BSCS 2019 - Neural Computation

VI - Prediction of neural activity

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

- What to predict about neural activity?
- Limits of measurability in the brain
- How to build, test and improve models?

- What to predict about neural activity?
- Limits of measurability in the brain
- How to build, test and improve models?

Importance of making falsifiable predictions

- in order to assess whether our ideas about neural representations, formalised in models are any good, we have to be able to compare them to actual data from physiological experiments

- strictly speaking, we cannot confirm the validity of a model, only reject it, if it's outright false about something we measure

- so the way to go is to squeeze out as many predictions about the models regarding phenomena we do have measurements about as possible

- we can also compare models in terms of predictive accuracy

- if we indeed falsify a prediction of a model, it doesn't (necessarily) mean that we did something wrong - we learn that some of the assumptions we made when building the model were not accurate

Hierarchy of hypotheses about the brain

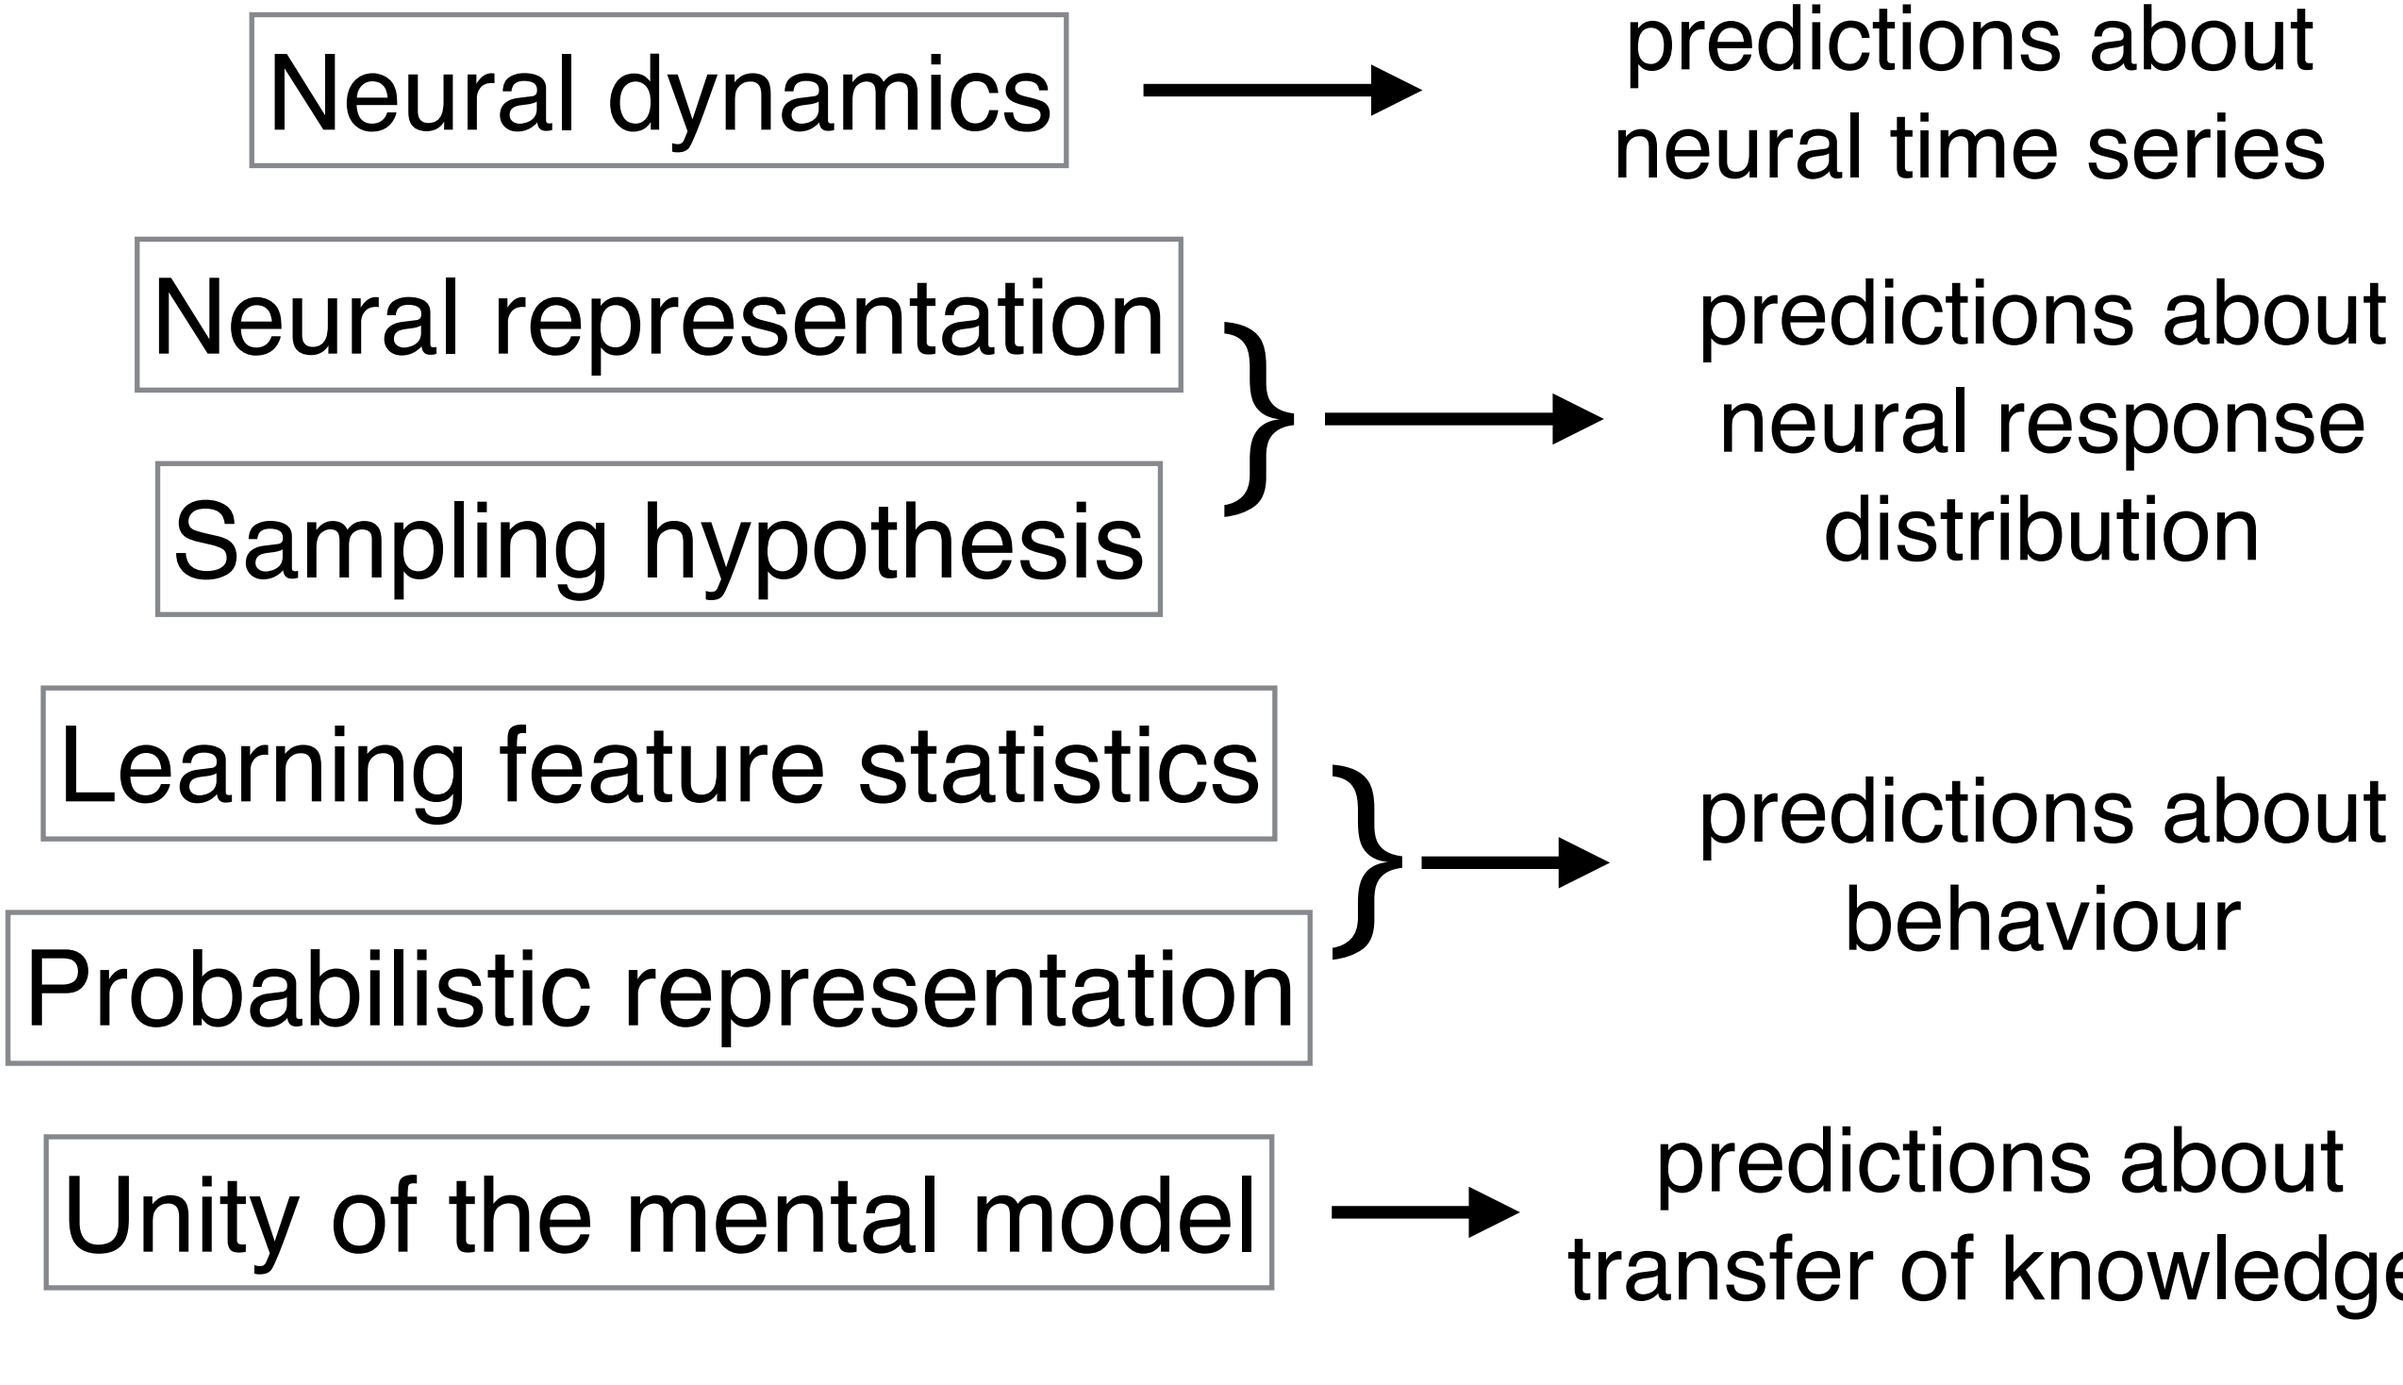

Using a specific model for prediction

- handling multiple observations in a graphical model
  - if we have several observations about the same variable, we would have to introduce each one of them as an independent observed value in the model - this would lead to giant models
  - there is a shorthand notation of this: the plate
    - $x_i \perp x_j \mid u \Leftrightarrow p(x|u) = \prod_i p(x_i|u)$
    - whatever is on the plate has a separate value for each observation, and whatever is off has a single value for all observations

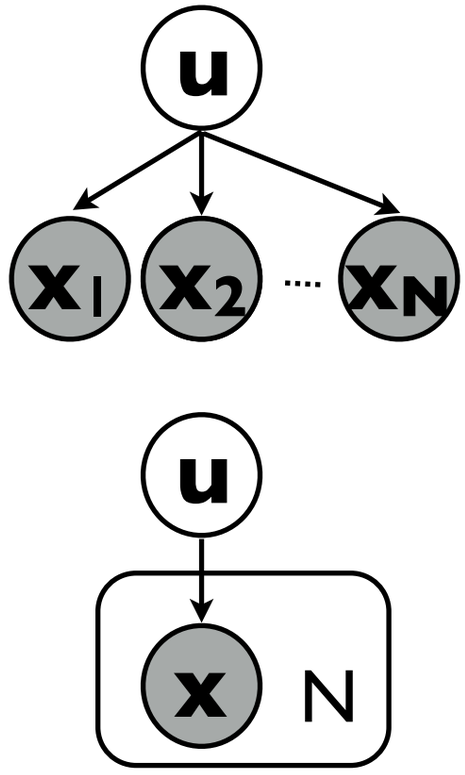

- predictions of response statistics per se are not strictly falsifiable, as we don't know what level of noise should be tolerable
- but we can produce predictions about how the response should change if the stimulus changes - and this is falsifiable

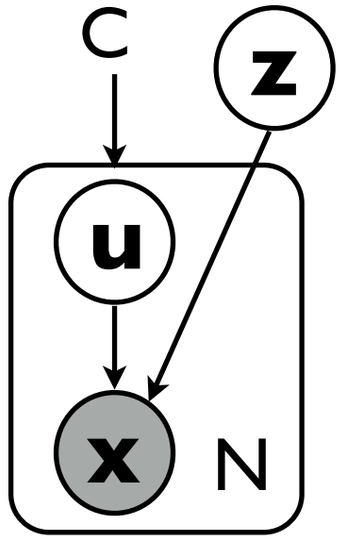

Functional intuition about brain regions

- in order to formalise a probabilistic model about what kind of inference a certain brain region implements, we have to have an idea about the computational problem it solves
    - the easiest targets of such intuition are sensory cortices
        - the visual cortex needs to perform object recognition
        - auditory cortex has to perform localisation & separation of sound sources
- this intuition suggests a formal computation that solves the same problem

- What to predict about neural activity?
- Limits of measurability in the brain
- How to build, test and improve models?

Three aspects of resolution of measurement techniques

- spatial - how small is the volume in which we can measure the neural activity separately from the rest of the brain?
    - EEG is very bad at this
    - fMRI is better, but still far from seeing individual cells
- temporal - how long does it take to make one measurement of the neural activity?
    - fMRI is very slow
    - electrophysiological methods are very good at this
- coverage - what portion of the whole brain can we monitor at the same time?
    - patch-clamp can only measure one cell
    - multielectrode arrays are better, but still only a couple of hundreds maximum
    - same with calcium imaging

Measuring from a behaving animal

- in an experiment with sensory cortex measurements, we have to control the stimulus the animal receives at each moment
    - in vision, we have to know what exactly the animal sees - not easy, as eyes move all the time
- one attempt to solve this is anaesthesia - primary sensory cortices keep working even if the animal is unconscious
    - but the anaesthetics have various, not completely known effects on the response distribution of neurons - results are not completely comparable to awake studies
- another possibility is training the animal to perform a task in which it has to fixate on the same point
    - only monkeys will do this, not cats or rats

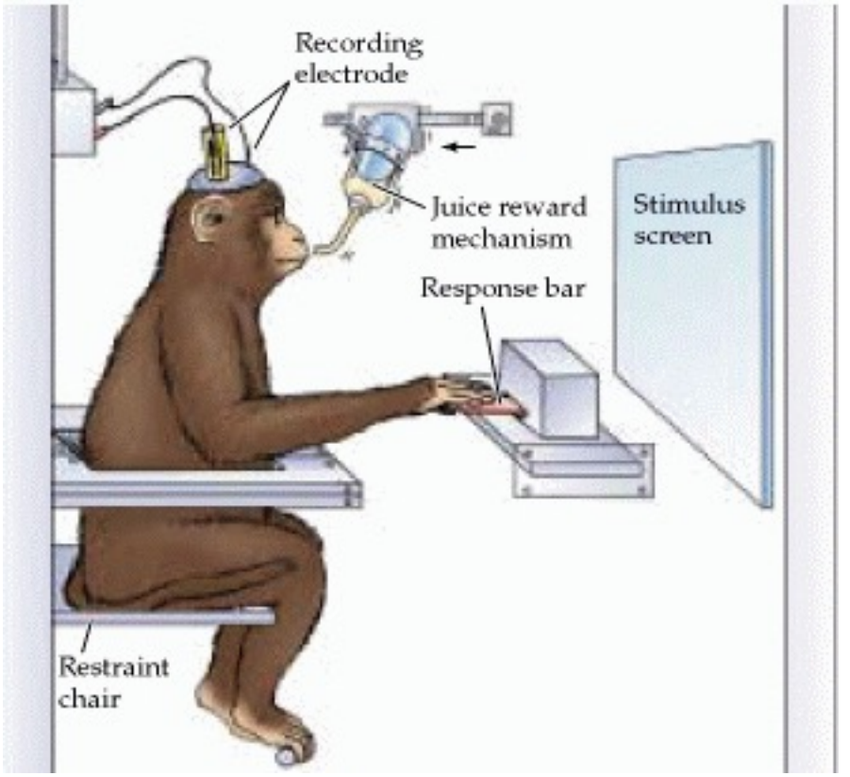

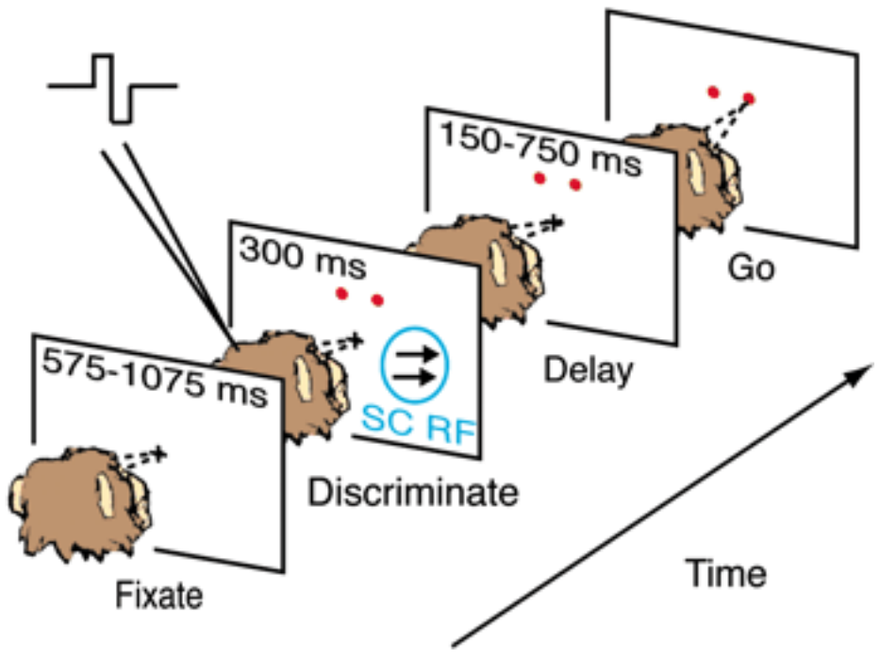

Microelectrode arrays

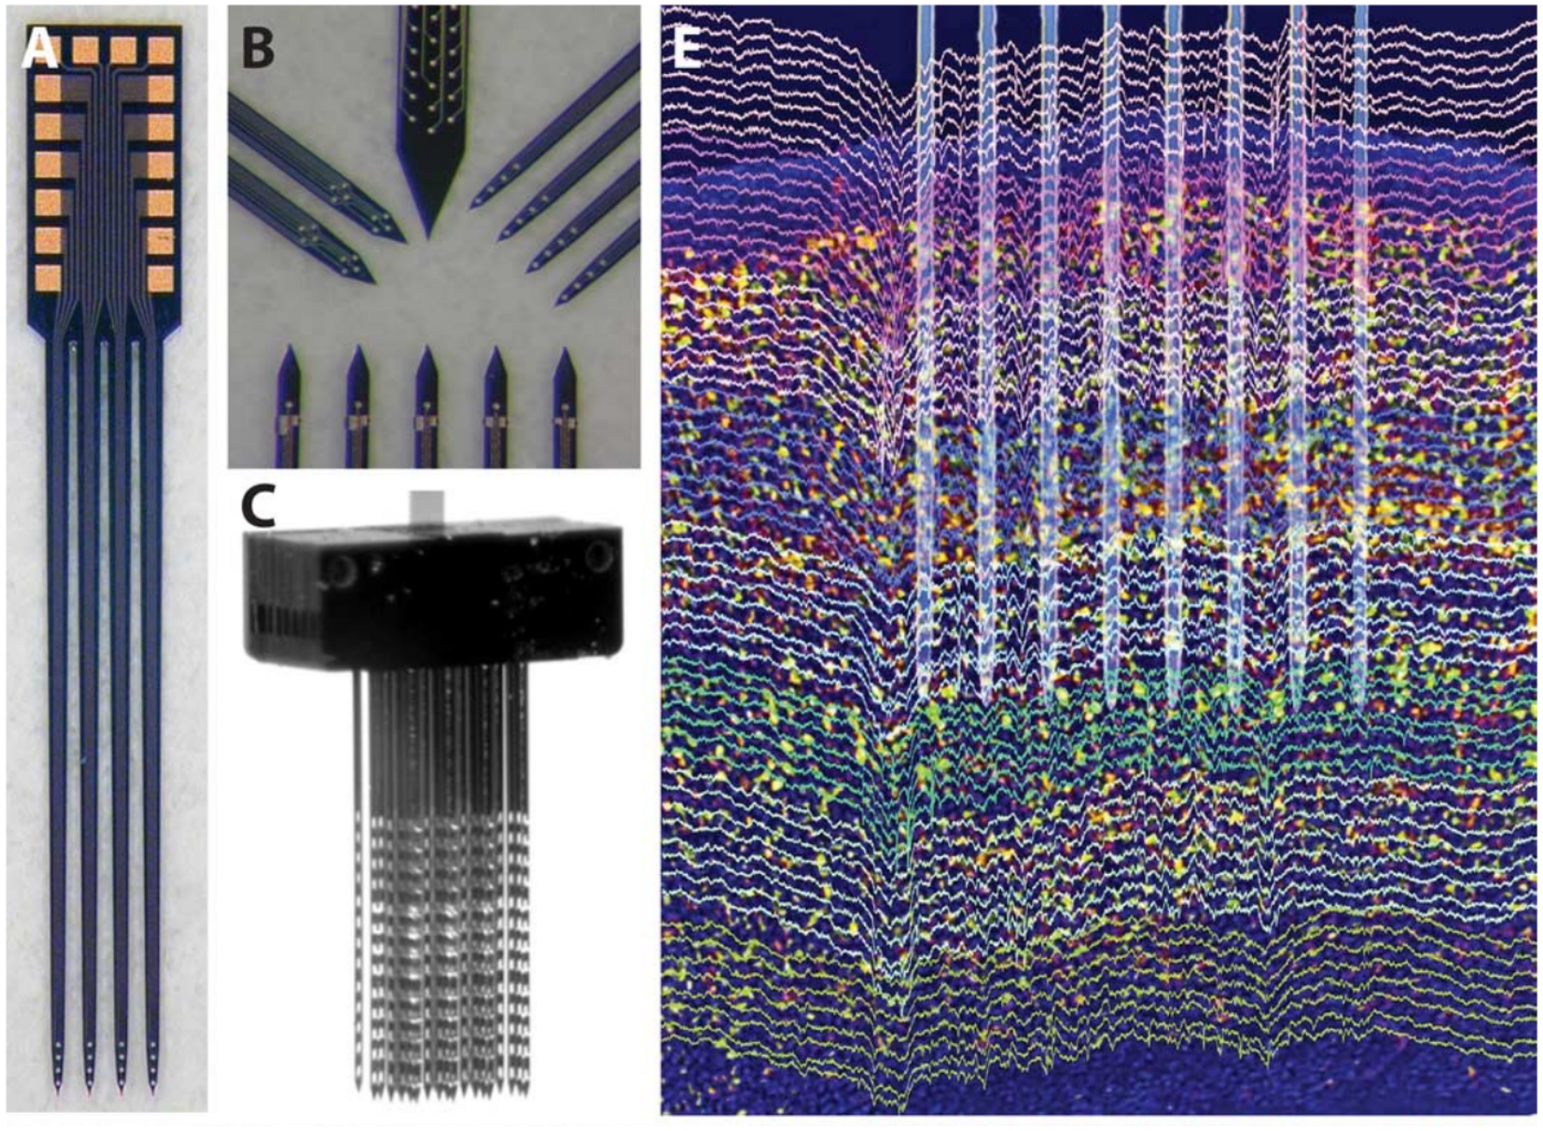

Extracellular recording

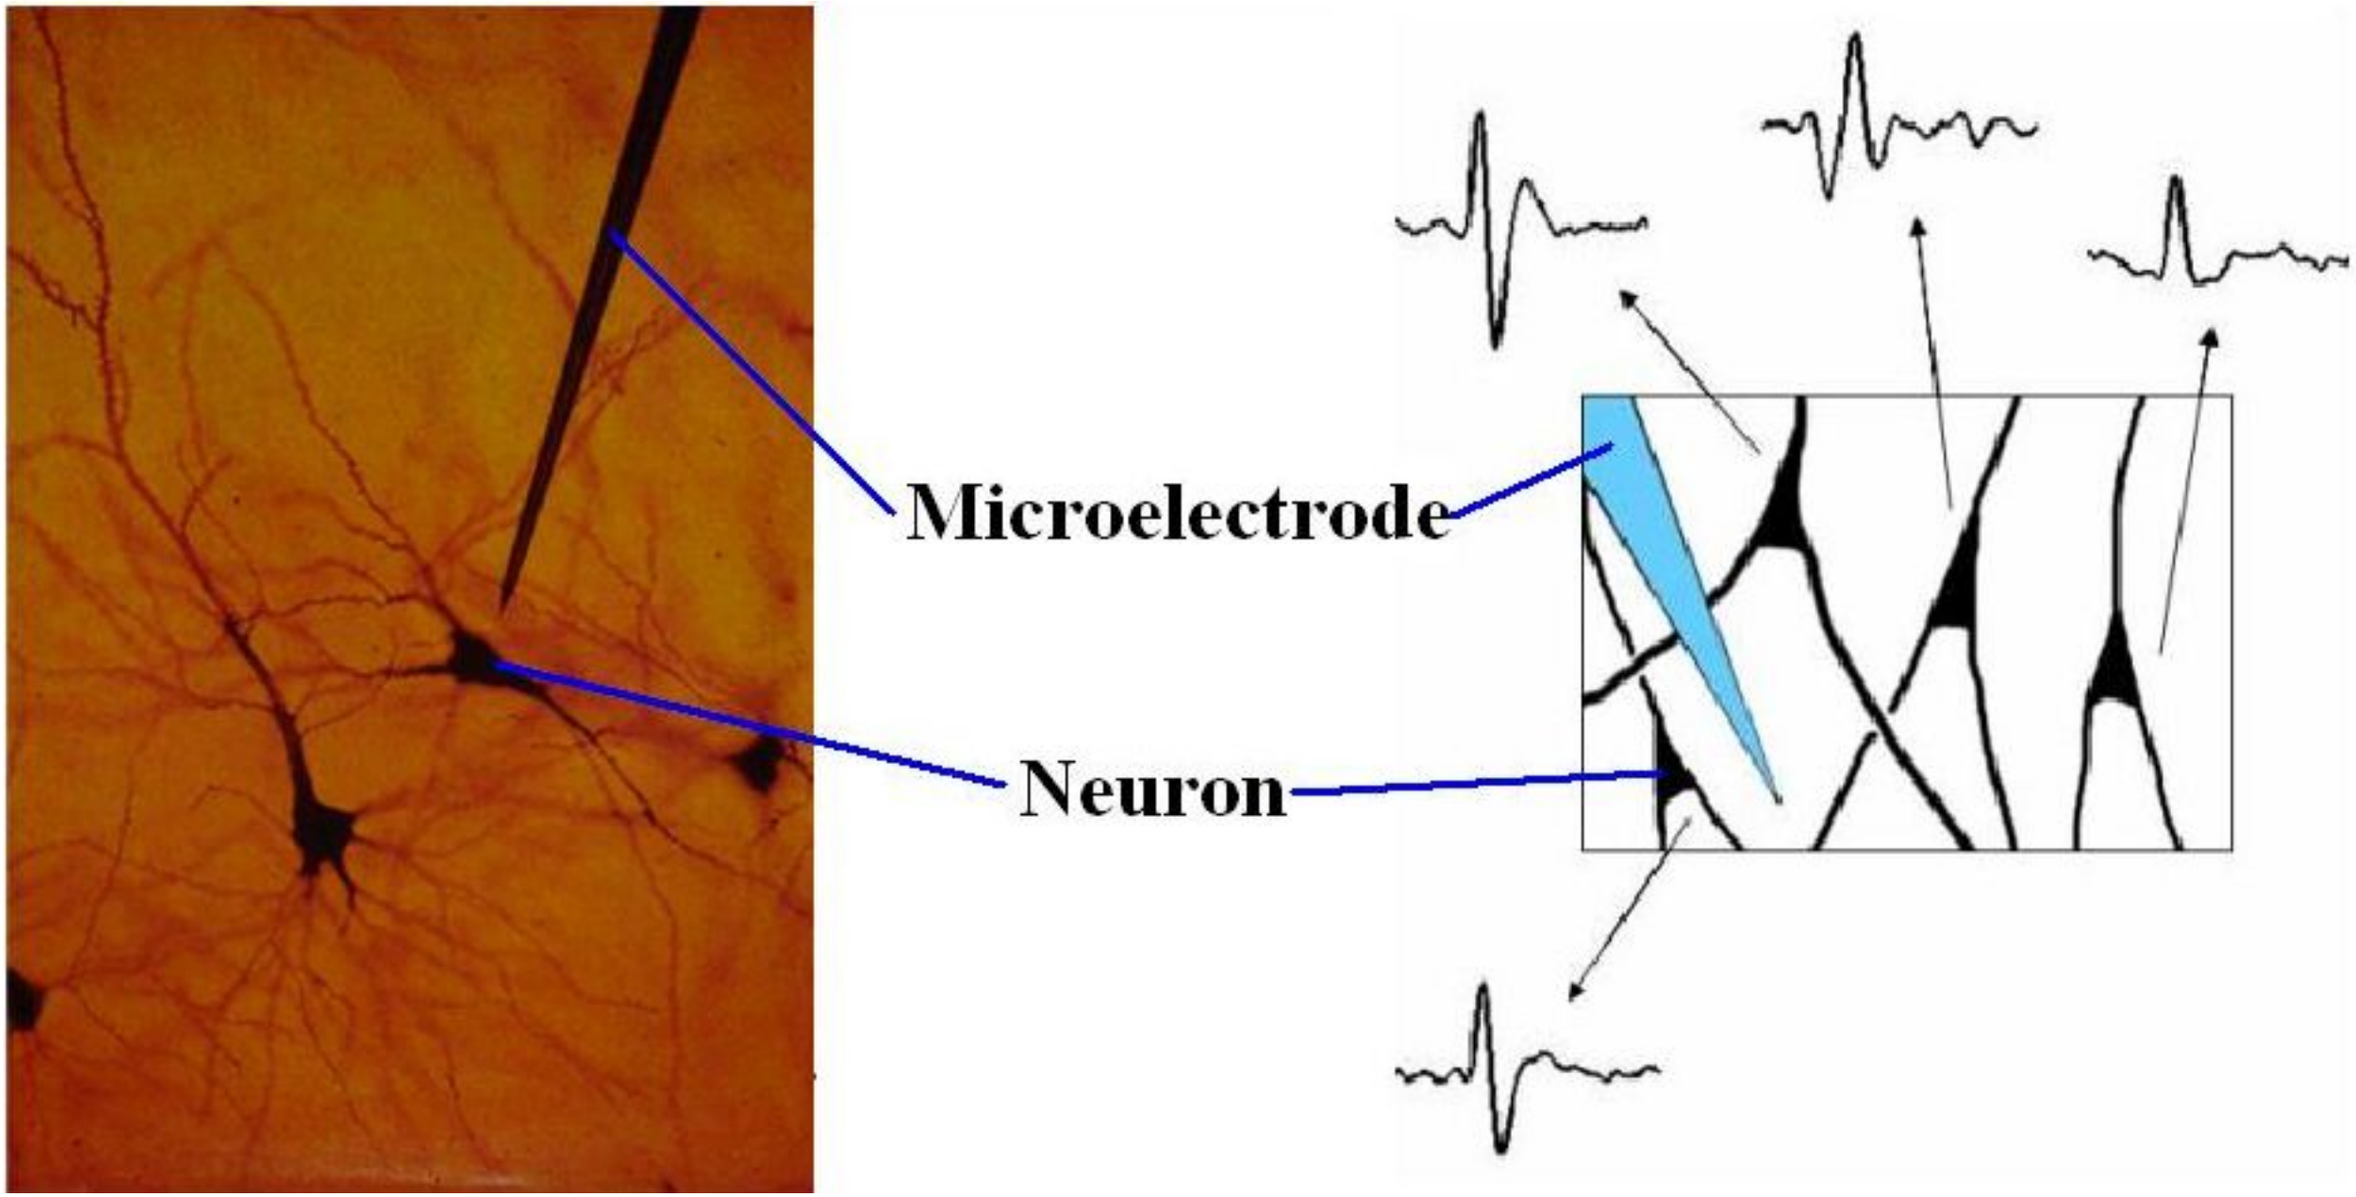

Extracellular recording

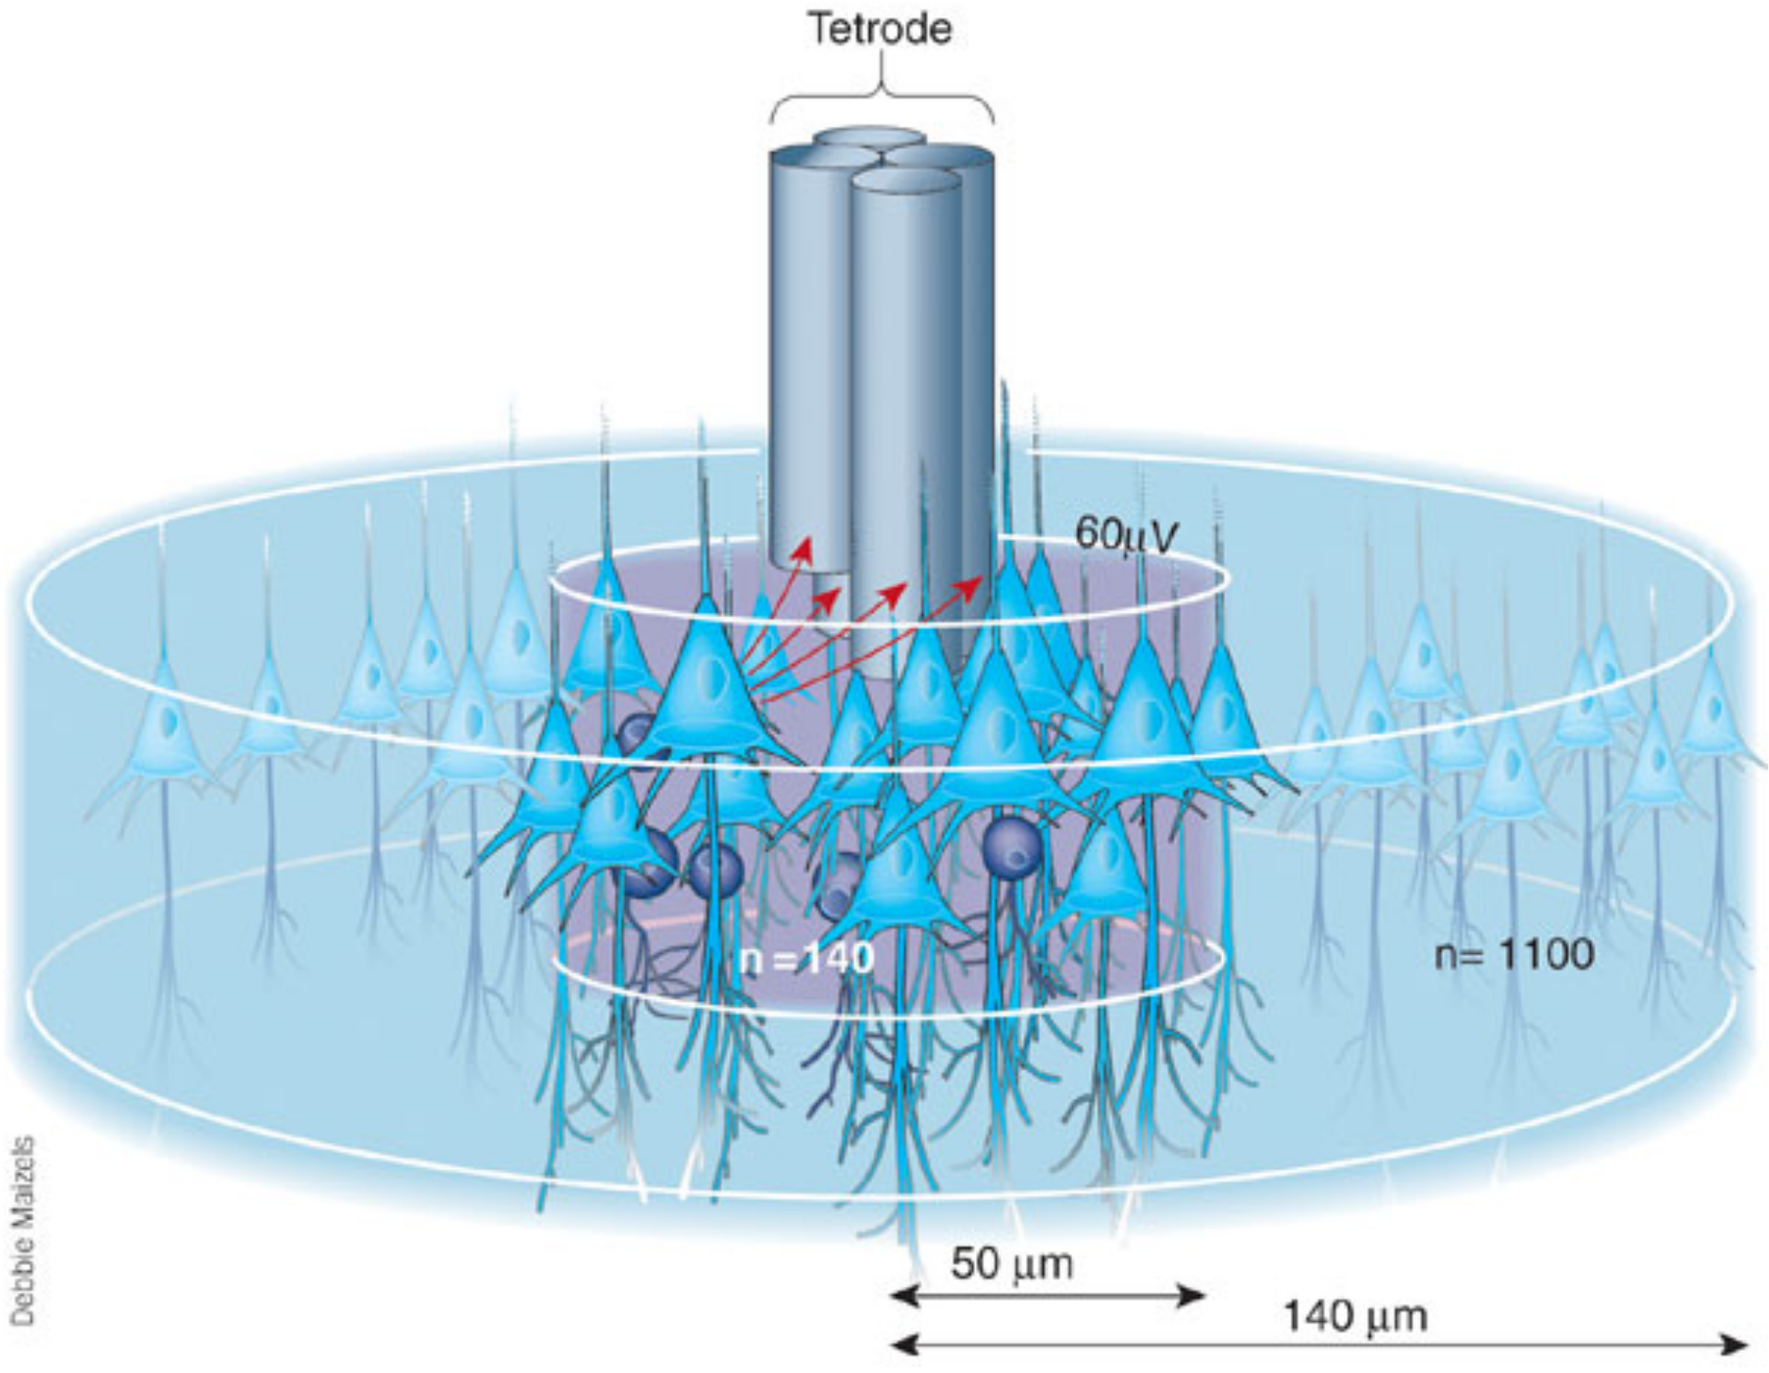

extracell recording

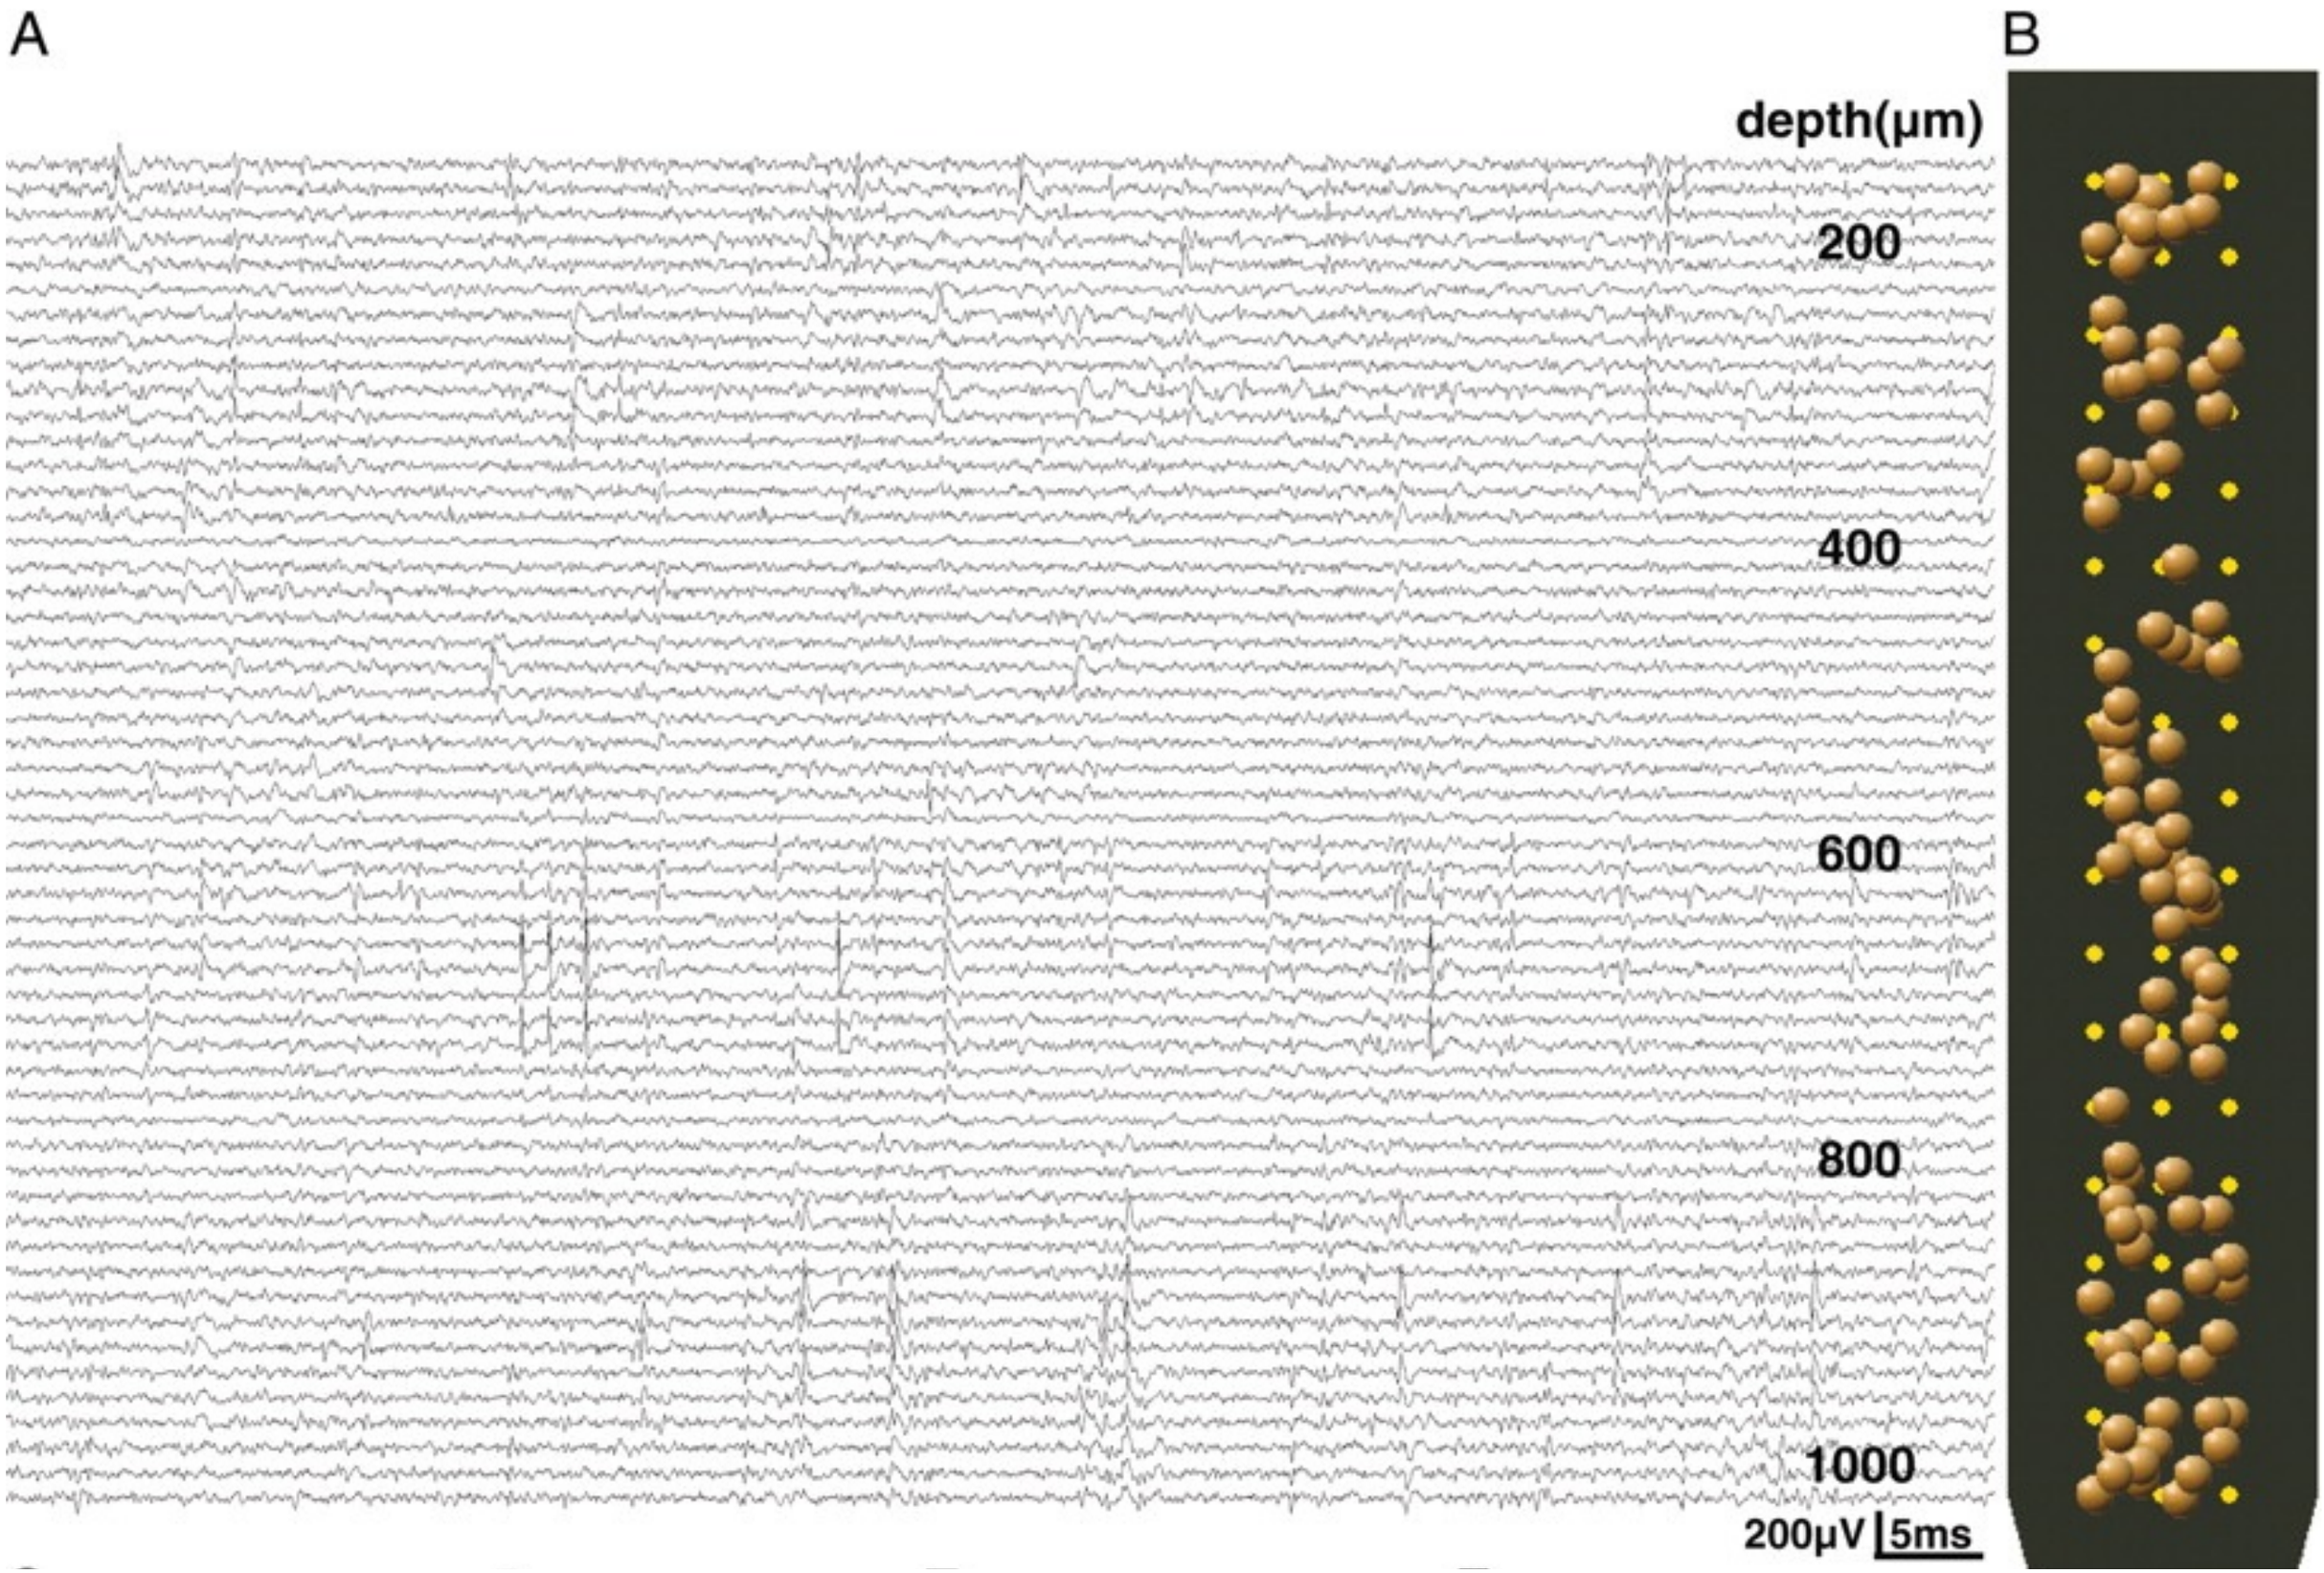

Spike sorting

A

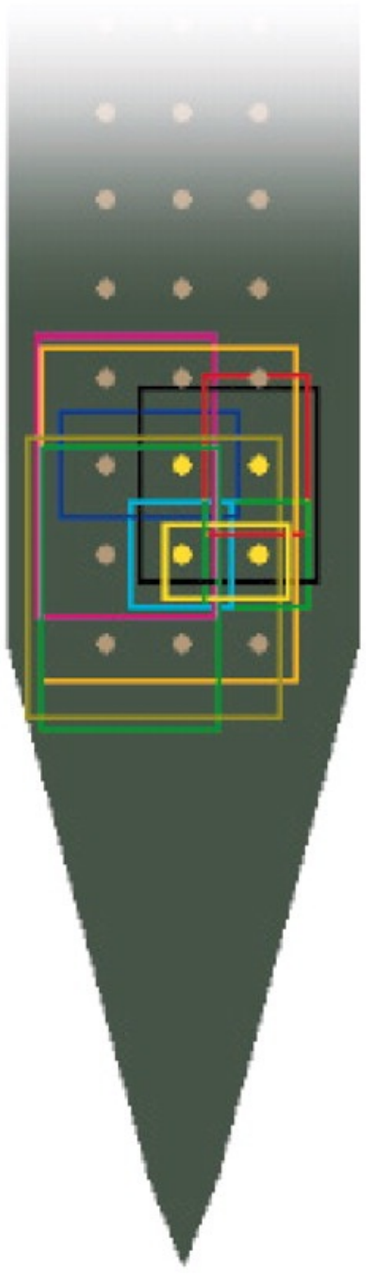

B

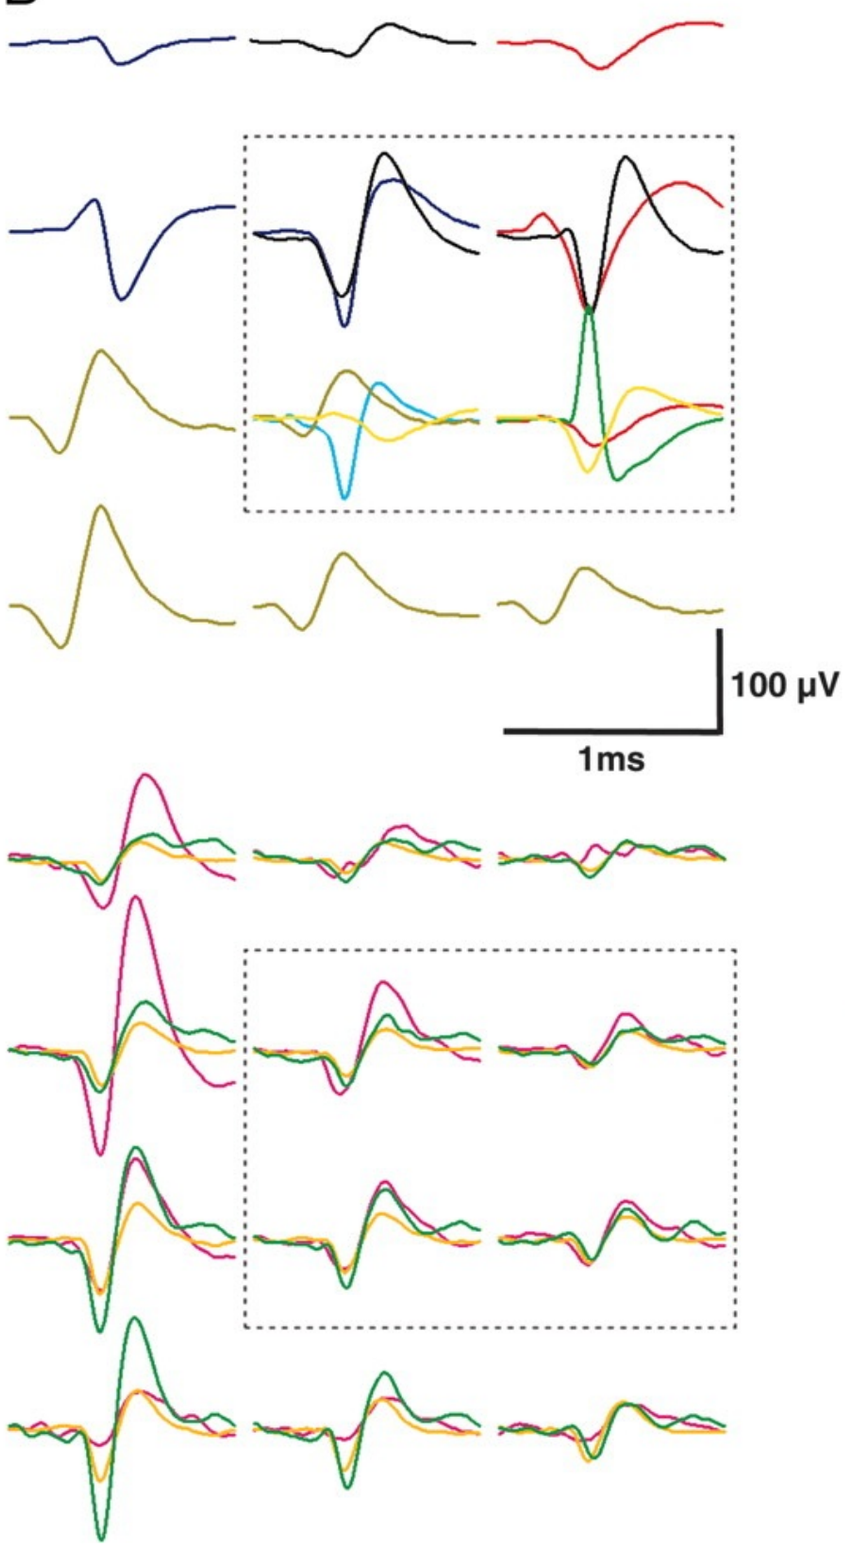

C

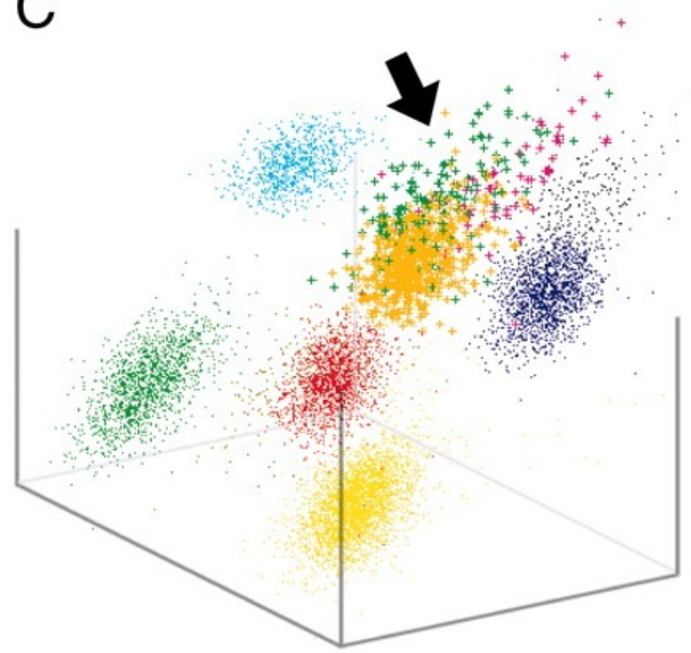

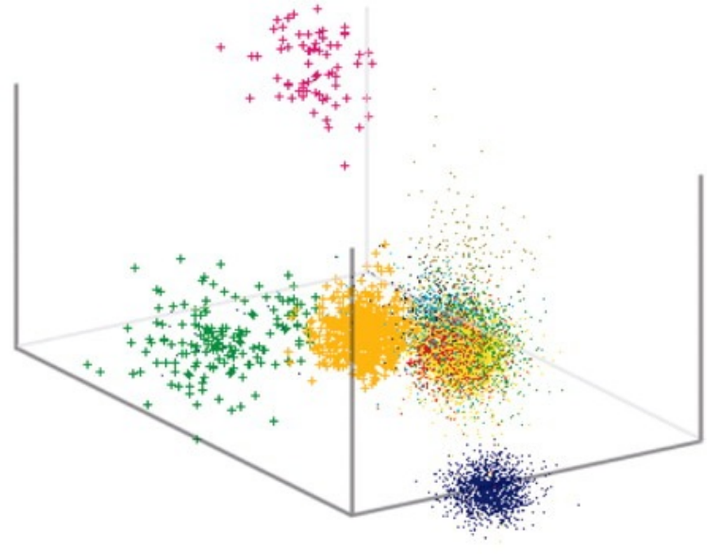

Spike sorting

- Manual and automatic methods to sort measured spikes into clusters that are assumed to correspond to cells

- Hard to evaluate the algorithms as there is little ground truth data

    - Technically quite difficult to measure simultaneously with extracellular electrodes and patch clamps from the same cells

- Introduces a noise of largely unknown properties

Calcium imaging

- genetically modified cells express fluorescent proteins that respond to the binding of calcium

- calcium levels in a cell change when spikes are generated

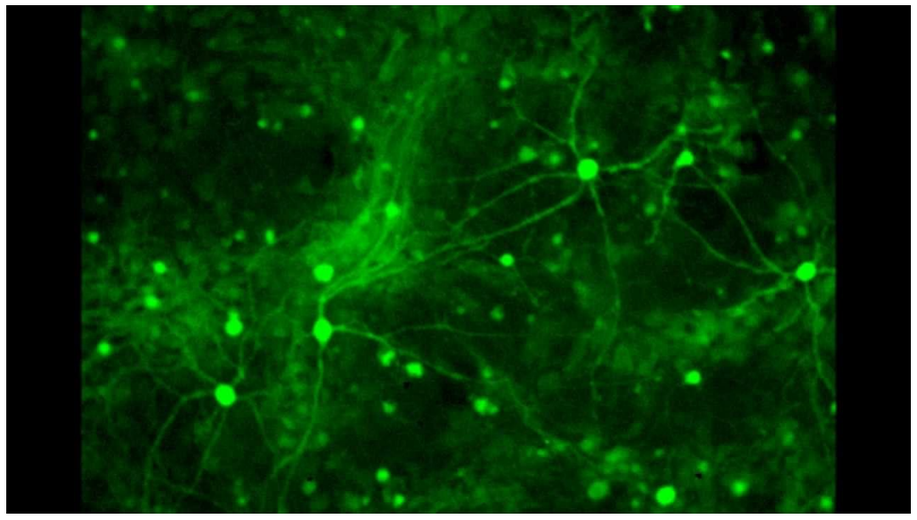

- the activity of cells can be recorded using fluorescent microscopy

- the calcium signal is slow - we can only tell wether there were any spikes in a certain time bin

- getting back the spike train for a cell is an inference problem

- we can sacrifice temporal resolution to record more cells up to ~ 10000

What kind of signals can we measure?

- How many cells at once - spatial scale
    - one ion channel -> whole brain
- With what level of detail - spatial resolution
    - one ion channel -> whole brain
- Temporal resolution
    - 20 kHz -> 1 Hz
- Trade-offs everywhere
- Neural population measurement with microscopic resolution from a behaving animal
    - ~100-1000 cells, depending on the temporal resolution, only spike trains

Publicly available datasets

- Many different datasets in the CRCNS database

- Some researchers make data available, e.g. Matthias Bethge

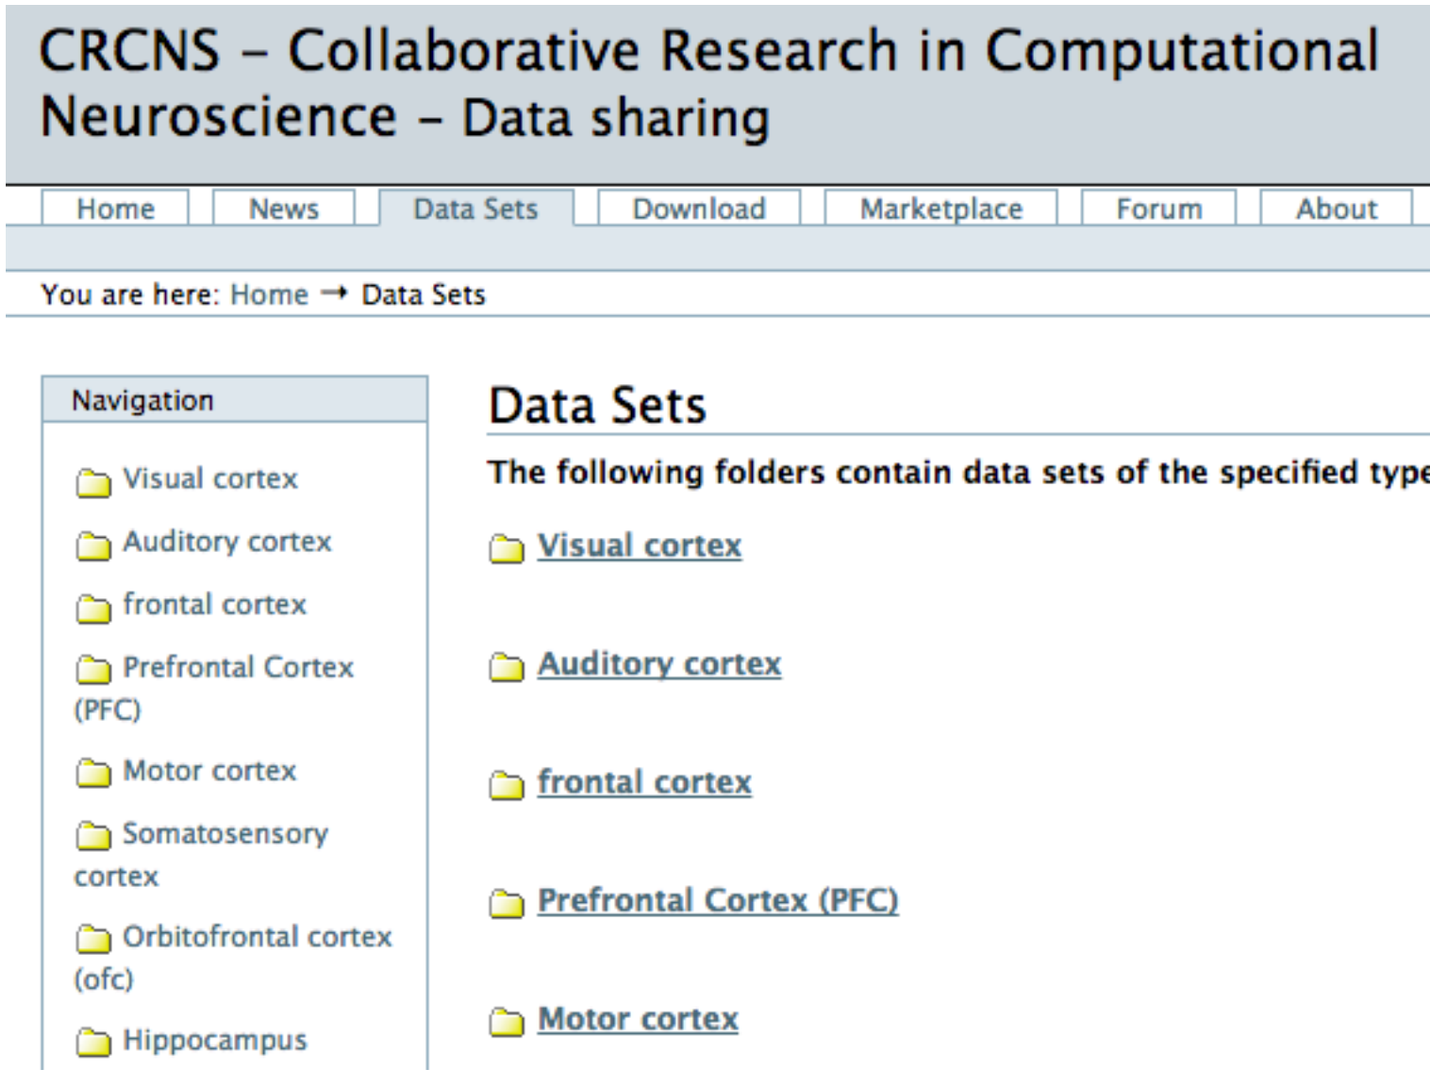

Pointer
http://bethgelab.org/datasets/v1gratings/

Pointer
http://crcns.org/

- What to predict about neural activity?
- Limits of measurability in the brain
- How to build, test and improve models?

Computation in the primary visual cortex

- in electrode recordings we can see that a lot of cells in V1 respond to bar segments placed in specific locations in the visual field

  - the strength of the response depends on how the bar is oriented

    - the **receptive field** of the cell is a localised, oriented edge

- these cells act like edge detectors

- their response can be crudely approximated by multiplying the stimulus with an edge filter

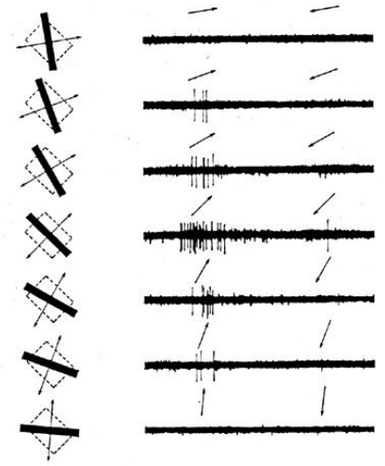

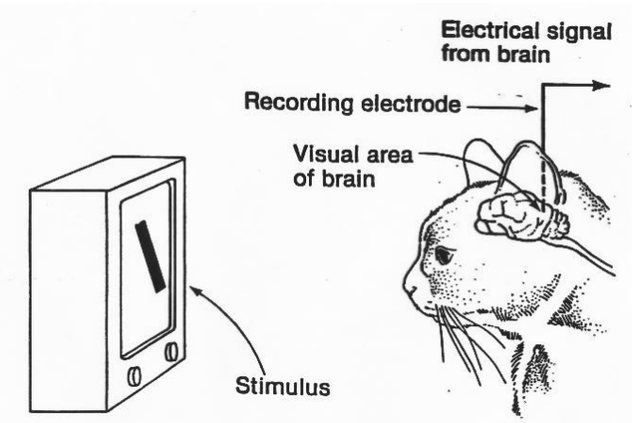

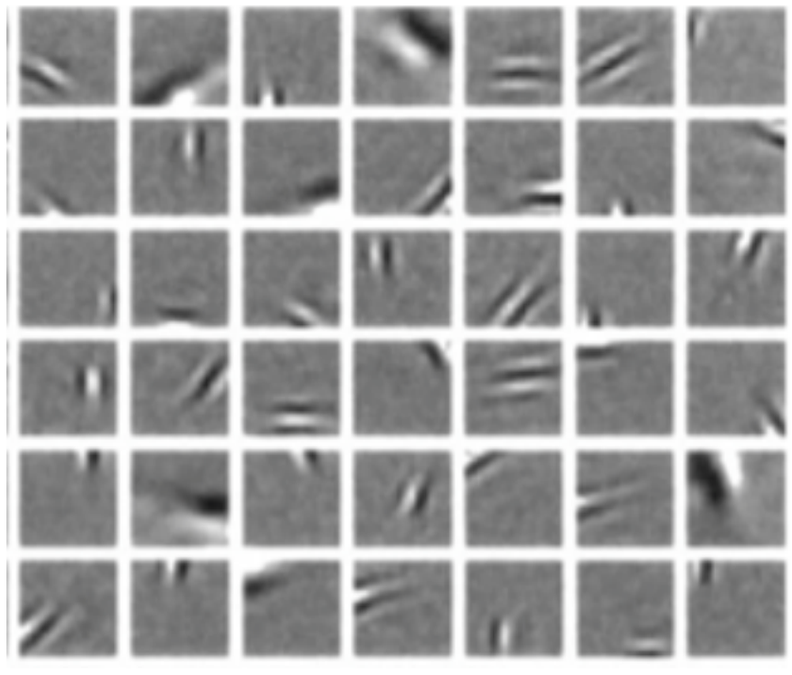

Modelling the variability of neural responses

- Can we relate it to perceptual uncertainty?

Step-by-step building of the graphical model of vision

- V1 responses are predicted by independent contributions of localised oriented edges to the stimulus
- stimulus content is independent of lighting conditions, thus contrast should be an independent modulation of the edge combinations
- As edges define object contours, they do not appear independently - a covariance matrix needs to be learned for them.

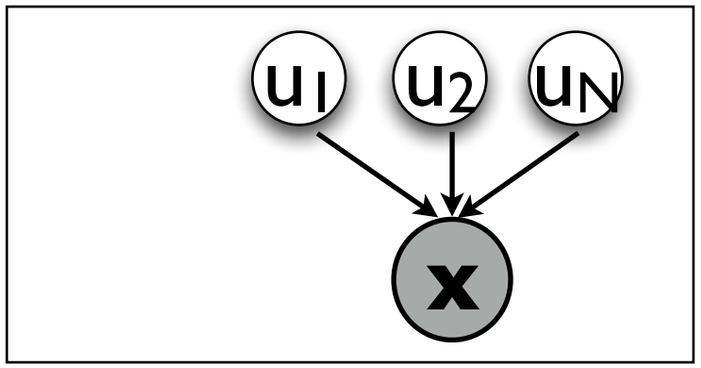

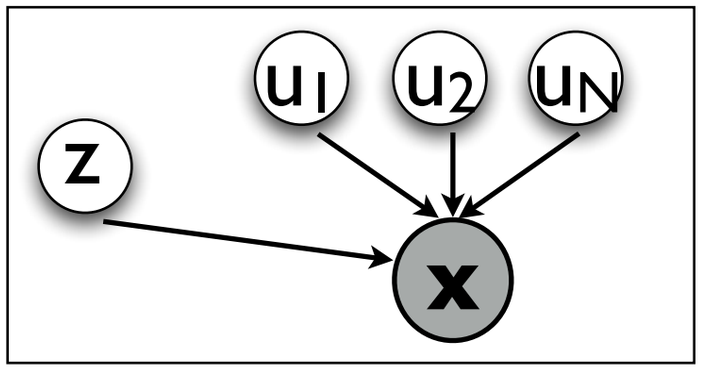

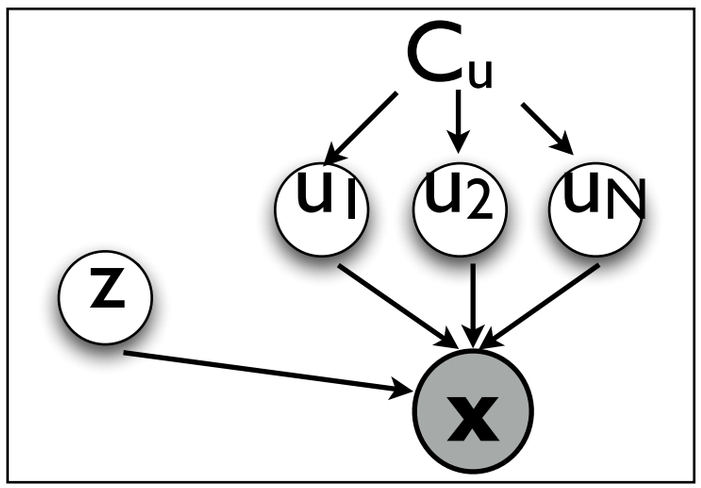

"Never attribute to stupidity that which is adequately explained by unstated assumptions."

Geert Bollen

Pointer
http://xcorr.net/2015/11/20/turing-machines-the-number-game-and-inference/

Contrast-normalised models of V1

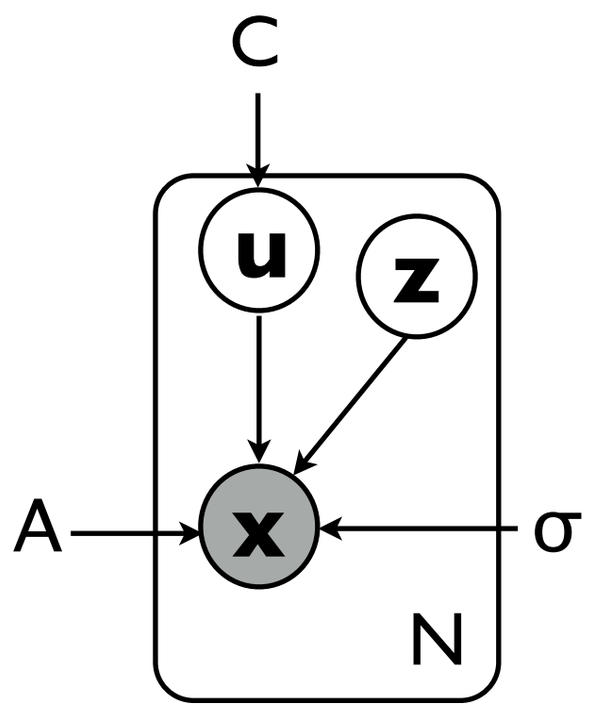

- the observed variable **x** encodes the pixels of a black&white image - an approximation of retinal rod sensor activations
- hidden variable **u** encodes how strongly specific localised edges contribute to the composition of the image - we will use this to predict the membrane potential of V1 neurons
- hidden variable **z** encodes the contrast of the image
- called Gaussian scale mixture models

$$\begin{aligned} \mathrm{P}(\mathbf{u}) &= \mathcal{N}(\mathbf{u}; \mathbf{0}, \mathbf{C}) \\ \mathrm{P}(z) &= \mathrm{Gamma}(z; k, \theta) \end{aligned}$$

$$p(\mathbf{x} \mid \mathbf{u}, z) = \mathcal{N}(\mathbf{x}; zA\mathbf{u}, \sigma^2 I)$$

$$\text{image} = \text{contrast} \times \left(a_1 \, \text{feature}_1 + a_2 \, \text{feature}_2 + \ldots + a_N \, \text{feature}_N \right) + \text{noise}$$

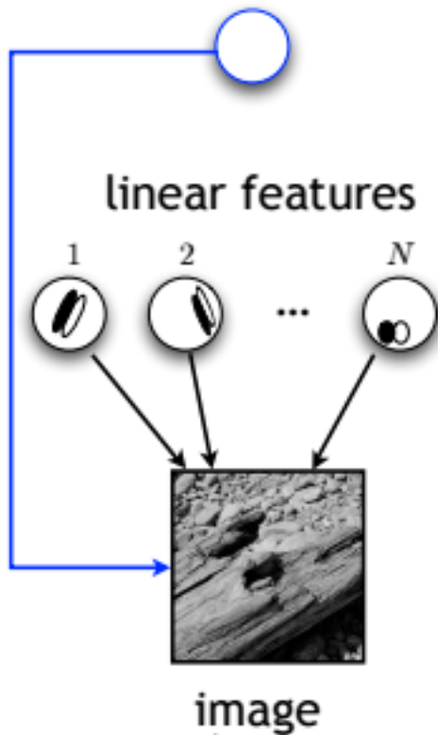

GSM predictions

- How do neural activations change with stimulus contrast?

Spontaneous and evoked activity

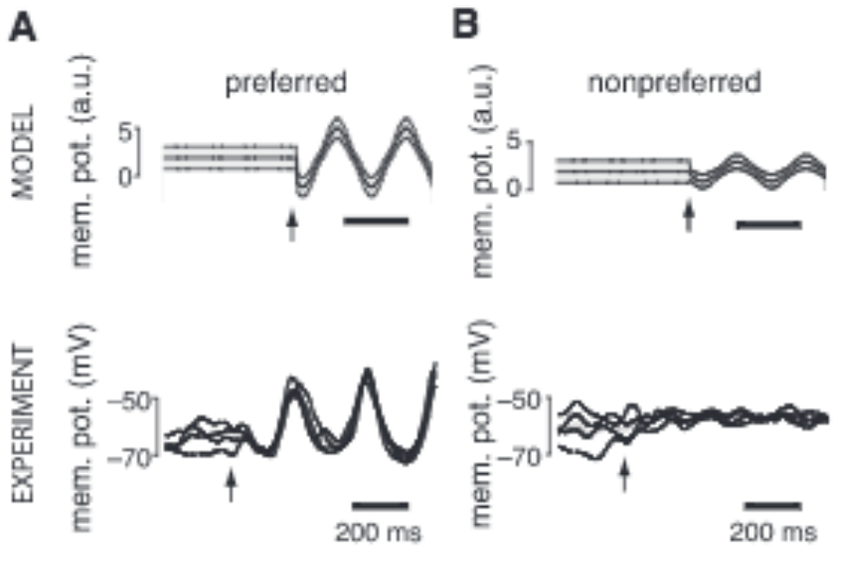

Orbán et al, 2016, Neuron

GSM predictions

- How does neural spiking change with stimulus contrast?

Response variance at different contrast levels

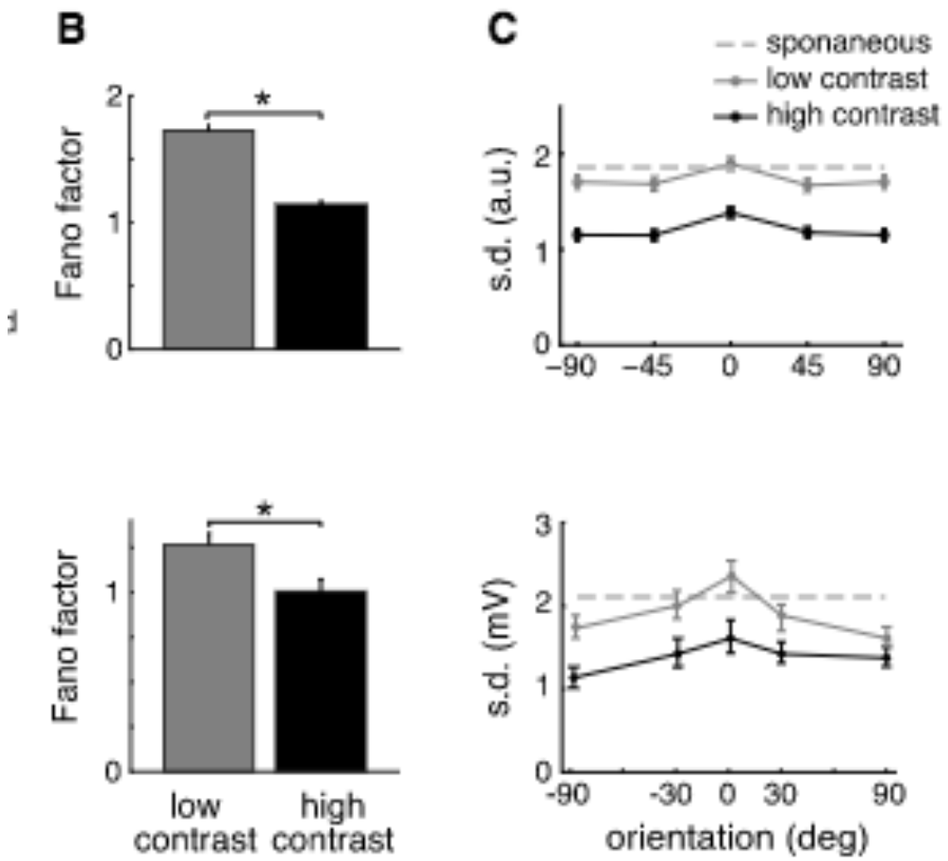

Orbán et al, 2016, Neuron

Hierarchy of hypotheses about the brain

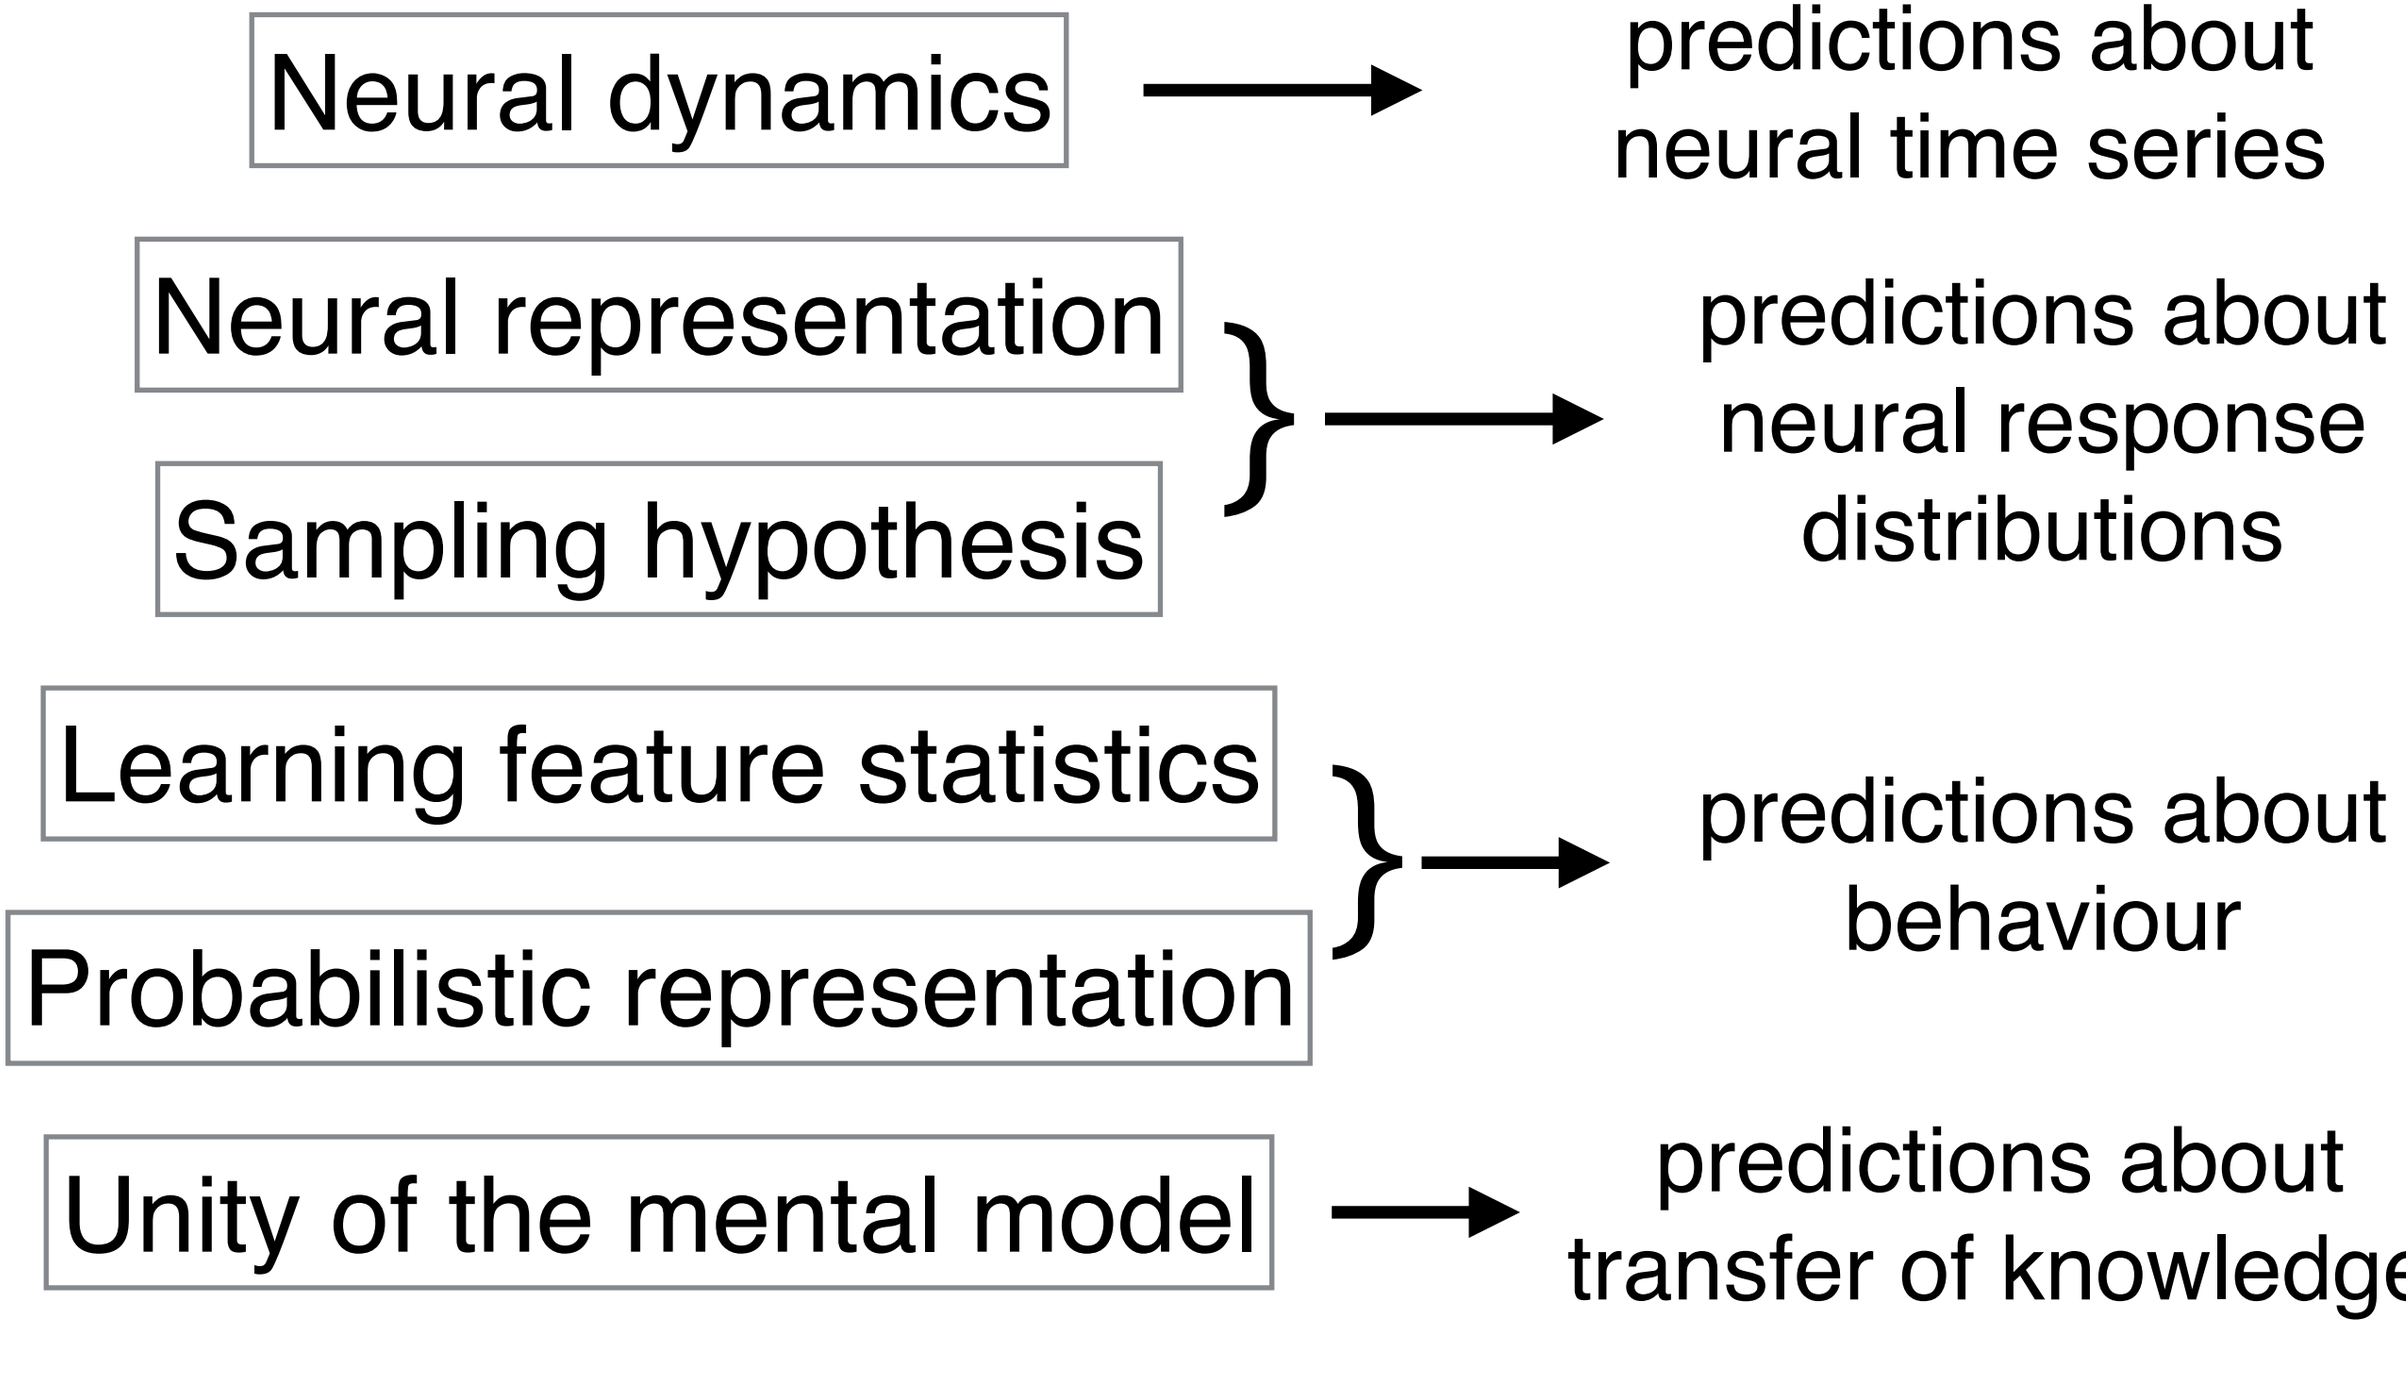

Moving towards the implementation level

- With the GSM model, we made predictions on the algorithmic level of abstraction, as we showed that evoked responses are compatible with probabilistic inference

- In order to really move to the implementation level, additional to the generative model, we have to assume a specific inference algorithm too

- Such an algorithm may give predictions not only about properties of response distributions, but actual time series of neural membrane potentials, that can be compared to measurements

- The sampling hypothesis suggests that the inference algorithm will be one that produces samples from the posterior distribution of the latent variables

- There are many such sampling algorithms with different properties, we choose one that is efficient and lends itself to implementation with neural networks

Measurement time series predictions with GSM

model

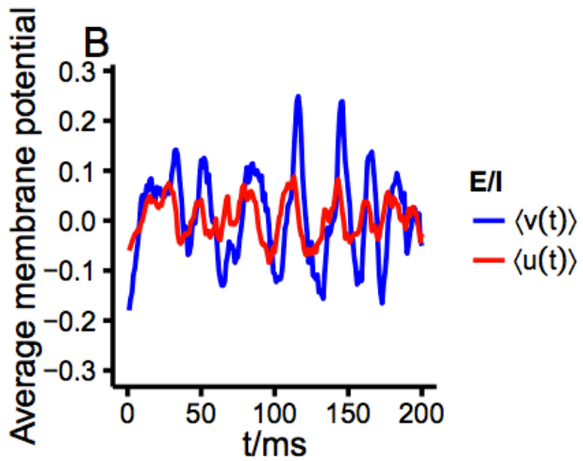

measurement

Ray & Maunsell, 2010

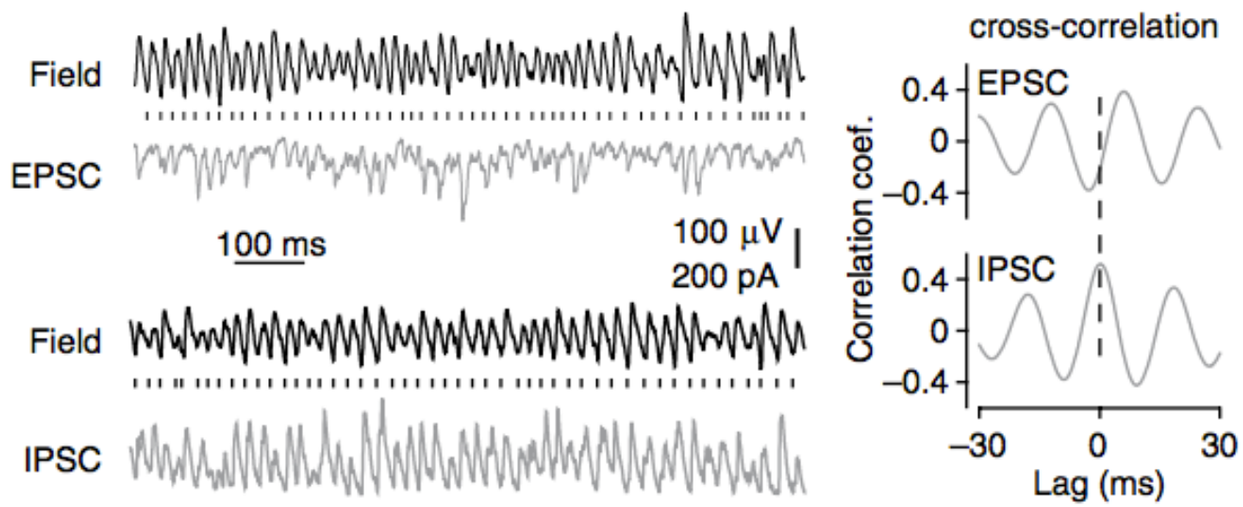

- implementation of an efficient sampling algorithm by a biophysically realistic neural network model

- inhibitory and excitatory subpopulations exhibit oscillatory dynamics at different amplitudes, and are shifted in phase relative to each other

- stimulus onset evokes a transient increase in firing rates

model

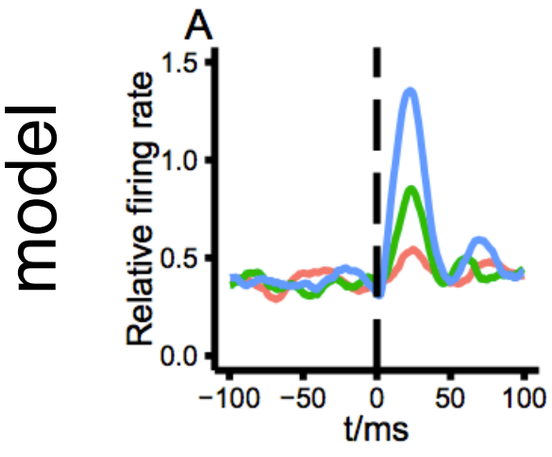

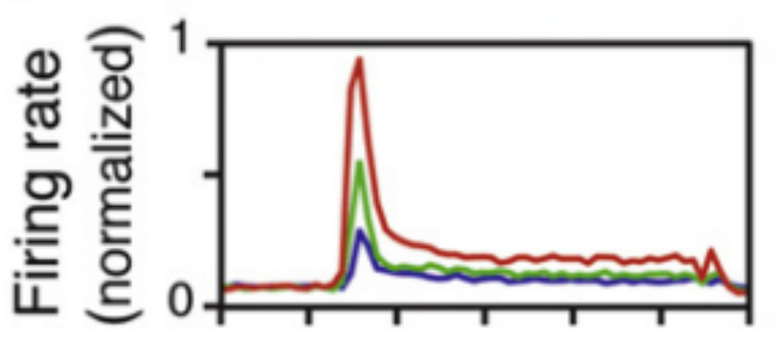

measurement

Mann & Mody, 2010

Aitchison, L., & Lengyel, M. (2014). The Hamiltonian Brain. arXiv preprint arXiv:1407.0973.

Modelling the correlation structure of neural responses

- Can we relate it to perceptual context?

Perception as hierarchical inference

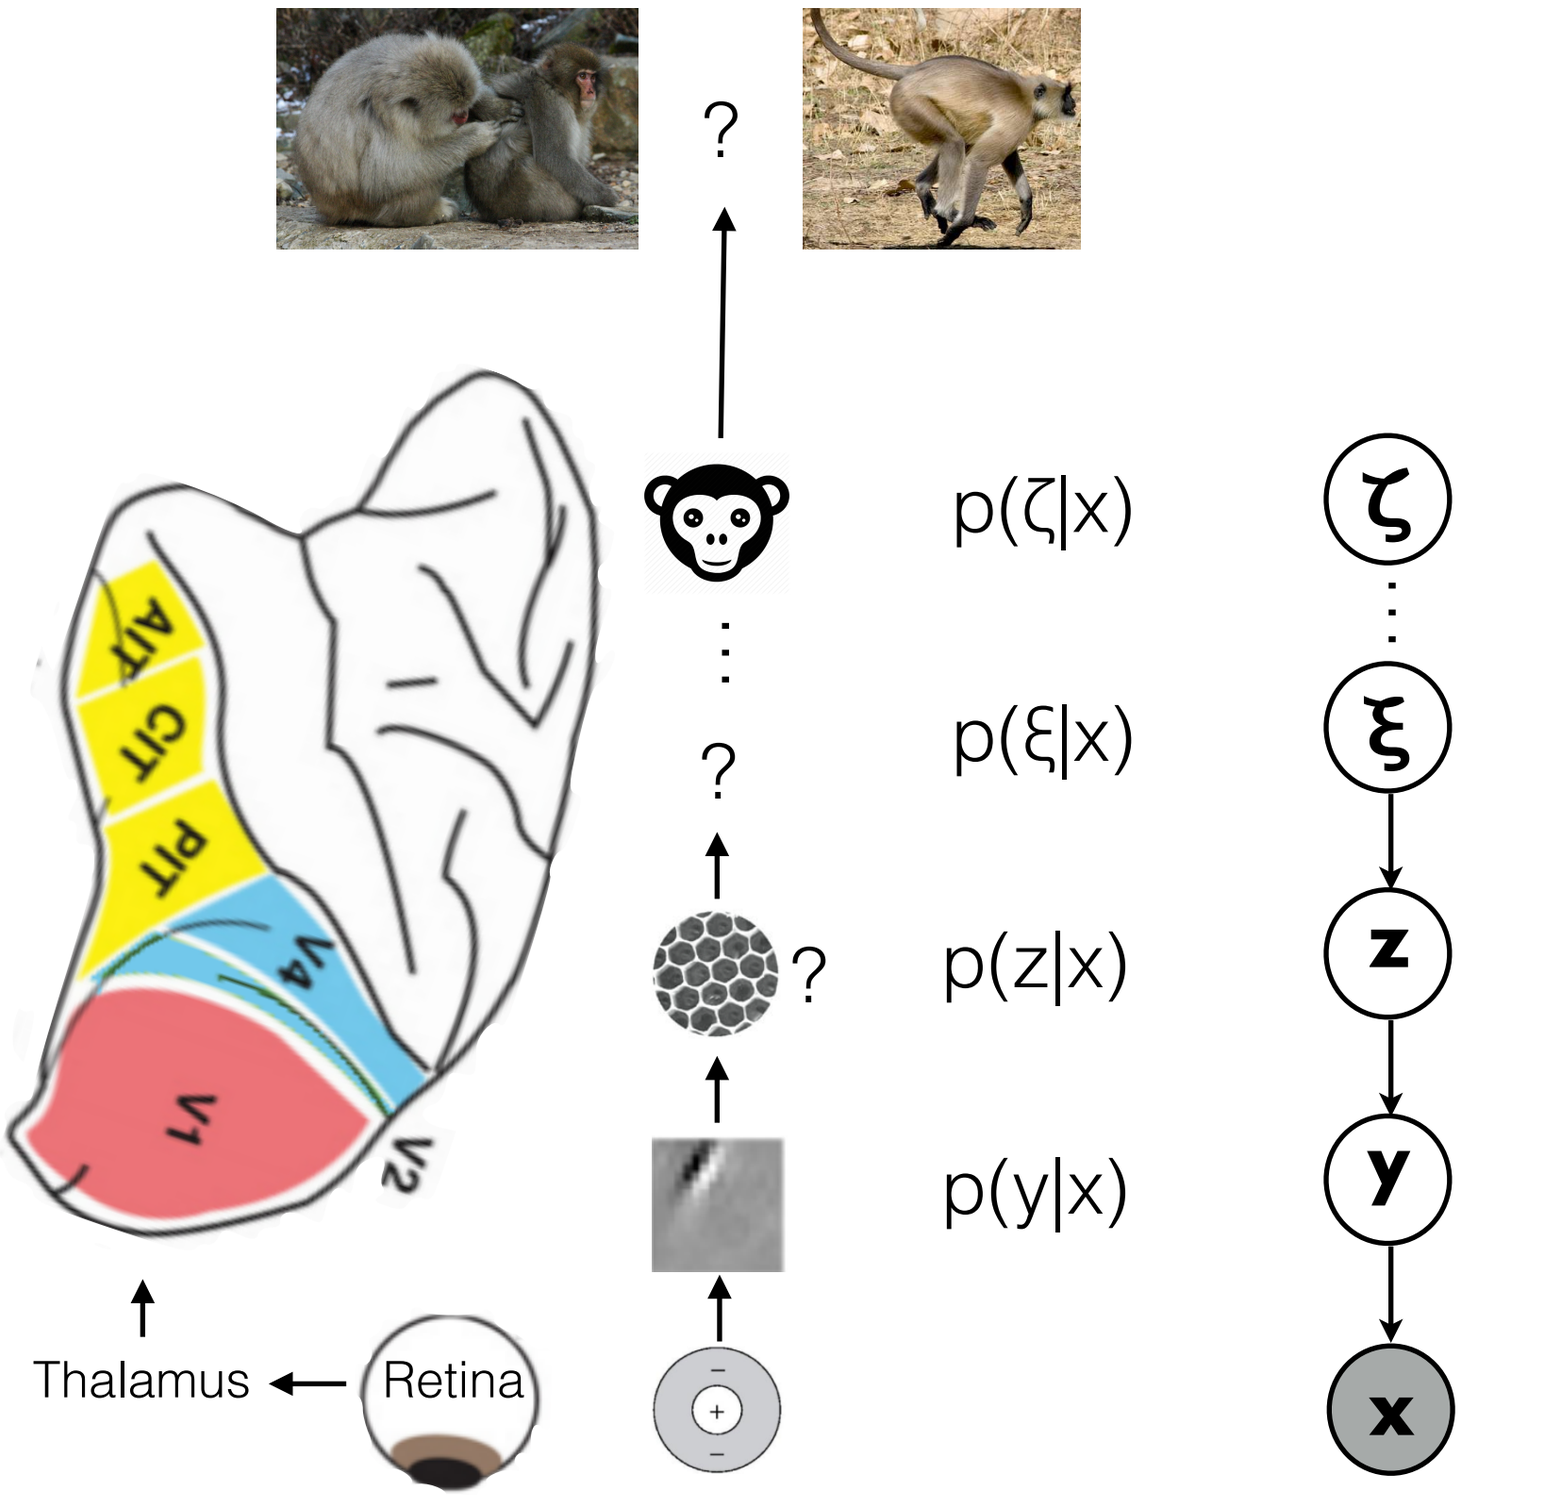

Measured top-down effect in mean responses

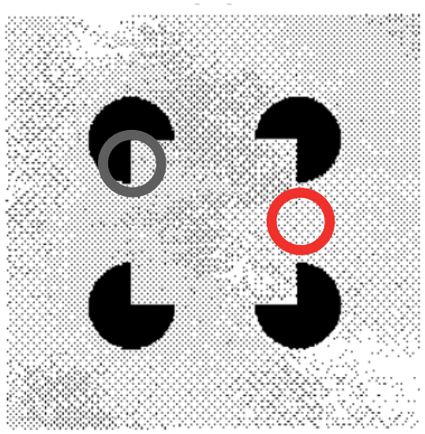

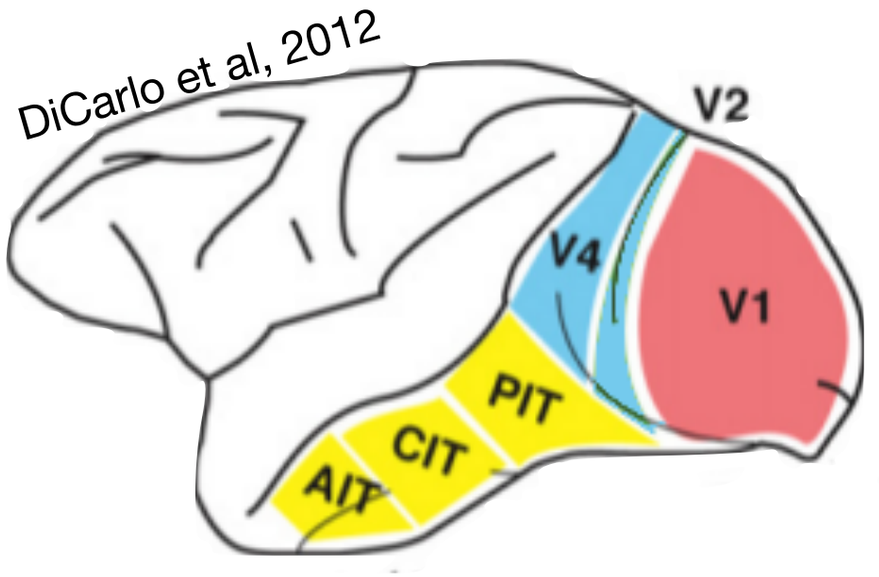

Responses of a V1 simple cell

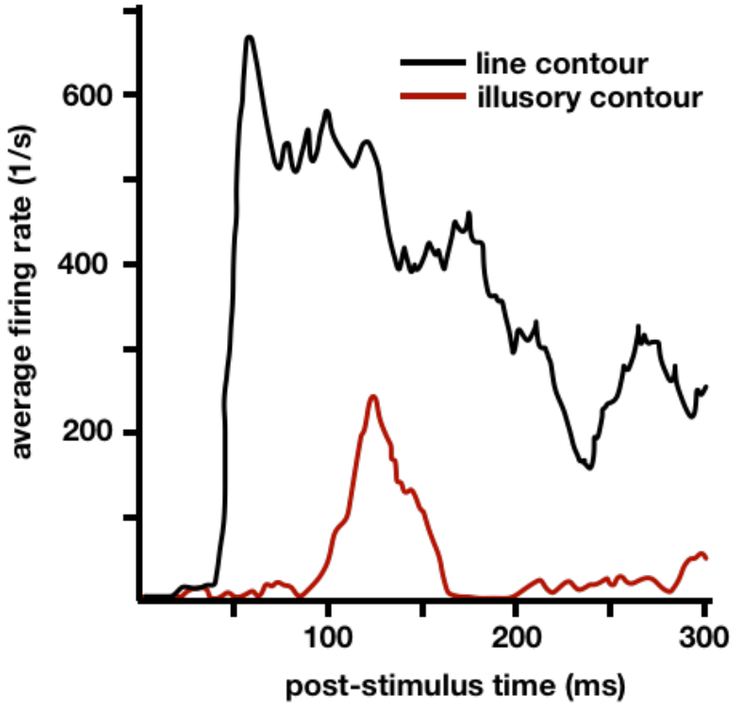

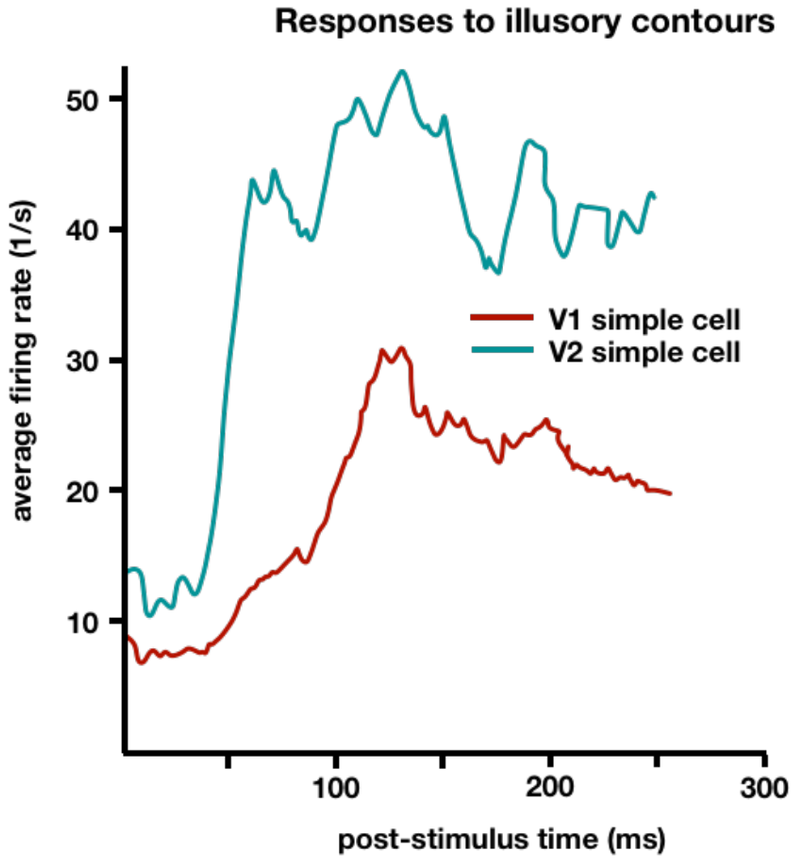

Lee & Nguyen, 2001
Lee & Mumford, 2003

Predictions of context-dependent models

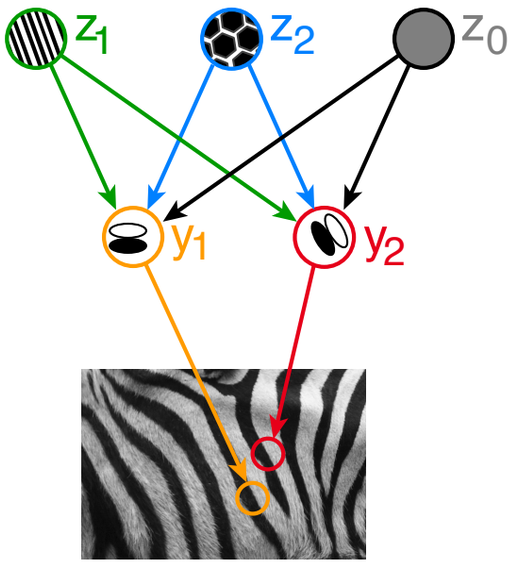

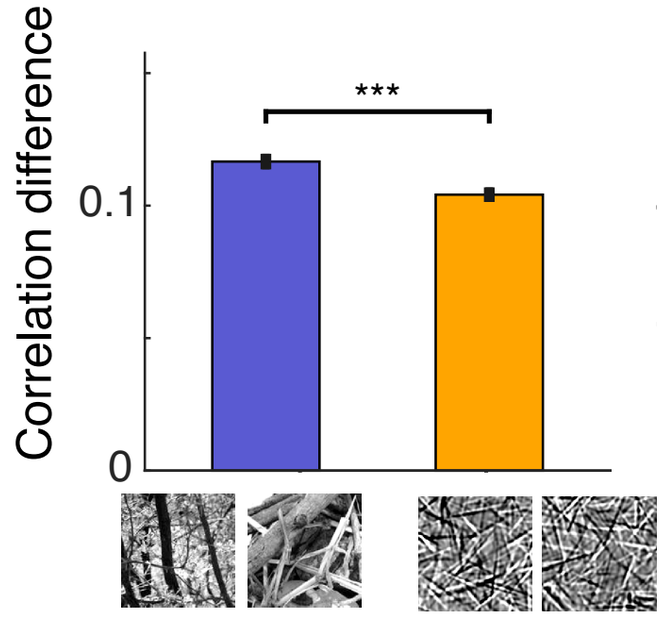

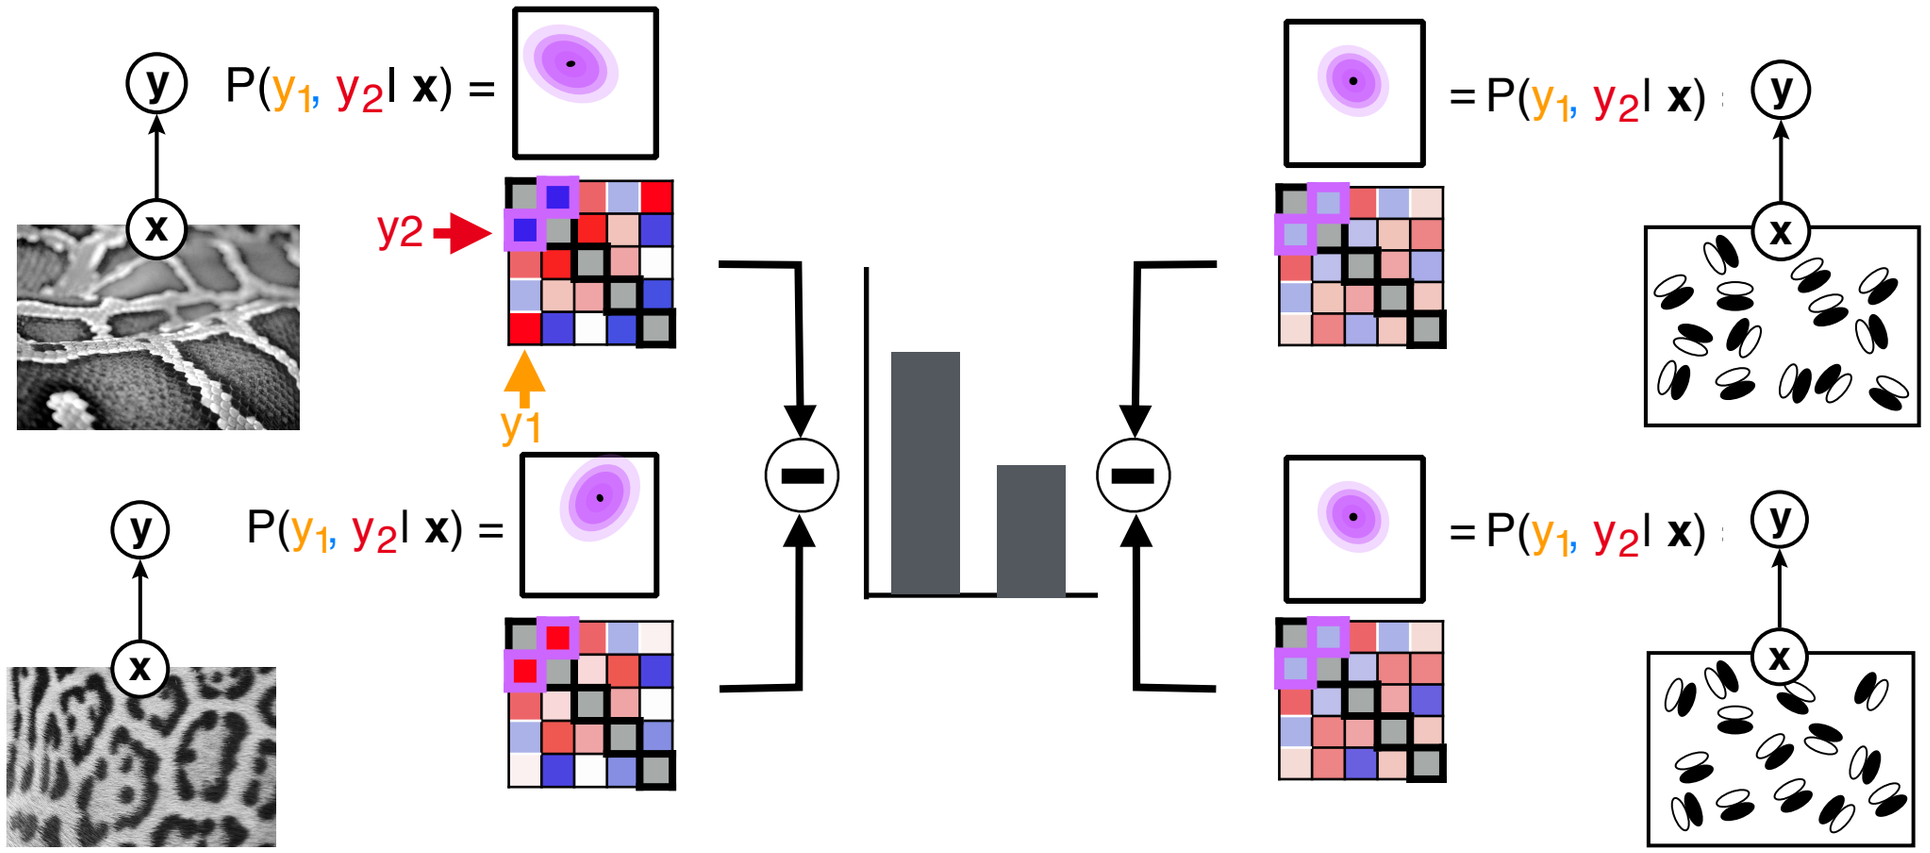

Bányai et al, 2019

A generative model for audition

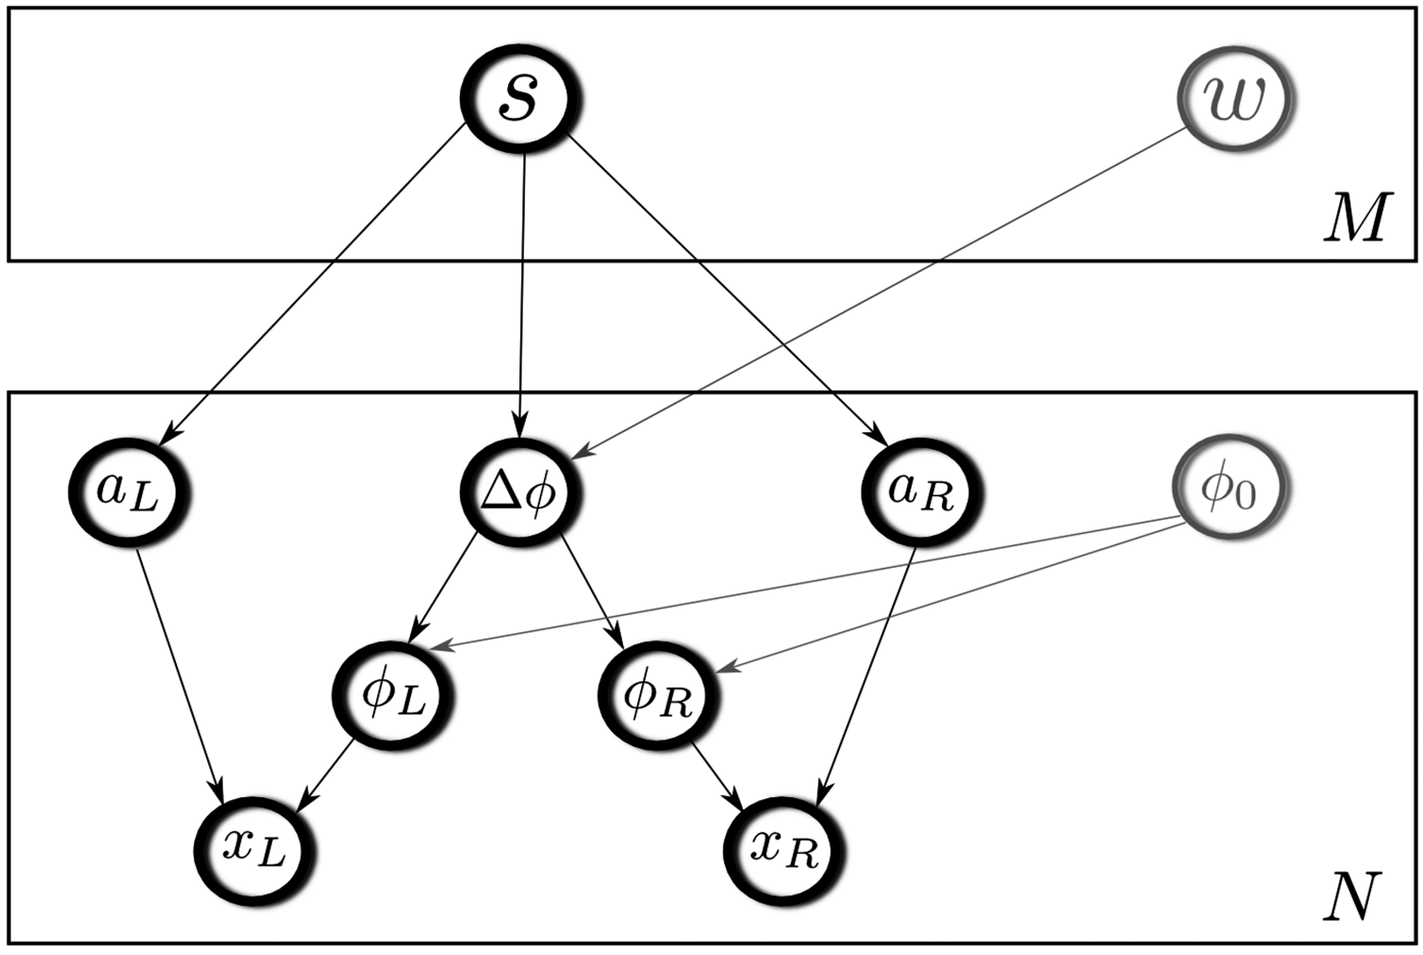

The lowest layer represents sound epochs perceived by the left and the right ear xL and xR. They are decomposed by a sparse coding algorithm into phase and amplitude vectors ϕL, ϕR and aL, aR. Phases are further subtracted from each other in order to obtain an intramural phase difference (IPD) vector Δϕ. The second layer encodes jointly monaural amplitudes and IPDs. Auxiliary variables (phase offset and the scaling factor w) are depicted in gray.

Wiktor Młynarski, 2015

Predictions of the auditory model

- Experimentally measured responses to sounds originating from different directions measured in the A1 area of the cat in two types of neurons.

- Model predictions sorted into two typical clusters of variables. Thin gray lines are values of single variables, while thick black lines depict cluster averages

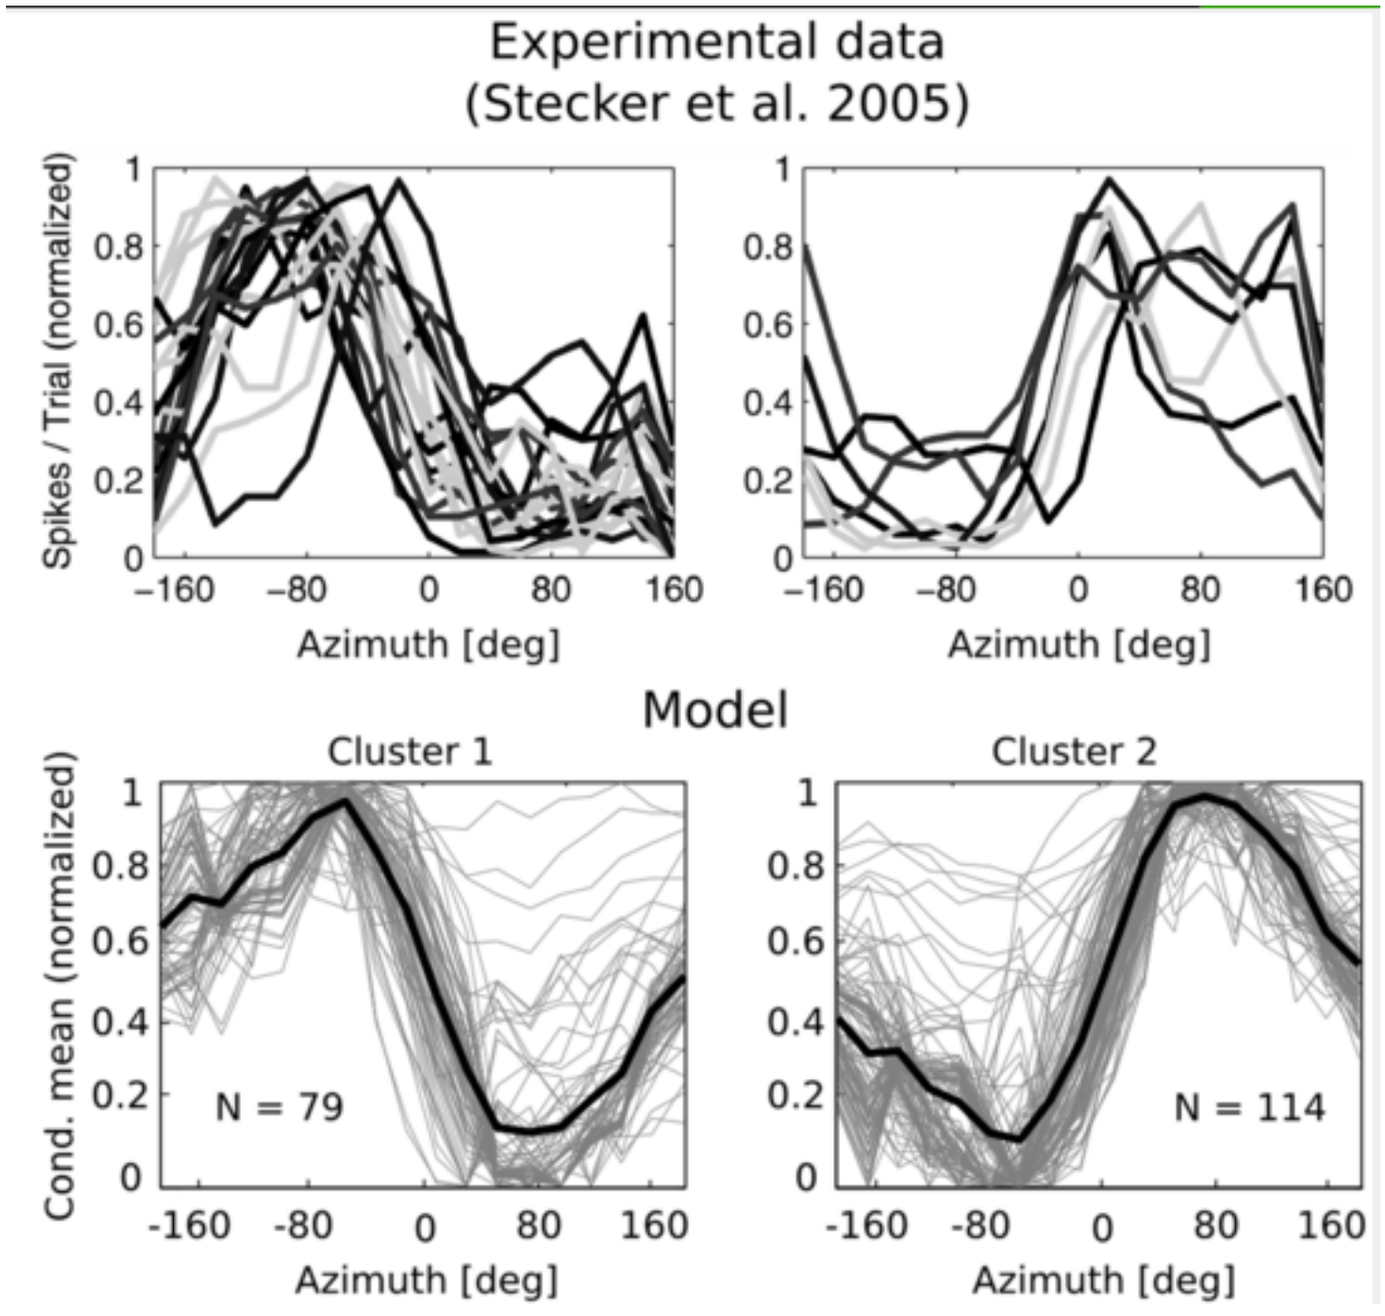

Wiktor Młynarski, 2015

Factorising relational and sensory information

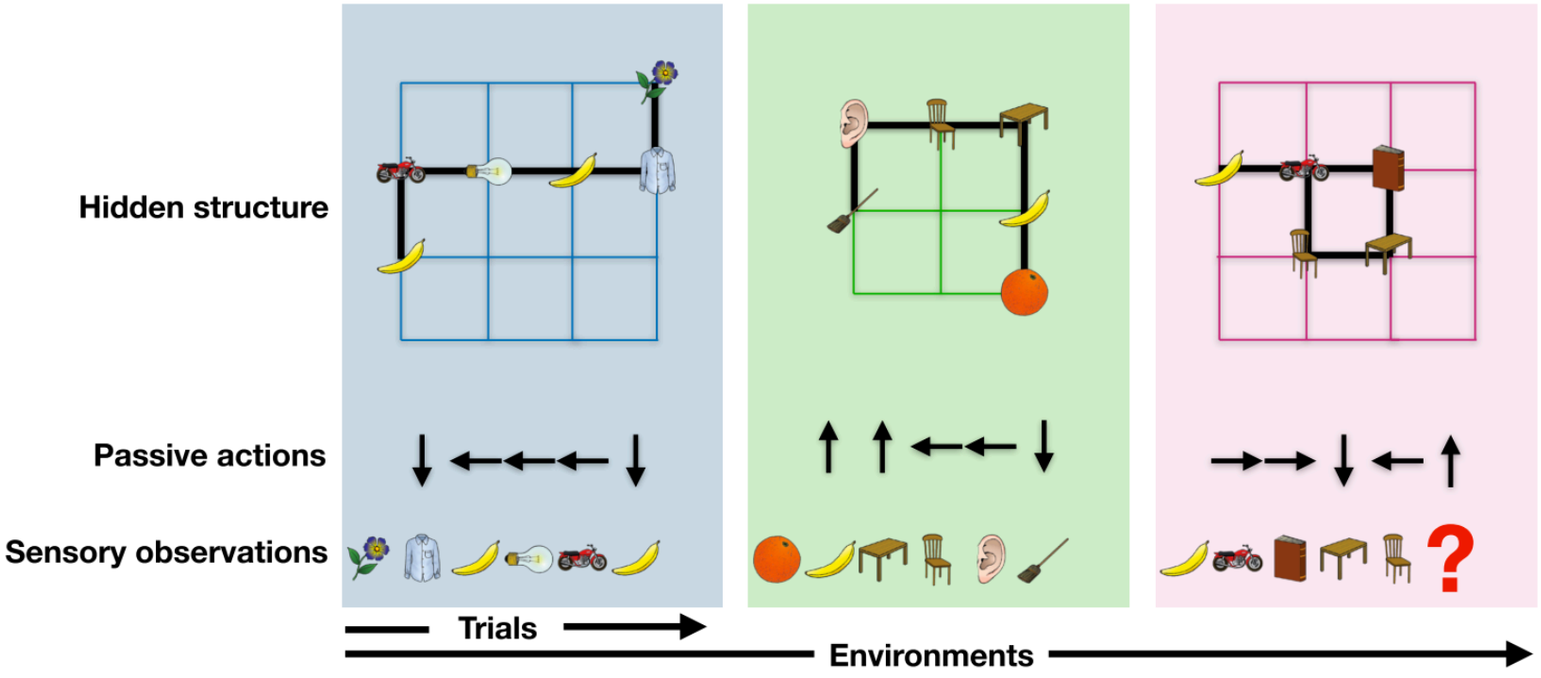

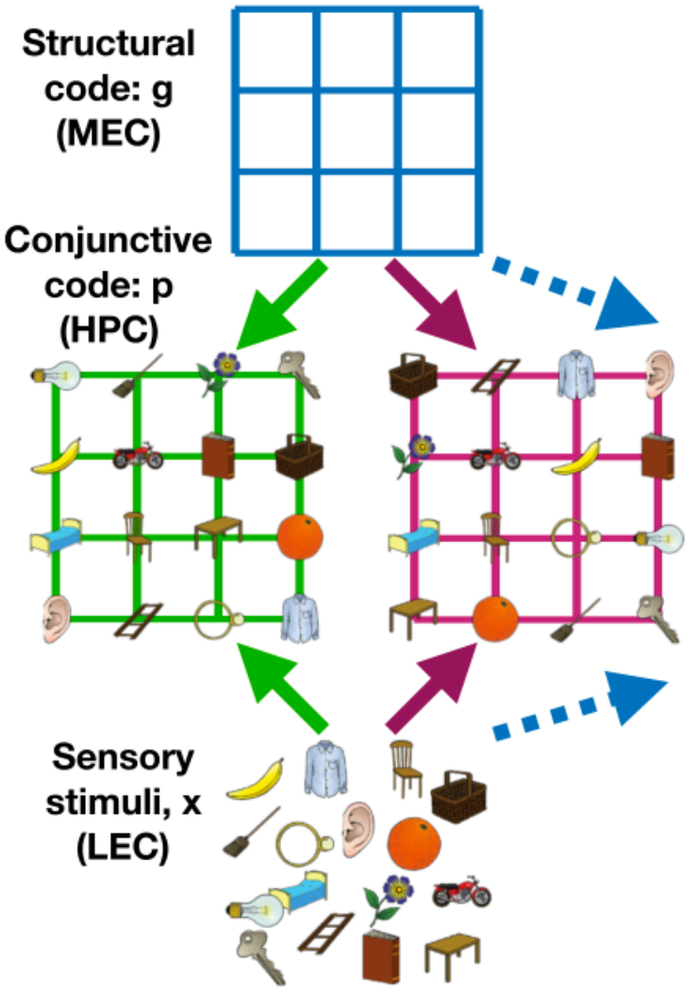

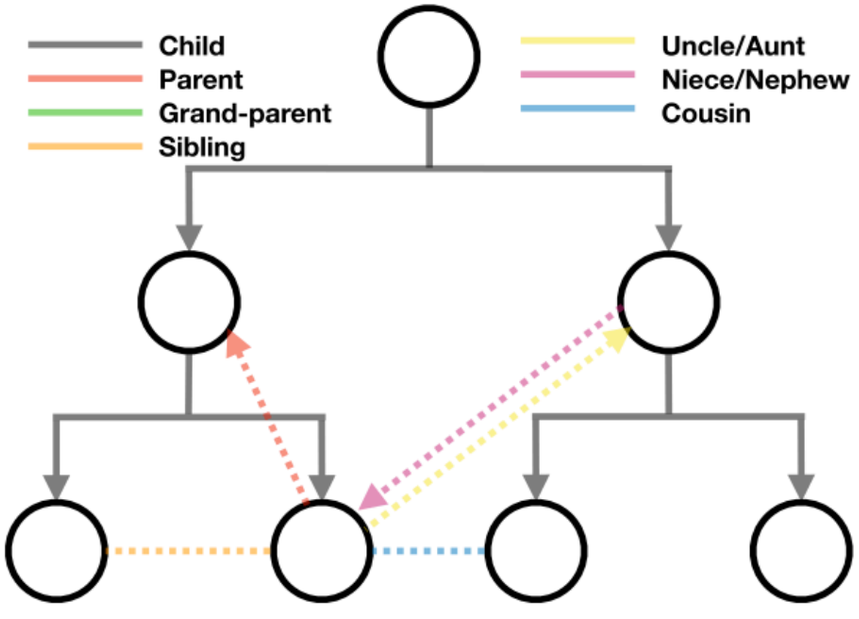

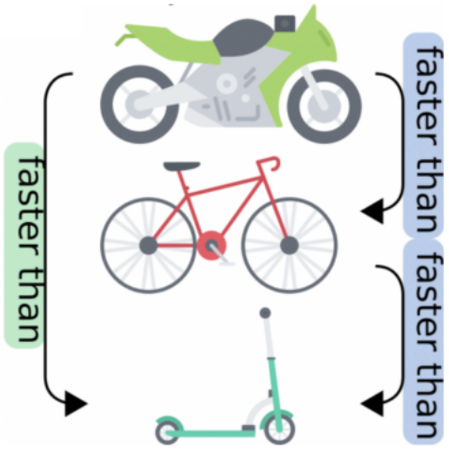

Whittington et al, 2019, bioRxiv

The Tolman-Eichenbaum machine

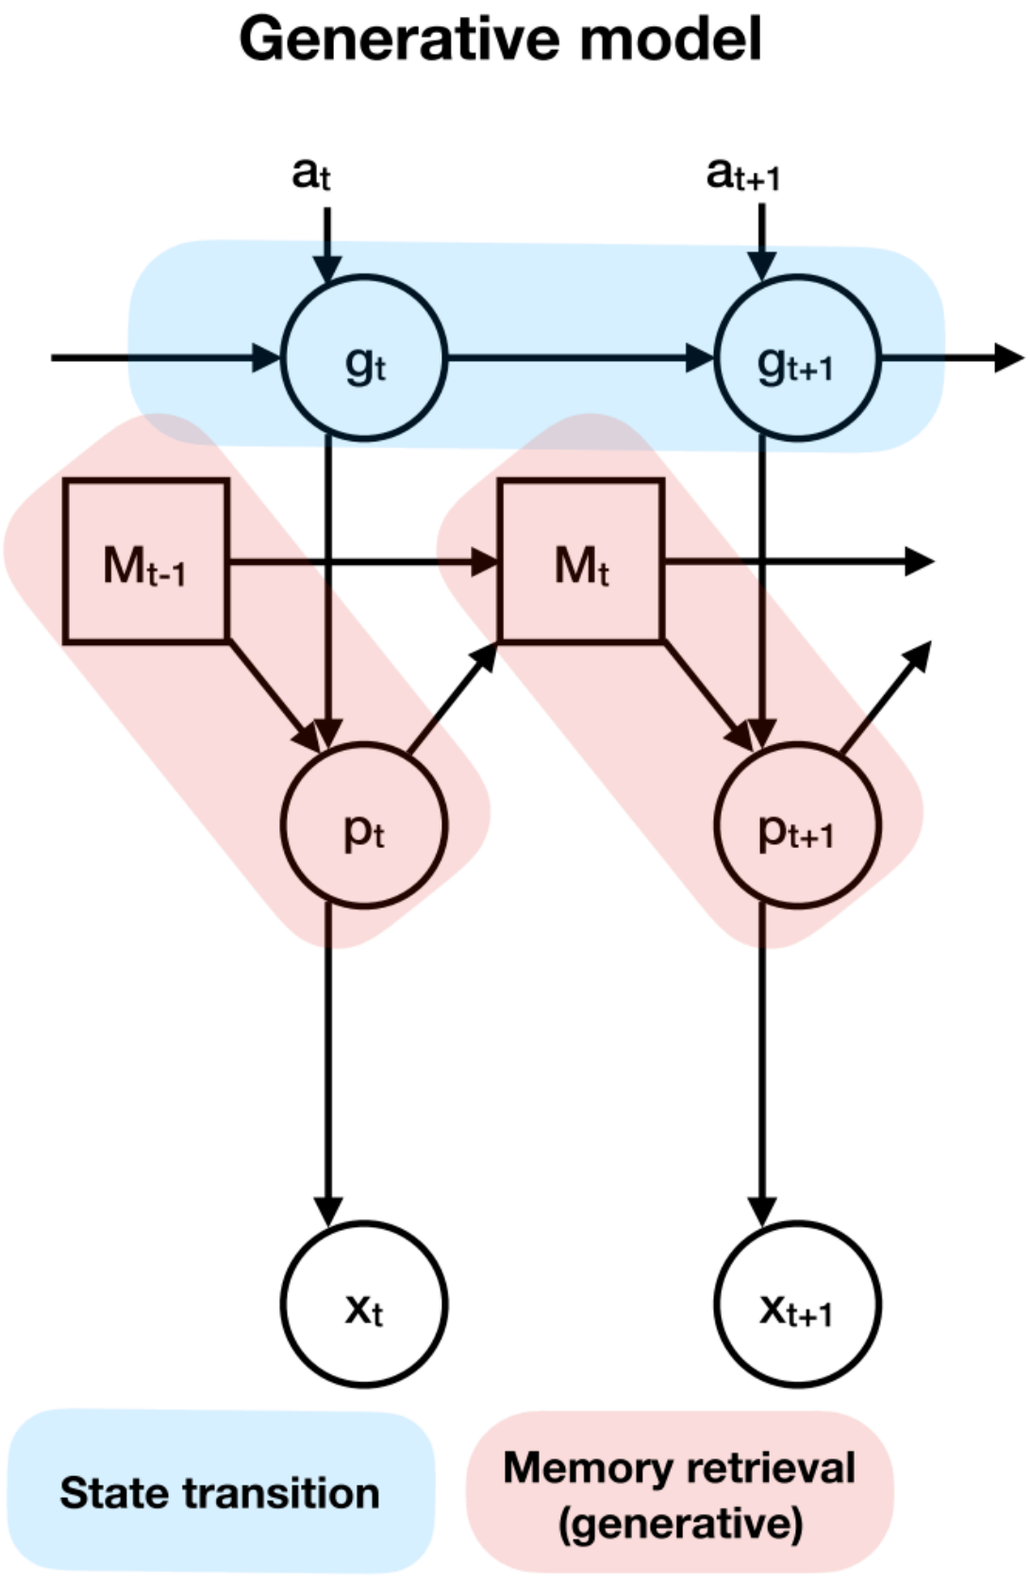

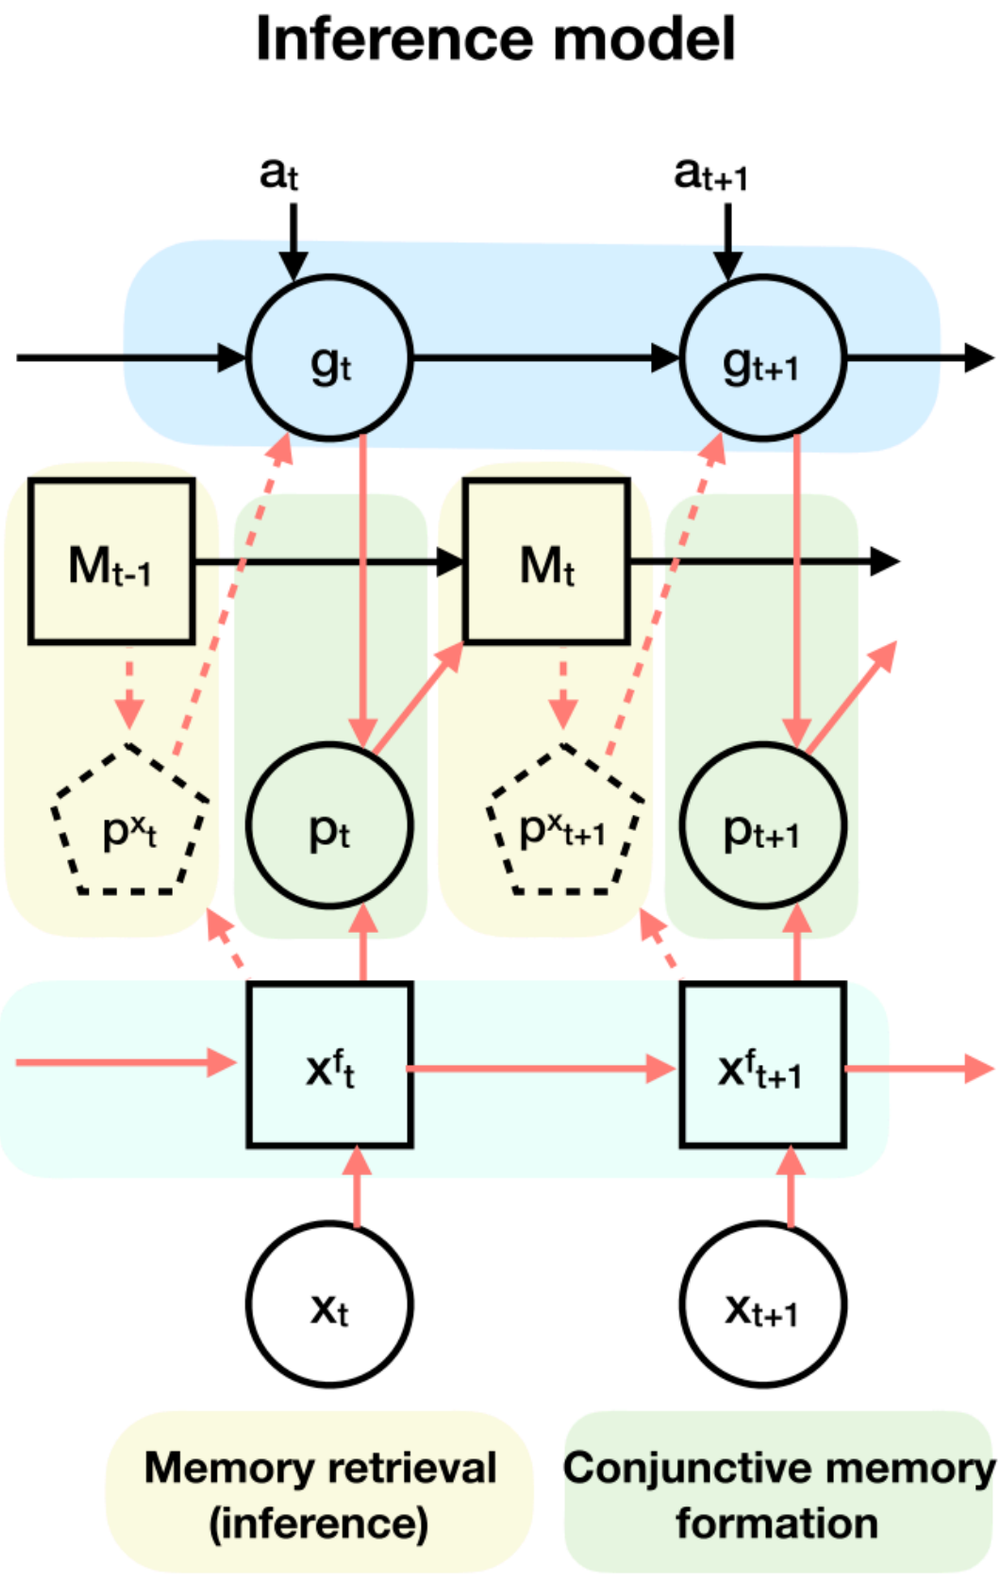

The TEM and the Brain

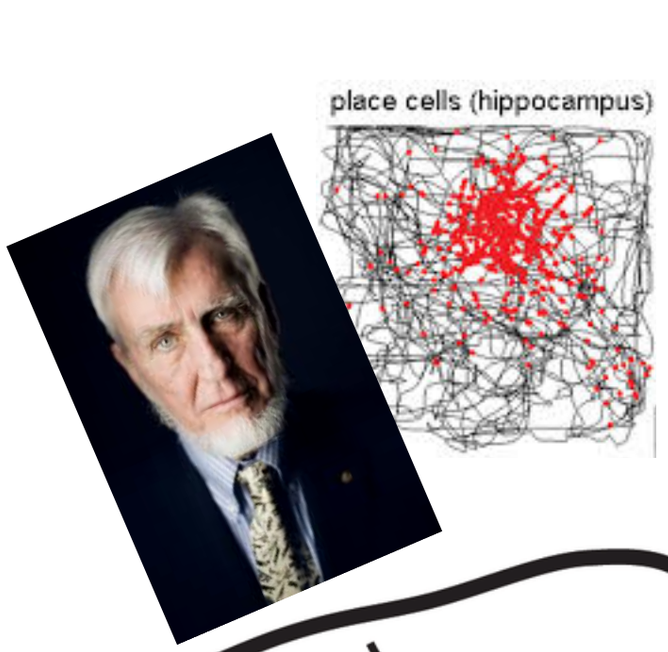

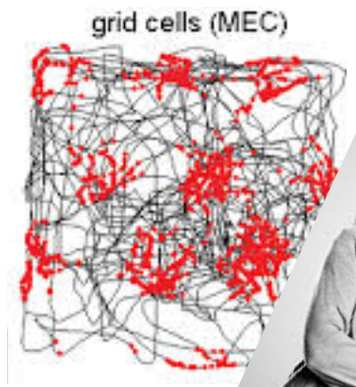

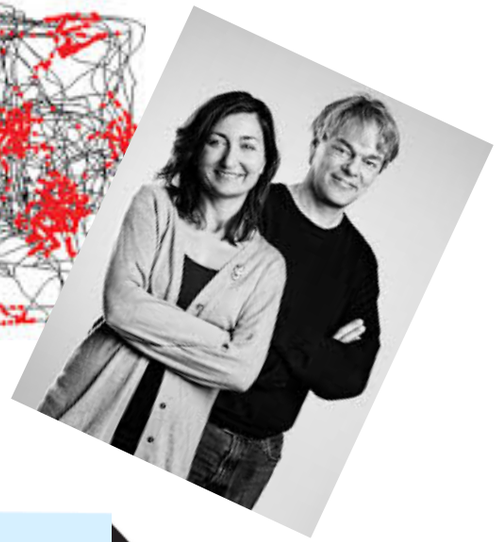

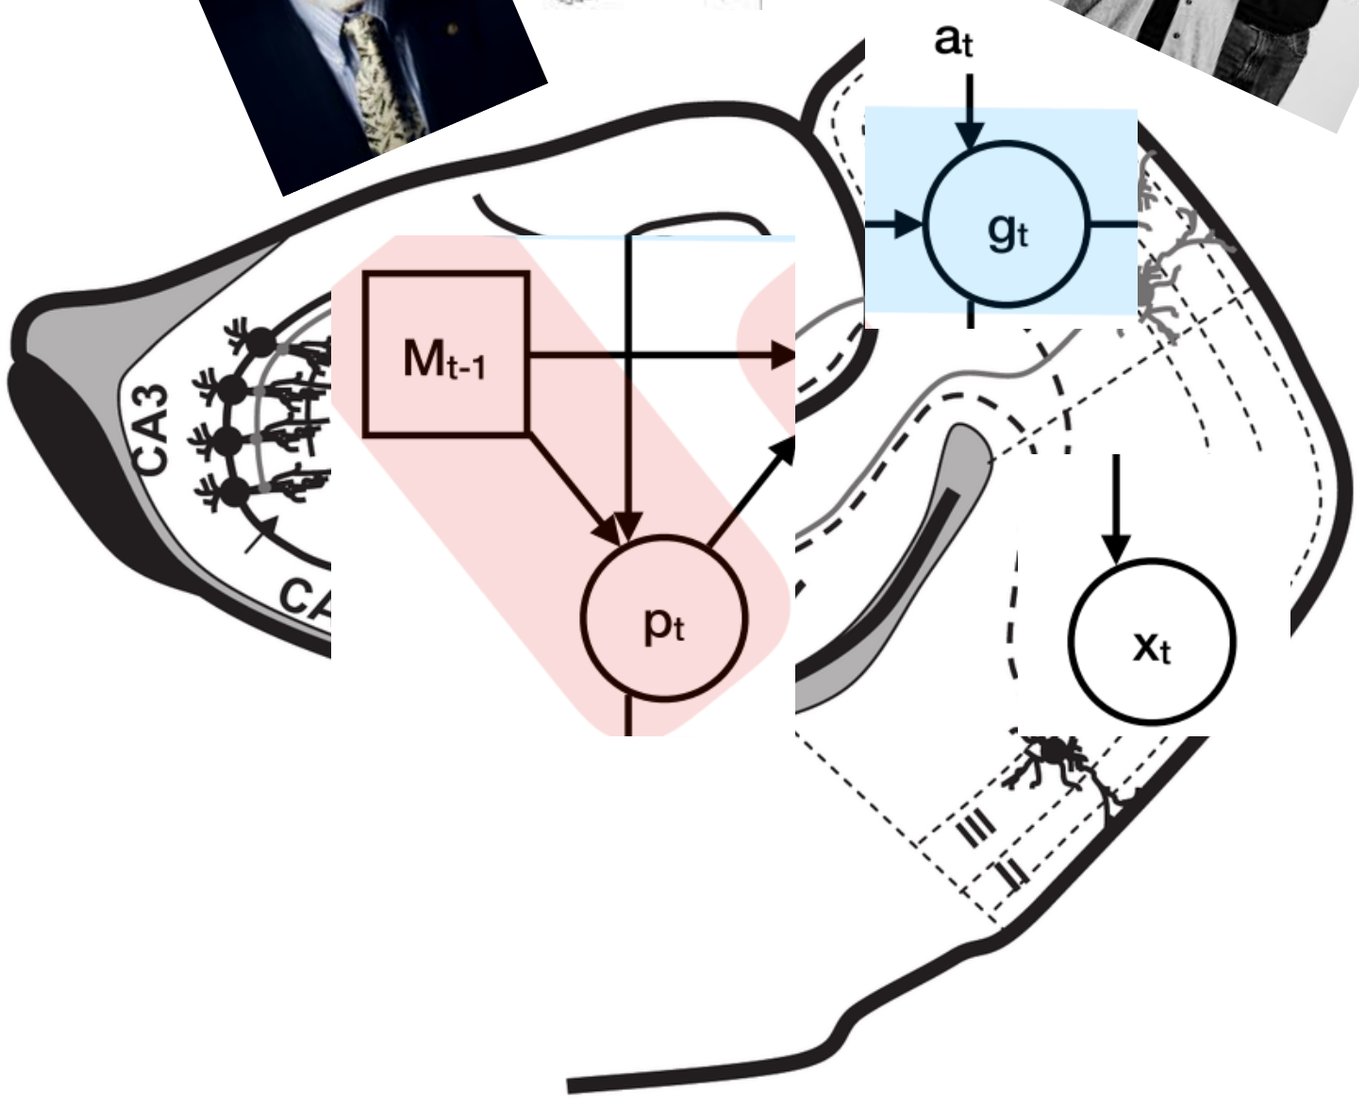

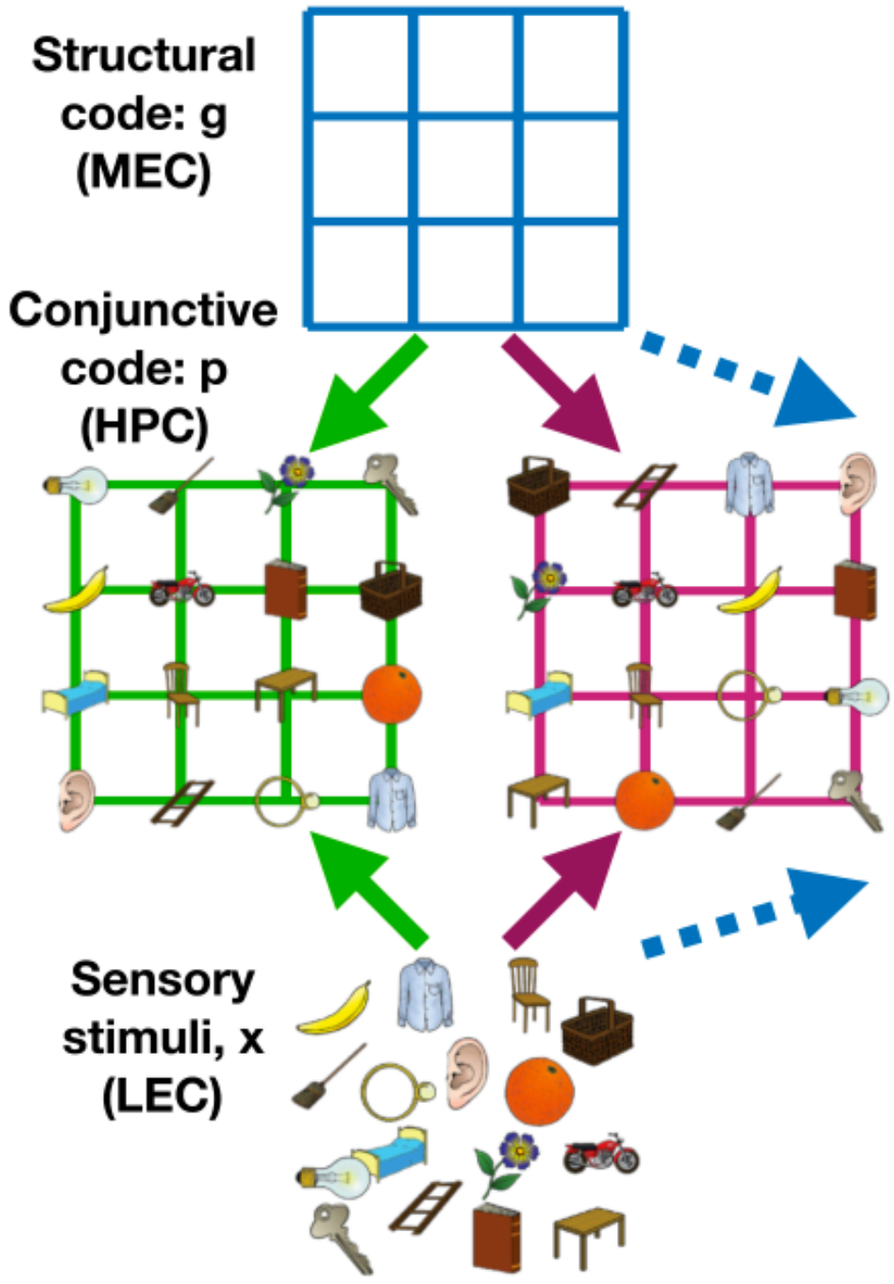

Random movement on

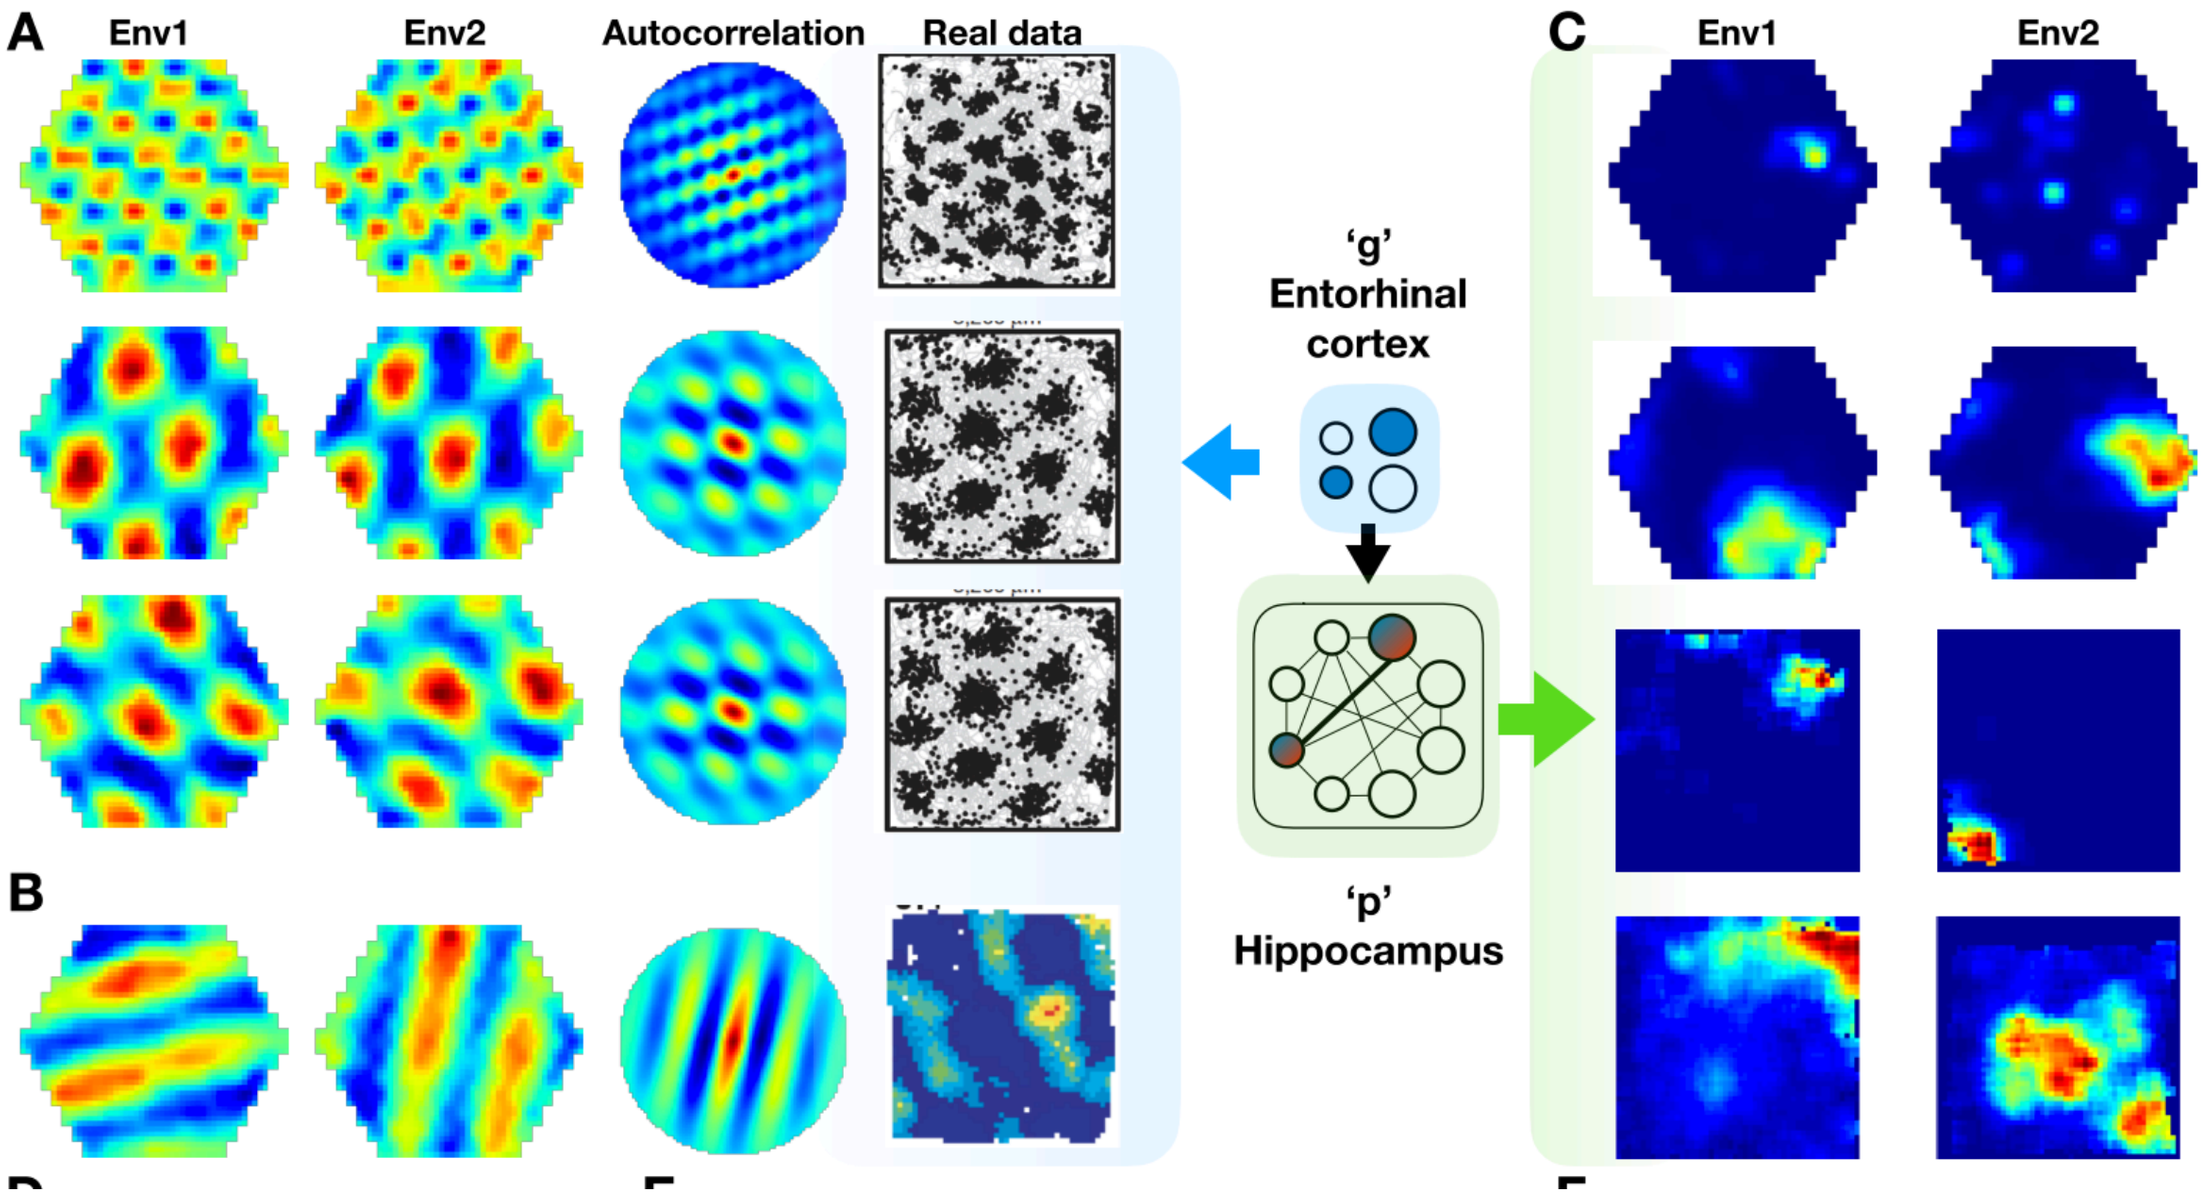

Prediction by parameter learning - A model of V1 receptive fields

- Let's try to model how V1 simple cells work
  - They respond to oriented edges
    - this is what the model should predict!
  - So the model of visual scenes (images of b&w pixels) on this level is that there are some edge-objects that translate to pixels (x) at different locations

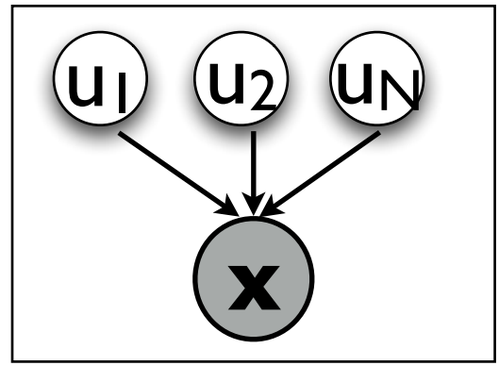

- Assumptions for a probabilistic model
  - there is a latent variable (u) for each possible edge
  - their prior distribution is sparse meaning that only a few of them will contribute to single scene
  - they are independent from each other (this is a strong simplification)
  - they mix linearly
  - pixels may deviate from the mix of images according to a Gaussian distribution (this means that there is some observation noise)

$$p(x \mid u) = \mathcal{N}(x; Au, \sigma_x I)$$

$$p(u) = Sparse(u; 0)$$

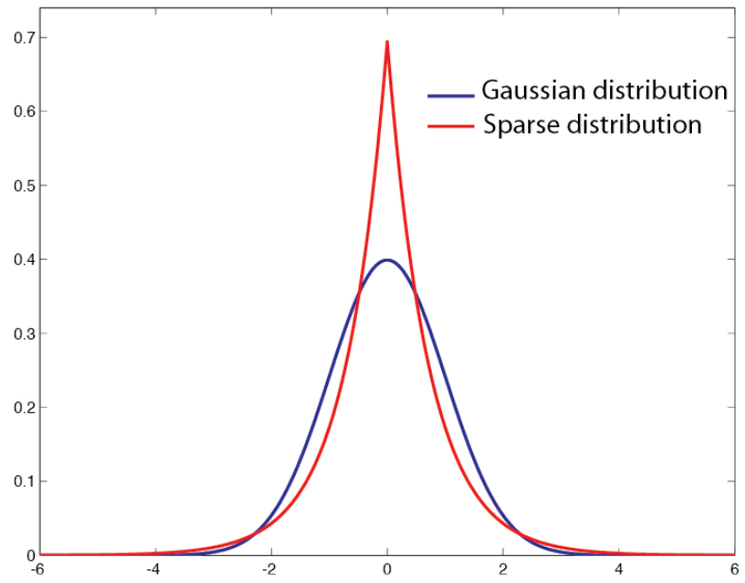

Learning edge filters from natural images

- a question is how to do the linear mixing of the latents so that they will follow the sparse distribution - that is, what should be the features (A) corresponding to latent variables
  - the algorithm that looks for linear filters for sparsely distributed, independent latent variables is Independent Component Analysis (ICA) - it is used in signal processing, more about similar things when we get to learning
  - If you apply such a procedure, you get features similar to Gabor filters used for edge detection in image processing (Olshausen & Field, 1996)
- The model predicts the shape of V1 simple cell receptive fields assuming sparsity and learning from natural images
  - Another prediction: the average activation (e.g. membrane potential or number of action potentials) of V1 simple cells in response to an image will be proportional with the latent values inferred by the model from the image

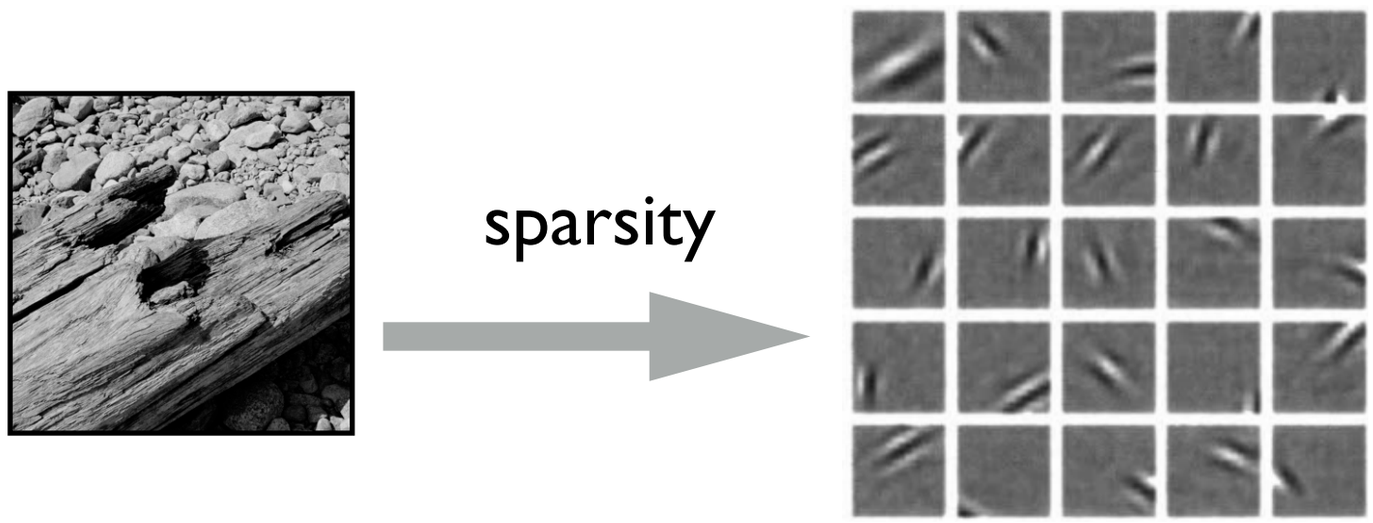

Olshausen & Field, 1996

The way forward

- We have seen that it is possible to produce falsifiable predictions regarding neural computation by using probabilistic models

- Some predictions of simple models of early vision hold up against measurements from behaving animals

- We have to build more complex models that can predict activity from higher-level visual processing areas, not just V1

    - we may use various state-of-the-art techniques from machine learning

- Step by step, we have to figure out what is the mental representation of the environment

What do we know about how the brain works?

- We know a lot about anatomy
  - but we still don't have the connectome - the blueprint of neural networks in the brain
  - local connectivity patterns within cortical regions are also only partially known
- We know a lot about dynamics
  - we can describe single neuron and network level electric behaviour patterns
  - but mostly without tying them to any function to prediction
- We know a lot about localisation
  - low-level perceptual functionality, motor areas, episodic memory, etc.
  - some hints about other functions
- We know a lot of receptive fields
  - if we looked for all the right quantities of the stimuli when we characterised them
  - for objects and concepts we have only hints
- We know a lot about how to solve problems the brain has to solve
  - specialised solutions exists, mostly to perceptual stuff, no general problem solvers
  - we have no idea how plausible these solutions are regarding biological implementation
- We know a little about the mental model and how to do inference in it