Statisztikus tanulás az idegrendszerben

ORBÁN GERGŐ

golab.wigner.mta.hu

WIGNER RESEARCH CENTRE FOR PHYSICS

TDK- és kutatási témáink

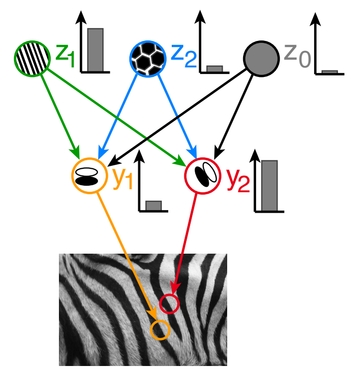

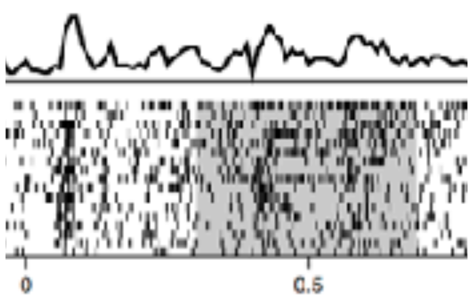

Vizuális kérgi reprezentációk modellezése

Az agykéreg vizuális rendszere az objektumfelismerést hierarchikus feldolgozás során valósítja meg, amelyben az idegsejtek együttes aktivitása a stimulust felépítő képelemek kompozícióját kódolja. E folyamat valószínűségi következtetésként való modellezése lehetővé teszi a neurális aktivitás predikcióját, és így az egymásra épülő agykérgi reprezentációk feltárását. Modelljeinket a gépi tanulás legújabb eredményeit felhasználva, a partnereink által feladatokat végrehajtó állatok agykérgéből elvezetett neurális aktivitáson tesztelve fejlesztjük adatelemzési módszerekkel párhuzamosan.

publikációk: Orbán+, 2016 Neuron; Bányai, Koman, Orbán, 2017 J Neurophys; Bányai+, 2017 biorXiv (preprint)

partnerek: Ernst Strüngmann Institute Frankfurt; University of Cambridge; UC Los Angeles

pozíciók: TDK, BSc, MSc, PhD, posztdoktori

kontakt: Orbán Gergő, Bányai Mihály, Stippinger Marcell

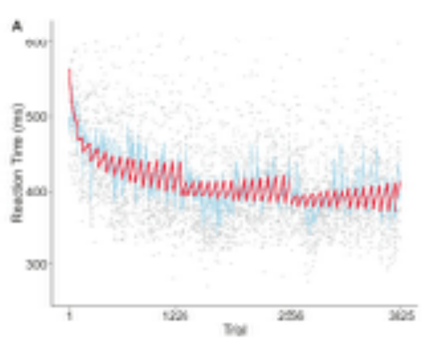

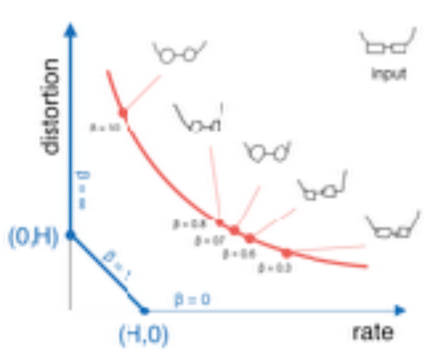

Memóriarendszerek és kognitív reprezentációk modellezése

Az emberi kognició során az agynak olyan információfeldolgozási folyamatokat kell megoldania mint az érzékelés, predikció, döntéshozás vagy a környezet modelljének felépítése tapasztalatok alapján. Csoportunkban ezen folyamatokat normatív megfontolások és a gépi tanulás eszközei segítségével vizsgáljuk. Futó projektjeinkben többek között az agy a környezetről és konkrét feladatokról alkotott belső reprezentációját következtetjük ki partnereink által rögzített viselkedéses adatokra alapozva, illetve tanulóágensek által használt memóriarendszerek optimális dinamikáját és ennek az emberi memóriakísérletekre való implikációit tárjuk fel.

publikációk: Orbán+, 2008 PNAS; Houlsby+, 2013; Nagy, Orbán, 2016 COGSCI; Török+, 2017 JEP

partnerek: ELTE PPK Pszichológia Intézet

pozíciók: TDK, BSc, MSc, PhD, posztdoktori

kontakt: Orbán Gergő, Nagy Dávid, Török Balázs

Előfeltételek: önálló programozási készség, lineáris algebra, analízis, valószínűségszámítás ismerete, angol nyelv használata

Kurzusaink I.

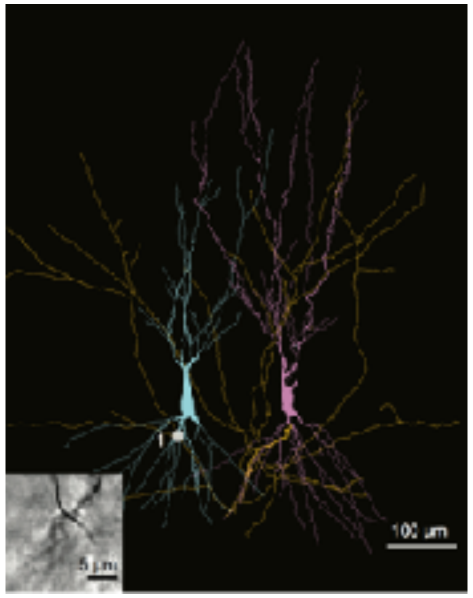

Idegrendszeri modellezés

A kurzus célja, hogy a hallgatók betekintést nyerjenek az idegrendszer működésébe. A modellek szerepe kettős: Egyrészt az idegrendszer egy komplex rendszer, a sokféle információ szintetizálásához pontos modellekre van szükség. Másrészt az idegrendszer modellezi a külvilágot, az aktivitását megfigyelve ezekről a modellekről is információt gyűjthetünk. A kurzushoz szabadon választható, egyszerű gyakorlati feladatok (programozás R nyelven) is kapcsolódnak.

feltételek: matematikai alapismeretek, magyar nyelv

tematika: http://cneuro.rmki.kfki.hu/education/neuromodel

meghirdetve: ELTE TTK, BME TTK; őszi félév

oktatók: Orbán Gergő, Somogyvári Zoltán, Ujfalussy Balázs

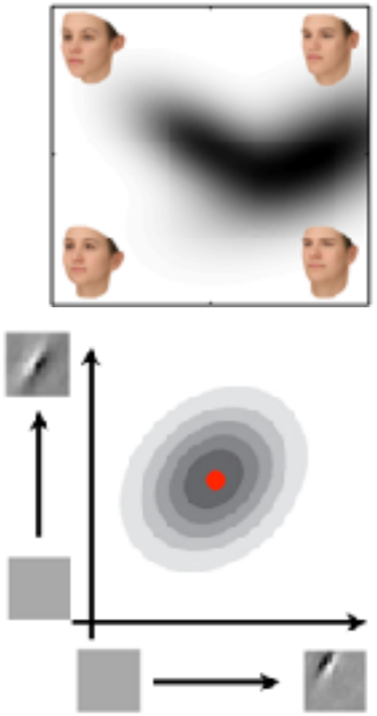

Statisztikai tanulás az idegrendszerben

Az agy működésének megértése a tudomány egyik legérdekesebb kérdése napjainkban. A kurzus a kérdést az agy feladatainak matematikai leírása és algoritmikus megoldása irányából közelíti meg, különös tekintettel a probabilisztikus módszerek alkalmazására. Foglalkozik a biológiai tanulás, reprezentáció és kódolás matematikai modelljeivel, illetve azzal is, hogy ezeket a számításokat hogyan implementálhatja az idegrendszer. A kurzus az általános elvek áttekintésén túl speciálisan a vizuális információ agykérgi feldolgozásának problémájára koncentrál.

feltételek: lineáris algebra, programozási tapasztalat, magyar nyelv

tematika: http://golab.wigner.mta.hu/teaching

meghirdetve: ELTE TTK, BME TTK; tavaszi félév

oktatók: Orbán Gergő, Bányai Mihály, Nagy Dávid, Török Balázs

Kurzusaink II.

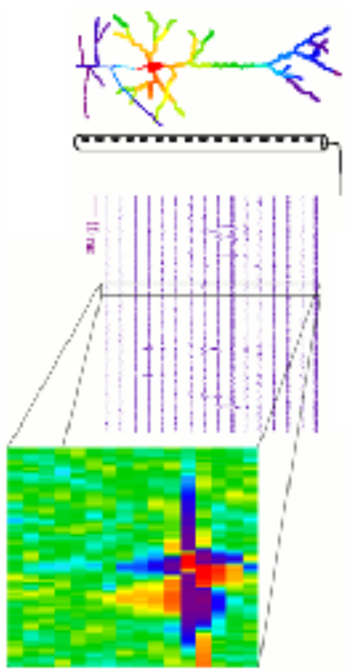

Neuroelektrofiziológiai adatelemzés

A kurzus azokat a matematikai elemzései eljárásokat mutatja be, amelyek az idegrendszerben mért - elsősorban elektromos – jelek elemzésében és értelmezésében jelenősek, kezdve a legegyszerűbb számításoktól egészen a legújabb módszerekig, akár nyitott kérdéseket is tárgyalva. Bár a kurzus során az idegi elektromos aktivitás elemzése lesz a vezérfonál és az ismertetett módszerek alkalmazási területe, az itt megismert matematikai eszközök a tudomány – és nem csak a tudomány - bármely területén használhatóak és hasznosak, ahol mért adatok alapján, egy komplex rendszer szerkezetének és működésének felderítése a cél.

feltételek: matematikai alapismeretek, magyar nyelv

tematika: http://cneuro.rmki.kfki.hu/education/statbio

meghirdetve: ELTE TTK; tavaszi félév

oktató: Somogyvári Zoltán

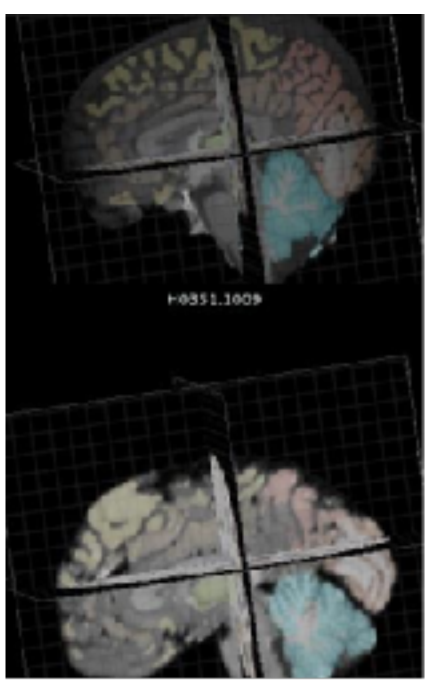

Neuroinformatika

A kurzus célja, hogy különböző élettudományi tudományterületek példáján keresztül a diákok elsajátítsák az adatfeldolgozás, beleértve a biostatisztika, jelfeldolgozás, modellezés, képfeldolgozás és más matematikai módszerek elméleti, és alkalmazásuk informatikai alapjait egy-egy közismert program használatába történő bevezetéssel úgy, hogy azt saját kutatási feladataikban is alkalmazni tudják.

feltételek: magyar nyelv

tematika: http://cneuro.rmki.kfki.hu/education/neuroinfo

meghirdetve: SE Szenthágotai DI; őszi félév, kétévente

oktatók: Négyessy László, Somogyvári Zoltán, Bányai Mihály, Bazsó Fülöp, Zalányi László és mások

Kurzusaink III.

Neocortex: from structure to function

feltételek: angol nyelv

tematika: http://sysneuro-semester.org/

meghirdetve: SE; nyári szünet, egy héten tömbösítve

oktató: Négyessy László

Computational Models in Systems Neuroscience

feltételek: angol nyelv

tematika: http://sysneuro-semester.org/

meghirdetve: SE; nyári szünet, egy héten tömbösítve

oktató: Bányai Mihály

Learning and Navigation

feltételek: angol nyelv

tematika: http://sysneuro-semester.org/

meghirdetve: SE; nyári szünet, egy héten tömbösítve

oktató: Somogyvári Zoltán

Statistics of the Brain

feltételek: angol nyelv

tematika: http://sysneuro-semester.org/

meghirdetve: SE; nyári szünet, egy héten tömbösítve

oktató: Orbán Gergő

Introduction

Knowledge representation

Probabilistic models

Bayesian behaviour

Approximate inference I (computer lab)

Vision I

Approximate inference II: Sampling

Measuring priors

Neural representation of probabilities

Structure learning

Vision II

Decision making and reinforcement learning

elméleti

- Introduction
- Knowledge representation
- Probabilistic models
- Bayesian behaviour
- Approximate inference I (computer lab)
- Vision I
- Approximate inference II: Sampling
- Measuring priors
- Neural representation of probabilities
- Structure learning
- Vision II
- Decision making and reinforcement learning

elméleti

kognitív

Introduction

Knowledge representation

Probabilistic models

Bayesian behaviour

Approximate inference I (computer lab)

Vision I

Approximate inference II: Sampling

Measuring priors

Neural representation of probabilities

Structure learning

Vision II

Decision making and reinforcement learning

Introduction

Knowledge representation

Probabilistic models

Bayesian behaviour

Approximate inference I (computer lab)

Vision I

Approximate inference II: Sampling

Measuring priors

Neural representation of probabilities

Structure learning

Vision II

Decision making and reinforcement learning

elméleti

kognitív

neurális

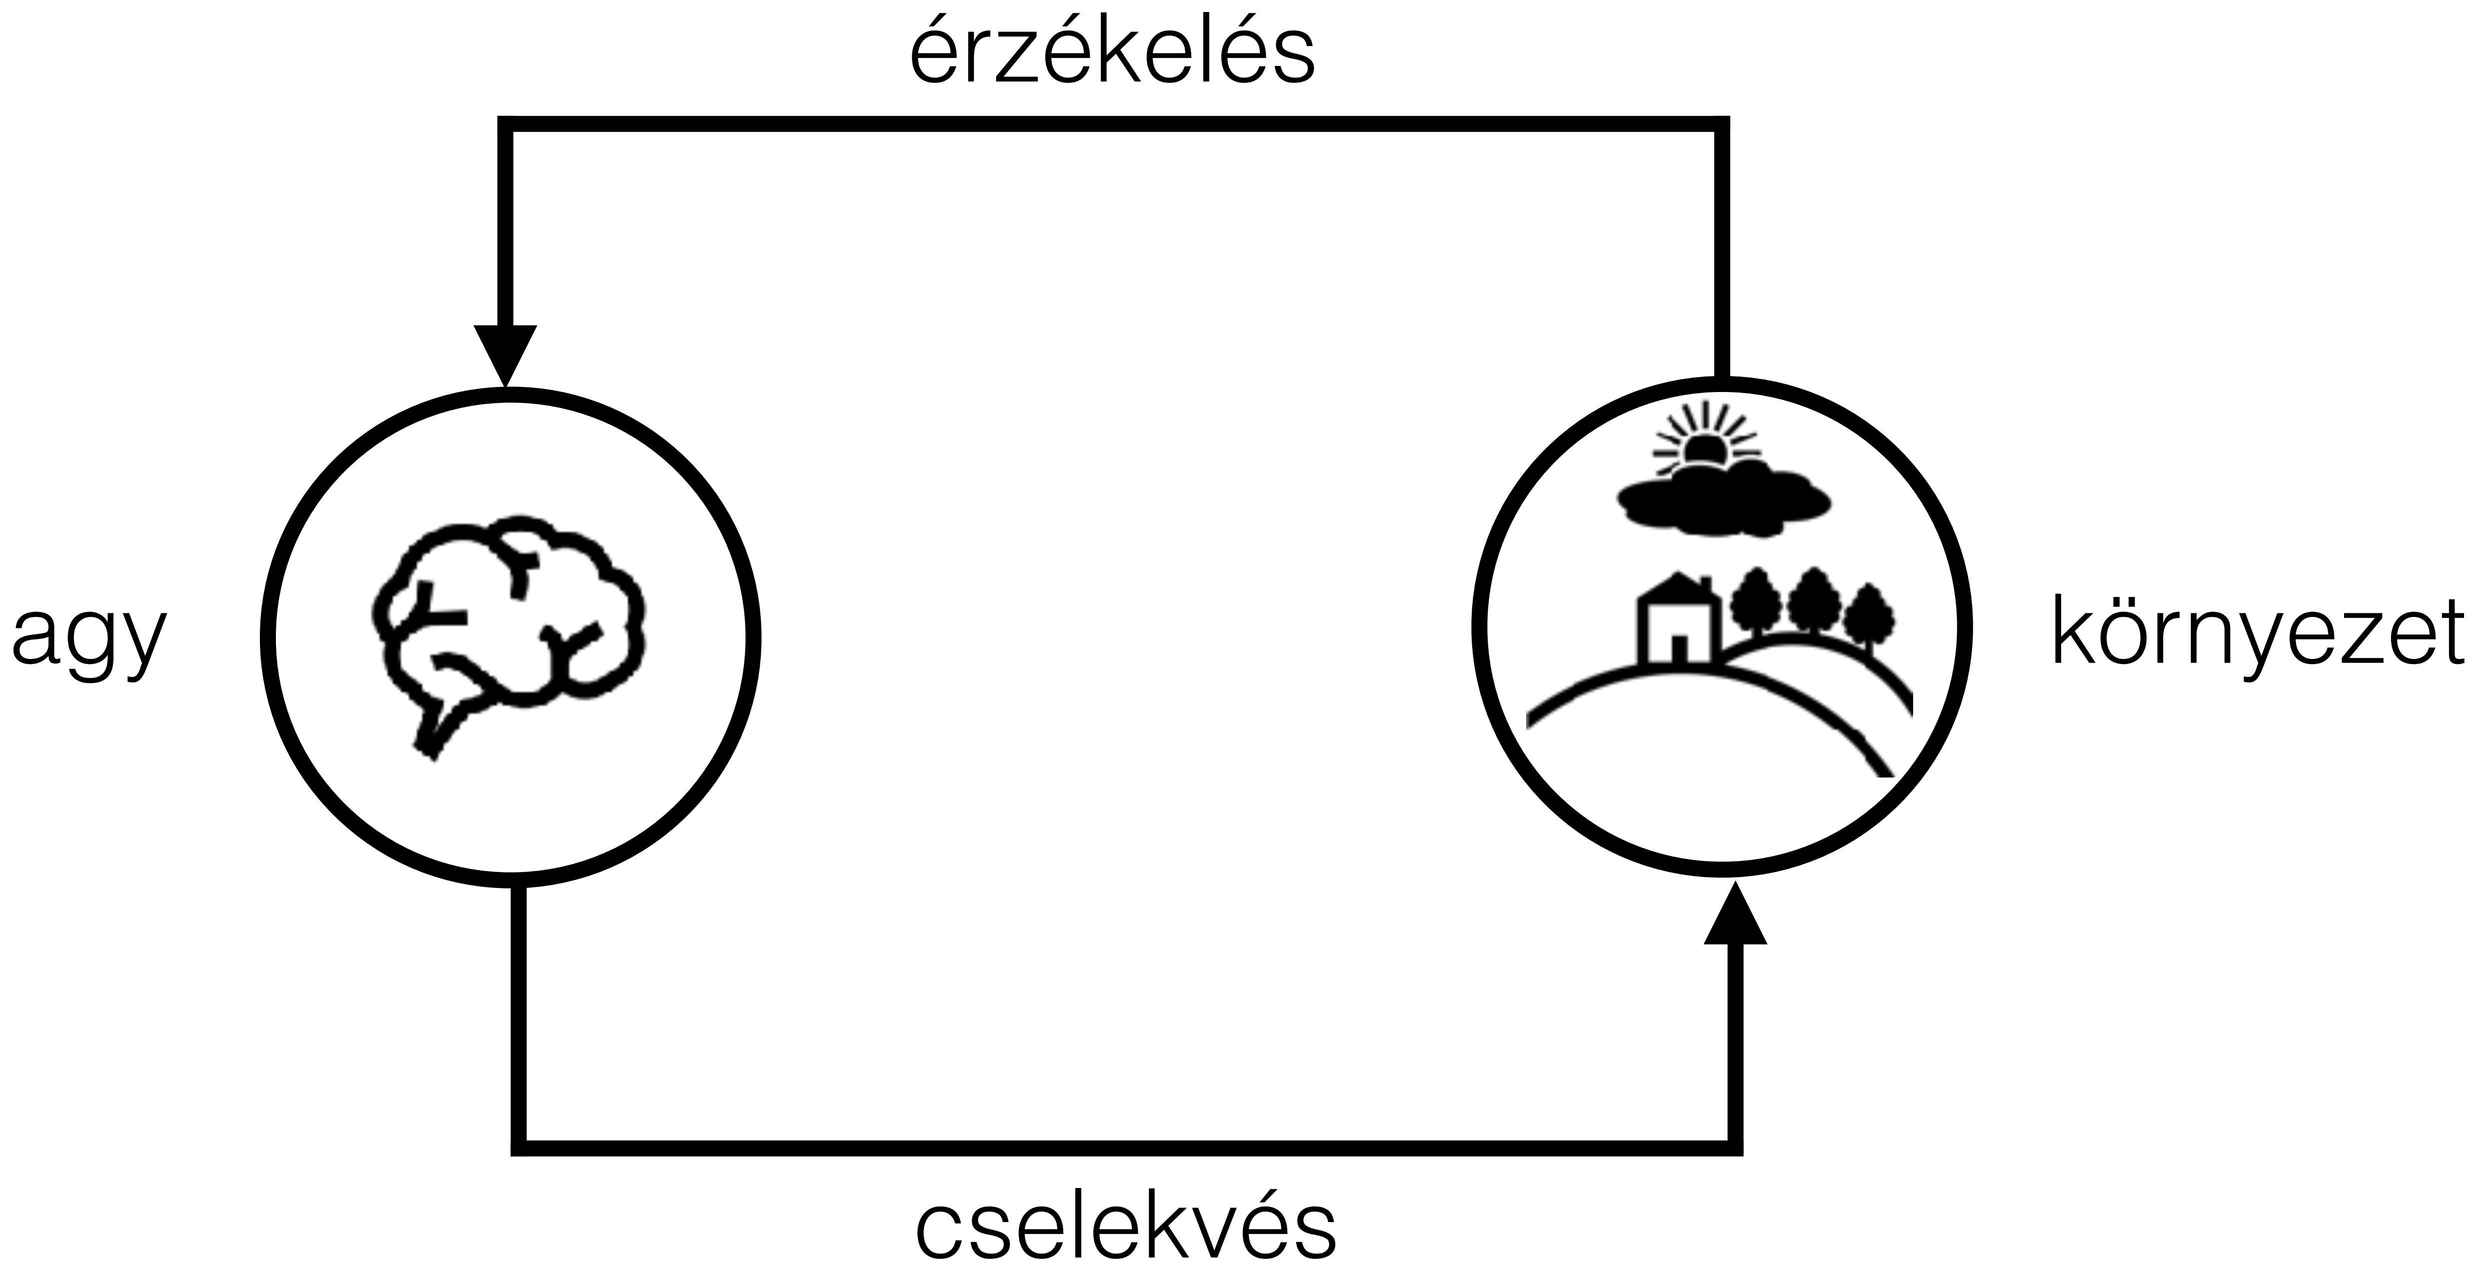
érzékelés
agy
környezet
cselekvés

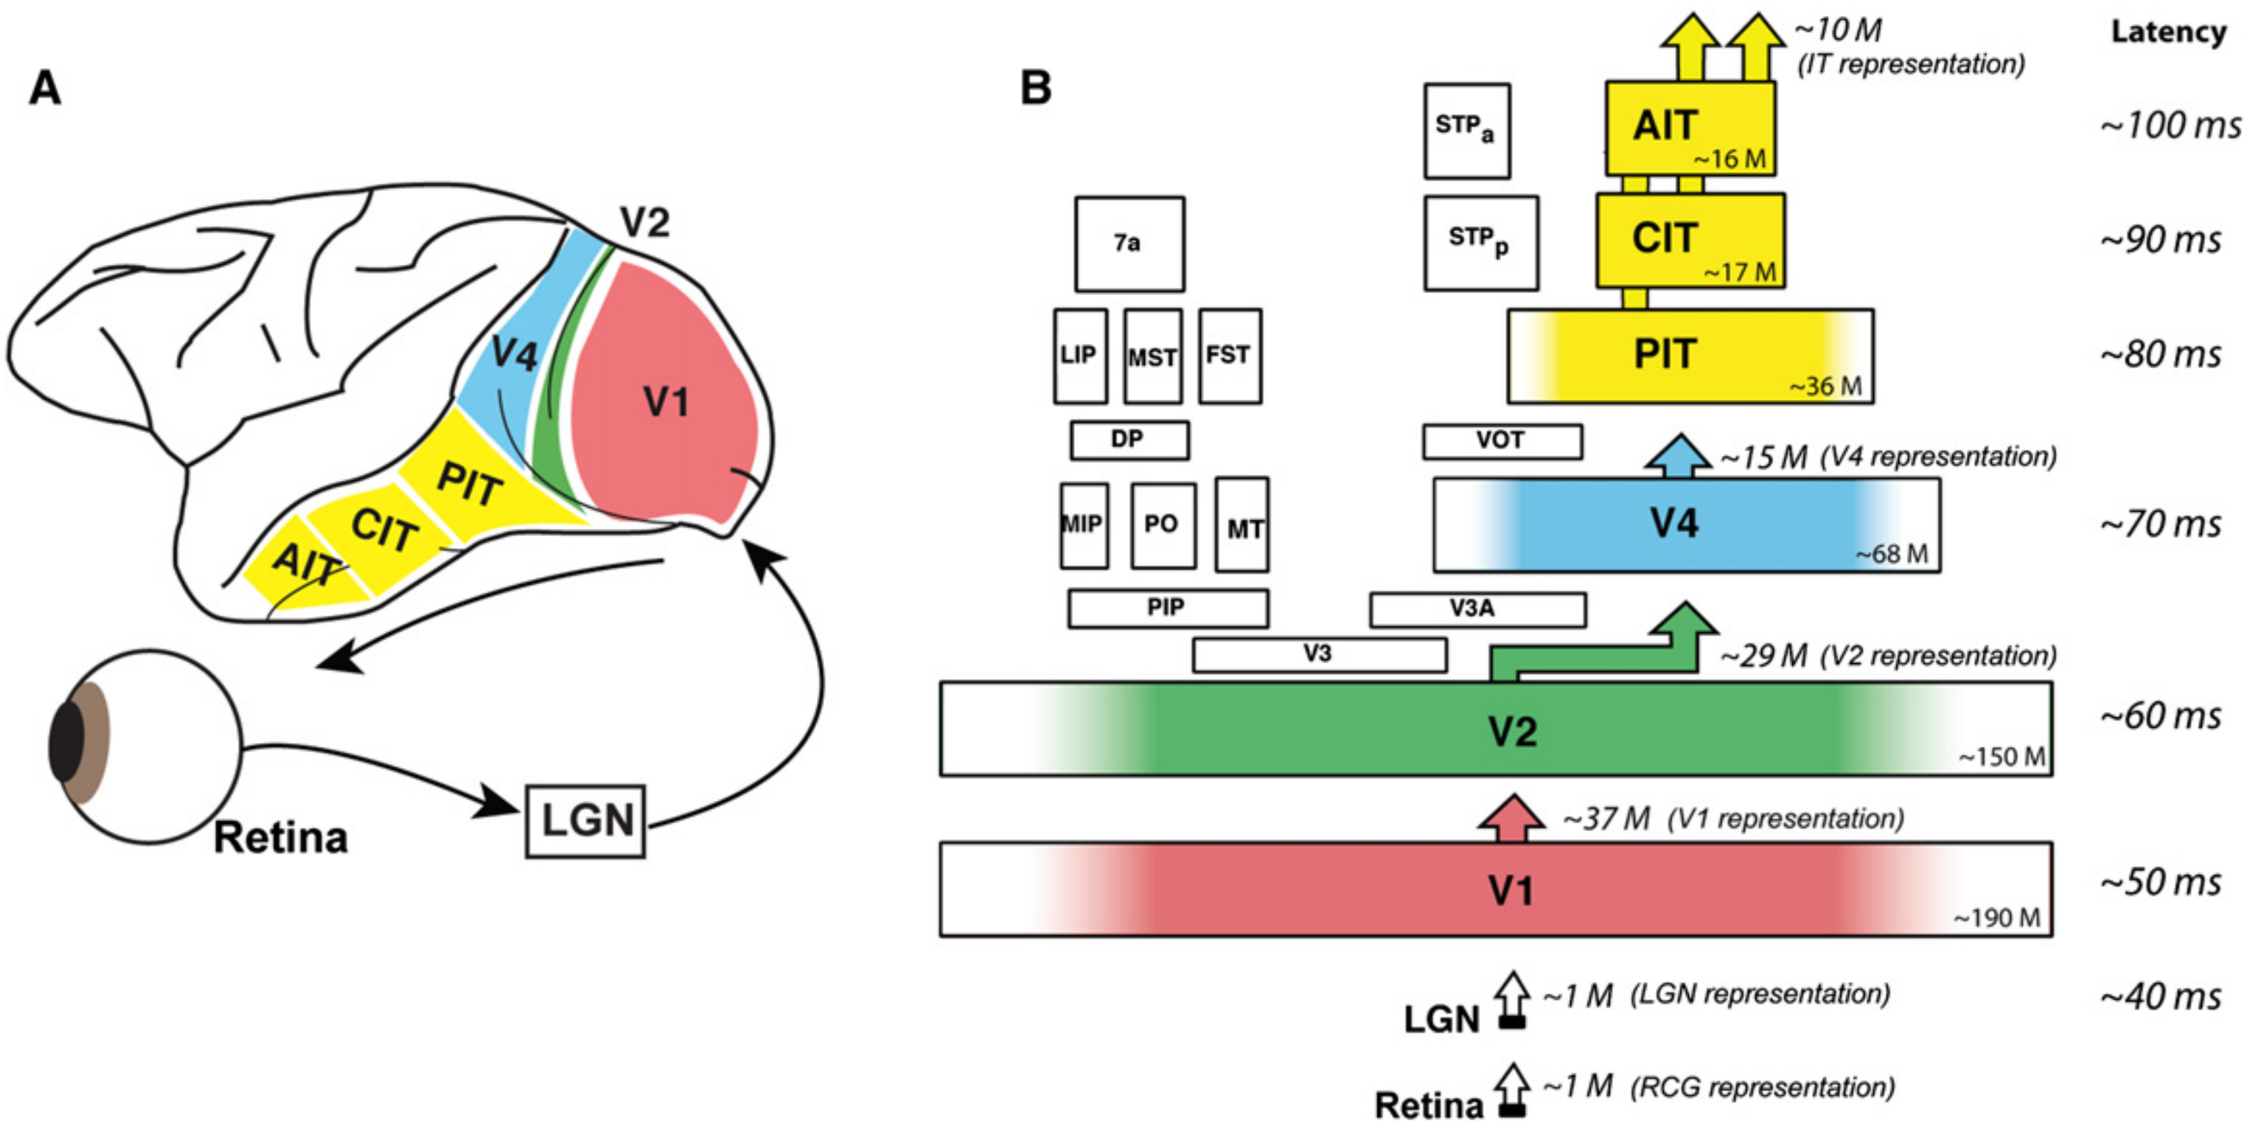
A
B
V2
V4
V1
PIT
CIT
AIT
Retina
LGN
Latency
~10 M
(IT representation)
STPa
AIT
~16 M
~100 ms
7a
STPp
CIT
~17 M
~90 ms
LIP
MST
FST
PIT
~36 M
~80 ms
DP
VOT
~15 M (V4 representation)
MIP
PO
MT
V4
~68 M
~70 ms
PIP
V3A
V3
~29 M (V2 representation)
V2
~150 M
~60 ms
~37 M (V1 representation)
V1
~190 M
~50 ms
LGN
~1 M (LGN representation)
~40 ms
Retina
~1 M (RCG representation)

Generatív/rekogniciós modell

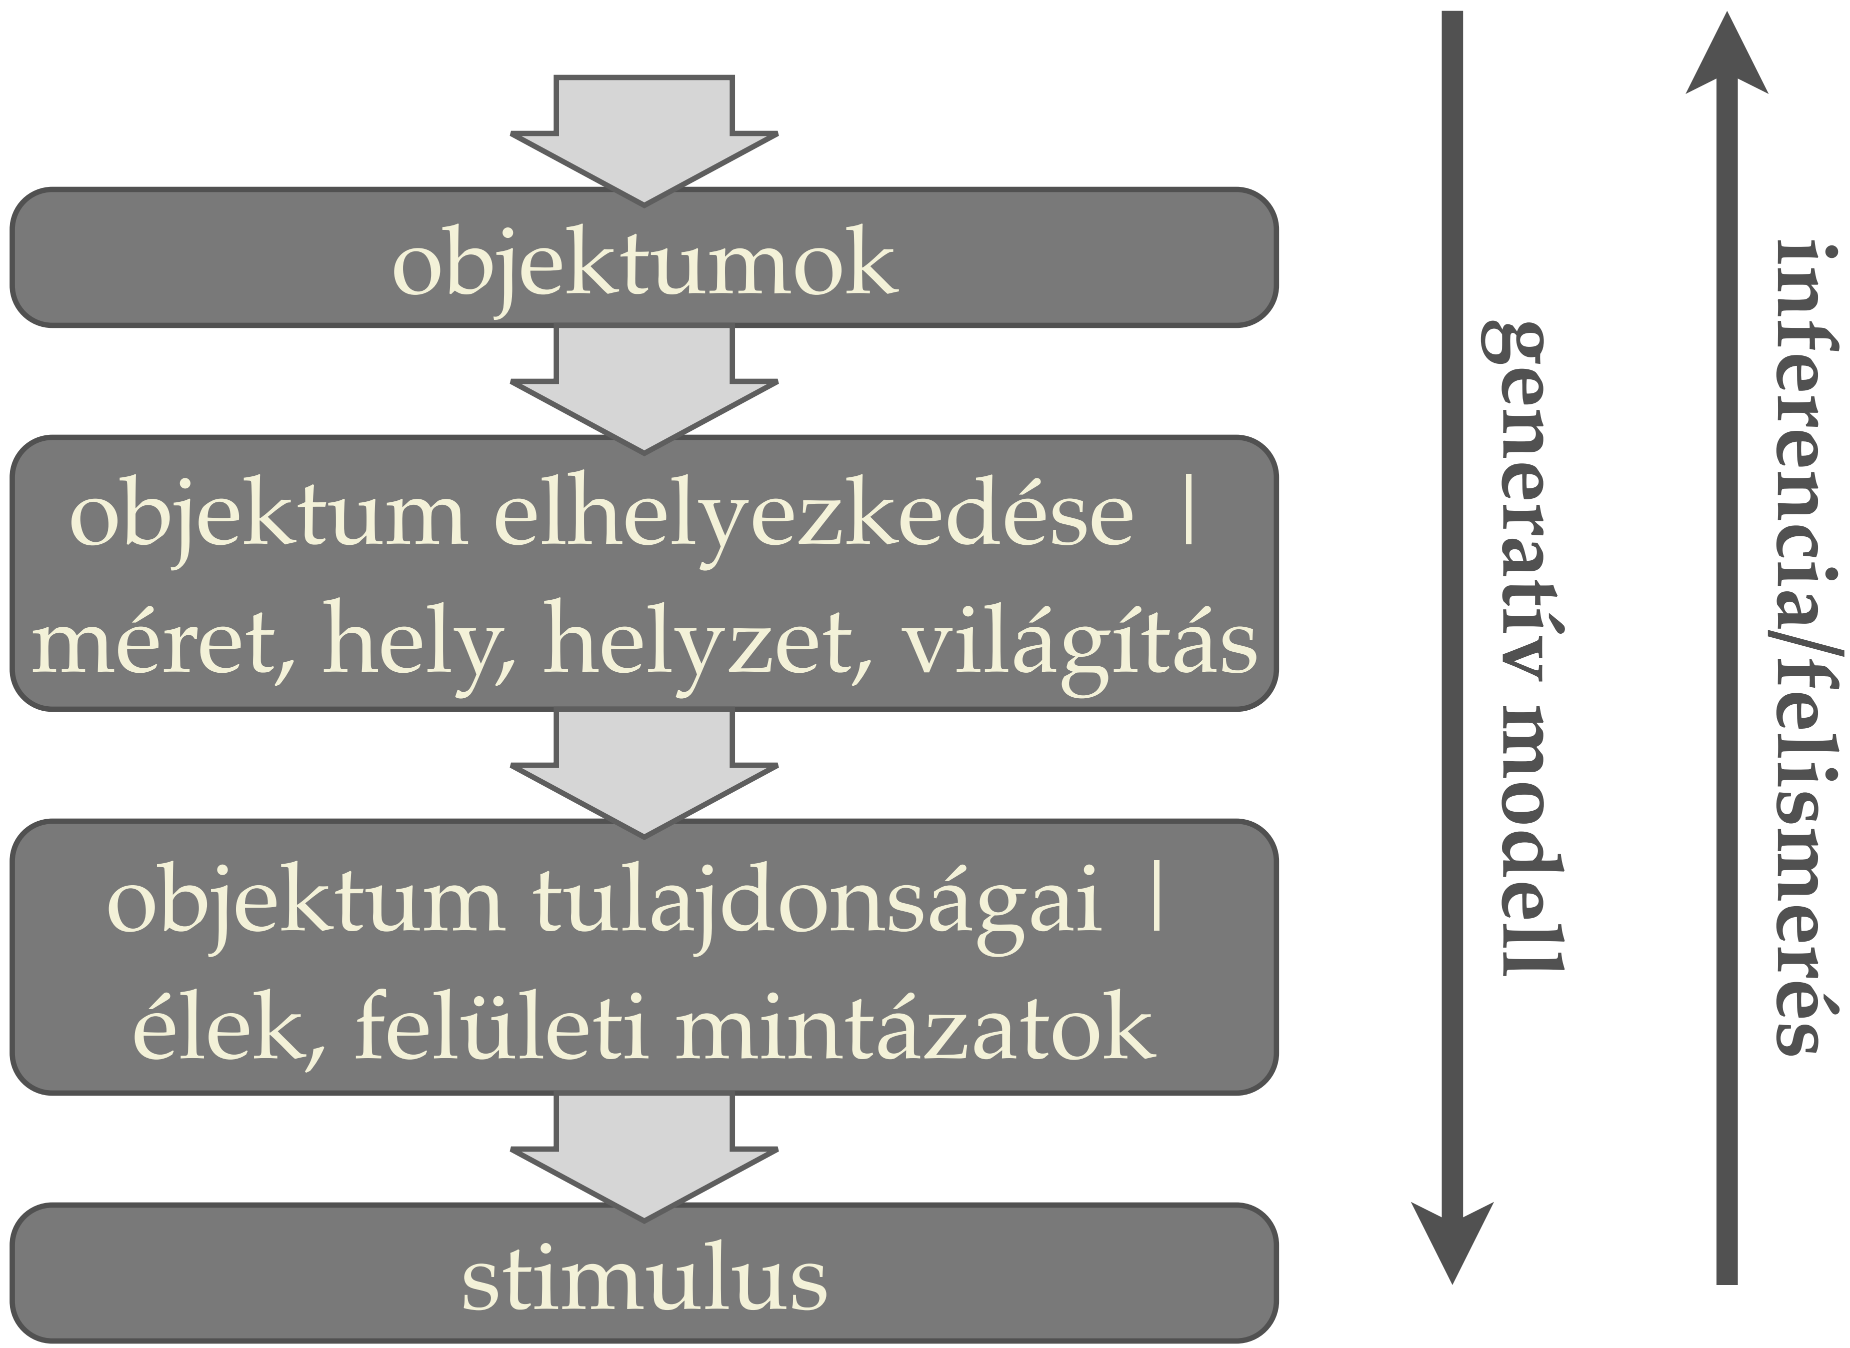

Generatív/rekogniciós modell

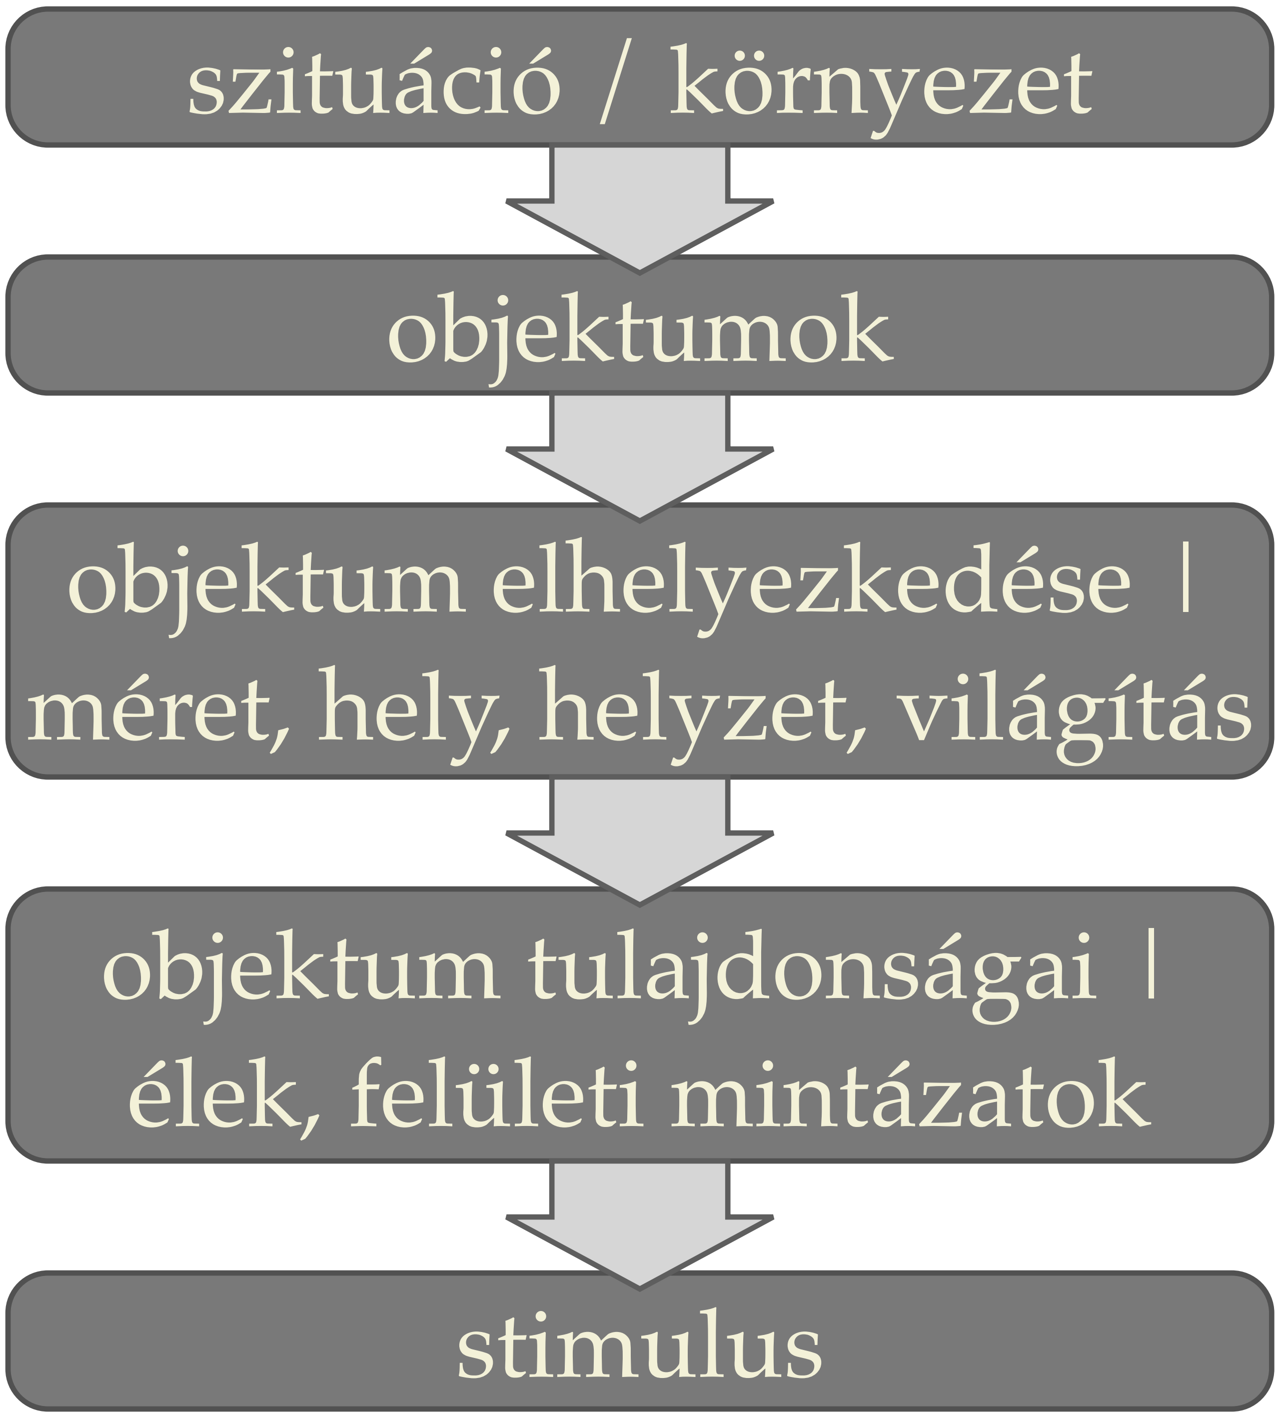

Modell definició -> generáció:
$P(x \mid z)$

Inferencia igénye -> rekogníció:
$P(z \mid x)$

Lineáris modellek

V1 receptive mezők:
- orientált
- sáváteresztő
- lokalizált

$$x = A\ z + eps$$

V1 stimulus-függés
- kontraszt invariancia
- extra-klasszikus receptív mezők

$$x = c\ (A\ z) + eps$$

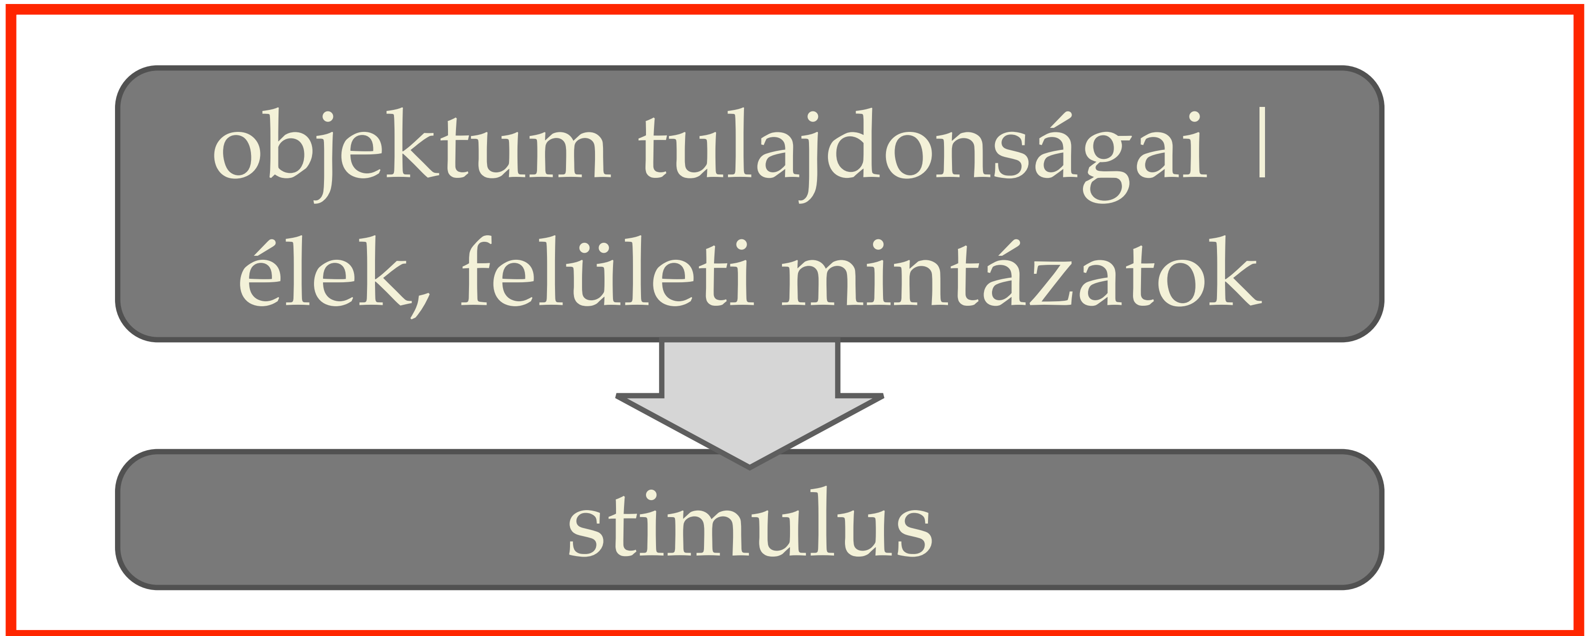

Scale Mixtures of Gaussians and the Statistics of Natural Images

Martin J. Wainwright
Stochastic Systems Group
Electrical Engineering & CS
MIT, Building 35-425
Cambridge, MA 02139
mjwain@mit.edu

Eero P. Simoncelli
Ctr. for Neural Science, and
Courant Inst. of Mathematical Sciences
New York University
New York, NY 10012
eero.simoncelli@nyu.edu

Abstract

The statistics of photographic images, when represented using multiscale (wavelet) bases, exhibit two striking types of non-Gaussian behavior. First, the marginal densities of the coefficients have extended heavy tails. Second, the joint densities exhibit variance dependencies not captured by second-order models. We examine properties of the class of Gaussian scale mixtures, and show that these densities can accurately characterize both the marginal and joint distributions of natural image wavelet coefficients. This class of model suggests a Markov structure, in which wavelet coefficients are linked by hidden scaling variables corresponding to local image structure. We derive an estimator for these hidden variables, and show that a nonlinear "normalization" procedure can be used to Gaussianize the coefficients.

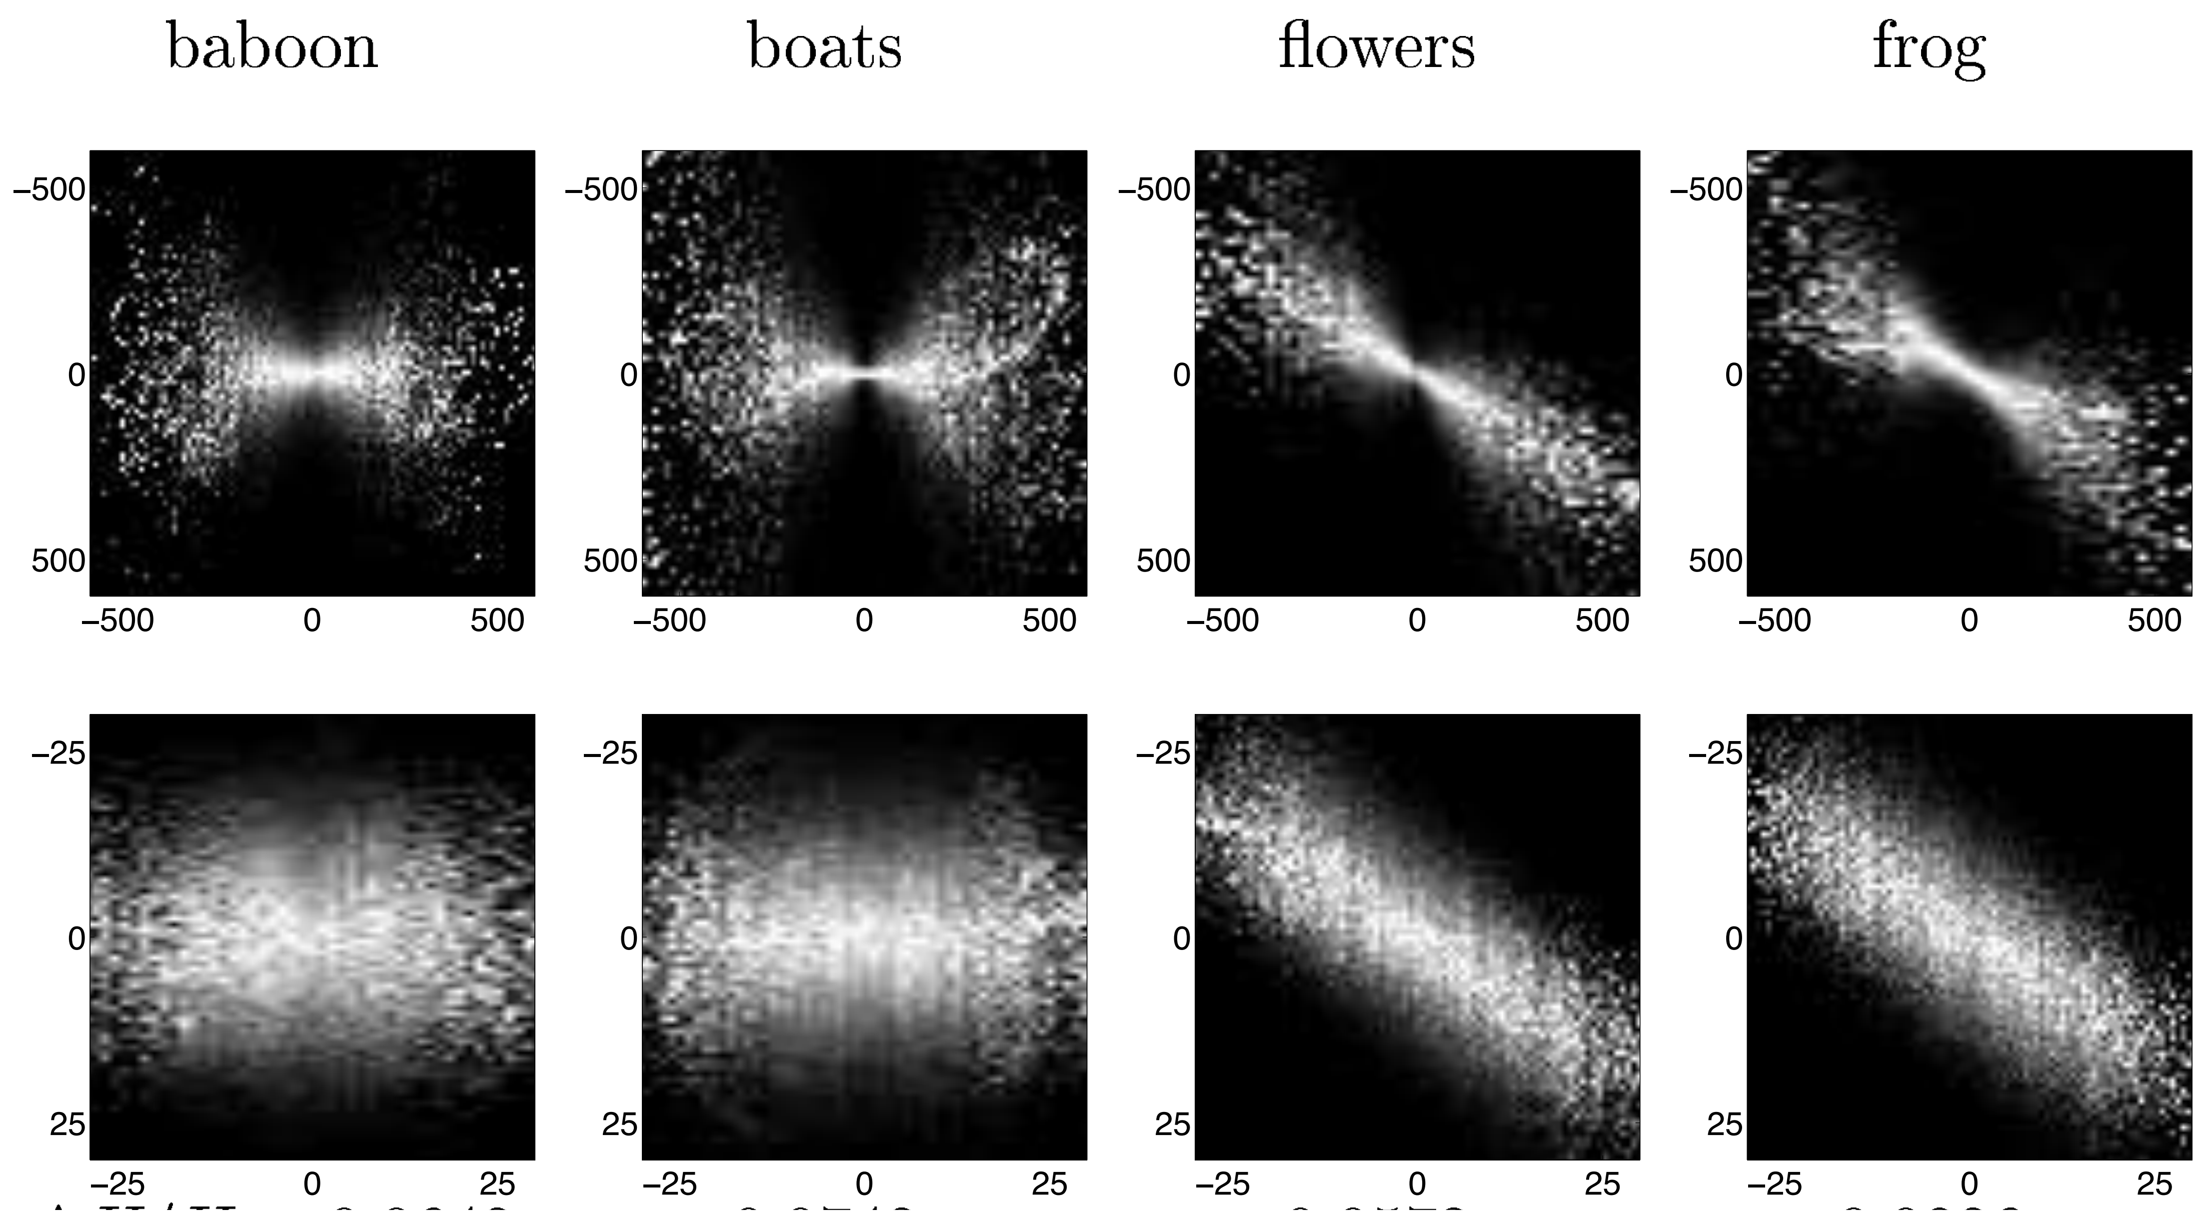
baboon
boats
flowers
frog
−500
0
500
−25
0
25

Gauging V2 responses so far

- gratings
- contours
- angles
- other forms of second order stats
- border ownership

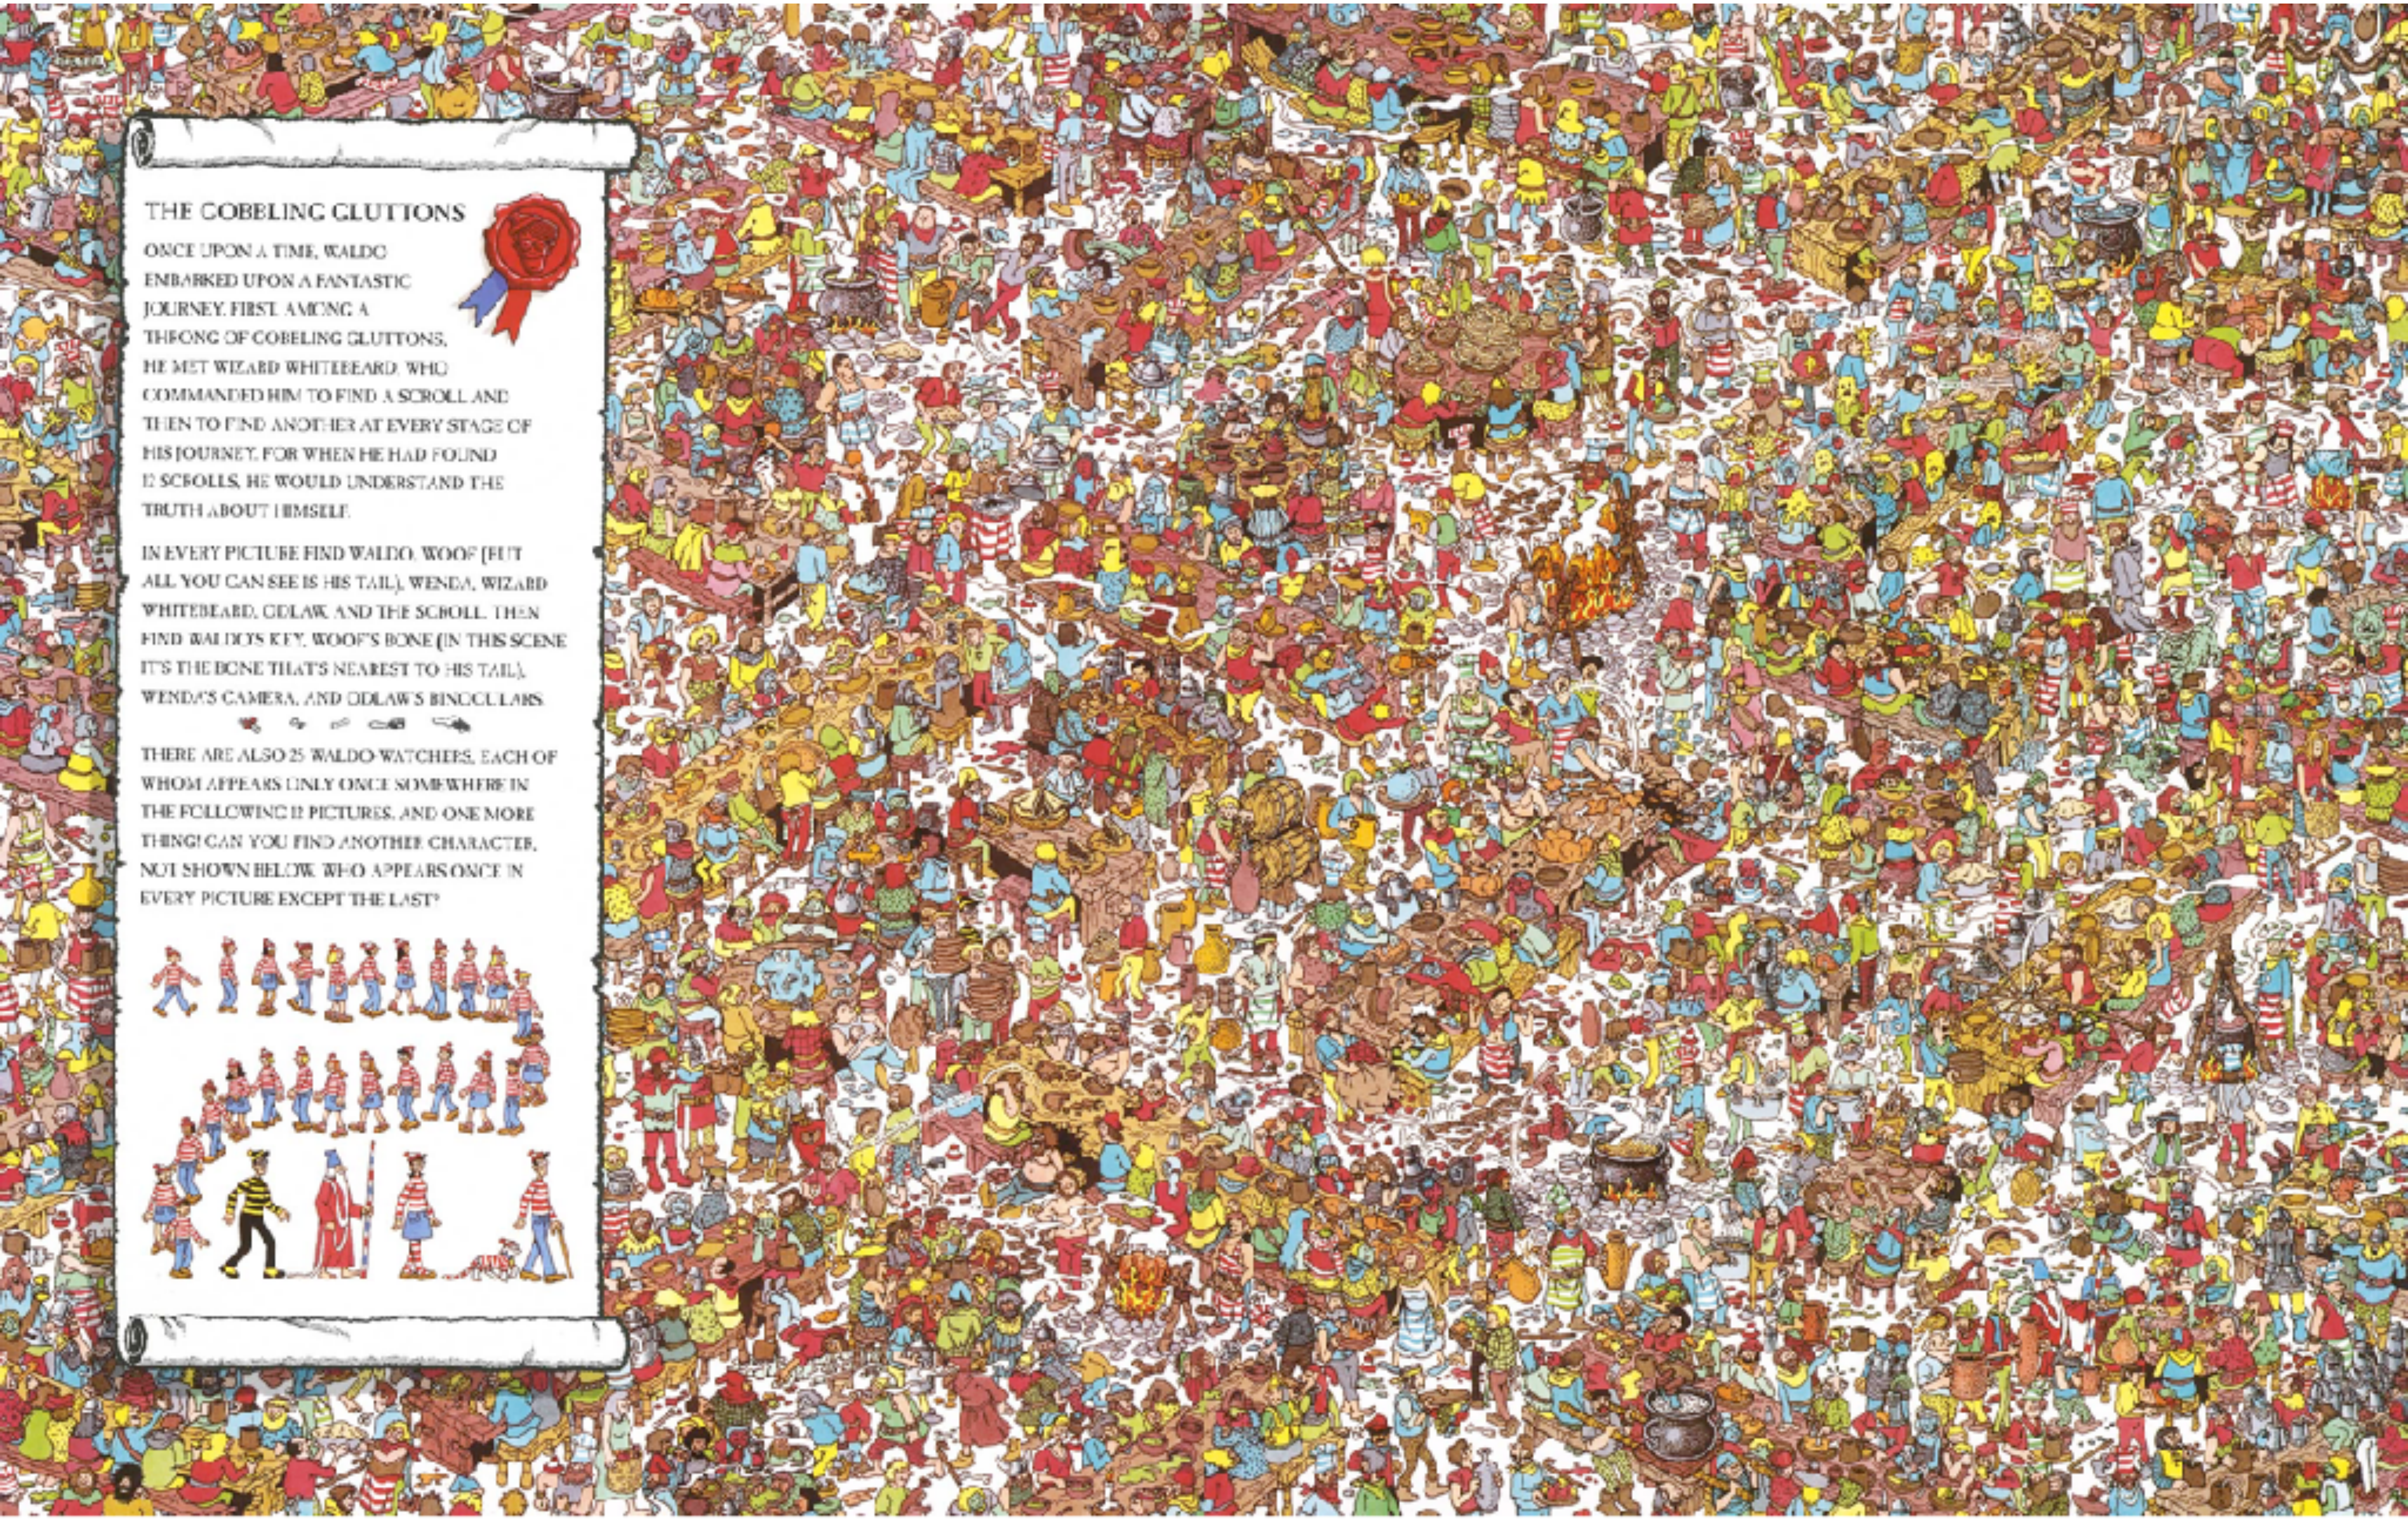
THE GOBBLING GLUTTONS
ONCE UPON A TIME, WALDO EMBARKED UPON A FANTASTIC JOURNEY. FIRST, AMONG A THRONG OF GOBBLING GLUTTONS, HE MET WIZARD WHITEBEARD, WHO COMMANDED HIM TO FIND A SCROLL AND THEN TO FIND ANOTHER AT EVERY STAGE OF HIS JOURNEY. FOR WHEN HE HAD FOUND 12 SCROLLS, HE WOULD UNDERSTAND THE TRUTH ABOUT HIMSELF.
IN EVERY PICTURE FIND WALDO, WOOF (BUT ALL YOU CAN SEE IS HIS TAIL), WENDA, WIZARD WHITEBEARD, ODLAW, AND THE SCROLL. THEN FIND WALDO'S KEY, WOOF'S BONE (IN THIS SCENE IT'S THE BONE THAT'S NEAREST TO HIS TAIL), WENDA'S CAMERA, AND ODLAW'S BINOCULARS.
THERE ARE ALSO 25 WALDO-WATCHERS, EACH OF WHOM APPEARS ONLY ONCE SOMEWHERE IN THE FOLLOWING 12 PICTURES. AND ONE MORE THING! CAN YOU FIND ANOTHER CHARACTER, NOT SHOWN BELOW, WHO APPEARS ONCE IN EVERY PICTURE EXCEPT THE LAST?

International Journal of Computer Vision 40(1), 49–71, 2000

A Parametric Texture Model Based on Joint Statistics of Complex Wavelet Coefficients

JAVIER PORTILLA AND EERO P. SIMONCELLI
Center for Neural Science, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10003, USA

Learning about the stats of an image

- Registering the responses of linear filters (simple cells)
- Registering the responses of energy filters (complex cells)
- Marginal statistics: variance, kurtosis, skewness

- Registering correlations between orientations
- Registering correlations between spatial frequencies
- Registering correlations across positions

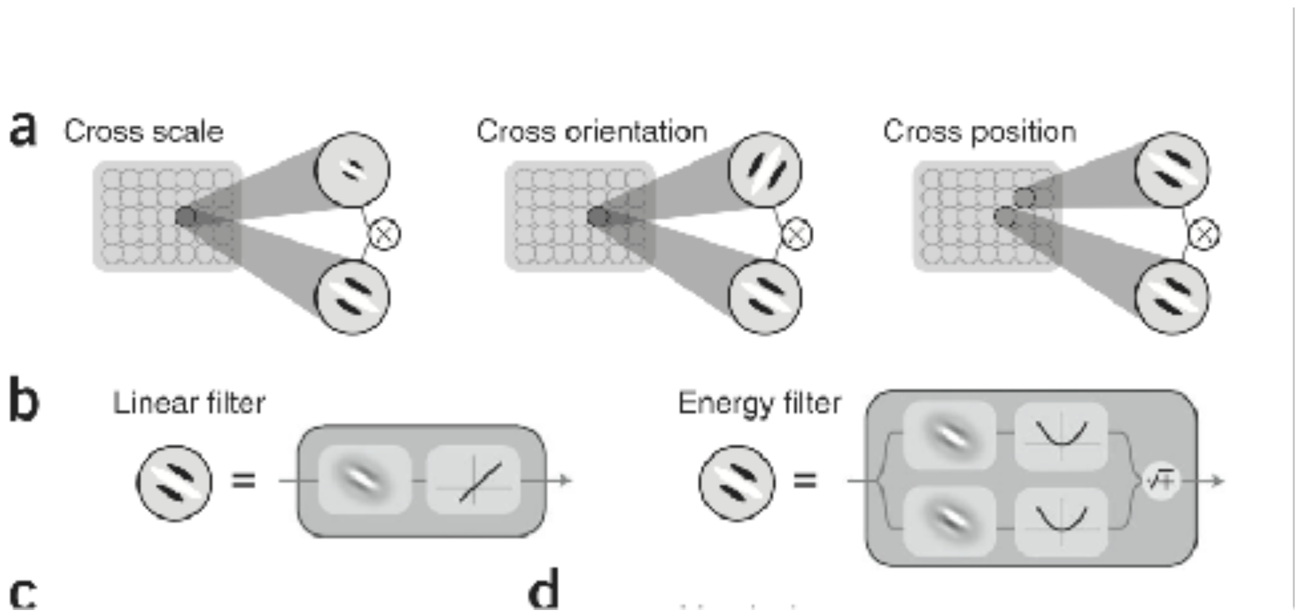

Synthetic textures

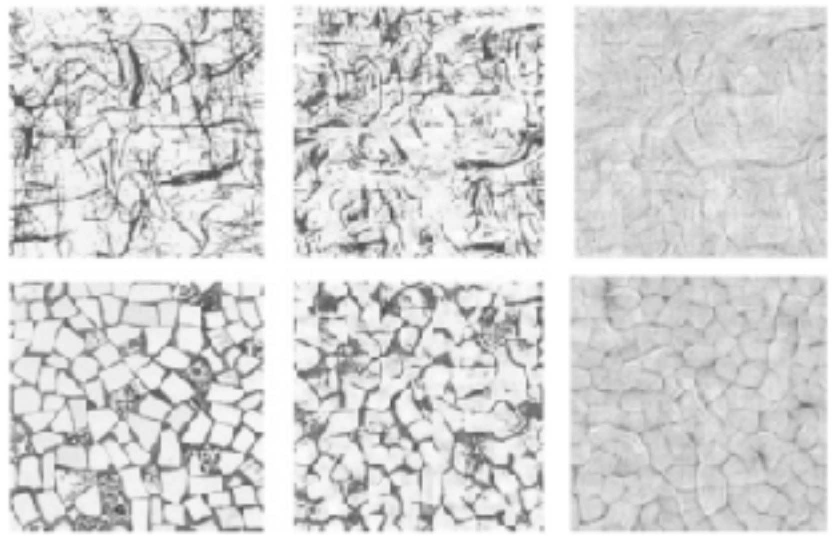

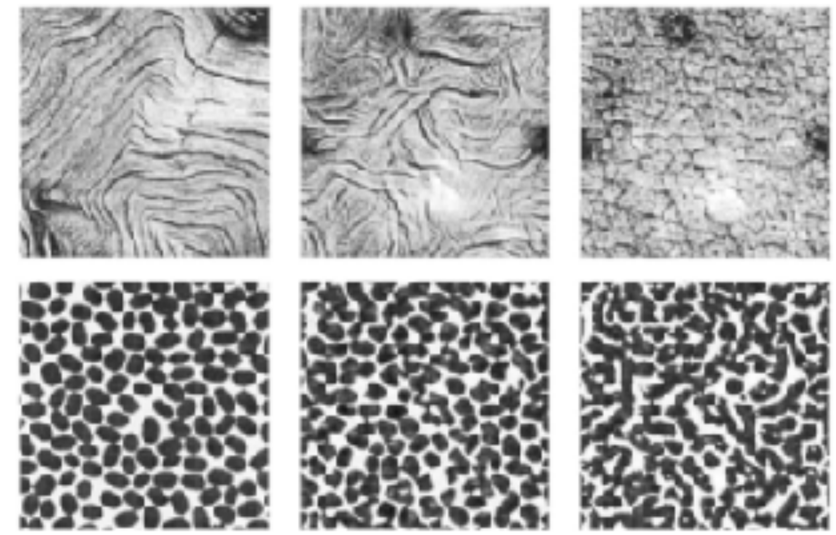

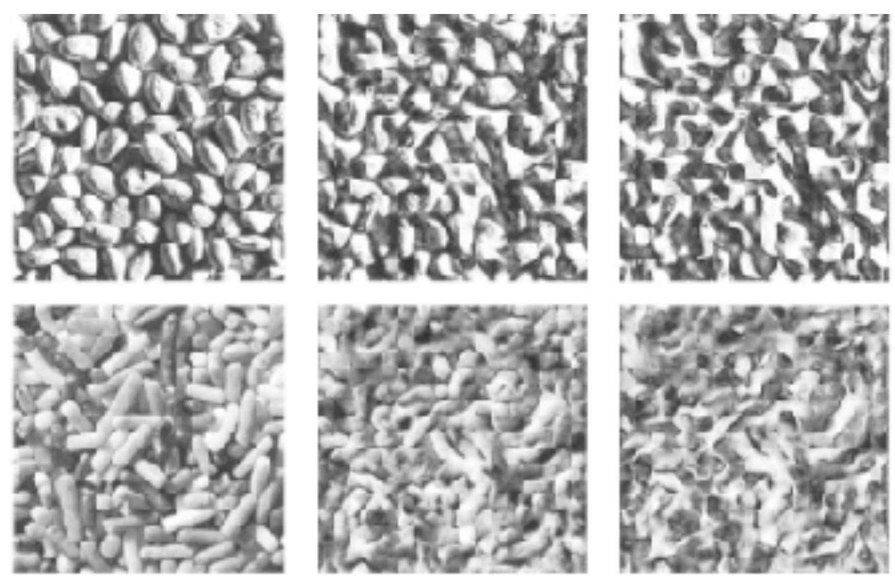

nature neuroscience

Metamers of the ventral stream

Jeremy Freeman[1] & Eero P Simoncelli[1–3]

The human capacity to recognize complex visual patterns emerges in a sequence of brain areas known as the ventral stream, beginning with primary visual cortex (V1). We developed a population model for mid-ventral processing, in which nonlinear combinations of V1 responses are averaged in receptive fields that grow with eccentricity. To test the model, we generated novel forms of visual metamers, stimuli that differ physically but look the same. We developed a behavioral protocol that uses metameric stimuli to estimate the receptive field sizes in which the model features are represented. Because receptive field sizes change along the ventral stream, our behavioral results can identify the visual area corresponding to the representation. Measurements in human observers implicate visual area V2, providing a new functional account of neurons in this area. The model also explains deficits of peripheral vision known as crowding, and provides a quantitative framework for assessing the capabilities and limitations of everyday vision.

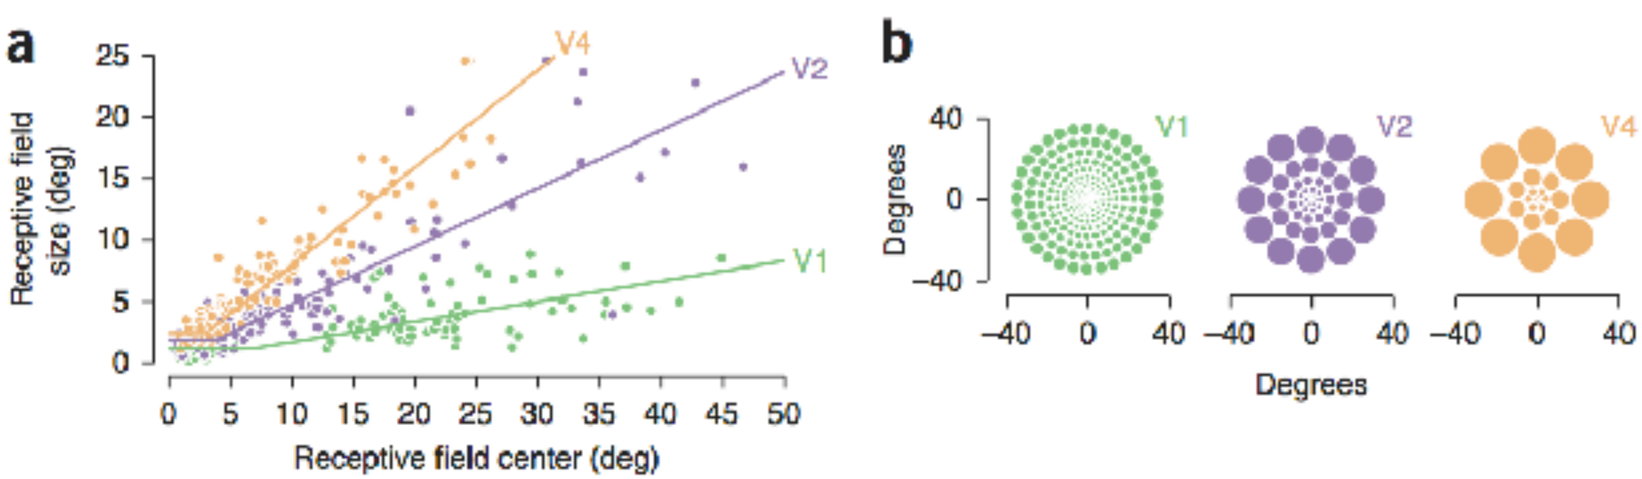

nature neuro

Meta

Jeremy F

The human ... , beginning
with prima ... ns of V1
responses ... metamers,
stimuli tha ... the
receptive f ... eam, our
behavioral ... cate visual
area V2, p ... own as
crowding, ...

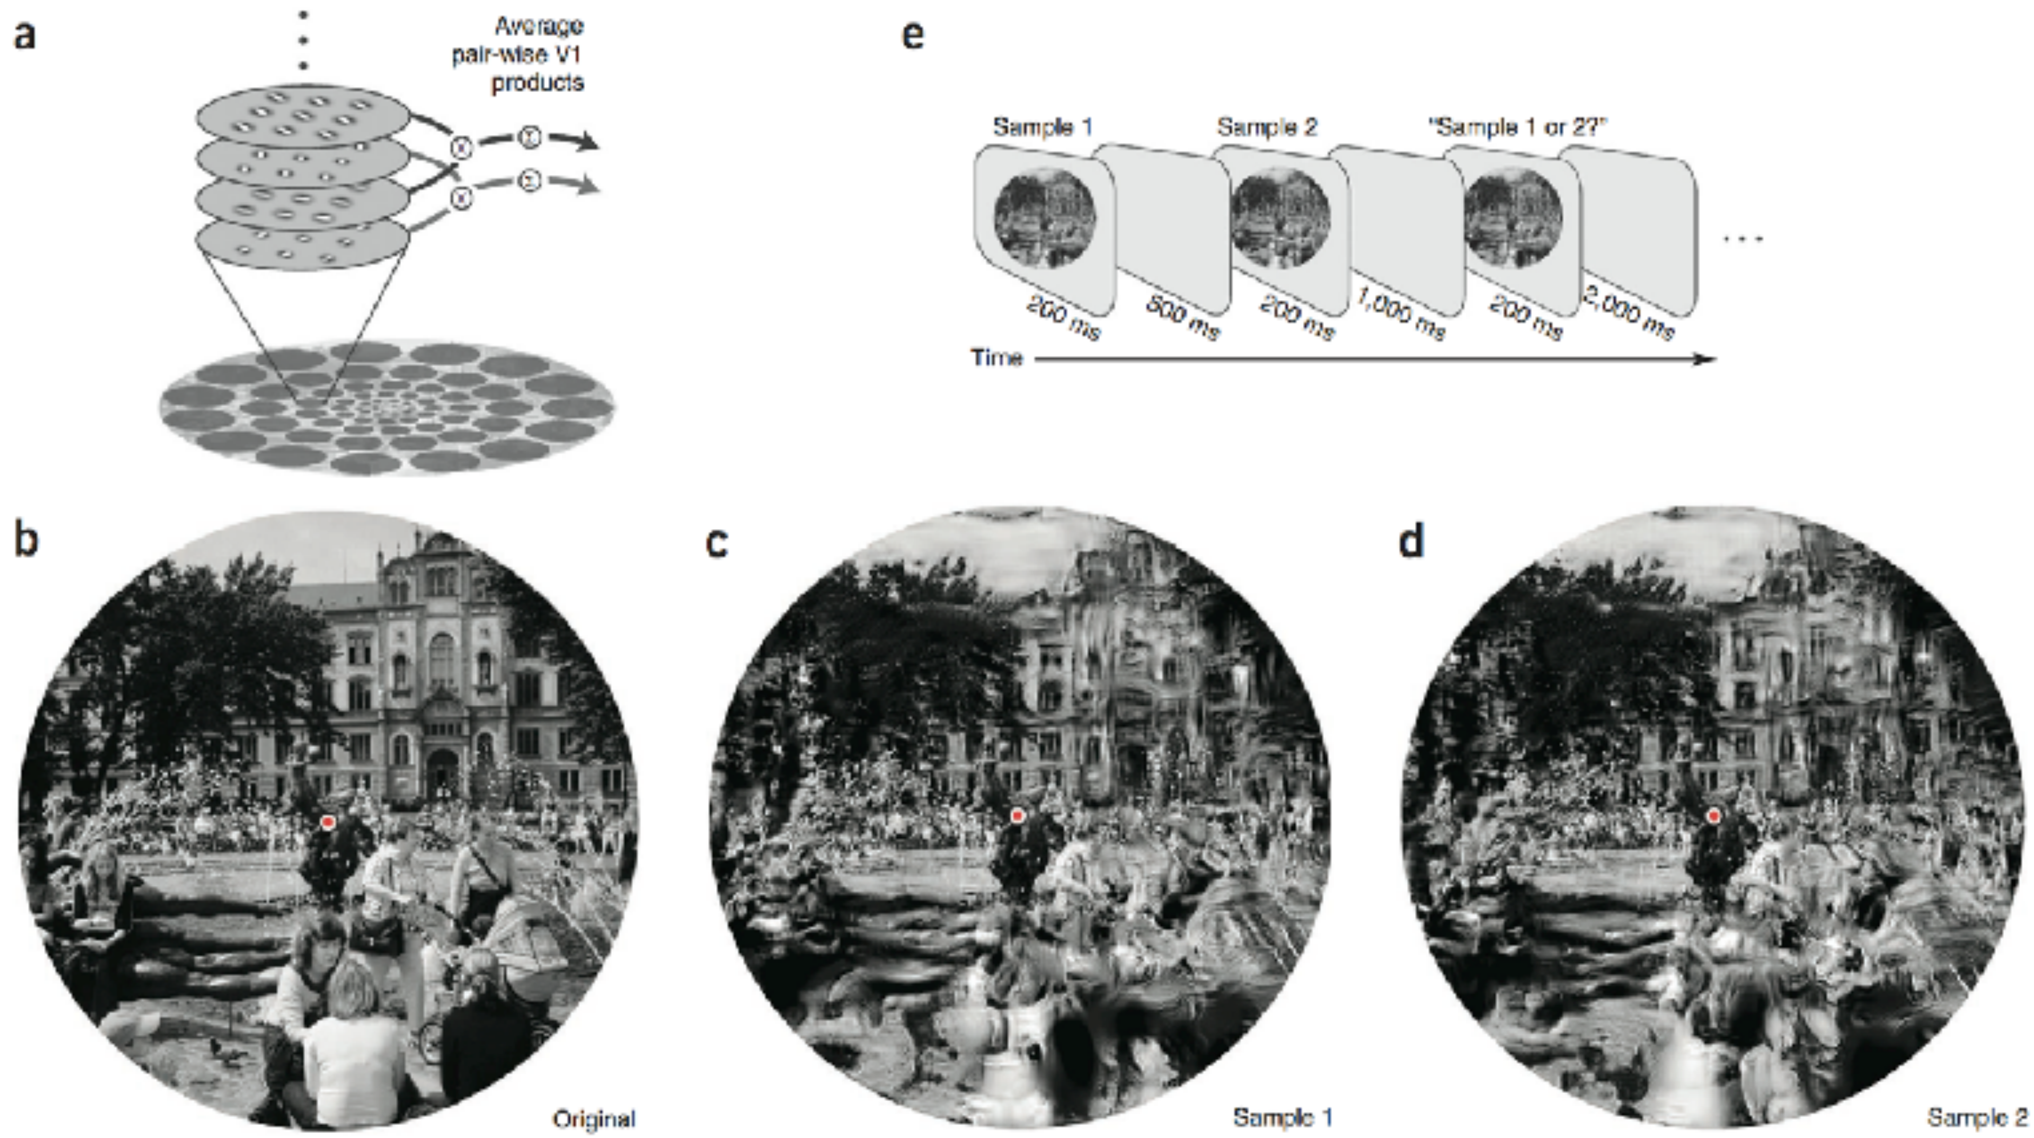

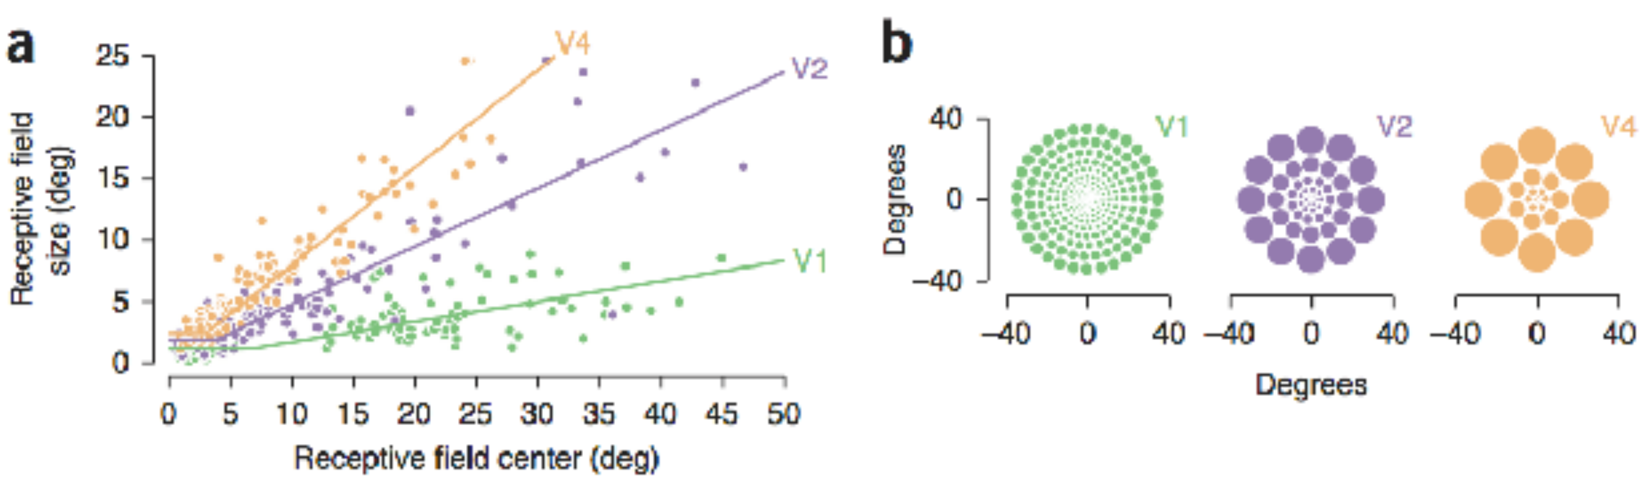

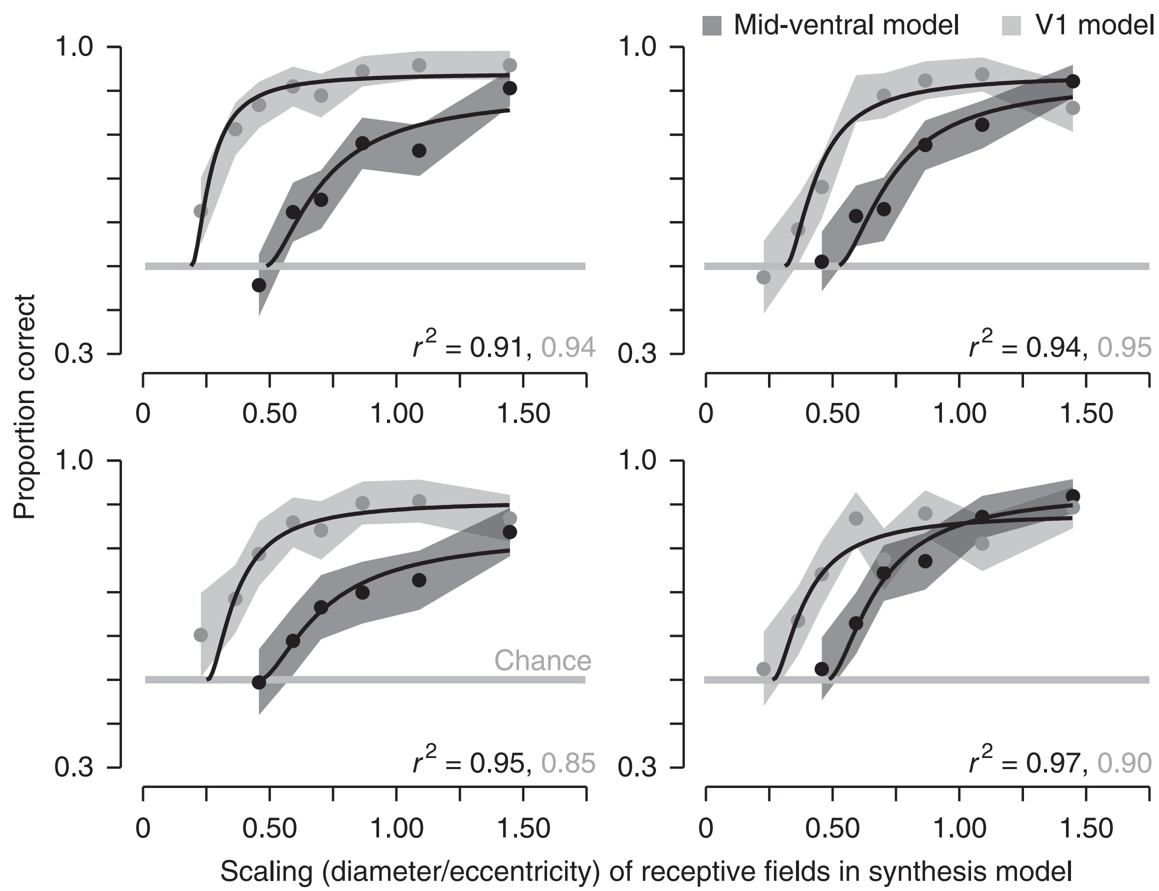
Mid-ventral model
V1 model
Proportion correct
1.0
0.3
r^2 = 0.91, 0.94
r^2 = 0.94, 0.95
r^2 = 0.95, 0.85
r^2 = 0.97, 0.90
Chance
0
0.50
1.00
1.50
Scaling (diameter/eccentricity) of receptive fields in synthesis model

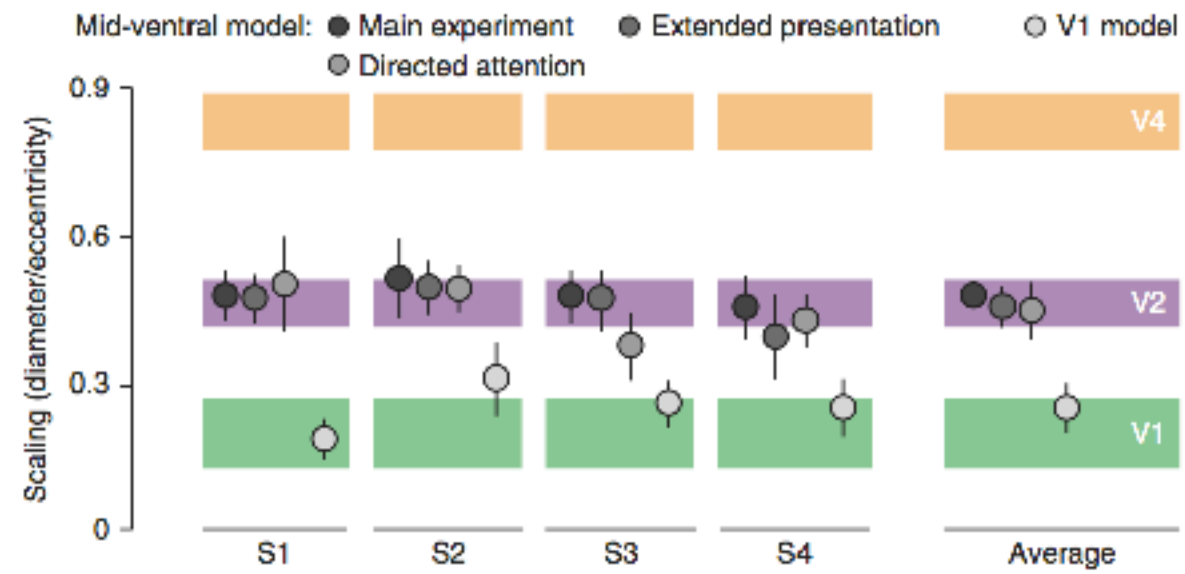
Mid-ventral model: Main experiment
Extended presentation
V1 model
Directed attention
Scaling (diameter/eccentricity)
0.9
0.6
0.3
0
V4
V2
V1
S1
S2
S3
S4
Average

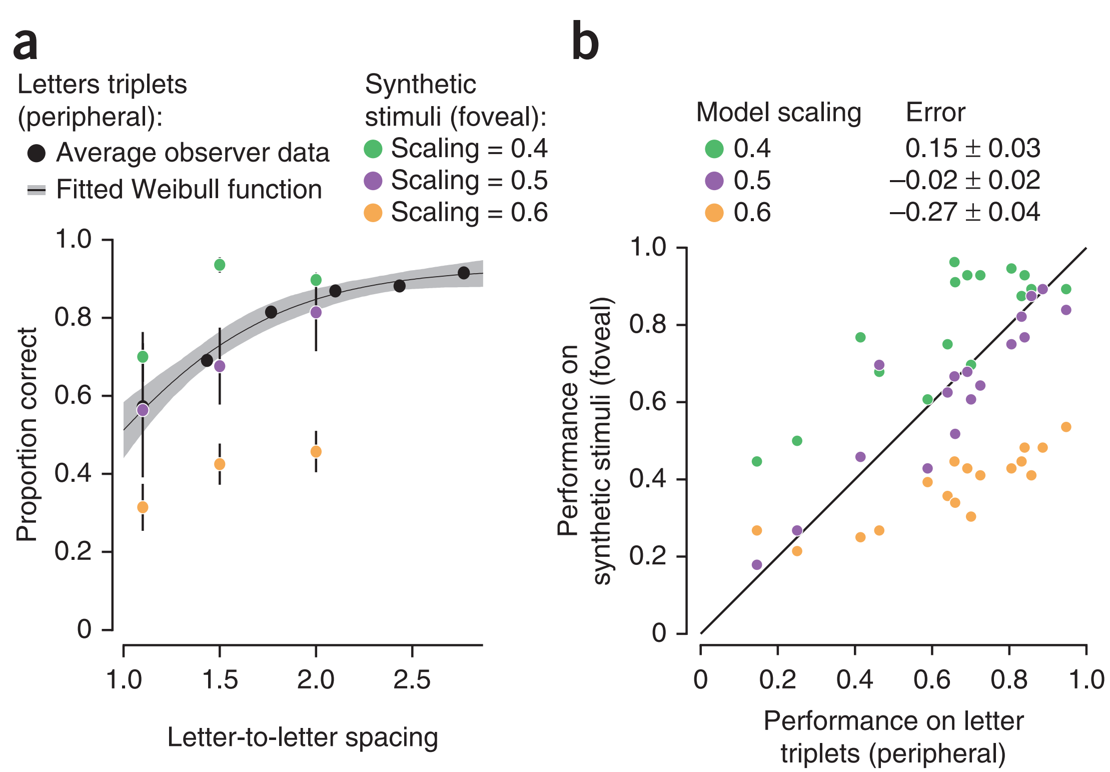
a
Letters triplets (peripheral):
Average observer data
Fitted Weibull function
Synthetic stimuli (foveal):
Scaling = 0.4
Scaling = 0.5
Scaling = 0.6
Proportion correct
Letter-to-letter spacing
b
Model scaling
Error
0.4
0.15 ± 0.03
0.5
−0.02 ± 0.02
0.6
−0.27 ± 0.04
Performance on synthetic stimuli (foveal)
Performance on letter triplets (peripheral)

nature neuroscience

A functional and perceptual signature of the second visual area in primates

Jeremy Freeman[1,5,7], Corey M Ziemba[1,5], David J Heeger[1,2], Eero P Simoncelli[1–4,6] & J Anthony Movshon[1,2,6]

There is no generally accepted account of the function of the second visual cortical area (V2), partly because no simple response properties robustly distinguish V2 neurons from those in primary visual cortex (V1). We constructed synthetic stimuli replicating the higher-order statistical dependencies found in natural texture images and used them to stimulate macaque V1 and V2 neurons. Most V2 cells responded more vigorously to these textures than to control stimuli lacking naturalistic structure; V1 cells did not. Functional magnetic resonance imaging (fMRI) measurements in humans revealed differences between V1 and V2 that paralleled the neuronal measurements. The ability of human observers to detect naturalistic structure in different types of texture was well predicted by the strength of neuronal and fMRI responses in V2 but not in V1. Together, these results reveal a particular functional role for V2 in the representation of natural image structure.

Synthetising images

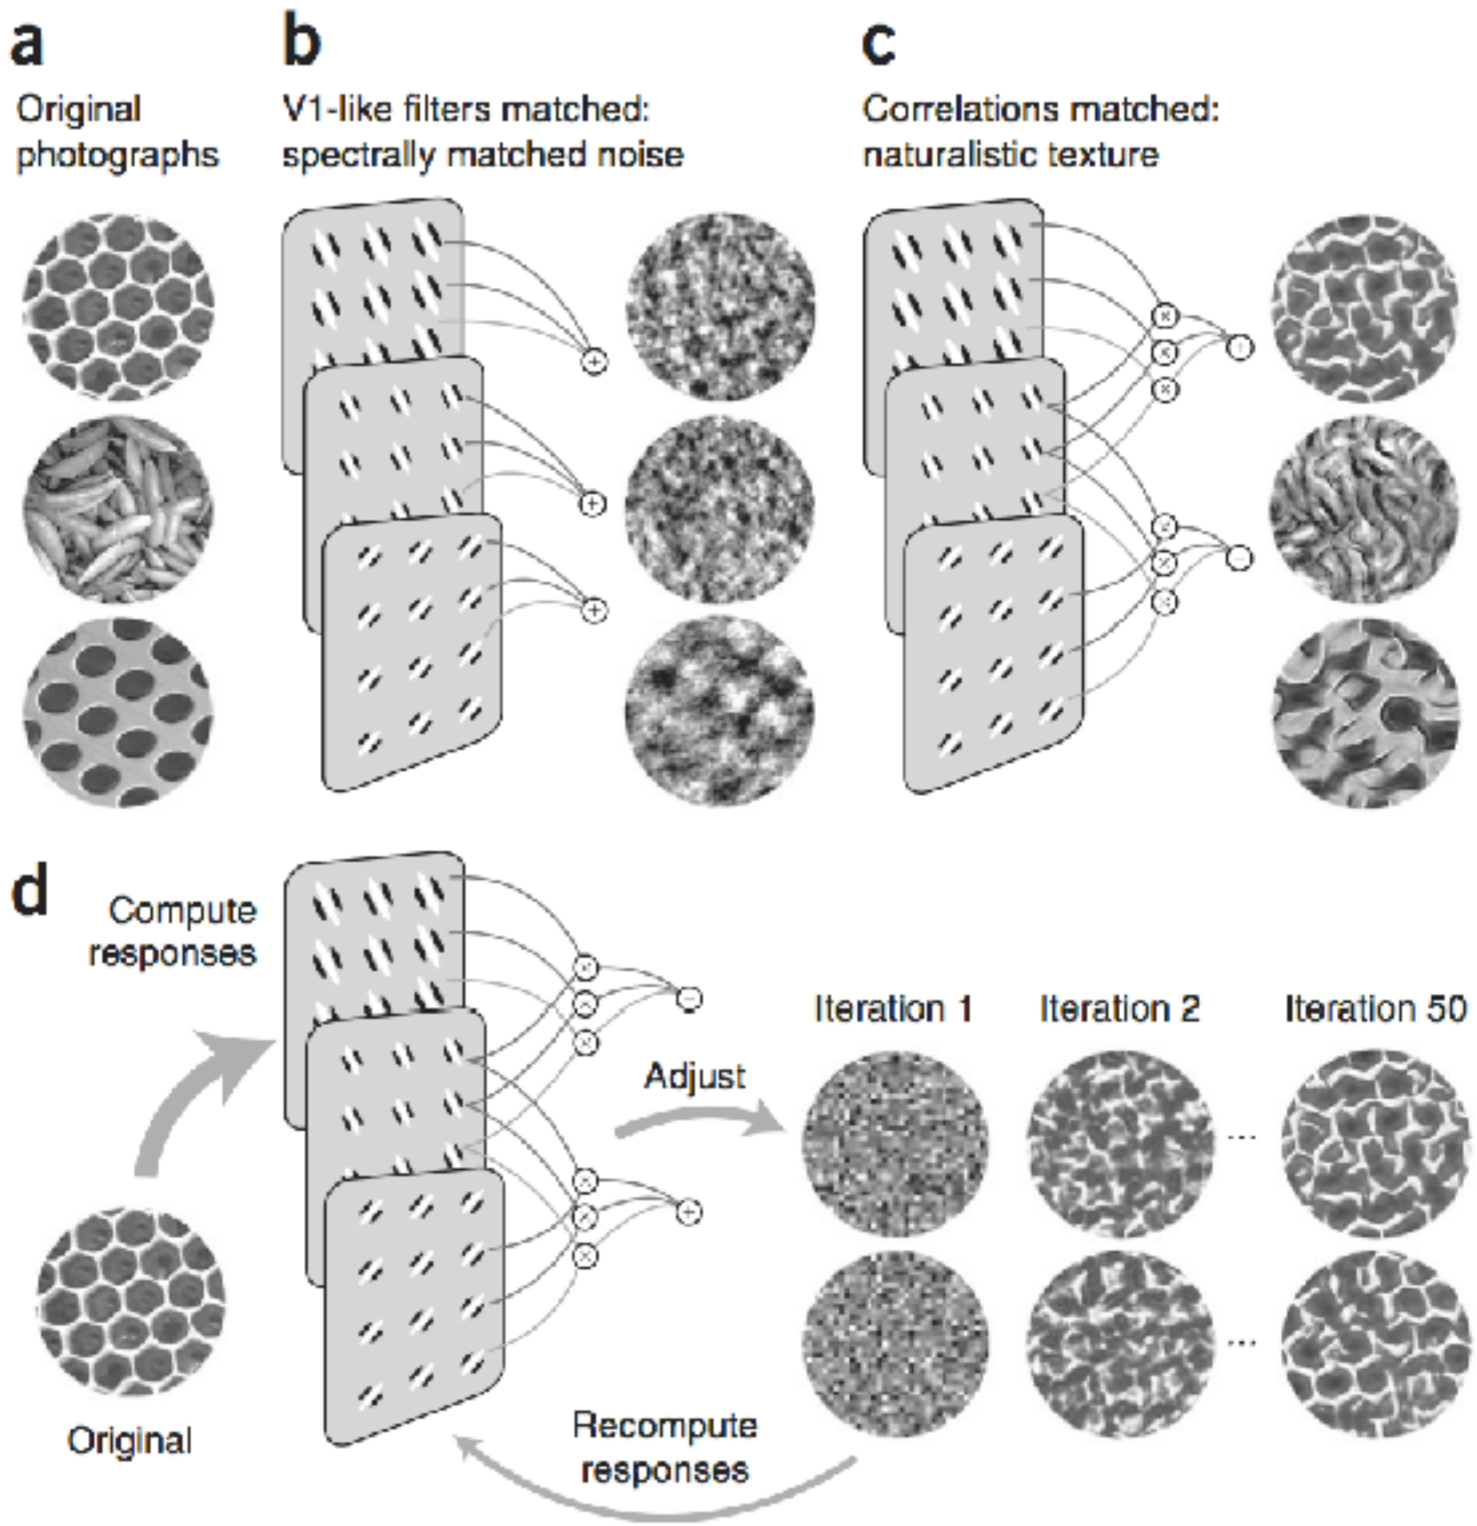

V2 responses to Portilla textures

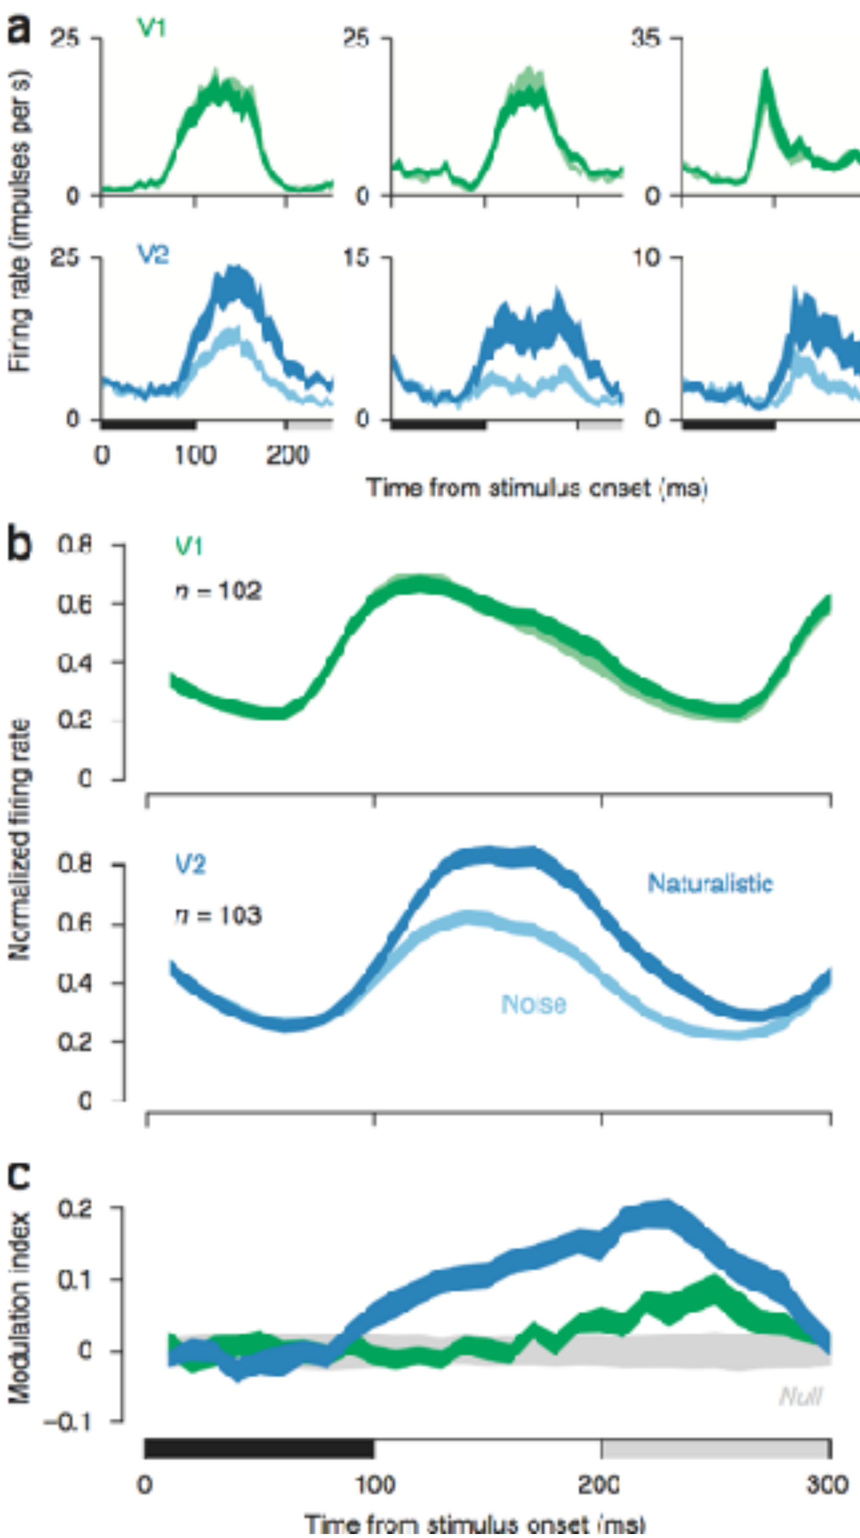

V2 responses to Portilla textures

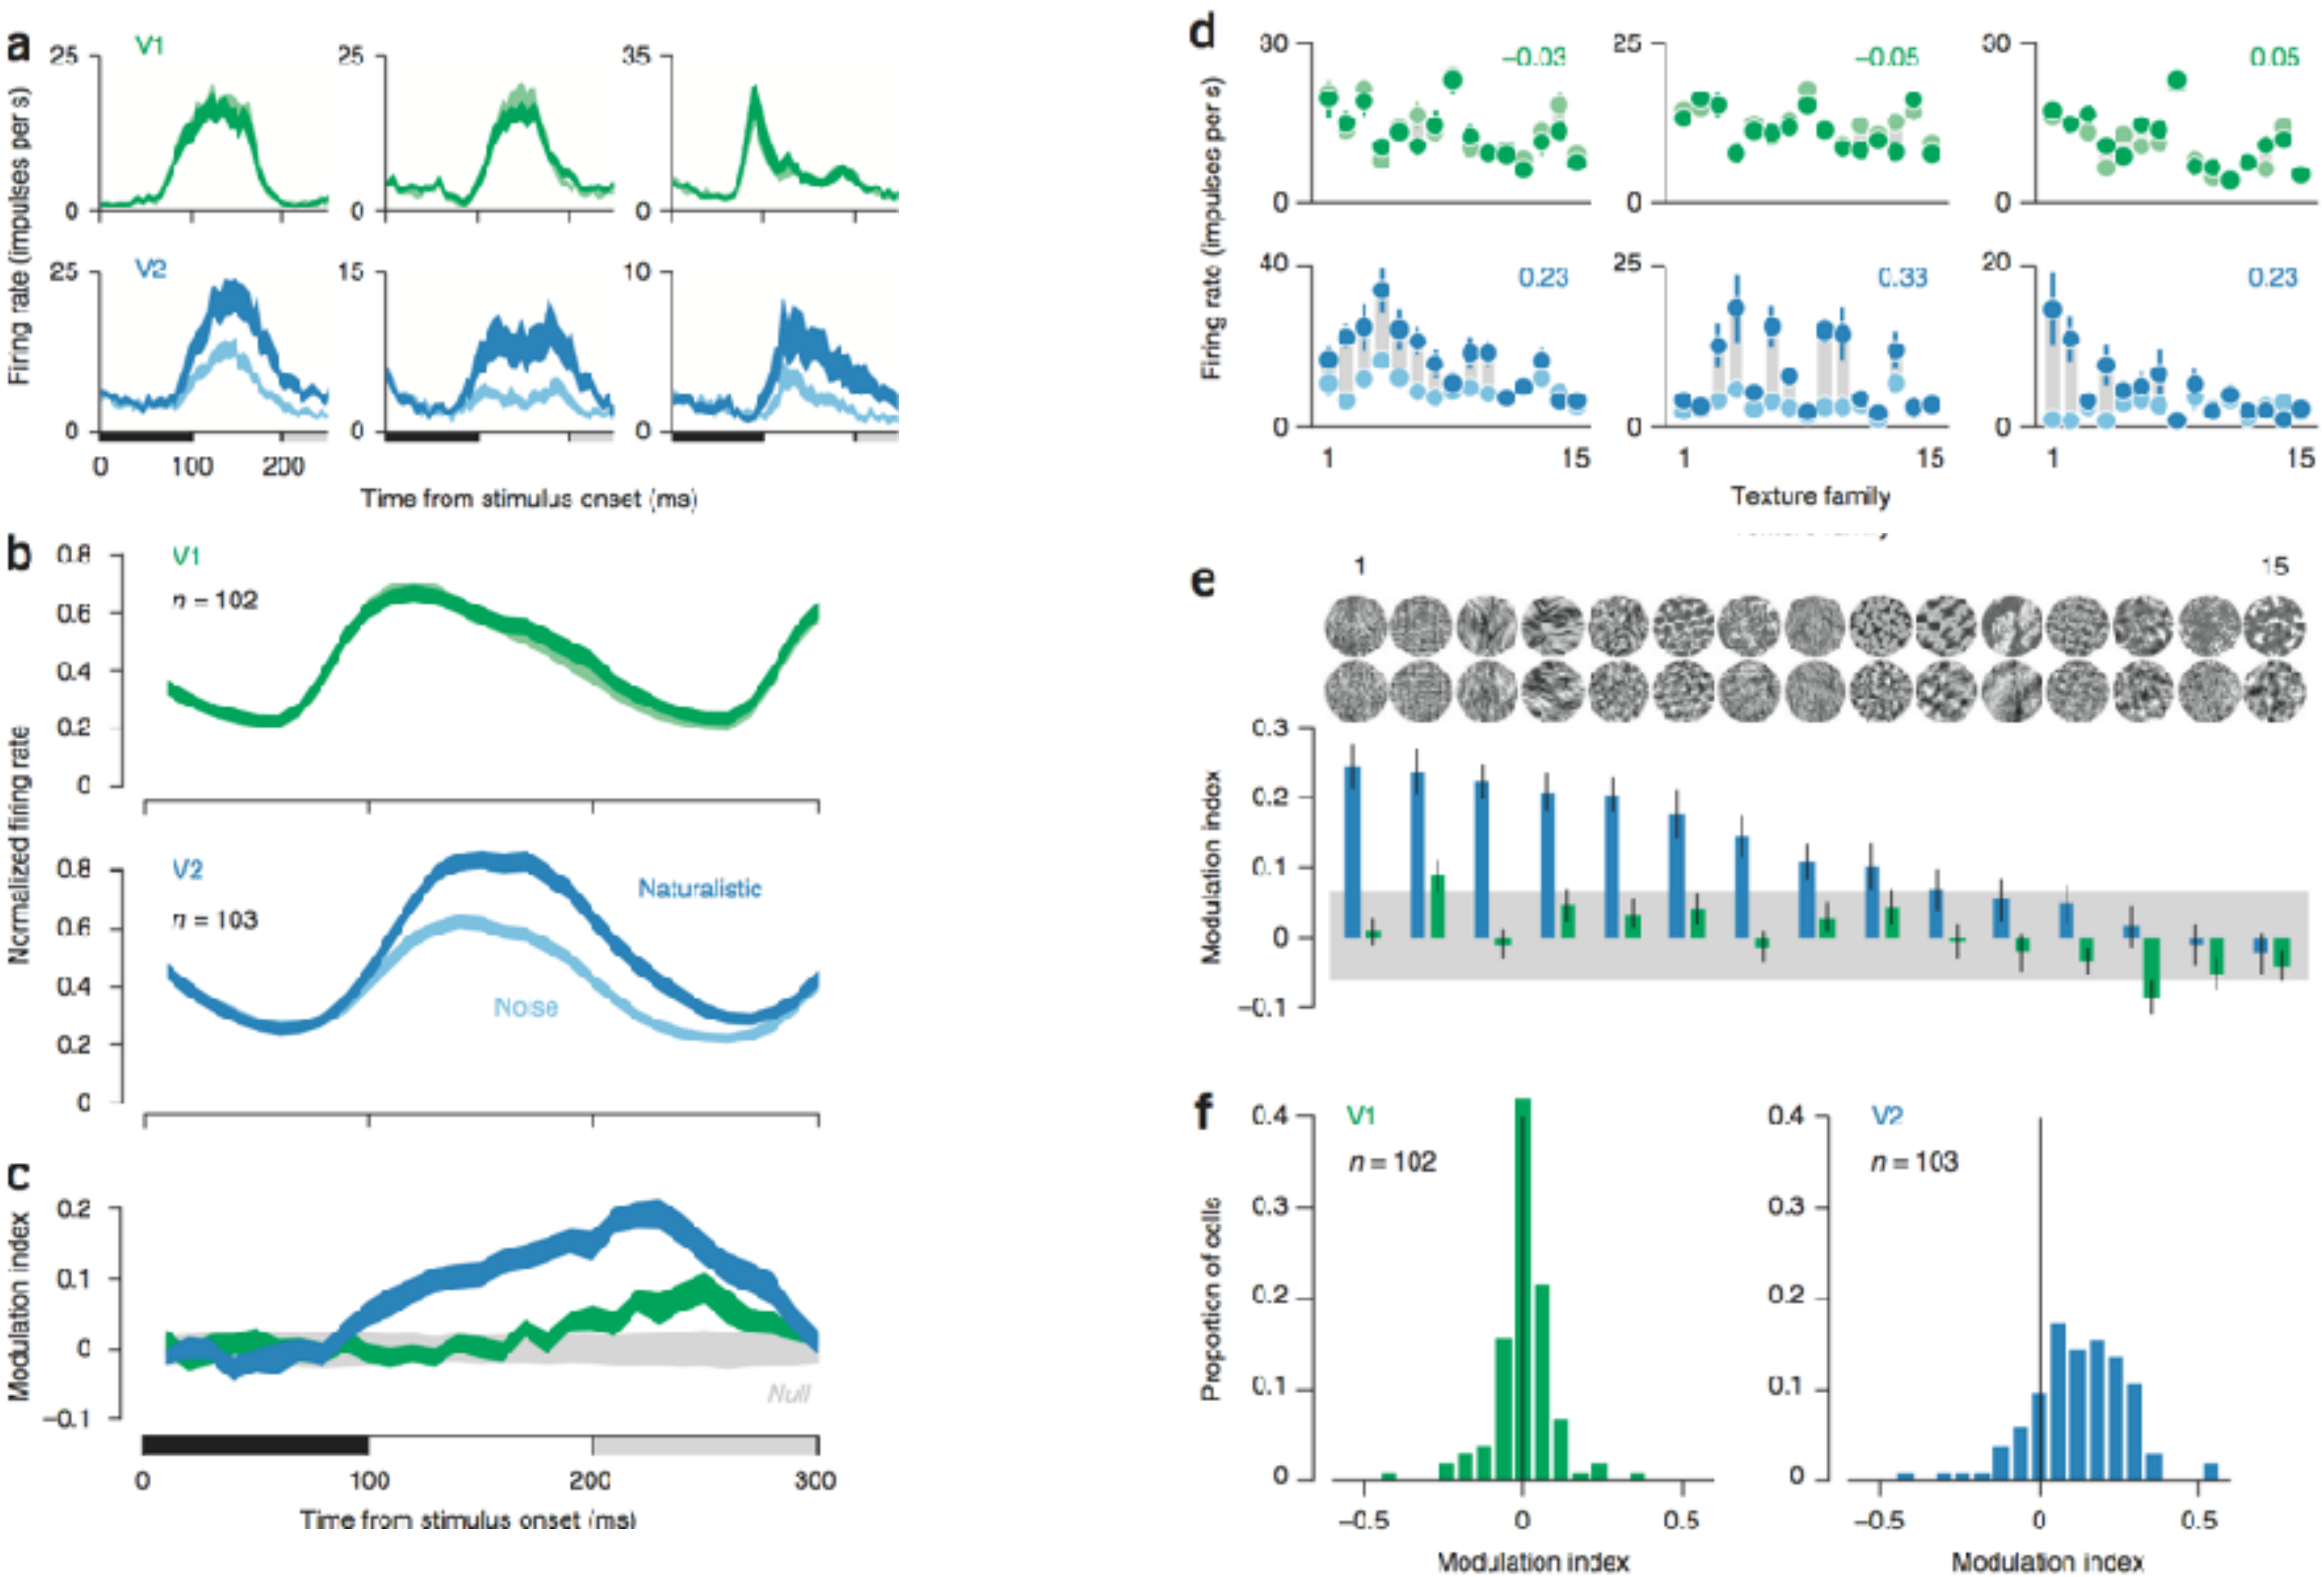

fMRI responses to Portilla textures

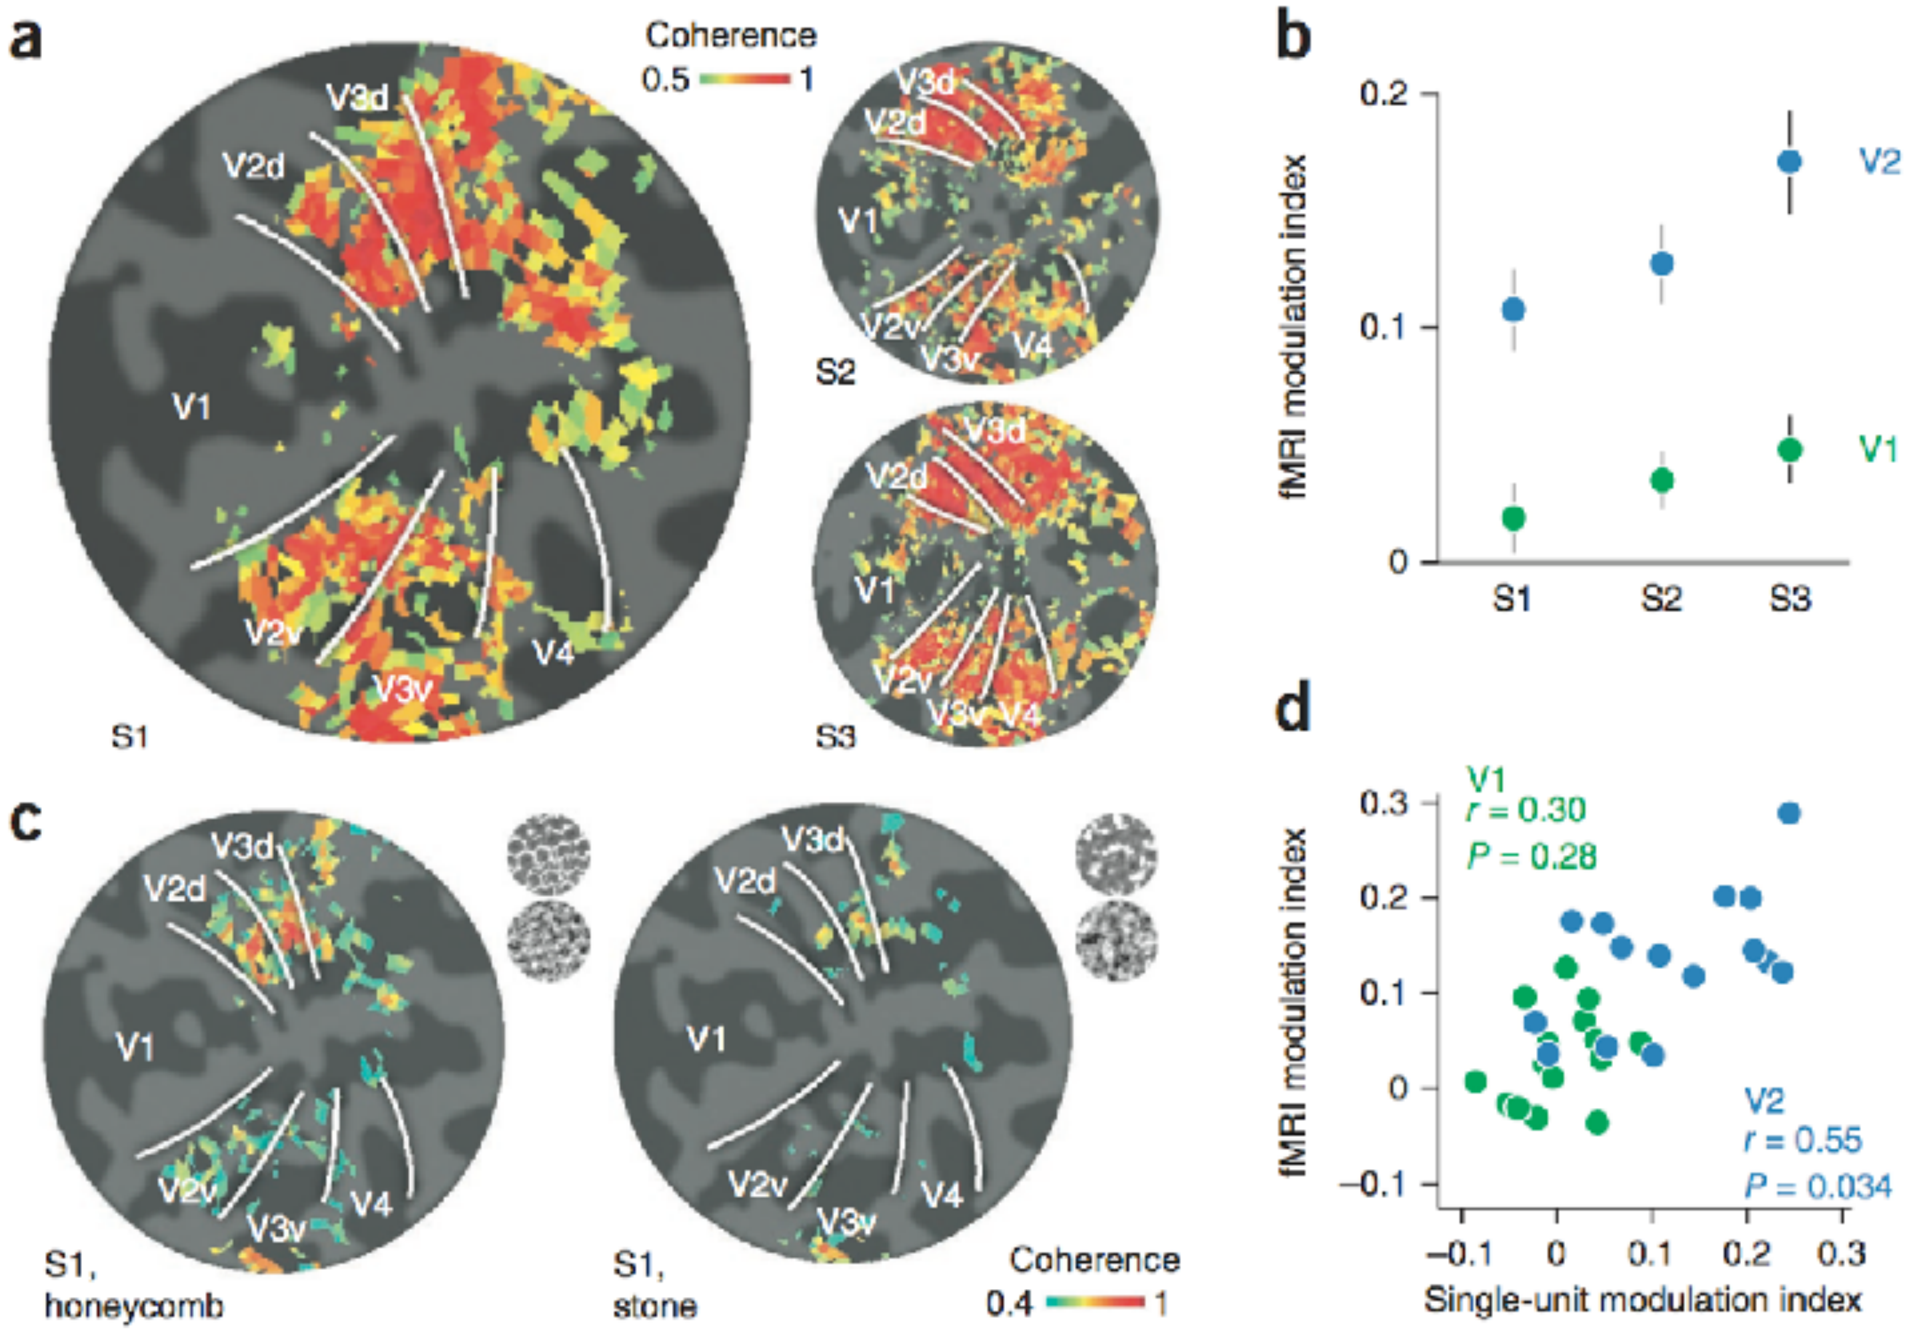

Perceptual correlates

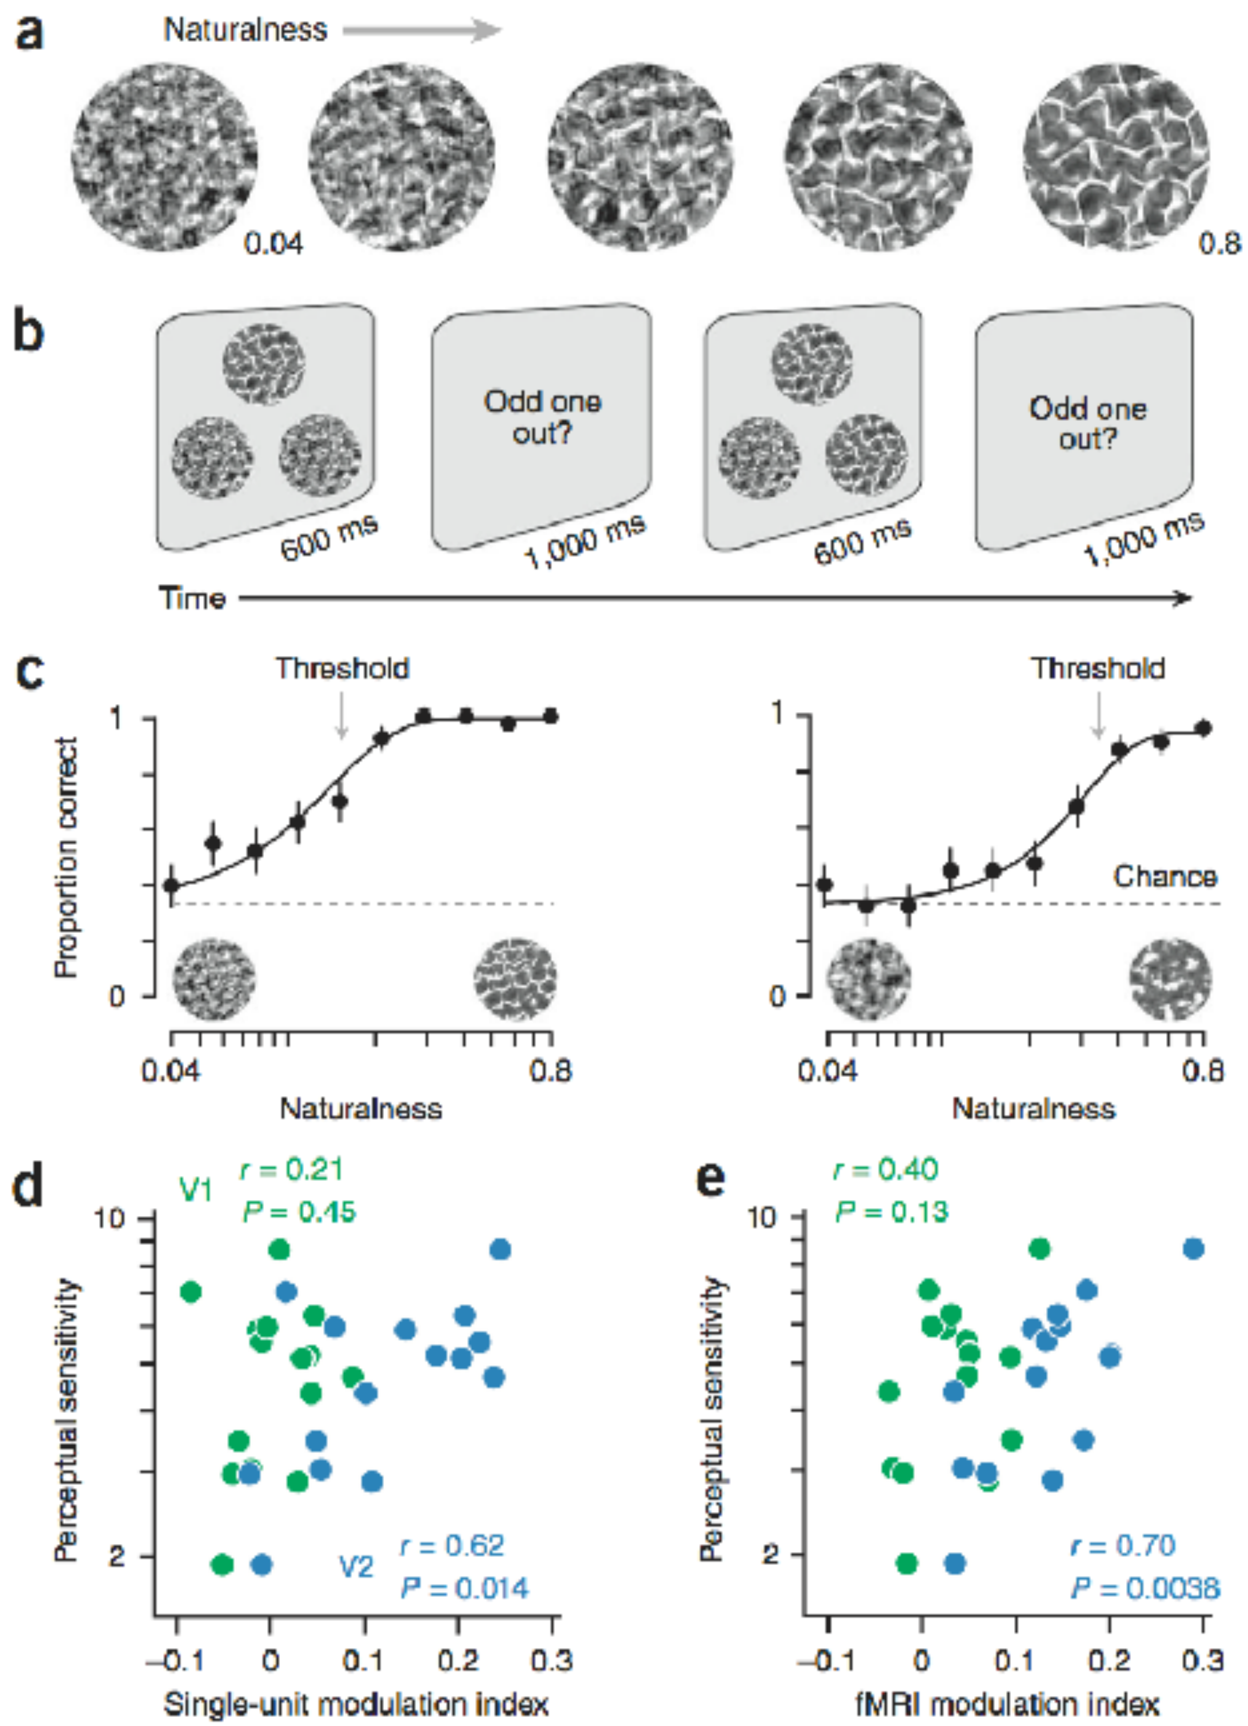

The mechanical Turk challenge

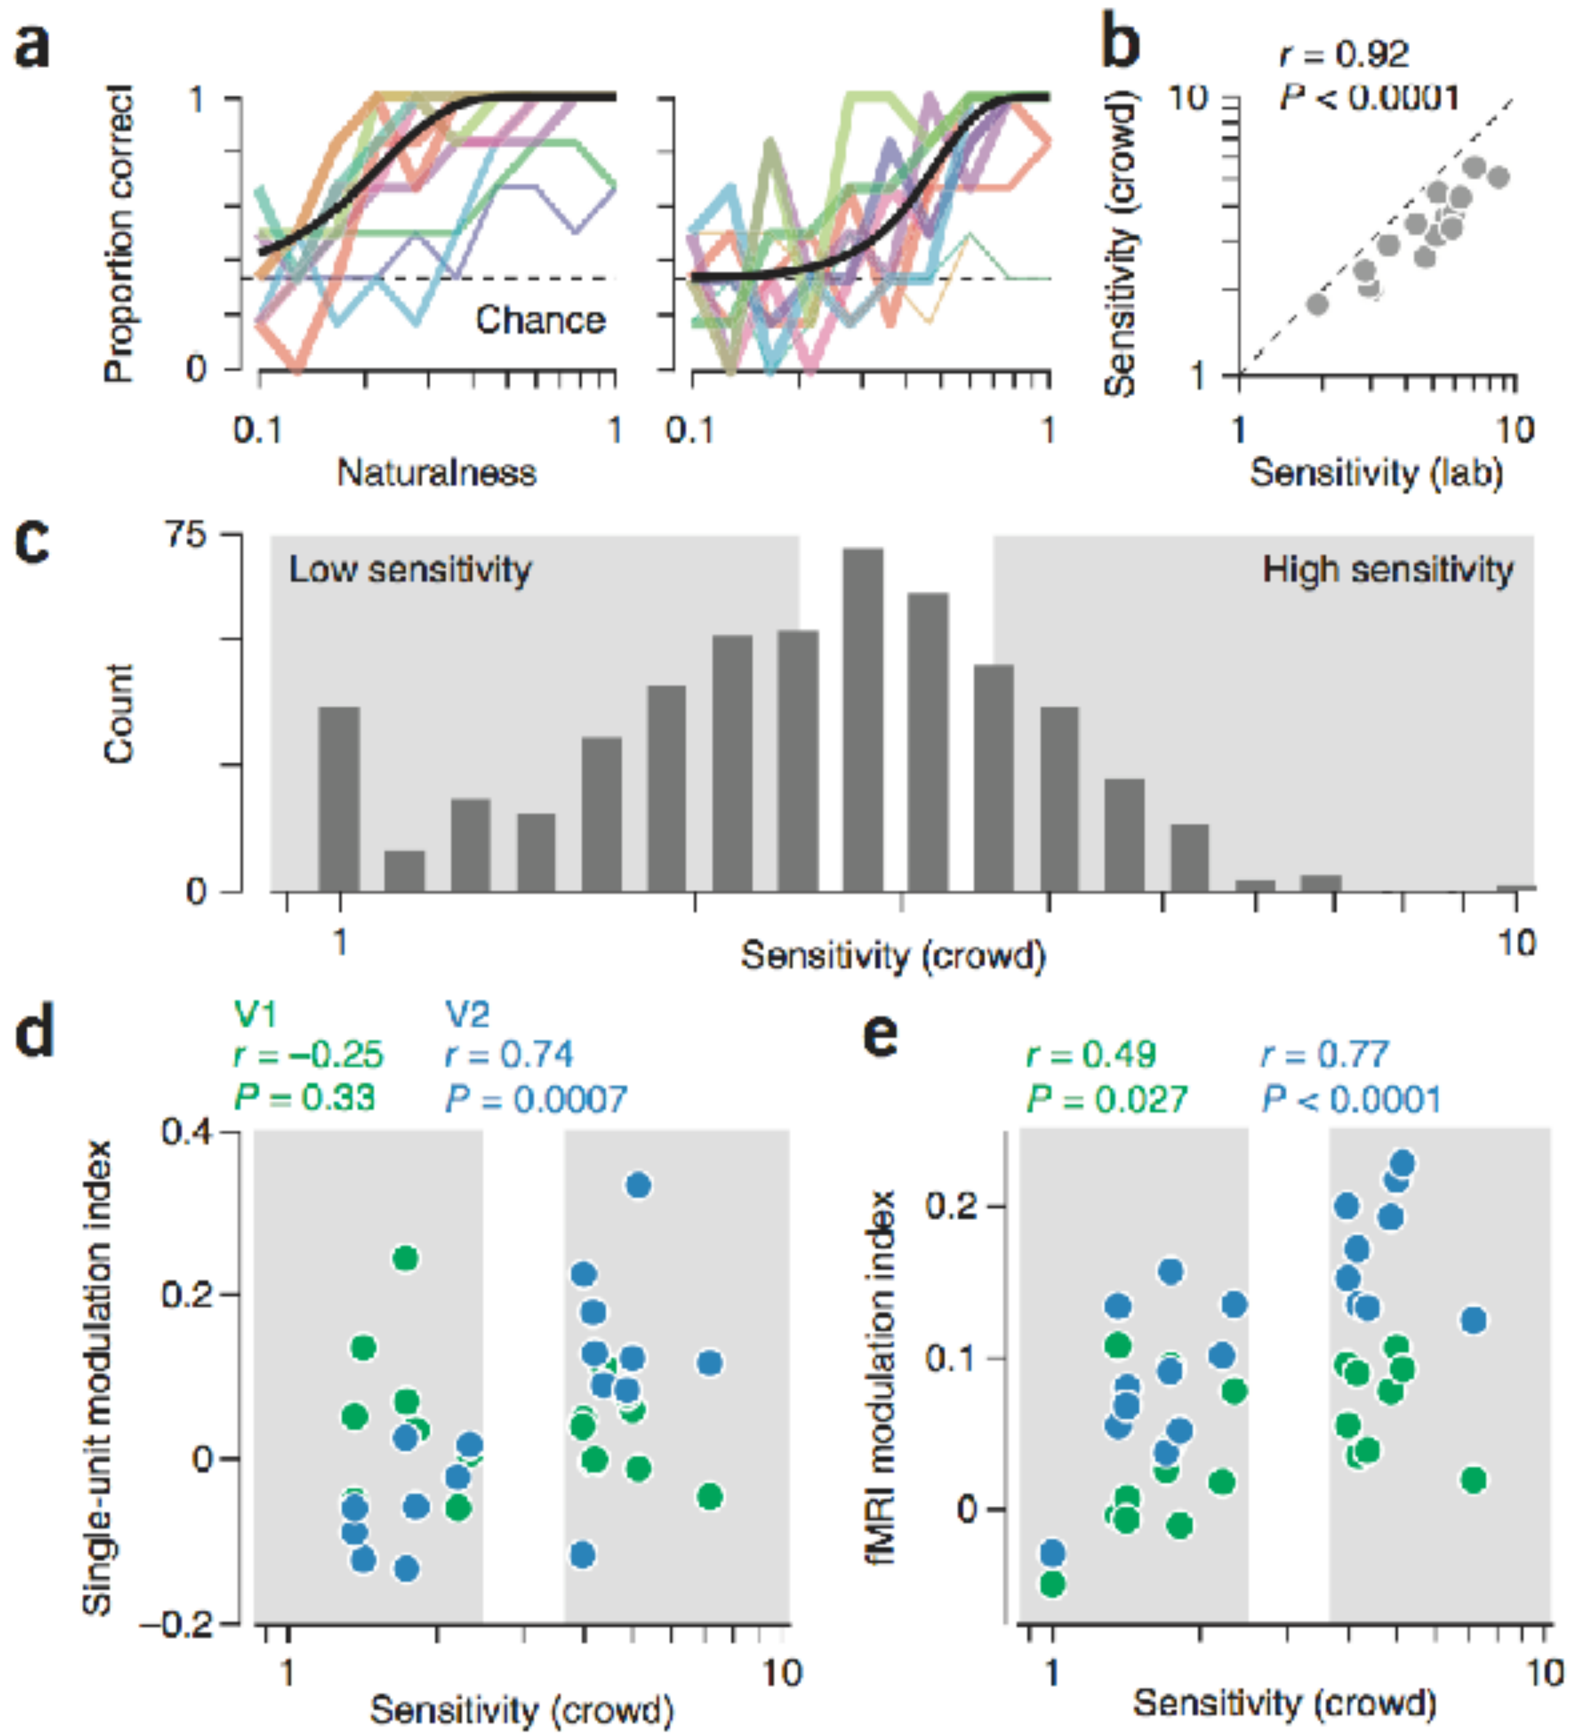

The mechanical Turk challenge

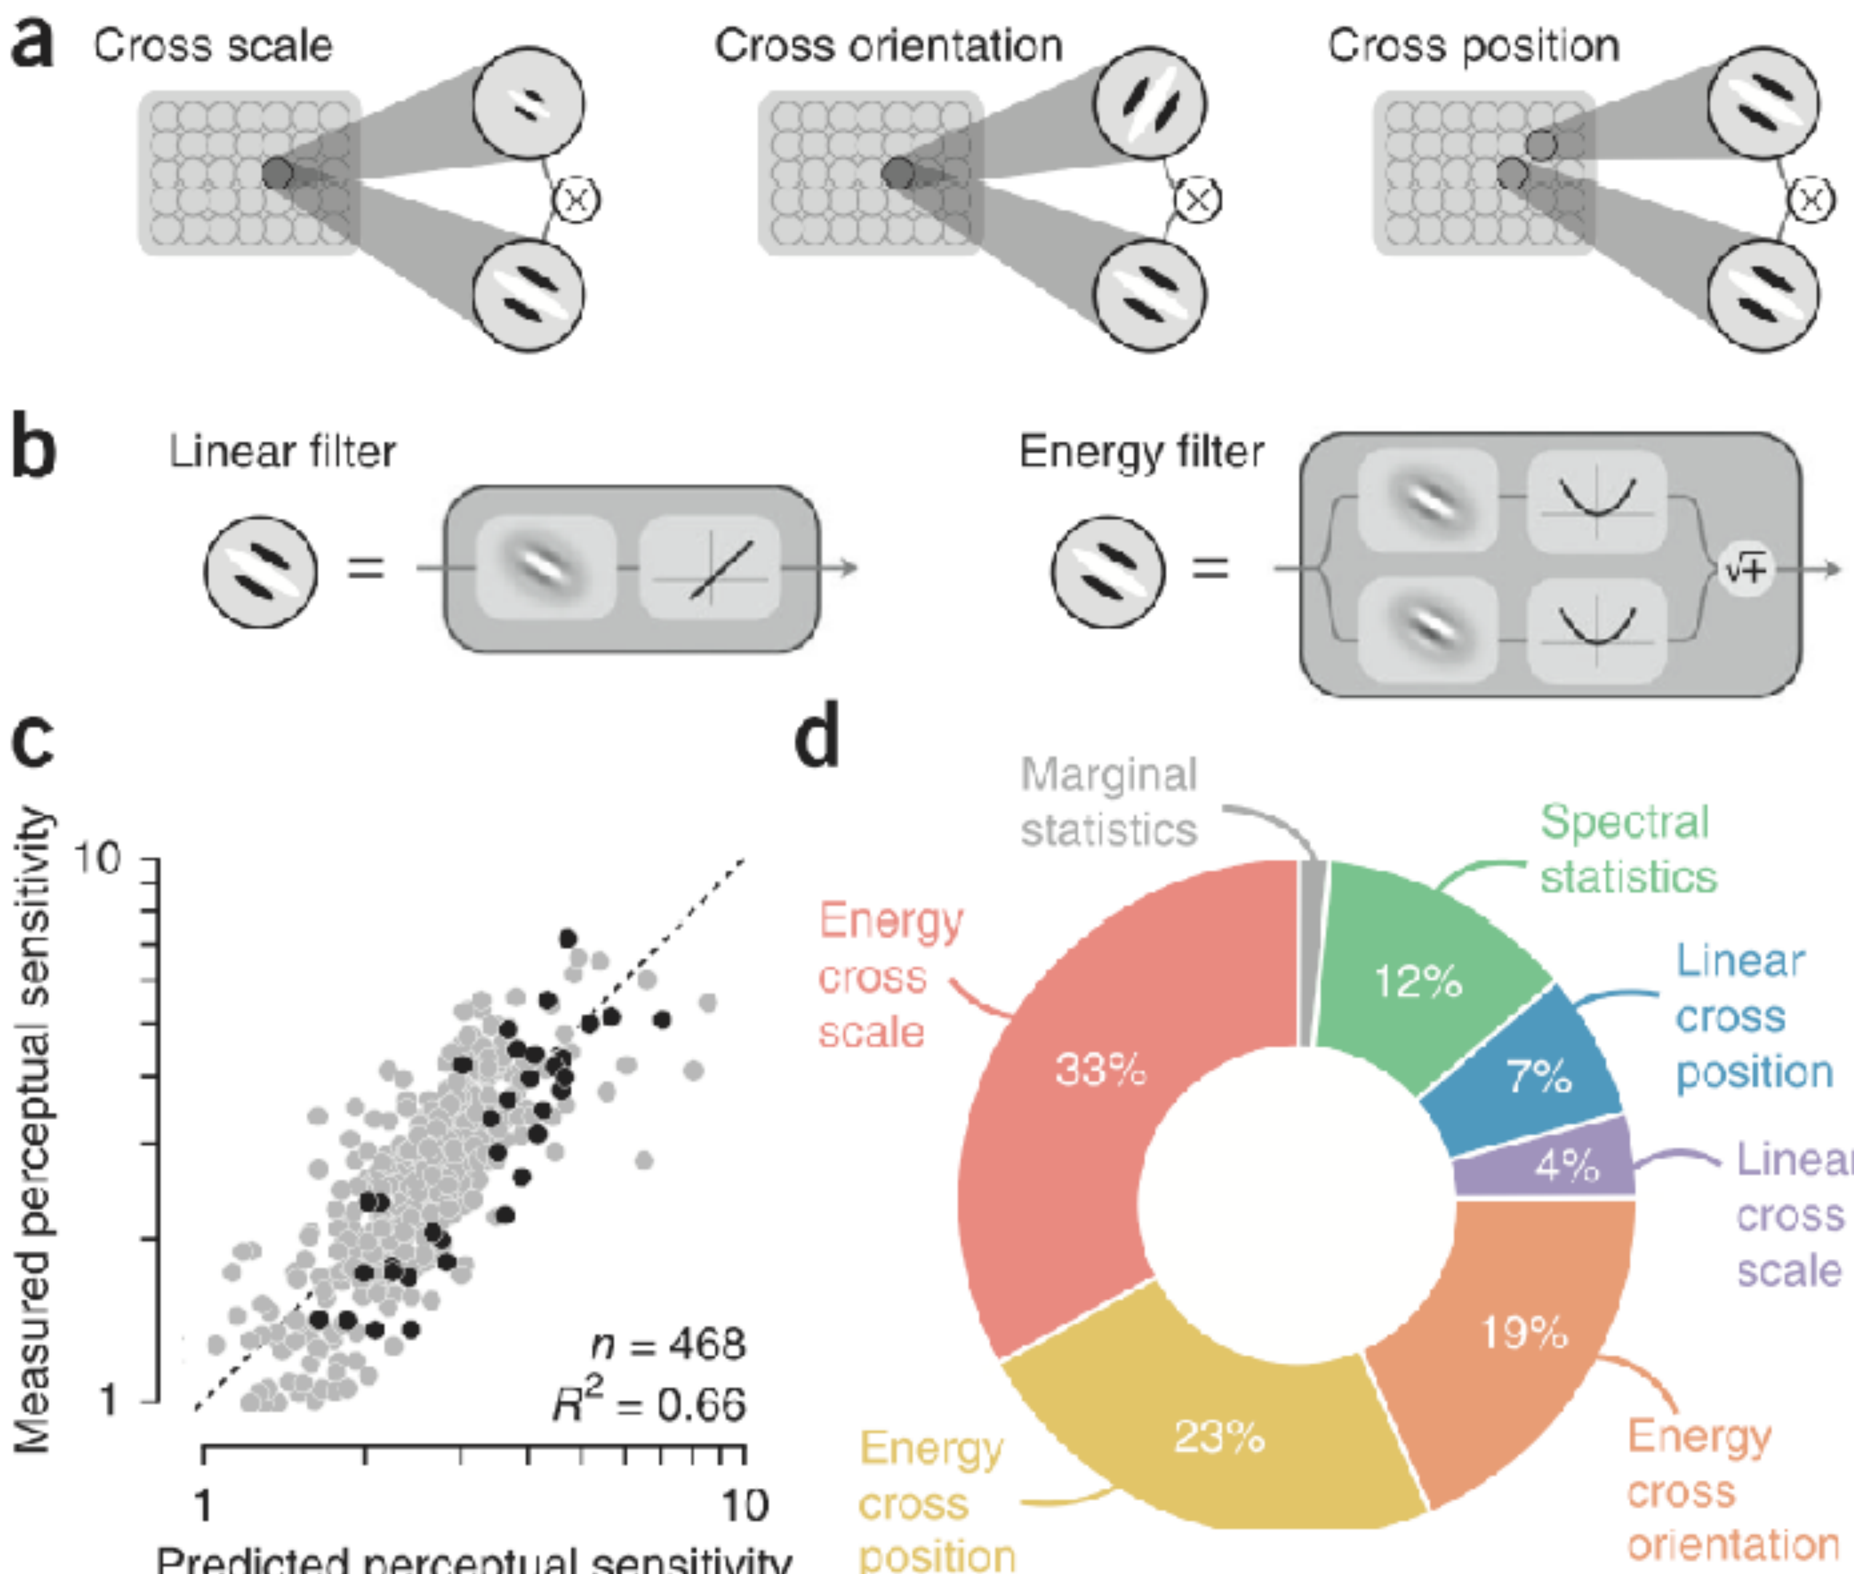

Representational untangling

Opinion

TRENDS in Cognitive Sciences Vol.11 No.8

Full text provided by www.sciencedirect.com

Untangling invariant object recognition

James J. DiCarlo and David D. Cox

Representational untangling

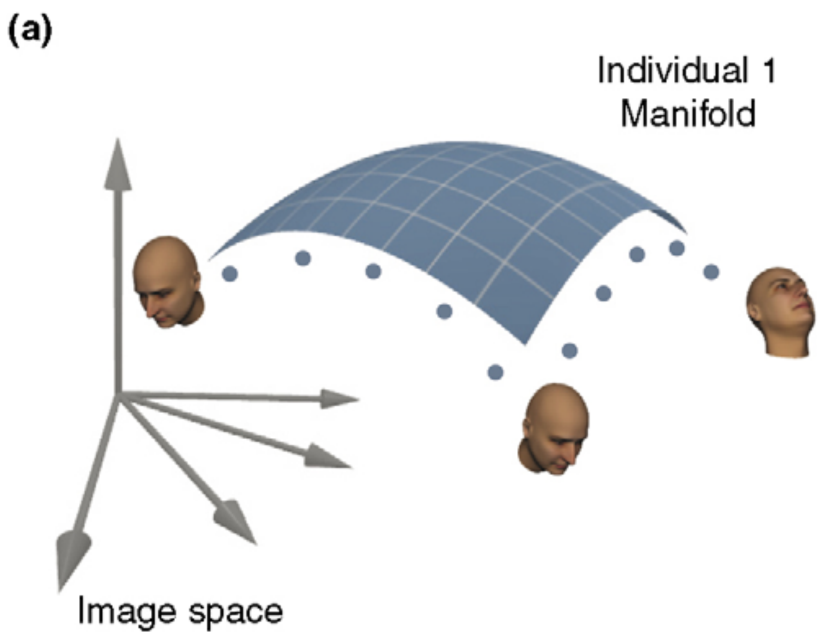

Representational untangling

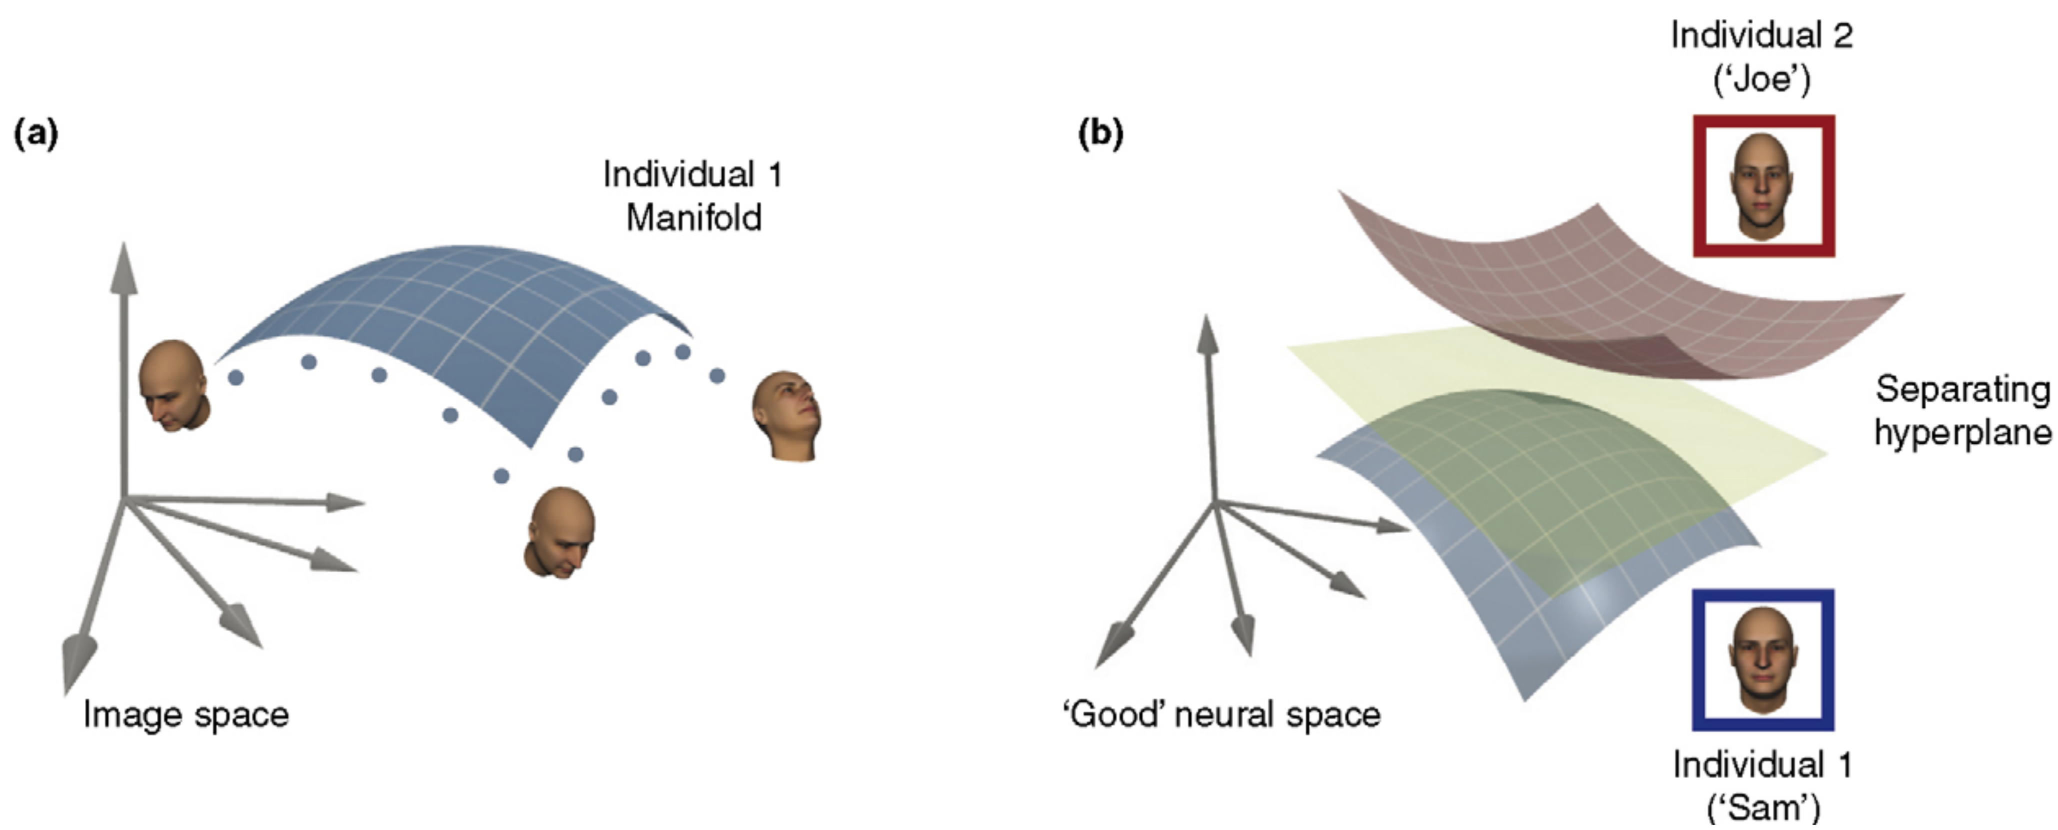

Representational untangling

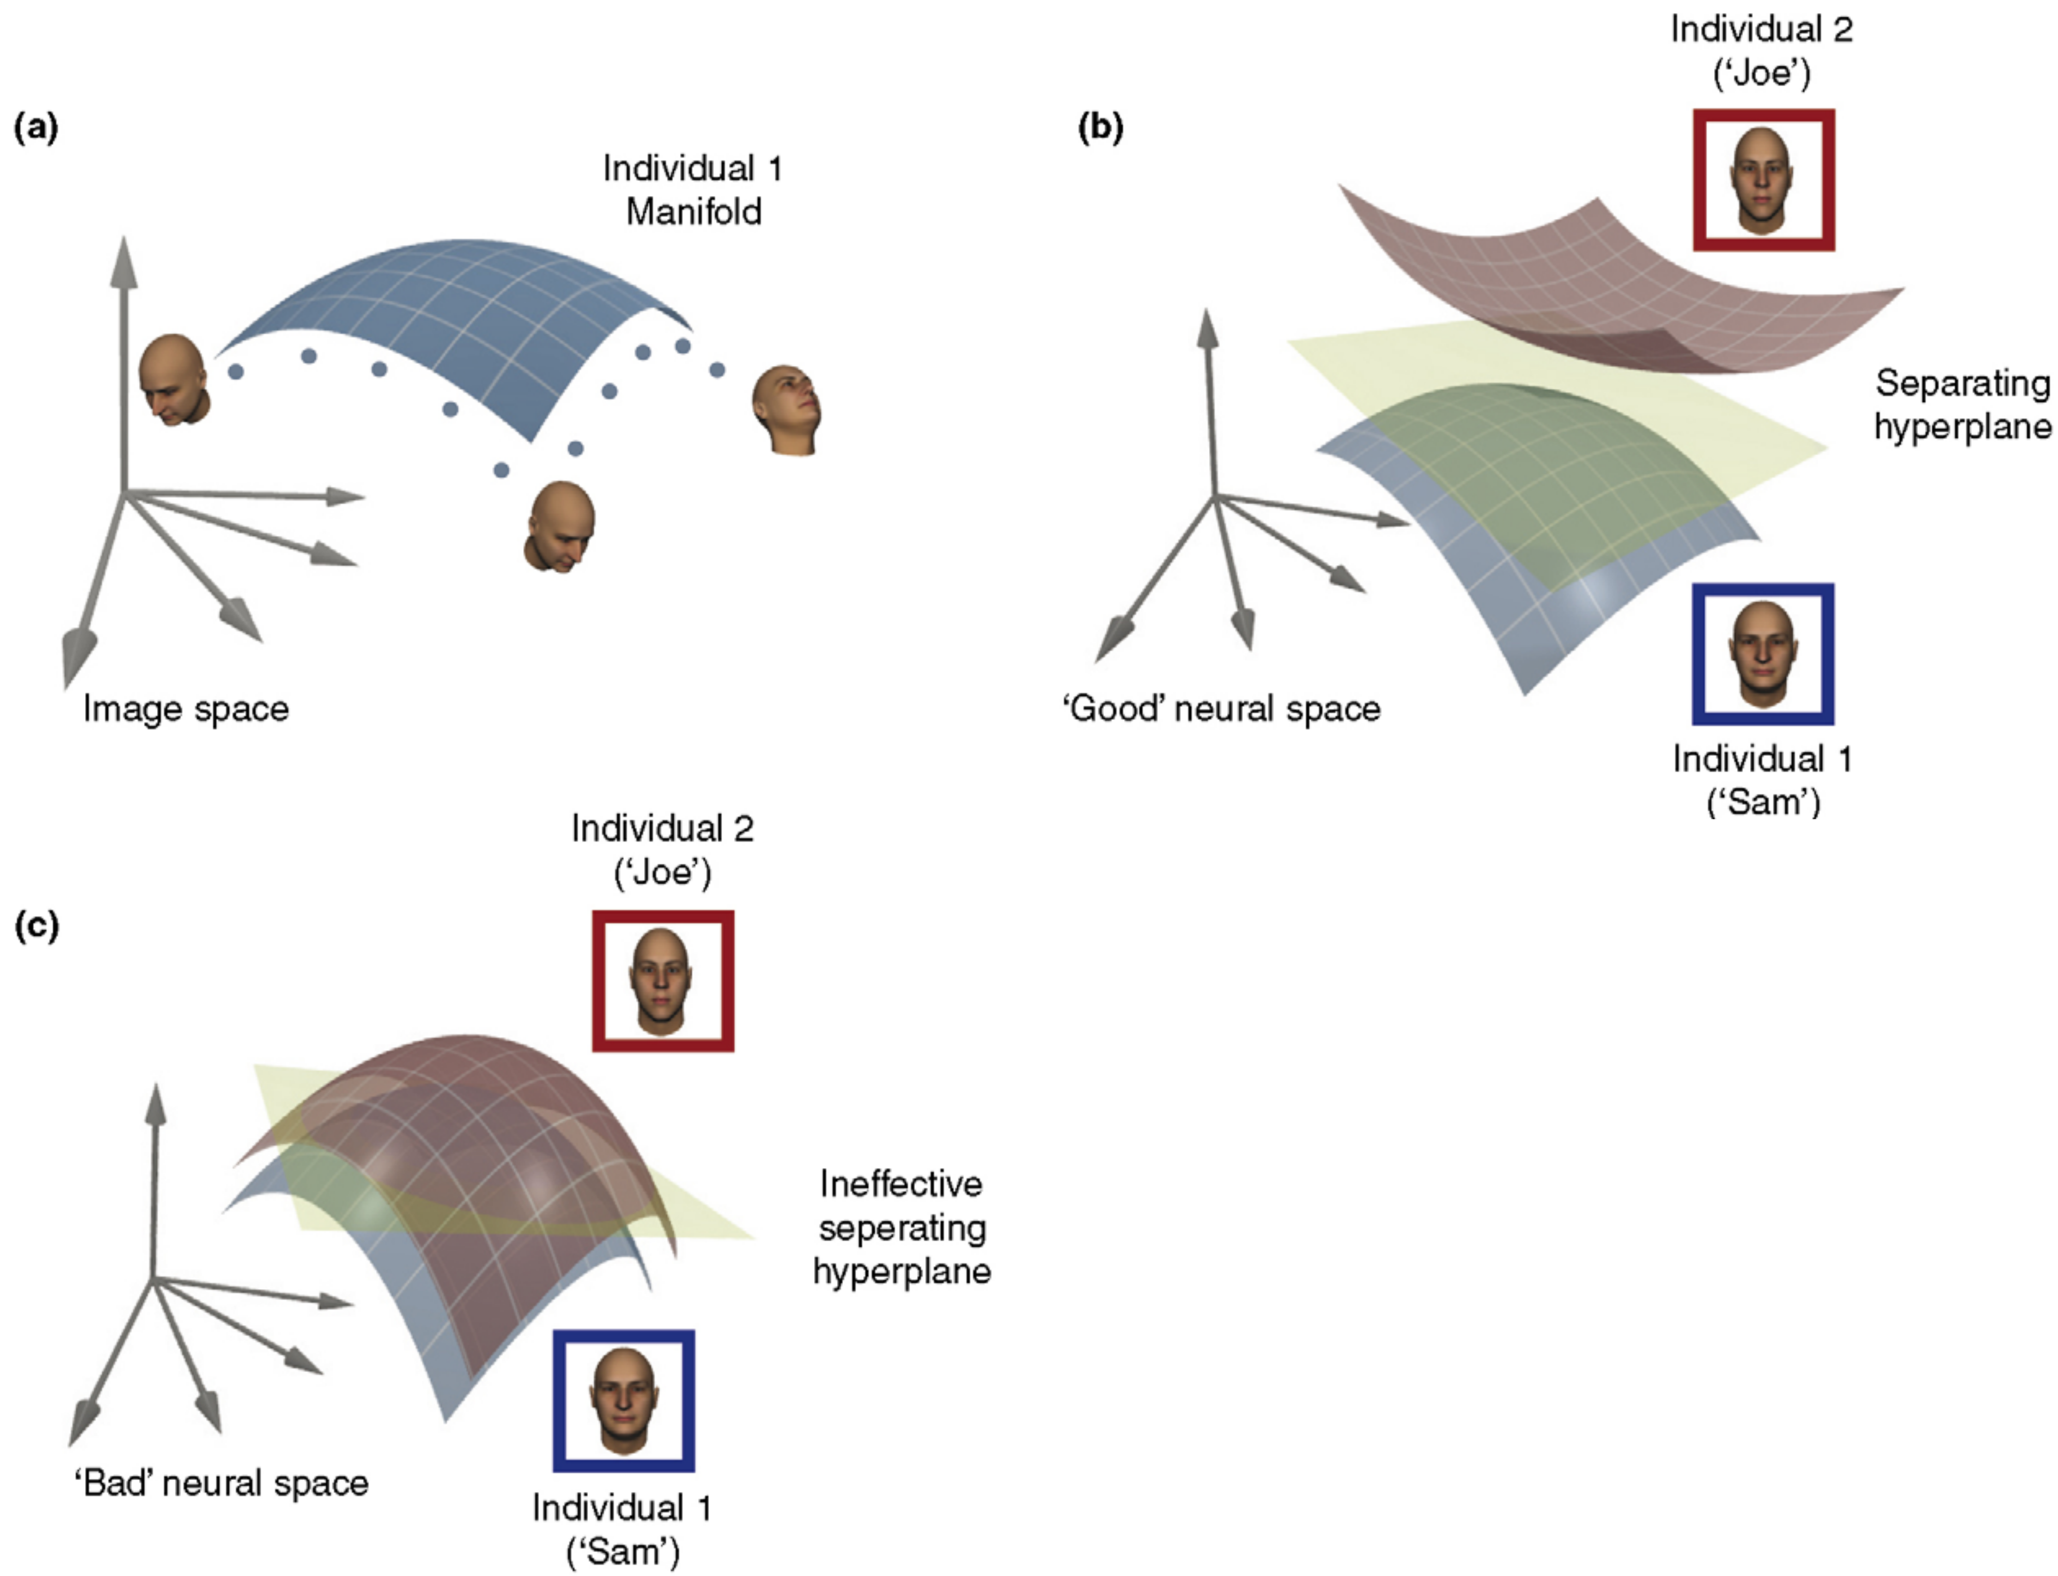

Representational untangling

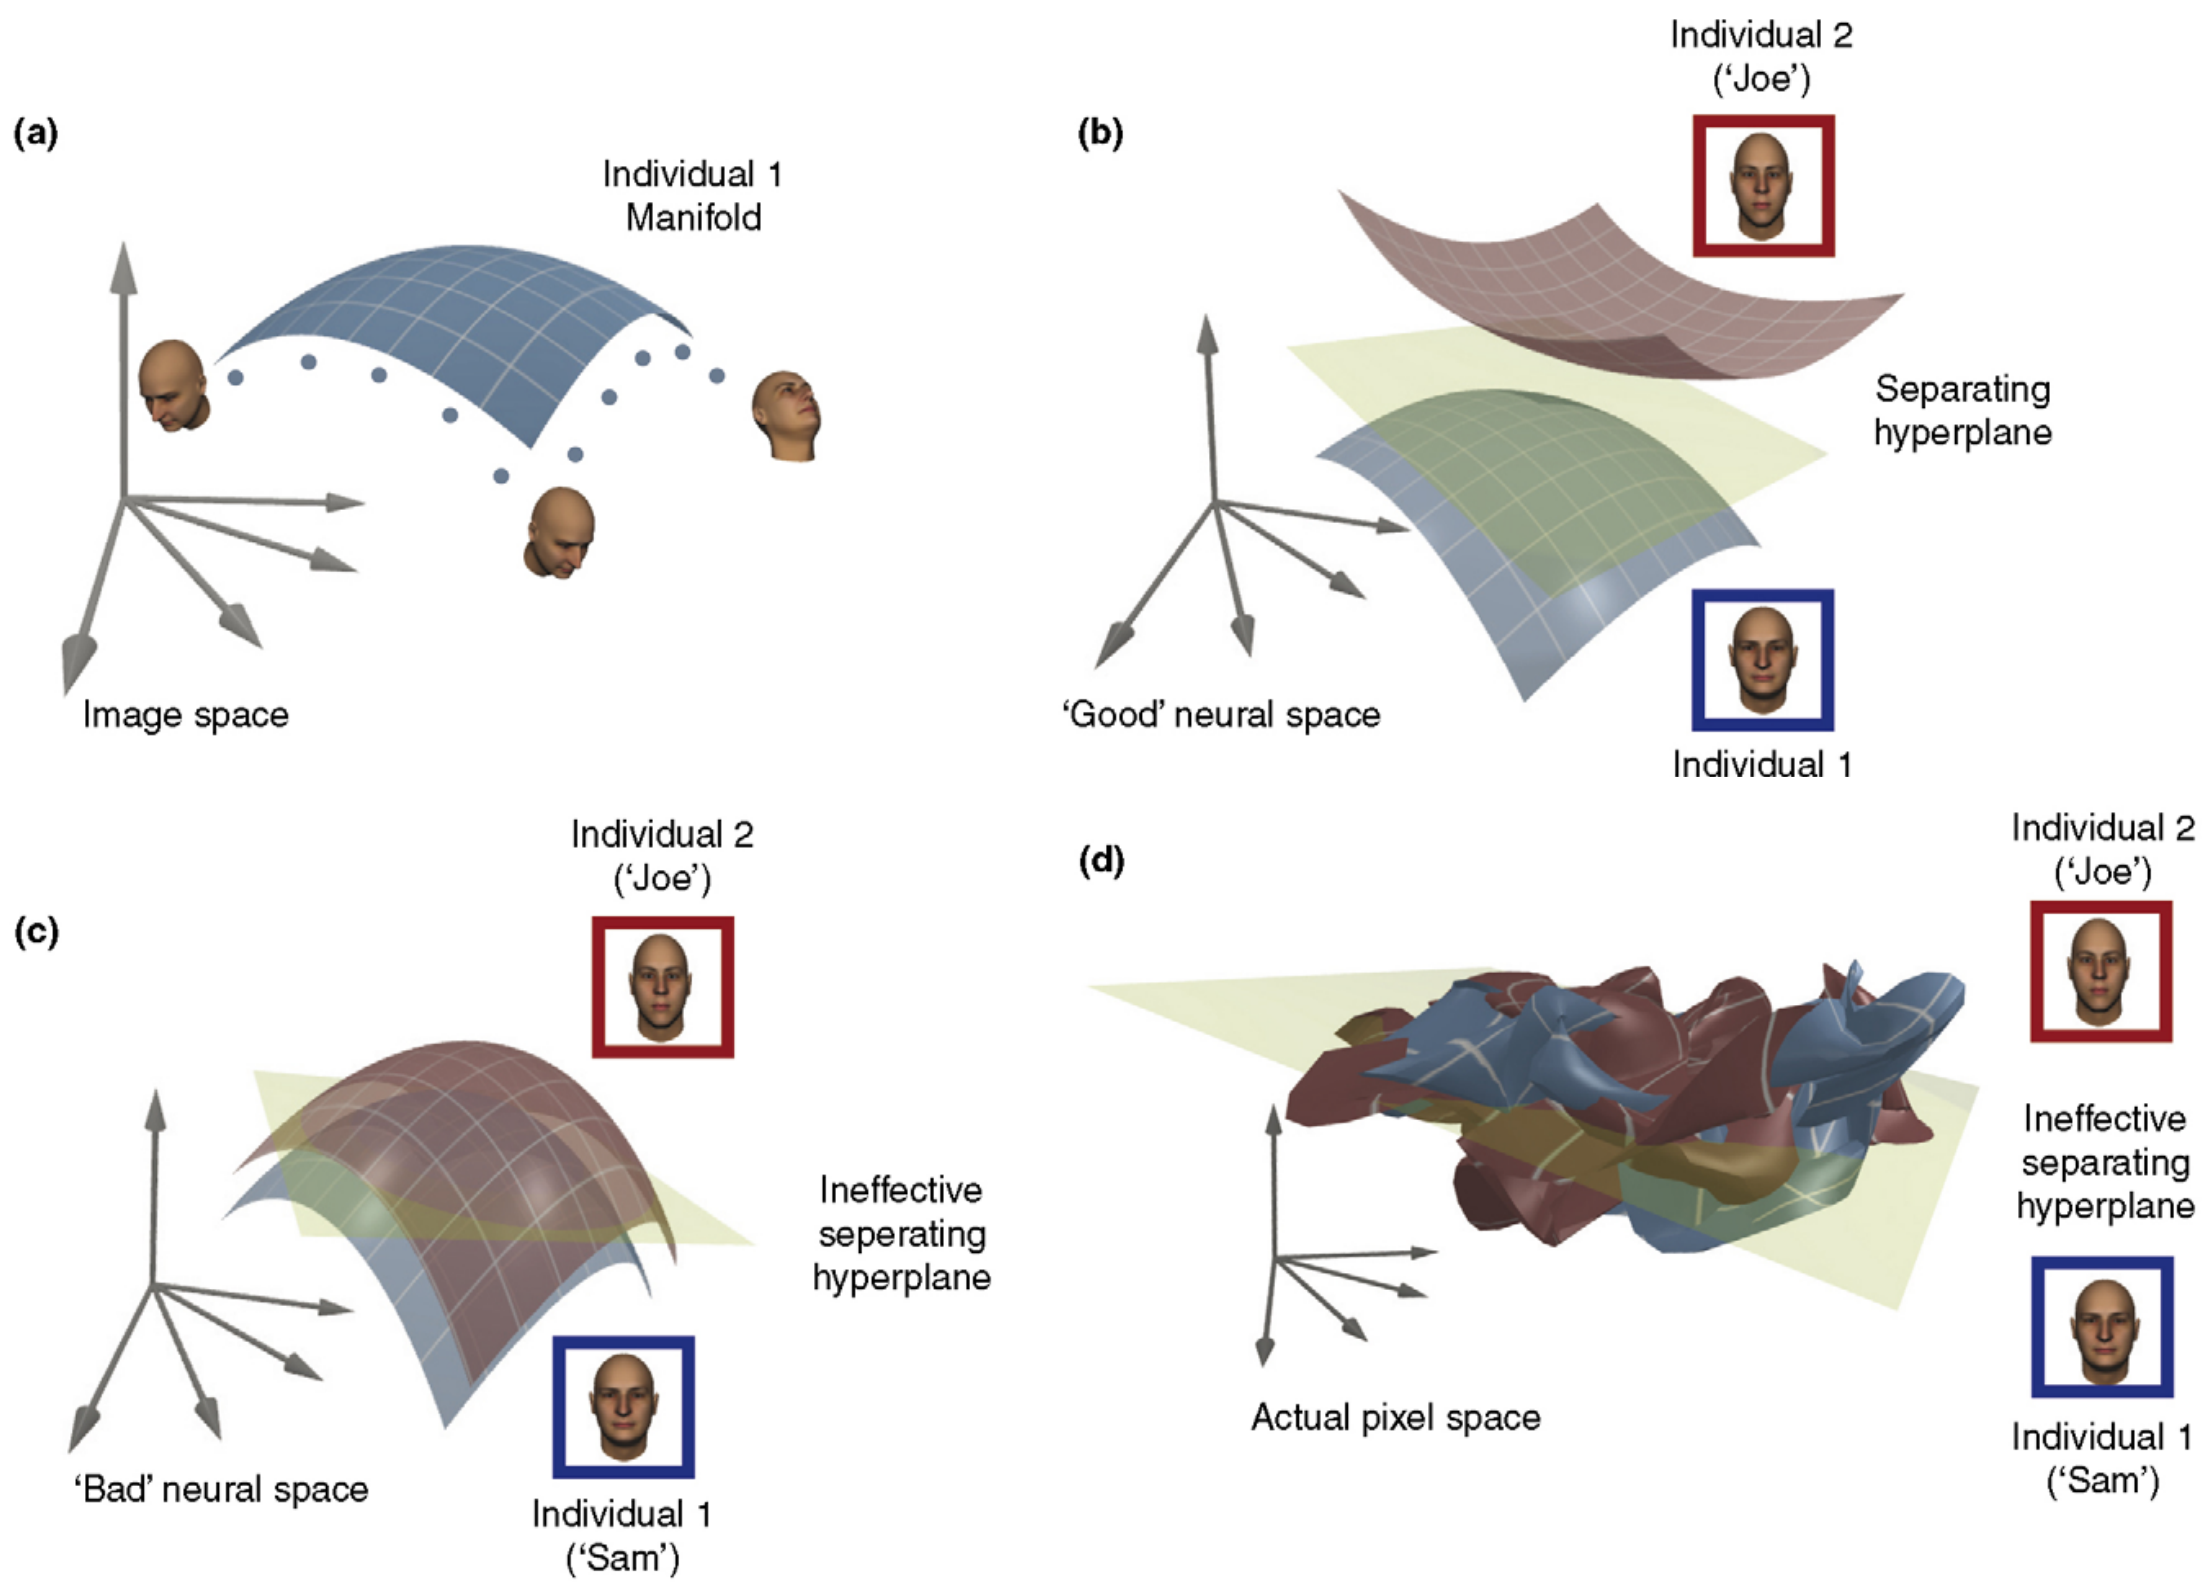

Invariance properties of V1/V2 neurons

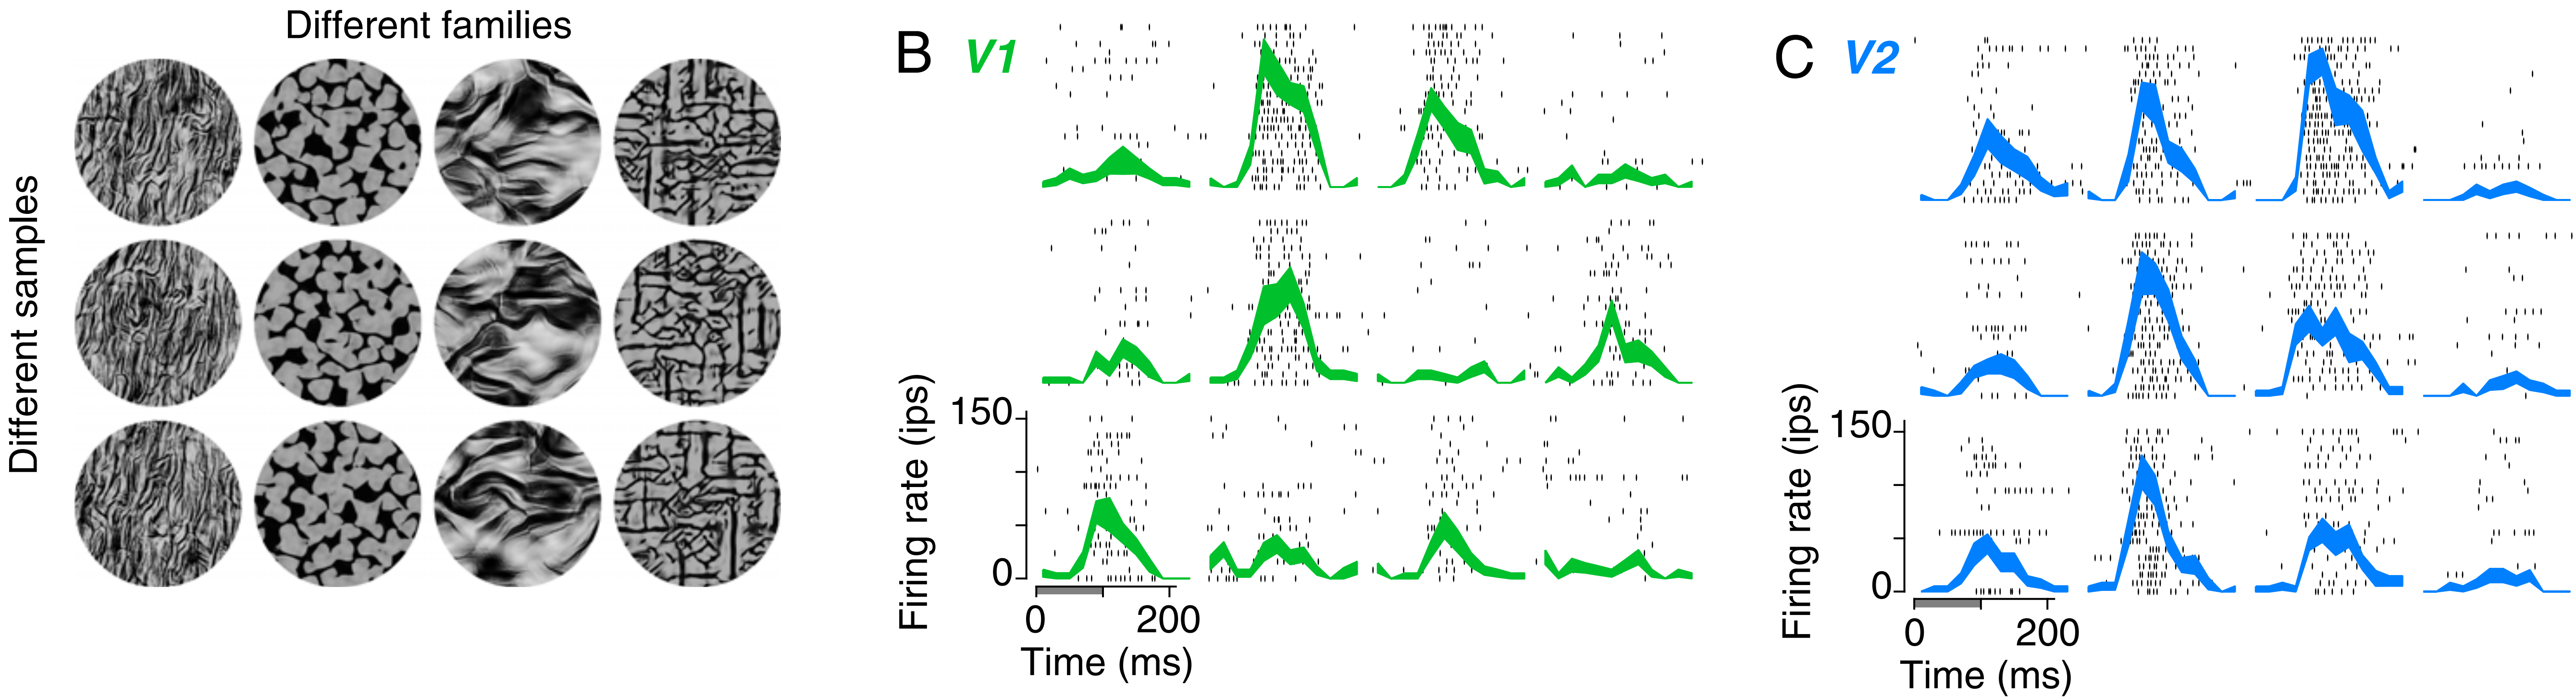

Representational untangling in the visual system

Texture family

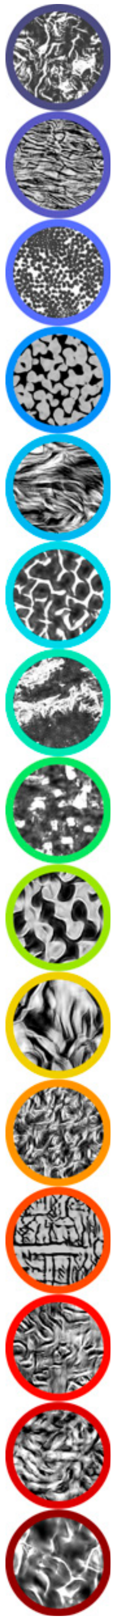

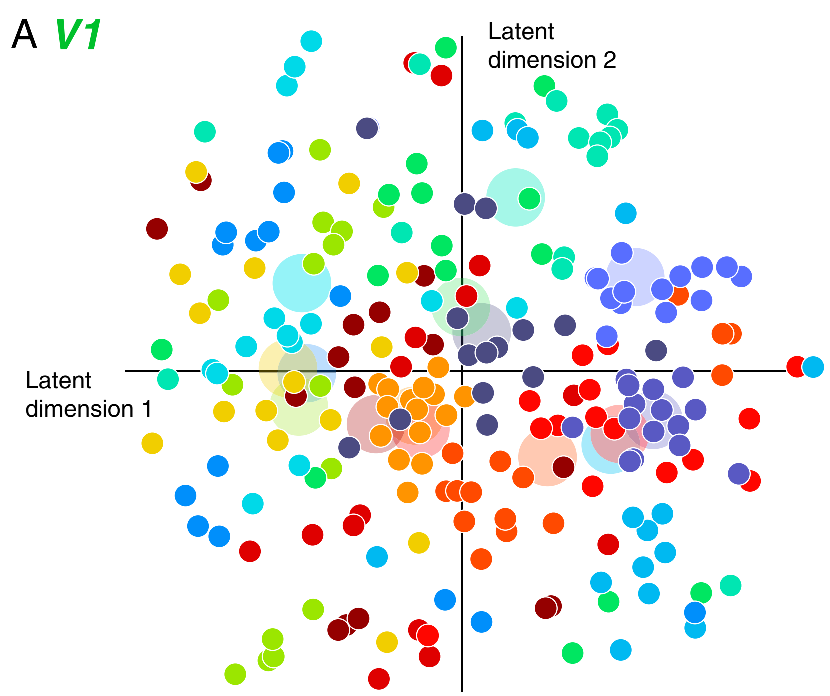

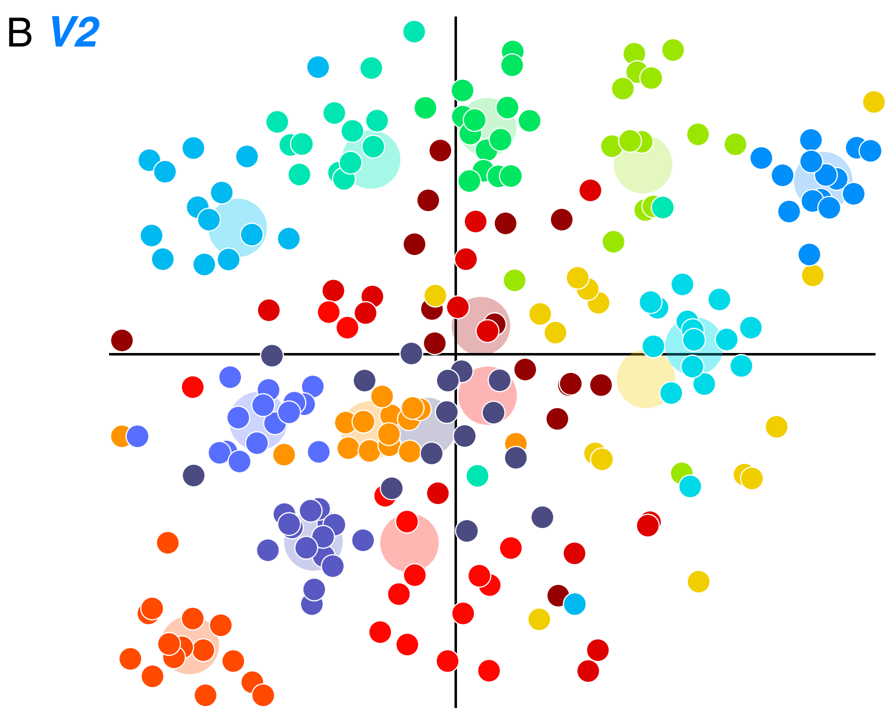

Decoding of stimulus information

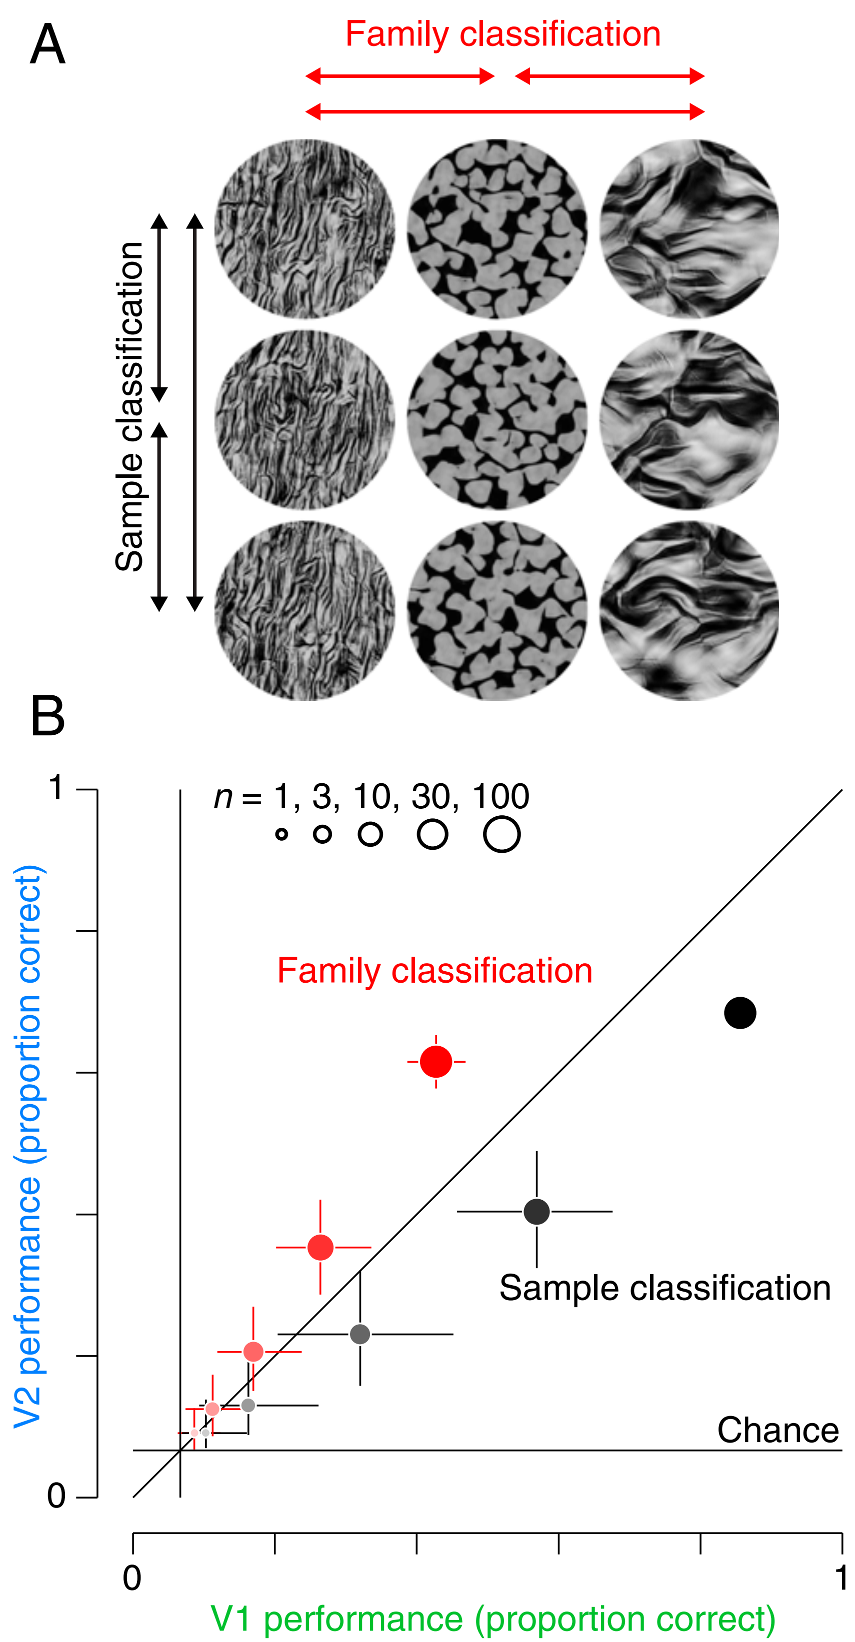

Hierarchical inference

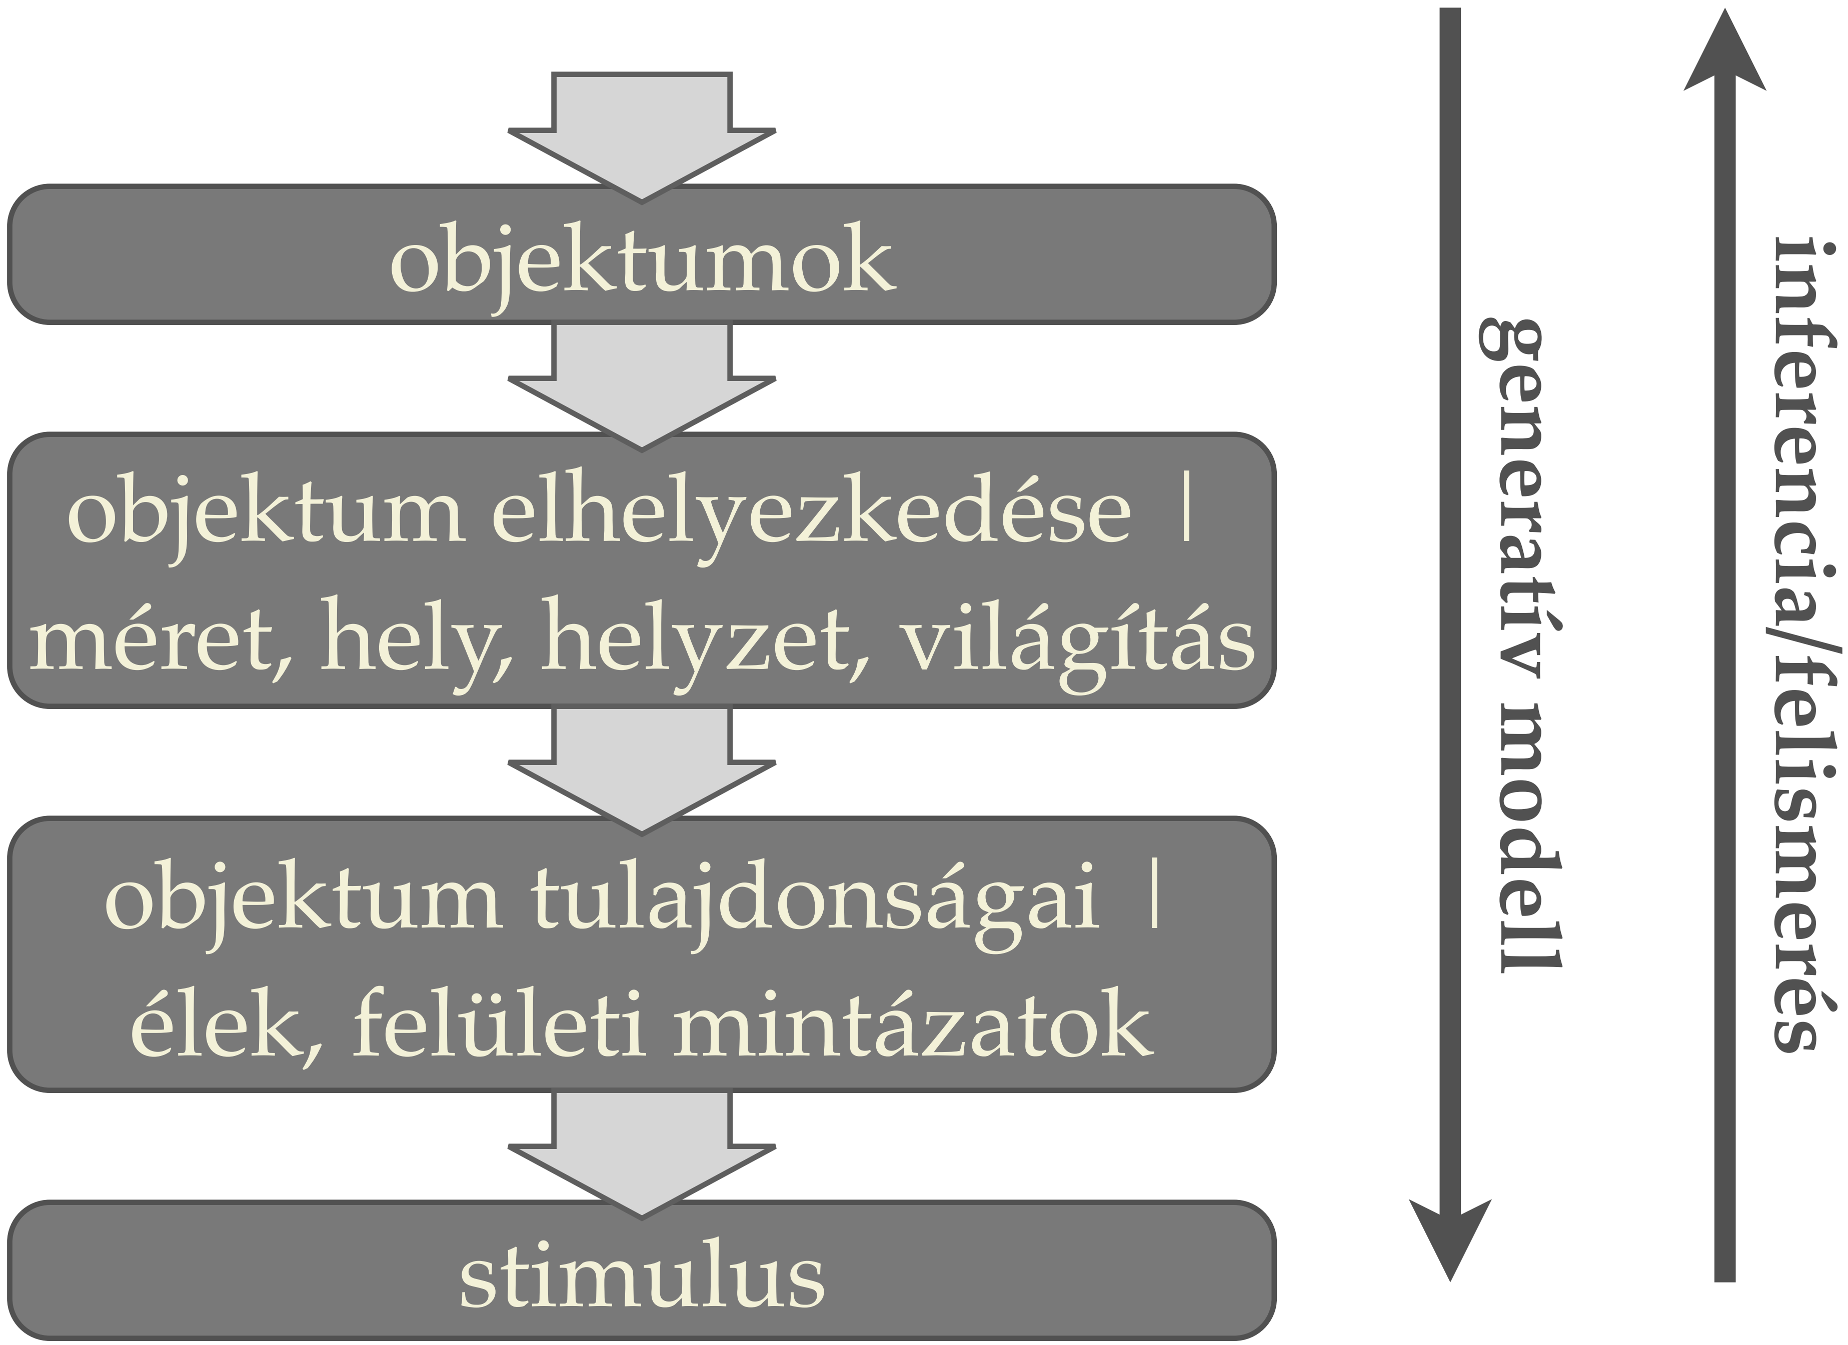

Hierarchical inference in a generative model

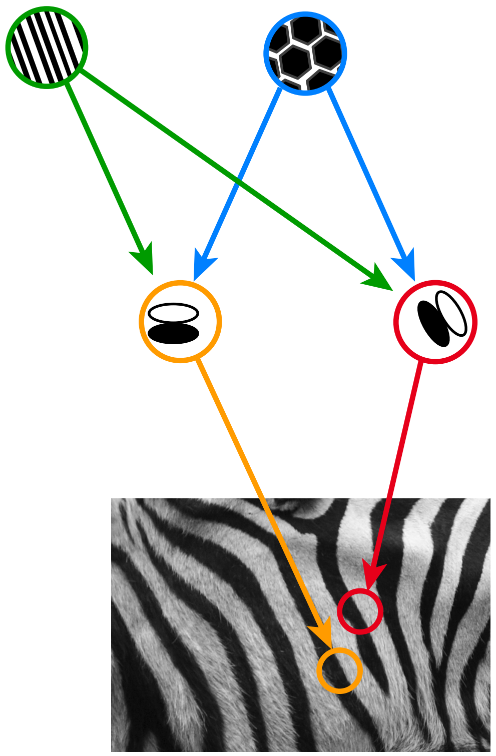

Hierarchical inference in a generative model

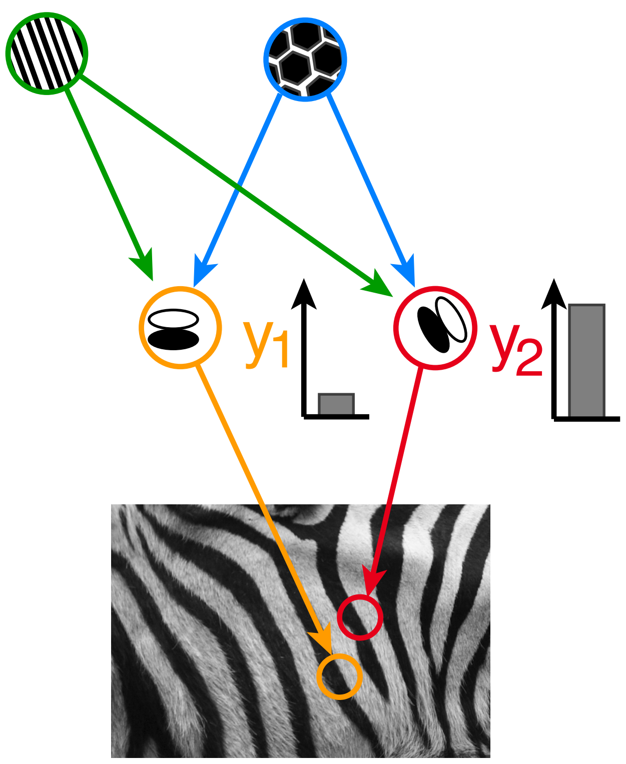

Hierarchical inference in a generative model

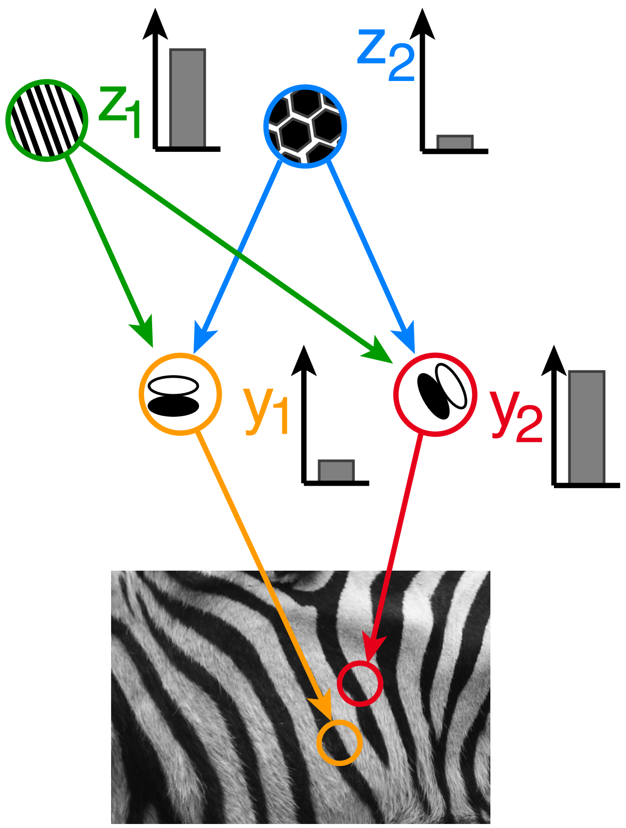

Hierarchical inference in a generative model

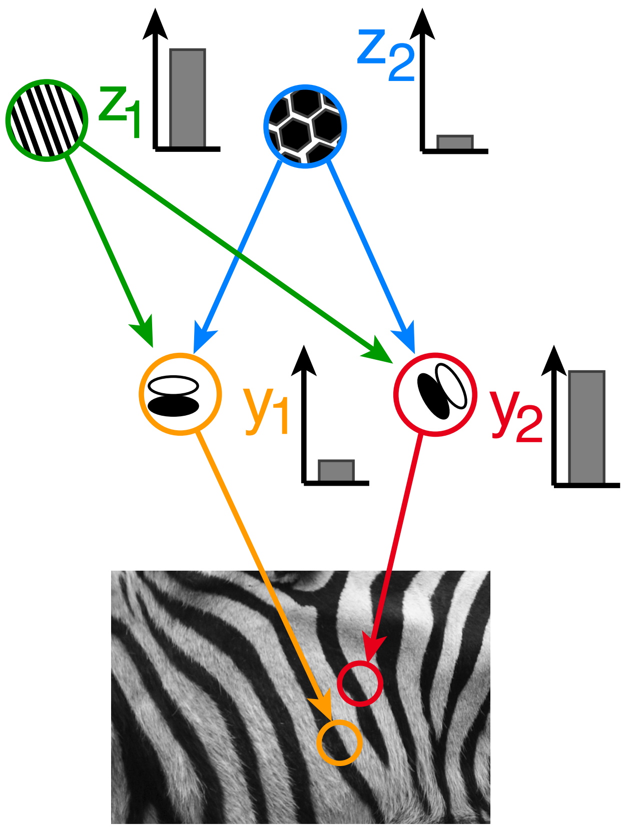

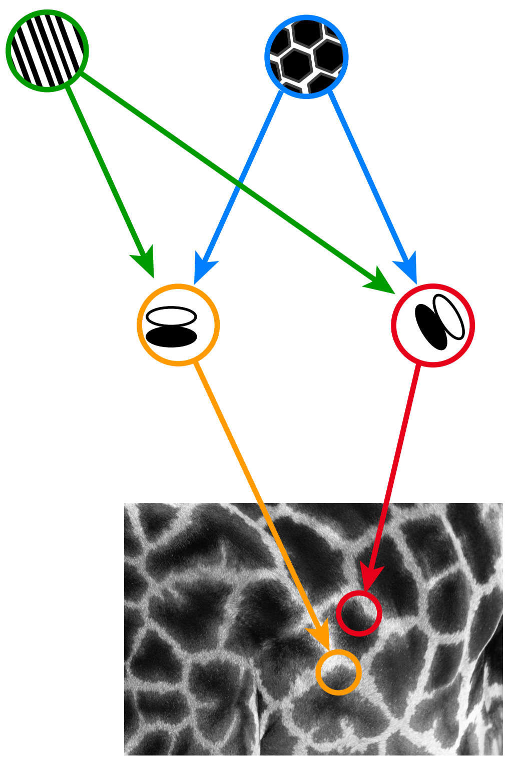

Hierarchical inference in a generative model

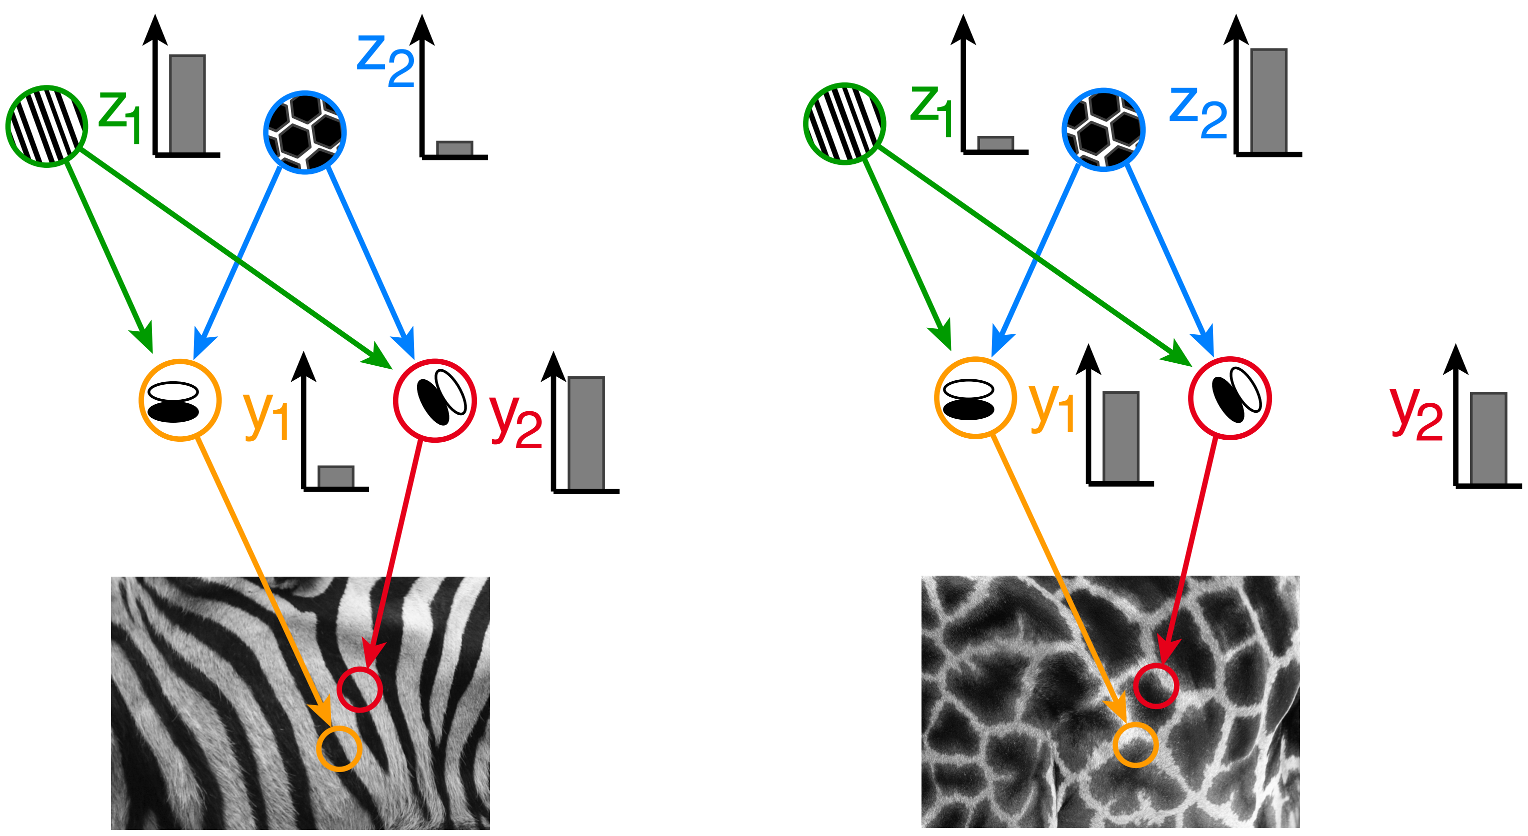

Hierarchical inference in a generative model

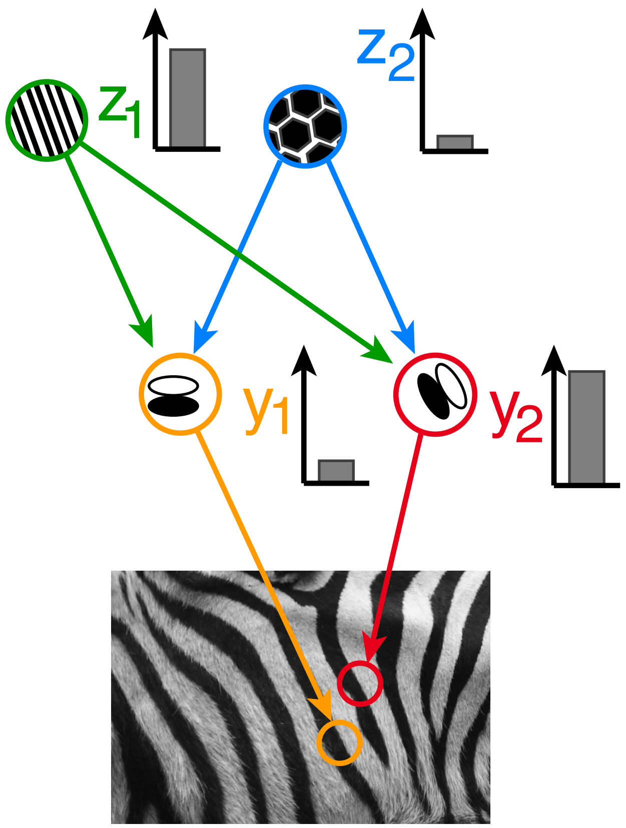

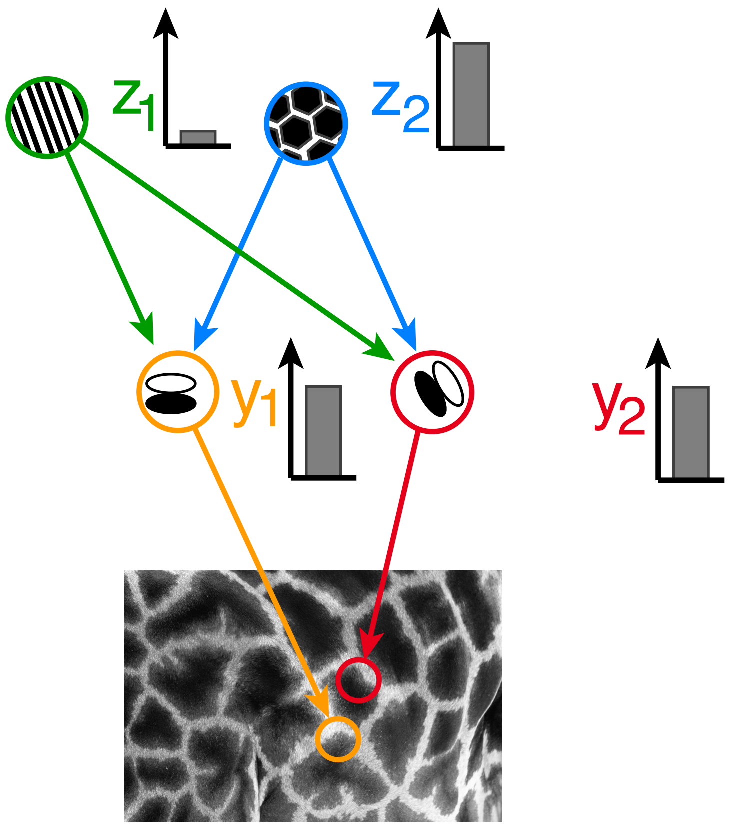

$$\mathrm{P}(\mathbf{x}, \mathbf{y}, \mathbf{z}) = \mathrm{P}(\mathbf{x} \,|\, \mathbf{y})\, \mathrm{P}(\mathbf{y} \,|\, \mathbf{z})\, \mathrm{P}(\mathbf{z})$$

$$\mathrm{P}(\mathbf{y} \,|\, \mathbf{x}, z_i) = \mathrm{P}(\mathbf{y} \,|\, \mathbf{x})\, \mathrm{P}(\mathbf{y} \,|\, z_i)$$

Hierarchical inference in a generative model

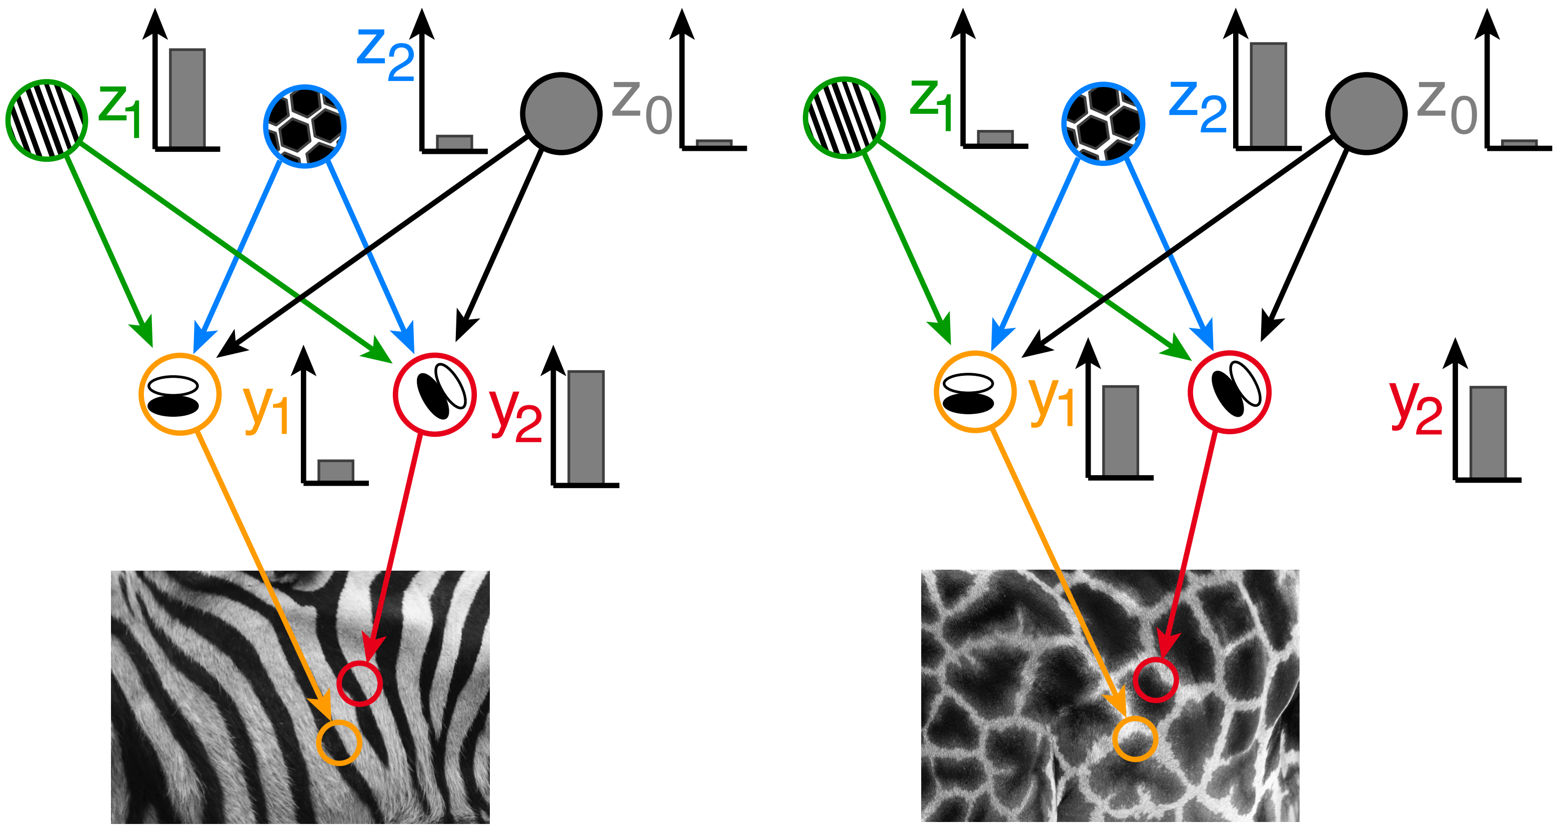

$$\mathrm{P}(\mathbf{x},\mathbf{y},\mathbf{z}) = \mathrm{P}(\mathbf{x}\,|\,\mathbf{y})\,\mathrm{P}(\mathbf{y}\,|\,\mathbf{z})\,\mathrm{P}(\mathbf{z})$$

$$\mathrm{P}(\mathbf{y}\,|\,\mathbf{x},z_i) = \mathrm{P}(\mathbf{y}\,|\,\mathbf{x})\,\mathrm{P}(\mathbf{y}\,|\,z_i)$$

Hierarchical inference in a generative model

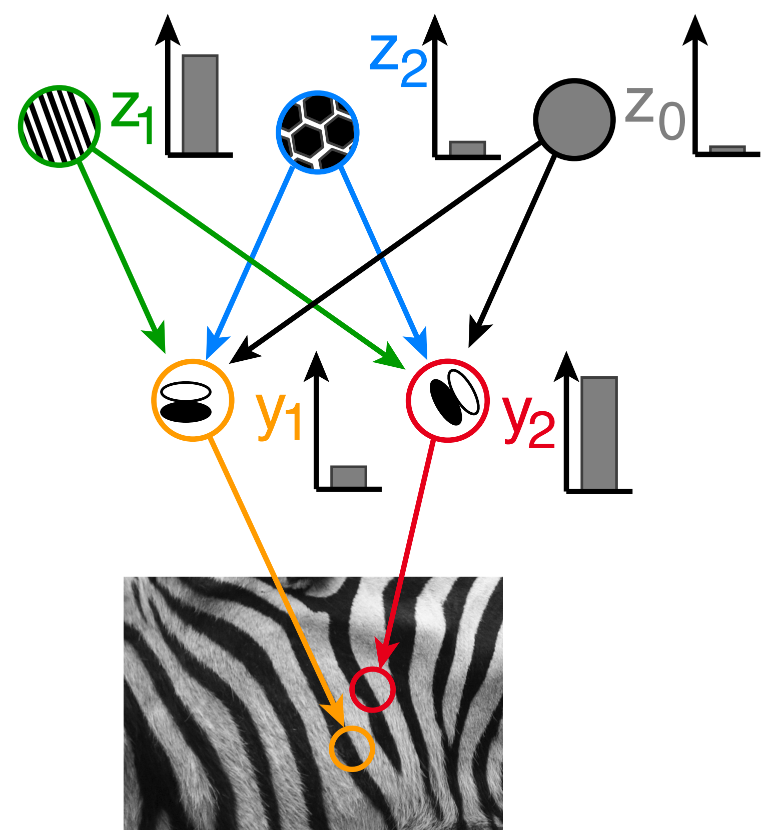

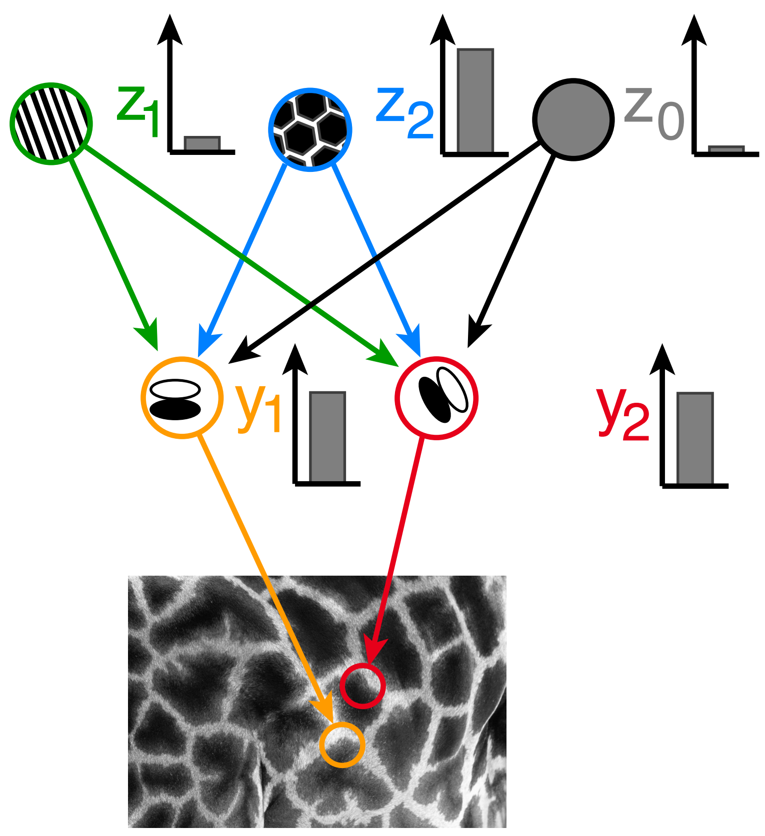

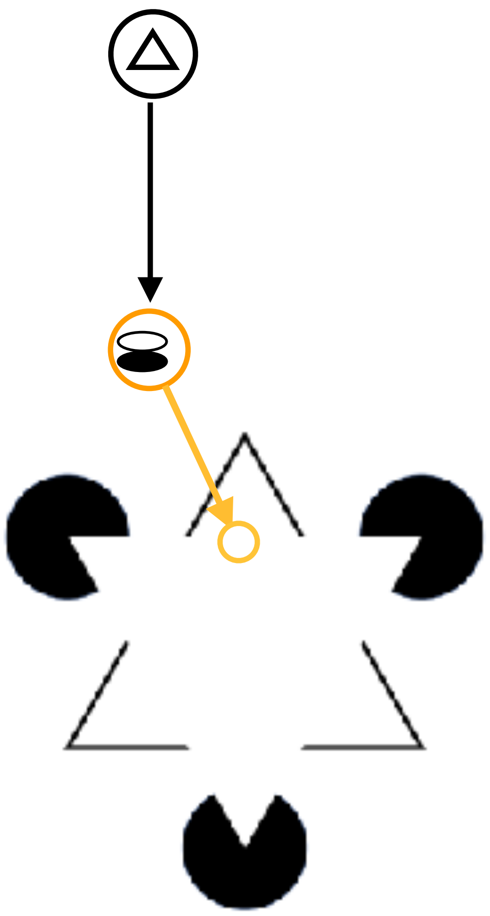

$$\mathrm{P}(\mathbf{x},\mathbf{y},\mathbf{z}) = \mathrm{P}(\mathbf{x}\,|\,\mathbf{y})\,\mathrm{P}(\mathbf{y}\,|\,\mathbf{z})\,\mathrm{P}(\mathbf{z})$$

$$\mathrm{P}(\mathbf{y}\,|\,\mathbf{x},z_i) = \mathrm{P}(\mathbf{y}\,|\,\mathbf{x})\,\mathrm{P}(\mathbf{y}\,|\,z_i)$$

Top-down effects of illusory contours

Dynamics of subjective contour formation in the early visual cortex

Tai Sing Lee*†‡ and My Nguyen*

PNAS | **February 13, 2001** | vol. 98 | no. 4 | **1907–1911**

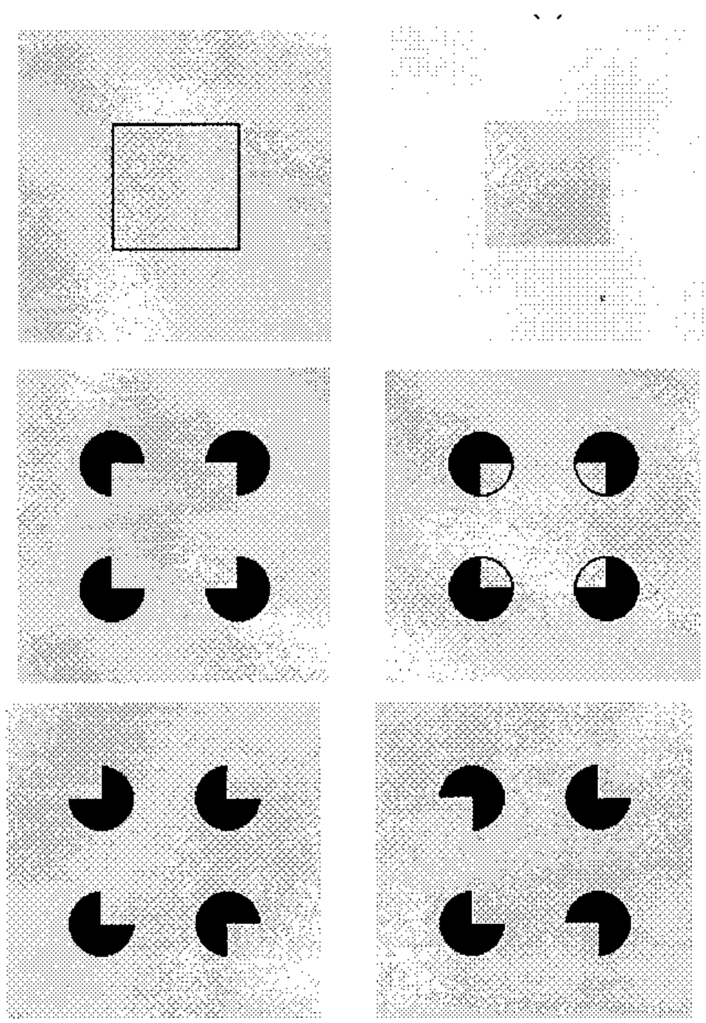

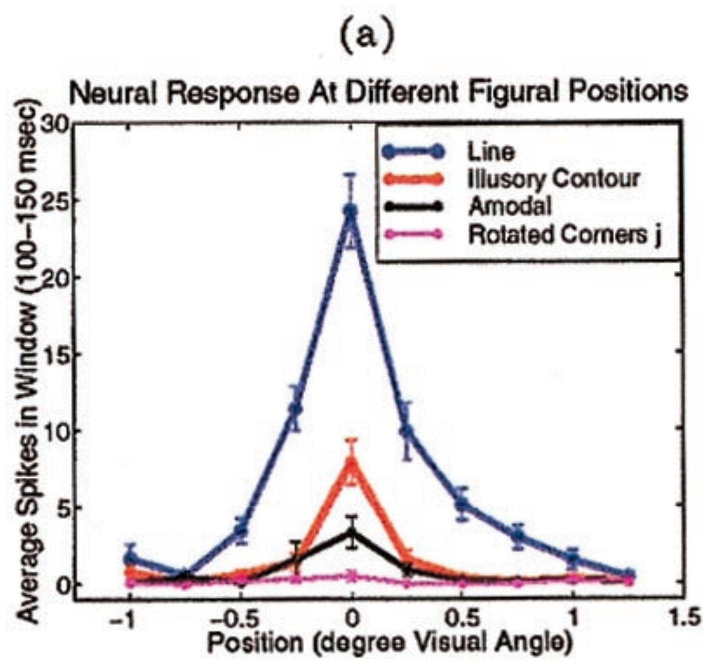

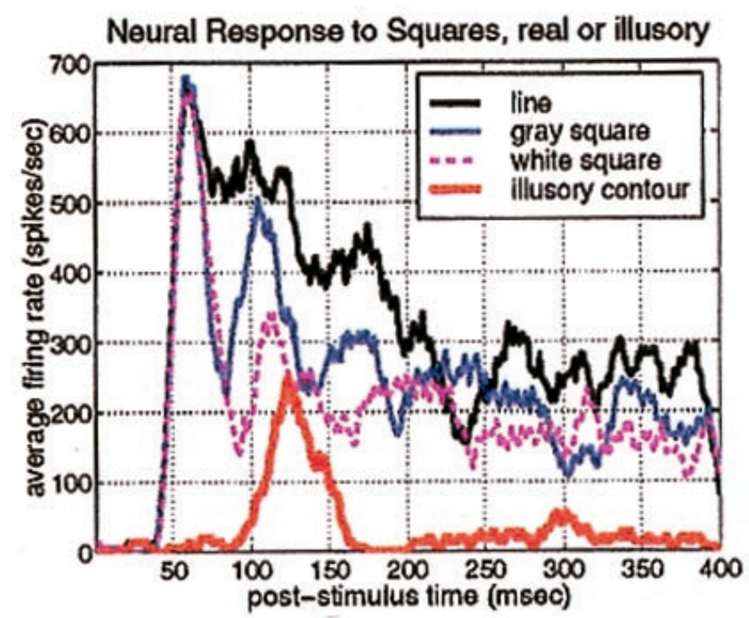

Hierarchical inference in a generative model

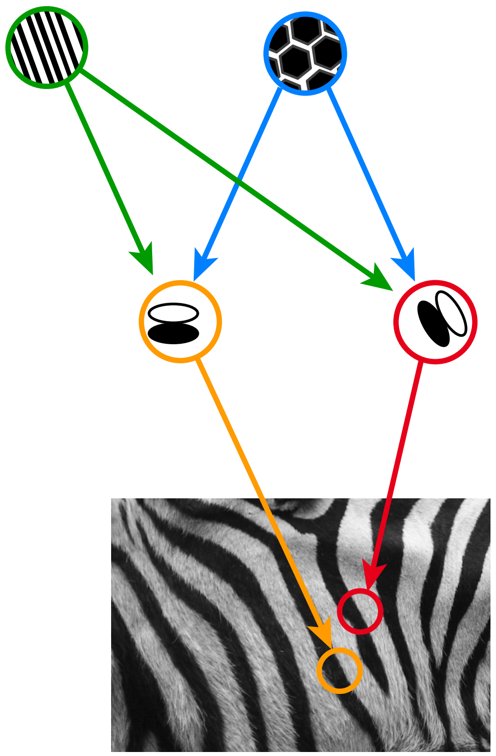

Hierarchical inference in a generative model

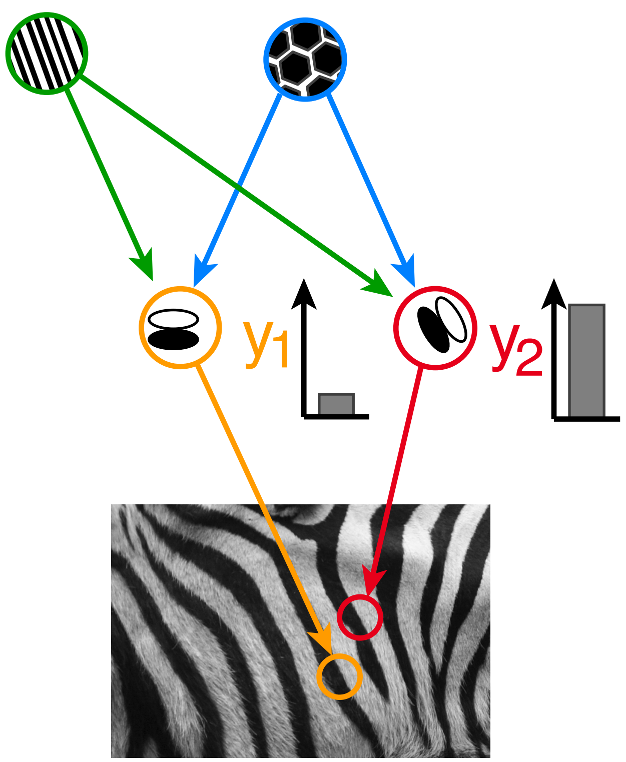

Hierarchical inference in a generative model

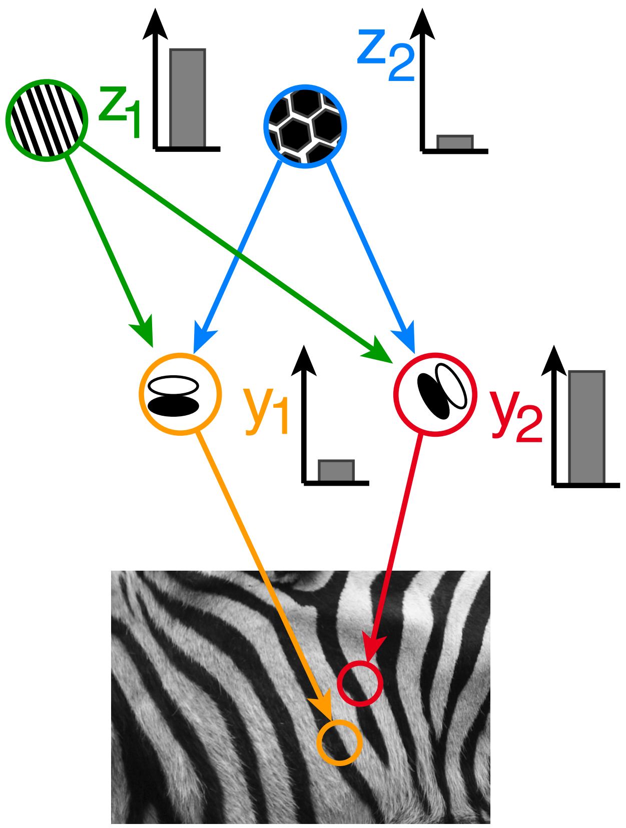

Hierarchical inference in a generative model

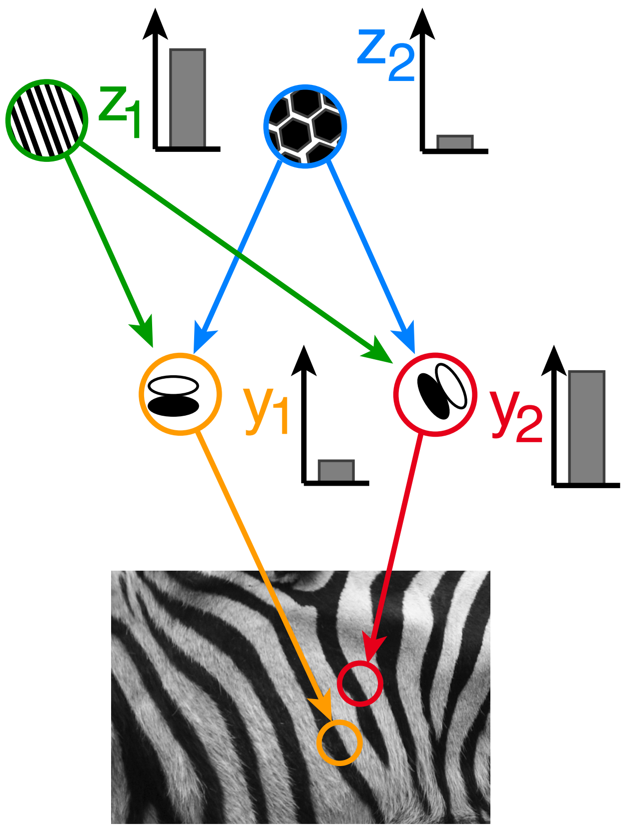

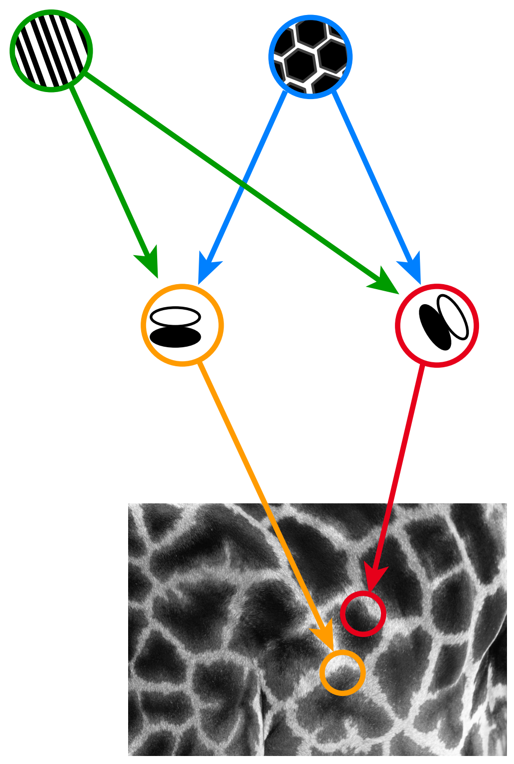

Hierarchical inference in a generative model

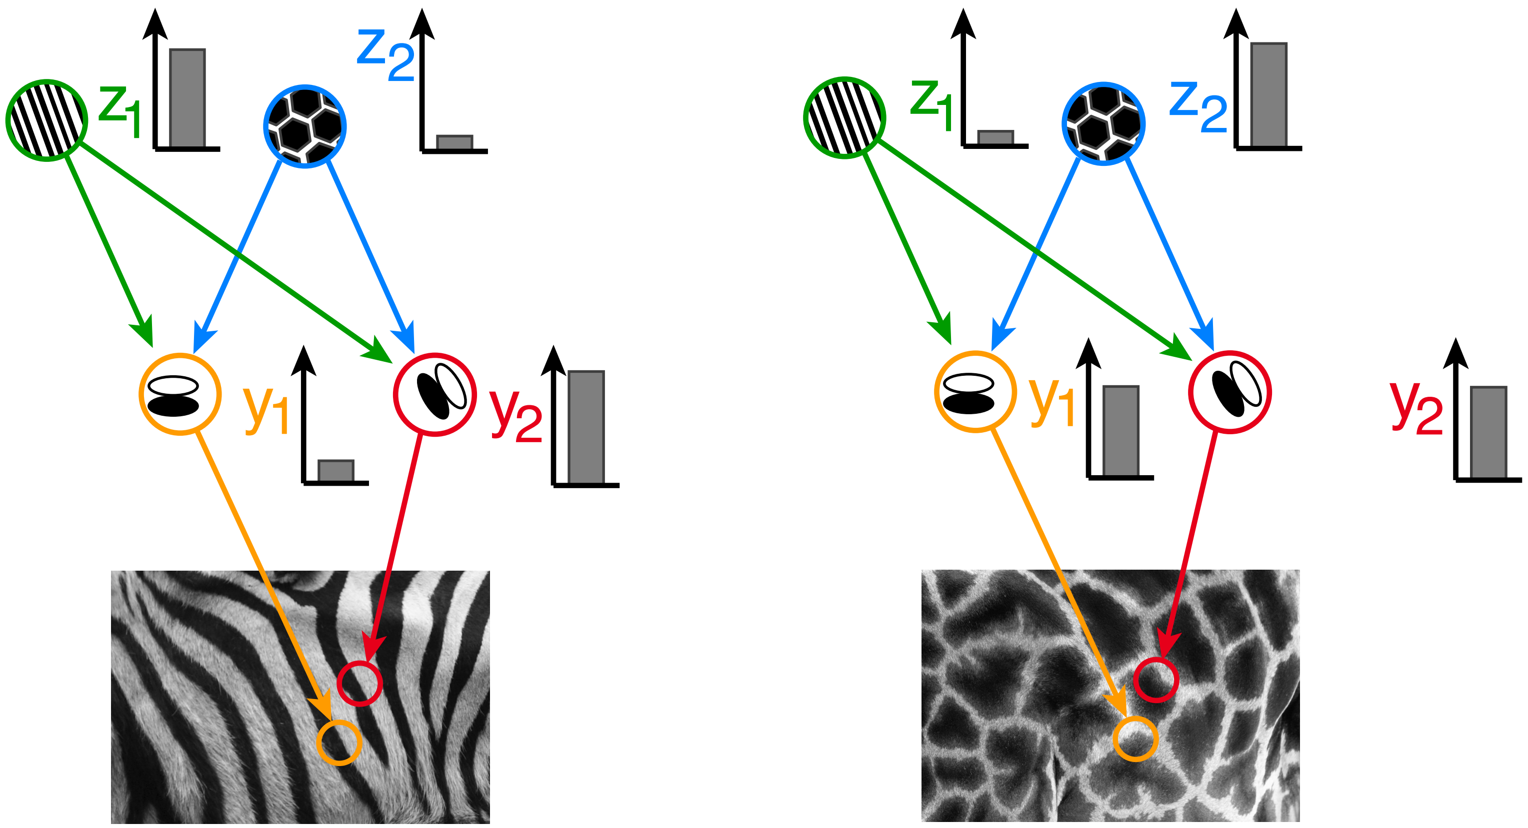

Hierarchical inference in a generative model

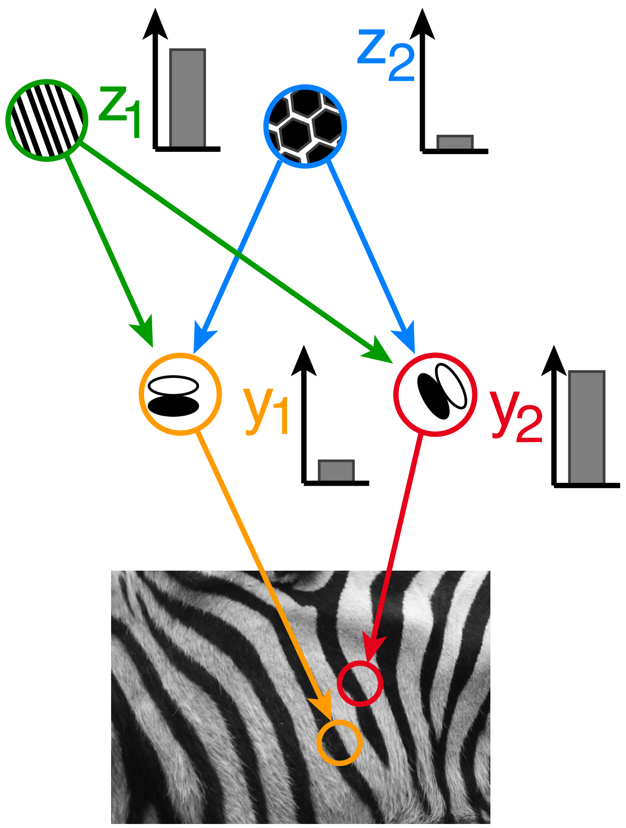

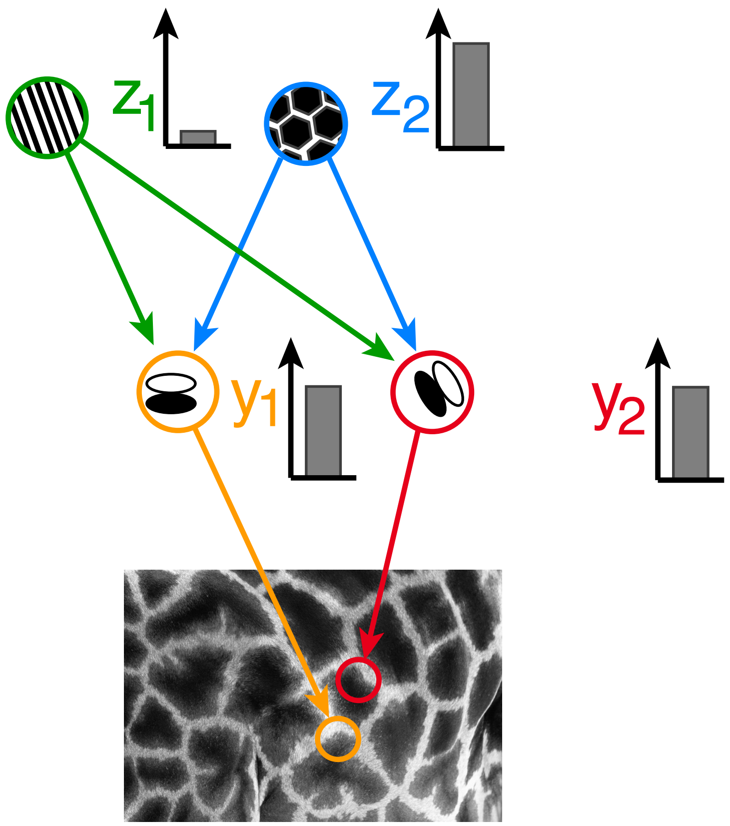

$$\mathrm{P}(\mathbf{x}, \mathbf{y}, \mathbf{z}) = \mathrm{P}(\mathbf{x} \,|\, \mathbf{y})\, \mathrm{P}(\mathbf{y} \,|\, \mathbf{z})\, \mathrm{P}(\mathbf{z})$$

$$\mathrm{P}(\mathbf{y} \,|\, \mathbf{x}, z_i) = \mathrm{P}(\mathbf{y} \,|\, \mathbf{x})\, \mathrm{P}(\mathbf{y} \,|\, z_i)$$

Hierarchical inference in a generative model

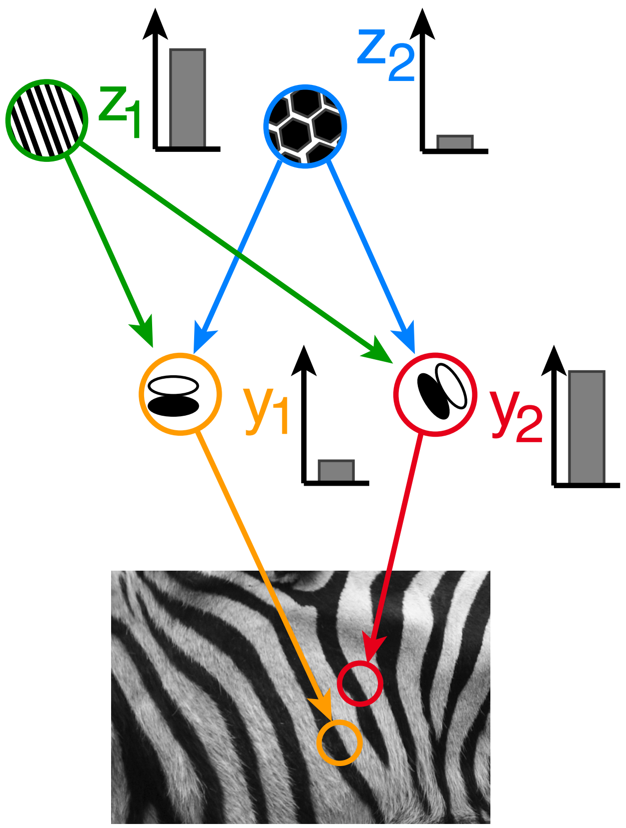

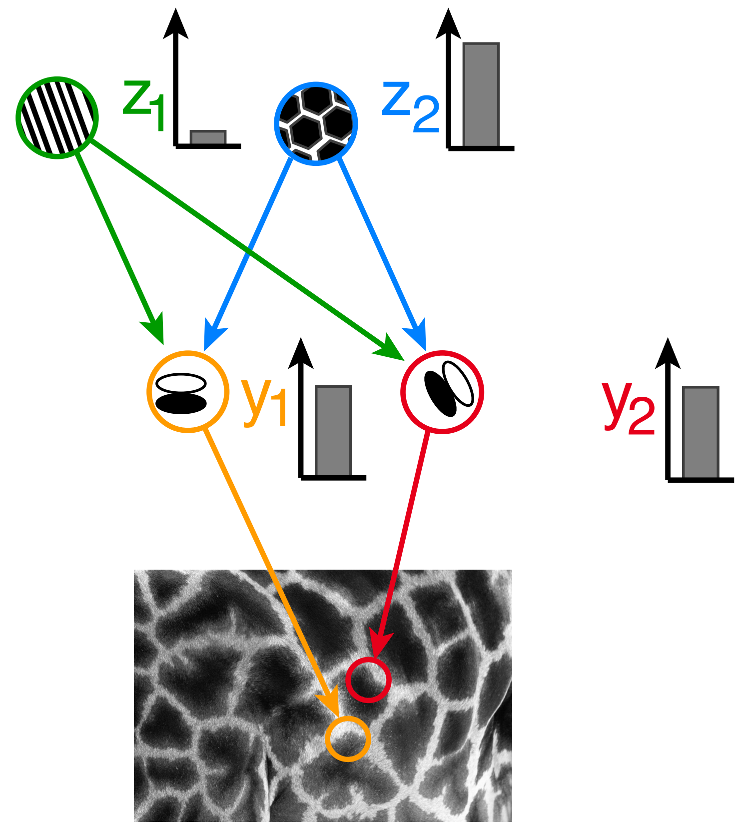

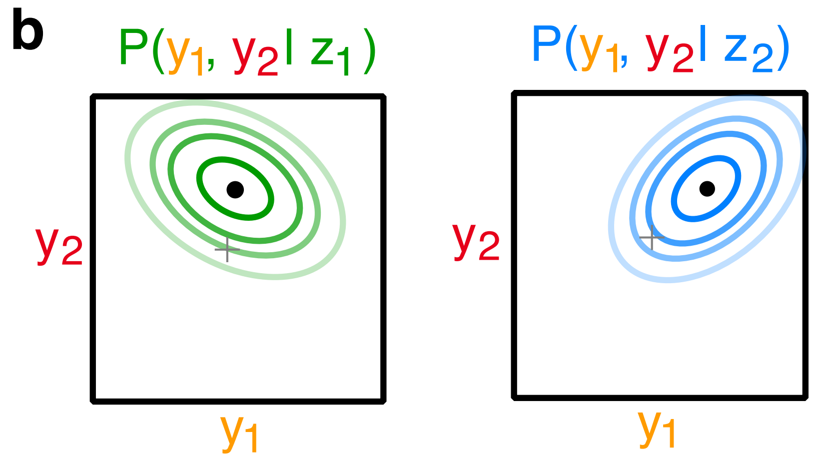

$$\mathrm{P}(\mathbf{x}, \mathbf{y}, \mathbf{z}) = \mathrm{P}(\mathbf{x} \,|\, \mathbf{y})\, \mathrm{P}(\mathbf{y} \,|\, \mathbf{z})\, \mathrm{P}(\mathbf{z})$$

$$\mathrm{P}(\mathbf{y} \,|\, \mathbf{x}, z_i) = \mathrm{P}(\mathbf{y} \,|\, \mathbf{x})\, \mathrm{P}(\mathbf{y} \,|\, z_i)$$

Hierarchical inference in a generative model

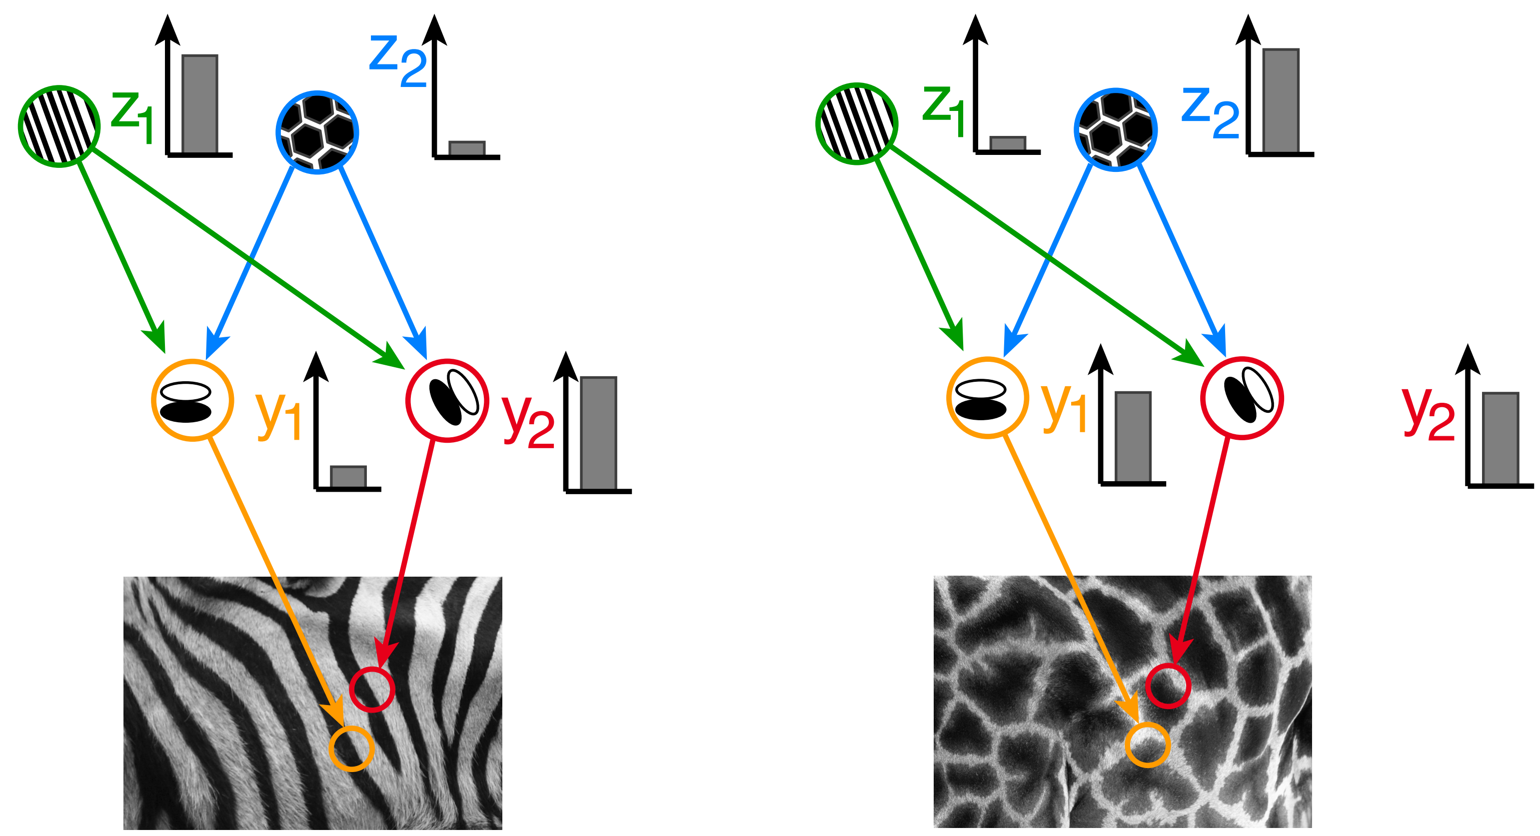

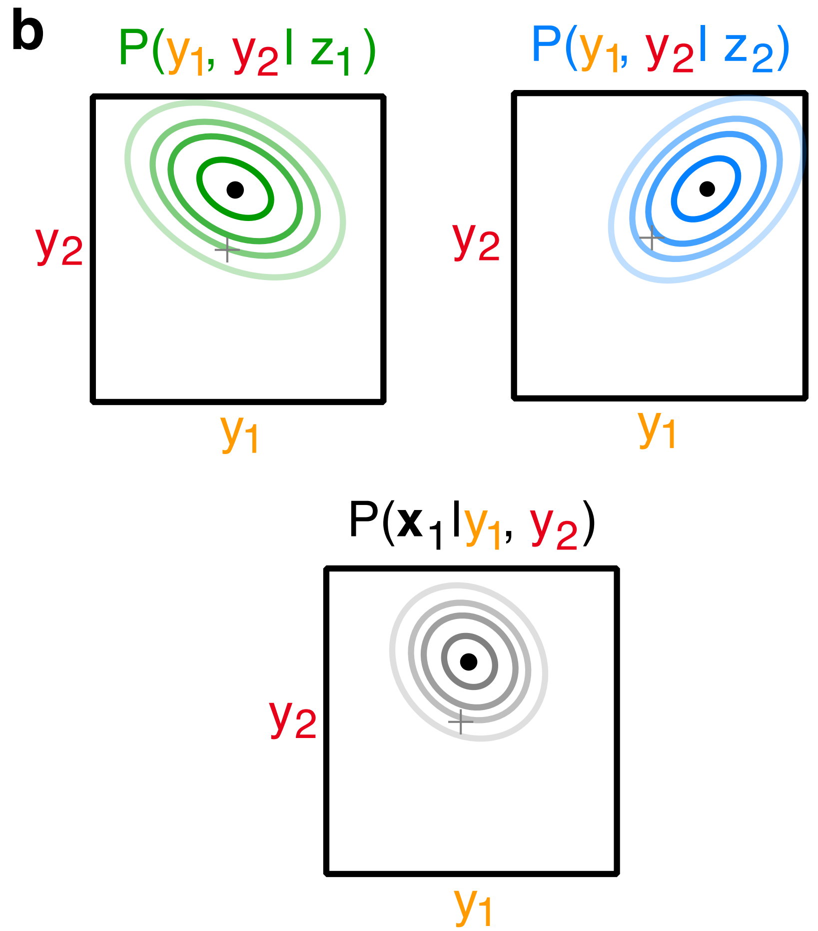

$$\mathrm{P}(\mathbf{x}, \mathbf{y}, \mathbf{z}) = \mathrm{P}(\mathbf{x} \,|\, \mathbf{y})\, \mathrm{P}(\mathbf{y} \,|\, \mathbf{z})\, \mathrm{P}(\mathbf{z})$$

$$\mathrm{P}(\mathbf{y} \,|\, \mathbf{x}, z_i) = \mathrm{P}(\mathbf{y} \,|\, \mathbf{x})\, \mathrm{P}(\mathbf{y} \,|\, z_i)$$

Hierarchical inference in a generative model

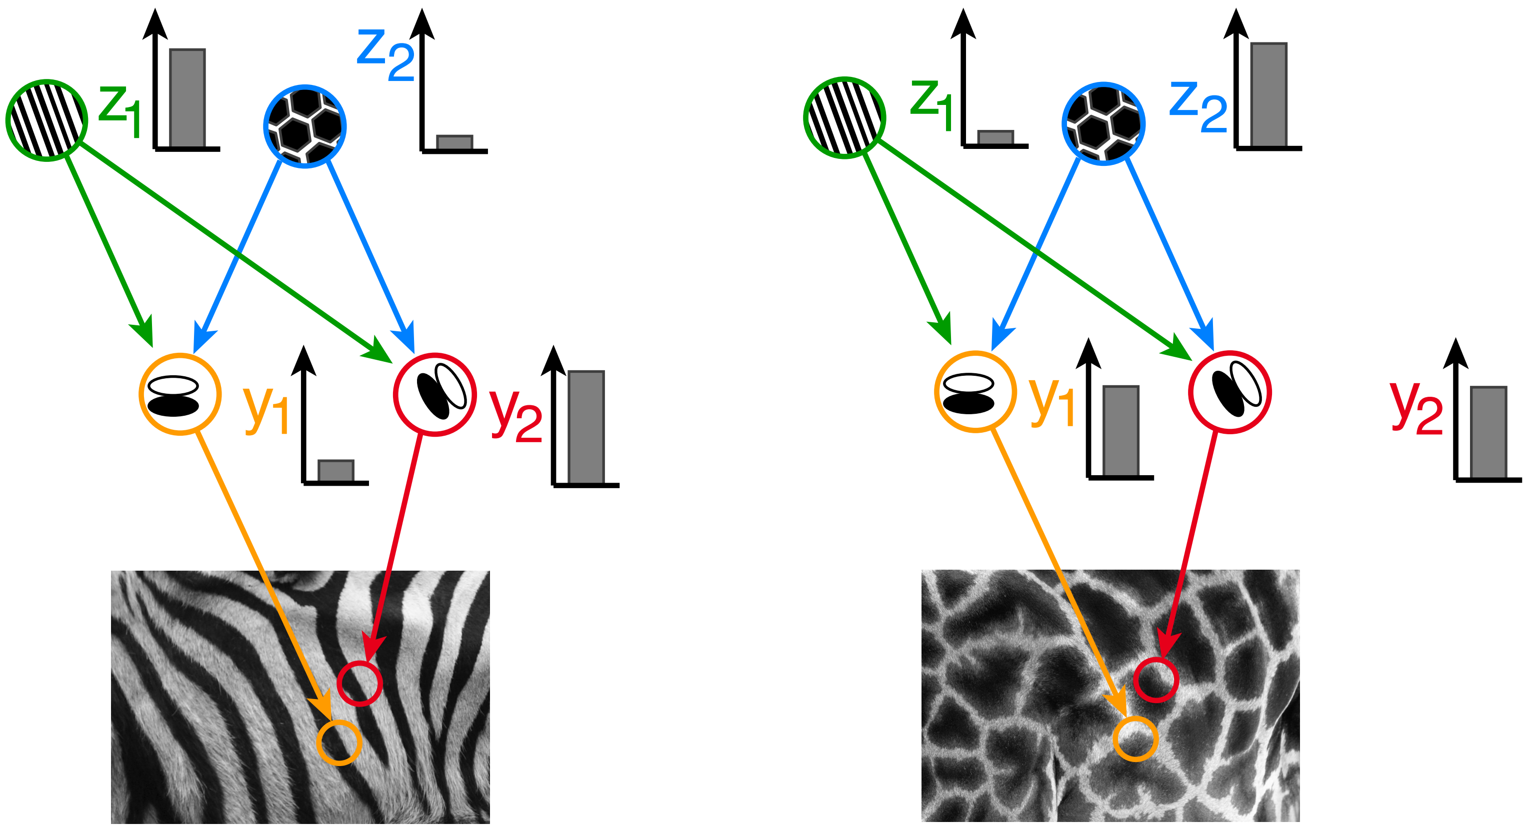

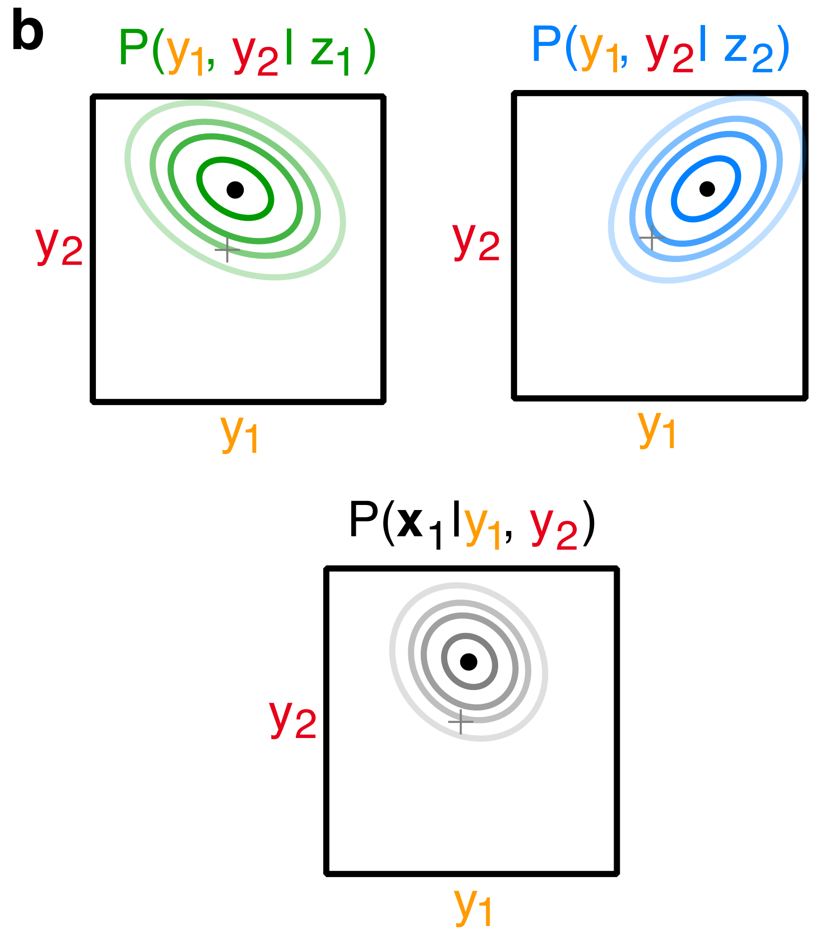

$$\mathrm{P}(\mathbf{x},\mathbf{y},\mathbf{z}) = \mathrm{P}(\mathbf{x}\,|\,\mathbf{y})\,\mathrm{P}(\mathbf{y}\,|\,\mathbf{z})\,\mathrm{P}(\mathbf{z})$$

$$\mathrm{P}(\mathbf{y}\,|\,\mathbf{x},z_i) = \mathrm{P}(\mathbf{y}\,|\,\mathbf{x})\,\mathrm{P}(\mathbf{y}\,|\,z_i)$$

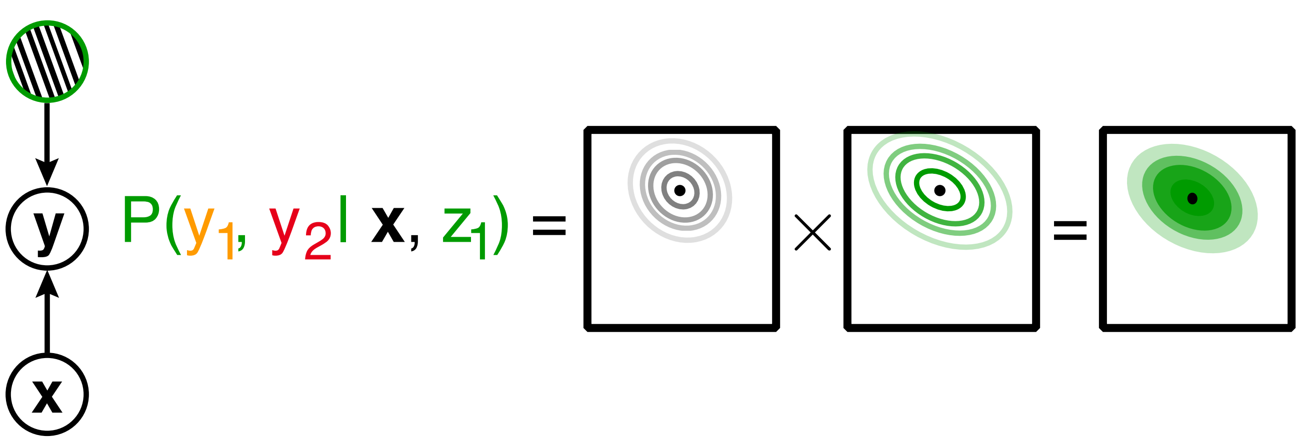

Hierarchical inference in a generative model

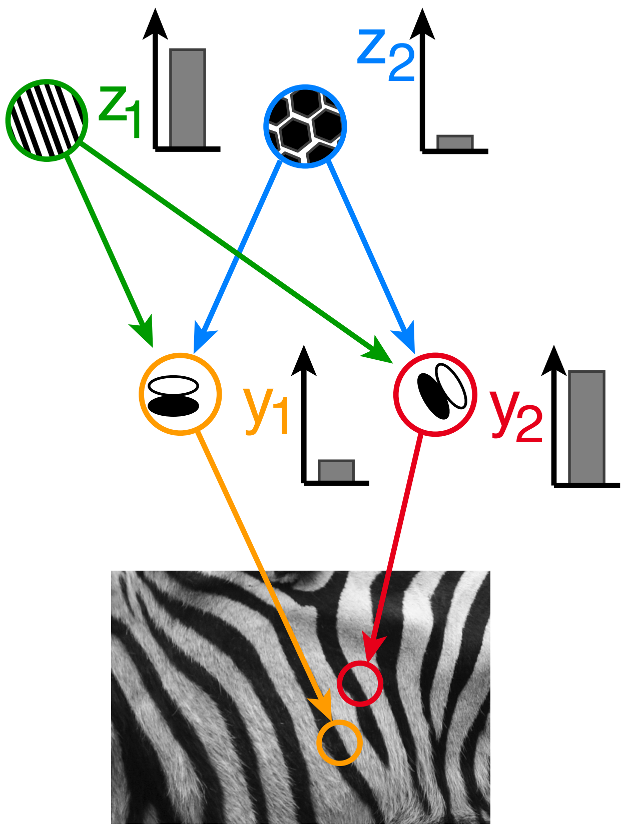

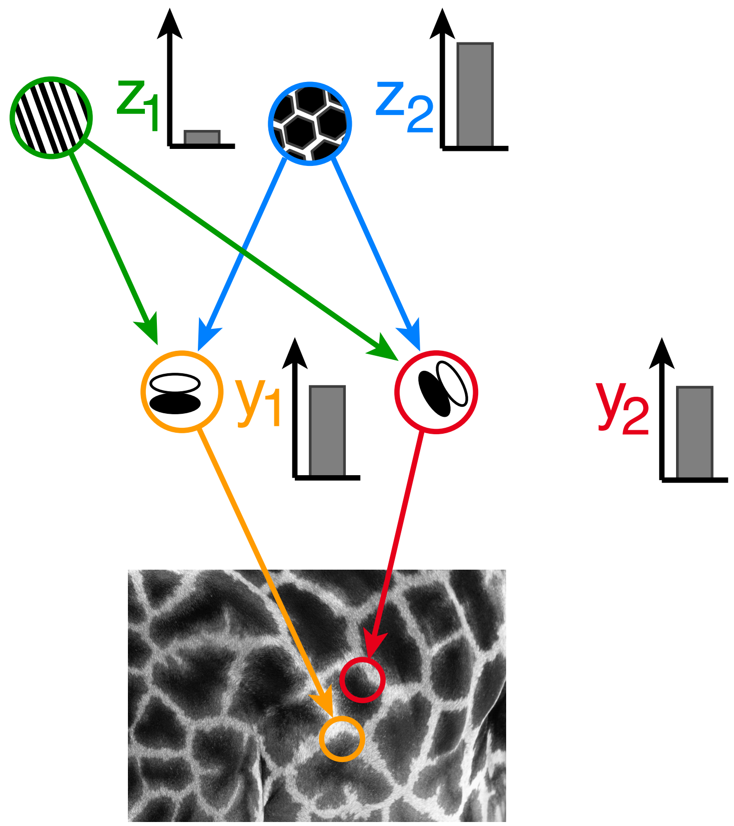

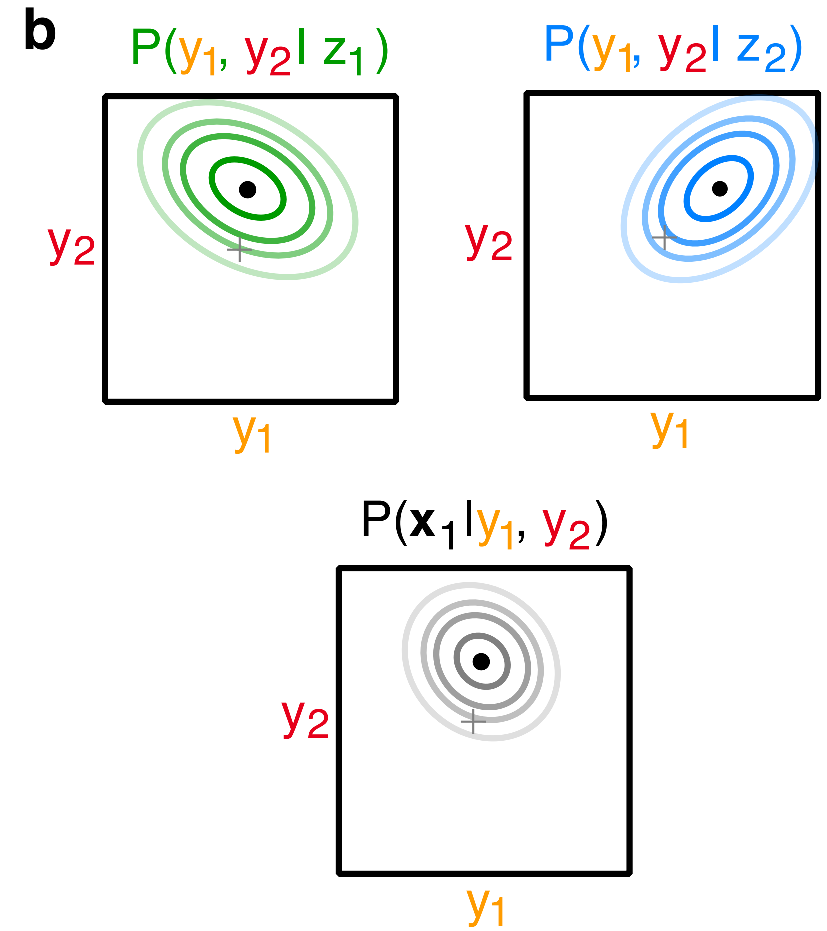

$$P(\mathbf{x}, \mathbf{y}, \mathbf{z}) = P(\mathbf{x} \,|\, \mathbf{y})\, P(\mathbf{y} \,|\, \mathbf{z})\, P(\mathbf{z})$$

$$P(\mathbf{y} \,|\, \mathbf{x}, z_i) = P(\mathbf{y} \,|\, \mathbf{x})\, P(\mathbf{y} \,|\, z_i)$$

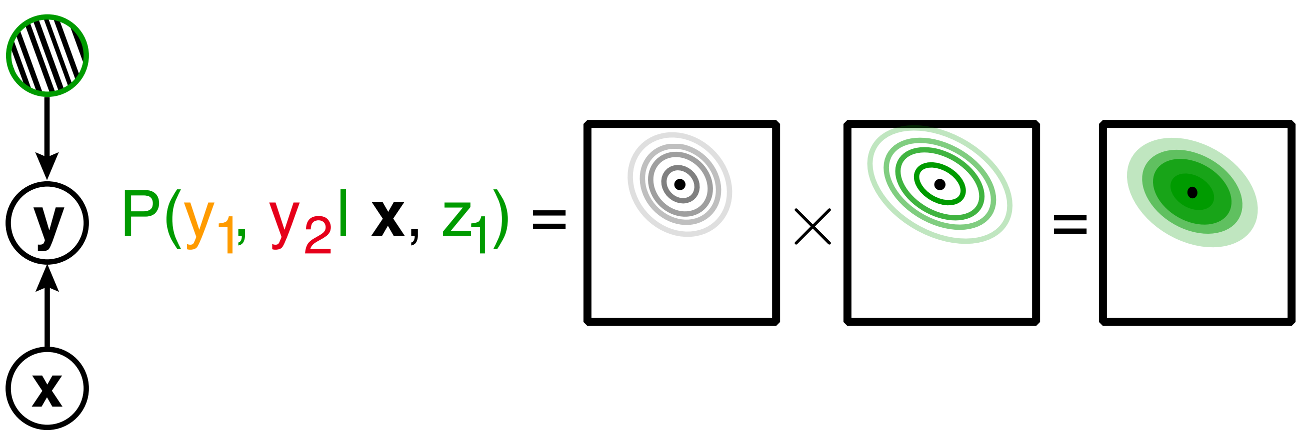

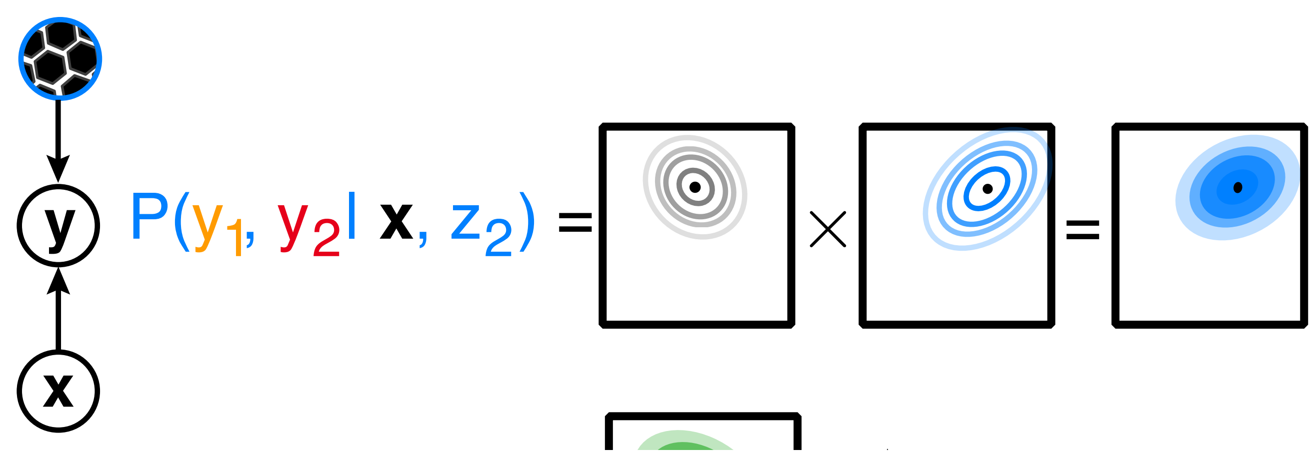

Hierarchical inference in a generative model

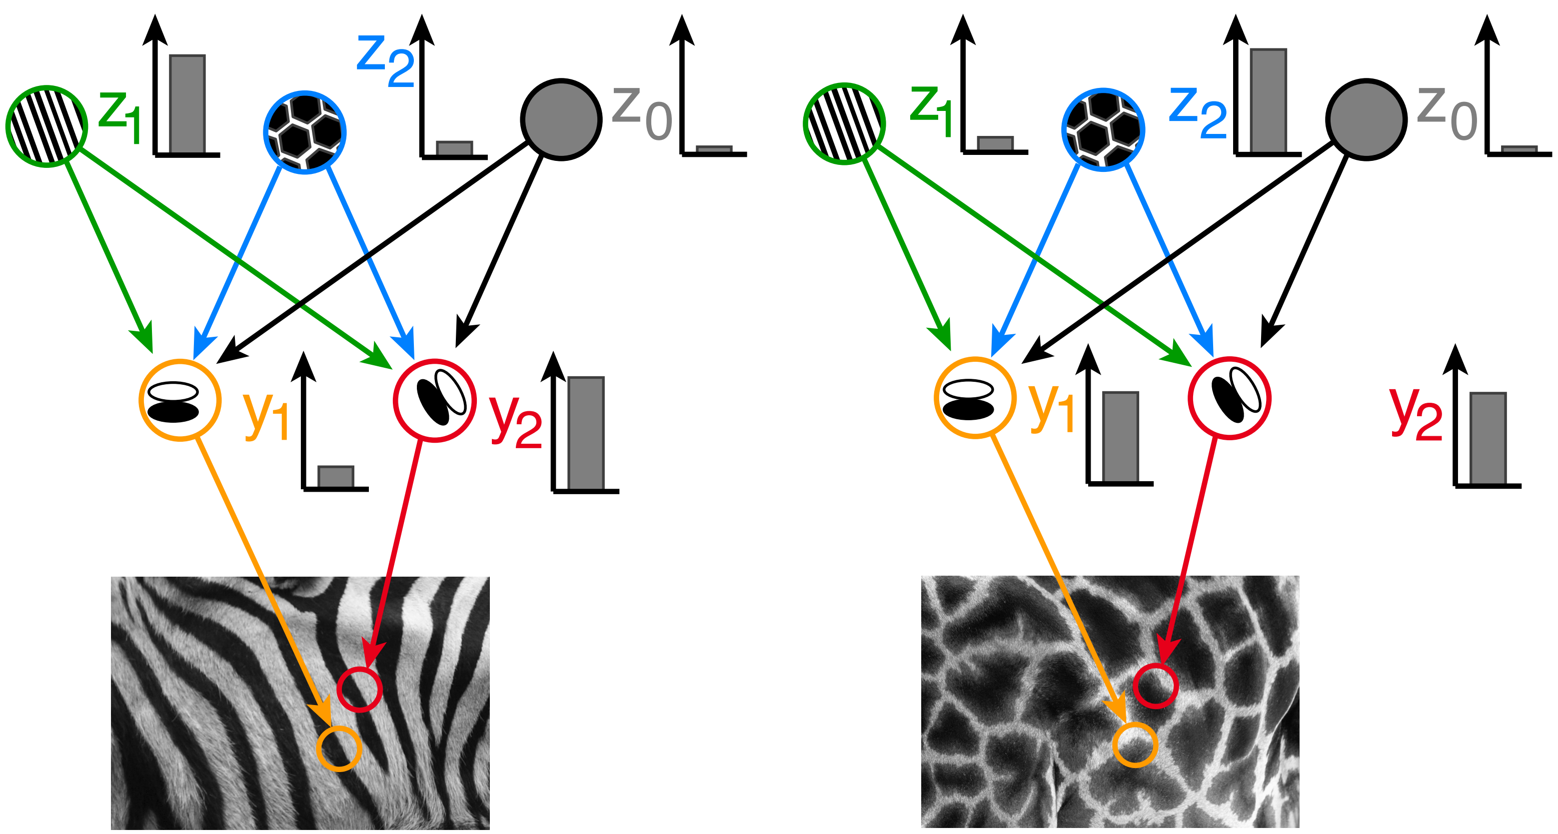

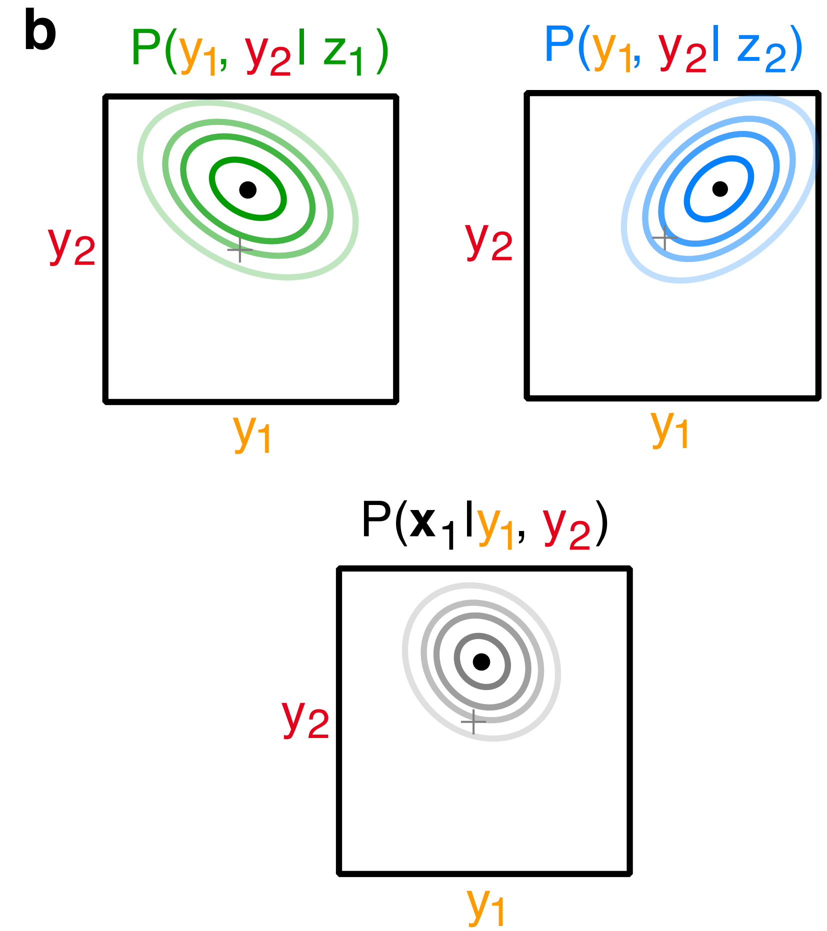

$$P(\mathbf{x},\mathbf{y},\mathbf{z}) = P(\mathbf{x}\,|\,\mathbf{y})\,P(\mathbf{y}\,|\,\mathbf{z})\,P(\mathbf{z})$$

$$P(\mathbf{y}\,|\,\mathbf{x},z_i) = P(\mathbf{y}\,|\,\mathbf{x})\,P(\mathbf{y}\,|\,z_i)$$

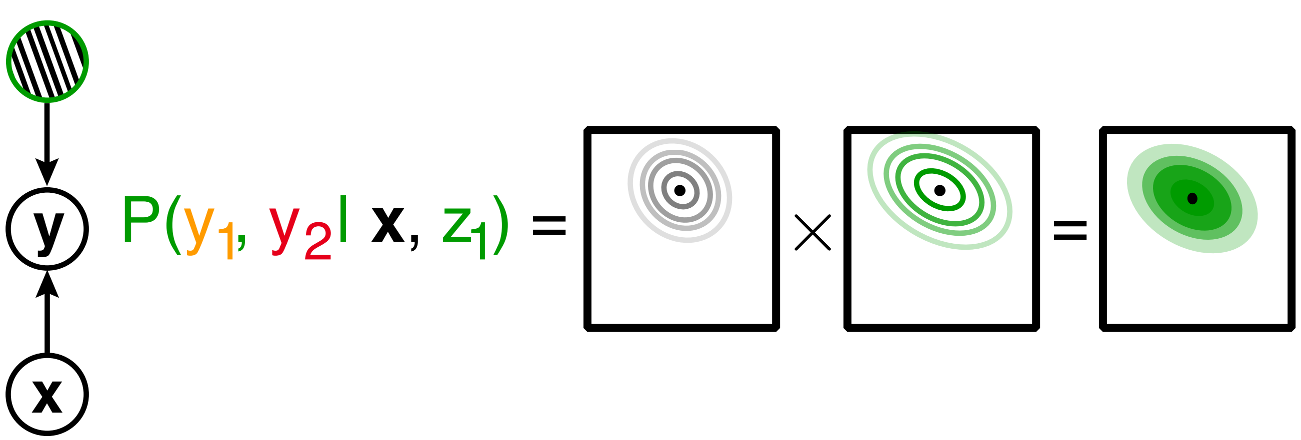

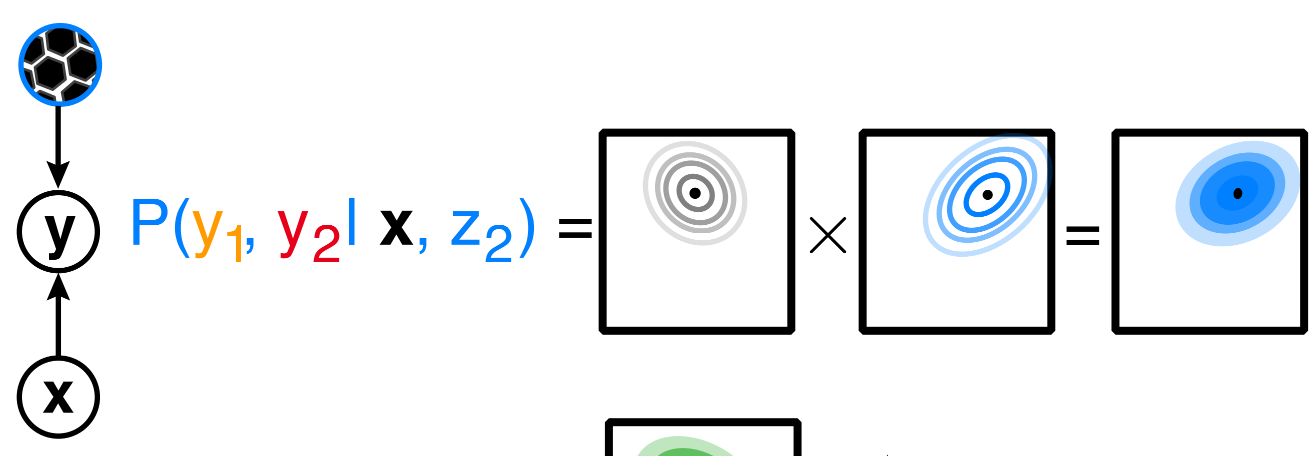

Hierarchical inference in a generative model

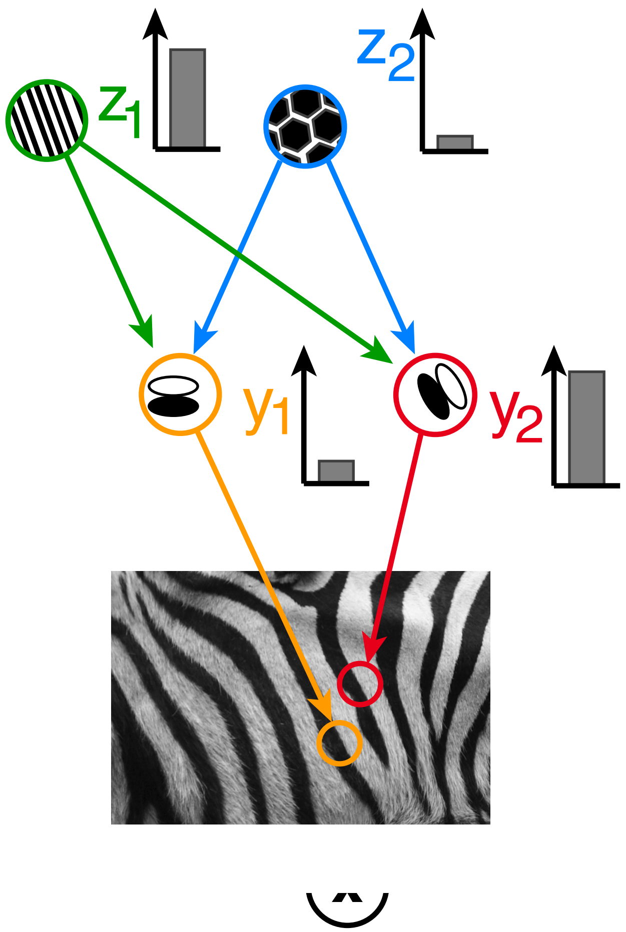

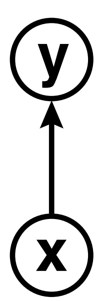

$P(y_1, y_2 | \mathbf{x}) =$

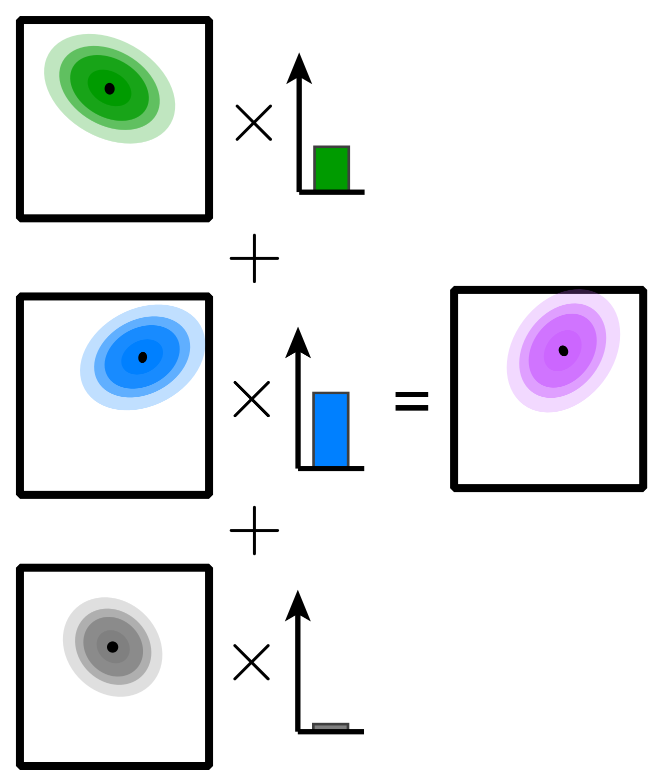

Hierarchical inference in a generative model

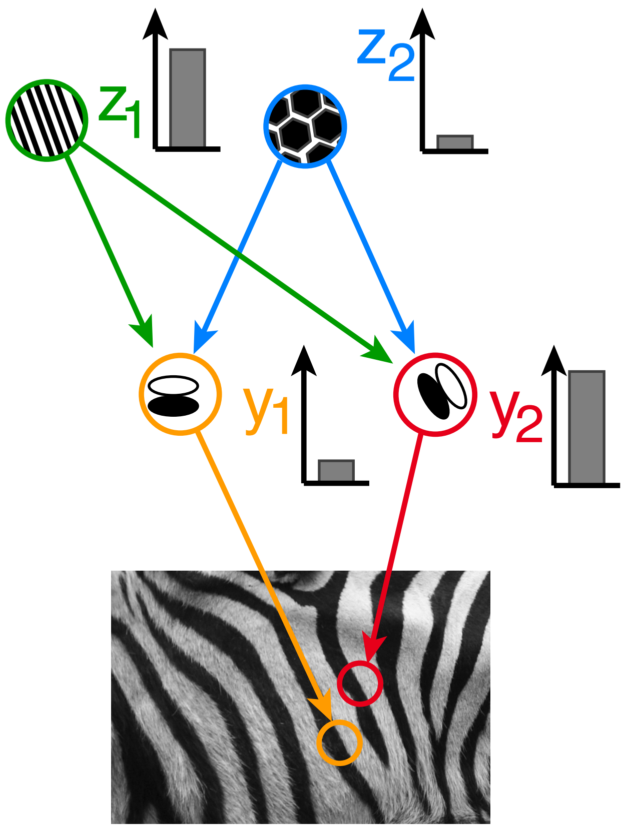

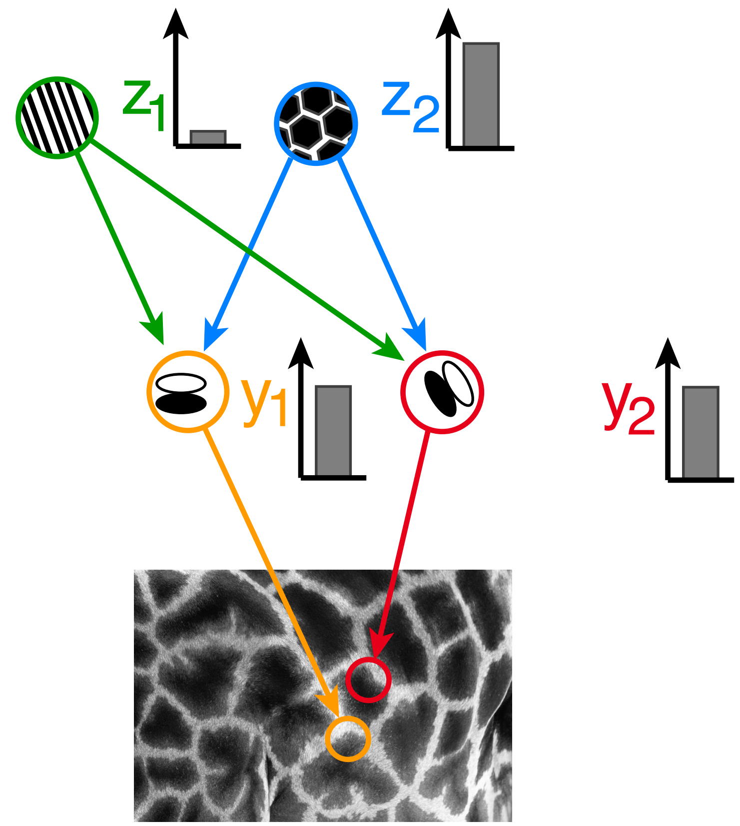

$P(y_1, y_2 | \mathbf{x}) =$

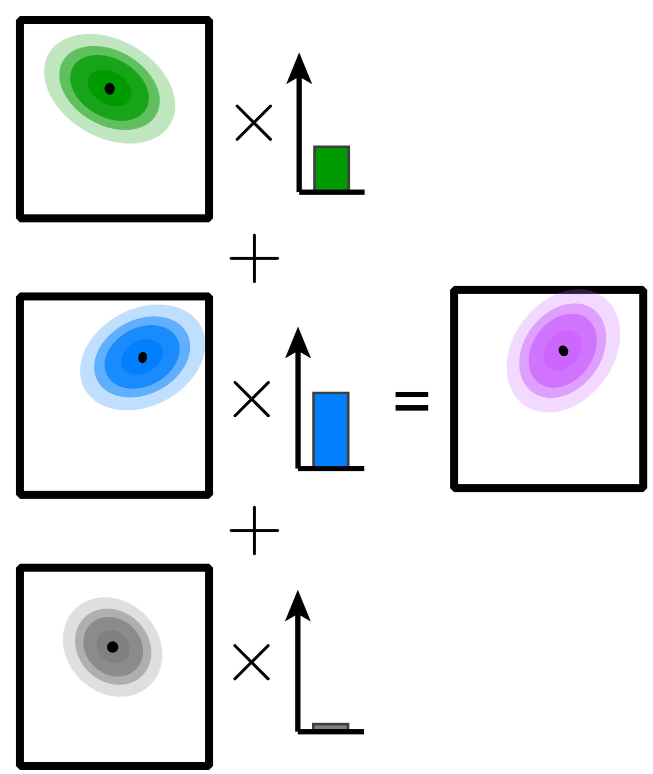

Hierarchical inference in a generative model

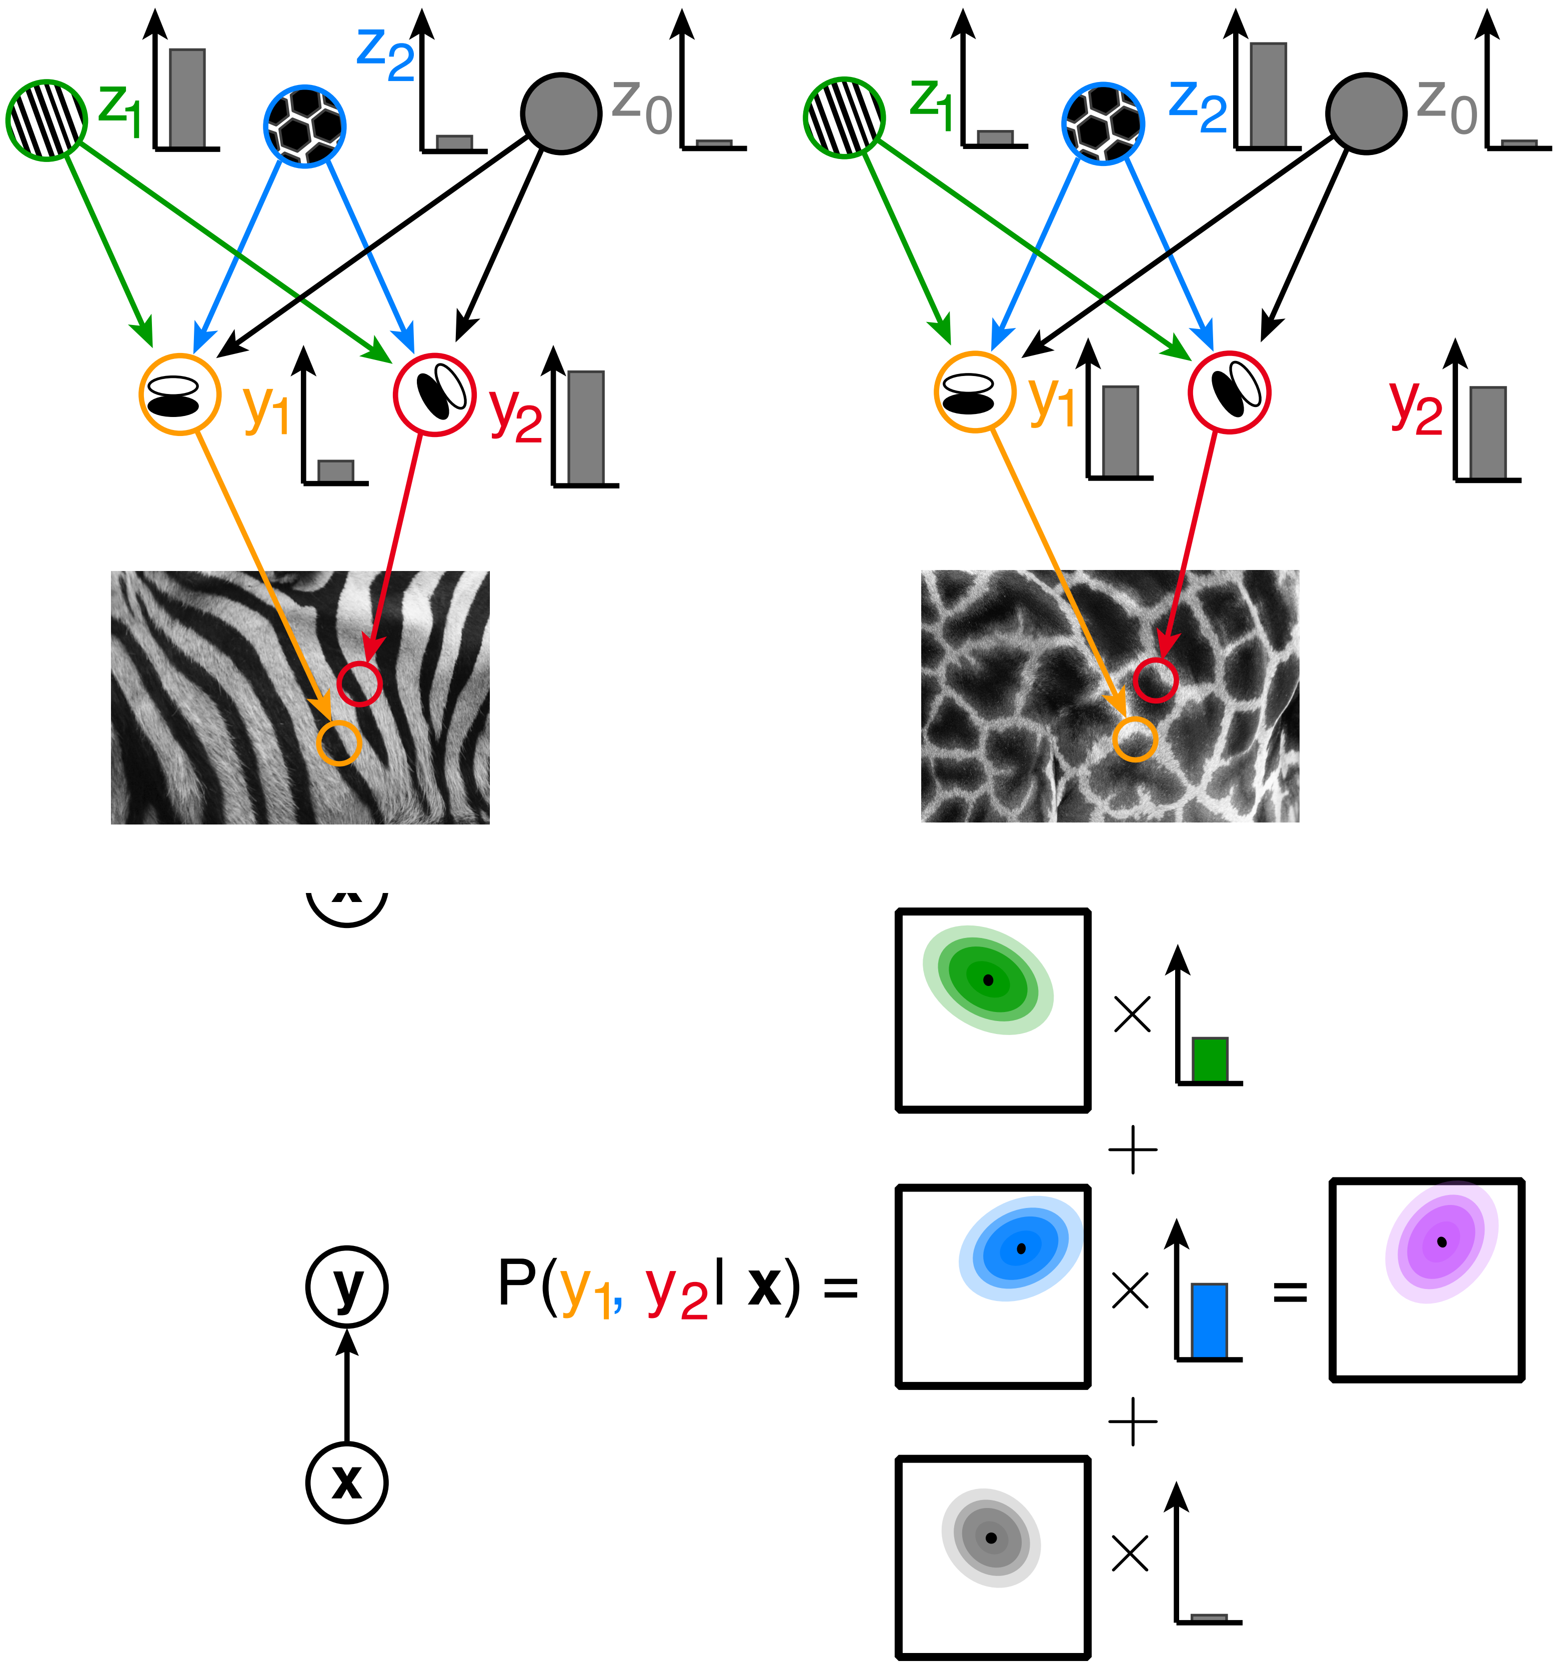

Hierarchical inference in a generative model

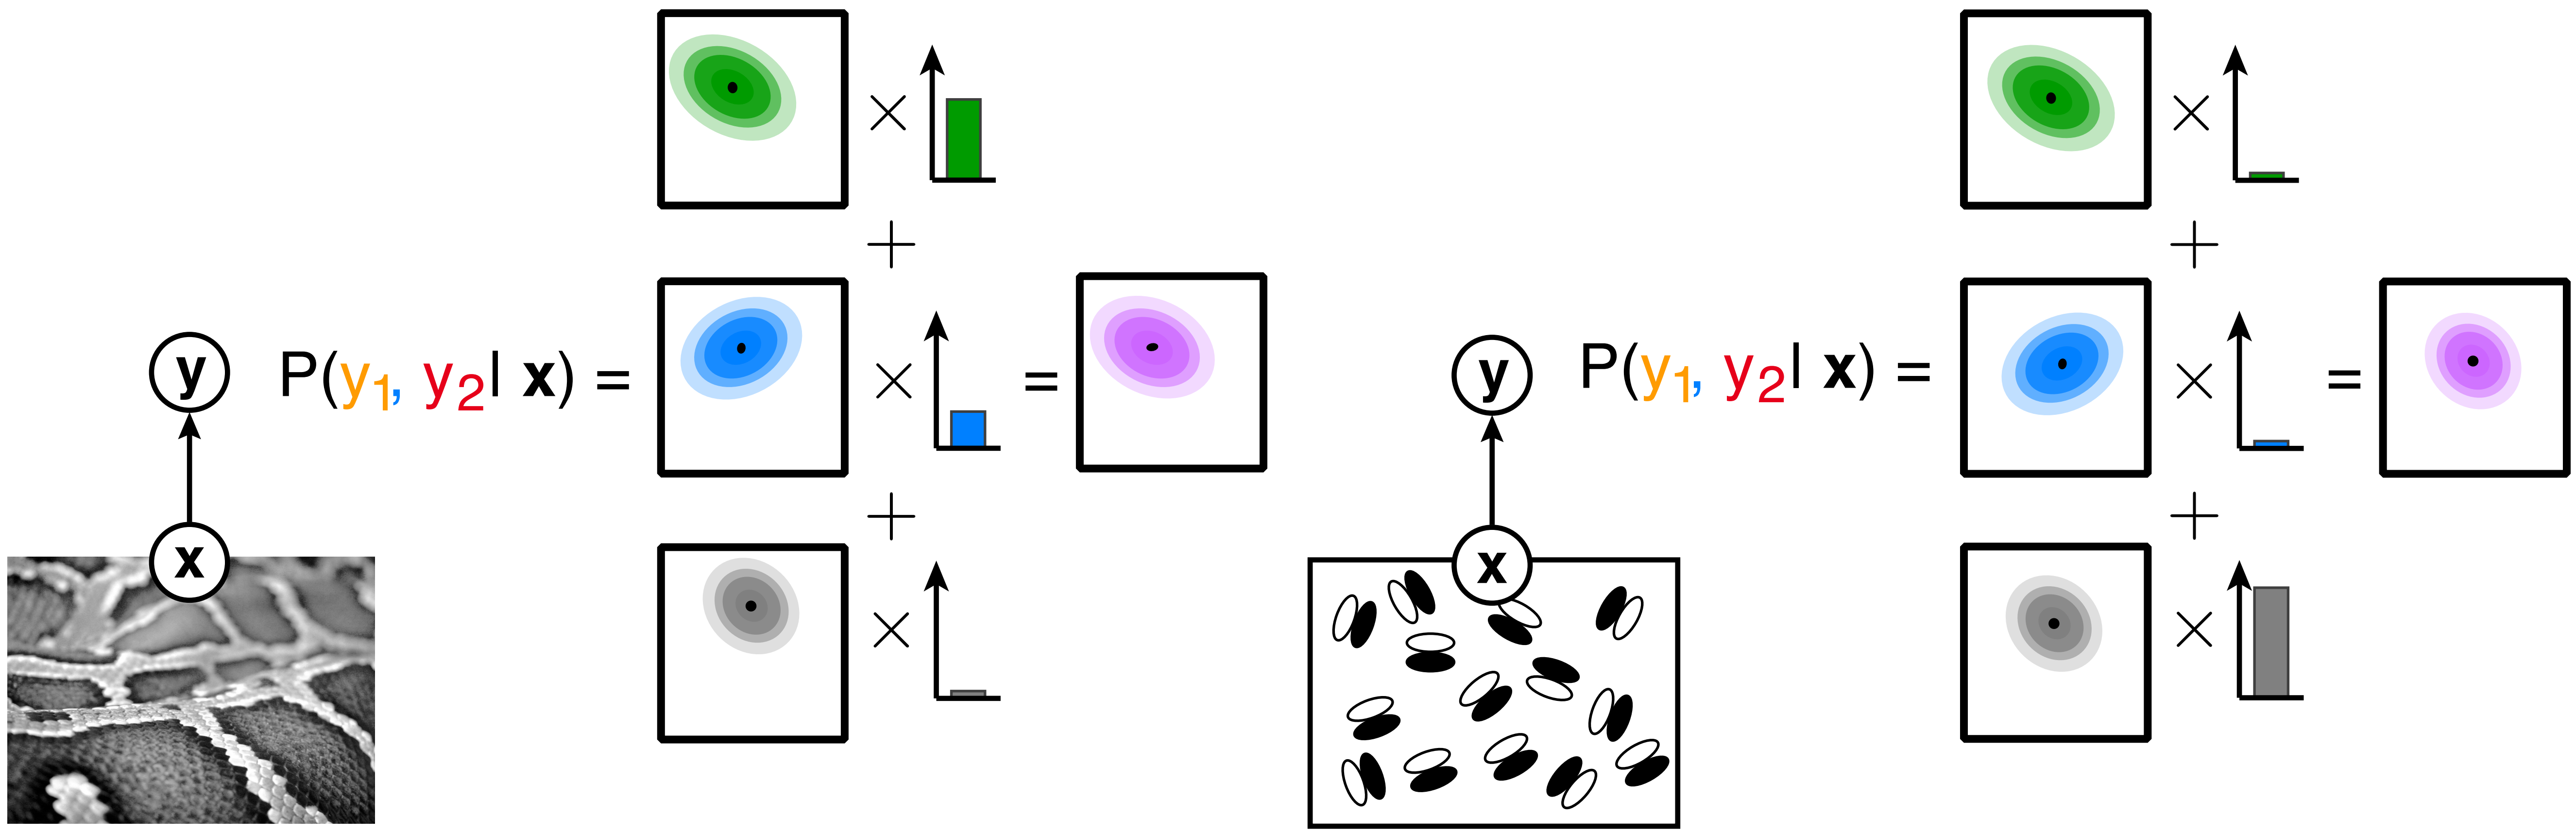

Hierarchical inference in a generative model

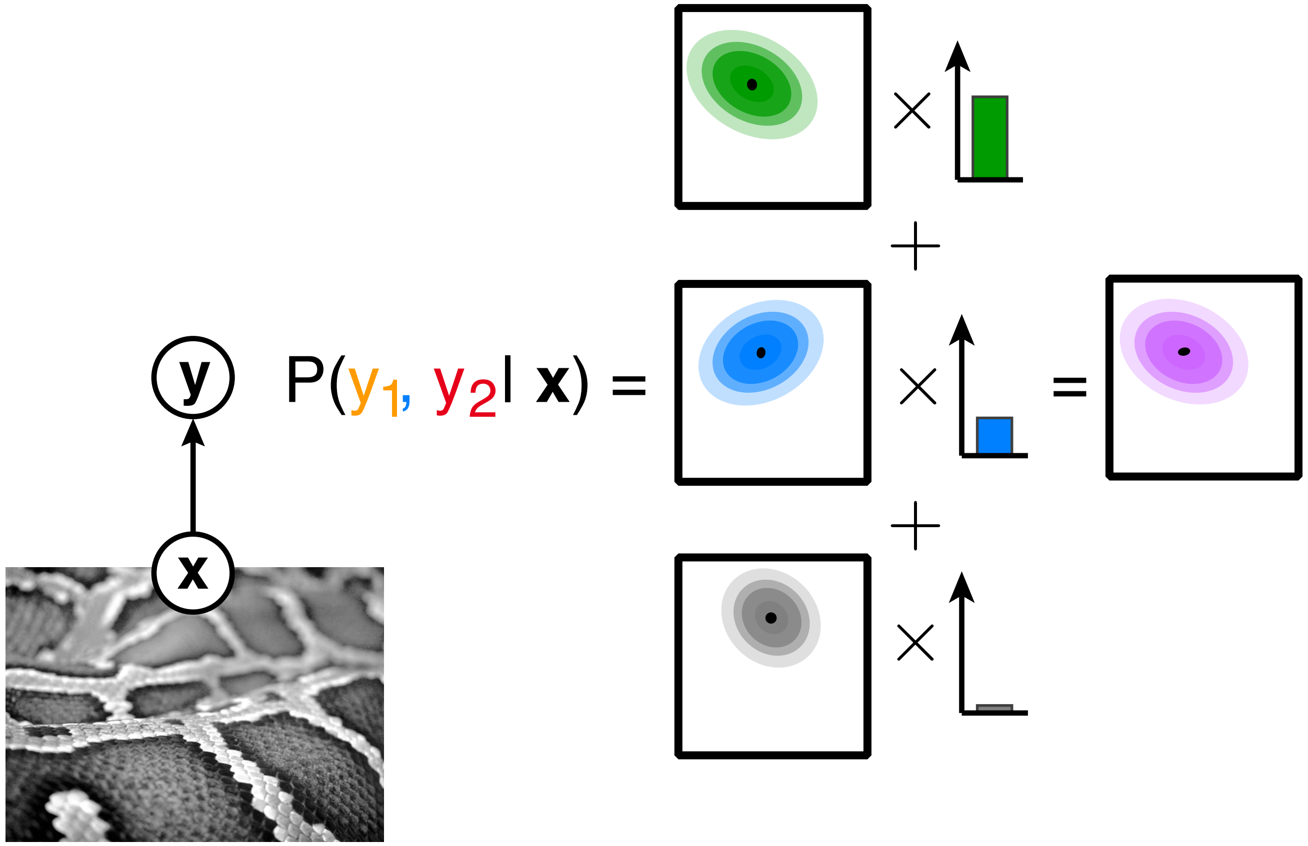

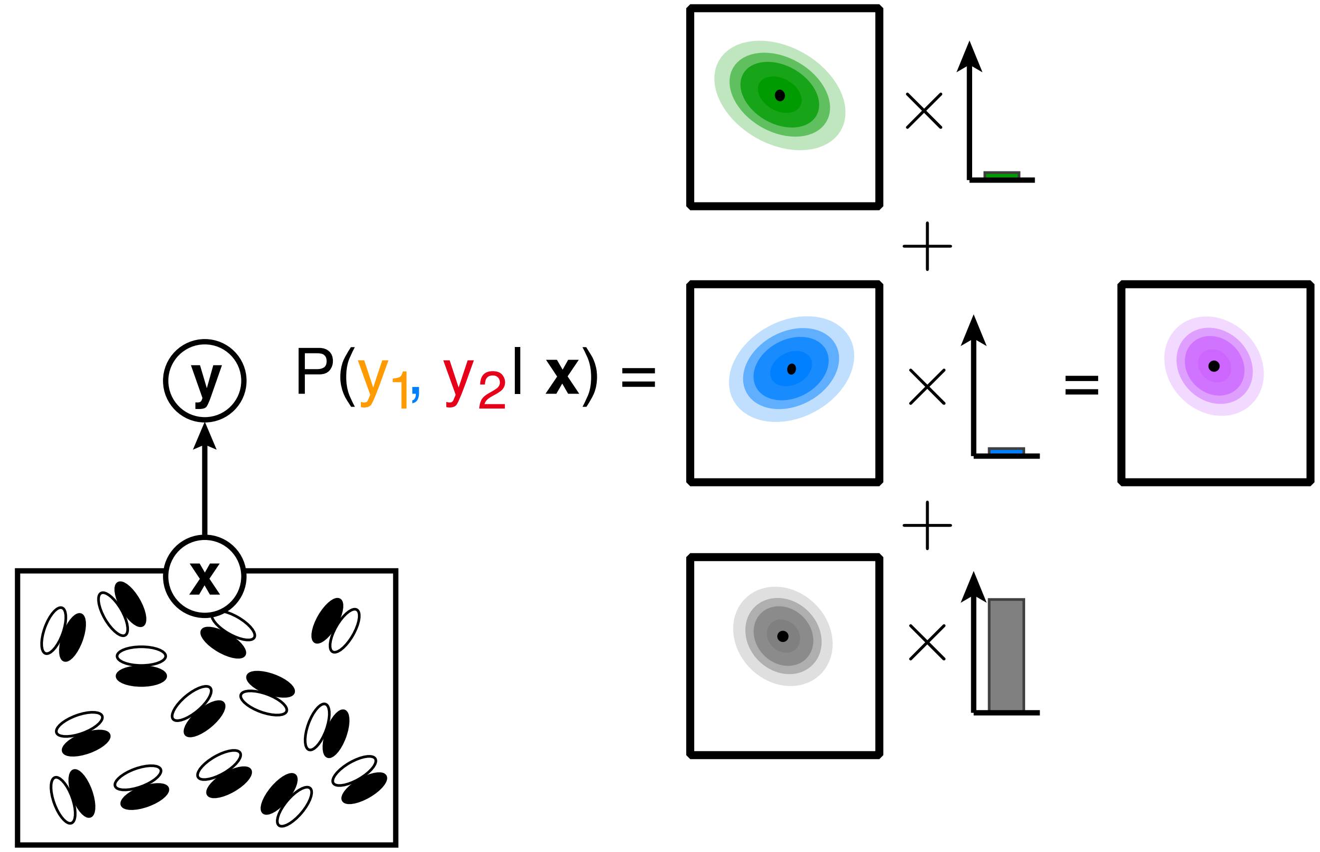

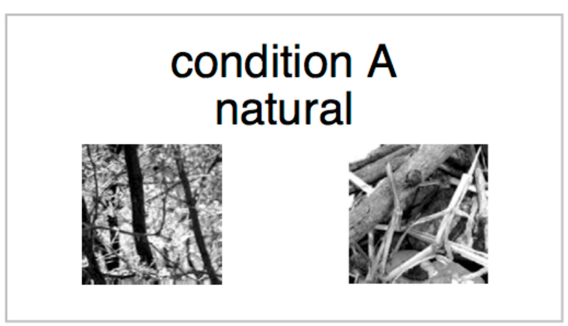

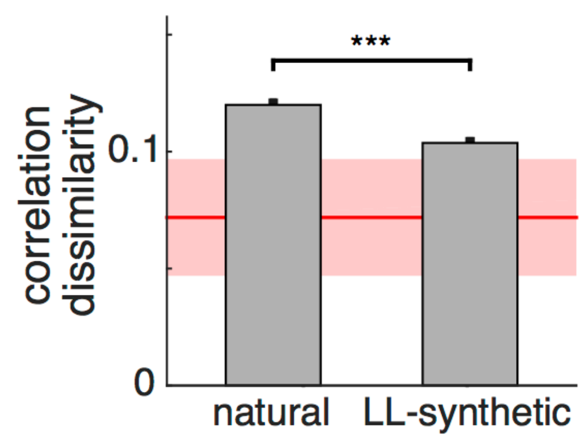

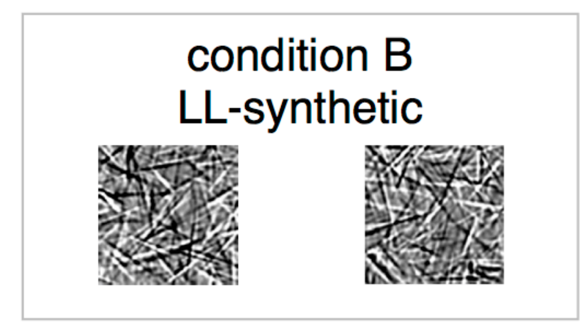

Hierarchical inference in a generative model