‘Bayesian’ theories of perception, cognition and mental illness

CCN Lecture 14

Peggy Seriès,
IML, University of Edinburgh
pseries@inf.ed.ac.uk

This series of lectures

1) **Do people behave as Bayesian Observers?**

   a - Evidence from multi-sensory integration (last time)

   → lab 4

   b - What priors does the brain use? (today)

2) **A new way to understand Mental Illness?**

   (lecture 15)

1) **What does this tell us about the Brain?**

   (lecture 16)

   Controversies and possible implementation ideas

1) Do People behave as Bayesian Observers?

a) Do brains take into account measurement uncertainty when combining different (simultaneous) information? Combine different sources optimally?
b) Do brains form a representation of the past statistics of the environment (priors) and combine it optimally with current information?

A Bayesian theory of the Brain: Priors

- How is the brain making use of previous knowledge? what priors?
- Prediction 1: the more uncertain the data, the more prior information should influence the interpretation.
- Prediction 2: The priors should reflect the statistics of the sensory world (on which time-scale?).

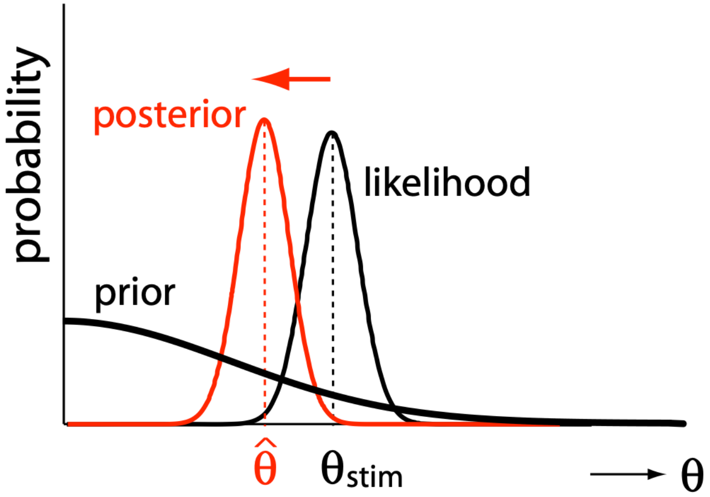


there is no lake in the photo.

Long-term "structural" priors

Visual illusions : insight into what sort of assumptions the visual system makes.

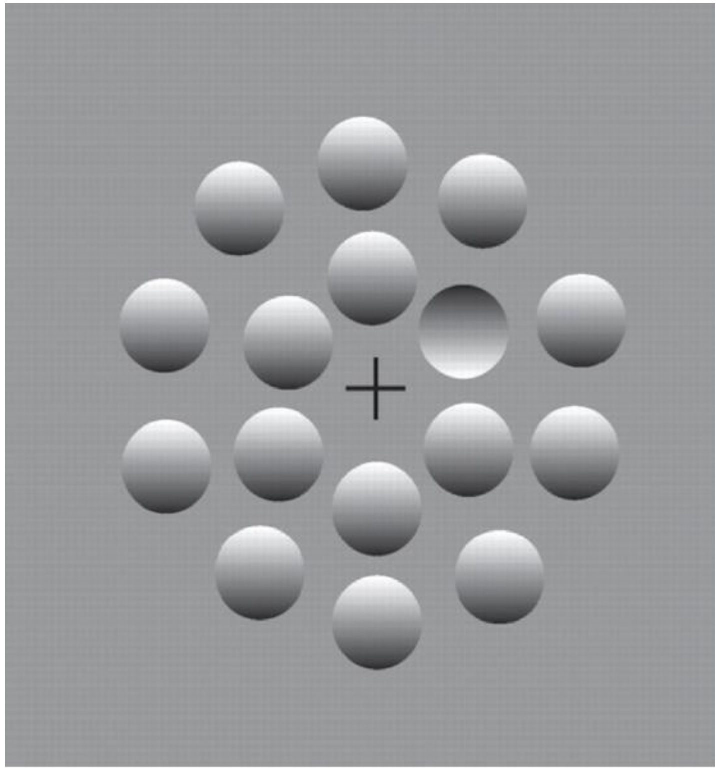

- Light comes from above
- Cardinal orientations are more frequent [Gershick et al 2011]
- smoothness [Geisler et al 2001]
- symmetry [Knill 2007]
- Objects don't move or only slowly [Weiss et al 2001; stocker & Simoncelli 2006]

... recently formalized in Bayesian terms

[T. Adelson, E. Simoncelli, O. Schwartz, Y. Weiss]

Interpreting motion : A Prior on Low Speeds (1)

- "Aperture Problem" : Motion shown in an aperture is fundamentally ambiguous; it can be interpreted in an infinite number of ways
- which one is chosen? why?

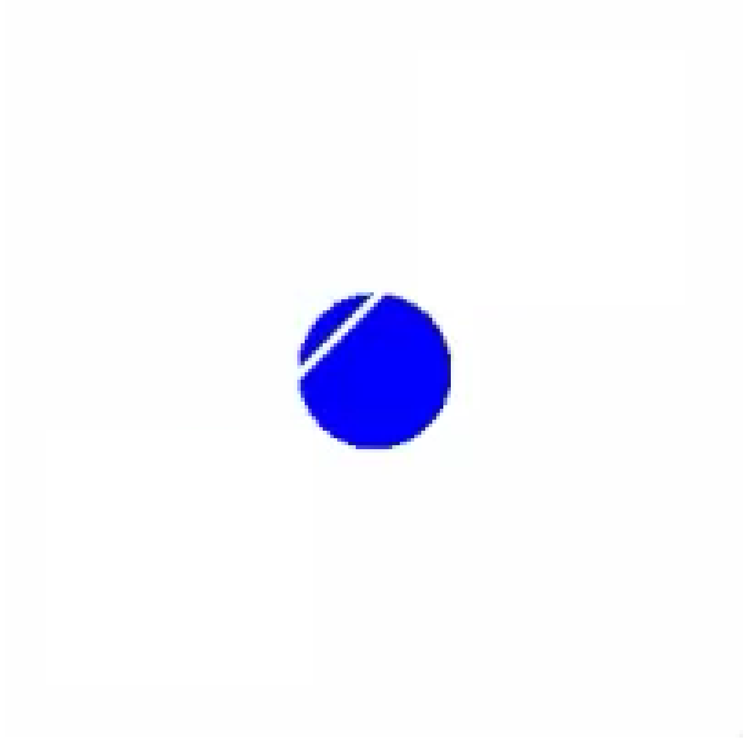

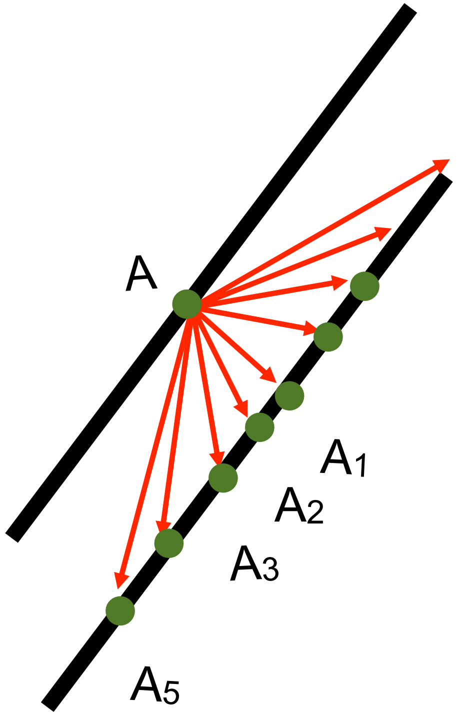

Interpreting motion : A Prior on Low Speeds (1)

- "Aperture Problem" : Motion shown in an aperture is fundamentally ambiguous; it can be interpreted in an infinite number of ways
- which one is chosen? why?

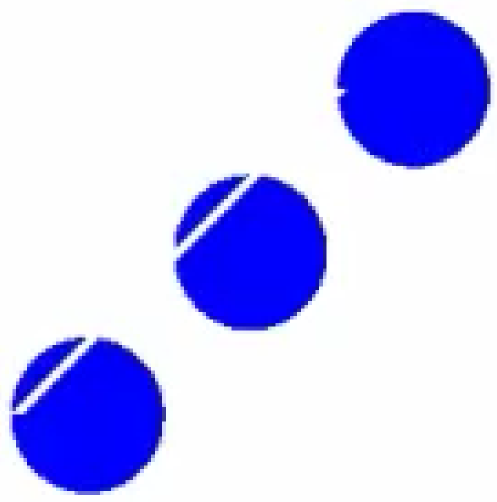

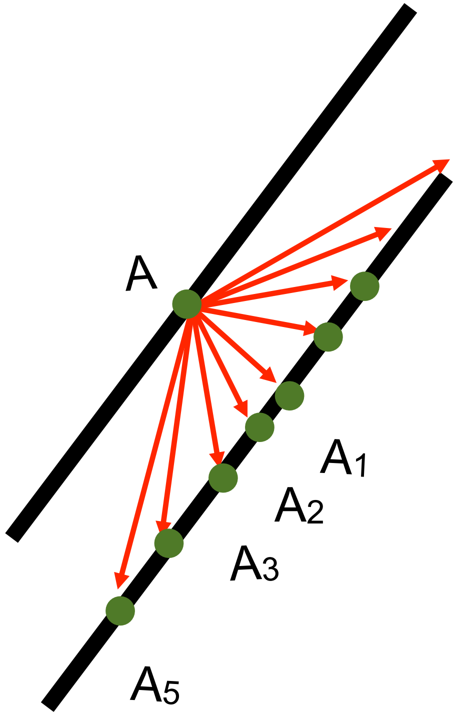

Interpreting motion : A Prior on Low Speeds (2)

- Hypothesis: humans tend to favour slower motions
- Use a (gaussian) prior on low speeds (centred at 0).
- Explain great variety of data -- elegant unifying explanation

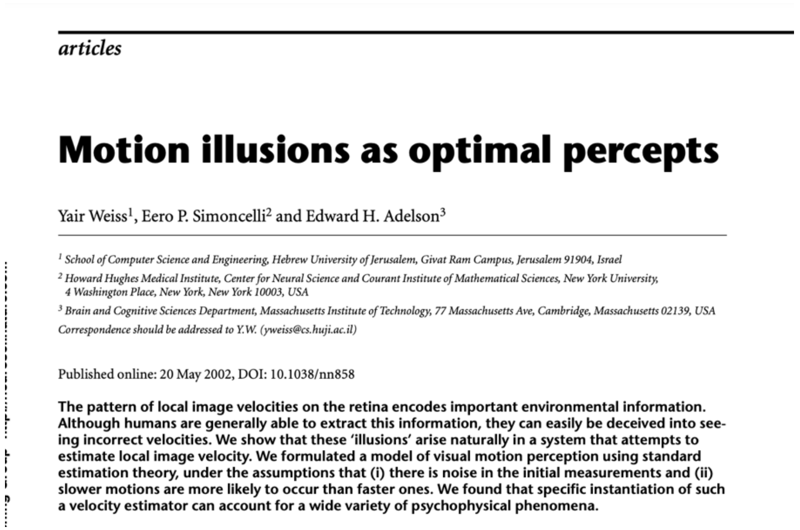

articles

Motion illusions as optimal percepts

Yair Weiss[1], Eero P. Simoncelli[2] and Edward H. Adelson[3]

[1] School of Computer Science and Engineering, Hebrew University of Jerusalem, Givat Ram Campus, Jerusalem 91904, Israel
[2] Howard Hughes Medical Institute, Center for Neural Science and Courant Institute of Mathematical Sciences, New York University, 4 Washington Place, New York, New York 10003, USA
[3] Brain and Cognitive Sciences Department, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, Massachusetts 02139, USA
Correspondence should be addressed to Y.W. (yweiss@cs.huji.ac.il)

Published online: 20 May 2002, DOI: 10.1038/nn858

The pattern of local image velocities on the retina encodes important environmental information. Although humans are generally able to extract this information, they can easily be deceived into seeing incorrect velocities. We show that these 'illusions' arise naturally in a system that attempts to estimate local image velocity. We formulated a model of visual motion perception using standard estimation theory, under the assumptions that (i) there is noise in the initial measurements and (ii) slower motions are more likely to occur than faster ones. We found that specific instantiation of such a velocity estimator can account for a wide variety of psychophysical phenomena.

Weiss, Adelson & Simoncelli, *Nat Neuro*, 2002

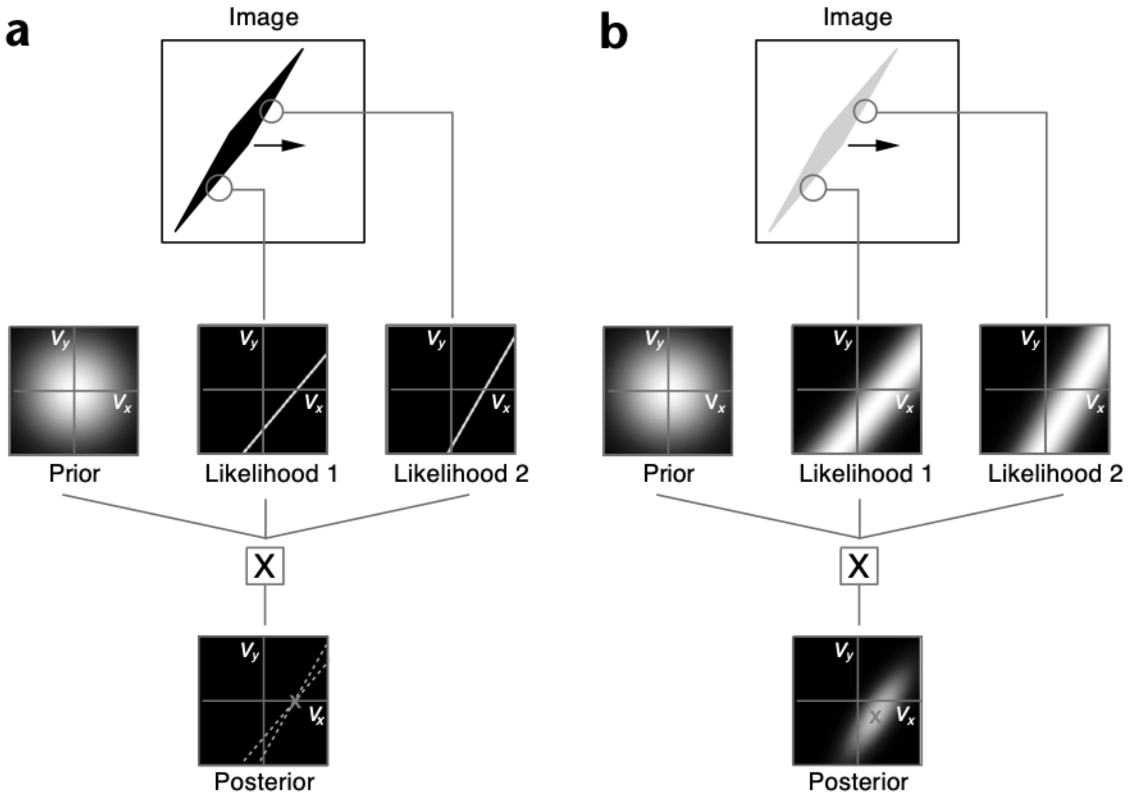

Interpreting motion : A Prior on Low Speeds (1)

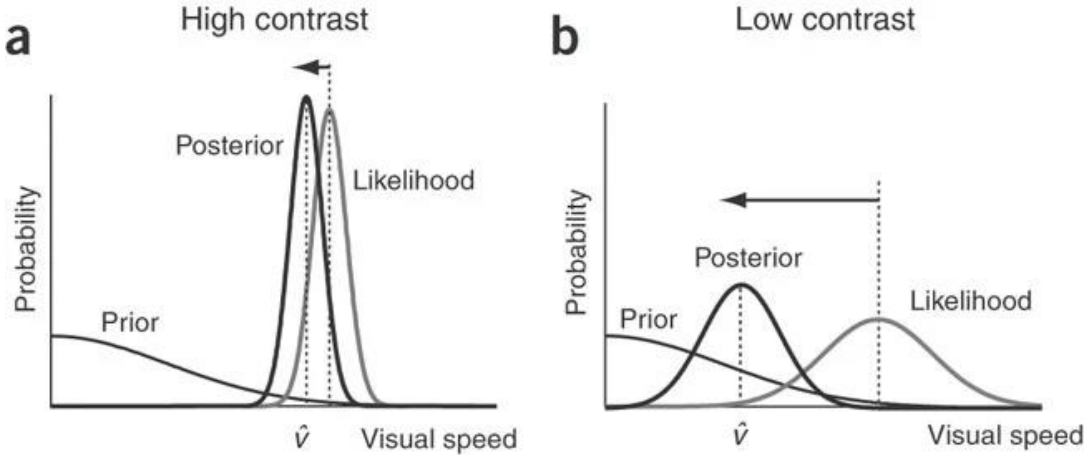

NATURE | VOL 392 | 2 APRIL 1998

Speed perception fogs up as visibility drops

Many horrendous vehicle accidents occur in foggy weather. Drivers know they should slow down because fog reduces visibility, but many still drive too quickly[1]. The 'blame' for many such accidents may be due to a perceptual quirk: it appears that drivers think they are driving far more slowly than they actually are in foggy conditions, and therefore increase their speed.

We used a virtual-environment driving

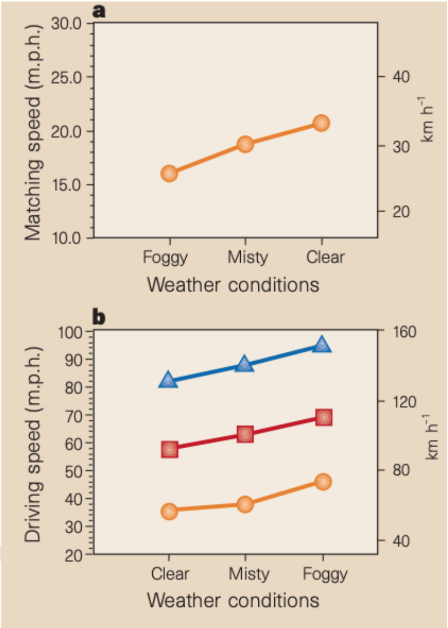

- The brain expects speed to be 0 or slow.
- Prior on low speed will influence the estimation of speed, mostly at very low contrast.
- This is proposed to be the explanation why drivers might misestimate their speed in the fog.

Can we measure people's prior experimentally?

- We can reverse engineer the shape of the prior from people's perceptual data
- Speed discrimination task at different contrast levels -- measure both bias and variability + fit Bayesian model --> recover speed prior and likelihood in individuals
- reveals inter-individual variability in the prior that different people use
- Speed prior not Gaussian

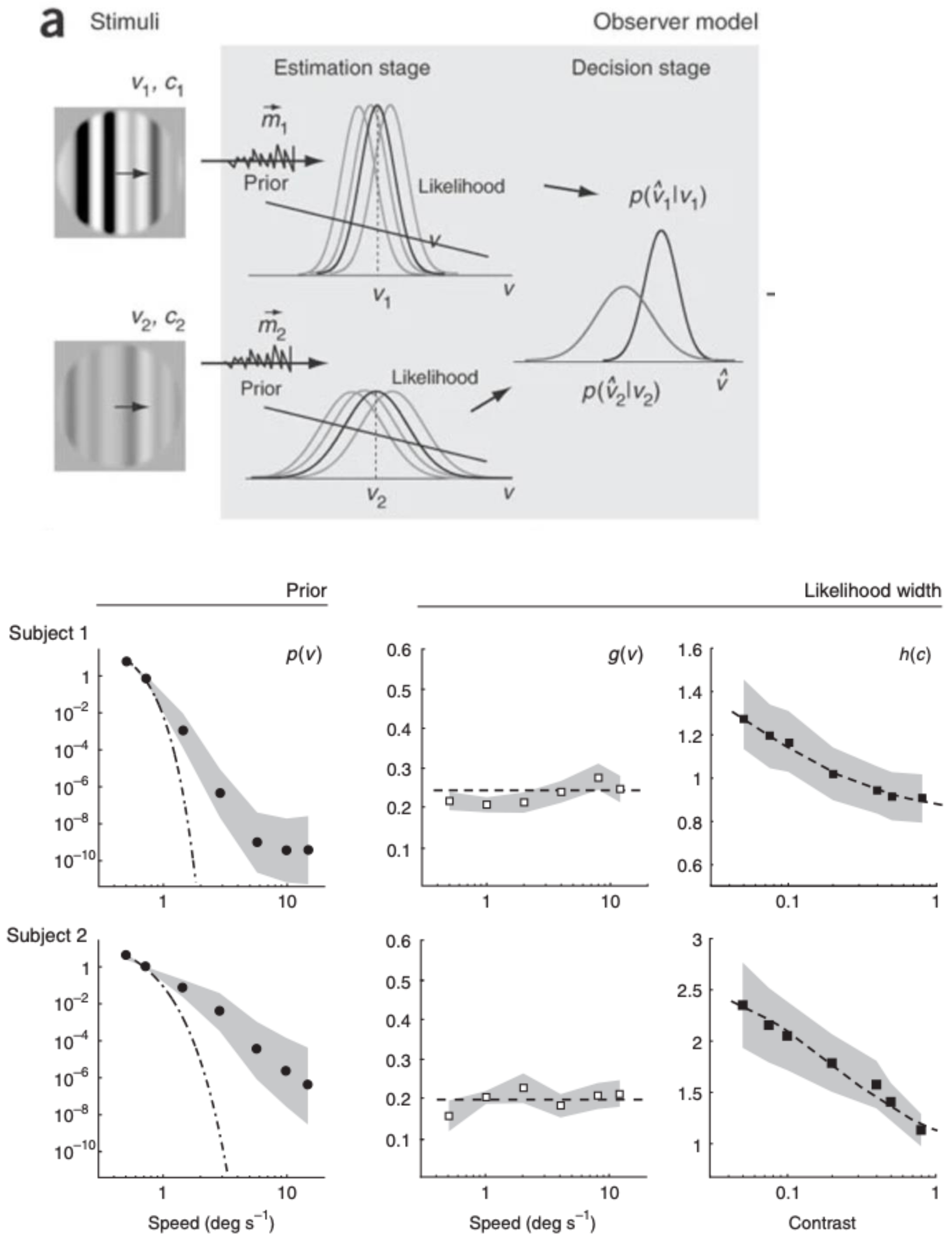

Stocker & Simoncelli, Nat Neuro, 2006

A way to discover people's beliefs?

- Reverse engineering Bayesian models as a way to discover people's priors/beliefs/expectations and measure then quantitatively

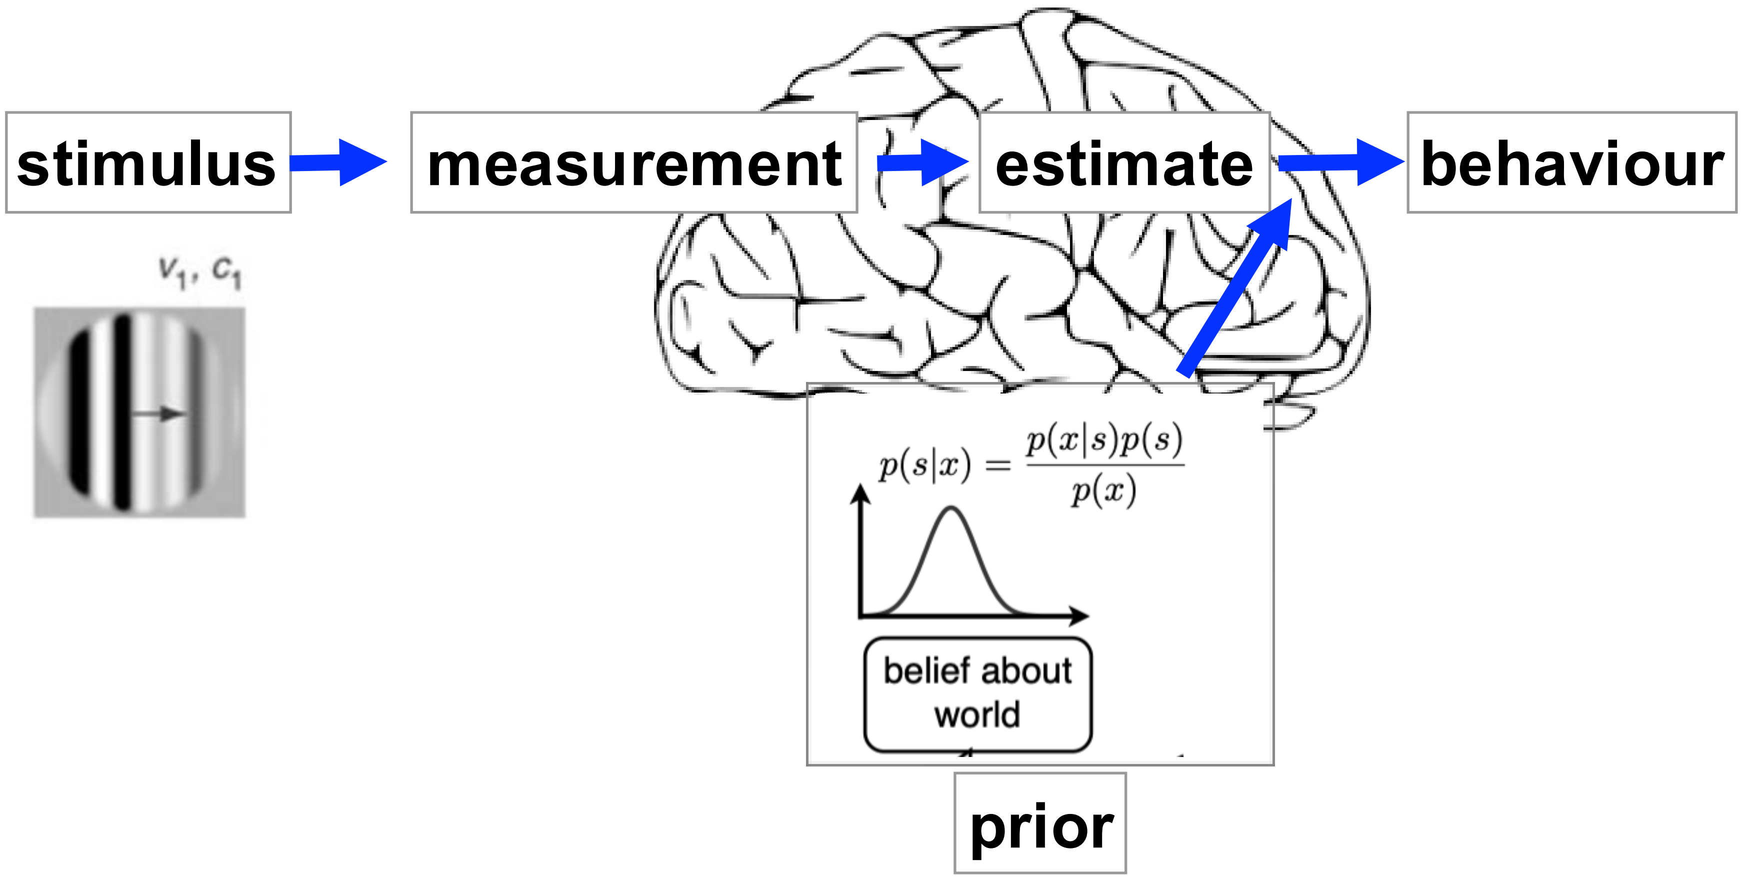

A way to discover people’s beliefs?

- Reverse engineering Bayesian models as a way to discover people’s priors/beliefs/expectations and measure then quantitatively

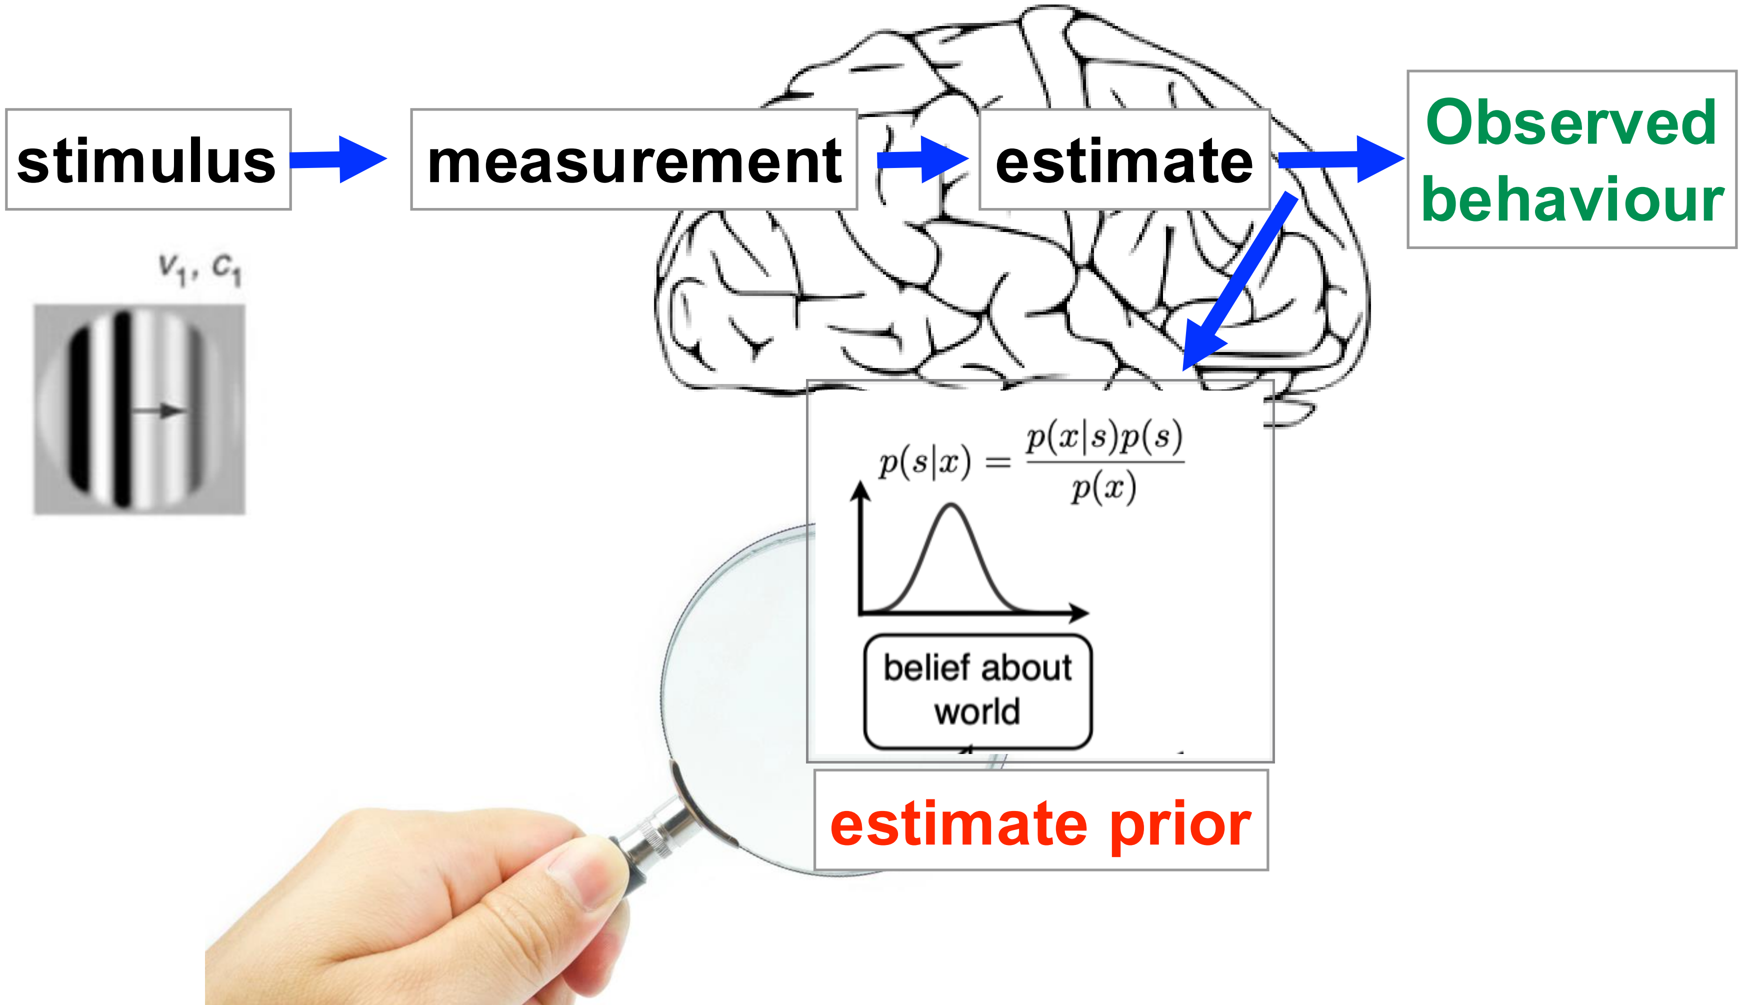

Do such priors correspond to the environment statistics?

- Difficult to assess for speed prior, but easier for orientation.

- Orientation judgments are more accurate at cardinal (horizontal and vertical) orientations.
- And biased toward cardinal orientations.

- Prior towards cardinal orientation, as estimated through reverse engineering behaviour, match orientation distribution measured in photographs.

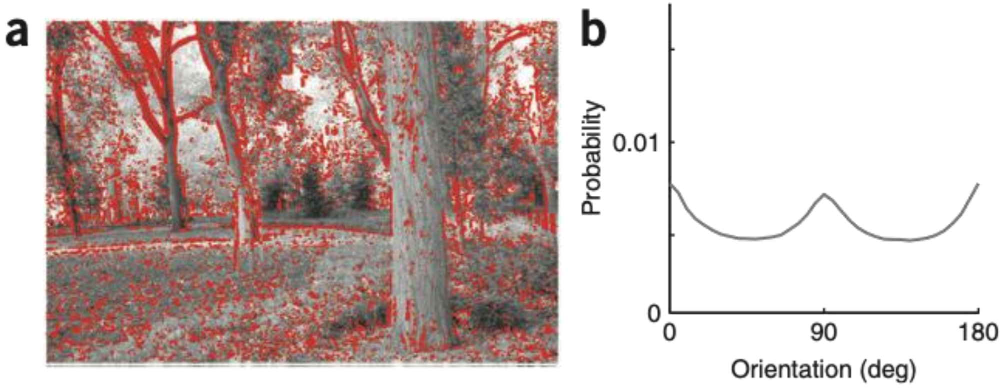

Cardinal rules: visual orientation perception reflects knowledge of environmental statistics

Ahna R Girshick[1,2], Michael S Landy[1,2] & Eero P Simoncelli[1–4]

Humans are good at performing visual tasks, but experimental measurements have revealed substantial biases in the perception of basic visual attributes. An appealing hypothesis is that these biases arise through a process of statistical inference, in which information from noisy measurements is fused with a probabilistic model of the environment. However, such inference is optimal only if the observer's internal model matches the environment. We found this to be the case. We measured performance in an orientation-estimation task and found that orientation judgments were more accurate at cardinal (horizontal and vertical) orientations. Judgments made under conditions of uncertainty were strongly biased toward cardinal orientations. We estimated observers' internal models for orientation and found that they matched the local orientation distribution measured in photographs. In addition, we determined how a neural population could embed probabilistic information responsible for such biases.

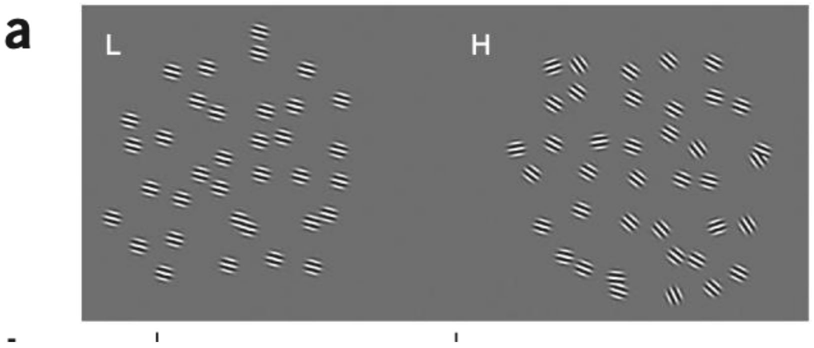

is L stimulus CW or CCW compared to H?

VOLUME 14 | NUMBER 7 | JULY 2011 NATURE NEUROSCIENCE

Do people form new priors for everything? How fast?

[Chalk, Seitz and Seriès, JOV 2010]

- On each trial, participants were presented with either a low contrast random dot motion stimulus (100% coherence) or a blank screen.
- Participants reported direction of motion (estimation), before reporting whether a stimulus was present (detection).

Behavioural Task

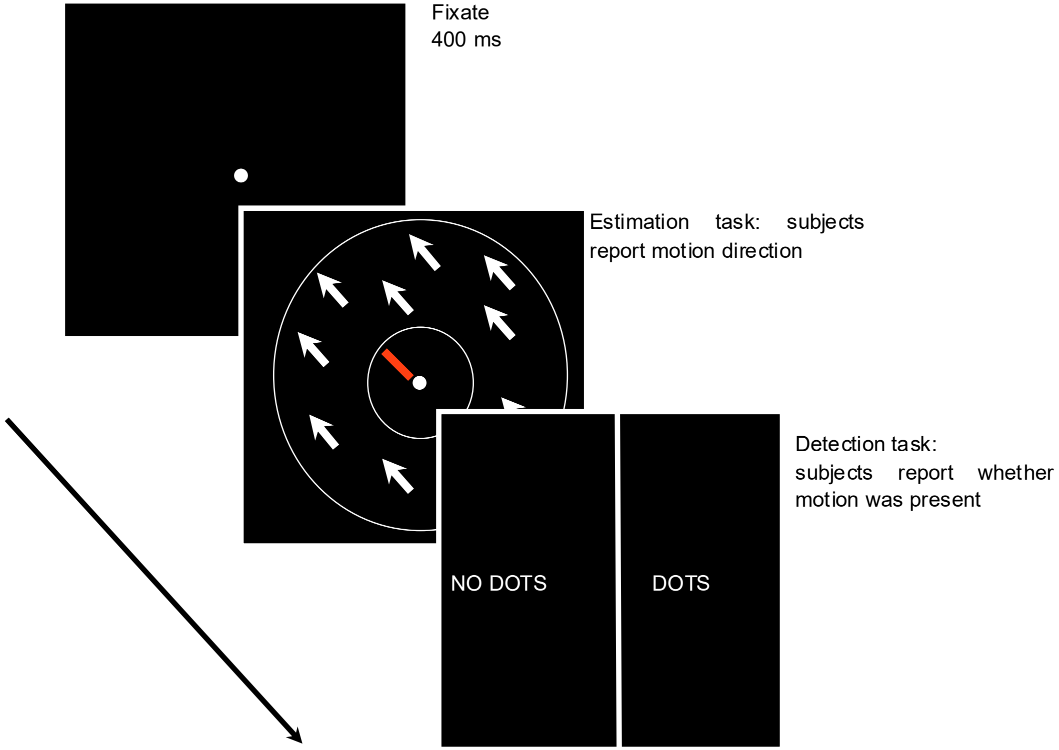

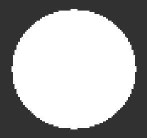

- Two motion directions were presented in a larger number of trials than other directions.

Stimulus distribution

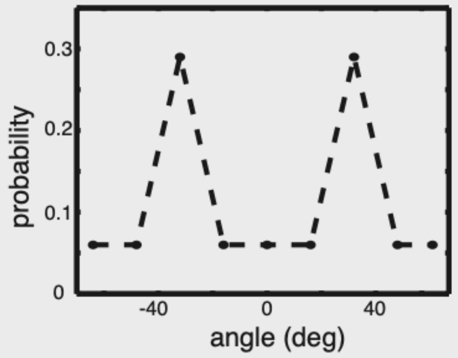

Questions

1. Are participants going to learn implicitly which directions are most likely to be presented?

2. How would these learned expectations bias their perception of subsequently presented motion stimuli?

Result 1/3: Detection is better and faster for the expected directions

- Detection performance was best for most frequently presented directions
- Reaction times were shorter
- Similar to the effects of selective attention (Posner et al. 1980) - suggesting that subjects were attending to expected directions.
- Knowledge about the statistics of the stimulus was however not conscious.

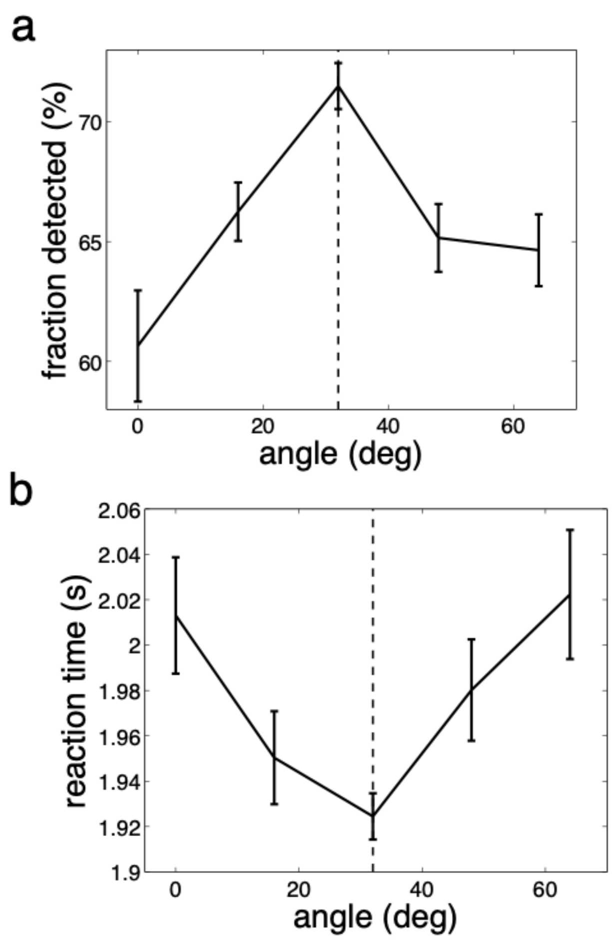

Result 2/3: Participants 'hallucinate' motion in expected directions

- On trials where no stimulus was presented, but where participants reported seeing a stimulus, they were strongly biased to report motion in the two most frequently presented directions.
- This effect was fast to develop, occurring in less than 200 trials / few minutes.

Distribution of estimates when no stimulus displayed

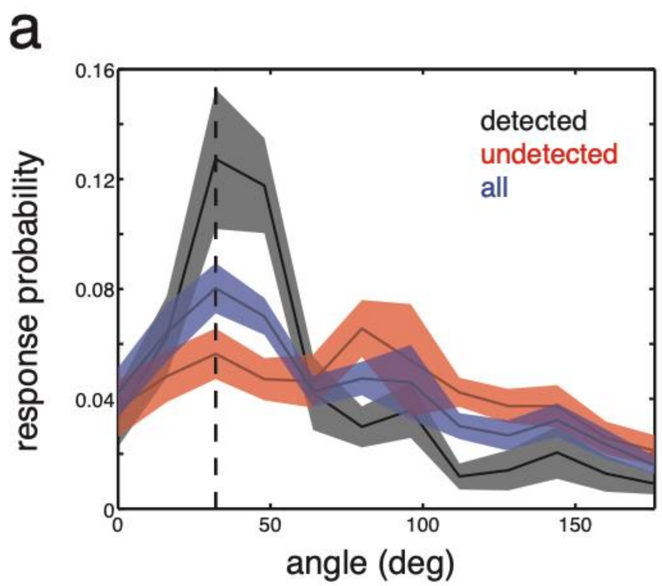

Result 3/3: Expectations bias perception of motion direction

[Chalk, Seitz, Seriès, JOV 2010]

Estimation bias

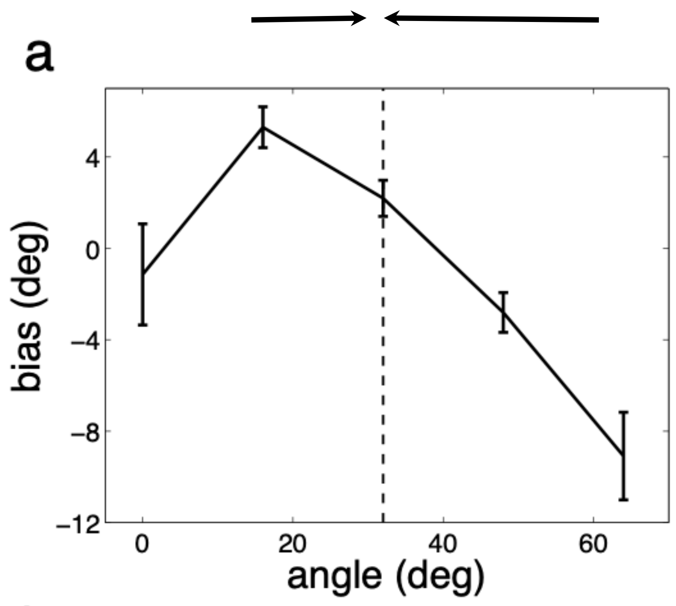

- Estimates of motion direction were biased towards most frequently presented directions:

  subjects perceive motion direction to be more similar to expected direction than it really is.

Modelling the estimation biases

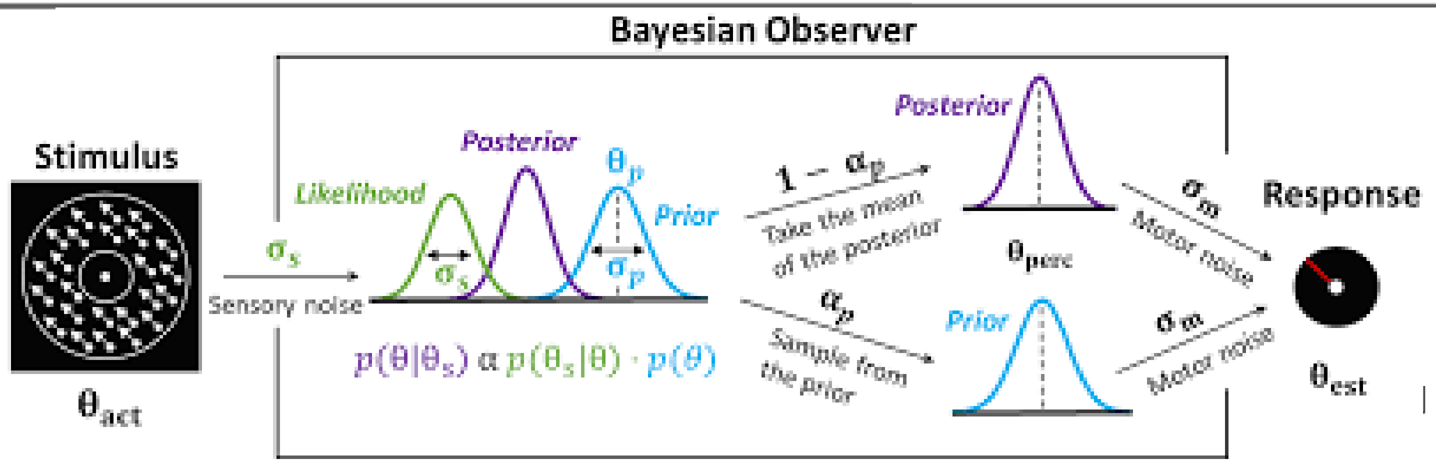

- Bayesian Modeling: subjects learn an expected distribution of the stimuli (prior) and combine it with sensory evidence
- Extract prior for each individual.

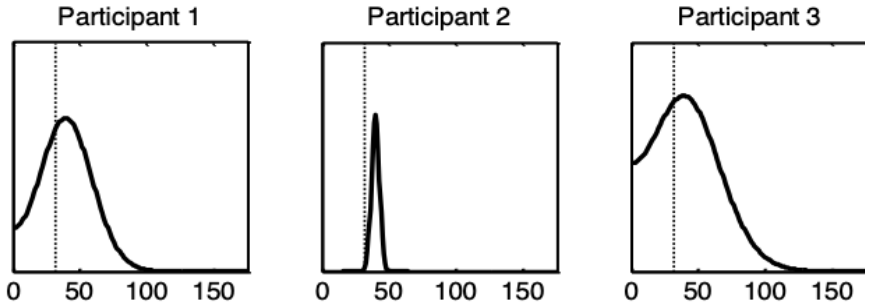

- Model Comparison: Bayesian model describes the data better than response strategy models. Individual priors look like approximation of stimulus

Model comparison

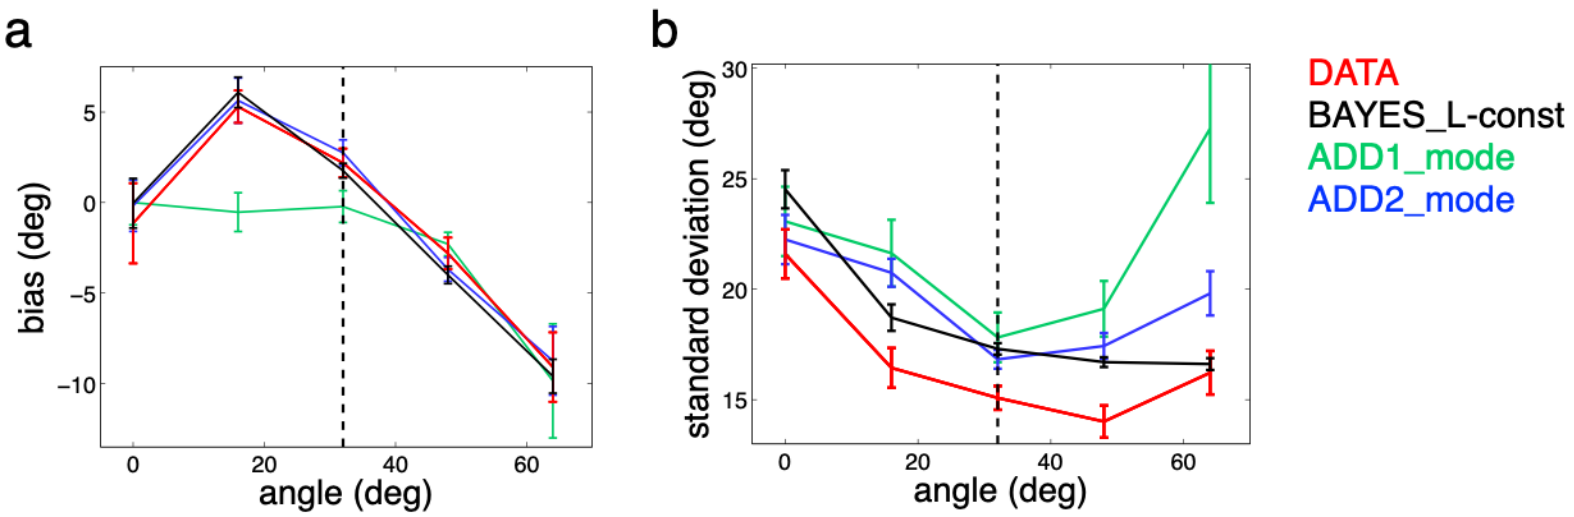

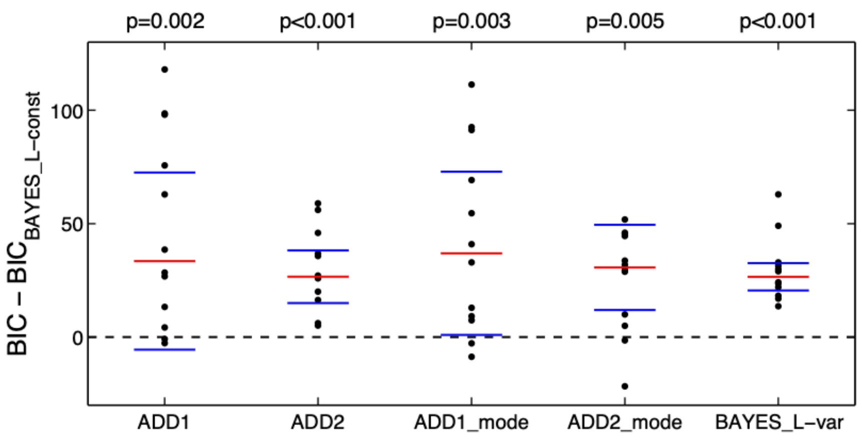

$$BIC = -2 \cdot \ln(L) + k \cdot \ln(n)$$

- Bayesian model describes the data better than response strategy models.

[Chalk, Seitz and Seriès, JOV 2010]

Conclusions: Fast learning of a Direction Prior

- Participants rapidly learn multimodal stimulus expectations (< 200 trials).
- These expectations bias their perception of simple motion stimuli, causing them to 'hallucinate' motion in the expected direction, and perceive motion stimuli as closer to the expected directions than they actually are.
- The biases we observed can be explained assuming that participants combine a 'learned prior' about the stimulus statistics with their sensory evidence in a probabilistically optimal way.
- Open questions: specificity of prior, time scale, neural implementation - substrate of expectation?
- in particular: can one learn any prior like this ? or are some priors fixed?

Are priors constantly updating? Even those supposedly corresponding to natural scene statistics? (1)

Journal of Comparative and Physiological Psychology
1970, Vol. 73, No. 3, 407–411

ATTACHED-SHADOW ORIENTATION PERCEIVED AS DEPTH BY CHICKENS REARED IN AN ENVIRONMENT ILLUMINATED FROM BELOW[1]

WAYNE HERSHBERGER[2]
Northern Illinois University

The depth perception of chickens reared in cages illuminated from below was tested using photographed dents with shadow orientation the relevant cue. As do humans, the chickens assumed an overhead source of illumination seeing dents shaded below and above as convex and concave, respectively. There appears to be an innate perceptual parameter corresponding to an "overhead source of illumination" in terms of which orientation of attached shadow is interpreted as depth.

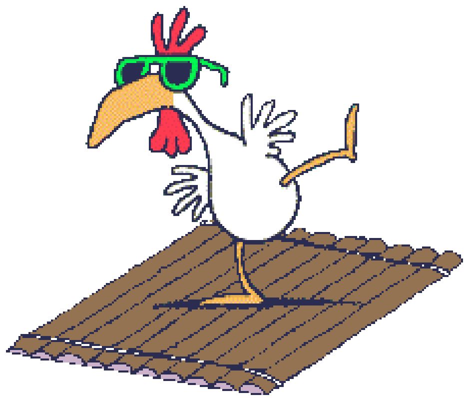

Experience can change the 'light-from-above' prior

Wendy J Adams[1], Erich W Graf[1] & Marc O Ernst[2]

To interpret complex and ambiguous input, the human visual system uses prior knowledge or assumptions about the world. We show that the 'light-from-above' prior, used to extract information about shape from shading is modified in response to active experience with the scene. The resultant adaptation is not specific to the learned scene but generalizes to a different task, demonstrating that priors are constantly adapted by interactive experience with the environment.

The circular patches in **Figure 1a** have competing interpretations. However, patches that are brighter at the top are generally seen as convex and the others as concave, consistent with an assumption of light from above[1,2]. The Bayesian approach has successfully described performance in many perceptual tasks where stimulus information is combined with prior assumptions[3–5]. However, whether visual priors are hard-wired or learned in response to environmental statistics is not known[6]. We investigate the adaptability of the 'light-from-above' prior by adding shape information via haptic (active touch) feedback.

a

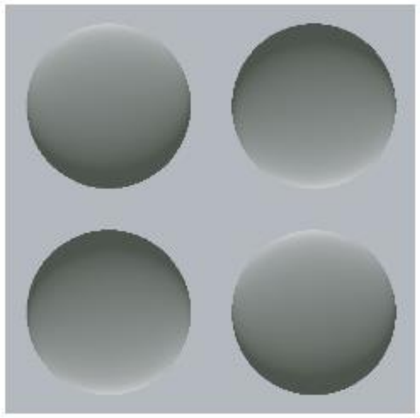

c

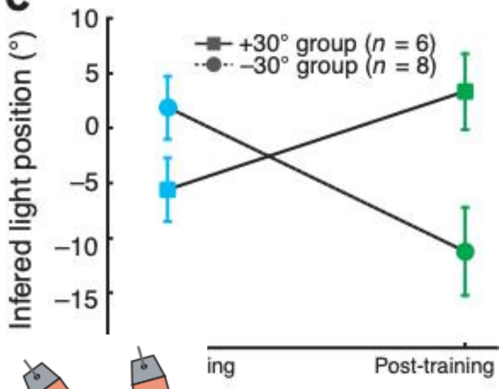

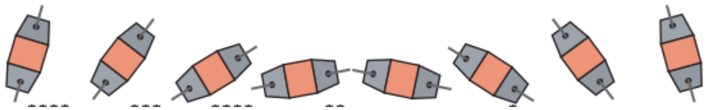

not in chickens, innate
Hershberger, W. (1970)

[Adams, Graf and Ernst Nature Neuroscience 2004]

Are priors constantly updating? Even those supposedly corresponding to natural scene statistics? (2)

The slow speed prior can be updated in a few sessions, just through exposure.
[Sotiropoulos, Seitz & Seriès (2011), Current Biology]

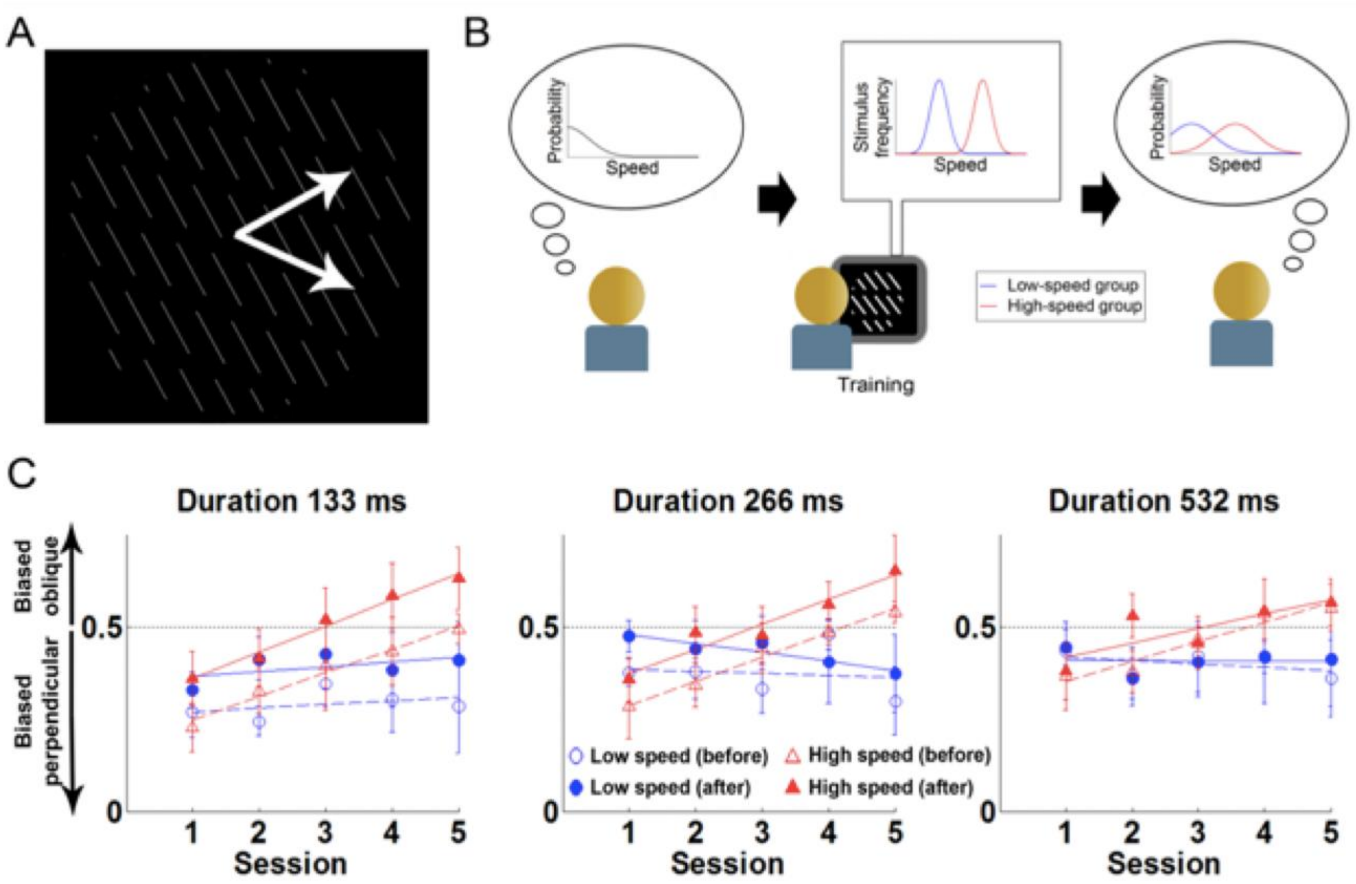

Extensions and open questions

- **What are the limits of prior learning? complexity?**
  [Gekas et al 2014; Acerbi et al 2014 ..]
- **How many priors can one learn simultaneously?**
  [Gekas et al 2014]
- **Are priors specific to learned conditions? stimulus? task? experimental context?**
  [Adams & Kerrigan 2013, Mamassian, Orban & Lengyel; Roach et al 2017]
- **Time scales of learning? unlearning?**
  [Lowenstein, Gekas et al 2015]
- **Heuristics or true Bayesian inference?**
  [Ravi & Loewenstein, Karvelis et al]

Perception as a "controlled hallucination"

- The brain uses an internal model/expectations to reconstruct the source of the input.
- Brain is **better** at processing data that is conform to expectations.
- Brain is **biased** towards perceiving the world as being more similar to its expectations that it really is.

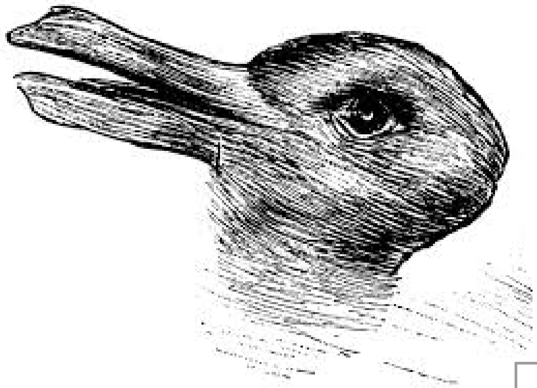

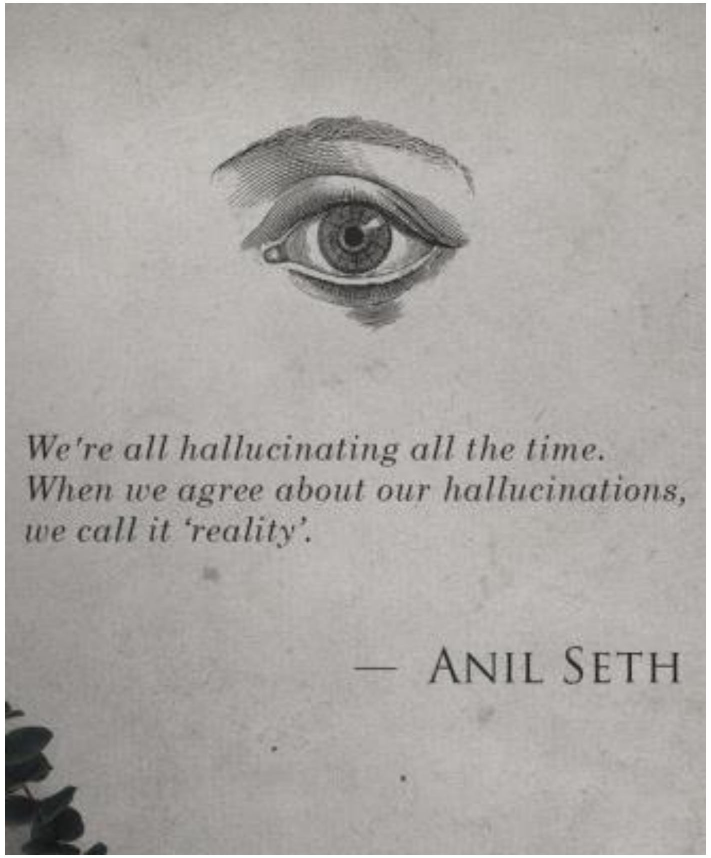

Approach and methods also extend to other domain of cognition

Example of a Cognitive Prior: Optimism

- It has been shown for example that trait Optimism could be formalised in terms of a prior belief about the probability of future reward.
- In doubt, optimistic people (according to the LOT-R questionnaire) tend to believe an uncertain target will be associated with reward, more so than pessimistic people.

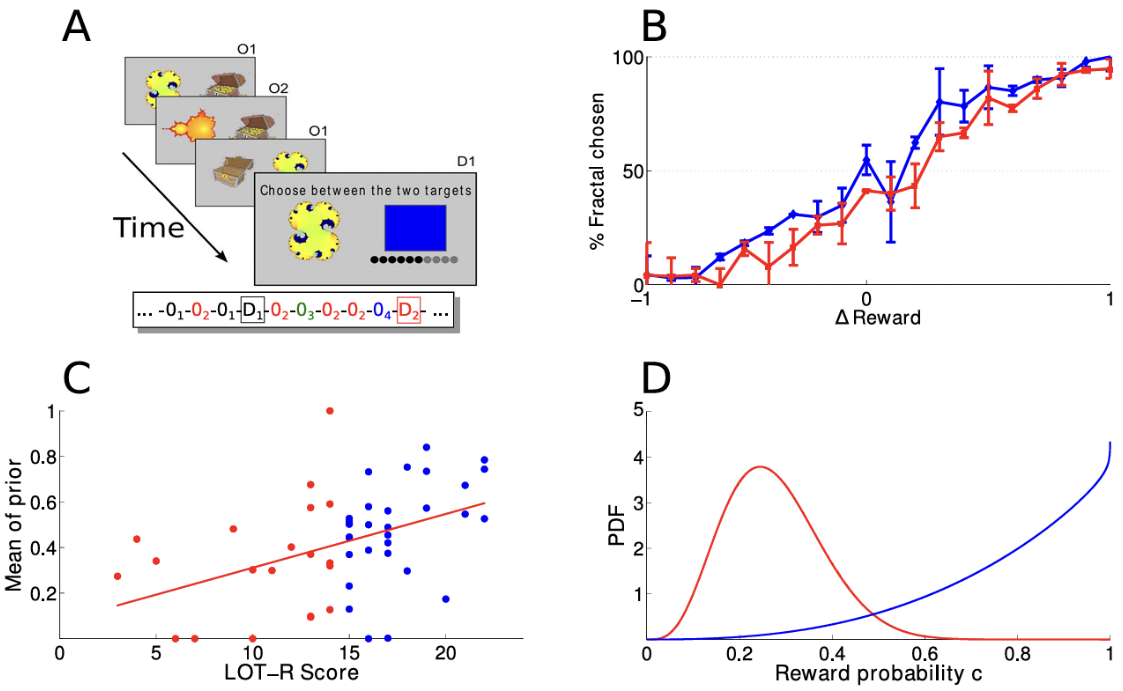

OPEN ACCESS Freely available online
PLOS | COMPUTATIONAL BIOLOGY

Optimism as a Prior Belief about the Probability of Future Reward

Aistis Stankevicius[1,§], Quentin J. M. Huys[2,3,4,§], Aditi Kalra[1], Peggy Seriès[1,*]

1 Institute for Adaptive and Neural Computation, University of Edinburgh, Edinburgh, United Kingdom, 2 Translational Neuromodeling Unit, Institute for Biomedical Engineering, University of Zürich and ETH Zürich, Zürich, Switzerland, 3 Department of Psychiatry, Psychotherapy and Psychosomatics, University Hospital of Psychiatry Zurich, Zürich, Switzerland, 4 Gatsby Computational Neuroscience Unit and Wellcome Trust Neuroimaging Centre, UCL, London, United Kingdom

Abstract

Optimists hold positive *a priori* beliefs about the future. In Bayesian statistical theory, *a priori* beliefs can be overcome by experience. However, optimistic beliefs can at times appear surprisingly resistant to evidence, suggesting that optimism might also influence how new information is selected and learned. Here, we use a novel Pavlovian conditioning task, embedded in a normative framework, to directly assess how trait optimism, as classically measured using self-report questionnaires, influences choices between visual targets, by learning about their association with reward progresses. We find that trait optimism relates to an *a priori* belief about the likelihood of rewards, but not losses, in our task. Critically, this positive belief behaves like a probabilistic prior, i.e. its influence reduces with increasing experience. Contrary to findings in the literature related to unrealistic optimism and self-beliefs, it does not appear to influence the iterative learning process directly.

Citation: Stankevicius A, Huys QJM, Kalra A, Seriès P (2014) Optimism as a Prior Belief about the Probability of Future Reward. PLoS Comput Biol 10(5): e1003605. doi:10.1371/journal.pcbi.1003605

Bayesian Brain - Summary

- Behavioural studies suggest that human observers behave in ways compatible with Bayesian inference
- Bayesian modelling offers a way to "reverse engineer" the internal models.
- Perception is a reconstruction or even a "controlled hallucination", based on the internal model of the brain
- The same tools can be applied for cognition more generally.
- A possible tool for understanding mental illness? A case where the internal models would be "abnormal"?

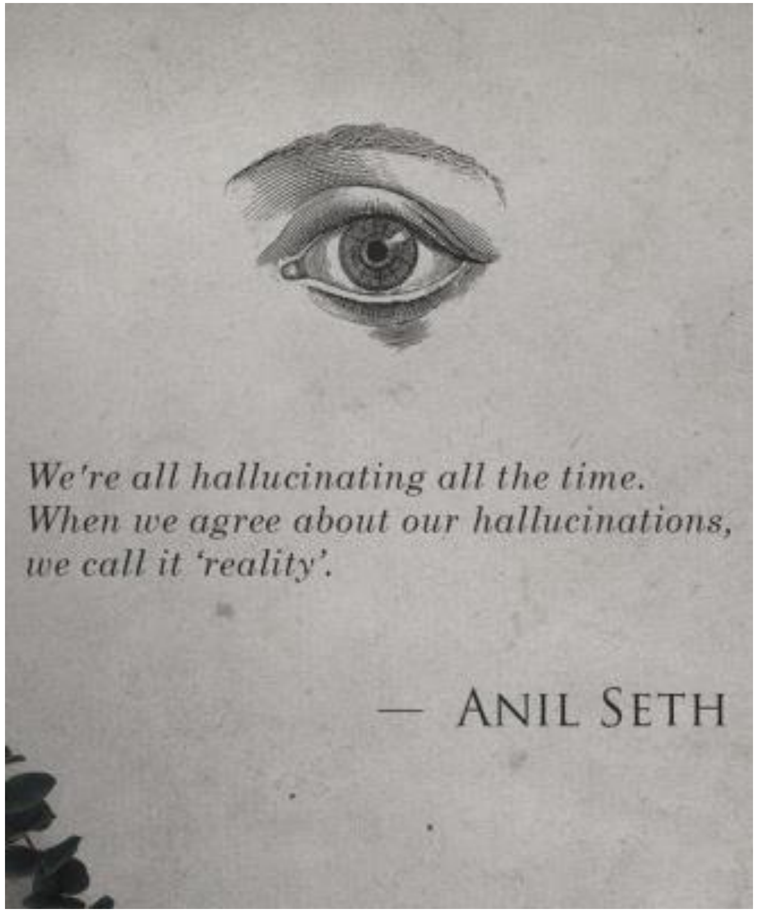

If you want to learn more

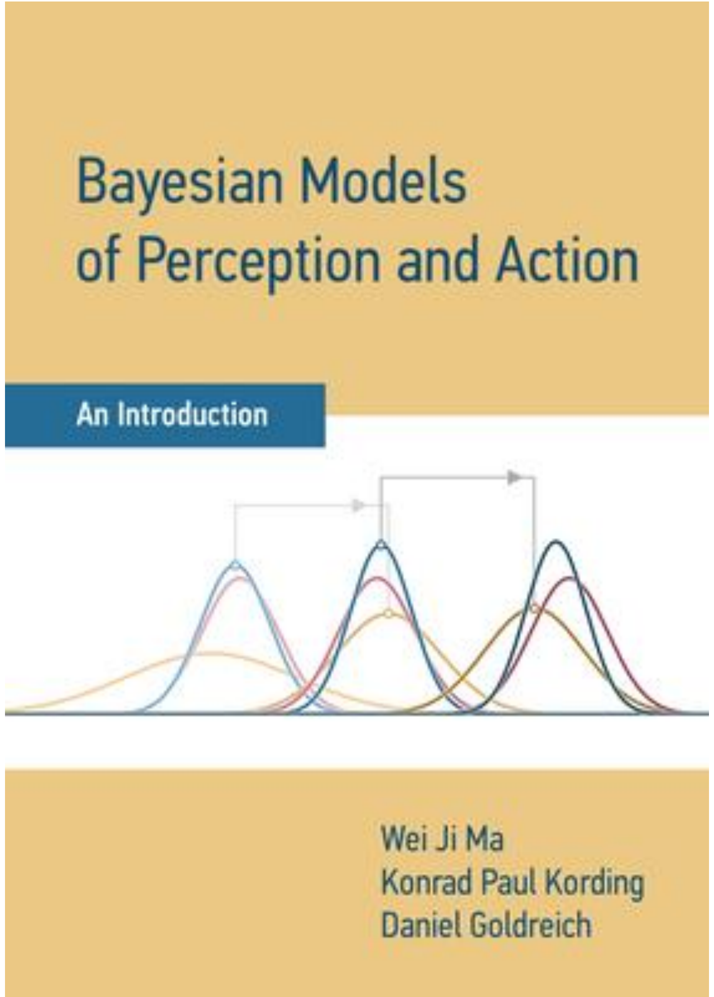

2023

- Learn all the steps of constructing and using a Bayesian model for studying perception or action
- Free pdf: https://www.cns.nyu.edu/malab/bayesianbook.html