Commentary on Rahnev & Denison (2018): [Suboptimality in Perceptual Decision Making](https://doi.org/10.1017/S0140525X18000936)
(<https://doi.org/10.1017/S0140525X18000936>)

Title: **The world is complex, not just noisy**

Author: Romain Brette

Sorbonne Universités, UPMC Univ Paris 06, INSERM, CNRS

Institut de la Vision, 17 rue Moreau, 75012 Paris, France

+33 1 53 46 25 36

romain.brette@inserm.fr

<http://romainbrette.fr>

Abstract

To deny that human perception is optimal is not to claim that it is suboptimal. Rahnev & Denison point out that optimality is often ill-defined. The fundamental issue is framing perception as a statistical inference problem. Outside the lab, the real perceptual challenge is to determine the lawful structure of the world, not variables of a predetermined statistical model.

Main text

Is human perception optimal? If we disagree, it would seem (as the title of the target article suggests) that we hold that human perception is suboptimal, perhaps because of our imperfect biological nature – we are only human. The authors, however, point out that optimality is “only well-defined in the context of a set of specific assumptions, rendering general statements about the optimality (or suboptimality) of human perceptual decisions meaningless”. If this is true, then the main issue is not so much whether humans are optimal or sub-optimal, as the scope of such claims.

Claims of optimality frame perception as a problem of statistical inference: sensory data are assumed to be produced by a statistical model that depends on hidden variables, and the perceptual problem is to estimate the value of those variables. There is generally a simple optimal solution given by Bayes’ formula.

The obvious conceptual issue is that variables are defined within a particular model, which captures the structure of the scene (objects and their relations), when a large part of the perceptual problem is precisely to capture that structure. We note for example that state-of-the-art computer vision algorithms (e.g. convolutional neural networks) excel at inferring the position of a cat in an image, but struggle at analyzing the structure of an image. For example, the question: “are there two identical objects in the image?” seems to pose a very challenging computational problem, even though this is a trivial problem for humans and animals (Ricci et al., 2018).

The only reason why such problems do not appear in accounts of optimal perception is that those accounts are built from results of constrained experiments, in which a few experimental variables are allowed to vary within a fixed structure. In other words, the focus

is on “known unknowns”: we do not know the value of the variables, but we know they exist and we know their probability distribution *a priori* (Rumsfeld, 2011). This is not generally the case in ecological settings, where the perceptual system has to deal with “unknown unknowns”. This makes the scope of optimality claims somewhat limited.

It should be stressed that building knowledge from observations is not generically a statistical inference problem. For example, Newtonian mechanics (and in fact, science in general) have not been derived by a process of statistical inference applied to the movements of bodies. To turn it into a statistical inference problem, one would first need to come with a few candidate models. This requires designing appropriate variables (e.g. acceleration of the center of mass) and postulate relations between them (acceleration equals gravity). Once this work has been done, what remains is relatively trivial.

In the same way, perception requires determining what constitutes an object, what are relevant object properties and what might be the relations between objects. Thus the optimality framework trivializes the problem of perception by focusing on statistical inference on a fixed model, leaving aside the most difficult questions, in particular object formation and scene analysis. The proposition of Rahnev & Denison does not seem to address this issue, as observer models are still cast within the Bayesian framework (Box 1), and the difficult questions appear to be hidden in point (1), the description of the generative model, which is a context-dependent model fixed by the scientist rather than resulting from the perceptual process itself.

This is not to deny that statistical inference can be a part of perceptual processes, but it constitutes only a part, arguably a small one. In this light, it seems difficult to make sense of broad claims such as “human perception is close to the Bayesian optimal” (Körding and Wolpert, 2006), given that perception as a whole cannot possibly be modeled by a Bayesian model. In the same way, the Bayesian brain hypothesis that “the world is represented by a conditional probability density function over the set of unknown variables” (Knill and Pouget, 2004) seems devoid of content, given that variables have no meaning by themselves, unless the model is also represented (which is not part of the hypothesis).

In brief, by casting perception as a statistical inference problem, claims of optimality miss the real computational problem of perception : the world is not noisy, it is complex.

References

Knill DC, Pouget A (2004) The Bayesian brain: the role of uncertainty in neural coding and computation. Trends Neurosci 27:712–719.

Körding KP, Wolpert DM (2006) Bayesian decision theory in sensorimotor control. Trends Cogn Sci 10:319–326.

Ricci M, Kim J, Serre T (2018) Not-So-CLEVR: Visual Relations Strain Feedforward Neural Networks. ArXiv180203390 Cs Q-Bio Available at: <http://arxiv.org/abs/1802.03390> [Accessed April 19, 2018].

Rumsfeld D (2011) Known and Unknown: A Memoir. Penguin.