highest-earning participants did. At very low signal-to-noise ratios, it is difficult to tell whether one should try harder or give up and guess. Further, although decreasing decision thresholds within trials are required by reward-rate-optimality in many circumstances, we found that fixed thresholds within trials gave the most parsimonious account of our data (Karşılar et al. 2014).

We conclude from these findings that when simple heuristics allow near-optimal computation, people can and often do implement those heuristics and adapt toward optimality. Time is a dimension in which people and animals routinely show evidence of reward-rate-optimal performance. Signal-to-noise ratio just is not. In fact, when it comes to timing tasks, people and animals show striking conformance with predictions of the reward-rate-optimality hypothesis (e.g., Balcı et al. 2009; 2011a; Çavdaroğlu et al. 2014; Simen et al. 2011).

We agree with the authors that the term "optimal" is widely used with different meanings. Bayesian modelers frequently describe models that incorporate new evidence according to Bayes' rule to be "optimal." Implicit in this assumption is that accurate evidence updating will automatically yield whatever sort of optimal outcome you ultimately choose to define. Yet added to the general computational intractability of Bayesian inference, the best rules for reading out that evidence into choices may themselves be intractable (e.g., with more than two options; McMillen & Holmes 2006). Surely then, heuristics are the best we can do, in general.

Nonetheless, optimality hypotheses can aid the development of suboptimal, heuristic models that adhere to plausible constraints. For example, reward-rate-optimality motivated our development of a heuristic neural network model that nearly optimizes speed-accuracy tradeoffs when tasks speed up or slow down, but that fails to adapt to changes in signal-to-noise ratio (Simen et al. 2006) – a behavioral pattern we later observed. Optimality analyses may therefore advance the standard observer model goal by establishing usable benchmarks for heuristic designers to exploit during the model-design phase. Hence, the scientific reward-rate-optimal theoretical stance may involve more optimality theory than the authors recommend.

Credo for optimality

Alan A. Stocker

Department of Psychology, University of Pennsylvania, Philadelphia PA 19104.
astocker@psych.upenn.edu https://www.sas.upenn.edu/~astocker

doi:10.1017/S0140525X18001346, e244

Abstract

Optimal or suboptimal, Rahnev & Denison (R&D) rightly argue that this ill-defined distinction is not useful when comparing models of perceptual decision making. However, what they miss is how valuable the focus on optimality has been in deriving these models in the first place. Rather than prematurely abandon the optimality assumption, we should refine this successful normative hypothesis with additional constraints that capture specific limitations of (sensory) information processing in the brain.

Scientific progress depends on our ability to formulate clear hypotheses that can be experimentally tested (i.e., models). The goal is to be able to explain why the data are what they are, which goes beyond a mere description of the data. *Optimal inference* (Helmholtz 1867) as a general hypothesis of perception has served the community extremely well in that regard, in particular in combination with the Bayesian framework (Knill & Richards 1996). It provided quantitative but nonetheless intuitive explanations for many fundamental characteristics of perception, such as how sensory information from different sources is combined (Ernst & Banks 2002), how prior information affects the percept (Körding & Wolpert 2004), and how stochastic choice behavior naturally emerges from a process of probabilistic inference (Stocker & Simoncelli 2006a). That some of these specific models are not universally valid (the listed examples of "suboptimalities" in Rahnev & Denison's [R&D] article) lies in the iterative nature of scientific progress: Parsimonious models must be successively refined as new data demand modifications. It seems rather foolish to question a very successful general hypothesis because specific assumptions of a particular model implementation turn out to be simplistic. We are far from having exhaustively explored the optimality hypothesis and therefore should not abandon it lightly; I elaborate on this in the following:

The discussion is premature. "Optimal" per definition refers to best possible with respect to some given limitations/constraints. What separates an "optimal observer" from an "ideal observer" is that the latter only considers limitations in terms of information provided to the observer and thus is well defined in an experimental setting, whereas the former also includes constraints that are internal to the observer. Because we just started to explore these constraints (such as limited representational resources), we are not yet in a position to make with any confidence a general assessment of whether perception is optimal or not. Hence, we should simply abstain from drawing any premature conclusions.

Optimality is a very valuable, normative hypothesis. The general hypothesis has proved extremely valuable in deriving models of perceptual decision making, in particular (but not only) in combination with the Bayesian formalism. Its normative nature allows us to formulate with relative ease an observer model for any specific task and thus to possess a quantitative model hypothesis *before* actually running the experiment and knowing the data. This is a substantial advantage as it empowers us to not only design experiments that are most efficient in validating the model, but also to cross-validate the model by making specific predictions for a new task based on the model parameters determined from data of a previous task. It is also important to realize that without optimality assumption, the Bayesian formalism would have been unlikely adopted by the community over the last 20 to 30 years, a formalism that undeniably was very successful by any rational metric. Bayesian decision theory and the optimality assumption are in many ways synonymous; without the latter, the former is not meaningful. Finally, optimality is by no means an arbitrary hypothesis but ultimately directly follows from the theory of evolution: A system's actions and behaviors are aimed to perform in the best way possible in order to optimize the chances of survival and reproduction in a competitive environment. Briefly, the optimality assumption is a well-supported, very useful assumption.

Is there an alternative? "To deny that we reason in a Bayesian way is to assert that we reason in a deliberately inconsistent [i.e., random] way" (Jaynes 1957/2003, p. 133). Clearly, Jaynes did not think that there was an alternative. But even with a more measured view, it is difficult to conceive of such a possibility. As R&D rightly state, "suboptimal" is not an alternative: Simply rejecting the null hypothesis is not a hypothesis, a common fallacy far too often encountered in the psychological sciences. Any alternative hypothesis of equal value must have a normative character (i.e., it must allow us to formulate quantitative models for specific perceptual decision-making tasks). Not only that, but it also must explain why under some conditions perceptual behavior is seemingly optimal and under others it is not. R&D's proposed "standard observer model" is as fuzzy as they described it (could be a "bag of tricks" or a "neural network"; sect. 5.2, para. 1) because they have no idea what alternative hypothesis it should represent – there simply is no equivalent alternative to the optimality assumption at the moment.

So, how do we best move forward? Let us simply follow proper scientific procedure: The current evidence in favor of optimality far outweighs the evidence in favor of any other potential hypothesis (whatever that might be). This does not mean that the experimental evidence suggesting "suboptimal" behavior should be ignored, on the contrary. But we should not prematurely abandon the optimality assumption either. Rather we should continue probing the general normative hypothesis that has been so good to us and try to refine and extend it to make it fit with new experimental data that require so. Recent work from my laboratory may serve as an example for this approach: We have noticed that perceptual biases (e.g., in perceived visual orientation) are frequently away from the peak densities of the expected prior distribution, which contradicts the predictions of a traditional optimal Bayesian observer (de Gardelle et al. 2010; Tomassini et al. 2010). But rather than claiming suboptimal behavior and calling it a day, we realized that the traditional Bayesian observer model relies on implicit assumptions that actually may be incorrect (in this case the noise distributions). Indeed, we showed that if we add an additional constraint to the observer model – namely, that the observer's representational resources are limited and must be used efficiently (the efficient coding hypothesis [Barlow 1961], another optimality assumption) – then such "suboptimal" behavior indeed turns out to be perfectly optimal (Wei & Stocker 2015). Not only that, but the additional constraint also allowed us to discover a new perceptual law, describing a functional relationship between perceptual bias and discrimination threshold (Wei & Stocker 2017). Many of the "suboptimalities" that R&D list (in their Table 1) can be thought of as observer-related constraints and limitations, and it seems more likely than not that eventually, they can all be described within an optimal model. There is much potential in using the optimal (Bayesian) observer model as our well-defined standard model and improving and extending it with additional constraints that we discover based on computational, psychophysical, and physiological considerations.

Eventually, however, I absolutely agree with R&D that we should not waste our energy with dogmatic battles (in particular because the topic of these battles seems not well defined) but rather focus on "building and testing detailed observer models that explain behavior across a wide range of tasks" (abstract) and "capture all the systematic weirdness of human behavior rather than preserve an aesthetic ideal" (sect. 6, para. 1). We are scientists, and as such, we should ideally value and judge different models solely based on their ability to account for and rightly predict the full richness of the data. However, in putting us in a position to do this, the value of the optimality hypothesis is currently unrivaled.

Perceptual suboptimality: Bug or feature?

Christopher Summerfield and Vickie Li

Department of Experimental Psychology, University of Oxford, Oxford OX2 6GG, United Kingdom.
christopher.summerfield@psy.ox.ac.uk chui.li@psy.ox.ac.uk

doi:10.1017/S0140525X18001437, e245

Abstract

Rahnev & Denison (R&D) argue that whether people are "optimal" or "suboptimal" is not a well-posed question. We agree. However, we argue that the critical question is why humans make suboptimal perceptual decisions in the first place. We suggest that perceptual distortions have a normative explanation – that they promote efficient coding and computation in biological information processing systems.

Rahnev & Denison (R&D) argue that psychologists and neuroscientists are unduly concerned with the question of whether perceptual decisions are "optimal" or "suboptimal." They suggest that this question is ill posed, and that researchers should instead use observer models to provide an idealised benchmark against which to compare human behaviour.

In large part, we agree. Nevertheless, we suggest that the article rather sidesteps the major conceptual issue that underpins this debate from the standpoint of cognitive science, neuroscience, and machine learning: Why do these suboptimalities occur in the first place? Here, we argue that paradoxically, perceptual distortions observed in the lab often have a sound normative basis. In other words, perceptual "suboptimality" is best seen as a "feature" rather than a "bug" in the neural source code that guides our behaviour.

The authors discuss how suboptimal behaviours arise from distortions in the prior or likelihood functions, or misconceptions about the relevant cost function or decision rule. As they show, the Bayesian framework offers an elegant means to characterise the sources of bias or variance that corrupt decisions. However, it does not offer principled insights into why perceptual distortions might occur. To illustrate why this question is pressing, consider the perspective of a researcher attempting to build an artificial brain. She needs to know whether a given behavioural phenomenon – for example, the sequential decision bias that R&D discuss – is something that the artificial system should embrace or eschew. Only by knowing why biological systems display this phenomenon can this question be addressed.

Over recent years, advances have been made towards addressing the "why" of perceptual distortion. One elegant example pertains to the oblique effect (Appelle 1972), which (as R&D allude to) can be brought under the umbrella of Bayesian inference by considering human priors over the natural statistics of visual scenes, in which cardinal orientations predominate. But here, the Bayesian notion of a "prior" is an oversimplification that does not explain how or why the effect arises. In fact, the oblique effect can be understood by considering the optimisation principle that allows visual representations to be formed in the first place. Various classes of unsupervised learning rule, such as Hebbian learning, encourage neural systems to form