‘Bayesian’ theories of perception, cognition and mental illness (part 4 - CCN Lecture 16)

Peggy Seriès,
IANC, University of Edinburgh

Behavioural studies: So what have we learned?

- Bayesian models offer parsimonious description of behaviour (descriptive tool)

- Transparent assumptions and emphasis on “why” question.
- Behaviour consistent with Bayesian hypothesis in that:
  - Brains take into account uncertainty, and combine sources of information combines information optimally (cue combination)
  - Use priors that are constantly updated
  - Those priors are consistent with (some approximation) of statistics of environment at different time scales. --> increase accuracy.

- Deviations from optimality are possibly informative about underlying biological constraints, or nature of approximations.
- Priors (+cost functions, likelihood) can be measured in individuals -- Bayesian modelling as a tool to describe the internal model used by individuals, possibly differentiating groups.

What does this tell us about the Brain ?

Will this change our understanding of neurobiology?

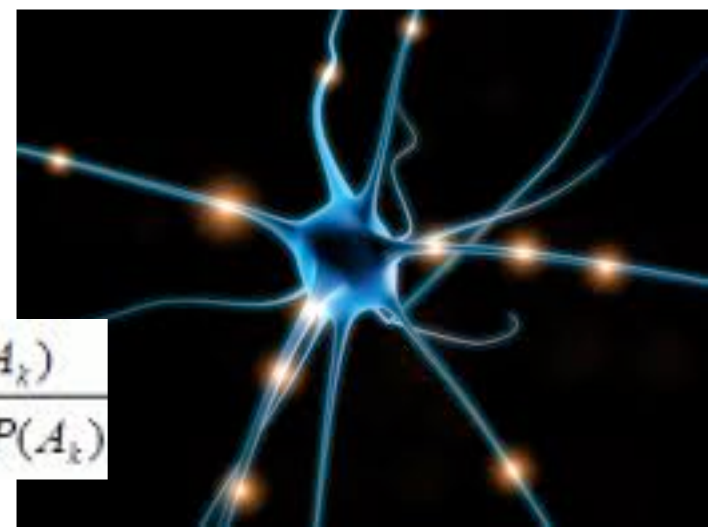

$$P(A_k \mid B) = \frac{P(B \mid A_k)P(A_k)}{\sum_{k=1}^{n} P(B \mid A_k)P(A_k)}$$

Is the Brain “Bayesian”? Debates

Psychological Bulletin
2012, Vol. 138, No. 3, 389–414

0033-2909/12/$12.00 DOI: 10.1037/a002645

Bayesian Just-So Stories in Psychology and Neuroscience

Jeffrey S. Bowers
University of Bristol

Colin J. Davis
Royal Holloway University of London

According to Bayesian theories in psychology and neuroscience, minds and brains are (near) optimal in solving a wide range of tasks. We challenge this view and argue that more traditional, non-Bayesian approaches are more promising. We make 3 main arguments. First, we show that the empirical evidence for Bayesian theories in psychology is weak. This weakness relates to the many arbitrary ways that priors, likelihoods, and utility functions can be altered in order to account for the data that are obtained, making the models unfalsifiable. It further relates to the fact that Bayesian theories are rarely better at predicting data compared with alternative (and simpler) non-Bayesian theories. Second, we show that the empirical evidence for Bayesian theories in neuroscience is weaker still. There are impressive mathematical analyses showing how populations of neurons could compute in a Bayesian manner but little or no evidence that they do. Third, we challenge the general scientific approach that characterizes Bayesian theorizing in cognitive science. A common premise is that theories in psychology should largely be constrained by a rational analysis of what the mind ought to do. We question this claim and argue that many of the important constraints come from biological, evolutionary, and processing (algorithmic) considerations that have no adaptive relevance to the problem per se. In our view, these factors have contributed to the development of many Bayesian “just so” stories in psychology and neuroscience; that is, mathematical analyses of cognition that can be used to explain almost any behavior as optimal.

Keywords: Bayes, Bayesian, optimal, heuristics, just-so stories

Is the Brain “Bayesian”? Debates

Psychological Bulletin
2012, Vol. 138, No. 3, 389–414

0033-2909/12/$12.00 DOI: 10.1037/a00264

Bayesian Just-So Stories in Psychology and Neuroscience

Jeffrey S. Bowers
University of Bristol

“Our main thesis is that Bayesian modeling, both in practice and in principle, is a misguided approach to studying the mind and brain”

Bowers & Davis, 2012.

Keywords: Bayes, Bayesian, optimal, heuristics, just-so stories

A Bit of Philosophy

- Marr's levels of analysis: computational / algorithmic / implementation.

Levels function mostly independently.

- *"Bayesian models are not intended to provide mechanistic or process accounts of cognition"*

[Jacobs and Kruschke, 2010]

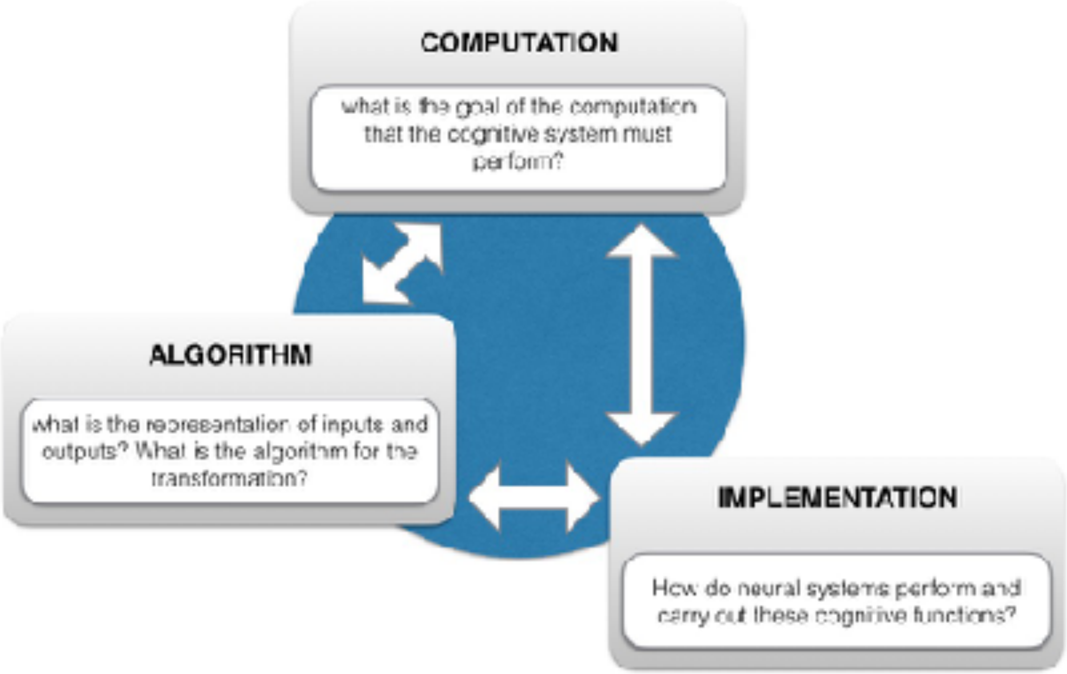

- only an approximation of Bayesian inference anyway.

•Bowers and Davis, 2012; O'Reilly et al., 2012

Debates: Criticism

- Confusion about optimality
- Falsifiability: Bayesian models are flexible enough to account for everything
- Rarely compared with alternative (non-Bayesian) hypotheses
- Integration with previous research knowledge (just a new vocabulary?)
- Lack of neurobiological predictions / evidence

Debates: Some Answers

- Optimality: claim is not that the system is optimally designed, but that given a potentially bad design, the combination of noisy inputs is optimal.

- Bayesian approach: a framework = typically not falsifiable only individual models are falsifiable.

- Rarely compared with alternative hypotheses: should be compared with hypotheses formulated at same level (computational).

- Not incompatible with mechanistic models, not even based on simple heuristics.

  "There need to be nothing intrinsically Bayesian about algorithms that approximate Bayesian inference"

  Griffith, Norris, Chater, Pouget (2012)

Neural implementation of Bayesian inference ?

1. How do populations of neurons represent uncertainty ? Does neural activity represent probabilities? (log probabilities?)

2. How can a prior be implemented? (baseline - spontaneous activity, number of neurons, gain, connectivity?). Can we distinguish stages where the likelihoods, priors, posterior could be 'measured' experimentally ?

3. Can networks of neurons implement optimal inference? How?

Recently, active topic of theoretical research (e.g. A. Pouget, S. Denève, P. Dayan, R. Rao, J. Fiser, M. Lengyel, W.J. Ma).

1) How could neurons represent probability distributions?

Ideas (explicit representations):

• neural activity of a given neuron with preferred stimulus s represents the probability that feature s is present

• or log probability

• or log probability that a feature takes on a particular value.

• Probabilities are functions: neural activity could represent the parameters of that function, possibly parameter in basis function parametrisation. Idea defended by the proponent of Probabilistic Population Codes (Pouget, Latham, Wei-Ji Ma, etc…)

1) How could neurons represent probability distributions?

• very few plausible computational models proposed for a neural implementation of probabilistic learning that would provide easily testable predictions. Offer proof of principle.

• 2 categories :

1.1) Probabilistic Population Codes (Pouget, Latham, Deneve, ..) Neural activities represent parameters of the probability distribution. A full probability distribution is represented (**implicitly**) at any moment in time.

1.1 Probabilistic population codes (PPC):

spiking rates could represent the coefficients of a basis function parametrisation of the log probability

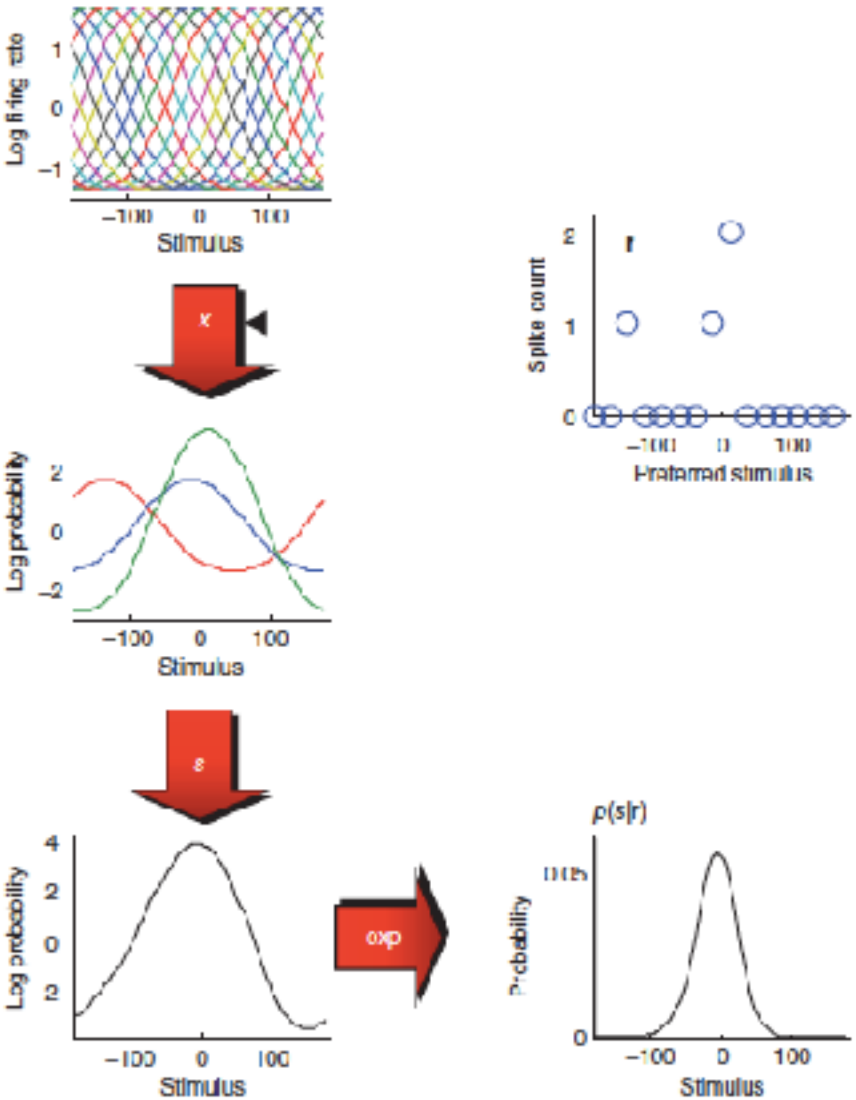

$$\log p(s\,|\,\mathbf{r}) = \sum_i r_i h_i(s) + \text{constant}$$

where $h_i(s)$ are the basis functions, e.g. log of tuning curves

Pouget et al 2013, Probabilistic brains: known and unknown

1.1 Optimal cue integration with PPC

A simple linear combination of the population patterns of activity guarantees optimal integration if neural variability is Poisson-like.

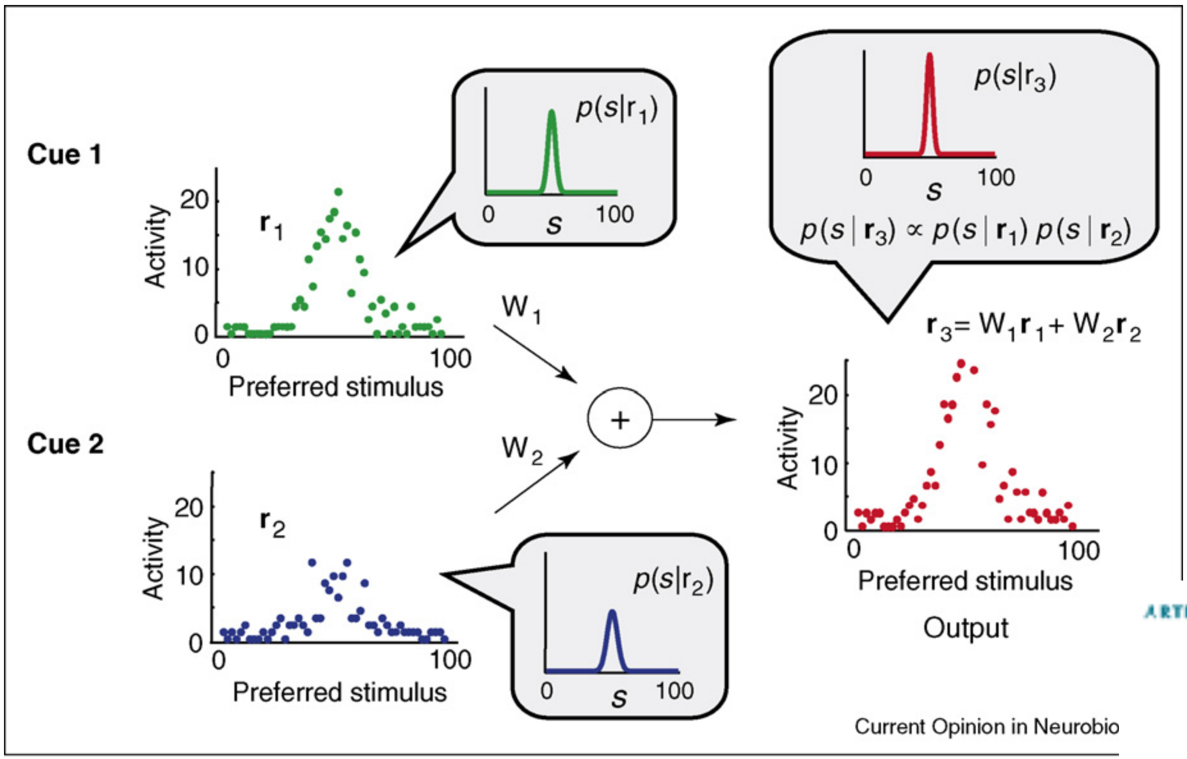


ARTICLES

nature neuroscience

Bayesian inference with probabilistic population codes

[Ma, Beck, Latham & Pouget, Nat Neuro 2006]

1) How could neurons represent probability distributions?

- very few plausible computational models proposed for a neural implementation of probabilistic learning that would provide easily testable predictions.

- 2 categories :

1.1) Probabilistic Population Codes (Pouget, Latham, Deneve, ..) Neural activities represent parameters of the probability distribution. A full probability distribution is represented (implicitly) at any moment in time.

1.2) Sampling Hypothesis (Fiser, Lengyel, ..): Neural activities represent the latent variables themselves, temporal variability represents uncertainty.

1.2 Sampling Hypothesis: Experimental Evidence

- What makes certain stimuli bistable ? (Necker Cube, Binocular Rivalry)
- Reflecting the fact that the posterior is bimodal?
- Hypothesis : the visual system draws a sequence of samples from the posterior over scene interpretations
- Gershman, Vul, Tenenbaum *NIPS* 2009

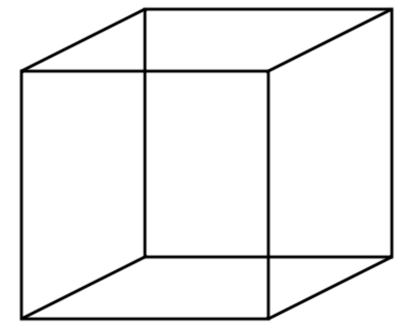

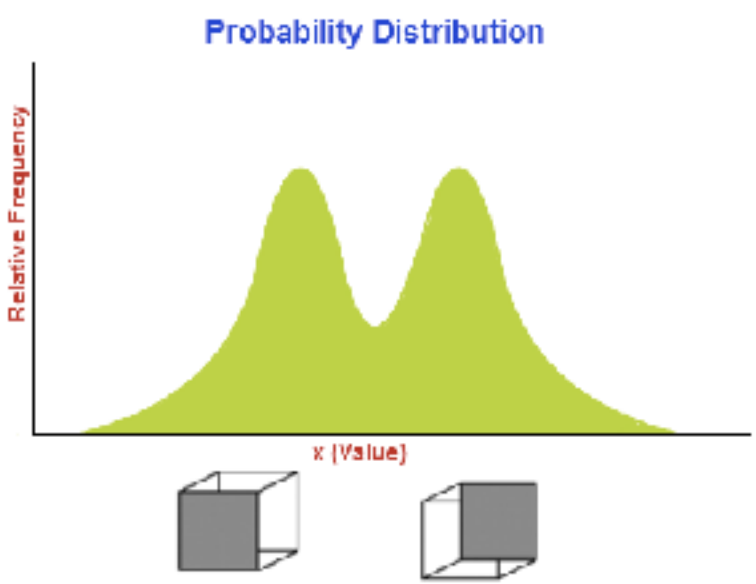

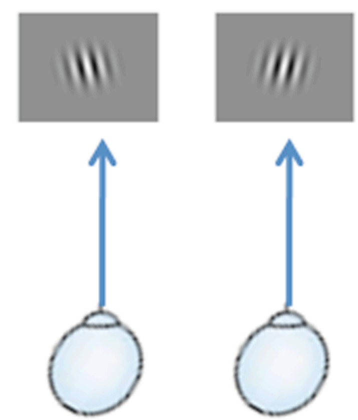

2) How could priors be implemented in the brain ?

- Priors: Where in the brain ?
- Top down inputs (predictive coding)
- Increase or decrease of activity ? [e.g. Summerfield & Egner 2009]
- in Tuning of neurons? [Gershick et al 2011; Fischer & Pena 2011]
- in Baseline activity? [Berkes et al 2010]
- The representation or the read-out?

- different time scale // different mechanisms

Can the effect of prior expectations be observed in fMRI ? (1)

The Journal of Neuroscience, October 9, 2013 • 33(41):15275–15284 • 15275

Behavioral/Cognitive

Prior Expectations Bias Sensory Representations in Visual Cortex

Peter Kok,[1] Gijs Joost Brouwer,[2] Marcel A.J. van Gerven,[1] and Floris P. de Lange[1]
[1]Radboud University Nijmegen, Donders Institute for Brain, Cognition and Behaviour, 6500 HE Nijmegen, Netherlands
Department of Psychology and Center for Neural Science, New York, New York 10003

Perception is strongly influenced by expectations. Accordingly, perception has sometimes been cast as
sensory inputs are combined with prior knowledge. However, despite a wealth of behavioral literature s
tion as probabilistic inference, the neural mechanisms underlying this process remain largely unkn
whether top-down expectation biases stimulus representations in early sensory cortex, i.e., whether the
and bottom-up inputs is already observable at the earliest levels of sensory processing. Alternatively,
unaffected by top-down expectations, and integration of prior knowledge and bottom-up input may take
areas that are proposed to be involved in perceptual decision-making. Here, we implicitly manipulated
tions about visual motion stimuli, and probed the effects on both perception and sensory representation
measured neural activity noninvasively using functional magnetic resonance imaging, and applied a
reconstruct the motion direction of the perceived stimuli from the signal in visual cortex. Our results s
bias representations in visual cortex, demonstrating that the integration of prior information and senso
stages of sensory processing.

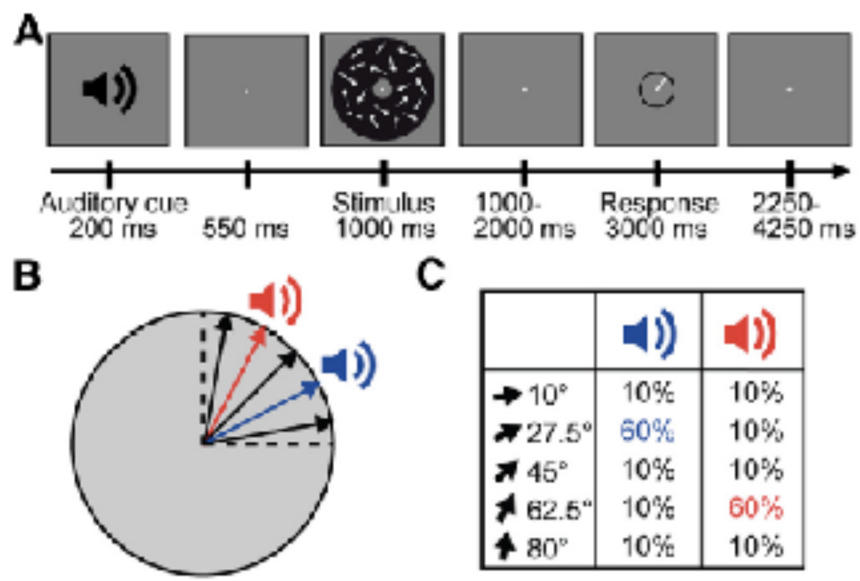

Can the effect of prior expectations be observed in fMRI ? (2)

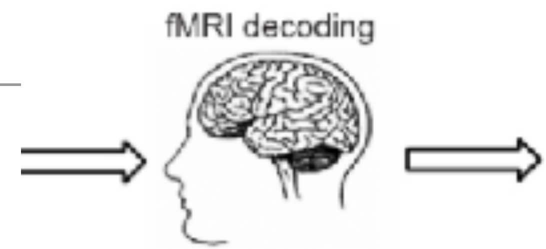

- Decoding from visual cortex : Does activity in visual cortex (V1, V2, V3, V4, MT) correspond to real stimulus or percept ? A: percept.
- Integration of prior expectations and sensory information in population activity is observed at the level of BOLD signals as early as in V1

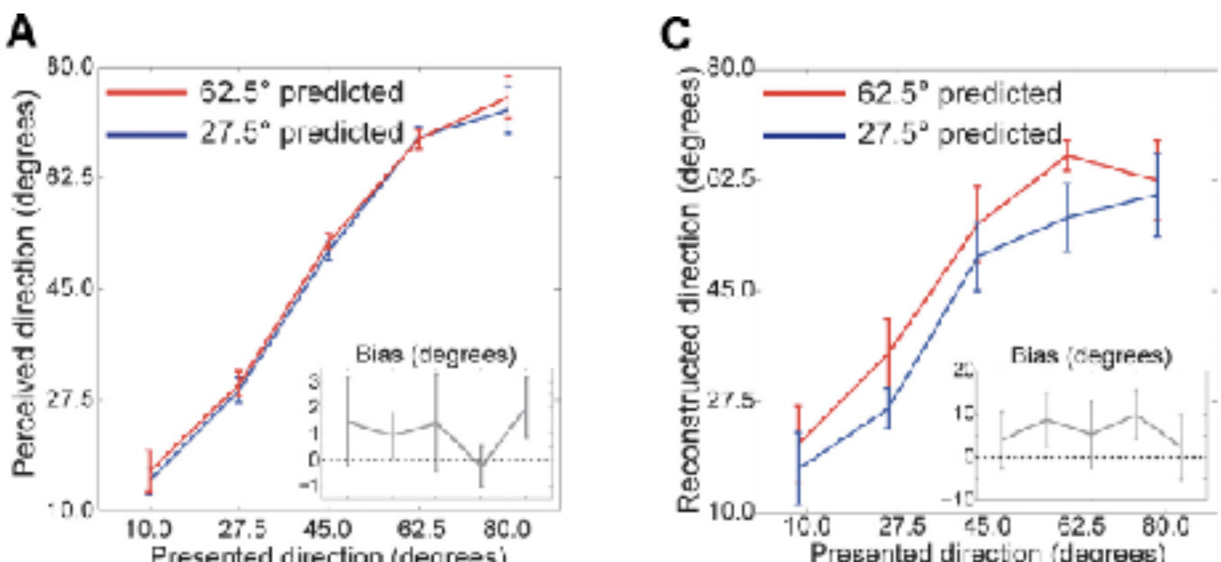

The Tuning of Neurons could implement a Prior

• The selectivities of neurons is a way by which (long-term) priors are implemented. e.g. selectivity to orientation

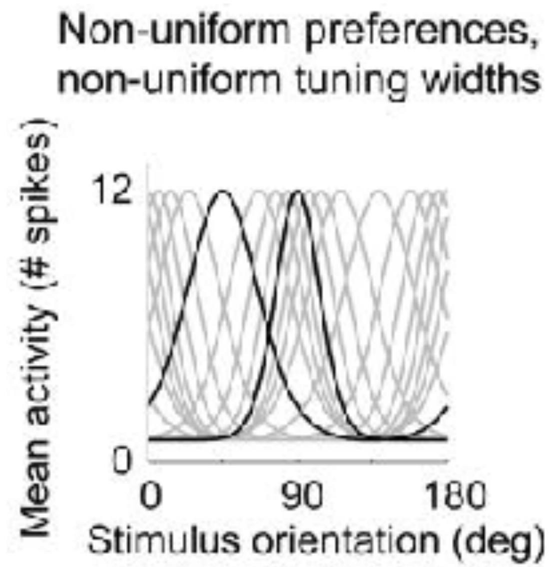

• Girshick and Simoncelli, Nat Neuro 2010.

Interpreting Orientation: A prior on Cardinal Directions

- Girshick and Simoncelli, Nat Neuro 2010.

- Orientation judgments are more accurate at cardinal (horizontal and vertical) orientations.
- Biased toward cardinal orientations.
- Prior towards cardinal orientation match orientation distribution measured in photographs.

a

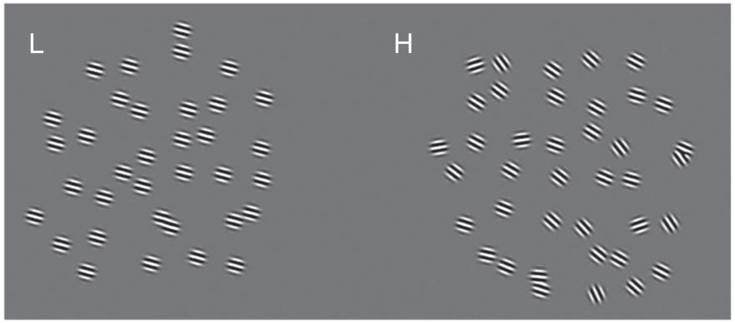

is L stimulus CW or CCW compared to H?

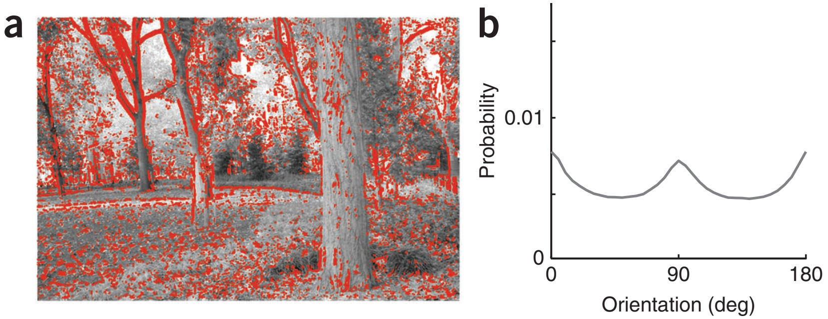

Simple Decoding Strategies

Population Vector

$$\mathbf{p} = \sum_i r_i \mathbf{e}_i$$

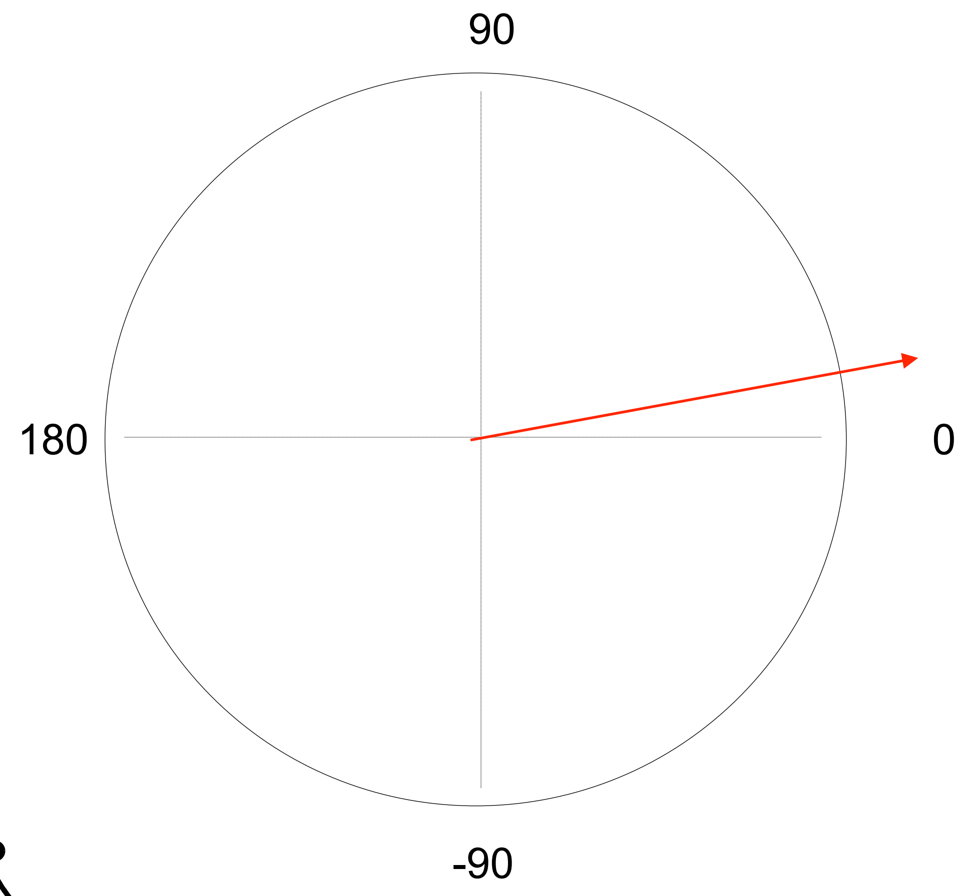

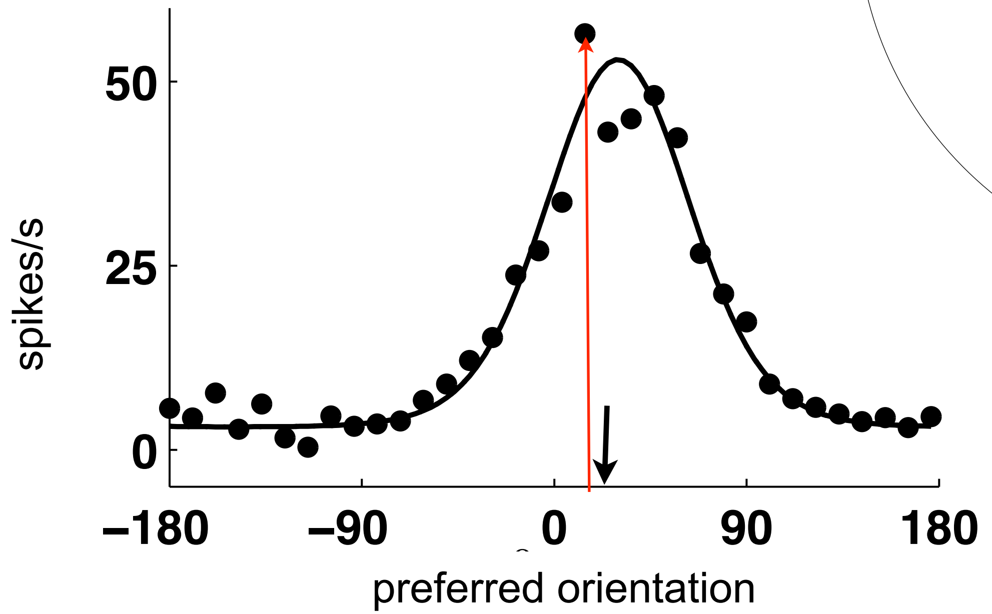

2. Simpler Decoding Strategies

Population Vector

$$\mathbf{p} = \sum_i r_i \mathbf{e}_i$$

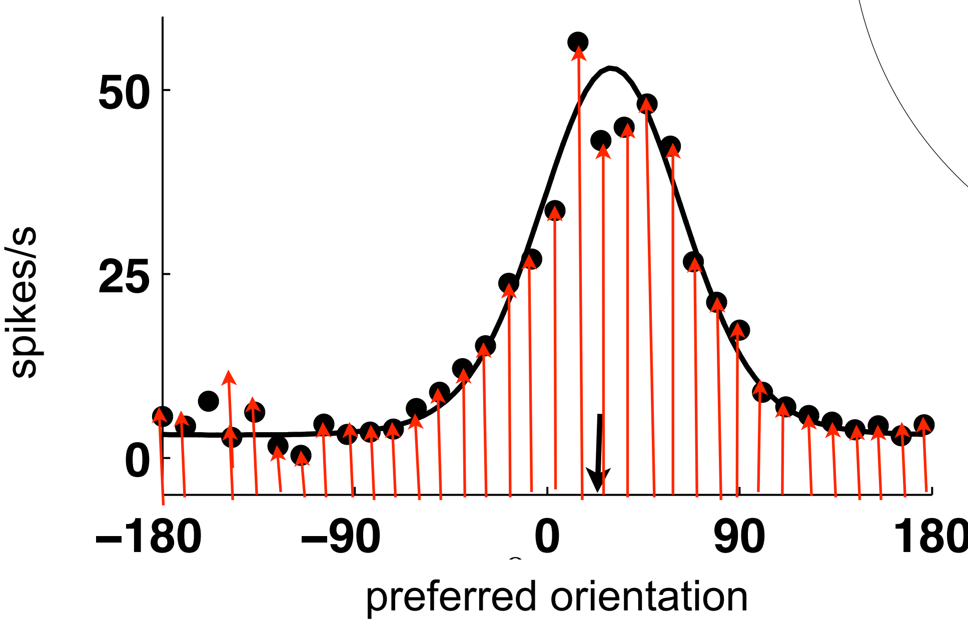

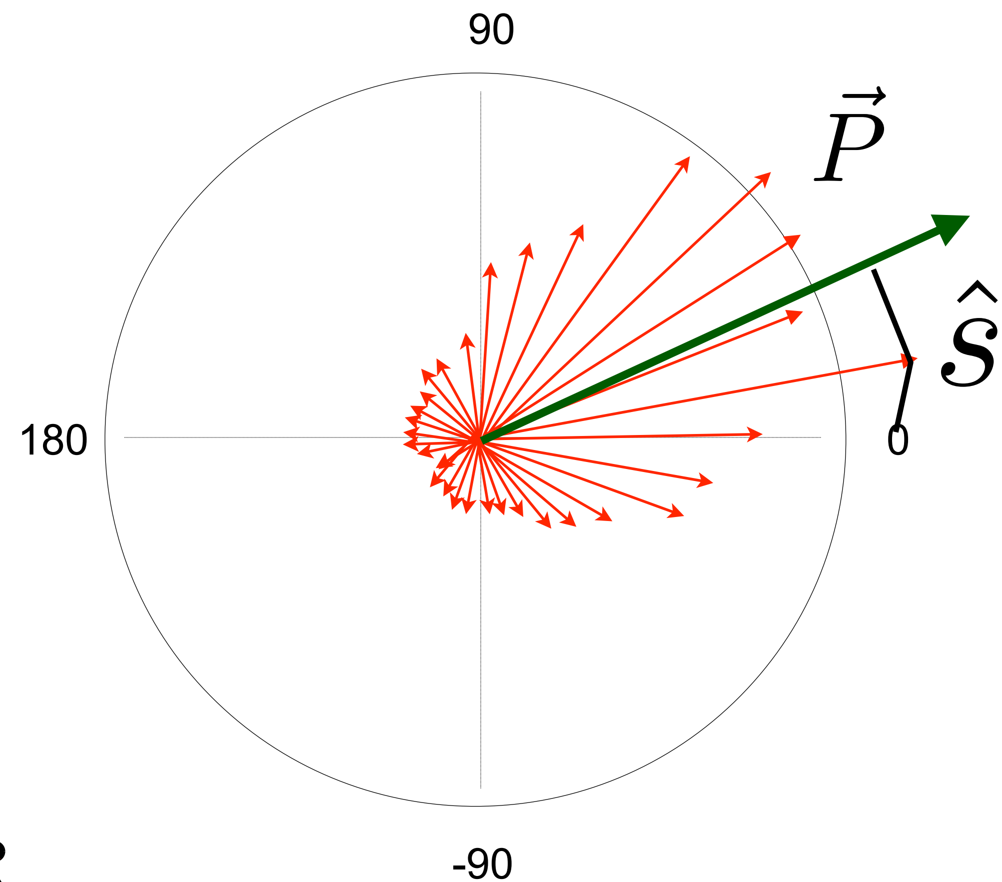

Interpreting Orientation: A prior on Cardinal Directions

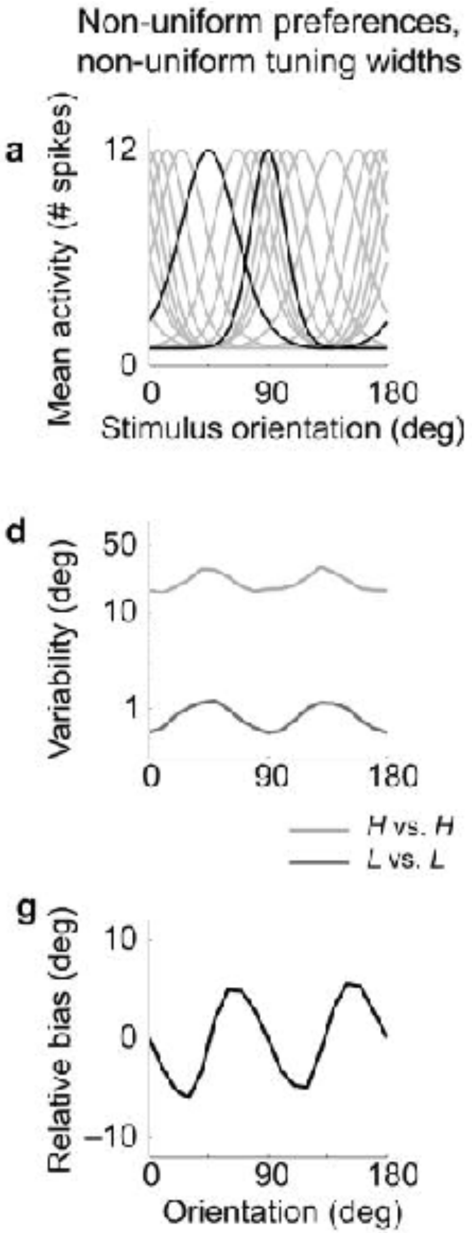

- Decoding from inhomogeneous population where cardinal orientations are over-represented with a generic decoder (population vector) produces biases and variability compatible Bayesian observer with a prior on the cardinals.

- Girshick and Simoncelli, Nat Neuro 2010.

Spontaneous activity represents the Prior

Spontaneous Cortical Activity Reveals Hallmarks of an Optimal Internal Model of the Environment

Pietro Berkes,[1]† Gergő Orbán,[1,2,3] Máté Lengyel,[3]* József Fiser[1,4,5]*

The brain maintains internal models of its environment to interpret sensory inputs and to prepare actions. Although behavioral studies have demonstrated that these internal models are optimally adapted to the statistics of the environment, the neural underpinning of this adaptation is unknown. Using a Bayesian model of sensory cortical processing, we related stimulus-evoked and spontaneous neural activities to inferences and prior expectations in an internal model and predicted that they should match if the model is statistically optimal. To test this prediction, we analyzed visual cortical activity of awake ferrets during development. Similarity between spontaneous and evoked activities increased with age and was specific to responses evoked by natural scenes. This demonstrates the progressive adaptation of internal models to the statistics of natural stimuli at the neural level.

Our percepts rely on an internal model of the environment, relating physical processes of the world to inputs received by our senses, and thus their veracity critically hinges upon how well this internal model is adapted to the statistical properties of the environment. For example, internal models in v extract the features, such as l edges or high-level objects, tha retinal image (*1*). This require model is adapted to the cooccu visual features in the environm they jointly determine natural ima of perception (*2*, *3*), motor co making (*5*, *6*), and higher cognit are governed by such statistical

[1]Volen National Center for Complex S versity, Waltham, MA 02454, USA. [2]Dep Research Institute for Particle and Nucl Academy of Sciences, H-1121 Budapes tional and Biological Learning Lab, Depa University of Cambridge, Cambridge C ment of Psychology and the Neuroscie University, Waltham, MA 02454, USA Institute for Advanced Study, Szenthárc H-1014, Hungary.

*These authors contributed equally t
†To whom correspondence should b berkes@brandeis.edu

Spontaneous activity is the statistical prior:

Berkes et al, *Science* 2011

- Evoked activity should represent the posterior for a given input image
- Spontaneous activity should represent the posterior for a blank stimulus
- This posterior should converge to prior distribution.

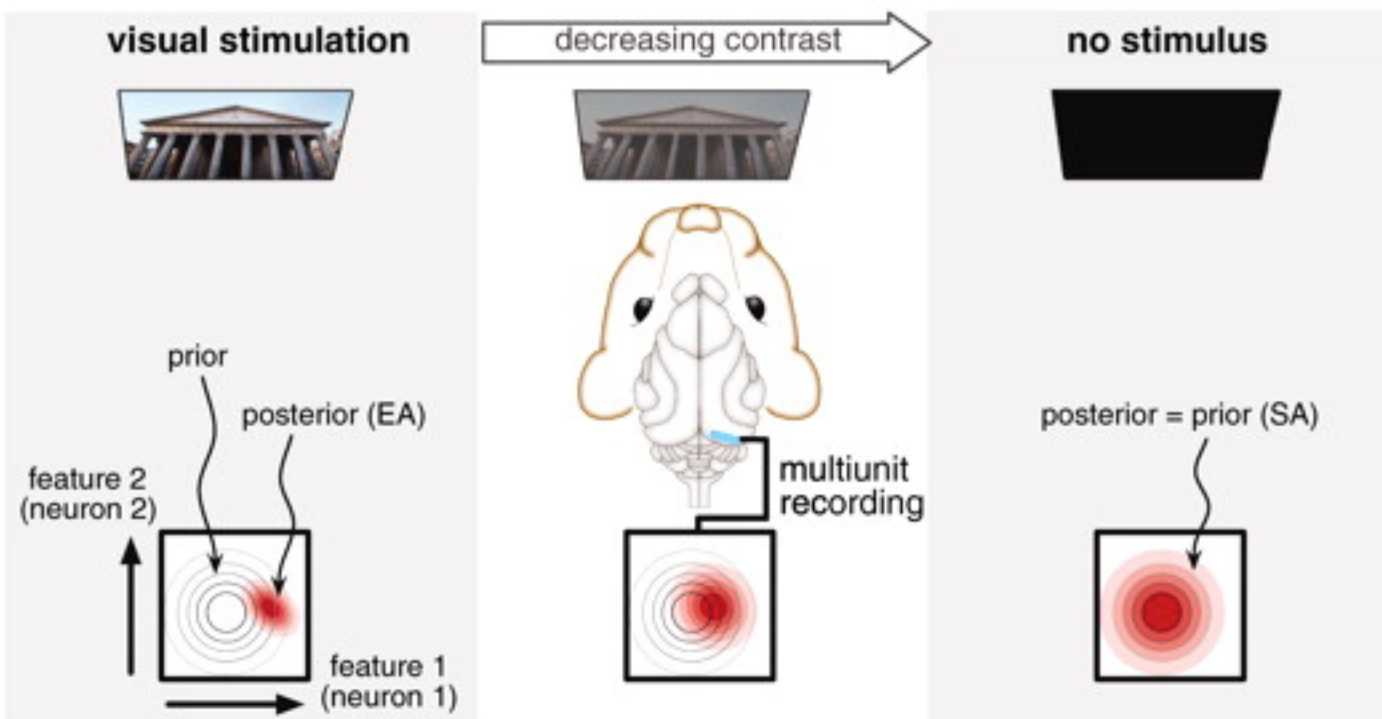

Spontaneous activity is the statistical prior:
Berkes et al, *Science* 2011

- Measured population activity within visual cortex of awake, freely viewing ferrets in response to natural scene movies and in darkness at different stages in development (postnatal P29, P44 and mature P83 and P129)
- Found that divergence between Evoked Activity and Spontaneous Activity decreases with age
- Similarity between EA and SA is specific to natural scenes
- Temporal correlations similar as well.

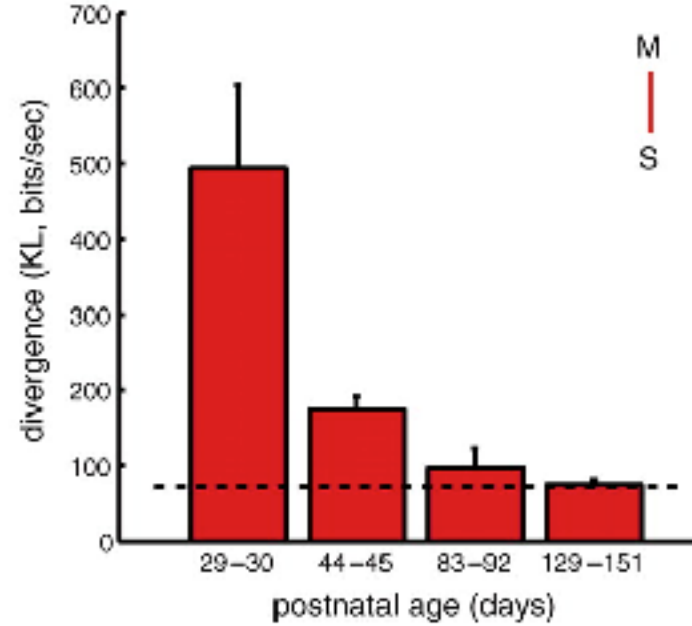

Neural Substrate of Priors

- Priors: Where in the brain ?
- Top down inputs (predictive coding)
- Increase or decrease of activity ? [e.g. Summerfield & Egner 2009]
- in Tuning of neurons? [Gershick et al 2011; Fischer & Pena 2011]
- in Baseline activity? [Berkes et al 2010]
- The representation or the read-out?

- different time scale // different mechanisms

3. How could approximate inference be implemented?

Machine learning informs us about possible approximate inference schemes:

- Sampling, Gibbs and MCMC;
- Deterministic approximation methods:

Laplace approximation and variational inference approximations

One type of variational inference approximation: Predictive Coding.

Priors as top-down inputs : Predictive Coding

• Perceptual inference: iterative matching process of top-down predictions against bottom-up evidence, along the visual cortical hierarchy.

• expectations or `representational units' that encode prediction, and error units that encode mismatch between sensory evidence and prediction and forward it to higher level.

• Mumford 1992, Rao & Ballard 1999; Lee & Mumford 2003; Friston 2005.

• experimental evidence still unclear

K. Friston

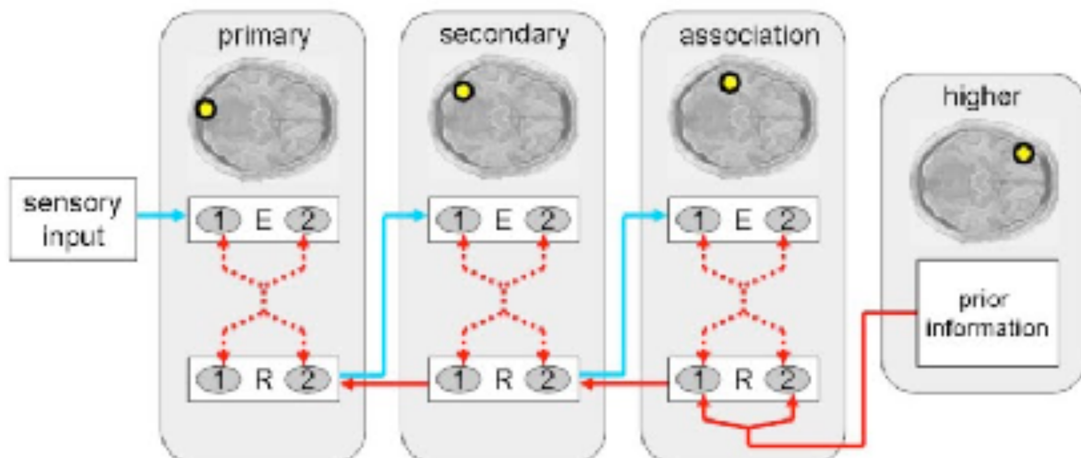

Predictive Coding:
Neural Implementation of Bayesian Inference

- Algorithms based on minimising prediction errors can approximate Bayesian inference.
- earning involves making the predictions more and more similar to the input: minimizing the prediction error.

prior p(h)

posterior
p(h|e)=p(e|h)p(h)

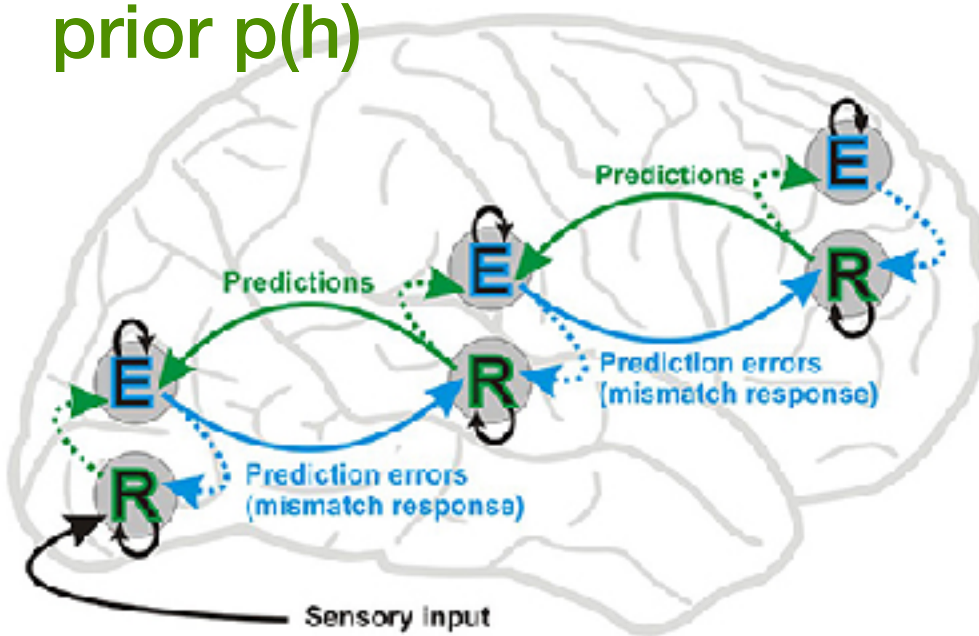

input p(e|h)

Conclusion

- Bayesian models successful at the behavioural level
- As as benchmark for performance, provide also constraints to more mechanistically models
- Much to do about: characterisation of internal models, and how they are learned, and the limits of learning.
- Applications to Psychiatry a promising avenue.
- Some confusion about the claims -- what exactly makes a neural model “Bayesian”.
- Neural implementation largely unknown. Many theories, little electrophysiological evidence, lots of very interesting questions.


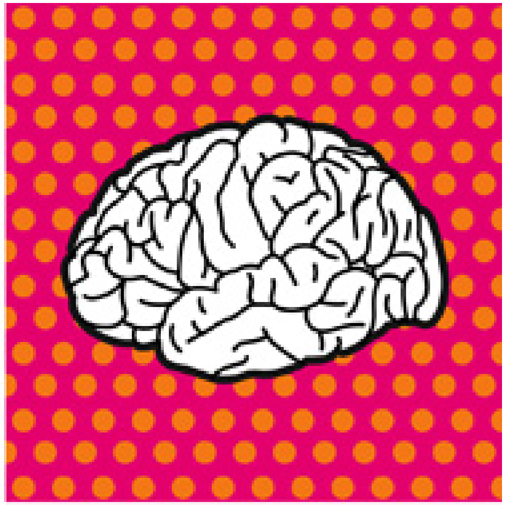
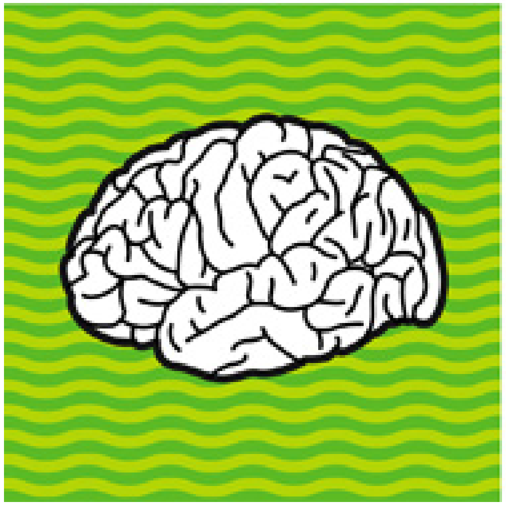

Thanks ！

This is the end of CCN lectures