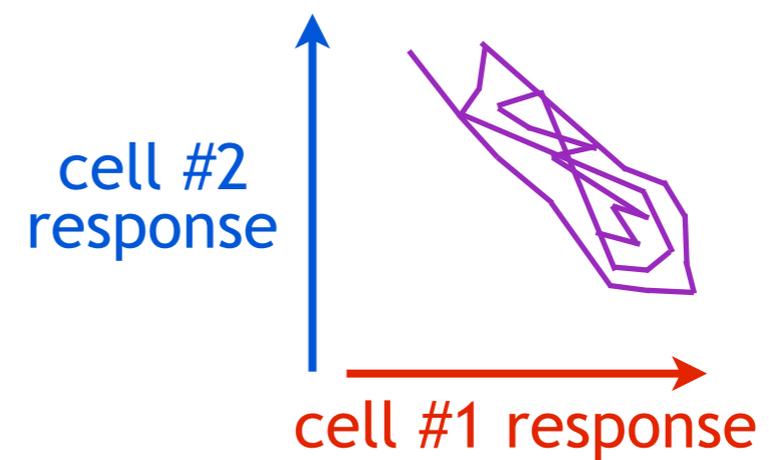
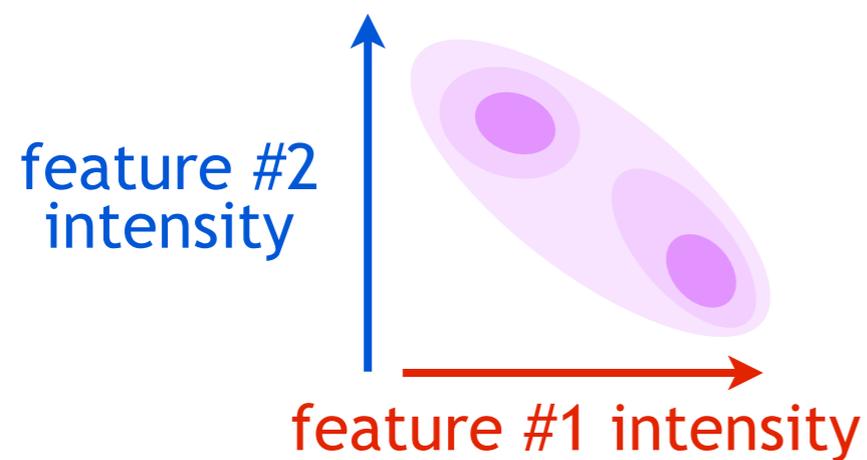
cell #1
response



time

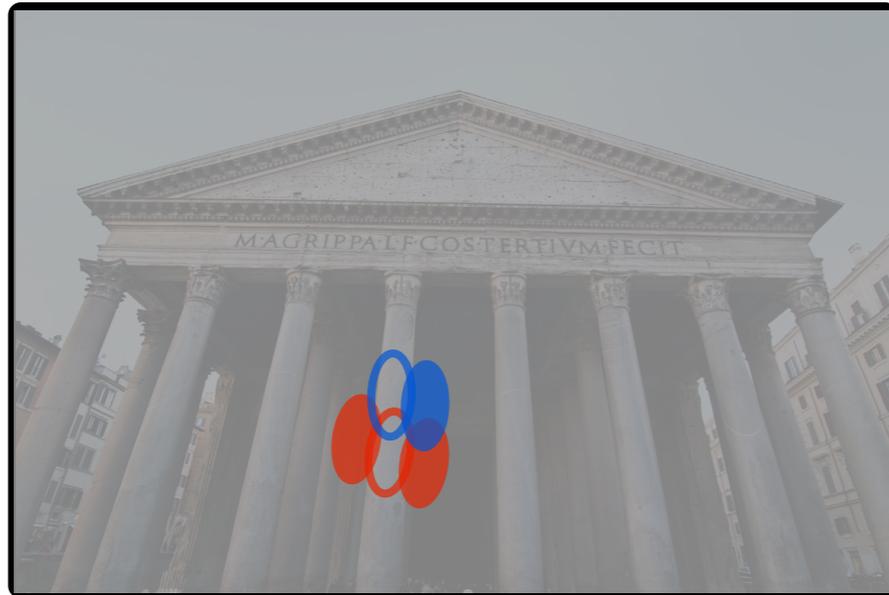


$$P(\text{feature 1, feature 2} \mid \text{stimulus})$$



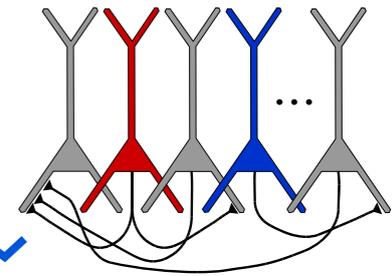
BAYESIAN INFERENCE IN THE BRAIN BY MCMC SAMPLING

ideal observer



stimulus

visual cortex



cell #2
response

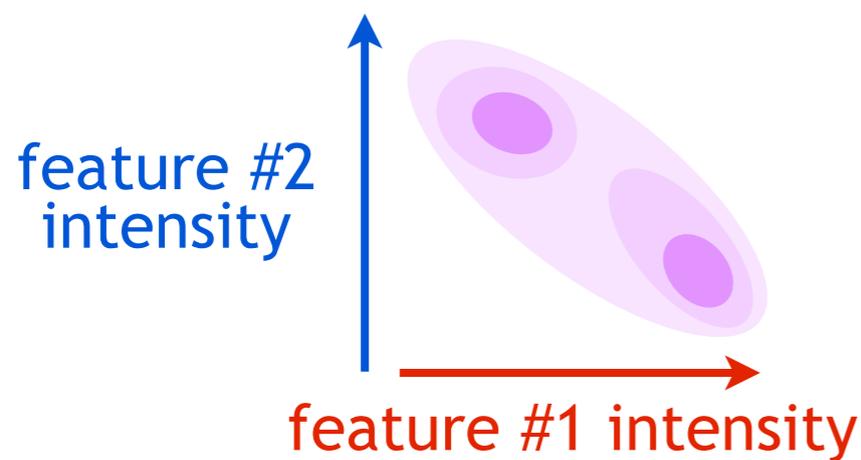


cell #1
response

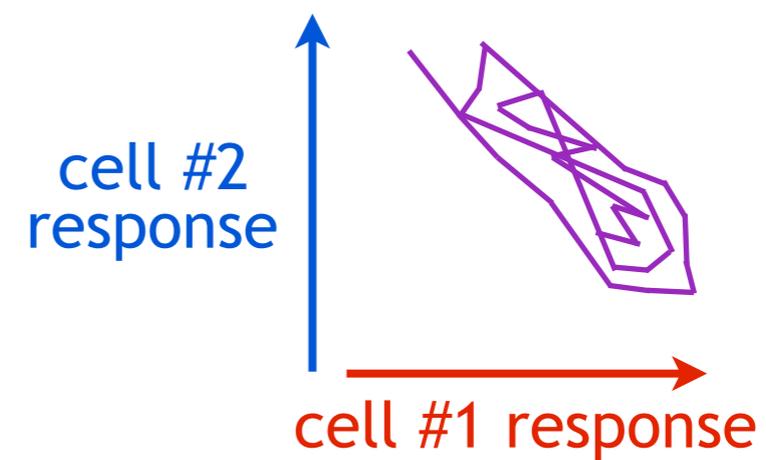


time

$$P(\text{feature 1, feature 2} \mid \text{stimulus})$$

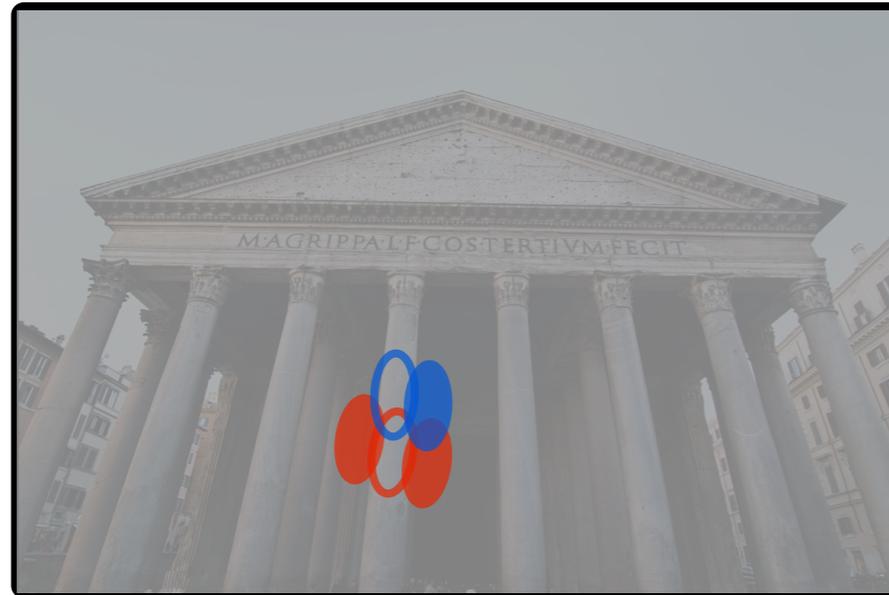


$$P(\text{response 1, response 2} \mid \text{stimulus})$$



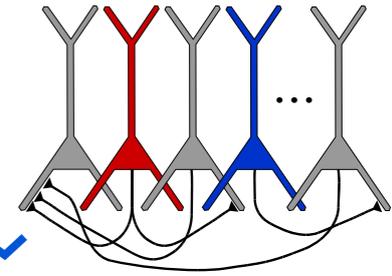
BAYESIAN INFERENCE IN THE BRAIN BY MCMC SAMPLING

ideal observer



stimulus

visual cortex



cell #2
response

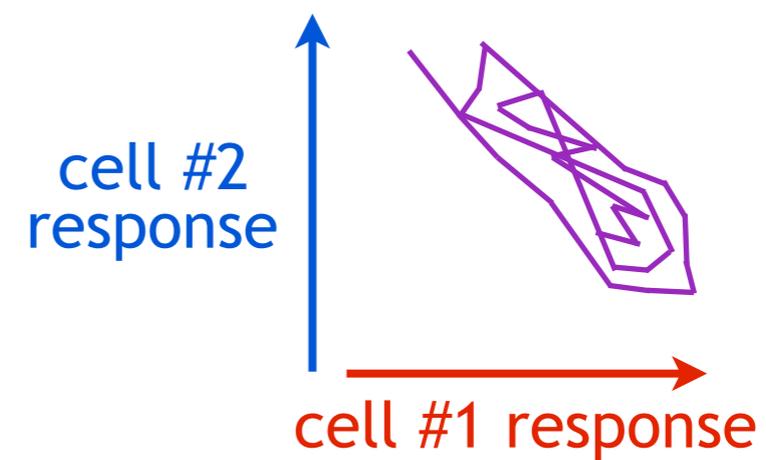
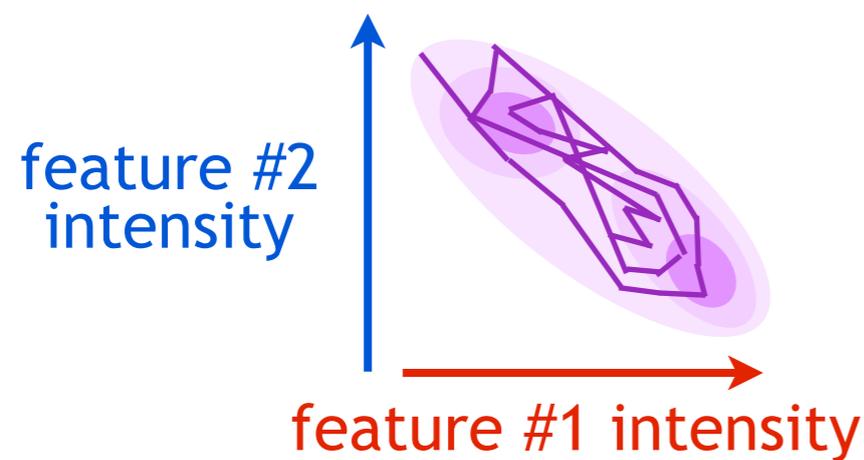


cell #1
response



time

$$P(\text{feature 1, feature 2} \mid \text{stimulus}) = P(\text{response 1, response 2} \mid \text{stimulus})$$



Fiser et al, TICS 2010

see also:

Hinton & Sejnowski, PDP 1986; Hinton et al, Science 1995; Dayan 1999; Hoyer & Hyvarinen, NIPS 2003, Lee & Mumford 2003

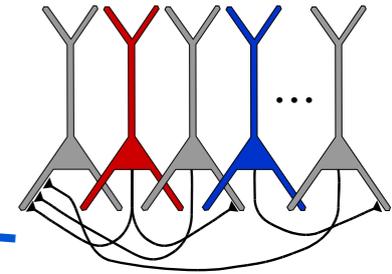
BAYESIAN INFERENCE IN THE BRAIN BY MCMC SAMPLING

ideal observer

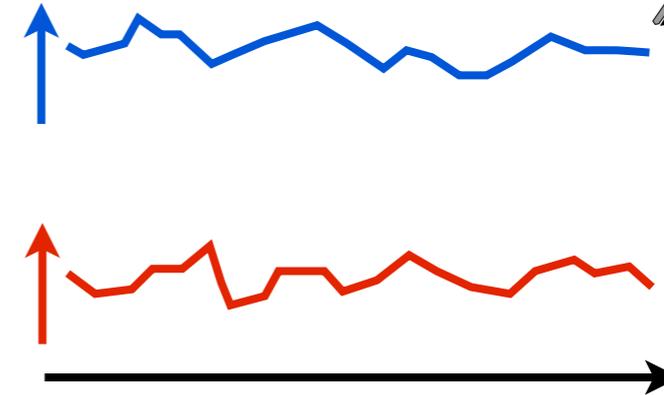


stimulus

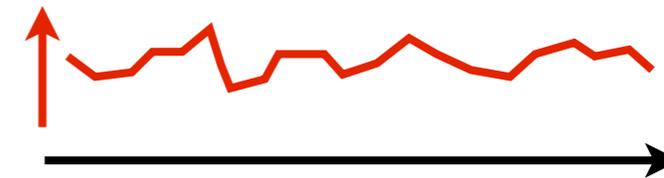
visual cortex



cell #2
response

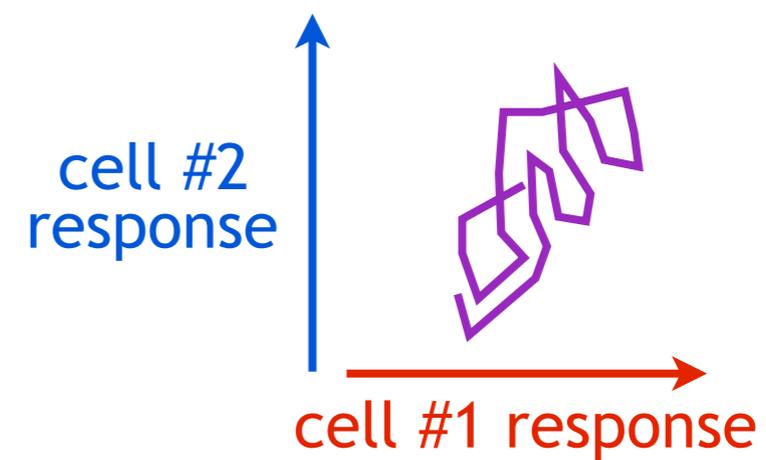
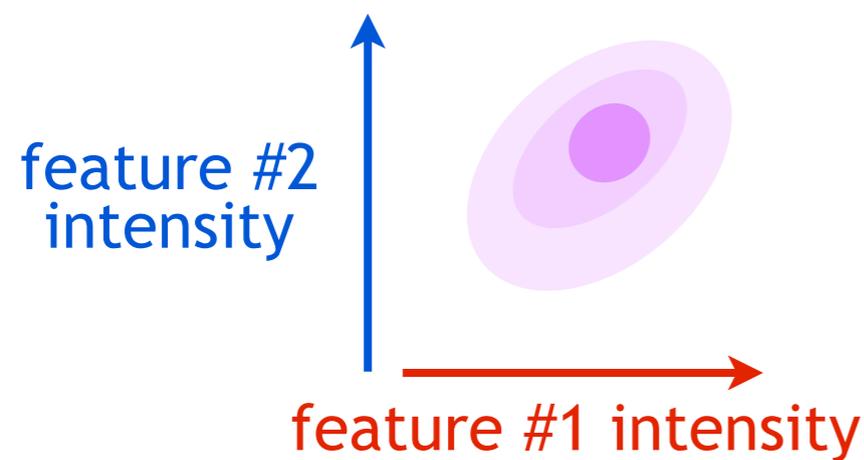


cell #1
response



time

$$P(\text{feature 1, feature 2} \mid \text{stimulus}) = P(\text{response 1, response 2} \mid \text{stimulus})$$



Fiser et al, TICS 2010

see also:

Hinton & Sejnowski, PDP 1986; Hinton et al, Science 1995; Dayan 1999; Hoyer & Hyvarinen, NIPS 2003, Lee & Mumford 2003

BAYESIAN INFERENCE IN THE BRAIN BY MCMC SAMPLING

ideal observer



stimulus

visual cortex

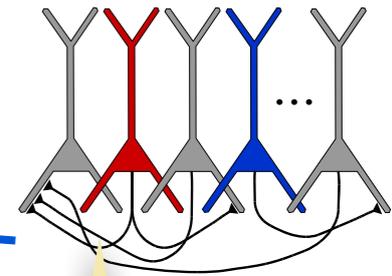
cell #2
response



cell #1
response

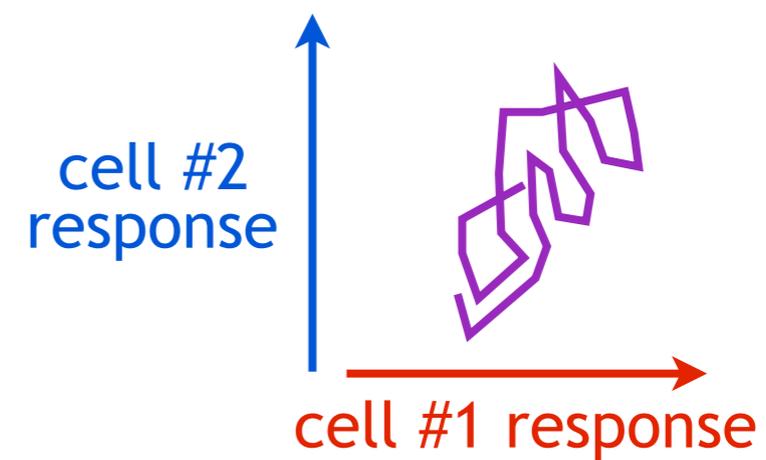
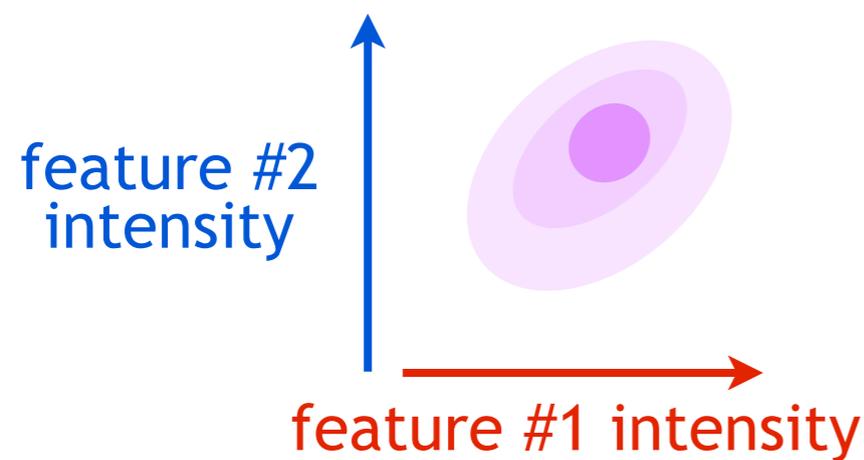


time



amortised inference:
network parameters
(synaptic weights, etc.)
remain the same

$$P(\text{feature 1, feature 2} \mid \text{stimulus}) = P(\text{response 1, response 2} \mid \text{stimulus})$$



Fiser et al, TICS 2010

see also:

Hinton & Sejnowski, PDP 1986; Hinton et al, Science 1995; Dayan 1999; Hoyer & Hyvarinen, NIPS 2003, Lee & Mumford 2003

NEURAL HALLMARKS OF SAMPLING

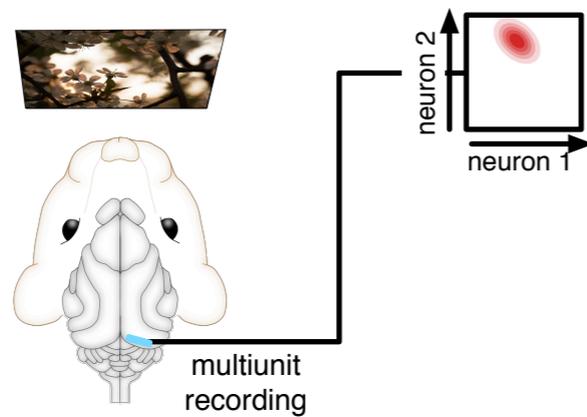
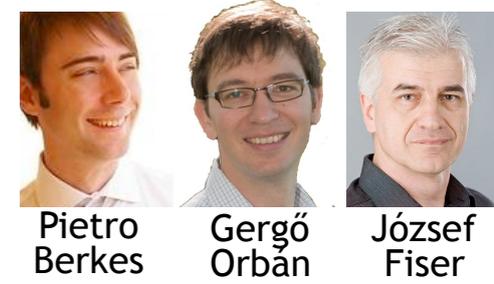


Pietro
Berkes

Gergő
Orbán

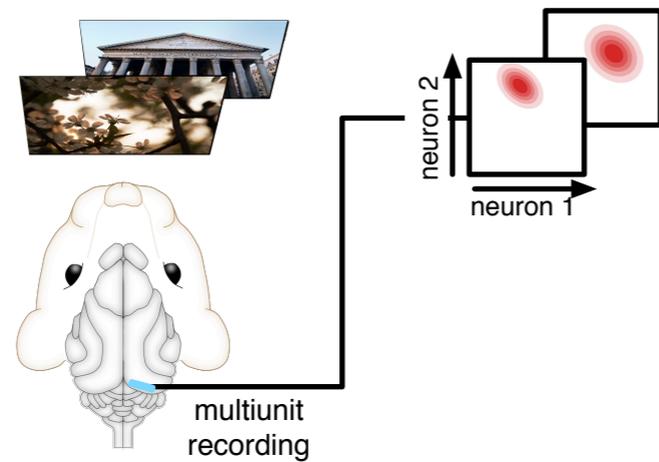
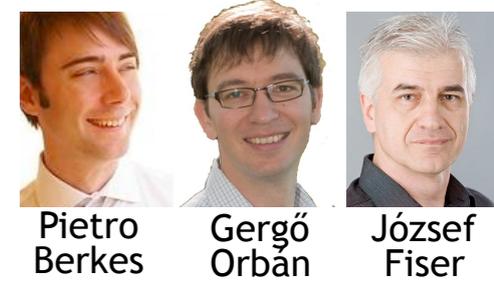
József
Fiser

NEURAL HALLMARKS OF SAMPLING



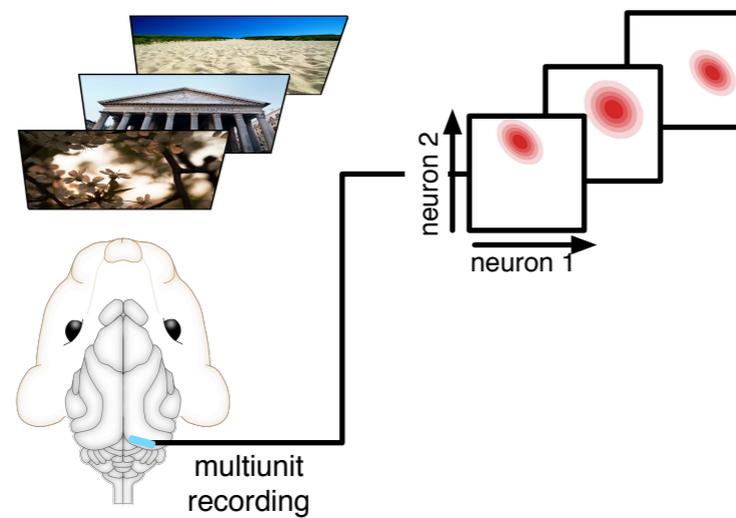
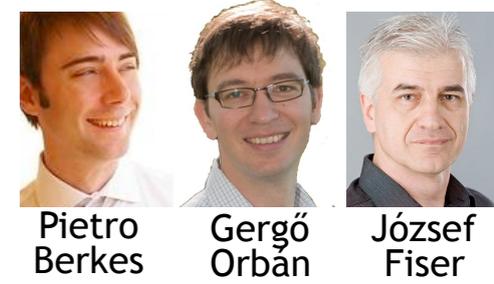
$$P_{\text{model}}(\text{feat}|\text{stim}_1)$$

NEURAL HALLMARKS OF SAMPLING



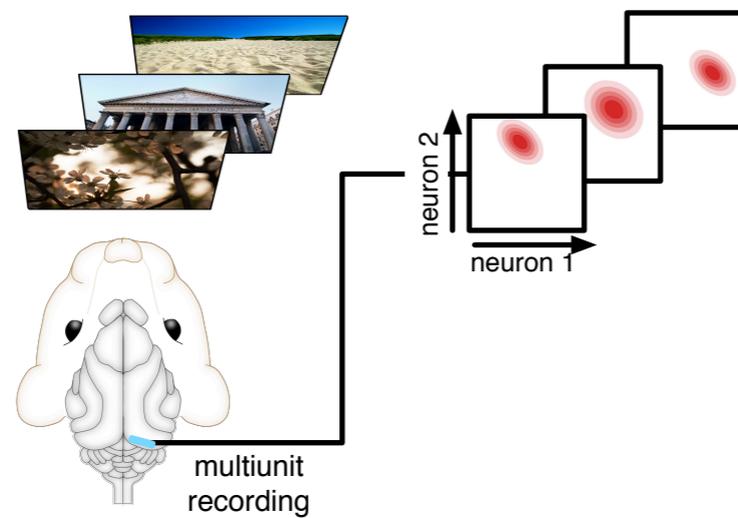
$$P_{\text{model}}(\text{feat}|\text{stim}_2)$$
$$P_{\text{model}}(\text{feat}|\text{stim}_1)$$

NEURAL HALLMARKS OF SAMPLING



$$P_{\text{model}}(\text{feat}|\text{stim}_3)$$
$$P_{\text{model}}(\text{feat}|\text{stim}_2)$$
$$P_{\text{model}}(\text{feat}|\text{stim}_1)$$

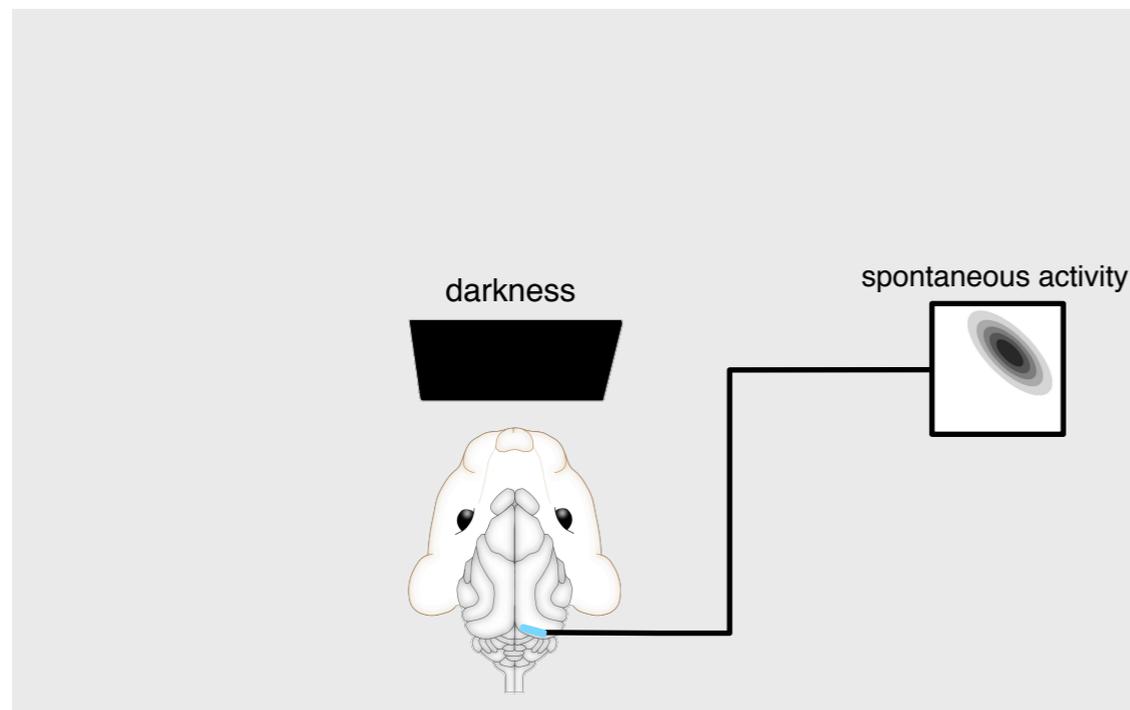
NEURAL HALLMARKS OF SAMPLING



$$P_{\text{model}}(\text{feat}|\text{stim}_3)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_2)$$

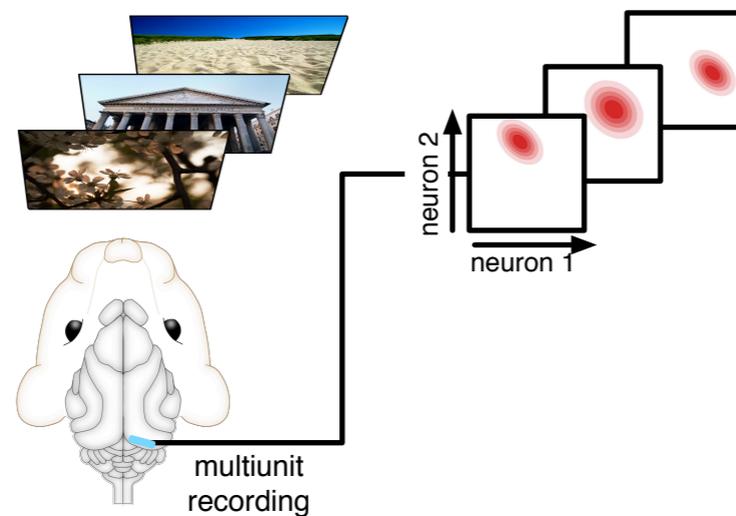
$$P_{\text{model}}(\text{feat}|\text{stim}_1)$$



$$P_{\text{model}}(\text{feat}|\text{stim} = 0)$$

Berkes et al, Science 2011

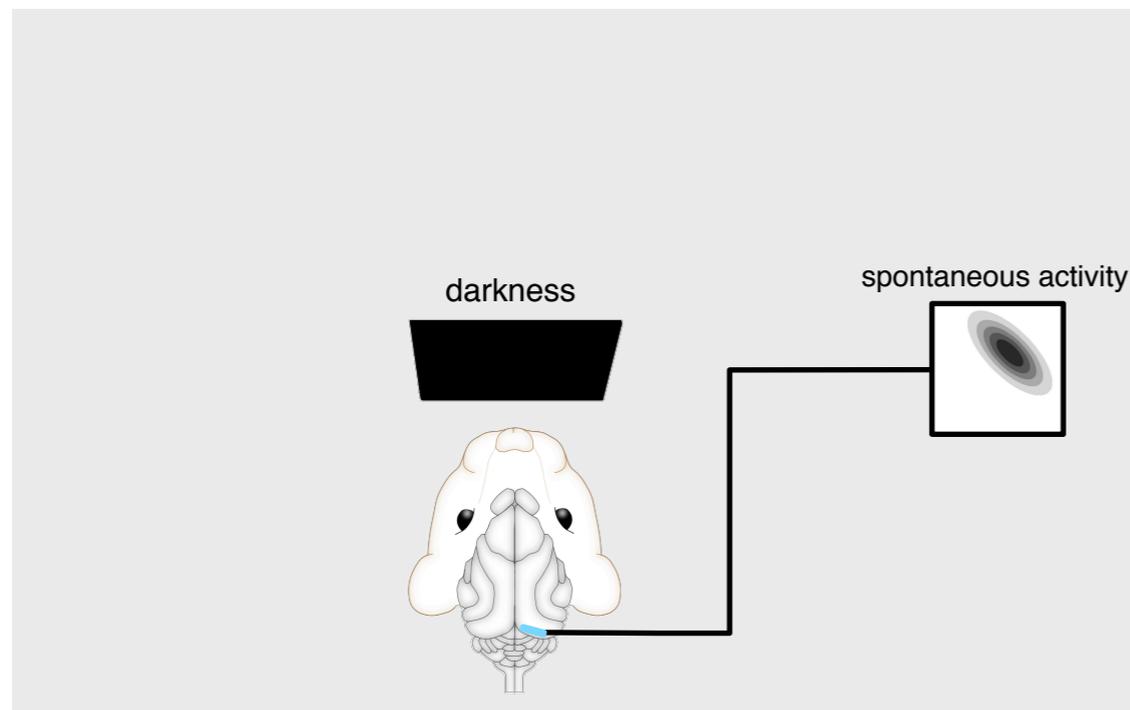
NEURAL HALLMARKS OF SAMPLING



$$P_{\text{model}}(\text{feat}|\text{stim}_3)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_2)$$

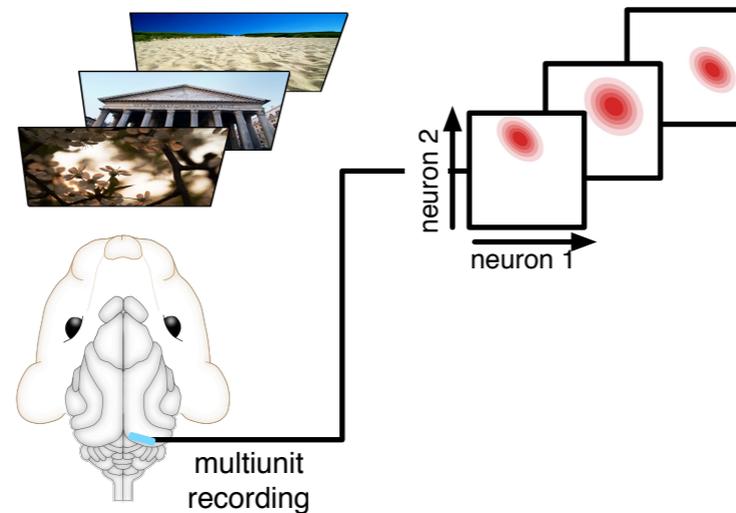
$$P_{\text{model}}(\text{feat}|\text{stim}_1)$$



$$P_{\text{model}}(\text{feat}|\text{stim} = 0) \simeq P_{\text{model}}(\text{feat})$$

Berkes et al, Science 2011

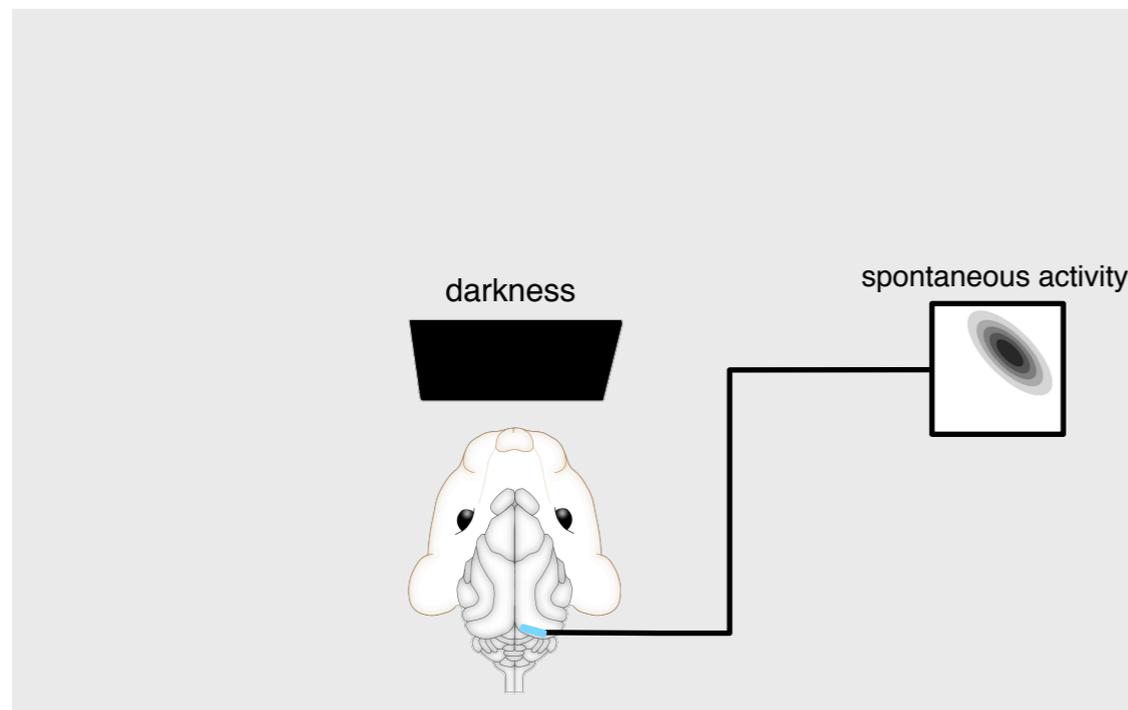
NEURAL HALLMARKS OF SAMPLING



$$P_{\text{model}}(\text{feat}|\text{stim}_3)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_2)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_1)$$

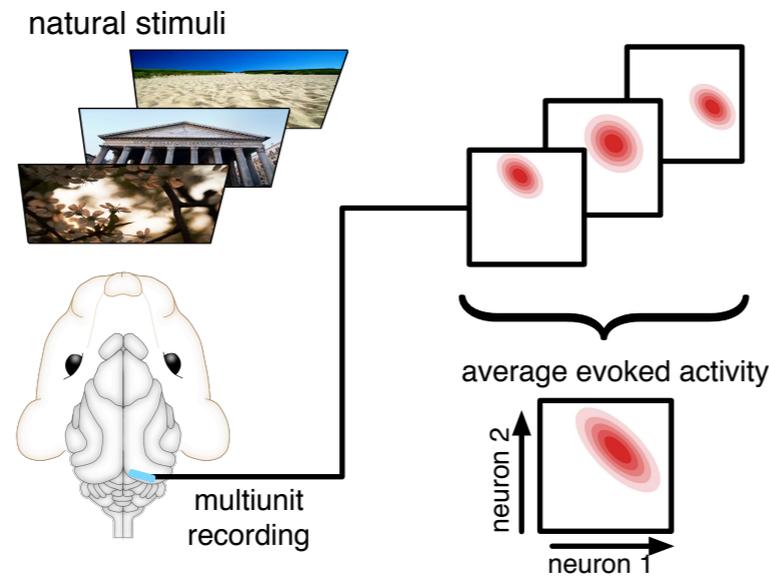
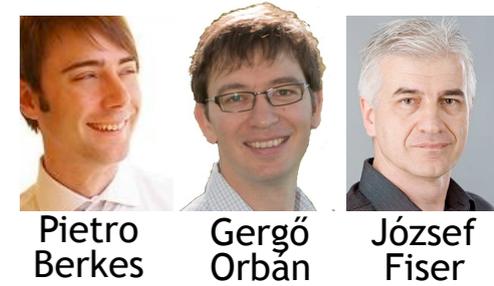


$$P_{\text{model}}(\text{feat}|\text{stim} = 0) \simeq P_{\text{model}}(\text{feat})$$

$$\simeq \int P_{\text{model}}(\text{feat}|\text{stim}) P_{\text{model}}(\text{stim}) d\text{stim}$$

Berkes et al, Science 2011

NEURAL HALLMARKS OF SAMPLING

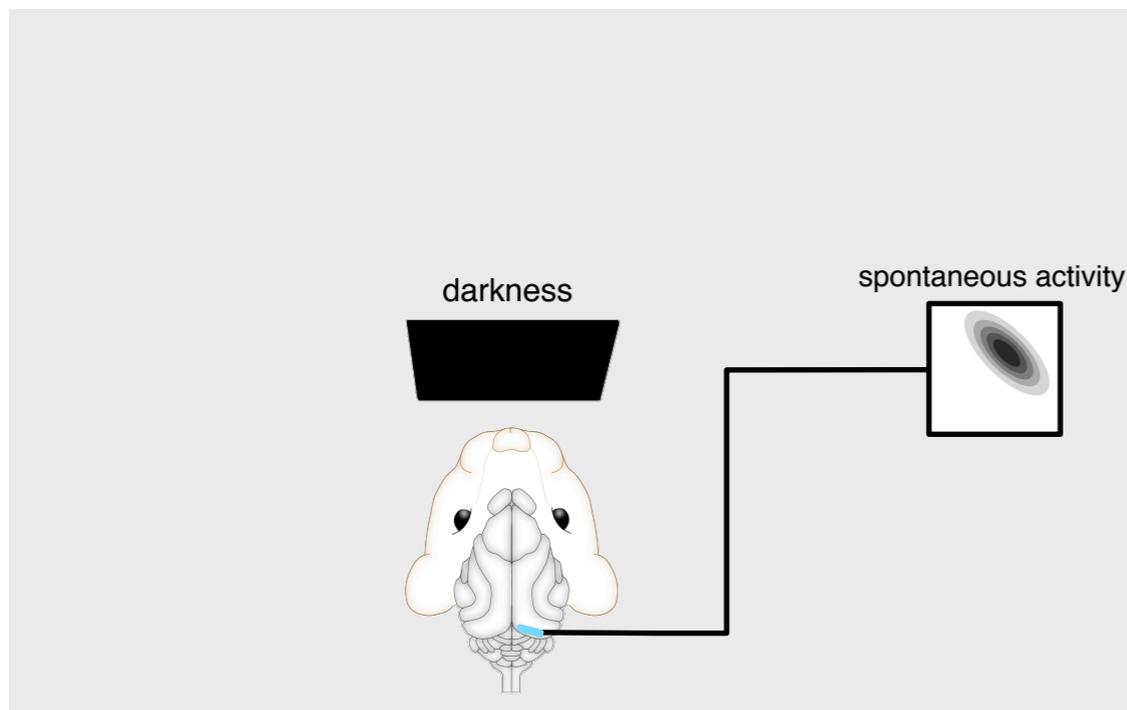


$$P_{\text{model}}(\text{feat}|\text{stim}_3)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_2)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_1)$$

$$\int P_{\text{model}}(\text{feat}|\text{stim}) P_{\text{natural}}(\text{stim}) d\text{stim}$$

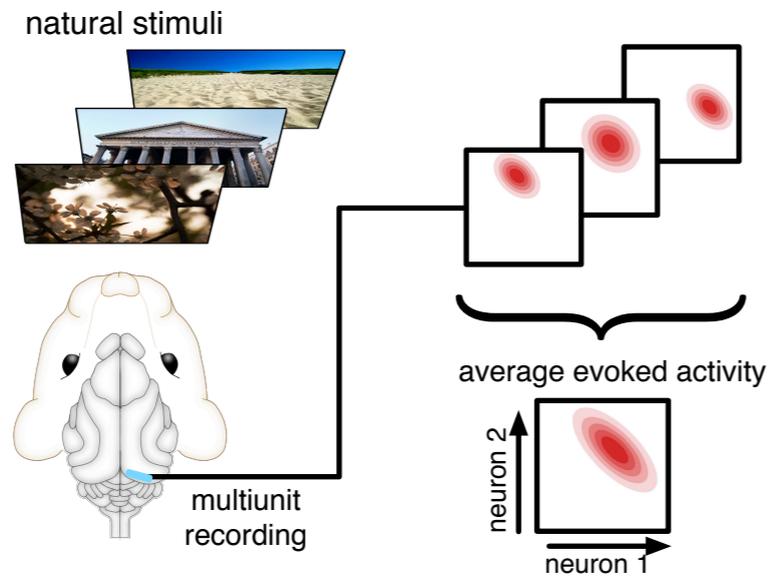
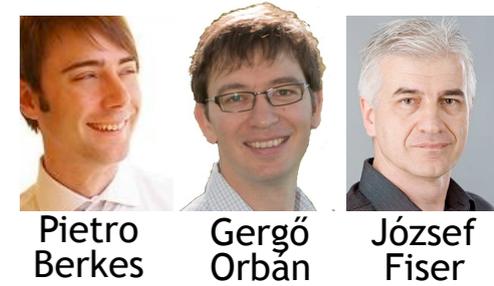


$$P_{\text{model}}(\text{feat}|\text{stim} = 0) \simeq P_{\text{model}}(\text{feat})$$

$$\simeq \int P_{\text{model}}(\text{feat}|\text{stim}) P_{\text{model}}(\text{stim}) d\text{stim}$$

Berkes et al, Science 2011

NEURAL HALLMARKS OF SAMPLING

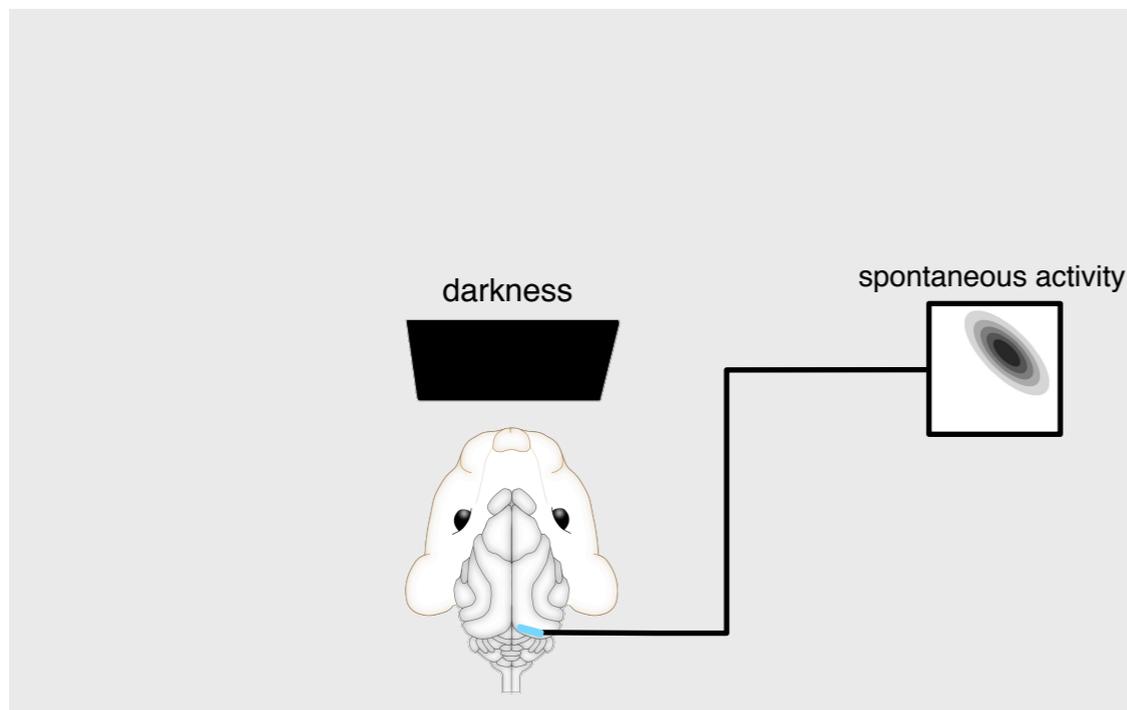


$$P_{\text{model}}(\text{feat}|\text{stim}_3)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_2)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_1)$$

$$\int P_{\text{model}}(\text{feat}|\text{stim}) P_{\text{natural}}(\text{stim}) d\text{stim}$$



$$P_{\text{model}}(\text{feat}|\text{stim} = 0) \simeq P_{\text{model}}(\text{feat})$$

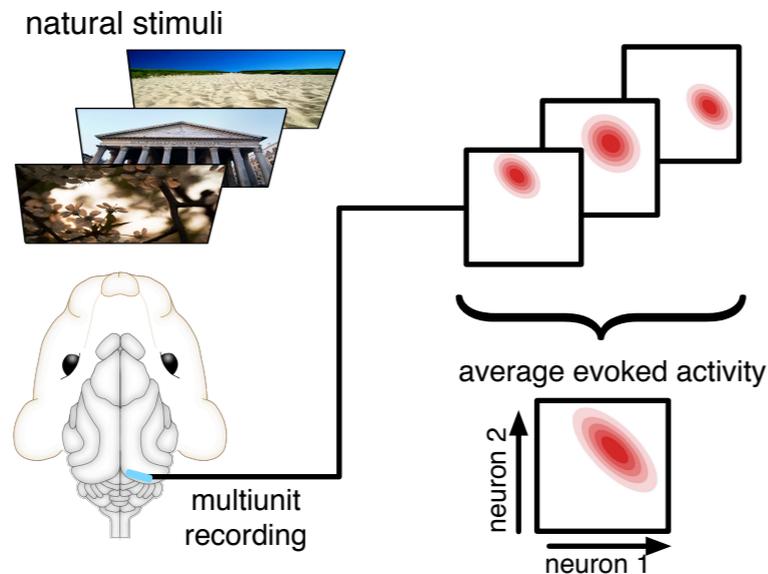
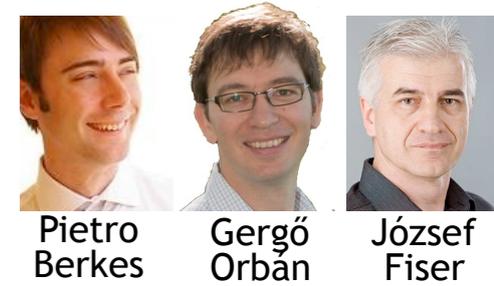
$$\approx \int P_{\text{model}}(\text{feat}|\text{stim}) P_{\text{model}}(\text{stim}) d\text{stim}$$

Berkes et al, Science 2011

NEURAL HALLMARKS OF SAMPLING

+ STATISTICALLY OPTIMAL ADAPTATION

$$P_{\text{model}}(\text{stim}) = P_{\text{natural}}(\text{stim})$$

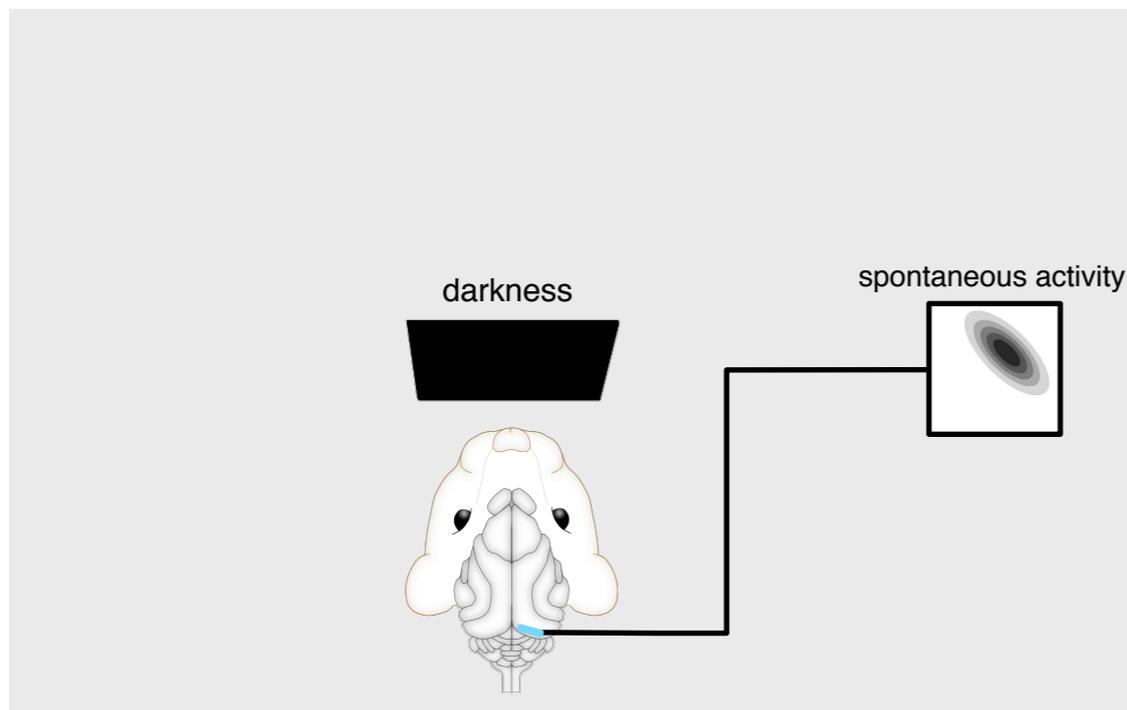


$$P_{\text{model}}(\text{feat}|\text{stim}_3)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_2)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_1)$$

$$\int P_{\text{model}}(\text{feat}|\text{stim}) P_{\text{natural}}(\text{stim}) d\text{stim}$$



$$P_{\text{model}}(\text{feat}|\text{stim} = 0) \simeq P_{\text{model}}(\text{feat})$$

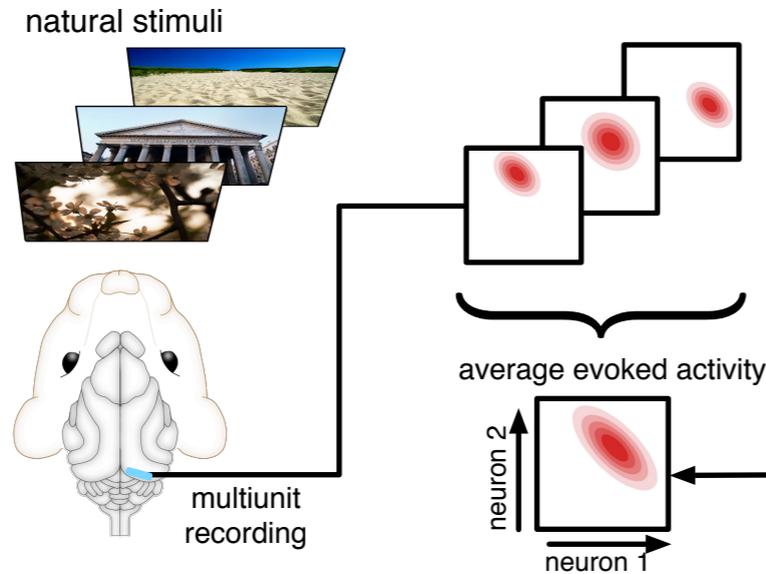
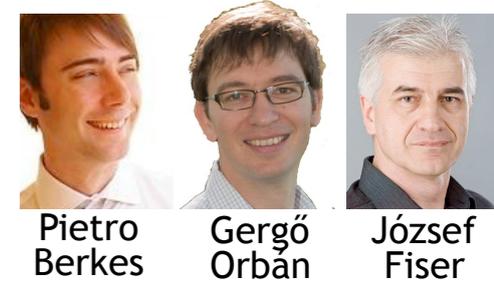
$$\simeq \int P_{\text{model}}(\text{feat}|\text{stim}) P_{\text{model}}(\text{stim}) d\text{stim}$$

Berkes et al, Science 2011

NEURAL HALLMARKS OF SAMPLING

+ STATISTICALLY OPTIMAL ADAPTATION

$$P_{\text{model}}(\text{stim}) = P_{\text{natural}}(\text{stim})$$

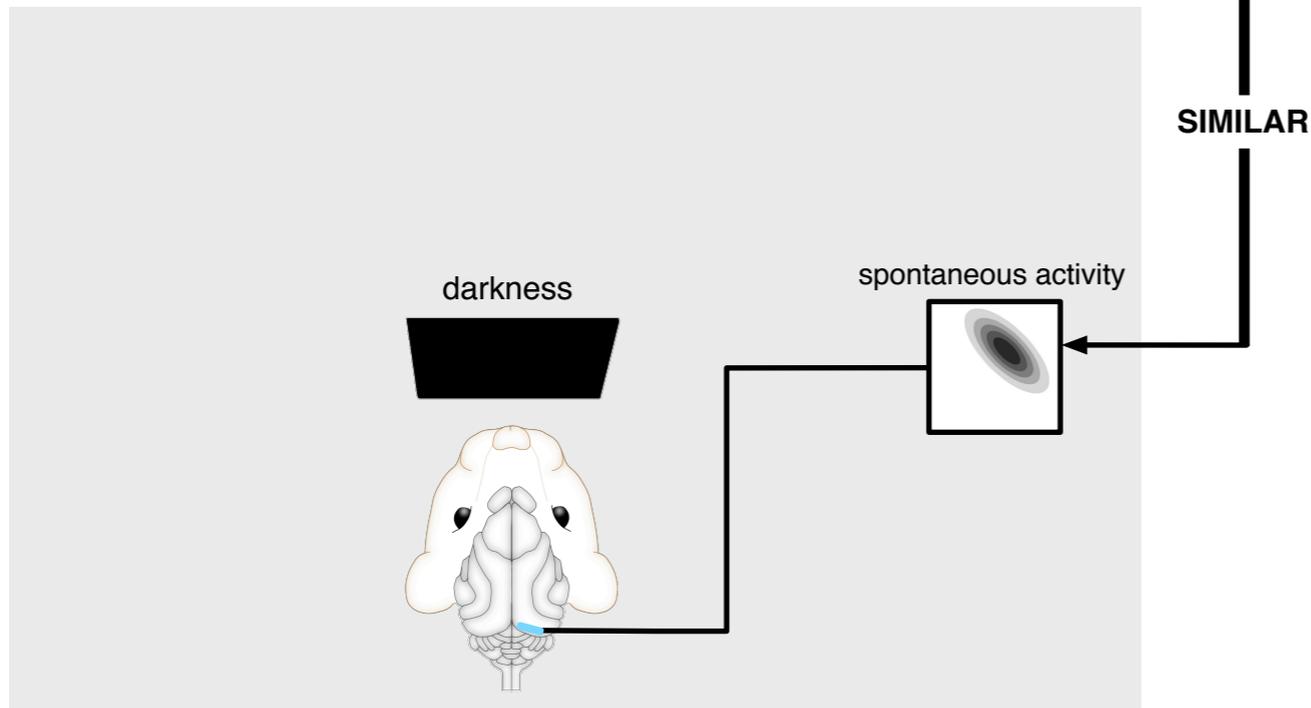


$$P_{\text{model}}(\text{feat}|\text{stim}_3)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_2)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_1)$$

$$\int P_{\text{model}}(\text{feat}|\text{stim}) P_{\text{natural}}(\text{stim}) d\text{stim}$$



$$P_{\text{model}}(\text{feat}|\text{stim} = 0) \simeq P_{\text{model}}(\text{feat})$$

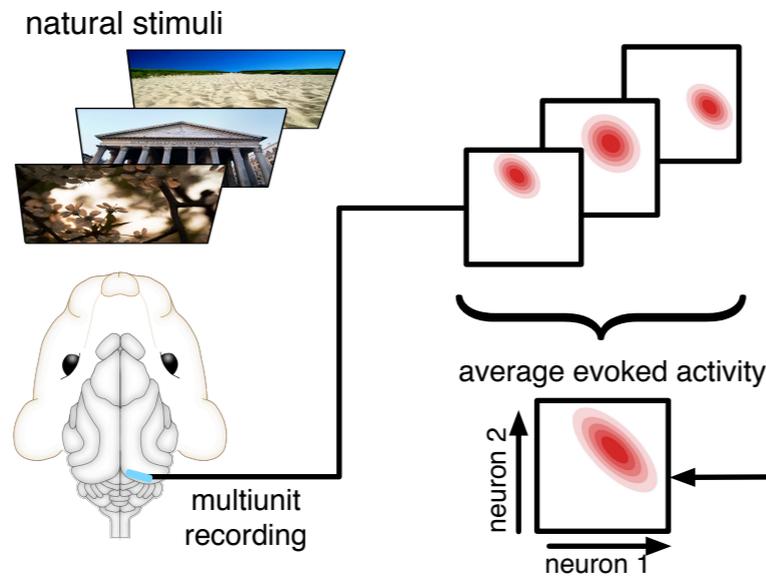
$$\simeq \int P_{\text{model}}(\text{feat}|\text{stim}) P_{\text{model}}(\text{stim}) d\text{stim}$$

Berkes et al, Science 2011

NEURAL HALLMARKS OF SAMPLING

+ STATISTICALLY OPTIMAL ADAPTATION

$$P_{\text{model}}(\text{stim}) = P_{\text{natural}}(\text{stim})$$



$$P_{\text{model}}(\text{feat}|\text{stim}_3)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_2)$$

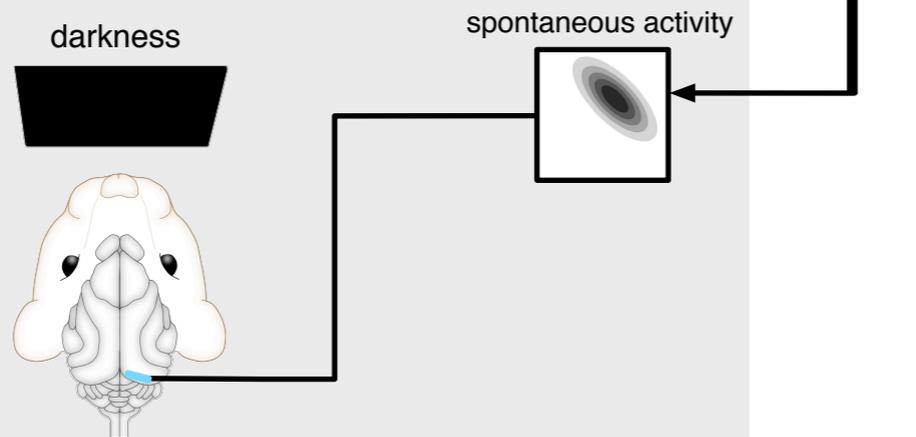
$$P_{\text{model}}(\text{feat}|\text{stim}_1)$$

$$\int P_{\text{model}}(\text{feat}|\text{stim}) P_{\text{natural}}(\text{stim}) d\text{stim}$$

average inference
=
prior expectations

$$P_{\text{model}}(\text{feat}|\text{stim} = 0) \simeq P_{\text{model}}(\text{feat})$$

$$\approx \int P_{\text{model}}(\text{feat}|\text{stim}) P_{\text{model}}(\text{stim}) d\text{stim}$$



Berkes et al, Science 2011

NEURAL HALLMARKS OF SAMPLING

+ STATISTICALLY OPTIMAL ADAPTATION



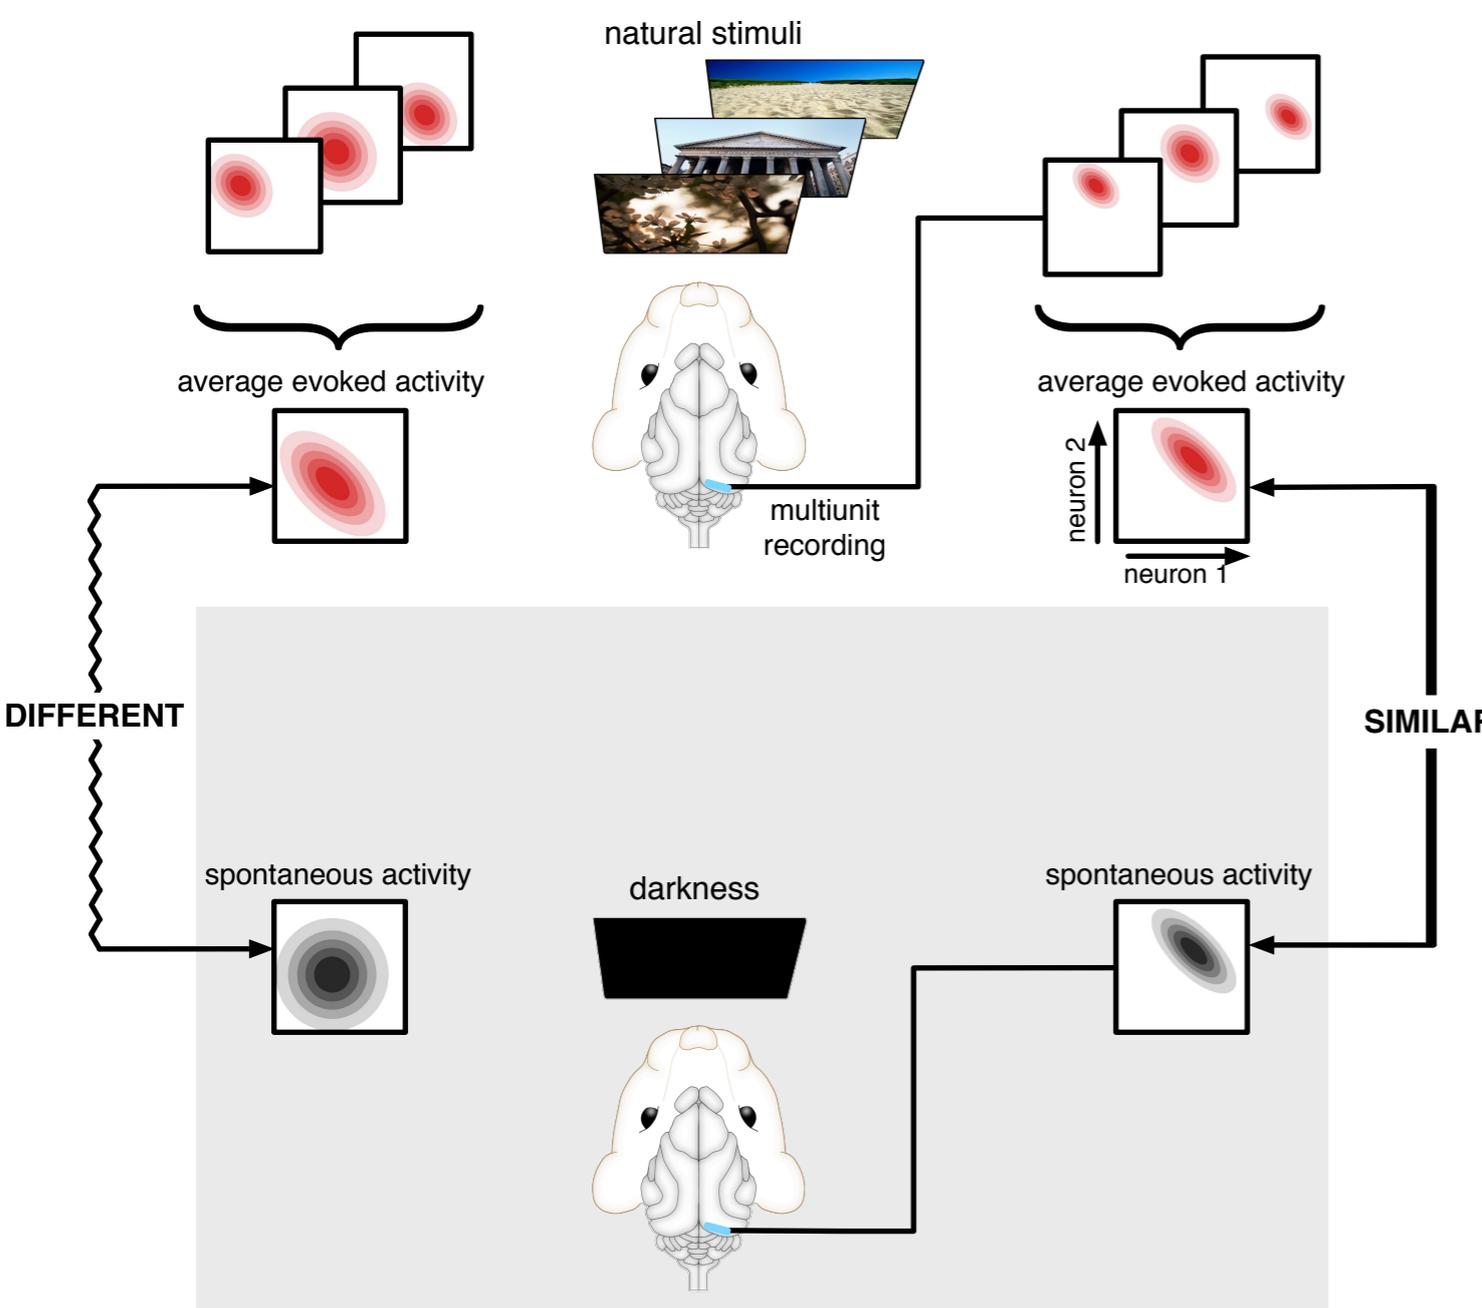
$$P_{\text{model}}(\text{stim}) \neq P_{\text{natural}}(\text{stim})$$

$$P_{\text{model}}(\text{stim}) = P_{\text{natural}}(\text{stim})$$

YOUNG

development

ADULT



$$P_{\text{model}}(\text{feat}|\text{stim}_3)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_2)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_1)$$

$$\int P_{\text{model}}(\text{feat}|\text{stim}) P_{\text{natural}}(\text{stim}) d\text{stim}$$

average inference
=
prior expectations

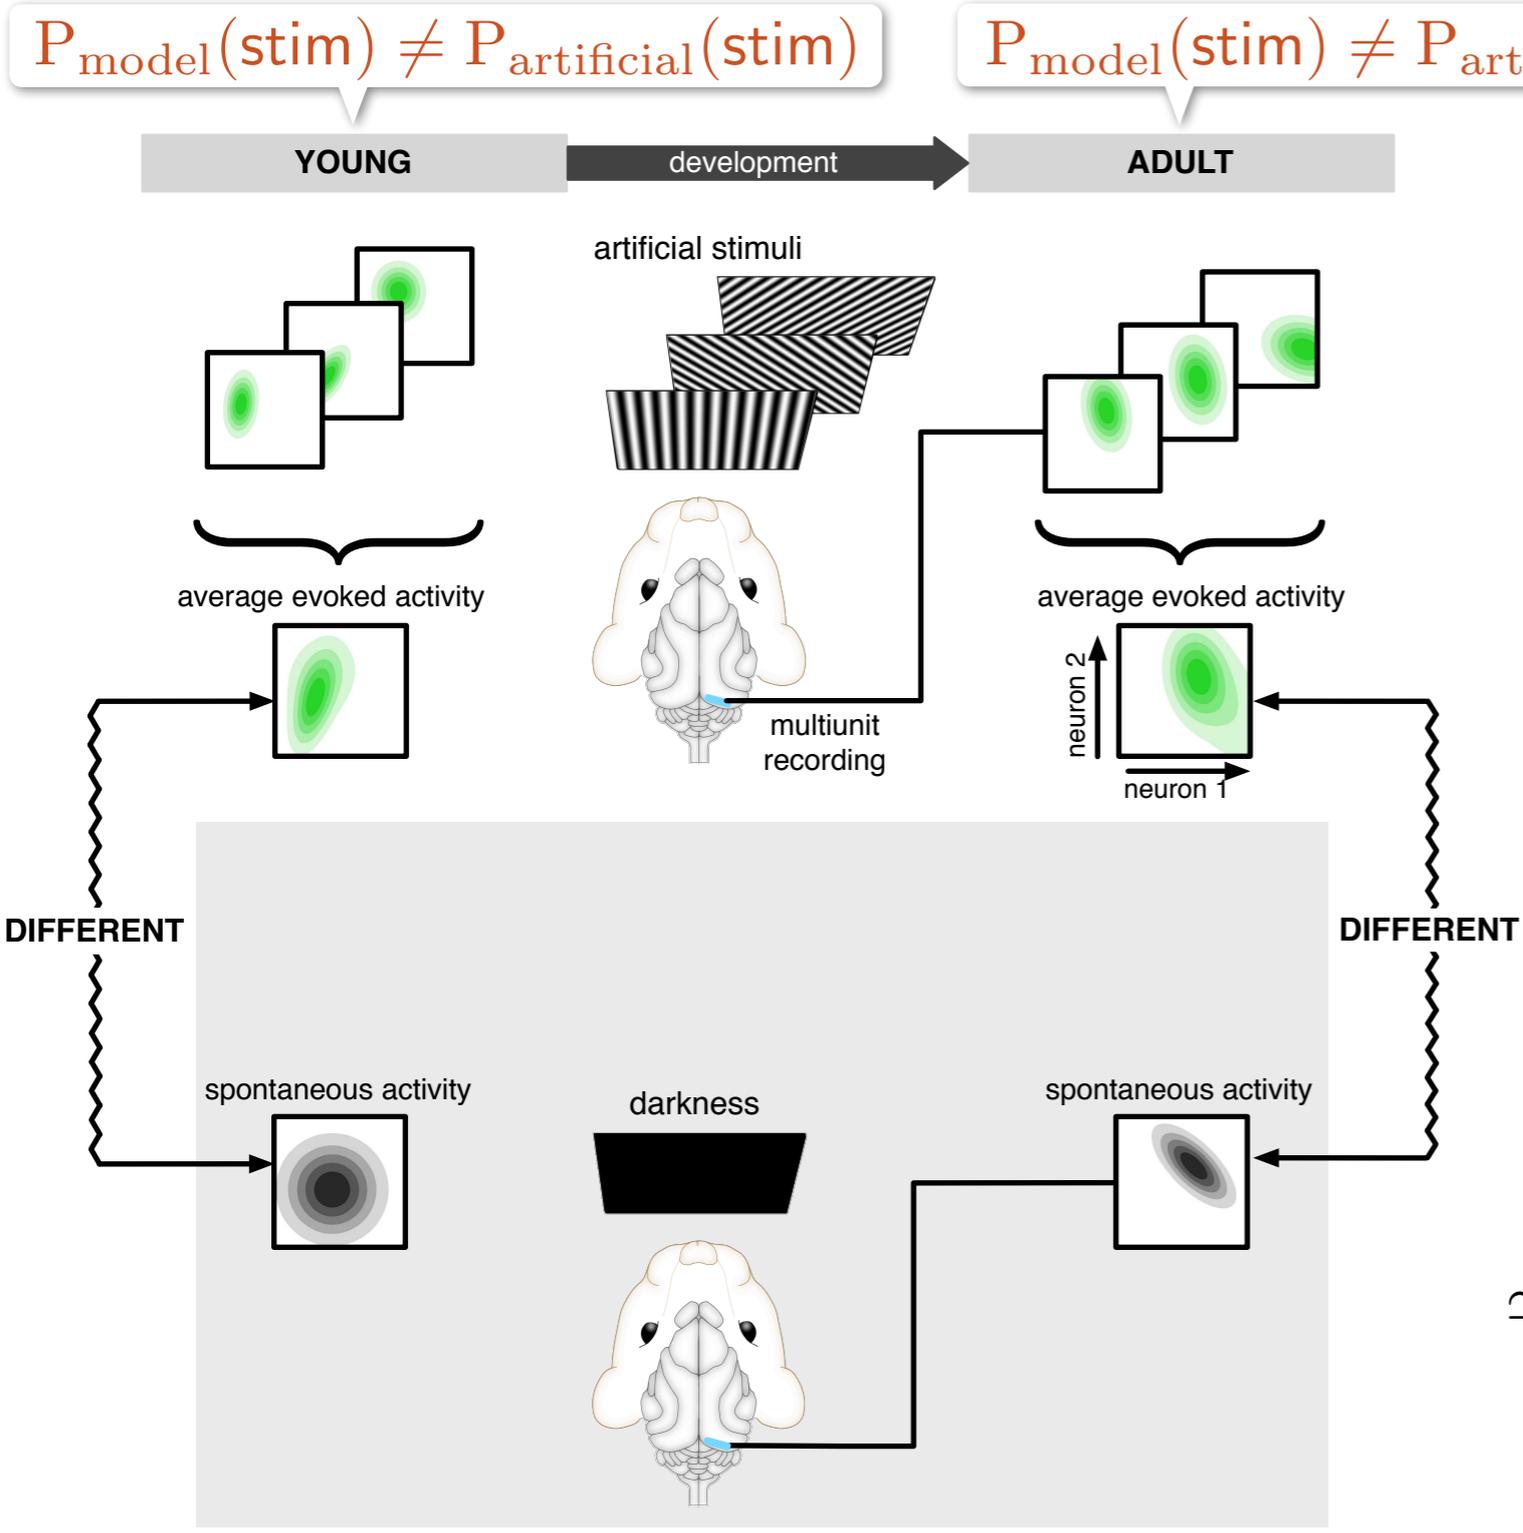
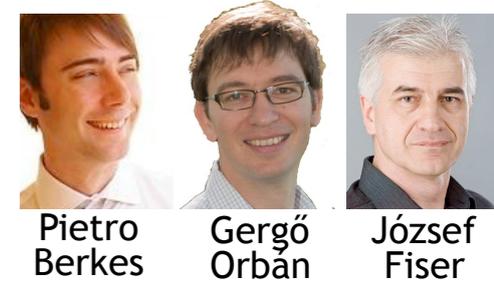
$$P_{\text{model}}(\text{feat}|\text{stim} = 0) \simeq P_{\text{model}}(\text{feat})$$

$$\simeq \int P_{\text{model}}(\text{feat}|\text{stim}) P_{\text{model}}(\text{stim}) d\text{stim}$$

Berkes et al, Science 2011

NEURAL HALLMARKS OF SAMPLING

+ STATISTICALLY OPTIMAL ADAPTATION



$$P_{\text{model}}(\text{stim}) \neq P_{\text{artificial}}(\text{stim})$$

$$P_{\text{model}}(\text{stim}) \neq P_{\text{artificial}}(\text{stim})$$

$$P_{\text{model}}(\text{feat}|\text{stim}_3)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_2)$$

$$P_{\text{model}}(\text{feat}|\text{stim}_1)$$

$$\int P_{\text{model}}(\text{feat}|\text{stim}) P_{\text{artificial}}(\text{stim}) d\text{stim}$$

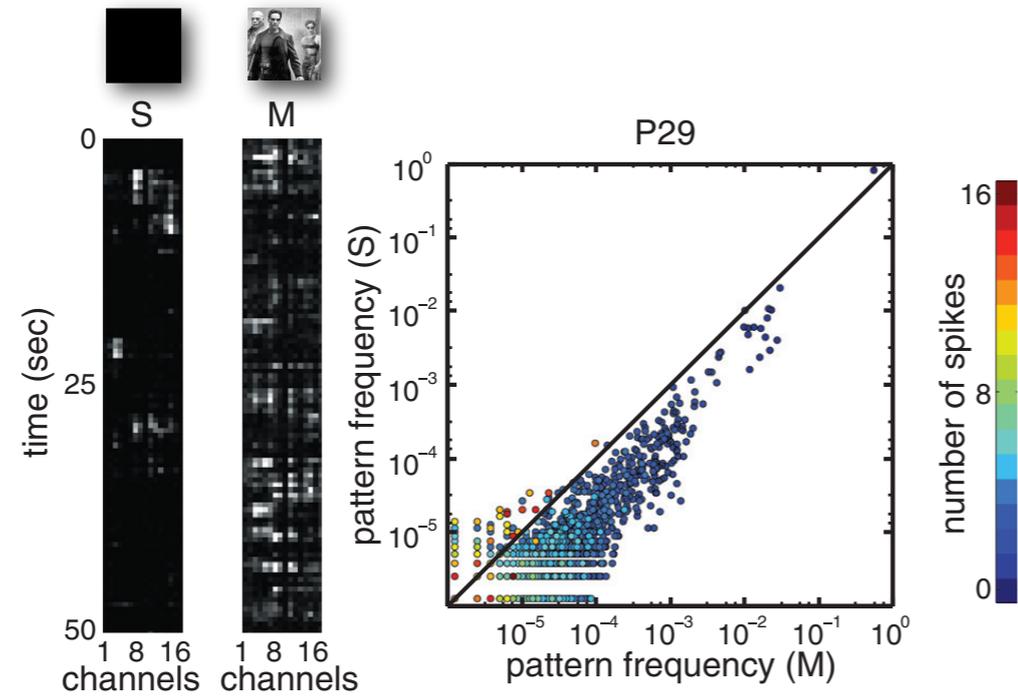
average inference
 \neq
prior expectations

$$P_{\text{model}}(\text{feat}|\text{stim} = 0) \simeq P_{\text{model}}(\text{feat})$$

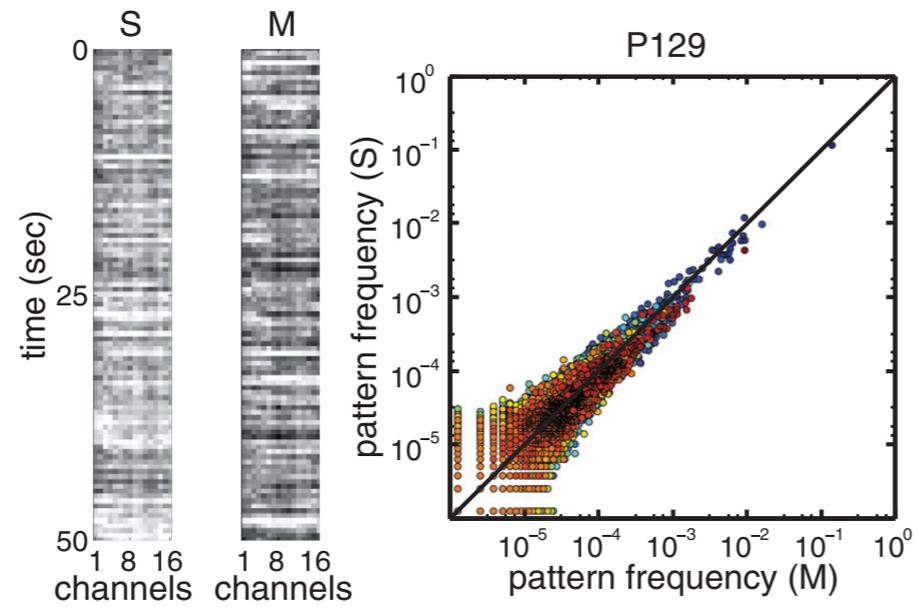
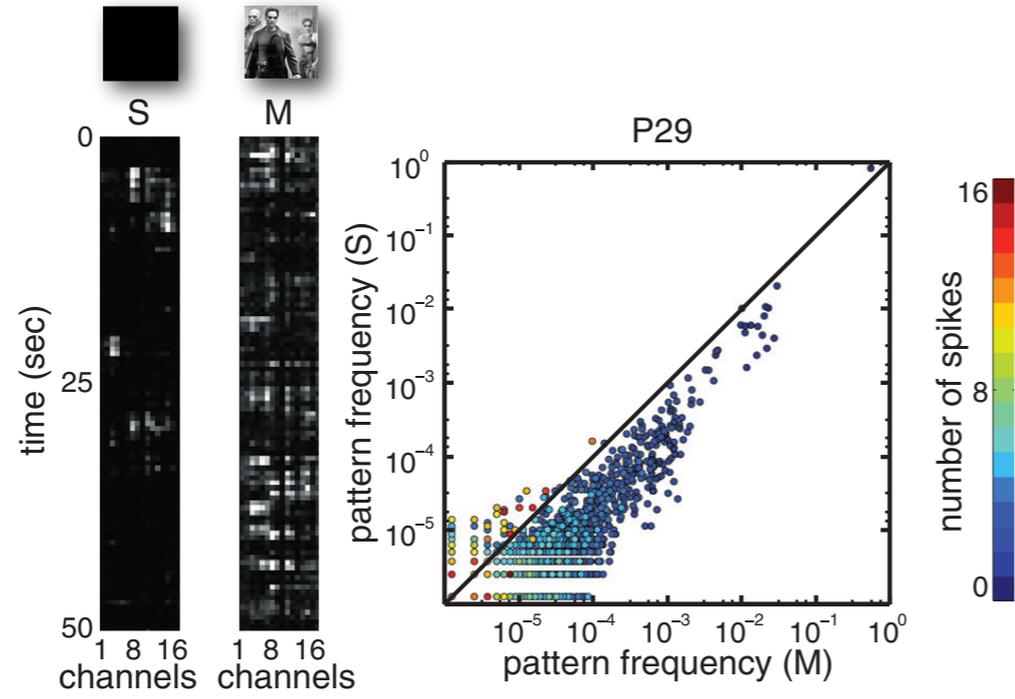
$$\simeq \int P_{\text{model}}(\text{feat}|\text{stim}) P_{\text{model}}(\text{stim}) d\text{stim}$$

Berkes et al, Science 2011

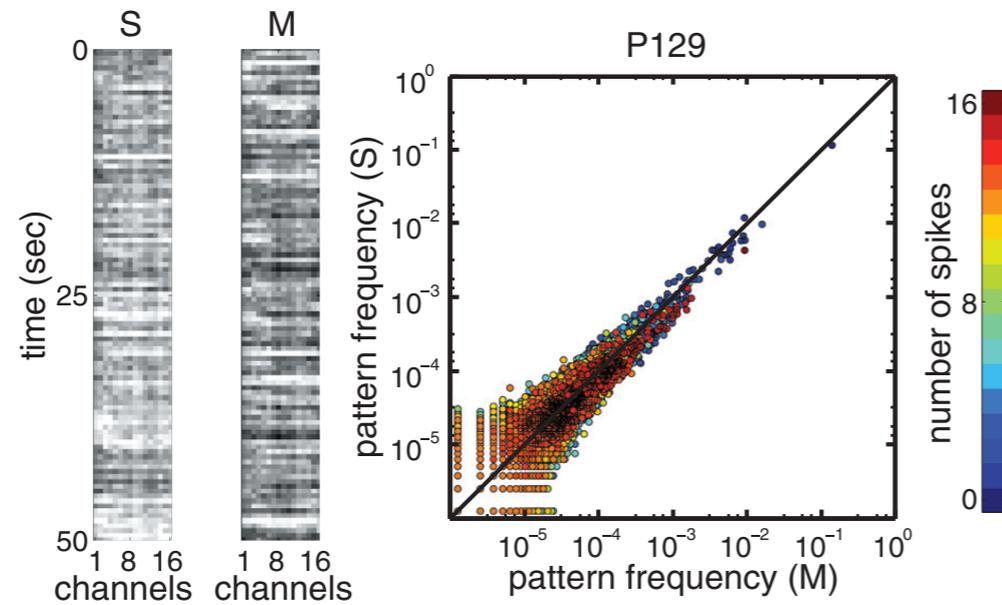
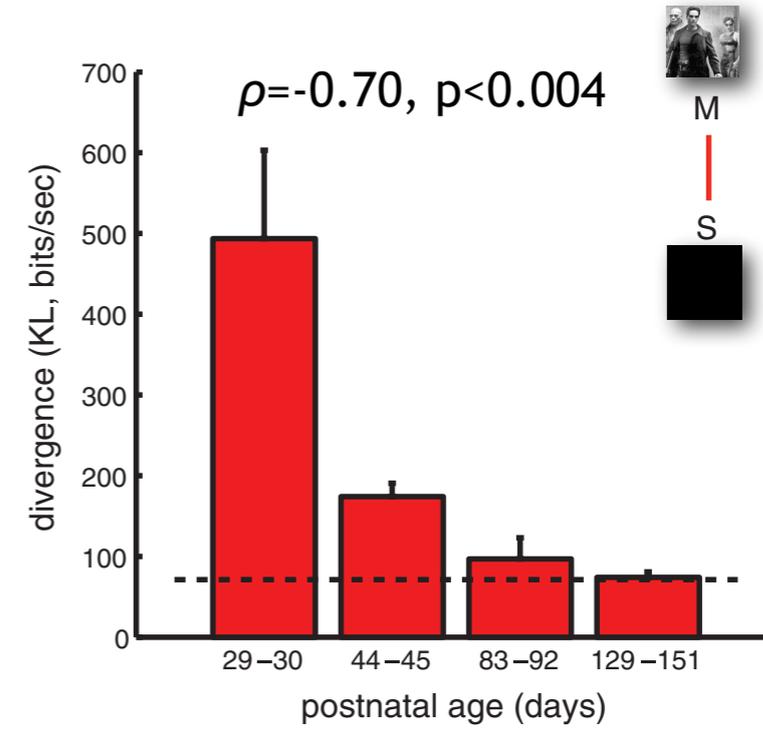
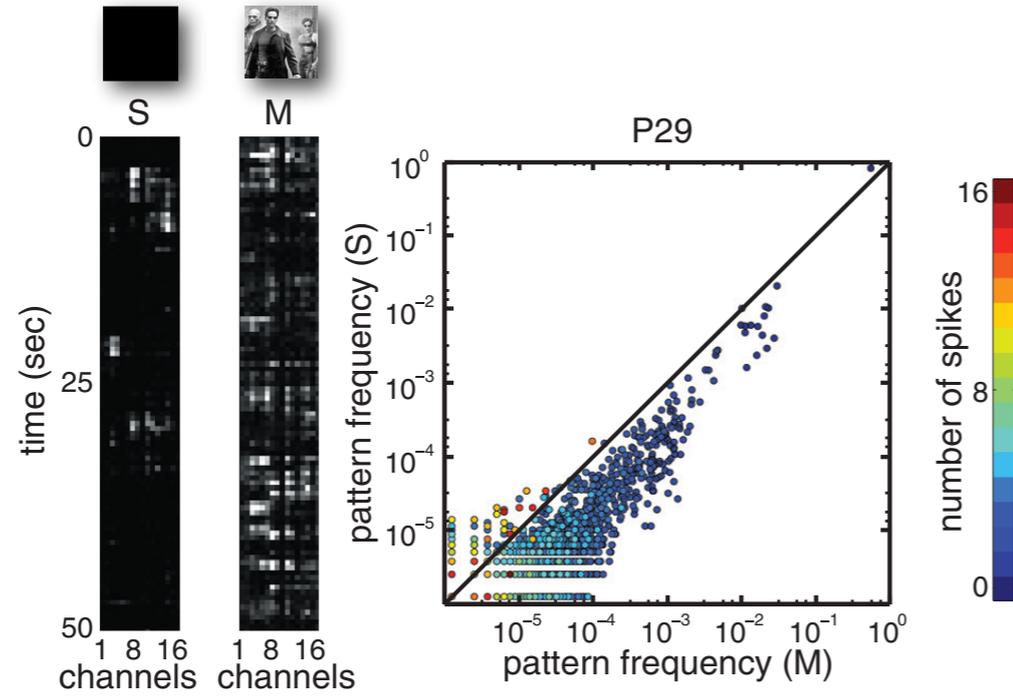
DEVELOPMENTAL CHANGES



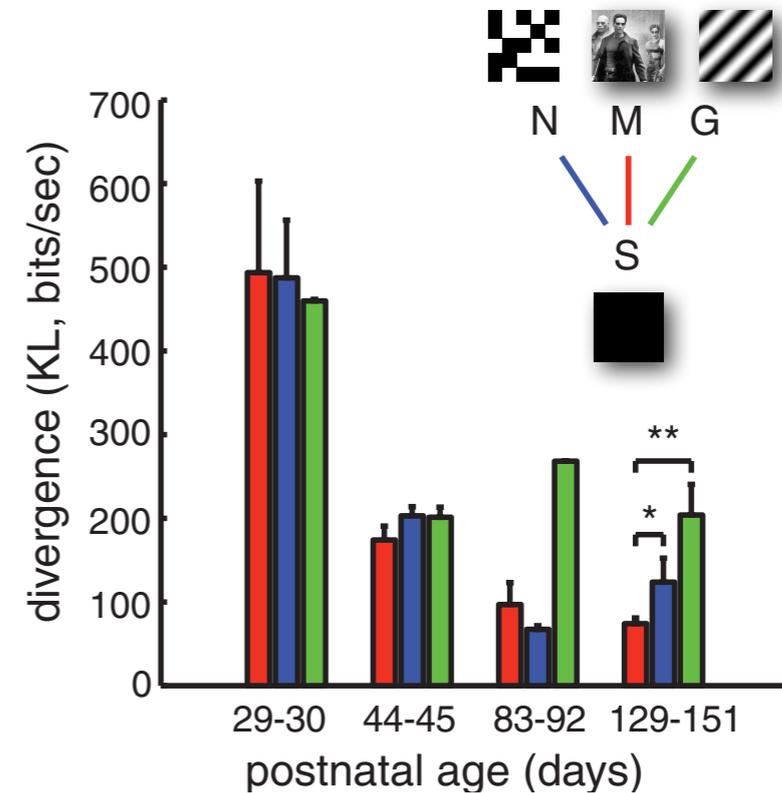
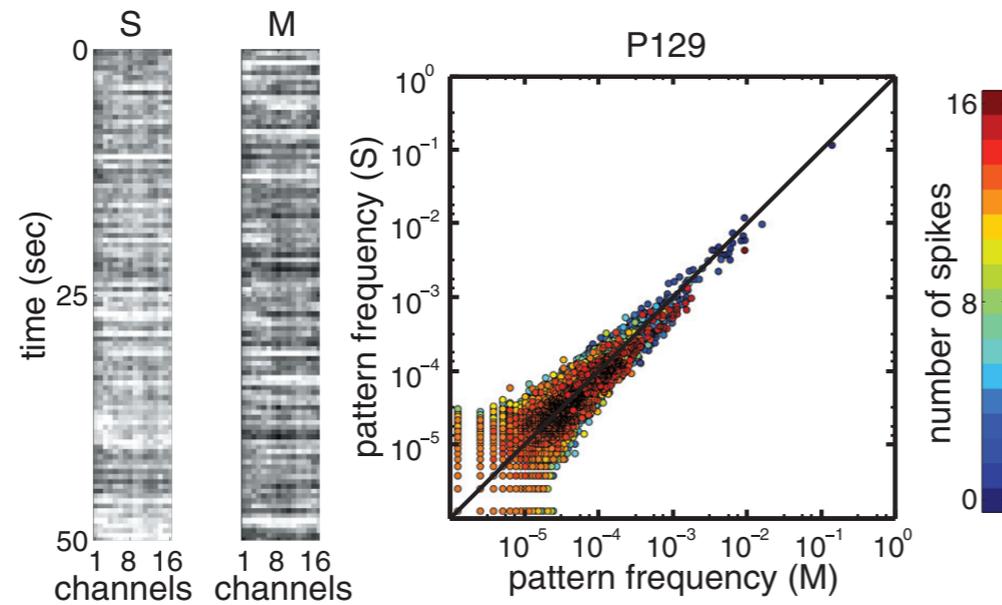
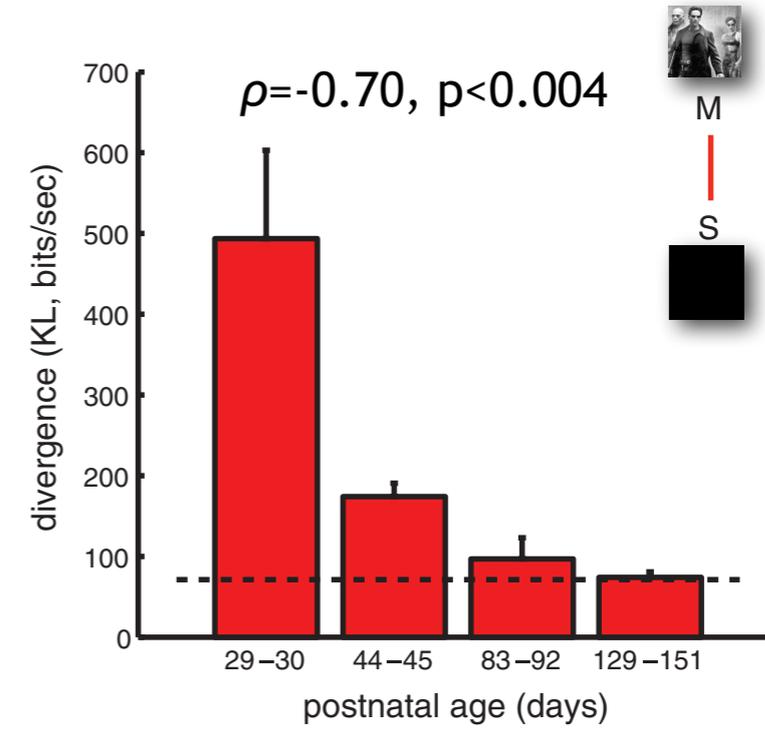
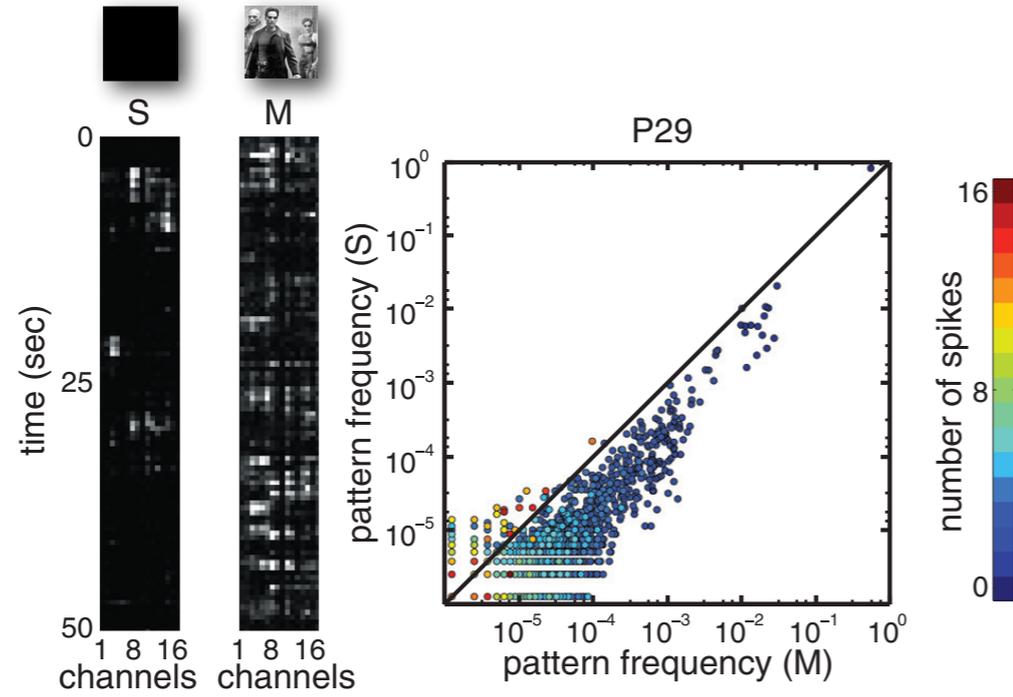
DEVELOPMENTAL CHANGES



DEVELOPMENTAL CHANGES

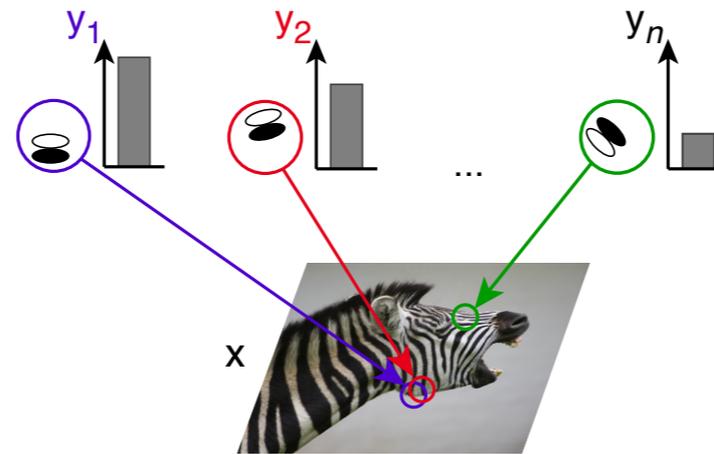


DEVELOPMENTAL CHANGES



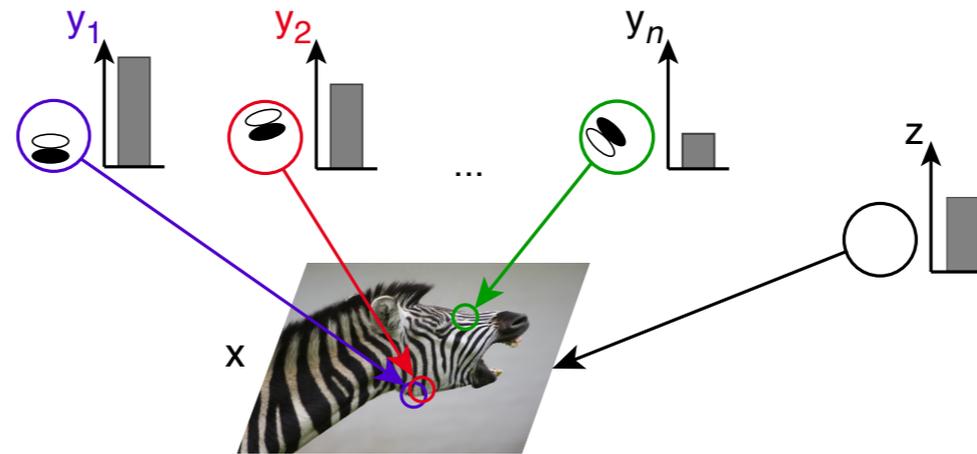
Berkes et al, Science 2011

AN IDEAL OBSERVER MODEL: GAUSSIAN SCALE MIXTURE



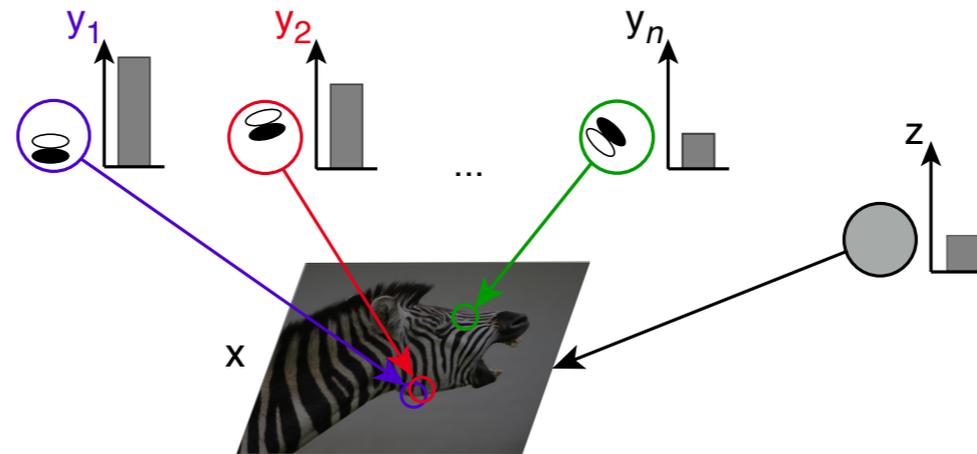
$$\mathbf{x} = y_1 \cdot \text{[purple ellipse]} + y_2 \cdot \text{[red ellipse]} + \dots + y_n \cdot \text{[green ellipse]}$$

AN IDEAL OBSERVER MODEL: GAUSSIAN SCALE MIXTURE



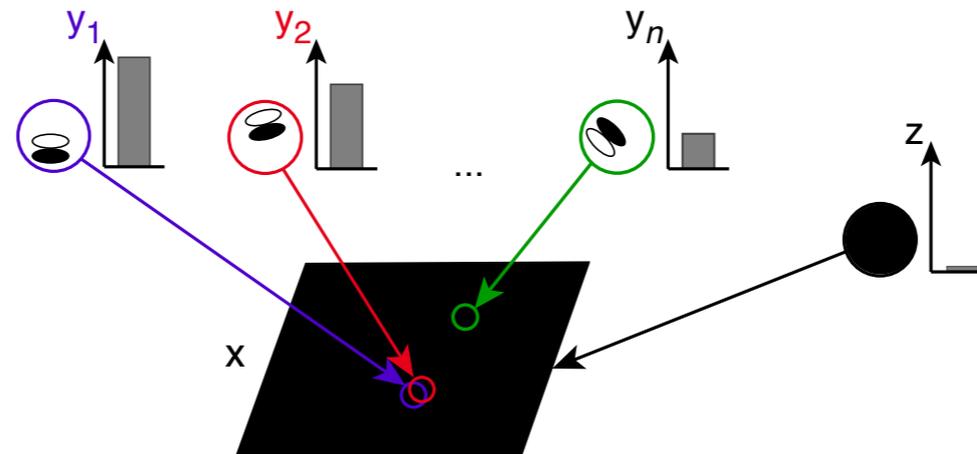
$$\mathbf{x} = (y_1 \cdot \text{zebra}_1 + y_2 \cdot \text{zebra}_2 + \dots + y_n \cdot \text{zebra}_n) \cdot z$$

AN IDEAL OBSERVER MODEL: GAUSSIAN SCALE MIXTURE



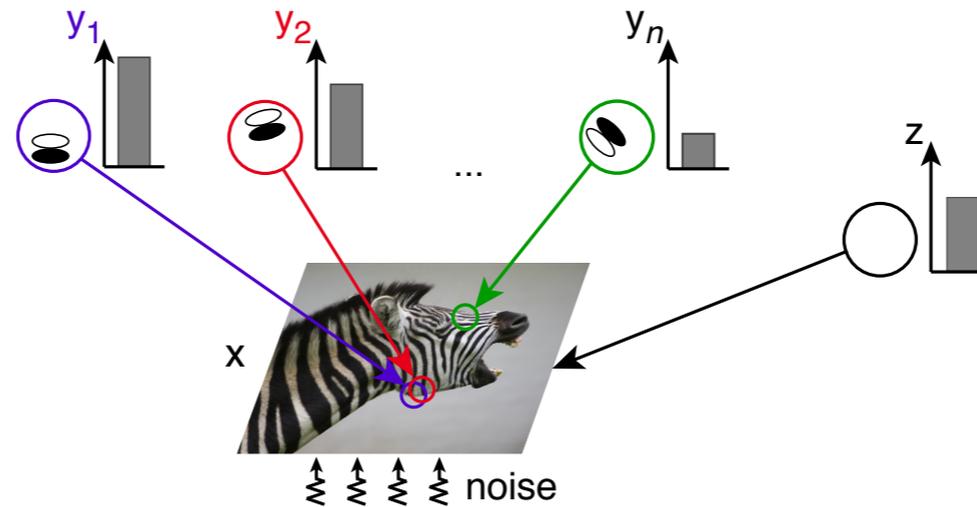
$$\mathbf{x} = (y_1 \cdot \text{zebra}_1 + y_2 \cdot \text{zebra}_2 + \dots + y_n \cdot \text{zebra}_n) \cdot z$$

AN IDEAL OBSERVER MODEL: GAUSSIAN SCALE MIXTURE



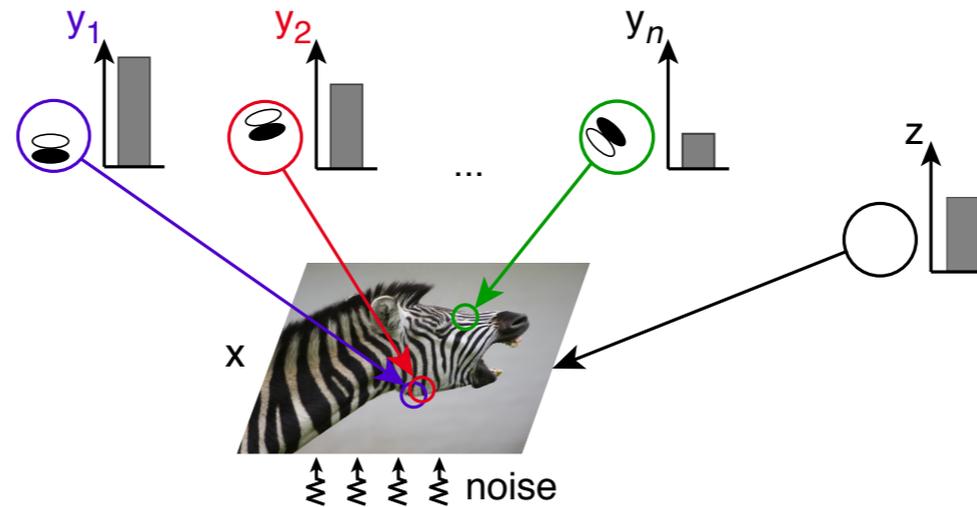
$$\mathbf{x} = (y_1 \cdot \text{bean}_1 + y_2 \cdot \text{bean}_2 + \dots + y_n \cdot \text{bean}_n) \cdot \mathbf{z}$$

AN IDEAL OBSERVER MODEL: GAUSSIAN SCALE MIXTURE



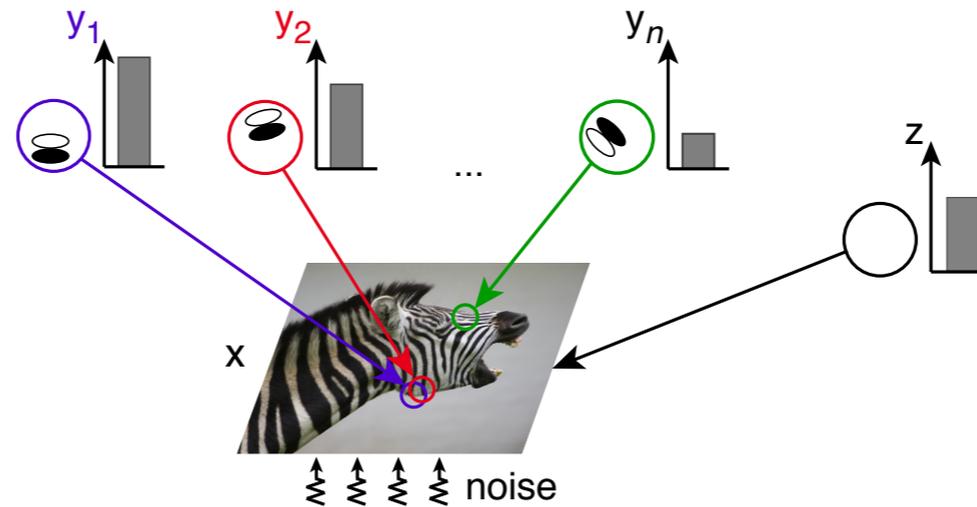
$$\mathbf{x} = (y_1 \cdot \text{node}_1 + y_2 \cdot \text{node}_2 + \dots + y_n \cdot \text{node}_n) \cdot z + \text{noise}$$

AN IDEAL OBSERVER MODEL: GAUSSIAN SCALE MIXTURE



$$\mathbf{x} = (y_1 \cdot \text{purple node} + y_2 \cdot \text{red node} + \dots + y_n \cdot \text{green node}) \cdot z + \text{noise}$$
$$\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}), z \sim \Gamma(\dots)$$

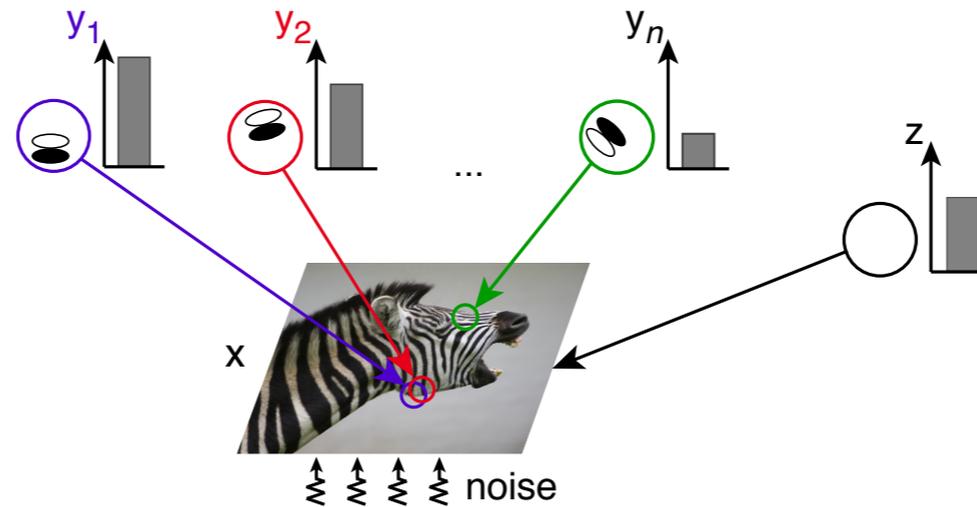
AN IDEAL OBSERVER MODEL: GAUSSIAN SCALE MIXTURE



$$\mathbf{x} = (y_1 \cdot \text{[purple circle]} + y_2 \cdot \text{[red circle]} + \dots + y_n \cdot \text{[green circle]}) \cdot z + \text{noise}$$
$$\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}), z \sim \Gamma(\dots)$$

Wainwright & Simoncelli 2000 → natural image statistics

AN IDEAL OBSERVER MODEL: GAUSSIAN SCALE MIXTURE



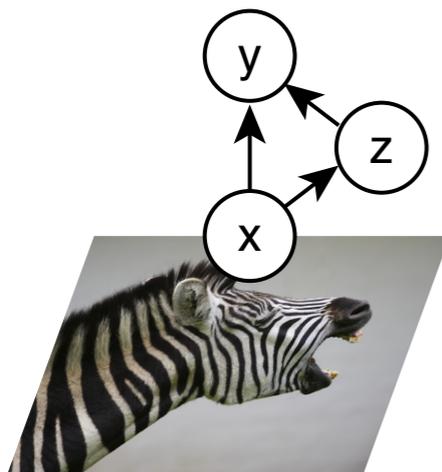
$$\mathbf{x} = (y_1 \cdot \text{[purple circle]} + y_2 \cdot \text{[red circle]} + \dots + y_n \cdot \text{[green circle]}) \cdot z + \text{noise}$$

$$\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}), z \sim \Gamma(\dots)$$

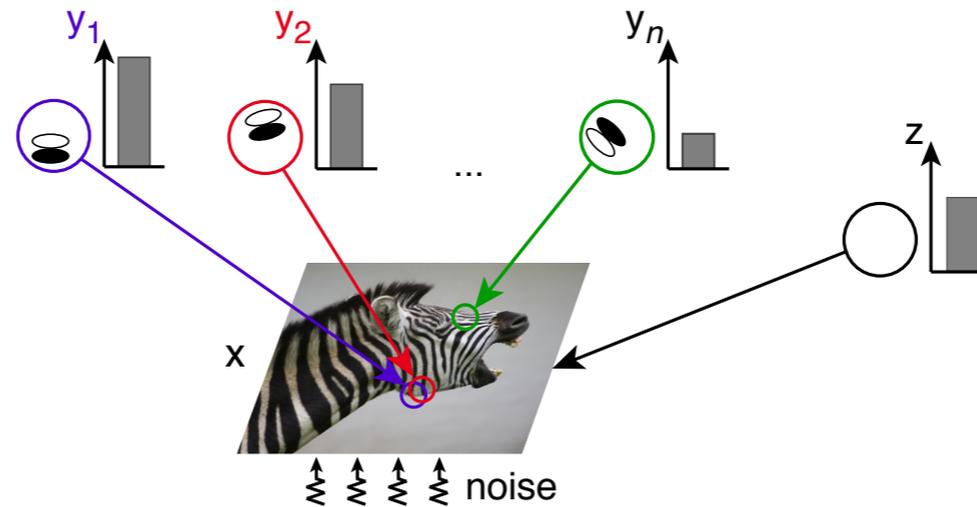
Wainwright & Simoncelli 2000 → natural image statistics

Schwartz & Simoncelli 2001
Coen-Cagli et al 2015 → trial-averaged neural responses

Schwartz & Dayan, J Vis 2009
Coen-Cagli et al, PLoS CB 2012 → psychophysics



AN IDEAL OBSERVER MODEL: GAUSSIAN SCALE MIXTURE



$$\mathbf{x} = (y_1 \cdot \text{[purple blob]} + y_2 \cdot \text{[red blob]} + \dots + y_n \cdot \text{[green blob]}) \cdot z + \text{noise}$$

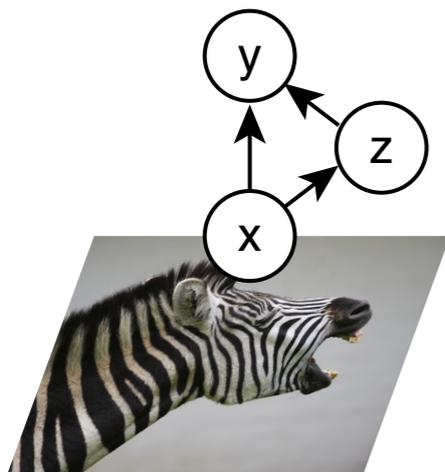
$$\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}), z \sim \Gamma(\dots)$$

Wainwright & Simoncelli 2000 → natural image statistics

Schwartz & Simoncelli 2001
Coen-Cagli et al 2015 → trial-averaged neural responses

Schwartz & Dayan, J Vis 2009
Coen-Cagli et al, PLoS CB 2012 → psychophysics

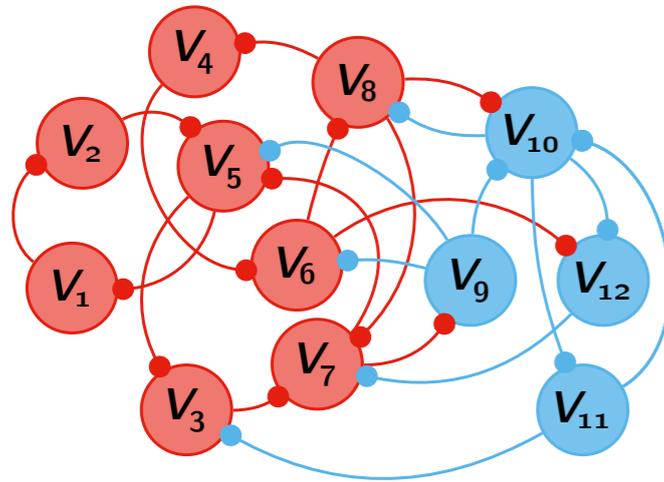
Orbán et al, Neuron 2016 → neural variability



THE STABILIZED SUPRALINEAR NETWORK



recurrently coupled E/I network



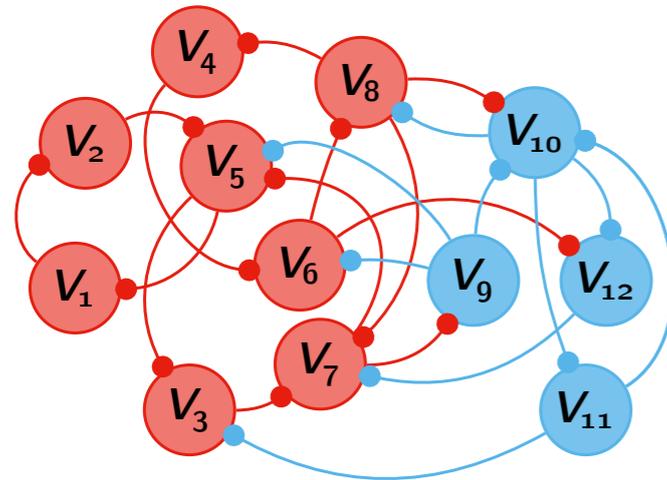
$$\tau_i \frac{dV_i}{dt} = -V_i + V_{\text{rest}} + \sum_j W_{ij} r_j(t) + h_i(t)$$

Ahmadian et al, *Neural Comput* 2013
Rubin et al, *Neuron* 2015

THE STABILIZED SUPRALINEAR NETWORK



recurrently coupled E/I network



low-pass filtered
membrane potential

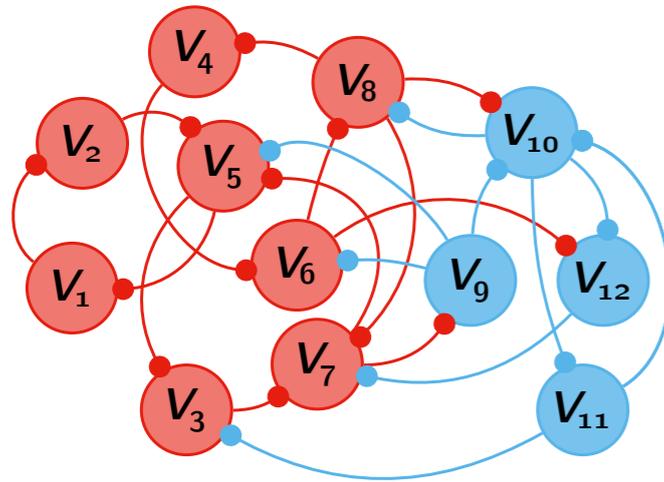
$$\tau_i \frac{dV_i}{dt} = -V_i + V_{\text{rest}} + \sum_j W_{ij} r_j(t) + h_i(t)$$

Ahmadian et al, *Neural Comput* 2013
Rubin et al, *Neuron* 2015

THE STABILIZED SUPRALINEAR NETWORK



recurrently coupled E/I network



low-pass filtered
membrane potential

$$\tau_i \frac{dV_i}{dt} = -V_i + V_{\text{rest}} + \sum_j W_{ij} r_j(t) + h_i(t)$$

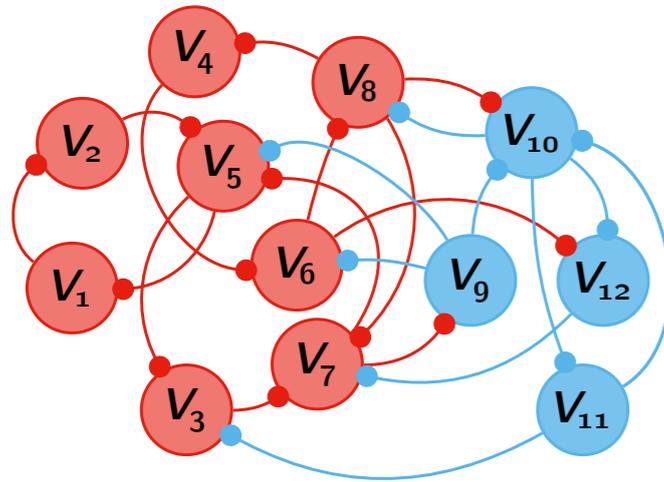
membrane
time constant

Ahmadian et al, *Neural Comput* 2013
Rubin et al, *Neuron* 2015

THE STABILIZED SUPRALINEAR NETWORK



recurrently coupled E/I network



low-pass filtered membrane potential

$$\tau_i \frac{dV_i}{dt} = -V_i + V_{\text{rest}} + \sum_j W_{ij} r_j(t) + h_i(t)$$

membrane time constant

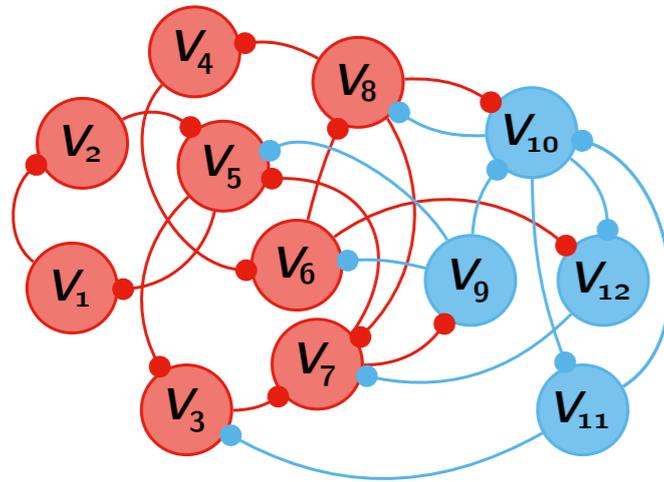
resting potential

Ahmadian et al, *Neural Comput* 2013
Rubin et al, *Neuron* 2015

THE STABILIZED SUPRALINEAR NETWORK



recurrently coupled E/I network



low-pass filtered membrane potential

$$\tau_i \frac{dV_i}{dt} = -V_i + V_{\text{rest}} + \sum_j W_{ij} r_j(t) + h_i(t)$$

membrane time constant

resting potential

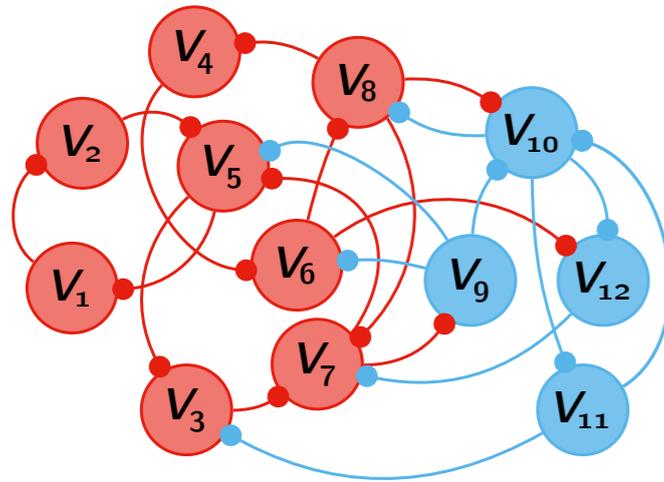
synaptic weights

Ahmadian et al, *Neural Comput* 2013
Rubin et al, *Neuron* 2015

THE STABILIZED SUPRALINEAR NETWORK



recurrently coupled E/I network



low-pass filtered membrane potential

$$\tau_i \frac{dV_i}{dt} = -V_i + V_{\text{rest}} + \sum_j W_{ij} r_j(t) + h_i(t)$$

firing rates

membrane time constant

resting potential

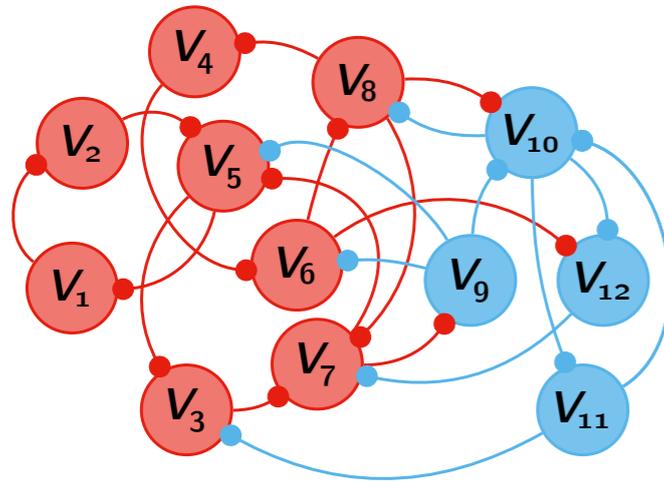
synaptic weights

Ahmadian et al, *Neural Comput* 2013
Rubin et al, *Neuron* 2015

THE STABILIZED SUPRALINEAR NETWORK



recurrently coupled E/I network



low-pass filtered membrane potential

$$\tau_i \frac{dV_i}{dt} = -V_i + V_{\text{rest}} + \sum_j W_{ij} r_j(t) + h_i(t)$$

firing rates

external input

membrane time constant

resting potential

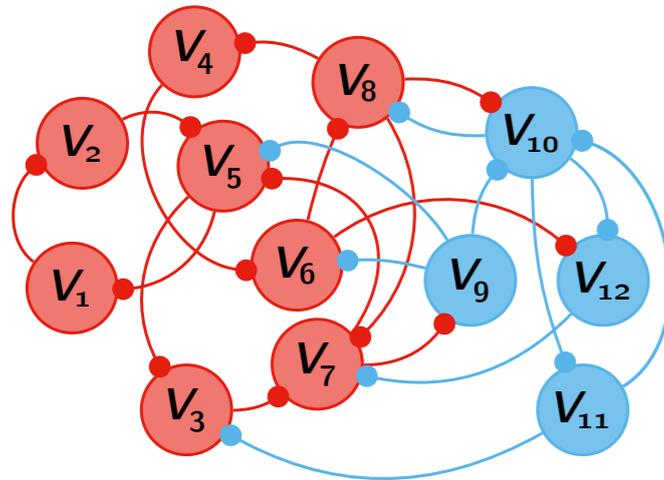
synaptic weights

Ahmadian et al, *Neural Comput* 2013
Rubin et al, *Neuron* 2015

THE STABILIZED SUPRALINEAR NETWORK



recurrently coupled E/I network



low-pass filtered membrane potential

$$\tau_i \frac{dV_i}{dt} = -V_i + V_{\text{rest}} + \sum_j W_{ij} r_j(t) + h_i(t)$$

firing rates

external input

membrane time constant

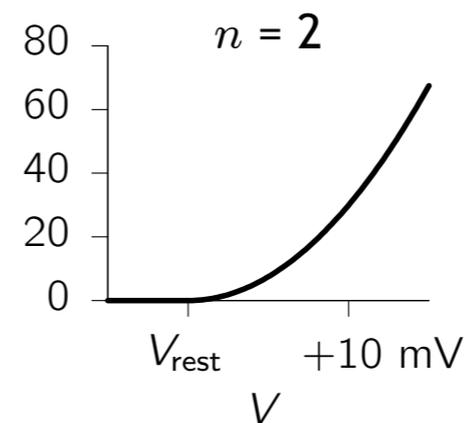
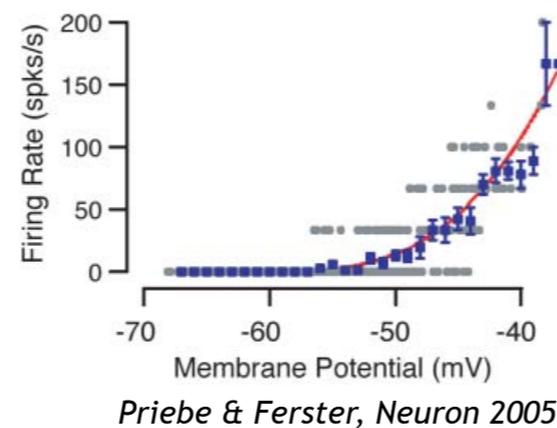
resting potential

synaptic weights

Ahmadian et al, *Neural Comput* 2013
Rubin et al, *Neuron* 2015

with expansive firing rate nonlinearities

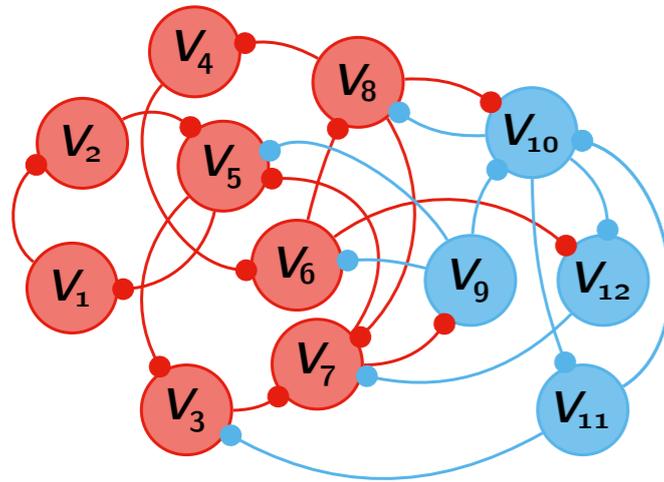
$$r_i(t) = k [V_i(t) - V_{\text{rest}}]_+^n$$



THE STOCHASTIC STABILIZED SUPRALINEAR NETWORK



recurrently coupled E/I network



low-pass filtered membrane potential

$$\tau_i \frac{dV_i}{dt} = -V_i + V_{\text{rest}} + \sum_j W_{ij} r_j(t) + h_i(t) + \text{noise}$$

firing rates

external input

membrane time constant

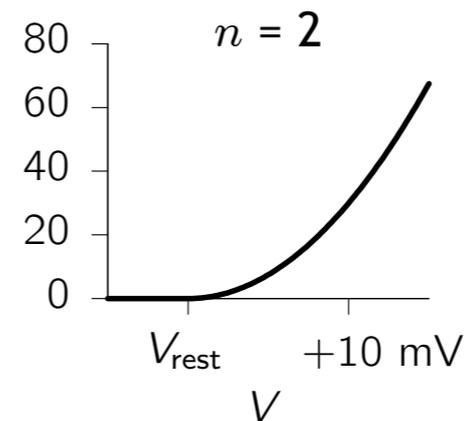
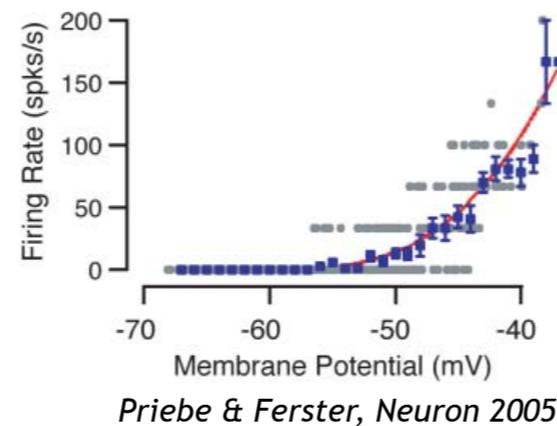
resting potential

synaptic weights

Ahmadian et al, *Neural Comput* 2013
 Rubin et al, *Neuron* 2015
 Hennequin et al, *accepted*

with expansive firing rate nonlinearities

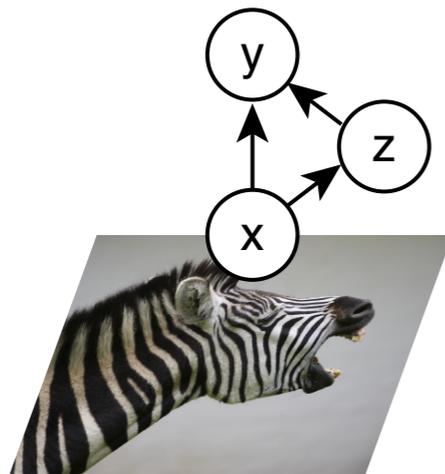
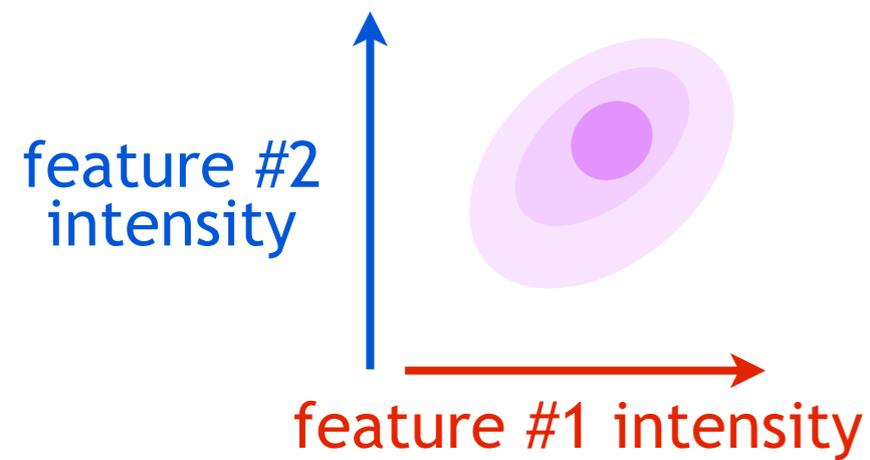
$$r_i(t) = k [V_i(t) - V_{\text{rest}}]_+^n$$



BRINGING THE TWO WORLDS TOGETHER

ideal observer: GSM

$$P(y_1, y_2 \mid x)$$



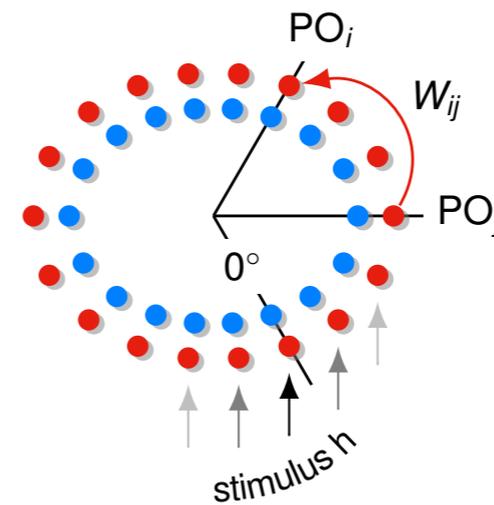
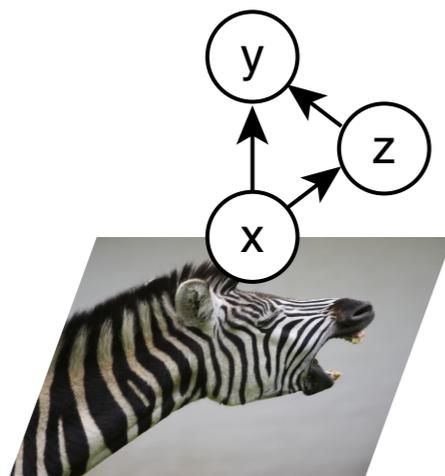
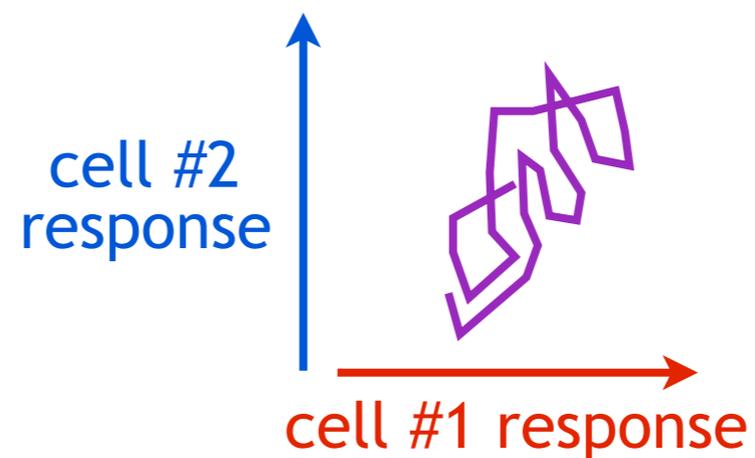
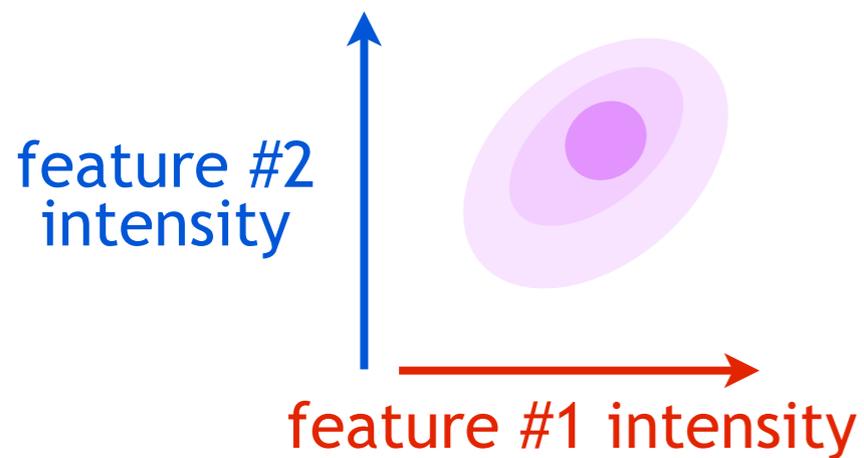
BRINGING THE TWO WORLDS TOGETHER

ideal observer: GSM

visual cortex: neural network

$$P(y_1, y_2 \mid x)$$

$$P(r_1, r_2 \mid x; \text{network parameters})$$



BRINGING THE TWO WORLDS TOGETHER

ideal observer: GSM

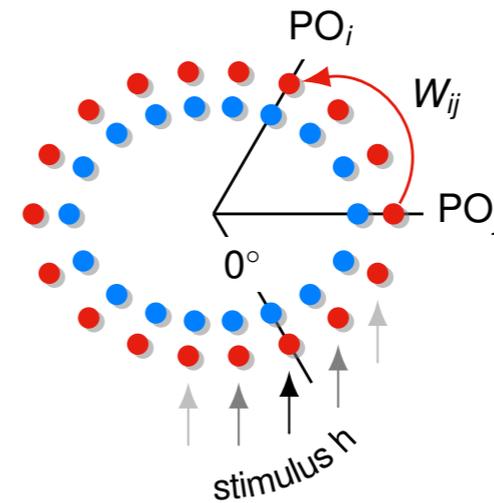
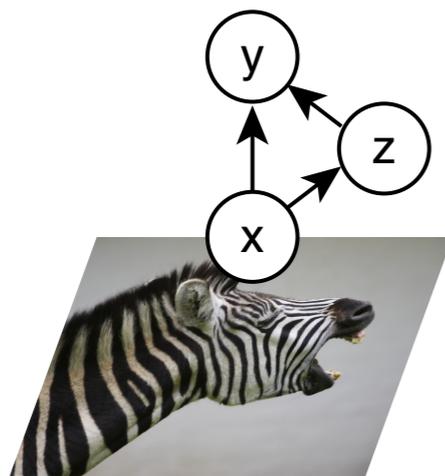
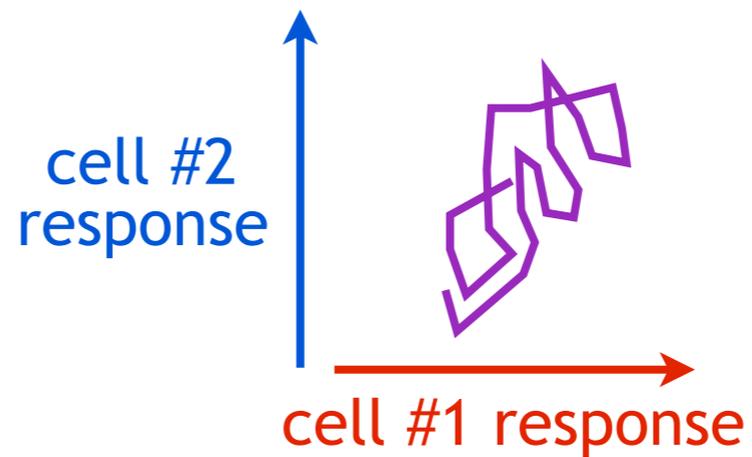
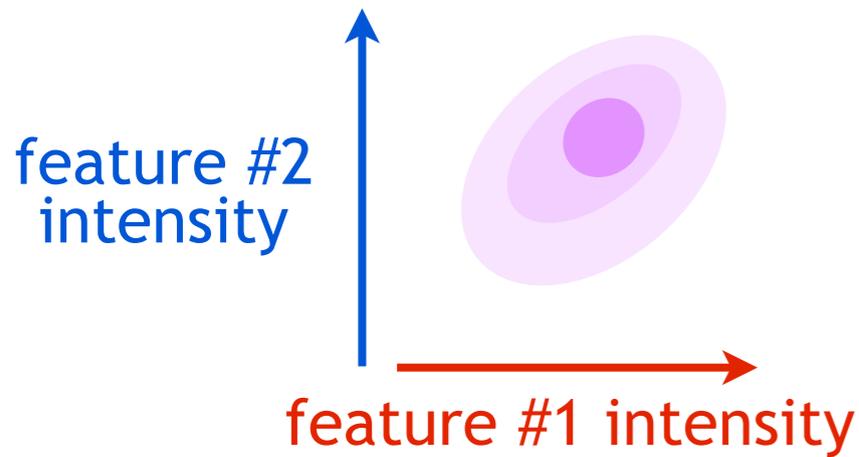
visual cortex: neural network



$$P(y_1, y_2 \mid x)$$

train

$$P(r_1, r_2 \mid x; \text{network parameters})$$



BRINGING THE TWO WORLDS TOGETHER

ideal observer: GSM

- ▶ full distribution (2nd order)

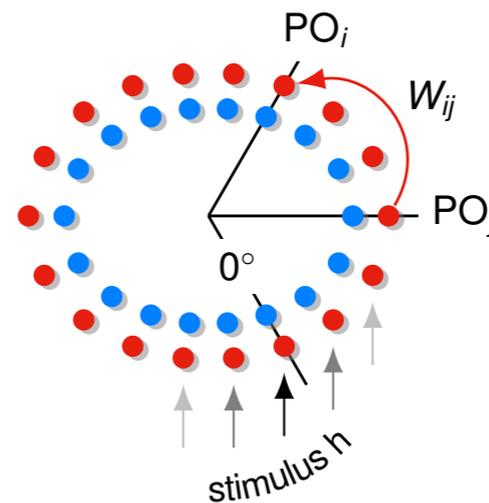
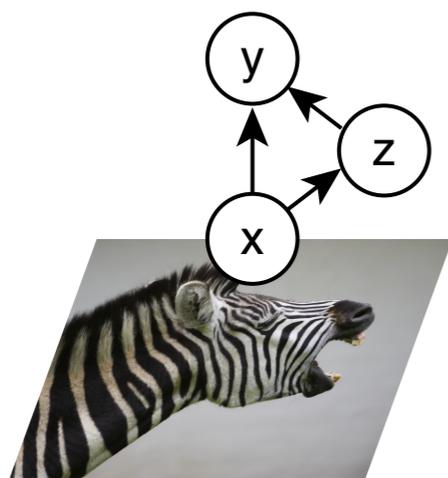
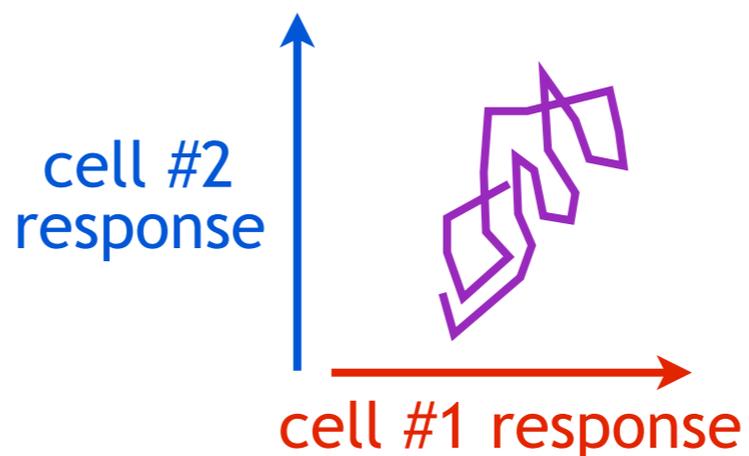
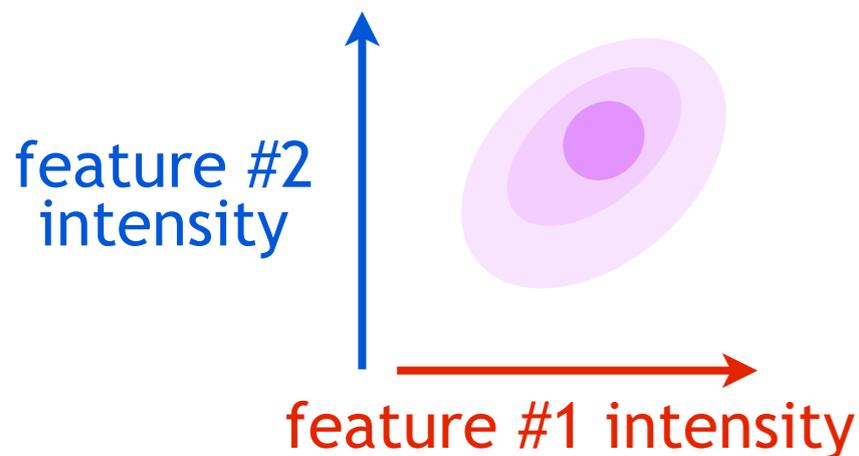
$$P(y_1, y_2 \mid x)$$



visual cortex: neural network

- ▶ recurrent
- ▶ E-I
- ▶ expansive nonlinearity

$$P(r_1, r_2 \mid x; \text{network parameters})$$



BRINGING THE TWO WORLDS TOGETHER

ideal observer: GSM

- ▶ full distribution (2nd order)

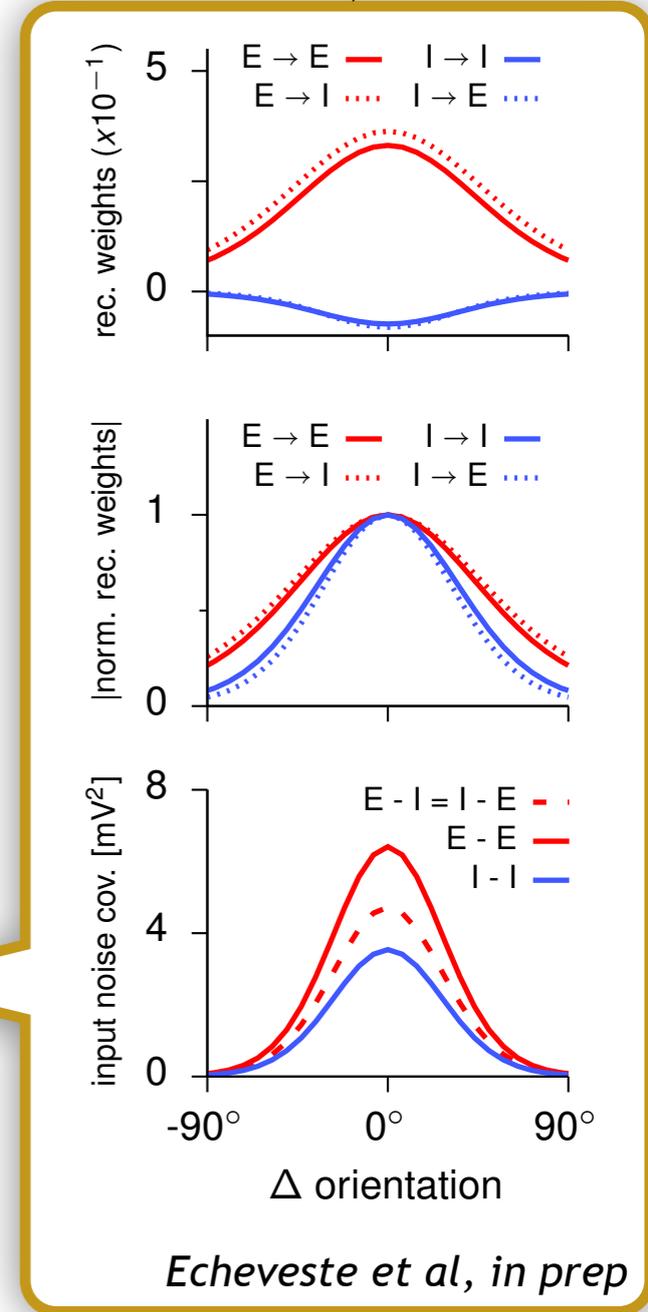
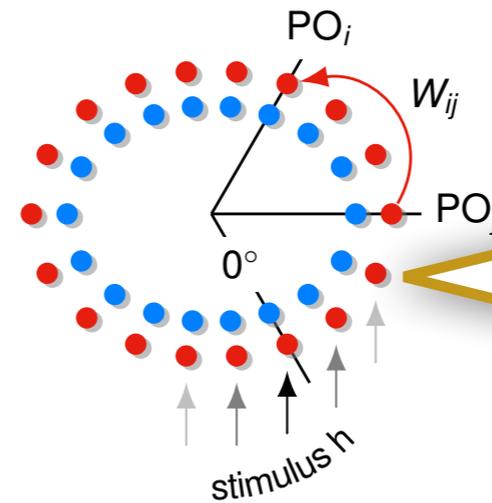
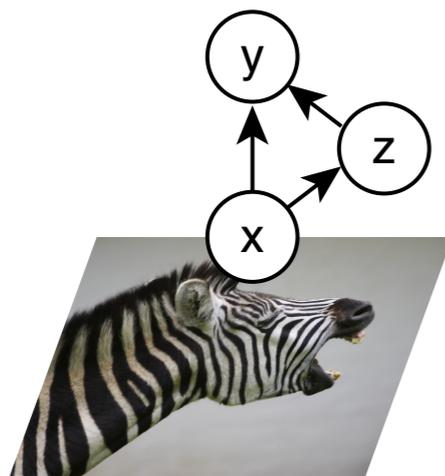
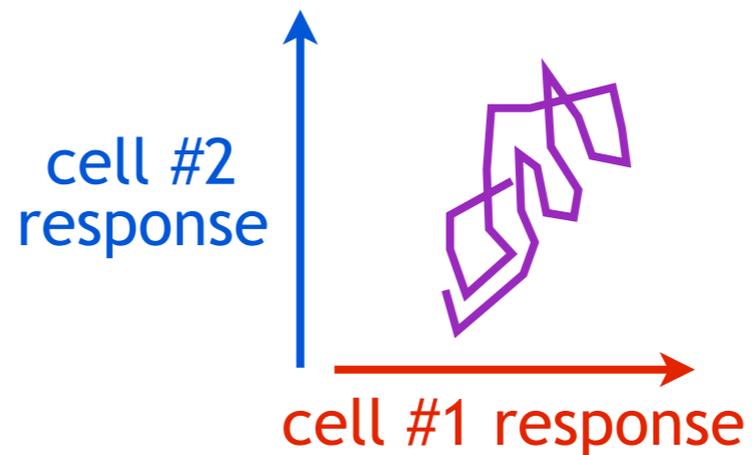
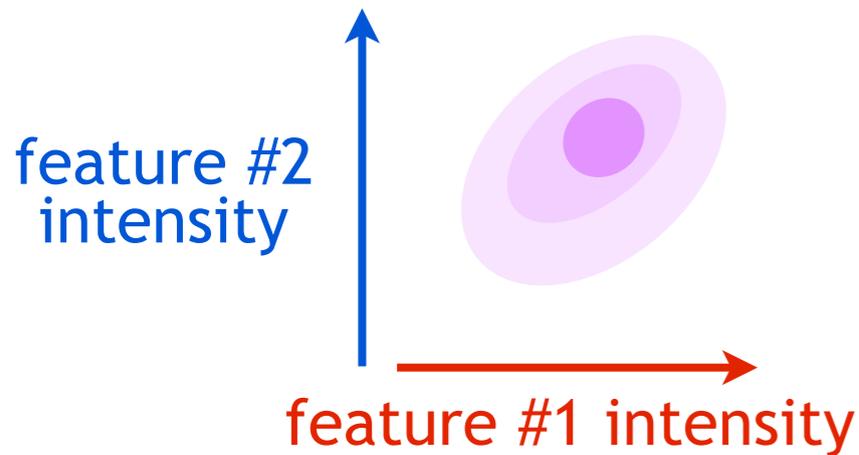
$$P(y_1, y_2 \mid x)$$



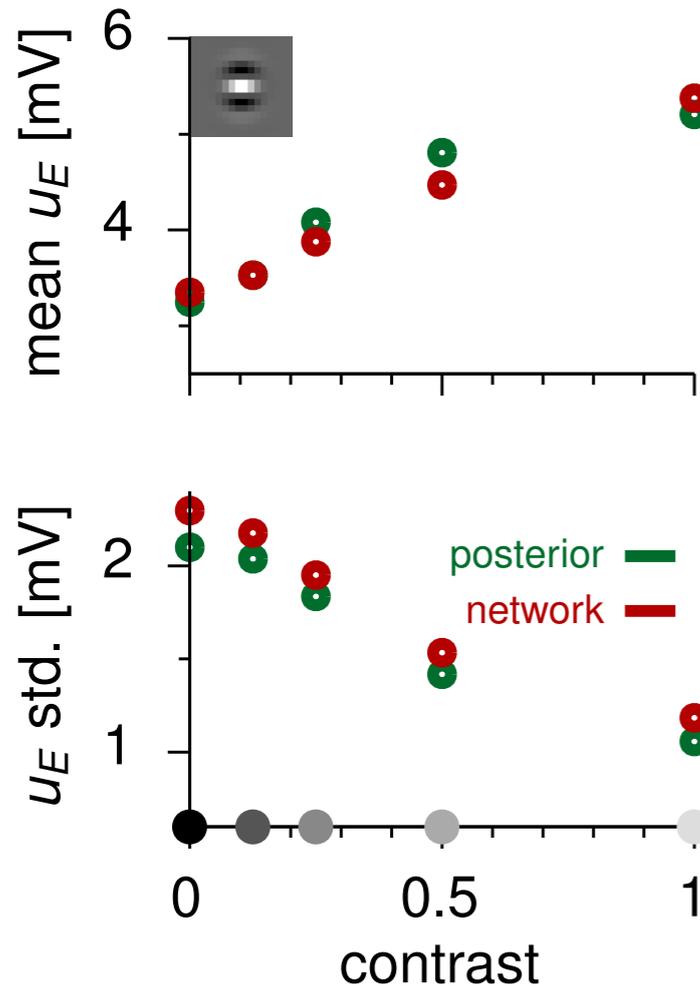
visual cortex: neural network

- ▶ recurrent
- ▶ E-I
- ▶ expansive nonlinearity

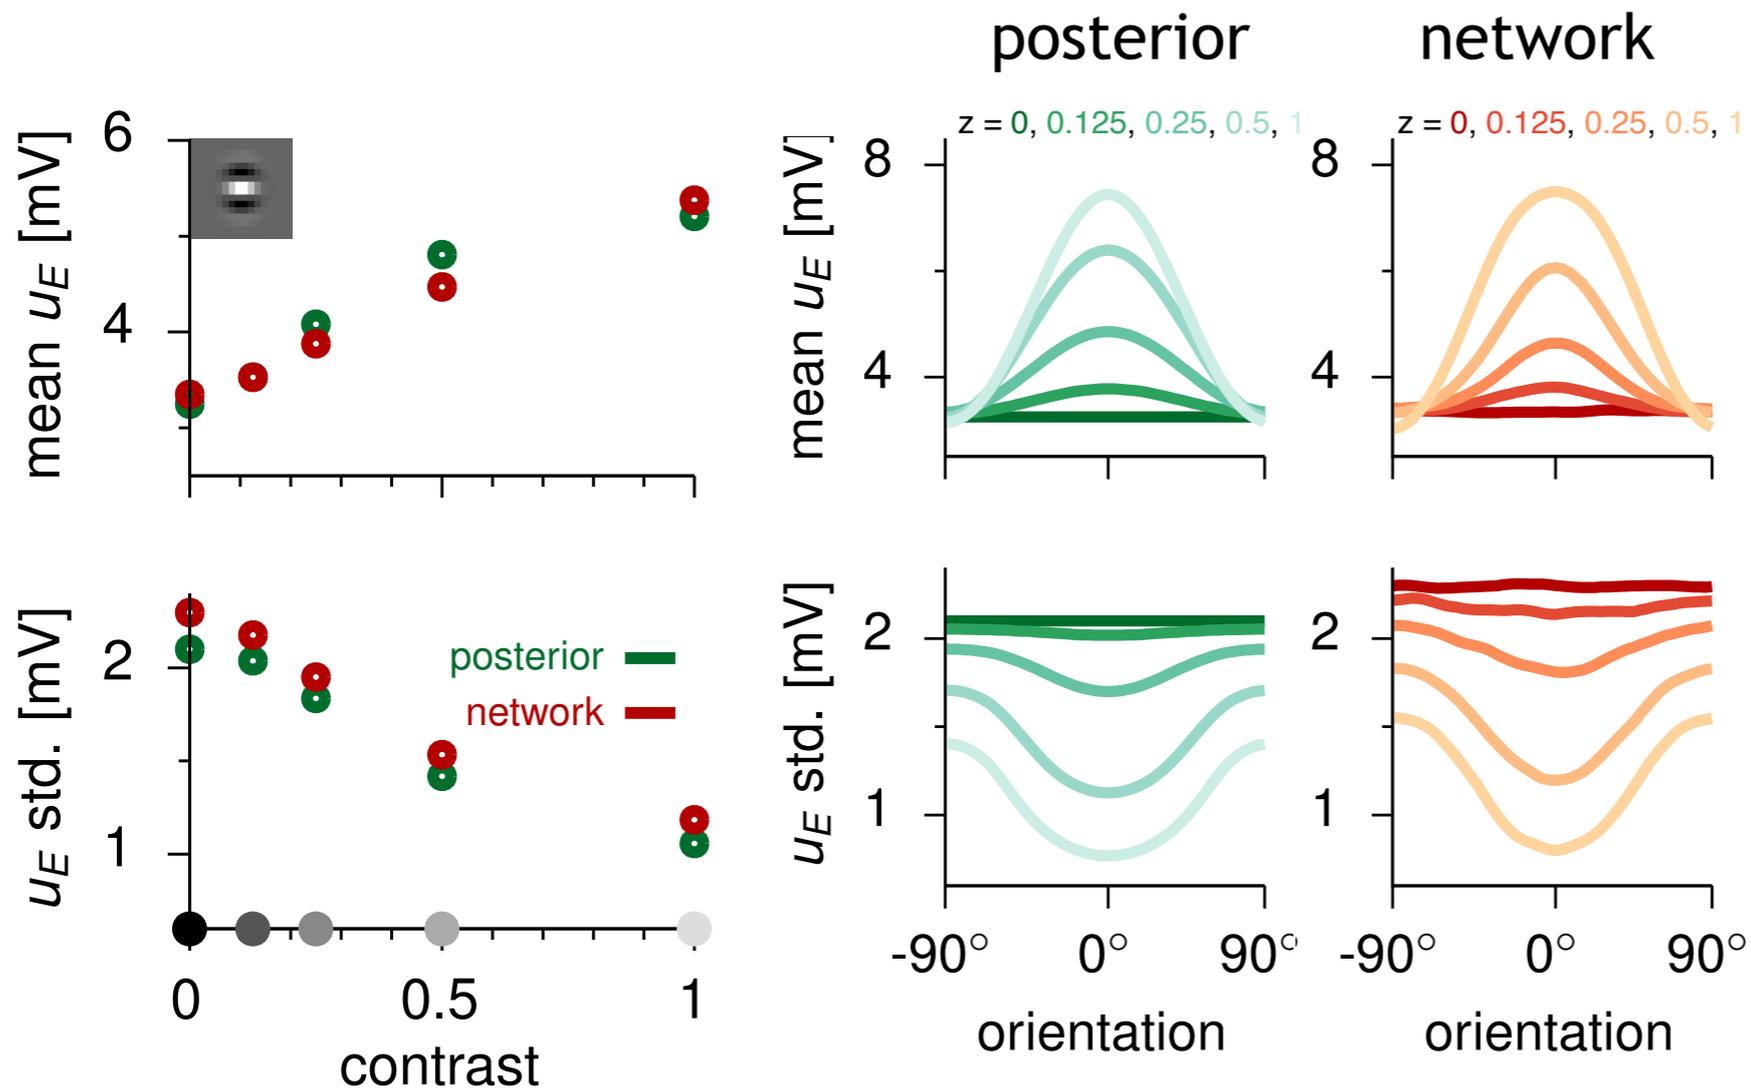
$$P(r_1, r_2 \mid x; \text{network parameters})$$



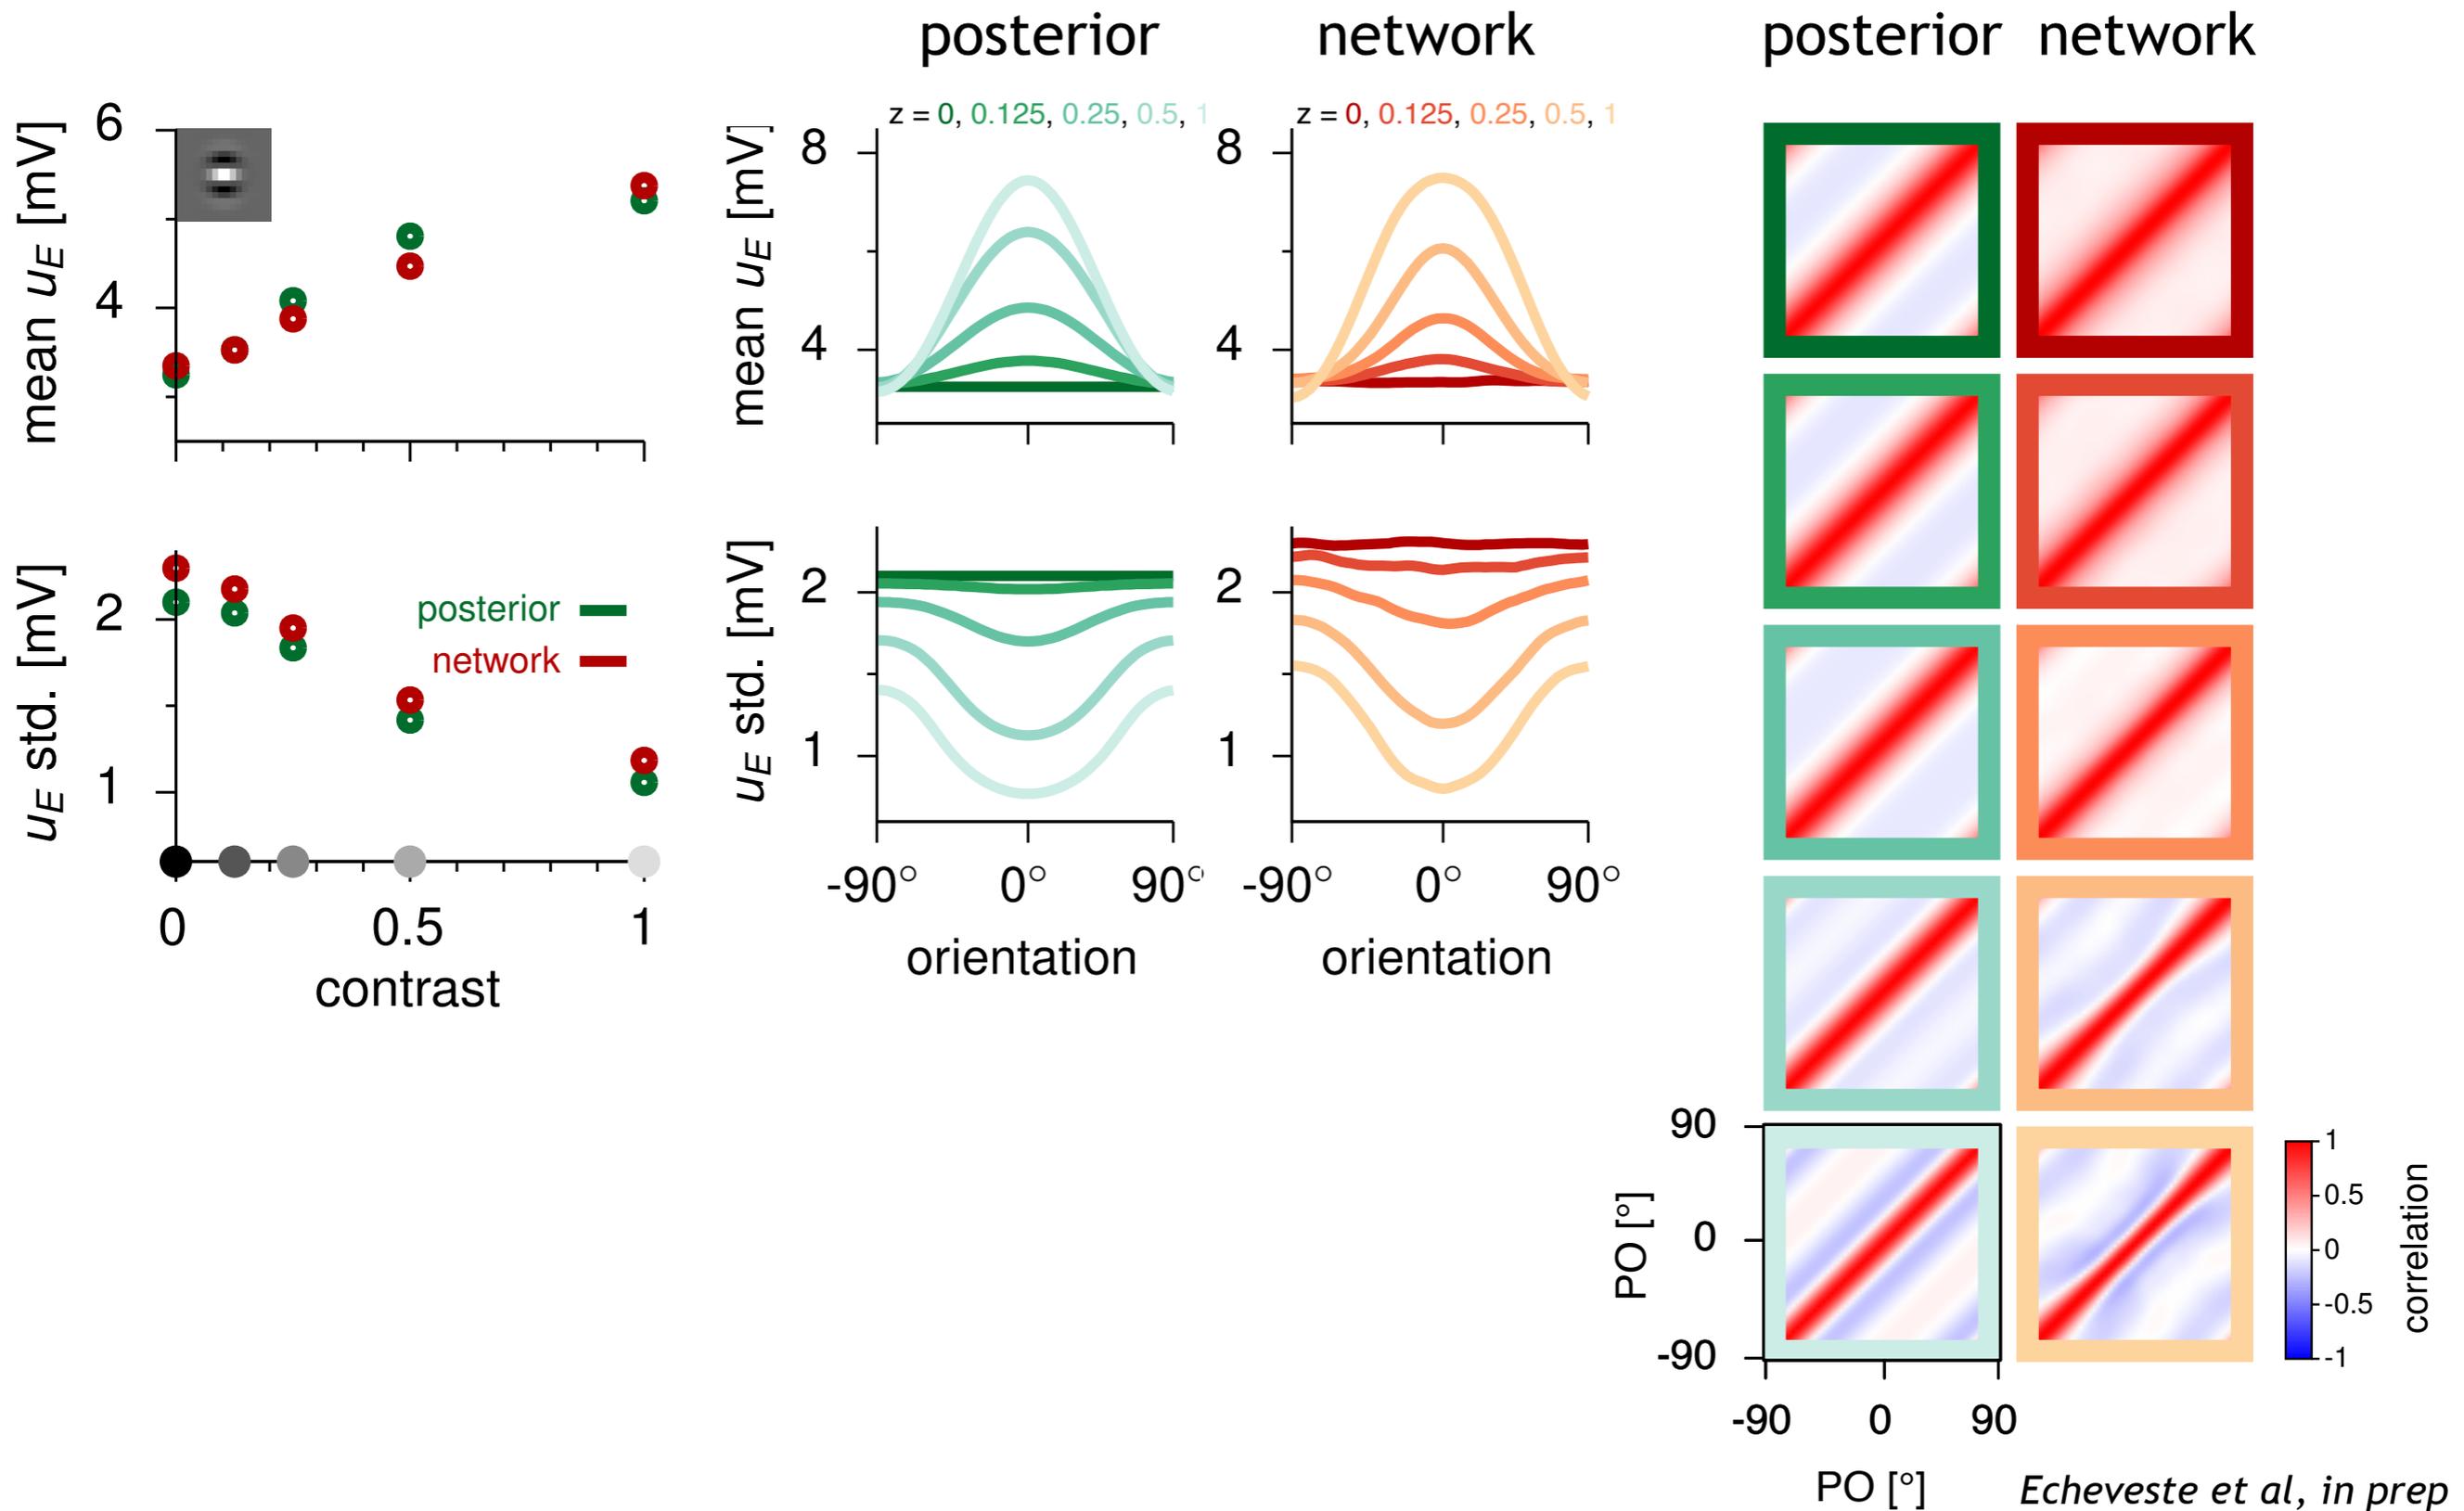
GSM POSTERIOR VS. NETWORK: TRAINING



GSM POSTERIOR VS. NETWORK: TRAINING

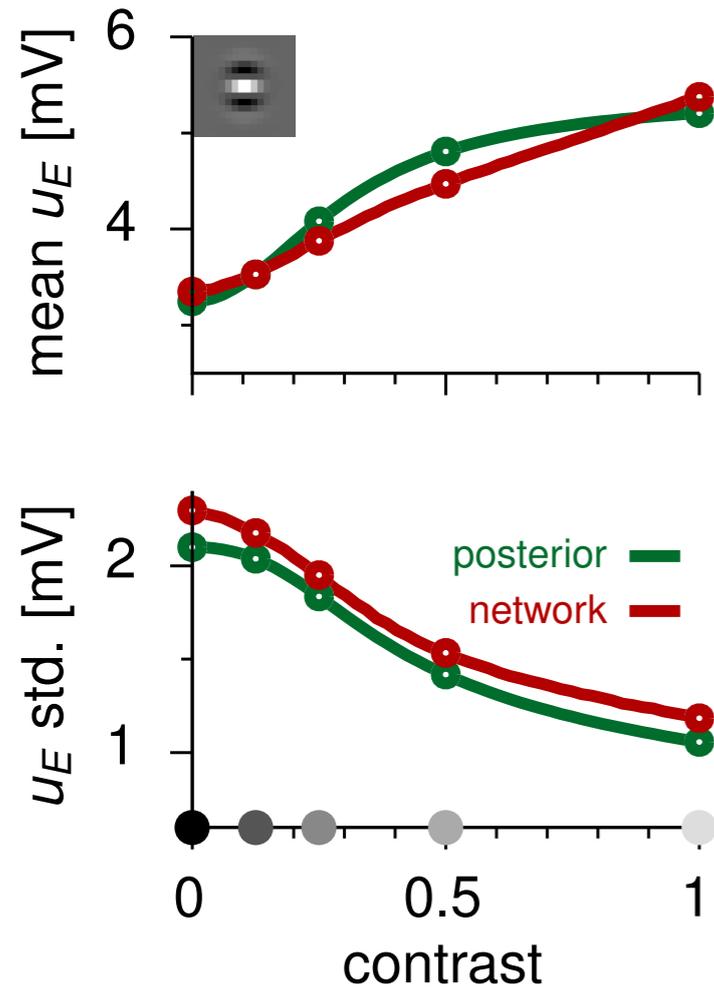


GSM POSTERIOR VS. NETWORK: TRAINING

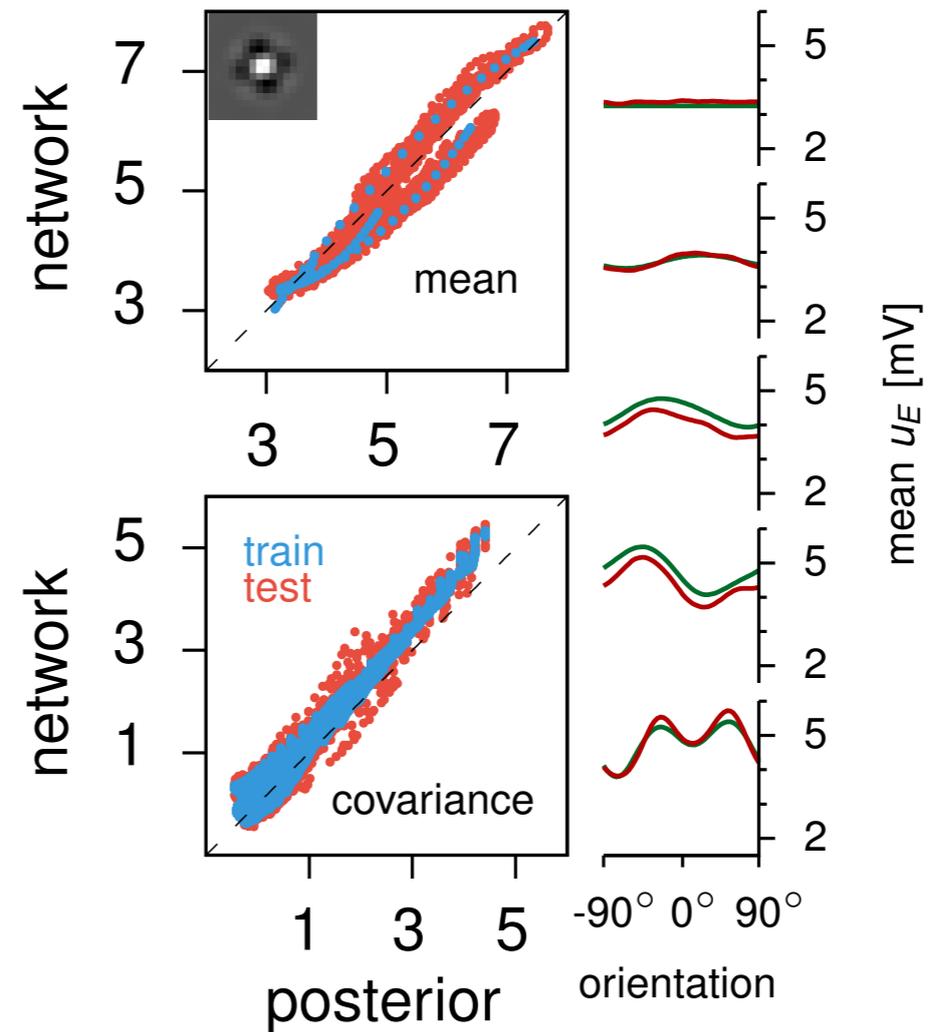
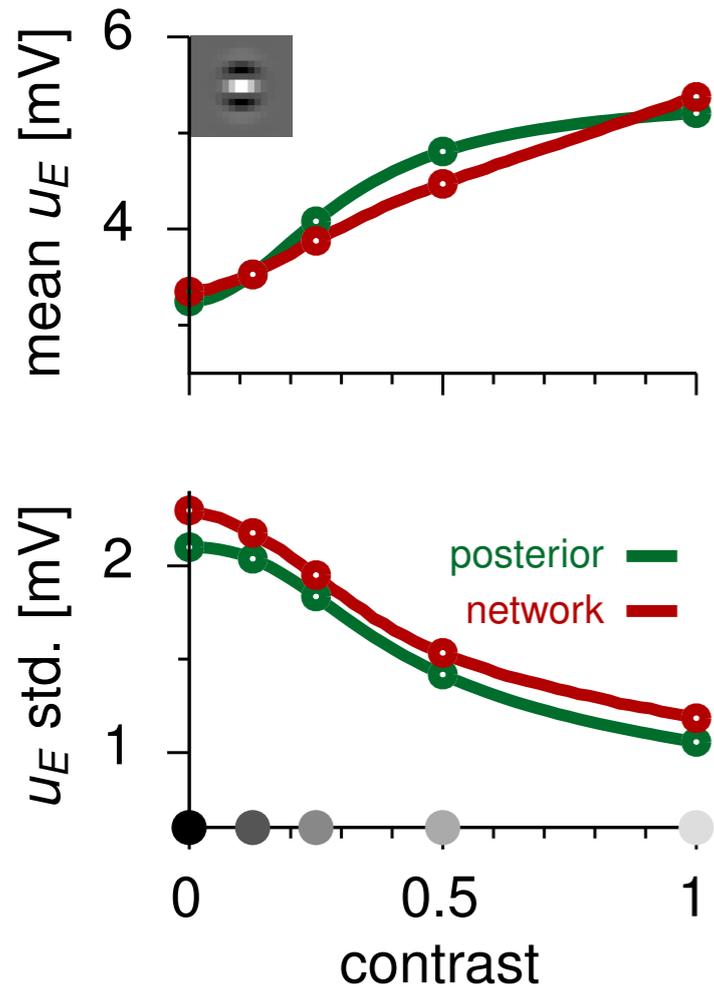


Echeveste et al, in prep

GSM POSTERIOR VS. NETWORK: GENERALISATION



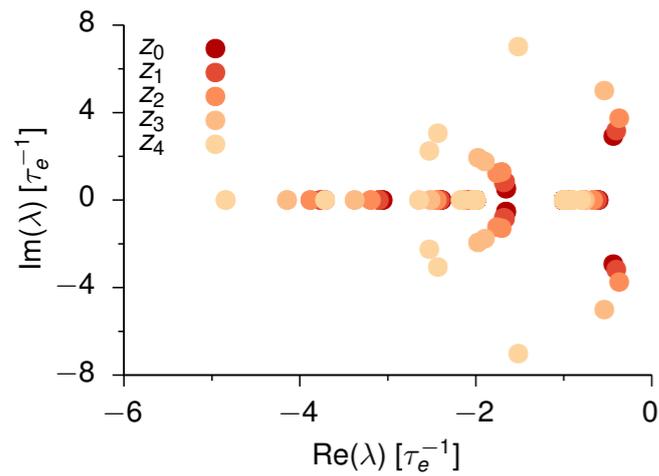
GSM POSTERIOR VS. NETWORK: GENERALISATION



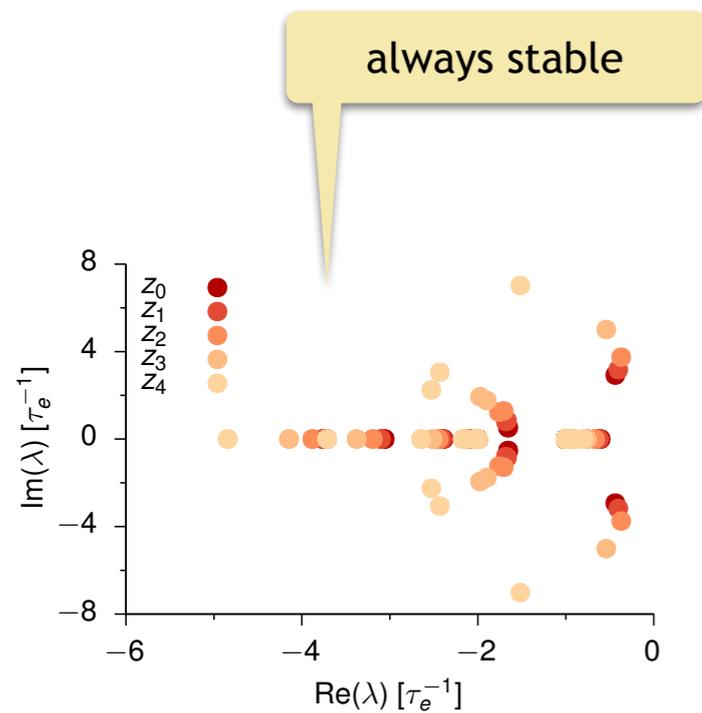
Echeveste et al, in prep

THE SAMPLING-OPTIMIZED NETWORK IS STABLE, FAST & UNBALANCED

THE SAMPLING-OPTIMIZED NETWORK IS STABLE, FAST & UNBALANCED

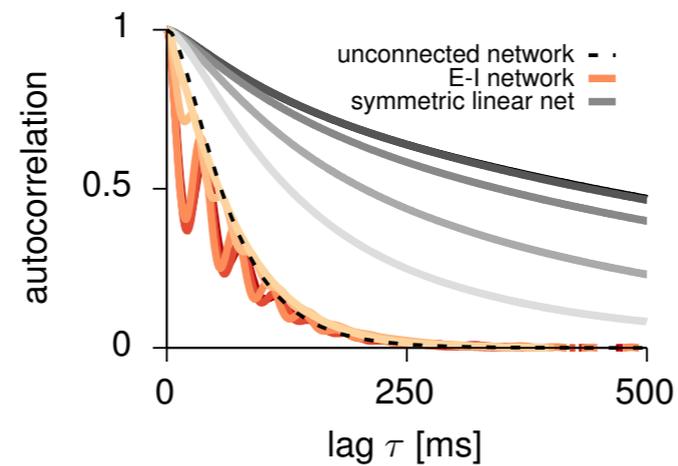
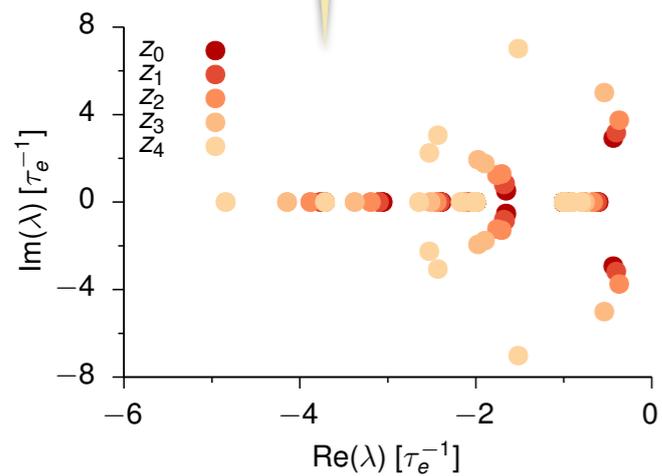


THE SAMPLING-OPTIMIZED NETWORK IS STABLE, FAST & UNBALANCED



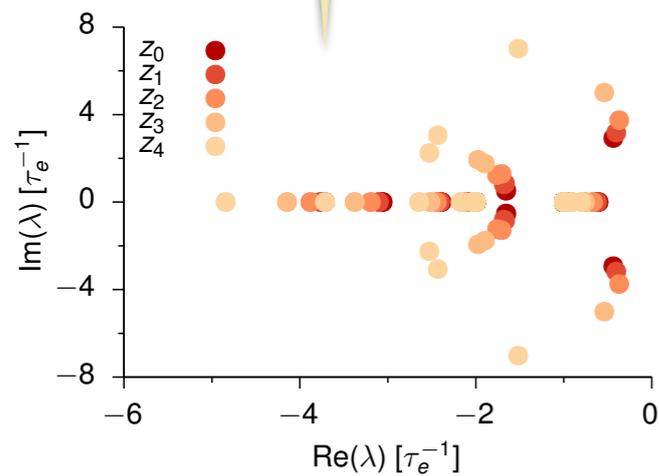
THE SAMPLING-OPTIMIZED NETWORK IS STABLE, FAST & UNBALANCED

always stable



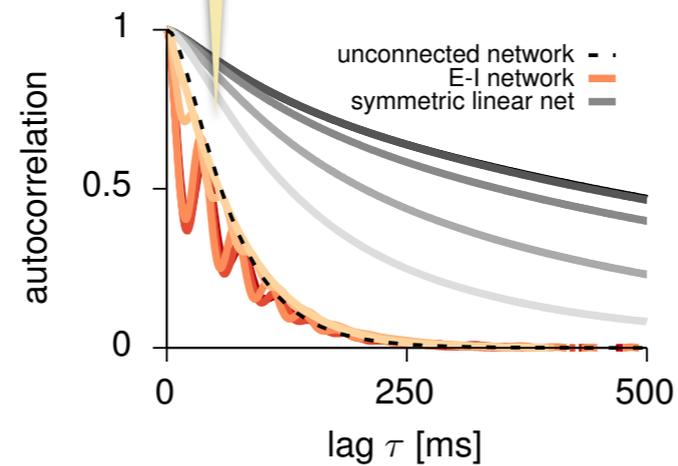
THE SAMPLING-OPTIMIZED NETWORK IS STABLE, FAST & UNBALANCED

always stable

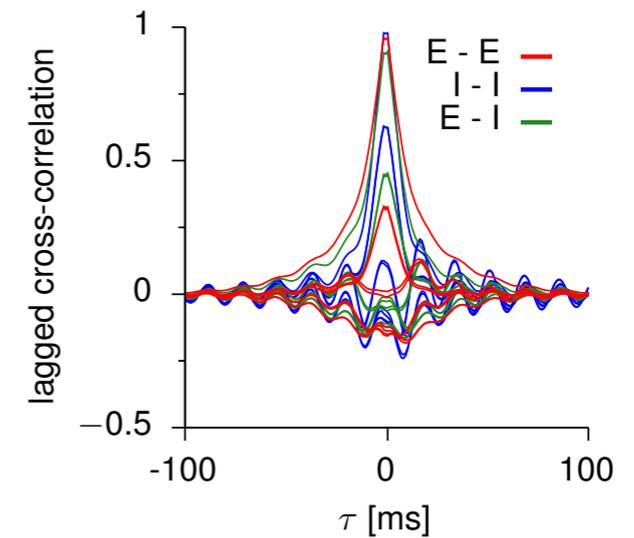
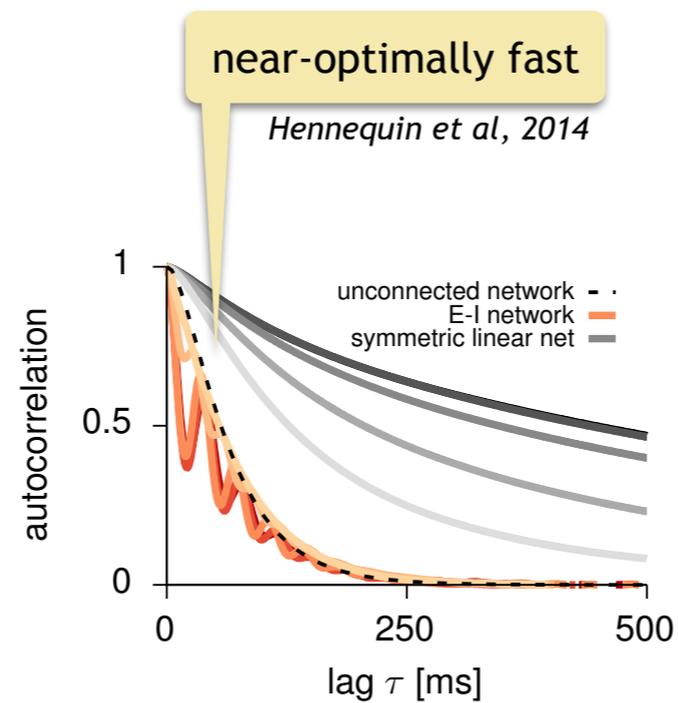
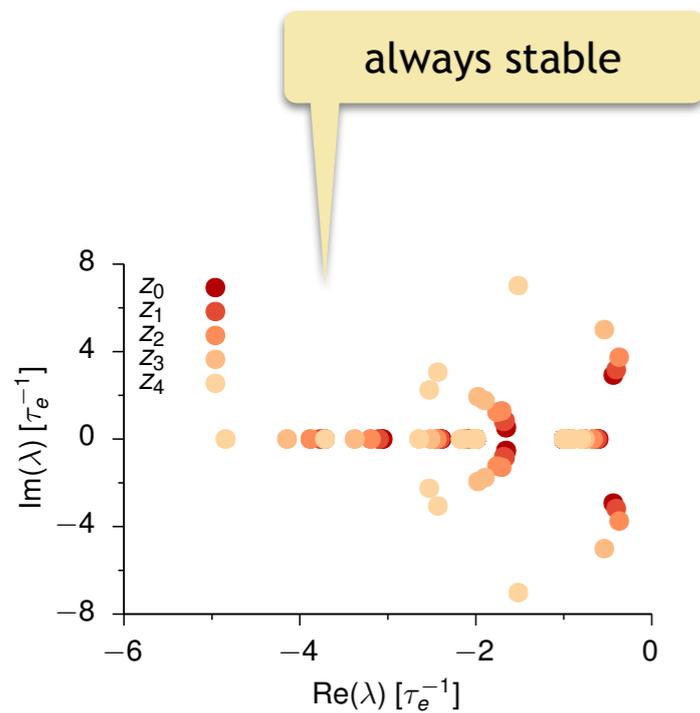


near-optimally fast

Hennequin et al, 2014

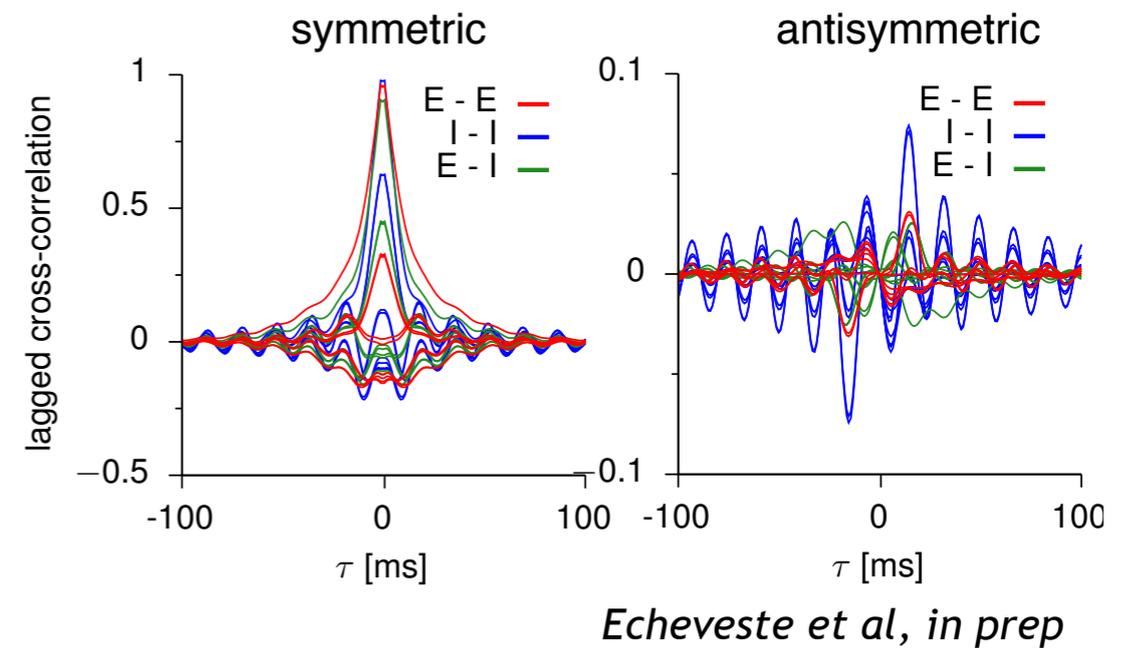
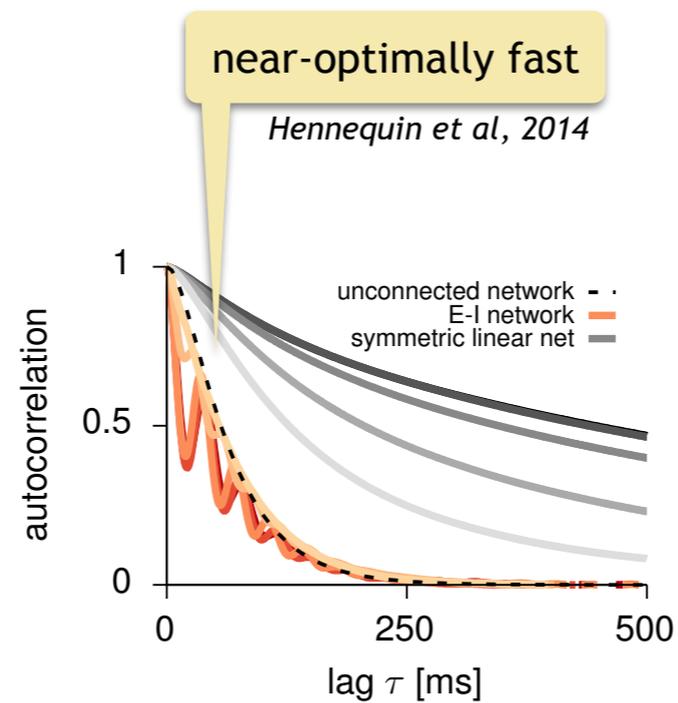
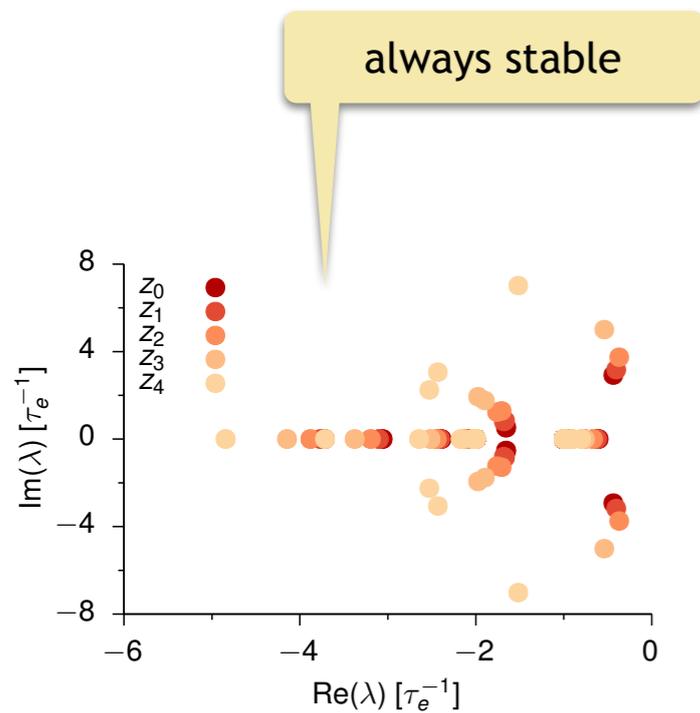


THE SAMPLING-OPTIMIZED NETWORK IS STABLE, FAST & UNBALANCED

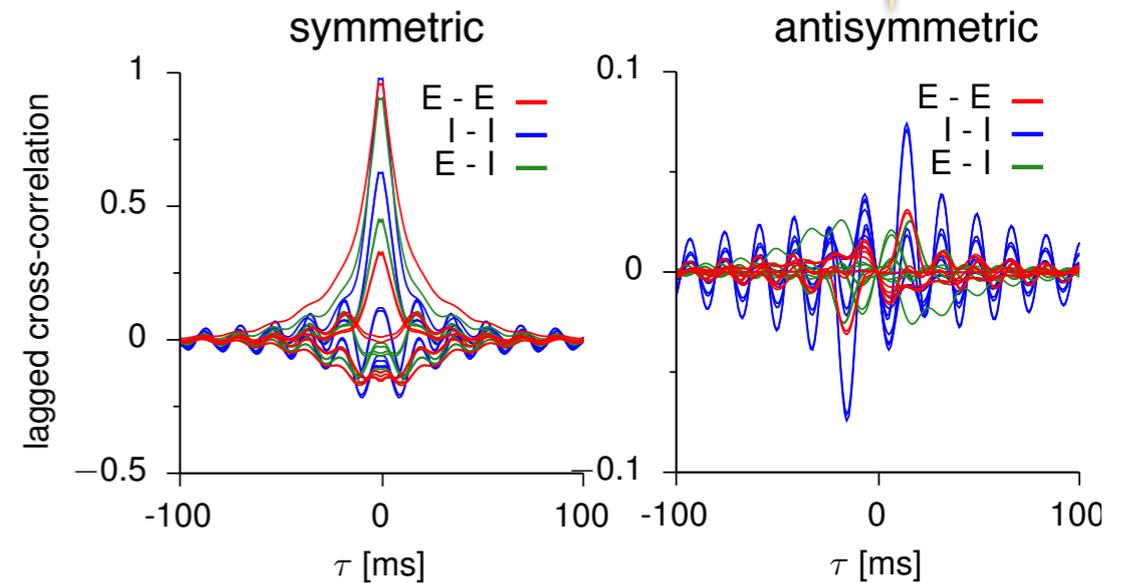
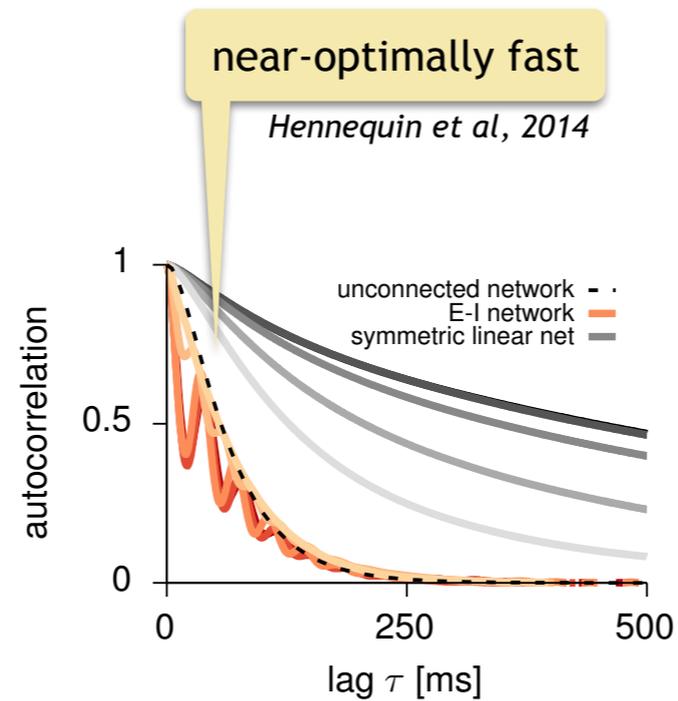
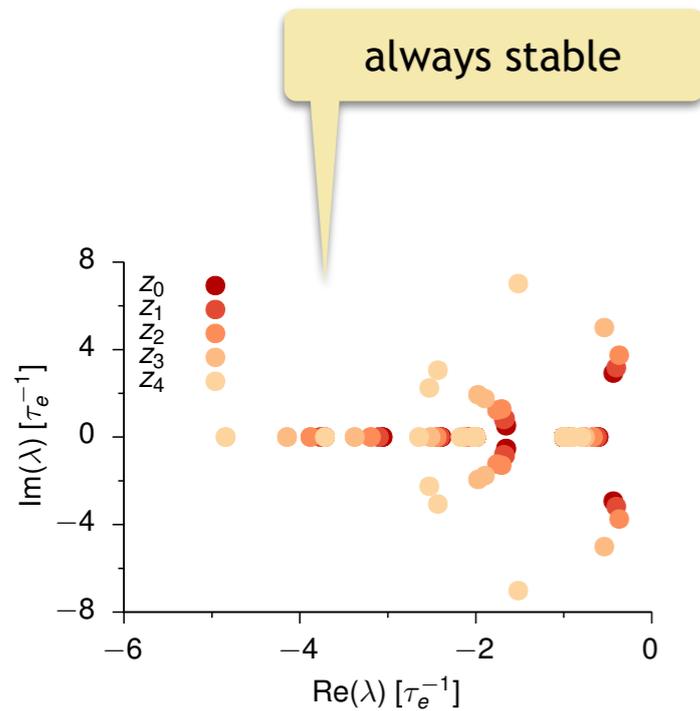


Echeveste et al, in prep

THE SAMPLING-OPTIMIZED NETWORK IS STABLE, FAST & UNBALANCED



THE SAMPLING-OPTIMIZED NETWORK IS STABLE, FAST & UNBALANCED



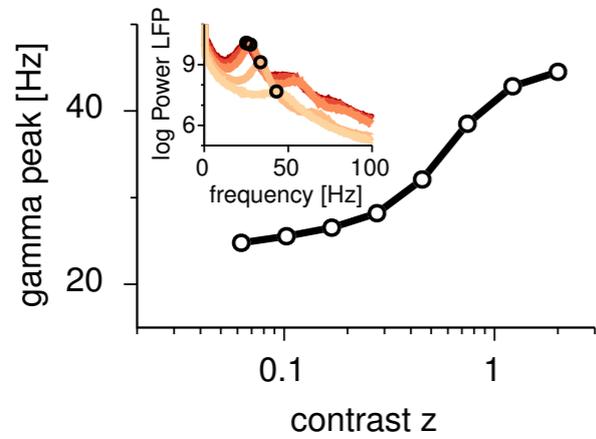
Echeveste et al, in prep

THE SAMPLING-OPTIMIZED NETWORK IS IN A REALISTIC DYNAMICAL REGIME

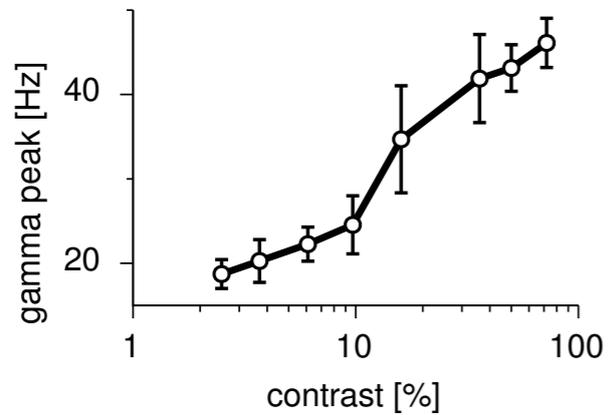
THE SAMPLING-OPTIMIZED NETWORK IS IN A REALISTIC DYNAMICAL REGIME

contrast-dependent oscillations

MODEL

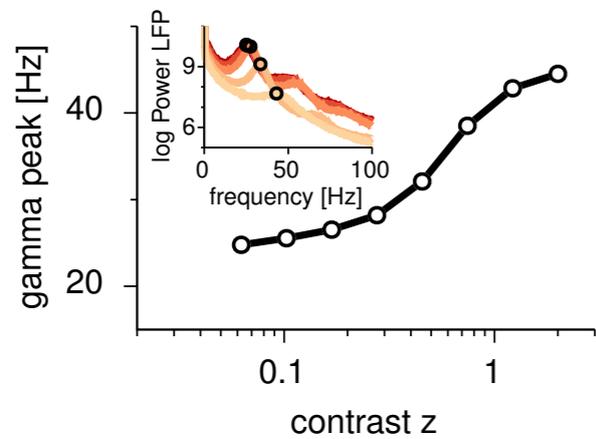


EXPERIMENT
Ray & Maunsell, 2010

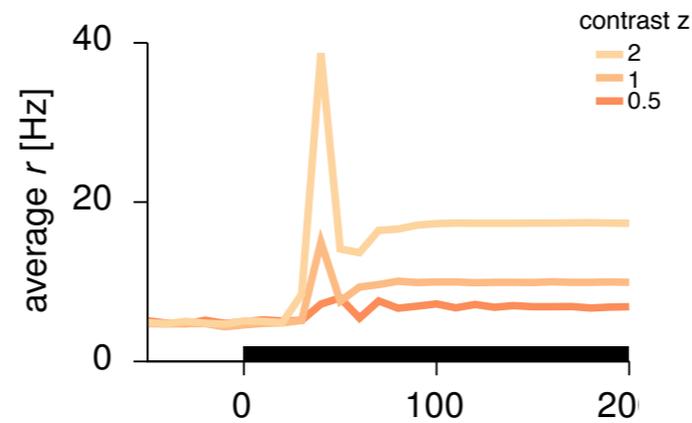


THE SAMPLING-OPTIMIZED NETWORK IS IN A REALISTIC DYNAMICAL REGIME

contrast-dependent oscillations

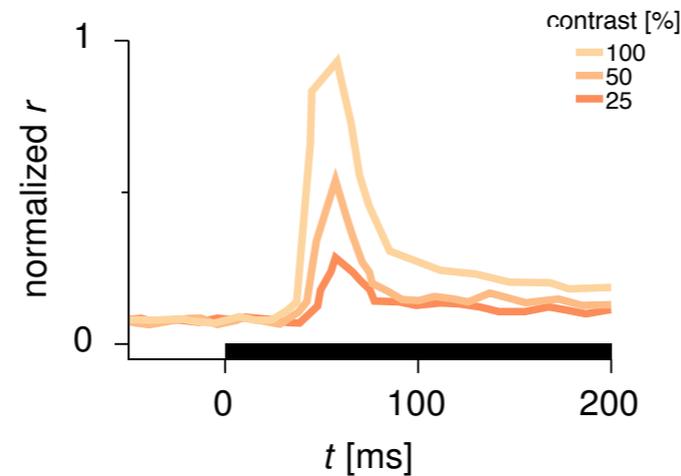
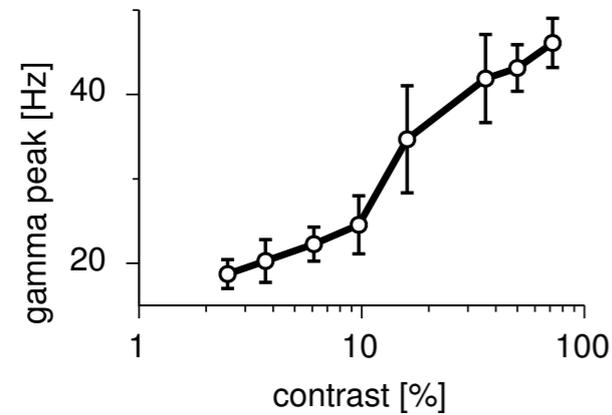


contrast-dependent stimulus-onset transients



MODEL

EXPERIMENT
Ray & Maunsell, 2010

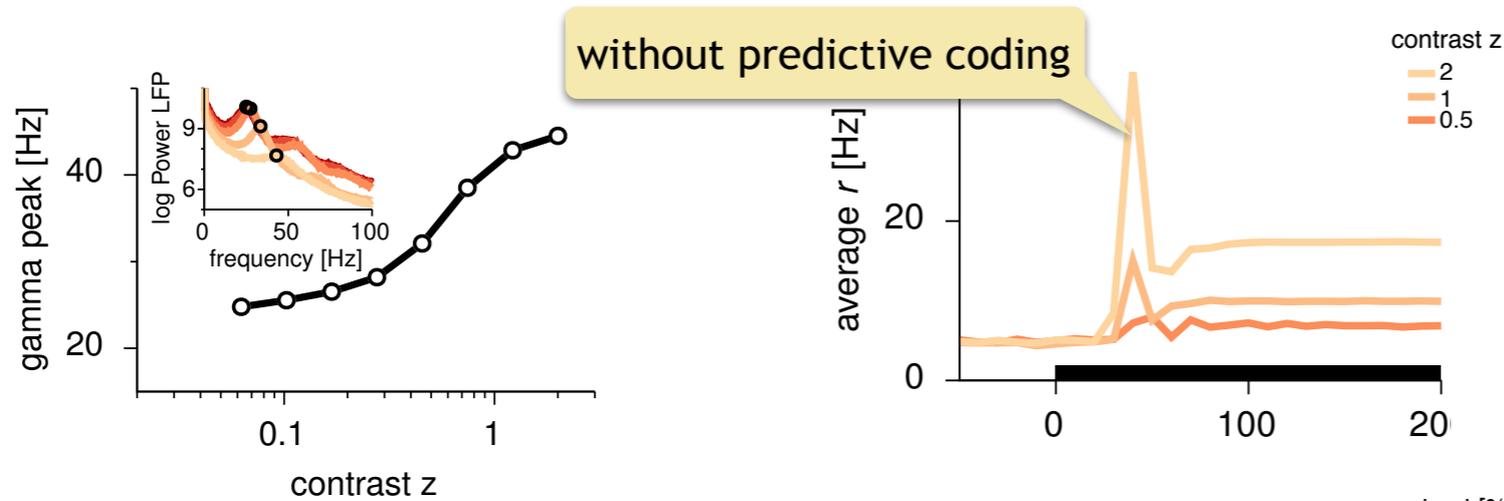


THE SAMPLING-OPTIMIZED NETWORK IS IN A REALISTIC DYNAMICAL REGIME

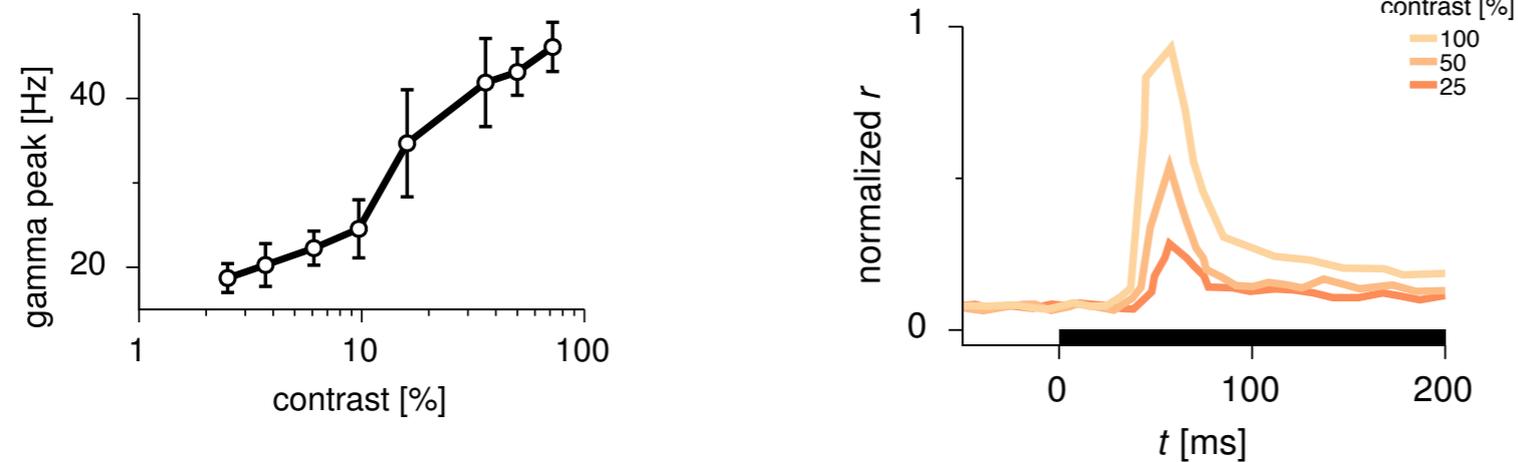
contrast-dependent oscillations

contrast-dependent stimulus-onset transients

MODEL

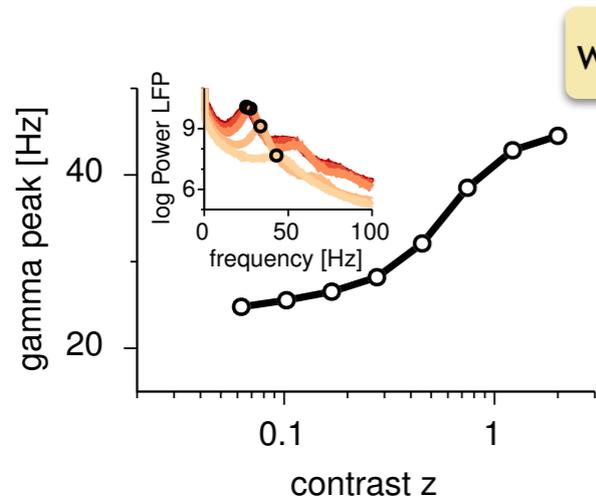


EXPERIMENT
Ray & Maunsell, 2010

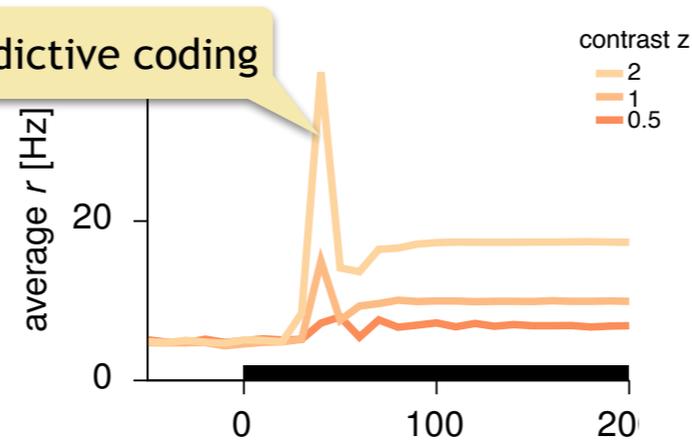


THE SAMPLING-OPTIMIZED NETWORK IS IN A REALISTIC DYNAMICAL REGIME

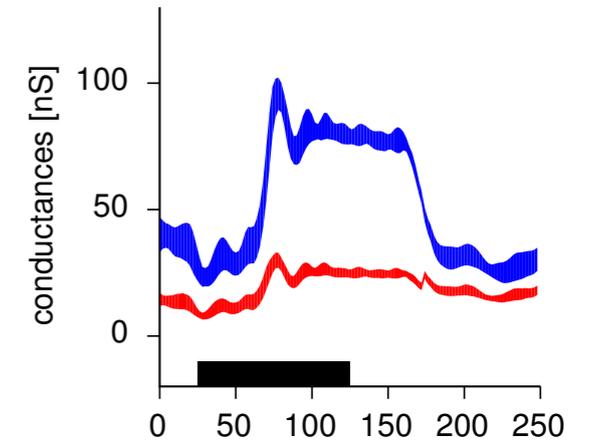
contrast-dependent oscillations



contrast-dependent stimulus-onset transients



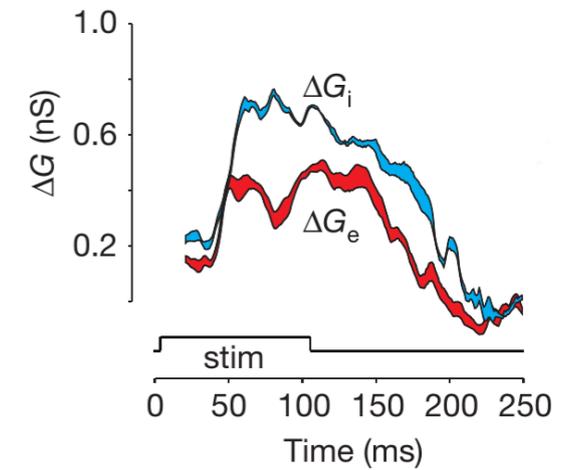
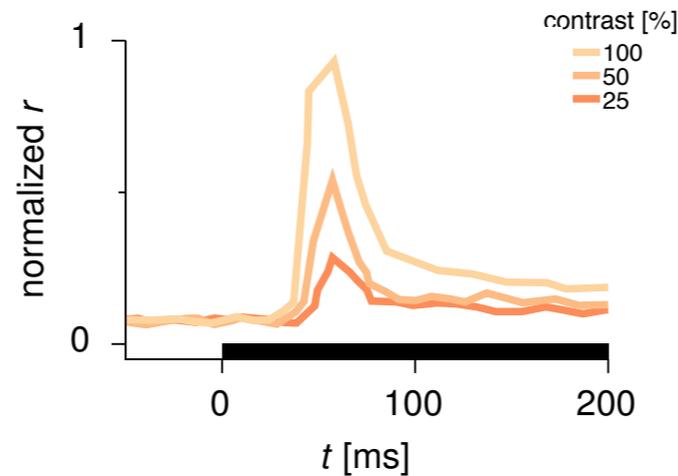
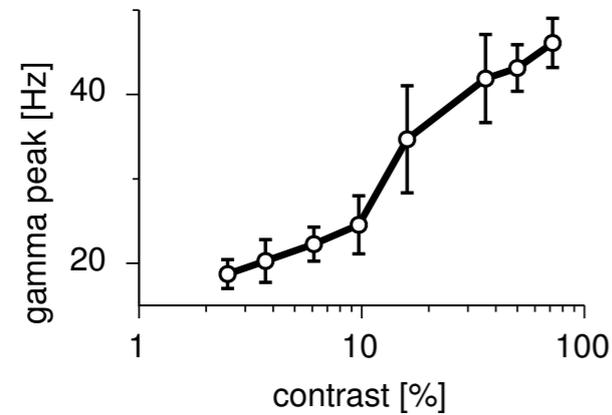
inhibition dominated



MODEL

EXPERIMENT

Ray & Maunsell, 2010
Haider et al, 2013



Echeveste et al, in prep

USING VARIABILITY IS CRITICAL

ideal observer: GSM

- ▶ full distribution
(2nd order)



train

visual cortex: neural network

- ▶ recurrent
- ▶ E-I
- ▶ expansive nonlinearity

USING VARIABILITY IS CRITICAL

ideal observer: GSM

- ▶ just the means
(1st order)



train

visual cortex: neural network

- ▶ recurrent
- ▶ E-I
- ▶ expansive nonlinearity

USING VARIABILITY IS CRITICAL

ideal observer: GSM

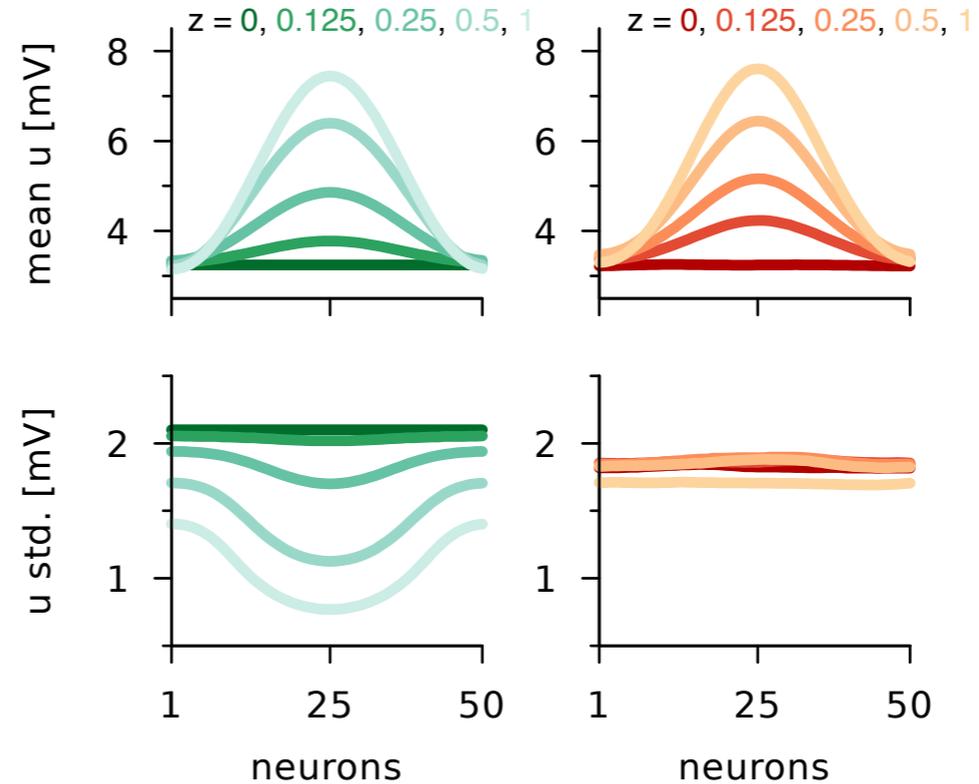
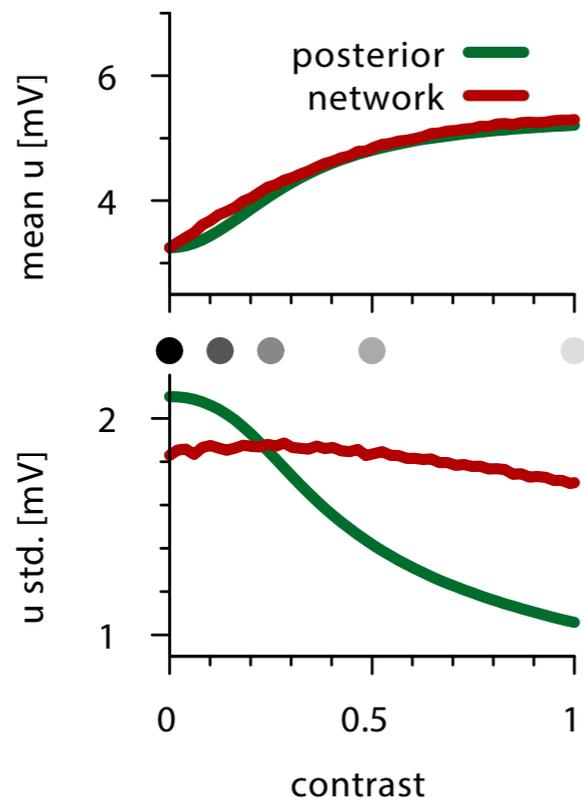
- ▶ just the means (1st order)



train

visual cortex: neural network

- ▶ recurrent
- ▶ E-I
- ▶ expansive nonlinearity



USING VARIABILITY IS CRITICAL

ideal observer: GSM

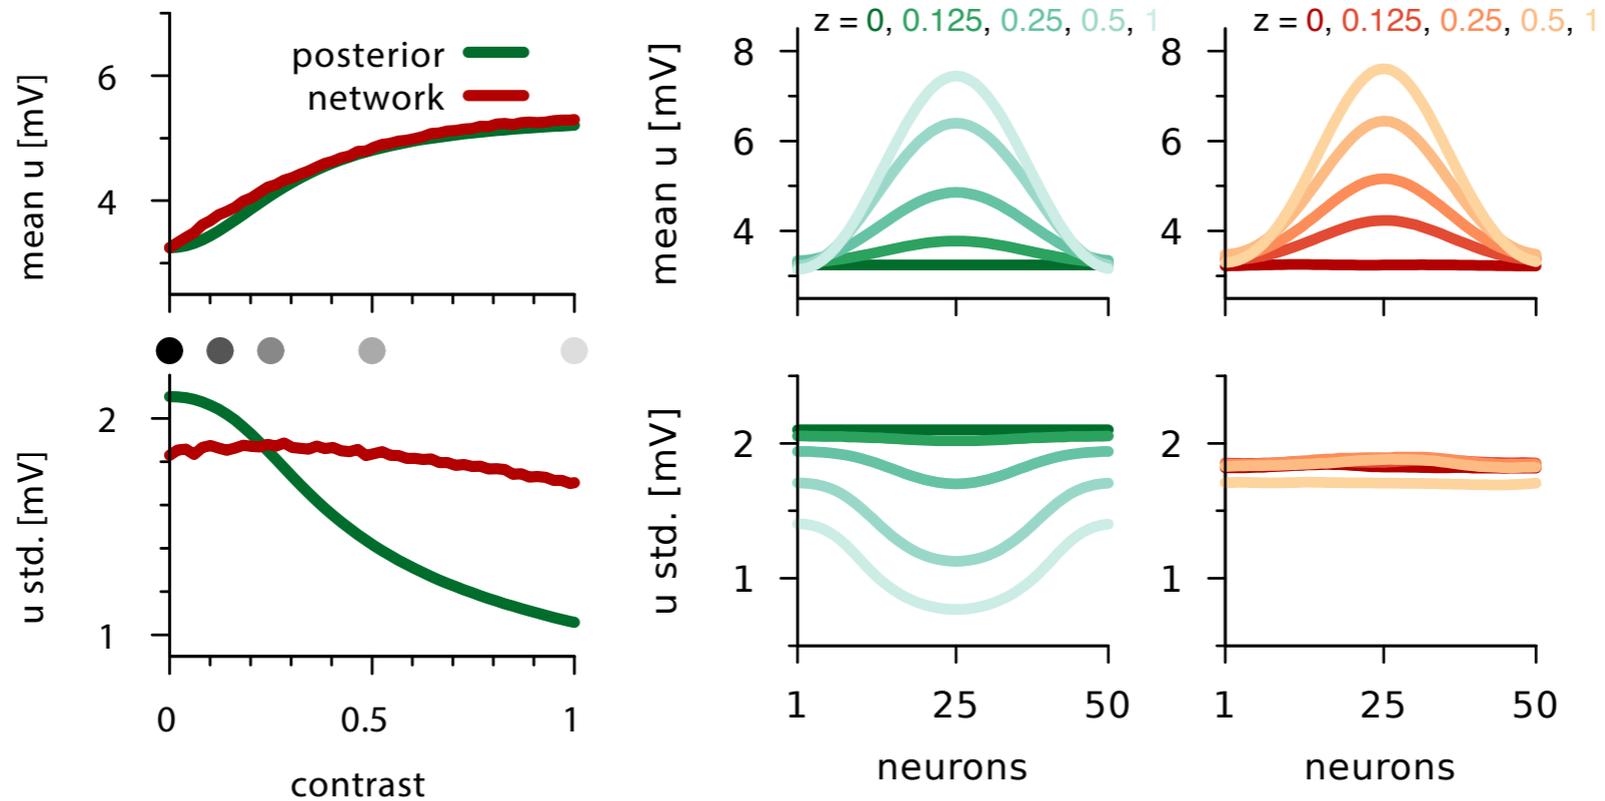
- ▶ just the means (1st order)



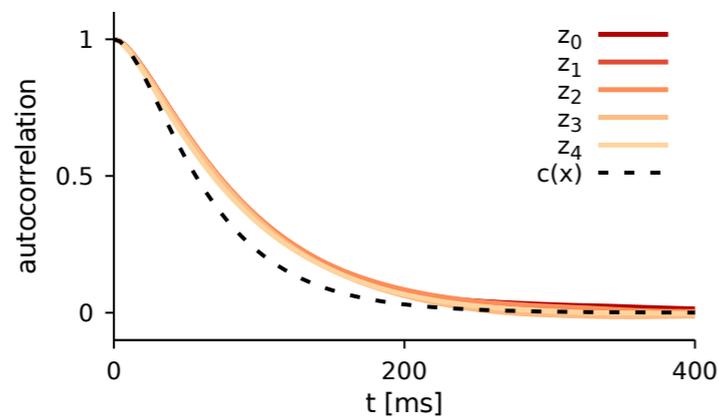
train

visual cortex: neural network

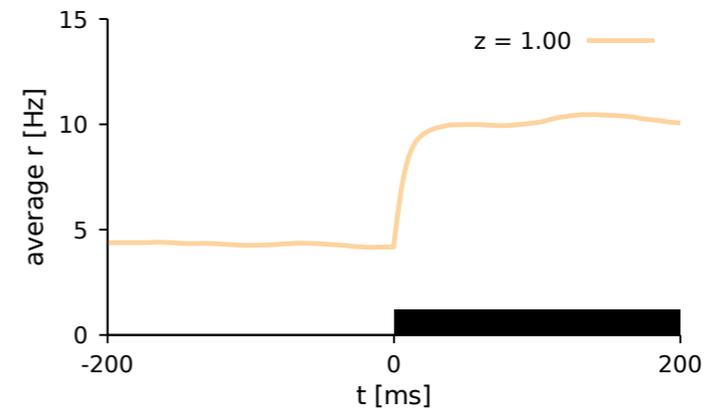
- ▶ recurrent
- ▶ E-I
- ▶ expansive nonlinearity



no fastness, no gamma



no transients



PROBABILISTIC INFERENCE AND LEARNING

cognitive science

neuroscience

theory

experiments

theory

experiments

✓

?

program \mathcal{P}

?

?

structure \mathcal{M}

✓

✓

?

?

parameters θ

✓

✓

✓✓

?

latent variables y

✓

✓

✓✓✓

?

data x

POSTDOC POSITION AVAILABLE

estimate humans' complex, high dimensional,
dynamically changing internal models from behaviour

collaborate with world-leading
experimental cognitive scientists

