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

Peggy Seriès,
IML, University of Edinburgh

Behavioural studies: 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 uncertainty into account, combine sources of information optimally
  - Use priors that are constantly updated
  - Those are consistent with (approximation) of statistics of environment at different time scales. > increase accuracy.

- 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?

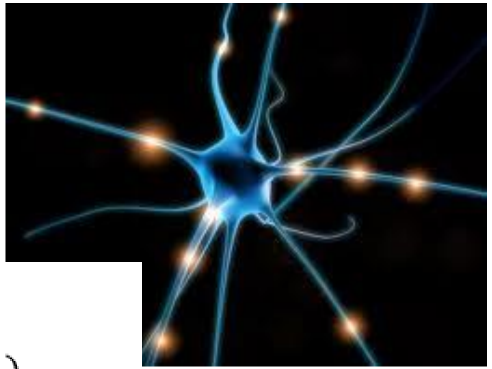

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

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

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/a002645

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]

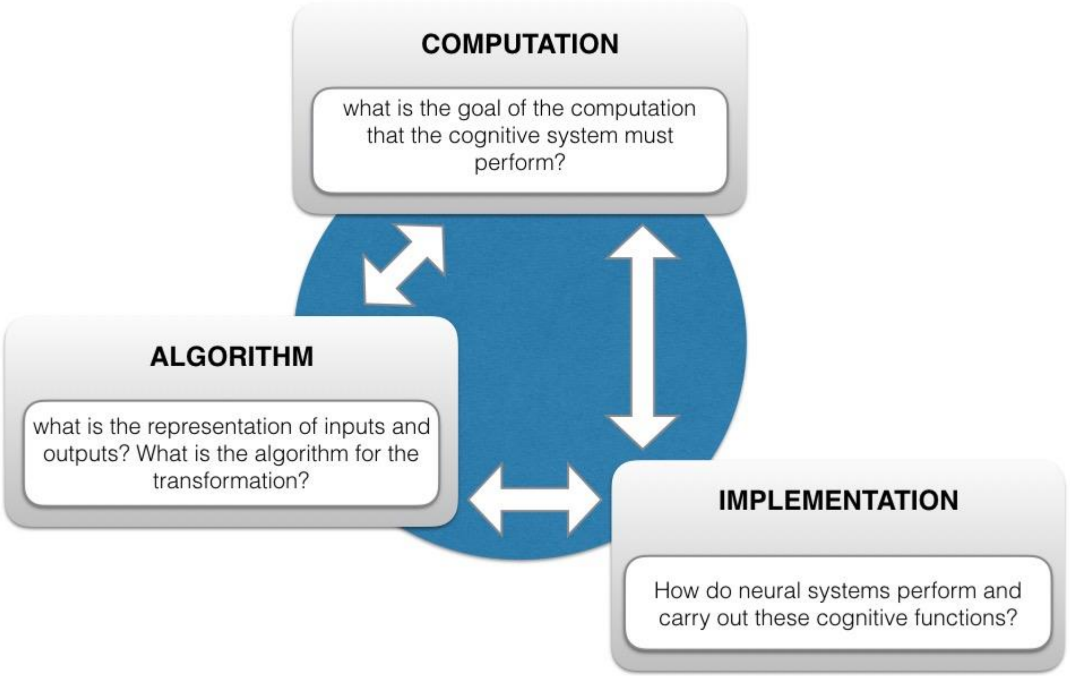

- only an approximation of Bayesian inference anyway.

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

Debates: Criticism

- Confusion about optimality

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.

Debates: Criticism

- Confusion about optimality
- Falsifiability: Bayesian models are flexible enough to account for everything

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.

Debates: Criticism

- Confusion about optimality
- Falsifiability: Bayesian models are flexible enough to account for everything
- Rarely compared with alternative (non-Bayesian) hypotheses

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)

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

Neurobiological predictions

Any plausible neural implementation must solve those computational problems:

- represent uncertainty, not just point estimates
- combine prior and likelihood
- update hierarchically
- weight updates by uncertainty/ precision
- learn statistical regularities on multiple timescales

Not ‘does the brain literally compute Bayes’ rule?’ but ‘what neural code, circuit architecture, and dynamical process could satisfy those computational requirements?”

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?

   Can we distinguish stages where the likelihoods, priors, posterior could be ‘measured’ experimentally ?

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

   Since the 2000s, 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 provide easily testable predictions. Offer proof of principle.

- 2 categories :

1.1) Probabilistic Population Codes (PPC, 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

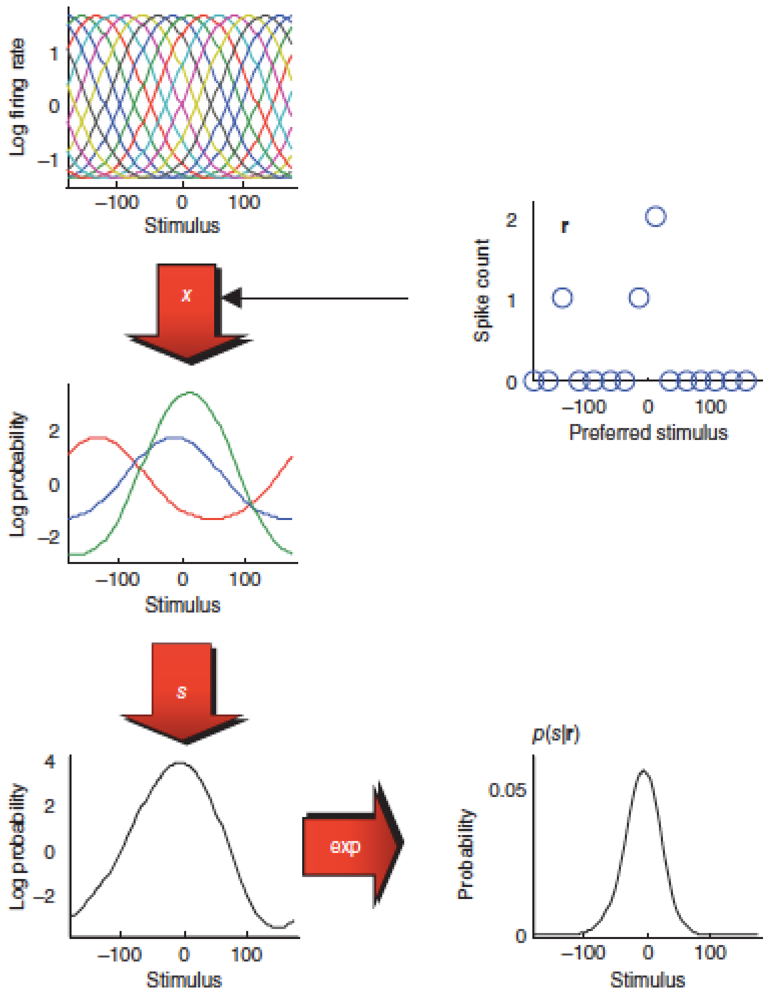

$$\log p(s \mid \mathbf{r}) = \sum_i r_i h_i(s) + \text{constant}$$

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.

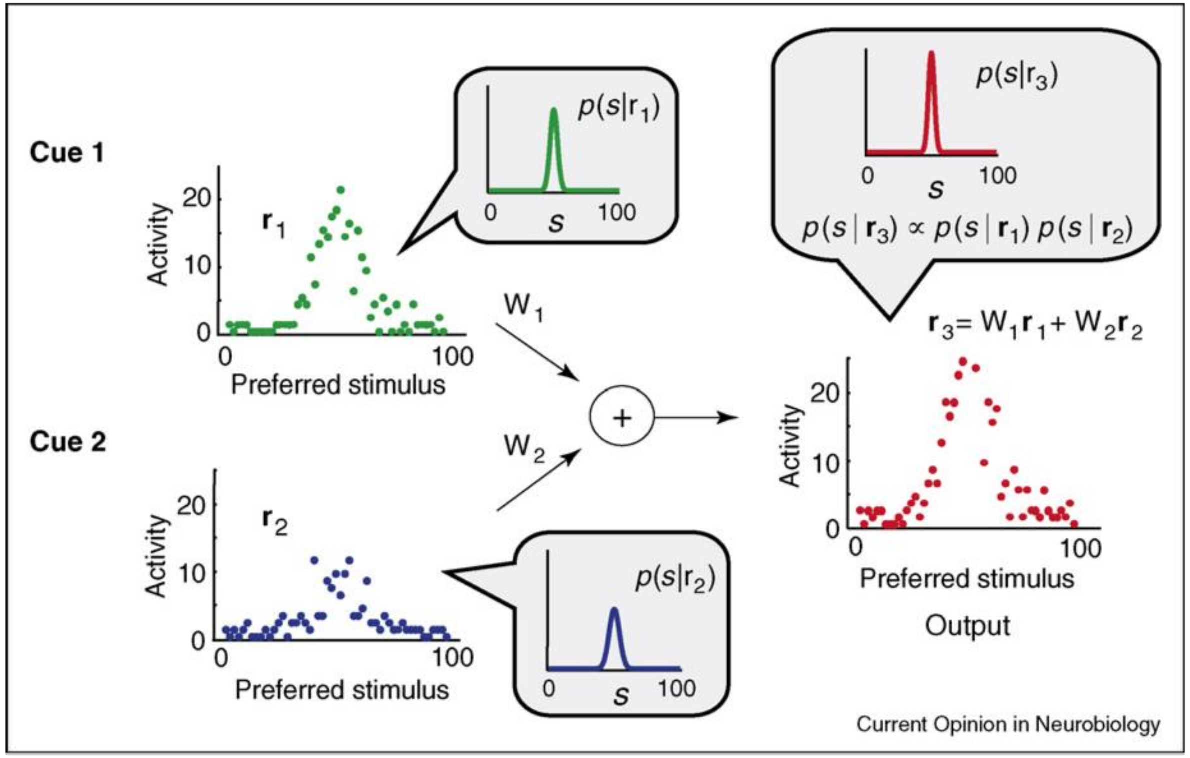

[Ma, Beck, Latham & Pouget, Nat Neuro 2006]

ARTICLES

nature neuroscience

Bayesian inference with probabilistic population codes

Wei Ji Ma[1,3], Jeffrey M Beck[1,3], Peter E Latham[2] & Alexandre Pouget[1]

Recent psychophysical experiments indicate that humans perform near-optimal Bayesian inference in a wide variety of tasks, ranging from cue integration to decision making to motor control. This implies that neurons both represent probability distributions and combine those distributions according to a close approximation to Bayes' rule. At first sight, it would seem that the high variability in the responses of cortical neurons would make it difficult to implement such optimal statistical inference in cortical circuits. We argue that, in fact, this variability implies that populations of neurons automatically represent probability distributions over the stimulus, a type of code we call probabilistic population codes. Moreover, we demonstrate that the Poisson-like variability observed in cortex reduces a broad class of Bayesian inference to simple linear combinations of populations of neural activity. These results hold for arbitrary probability distributions over the stimulus, for tuning curves of arbitrary shape and for realistic neuronal variability.

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

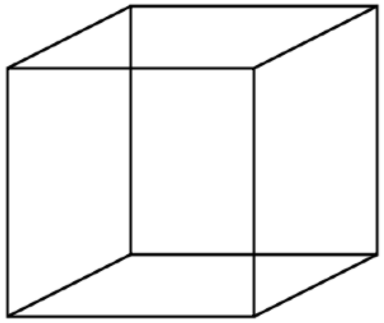

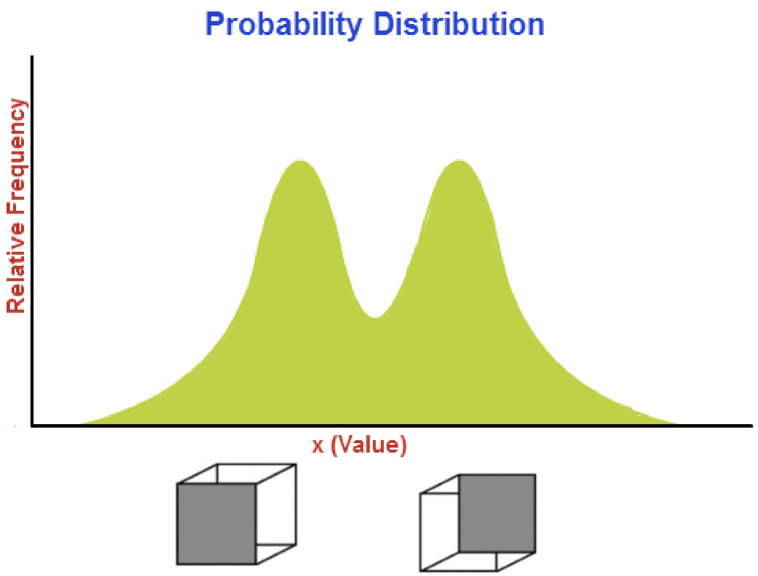

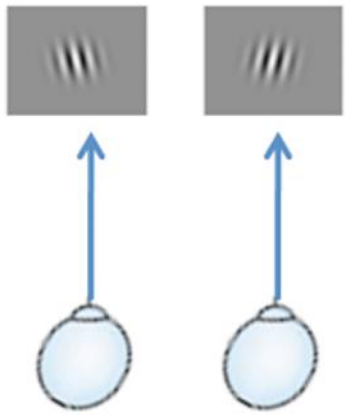

2) How could priors be implemented in the brain ?

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):16275–16284 • 16275

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, and [2]New York University, 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 a process of inference, whereby sensory inputs are combined with prior knowledge. However, despite a wealth of behavioral literature supporting an account of perception as probabilistic inference, the neural mechanisms underlying this process remain largely unknown. One important question is whether top-down expectation biases stimulus representations in early sensory cortex, i.e., whether the integration of prior knowledge and bottom-up inputs is already observable at the earliest levels of sensory processing. Alternatively, early sensory processing may be unaffected by top-down expectations, and integration of prior knowledge and bottom-up input may take place in downstream association areas that are proposed to be involved in perceptual decision-making. Here, we implicitly manipulated human subjects' prior expectations about visual motion stimuli, and probed the effects on both perception and sensory representations in visual cortex. To this end, we measured neural activity noninvasively using functional magnetic resonance imaging, and applied a forward modeling approach to reconstruct the motion direction of the perceived stimuli from the signal in visual cortex. Our results show that top-down expectations bias representations in visual cortex, demonstrating that the integration of prior information and sensory input is reflected at the earliest stages of sensory processing.

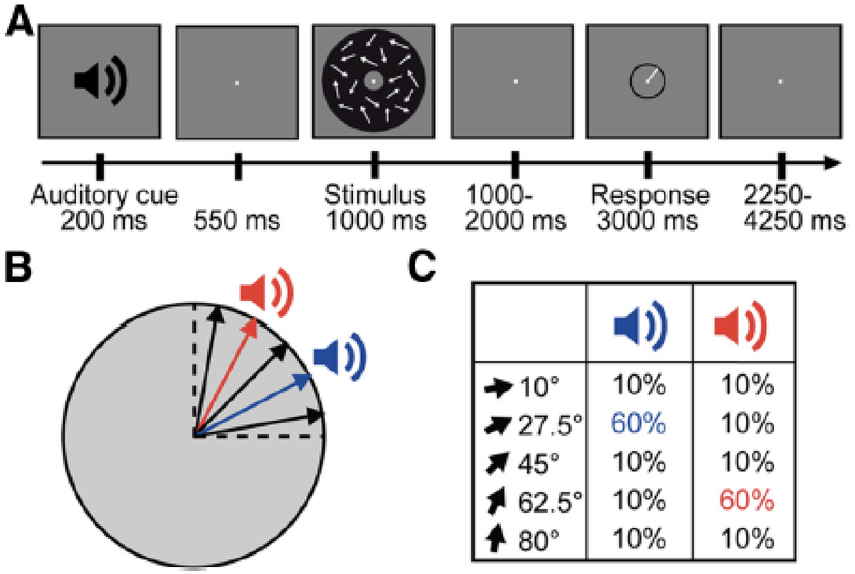

Can the effect of prior expectations be observed in fMRI ? (1)

- 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

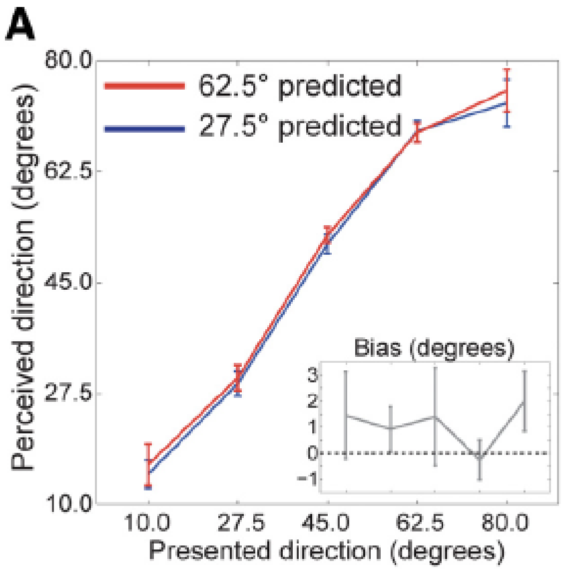

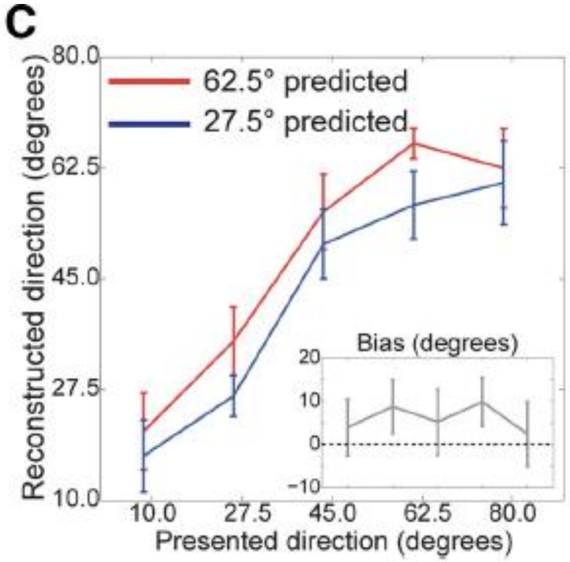

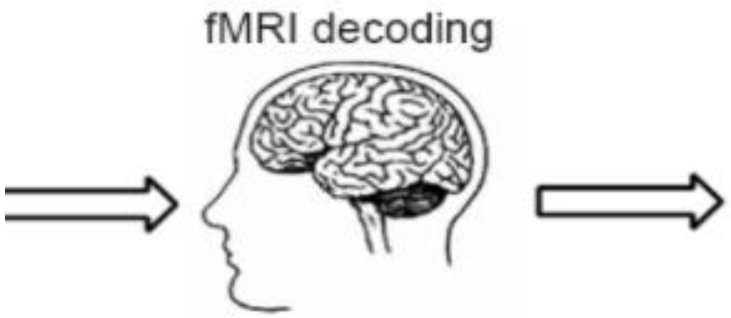

The Tuning of Neurons could implement a Prior

- The selectivities of neurons is a way by which (long-term) priors are implemented.

e.g. selectivity to orientation

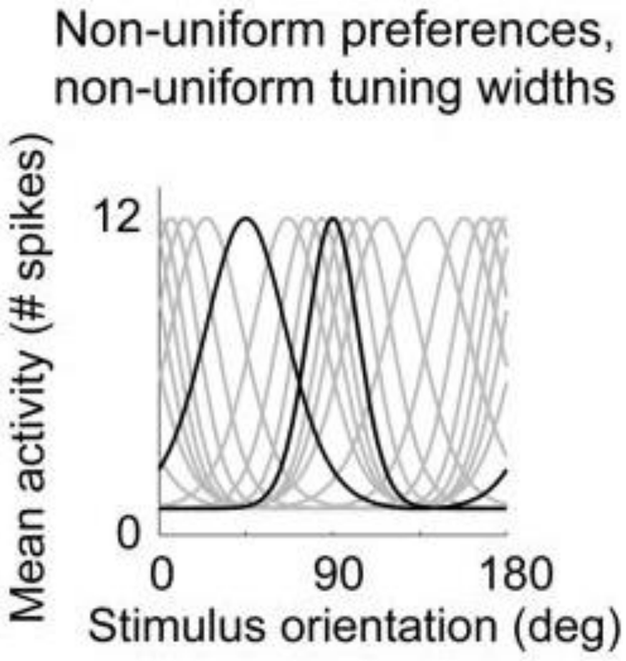

- Girshick and Simoncelli, Nat Neuro 2010.

Interpreting Orientation: A prior on Cardinal Directions

- Girshick and Simoncelli, Nat Neuro 2010.

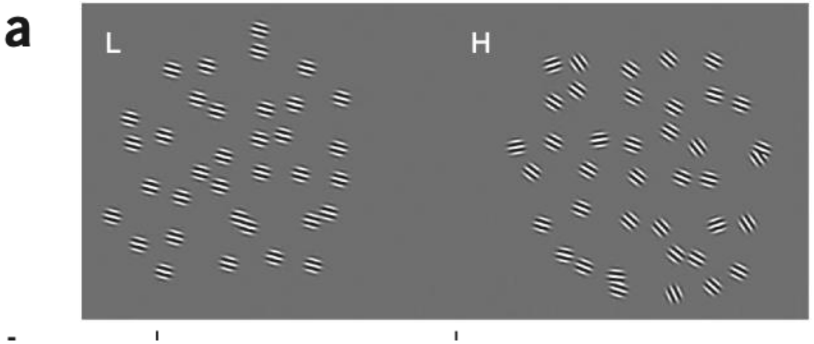

is L stimulus CW or CCW compared to H?

- 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.

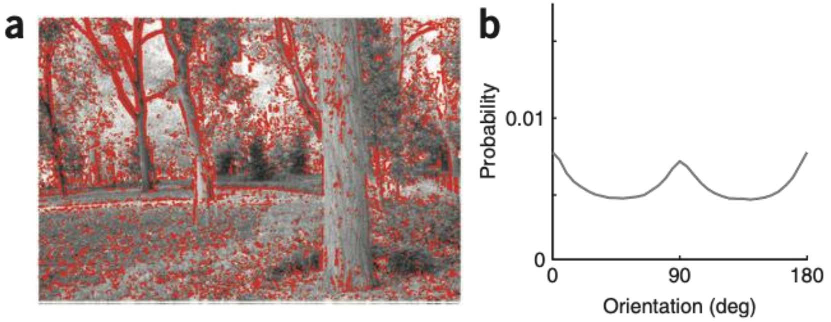

2. Simpler Decoding Strategies

Population Vector

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

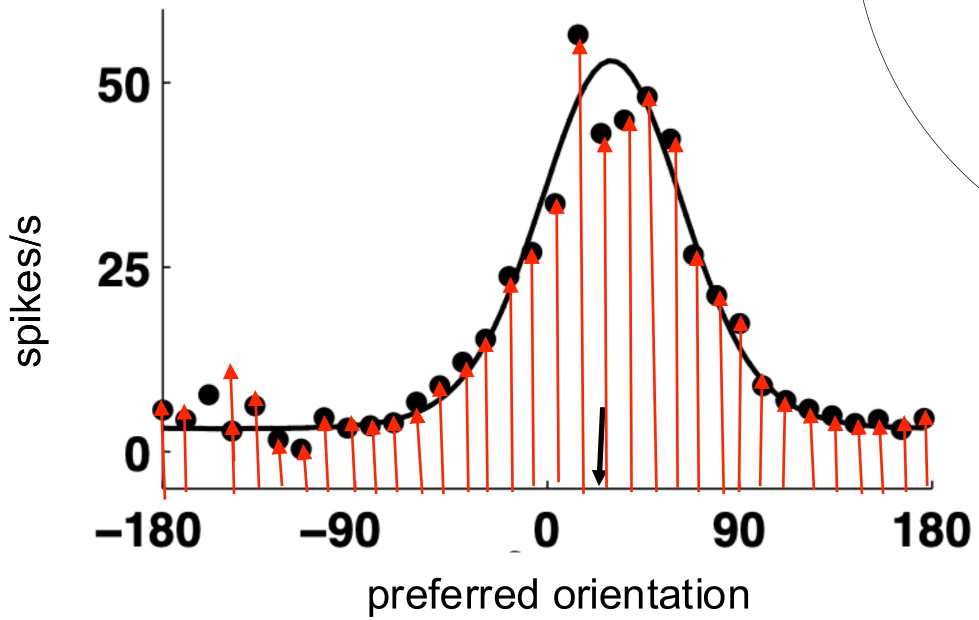

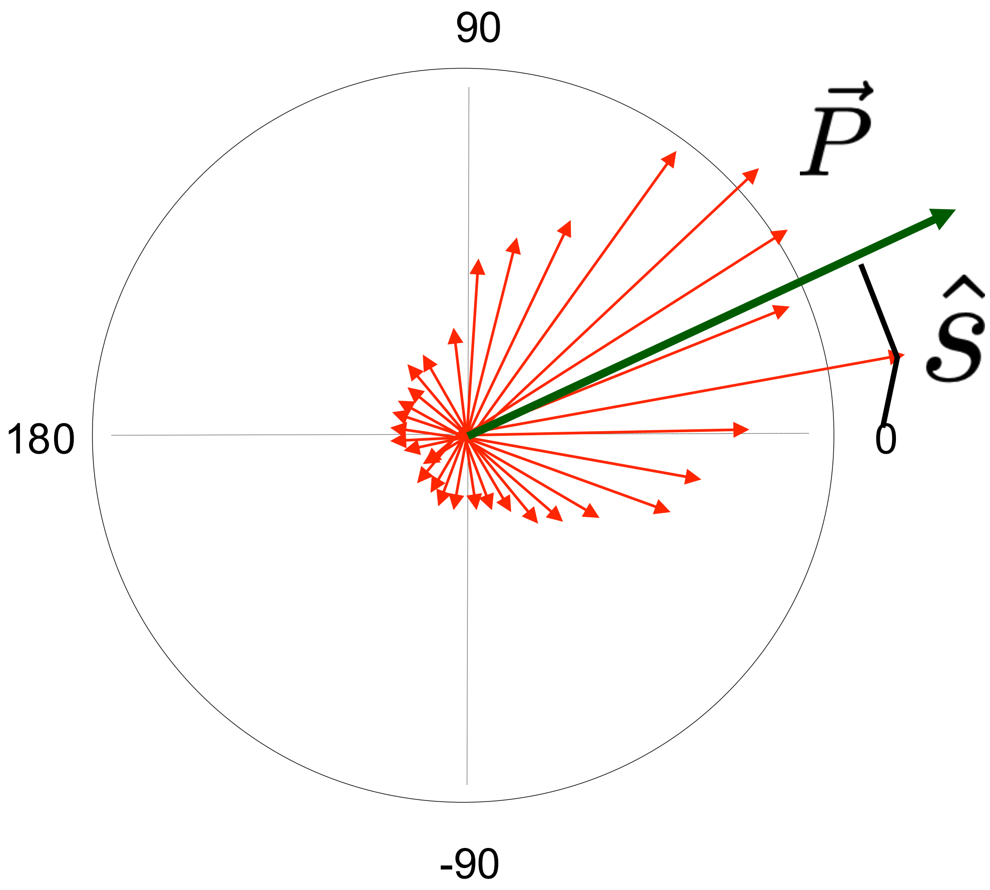

Interpreting Orientation: A prior on Cardinal Directions

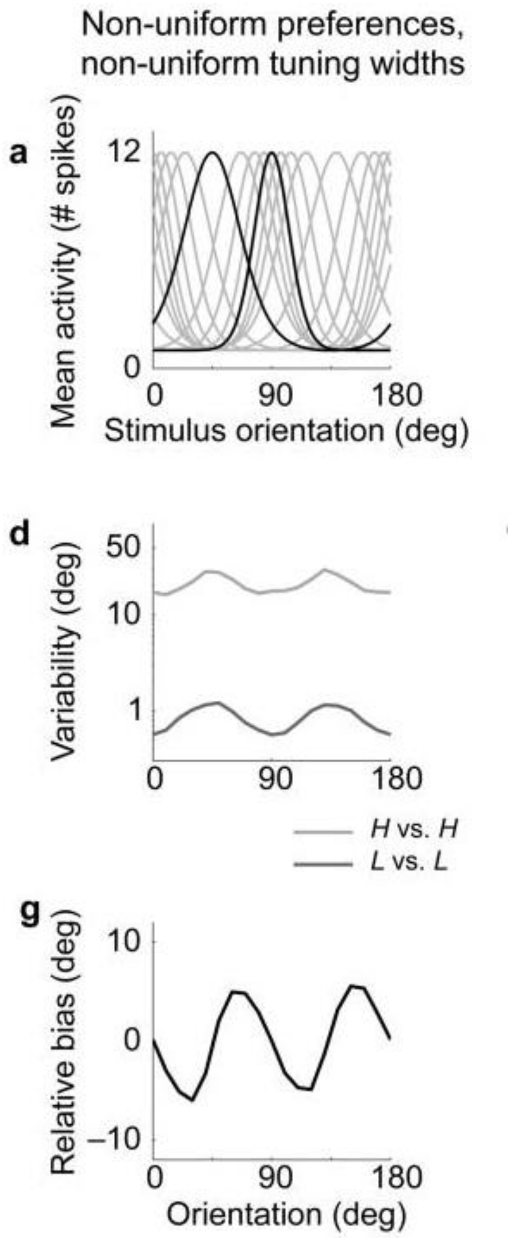

- Decoding from inhomogeneous population where cardinal orientations are over-represented with a generic decoder (population vector) produces biases and variability compatible Bayesian observer with a prior on the cardinals.

- 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 vision are used to extract the features, such as low-level oriented edges or high-level objects, that gave rise to the retinal image (*1*). This requires that the internal model is adapted to the cooccurrence statistics of visual features in the environment and the way they jointly determine natural images. Several aspects of perception (*2*, *3*), motor control (*4*), decision making (*5*, *6*), and higher cognitive reasoning (*7*, *8*) are governed by such statistically optimal internal

[1]Volen National Center for Complex Systems, Brandeis University, Waltham, MA 02454, USA. [2]Department of Biophysics, Research Institute for Particle and Nuclear Physics, Hungarian Academy of Sciences, H-1121 Budapest, Hungary. [3]Computational and Biological Learning Lab, Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK. [4]Department of Psychology and the Neuroscience Program, Brandeis University, Waltham, MA 02454, USA. [5]Collegium Budapest Institute for Advanced Study, Szentháromság utca 2, Budapest H-1014, Hungary.

*These authors contributed equally to this work.
†To whom correspondence should be addressed. E-mail: 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.

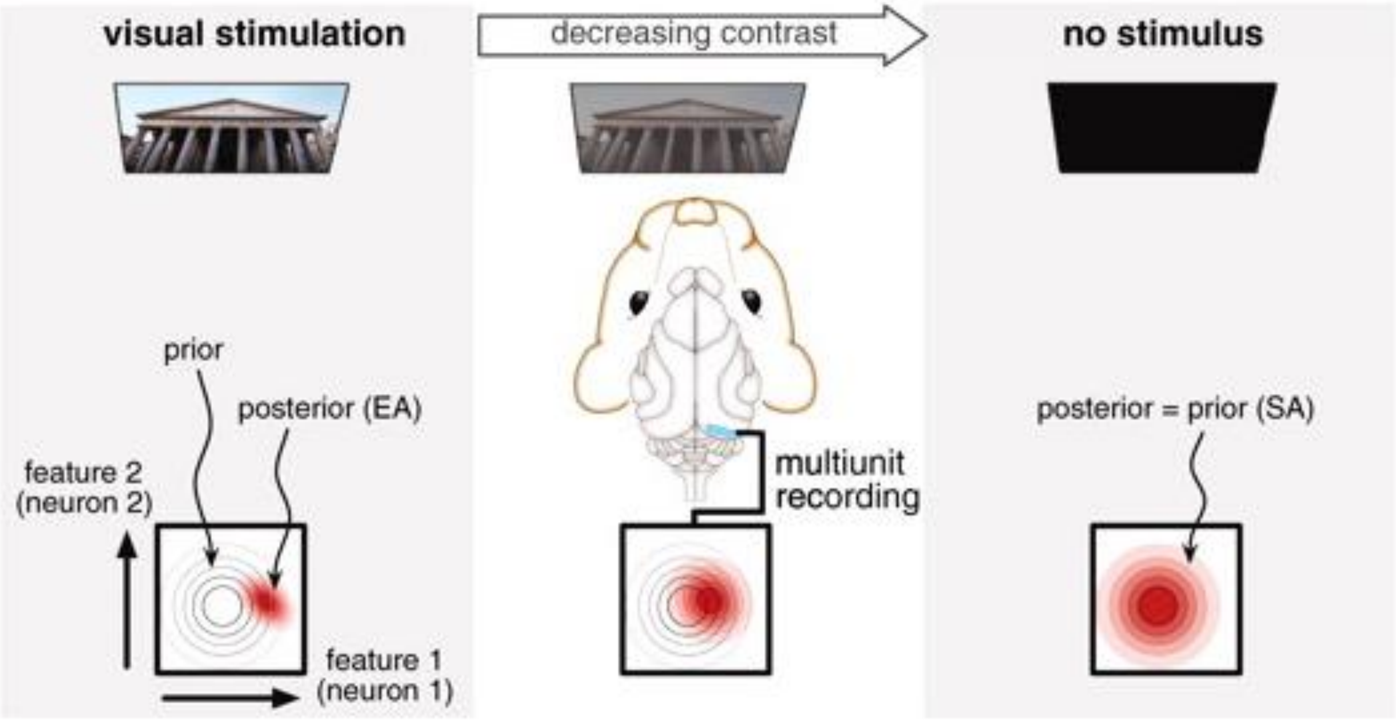

Spontaneous activity is the statistical prior: Berkes et al, *Science* 2011

- 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.

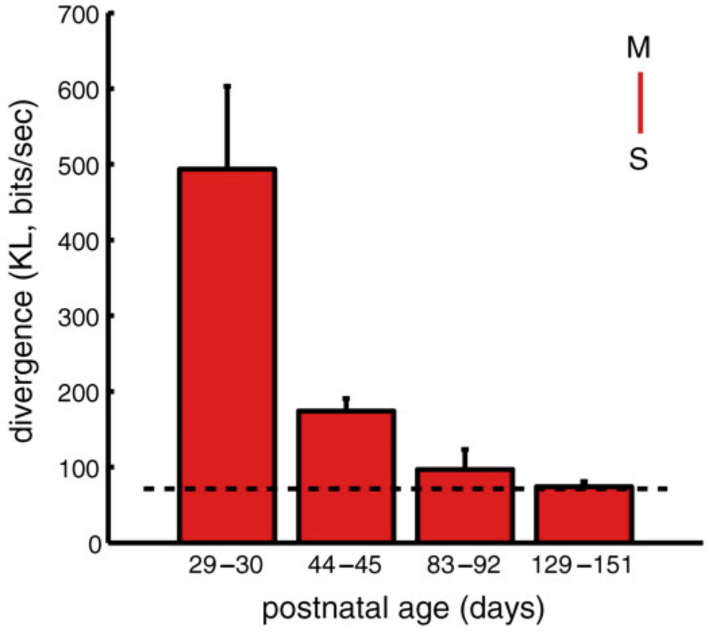

2) How could priors be implemented in the brain ?

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?

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.

K. Friston

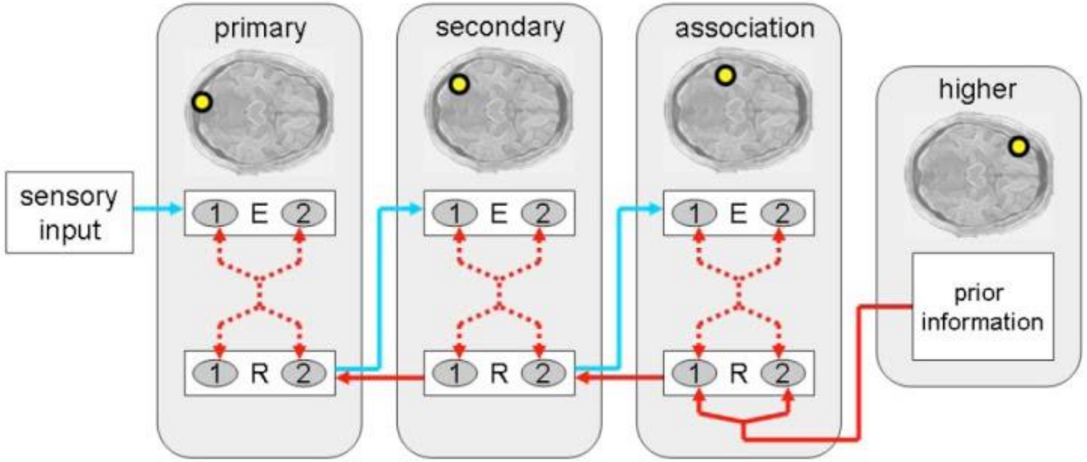

Predictive Coding: Neural Implementation of Bayesian Inference

- Algorithms based on minimising prediction errors can approximate Bayesian inference.
- Learning 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)

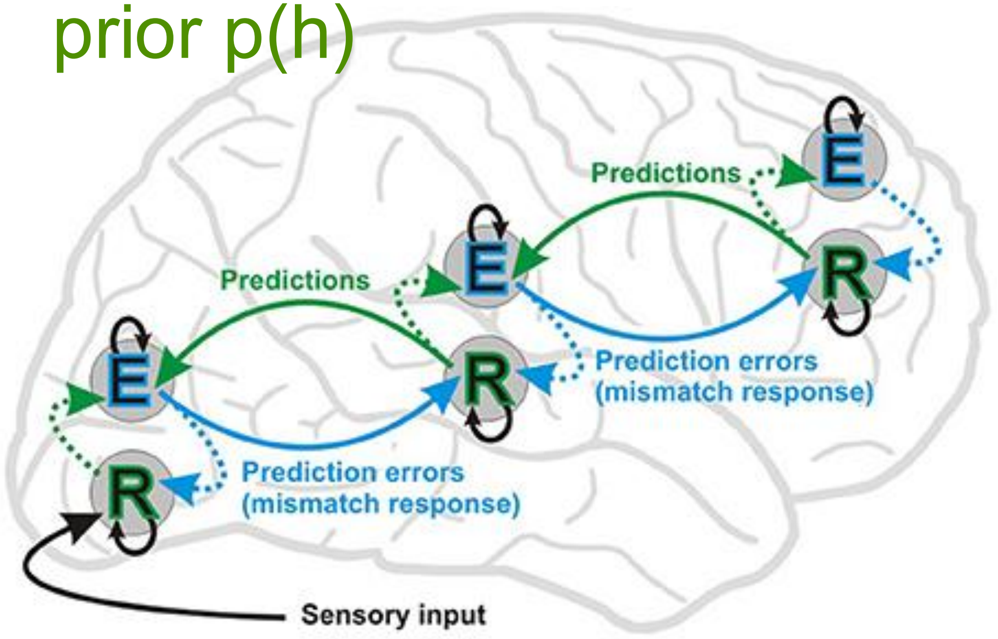

input p(e|h)

How would Predictive Coding be realized? Layers and Rhythms

- Bastos et al 2012, 2020
- Superficial layers: candidate feedforward error signals
- Deep layers: candidate feedback prediction signals
- Local recurrence computes iterative updates

- Feedforward, error-like signals might be relatively associated with faster rhythms (Gamma 30-100 Hz)
- While feedback, prediction-like signals might be associated with slower rhythms (Beta 13-30 Hz/Alpha 8-13Hz).

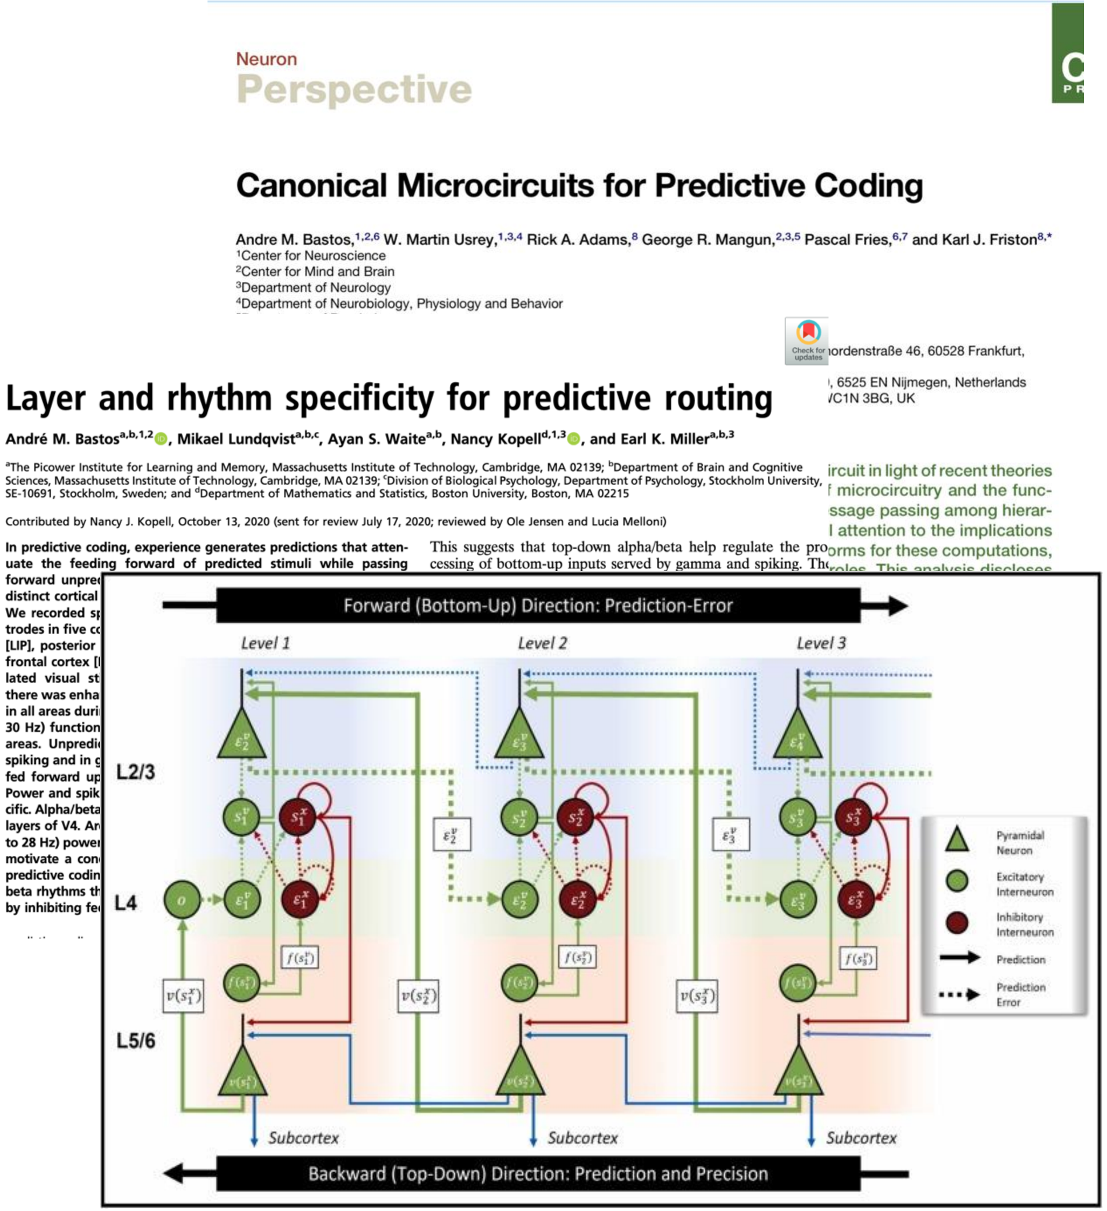

How would Predictive Coding be realized? Layers

Thomas et al. (2024), 7T laminar fMRI

- Expected orientations decoded across layers

- Unexpected orientations decoded mainly in superficial laminae

- Broadly consistent with superficial error-like signalling

Report

Current Biology

Predictions and errors are distinctly represented across V1 layers

Highlights

- A 7T fMRI study presents expected and unexpected Gabor orientations
- Expectation differentially modulates decoding across primary visual cortex layers
- This pattern supports predictive processing accounts of the brain and mind

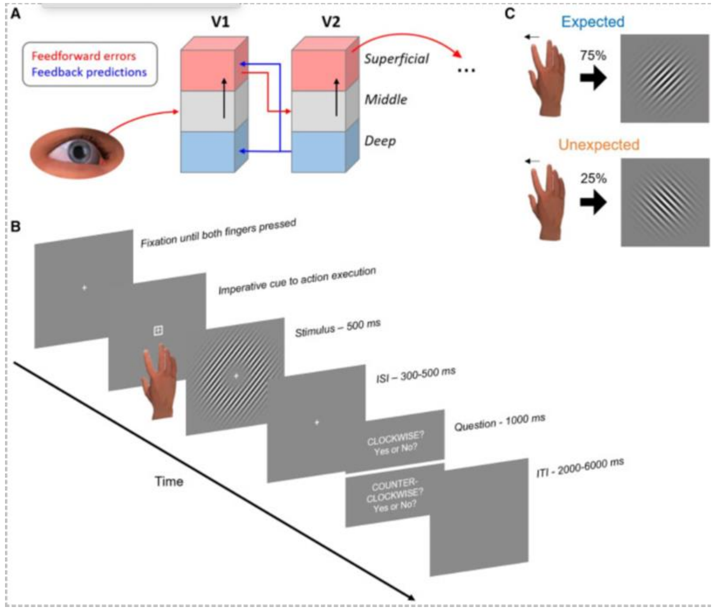

Authors

Emily R. Thomas, Joost Haarsma, Jessica Nicholson, Daniel Yon, Peter Kok, Clare Press

Correspondence

emilyrosethomas@outlook.com (E.R.T.), c.press@ucl.ac.uk (C.P.)

In brief

Thomas et al. use 7T fMRI to find that expected events are represented similarly across deep, middle, and superficial layers of the primary visual cortex, while unexpected events are only robustly represented in superficial layers. These findings support accounts of sensory processing requiring distinct

How would Predictive Coding be realized? The importance of precision

- Predictive coding is not just about prediction error, but precision-weighted prediction error: The brain can up- or down- modulate the weight given to the prediction error for updating (corresponding to changing the relative balance between likelihood and priors).

- Biologically plausible implementation = gain control

- Neuromodulators might implement precision, e.g. Acetylocholine (Ach), DA.

The Journal of Neuroscience, May 8, 2013 • 33(19):8227–8236 • 8227

Behavioral/Cognitive

Free Energy, Precision and Learning: The Role of Cholinergic Neuromodulation

Rosalyn J. Moran,[1,2]
[1]Wellcome Trust Centre fo
Carilion Research Institute
Psychology, Autonoma Un
Zurich and Swiss Federal I
Department of Economics,

Acetylcholine (ACh) is
computational simulat
the free energy princip
hierarchies by optimizi

eLife | Home Browse Magazine Community About Searc

Research Article
Neuroscience

Acetylcholine modulates the precision of prediction error in the auditory cortex

David Pérez-González, Ana Belén Lao-Rodríguez, Cristian Aedo-Sánchez, Manuel S Malmierca
n, Calle Pintor Fernando
of Basic Psychology,
ad Jardín, University of
uel de Unamuno, University of

Molecular Psychiatry (2021) 26:5320–5333
https://doi.org/10.1038/s41380-020-0803-8

ARTICLE

Precision weighting of cortical unsigned prediction error signals benefits learning, is mediated by dopamine, and is impaired in psychosis

J. Haarsma[1] · P. C. Fletcher[1,2,3] · J. D. Griffin[1] · H. J. Taverne[1] · H. Ziauddeen[1,2,3] · T. J. Spencer[4] · C. Miller[1] · T. Katthagen[5] · I. Goodyer[1,3] · K. M. J. Diederen[1,4] · G. K. Murray[1,3]

Received: 11 April 2019 / Revised: 24 April 2020 / Accepted: 28 May 2020 / Published online: 24 June 2020

Abstract
Recent theories of cortical function construe the brain as performing hierarchical Bayesian inference. According to these theories, the precision of prediction errors plays a key role in learning and decision-making, is controlled by dopamine and contributes to the pathogenesis of psychosis. To test these hypotheses, we studied learning with variable outcome-precision in healthy individuals after dopaminergic modulation with a placebo, a dopamine receptor agonist bromocriptine or a

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.
- Some confusion about the claims -- what exactly makes a neural model “Bayesian”.
- Applications to Psychiatry a promising avenue to understanding individual differences

- Neural implementation is still to be established. Predictive coding promising but not yet completely established.

- Many theories, electrophysiological evidence still limited, lots of very interesting questions.

Thanks !

This is the end of CCN lectures

Next time: exam revision