Gaps between the Observed Stochasticity of Neural Systems and Probabilistic Inference Models

Wolfgang Maass
Institut für Grundlagen der Informationsverarbeitung
Technische Universität Graz, Austria

Structure of my talk

I will address three gaps in our knowledge:

- Which algorithms for probabilistic inference could be implemented in neural systems, given their strong variability?
- How can neural systems **learn** to create internal models for salient probability distributions?
- How are functionally relevant random variables (RVs) represented by neurons?

Virtual all neural systems exhibit substantial trial-to-trial variability

Example: Variability of spike responses in area V1 of cat: each column shows 3 trials with the same stimulus

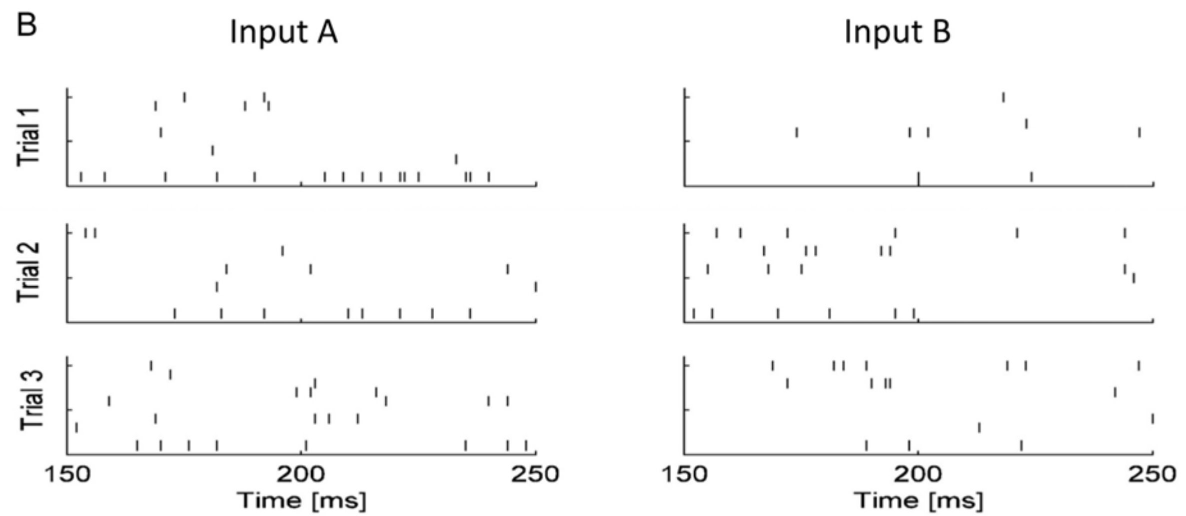

[Nikolic, Haeusler, Singer, Maass, PLoS Biol. 2009]

A major source of variability in neural circuits: probabilistic vesicle release at synapses

Common estimates of release probability of a vesicle in response to a presynaptic spike are around **0.5** (for neocortex), see e.g.

(Branco, Staras, Nat. Rev. in Neurosci, 2009)

In addition vesicles are frequently released without a presynaptic spike

(Kavalali, Nat. Rev. in Neurosci., 2015)

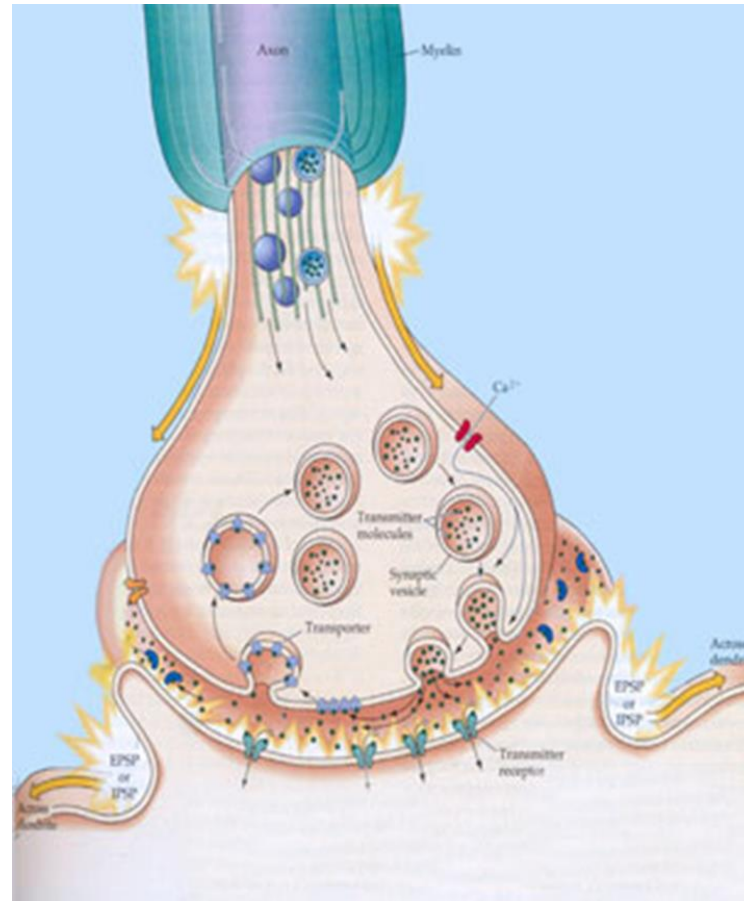

A note on concepts

- Trial-to-trial variability is observed in the brain at virtually all spatial and temporal scales
- It is not known to what extent this variability is due to deterministic dynamics of hidden variables
- From the perspective of the models for computation and learning that I will discuss, it does not matter whether this variability results from true stochasticity, as long as it can be described well enough by probabilistic rules

How deterministic can neural responses get ?

Bowers, Jeffrey S., and Colin J. Davis. "Bayesian just-so stories in psychology and neuroscience." Psychological Bulletin, 2012

tried to summarize **the most reliable neural responses that have been found:**

- *Gur, M., & Snodderly, D.M. (2006). High response reliability of neurons in primary visual cortex (V1) of alert, trained monkeys. Cerebral Cortex*

But: these are results for artificial stimuli (and only for selected neurons)

- *DeWeese, M. R., Wehr, M., & Zador, A. M. (2003). Binary spiking in auditory cortex. Journal of Neuroscience*

But: these reliable responses in anaest. rats were not duplicated in their subsequent study of awake rats:

Hromádka, T., DeWeese, M. R., & Zador, A. M. (2008). Sparse representation of sounds in the unanesthetized auditory cortex. PLoS Biol.

Which algorithms for probabilistic inference could be implemented in neural systems, given their strong variability?

- Belief propagation (message passing) needs to implement complex arithmetical calculations, of high arithmetical depth. How could a neural system make this calculation noise robust?

- Are there besides stochastic sampling approaches other algorithmic/neural implementation approaches for probabilistic inference that are compatiable with strong variability?

How can neural system *learn* to create internal models for salient probability distributions?

One step into this direction by (Pecevski et al., under review)

Assumption: Some external distribution p* generates examples **y.**

Goal: Learn an internal model of p*.

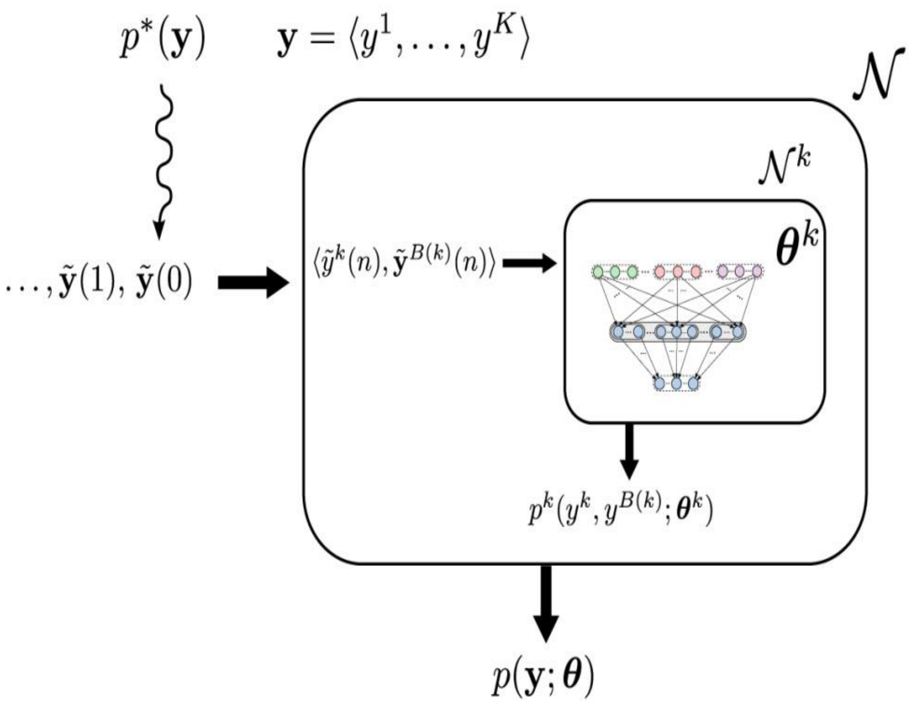

The network consists of 3-layer moduls, that each learn the probability table for one RV, conditioned on the RVs in its Markov blanket:

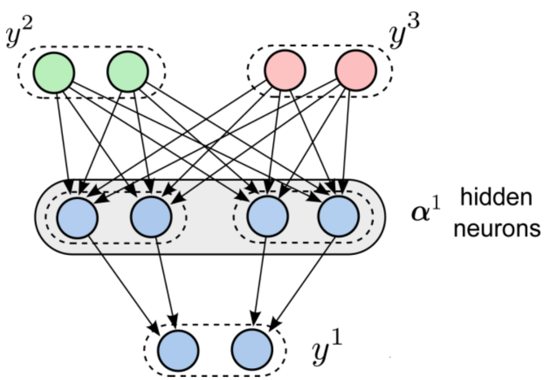

Learning takes place through STDP on synapses to hidden layer neurons, that are split into several WTA circuits.

A closer look at plasticity in the network modules

Hidden neurons are needed, because the distribution of value assignments to the Markov blanket **x** that causes e.g. z=2 may be a mixture of multinomials.

Each such mixture of multinomials for a given value of z is learnt through STDP by one WTA circuit of hidden neurons

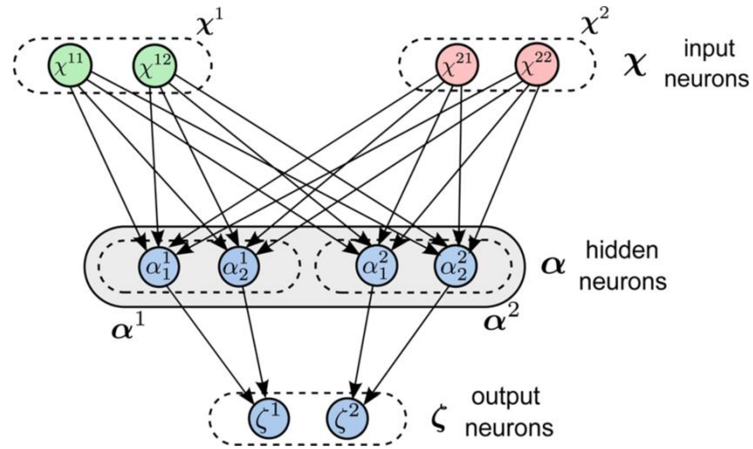

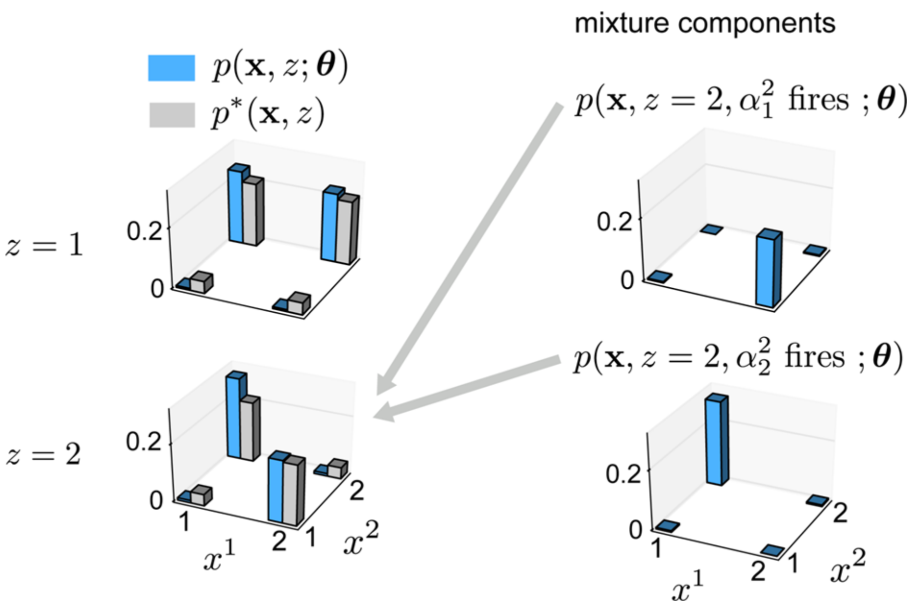

Ideally each WTA circuit should have as many neurons as there are multinomials; hence in this case 4 neurons instead of 2 ; but 4 neurons do not improve performance:

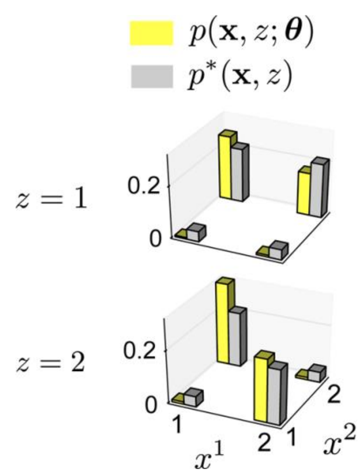

Application to a standard *(but nontrivial)* **probabilistic inference task: Explaining away in visual perception (***Knill, Kersten, Nature 1991)*

The observed contour influences our perception of relative reflectance:

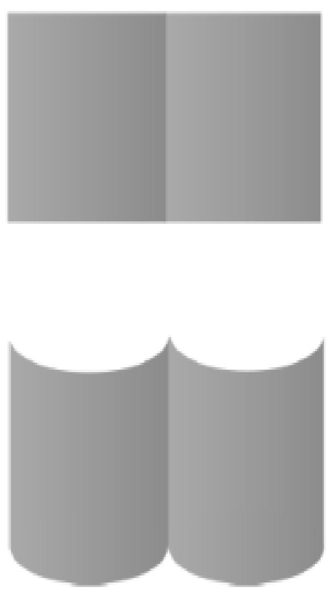

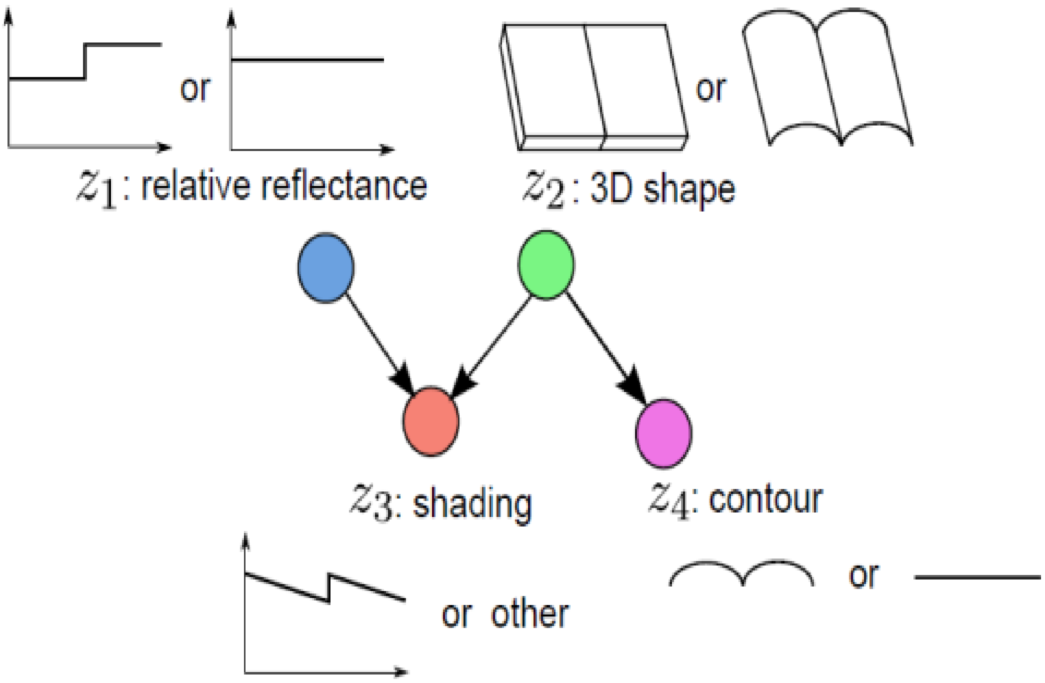

When the curved contour is occluded (z_4 changes from 1 to 0), we suddenly perceive **different** relative reflectance of the 2 parts (z_1 changes from 0 to 1).

A rigorous learning theory (*EM= Expectation Maximization*) implies that the Kullback-Leibler divergence between ideal and learnt conditional probabilities is locally optimized through STDP (this *requires* a stochastic neural network)

The learning network consists for the example of Knill/Kersten of 4 learning modules:

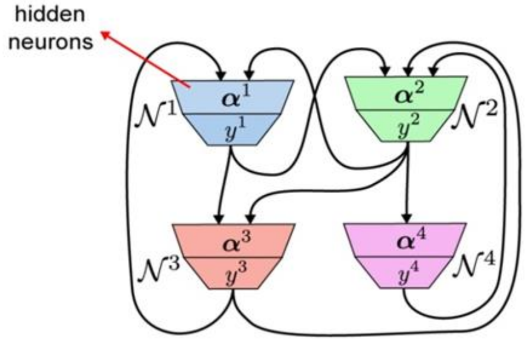

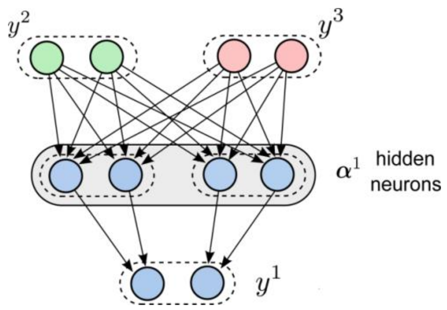

Time courses of the Kullback-Leibler divergences for the 4 learning modules:

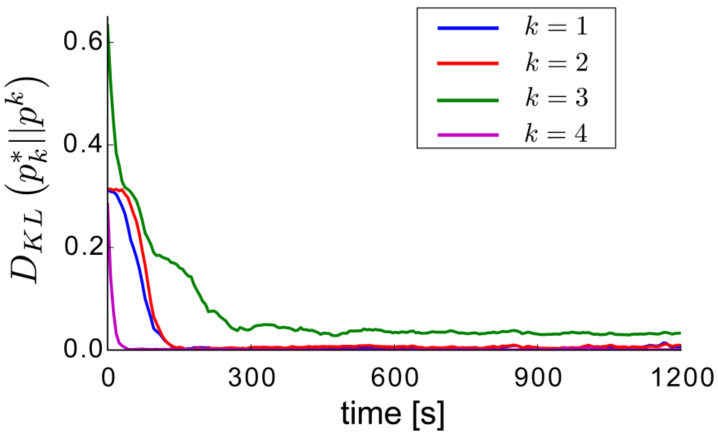

Functional consequences of this learning

Spontaneous firing activity after learning approximates the distribution p* that generated the examples

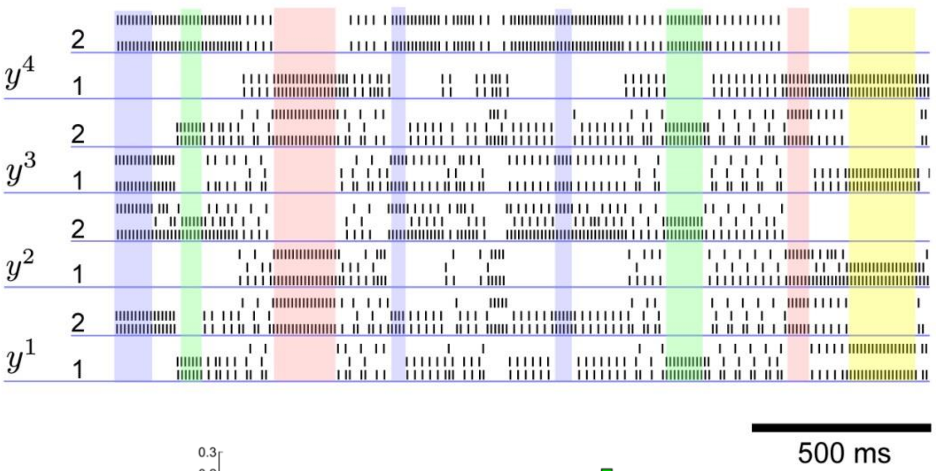

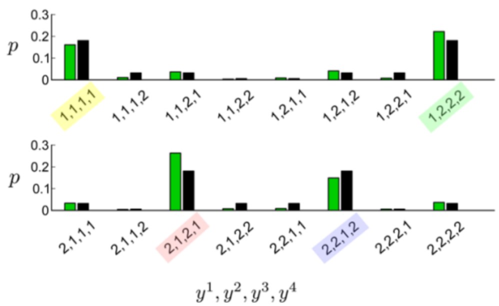

State probabilities of learnt joint distribution (green) and of target distribution p* (black)

Learnt „explaining away“: Observed contour changes to „straight“ after 3s:

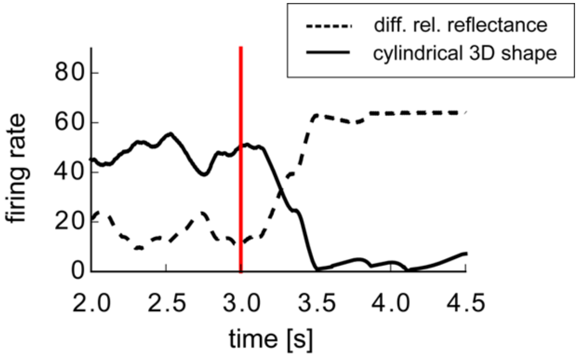

Open problems regarding learning of probability distributions, and probabilistic inference

- How can the learning be sped-up?
- How can transfer of information from previously learnt distributions be implemented?
- Which other approaches for neural implementations of probabilistic inference are amenable to learning from examples?

  (see work of Sophie Deneve et al).

Addressing a 3rd gap in our understanding:
How are functionally relevant random variables represented by neurons?

Most neural sampling models let relevant RVs be encoded by single neurons. This is problematic for several reasons, e.g.

- lack of robustness against disfunction of single neurons
- a single neuron cannot cause other neurons to fire with high probability

In addition, a number of experimental data suggest that behaviourally relevant random variables are encoded by **assemblies** of neurons:

Data from premotor cortex in monkey suggest that many neurons together encode a RV, each voting for a different value of it

(Cisek, Kalaska, Neuron 2005)

Experimental setup: Two subsequent cues provide information which saccade goal gets rewarded:

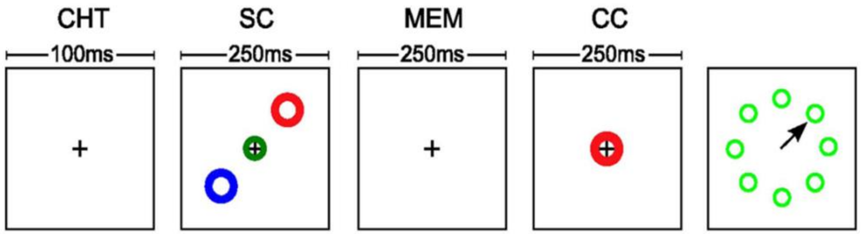

Recorded neural activity in premotor cortex (neurons ordered according to preferred saccade goal):

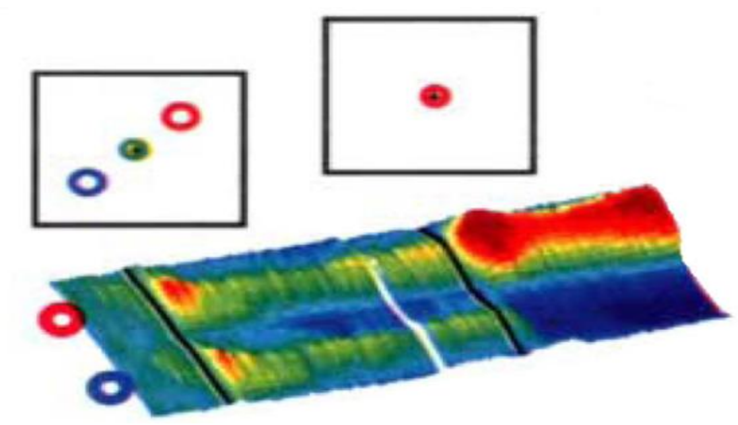

These data suggest that neural circuits encode a RV through spatial coding, like a particle filter (Bayesian filter)

It was recently shown that networks of stochastically spiking neurons can in fact approximate basically any particle filter. In particular, the data of (Cisek and Kalaska, 2005) can be reproduced in such model:

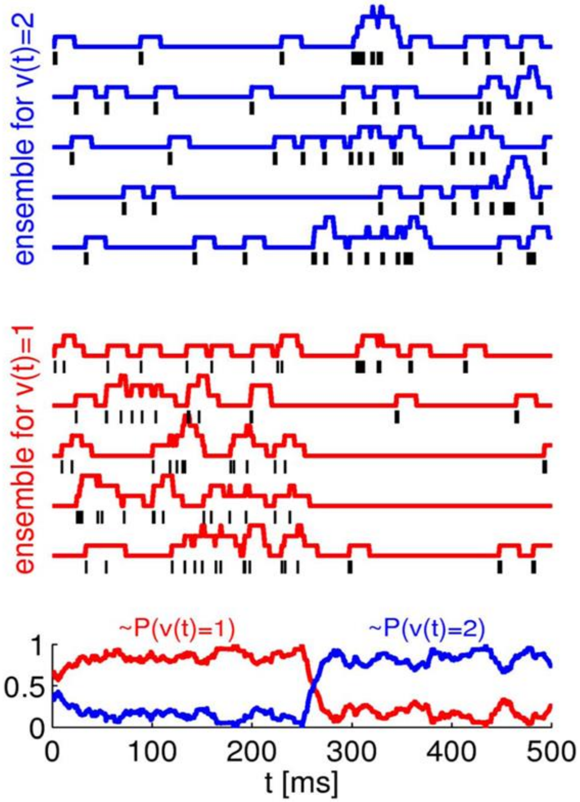

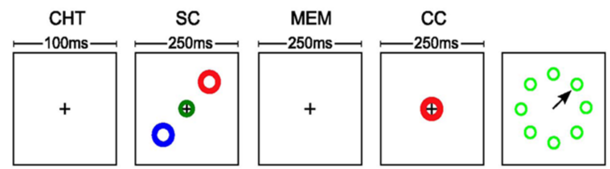

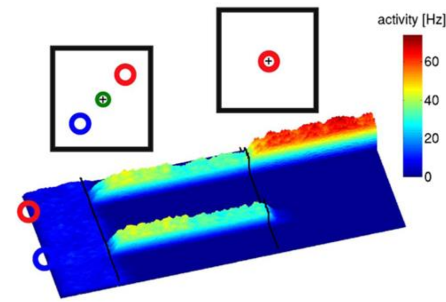

General functional advantages of this approach:
--more stable and fast representation of estimates
--applicable to time-varying probabilities.

(Legenstein, Maass, PLOS Comp. Biol. 2014),
(Savin, Deneve, NIPS 2014)

In this neural code also confidence of estimates of probabilities can be expressed

This issue was raised yesterday in the talk of Florent Meyniel (and in his paper in PLOS Comp. Biology).

If each neuron votes for a particular value of a RV, the variance of votes encodes the confidence of the estimate.

Another hint from experimental data: Multi-unit recording show stereotypical firing patterns of many neurons, both spontaneously and stimulus-evoked

Assemblies of neurons in rodent area A1 (awake); (Harris et al., 2008):

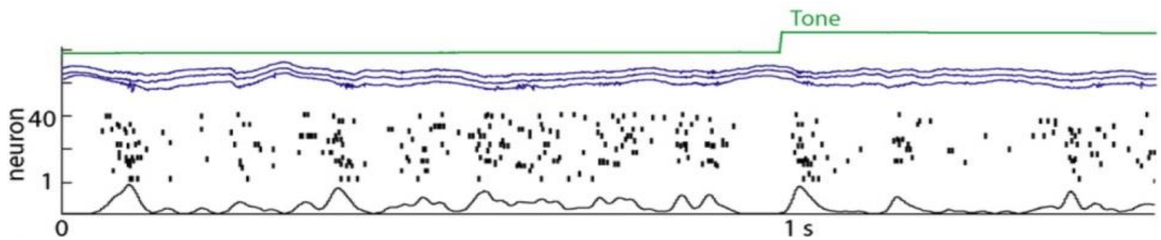

Similar firing patterns emerge spontaneously and stimulus-induced ([Luczak, Bartho, Harris, Neuron 2009].

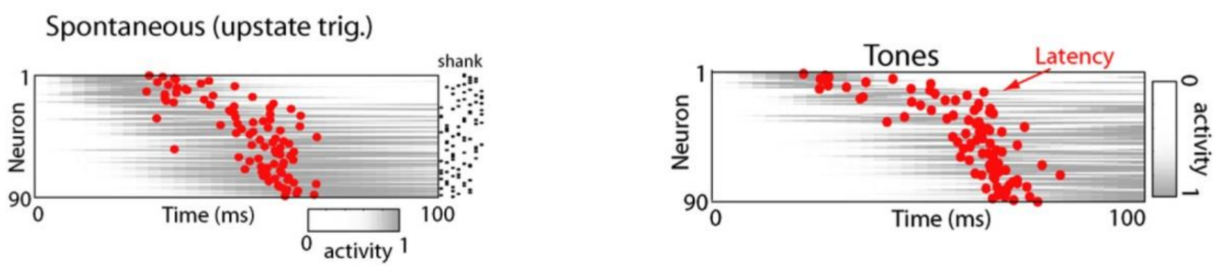

Data from Ca-imaging in rodent area V1 (Miller, Ayzenshtat, Carillo-Reid, Yuste, PNAS 2014):

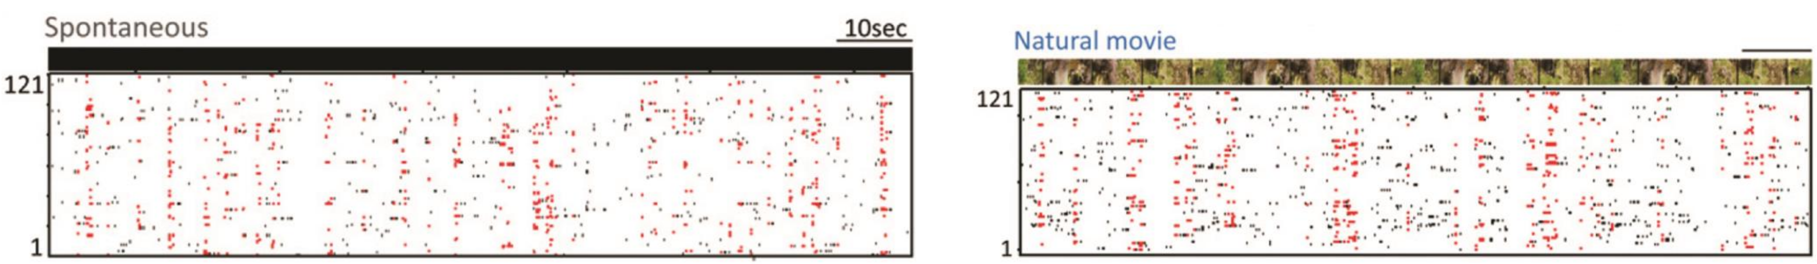

These data suggest a different role of WTA-like circuit motifs in building neural representations

Roughly: Interacting ensembles of excitatory and inhibitory neurons behave more like adaptive k-WTA circuits, rather than WTA circuits.

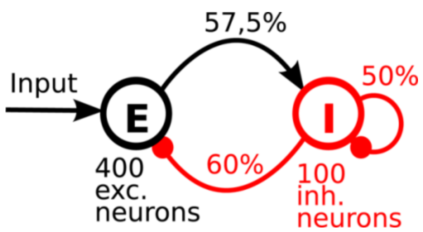

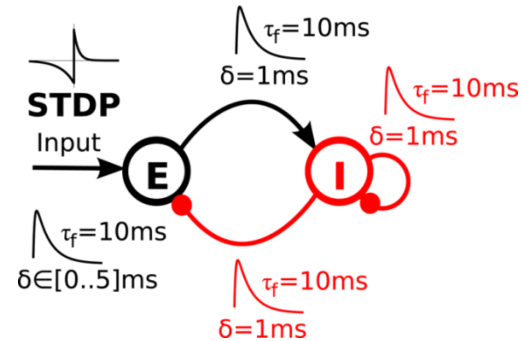

Instead of a mixture model as generative model (as suggested by WTA-circuits) STDP fits in these more realisitc models a product distribution (Noisy OR) to input spike streams:

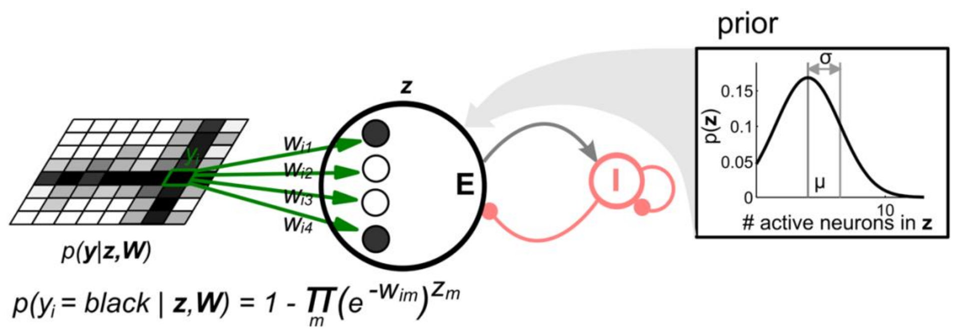

(Zonke, Legenstein, Maass, in preparation)

Resulting new coding properties that emerge under STDP:

*Neural codes emerge for pattern **components***

If high-D spike inputs contain repeatedly occurring components, each component gets represented by an assemblies of neurons in the EI motif:

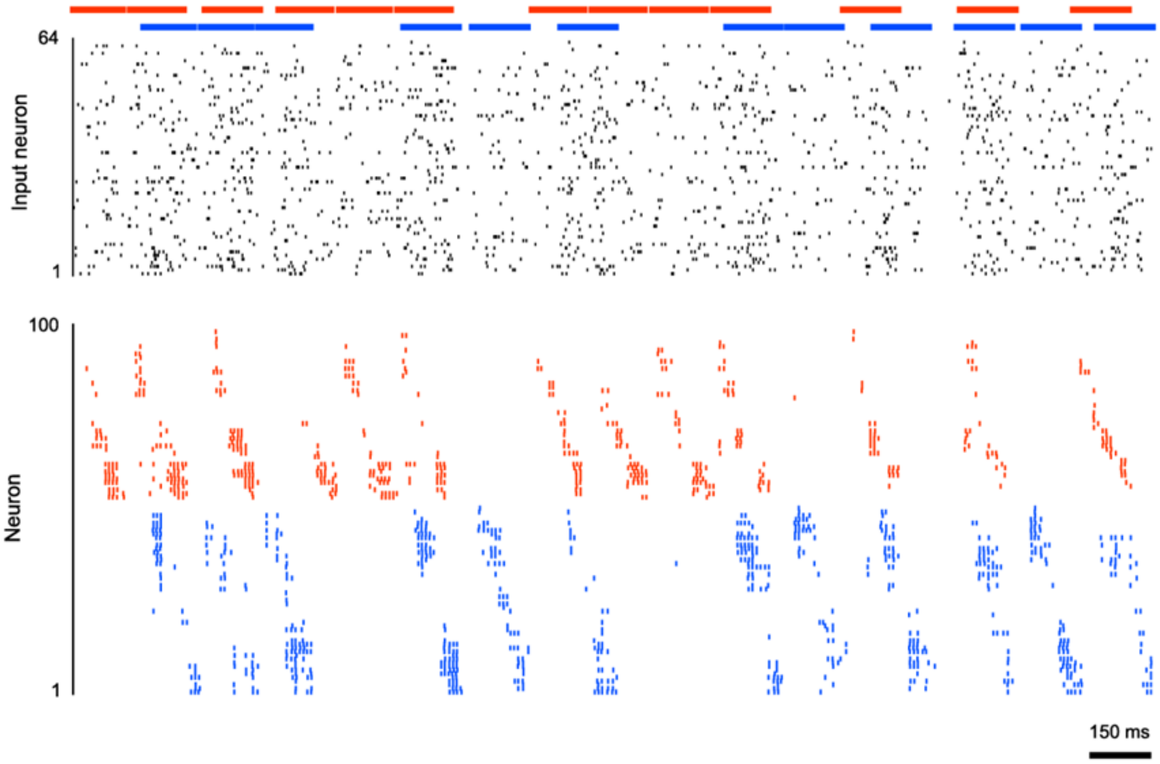

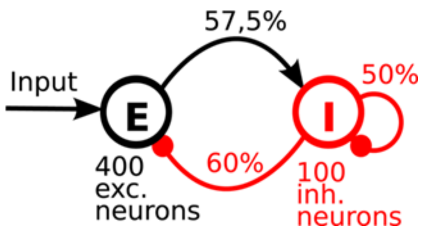

Open problems regarding neural representations of salient RVs for probabilistic inference

- How can these emergent loose assembly codes be used for probabilistic inference through neural sampling?

- Can such soft WTA-circuits serve as elementary modules for larger networks that learn complex Bayesian networks (distributions) from examples ?

Collaborators whose work I discussed

Dejan Pecevski
(UBIMED, Wien)

Zeno Jonke

Stefan Habenschuss
(GLN Graz)

Lars Büsing
(Columbia Univ)

Robert Legenstein

Summary

I have addressed three gaps in models for probabilistic inference through neural sampling:

- Which algorithms for probabilistic inference could be implemented in neural systems, given their strong variability?

- How can neural system learn to create internal models for salient probability distributions?

--Forthcoming paper (Pecevski, Maass, 2015) provides a first model based on STDP in simple WTA-motifs.

--Further work is needed to close the gap to more detailed models for neural circuits..

- How are functionally relevant random variables represented by neurons?

--Biological data show that loose assemblies (rather than single neurons) emerge as tokens of stochastic network dynamics.

--We need to understand whether and how they could support neural sampling.

--Such loose assemblies emerge through STDP in ensembles of pyramidal cells with data-based lateral inhibition.