BSCS 2019 - Neural Computation

VI - Neural coding

Mihály Bányai
banyai.mihaly@wigner.mta.hu
http://golab.wigner.mta.hu/people/mihaly-banyai/

- Tying perceptual inference to biophysics
- How to represent a probability distribution?
- Predictions of the sampling hypothesis

- Tying perceptual inference to biophysics
- How to represent a probability distribution?
- Predictions of the sampling hypothesis

Aiming to predict neural responses

- If we have a computational theory of mind, we may reproduce behaviour without having to look at biology
- But then we only know that we solved the same problem the brain has to
- We also want to try to understand how all this might be implemented
- Specific constraints introduced by the neural hardware may shed light on some properties of human thinking
- Disorders of mental function probably also have to be addressed on the neural level

> "After the great positive contribution of Turing-cum-Pitts-and-McCulloch is assimilated, the situation is rather worse than better than before. Indeed these authors have demonstrated in absolute and hopeless generality that anything and everything (…) can be done by an appropriate mechanism, and specifically by a neural mechanism—and that even one, definite mechanism can be 'universal.' Inverting the argument: Nothing that we may know or learn about the functioning of the organism can give, without 'microscopic,' cytological work any clues regarding the further details of the neural mechanism."
>
> *John von Neumann*

What is the basic anatomical element of computation/representation?

- What corresponds to one variable?
- Neurons
- Dendritic branches
- Microcircuits
- Populations of neurons

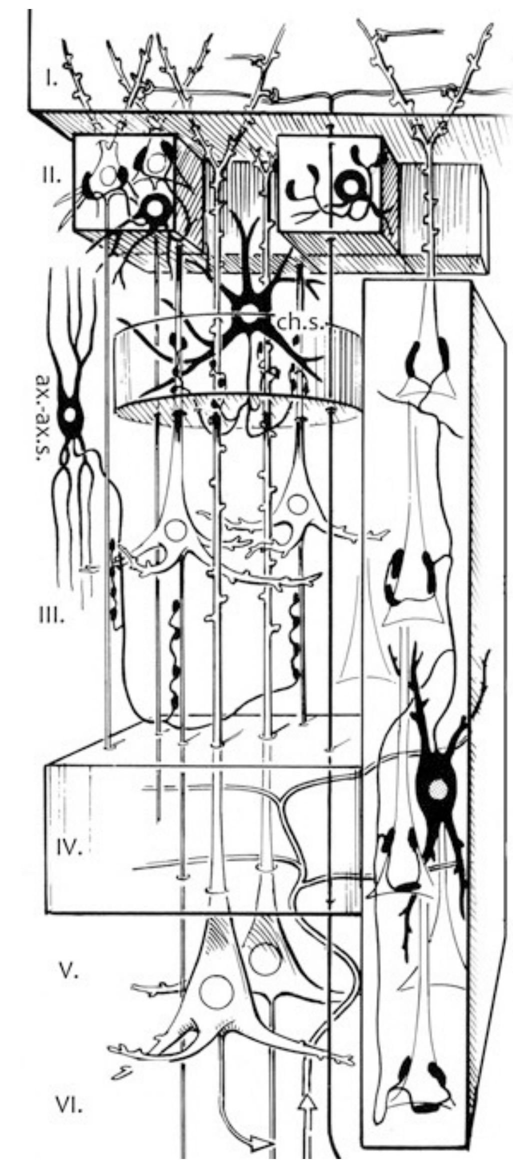

What is the biophysical quantity that implements computation?

- what corresponds to the value or probability of the variable?
- action potentials?
  - with spike timing having a role in computation
  - or only the rate of firing counts
- membrane potential in general?
- it looks relatively sure that long-term memory is implemented in synaptic plasticity

What is measurable in the brain?

- cellular level recordings from the cortex of behaving animals, simultaneously from many neurons
    - implanted electrode arrays
    - calcium imaging with microscope
- both see a couple of hundred cells maximum
- we can measure spike trains, not membrane potentials directly

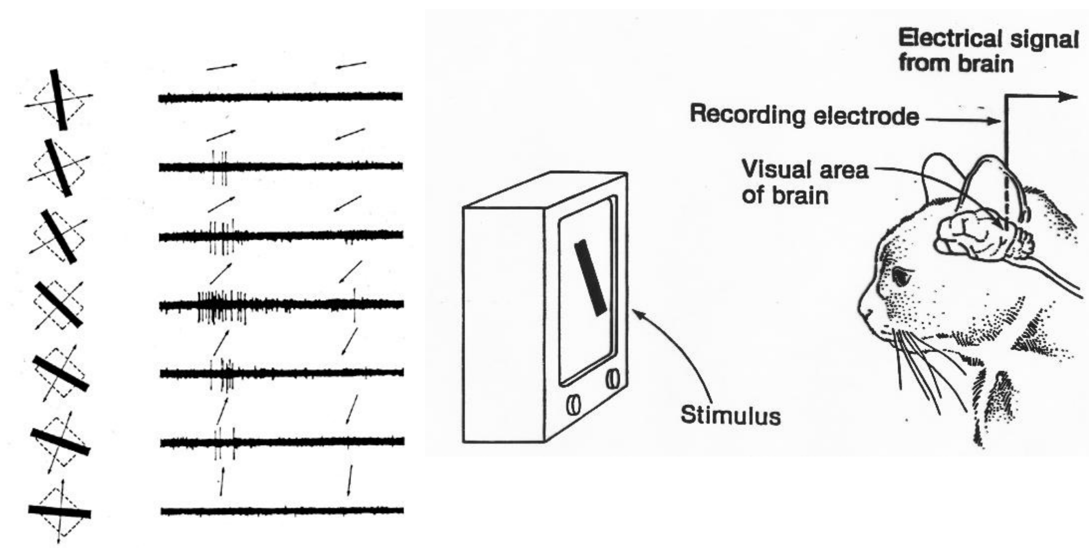

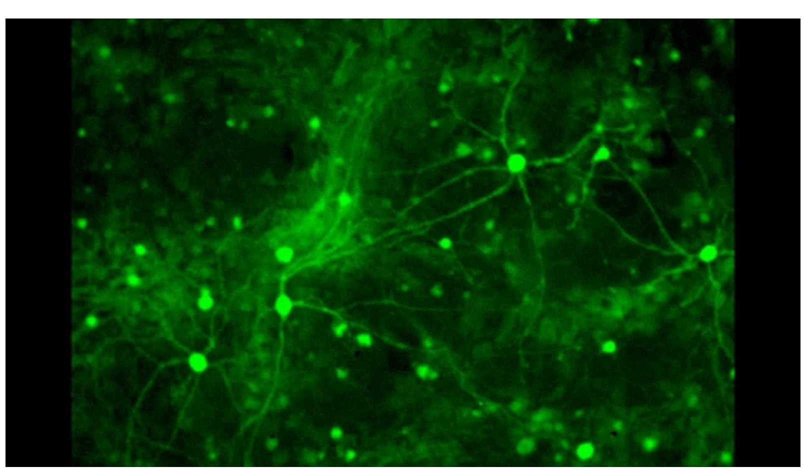

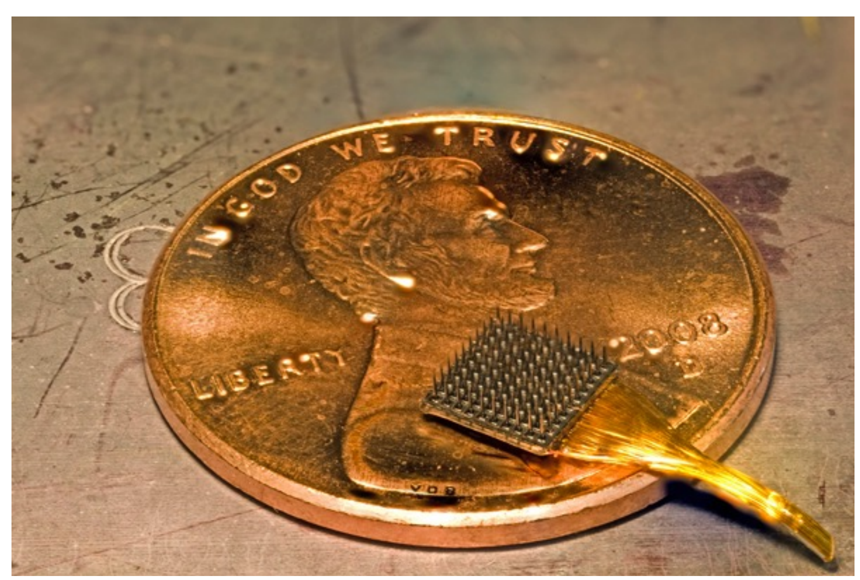

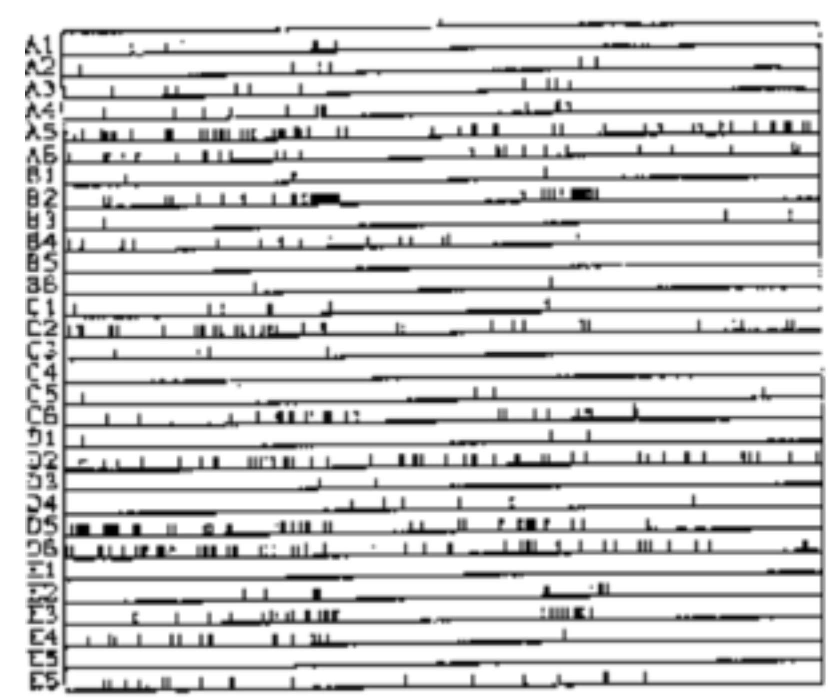

- Tying perceptual inference to biophysics
- How to represent a probability distribution?
- Predictions of the sampling hypothesis

How can we represent probability distributions?

- by specifying a PMF with probability values or a PDF in a parametric form, as we did so far

- or alternatively, by giving a lot of values of the variables that are distributed according to the PDF or PMF - **samples** from the distribution

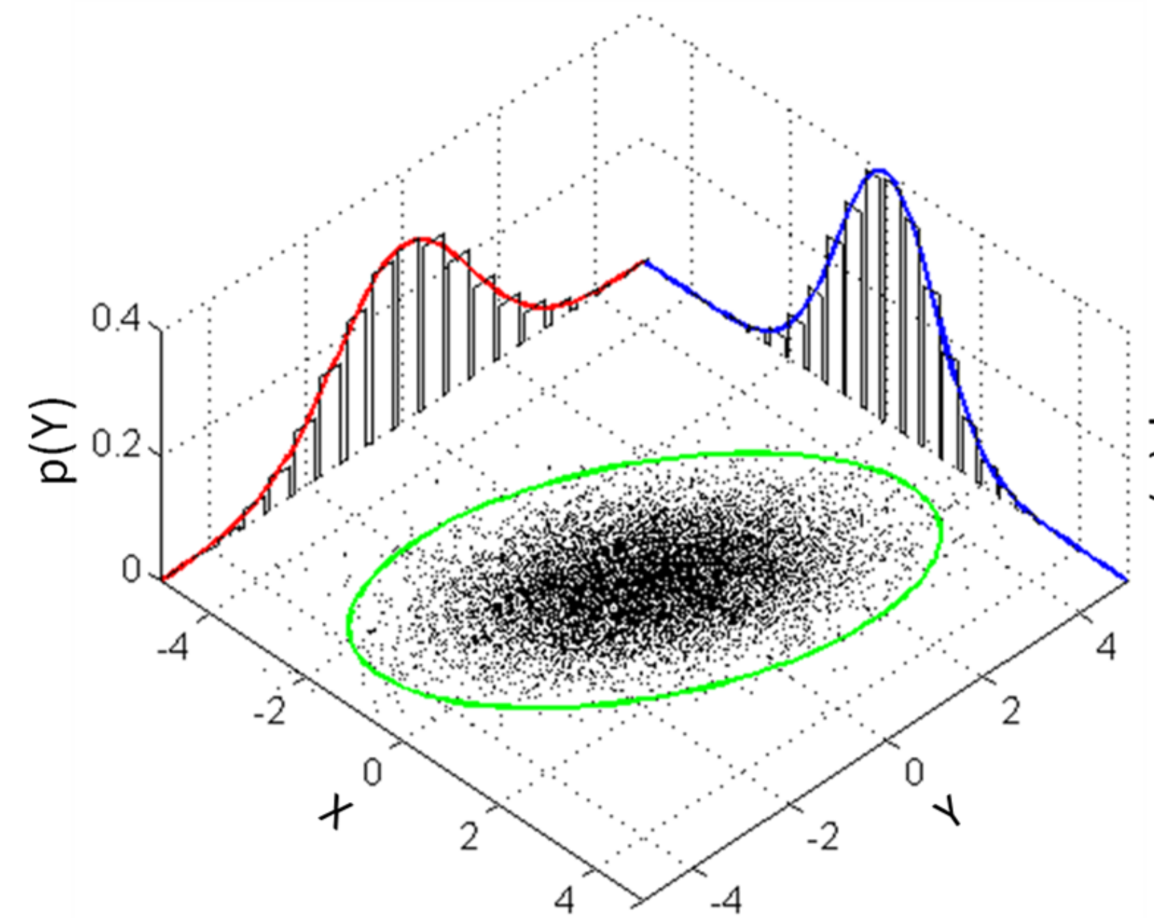

Samples of a variable

- generating samples from a distribution means that we have to **generate random numbers** that are distributed according to it
    - there are standard algorithms that can generate numbers that look like they come from a large number of possible values with equal probability independently after each other - pseudo-random samples from a **uniform** distribution
- these number sequences can be **transformed** to match our probability distribution by different methods, depending on the mathematical properties of the PDF or PMF
- “fake data”

Generative models

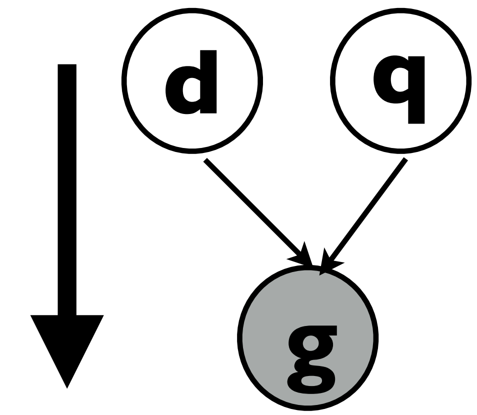

- a probabilistic model has a special interpretation: it describes a procedure by which the observations are **generated**
- we can use the model to create synthetic data, i.e. artificial observation sequences that reflect the statistical assumptions we made when quantifying the PMFs of the model
- to do so, we have to first **sample** values for the hidden variables from their marginal PMFs
- then we can sample from the conditional PMF of the observable to obtain observable values
- the generative procedure can be thought of like the model *imagines* the kind of observations that fit its beliefs

The sampling hypothesis

- the brain represents probability distributions by samples
- the membrane potential of a cell at a given time point represents a sample from the posterior probability distribution of a hidden variable of the mental model, conditioned on the sensory stimulus actually observed by the animal or human
- tuning distributions to match observed statistics is implemented by synaptic plasticity

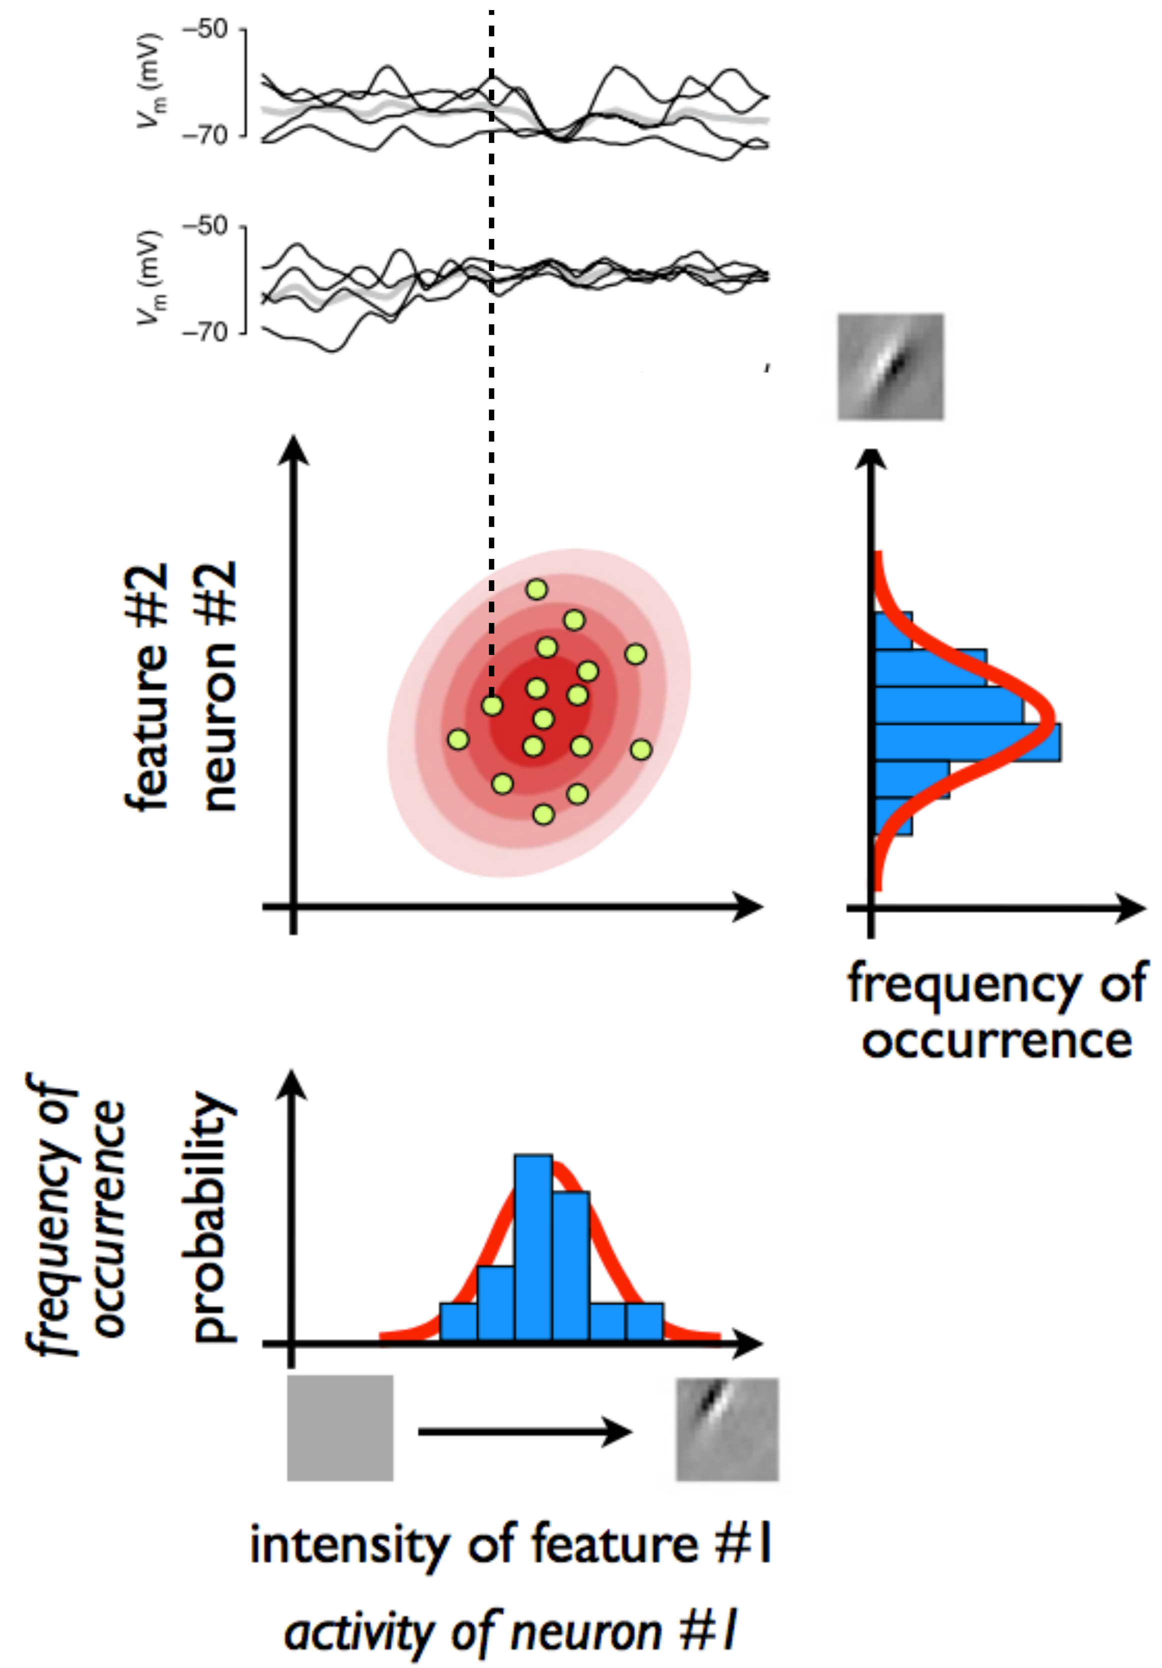

Sources of randomness in the brain

- Producing random samples requires something like a random number generator
- Sensory inputs introduce a variability that from the perspective of the brain as observer, is random
  - but it's not enough, we need randomness with fixed inputs too
- There is some inherent noise in sensors, ion channels, synaptic conductance, etc.
- If we have a complex deterministic dynamics, we don't necessarily need a truly random number - e.g. a chaotic system can produce values that are "random" enough for any practical purpose, like sampling
  - computers don't use real randomness either, only pseudo-random numbers, and they work well
- I'm not speaking about quantum phenomena on purpose

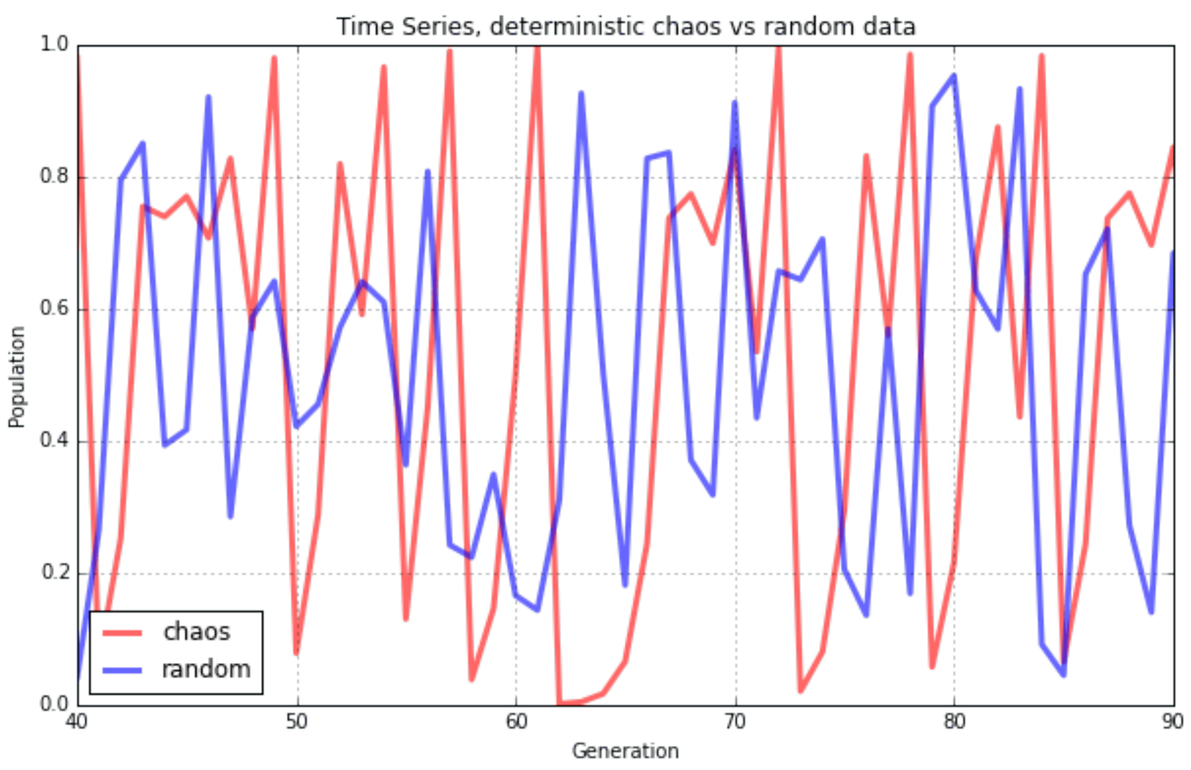

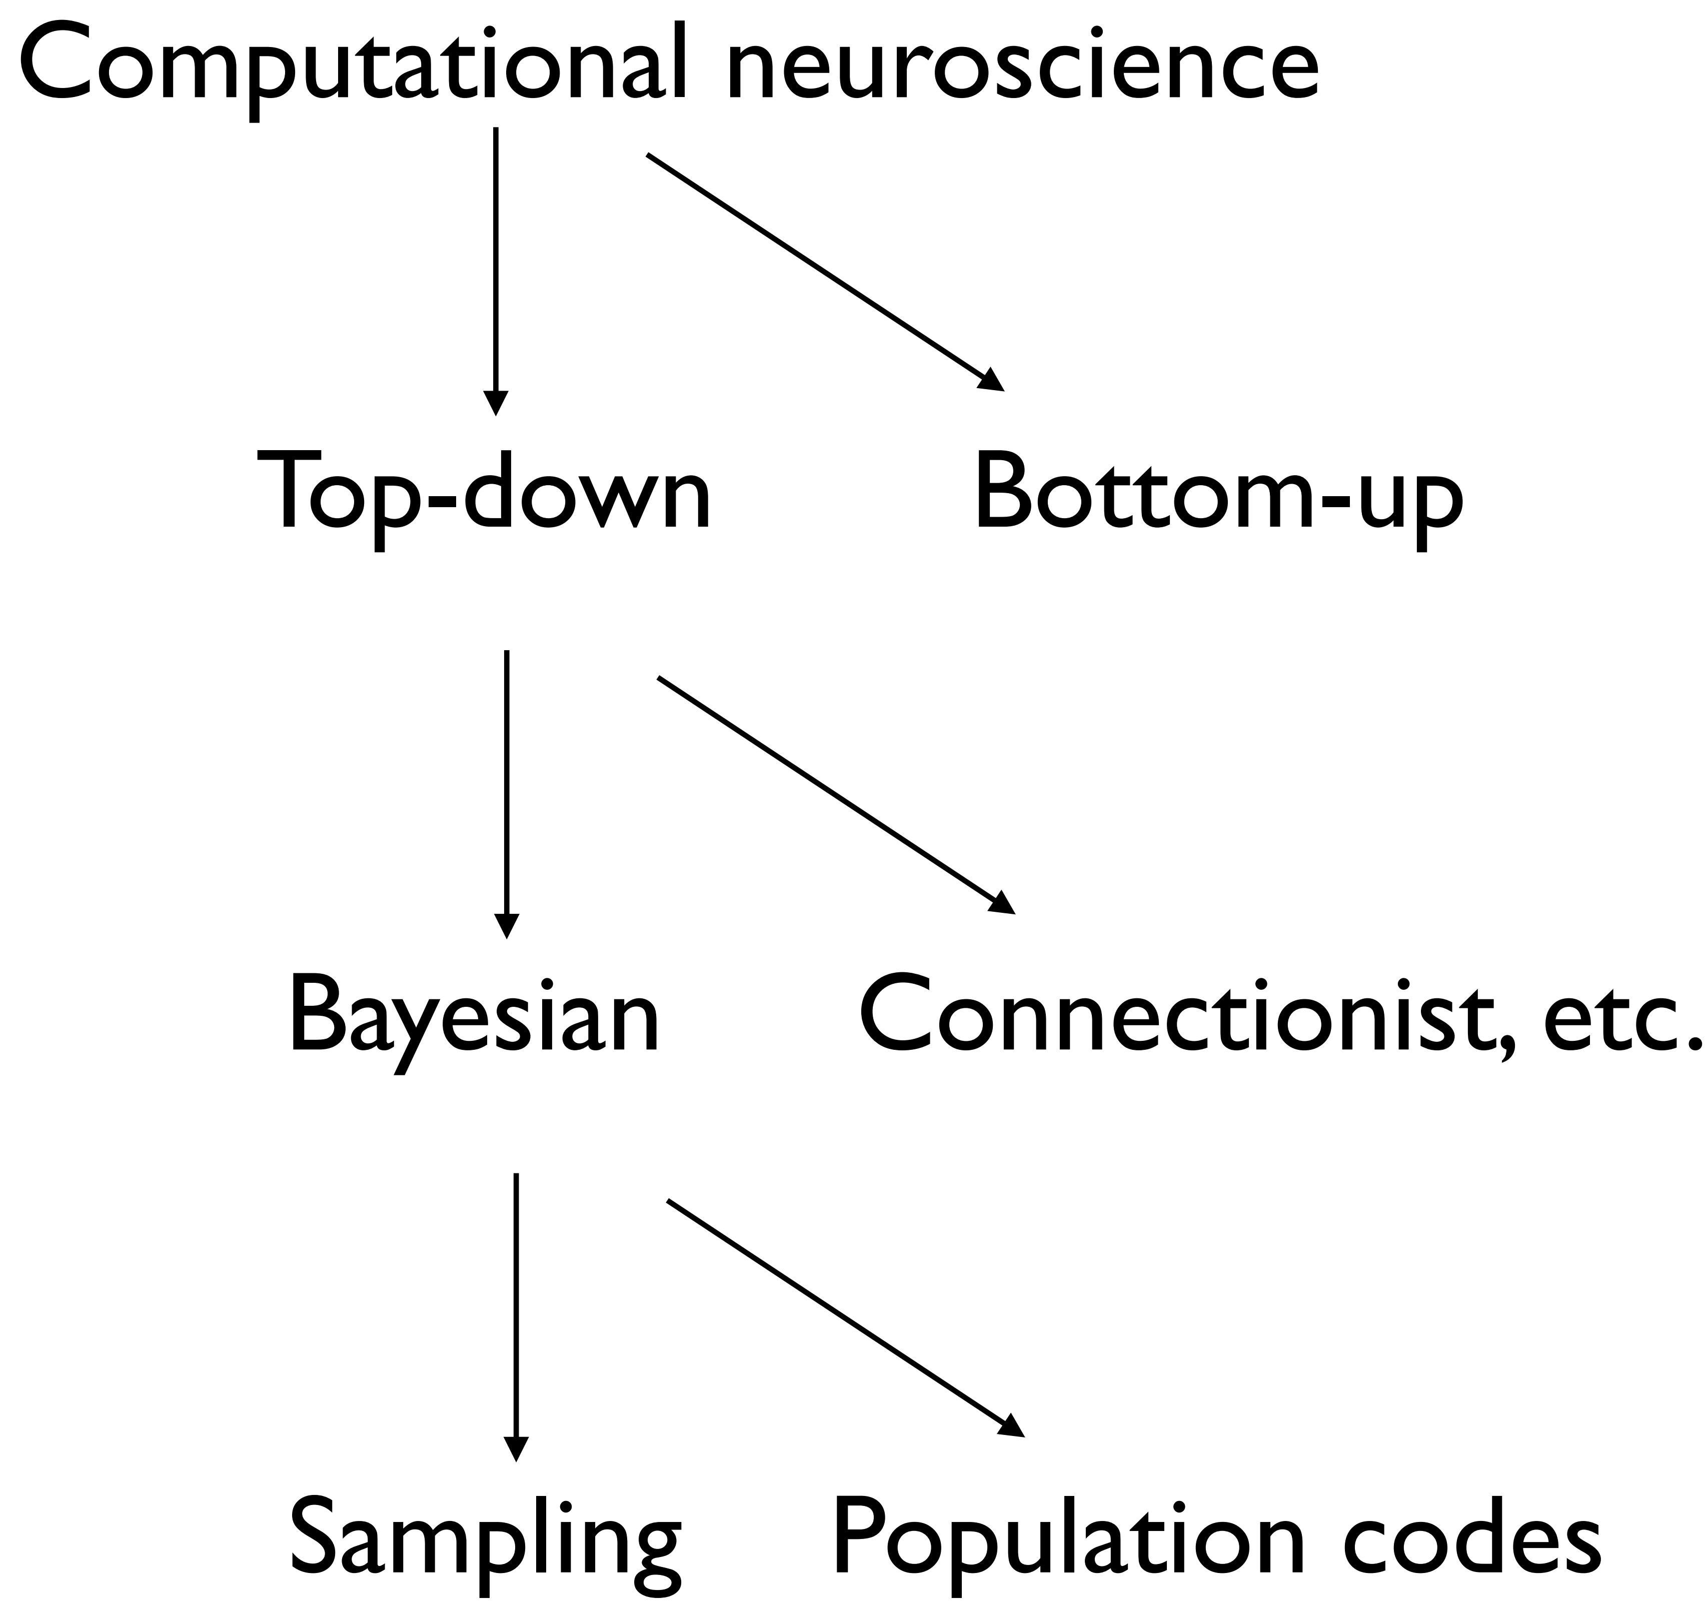
Computational neuroscience
Top-down
Bottom-up
Bayesian
Connectionist, etc.
Sampling
Population codes

- Tying perceptual inference to biophysics
- How to represent a probability distribution?
- Predictions of the sampling hypothesis

Bistability of perceptions

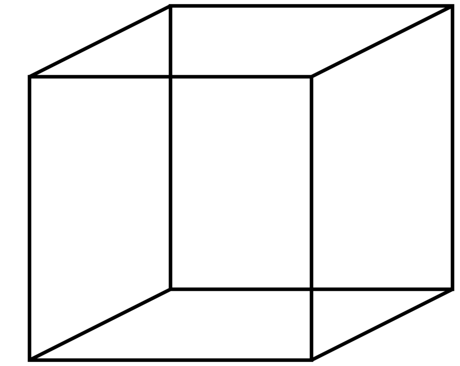

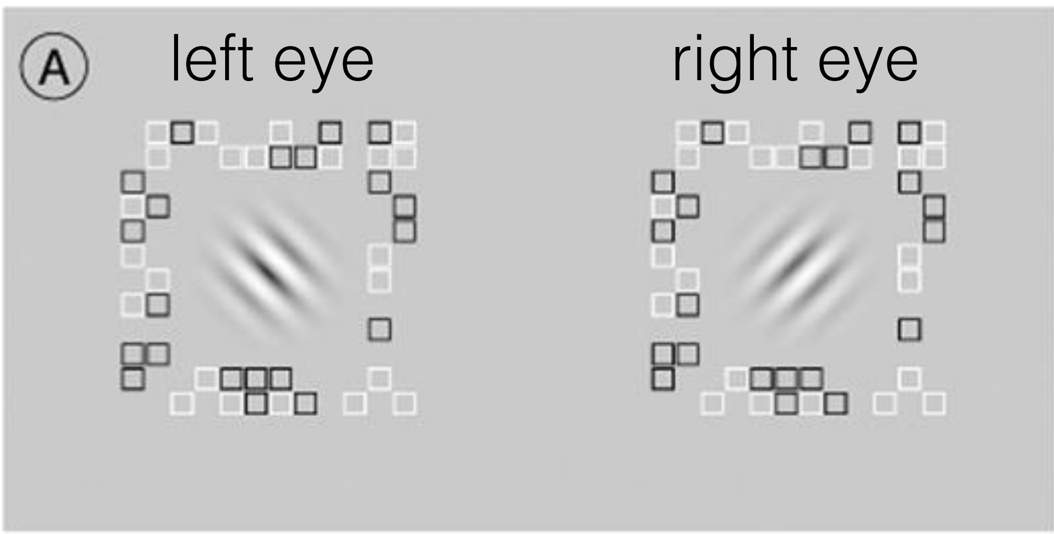

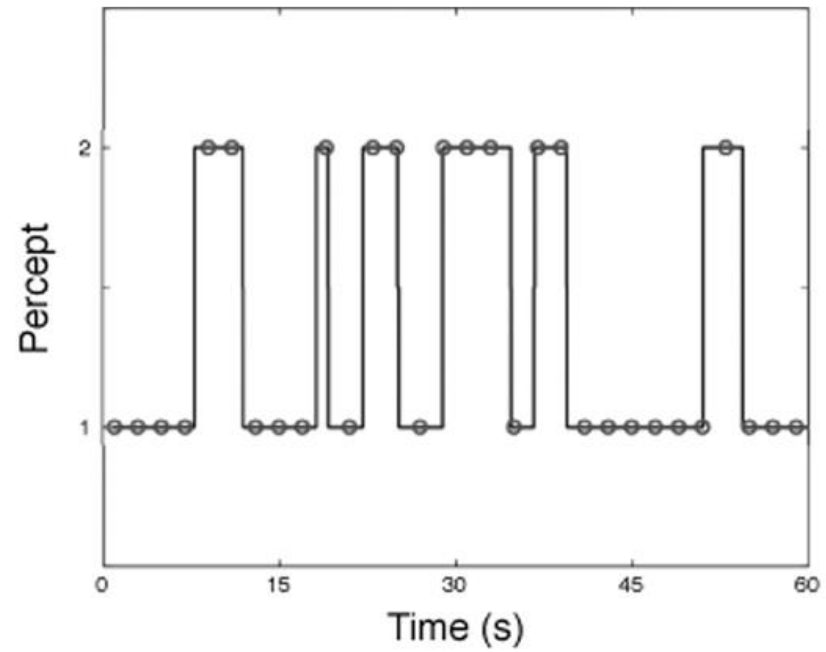

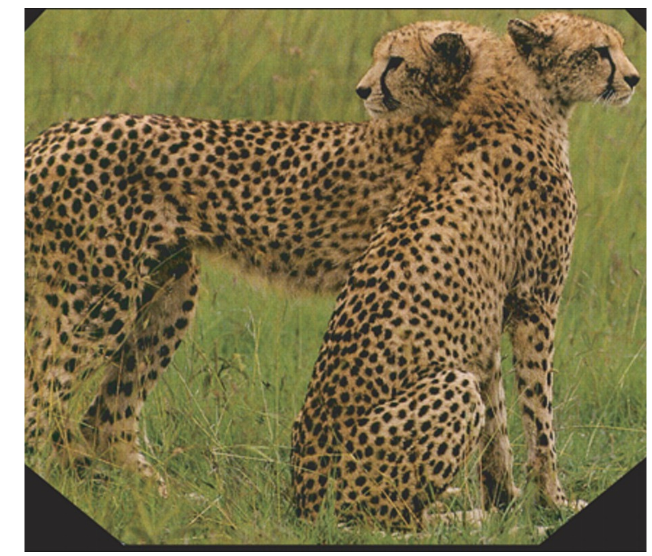

- perception shifts from one interpretation to another repeatedly
- suggests an underlying stochastic process

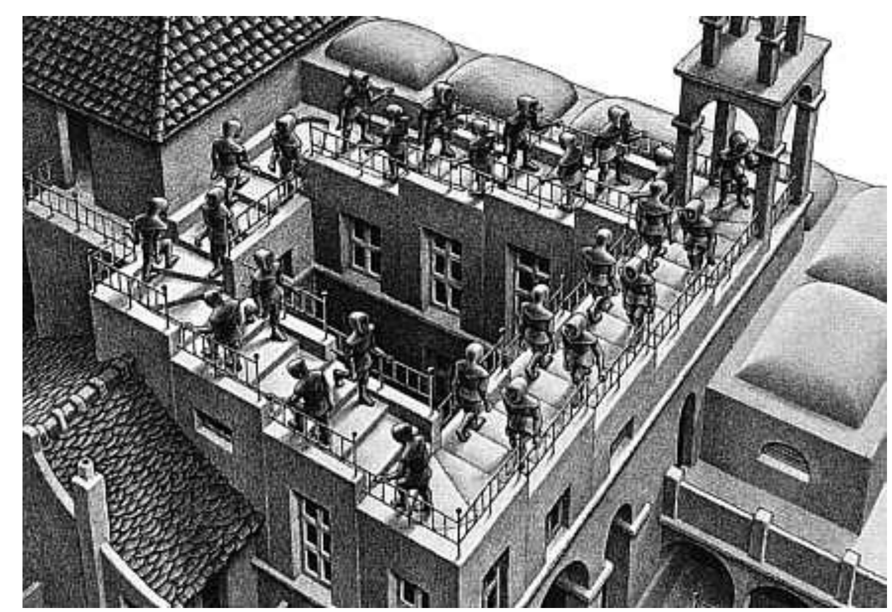

Mamassian, P., & Goutcher, R. (2005). Temporal dynamics in bistable perception. Journal of Vision, 5(4), 7.

What is spontaneous activity according to the SH?

- evoked activity is measured in response to a stimulus, and is sampling from a posterior distribution of the model according to the SH
- spontaneous activity is measured when there is no stimulus presented
- the posterior without any contribution of an observation through the likelihood is proportional to the prior
- thus, according to the SH, spontaneous activity samples from the prior distribution of the variables of the mental model
- this is equivalent to generating synthetic observations of the model
- may describe how we imagine stimuli that we do not see

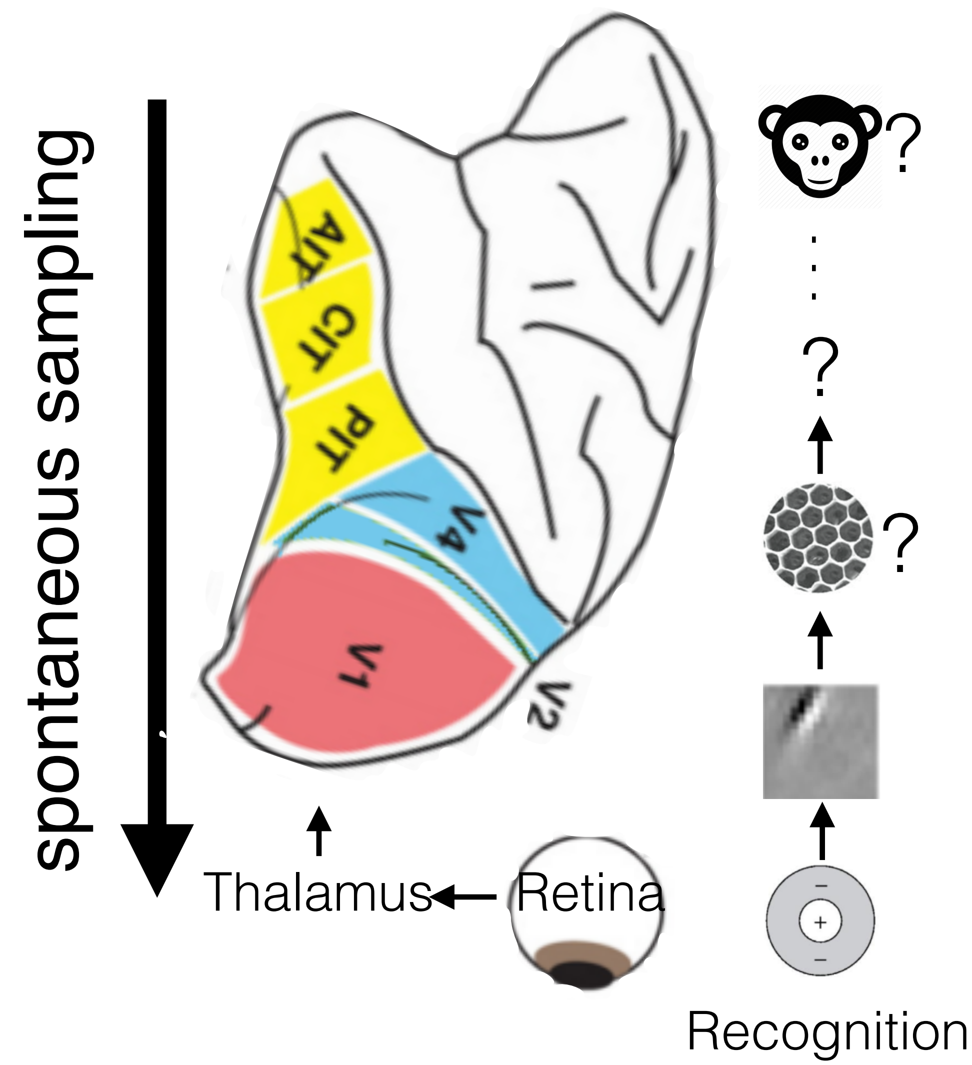

A prediction of the SH about variability

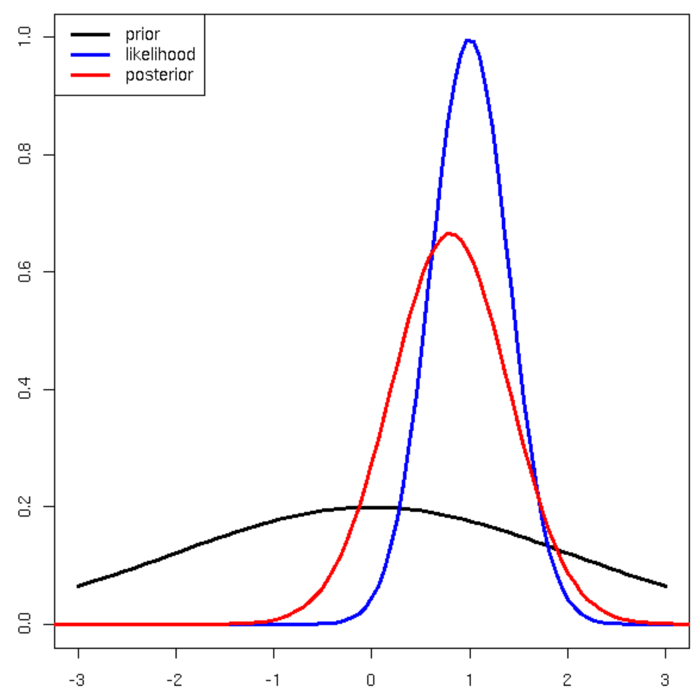

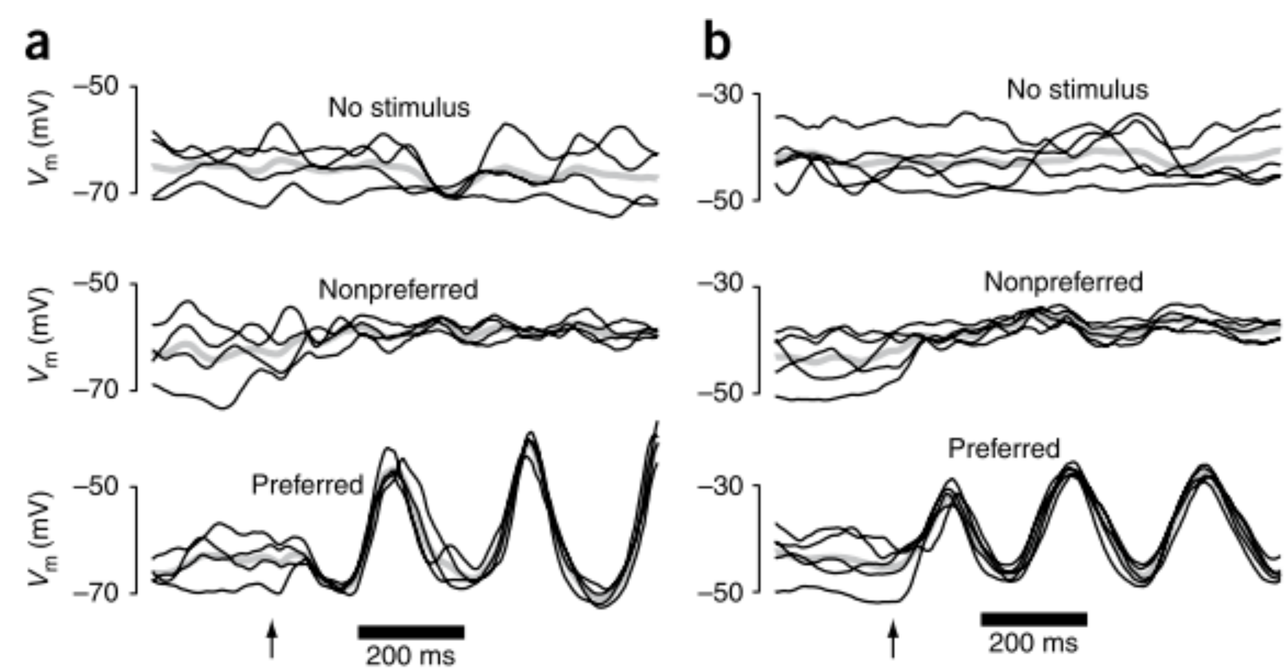

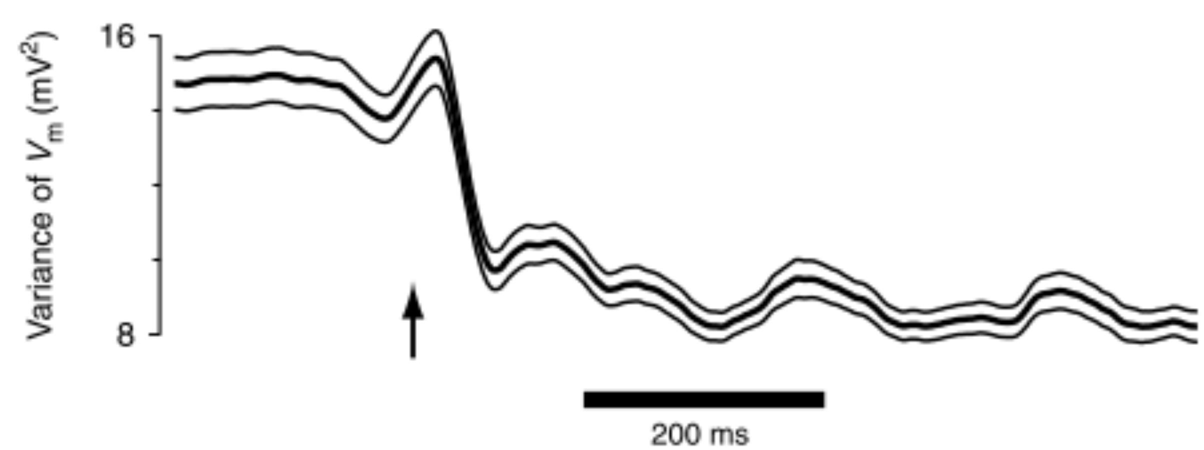

- variability of spontaneous activity should be larger than that of evoked activity
- because the prior distribution of a quantity if always wider than the posterior

Churchland, Mark M., et al. "Stimulus onset quenches neural variability: a widespread cortical phenomenon." Nature neuroscience 13.3 (2010): 369-378.

A prediction of the SH about average response statistics

- If it is true that
  - the brain builds a mental model of the environment
  - that is tuned to reproduce the feature statistics of stimuli observed throughout the lifetime of the animal or human
  - and represents the probability distributions of the model through sampling from them by the neural dynamics
  - evoked activity being sampled from the posterior distribution of the quantity represented by the given neuron, conditioned on the given observation
  - and spontaneous activity being sampled from the prior distribution of the quantity represented by the given neuron
- then it has to be true that
  - after stimulus-tuning (learning) takes place, the average evoked activity conditioned on many stimuli that are similar to those observed during learning (i.e. that are natural) should be similar to spontaneous neural activity
  - but probably not before observing anything from the world - though it's an interesting question how much of the neural representation is innate
  - and for stimuli that are very different from the ones seen during learning (i.e. artificial) it should not be true, or at least less so

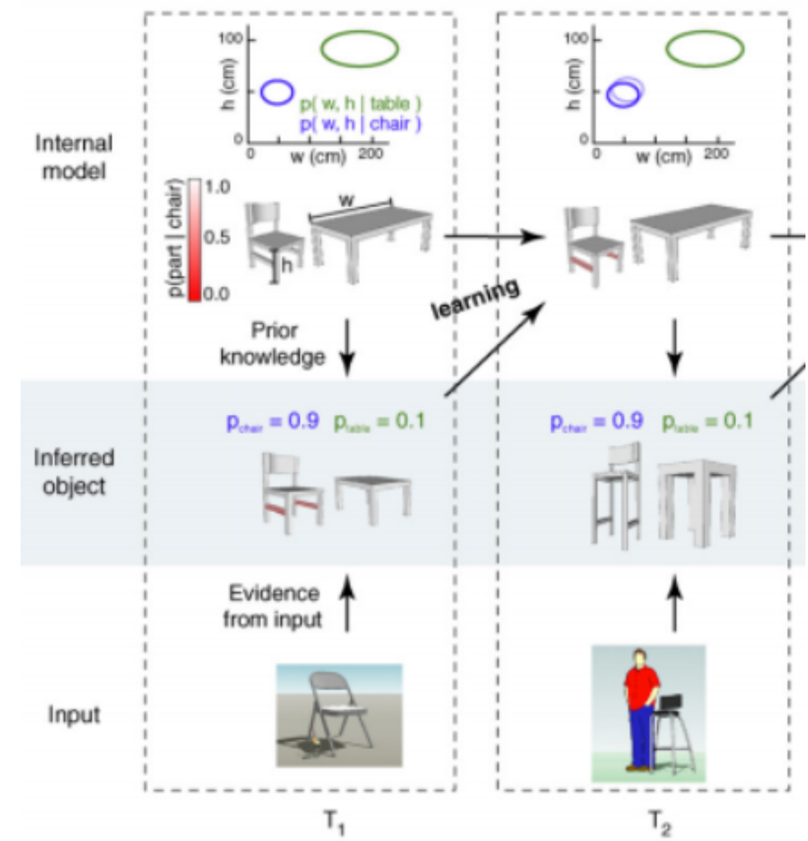

Measuring neural activity statistics

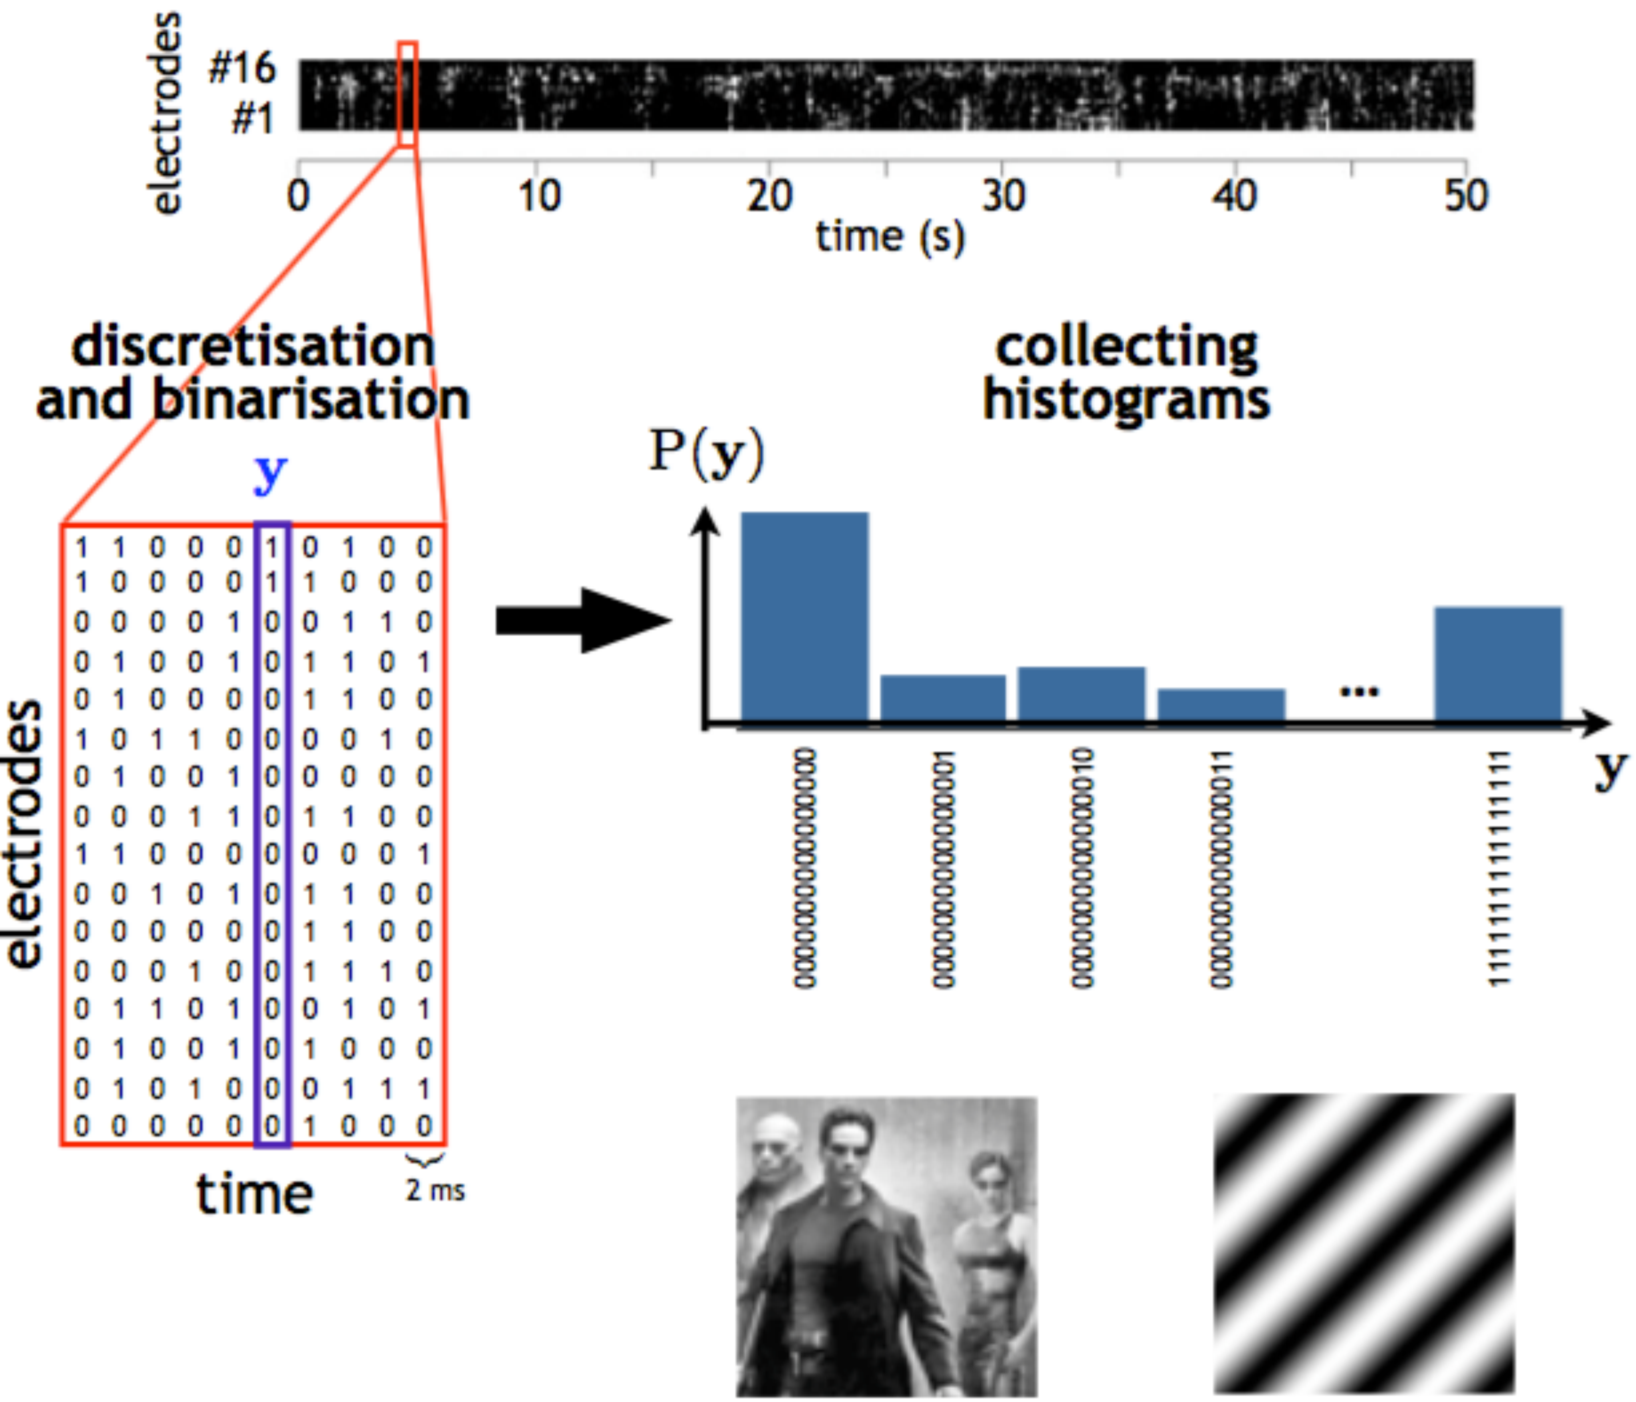

- electrodes implanted to ferret visual cortex
- the ferrets watch stimulus with natural image statistics (a movie)
- stimuli with artificial statistics
- spontaneous activity in darkness is also recorded

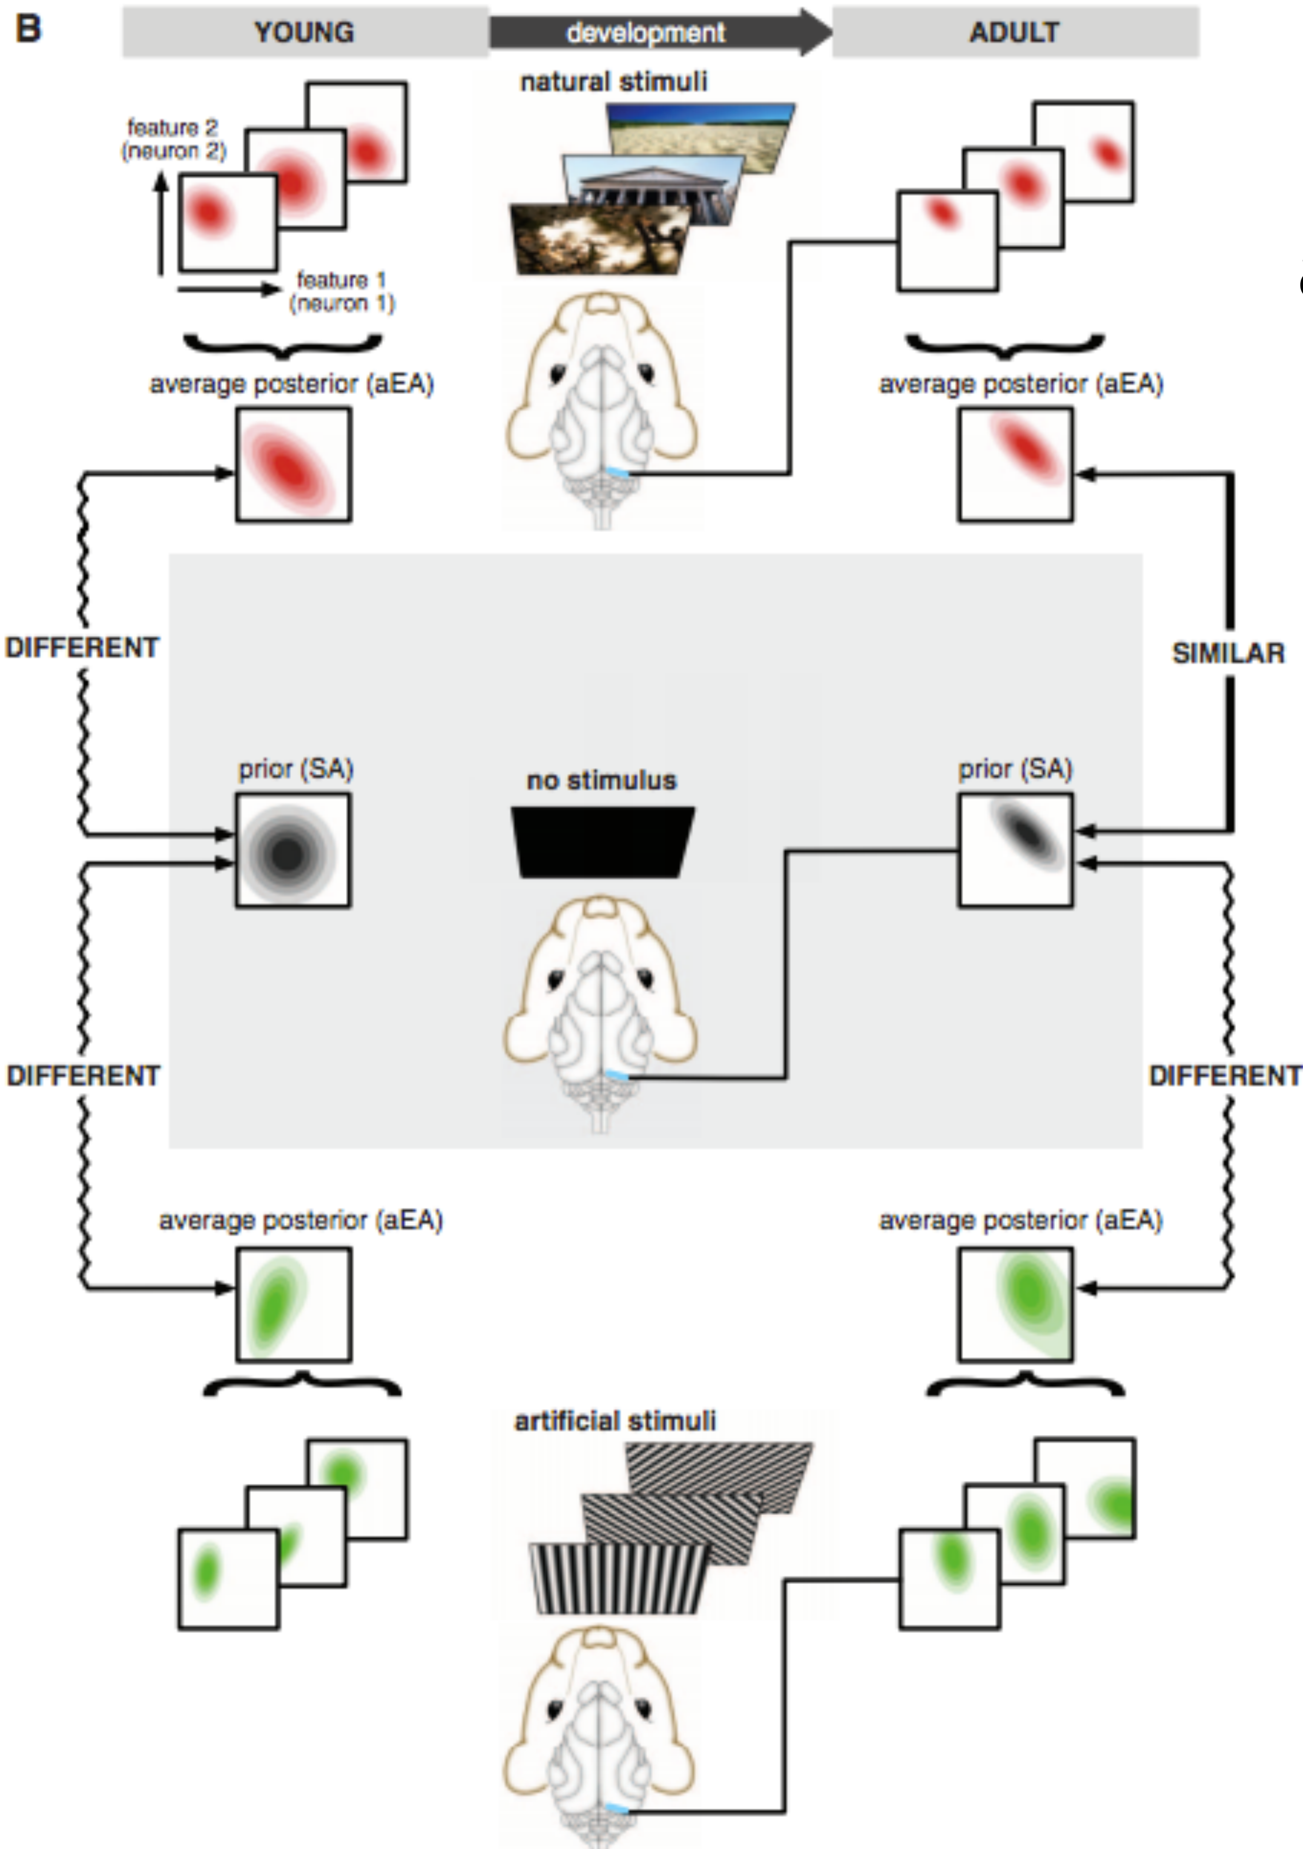

Relationship of evoked and spontaneous activity during development

- evoked activity averaged over natural stimuli is similar to spontaneous activity
    - only in the adult animal, not at eye opening
- for artificial stimuli, averaged evoked activity is not similar to spontaneous activity even in the adult ferrets
- this finding confirms the prediction of SA about spontaneous neural statistics

Berkes, P., Orbán, G., Lengyel, M., & Fiser, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. Science, 331(6013), 83-87.

The way forward

- now we have a possible way to represent probability distributions with neural activity

- we have seen some predictions of the sampling hypothesis compared to neural measurements

- we will move on to see what kind of neural response properties can we predict with probabilistic generative models of stimuli together with a sampling-based neural code

- we also want to explore the possible ways to constrain the set of possible generative models implemented by the cortex