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

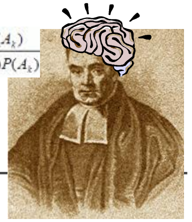

The `Bayesian' approach to perception, cognition and disease (3)

Peggy Seriès,
IANC, University of Edinburgh

A Bayesian theory of the Brain

- 1990s- Purpose of the brain: infer state of the world from noisy and incomplete data [G. Hinton, P. Dayan, A. Pouget, R. Zemel, R. Rao, etc..]

- Perception often modelled using the framework of Bayesian Inference

$$P(h_1|e) = \frac{P(e|h_1)P(h_1)}{P(e)}$$

$$\text{posterior} = \frac{\text{likelihood} \times \text{prior}}{\text{normalizing constant}}$$

manipulating probabilities -- degree of belief.

"Instead of trying to come up with an answer to a question, the brain tries to come up with a probability that a particular answer is correct," Alex Pouget.

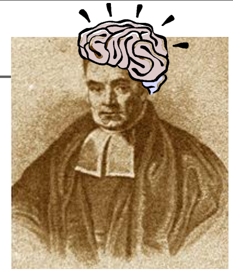

Reverend Thomas Bayes, 1702- 1761

Behavioural studies: So what have we learned?

- Bayesian model offer elegant/ 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.
- Those priors (but also 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 ?

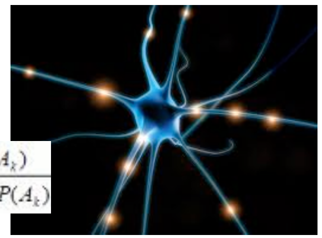

$$P(A_k \mid B) = \frac{P(B \mid A_k)P(A_k)}{\sum_{i=1}^{n} P(B \mid A_i)P(A_i)}$$

Is the Brain "Bayesian"? Debates

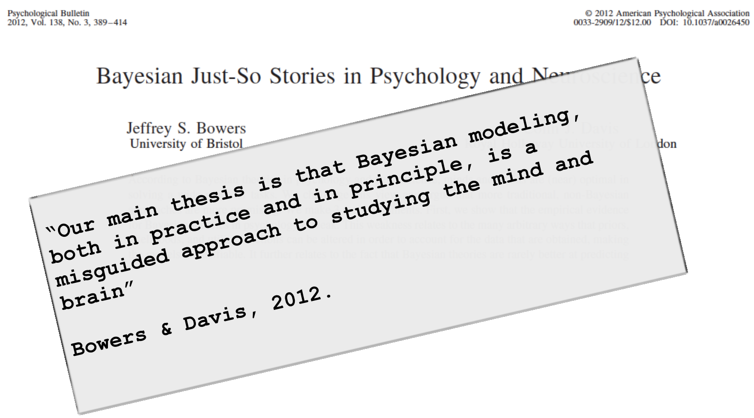

Psychological Bulletin
2012, Vol. 138, No. 3, 389–414

0033-2909/12/$12.00 DOI: 10.1037/a0026450

Bayesian Just-So Stories in Psychology and Ne[illegible]ce

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.

A Bit of Philosophy

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

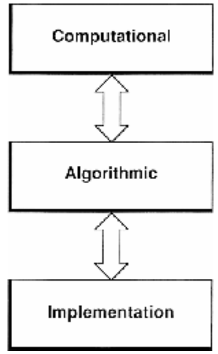

- *"Bayesian models are not intended to provide mechanistic or process accounts of cognition"* *[Jacobs and Kruschke, 2010]*

- only an approximation of Bayesian inference anyway.

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

Debates: criticism

- Confusion about optimality
- Falsifiability: 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 the 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 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 ?

- How do populations of neurons represent uncertainty ?
- Does neural activity represent probabilities? (log probabilities?)
- Can we distinguish stages where the likelihoods, priors, posterior could be 'measured' experimentally ?
- Can networks of neurons implement optimal inference?
- How can we discover the priors used by the brain?
- How can a prior be implemented? (baseline - spontaneous activity, number of neurons, gain, connectivity?).
- Recently, active topic of theoretical research (e.g. A. Pouget, S. Deneve, P. Dayan, R. Rao, J. Fiser, M. Lengyel).

1) A question about Representation: how do 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[38–42].

- probabilities are functions: neural activity represents the parameters of that function, possibly parameter in basis function parametrisation

Probabilistic population codes:
spiking rates could represent the coefficients of a basis function parametrisation of the log probability

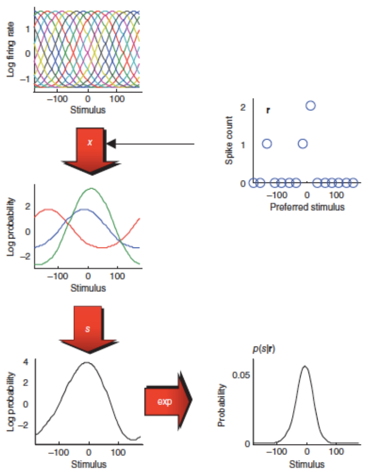

$$\log p(s \mid \mathbf{r}) = \sum_i r_i h_i(s) + \text{constant}$$

Pouget et al 2013,
Probabilistic brains: known and unknown

11

Optimal cue integration with PPC

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

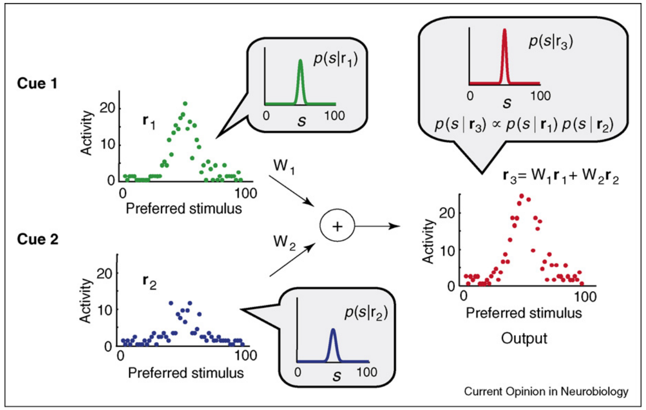

[Ma, Beck, Latham & Pouget, Nat Neuro 2006]

A question about Representation: how do 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) 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.

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

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

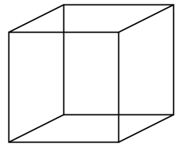

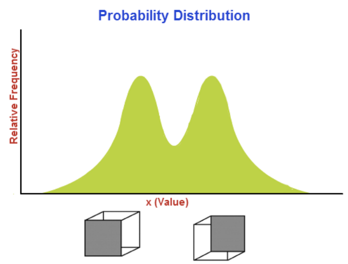

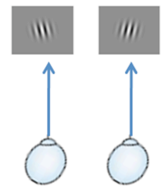

Neural Substrate of Priors

- 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

Experimental investigation of expectations

- is problematic, due to confounds between expectations and attention, adaptation and learning.
- Attention = enhancement of responses
- expectations expressed as suppression of activity?

[Summerfield & Egner, 2009]

- e.g. mismatch negativity: response to odd stimulus in a sequence -- consistent with predictive coding.

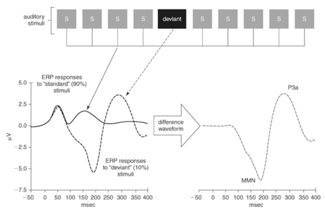

Can the effect of prior expectations be observed in fMRI ?

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, Neth
Department of Psychology and Center for Neural Science, New York, New York 10003

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

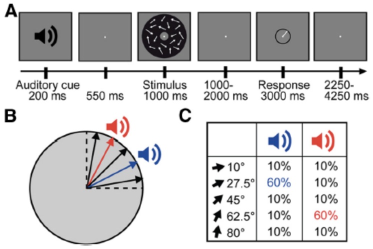

Can the effect of prior expectations be observed in fMRI ?

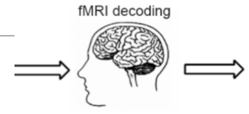

- Decoding from visual cortex : Does activity in visual cortex (V1, V2, V3, V4, MT) correspond to real stimulus or percept ? A: percept.
- Integration of prior expectations and sensory information in population activity is observed at the level of BOLD signals as early as in V1

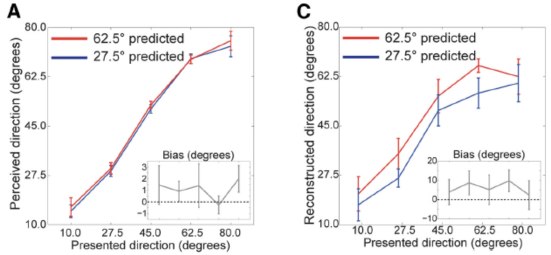

2. Neural Implementation of probabilities and priors?

- It has also been proposed that the selectivities of neurons could be a way by which priors are implemented.

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.

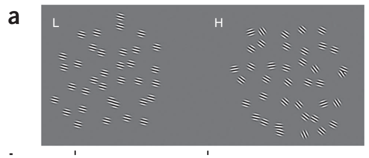

is L stimulus CW or CCW compared to H?

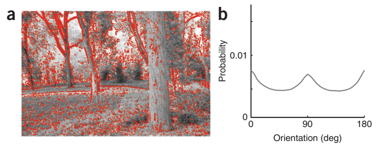

Spontaneous activity represents the prior

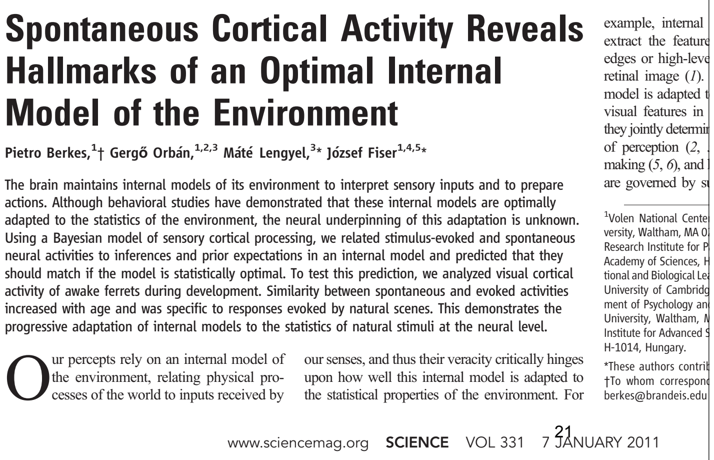

Spontaneous Cortical Activity Reveals Hallmarks of an Optimal Internal Model of the Environment

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 extract the feature edges or high-leve retinal image (1). model is adapted t visual features in they jointly determi of perception (2, making (5, 6), and are governed by s

[1]Volen National Cente versity, Waltham, MA 0 Research Institute for P Academy of Sciences, H tional and Biological Le University of Cambridg ment of Psychology an University, Waltham, M Institute for Advanced S H-1014, Hungary.
*These authors contrib
†To whom correspond berkes@brandeis.edu

www.sciencemag.org SCIENCE VOL 331 7 JANUARY 2011

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.

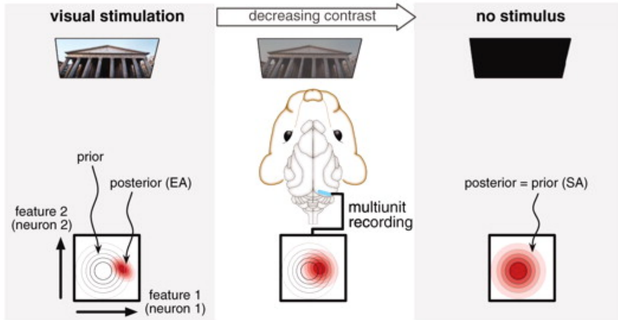

Spontaneous activity is the statistical prior: Berkes et al, *Science* 2011

- 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.

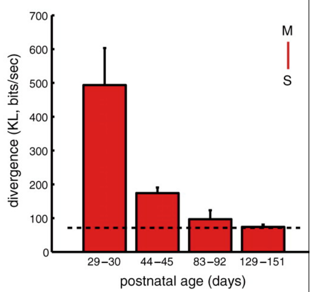

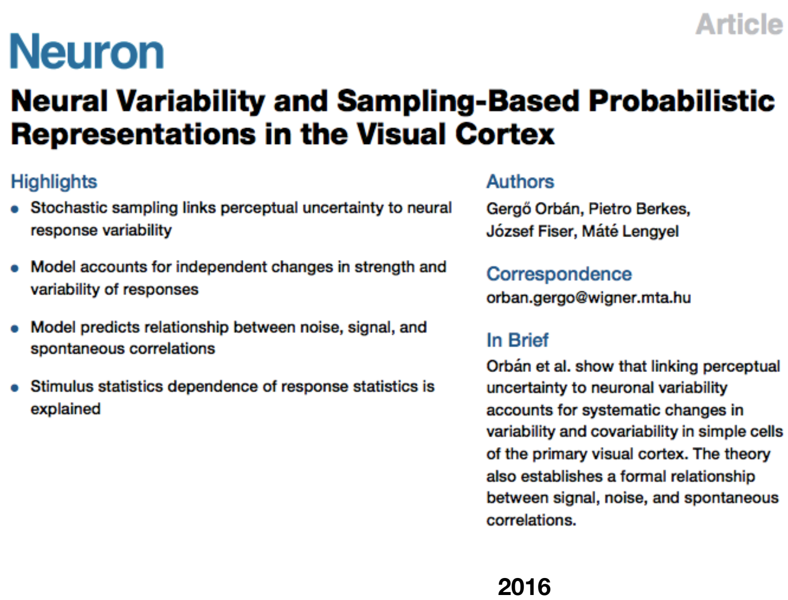

Neuron

Article

Neural Variability and Sampling-Based Probabilistic Representations in the Visual Cortex

Highlights

- Stochastic sampling links perceptual uncertainty to neural response variability
- Model accounts for independent changes in strength and variability of responses
- Model predicts relationship between noise, signal, and spontaneous correlations
- Stimulus statistics dependence of response statistics is explained

Authors

Gergő Orbán, Pietro Berkes, József Fiser, Máté Lengyel

Correspondence

orban.gergo@wigner.mta.hu

In Brief

Orbán et al. show that linking perceptual uncertainty to neuronal variability accounts for systematic changes in variability and covariability in simple cells of the primary visual cortex. The theory also establishes a formal relationship between signal, noise, and spontaneous correlations.

2016

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 approximations

On type of variational approximation: predictive coding.

Priors as top-down inputs : Predictive Coding

- perceptual inference as an iterative matching proces 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

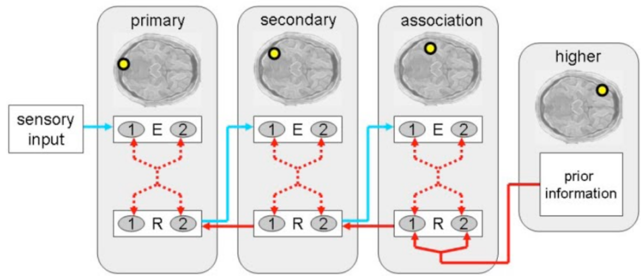

Predictive Coding: Neural Implementation of Bayesian Inference

- Learning involves making the predictions more and more similar to the input: minimizing the prediction error.

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

prior p(h)

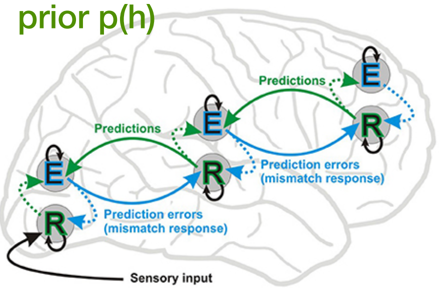

input p(e|h)

Predictive Coding: Neural Implementation of Bayesian Inference

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

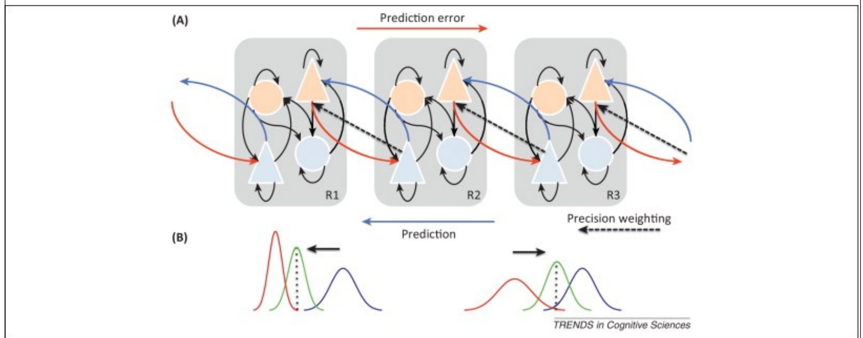

sensory signals: high (left) or low precision (right)
prior
posterior

Evidence for Predictive Coding

15164–15169 | PNAS | November 12, 2002 | vol. 99 | no. 23

Shape perception reduces activity in human primary visual cortex

Scott O. Murray*†, Daniel Kersten‡, Bruno A. Olshausen*§, Paul Schrater‡¶, and David L. Woods‖**

*Center for Neuroscience, §Department of Psychology, and ‖Department of Neurology, University of California, Davis, CA 95616; Departments of ‡Psychology and ¶Computer Science and Engineering, University of Minnesota, Minneapolis, MN 55455; and **Neurology Service (127E), Department of Veterans Affairs Northern California Health Care System, 150 Muir Road, Martinez, CA 94553

Communicated by David Mumford, Brown University, Providence, RI, September 24, 2002 (received for review April 25, 2002)

Visual perception involves the grouping of individual elements into coherent patterns that reduce the descriptive complexity of a visual scene. The physiological basis of this perceptual simplification remains poorly understood. We used functional MRI to measure activity in a higher object processing area, the lateral occipital complex, and in primary visual cortex in response to visual elements that were either grouped into objects or randomly arranged. We observed significant activity increases in the lateral occipital complex and concurrent reductions of activity in primary visual cortex when elements formed coherent shapes, suggesting that activity in early visual areas is reduced as a result of grouping processes performed in higher areas. These findings are consistent with predictive coding models of vision that postulate that inferences of high-level areas are subtracted from incoming sensory information in lower areas through cortical feedback.

response to images of objects versus scrambled versions of the same images and textures (9, 10). More recent studies have shown that this area increases in activity whenever individual features are grouped into an object or a coherent scene (11). Thus, the LOC may subserve high-level grouping of low-level image features. In the present study, we examined the effect of perceived shape on activity in V1 and in the LOC in a series of functional MRI experiments where visual elements were either perceived as coherent shapes or as random elements. We observed reduced activity in V1 and increased activity in the LOC when elements were grouped into coherent shapes, consistent with the view that higher visual areas "explain away" activity in lower areas through feedback processes.††

Materials and Methods

Experiment 1. Drawings were presented of (*i*) random lines, (*ii*)

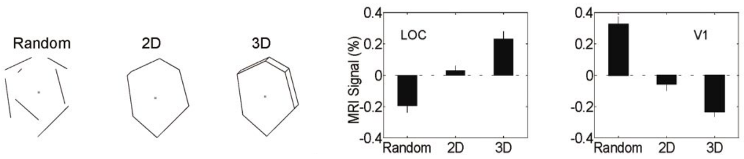

2017

Available online at www.sciencedirect.com
ScienceDirect
Current Opinion in Neurobiology

With or without you: predictive coding and Bayesian inference in the brain

Laurence Aitchison[1] and Máté Lengyel[1,2]

Two theoretical ideas have emerged recently with the ambition to provide a unifying functional explanation of neural population coding and dynamics: predictive coding and Bayesian inference. Here, we describe the two theories and their combination into a single framework: Bayesian predictive coding. We clarify how the two theories can be distinguished, despite sharing core computational concepts and addressing an overlapping set of empirical phenomena. We argue that predictive coding is an algorithmic/representational motif that can serve several different computational goals of which Bayesian inference is but one. Conversely, while Bayesian inference can utilize predictive coding, it can also be realized by a variety of other representations. We critically evaluate the experimental evidence supporting Bayesian predictive coding and discuss how to test it more directly.

Addresses

the input but must be *inferred* from context [6]. Overall, there is much evidence that perception and, correspondingly, neural responses in sensory cortical areas are as influenced by predictions and expectations about stimuli as by the actual stimuli themselves [7,8*]. Indeed, while ascending feed-forward connections convey stimulus-related information [9], long-range horizontal and feedback connections within and between different cortical areas provide a natural anatomical substrate for conveying such 'contextual' effects. The principles for how these contextual signals are computed, integrated with sensory information and represented in neural activities have been formalised in two different, though closely related theoretical frameworks: predictive coding and Bayesian inference.

Predictive coding

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.

• some confusion about the claims -- what exactly makes a neural model "Bayesian"

• Neural implementation largely unknown.

• Looking at update of priors / expectations for simple features (motion, speed) might be a good way to start.