Suboptimality in perceptual decision making

Dobromir Rahnev[a] and Rachel N. Denison[b]

[a]School of Psychology, Georgia Institute of Technology, Atlanta, GA 30332 and [b]Department of Psychology and Center for Neural Science, New York University, New York, NY 10003.
drahnev@gmail.com   rahnevlab.gatech.edu
rachel.denison@nyu.edu   racheldenison.com

What is Open Peer Commentary? What follows on these pages is known as a Treatment, in which a significant and controversial Target Article is published along with Commentaries (p. 17) and an Author's Response (p. 46). See bbsonline.org for more information.

CAMBRIDGE
UNIVERSITY PRESS

Abstract. Human perceptual decisions are often described as optimal. Critics of this view have argued that claims of optimality are overly flexible and lack explanatory power. Meanwhile, advocates for optimality have countered that such criticisms single out a few selected papers. To elucidate the issue of optimality in perceptual decision making, we review the extensive literature on suboptimal performance in perceptual tasks. We discuss eight different classes of suboptimal perceptual decisions, including improper placement, maintenance, and adjustment of perceptual criteria; inadequate tradeoff between speed and accuracy; inappropriate confidence ratings; misweightings in cue combination; and findings related to various perceptual illusions and biases. In addition, we discuss conceptual shortcomings of a focus on optimality, such as definitional difficulties and the limited value of optimality claims in and of themselves. We therefore advocate that the field drop its emphasis on whether observed behavior is optimal and instead concentrate on building and testing detailed observer models that explain behavior across a wide range of tasks. To facilitate this transition, we compile the proposed hypotheses regarding the origins of suboptimal perceptual decisions reviewed here. We argue that verifying, rejecting, and expanding these explanations for suboptimal behavior – rather than assessing optimality per se – should be among the major goals of the science of perceptual decision making.

1. Introduction

How do people make perceptual judgments based on the available sensory information? This fundamental question has been a focus of psychological research from the nineteenth century onward (Fechner 1860; Helmholtz 1856). Many perceptual tasks naturally lend themselves to what has traditionally been called "ideal observer" analysis, whereby the optimal behavior is mathematically determined given a set of assumptions such as the presence of sensory noise, and human behavior is compared to this standard (Geisler 2011; Green & Swets 1966; Ulehla 1966). The extensive literature on this topic includes many examples of humans performing similarly to an ideal observer but also many examples of suboptimal behavior. Perceptual science has a strong tradition of developing models and theories that attempt to account for the full range of empirical data on how humans perceive (Macmillan & Creelman 2005).

Recent years have seen an impressive surge of Bayesian theories of human cognition and perception (Gershman et al. 2015; Griffiths et al. 2015; Tenenbaum et al. 2011). These theories often depict humans as optimal decision makers, especially in the area of perception. A number of high-profile papers have shown examples of human perceptual behavior that is close to optimal (Ernst & Banks 2002; Körding & Wolpert 2004; Landy et al. 1995; Shen & Ma 2016), whereas other papers have attempted to explain apparently suboptimal behaviors as being in fact optimal (Weiss et al. 2002). Consequently, many statements by researchers in the field leave the impression that humans are essentially optimal in perceptual tasks:

> Psychophysics is providing a growing body of evidence that human perceptual computations are "Bayes' optimal." (Knill & Pouget 2004, p. 712)
> Across a wide range of tasks, people seem to act in a manner consistent with optimal Bayesian models. (Vul et al. 2014, p. 1)
> These studies with different approaches have shown that human perception is close to the Bayesian optimal. (Körding & Wolpert 2006, p. 321)

Despite a number of recent criticisms of such assertions regarding human optimality (Bowers & Davis 2012a; 2012b; Eberhardt & Danks 2011; Jones & Love 2011; Marcus & Davis 2013; 2015), as well as statements from some of the most prominent Bayesian theorists that their goal is not to demonstrate optimality (Goodman et al. 2015; Griffiths et al. 2012), the previous quotes indicate that the view that humans are (close to) optimal when making perceptual decisions has taken a strong foothold.

The main purpose of this article is to counteract assertions about human optimality by bringing together the extensive literature on suboptimal perceptual decision making.

Although the description of the many findings of suboptimality will occupy a large part of the article, we do not advocate for a shift of labeling observers from "optimal" to "suboptimal." Instead, we will ultimately argue that we should abandon any emphasis on optimality or suboptimality and return to building a science of perception that attempts to account for all types of behavior.

The article is organized into six sections. After introducing the topic (sect. 1), we explain the Bayesian approach to perceptual decision making and explicitly define a set of standard assumptions that typically determine what behavior is considered optimal (sect. 2). In the central section of the article, we review the vast literature of suboptimal perceptual decision making and show that suboptimalities have been reported in virtually every class of perceptual tasks (sect. 3). We then discuss theoretical problems with the current narrow focus on optimality, such as difficulties in defining what is truly optimal and the limited value of optimality claims in and of themselves (sect. 4). Finally, we argue that the way forward is to build observer models that give equal emphasis to all components of perceptual decision making, not only the decision rule (sect. 5). We conclude that the field should abandon its emphasis on optimality and instead focus on thoroughly testing the hypotheses that have already been generated (sect. 6).

2. Defining optimality

Optimality can be defined within many frameworks. Here we adopt a Bayesian approach because it is widely used in the field and it is general: other approaches to optimality can often be expressed in Bayesian terms.

2.1. The Bayesian approach to perceptual decision making

The Bayesian approach to perceptual decision making starts with specifying the generative model of the task. The model defines the sets of world states, or stimuli, $\mathcal{S}$, internal responses X, actions $\mathcal{A}$, and relevant parameters Θ (such as the sensitivity of the observer). We will mostly focus on cases in which two possible stimuli s_1 and s_2 are presented, and the possible "actions" a_1 and a_2 are reporting that the corresponding stimulus was shown. The Bayesian approach then specifies the following quantities (see Fig. 1 for a graphical depiction):

DOBROMIR RAHNEV is Assistant Professor of Psychology at the Georgia Institute of Technology. He earned a B.S. in psychology from Harvard University and a Ph.D. in psychology from Columbia University, after which he completed a postdoctoral training in cognitive neuroscience at University of California, Berkeley. Rahnev's research focuses on perceptual decision making and confidence generation with an emphasis on creating computational models that explain human behavior.

RACHEL N. DENISON, Postdoctoral Associate in the Department of Psychology and Center for Neural Science at New York University, studies visual perception, attention, and decision making. She holds a B.A. in cognitive science from Yale University, an M.Sc. in neuroscience from the University of Oxford, and a Ph.D. in neuroscience from the University of California, Berkeley. Her work has been supported by a Marshall Scholarship and fellowships from the National Science Foundation and National Institutes of Health.

Likelihood function. An external stimulus can produce a range of internal responses. The measurement density, or distribution, $p(x|s, \theta)$ is the probability density of obtaining an internal response x given a particular stimulus s. The likelihood function $l(s|x, \theta)$ is equal to the measurement density but is defined for a fixed internal response as opposed to a fixed stimulus.

Prior. The prior $\pi(s)$ describes one's assumptions about the probability of each stimulus s.

Cost function. The cost function $\mathcal{L}(s, a)$ (also called *loss function*) specifies the cost of taking a specific action for a specific stimulus.

Decision rule. The decision rule $\delta(x)$ indicates under what combination of the other quantities you should perform one action or another.

We refer to the likelihood function, prior, cost function, and decision rule as the LPCD components of perceptual decision making.

According to Bayesian decision theory (Körding & Wolpert 2006; Maloney & Mamassian 2009), the optimal decision rule is to choose the action a that minimizes the expected loss over all possible stimuli. Using Bayes' theorem, we can derive the optimal decision rule as a function of the likelihood, prior, and cost function:

$$\delta(x) = \operatorname{argmin}_{a \in A} \sum_{s \in \mathcal{S}} l(s|x, \theta)\ \pi(s)\ \mathcal{L}(s, a).$$

2.2. Standard assumptions

Determining whether observers' decisions are optimal requires the specification of the four LPCD components. How do researchers determine the quantitative form of each component? The following is a typical set of *standard assumptions* related to each LPCD component:

Likelihood function assumptions. The standard assumptions here include Gaussian measurement distributions and stimulus encoding that is independent from other factors such as stimulus presentation history. Note that the experimenter derives the likelihood function from the assumed measurement distributions.

Prior and cost function assumptions. The standard assumption about observers' internal representations of the prior and cost function is that they are identical to the quantities defined by the experimenter. Unless specifically mentioned, the experiments reviewed subsequently here present s_1 and s_2 equally often, which is equivalent to a uniform prior (e.g., $\pi(s_i) = \dfrac{1}{2}$ when there are two stimuli), and expect observers to maximize percent correct, which is equivalent to a cost function that punishes all incorrect responses, and rewards all correct responses, equally.

Decision rule assumptions. The standard assumption about the decision rule is that it is identical to the optimal decision rule.

Finally, additional general standard assumptions include expectations that observers can perform the proper computations on the LPCD components. Note that as specified, the standard assumptions consider Gaussian variability at encoding as the sole corrupting element for perceptual decisions. Section 3 assembles the evidence against this claim.

The attentive reader may object that the standard assumptions cannot be universally true. For example, assumptions related to
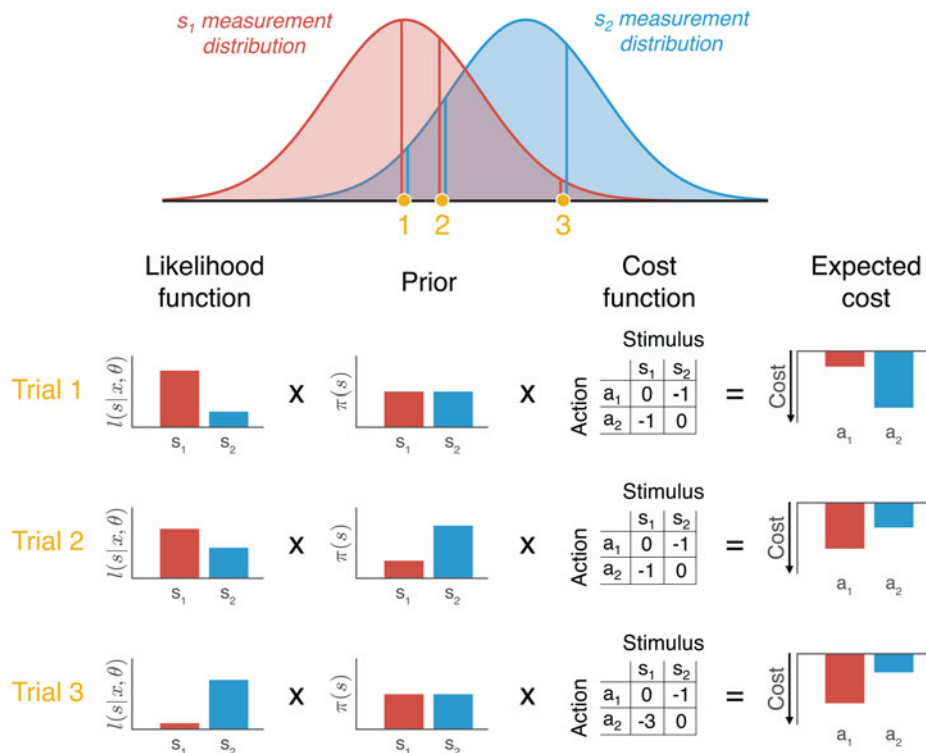
Figure 1. Graphical depiction of Bayesian inference. An observer is deciding between two possible stimuli – s_1 (e.g., leftward motion) and s_2 (e.g., rightward motion) – which produce Gaussian measurement distributions of internal responses. The observer's internal response varies from trial to trial, depicted by the three yellow circles for three example trials. On a given trial, the likelihood function is equal to the height of each of the two measurement densities at the value of the observed internal response (lines drawn from each yellow circle) – that is, the likelihood of each stimulus given an internal response. For illustrative purposes, a different experimenter-provided prior and cost function are assumed on each trial. The action a_i corresponds to choosing stimulus s_i. We obtain the expected cost of each action by multiplying the likelihood, prior, and cost corresponding to each stimulus and then summing the costs associated with the two possible stimuli. The optimal decision rule is to choose the action with the lower cost (the bar with less negative values). In trial 1, the prior and cost function are unbiased, so the optimal decision depends only on the likelihood function. In trial 2, the prior is biased toward s_2, making a_2 the optimal choice even though s_1 is slightly more likely. In trial 3, the cost function favors a_1, but the much higher likelihood of s_2 makes a_2 the optimal choice.

the likelihood function are likely false for specific paradigms (e.g., measurement noise may not be Gaussian), and assumptions about observers adopting the experimentally defined prior and cost function are likely false for complex experimental designs (Beck et al. 2012). Nevertheless, we take the standard assumptions as a useful starting point for our review because, explicitly or implicitly, they are assumed in most (although not all) studies. In section 3, we label all deviations from behavior prescribed by the standard assumptions as examples of suboptimality. We discuss alternative ways of defining optimality in section 4 and ultimately argue that general statements about the optimality or suboptimality of perceptual decisions are meaningless.

3. Review of suboptimality in perceptual decision making

We review eight categories of tasks for which the optimal decision rule can be determined. For each task category, we first note any relevant information about the measurement distribution, prior, or cost function. We plot the measurement distributions together with the optimal decision rule (which we depict as a criterion drawn on the internal responses X). We then review specific suboptimalities within each task category. For each explanation of apparently suboptimal behavior, we indicate the standard LPCD components proposed to have been violated using the notation [*LPCD component*], such as [*decision rule*]. Note that violations of the assumed measurement distributions result in violations of the assumed likelihood functions. In some cases, suboptimalities have been attributed to issues that apply to multiple components (indicated as [*general*]) or issues of methodology (indicated as [*methodological*]).

3.1. Criterion in two-choice tasks

In the most common case, observers must distinguish between two possible stimuli, s_1 and s_2, presented with equal probability

and associated with equal reward. In Figure 2, we plot the measurement distributions and optimal criteria for the cases of equal and unequal internal variability. The criterion used to make the decision corresponds to the decision rule.

3.1.1. Detection criteria

Many tasks involve the simple distinction between noise (s_1) and signal + noise (s_2). These are usually referred to as *detection tasks*. In most cases, s_1 is found to produce smaller internal variability than s_2 (Green & Swets 1966; Macmillan & Creelman 2005; Swets et al. 1961), from which it follows that an optimal observer would choose s_1 more often than s_2 even when the two stimuli are presented at equal rates (Fig. 2). Indeed, many detection studies find that observers choose the noise distribution s_1 more than half of the time (Gorea & Sagi 2000; Green & Swets 1966; Rahnev et al. 2011b; Reckless et al. 2014; Solovey et al. 2015; Swets et al. 1961). However, most studies do not allow for the estimation of the exact measurement distributions for individual observers, and hence it is an open question how optimal observers in those studies actually are. A few studies have reported conditions in which observers choose the noise stimulus s_1 *less* than half of the time (Morales et al. 2015; Rahnev et al. 2011b; Solovey et al. 2015). Assuming that the noise distributions in those studies also had lower variability, such behavior is likely suboptimal.

3.1.2. Discrimination criteria

Detection tasks require observers to distinguish between the noise versus signal + noise stimuli, but other tasks require observers to discriminate between two roughly equivalent stimuli. For example, observers might discriminate leftward versus rightward motion or clockwise versus counterclockwise grating orientation. For these types of stimuli, the measurement distributions for each stimulus category can be assumed to have approximately equal variability (Macmillan & Creelman 2005; See et al. 1997). Such

Figure 2. Depiction of the measurement distributions (colored curves) and optimal criteria (equivalent to the decision rules) in two-choice tasks. The upper panel depicts the case when the two stimuli produce the same internal variability ($\sigma_1 = \sigma_2$, where σ is the standard deviation of the Gaussian measurement distribution). The gray vertical line represents the location of the optimal criterion. The lower panel shows the location of the optimal criterion when the variability of the two measurement distributions differs ($\sigma_1 < \sigma_2$, in which case the optimal criterion results in a higher proportion of s_1 responses).

studies find that the average criterion location across the whole group of observers is usually close to optimal, but individual observers can still exhibit substantial biases (e.g., Whiteley & Sahani, 2008). In other words, what appears as an optimal criterion on average (across observers) may be an average of suboptimal criteria (Mozer et al. 2008; Vul et al. 2014). This issue can appear within an individual observer, too, with suboptimal criteria on different trials averaging out to resemble an optimal criterion (see sect. 3.2). To check for criterion optimality within individual observers, we re-analyzed the data from a recent study in which observers discriminated between a grating tilted 45 degrees clockwise or counterclockwise from vertical (Rahnev et al. 2016). Seventeen observers came for four sessions on different days completing 480 trials each time. Using a binomial test, we found that 57 of the 68 total sessions exhibited significant deviation from unbiased responding. Further, observers tended to have relatively stable biases as demonstrated by a positive criterion correlation across all pairs of sessions (all p's < .003). Hence, even if the performance of the group appears to be close to optimal, individual observers may deviate substantially from optimality.

3.1.3. Two-stimulus tasks
The biases observed in detection and discrimination experiments led to the development of the two-alternative forced-choice (2AFC) task, in which both stimulus categories are presented on each trial (Macmillan & Creelman 2005). The 2AFC tasks separate the two stimuli either temporally (also referred to as two-interval forced-choice or 2IFC tasks) or spatially. Note that, in recent years, researchers have begun to use the term "2AFC" for two-choice tasks in which only one stimulus is presented. To avoid confusion, we adopt the term "two-stimulus tasks" to refer to tasks where two stimuli are presented (the original meaning of 2AFC) and the term "one-stimulus tasks" to refer to tasks like single-stimulus detection and discrimination (e.g., the tasks discussed in sects. 3.1.1 and 3.1.2).

Even though two-stimulus tasks were designed to remove observer bias, significant biases have been observed for them, too. Although biases in spatial 2AFC tasks have received less attention, several suboptimalities have been documented for 2IFC tasks. For example, early research suggested that the second stimulus is more often selected as the one of higher intensity, a phenomenon called *time-order errors* (Fechner 1860; Osgood 1953). More recently, Yeshurun et al. (2008) re-analyzed 2IFC data from 17 previous experiments and found significant interval biases. The direction of the bias varied across the different experiments, suggesting that the specific experimental design has an influence on observers' bias.

3.1.4. Explaining suboptimality in two-choice tasks
Why do people appear to have trouble setting appropriate criteria in two-choice tasks? One possibility is that they have a tendency to give the same fixed response when uncertain [*decision rule*]. For example, a given observer may respond that he saw left (rather than right) motion every time he got distracted or had very low evidence for either choice. This could be because of a preference for one of the two stimuli or one of the two motor responses. Re-analysis of another previous study (Rahnev et al. 2011a), where we withheld the stimulus-response mapping until after the stimulus presentation, found that 12 of the 21 observers still showed a significant response bias for motion direction. Therefore, a preference in motor behavior cannot fully account for this type of suboptimality.

Another possibility is that for many observers even ostensibly "equivalent" stimuli such as left and right motion give rise to measurement distributions with unequal variance [*likelihood function*]. In that case, an optimal decision rule would produce behavior that appears biased. Similarly, in two-stimulus tasks, it is possible that the two stimuli are not given the same resources or that the internal representations for each stimulus are not independent of each other [*likelihood function*]. Finally, in the case of detection tasks, it is possible that some observers employ an idiosyncratic cost function by treating misses as less costly than false alarms because the latter can be interpreted as lying [*cost function*].

3.2. Maintaining stable criteria
So far, we have considered the optimality of the decision rule when all trials are considered together. We now turn our attention to whether observers' decision behavior varies across trials or conditions (Fig. 3).

3.2.1. Sequential effects
Optimality in laboratory tasks requires that judgments are made based on the evidence from the current stimulus independent of previous stimuli. However, sequential effects are ubiquitous in perceptual tasks (Fischer & Whitney 2014; Fründ et al. 2014; Kaneko & Sakai 2015; Liberman et al. 2014; Norton et al. 2017; Tanner et al. 1967; Treisman & Faulkner 1984; Ward & Lockhead 1970; Yu & Cohen 2009). The general finding is that observers' responses are positively autocorrelated such that the response on the current trial is likely to be the same as on the previous trial, though in some cases negative autocorrelations have also been reported (Tanner et al. 1967; Ward & Lockhead 1970). Further, observers are able to adjust to new trial-to-trial statistics, but this adjustment is only strong in the direction of default biases and weak in the opposite direction (Abrahamyan

Figure 3. Depiction of a failure to maintain a stable criterion. The optimal criterion is shown in Figure 2, but observers often fail to maintain that criterion over the course of the experiment, resulting in a criterion that effectively varies across trials. Colored curves show measurement distributions.

et al. 2016). Similar effects have been observed in other species such as mice (Busse et al. 2011).

3.2.2. Criterion attraction

Interleaving trials that require different criteria also hinders optimal criterion placement. Gorea and Sagi (2000) proposed that when high-contrast stimuli (optimally requiring a relatively conservative detection criterion) and low-contrast stimuli (optimally requiring a relatively liberal detection criterion) were presented simultaneously, observers used the same compromised detection criterion that was suboptimal for both the high- and low-contrast stimuli. This was despite the fact that, on each trial, they told observers with 100% certainty which contrasts might have been present in each location. Similar criterion attraction has been proposed in a variety of paradigms that involved using stimuli of different contrasts (Gorea & Sagi 2001; 2002; Gorea et al. 2005; Zak et al. 2012), attended versus unattended stimuli (Morales et al. 2015; Rahnev et al. 2011b), and central versus peripheral stimuli (Solovey et al. 2015). Although proposals of criterion attraction consider the absolute location of the criterion on the internal decision axis, recent work has noted the methodological difficulties of recovering absolute criteria in signal detection tasks (Denison et al. 2018).

3.2.3. Irrelevant reward influencing the criterion

The optimal decision rule is insensitive to multiplicative changes to the cost function. For example, rewarding all correct responses with $0.01 versus $0.03, while incorrect responses receive $0, should not alter the decision criterion; in both cases, the optimal decision rule is the one that maximizes percent correct. However, greater monetary rewards or punishments lead observers to adopt a more liberal detection criterion such that more stimuli are identified as targets (Reckless et al. 2013; 2014). Similar changes to the response criterion because of monetary motivation are obtained in a variety of paradigms (Henriques et al. 1994; Taylor et al. 2004). To complicate matters, observers' personality traits interact with the type of monetary reward in altering response criteria (Markman et al. 2005).

3.2.4. Explaining suboptimality in maintaining stable criteria

Why do people appear to shift their response criteria based on factors that should be irrelevant for criterion placement? Sequential effects are typically explained in terms of an automatic tendency to exploit the continuity in our normal environment, even though such continuity is not present in most experimental setups (Fischer & Whitney 2014; Fritsche et al. 2017; Liberman et al. 2014). The visual system could have built-in mechanisms that bias new representations toward recent ones [*likelihood function*], or it may assume that a new stimulus is likely to be similar to a recent one [*prior*]. (Note that the alternative likelihoods or priors would need to be defined over pairs or sequences of trials.) Adopting a prior that the environment is autocorrelated may be a good strategy for maximizing reward: Environments typically are autocorrelated and, if they are not, such a prior may not hurt performance (Yu & Cohen 2009).

Criterion attraction may stem from difficulty maintaining two separate criteria simultaneously. This is equivalent to asserting that in certain situations observers cannot maintain a more complicated decision rule (e.g., different criteria for different conditions) and instead use a simpler one (e.g., single criterion for all conditions) [*decision rule*]. It is harder to explain why personality traits or task features such as increased monetary rewards (that should be irrelevant to the response criterion) change observers' criteria.

3.3. Adjusting choice criteria

Two of the most common ways to assess optimality in perceptual decision making are to manipulate the prior probabilities of the stimulus classes and to provide unequal payoffs that bias responses toward one of the stimulus categories (Macmillan & Creelman 2005). Manipulating prior probabilities affects the prior $\pi(s)$, whereas manipulating payoffs affects the cost function $\mathcal{L}(s, a)$. However, the two manipulations have an equivalent effect on the optimal decision rule: Both require observers to shift their decision criterion by a factor dictated by the specific prior probability or reward structure (Fig. 4).

3.3.1. Priors

Two main approaches have been used to determine whether observers can optimally adjust their criterion when one of two stimuli has a higher probability of occurrence. In base-rate manipulations, long blocks of the same occurrence frequency are employed, and observers are typically not informed of the probabilities of occurrence in advance (e.g., Maddox 1995). Most studies find that observers adjust their criterion to account for the unequal base rate, but this adjustment is smaller than what is required for optimal performance, resulting in a conservative criterion placement (Bohil & Maddox 2003b; Green & Swets 1966; Maddox & Bohil 2001; 2003; 2005; Maddox & Dodd 2001; Maddox et al. 2003; Tanner 1956; Tanner et al. 1967; Vincent 2011). Some studies have suggested that observers become progressively more suboptimal as the base rate becomes progressively more extreme (Bohil & Maddox 2003b; Green & Swets 1966). However, a few studies have reported that certain conditions result in extreme criterion placement such that observers rely more on base rate information than is optimal (Maddox & Bohil 1998b).

A second way to manipulate the probability of occurrence is to do it on a trial-by-trial basis and explicitly inform observers about the stimulus probabilities before each trial. This approach also leads to conservative criterion placement such that observers do not shift their criterion enough (Ackermann & Landy 2015; de Lange et al. 2013; Rahnev et al. 2011a; Summerfield & Koechlin 2010; Ulehla 1966).

Figure 4. Depiction of optimal adjustment of choice criteria. In addition to the s_1 and s_2 measurement distributions (in thin red and blue lines), the figure shows the corresponding posterior probabilities as a function of x assuming uniform prior (in thick red and blue lines). The vertical criteria depict optimal criterion locations on x (thin gray lines) and correspond to the horizontal thresholds (thick yellow lines). Optimal criterion and threshold for equal prior probabilities and payoffs are shown in dashed lines. If unequal prior probability or unequal payoff is provided such that s_1 ought to be chosen three times as often as s_2, then the threshold would optimally be shifted to 0.75, corresponding to a shift in the criterion such that the horizontal threshold and vertical criterion intersect on the s_2 posterior probability function. The y-axis is probability density for the measurement distributions and probability for the posterior probability functions; the y-axis ticks refer to the posterior probability.

3.3.2. Payoffs

The decision criterion can also be manipulated by giving different payoffs for different responses. The general finding with this manipulation is that observers, again, do not adjust their criterion enough (Ackermann & Landy 2015; Bohil & Maddox 2001; 2003a; 2003b; Busemeyer & Myung 1992; Maddox & Bohil 1998a; 2000; 2001; 2003; 2005; Maddox & Dodd 2001; Maddox et al. 2003; Markman et al. 2005; Taylor et al. 2004; Ulehla 1966) and, as with base rates, become more suboptimal for more extreme payoffs (Bohil & Maddox 2003b). Nevertheless, one study that involved a very large number of sessions with two monkeys reported extreme criterion changes (Feng et al. 2009).

Criterion adjustments in response to unequal payoffs are usually found to be more suboptimal compared with adjustments in response to unequal base rates (Ackermann & Landy 2015; Bohil & Maddox 2001; 2003a; Busemeyer & Myung 1992; Healy & Kubovy 1981; Maddox 2002; Maddox & Bohil 1998a; Maddox & Dodd 2001), though the opposite pattern was found by Green and Swets (1966).

Finally, the exact payoff structure may also influence observers' optimality. For example, introducing a cost for incorrect answers leads to more suboptimal criterion placement compared with conditions with the same optimal criterion shift but without a cost for incorrect answers (Maddox & Bohil 2000; Maddox & Dodd 2001; Maddox et al. 2003).

3.3.3. Explaining suboptimality in adjusting choice criteria

Why do people appear not to adjust their decision criteria optimally in response to priors and rewards? One possibility is that they do not have an accurate internal representation of the relevant probability implied by the prior or reward structure [*general*]

(Acerbi et al. 2014b; Ackermann & Landy 2015; Zhang & Maloney 2012). For example, Zhang and Maloney (2012) argued for the presence of "ubiquitous log odds" that systematically distort people's probability judgments such that small values are overestimated and large values are underestimated (Brooke & MacRae 1977; Juslin et al. 2009; Kahneman & Tversky 1979; Varey et al. 1990).

A possible explanation for the suboptimality in base-rate experiments is the "flat-maxima" hypothesis, according to which the observer adjusts the decision criterion based on the change in reward and has trouble finding its optimal value if other criterion positions result in similar reward rates [*methodological*] (Bohil & Maddox 2003a; Busemeyer & Myung 1992; Maddox & Bohil 2001; 2003; 2004; 2005; Maddox & Dodd 2001; Maddox et al. 2003; von Winterfeldt & Edwards 1982). Another possibility is that the prior observers adopt in base-rate experiments comes from a separate process of Bayesian inference. If observers are uncertain about the true base rate, a prior assumption that it is likely to be unbiased would result in insufficient base rate adjustment [*methodological*]. A central tendency bias can also arise when observers form a prior based on the sample of stimuli they have encountered so far, which are unlikely to cover the full range of the experimenter-defined stimulus distribution (Petzschner & Glasauer 2011). We classify these issues as methodological because if the observers have not been able to learn a particular likelihood, prior, and cost function (LPC) component, then they cannot adopt the optimal decision rule.

Finally, another possibility is that observers also place a premium on being correct rather than just maximizing reward [*cost function*]. Maddox and Bohil (1998a) posited the competition between reward and accuracy maximization (COBRA) hypothesis according to which observers attempt to maximize reward but also place a premium on accuracy (Maddox & Bohil 2004; 2005). This consideration applies to manipulations of payoffs but not of prior probabilities and may explain why payoff manipulations typically lead to larger deviations from optimality than priors.

3.4. Tradeoff between speed and accuracy

In the previous examples, the only variable of interest has been observers' choice irrespective of their reaction times (RTs). However, if instructed, observers can provide responses faster at lower accuracy, a phenomenon known as speed-accuracy tradeoff (SAT; Fitts 1966; Heitz 2014). An important question here is whether observers can adjust their RTs optimally to achieve maximum reward in a given amount of time (Fig. 5). A practical difficulty for studies attempting to address this question is that the accuracy/RT curve is not generally known and is likely to differ substantially between different tasks (Heitz 2014). Therefore, the only standard assumption here is that accuracy increases monotonically as a function of RT. Precise accuracy/RT curves can be constructed by assuming one of the many models from the sequential sampling modeling framework (Forstmann et al. 2016), and there is a vibrant discussion about the optimal stopping rule depending on whether signal reliability is known or unknown (Bogacz 2007; Bogacz et al. 2006; Drugowitsch et al. 2012; 2015; Hanks et al. 2011; Hawkins et al. 2015; Thura et al. 2012). However, because different models predict different accuracy/RT curves, in what follows we only assume a monotonic relationship between accuracy and RT.

Figure 5. (A) Depiction of one possible accuracy/reaction time (RT) curve. Percent correct responses increases monotonically as a function of RT and asymptotes at 90%. (B) The total reward/RT curve for the accuracy/RT curve from panel A with the following additional assumptions: (1) observers complete as many trials as possible within a 30-minute window, (2) completing a trial takes 1.5 seconds on top of the RT (because of stimulus presentation and between-trial breaks), and (3) each correct answer results in 1 point, whereas incorrect answers result in 0 points. The optimal RT – the one that maximizes the total reward – is depicted with dashed lines.

3.4.1. Trading off speed and accuracy

Although observers are able to adjust their behavior to account for both accuracy and RT, they cannot do so optimally (Balcı et al. 2011b; Bogacz et al. 2010; Simen et al. 2009; Starns & Ratcliff 2010; 2012; Tsetsos et al. 2015). In most cases, observers take too long to decide, leading to slightly higher accuracy but substantially longer RTs than optimal (Bogacz et al. 2010; Simen et al. 2009; Starns & Ratcliff 2010; 2012). This effect occurs when observers have a fixed period of time to complete as many trials as possible (Bogacz et al. 2010; Simen et al. 2009; Starns & Ratcliff 2010; 2012) and in the more familiar design with a fixed number of trials per block (Starns & Ratcliff 2010; 2012). Further, observers take longer to decide for more difficult compared with easier conditions, even though optimizing the total reward demands that they do the opposite (Oud et al. 2016; Starns & Ratcliff 2012). Older adults are even more suboptimal than college-age participants by this measure (Starns & Ratcliff 2010; 2012).

3.4.2. Keeping a low error rate under implicit time pressure

Even though observers tend to overemphasize accuracy, they are also suboptimal in tasks that require an extreme emphasis on accuracy. This conclusion comes from a line of research on visual search in which observers are typically given an unlimited amount of time to decide whether a target is present or not (Eckstein 2011). In certain situations, such as airport checkpoints or detecting tumors in mammograms, the goal is to keep a very low miss rate irrespective of RT, because misses can have dire consequences (Evans et al. 2013; Wolfe et al. 2013). The optimal RT can be derived from Figure 5A as the minimal RT that results in the desired accuracy rate. A series of studies by Wolfe and colleagues found that observers, even trained doctors and airport checkpoint

screeners, are suboptimal in such tasks in that they allow overly high rates of misses (Evans et al. 2011; 2013; Wolfe & Van Wert 2010; Wolfe et al. 2005; 2013). Further, this effect was robust and resistant to a variety of methods designed to help observers take longer in order to achieve higher accuracy (Wolfe et al. 2007) or reduce motor errors (Van Wert et al. 2009). An explanation of this suboptimality based on capacity limits is rejected by two studies that found that observers can be induced to take longer time, and thus achieve higher accuracy, by first providing them with a block of high prevalence targets accompanied by feedback (Wolfe et al. 2007; 2013).

3.4.3. Explaining suboptimality in the speed-accuracy tradeoff

Why do people appear to be unable to trade off speed and accuracy optimally? Similar to explanations from the previous sections, it is possible to account for overly long RTs by postulating that, in addition to maximizing their total reward, observers place a premium on being accurate [*cost function*] (Balcı et al. 2011b; Bogacz et al. 2010; Holmes & Cohen 2014). Another possibility is that observers' judgments of elapsed time are noisy [*general*], and longer-than-optimal RTs lead to a higher reward rate than RTs that are shorter than optimal by the same amount (Simen et al. 2009; Zacksenhouse et al. 2010). Finally, in some situations, observers may also place a premium on speed [*cost function*], preventing a very low error rate (Wolfe et al. 2013).

3.5. Confidence in one's decision

The Bayesian approach prescribes how the posterior probability should be computed. Although researchers typically examine the question of whether the stimulus with highest posterior probability is selected, it is also possible to examine whether observers can report the actual value of the posterior distribution or perform simple computations with it (Fig. 6). In such cases, observers are asked to provide "metacognitive" confidence ratings about the accuracy of their decisions (Metcalfe & Shimamura 1994; Yeung & Summerfield 2012). Such studies rarely provide subjects with an explicit cost function (but see Kiani & Shadlen 2009; Rahnev et al. 2013) but, in many cases, reasonable assumptions can be made in order to derive optimal performance (see sects. 3.5.1–3.5.4).

3.5.1. Overconfidence and underconfidence (confidence calibration)

It is straightforward to construct a payoff structure for confidence ratings such that observers gain the most reward when their confidence reflects the posterior probability of being correct (e.g., Fleming et al. 2016; Massoni et al. 2014). Most studies, however, do not provide observers with such a payoff structure, so assessing the optimality of the confidence ratings necessitates the further assumption that observers create a similar function internally. To test for optimality, we can then consider, for example, all trials in which an observer has 70% confidence of being correct and test whether the average accuracy on those trials is indeed 70%. This type of relationship between confidence and accuracy is often referred to as confidence calibration (Baranski & Petrusic 1994). Studies of confidence have found that for certain tasks observers are overconfident (i.e., they overestimate their accuracy) (Adams 1957; Baranski & Petrusic 1994; Dawes 1980; Harvey 1997; Keren 1988; Koriat 2011), whereas for other tasks observers are underconfident (i.e., they underestimate their accuracy) (Baranski & Petrusic 1994; Björkman et al. 1993; Dawes 1980;

Figure 6. Depiction of how an observer should give confidence ratings. Similar to Figure 4, both the measurement distributions and posterior probabilities as a function of *x* assuming uniform prior are depicted. The confidence thresholds (depicted as yellow lines) correspond to criteria defined on *x* (depicted as gray lines). The horizontal thresholds and vertical criteria intersect on the posterior probability functions. The *y*-axis is probability density for the measurement distributions and probability for the posterior probability functions; the *y*-axis ticks refer to the posterior probability.

Harvey 1997; Winman & Juslin 1993). One pattern that emerges consistently is that overconfidence occurs in difficult tasks, whereas underconfidence occurs in easy tasks (Baranski & Petrusic 1994, 1995, 1999), a phenomenon known as the *hard-easy effect* (Gigerenzer et al. 1991). Similar results are seen for tasks outside of the perceptual domain such as answering general knowledge questions (Griffin & Tversky 1992). Overconfidence and underconfidence are stable over different tasks (Ais et al. 2015; Song et al. 2011) and depend on non-perceptual factors such as one's optimism bias (Ais et al. 2015).

3.5.2. Dissociations of confidence and accuracy across different experimental conditions

Although precise confidence calibration is computationally difficult, a weaker test of optimality examines whether experimental conditions that lead to the same performance are judged with the same level of confidence (even if this level is too high or too low). This test only requires that observers' confidence ratings follow a consistent internal cost function across the two tasks. Many studies demonstrate dissociations between confidence and accuracy across tasks, thus showing that observers fail this weaker optimality test. For example, speeded responses can decrease accuracy but leave confidence unchanged (Baranski & Petrusic 1994; Vickers & Packer 1982), whereas slowed responses can lead to the same accuracy but lower confidence (Kiani et al. 2014). Dissociations between confidence and accuracy have also been found in conditions that differ in attention (Rahnev et al. 2012a; Rahnev et al. 2011b; Wilimzig et al. 2008), the variability of the perceptual signal (de Gardelle & Mamassian 2015; Koizumi et al. 2015; Samaha et al. 2016; Song et al. 2015; Spence et al. 2016; Zylberberg et al. 2014), the stimulus-onset asynchrony in metacontrast masking (Lau & Passingham 2006), the presence of unconscious information (Vlassova et al. 2014), and the relative timing of a concurrent saccade (Navajas et al. 2014). Further, some of these biases seem to arise from individual differences that are stable across multiple sessions (de Gardelle & Mamassian 2015). Finally, dissociations between confidence and accuracy have been found in studies that applied transcranial

magnetic stimulation (TMS) to the visual (Rahnev et al. 2012b), premotor (Fleming et al. 2015), or frontal cortex (Chiang et al. 2014).

3.5.3. Metacognitive sensitivity (confidence resolution)

The previous sections were concerned with the average magnitude of confidence ratings over many trials. Another measure of interest is the degree of correspondence between confidence and accuracy on individual trials (Metcalfe & Shimamura 1994), called *metacognitive sensitivity* (Fleming & Lau 2014) or *confidence resolution* (Baranski & Petrusic 1994). Recently, Maniscalco and Lau (2012) developed a method to quantify how optimal an observer's metacognitive sensitivity is. Their method computes *meta-d′*, a measure of how much information is available for metacognition, which can then be compared with the actual *d′* value. An optimal observer would have a *meta-d′/d′* ratio of 1. Maniscalco and Lau (2012) obtained a ratio of 0.77, suggesting a 23% loss of information for confidence judgments. Even though some studies that used the same measure but different perceptual paradigms found values close to 1 (Fleming et al. 2014), many others arrived at values substantially lower than 1 (Bang et al. in press; Maniscalco & Lau 2015; Maniscalco et al. 2016; Massoni 2014; McCurdy et al. 2013; Schurger et al. 2015; Sherman et al. 2015; Vlassova et al. 2014). Interestingly, at least one study has reported values significantly greater than 1, suggesting that in certain cases the metacognitive system has more information than was used for the primary decision (Charles et al. 2013), thus implying the presence of suboptimality in the perceptual decision.

3.5.4. Confidence does not simply reflect the posterior probability of being correct

Another way of assessing the optimality of confidence ratings is to determine whether observers compute confidence in a manner consistent with the posterior probability of being correct. This is also a weaker condition than reporting the actual posterior probability of being correct, because it does not specify how observers should place decision boundaries between different confidence ratings, only that these boundaries should depend

on the posterior probability of being correct. Although one study found that confidence ratings are consistent with computations based on the posterior probability (Sanders et al. 2016; but see Adler & Ma 2018b), others showed that either some (Aitchison et al. 2015; Navajas et al. 2017) or most (Adler & Ma 2018a; Denison et al. 2018) observers are described better by heuristic models in which confidence depends on uncertainty but not on the actual posterior probability of being correct.

Further, confidence judgments are influenced by a host of factors unrelated to the perceptual signal at hand and thus in violation of the principle that they should reflect the posterior probability of being correct. For example, emotional states, such as worry (Massoni 2014) and arousal (Allen et al. 2016), affect how sensory information relates to confidence ratings. Other factors, such as eye gaze stability (Schurger et al. 2015), working memory load (Maniscalco & Lau 2015), and age (Weil et al. 2013), affect the relationship between confidence and accuracy. Sequential effects have also been reported for confidence judgments such that a high confidence rating is more likely to follow a high, rather than low, confidence rating (Mueller & Weidemann 2008). Confidence dependencies exist even between different tasks, such as letter and color discrimination, that depend on different neural populations in the visual cortex (Rahnev et al. 2015). Inter-task confidence influences have been dubbed "confidence leak" and have been shown to be negatively correlated with observers' metacognitive sensitivity (Rahnev et al. 2015).

Confidence has also been shown to exhibit a "positive evidence" bias (Maniscalco et al. 2016; Zylberberg et al. 2012). In two-choice tasks, one can distinguish between sensory evidence in a trial that is congruent with the observer's response on that trial (positive evidence) and sensory evidence that is incongruent with the response (negative evidence). Even though the perceptual decisions usually follow the optimal strategy of weighting equally both of these sources of evidence, confidence ratings are suboptimal in depending more on the positive evidence (Koizumi et al. 2015; Maniscalco et al. 2016; Samaha et al. 2016; Song et al. 2015; Zylberberg et al. 2012).

3.5.5. Explaining suboptimality in confidence ratings
Why do people appear to give inappropriate confidence ratings? Some components of overconfidence and underconfidence can be explained by inappropriate transformation of internal evidence into probabilities [*general*] (Zhang & Maloney 2012), methodological considerations such as interleaving conditions with different difficulty levels, which can have inadvertent effects on the prior [*methodological*] (Drugowitsch et al. 2014b), or even individual differences such as shyness about giving high confidence, which can be conceptualized as extra cost for high-confidence responses [*cost function*]. Confidence-accuracy dissociations are often attributed to observers' inability to maintain different criteria for different conditions, even if they are clearly distinguishable [*decision rule*] (Koizumi et al. 2015; Rahnev et al. 2011b). The "positive evidence" bias [*decision rule*] introduced in the end of section 3.5.4 can also account for certain suboptimalities in confidence ratings.

More generally, it is possible that confidence ratings are not only based on the available perceptual evidence as assumed by most modeling approaches (Drugowitsch & Pouget 2012; Green & Swets 1966; Macmillan & Creelman 2005; Ratcliff & Starns 2009; Vickers 1979). Other theories postulate the existence of either different processing streams that contribute differentially to the perceptual decision and the subjective confidence judgment (Del

Cul et al. 2009; Jolij & Lamme 2005; Weiskrantz 1996) or a second processing stage that determines the confidence judgment and that builds on the information in an earlier processing stage responsible for the perceptual decision (Bang et al. in press; Fleming & Daw 2017; Lau & Rosenthal 2011; Maniscalco & Lau 2010, 2016; Pleskac & Busemeyer 2010; van den Berg et al. 2017). Both types of models could be used to explain the various findings of suboptimal behavior and imply the existence of different measurement distributions for decision and confidence [*likelihood function*].

3.6. Comparing sensitivity in different tasks
The previous sections discussed observers' performance on a single task. Another way of examining optimality is to compare the performance on two related tasks. If the two tasks have a formal relationship, then an optimal observer's sensitivity on the two tasks should follow that relationship.

3.6.1. Comparing performance in one-stimulus and two-stimulus tasks
Visual sensitivity has traditionally been measured by employing either (1) a one-stimulus (detection or discrimination) task in which a single stimulus from one of two stimulus classes is presented on each trial or (2) a two-stimulus task in which both stimulus classes are presented on each trial (see sect. 3.1.3). Intuitively, two-stimulus tasks are easier because the final decision is based on more perceptual information. Assuming independent processing of each stimulus, the relationship between the sensitivity on these two types of tasks can be mathematically defined: The sensitivity on the two-stimulus task should be $\sqrt{2}$ times higher than on the one-stimulus task (Macmillan & Creelman, 2005; Fig. 7). Nevertheless, empirical studies have often contradicted this predicted relationship: Many studies have found sensitivity ratios smaller than $\sqrt{2}$ (Creelman & Macmillan 1979; Jesteadt 1974; Leshowitz 1969; Markowitz & Swets 1967; Pynn 1972; Schulman & Mitchell 1966; Swets & Green 1961; Viemeister 1970; Watson et al. 1973; Yeshurun et al. 2008), though a few have found ratios larger than $\sqrt{2}$ (Leshowitz 1969; Markowitz & Swets 1967; Swets & Green 1961).

3.6.2. Comparing performance in other tasks
Many other comparisons between tasks have been performed. In temporal 2IFC tasks, observers often have different sensitivities to the two stimulus intervals (García-Pérez & Alcalá-Quintana 2010; 2011; Yeshurun et al. 2008), suggesting an inability to distribute resources equally. Other studies find that longer inter-stimulus intervals in 2IFC tasks lead to decreases in sensitivity (Berliner & Durlach 1973; Kinchla & Smyzer 1967; Tanner 1961), presumably because of memory limitations. Further, choice variability on three-choice tasks is greater than what would be predicted by a related two-choice task (Drugowitsch et al. 2016). Creelman and Macmillan (1979) compared the sensitivity on nine different psychophysical tasks and found a complex pattern of dependencies, many of which were at odds with optimal performance. Finally, Olzak (1985) demonstrated deviations from the expected relationship between detection and discrimination tasks.

An alternative approach to comparing an observer's performance on different tasks is allowing observers to choose which tasks they prefer to complete and analyzing the optimality of these decisions. In particular, one can test for the presence of transitivity: If an observer prefers task A to task B and task B to task C, then the observer should prefer task A to task C. Several studies

Figure 7. Depiction of the relationship between one-stimulus and two-stimulus tasks. Each axis corresponds to a one-stimulus task (e.g., Fig. 2). The three sets of concentric circles represent two-dimensional circular Gaussian distributions corresponding to presenting two stimuli in a row (e.g., s_2,s_1 means that s_2 was the first stimulus and s_1 was the second stimulus). If the discriminability between s_1 and s_2 is d' (one-stimulus task; gray lines in triangle), then the Pythagorean theorem gives us the expected discriminability between s_1,s_2 and s_2,s_1 (two-stimulus task; blue line in triangle).

Figure 8. Optimal cue combination. Two cues that give independent information about the value of a sensory feature (red and blue curves) are combined to form a single estimate of the feature value (yellow curve). For Gaussian cue distributions, the combined cue distribution is narrower than both individual cue distributions, and its mean is closer to the mean of the distribution of the more reliable cue.

suggest that human observers violate the transitivity principle both in choosing tasks (Zhang et al. 2010) and choosing stimuli (Tsetsos et al. 2016a), though there is considerable controversy surrounding such findings (Davis-Stober et al. 2016; Kalenscher et al. 2010; Regenwetter et al. 2010; 2011, 2017).

3.6.3. Explaining suboptimality in between-task comparisons

Why does human performance on different tasks violate the expected relationship between these tasks? One possibility is that observers face certain capacity limits in one task, but not the other, that alter how the stimuli are encoded [*likelihood function*]. For example, compared to a one-stimulus task, the more complex two-stimulus task requires the simultaneous processing of two stimuli. If limited resources hamper the processing of the second stimulus, then sensitivity in that task will fall short of what is predicted based on the one-stimulus task.

In some experiments, observers performed worse than expected on the one-stimulus task, rather than on the two-stimulus task. A possible explanation of this effect is the presence of a larger "criterion jitter" in the one-stimulus task (i.e., a larger variability in the decision criterion from trial to trial). Because two-stimulus tasks involve the comparison of two stimuli on each trial, these tasks are less susceptible to criterion jitter. Such criterion variability, which could stem from sequential dependencies or even random criterion fluctuations (see sect. 3.2), decreases the estimated stimulus sensitivity (Mueller & Weidemann 2008). The criterion jitter could also be the result of computational imprecision [*general*] (Bays & Dowding 2017; Beck et al. 2012; Dayan 2014; Drugowitsch et al. 2016; Renart & Machens 2014; Whiteley & Sahani 2012; Wyart & Koechlin 2016). Such imprecision could arise from constraints at the neural level and may account for a large amount of choice suboptimality (Drugowitsch et al. 2016).

3.7. Cue combination

Studies of cue combination have been fundamental to the view that sensory perception is optimal (Trommershäuser et al. 2011). Cue combination (also called "cue integration") is needed whenever different sensory features provide separate pieces of information about a single physical quantity. For example, auditory and visual signals can separately inform about the location of an object. Each cue provides imperfect information about the physical world, but different cues have different sources of variability. As a result, integrating the different cues can provide a more accurate and reliable estimate of the physical quantity of interest.

One can test for optimality in cue combination by comparing the perceptual estimate formed from two cues with the estimates formed from each cue individually. The optimal estimate is typically taken to be the one that maximizes precision (minimizes variability) across trials (Fig. 8). When the variability for each cue is Gaussian and independent of the other cues, the maximum likelihood estimate (MLE) is a linear combination of the estimates from each cue, weighted by their individual reliabilities (Landy et al. 2011). Whether observers conform to this weighted sum formula can be readily tested psychophysically, and a large number of studies have done exactly this for different types of cues and tasks (for reviews, see Ma 2010; Trommershäuser et al. 2011).

In particular, the optimal mean perceptual estimate (x) after observing cue 1 (with feature estimate x_1 and variance σ_1^2) and cue 2 (with feature estimate x_2 and variance σ_2^2) is

$$x = \frac{\frac{x_1}{\sigma_1^2} + \frac{x_2}{\sigma_2^2}}{\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}},$$

such that the feature estimate x_i is weighted by its reliability $\frac{1}{\sigma_i^2}$ and the whole expression is normalized by the sum of the reliabilities. The optimal variance of the perceptual estimate (σ^2) is

$$\sigma^2 = \frac{\sigma_1^2 \sigma_2^2}{\sigma_1^2 + \sigma_2^2}.$$

3.7.1. Examples of optimality in cue combination

A classic example of cue combination is a study of visual-haptic cue combination by Ernst and Banks (2002). In this study, observers estimated the height of a rectangle using (1) only sight, (2) only touch, or (3) both sight and touch. Performance in the visual-haptic condition was well described by the MLE formula: The single cue measurements predicted both the reliability of the combined estimates and the weights given to each cue. Many studies have observed similar optimal cue combination behavior in a range of tasks estimating different physical quantities (Trommershäuser et al. 2011). These studies have investigated integration across two modalities (including vision, touch, audition, the vestibular sense, and proprioception; e.g., Alais & Burr, 2004; Ernst & Banks, 2002; Gu et al. 2008; van Beers et al. 1996) and across two features in the same modality, such as various visual cues to depth (e.g., Jacobs 1999; Landy et al. 1995). Common among these experiments is that trained observers complete many trials of a psychophysical task, and the two cues provide similar estimates of the quantity of interest. Optimal cue combination has also been observed during sensory-motor integration (Maloney & Zhang 2010; Trommershäuser 2009; Wei & Körding 2011; Yeshurun et al. 2008).

3.7.2. Examples of suboptimality in cue combination

Because optimality is often the hypothesized outcome in cue combination studies, findings of suboptimality may be underreported or underemphasized in the literature (Rosas & Wichmann 2011). Still, a number of studies have demonstrated suboptimal cue combination that violates some part of the MLE formula. These violations fall into two categories: (1) those in which the cues are integrated but are not weighted according to their independently measured reliabilities, and (2) those in which estimates from two cues are no better than estimates from a single cue.

In the first category are findings from a wide range of combined modalities: visual-auditory (Battaglia et al. 2003; Burr et al. 2009; Maiworm & Röder 2011), visual-vestibular (Fetsch et al. 2012; Prsa et al. 2012), visual-haptic (Battaglia et al. 2011; Rosas et al. 2005), and visual-visual (Knill & Saunders 2003; Rosas et al. 2007). For example, auditory and visual cues were not integrated according to the MLE rule in a localization task; instead, observers treated the visual cue as though it were more reliable than it really was (Battaglia et al. 2003). Similarly, visual and haptic texture cues were integrated according to their reliabilities, but observers underweighted the visual cue (Rosas et al. 2005). Suboptimal integration of visual and auditory cues was also found for patients with central vision loss, but not for patients with peripheral vision loss (Garcia et al. 2017).

In some of these studies, cue misweighting was restricted to low-reliability cues: In a visual-vestibular heading task, observers overweighted vestibular cues when visual reliability was low (Fetsch et al. 2012), and in a visual-auditory temporal order judgment task, observers overweighted auditory cues when auditory reliability was low (Maiworm & Röder 2011). However, overweighting does not only occur within a limited range of reliabilities (e.g., Battaglia et al. 2003; Prsa et al. 2012).

Several studies have failed to find optimal cue combination in the temporal domain. In an audiovisual rate combination task, observers only partially integrated the auditory and visual cues, and they did not integrate them at all when the rates were very different (Roach et al. 2006). Observers also overweighted auditory cues in temporal order judgment tasks (Maiworm & Röder 2011) and temporal bisection tasks (Burr et al. 2009). It is well established that when two cues give very different estimates, observers tend to discount one of them (Gepshtein et al. 2005; Jack & Thurlow 1973; Körding et al. 2007; Roach et al. 2006; Warren & Cleaves 1971), an effect which has been called "robust fusion" (Maloney & Landy 1989), which may arise from inferring that the two cues come from separate sources (Körding et al. 2007). However, in most of the studies just described, suboptimal cue combination was observed even when the cues gave similar estimates.

In the second category of suboptimal cue combination findings, two cues are no better than one (Chen & Tyler 2015; Drugowitsch et al. 2014a; Landy & Kojima 2001; Oruç et al. 2003; Rosas et al. 2005; 2007). (Note that some of these studies found a mix of optimal and suboptimal observers.) Picking the best cue is known as a "veto" type of cue combination (Bülthoff & Mallot 1988) and is considered a case of "strong fusion" (Clark & Yullie 1990; Landy et al. 1995). This is an even more serious violation of optimal cue combination, because it is as though no integration has taken place at all – the system either picks the best cue or, in some cases, does *worse* with two cues than with one.

Cues may also be mandatorily combined even when doing so is not suitable for the observer's task. For example, texture and disparity information about slant was subsumed in a combined estimate, rendering the single cue estimates unrecoverable (Hillis et al. 2002). Interestingly, the single cue estimates were not lost for children, allowing them to outperform adults when the cues disagreed (Nardini et al. 2010). In a related finding, observers used multiple visual features to identify a letter even when the optimal strategy was to use only a single, relevant feature (Saarela & Landy 2015).

3.7.3. Combining stimuli of the same type

So far, we have only considered cue combination studies in which the two cues come from different sensory modalities or dimensions. Suboptimal behavior has also been observed when combining cues from the same dimension. For example, Summerfield and colleagues have shown that observers do not weight every sample stimulus equally in a decision (Summerfield & Tsetsos 2015). For simultaneous samples, observers underweighted "outlier" stimuli lying far from the mean of the sample (de Gardelle & Summerfield 2011; Michael et al. 2014; 2015; Vandormael et al. 2017). For sequential samples, observers overweighted stimuli toward the end of the sequence (a recency effect) as well as stimuli that are similar to recently presented items (Bang & Rahnev 2017; Cheadle et al. 2014; Wyart et al. 2015). Observers also used only a subset of a sample of orientations to estimate the mean orientation of the sample (Dakin 2001). More generally, accuracy on tasks with sequential samples is substantially lower than what would be predicted by sensory noise alone (Drugowitsch et al. 2016).

3.7.4. Combining sensory and motor cues

Suboptimal cue integration has also been found in sensory-motor tasks. For example, when integrating the path of a pointing movement with online visual feedback, observers underestimated the uncertainty indicated by the feedback (Körding & Wolpert 2004). In a pointing task in which observers were rewarded for physically touching the correct visual target, observers underweighted the difficulty of the motor task by aiming for a small target, even though the perceptual information indicating the target was also uncertain (Fleming et al. 2013). Similar biases were reported in a related task (Landy et al. 2007). Within the action domain (and so beyond our focus on perception), Maloney and Zhang (2010) have reviewed studies showing both optimal and suboptimal behavior.

3.7.5. Cue combination in children

Optimal cue integration takes time to develop. Children are suboptimal until around 10 years of age when combining multisensory (Gori et al. 2008; Nardini et al. 2008; Petrini et al. 2014) or visual (Dekker et al. 2015; Nardini et al. 2010) cues.

3.7.6. Explaining suboptimal cue combination

Why do people sometimes appear to combine cues suboptimally? One possible explanation is that observers do not have accurate representations of the reliability of the cues (Knill & Saunders 2003; Rosas et al. 2005) because learning the reliability is difficult [*methodological*]. This methodological issue is particularly acute when the cues are new to the observer. For example, in one task for which cue combination was suboptimal, observers haptically explored a surface with a single finger to estimate its slant. However, observers may have little experience with single-finger slant estimation, because multiple fingers or the whole hand might ordinarily be used for such a task (Rosas et al. 2005). Alternatively, cue combination may be suboptimal when one cue provides all information in parallel but the other cue provides information serially (Plaisier et al. 2014). Reliability estimation might also be difficult when the reliability is very low. This possibility may apply to studies in which observers were optimal within a range of sensory reliabilities, but not outside it (Fetsch et al. 2012; Maiworm & Röder 2011).

Some authors suggest that another reason for overweighting or underweighting a certain cue could be prior knowledge about how cues ought to be combined [*prior*]. This could include a prior assumption about how likely a cue is to be related to the desired physical property (Battaglia et al. 2011; Ganmor et al. 2015), how likely two cue types are to correspond to one another (and thus be beneficial to integrate) (Roach et al. 2006), or a general preference to rely on a particular modality, such as audition in a timing task (Maiworm & Röder 2011).

For certain tasks, some researchers question the assumptions of the MLE model, such as Gaussian noise [*likelihood function*] (Burr et al. 2009) or the independence of the neural representations of the two cues [*likelihood function*] (Rosas et al. 2007). In other cases, it appears that observers use alternative cost functions by, for example, taking RT into account [*cost function*] (Drugowitsch et al. 2014a).

"Robust averaging," or down-weighting of outliers, has been observed when observers must combine multiple pieces of information that give very different perceptual estimates. Such down-weighting can stem from adaptive gain changes [*likelihood function*] that result in highest sensitivity to stimuli close to the mean of the sample (or in the sequential case, the subset of the sample that has been presented so far; Summerfield & Tsetsos, 2015). This adaptive gain mechanism is similar to models of sensory adaptation (Barlow 1990; Carandini & Heeger 2012; Wark et al. 2007). By following principles of efficient coding that place the largest dynamic range at the center of the sample (Barlow 1961; Brenner et al. 2000; Wainwright 1999), different stimuli receive unequal weightings. Psychophysical studies in which stimulus variability is low would not be expected to show this kind of suboptimality (Cheadle et al. 2014).

It is debated whether suboptimal cue combination in children reflects a switching strategy (Adams 2016) or immature neural mechanisms for integrating cues, or whether the developing brain is optimized for a different task, such as multisensory calibration or conflict detection (Gori et al. 2008; Nardini et al. 2010).

3.8. Other examples of suboptimality

Thus far we have specifically focused on tasks where the optimal behavior can be specified mathematically in a relatively uncontroversial manner (though see sect. 4.2). However, the issue of optimality has been discussed in a variety of other contexts.

3.8.1. Perceptual biases, illusions, and improbabilities

A number of basic visual biases have been documented. Some examples include repulsion of orientation or motion direction estimates away from cardinal directions (Fig. 9A; Jastrow 1892; Rauber & Treue 1998), a bias to perceive speeds as slower than they are when stimuli are low contrast (Stone & Thompson 1992; Thompson 1982; but see Thompson et al. 2006), a bias to perceive surfaces as convex (Langer & Bülthoff 2001; Sun & Perona 1997), and a bias to perceive visual stimuli closer to fixation than they are (whereas the opposite is true for auditory stimuli; Odegaard et al. 2015).

When biases, context, or other factors lead to something looking dramatically different from its physical reality, we might call it a visual illusion. A classic example is the brightness illusion (Fig. 9B) in which two squares on a checkerboard appear to be different shades of gray even though they actually have the same luminance (Adelson 1993). Perceptual illusions persist even when the observer knows about the illusion and even after thousands of trials of exposure (Gold et al. 2000).

Some illusions are difficult to reconcile with existing theories of optimal perception. Anderson et al. (2011), for example, reported strong percepts of illusory surfaces that were improbable according to optimal frameworks for contour synthesis. In the size-weight illusion, smaller objects are perceived as heavier than larger objects of the same weight, even though the prior expectation is that smaller objects are lighter (Brayanov & Smith 2010; Peters et al. 2016).

3.8.2. Adaptation

Adaptation is a widespread phenomenon in sensory systems in which responsiveness to prolonged or repeated stimuli is reduced (Webster 2015). As some researchers have discussed (Wei & Stocker 2015), adaptation could be seen as suboptimal from a Bayesian perspective because subsequent perceptual estimates tend to diverge from rather than conform to the prior stimulus. For example, after prolonged viewing of a line tilted slightly away from vertical, a vertical line looks tilted in the opposite direction (the "tilt aftereffect," Fig. 9C; Gibson & Radner 1937). Or, after viewing motion in a certain direction, a stationary stimulus appears to drift in the opposite direction (Wohlgemuth 1911).
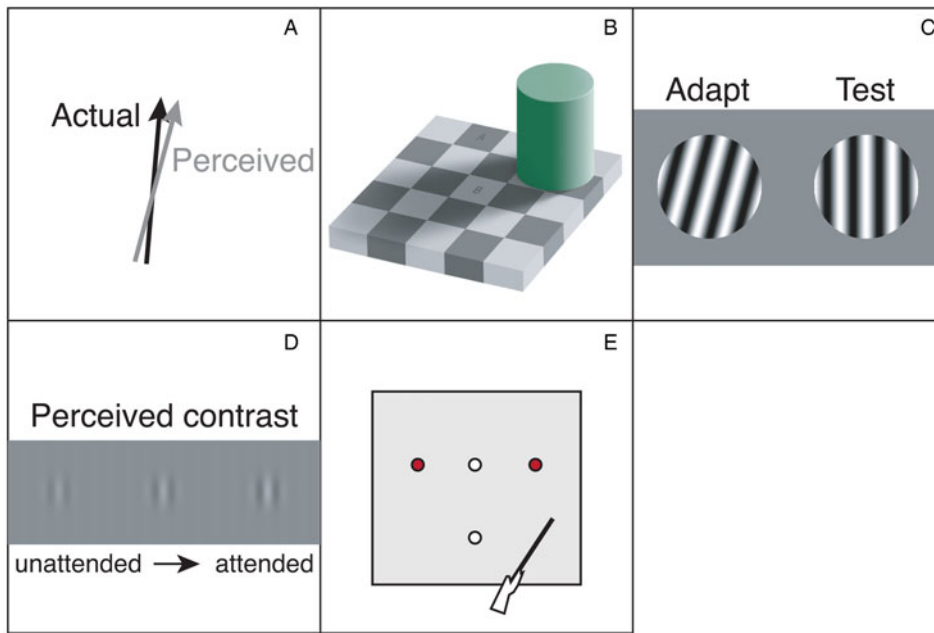
Figure 9. Examples of illusions and biases. (A) Cardinal repulsion. A nearly vertical (or horizontal) line looks more tilted away from the cardinal axis than it is. (B) Adelson's checkerboard brightness illusion. Square B appears brighter than square A, even though the two squares have the same luminance. (Image ©1995, Edward H. Adelson) (C) Tilt aftereffect. After viewing a tilted adapting grating (left), observers perceive a vertical test grating (right) to be tilted away from the adaptor. (D) Effects of spatial attention on contrast appearance (Carrasco et al. 2004). An attended grating appears to have higher contrast than the same grating when it is unattended. (E) Effects of action affordances on perceptual judgments (Witt 2011). Observers judge an object to be closer (far white circle compared to near white circle) relative to the distance between two landmark objects (red circles) when they are holding a tool that allows them to reach that object than when they have no tool.

After adapting to a certain color, perception is biased toward the complementary color (Sabra 1989; Turnbull 1961), and after adapting to a specific face, another face appears more different from that face than it would have otherwise (Webster & MacLeod 2011; Webster et al. 2004). In all of these examples, perception is repelled away from the prior stimulus, which, at least on the surface, appears suboptimal (but see sect. 3.8.5).

3.8.3. Appearance changes with visual attention
The same physical stimulus can also be perceived in different ways depending on the state of visual attention. Directing spatial attention to a stimulus can make it appear larger (Anton-Erxleben et al. 2007), faster (Anton-Erxleben et al. 2013; Fuller et al. 2009; Turatto et al. 2007), and brighter (Tse 2005), and to have higher spatial frequency (Abrams et al. 2010; Gobell & Carrasco 2005) and higher contrast (Fig. 9D; Carrasco et al. 2004; Liu et al. 2009; Störmer et al. 2009) than it would otherwise. Often attention improves performance on a visual task, but sometimes it makes performance worse (Ling & Carrasco 2006; Yeshurun & Carrasco 1998), demonstrating inflexibility in the system.

3.8.4. Cognition-based biases
Other studies have documented visual biases associated with more cognitive factors, including action affordances (Witt 2011), motivation (Balcetis 2016), and language (Lupyan 2012). For example, when people reach for an object with a tool that allows them to reach farther, they report the object as looking closer than when they reach without the tool (Fig. 9E; Witt et al. 2005). In the linguistic domain, calling an object a "triangle" leads observers to report the object as having more equal sides than when the object is called "three sided" (Lupyan 2017). How much these more cognitive factors affect perception per se, as opposed to post-perceptual judgments, and to what extent the observed visual biases are mediated by attention remain controversial questions (Firestone & Scholl 2016).

3.8.5. Explaining these other examples of apparent suboptimality
Why are people prone to certain biases and illusions? Some biases and illusions have been explained as arising from priors in the visual system [prior]. Misperceptions of motion direction (Weiss et al. 2002) and biases in reporting the speed of low-contrast stimuli (Stocker & Simoncelli 2006a; Thompson 1982; Vintch & Gardner 2014) have been explained as optimal percepts for a visual system with a prior for slow motion (Stocker & Simoncelli 2006a; Weiss et al. 2002). Such a prior is motivated by the fact that natural objects tend to be still or move slowly but has been empirically challenged by subsequent research (Hammett et al. 2007; Hassan & Hammett 2015; Thompson et al. 2006; Vaziri-Pashkam & Cavanagh 2008). Priors have been invoked to explain many other biases and illusions (Brainard et al. 2006; Girshick et al. 2011; Glennerster et al. 2006; Raviv et al. 2012). The suggestion is that these priors have been made stable over a lifetime and influence perception even when they do not apply (e.g., in a laboratory task).

Optimal decoding of sensory representations in one task can be accompanied by suboptimal biases in another task using the same stimuli. For example, in a fine-motion discrimination task, observers seem to weight the neurons tuned away from the discrimination boundary more strongly, because these neurons distinguish best between the two possible stimuli. This weighting could explain why motion direction judgments in an interleaved estimation task are biased away from the boundary (Jazayeri & Movshon 2007). Another interpretation of these results is in terms of an improper decision rule (Zamboni et al. 2016). Specifically, observers may discard sensory information related to the rejected decision outcome [decision rule] (Bronfman et al. 2015; Fleming et al. 2013; Luu & Stocker 2016), an effect known as self-consistency bias (Stocker & Simoncelli 2008).

Various efforts have been made to explain adaptation in the framework of Bayesian optimality (Grzywacz & Balboa 2002; Hohwy et al. 2008; Schwiedrzik et al. 2014; Snyder et al. 2015). One of the most well-developed lines of work explains the

repulsive effects of adaptation as a consequence of efficient coding [*likelihood function*] (Stocker & Simoncelli 2006b). In this framework, a sensory system adapts to maximize its dynamic range around the value of previous input. This change in coding does not affect the prior (as might be expected in a Bayesian treatment of adaptation) but rather affects the likelihood function. Specifically, it skews new observations away from the adapted stimulus, giving rise to repulsive aftereffects. A similar principle has been suggested to explain why perceptual estimates are repelled from long-term priors, such as those determined by the statistics of natural images (Wei & Stocker 2013; 2015).

4. Assessing optimality: Not a useful goal in itself

The extensive review in the previous section demonstrates that general claims about the optimality of human perceptual decision making are empirically false. However, there are also theoretical reasons to turn away from assessing optimality as a primary research goal.

4.1. Challenges in defining optimality

Section 2 introduced a formal definition of optimality based on Bayesian decision theory. However, the question of what phenomena should be considered optimal versus suboptimal quickly becomes complicated in many actual applications. There are at least two issues that are not straightforward to address.

The first issue concerns the exact form of the cost function. Bayesian decision theory postulates that observers should minimize the expected loss. However, observers may reasonably prefer to minimize the maximum loss, minimize the variability of the losses, or optimize some other quantity. Therefore, behavior that is suboptimal according to standard Bayesian decision theory may be optimal according to other definitions. A related, and deeper, problem is that some observers may also try to minimize other quantities such as time spent, level of boredom, or metabolic energy expended (Lennie 2003). What appears to be a suboptimal decision on a specific task may be optimal when all of these other variables are taken into account (Beck et al. 2012; Bowers & Davis 2012a). Even the clearest cases of suboptimal decision rules (e.g., the self-consistency bias) could be construed as part of a broader optimality (e.g., being self-consistent may be important for other goals). In a Bayesian framework, taking into account extra variables requires that each of the LPCD components is defined over all of these variables. If one pursues this logic, it leads to a cost function that operates over our entire evolutionary history. We do not think efforts to explore such cost functions should be abandoned, but specifying them quantitatively is impossible given our current knowledge.

The second issue concerns whether optimality should depend on the likelihood, prior, and cost function adopted by the observer. In order to be able to review a large literature using consistent assumptions, we defined a set of standard assumptions and labeled any deviation from these assumptions as suboptimal. This approach is by no means uncontroversial. For example, priors based on a lifetime of experience may be inflexible, so one could consider the standard assumption about following the experimenter-defined prior overly restrictive. An alternative view could be that suboptimal behavior concerns only deviations from the experimenter-defined quantities that are under observers' control (Tenenbaum & Griffiths 2006; Yu & Cohen 2009). The problem with this definition is that it introduces a new variable to consider – what exactly is truly under observers' control – which is often hard to determine. A third approach is to define optimality exclusively in terms of the decision rule regardless of what likelihood, prior, and cost function the observer adopts. In this view, observers are under no obligation to follow the experimenter's instructions (e.g., they are free to bring in their own priors and cost function). The problem with this approach is that failing to adopt the proper prior or cost function can result in just as much missed objective reward as adopting an improper decision rule. Similar problems apply to "improper" likelihood functions: As an extreme example, a strategy in which the observer closes her eyes (resulting in a non-informative likelihood function) and chooses actions randomly has to be labeled "optimal" because the decision rule is optimal. The ambiguity regarding the role of the likelihood, prior, or cost function points to the difficulties in constructing a general-purpose definition of optimality.

In short, optimality is impossible to define in the abstract. It 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.

4.2. Optimality claims in and of themselves have limited value

The current emphasis on optimality is fueled by the belief that demonstrating optimality in perception provides us with important insight. On the contrary, simply stating that observers are optimal is of limited value for two main reasons.

First, it is unclear when a general statement about the optimality of perceptual decisions is supposed to apply. Although most experimental work focuses on very simple tasks, it is widely recognized that the computational complexity of many real-world tasks makes optimality unachievable by the brain (Bossaerts & Murawski 2017; Cooper 1990; Gershman et al. 2015; Tsotsos 1993; van Rooij 2008). Further, in many situations, the brain cannot be expected to have complete knowledge of the likelihood function, which all but guarantees that the decision rule will be suboptimal (Beck et al. 2012). (Attempting to incorporate observers' computational capacities or knowledge brings back the problems related to how one defines optimality discussed in sect. 4.1.) Therefore, general statements about optimality must be intended only for the simplest cases of perceptual decisions (although, as sect. 3 demonstrated, even for these cases, suboptimality is ubiquitous).

Second, even for a specific task, statements about optimality alone are insufficient to predict behavior. Instead, to predict future perceptual decisions, one needs to specify each part of the process underlying the decision. Within the Bayesian framework, for example, one needs to specify each LPCD component, which goes well beyond a statement that "observers are optimal."

Is it useless to compare human performance to optimal performance? Absolutely not. *Within the context of a specific model*, demonstrating optimal or suboptimal performance is immensely helpful (Goodman et al. 2015; Tauber et al. 2017). Such demonstrations can support or challenge components of the model and suggest ways to alter the model to accommodate actual behavior. However, the critical part here is the model, not the optimality.

5. Toward a standard observer model

If there are so many empirical examples of suboptimality (sect. 3) and optimality can be challenging even to define (sect. 4), then what is the way forward?

5.1. Creating and testing observer models

Psychophysics has a long history of creating ideal observer models (Geisler 2011; Green & Swets 1966; Ulehla 1966). These models specify a set of assumptions about how sensory information is represented internally and add an optimal decision rule in order to generate predictions about behavior. The motivation behind these models has been to test the collective set of assumptions incorporated into the model. However, over time, the "ideal" part of ideal observer models has become dominant, culminating in the current outsized emphasis on demonstrating the optimality of the decision rule – what we call the *optimality approach*. Even frameworks such as "bounded rationality" (Gigerenzer & Selten 2002; Simon 1957) or "computational rationality" (Gershman et al. 2015), which explicitly concern themselves with the limitations of the decision-making process, still place the greatest emphasis on the optimality of the decision rule.

The emphasis on the decision rule in the optimality approach has led to an overly flexible treatment of the other LPCD components (Bowers & Davis 2012a). This issue is especially problematic because of the inherent degeneracy of Bayesian decision theory (Acerbi 2014): Different combinations of the likelihood, prior, cost function, and decision rule can lead to the same expected loss. Further, for any likelihood, cost function, and decision rule, a prior can be found for which that decision rule is optimal (complete class theorem) (Berger 1985; Jaynes 1957/2003).

To eliminate the flexibility of the optimality approach, the field should return to the original intention of building ideal observer models – namely, to test the collective set of assumptions incorporated into such models. To this end, we propose that researchers drop the "ideal" and shift emphasis to building, simply, "observer models." Creating observer models should differ from the current optimality approach in two critical ways. First, whether or not the decision rule is optimal should be considered irrelevant. Second, the nature of the decision rule should not be considered more important than the nature of the other components.

These two simple changes address the pitfalls of the optimality approach. Within the optimality approach, a new finding is often modeled using flexibly chosen LPCD components (Bowers & Davis 2012a). Then, depending on the inferred decision rule, a conclusion is reached that observers are optimal (or suboptimal). At this point, the project is considered complete and a general claim is made about optimality (or suboptimality). As others have pointed out, this approach has led to many "just-so stories" (Bowers & Davis 2012a), because the assumptions of the model are not rigorously tested. On the contrary, when building observer models (e.g., in the Bayesian framework), a new finding is used to generate *hypotheses* about a particular LPCD component (Maloney & Mamassian 2009). Hypotheses about the likelihood, prior, or cost function are considered as important as hypotheses about the decision rule. Critically, unlike in the optimality approach, this step is considered just the beginning of the process! The hypotheses are then examined in detail while evidence is gathered for or against them. Researchers can formulate alternative hypotheses to explain a given data set and evaluate them using model comparison techniques. In addition, researchers can conduct follow-up experiments in which they test their hypotheses using different tasks, stimuli, and observers. There are researchers who already follow this approach, and we believe the field would benefit from adopting it as the standard practice. In Box 1, we list specific steps for implementing observer models

Box 1. Implementing observer models within a Bayesian framework

1. Describe the complete generative model, including assumptions about what information the observer is using to perform the task (e.g., stimulus properties, training, experimenter's instructions, feedback, explicit vs. implicit rewards, response time pressure, etc.).
2. Specify the assumed likelihood function, prior, and cost function. If multiple options are plausible, test them in different models.
3. Derive both the optimal decision rule and plausible alternative decision rules. Compare their abilities to fit the data.
4. Interpret the results with respect to what has been learned about each LPCD component, not optimality per se. Specify how the conclusions depend on the assumptions about the other LPCD components.
5. Most importantly, follow up on any new hypotheses about LPCD components with additional studies in order to avoid "just-so stories."
6. New hypotheses that prove to be general eventually become part of the standard observer model (see sect. 5.2).

within a Bayesian framework (the steps will be similar regardless of the framework).

Two examples demonstrate the process of implementing observer models. A classic example concerns the existence of Gaussian variability in the measurement distribution. This assumption has been extensively tested for decades (Green & Swets 1966; Macmillan & Creelman 2005), thus eventually earning its place among the standard assumptions in the field. A second example comes from the literature on speed perception (sect. 3.8.5). A classic finding is that reducing the contrast of a slow-moving stimulus reduces its apparent speed (Stone and Thompson 1992; Thompson 1982). A popular Bayesian explanation for this effect is that most objects in natural environments are stationary, so the visual system has a prior for slow speeds. Consequently, when sensory information is uncertain, as occurs at low contrasts, slow-biased speed perception could be considered "optimal" (Weiss et al. 2002). Importantly, rather than stopping at this claim, researchers have investigated the hypothetical slow motion prior in follow-up studies. One study quantitatively inferred observers' prior speed distributions under the assumption of a Bayesian decision rule (Stocker & Simoncelli 2006a). Other researchers tested the slow motion prior and found that, contrary to its predictions, high-speed motion at low contrast can appear to move *faster* than its physical speed (Hammett et al. 2007; Hassan & Hammett 2015; Thompson et al. 2006). These latter studies challenged the generality of the slow motion prior hypothesis.

5.2. Creating a standard observer model

We believe that an overarching goal of the practice of creating and testing observer models is the development of a standard observer model that predicts observers' behavior on a wide variety of perceptual tasks. Such a model would be a significant achievement for the science of perceptual decision making. It is difficult – perhaps impossible – to anticipate what form the standard observer

Table 1. Summary of hypotheses proposed to account for suboptimal decisions.

LPCD component	Hypothesis	Description of the hypothesis	Relevant sections
Likelihood function	Capacity limitations	All stimuli may not be processed fully because of limited resources.	3.1, 3.2, 3.6, 3.7, 3.8
	Incorrect likelihood function	The experimenter may make wrong assumptions about the likelihood function.	3.1, 3.5, 3.7
Prior	Inappropriate priors	Priors may not be appropriate to the experimental setting (may instead reflect habitual assumptions).	3.2, 3.7, 3.8
Cost function	Placing a premium on accuracy	Observers may sacrifice total reward in order to have higher accuracy.	3.3, 3.4
	Idiosyncratic cost functions	Observers may place a premium on speed, avoid false alarms, or avoid or high/low confidence ratings.	3.1, 3.4, 3.5, 3.7
Decision rule	Non-random guesses	Observers may have a tendency to give the same response when uncertain.	3.1
	Inability to employ a complex decision rule	Observers may adopt a simpler, suboptimal decision rule when the optimal one is relatively complex.	3.2, 3.5
	Ignoring information	Observers may not consider all information relevant to the decision.	3.5, 3.8
General	Computational imprecision	Internal computations may carry inherent imprecision leading to behavioral variability.	3.4, 3.6
	Incorrect probabilities	Observers may represent or transform probabilities incorrectly.	3.3
Methodological	Incomplete learning	If observers could not learn a specific LPCD component, then they cannot perform optimally.	3.3, 3.7
	Task specification	"Irrelevant" experimental details may lead observers to alter their assumptions or strategy.	3.5
	Optimal and suboptimal behavior too similar	Optimal and suboptimal decision rules cannot be distinguished if they result in similar behavior.	3.3

model will take. It may be a Bayesian model (Maloney & Mamassian 2009), a "bag of tricks" (Ramachandran 1990), a neural network (Yamins et al. 2014), and so forth. However, regardless of the framework in which they were originally formulated, hypotheses with overwhelming empirical support will become part of the standard observer model. In this context, perhaps the most damaging aspect of the current outsized emphasis on optimality is that although it has generated many hypotheses, few of them have received sufficient subsequent attention to justify inclusion in (or exclusion from) the eventual standard observer model.

We suggest that immediate progress can be made by a concerted effort to test the hypotheses that have already been proposed to explain suboptimal decisions. To facilitate this effort, here we compile the hypotheses generated in the course of explaining the findings from section 3. Within a Bayesian framework, these hypotheses relate to the likelihood function, prior, cost function, or decision rule (the LPCD components). Further, a few of them are general and apply to several LPCD components, and a few are methodological considerations. In some cases, essentially the same hypothesis was offered in the context of several different empirical effects. We summarize these hypotheses in Table 1. Note that the table by no means exhaustively covers all existing hypotheses that deserve to be thoroughly tested.

Table 1 classifies instances of deficient learning as methodological issues. This choice is not to downplay the problem of learning. Questions of how observers acquire their priors and cost functions are of utmost importance, and meaningful progress

has already been made on this front (Acerbi et al. 2012; 2014b; Beck et al. 2012; Geisler & Najemnik 2013; Gekas et al. 2013; Seriès & Seitz 2013). Here we categorize deficient learning as a methodological issue when, because of the experimental setup, an observer cannot acquire the relevant knowledge even though she has the capacity to do so.

Future research should avoid the methodological issues from Table 1. In particular, great care must be taken to ensure that observers' assumptions in performing a task match exactly the assumptions implicit in the analysis.

We have stated the hypotheses in Table 1 at a fairly high level to succinctly capture the broad categories from our review. Much of the work ahead will be to break each high-level hypothesis down into multiple, specific hypotheses and incorporate these hypotheses into observer models. For example, statements about "inappropriate priors" or "capacity limitations" prompt more fine-grained hypotheses about specific priors or limitations whose ability to predict behavior can be tested. Some hypotheses, like capacity limitations, have already been investigated extensively – for example, in studies of attention and working memory (e.g., Carrasco 2011; Cowan 2005). Turning our existing knowledge of these phenomena into concrete observer models that predict perceptual decisions is an exciting direction for the field. Other hypotheses, like placing a premium on accuracy, have not been tested extensively and therefore should still be considered "just-so stories" (Bowers & Davis 2012a). Hence, the real work ahead lies in verifying, rejecting, and expanding the hypotheses generated from findings of suboptimal perceptual decisions.

5.3. Implications of abandoning the optimality approach

Abandoning the optimality approach has at least two immediate implications for research practices.

First, researchers should stop focusing on optimality. What should be advertised in the title and abstract of a paper is not the optimality but what is learned about the components of the perceptual process. One of the central questions in perceptual decision making is how best to characterize the sources that corrupt decisions (Beck et al. 2012; Drugowitsch et al. 2016; Hanks & Summerfield 2017; Wyart & Koechlin 2016). By shifting attention away from optimality, the effort to build complete observer models sharpens the focus on this question.

Second, new model development should not unduly emphasize optimal models. According to some Bayesian theorists, models that assume optimal behavior are intrinsically preferable to models that do not. This preference stems from the argument that because people can approximate optimal behavior on some tasks, they must possess the machinery for fully optimal decisions (Drugowitsch & Pouget 2012). Many models have been judged positively for supporting optimal decision rules: probabilistic population codes for allowing optimal cue combination (Ma et al. 2006), neural sampling models for allowing marginalization (which is needed in many optimal decision rules) (Fiser et al. 2010), and drift diffusion models for allowing optimal integration of information across time (Bogacz 2007). The large body of findings of suboptimality reviewed here, however, should make this reasoning suspect: If the brain is built to make optimal decisions, then why does it produce so many suboptimal ones? It is also important to remember that close-to-optimal behavior can also be produced by suboptimal decision rules (Bowers & Davis 2012a; Maloney & Mamassian 2009; Shen & Ma 2016). Influential theories postulate that evolutionary pressures produced heuristic but useful, rather than normative, behavior (Gigerenzer and Brighton 2009; Juslin et al. 2009; Simon 1956). Therefore, a model should be judged only on its ability to describe actual behavior, not on its ability to support optimal decision rules.

6. Conclusion

Are perceptual decisions optimal? A substantial body of research appears to answer this question in the affirmative. Here we showed instead that every category of perceptual tasks that lends itself to optimality analysis features numerous findings of suboptimality. Perceptual decisions cannot therefore be claimed to be optimal in general. In addition, independent of the empirical case against optimality, we questioned whether a focus on optimality per se can lead to any real progress. Instead, we advocated for a return to building complete observer models with an equal focus on all model components. Researchers should aim for their models to capture all of the systematic weirdness of human behavior rather than preserve an aesthetic ideal. To facilitate this effort, we compiled the hypotheses generated in the effort to explain the findings of suboptimality reviewed here. The real work ahead lies in testing these hypotheses, with the ultimate goal of developing a standard observer model of perceptual decision making.

Open Peer Commentary

Excess of individual variability of priors prevents successful development of general models

Talis Bachmann

School of Law and Institute of Psychology, Faculty of Social Sciences, University of Tartu (Tallinn branch), 10119 Tallinn, Estonia.
talis.bachmann@ut.ee

Abstract

Perceptual judgments are influenced by a multitude of factors in addition to the perceptual input. Particularly, the widely varying individual neurobiological endophenotypes and individual differences in the propensity for expectation-based illusory percepts make it unlikely that optimality is possible to define and defend by the type of abstract modeling approach criticized by Rahnev & Denison (R&D).

There are two main interrelated perspectives in the measurement of perception – "outer psychophysics" with its behavioral responses and "inner psychophysics" with its neurobiological responses to stimulation (Fechner 1882). The measurements in outer psychophysics depend on the objective stimuli and the neural processes dealing with the signals from the stimuli. These neural processes as the basis for perceptual judgments not only depend on the input signals but, as Rahnev & Denison (R&D) rightfully seem to accept, they are also heavily influenced by a host of factors unrelated to sensory signal-to-noise ratio. There are difficult problems with valid setting of the posterior probability and knowing sufficiently the priors related actual factors. Determining whether observers' decisions are optimal necessitates sufficiently detailed specification of all likelihood function, prior, cost function, and decision rule (LPCD) components. It must be stressed that substantial individual differences in the behavioral phenotypes and endophenotypes impacting perceptual judgments make a major source of obstacles for developing a powerful general model founded on the optimal decision rule.

Inner psychophysics variability as dependent on common genetic variability cannot be accounted for by any strict generalized formal models expected to be valid for the majority of experimental subjects. Given the huge number of genes implicated in neurobiological mechanism–based perceptual responses, any general models of perceptual decisions that would support the optimality theory are impossible. Indeed, even when only a few genes with variants associated with different serotonergic or dopaminergic endophenotypes are used as experimental independent variables, metacontrast masking in two-alternative discrimination shows considerable individual differences (Maksimov et al. 2013;

2015a; 2015b). Considering that there are a multitude of genes with effects expressed also in cholinergic, noradrenergic, glutamatergic, and other functions (obviously impacting perceptual behavior) and that already any optional relatively small observers' sample can invoke covert neurobiological variability with even more varying perceptual effects, any model aspiring to be general cannot be meaningfully established. Therefore, the recommendation by R&D to test more specific models with specific priors and individual limitations (as the sources that corrupt decisions) seems right. As the number of free parameters in genetic variability is very large, the task of distilling and grouping individualized endophenotypal priors as independent experimental variables is not quick and easy. But it may have importance considering that many genes implicated in perceptual behavior are also the so-called risk genes in terms of their relation to health-related vulnerability.

R&D argue that dissociations and low degree of correspondence between confidence and accuracy make general statements about the optimality or suboptimality of perceptual decisions meaningless; similar problems are related to perceptual biases, illusions, and improbabilities (sect. 3.5.2 and 3.8.1). However, the concepts of noise and bias are ambiguous when we move from general information-theoretic perceptual decision models to real psychophysiological processes and mechanisms and want to use these models in consciousness research. For example, associative learning–based expectations tend to produce hallucination-like illusions such that observers experience auditory or visual stimuli that are objectively absent (Aru & Bachmann 2017; Mack et al. 2016; Powers et al. 2017). Importantly, the proneness to such illusions is individually variable. Bayesian prior related top-down effects can produce perceptual illusions or hallucinations that are technically noise (and must be included in noise distribution), but they are at the same time phenomenal perceptual experiences taken as veridical by the observers. Hence, this "technical noise" becomes an important factor of real awareness based behavior. On the other hand, noise can be interpreted traditionally as the factor countering sensitivity simply by decreasing the signal-to-noise ratio without any illusory non-veridical experiences. Similarly, hallucinatory experiences are technically indications of bias (type 1 bias) but are at the same time factors determining behavioral choices. The other version of the bias effect relates to its impact on decision without illusory experiences of the imperative stimuli being involved (type 2 bias). In the case of illusory non-veridical percepts, high confidence ratings are expected, but in the case of no illusory experience (but low signal-to-noise ratio), low confidence ratings are expected. Both the concept of noise and the concept of bias in the traditional perceptual decision models used for advocating optimality are underspecified or even misleading in the light of experimental data on consciousness contents and qualia of perception, including illusions and hallucinations. Behaviorally incorrect perceptual responses can be subjectively right, leading to high confidence.

With regard to the domain of decision rules, R&D note that observers may not consider all information relevant to the decision. Moreover, in most of the stimulation that is present for evaluation by the observer, there are several cues on which the decision can be based. When perceptual decision models (and also methods of scaling subjective experiences such as the Perceptual Awareness Scale [PAS]; Sandberg et al. 2010) use too general, underspecified evaluation criteria such as level of confidence or subjective clarity, difficulties arise in interpreting the

results and explaining the underlying mechanisms and processes. The same behavioral outcome (e.g., successful discrimination or recognition) can be founded on different aspects of subjective contents of perception as different criterion contents (Jannati & Di Lollo 2012; Kahneman 1968). For behaviorist models of perceptual decisions, this may not be a problem, but for models to be used for consciousness studies, the excess generality of scales and criteria is a shortcoming. The cure would be development and use of specific rating scales for different subjective aspects of the same stimulation. This would be useful also for accounting for the effects of cue dissociation, which is an alternative source of ambiguities in addition to redundancy effects from cue combination.

To conclude, I tend to agree with R&D in that general statements about the optimality or suboptimality of perceptual decisions are meaningless unless detailed enough observer models will be put together.

Suboptimality in perceptual decision making and beyond

Hilary C. Barth,[a] Sara Cordes,[b] and Andrea L. Patalano[a]

[a]Department of Psychology, Wesleyan University, Middletown, CT 06459 and [b]Department of Psychology, Boston College, Chestnut Hill, MA 02467.
hbarth@wesleyan.edu　http://hbarth.faculty.wesleyan.edu
cordess@bc.edu　https://www2.bc.edu/sara-cordes/lab/
apatalano@wesleyan.edu　http://apatalano.faculty.wesleyan.edu

Abstract

We concur with the authors' overall approach and suggest that their analysis should be taken even further. First, the same points apply to areas beyond perceptual decision making. Second, the same points apply beyond issues of optimality versus suboptimality.

Rahnev & Denison (R&D) make the convincing case, with a lengthy review of examples of suboptimal perceptual performance, that broad claims of optimality in perceptual decision making are false. We applaud the overall spirit of R&D's analysis. They make the point that the pursuit of demonstrations of optimality is unproductive because "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" (sect. 4.1, para. 4). We concur also with the authors' call to test models, something that many of the cited researchers are already doing, rather than telling optimality-based "just-so stories." R&D's proposed approach is likely to do part of the work in avoiding some of the pitfalls laid out by Bowers and Davis (2012a).

In our view, R&D's main points should be taken even further, in at least two respects. First, R&D's arguments apply well beyond the realm of perceptual decision making, with relevance for other areas of mind and brain science. The article mentions that cue combination studies have been fundamental to the view that

perception is optimal. In addition to the given examples of cue combination across multiple senses, or sensorimotor integration, similar claims of optimality have been made, across development, for a wide range of tasks involving judgments of stimuli from memory (Cicchini et al. 2012; Duffy et al. 2006; Huttenlocher et al. 2000; Sciutti et al. 2015; see also Petzschner et al. 2015). Claims of optimality are widespread in this literature, and just as in perceptual decision making, because "optimality is often the hypothesized outcome … findings of suboptimality may be underreported or underemphasized" (R&D, sect. 3.7.2, para. 1). For example, this literature emphasizes models developed to explain central tendency biases in the reproduction of temporal and spatial magnitudes and characterizes those biases as optimal, yet there are multiple instances of data that violate the assumptions of these models (e.g., Barth et al. 2015; see also Crawford & Duffy 2010; Duffy & Smith 2017).

Second, seeking to explain suboptimal, as well as optimal, behavior may not be sufficient to avoid the just-so story dilemma. R&D focus on the idea that broad claims about perceptual optimality are meaningless and untestable, but this objection does not apply only to optimality per se: It applies to broad claims about perception as Bayesian inference more generally. As many critics have pointed out, both within and beyond the context of perceptual decision making, "in almost all published cases, some Bayesian model can be found that approximates or parallels human performance" (Anderson 2015, p. 280). To use an example mentioned briefly by R&D, some visual illusions are particularly difficult to explain in the context of theories of optimal perception. Usually, illusory surfaces or contours are perceived when the visual system needs to "explain" the disappearance or occlusion of an image element. The percept of an illusory occluding surface (a surface that is not defined by contrast edges) is often thought to occur because it is the most probable explanation of that disappearance. But in one class of dynamic displays, an illusory occluding surface is perceived even though a clearly visible occluding surface is already present. The disappearance of image elements is already fully explained by the presence of the visible occluder, yet our perceptual systems construct a second highly improbable occluding surface – the illusory one – generating a redundant "explanation." The phenomenon has therefore been taken as evidence against claims of visual contour synthesis as a case of rational or Bayesian inference (Anderson et al. 2011). However, as other researchers have pointed out, despite the fact that "the presence of the spurious contours results from higher-level constraints than those captured by a single Bayesian model concerned only with contour completion," the phenomenon could still be explained within a broader Bayesian view of the brain (Fleming 2011, p. R261). This is one of many examples of the ultimately untestable nature of claims about perception as Bayesian inference (for a deeper analysis of theoretical challenges to the Bayesian program in visual perception, see Anderson [2015]).

Are we trying to make a general argument against the adoption of Bayesian perspectives? No. We argue for conceptual clarity regarding the status of overarching theoretical perspectives, and against the common practice of having it both ways. Circularity is unavoidable if researchers work entirely within Bayesian frameworks while, at the same time, specific findings are taken as evidence that Bayesian perception or cognition are in play. That is, researchers who choose to take a broadly Bayesian perspective as a matter of a priori theoretical preference, such that only Bayesian explanations of observed phenomena are entertained,

must resist the temptation to claim that correspondences between their data and their models *provide evidence for* their broad perspective. Some researchers are careful to limit their claims, presenting supporting evidence only for particular models. Others are not, creating highly flexible Bayesian models to fit patterns of behavior and finding – unsurprisingly – that their data match the predictions of these models and then using that match to make strong general claims about Bayesian brains. But the perspective itself – although it may be a productive environment for thinking about how the mind might work – is not falsifiable, and so evidentiary claims at the level of that perspective are unsupportable.

There is something to be lost when any one perspective dominates the literature. The problem is not just limiting inquiry to the chase for the aesthetic ideal of optimality; it is also concerning if we limit inquiry to the use of models that fall within particular theoretical frameworks. Testing very different kinds of models, rather than working within a single theoretical perspective, is important too. The authors conclude that researchers "should aim for their models to capture all of the systematic weirdness of human behavior rather than preserve an aesthetic ideal" (R&D, sect. 6, para. 1). The quest to capture that systematic weirdness will undoubtedly benefit from theoretical heterodoxy in the field.

How did that individual make that perceptual decision?

David A. Booth

School of Psychology, University of Sussex, Falmer, Brighton BN1 9QH, United Kingdom.
d.a.booth@sussex.ac.uk   www.sussex.ac.uk/profiles/335100

Abstract

Suboptimality of decision making needs no explanation. High-level accounts of suboptimality in diverse tasks cannot add up to a mechanistic theory of perceptual decision making. Mental processes operate on the contents of information brought by the experimenter and the participant to the task, not on the amount of information in the stimuli without regard to physical and social context.

Belief in Bayesian optimality is an example of recurring efforts to escape from the study of basic mechanisms into a world of ideals. Reality takes its revenge as more and more departures from ideal are found and attempts to explain them are refuted or forced into extra detail, as Rahnev and Denison (R&D) show.

Departures from optimality do not need explaining nor can they illuminate mechanisms of perceptual performance. The Bayesian programme fails to reckon with Claude Shannon's insistence that the quantity of information tells us nothing about what the information contains (Shannon & Weaver 1949). Success or failure at meeting a criterion of optimal use of the amount of information in experimenter's stimuli is irrelevant to what is actually going on in making a perceptual decision. First, any

experiment is rich in information of which optimality calculations take no account. Crucial contextual information is explicit as physical arrangements and social signals such as verbal instructions, and implicit in the cultural and material memory that the individual participant brings to each response. Second, the contextual information contents can make some of the content of the experimenter's stimuli dispensable for the perceptual decision. Hence, calculations that do not take context into account can yield an illusory suboptimality. In fact, a substantial number of participants use the whole of the information that each processes from the present and past (e.g., Booth et al. 2011a; 2011b).

R&D identify individuality as one source of suboptimality. Far more than that, disregard of individuality prevents mechanistic understanding. Every perceptual decision is determined by an individual's use of information contained in the cultural and material environment of the test. This causal mediation is the transient structure in discrimination-scaled distances between the individual's present and past output/input values (Booth & Freeman 1993; Booth et al. 2011b). The processed information content varies across individuals and circumstances, and even between particular occasions of the same situation in the same person. Therefore, raw data from individuals should never be averaged before testing a mechanistic hypothesis (Booth 2017; Booth & Freeman 1993; Conner et al. 1988), as is now becoming more widely acknowledged (e.g., Luce 2013). The standard observer models sought by R&D also neglect the idiosyncrasies of information content in the actual mechanisms of perceptual decision making.

Many of results cited by R&D indicate that physical stimuli and context provided by the investigators interact with social context brought by the participant. One of the paradigms reviewed by R&D is psychophysical judgement. In the usual design, the experimenter uses one of a pair of stimuli as a standard of comparison with the other stimulus, which is varied. In fact, each stimulus presentation, whether test or standard, is compared with memory of previous stimuli. The comparative decision is determined by the difference in distances of test and standard from memory of previous exposures in a similar context (e.g., Booth & Freeman 1993; Stewart et al. 2005). The standard stimulus is at best redundant and may even be a source of range-frequency bias (Poulton 1989; cp. Conner et al. 1987 and Booth 2015).

In other words, the experimenter's standard is part of the physical context for the test stimulus on which the perceptual decision is supposedly made. Far from the memory of the first stimulus decaying, as R&D cite, long-term memory is updated at each presentation. That is how pretreatment with a high incidence of positive stimuli reduces a bias to making negative responses, also cited by R&D. Accurate diagnosis of the causal structure of a session of perceptual decisions depends on personal tailoring of stimulus levels to be balanced around the familiar level within the range of Weber fraction constancy (Booth et al. 2010; 2011a; Conner et al. 1988).

R&D review a number of the paradigms showing effects of social context, disguised as personality score. Personality inventories are designed to obscure differences in behaviour between situations to create a stable trait, but they vary with state to unknown extents. To permit mechanistic analysis, each relevant social signal has to be presented at two or more levels, unconfounded with other signals within a session. For example, anxiety about being a lying witness in a detection task might be manipulated by the experimenter indicating that some stimuli test for absence of the signal.

Confidence ratings merely express optimistic behaviour or other habits, rather than giving introspective access to mechanisms of perception. Whether the causation is conscious or unconscious, the only access is through output-input relationships placed on a universal scale of discrimination between present and past.

R&D discuss the variations in the tradeoff between speed and accuracy in reaction times induced by direct instructions or time limits on massed tests, without considering these designs as social pressures. If stimuli provided more scope for using past experience, then the mechanisms of interaction with social context could be investigated. For example, conventional demands for a fast decision or a right answer could be presented at different levels.

Similarly, deficiencies in signal detection cited by R&D could be reduced by more ecological validity of the random background provided for the test signal. If a familiar enough context were provided throughout, the variance of the response distribution would less likely be higher in the presence of the signal. Furthermore, the line of investigation could be relevant to theoretical and practical issues in such contexts (Booth 2015).

R&D review evidence that improper perceptual criteria and supposed misweightings in cue combinations account for suboptimality. If the experiments were designed to be analogs of familiar real life scenarios, personally relevant content of the cues could be tested as mechanisms to explain performance.

Finally, it should be noted that effect of unidentified contextual factors can be measured from the individual's causal structure of discrimination-scaled content in a session of tests. First, the most successful combination of known output/input relationships may account for substantially less than the total variance in the perceptual response. Second, the discrimination distances between present and past of observed features of the situation may not interpolate through the zero from the past. The eccentricity measures the contextual defect in discrimination units or response quantity (Booth et al. 2011b).

The world is complex, not just noisy

Romain Brette

Sorbonne Universités, UPMC Univ Paris 06, INSERM, CNRS, Institut de la Vision, 75012 Paris, France.
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 (R&D) point out that optimality is often ill defined. The fundamental issue is framing perception as a statistical inference problem. Outside of the lab, the real perceptual challenge is to determine the lawful structure of the world, not variables of a predetermined statistical model.

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" (Rahnev & Denison [R&D], sect. 4.1, para. 4). If this is true, then the main issue is not so much whether humans are optimal or suboptimal, 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 up 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. Hence, 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 R&D does not seem to address this issue, as observer models are still cast within the Bayesian framework (their 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 & 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 & Pouget 2004, p. 712) 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.

Although optimal models are useful, optimality claims are not that common

Claire Chambers and Konrad Paul Kording ⓘ

Department of Bioengineering and Department of Neuroscience, University of Pennsylvania, PA 19104.
clairenc@seas.upenn.edu   kording@seas.upenn.edu   http://kordinglab.com/

Abstract

Rahnev & Denison (R&D) argue that human behavior is often described as "optimal," despite many previous findings of suboptimality. We address how the literature handles these concepts and discuss our own findings on suboptimality. Although we agree that the field should embrace the "systematic weirdness of human behavior" (sect. 6, para. 1), this does not detract from the value of the Bayesian approach.

Rahnev & Denison (R&D) rightly point out a fact endorsed by many prominent Bayesian theorists, that human behavior often deviates from the predictions of statistically optimal models (Acerbi et al. 2014b; Beck et al. 2012; Summerfield & Tsetsos 2015). Being data scientists, we want to first trace the way the literature handles the concepts discussed by the authors. Specifically, we want to establish whether there is a "current narrow focus on optimality," as argued by R&D (sect. 1, para. 5). We searched PubMed for journal articles published between January 1, 1995, and April 1, 2018, that addressed optimality in human behavior. We found that although the number of articles that apply Bayesian modeling to human behavior has increased (Fig. 1), only a small proportion of the articles that we analyzed make optimality or near-optimality claims (0.30), and that claims of suboptimality are in roughly equal proportion (0.25). Nor have optimality claims drastically increased relative to the overall number of articles that use Bayesian modeling over the last two decades: 2000s (0.32) and 2010s (0.30). The majority of publications that we analyzed use Bayesian modeling as a tool to understand behavior at the computational level and/or emphasize the suboptimality of behavior. Therefore, optimality claims may not be as prevalent as the authors contend.

We endorse the view of R&D that humans deviate from statistically optimal behavior, and some of our recent perceptual

Figure 1 (Chambers). Number of journal articles as a function of year of publication that use Bayesian models to describe behavior, that make (near) optimality claims, and that make suboptimality claims. We used the following search terms: "(sub)optimality perception," "(sub)optimal perceptual decision making," "Bayesian psychophysics," and "Bayesian observer model behavior/perception." We included in our analysis the 185 articles that reported human psychophysics data and modeling. We manually coded the type of the claim made by each paper.

work has added new ways in which behavior is suboptimal. For example, prior knowledge acquired in a sensorimotor task does not fully generalize to a perceptual decision-making task that requires the same probabilistic treatment of sensory and prior information (Chambers et al. 2017a). Also, we show that young children fail to incorporate experimentally imposed priors into their sensorimotor estimates and learn to approximate efficient use of statistical information during development (Chambers et al. 2017b). If the underlying neural processes are innately "Bayes optimal," then behavior should reflect this under different conditions and without learning. We have observed this not to be the case when the decision modality changes in decision making in adults and during child development. These findings underline the importance of assessing the generality of claims concerning the efficient use of statistical information in human decision making and add to the literature on suboptimal perceptual decision making discussed by R&D.

R&D advocate for a change in the culture of how behavioral data are modeled. They argue that much existing work is limited because it includes one ad hoc model for one set of experiments. We agree that this is problematic because based on this ad hoc approach we cannot know how well optimal models generalize or "transfer" to other situations (Maloney & Mamassian 2009). The authors recommend that in cases of findings of suboptimality, the assumptions of each model component (likelihood, prior, cost function, decision rule) should be examined, and that in cases of good agreement with optimal models, model performance should be examined under different conditions and tasks. We agree with the authors' prescriptions on a more rigorous approach to following up on previous claims and believe that this practice will help the field to develop models that provide more complete accounts of human behavior.

However, it should be emphasized that Bayesian optimal models remain an important tool for building computational models of behavior. Uncertainty exists in the outside world, and we do well in everyday tasks by taking uncertainty into account (Kersten et al. 2004; Körding & Wolpert 2006; Vilares & Kording 2011). Alternative non-optimal frameworks, like the

"bag of tricks" (Ramachandran 1990) mentioned by the authors and heuristic models (Gigerenzer & Gaissmaier 2011; Kahneman et al. 1982a) do not contain explicit formulations of how probabilistic information should be combined and do not as easily capture how we deal with uncertainty as Bayesian models do. The Bayesian framework has other advantages as a computational framework. It provides an important benchmark for human performance that can constrain non-optimal models. It is a compact mathematical framework that avoids assumptions about implementational details. It is transparent: The components of the model (prior, likelihood, cost function, decision) must be declared (Griffiths et al. 2012). Although we agree that the usefulness of labeling of behavior as "optimal" is questionable, we maintain that Bayesian modeling is still an important tool for the computational understanding of behavior.

Serial effects are optimal

Guido Marco Cicchini[a] and David C. Burr[b]

[a]CNR Institute of Neuroscience, Pisa, 56124, Italy and [b]Department of Neuroscience, Psychology, Pharmacology and Child Health, University of Florence, 50139, Italy.
cicchini@in.cnr.it   www.pisavisionlab.org/cicchini
dave@in.cnr.it   www.pisavisionlab.org/burr

Abstract

In the target article, Rahnev & Denison (R&D) use serial effects as an example of suboptimality. We show here that serial effects can be beneficial to perception, serving to reduce both error and response times in a near-optimal fashion. Furthermore, serial effects for stable attributes are positive, whereas those for changeable attributes are negative, demonstrating that they are engaged flexibly to optimize performance.

We read with great interest the article by Rahnev & Denison (R&D), reporting both a wide coverage of the issue of optimality in perception, as well as the many instances in which optimality has been hard to prove. One example of non-optimality for the authors is *serial dependence*, the influence of previous stimuli on current responses, in a sequential task (Burr & Cicchini 2014; Cicchini et al. 2014; Cicchini & Kristjánsson 2015; Fischer & Whitney 2014; Frund et al. 2014; Liberman et al. 2014). The authors speculate that because the experimental setting prescribes independence between trials, it is suboptimal to carry over information from the previous trial. The only possibility they see is that perhaps the perceptual system is attuned to a general rule of continuity, which accidentally spills over into laboratory performance.

We encountered serial effects while studying perception of numerosity, finding that subjects were strongly biased toward the previous estimate, by up to 20% (Cicchini et al. 2014). Importantly, at higher numerosities – where sensory resolution is lower – the serial effects were larger. This prompted us to investigate the effect with a standard Bayesian model in which the previous sensory experience can be considered an extra source of information (Fig. 1).
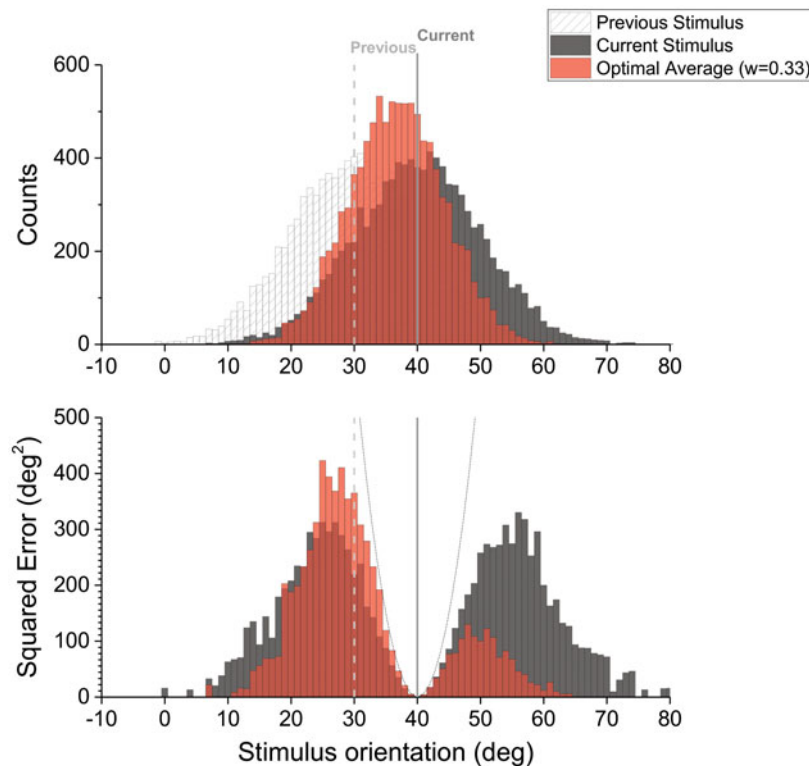
Figure 1 (Cicchini & Burr). Serial dependence can be optimal. We illustrate the behavior of a noisy observer (σ = 10°) who is bound to estimate a stimulus at 40° orientation and can make optimal use of the previous trial which was at 30°. Response distributions are displayed in the top panel (black, memoryless model; red, model with serial dependence) and show a slight shift of responses toward the previous trial with a tightening of the distribution. Bottom panel shows the histograms of squared error cost: The overall error of the model taking advantage of serial dependencies is smaller than that of the memoryless model.

Simulations show that in the case of similar successive stimuli, this strategy is beneficial, as the uncertainty associated with the current judgment may benefit from integrating information from the past. According to our model, previous sensory experience should be weighted, taking into account the sensory resolution of the current and previous stimulus (σ_{curr}, σ_{prev}) along with the difference in intensity between two presentations (*d*): $w_{prev} = \frac{\sigma^2_{curr}}{\sigma^2_{curr} + \sigma^2_{curr} + d^2}$. This simple rule states that whenever there is multiple, congruent information, it is beneficial to blend it, without needing assumptions of continuity or meta-priors.

This model of optimal performance provided a good fit to the numerosity data set (Cicchini et al. 2014). We went on to apply it successfully to orientation reproduction tasks (Cicchini et al. 2017; Cicchini et al. 2018). In this experiment, we collected individual measures of precision and predicted (with zero degree of freed model and no further assumptions) subject behavior, obtaining an excellent fit both of the amount of serial dependence, as well as the range of orientation differences over which the effect occurs (Cicchini et al. 2018). With the same data set, we also measured the benefit of serial dependence. We compared trials that were preceded by an identical stimulus (maximal dependence) and those preceded by a stimulus 45 degrees apart (when serial dependence waned). As predicted, we found a reduction of the squared error, about 45%, when serial dependence was maximal. We also compared response times and found that they were up to 60 ms faster for identical than for differing preceding stimuli. Overall, the two results show that the serial presentation of similar stimuli led to a genuine increase of information in the system. This latter result was totally unexpected as our model was developed starting from optimal cue integration literature (Alais & Burr 2004; Beierholm et al. 2008; Cicchini et al. 2012; Ernst & Banks 2002; Jazayeri & Shadlen 2010; Roach et al. 2006) and was meant to optimize response error without considering time limits.

A final example that serial dependence does not result merely from a passive, no-optimal stickiness of the system is the demonstration that stimuli can induce either positive or negative serial dependencies, depending on their usefulness to the task in hand. Taubert et al. (2016) asked subjects to judge both the gender and expression (happy/sad) of sequentially presented faces. Strong positive serial dependence was found for gender – a stable attribute of a face that does not change over time. In the same experiment, with the same face stimuli, negative serial dependence was observed for expression, a labile attribute that changes frequently and rapidly, where the information is often in the change. Carrying over signals of expression from previous exposure would be of little help for the task in hand, and the system does not do it. This is a very clear demonstration that serial dependence is not an automatic result of the sluggishness of the system, but an active and flexible strategy to improve, and possibly optimize, perception.

Overall, we believe that the Bayesian framework has several merits for brain science. First, it encourages researchers to think of the brain as a statistical observer who accumulates information, and second, because it has helped to discover several strategies to obtain the same performance with fewer resources, such as removal of redundancy from retinal images to be transmitted via a small optic nerve. We agree with the authors that optimality is a loose concept and cannot be the only principle working in the brain; however, it has proved to provide an excellent framework with which to uncover the mechanisms of an organism that started evolving 200 million years ago.

LPCD framework: Analytical tool or psychological model?

David Danks

Departments of Philosophy and Psychology, Carnegie Mellon University,
Pittsburgh, PA 15213.
ddanks@cmu.edu
https://www.cmu.edu/dietrich/philosophy/people/faculty/danks.html

Abstract

The target article uses a mathematical framework derived from Bayesian decision making to demonstrate suboptimal decision making but then attributes psychological reality to the framework components. Rahnev & Denison's (R&D) positive proposal thus risks ignoring plausible psychological theories that could implement complex perceptual decision making. We must be careful not to slide from success with an analytical tool to the reality of the tool components.

Throughout the target article, Rahnev & Denison (R&D) employ the likelihood function, prior, cost function, and decision rule (LPCD) framework, as they analyze the (sub)optimality of perceptual decision making in terms of the *likelihood* of the world state (i.e., stimulus) given the internal "signal" (i.e., perceptual experience); the *prior probability* of that world state; the *cost function* of an action given the world state; and the *decision rule* used to determine what action to take in light of the other three components. This framework is mathematically derived from models of Bayesian decision making and mirrors much of the language – mathematical and conceptual – used by proponents of optimality in perceptual decision making. Over the course of their target article, however, R&D shift between two different interpretations of this mathematical framework and thereby provide a positive proposal that does not follow from their empirical survey.

For the first four sections of the target article, R&D use the LPCD framework as an analytical tool to decompose the complex behavioral functions identified in the many experiments that they survey. In general, people's decisions are a complicated function of the current environment, task demands, internal constraints (e.g., memory or attention), perceptual history, and more. R&D briefly discuss the complexity, and so possible computational intractability, of the *optimal* function, but of course, our *actual* perceptual decision functions are also incredibly complicated. The LPCD framework provides a way to decompose this complicated behavioral function into tractable, interpretable components.

This use of the LPCD framework is analogous to a Fourier decomposition of a complex waveform into a set (possibly infinite) of component sine waves of varying frequency and amplitude. Importantly, those component sine waves need not correspond to any underlying causal or generative processes; they are simply a more tractable way to understand the input waveform. Similarly, we can use the LPCD framework as an analytical tool to decompose people's perceptual decision functions into those four components, without thereby thinking that the LPCD components correspond to anything in the perceptual or cognitive system. Those mathematical components are only relatively simpler functions that jointly provide accurate (and sometimes, rationally defensible) predictions.

In section 5.1, however, R&D shift to using LPCD in a very different way. In particular, they start to interpret the four components realistically, as elements that are really present in the mind/brain. For example, they propose that we should search for hypotheses about the content, functional form, or source of particular LPCD components, where those are understood to be represented internally (LPC) or actually used to make a decision (D). That is, R&D no longer use LPCD solely as an analytical tool, but rather interpret the components as psychologically real. For example, many of the hypotheses in their Table 1 treat the LPCD components as discrete psychological or neurally encoded elements that can be independently measured and manipulated by researchers. They are not simply a useful way of mathematically decomposing a complex behavioral function, but are presumed to be legitimate targets of direct scientific inquiry.

R&D might well be right that the LPCD components are psychologically real and can be shaped or influenced in various ways. For example, some of the cited papers that develop optimality models of perceptual decision making also provide evidence for the psychological reality (and modularity) of one or another element in a Bayesian model. However, R&D do not base their realist commitments on those studies, but rather on extrapolation from the structure they propose for observer models. But if we do not have a background assumption of optimality (or near optimality), then we need not assume that perceptual decision making decomposes, in any realist sense, into the LPCD components. Mathematical decompositions do not necessarily provide a guide to underlying causal or psychological processes. For example, the function $f(x, y) = x^2 - y^2$ can be mathematically decomposed into x^2 and y^2 "modules," but that does not imply that the underlying system actually computes x^2 and y^2 in distinct components and then subtracts the results. The system could just as easily have components that compute $(x + y)$ and $(x - y)$ and then multiply the results. Or the function could be actually computed in many other ways.

Moreover, the realist commitments in section 5 are not idle, but rather can have real impacts. As R&D note, there are many different types of models that we might consider, not only Bayesian or LPCD-based models. Nonetheless, R&D's positive proposal involves viewing all of these proposals through the LPCD lens and, therefore, privileging that particular decomposition of the complex perceptual decision-making function, even if the mind/brain does not actually work that way. Of course, all science requires some theoretical commitments, but there is well-known value in having variation across the scientific community. In contrast, R&D propose that the LPCD components, in the form of standard observer models, should be the default theoretical framework for this research domain.

R&D demonstrate the utility of the LPCD framework as a tool for identifying suboptimal behaviors. They do not, however, provide reason to prefer substantive psychological hypotheses that are expressed with LPCD components, and so their positive proposal about how to advance the science risks placing unnecessary constraints on our theorizing.

Optimality is critical when it comes to testing computation-level hypotheses

Laura S. Geurts, Andrey Chetverikov†,
Ruben S. van Bergen†, Ying J. Zhou, Andrea Bertana
and Janneke F. M. Jehee

Donders Institute for Brain, Cognition and Behavior, Radboud University
Nijmegen, 6525 EN Nijmegen, The Netherlands.
www.jeheelab.org a.bertana@donders.ru.nl a.chetverikov@donders.ru.nl
janneke.jehee@donders.ru.nl l.geurts@donders.ru.nl
r.vanbergen@donders.ru.nl y.zhou@donders.ru.nl

Abstract

We disagree with Rahnev & Denison (R&D) that optimality should be abandoned altogether. Rather, we argue that adopting a normative approach enables researchers to test hypotheses about the brain's computational goals, avoids just-so explanations, and offers insights into function that are simply inaccessible to the alternatives proposed by R&D.

Although we concur with Rahnev & Denison (R&D) that proving the optimality of human behavior should never be a goal in and of itself, we disagree that optimality should be discarded altogether. Rather, we argue that an ideal observer approach is the only valid approach for understanding the nature of the computational problems that the brain is trying to solve (Marr 1982).

Understanding the brain's computational goals is key to understanding human behavior. Knowing what the system is trying to do, and why, can inform hypotheses about the algorithms that the brain uses to achieve these goals, and how these algorithms are implemented in neurons and connectivity (see Marr 1982). To test computational theories, researchers typically translate hypotheses into predictions that can be compared against human data. Although this translation can take many forms, we argue that the best recipe for converting computational goals into predictions is provided by the normative approach, because this approach rightly specifies what kind of behavior participants ought to display under the hypothesized goals. For any other description, the translation from hypothesis into prediction may be imprecise, making it impossible to determine whether the hypothesis or rather its translation was incorrect, when predicted behavior fails to match human data.

Consider, for example, visual orientation perception. Human orientation perception tends to be biased toward vertical and horizontal (cardinal) orientations, such that a tilted (oblique) stimulus appears slightly more cardinal than it really is (Tomassini et al. 2010). In a seminal study, Girshick et al. (2011) hypothesized that such biases arise because observers, when judging stimulus orientation, use knowledge about the distribution of orientations in the natural environment (where cardinal orientations are predominant). To address this hypothesis, the authors presented human participants with random orientation stimuli, asking them to estimate the orientations they had seen. Girshick et al. (2011) then specified the behavior of an ideal observer who combines noisy sensory measurements of a stimulus with knowledge about its prior probability. By using this approach to model human behavior, the authors were able to characterize the prior beliefs about orientation that the participants applied in their perceptual estimates. Interestingly, this prior distribution resembled the actual distribution of orientations in the environment, corroborating the theoretical predictions.

This study nicely illustrates how the ideal observer framework can be exploited to not only describe, but also explain behavior, by starting with a specific hypothesis about what the system is trying to do (minimizing perceptual error in a natural environment by taking into account natural orientation statistics) and translating this hypothesis into a computational model that predicts behavior. Precisely because human behavior matched that of the ideal observer, rather than some arbitrary formulation, the findings provided strong evidence for the computational theory and offered insights into function that are simply inaccessible to the non-normative alternatives (e.g., "bag of tricks" or "neural network" models; sect. 5.2, para. 1) proposed by R&D.

Note that in the example discussed previously, human behavior was, in fact, suboptimal for the experimental situation in which the participants were tested: When presented with a uniform distribution of orientation stimuli, it would make little sense for an observer to show biased judgments toward vertical or horizontal orientations. Yet, the ideal observer framework enabled the researchers to address what kinds of knowledge the observers did bring to the task, by comparing human task behavior against that of an ideal observer using prior knowledge of the orientation statistics of the natural environment.

We argue that it is precisely these kinds of situations that are informative. The ideal observer framework provides a powerful tool to implement hypotheses and, as such, offers insight into the knowledge and goals that a human observer brings to the task at hand. When human behavior does not conform to the theory's predictions, one should rather conclude that the theory makes incorrect assumptions about the computational goals of the system and try to improve the theoretical assumptions, instead of refuting the normative tools that merely implemented the theory.

The role of (bounded) optimization in theory testing and prediction

Andrew Howes[a] and Richard L. Lewis[b]

[a]School of Computer Science, University of Birmingham, Edgbaston, Birmingham B15 3TT, United Kingdom and [b]Department of Psychology, University of Michigan, Ann Arbor, MI 48109.
HowesA@bham.ac.uk https://www.cs.bham.ac.uk/~howesa/
rickl@umich.edu http://www-personal.umich.edu/~rickl/

Abstract

We argue that a radically increased emphasis on (bounded) optimality can contribute to cognitive science by supporting prediction. Bounded optimality (computational rationality), an idea

that borrowed from artificial intelligence, supports a priori behavioral prediction from constrained generative models of cognition. Bounded optimality thereby addresses serious failings with the logic and testing of descriptive models of perception and action.

The target article claims that cognitive science "should abandon any emphasis on optimality or suboptimality" (Rahnev & Denison [R&D], sect. 1, para. 4). In contrast, we argue that a radically increased emphasis on (bounded) optimality is crucial to the success of cognitive science. This commentary draws on a significant literature on bounded optimality, an idea that is borrowed from artificial intelligence (Russell and Subramanian 1995). It argues that comparing the behavior of bounded optimal (also known as computationally rational) models to human behavior is a better way to progress the science of the mind than the authors' "observer models." Observer models are a form of descriptive model. In contrast, bounded optimal models can be predictive and explanatory (Howes et al. 2009; 2016; Lewis et al. 2014).

Lewis et al. (2014) proposed computational rationality as an alternative to the standard use of optimality (rational analysis) in cognitive science. Computational rationality is a framework for testing mechanistic theories of the mind. The framework is based on the idea that behaviors are generated by cognitive mechanisms that are adapted to the structure of the external environment and also to the structure of the mind and brain itself. In this framework, theories of vision, cognition, and action are specified as "optimal program problems," defined by an adaptation environment, a bounded machine, and a utility function. This optimal program problem is then solved by optimization (one of the utility maximizing programs is identified and selected), and the resulting behavior is compared to human behavior. Success is not used to conclude that people are either optimal or suboptimal, but rather, success indicates evidence in favor of the theory of the environment, bounded machine, and utility function.

One example of how computational rationality can be used to test theories was provided by Howes et al. (2009). Consider an elementary dual-task scenario in which a manual response must be given to a visual pattern and a verbal response to the pitch of a tone. This task, known as a psychological refractory period (PRP) task, has been used extensively in an effort to understand whether cognition is strictly serial or whether it permits parallel processing (Meyer and Kieras 1997). Although many theories had been proposed prior to Howes et al.'s work, they were sufficiently flexible that both serial and parallel models could be fitted to a large range of PRP data. A key source of the flexibility was the cognitive program (the strategy) by which elementary cognitive, perceptual, and motor processes were organized. Before Howes et al. (2009), different (but plausible) programs were used to fit models to a wide range of data irrespective of whether cognition was assumed to be serial or parallel. Similarly, in perceptual decision-making tasks, the decision rule (criterion or threshold) is a simple example of a strategy or program. Freed from the requirements that the decision rule is optimal for the defined problem, almost any data might be fitted.

For the PRP task, Howes et al. used separate computational theories of the serial and parallel bounded machine and derived the optimal program for each. The optimal program was used to predict behavior. Optimality was not under test, but rather it was used as a principled method of selecting cognitive programs, independently of the observed data. What was under test was whether either serial and/or a parallel cognition could predict the observed behavior. On the basis of detailed quantitative analysis, their conclusion was that it was the serial theory that offered a better explanation of the data. Similarly, Myers et al. (2013) tested the implications of noise in peripheral vision for human visual search. They found that a particular model of noise in peripheral vision predicts well-known visual search effects.

Illustrated with the previous examples and supported by the analysis of Lewis et al. (2014), we can see that computational rationality has the following similarities and differences to the authors' observer models:

Neither computationally rational models nor observer models set out to test whether people are optimal or suboptimal. In both cases, the aim is to test the collective set of theoretical assumptions in the model.

Unlike observer models, the computational rationality framework assumes that a program (a strategy or decision rule) will be determined, by the scientist, using an optimization algorithm. In so doing, it allows a quantitative and causal relationship to be established between theoretical assumptions and behavior. In contrast, with observer models, the analyst is permitted to pick any "plausible" decision rule (step 2 of Box 1 of R&D). As a consequence, despite their desire to reduce the perceived flexibility of the optimality approach (Bowers & Davis's [2012a] just-so stories), R&D permit potentially extensive flexibility through the informal notion of plausibility.

By virtue of the fact that programs are determined through optimization, computational rationality supports prediction, whereas observer models are descriptive. For example, in Howes et al. (2009), model parameters (e.g., noise level) were calibrated to single-task scenarios; optimal strategies were determined for dual-task scenarios; and test variables, including dual-task duration, were predicted by executing the optimized model. In contrast, observer models are descriptive by virtue of the admission of plausible decision rules. The potential arbitrariness of plausible decision rules dooms step 4 (Box 1), "specify how the conclusions depend on the assumptions," to be an informal process of constructing just-so stories. In other words, observer models cannot be said to make predictions if the analyst must intervene to determine what is and what is not plausible.

In summary, although we agree with R&D that the focus of the behavioral sciences should be on testing theories of how the mind processes information, we believe that optimization, used by a scientist as a tool to determine the adaptive consequences of theoretical assumptions, offers a better way forward for psychological science than the proposed (descriptive) observer models.

Model comparison, not model falsification

Bradley C. Love ⬤

Experimental Psychology, University College London, London WC1H 0AP, United Kingdom.
b.love@ucl.ac.uk   http://bradlove.org

Abstract

Systematically comparing models that vary across components can be more informative and explanatory than determining whether behaviour is optimal, however defined. The process of model comparison has a number of benefits, including the possibility of integrating seemingly disparate empirical findings, understanding individual and group differences, and drawing theoretical connections between model proposals.

Determining whether behaviour is optimal can be difficult because what is optimal is often a matter of debate. For example, optimality can be defined in terms of task, related real-world environments, hypothesised evolutionary environments, short- versus long-term rewards, and so forth. Furthermore, notions of optimality can be expanded to respect resource limitations, such as constraints specified in terms of energy, time, effort, or cognitive resources. More thought may go into choosing a measure of optimality than in evaluating how people compare to the chosen yardstick. Many of these issues recapitulate criticisms of rational approaches to understanding perception and cognition (Jones & Love 2011).

Sensibly, Rahnev & Denison (R&D) argue for moving away from notions of optimality. Instead, they specify various ways in which people can be suboptimal. Although one can argue about the particular set of components identified as sources of suboptimal decision making, the basic approach is promising. Careful comparison of models has the potential to identify the root causes of behaviour as opposed to making a blanket statement about a debatable notion of optimality.

Indeed, one could go further and simply advocate for model comparison without considering optimality. Although thinking about optimality can be a useful starting point for developing models and evaluating human performance, a strong focus can be restrictive. The question of whether people are optimal invites a Popperian odyssey to falsify the claim. Unfortunately, accepting or rejecting a hypothesis in isolation is usually not very informative or explanatory. Alternatively and perhaps more productively, one could specify a rich set of hypotheses, formalise these hypotheses as models, and perform a proper model comparison. The outcome of such a process is the best available explanation (i.e., model) of the data.

Model comparison offers a route for model and theory development. New model proposals can draw on past models that have enjoyed success. For example, in the category learning literature, the lineage of models stretches across decades. Past work has influenced my own proposals (e.g., Love et al. 2004). As new sources of data become available, such as brain imaging data, model comparison approaches can embrace these new data sources (Mack et al. 2013).

Finally, model comparison offers a number of advantages for our science. Model comparison requires specifying what the relevant data are. In doing so, the scope of models becomes clearer. Formalising theories as models of course has its own advantages in terms of making assumptions clearer, enabling quantitative prediction in novel circumstances, characterising individual and group differences in terms of fitted parameter values, directing future experimentation, and identifying broad principles that span data sets and models. Overall, model comparison offers a path to explain behavioural phenomena that can be more integrative and explanatory than blanket statements about optimality.

Identifying suboptimalities with factorial model comparison

Wei Ji Ma

Center for Neural Science and Department of Psychology, New York University, New York, NY 10003.
weijima@nyu.edu

Abstract

Given the many types of suboptimality in perception, I ask how one should test for multiple forms of suboptimality at the same time – or, more generally, how one should compare process models that can differ in any or all of the multiple components. In analogy to factorial experimental design, I advocate for factorial model comparison.

Rahnev & Denison (R&D) identify many possible sources of (apparent) suboptimality in behavior, including capacity limitations, incorrect observer assumptions about stimulus statistics, heuristic decision rules, and decision noise or criterion jitter (their Table 1). They urge the field to test the collective set of these hypotheses. I strongly support this message; for example, in recent work, we tested no fewer than 24 alternatives to the optimal decision rule (Shen & Ma 2016). However, the research agenda as a whole encounters a practical challenge: It is in most cases impossible to test one hypothesis at a time, as it would require the experimenter to make arbitrary choices regarding the other hypotheses.

The rational approach, which is taken surprisingly rarely in the study of perception and cognition, is to consider each hypothesis as a factor in a model (Keshvari et al. 2012; van den Berg et al. 2014). Each factor could be binary or multivalued; for example, the factor "decision rule" could take the values "optimal," "heuristic 1," and "heuristic 2." Just like factorial design is a cherished tool in experimentation (Fisher 1926), models can also be combined in a factorial manner: Every logically consistent combination of values of the different factors constitutes a model that should be tested. If there are n binary factors, there will be up to 2^n models. The goodness of fit of each model is then evaluated using one's favorite metric, such as the Akaike information criterion (AIC; Akaike 1974) or leave-one-out cross validation (Vehtari et al. 2017).

In factorial model comparison, ranking all individual models is usually not the end of the road: often, one is less interested in the evidence for individual models than in the evidence for the different levels of a particular factor. This can be obtained by aggregating across "model families." A model family consists of all models that share a particular value for a particular factor (e.g., the value "heuristic 1" for the factor "decision rule"). In the binary example with 2^n models, each model family would have 2^{n-1} members. The goodness of fit of a model family is then a suitably chosen average of the goodness of fit of its members. In a fully Bayesian treatment, this averaging is marginalization. If AIC is the metric of choice, one could average the AIC weights (Wagenmakers & Farrell 2004) of the family members (Shen & Ma, in press). Finally, one could represent each family by its best-performing member (van den Berg et al. 2014).

An alternative to factorial model comparison is to construct a "supermodel," of which all models of interest are special cases (Acerbi et al. 2014a; Pinto et al. 2009). For example, an observer's belief about a Gaussian stimulus distribution with fixed mean and variance could be modeled using a Gaussian distribution with free mean and variance. Then, all inference amounts to parameter estimation, for which one can use Bayesian methods. In some cases, however, a factor is most naturally considered categorical – for example, when comparing qualitatively distinct decision rules.

My lab performed factorial model comparison for the first time in a study of change detection, crossing hypotheses about the nature of encoding precision with ones about observer assumptions about encoding precision, and with ones about the decision rule (Keshvari et al. 2012). In this case, an optimal-observer model with variable encoding precision and right observer assumptions won convincingly.

Factorial model comparison is no silver bullet. It is easy to end up with statistically indistinguishable models in the full ranking. In a study of the limitations of visual working memory, we crossed hypotheses about the number of remembered items with ones about the nature of encoding precision, and with ones about the presence of non-target reports (van den Berg et al. 2014). This produced 32 models, many of which were indistinguishable from others in goodness of fit. Family-wise aggregation helped to draw conclusions, but even that might not always be the case. Such non-identifiability, however, is not a problem of the method but a reflection of the difficulty of drawing inferences about multicomponent processes based on limited behavioral data. As we wrote in van den Berg et al. (2014, p. 145), "Factorially comparing models using likelihood-based methods is the fairest and most objective method for drawing conclusions from psychophysical data. If that forces researchers to reduce the level of confidence with which they declare particular models to be good representations of reality, we would consider that a desirable outcome."

Satisficing as an alternative to optimality and suboptimality in perceptual decision making

Antonio Mastrogiorgio[a] and Enrico Petracca[b,c]

[a]Department of Neurosciences, Imaging and Clinical Sciences & CeSI-MeT, University of Chieti-Pescara, 66100 Chieti Scalo (CH), Italy; [b]Department of Economics (IRENE), University of Neuchâtel, 2000 Neuchâtel, Switzerland and [c]School of Economics, Management and Statistics, University of Bologna, 40126 Bologna, Italy.
mastrogiorgio.antonio@gmail.com   enrico.petracca2@unibo.it
https://sites.google.com/site/embodiedrationality/

Abstract

Rahnev & Denison's (R&D) critique of optimality in perceptual decision making leads either to implicitly retaining optimality as a normative benchmark or disregarding the normative approach altogether. We suggest that "bounded rationality," and particularly the "satisficing" criterion, would help dispense with optimality while salvaging normativity. We also suggest that satisficing would provide a parsimonious and robust explanation for perceptual behavior.

For decades, the field of judgment and decision making (under the constant impulse coming from economics) has been struggling over how to overcome optimality as a default framework for decision making. With the benefit of this experience, we comment on Rahnev & Denison's (R&D) article, which may be traversing through similar hurdles in the field of perceptual decision making. Symptomatic of the fact that R&D (and the whole field of perceptual decision making) are in the early stages of a struggle with optimality is that they are inevitably led to call perceptual behavior that violates "optimality" suboptimal. By doing so, they compel themselves within the strictures of the optimality/suboptimality dichotomy, making their subsequent plea to overcome optimality – although repeatedly stated – methodologically unconvincing. R&D's plea for the construction of a standard observer model closely resembles well-known attempts to descriptively amend optimality frameworks in the face of evidence of optimality violations (e.g., Kahneman & Tversky 1979), maintaining, however, optimality as a normative benchmark, even if only for diagnostic purposes. On a more radical reading, R&D would reject normativity in perceptual decision making altogether, in favor of a purely descriptive account of both optimal and suboptimal behavior. We argue that both of these ways would methodologically result in a dead end. This would mainly be attributable to R&D's reliance on the hidden assumption that conflates optimality and normativity (i.e., the assumption that considers optimality as the sole possible normative benchmark for perceptual decision making). In this commentary, we advance the idea that "bounded rationality," and particularly the notion of "satisficing" (Simon 1955; 1956), are able to altogether overcome optimality by (1) providing a benchmark that rejects optimality but salvages normativity in perceptual decision making and (2) proposing a more parsimonious and robust explanation of perceptual behavior.

With regard to the first point, R&D's contention that "[bounded rationality models] still place the greatest emphasis on the optimality of the decision rule" (sect. 5.1, para. 1), while mentioning Herbert Simon (1955) and Gigerenzer and Selten (2001) in support of their statement, is particularly striking. Simon (1956) was resolutely against any form of optimality, so much that he founded the notion of "bounded rationality" on a completely new criterion called "satisficing" (a neologism conflating "satisfy" and "suffice"). As Simon (1996) stated, "Many … have argued that the gap between satisfactory and best is of no great importance, hence the unrealism of the assumption that the actors optimize does not matter; others, including myself, believe that it does matter, and matters a great deal" (p. 29). Given the naturally limited availability of time, information, and computational capacity to make decisions in real-world environments, agents' decision rules cannot optimize, but rather obey a criterion of satisfaction, operationally meaning that satisfaction is achieved when a certain threshold, or "aspiration level," is reached. More recently, Gigerenzer et al. (1999), building upon Simon's framework, have emphasized its ecological traits, maintaining that a decision rule is rational "to the degree that it is adapted to the structure of an environment" (p. 13). In both Simon's and Gigerenzer's frameworks, "adaptation" is the keyword, as it provides an alternative normative framework to decision making (e.g., Hands 2014). Notably, this adaptation framework rejects the very idea that adaptive criteria, such as satisficing, could be reduced to some form of optimization (typically, optimization

under constraints) (Gigerenzer 2004). R&D's misconception of Simon's and Gigerenzer's theories – particularly when they say that "influential theories [of bounded rationality] postulate that evolutionary pressures produced heuristic but useful, rather than normative, behavior" (sect. 5.3, para. 3) – has to be traced back to a failure to figure out that bounded rationality's and satisficing's normative content lies in adaptation. This is most unfortunate, as bounded rationality's normative core would greatly help perceptual decision making overcome optimality in any residual form, without abdicating to normativity.

With regard to the second point, satisficing could also provide a more parsimonious and robust explanatory framework for perceptual behavior with respect to optimality and suboptimality. To see how, we specifically consider the "diffusion model," typically used to explain the speed/accuracy tradeoff in perceptual decisions. As it is commonly described, "the diffusion model assumes that decisions are made by a noisy process that accumulates information over time from a starting point toward one of two response criteria or boundaries.… When one of the boundaries is reached, a response is initiated" (Ratcliff & McKoon 2008, p. 875). It is surprising how closely this model's mechanism resembles satisficing models of decision making introduced by Simon (1955). The descriptive consistency of these two classes of models can be used, we suggest, to parsimoniously explain the puzzling evidence that some subjects set decision thresholds optimally while others only suboptimally in the same task (Bogacz et al. 2010). In principle, evidence of individual differences in threshold setting is consistent with the idea that individual variability falls within a common adaptive interval. In this adaptive interpretation, the main question addressed by diffusion models should not be whether thresholds are optimal or suboptimal, but whether they are adapted or not to a given perceptual task. In this line of argument, although ecology is mentioned as an explanation of optimality in the speed/accuracy tradeoff (e.g., Bogacz 2007), ecological arguments seem to be missing in R&D's discussion of suboptimality. As certain perceptual tasks are distinctly oriented to accuracy and others distinctly oriented to speed, ecological arguments dictate whether speed and accuracy should be treated together or separately (Todd & Gigerenzer 2003, p. 151). The usual explanation that first assumes the speed/accuracy tradeoff and then describes either speed-oriented or accuracy-oriented behavior as "limiting cases" can be unnecessarily complex from an ecological point of view. More generally, as analysis of optimality/suboptimality typically explains perceptual behavior in tasks that are relatively simple, it can say little regarding more complex tasks, such as those considered by Simon (Bogacz et al. 2010, p. 888). In these latter cases, there are ecological reasons to argue that satisficing procedures may provide a robust explanation for a wider range of perceptual tasks (e.g., Martignon & Hoffrage 2002).

Optimality is both elusive and necessary

Joachim Meyer

Department of Industrial Engineering, Tel Aviv University, Tel Aviv 6997801, Israel.
jmeyer@tau.ac.il   jmeyer.tau.ac.il

Abstract

Optimality of any decision, including perceptual decisions, depends on the criteria used to evaluate outcomes and on the assumptions about available alternatives and information. In research settings, these are often difficult to define, and therefore, claims about optimality are equivocal. However, optimality is important in applied settings when evaluating, for example, the detection of abnormalities in medical images.

A long history of research supports the notion that human decisions may not be optimal. Many of the early studies dealt with probability learning and "probability matching." In these studies, decision makers often choose the alternative with the higher payoff probability according to this probability (e.g., when there is a 70% chance to get a positive payoff from choosing an alternative, participants tend to choose it 70% of the time). The optimal strategy in such experiments is always to choose the alternative with the higher payoff probability. Kenneth Arrow (1958, p. 14) stated 60 years ago, referring to probability matching, that "the remarkable thing about this is that the asymptotic behavior of the individual, even after an indefinitely large amount of learning, is not the optimal behavior."

However, probability matching may actually be optimal. This is the case if it is used in an environment with possible changes and competition for resources (Gallistel 2005). This description probably characterizes the vast majority of environments outside the experimental psychology lab. Therefore, probability matching may be the optimal strategy in most settings, and using it in a lab experiment does not imply non-optimality of human decisions. Hence, we need to define the criteria according to which one evaluates a decision. A decision may be optimal under some assumptions and will be non-optimal under different assumptions.

The optimality of decisions does not only depend on the assumptions on which evaluations are based. Optimality is also always judged from a particular point of view. Perceptual decisions may seem non-optimal when evaluated from a "god's eye view," knowing the true probabilities of events. However, as Rahnev & Denison (R&D) point out, prior expectations affect judgments. If a person believes certain events are more likely than others and bases her decisions on this belief, the decisions may very well be optimal, considering the information that is available when the decision is made.

Similarly, the decision maker's experience in the experiment determines likelihood estimates. In binary classification experiments, the probability of events is binomially distributed. As long as the number of observations is relatively small, assessments of the probabilities of events may differ greatly from the "true" probabilities. Furthermore, even if the person observed a large number of events, the assessed likelihood depends on the memory for the events (the relative number of true and false positive and true and false negative classifications). Not all of these events may be equally salient in memory, leading to possibly biased likelihood estimates. An optimal response to biased likelihood estimates will seem non-optimal.

Hence, the notion of optimality is not very informative. Instead, an attempt should be made to model the cognitive processes that lead to the perceptual decisions, in line with R&D's suggestions. Such models should consider the context in which decisions are made, the person's prior expectations, and the

events the person encounters. They should also take into account the properties of the memories these events leave. The "optimality" criterion can be some benchmark to which one can compare decisions, but unless we believe people have some supernatural ability for clairvoyance, we should not expect people to reach this optimality criterion.

However, in some conditions, optimality is critical, and the criteria for optimality are relatively clearly defined. This is often the case outside the laboratory, as when clinicians make decisions regarding the existence of malignancies in medical images. The optimality of decisions in such tasks depends on the costs and benefits of different outcomes and the likelihood of malignancies in a population. If such decisions deviate systematically from optimality, some steps can be taken to lower the discrepancies (e.g., provide training, change procedures). Alternatively, it may be possible to provide human decision makers with aids that help in the detection process (based on image analyses, etc.). If such aids are available, the question of whether people assign optimal weights to the information from these aids is of major importance. In fact, it turns out that people often assign insufficient weight to valid aids, and they over-rely on their own perceptions (as, e.g., in Meyer et al. 2014). If people clearly deviate from optimality, and the perceptual decisions can be made without involving humans (e.g., by employing some computer vision and artificial intelligence [AI] mechanisms), then perhaps we should not include people in these tasks. Hence, the optimality of human decisions can be a factor in the design and evaluation of human-computer systems.

To conclude, in a research context one may aim to predict human decisions from a detailed understanding of the evolving situation in which the decisions are made, without committing oneself to the elusive notion of optimality. In applied settings, it is important to analyze performance and to compare it to optimality criteria when evaluating a system that is used to achieve some goal. These two statements do not contradict each other. To achieve both goals, we should develop models of the task, of the way the human performs the task, and of the implications this task performance may have. These models can help us to understand human decisions, whether these are optimal or not. The models can also serve to predict the overall performance of a system in which humans use technology to perform some task involving perceptual decisions.

The standard Bayesian model is normatively invalid for biological brains

Rani Moran[a,b] and Konstantinos Tsetsos[c]

[a]Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University College London, London WC1B 5EH, United Kingdom; [b]Wellcome Centre for Human Neuroimaging, University College London, London WC1N 3BG, United Kingdom and [c]Department of Neurophysiology and Pathophysiology, University Medical Center Hamburg-Eppendorf, 20246 Hamburg, Germany.
rani.moran@gmail.com
https://www.mps-ucl-centre.mpg.de/en/people/rani-moran
k.tsetsos62@gmail.com   https://sites.google.com/site/konstantinostsetsos/

Abstract

We show that the benchmark Bayesian framework that Rahnev & Denison (R&D) used to assess optimality is actually suboptimal under realistic assumptions about how noise corrupts decision making in biological brains. This model is therefore invalid qua normative standard. We advise against generally forsaking optimality and argue that a biologically constrained definition of optimality could serve as an important driver for scientific progress.

Rahnev & Denison's (R&D) extensive review of the perceptual decision-making literature points to the fact that human behavior substantially deviates from optimality. Notably, R&D define optimality according to a standard Bayesian framework (henceforth, the "benchmark model"). Here, we would like to support and elaborate on the claim that this benchmark model has limited validity in describing perceptual decisions. However, rather than addressing the *descriptive* (in)ability of the model to account for behavioral data, we would like to go a step further and postulate that this model has limited *normative* validity once constraints of information processing in biological brains are taken into account.

Specifically, we argue that the assumptions of the benchmark model are over-simplistic, considering how perceptual processing is implemented in biological brains, and that under more realistic assumptions the benchmark model ceases to be optimal. A key assumption in the benchmark model is that encoding noise is the sole corrupting element of perceptual decisions (sect. 2.2). Alas, decision-making operations are never limited to stimulus encoding. Rather, they involve a sequential cascade of processes past the encoding stage, such as integration of information across different sources, mental inference, decision formation, response selection, and motor execution. In biological brains, these processes are performed across a hierarchy of cortical layers, which are prone to different sources of noise (Servan-Schreiber et al. 1990). Hence, perceptual choices are unavoidably corrupted by "late," post-encoding, noise. This notion was corroborated in a recent study that dissociated the contributions of three noise sources to decision suboptimality: sensory encoding, response selection, and mental inference (Drugowitsch et al. 2016). Strikingly, noise in mental inference, rather than encoding, was found to be the main contributor to suboptimality.

Here, we argue that in the presence of post-encoding sources of noise, the benchmark model ceases to be optimal. As we next show, higher accuracy can be obtained by down-weighting some aspects of the available information via a "selective integration" process (Tsetsos et al. 2016a; 2016b; see Osmmy et al. [2013] for an illustration of how a different form of information downweighting can facilitate signal detection performance).

To illustrate the superiority of the selective integration model, we consider a simple binary choice scenario, as in R&D's Figure 1. We assumed that the measurement distributions of these two stimuli are $N(\pm 1, 0.7^2)$, that both stimuli are presented with equal prior probabilities, and that observers are awarded or penalized one point for each right or erroneous choice, respectively. Critically, unlike R&D, we assumed that predicted action costs calculated by observers are prone to an additional source of late Gaussian noise denoted by $N(0, \xi^2)$. Note that this late noise affects only the costs that observers predict but not the actual rewards and penalties they receive (± 1). We assumed that late noise is independent across actions and independent of the

Figure 1 (Moran & Tsetsos). Comparison between the Bayesian "benchmark" and the selective integration models. Average actual cost is displayed as a function of late noise (ξ). For simulation code, see https://osf.io/gexrd/?view_only=aa02df150be94beebf310f1e56cec16f.

encoding noise. Using R&D's notations, the predicted cost of each action is as follows:

$$cost_{predicted}(a) = \sum_{s \in S} [l(s|x, \theta)\pi(s)\mathcal{L}(s, a)] + N(0, \xi^2).$$

According to the benchmark model, the chosen action on each trial minimizes predicted cost:

$$\delta(x) = argmin_{a \in A}\{cost_{predicted}(a)\}.$$

Unsurprisingly, the actually realized cost of the model-selected action increases as a function of late noise, ξ (Fig. 1). Focal to our interest, however, is the comparison between the benchmark and the selective integration models. In the selective integration model, cost predictions are based on "pseudo-likelihoods" l' defined as follows:

$$l'(x|s_i, \theta) = \begin{cases} l(x|s_i, \theta), & if\ l(x|s_i, \theta) \geq l(x|s_j, \theta) \\ l(x|s_i, \theta)w, & if\ l(x|s_i, \theta) < l(x|s_j, \theta), \end{cases}$$

where $w < 1$ (in Fig. 1, $w = 0.1$). Note that the pseudo-likelihood equals the standard likelihood for the more likely of the two stimuli. However, the likelihood of the less likely stimulus is selectively down-weighted. Action selection is determined by minimization of the predicted pseudo-costs:

$$cost'_{predicted}(a) = \sum_{s \in S} [l'(s|x, \theta)\pi(s)\mathcal{L}(s, a)] + N(0, \xi^2)$$

$$\delta'(x) = argmin_{a \in A}\{cost'_{predicted}(a)\}.$$

By down-weighting the likelihood of the less likely stimulus, observers attenuate the harmful influences that late noise exerts on performance (see Tsetsos et al. [2016a] for a detailed discussion). Indeed, the selective integration model strikingly outperforms the benchmark model by achieving lower actual costs (Fig. 1; negative costs correspond to positive gains). Therefore, the benchmark model is suboptimal when post-encoding noise is present and is thus an inadequate standard for assessing optimality in behavior. We recently showed that an intriguing

violation of rational choice theory, intransitivity of choices, was a by-product of adaptive selective integration processes (Tsetsos et al. 2016a). Future research should investigate whether and which of the sundry "suboptimal" behaviors reviewed by R&D actually reflect biologically constrained adaptive processes.

Next, we wish to qualify R&D's arguments about the overall utility of assessing optimality (sect. 4). Although we agree that addressing optimality should not be in itself the ultimate *goal* of the study of perception, we still think that the notion of optimality can serve as an important driver to scientific progress. From an evolutionary perspective, perceptual processing reflects an extended adaptation process and, as such, is ex ante expected to be optimal (Moran 2015). Therefore, when theories of optimality are defined appropriately, taking into account the biological constraints of human information processing, they can provide an invaluable benchmark, guiding both theory development and behavioral assessments. On the one hand, when organisms are found to behave optimally, it raises questions about how optimality is achieved at the algorithmic and implementation levels. When, on the other hand, suboptimal behaviors are found, it raises questions pertaining to why these alleged deviations from optimality occur. Such questions can lead to a better understating of the constraints and limitations of human information processing (Tsetsos et al. 2016a), of the cost function that neural systems strive to minimize (Soltani et al. 2012), and of the statistical structure of the environment in which cognitive processes evolved (Fawcett et al. 2014). This understating often leads to a subtler definition of optimality, which supports novel behavioral hypotheses and assessments. In sum, we conceive of scientific progress as an iterative process, in which the notion of optimality, rather than being relinquished, continually evolves and undergoes refinement. It is challenging to define optimality, but we nevertheless think that such attempts are instrumental in scientific progress.

Observer models of perceptual development

Marko Nardini[a] and Tessa M. Dekker[b]

[a]Department of Psychology, Durham University, Durham DH1 3LE, United Kingdom and [b]Department of Experimental Psychology and Institute of Ophthalmology, University College London, London WC1E 6BT, United Kingdom.
marko.nardini@durham.ac.uk   http://community.dur.ac.uk/marko.nardini/
t.dekker@ucl.ac.uk   http://www.ucl.ac.uk/~ucjttb1/

Abstract

We agree with Rahnev & Denison (R&D) that to understand perception at a process level, we must investigate *why* performance sometimes deviates from idealised decision models. Recent research reveals that such deviations from optimality are pervasive during perceptual development. We argue that a full understanding of perception requires a model of how perceptual systems become increasingly optimised during development.

Perceptual abilities undergo major development during infancy and childhood – for example, for detecting low-contrast stimuli (Adams & Courage 2002) and noisy patterns of motion (Hadad et al. 2011) or recognising complex stimuli such as faces (Mondloch et al. 2002). Classically, the focus of perceptual development research has been on improvements in sensitivity (likelihoods). As reviewed in the target article, decades of adult research show how sensitivity changes can result from changes within a decision-model framework that incorporates likelihoods, priors, cost functions, and decision rules. Applying this framework to development, we argue that *perceptual improvements must be explained in terms of changes to these components.* This will lead to a new understanding of how perceptual systems attain their more highly optimised mature state.

Specifically, we need to know the following:

(1) *Which elements of the observer model are changing (developing), leading to improvements in perceptual function?* Recent evidence suggests that multiple components of the decision model are developing significantly during childhood. Until late into childhood, observers are still using *decision rules* less efficiently: misweighting informative cues (Gori et al. 2008; Manning et al. 2014; Sweeny et al. 2015) or using qualitatively different decision rules altogether (Jones & Dekker 2017; Nardini et al. 2008; 2010). Other studies show abilities to learn and use *priors* and *costs* also to be developing late into childhood (e.g., Dekker & Nardini 2016; Stone 2011; Thomas et al. 2010). The new, model-based approach to development pioneered in these studies paves the way for understanding how likelihoods, priors, cost functions, and decision rules are shaped as children learn, and for testing which common processes can explain perceptual development across a range of different tasks. Studies to date have successfully captured developmental changes in performance by fitting how parameters of specific components of the decision model change with age on single tasks. This usefully sets quantitative bounds on potential changes in these processes, but the data are often compatible with more than one account. For example, in a rewarded reaching task (Dekker & Nardini 2016), children up to the age of 11 years aim too close to a penalty region to maximise their score, reflecting overconfidence in likelihood of hitting the target, underestimation of cost, or a central pointing prior. An important way forward is therefore to evaluate the fit of developmental models to multiple tasks and to test their predictions on new tasks.

(2) *How are more efficient and adult-like decision rules, priors, and cost functions acquired during development?* Beyond characterising the changes in decision-model components underlying perceptual development, the ultimate aim is to understand the mechanisms driving these changes. A major contributing factor is likely to be experience, which shapes the sensitivity of neuronal detectors, determining likelihoods (Blakemore & Van Sluyters 1975), changes priors (Adams et al. 2004), and is needed to learn the potential consequences of actions (cost factors). It is not clear in which circumstances such experience is generalizable (e.g., priors or costs learned during one task applied to another), how experience drives learning of decision rules, or whether there are sensitive periods like those for sensitivities (likelihoods) in other parts of the decision model (e.g., for learning priors). A useful approach is investigating the neural changes supporting improvements in decision-model components as perception becomes more optimised, such as more precise representation of likelihoods (Van Bergen et al. 2015) and values (Wu et al. 2011), or more precise computing of weighted averages, perhaps implemented via divisive normalisation (Ohshiro et al. 2011). The power of this approach is demonstrated by recent studies of developmental disorders, in which there are exciting developments in linking components of observer models to specific neural mechanisms (Rosenberg et al. 2015). For example, in autism, tasks that involve combining new evidence with prior knowledge are disproportionally affected, and this has recently been linked to the overweighting of sensory likelihoods versus priors, possibly because of altered neural operations mediated by noradrenaline and acetylcholine (Lawson et al. 2017). In addition, a new, model-based approach to developmental neuroimaging lets us disentangle components of the developing decision model across different neural processing stages. We recently showed that development of cue integration during depth perception was linked to a shift from using depth cues independently to combining them, by neural detectors in sensory cortex (adopting a "fusion" rule; Dekker et al. 2015). This suggests that the late development of cue integration is driven by a change in how sensory information is combined (sensory decision rule), rather than improved readout of the fused estimate during task performance (higher-order decision rule or cost function). These studies demonstrate how a developmental approach can provide computational-level understanding of the crucial ingredients for building a mature optimised observer.

The end goal of this approach is an observer model incorporating processes of learning and development: a developing standard observer model. This will provide a more complete understanding of perceptual systems and a basis for developing intelligent machines that can learn to perceive in novel environments. For example, understanding the structure of experience that scaffolds our ability to transfer previous likelihoods, cost functions, and decision rules from one task to another can inform the development of more flexible artificial intelligence (AI) agents (Wang et al. 2017). Similarly, significant improvements in robotic grasp performance have been gained from incorporating developmental stages such as motor babbling and gradual improvements in visual acuity into the training regime (Cangelosi et al. 2015). In addition, understanding which developmental changes in the decision model (e.g., sensitivity vs. decision rule) drive perceptual improvements at different ages will provide a crucial basis for better training of perception and action in patients with sensory loss.

Supra-optimality may emanate from suboptimality, and hence optimality is no benchmark in multisensory integration

Jean-Paul Noel

Vanderbilt Brain Institute, Vanderbilt University, Nashville, TN 37240.
jean-paul.noel@vanderbilt.edu   http://jeanpaulnoel.com/

Abstract

Within a multisensory context, "optimality" has been used as a benchmark evidencing interdependent sensory channels. However, "optimality" does not truly bifurcate a spectrum from suboptimal to supra-optimal – where optimal and supra-optimal, but not suboptimal, indicate integration – as supra-optimality may result from the suboptimal integration of a present unisensory stimuli and an absent one (audio = audio + absence of vision).

Arguably the study of multisensory integration was born from the recording of spikes in the feline superior colliculus (Stein & Meredith 1993). These early studies presented animals with simple visual (V) flashes and auditory (A) beeps and held the occurrence of supra-additive responses (i.e., audiovisual [AV] responses greater than the sum of auditory and visual responses) as the hallmark for multisensory integration. However, this phenomenon is not common in the neocortical mantle (vs. subcortex; Frens & Van Opstal 1998), nor when multisensory integration is indexed via behavior or by measuring ensembles of neurons (e.g., local field potentials, electroencephalography [EEG], functional magnetic resonance imaging [fMRI]; Beauchamp 2005). Hence, over the last two decades there has been a greater appreciation for sub-additive responses as equally demonstrating an interesting transformation from input (i.e., A + V) to output (i.e., AV), and thus highlighting the synthesis of information across senses. That is, arguably the classic study of multisensory integration has grown to conceive of sub- and supra-additivity as being on extremes of a spectrum where both ends are interesting and informative.

In parallel, the originally described "principles of multisensory integration" (e.g., information that is in close spatial and temporal proximity will be integrated) have been translated to a computational language that is seemingly applicable throughout the cortex and widely observed in behavior. As Rahnev & Denison (R&D) underline in their review, this computational framework dictating much of the current work within the multisensory field is that of Bayesian decision theory. Indeed, among others, audiovisual (Alais & Burr 2004), visuo-tactile (Ernst & Banks 2002), visuo-vestibular (Fetsch et al. 2009), and visuo-proprioceptive (van Beers et al. 1999) pairings have been demonstrated to abide by maximum likelihood estimation (MLE) – the weighting of likelihoods by relative reliabilities and concurrent reduction in integrated (vs. unisensory) variance. Given this extensive body of literature, I believe the gut reaction of many multisensory researchers – mine included – to this review and the thesis that assessing optimality is not useful was that we must acknowledge the limitations of solely considering "optimality" without examining the underlying components (e.g., prior, cost function), but that this construct is nevertheless valuable. If subjects behave optimally (i.e., reduction of uncertainty), then at minimum, there is evidence for interdependent channels. Namely, the reduction of variance in multisensory cases (vs. unisensory) is evidence for the fact that at some point, unisensory components are fused; the next step is to understand exactly how these channels are fused. Furthering this argument, it could be conceived that supra- and suboptimality exist on a continuum where evidence for supra-optimality or optimality is evidence for multisensory integration (admittedly without providing much mechanistic insight given the points raised by R&D), while suboptimality does not bear evidence of a synthesis across the senses. In other words, indexing optimality as a benchmark for integration is useful because Bayesian computations are ubiquitous in the brain and behavior, and in that it reduces the state space of integration from "anything apart from linear summation" (i.e., from sub-additive to supra-additive excluding additive) to "anything greater than or equal to optimal" (i.e., from optimal to supra-optimal but not suboptimal).

However, upon further consideration, I believe this reasoning to be erroneous (and therefore I agree with the thesis put forward by R&D). In short, contrarily to the case of additivity, optimality does not lie on a spectrum from sub- to supra-optimal, and hence optimality per se is no benchmark.

Traditionally, supra-optimality (an apparent impossibility) within multisensory systems has been hypothesized to emerge from a process of "active sensing" (Schroeder et al. 2010). That is, the presence of a second sensory stimulus (e.g., A) may sharpen the representation of a first unisensory stimulus (e.g., V) so that when these are combined (e.g., AV), sharper unisensory estimates than originally considered are combined, resulting in apparently supra-optimality. Nonetheless, as Shalom and Zaidel (2018) have recently highlighted, somewhat paradoxically, it could additionally be the case that supra-optimality results from suboptimal integration. Namely, researchers typically take unisensory likelihoods at face value. However, within a multisensory (e.g., AV) context, the presentation of auditory stimuli alone is in fact not auditory alone (e.g., A), but instead the presence of auditory information and the absence of visual information (e.g., A + no V). Therefore, in this example, researchers are underestimating the reliability of the auditory channel (which is truly A-likelihood + a flat visual likelihood), which will ultimately result in claims of supra-optimal multisensory integration. This second observation (by Shalom & Zaidel 2018) is similar to the case of active sensing, in that the sharpness of unisensory likelihoods is underestimated. However, the perspective is quite different in that supra-optimality is not the result of cross-modal feedback enhancing unisensory representation solely when presented in a multisensory context, but in fact, in this latter case, supra-optimality is merely an experimental construct that results from the erroneous underestimation of a unisensory likelihoods; the world is by nature multisensory, and hence unisensory estimates are impoverished estimates wherein a cue has been artificially removed. That is, supra-optimality can result from the non-optimal integration of a signal (e.g., A) and noise (e.g., a non-present V signal). In turn, there is no true gradient between supra- and suboptimality, and hence positioning optimality as a benchmark bifurcating between multisensory fusion and fission is ill advised. Instead, as highlighted by R&D, we ought to conceive of (multisensory) perception as a dynamic system where likelihoods, priors, cost functions, and decision criteria all interact interdependently in both feedforward and feedback manners.

When the simplest voluntary decisions appear patently suboptimal

Emilio Salinas [ORCID], Joshua A. Seideman and Terrence R. Stanford

Department of Neurobiology & Anatomy, Wake Forest School of Medicine, Winston-Salem, NC 27157-1010.

esalinas@wakehealth.edu   jseidema@wakehealth.edu
stanford@wakehealth.edu

Abstract

Rahnev & Denison (R&D) catalog numerous experiments in which performance deviates, often in subtle ways, from the theoretical ideal. We discuss an extreme case, an elementary behavior (reactive saccades to single targets) for which a simple contextual manipulation results in responses that are dramatically different from those expected based on reward maximization – and yet are highly informative and amenable to mechanistic examination.

The conclusions drawn by Rahnev & Denison (R&D) rely on analyses spanning many tasks and experimental conditions in which perceptually guided decisions deviate, for a variety of reasons, from those of an ideal observer model. Indeed, they exhaustively build a convincing argument. But sometimes a single, powerful example can illustrate a general result with great eloquence. That is the case with an elegant paradigm known as the one-direction-rewarded, or 1DR, task.

The 1DR task is deceptively simple (Hauser et al. 2018; Hikosaka et al. 2006; Lauwereyns et al. 2002). The subject (a monkey, in this case) is instructed to perform an elementary action: to look at a lone, clearly visible stimulus. Each trial starts with the monkey briefly fixating on a central spot on an otherwise blank screen. Then the fixation spot disappears, and, at the same time, a target stimulus appears at one of four possible symmetric locations (or one of two locations, depending on the study). The stimulus location varies randomly across trials. The monkey is rewarded for making a quick eye movement, a saccade, to the target – a reaction that is, in fact, quite natural.

But there is a catch. Correct saccades to one location yield a large reward, whereas correct saccades to the other locations yield either a small reward or no reward (this varies by monkey, but, importantly, the results are the same). The rewarded location stays constant for a block of trials. The spatial asymmetry in reward expectation leads to a conflict: The monkey wants to look in one direction but is often instructed to look elsewhere. Nevertheless, all trials must be completed, whether the reward on offer is large or small. There is no strategic advantage to responding differently in one condition compared with the other. Only one alternative is available, so deliberating is unnecessary. To maximize the reward rate, the monkey should look at the target as quickly and as accurately as possible each time, regardless of where it appears.

However, the observed behavior diverges quite drastically from this prescription. Saccades in congruent trials, in which the target and the highly rewarded locations coincide, are initiated more quickly and are more accurate than those in incongruent trials, in which the target and the highly rewarded locations differ. The effects are huge. For example, in our own data (Hauser et al. 2018), we found that the reaction time (RT) went from about 150 ± 25 ms (mean ± standard deviation) to about 250 ± 80 ms, with the error rate changing from virtually zero (99.7% right) to about 10% of incorrect saccades. The extraordinary sensitivity of the monkeys to reward asymmetry also manifests in other, low-level behavioral metrics, such as the peak saccade

velocity, as well as in the swiftness with which the animals respond to changes in the asymmetry over time. When the rewarded location changes, which happens without warning, it takes a single trial for the spatial bias to switch accordingly (when only two locations are used). This rich phenomenology is highly consistent between animals, laboratories, and task variants, and it remains stable for months, even after many thousands of trials of practice (Hauser et al. 2018; Hikosaka et al. 2006; Takikawa et al. 2002; Watanabe et al. 2001).

Such behavior runs counter to the expectation based on reward maximization, as outlined previously. Within the behavioral repertoire discussed by R&D, the spatial bias represents a particularly drastic breakdown of the speed-accuracy tradeoff (sect. 3.4), because one condition (congruent) leads to more accurate and much faster responses than the other (incongruent). The 1DR behavior can also be considered as a limit case of a choice task in which different responses have different payoffs (sect. 3.3). Normally, in monkeys, such asymmetry produces a shift in criterion (Feng et al. 2009; Stanford et al. 2010). Here, the perceptual uncertainty about the right option is eliminated, and the adjustment in criterion is grossly inappropriate. Either way, the underlying "cost function" guiding the behavior must be radically different from those that may be naively construed as optimal.

It is not difficult to imagine why such a discrepancy arises. The capacity to discriminate and seek rewarding events must be critical for survival, so it is not surprising that reward drives or modulates numerous cognitive processes. In particular, reward expectation is intimately linked to attentional deployment and oculomotor control (Hikosaka et al. 2006; Maunsell 2004; Peck et al. 2009; Preciado et al. 2017). The conditions in the 1DR task likely set up a cognitive trap of sorts – the illusion of a choice – such that the monkeys never cease to strongly prioritize the rewarded location. In essence, they demonstrate persistent wishful thinking.

Regardless, the 1DR paradigm has been extremely useful, even though it does not adhere to a normative theory. For many years, Hikosaka and colleagues have exploited it to investigate how cognition and motivation interact, seeking to identify and functionally characterize the oculomotor and reward-encoding neural circuits that mediate the biasing effects and their motor expression. Theirs is an impressive research program that has uncovered many such contributions and mechanistic components (e.g., Ding & Hikosaka 2006; Ikeda & Hikosaka 2003; Isoda & Hikosaka 2008; Tachibana & Hikosaka 2012; Takikawa et al. 2004; Yasuda & Hikosaka 2017). In this context, justifying the animals' behavior on the basis of an optimality principle or ideal observer model seems rather unnecessary. Furthermore, in our own laboratory, we recently developed a mechanistic model that replicates the monkeys' RT distributions as well as single-neuron activity in the frontal eye field (FEF) during performance of the 1DR task (Hauser et al. 2018). This model explains the observed behavior in great quantitative detail based on dynamical interactions found in FEF.

In summary, the results in the 1DR task exemplify one of the main conclusions drawn by R&D – that although a normative benchmark may provide useful interpretive guidance in many cases, it is by no means necessary for understanding a particular behavior, or for generating a complete mechanistic description of it.

Characterising variations in perceptual decision making

Johannes Schultz ⬤ and René Hurlemann ⬤

Department of Psychiatry and Division of Medical Psychology, University of Bonn Medical Center, 53105 Bonn, Germany.
johannes.schultz@gmail.com   http://sites.google.com/site/johannesschultz/
renehurlemann@icloud.com   http://renehurlemann.squarespace.com/welcome/

Abstract

Current perspectives propose that observer models accounting for both optimal and suboptimal behaviour may yield real progress in understanding perception. We propose that such models could, in addition, be very useful for precisely characterising the variation in perception across healthy participants and those affected by psychiatric disorders, as well as the effects of neuromodulators such as oxytocin.

In their thought-provoking target article, Rahnev & Denison (R&D) argue that real progress in understanding perception could be achieved by observer models that account for optimal and suboptimal behaviour. We believe that such models could furthermore be very useful for characterising variations in perception across healthy participants and those affected by psychiatric disorders. Inter-individual variations in perception (e.g., Grzeczkowski et al. 2017; Partos et al. 2016; Schultz & Bülthoff 2013; van Boxtel et al. 2017) and perceptual decision making (Ratcliff et al. 2010; 2011; Schmiedek et al. 2007) have been widely reported. An established approach to investigate these processes and their variations has been to model accuracy and response times using diffusion models (Ratcliff 1978; Ratcliff et al. 2016). Comparing parameters of these models with personality traits across healthy participants, or between healthy participants and patients, provides insight into the origins of the variability. This has allowed researchers to relate individual differences in perceptual decision making to individual differences in IQ, working memory, and reading measures (Ratcliff et al. 2010; 2011; Schmiedek et al. 2007) and to characterise deficits in participants with aphasia (Ratcliff et al. 2004), dyslexia (McKoon & Ratcliff 2016), attention-deficit/hyperactivity disorder (Mulder et al. 2010), schizophrenia (Moustafa et al. 2015), depression, and anxiety (White et al. 2009).

As part of the Research Domain Criteria (RDoC) project (Insel et al. 2010) aiming to incorporate genetics, neuroimaging, and cognitive science into future psychiatric diagnostic schemes, applying neurobiologically plausible models of value-based decision making to characterise deficits observed in psychiatric disorders (Collins et al. 2017; Huys et al. 2015) has led to the development of computational psychiatry (Maia et al. 2017; Wiecki et al. 2014). This approach promises mechanistic explanations of how psychiatric symptoms such as cognitive biases may result from failures of decision variable evaluation (Huys et al. 2015). Bayesian models combining prior information with sensory evidence are particularly promising in yielding insight into pathophysiological mechanisms of perceptual distortions observed in schizophrenia. For example, information processing favouring prior knowledge over incoming sensory evidence can account for differences in visual illusion perception observed in early psychosis and schizotypy (Partos et al. 2016; Teufel et al. 2015). The "jumping-to-conclusions" bias in event probability estimation typical of schizophrenia can be characterised by increased circular inference – that is, the corruption of sensory data by prior information, with feedforward and feedback loops of the model correlating with negative and positive symptoms, respectively (Jardri et al. 2017). In time, such approaches may allow us to develop specific therapeutic approaches, such as metacognitive training (e.g., see Moritz & Woodward 2007).

Observer models may allow similar progress in understanding the mechanisms underlying dysfunctions of social perception and interaction. Parameterizable social stimuli may prove very helpful in this regard; for example, point-light motion stimuli and tasks assessing different levels of processing have allowed researchers to better understand how autistic traits affect certain aspects of biological motion perception (van Boxtel et al. 2017). The response to others' gaze is also affected in autism (Leekam et al. 1998; Wallace et al. 2006); here, a recently developed computational model of the perception of gaze direction (Palmer & Clifford 2017) has yielded insight into the origin of those dysfunctions: It has been proposed that autism is associated with reduced divisive normalisation of sensory responses, attributable to an increased ratio of cortical excitation to inhibition (Rosenberg et al. 2015). Interestingly, both divisive normalisation and sensory adaptation occur robustly in autism in the context of gaze processing (Palmer et al. 2018). This suggests that the differences in response to others' gaze may instead be related to differences in the interpretation of gaze direction or the spontaneous following of others' gaze (Senju et al. 2009). Similar work could be undertaken for elucidating other essential social cognitive functions, such as face recognition. Face recognition capacities widely vary across healthy participants (Wilmer et al. 2012), ranging from congenital prosopagnosia (Behrmann & Avidan 2005; McConachie 1976) to "super-recognition" (Russell et al. 2009). Although progress towards understanding the cognitive and neural underpinnings of congenital prosopagnosia is being made (Susilo & Duchaine 2013), the most widely used tests may not capture the alternative perceptual strategies adopted by people afflicted by prosopagnosia (Esins et al. 2016). Parameterizable face stimuli (Dobs et al. 2014; Esins et al. 2014; 2016) may allow us to better characterise those strategies by allowing direct comparisons between human and ideal observer performance (Dobs et al. 2016; 2017). Such approaches may be instrumental in identifying alternative heuristics used by participants with congenital prosopagnosia and other impairments of social perception.

Recent studies have demonstrated that exogenous administration of the neuropeptide oxytocin (OT) influences the perception of social stimuli such as facial emotions in a dose-dependent manner (Spengler et al. 2017b). Furthermore, OT modulates attractiveness judgements of faces (Hurlemann et al. 2017; Striepens et al. 2014), alters the sensory quality of social touch (Kreuder et al. 2017; Scheele et al. 2014) and body odours (Maier et al. 2018), increases a tendency to anthropomorphise (Scheele et al. 2015), or, in rats, boosts the salience of acoustic social stimuli (Marlin et al. 2015). At present, it is still unclear whether the behavioural effects of OT result from perceptual changes, such as increased attention to the socially informative eye region (Guastella et al. 2008), improved recognition of cues

about sex and relationship (Scheele et al. 2013), and/or facilitated sensing of and responding to emotional stimuli (Spengler et al. 2017a). Analysing these effects using observer models may help identify which aspect of the perceptual decision-making process is influenced by OT. As OT is also a promising therapeutic (Hurlemann 2017; Palmer & Clifford 2017), understanding its mode of action may be informative in order to specifically target dysfunctional perceptual processes particularly amenable to OT treatment.

Inclusion of neural effort in cost function can explain perceptual decision suboptimality

Yury P. Shimansky and Natalia Dounskaia

Kinesiology Program, Arizona State University, Phoenix, AZ 85004.
yury.shimansky@asu.edu   natalia.dounskaia@asu.edu

Abstract

A more general form of optimality approach applied to the entire behavioral paradigm should be used instead of abandoning the optimality approach. Adding the cost of information processing to the optimality criterion and taking into account some other recently proposed aspects of decision optimization could substantially increase the explanatory power of an optimality approach to modeling perceptual decision making.

Use of an optimality approach has been remarkably successful in many different domains of the natural sciences. Its application to biological organisms equipped with a well-developed central nervous system (CNS) is generally based on the fact that these organisms learn, thus optimizing their behavior. Hence, a well-learned behavior can be considered optimized, and various features of such behavior can be explained as a result of optimization with respect to a specific criterion (also known as cost function). The explanatory power is perhaps the main advantage of an optimality approach. To fully benefit from that advantage, an optimality approach should be applied to the entire behavioral paradigm. However, the suggestion of Rahnev & Denison (R&D) to abandon the optimality approach is based on examples of suboptimality obtained through application of an optimality approach to the decision rule only. Here we argue that this suboptimality is indicative of incompleteness of the used optimality criterion, and therefore, instead of abandoning the optimality approach, it is more productive to focus on the identification of the important aspects of the entire behavioral paradigm in addition to the perceptual decision rule. Here, we will demonstrate that optimization of behavior as a whole accounts for the suboptimality of the decision rule on two examples discussed by the authors – namely, inter-observer/trial variability, while decision rule optimality is observed on average (sect. 3.1.2 and 3.2) and perceptual biases (sect. 3.8.1 and 3.8.4).

The inter-observer/trial variability can be explained as a result of overlooking important components of the optimality criterion.

In particular, the cost of neural effort for information processing should be taken into account (Dounskaia & Shimansky 2016; Shimansky & Rand 2013). Decision making is based on processing sensory information and integrating it with internal representations of past experience. These steps of information processing are required for estimating the state of the relevant constituents of the environment (perhaps including own body) to which the decision rule is applied. These steps are an essential component of the behavioral paradigm, and therefore, the cost of the corresponding neural effort needs to be included in the optimality criterion. In the field of movement control, where an optimality approach has successfully accounted for vast experimental data (for reviews, see Shimansky et al. 2004; Todorov 2004), the consideration of the cost of information processing was critical for understanding even relatively simple motor behaviors such as reaching to grasp (Shimansky & Rand 2013) and point-to-point movements (Rand & Shimansky 2013). Also, the "trailing" pattern of joint coordination typically observed during well-learned arm movements can be fully understood only if neural effort for joint coordination is considered as a primary component of the optimality criterion (Dounskaia & Shimansky 2016; Goble et al. 2007).

The consideration of the cost of neural effort for information processing implies that the total cost function is a weighted sum of this cost and the cost of deviations from decision optimality. The brain is therefore required to perform a tradeoff between the two costs. Disregarding this tradeoff and focusing on decision optimality only would lead to a conclusion that experimentally observed decisions are suboptimal. However, deviations of the decisions from optimality are predicted by a tradeoff between the two costs: Decision optimality often requires a neural effort of excessive cost, thus making the total cost greater than optimal. Hence, perceptual decision suboptimality can be explained by applying an optimality approach to the entire behavior instead of to the decision rule only. Similar considerations were employed to explain the variability of hand motion during reach-to-grasp movements (Shimansky & Rand 2013).

Experimentally observed perceptual biases may also be consistent with perceptual decision optimality. Shimansky (2011) used an optimality approach to predict perceptual biases in experimental conditions that included a combination of perceptual uncertainty with loss asymmetry with respect to the direction of decision error. An example can be jumping over an obstacle under poor visibility conditions. As Shimansky (2011) demonstrated, an optimality approach predicts a tendency to overestimate the size of the obstacle under these conditions.

In addition to suboptimality of experimentally observed perceptual decisions, another criticism of the optimality approach formulated by R&D is "flexibility" of the optimization criterion because of uncertainty about its exact composition. Similar arguments against the optimality approach were formulated in the field of movement control. Namely, it was noted that use of an optimality approach leads to "circulatory" reasoning, meaning that experimental data are used to determine the optimization criterion, and then the optimality approach is used for explaining the experimental results (e.g., Diedrichsen et al. 2010). A solution to this seeming paradox has been proposed by Shimansky and Rand (2013). In brief, it consists of using a relatively small subset of the total amount of collected experimental data for determining unknown parameters of the optimality criterion (e.g., weights of the cost function components), with subsequent testing of the determined parameters on the rest of the experimental data.

Although validity of this method of establishing the optimality criterion was demonstrated for data obtained in experiments on reach-to-grasp and point-to-point movements, this method can be extended to the field of perceptual decision making.

It is also noteworthy that even though R&D suggest abandoning the optimality approach, the specific hypotheses they formulate (their Table 1) are indirectly based on a certain model of optimal behavior. For example, the terms "corrupt," "weird," or "'suboptimal" behavior (used in Table 1) make sense only with respect to a certain criterion of behavior optimality. Therefore, an optimality approach is needed to measure the extent of optimality and, in case of suboptimality, help identifying possible factors causing it.

In conclusion, the suboptimality of perceptual decisions described in the target article does not warrant abandoning the usage of an optimality approach. A more general form of optimality approach is required, where an assumption of optimality is applied to the entire paradigm instead of the decision rule alone, which would require use of a more complex form of cost function. Specific hypotheses regarding possible reasons for decision rule suboptimality could be formulated in terms of additional cost function components.

Discarding optimality: Throwing out the baby with the bathwater?

Patrick Simen[a] and Fuat Balcı[b]

[a]Neuroscience Department, Oberlin College, OH 44074 and [b]Department of Psychology & Research Center for Translational Medicine, Koç University, 34450 Sarıyer/Istanbul, Turkey.
psimen@oberlin.edu   http://www.oberlin.edu/faculty/psimen/
fbalci@ku.edu.tr   https://tmdmlab.ku.edu.tr/

Abstract

Rahnev & Denison (R&D) argue against normative theories and in favor of a more descriptive "standard observer model" of perceptual decision making. We agree with the authors in many respects, but we argue that optimality (specifically, reward-rate maximization) has proved demonstrably useful as a hypothesis, contrary to the authors' claims.

Rahnev & Denison (R&D) have produced a comprehensive survey and incisive critique of optimality hypotheses in perceptual decision research. It is an important and corrective contribution. For us, as for the authors, whether an organism behaves optimally or suboptimally is largely beside the point. What we really want to know is: What are the neural mechanisms underlying perception, cognition and action, and how do they work?

One thing seems clear: Perceptual decision making is fundamentally adaptive. Animal perception is alive with adaption, from performance improvements with practice, to the biasing effects of rewards. But inherent in the concept of adaptation is some form of objective or utility function, such that changes in an organism's behavior tend to increase its value. The authors advocate the development of a standard observer model of perceptual decision

making: a model that describes simply what is – not what is ideal. We applaud this proposal, but we venture to guess that a well-defined utility function will be found at the core of any future standard observer model, and that brain mechanisms of adaptation will be most readily uncovered if aided by a well-developed theory of optimality.

The authors advocate ditching the optimal-suboptimal dichotomy altogether, based on their strong position that the optimality concept itself is essentially meaningless. They also cite widespread findings of suboptimal behavior, and they note the problematic coexistence of multiple meanings of optimality in the literature. We agree that terminology abuse is truly a problem. We disagree, however, when the authors cite our own work as evidence of suboptimality in speed-accuracy tradeoff adjustment, even according to an explicitly stated definition – namely, that behavior is optimal if it maximizes the rate of rewards earned for performance (Bogacz et al. 2006).

Our interpretation of our work on speed-accuracy tradeoff adjustment and the reward-rate-optimality hypothesis is that there is strong evidence for reward-rate-optimal tradeoffs in some circumstances. Further, we argue that, far from being a useless concept, optimality analyses have generated interesting and falsifiable predictions about two-choice perceptual decisions. We tested them and found that people, as a group, appeared to be distributed in their behavior around the reward-rate-maximizing speed-accuracy tradeoff (Simen et al. 2009). We also found that individually, people exhibited striking transitions between fast guessing and evidence accumulation in a manner consistent with some of the most intuition-defying predictions, in Bogacz et al. (2006), of a diffusion model with reward-rate-optimal boundary and starting point parameters

Most importantly, we found that there are some decision-task features that people can handle almost reward rate optimally with sufficient training, whereas there are others for which they remain notably suboptimal. In short, people tend to shift toward the reward-rate-optimal speed-accuracy tradeoff when changes in a task have to do with *time* or *rate*. For example, the reward-rate-optimal speed-accuracy tradeoff favors speed for shorter inter-trial delays but accuracy for longer delays. There is good evidence that people not only adapt in the direction predicted by reward-rate-optimality, but that they (at least young participants) come very close to achieving optimality when the incentive structure is clear (Evans & Brown 2017; Simen et al. 2009; Starns & Ratcliff 2010). Our participants collectively appeared to be clustered in their boundary and starting point parameter settings around the ideal observer values, earning 97%–99% of the maximum possible reward. Participants did exhibit a slight, but consistent over-emphasis on accuracy. This is evidence for suboptimality, but only in the sense that 3.14 does not equal pi.

On the other hand, our participants tended to fail miserably to approximate optimal performance when it was the signal-to-noise ratio that changed across conditions (Balcı et al. 2011b). That failure may have an explanation in terms of the reward-rate-optimality hypothesis. Contrary to what is stated in the first paragraph of section 3.4, the reward-rate-optimal tradeoff is non-monotonic over the range of possible signal-to-noise ratios. For very high and low signal-to-noise ratios, one should decide quickly; only when the signal-to-noise ratio is at an intermediate level should one slow down to boost accuracy (Bogacz et al. 2006). Only an optimality analysis reveals this prediction, and tests of it reveal that most people do not follow it when the signal-to-noise ratio is low, though our

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 [ID]

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

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 [ORCID]

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

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

representations whose statistics match those of the external world (Simoncelli 2003). This gives rise to an efficiency principle: Neural coding is distributed in a way that ensures maximal resources to be devoted to those features that are most likely to be encountered in natural environments (Girshick et al. 2011; Wei & Stocker 2015; 2017). The "why" of the oblique effect has an answer: It arises because of a neural coding scheme that has evolved to be maximally efficient.

Another way of understanding why perceptual distortions might arise is via consideration of the sources of uncertainty that corrupt decisions. When judging visual stimuli, noise arising during sensory encoding limits performance – for example, low-contrast stimuli are hard to see. However, for a capacity-limited system (such as a biological agent), noise that arises "late" – that is, during inference itself – place a further constraint on the fidelity of information processing (Drugowitsch et al. 2016). Recently, categorisation tasks that require the integration of information in space and time have revealed perceptual distortions in humans, such as the "robust averaging" of visual features to which R&D refer (de Gardelle & Summerfield 2011). The compressive non-linearity that produces this effect leads to performance reductions for an observer with limitless capacity. However, simulations show that distorted transduction can paradoxically maximise reward when decisions are corrupted by "late" noise – that is, noise that arises during inference, rather than at the level of sensory encoding. This is again because of an efficiency principle – when computational resources are limited, the best policy may be to transduce perceptual information nonlinearly, allowing gain to be allocated by preference to some features over others (Li et al. 2017). In fact, the precise form of the reward-maximising distortion varies according to the overall distribution of stimuli observed, and both behavioural and neural data suggest that humans shift from a compressive to an anticompressive form of distortion in a way that consistently maximises their performance (Spitzer et al. 2017).

Although the details are only emerging, we think it is likely that the wide range of perceptual "suboptimalities" that R&D highlight – sequential trial history effects, central tendency biases, sluggish belief change, and adaptation and/or normalisation processes – are all hallmarks of a cognitive system that has evolved to perform efficient computation in natural environments that exhibit stereotyped statistics and autocorrelation both in space and time. Indeed, a similar efficiency principle has been shown to hold for decisions in other domains, including the well-described decision biases in economic tasks, such as deciding among prospects with differing value. In one example, policies that lead to violations of axiomatic rationality can be shown to be optimal under late noise, providing an "optimal" explanation for economic irrationality (Tsetsos et al. 2016a).

More generally, biological information processing systems have intrinsic costs and constraints that place a premium on computational efficiency. In other words, brains have evolved to minimise both a behavioural cost function (maximising reward) and a neural cost function (minimising computational load). Rather than being ad hoc failure modes in biological brains, "suboptimalities" in perception expose how computation has adapted efficiently to the structure of the natural environments in which biological organisms exist. A principled research agenda, rather than merely documenting deviations from optimality, should attempt to explain them.

Bayesian statistics to test Bayes optimality

Brandon M. Turner,[a] James L. McClelland,[b] and Jerome Busemeyer[c]

[a]Psychology Department, The Ohio State University, Columbus, OH 43210;
[b]Psychology Department, Stanford University, Stanford, CA 94305; and
[c]Psychology Department, Indiana University, Bloomington, IN 47405.
turner.826@gmail.com   https://turner-mbcn.com/
jlmcc@stanford.edu   https://stanford.edu/~jlmcc/
jbusemey@indiana.edu   http://mypage.iu.edu/~jbusemey/home.html

Abstract

We agree with the authors that putting forward specific models and examining their agreement with experimental data are the best approach for understanding the nature of decision making. Although the authors only consider the likelihood function, prior, cost function, and decision rule (LPCD) framework, other choices are available. Bayesian statistics can be used to estimate essential parameters and assess the degree of optimality.

Rahnev & Denison (R&D) perform a useful service by reminding us that human behavior often deviates from what can be construed as optimal, once some target variable (such as expected reward) has been specified. We would like to remind readers that Bayesian statistics can be used to fit models and to contribute to our assessment of the degree of optimality. This differs from the usual Bayes rational approach (BRA). Unlike BRA, Bayesian statistics can be used to assess optimality by fitting a model to data and assessing how close the resulting parameter estimates are to optimal settings. For example, if an observer model is specified in such a way that one parameter combination results in optimal performance, one can use posterior estimates to quantify (i.e., in terms of both the parameter estimates and their uncertainty) how likely subjects were to have been performing optimally.

We have used such techniques in our own work to further support the author's conclusion that optimality is rarely observed in human decision making. The models we use are often process models that are not formulated within the likelihood function, prior, cost function, and decision rule (LPCD) framework, although variables or relationships among them within the models could be used to derive LPCD quantities. Hence, LPCD may not characterize the actual computations performed during decision making. We illustrate these points through consideration of two topics, as discussed subsequently.

Although the authors note many suboptimal results in the signal detection theory model, it may be useful to separate the question of whether people are optimal from the question of whether their behavior is best understood in terms of the elements of LPCD. Some other process models might allow the constructs of LPCD to be seen as descriptions of the outcome of a learning process that does not employ these constructs directly. Turner et al. (2011) developed a category-learning model that encodes feature information about categories – the more general case of signal detection – and gradually forms a representation of

category structure that resembles the learning process observed in several experiments. The model has explicit mechanisms that bind feature information to category structure, but also allows the degree of association through repeated presentations to affect the acquisition of new or consistent knowledge. When feature-to-category associations are first being formed, changes in the strength of associations are rapid, enabling a general sense of category structure to be established quickly. However, with continued and consistent associations, the changes in associative strengths asymptote, meaning that further stimulus exposure does not necessarily change the category representation. In this model, both prior and likelihood information are latent in the category knowledge distributions. When the model is fitted to human data, it only conforms to optimal Bayesian integration under some parameterizations (Anderson 1991; McClelland 2013; Sanborn et al. 2010; Shi et al. 2010; Turner, under review; Turner & Van Zandt 2014; Turner et al. 2011). Turner (under review) investigated the degree to which the Bayes optimal version of the model could account for data across three experiments that exploited unannounced shifts in the properties of the feature-to-category maps over time. Across all experiments, the optimal version of the model only accounted for 11 out of 135 subjects, suggesting generally weak evidence for optimal category learning across different manipulations.

We have also observed suboptimal performance in perceptual decision making when subjects are asked to integrate multiple cues to form a decision. In the context of multimodal integration, Turner et al. (2017) investigated how subjects integrate two sources of information – auditory and visual cues – to arrive at a single perceptual representation. In their experiment, subjects completed three conditions, varying on whether subjects experienced (1) visual information alone, (2) auditory information alone, or (3) both streams of information together (the bimodal condition). Having the two unimodal conditions allowed Turner et al. to isolate the individual contribution of each modality in the decisions observed in the bimodal condition. Following the literature on optimal decision making, Turner et al. examined three different methods for integrating the two streams of information. The first two considered different forms of optimal integration, whereas the third method simply allowed the weights assigned to each cue to be freely estimated. By allowing the modality weighting parameters to freely vary, the third model could be compared to the two optimal models. If the (posterior estimate of the) weight parameter for the third model overlapped with the weights derived from either optimal model, then subjects could conceivably be combining the cue information together in an optimal fashion. However, of the six subjects in Turner et al., only two subjects showed any evidence of such an overlap, indicating that auditory and visual cues were not being integrated optimally by most of the participants.

In another multiple source integration study of perceptual decision making, Hotaling et al. (2015) examined optimality of inference with an experiment in which participants categorized faces based on resemblance to a family patriarch. Each face was split into two parts, and the amount of evidence contained in the top and bottom halves of each test face was independently manipulated. In one condition, the two halves were aligned to appear as a configural whole face; in another condition, the two halves were misaligned into separate parts. These data allow us to investigate a canonical example of suboptimal information integration from the judgment and decision-making literature, the dilution effect, which refers to the finding that adding weak

but positive evidence to strong positive evidence for a hypothesis decreases the estimated support participants derive for that hypothesis, rather than increasing it as required by Bayesian inference. Hotaling et al. used a multicomponent information accumulation model, a hybrid optimal/averaging model of information integration, to successfully account for key accuracy, response time, and dilution effects. The model weights represented the probability of attending to optimal or suboptimal information, which provided a parameter to assess an individual's approximation to optimal integration. Splitting the top and bottom halves of a face, a manipulation meant to encourage controlled integration of information, produced farther from optimal behavior and larger dilution effects.

Suboptimalities for sure: Arguments from evolutionary theory

Rob Withagen,[a] John van der Kamp,[b]
and Matthieu de Wit[c]

[a]Center for Human Movement Sciences, University of Groningen, University Medical Center Groningen, 9713 AV Groningen, The Netherlands; [b]Amsterdam Movement Sciences, Faculty of Behavioral and Movement Sciences, Vrije Universiteit Amsterdam, 1081BT Amsterdam, The Netherlands and [c]Department of Neuroscience, Muhlenberg College, Allentown, PA 18104.
r.g.withagen@umcg.nl   https://www.rug.nl/staff/r.g.withagen/
j.vander.kamp@vu.nl   https://vu-nl.academia.edu/johnvanderkamp
matthieudewit@muhlenberg.edu   https://matthieumdewit.wordpress.com/

Abstract

Rahnev & Denison (R&D) addressed the issue of (sub)optimalities in perception but only made a passing reference to evolutionary thinking. In our commentary, we concur with the authors' claim that evolution does not work toward optimalities, but argue that an evolutionary perspective on perception questions the Bayesian approach that the authors adopted.

Although Rahnev & Denison (R&D) addressed the question of (sub)optimalities in perception, they mentioned evolutionary thinking only briefly. However, the theory of evolution can shed light on this important issue, suggesting that suboptimalities in perception (and other biological functions) are to be expected. The reason for this is twofold.

First, evolution by natural selection works if and only if three conditions are met: There is variability among the members of a species; the variability is hereditable; and there is differential reproduction based on this variability (e.g., Lewontin 1970). If we apply this line of thinking to the perceptual realm, it implies that the perceptual systems of members of a species vary *and* differ in degree of fitness – some members perceive the environment more adaptively than other members. And although the power of "natural selection in the wild" (Endler 1986) should not be underestimated, it is unlikely that it will eliminate *all* the variability in the evolving population (e.g., Mayr 2004). After all, many perceptual tasks are not crucial for survival and reproduction, allowing

for a fair amount of variability *and* suboptimality among the members of a species (e.g., Withagen & Chemero 2009).

A second reason why suboptimalities are to be expected is that there are always multiple constraints that have a severe influence on evolving organisms. As Gould and Lewontin (1978) noted, organisms are "integrated wholes, with *Baupläne* […] constrained by phyletic heritage, pathways of development, and general architecture" (p. 581; emphasis in original). These constraints (e.g., Dawkins 1982; Gould & Lewontin 1978; Mayr 2002) imply that it is unlikely that the best possible (i.e., optimal) perceptual system is always available in the population that natural selection is working on (Withagen 2004).

Although an evolutionary analysis of perception predicts the suboptimalities that R&D discussed and defended to some extent, it questions the overall Bayesian approach that they adopted. As Darwin (1859/1985) had already claimed toward the end of *On the Origin of Species,* his theory of evolution by natural selection provides psychology with "a new foundation" (p. 458). Yet when it comes to theorizing about perception, many post-Darwinian theories of perception follow pre-Darwinian theories both in the conceptualization of what perception is (i.e., a mental state residing in the head) and how it comes about (i.e., by means of inferential processes). The Bayesian account that R&D adopted is no exception. Although this account is often connected to Helmholtz's perspective, ultimately the gist of both accounts can be traced back to Descartes' theory of perception—the stimulus information that reaches the senses is impoverished; hence, inferential processes are needed to gain knowledge of the environment (Reed 1982). In fact, when theorists of perception adopt evolutionary thinking, they often consider it as an "afterthought," one that does not affect the fundamental principles of their theorizing (Heft 2007, p. 92).

Gibson was one of the first to complain about this situation. As he put it in his book *The Senses Considered as Perceptual Systems*, "[t]he classics of vision were unaffected by evolutionary considerations or by knowledge of animal behavior but nevertheless they dominate the theories of perception" (Gibson 1966, p. 155). Gibson took evolutionary considerations seriously and argued that they require us to rethink the fundamentals of the science of perception. We agree. After all, from an evolutionary perspective, the main function of perception is not to *gain knowledge of the environment*, as many theories of perception hold, but to *guide our actions in the environment* (e.g., Anderson 2014; Chemero 2009). Animals have to perceive what Gibson called affordances, the action possibilities the environment provides them.

Crucially for present purposes, Gibson claimed that an evolutionary approach questions the dominant assumption of impoverished stimulus information that, as mentioned earlier, also underlies the Bayesian approach. Fundamentally, this assumption implies that the animal is in a state of uncertainty about the environment. As Müller (1837–1840/1938), Helmholtz' mentor, put it, "[i]n our intercourse with external nature it is always our own sensations that we become acquainted with, and from them we *form conceptions* of the properties of external objects, *which may be relatively correct*" (p. 1068; emphases added). If animals are in such a state of uncertainty about what is "out there" in the environment, it is hard to explain how animals can generally adaptively cope with the environment. Hence, Gibson replaced the representational theory of perception with a so-called contact theory of perception (Dreyfus & Taylor 2015). In his view, perception is not a mental state in the head, but a "keeping-in-touch with the world" (Gibson 1979/1986, p. 239). Gibson argued that this direct contact can be established because there is information

in the ambient energy arrays that specifies the affordances, that is, that informs about them. And if animals detect these specifying variables, a direct and adaptive perceptual contact with the affordances in their environment is established. Notably, the availability of such rich information obviates the need for inference, be it Helmholtzian or Bayesian.

One might argue that this focus on specifying variables is not in keeping with the suboptimal performances that have been demonstrated in empirical studies and are implied by the above evolutionary analysis. Indeed, to be a truly evolutionary view, the Gibsonian perspective needs to recognize that animals (occasionally) behave and perceive suboptimally (e.g., Chemero 2009; Withagen 2004). However, suboptimality (or optimality, for that matter) need not be explained in Bayesian terms of uncertainty, but can also be accounted for in terms of the informational variables that animals detect (e.g., Jacobs & Michaels 2007). The numerous patterns in the ambient energy arrays that animals might possibly detect differ in degree of usefulness, and the adaptiveness of the perceptual grip on, and actions in the environment is determined by which of these variables is picked up (e.g., de Wit et al. 2015; Withagen 2004; Withagen & Chemero 2009). Contrary to the Bayesian approach that R&D adopted, such a perspective stays closer to the fundamental change in thinking that Darwin's evolutionary program implies.

Leveraging decision consistency to decompose suboptimality in terms of its ultimate predictability

Valentin Wyart [ORCID]

Laboratoire de Neurosciences Cognitives et Computationnelles, Institut National de la Santé et de la Recherche Médicale, Département d'Etudes Cognitives, Ecole Normale Supérieure, PSL University, 75005 Paris, France.
valentin.wyart@ens.fr   http://lnc2.dec.ens.fr/inference-and-decision-making

Abstract

Although the suboptimality of perceptual decision making is indisputable in its strictest sense, characterizing the nature of suboptimalities constitutes a valuable drive for future research. I argue that decision consistency offers a rarely measured, yet important behavioral metric for decomposing suboptimality (or, more generally, deviations from any candidate model of decision making) into ultimately predictable and inherently unpredictable components.

The function of perceptual decision making is to make sense of an uncertain environment whose current state is only partially observable through imperfect sensory measurements. At this "computational" level of description (Marr 1982), the question of whether human observers process information available in their environment as accurately as possible has been the subject of a large body of work in the recent years. Rahnev & Denison (R&D) make an important case that the labeling of perceptual decisions as suboptimal does not yield much insight regarding the precise nature of suboptimalities. Rather than focusing on which aspects of decision making are suboptimal in a particular

task, R&D propose an alternative road map for future research, which consists in developing a general "observer model" of perceptual decision making across many different tasks.

This proposition is particularly attractive because optimality is often undefined for certain aspects of a task. For example, human observers carry priors, which can be suboptimal for a particular laboratory experiment, but optimal when considering the overall statistics of natural environments (Girshick et al. 2011). Similarly, the cost function assigned to most perceptual decisions is unknown, such that seemingly suboptimal biases in the strictest statistical sense can be seen as optimal in terms of efficient coding (Wei & Stocker 2015). However, the endeavor proposed by R&D is likely to face challenges for which the framework outlined in their Box 1 will be of little help.

Perhaps most strikingly, the long list of suboptimalities summarized by R&D in Table 1 is by definition non-exhaustive. Therefore, it remains unknown how much an observer model fails to capture unspecified suboptimalities in any given task. The approach proposed by R&D, which consists in specifying additional forms of suboptimalities and then testing whether they improve model fits, sounds a bit like fumbling in the dark. When will one know that a current "observer model" captures a dominant fraction of suboptimalities? Quality-of-fit metrics are only meaningful in a relative sense – that is, for comparing candidate models (Palminteri et al. 2017) – and they are thus blind to "how wrong" a given model is in an absolute sense.

To address this difficult question, it is important to consider not which aspects of the decision process may be suboptimal, but whether suboptimalities produce random or deterministic variability in behavior – a decomposition known as the "bias-variance tradeoff" in statistics. These two forms of suboptimalities map onto the classical distinction between noise and bias – for example, sensitivity and criterion in signal detection theory (Green & Swets 1966). Independently of any specific theory, the difference between random and deterministic suboptimalities is important in this context because biases trigger suboptimal decisions, which are ultimately predictable, whereas noise triggers suboptimal decisions, which are inherently unpredictable. If the long-term goal of R&D's framework is to predict decision behavior across tasks, then knowing the upper bound on the predictability of decision making in any given task is indispensable.

Although the theoretical distinction between random and deterministic suboptimalities may seem at first abstract and distant from behavior, the two produce antagonistic effects on a simple behavioral metric that can be easily measured in most perceptual tasks: the consistency of decisions across two repetitions of the exact same trial/condition (Wyart & Koechlin 2016). Indeed, deterministic biases tend to increase the consistency of decisions, whereas random noise tends to decrease the same quantity. Therefore, I propose to use decision consistency to decompose suboptimality (or, more generally, deviations from any candidate model of decision making) into a bias (predictable) term and a variance (unpredictable) term. In practice, the only modification that needs to be made to existing tasks is that the same trial/condition has to be presented at least twice, in order to measure the fraction of repeated trial pairs for which decisions are matched – irrespectively of whether they are right or not.

In terms of modeling, the approach consists in comparing behavior to simulations of a candidate model of decision making in terms of decision consistency. If simulated decisions are less consistent across repeated trial pairs than human decisions are,

then a fraction of the noise fitted by the model is attributable to unspecified biases – in other words, to unknown sources of suboptimalities that have not been captured by the model rather than to true randomness in the decision process. This discrepancy can be quantified in terms of the fraction of random variance in the model that can be pinned down to unknown biases. As an example, we obtained a value of 32% in a canonical probabilistic reasoning task when fitting an optimal model corrupted by noise to human decisions (Drugowitsch et al. 2016). This indicates that about a third of deviations from optimality are attributable to deterministic, predictable biases. This decomposition of suboptimality into ultimately predictable biases and unpredictable noise can serve not only to measure the effective precision of the decision process in a given task (i.e., the absolute variance of the unpredictable noise term), but also to determine how much a candidate model of decision making lacks additional, to-be-specified biases.

Like any approach, this bias-variance decomposition of suboptimalities has its limits. First, the bias term will by definition capture only within-trial biases, not sequential biases that propagate across successive trials (Wyart & Koechlin 2016). Sequential biases should therefore be specified in the model to be accounted for in the analysis of decision consistency. Second, biases that change over the course of the experiment will spill into the variance term. To control for the existence of such time-varying biases, the experimental design can be made such that the distance between repeated trials is varied across trial pairs (Drugowitsch et al. 2016).

An important corollary of R&D's road map is to build an observer model that provides an accurate split between suboptimal biases and true randomness in decision making. Therefore, analyzing decision consistency should become standard practice to determine how much a candidate model approximates the decision process as good as it possibly can.

Descending Marr's levels: Standard observers are no panacea

Carlos Zednik[a] and Frank Jäkel[b]

[a]Otto-von-Guericke-Universität Magdeburg, D-39016 Magdeburg, Germany and [b]Technische Universität Darmstadt, Centre for Cognitive Science, D-64283 Darmstadt, Germany.
carlos.zednik@ovgu.de   https://sites.google.com/site/czednik/
jaekel@psychologie.tu-darmstadt.de

Abstract

According to Marr, explanations of perceptual behavior should address multiple levels of analysis. Rahnev & Denison (R&D) are perhaps overly dismissive of optimality considerations at the computational level. Also, an exclusive reliance on standard observer models may cause neglect of many other plausible hypotheses at the algorithmic level. Therefore, as far as explanation goes, standard observer modeling is no panacea.

Rahnev & Denison (R&D) argue that "we should abandon any emphasis on optimality or suboptimality and return to building

a science of perception that attempts to account for all types of behavior" (sect. 1, para. 4). We agree that the current fixation on optimality is unhealthy. At the same time, however, we question whether standard observers are really sufficient to "account for" perceptual behavior. Because they cut across different tasks, they may provide some much-needed unification (Colombo & Hartmann 2017). Nevertheless, they are by themselves unlikely to constitute full-fledged explanations. Following Marr (1982), explanations of perceptual behavior should answer questions at three distinct levels of analysis. Alas, it is not clear how standard observers help descend Marr's levels from the computational level to the algorithmic and implementational levels.

At the computational level, investigators ask "what" a perceptual system is doing and "why." The popularity of ideal observers (Swets et al. 1961) stems in part from answering both of these questions. Because ideal observer models are tweaked to fit behavioral data, they provide mathematical descriptions of "what" a perceptual system is doing. They also answer questions about "why": A perceptual system behaves as it does because that behavior is optimal for the task (Bechtel & Shagrir 2015).

Like ideal observers, standard observers address "what" questions at the computational level by fitting behavioral data. Whereas ideal observer models are often criticized for failing to address questions below the computational level (Jones & Love 2011), R&D's standard observer models also address "how" questions at the algorithmic level. Many (but not all) of the hypotheses in Table 1 of the target article emphasize algorithmic-level features such as capacity limitations, imprecisions, ignorance, or the inability to employ complex decision rules. These algorithmic-level aspects are easily accommodated once optimality is given up. In other words, R&D trade in the ability to answer questions about "why" for an improved ability to answer questions about "how."

We applaud this shift in emphasis from "why" to "how." However, we feel that (a) "why" questions should not be dismissed quite so quickly, and that (b) properly answering "how" questions may require taking into account hypotheses that are unlikely to be considered within the standard observer approach.

Regarding (a), R&D's dismissive attitude toward optimality is understandable insofar as the explanatory value of "why" questions remains unclear (Danks 2008; but cf. Shagrir 2010). Nevertheless, such questions can still have pragmatic import; considering what a system is supposed to be doing may lead to an improved understanding of what it is actually doing. R&D admit as much in section 4.2 but do not go far enough. Many historical attempts to uncover mechanisms in biology and neuroscience begin by specifying these mechanisms' roles in the containing environment: The heart is viewed as a pump for the circulatory system (Bechtel 2009), and dopamine is known to contribute to the regulation of emotions (Craver 2013). In this vein, Swets et al. (1961, p. 311) argue that ideal observers should be used not only to describe optimal behavior, but also as a "convenient base from which to explore the complex operations of a real organism." In line with this view, we believe that perceptual scientists may productively tweak an ideal observer's optimal solution so as to eventually arrive at an organism's actual solution (see also Zednik & Jäkel 2016). Hence, although we agree that it is a mistake to rely too heavily on the unclear explanatory value of optimality considerations, we believe that it would be a mistake to dismiss these considerations altogether.

Regarding (b), more should be said about the transition from "what" and "why" questions at the computational level to "how" questions at the algorithmic level. We have previously argued that Marr's hierarchy can be descended by applying heuristic strategies to identify candidate hypotheses at lower levels of analysis (Zednik & Jäkel 2014; 2016). Many of the hypotheses summarized in Table 1 result from the "push-down" and "plausible-algorithms" heuristics: Whereas the former involves hypothesizing that an ideal observer's computational-level structure reflects an algorithmic-level description of the underlying mechanism, the latter involves adapting this description according to established psychological principles about, for example, capacity limitations. Additionally, R&D's plea for standard observers that can unify models across different tasks attaches great importance to what we have called the "unification" heuristic. Many other useful heuristics are not considered in the target article, however. In particular, some of the most promising recent work is driven by the "tools-to-theories" heuristic (cf. Gigerenzer 1991), in which algorithms developed in, for example, machine learning and Bayesian statistics are co-opted as algorithmic-level hypotheses for explaining how real organisms approximate (or fail to approximate) ideal observers. In particular, Sanborn et al. (2010) suggest that particle filters – a class of algorithms for approximating Bayesian inference – accurately describe the algorithms that humans deploy to learn categories. Interestingly, these algorithms approximate priors and posteriors through samples and thereby suggest very different components and processes than the original ideal observers. Hence, whereas developing standard observers may be one viable way of addressing "how" questions at the algorithmic level, other approaches may lead to different answers that also merit consideration.

In summary, although we agree that perceptual scientists should in fact shift from questions about "what" and "why" to questions about "how," we warn against thinking of the standard observer framework as a panacea. For one, "why" questions may continue to play an important role in the process of scientific discovery at the computational level and should not be dismissed prematurely. For another, although standard observers may be one promising way to answer "how" questions at the algorithmic level, other approaches might yield diverging and even incompatible answers. Finally, very little has yet been said about "where" questions at the implementational level (Stüttgen et al. 2011; Zednik 2017). Therefore, although standard observer models may play an important role in explanations of perceptual behavior, until we have satisfactory explanations on all three of Marr's levels, we should be patient and let different research strategies run their course.

Non-optimal perceptual decision in human navigation

Mintao Zhao[a] and William H. Warren[b]

[a]School of Psychology, University of East Anglia, Norwich NR4 7TJ, United Kingdom and [b]Department of Cognitive, Linguistic, and Psychological Sciences, Brown University, Providence, RI 02910.

Mintao.Zhao@uea.ac.uk
https://www.uea.ac.uk/psychology/people/profile/mintao-zhao
Bill_Warren@brown.edu   http://www.cog.brown.edu/Research/ven_lab/

Abstract

We highlight that optimal cue combination does not represent a general principle of cue interaction during navigation, extending Rahnev & Denison's (R&D) summary of nonoptimal perceptual decisions to the navigation domain. However, we argue that the term "suboptimality" does not capture the way visual and non-visual cues interact in navigational decisions.

We appreciate Rahnev & Denison's (R&D) brave target article for both its comprehensive summary of non-optimal perceptual decisions in various behaviors and its stringent critique of the conceptual shortcoming of optimality in characterizing human perception. Nonetheless, R&D's description of non-optimal perceptual decisions as *suboptimal* suggests that they are still trapped by the "optimality doctrine," rather than abandoning it. Taking studies of cue combination in navigation as an example, we argue (1) that perceptual decisions in navigation are not optimal in the sense of Bayesian theory, and (2) that suboptimality does not capture the nature of cue interaction in navigation.

Within the framework of the "Bayesian brain" (e.g., Knill & Pouget 2004), researchers have argued that perceptual decisions in navigation are statistically optimal (Cheng et al. 2007; Nardini et al. 2008). According to this view, when independent sources of spatial information (e.g., visual landmarks and idiothetic information about self-motion) are available for judging one's location or orientation, they are combined based on the reliability of each source. The greater the reliability of a source, the more heavily it is weighted in determining the navigator's decision. Under certain circumstances, the relative weighting of visual and self-motion cues in human navigational decisions conforms nicely to the prediction of Bayesian integration (e.g., Chen et al. 2017; Nardini et al. 2008; Zhao & Warren 2015b; see also Xu et al. [2017] for cue integration in spatial reorientation).

However, optimal cue combination does not represent a general principle of cue interaction in navigational decisions. For example, it has difficulty accounting for the competition among spatial cues in determining the direction of locomotion. Although visual and self-motion cues may be optimally integrated to reduce the variability of spatial judgments (e.g., Chen et al. 2017; Nardini et al. 2008), these cues often compete to determine the direction in which a navigator should go (Tcheang et al. 2011; Zhao & Warren 2015b). Visual cues often "veto" self-motion cues when they provide conflicting estimates of orientation or location; when such conflict becomes substantially large, the dominance reverts to self-motion cues (Foo et al. 2005; Mou & Zhang 2014; Zhang & Mou 2017; Zhao & Warren 2015b; see Cheng et al. 2007 for a review). This competition between visual and self-motion information occurs in both human and nonhuman animal navigation and manifests in terms of both behavioral and neurophysiological responses (e.g., Etienne & Jeffery 2004; Yoder et al. 2011). Such cue dominance in navigation indicates that spatial cues are not generally combined in a statistically

optimal or even suboptimal fashion, posing a challenge to Bayesian optimality in navigation. Without additional assumptions, the reliability-based theories of optimal cue combination predict neither the dominance of less reliable cues nor the coexistence of cue combination and cue competition in the same spatial judgment (Zhao & Warren 2015b).

Another challenge to optimal cue combination in navigation is that many factors irrelevant to cue reliability also modulate cue interactions. One such factor is feedback about performance. Distorted feedback can change the reliability of visual or self-motion cues and their combination during navigation (Chen et al. 2017). Therefore, in addition to cue reliability per se, subjective evaluation of cue reliability also contributes to the weighting of spatial cues in navigation. Another factor is related to previous experience. Exposure to a stable visual environment can completely "silence" the contribution of self-motion cues to navigation (Zhao & Warren 2015a), whereas experience with an unstable visual world can reduce or "switch off" the reliance on visual cues (Chen et al. 2017; Zhao & Warren 2015a). Such experience-dependent cue interaction is observed in both human and nonhuman animal navigation (e.g., Knight et al. 2014) but is rarely considered in formulating optimal cue combination in navigation. The last factor we want to highlight here is individual differences. Optimal cue combination is often demonstrated at the group level. However, whether spatial cues are combined and, if so, the optimality of integration can vary substantially between individuals (Chen et al. 2017; Cheng et al. 2007; Nardini et al. 2008; Zhao & Warren 2015b).

As R&D mention, these challenges to Bayesian optimality might be addressed by adjusting assumptions about the likelihood, prior, cost function, and decision rules (LPCD), and their combinations – although this renders Bayesian models unconstrained and unfalsifiable (Bowers & Davis 2012a; Jones & Love 2011). But before determining which components of LPCD are responsible for nonoptimal decisions, a prior question is why they should be optimal in the first place. If perceptual decisions need not to be statistically optimal, then seeking the causes of suboptimality will not help us to build models of perception and cognition. We see little evidence to justify such necessity. For example, optimal perceptual decisions assume that humans are rational decision makers, which is often not the case (Kahneman et al. 1982b). In navigation, when two spatial cues point in different directions, optimally integrating them would lead one to walk somewhere in between, guaranteeing that one gets lost. Ultimately, evolution does not necessarily produce optimal solutions, given the rates of natural selection and environmental change, pleiotropy and other structural constraints, the heterogeneity of populations, and the random effects of genetic drift.

Without establishing the necessity of optimal cue combination in navigation, referring to the over- or underweighting of cues as "suboptimal" still buys into the optimality approach. It implies that spatial cues *should* interact in a Bayesian optimal manner, and if they do not, some aspects of LPCD need to be better specified. This approach runs the risk of overlooking the cognitive and neural processes that actually underlie cue interactions (see also Jones & Love 2011). In fact, decades of research has shown that navigational decisions in mind and brain are often captured by one of two cues rather than their optimal – or suboptimal – combination (Etienne & Jeffery 2004; Yoder et al. 2011).

Authors' Response

Behavior is sensible but not globally optimal: Seeking common ground in the optimality debate

Dobromir Rahnev[a] ⬤ and Rachel N. Denison[b] ⬤

[a]School of Psychology, Georgia Institute of Technology, Atlanta, GA 30332 and [b]Department of Psychology and Center for Neural Science, New York University, New York, NY 10003.
drahnev@gmail.com   rachel.denison@nyu.edu
rahnevlab.gatech.edu   racheldenison.com

doi:10.1017/S0140525X18002121, e251

Abstract

The disagreements among commentators may appear substantial, but much of the debate seems to stem from inconsistent use of the term *optimality*. Optimality can be used to indicate *sensible* behavior (adapted to the environment), *globally optimal* behavior (fully predicted from optimality considerations alone), *locally optimal* behavior (conforming to a specific model), and *optimality as an empirical strategy* (a tool for studying behavior). Distinguishing among these different concepts uncovers considerable common ground in the optimality debate.

R1. Introduction

Is perceptual decision making optimal or not? At a superficial level, our commentators could not disagree more. On one hand, many commentators strongly defended optimality. **Stocker** wrote a "credo for optimality," **Howes & Lewis** proposed that "a radically increased emphasis on (bounded) optimality is crucial to the success of cognitive science" (para. 1), and **Shimansky & Dounskaia** suggested that examples of suboptimality are in fact optimal if neural effort is included in the cost function. On the other hand, **Withagen, van der Kamp, & de Wit** (**Withagen et al.**) claimed that there are "suboptimalities for sure," **Chambers & Kording** endorsed the view that "humans deviate from statistically optimal behavior" (para. 2), and **Salinas, Seideman, & Stanford** (**Salinas et al.**) argued that "a normative benchmark … is by no means necessary for understanding a particular behavior" (para. 8). Hence, it may appear that we, as a field, are hopelessly divided.

The division extended even to the commentators' treatment of evolution. Whereas some commentaries expressed the view that optimality "directly follows from the theory of evolution" (**Stocker**, para. 3; also endorsed by **Moran & Tsetsos** and possibly by **Shimansky & Dounskaia**), two other commentaries argued that "evolution does not work toward optimalities" (**Withagen et al.**, abstract; also endorsed by **Zhao & Warren**).

Nevertheless, a close reading of the commentaries convinced us that there is perhaps more common ground than it appears on the surface. In fact, a large part of the disagreements can be traced to a failure to clearly distinguish substantive issues (i.e., issues relating to the nature of human behavior) from tools (i.e., the methods we use to uncover the nature of behavior). We

believe that appreciating this distinction and maintaining it zealously will go some way toward resolving disagreements in the optimality debate.

We organize this response in two parts. In the first part (sects. R2–R4), we discuss three main topics of disagreement:

(1) Are people globally optimal? (sect. R2)
(2) Is assuming that people are optimal a fruitful empirical strategy? (sect. R3)
(3) Should we adopt Bayesian approaches or not? (sect. R4)

In the second part of this response, we discuss a number of topics that arose from our target article, such as the merits of focusing on building a standard observer model (sect. R5) and specific comments regarding individual approaches or findings (sect. R6). We conclude (sect. R7) on a hopeful note that rather than entrenching fissures in the field, the current collection of articles would contribute to a deep level of agreement and give us the tools to express and recognize this agreement.

We include a Glossary to facilitate communication and as a primer to those new to the field. Although our target article was relatively narrowly focused on perceptual decision making, this response is intended to generalize more broadly to all areas of study where optimality considerations are relevant (**Barth, Cordes, & Patalano** [**Barth et al.**]).

R2. Optimality as a substantive claim: *Sensibleness* versus *global optimality*

The first area of disagreement among the commentators, and the topic of our target article, is whether people's behavior is optimal. One set of commentators considered optimality a foregone conclusion (**Geurts, Chetverikov, van Bergen, Zhou, Bertana, & Jehee** [**Geurts et al.**]; **Howes & Lewis**; **Moran & Tsetsos**; **Shimansky & Dounskaia**; **Stocker**). For example, one of these commentators argued that there is no alternative to optimality because non-optimality means that "we reason in a deliberately inconsistent [i.e., random] way" (Jaynes 1957/2003, as cited by Stocker, para. 4). The surprise and horror could be sensed in some of these commentaries: Could so many scientists really believe that people's behavior comes out of thin air and makes no sense? Do they really believe that humans act randomly rather than in ways that are shaped by our environments?

At the same time, another set of commentators considered lack of optimality a foregone conclusion (**Bachmann**; **Booth**; **Brette**; **Withagen et al.**). These commentators argued that optimality in its strictest sense is either not well defined or clearly unachievable. A similar level of surprise and horror could be sensed in some of these commentaries: Could so many scientists really believe that people's behavior is optimal given that it was produced by the messy process of evolution? How can anyone claim that optimality is ever achievable in a biological system?

Although there is likely a substantive disagreement between these two groups of commentators, we believe that a large part of the disagreement boils down to differences in terminology.

R2.1. The concept of sensibleness

Some pro-optimality commentators (e.g., **Stocker**) saw the alternative to optimal behavior to be "deliberately random" behavior. However, neither we nor a single commentator endorsed a view that behavior is random or makes no sense.

GLOSSARY. We provide our definitions of new terms introduced in the target article and this response, as well as standard terms that are in wide use in the literature.

New terms
Global optimality – the view that behavior can be fully predicted by optimality considerations alone after taking into account an organism's environment and evolutionary history. Such behavior is **globally optimal**. Researchers who hold this view are **global optimalists**.
Local optimality – optimality as defined for a particular model, regardless of the accuracy of that model.
LPCD – the four components of all Bayesian models: likelihood (L), prior (P), cost function (C), decision rule (D).
Optimality as an empirical strategy – an approach to studying behavior based on assuming optimality a priori. When suboptimal behavior is found, this strategy dictates that the model should be changed. Precludes inferences that behavior is optimal, because optimality is assumed.
Sensibleness – the view that behavior is generally adaptive and functional, motivated by a combination of an organism's goals and constraints, which are in turn driven by its environment and evolutionary history. Such behavior is **sensible**. Researchers who hold this view are **sensibilists**.

Standard terms
Bayesian statistics – a field of statistics in which Bayes' theorem is used to draw inferences from data.
Bayesian theories – scientific claims that the brain represents Bayesian variables and performs Bayesian computations (e.g., Bayesian components [LPCD] are explicitly represented by the brain; people actually calculate posteriors using Bayes' rule).
Bayesian tools – a mathematical formalism for decomposing decision behavior into LPCD components. Specifies how to combine a likelihood and prior to form a posterior, and how to derive a locally optimal decision rule from the posterior and a cost function.
Bounded rationality – a theoretical framework in which decision rules are selected to achieve good enough levels of performance, rather than the best possible performance.
Ideal observer model – a model that includes a full generative model of an observer's internal responses, mathematically specifies all of the components of the decision-making process, and postulates a locally optimal decision rule.
Normativity – having some standard of evaluation. Behavior that is normative is as it "should" be.
Optimality – the minimization of costs given constraints; "as good as possible."
Optimal model – equivalent to **ideal observer model**.
Process model – a general term for a mechanistic model of the processing underlying a specific behavioral task.
Rationality – often used in the context of economic decision making; can refer to internal consistency, pursuit of self-interest, or perfect strategic foresight, as defined by some mathematical model of behavior.

In fact, impassioned arguments against optimality were followed by statements about how animals "adaptively cope with the environment" (**Withagen et al.**, para. 6) or that "it is not difficult to imagine why" a particular case of suboptimality occurs (thus indicating that the suboptimality is not random [**Salinas et al.**, para. 6]).

We think it is fair to say that every single commentator agreed that behavior "makes sense" on some level. To highlight the substantive agreement, we find it useful to invent a term unburdened by previous inconsistent usage in order to communicate the concept that behavior "makes sense." We propose to use the term *sensibleness* (see Glossary).

Sensibleness is a weak concept. We define sensible behavior as behavior that is motivated by a combination of our goals and constraints, driven by the environment and evolutionary history. It is generally accepted that natural selection produces sensible behavior in that behavior is adaptive and functional (Rosenberg & McShea 2008) – though it could be maladaptive in particular circumstances or environments. Sensible behavior may be suboptimal according to some definitions while remaining a "feature" and not a "bug" (**Summerfield & Li**).

R2.2. The concept of global optimality

Many commentators embraced sensibleness but argued against optimality. Rejecting optimality for these commentators does not mean that behavior is random, but that behavior is not always as optimized as it possibly could be. Another way to put this is that behavior cannot be predicted by optimality considerations alone. However, statements like "people are optimal" or "perception is optimal" are not uncommon in the literature. These statements give the impression that optimality can in fact be used to predict all of our behavior.

Just as with the concept of sensibleness, we find it useful to introduce an unburdened term for this interpretation of optimality. We propose the term *global optimality* (see Glossary). Global optimality is the view that behavior can be fully predicted by optimality considerations alone after taking into account our environment and evolutionary history, which together determine our goals and the constraints on our behavior. Concretely specifying globally optimal behavior would require full knowledge of these goals and constraints. At an abstract level, global optimality entails that (1) given enough information, we can unambiguously specify the globally optimal behavior, and (2) people's behavior perfectly matches the globally optimal behavior.

Let us illustrate this position with two examples. In the famous marshmallow task, people are allowed to have one marshmallow immediately or two after a certain delay. The global optimality assumption holds that in this task one can identify the globally optimal course of action for every individual and that everybody takes this course of action (though what is globally optimal may be different for different people). Both behaviors are sensible because they can easily be justified, but only one of them can be globally optimal for a given individual. A second example, from perception, is the following: Should different images presented to the two eyes be merged, should they alternate, or should one image dominate in conscious perception? Again, each possibility is sensible (can be reasonably justified), but for every pair of images only one of these possibilities can be globally optimal. The global optimality assumption holds that it is theoretically achievable to identify the best possible percept and that people actually form that percept.

Global optimality is a strong concept. It implies that each organism is constantly at its most optimized state. In other words, global optimality holds that complex organisms are able to optimize jointly the thousands of behaviors in which they engage.

R2.3. The case against global optimality

There are both theoretical and practical arguments against global optimality. Such arguments are featured in the target article and in several of the commentaries.

Theoretically, it is unlikely that any optimization algorithm could find the global optimum in an optimization task as complex as jointly optimizing all behaviors, especially considering that the algorithm would also need to optimize resource allocation to different behaviors (**Mastrogiorgio & Petracca**). Given the trial-and-error process behind natural selection, genetic drift, and random genetic variation across individuals, it is doubtful that all behaviors are as optimized as they possibly could be (**Bachmann**; **Withagen et al.**; **Zhao & Warren**).

More practically, our target article illustrates that a global optimalist view is not appealing given existing data. Our approach was to define optimality in the most "standard" way and then survey hundreds of instances of suboptimality. As might be expected, several commentators took issue with the assumptions we used to define optimality. For example, **Shimansky & Dounskaia** argued that "inclusion of neural effort in cost function can explain perceptual decision suboptimality" (similar sentiments were expressed by **Howes & Lewis**; **Stocker**), whereas **Moran & Tsetsos** showed that assuming late noise in the system leads to qualitatively different predictions for optimal behavior. We anticipated such criticisms and in the second part of our target article argued that it is impossible to specify what is optimal in a global sense.

Indeed, although we define global optimality in the abstract, we doubt it can be specified in a way that would make it empirically testable. The commentators who defended optimality unfortunately did not seriously grapple with the issue of whether it is possible to define global optimality unambiguously – though this issue was addressed by many others such as **Bachmann**; **Brette**; **Love**; **Meyer**; **Noel**; and **Wyart**. It is easy to argue that a conclusion of suboptimality in some task is invalidated because it is based on wrong assumptions. However, that argument also invalidates all of the conclusions of optimality that are based on the same assumptions.

If we follow this line of reasoning, then we must conclude that there is no firm evidence that people are either globally optimal or suboptimal. This is in fact the conclusion we reached in our target article and is precisely why we argued that the optimal/suboptimal distinction is not useful. We were especially surprised to read **Stocker**'s admission that this distinction is "ill-defined," as it was followed by a fervent defense of optimality. Later on, Stocker claimed that optimality allows us to build "a quantitative model hypothesis *before* actually running the experiment" (para. 3). We see a contradiction here, which is at the heart of the issue at hand: If we do not have enough information to specify what is globally optimal, how can we generate any firm predictions based on the concept of optimality? All predictions can only be based on the assumptions built into a specific model and are therefore tentative and dependent on the validity of the model's assumptions.

We do not think that at present there is evidence for a global optimalist view and, given the definitional difficulties, are doubtful that there ever could be. We therefore consider the argument between sensibilists and global optimalists to be philosophical rather than scientific.

R2.4. Pinning down the commentators' philosophical commitments

The majority of our commentators (e.g., **Booth**; **Chambers & Kording**; **Cicchini & Burr**; **Love**; **Ma**; **Meyer**; **Salinas et al.**; **Simen & Balcı**; **Summerfield & Li**; **Wyart**; **Zednik & Jäkel**) appeared to be sensibilists but not global optimalists. Although it is impossible to prove that all behavior is sensible, both we and all commentators consider behavior to be generally sensible.

However, it was more difficult to discern whether any commentator is a global optimalist. In particular, the commentaries of **Geurts et al.**; **Howes & Lewis**; **Moran & Tsetsos**; **Shimansky & Dounskaia**; and **Stocker** could be interpreted as supporting a global optimalist position. Another interpretation, however, is that these commentators only argued that optimality is a fruitful empirical strategy, without committing themselves to a global optimalist view. We explore this position in section R3. We hope researchers can use the sensibilist/global optimalist distinction to clarify their positions regarding the nature of human behavior.

Even if most (or all) researchers turn out to agree that human behavior is sensible but not globally optimal, it is still likely that there are remaining differences in opinion about the substantive question of how close behavior is to global optimality. There is potentially a vast space between a sensibilist and a global optimalist. The disagreement is perhaps best captured by the simple question: "Assuming perfect knowledge of an organism, what percent of its behavior can be predicted using optimality considerations alone?" This question can be seen as defining a "global optimality scale." To be a global optimalist is to answer 100%. It is likely that researchers who are moved to defend optimality are high on the global optimality scale, whereas researchers who are moved to attack optimality are low on the global optimality scale. It should be stressed that where one places oneself on the global optimality scale is based on a philosophical perspective rather than empirical findings; in fact, it is unclear that we can ever find the correct answer or that such an answer even exists. We think of the global optimality scale as a useful shorthand for highlighting differences in opinion rather than as a topic for empirical investigation.

Hence, appreciating the distinction between sensibleness and global optimality does not necessarily resolve all our differences. However, it (1) forces us to clarify our positions, which may otherwise slide back and forth between implying global optimality or only sensibleness; (2) highlights agreements, which we think are more extensive than they appear on the surface; and (3) discourages arguments that misrepresent other researchers' views as belonging to 0% or 100% on the global optimality scale. We hope this distinction provides the language to express more precisely one's position and escape false either/or dichotomies. We are all optimalists when that means sensibleness; few (if any) of us are optimalists when that means 100% global optimality.

R2.5. Relating sensibleness and global optimality to other optimality-related terminology

We defined the terms sensibleness and global optimality to clarify areas of agreement and disagreement in the optimality debate. We think these terms are useful, because careful consideration of the commentaries convinced us that common optimality-related

terms are used inconsistently. Here we briefly review how our commentators used such terms and relate their meaning to our sensibleness/global optimality distinction.

R2.5.1. Optimality
Some authors expressed the view that optimality is an ill-defined concept (**Bachmann; Brette; Cicchini & Burr; Love; Meyer; Noel; Stocker; Summerfield & Li; Wyart**). Others talked about optimality as an a priori assumption (**Danks; Geurts et al.; Howes & Lewis; Moran & Tsetsos; Shimansky & Dounskaia;** Stocker). Still others described optimality as defined in the context of a specific model (Bachmann; **Barth;** Brette; **Chambers & Kording;** Geurts et al.; Howes & Lewis; **Ma; Mastrogiorgio & Petracca;** Moran & Tsetsos; **Nardini & Dekker;** Summerfield & Li, Stocker; **Turner, McClelland, & Busemeyer [Turner et al.];** Wyart; **Zednik & Jäkel; Zhao & Warren**). Importantly, there is significant overlap in these lists showing that the same people sometimes talk about optimality in different ways. This is not a criticism of our commentators; in fact, we made each of these three claims about optimality in our target article. Rather than contradicting each other, these different uses of the term *optimality* refer to different concepts. Statements about optimality as an ill-defined concept or as an a priori assumption refer to a "global" sense of optimality, which is best captured by our term *global optimality*. Statements about optimality in the context of a specific model have a "local" sense, which we call *local optimality* (see Glossary).

Note that all empirical studies test for optimality in the "local" sense. For all such studies, it should be understood that because the model itself is not definitely correct, findings of optimality/suboptimality do not allow statements about global optimality. That is, optimality/suboptimality according to a particular model only has implications about global optimality if the model faithfully captures all of the goals and constraints of decision making, including how they developed over evolutionary history. It can be safely assumed that no current model does, and therefore "local" and "global" optimality remain unbridgeable.

These considerations demonstrate that the word *optimality* can be used in a global sense to mean global optimality or in a local sense to refer to a specific decision rule in the context of a specific model. Because these concepts are very different, we advise that the term *optimality* be used only in the local sense. When a global meaning is desired, *optimality* should be replaced by *sensibility* or *global optimality*, depending on the intended interpretation. We adopt this practice in the rest of this response.

R2.5.2 Normativity, rationality, and ideal observer
When discussing the topic of optimality, many commentators also used related terms such as *normativity*, *rationality*, and *ideal observer*.

The *Stanford Encyclopedia of Philosophy* defines a normative (see Glossary) theory of decision making as "a theory of how people *should* make decisions" (Briggs 2017). Defined this way, normative implies that some decisions are better than others according to some standard, but it leaves entirely open how one determines better and worse. This definition of normative is closest to our term *sensible*. Indeed, some commentators (**Mastrogiorgio & Petracca; Summerfield & Li**) used the term in this way. However, the majority of the commentators (**Geurts et al.; Howes & Lewis; Moran & Tsetsos; Salinas et al.; Simen & Balcı; Stocker**) used the term seemingly synonymously with either local or global optimality, as they asserted

that a normative approach allows us to make specific predictions. In fact, we also equated normativity with optimality when we wrote that according to bounded rationality (see Glossary), evolution does not produce normative behavior. What we meant is that according to bounded rationality, evolution does not produce globally optimal behavior. However, Mastrogiorgio & Petracca criticized us as they interpreted this statement to mean that bounded rationality does not produce sensible behavior. Clearly, then, different researchers use normativity with different intended meanings.

Another common optimality-related term is *rationality* (see Glossary). Rationality was mostly used in the commentaries in the context of economic decision making (**Moran & Tsetsos; Summerfield & Li; Zhao & Warren**) to refer to concepts such as intransitivity (the notion that if you prefer A over B and B over C, then you should prefer A over C). However, rationality was also a component of two other terms, *computational rationality* and *bounded rationality*. **Howes & Lewis** seemed to use the term *computational rationality* as synonymous with global optimality, whereas **Mastrogiorgio & Petracca** described the long tradition of *bounded rationality* theories that use this term in ways that are closest to what we call sensibleness. Again, the term rationality does not have a single meaning.

A final optimality-related term that appeared in many commentaries was *ideal observer*. Most authors saw ideal observer models as examples of optimal models (**Geurts et al.; Salinas et al.; Schultz & Hurlemann; Simen & Balcı; Zednik & Jäkel**). That is, an ideal observer model includes a full generative model of an observer's internal responses and mathematically specifies all of the components of the decision-making process. The decision rule that minimizes costs can then be determined, and this is the rule an "ideal observer" uses. Because the optimal decision rule here is based on a specific model, the concept of *ideal observer* is equivalent to local optimality. **Stocker**, however, drew a distinction between ideal observer models and optimal models, where an ideal observer "only considers limitations in terms of information provided to the observer," whereas an optimal observer "also includes constraints that are internal to the observer" (para. 2). This is an interesting distinction, which may reflect early uses of the term ideal observer (e.g., Siegert's ideal observer, described in Lawson & Uhlenbeck 1950), but in the current literature, ideal observer models typically include at least internal noise and often additional constraints (Geisler 2011), which makes them indistinguishable from optimal models. For the most part, the term *ideal observer* does seem to be used in a consistent way.

R2.5.3. Adopting unambiguous terminology
Common optimality-related terms, then, have multiple interpretations and are used inconsistently. This is especially true for *optimality*, *normativity*, and *rationality*. On the other hand, terms such as *bounded rationality* and *computational rationality* have well-defined meanings but are more specialized. Finally, the term *ideal observer* is a good synonym for what we have been calling *local optimality*. However, when global statements are intended, we suggest the use of the terms *sensibleness* and *global optimality*.

We believe that the use of more unambiguous terminology will help the field find common ground in the optimality debate. For example, the disagreement regarding whether evolution makes us (globally) optimal – **Moran & Tsetsos; Shimansky & Dounskaia;** and **Stocker** argued yes; **Withagen et al.** and **Zhao & Warren**

argued no – could perhaps be settled using unambiguous terminology: Evolution produces sensible but not globally optimal behavior. Better terminology would not dissolve the substantive areas of disagreement, but it would help us form more nuanced positions. The distinction between global optimality and local optimality (ideal observer models) might help distinguish between the rather philosophical issue of the global nature of human behavior and the more concrete issue of how close behavior is to a specific ideal based on a particular set of assumptions.

R3. Optimality as an empirical strategy: Studying behavior by assuming optimality

Several commentators we identified as potential global optimalists (**Geurts et al.; Howes & Lewis; Moran & Tsetsos; Shimansky & Dounskaia; Stocker**) may instead have been defending the assumption of optimality as a useful *empirical strategy*, rather than as a fact about human behavior. Other commentators were skeptical of assuming optimality, given doubts about the substantive claim of global optimality (**Bachmann; Brette; Love; Meyer; Noel**), though they did not explicitly discuss optimality as an empirical strategy.

R3.1. What is optimality as an empirical strategy?

Cicchini & Burr; Geurts et al.; Howes & Lewis; Moran & Tsetsos; Shimansky & Dounskaia; Simen & Balcı; Stocker; and **Zednik & Jäkel** argued that, at least in some cases, optimality is useful as an a priori assumption for studying behavior. In this empirical strategy, it is assumed that the correct model of a task will have a locally optimal decision rule. The idea is that assuming local optimality (often along with various other standard assumptions, as described in the target article) constrains the problem of generating a model of a task. Critically, a finding of suboptimality leads to proposals about new constraints that will preserve the presumed optimality, and these new constraints are then independently validated.

For example, in a Bayesian framework, one specifies the relevant likelihood function, prior, cost function, and decision rule – what we call the *LPCD components* (see Glossary). *Optimality as an empirical strategy* (see Glossary) consists of assuming that the decision rule (D) should be optimal and adjusting the LPC components to make it so. These LPC components are new hypotheses that are then independently tested and either supported – leading to new discoveries – or disconfirmed – prompting new LPC hypotheses to be generated. Several commentaries (e.g., **Cicchini & Burr; Geurts et al.; Stocker**) provided compelling examples of this strategy at its best.

R3.2. Pitfalls of optimality as an empirical strategy

Although optimality as an empirical strategy is fruitful when used carefully, it invites a number of pitfalls that require consideration by researchers who use this strategy.

R3.2.1. Using the strategy to infer optimality

Unfortunately, there is a tendency to conflate optimality as an empirical strategy – a tool for generating hypotheses – and the substantive claim about the local optimality of human behavior. For example, **Stocker** wrote that "the optimality assumption is a well-supported, very useful assumption" (para. 3). There are

two separate statements here: (1) (global?) optimality is well supported, and (2) optimality as an empirical strategy is useful. We disagree with (1) but agree with (2). We think that the conflation of these two statements contributes to unnecessary disagreement in the field.

For researchers who use optimality as an empirical strategy, we urge clear separation of tool from substance. If local optimality is already assumed, it cannot be falsified (**Barth et al.**); it is an a priori commitment (as can be seen, for example, in **Geurts et al.; Howes & Lewis; Moran & Tsetsos**). The only things that can be uncovered are relevant LPC components, because the decision rule is already assumed to be optimal. Therefore, what should be advertised is the nature of the LPC components and not optimality, which was assumed all along.

R3.2.2. Accepting hypotheses without independent validation

It is inappropriate to accept hypotheses generated using optimality as an empirical strategy without independently validating them.

A tendency to overlook the independent validation step can be seen in some commentators' statements about what makes a model believable. For example, **Geurts et al.** contended that "precisely because human behavior matched that of the ideal observer, rather than some arbitrary formulation, the findings provided strong evidence for the computational theory" (para. 4). This statement seems to suggest that optimal behavior implies the veracity of the model even without independent verification. An apparent underemphasis of independent validation could be detected in statements by commentators who maintained that optimality as an empirical strategy can be used to "explain" behavior. For example, **Shimansky & Dounskaia** argued that "the explanatory power is perhaps the main advantage of an optimality approach" (para. 1), and **Stocker** contended that optimality has "provided quantitative but nonetheless intuitive explanations" (para. 1).

However, as we and others have pointed out, one can always find a set of LPC components that makes the decision rule locally optimal. Therefore, the ability of optimality as an empirical strategy to "explain" behavior is trivial; it is only useful as a tool for generating hypotheses. The strategy succeeds not when one finds a model that makes behavior locally optimal but when new hypotheses expressed as components of the model have been independently validated.

R3.2.3. Overemphasizing the decision rule

The optimality of the decision rule should not be emphasized over the other model components. In the target article, we advocated for moving away from such overemphasis, which we called "the optimality approach." In retrospect, this wording was unfortunate, as this phrase was interpreted by many commentators to mean what we now call "optimality as an empirical strategy." Optimality as an empirical strategy can invite overemphasis of the decision rule, because it is focused on making the decision rule optimal. But used carefully, it will not do so. In fact, optimality as an empirical strategy cannot discover anything about the decision rule, because the exact form of the decision rule is already assumed. Therefore, researchers interested in the decision rule should *not* use this strategy. Instead, they should test a variety of decision rules, which could be optimal or suboptimal (**Love; Ma**).

R3.3. Is optimality as an empirical strategy a fruitful strategy?

In the commentaries, it seemed that researchers who thought optimality as an empirical strategy was useful were high on the global optimality scale, and vice versa. However, views on tools and substance need not be this tightly linked. We believe that as long as behavior is sensible, optimality as an empirical strategy is likely to be fruitful.

At the same time, other research strategies can also be fruitful. A great many discoveries have been made by explicitly rejecting optimality (e.g., Kahneman & Tversky 1979) or not considering optimality (e.g., Gibson & Radner 1937; Hubel & Wiesel 1970; Müller-Lyer 1889). Hence, we agree with **Barth et al.; Danks**; and **Zednik & Jäkel** on the importance of maintaining a diversity of strategies for understanding human behavior. What is critical is to uncover phenomena that are replicable and generalizable, not the strategies that we use to uncover them.

R4. Bayesian approaches: Bayesian tools versus Bayesian theories

The third area of disagreement among our commentators relates to the usefulness of Bayesian approaches to perceptual decision making. A number of commentators appeared strongly either pro- or anti-Bayesian. Here we examine the areas of agreement and disagreement, as well as the relationship between Bayesian approaches and optimality.

As discussed in sections R2 and R3, a considerable part of the optimality disagreement could be attributed to whether one focuses on the nature of behavior or optimality as an empirical strategy. In a similar fashion, the conflict between pro-Bayesian and anti-Bayesian commentators may be largely attributable to whether a commentator focused on Bayesian tools or Bayesian theories. *Bayesian tools* (see Glossary) refer only to the mathematical formalism for decomposing decision behavior into likelihood, prior, cost function, and decision rule (the LPCD components); they make no scientific claims. *Bayesian theories* (see Glossary), on the other hand, entail scientific claims, such as that Bayesian components (LPCD) are explicitly represented by the brain and people actually calculate posteriors and decision rules using Bayes' rule.

Note that what we are calling Bayesian tools are different from Bayesian statistics (see Glossary), which were discussed by **Turner et al.** The term *Bayesian statistics* refers to a statistical analysis method that uses Bayes' theorem to draw inferences from data and is often contrasted with frequentist statistics.

R4.1. Bayesian tools

Multiple commentators argued for the usefulness of Bayesian tools in research on perceptual decision making (**Chambers & Kording; Cicchini & Burr; Geurts et al.; Howes & Lewis; Nardini & Dekker; Schultz & Hurlemann; Stocker**). Commentators noted that Bayesian tools provide a compact, transparent, and explicit formulation of how to integrate different pieces of information. The Bayesian formulation can also provide a benchmark for human performance (but see **Noel** for an argument against this claim). Commentators who seemed to adopt optimality as an empirical strategy also noted that Bayesian tools are a natural analytic framework for this strategy (Geurts et al.; Howes & Lewis; Stocker); optimality is naturally expressed in Bayesian terms. We agree with all of these arguments for the usefulness of Bayesian tools, and we used these tools extensively in our target article.

At the same time, it is important to acknowledge the limitations of Bayesian tools. Several commentators pointed out that Bayesian approaches do not provide insight into some of the hardest problems of perceptual decision science. **Summerfield & Li** noted that the LPCD framework does not automatically provide meaningful explanations of behavior and pointed to the idea of efficient coding as a more satisfying kind of explanation. **Brette** observed that a major challenge in developing a standard observer model under the LPCD framework is specifying the generative model. Perhaps another way to say this is that the hardest part of understanding perceptual decision making is understanding perception itself. In fact, according to Brette, an overemphasis on using Bayesian approaches can obscure the difficulties that perception poses and can lead researchers to ignore or underemphasize these difficulties. Finally, **Bachmann** argued that Bayesian models may not map well onto the phenomenology of conscious perception. We agree with these perspectives, too.

In the end, this is almost always the nature of a single tool (e.g., a hammer): It is useful in certain contexts (e.g., putting a nail in a wall) but is virtually never sufficient by itself to solve a larger problem (e.g., building a house). Overly high expectations of Bayesian tools can lead to both excessive reliance on them and excessive criticism of them.

R4.2. Bayesian theories

In the field at large, there is a substantive debate about Bayesian theories (i.e., theories that hold that the brain explicitly represents LPCD components and explicitly computes using Bayes' theorem). However, no commentator defended Bayesian theories outright, so here we focus on their limitations. **Simen & Balcı** reminded us of the computational demands of exact Bayesian inference, which they argued make Bayesian theories unlikely to be true. "Surely then," they remarked, "heuristics are the best we can do, in general" (para. 8). **Turner et al.** argued that sometimes behavior is not "best understood in terms of the elements of LPCD" (para. 3) and suggested that *process models* (see Glossary) could better capture the true mechanisms underlying behavior. We agree that these are both serious difficulties for Bayesian theories. **Withagen et al.** gave an evolutionary argument for rejecting a Bayesian approach and questioned a foundation of Bayesian theories, the notion of inference, arguing that the primary challenge for human perception does not stem from uncertainty about the environment. We agree that more attention should be given to the interaction between perception and action, but we consider dealing with uncertainty to be a central issue in perceptual decision making.

R4.3. Conflating Bayesian tools and Bayesian theories

In section R3, we distinguished between optimality as a substantive hypothesis and optimality as an empirical strategy. In a similar fashion, it is important to distinguish between Bayesian theories (which are substantive claims about the nature of the internal representation) versus Bayesian tools (which simply refer to an analytical framework). **Barth et al.** discussed this issue extensively and pointed out that successful use of Bayesian tools does not imply that Bayesian theories are true. **Danks** emphasized a similar distinction. He referred to Bayesian theories as a "realist" interpretation and to Bayesian tools as an "as-if"

interpretation of the Bayesian formalism. Danks explained the difference between Bayesian theories and Bayesian tools and the importance of not conflating them.

In fact, **Danks** criticized our target article for starting out by using Bayesian approaches as tools and then switching to using them as substantive theories. Let us clarify that we intended to use the LPCD framework purely as a tool. The positive proposal in section 5 of our target article was intended to be completely approach-agnostic. This is why we stated that the standard observer model may eventually take the form of "a Bayesian model … a 'bag of tricks' … a neural network … and so forth" (sect. 5.2, para. 1). The hypotheses that we thought might generalize across tasks, such as "placing a premium on accuracy" (sect. 5.2, para. 5) were not intimately linked to the Bayesian approach (we classified them under different LPCD components purely for convenience). We see the question of whether such hypotheses generalize across tasks as independent from whether one supports Bayesian theories. Nevertheless, Danks's criticism of our target article shows the practical difficulties involved in keeping Bayesian tools clearly separated from Bayesian theories and indicates the need for enhanced conceptual clarity on this issue.

R4.4. Bayesian approaches and optimality

Continuing the theme of conflating concepts that should be kept separate, the term *Bayesian* was sometimes equated with *optimal* in the commentaries, a common issue we addressed in the target article. For example, **Stocker** argued that "Bayesian decision theory and the optimality assumption are in many ways synonymous; without the latter, the former is not meaningful" (para. 3). This view was also endorsed by **Zhao & Warren**. We agree with these commentators that Bayesian tools are a natural choice if one adopts the use of optimality as an empirical strategy. However, optimality as an empirical strategy can be pursued independently of Bayesian tools (e.g., maximizing a non-Bayesian utility function), and Bayesian tools can be used without assuming optimality (e.g., an LPCD model with a suboptimal decision rule) – so the two should not be equated.

Bayesian tools and theories should also be clearly separated from claims about optimality (Ma [2012]). As just noted, the Bayesian formalism can be used to model decision behavior, whether that behavior is optimal or not, as with a suboptimal decision rule (**Chambers & Kording**). Similarly, Bayesian theories could be substantively correct in that the brain explicitly represents LPCD components and performs Bayesian computations, but if the brain represents the wrong components, then behavior could still be suboptimal. Hence, claims about both local and global optimality should be kept separate from Bayesian tools and theories.

R5. The standard observer model

In the first part of this response (sects. R2–R4), we addressed the three main areas of disagreement in the commentaries. In this second part, we discuss issues related to the standard observer model (sect. R5) and specific suggestions and topics (sect. R6).

R5.1. The benefits and limitations of a standard observer model

In our target article, we communicated a vision of a standard observer model that will predict observers' behavior on a wide variety of perceptual tasks. We then urged researchers to actively work toward building such a model. A number of commentators reflected on the benefits and limitations of this type of effort.

The benefits of building a standard observer model were recognized by a number of commentators (e.g., **Nardini & Dekker; Meyer; Simen & Balcı; Schultz & Hurlemann; Turner et al.; Wyart**). They noted the utility of having a model that applies across a variety of tasks and quantitatively predicts novel behavior. Further, the ability to make quantitative predictions would also make the standard observer model readily falsifiable.

An observer model is a prerequisite for making any new predictions. **Stocker** appeared to disagree, contending that the principle of optimality could generate predictions from first principles, because optimality "allows us to formulate with relative ease an observer model for any specific task" (para. 3). However, neither global nor local optimality lead to specific models. Global optimality is not specified in sufficient detail to model specific tasks, and local optimality can only determine the decision rule after the rest of the model is formulated (e.g., as **Moran & Tsetsos** showed, local optimality makes different predictions depending on whether observer models feature late noise or not). At the same time, we agree with **Simen & Balcı** that whatever form the future standard observer model takes, some "utility function will be found at [its] core" (para. 2). In other words, the model would imply that behavior is sensible.

Some commentators were concerned that a focus on building a standard observer model eschews the question of why behavior is the way it is (**Summerfield & Li; Zednik & Jäkel**) and that such a model is too descriptive (**Howes & Lewis; Mastrogiorgio & Petracca; Simen & Balcı**). We are sympathetic to these concerns, but at the same time, we doubt one could satisfactorily answer why behavior is the way it is before understanding in much more detail what people actually do and how they do it. Nevertheless, we think that various components of the standard observer model can and will be inspired by normative (i.e., sensibilist) considerations like the ones Summerfield & Li discussed. In the end, we support any approach that yields predictive, generalizable models of human behavior.

Finally, **Howes & Lewis** criticized observer models for being overly flexible because they are not constrained by a locally optimal decision rule. However, it is important to point out that a locally optimal decision rule is no solution for excessive flexibility either (Bowers & Davis [2012a]). Avoiding excessive flexibility is why the components of any model need to be independently verified, and this is exactly the principle that we advocated for building the standard observer model.

R5.2. What to do about individual differences?

Perception science has traditionally focused on what is common among all observers rather than what is different. However, as several commentators pointed out, there are meaningful individual differences in how people perform perceptual decision-making tasks (**Bachmann; Booth; Love; Mastrogiorgio & Petracca; Schultz & Hurlemann; Withagen et al.; Zhao & Warren**). Such differences create a big additional hurdle for the global optimalist view (Bachmann; Booth; Withagen et al.), which must explain why such differences should be globally optimal in the first place. Individual differences create challenges for observer models, too. Such differences are typically accommodated by free parameters that can vary across individuals (such as the level of internal noise), but it is possible that in some

cases different model components may be needed for different people (e.g., different observers may use qualitatively different decision rules). In our view, individual differences ultimately bring opportunities to develop and test observer models. The predictive validity of a model can be tested not only by generalizing to a new stimulus or task but also to the individual differences that could be expected for the stimulus or task.

R5.3. Finding the best approach for building a standard observer model

One of the most pressing questions for a researcher who wishes to contribute to building a standard observer model is what form this model should take. **Stocker** criticized our vision as "fuzzy," which is not wrong. We certainly do not claim to know what the final model will be. As some commentators thought we were advocating for bag of tricks or neural network models (**Geurts et al.**; Stocker) and others thought we were advocating for LPCD models (**Danks**), we realized we did not make clear that our position is one of agnosticism, not advocacy. We support a diversity of modeling approaches in our field (**Barth et al.**; Danks; **Zednik & Jäkel**), with the common theme that all model components should be tested in multiple ways. There are many types of modeling efforts in the field right now, and this seems to us like a good thing.

Several broad views of how to proceed in studying perceptual decision making were brought up in the commentaries. None of them are specific observer models – they do not generate predictions for behavior – but they provide overarching frameworks and philosophies that guide research and can lead to the development of observer models. We have already discussed optimality as an empirical strategy (**Geurts et al.; Howes & Lewis; Moran & Tsetsos; Shimansky & Dounskaia; Stocker**) and the LPCD framework as potential approaches. Howes & Lewis argued for "computational rationality" (or "bounded optimality"), which is a particular take on optimality as an empirical strategy. **Mastrogiorgio & Petracca** advocated for "bounded rationality," in which decision rules are selected to reach a good enough level of performance, rather than the best possible performance. Such "satisficing" is a clear example of sensibleness without global optimality. (In the target article, we mischaracterized bounded rationality as emphasizing "the optimality of the decision rule"; bounded rationality does not assert that decision rules should be globally or locally optimal.) **Turner et al.** and **Salinas et al.** discussed the option of implementing process models at Marr's algorithmic and implementational levels, which need not be formulated using Bayesian tools. We see value in all of these approaches and perspectives.

R6. Addressing specific suggestions and topics

We devoted most of this response to tackling the larger and more theoretical issues that arose in the commentaries. However, the commentaries were full of additional insights about research methods and specific research topics. Though we only have space to discuss these topics briefly here, we encourage readers to read the relevant commentaries for further information.

R6.1. Modeling suggestions

A few commentators made specific methodological suggestions related to modeling human behavior, which may prove useful in the development of a standard observer model. **Ma** and **Love** both discussed the importance of model comparison, a point with which we wholeheartedly agree. Ma introduced the notion of factorial model comparison, which allows researchers to systematically explore a large number of factors that could jointly affect behavior. **Turner et al.** advocated for the use of Bayesian statistics to compute posterior distributions for each model parameter rather than a single point estimate of the parameter's most likely value. **Shimansky & Dounskaia** proposed that to validate models, a useful strategy is to use a subset of the data to determine unknown parameters and test the resulting model using the rest of the data. Such cross-validation is already common in many fields and should indeed be adopted more often in perceptual decision making. Finally, **Wyart** identified a key issue in evaluating observer models: Given that behavioral data are always noisy, how will we know when a model predicts the data well enough? Wyart proposed a strategy for determining this "good enough" level, also called a "noise ceiling," for an individual observer by measuring the observer's decision consistency across repeated trials.

R6.2. Specific topics

Many commentators gave examples of suboptimality not included in our target article: **Salinas et al.** on reward processing; **Noel; Turner et al.**; and **Zhao & Warren** on cue combination; **Chambers & Kording** on sensorimotor behavior; and Turner et al. on category learning. We appreciated these additional examples of the inability of "standard" optimality considerations to predict behavior. Though we strove to be comprehensive in our survey, there are likely many more such examples in the published literature.

Some commentators argued that a specific effect we cited as locally suboptimal is in fact locally optimal if further factors are taken into account. **Cicchini & Burr** contended that serial dependence is optimal if the previous perceptual information can be used to reduce the uncertainty about the current stimulus. **Summerfield & Li** suggested that many of the surveyed suboptimalities may arise from the principles of efficient coding. We find both of these to be promising hypotheses (efficient coding considerations were discussed in the target article). Finally, **Simen & Balcı** pointed out that our list of suboptimal findings in speed-accuracy tradeoff (SAT) tasks included two articles (Balcı et al. 2011b; Simen et al. 2009) that in fact primarily showed optimal behavior. Their commentary provided a balanced view of findings of optimality and suboptimality in SAT tasks.

Other commentators showed how our target article can be extended to development (**Nardini & Dekker**), computational psychiatry (**Schultz & Hurlemann**), detection of abnormalities in medical images (**Meyer**), and other fields beyond perceptual decision making (**Barth et al.**). Although we focused our target article on perceptual decision making in healthy adults, we strongly support such extensions.

Finally, **Chambers & Kording** criticized our claim that there is a "current narrow focus on optimality" (target article, sect. 1, para. 5). They surveyed articles from the last 23 years and found a similar proportion of optimality and suboptimality claims. However, their data, which they graciously shared with us, suggested a potential difference in the visibility of optimality versus suboptimality claims as measured by the impact factors (IFs) of the journals in which they appeared. Indeed, the numbers of optimality/"near optimality"/suboptimality claims were: 4/0/0

for IF > 10, 7/2/1 for IF > 8, and 9/2/1 for IF > 6. We note that it is difficult to know what the unbiased proportions of optimality and suboptimality claims should be.

R7. Conclusion

Reading the collection of 27 commentaries has been a tremendous source of insight and inspiration. Commentators eloquently described points of view addressed too briefly in the target article and introduced us to new perspectives and empirical findings. Finally, we appreciated the many balanced views expressed. Optimality is a complicated topic, and nuance is paramount. In our view, much disagreement will dissolve if three distinctions are zealously maintained.

First, researchers should not conflate local and global optimality. Every single empirical finding is a finding of local optimality or local suboptimality (i.e., it depends on the assumptions of a particular model). Every given behavior is simultaneously locally optimal and locally suboptimal according to different models. Arguably, "all models are wrong" (Box 1979, p. 202), so local findings of optimality or suboptimality virtually never license statements about global optimality. We will likely never know how globally optimal human behavior actually is.

Second, researchers should recognize the philosophical distinction between sensibleness and global optimality. We find broad agreement that human behavior is sensible but not globally optimal, with perhaps some researchers adopting a global optimalist position. Researchers who identify as global optimalists should define and defend the view that behavior is globally optimal rather than the easier-to-defend and already widely accepted view that behavior is sensible. However, we ultimately consider this distinction to be scientifically unresolvable and therefore likely to remain a matter of opinion.

Third, researchers should clearly separate tools from substantive claims. Both the a priori assumption of optimality as an empirical strategy and the adoption of the Bayesian formalism are tools; they are not in themselves correct or incorrect. However, there are correct and incorrect ways to use them. These tools should be kept separate from substantive claims about global optimality or Bayesian theories.

Maintaining these three distinctions depends on using consistent terminology (see the Glossary). Being more precise with the language we use to talk about optimality and related concepts will help us to identify our common ground – which we believe is more extensive than it might appear on the surface.

We end on the same note as in our target article, with a plea to shift focus away from the optimality of behavior for its own sake and to independently validate all of our hypotheses and assumptions, and all components of our models. We see this practice as the best way to advance our understanding of human behavior and work toward the major goal of developing a standard observer model for perceptual decision making.

References

[The letters "a" and "r" before author's initials stand for target article and response references, respectively]

Abrahamyan A., Luz Silva L., Dakin S. C., Carandini M. & Gardner J. L. (2016) Adaptable history biases in human perceptual decisions. *Proceedings of the National Academy of Sciences of the United States of America* 113(25):E3548–57. Available at: http://www.pnas.org/lookup/doi/10.1073/pnas.1518786113. [aDR]

Abrams J., Barbot A. & Carrasco M. (2010) Voluntary attention increases perceived spatial frequency. *Attention, Perception, & Psychophysics* 72(6):1510–21. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=20675797&retmode=ref&cmd=prlinks. [aDR]

Acerbi L. (2014) Complex internal representations in sensorimotor decision making: A Bayesian investigation. University of Edinburgh. Available at: https://www.era.lib.ed.ac.uk/bitstream/handle/1842/16233/Acerbi2015.pdf?sequence=1&isAllowed=y. [aDR]

Acerbi L., Ma W. J. & Vijayakumar S. (2014a) A framework for testing identifiability of Bayesian models of perception. Paper presented at Advances in Neural Information Processing Systems 27 (NIPS 2014). [WJM]

Acerbi L., Vijayakumar S. & Wolpert D. M. (2014b) On the origins of suboptimality in human probabilistic inference. *PLoS Computational Biology* 10(6):e1003661. Available at: https://doi.org/10.1371/journal.pcbi.1003661. [CC, aDR]

Acerbi L., Wolpert D. M. & Vijayakumar S. (2012) Internal representations of temporal statistics and feedback calibrate motor-sensory interval timing. *PLoS Computational Biology* 8(11):e1002771. Available at: http://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1002771. [aDR]

Ackermann J. F. & Landy M. S. (2015) Suboptimal decision criteria are predicted by subjectively weighted probabilities and rewards. *Attention, Perception & Psychophysics* 77(2):638–58. Available at: http://www.ncbi.nlm.nih.gov/pubmed/25366822. [aDR]

Adams J. K. (1957) A confidence scale defined in terms of expected percentages. *American Journal of Psychology* 70(3):432–36. [aDR]

Adams R. J. & Courage M. L. (2002) Using a single test to measure human contrast sensitivity from early childhood to maturity. *Vision Research* 42(9):1205–10. Available at: https://doi.org/10.1016/S0042-6989(02)00038-X. [MN]

Adams W. J. (2016) The development of audio-visual integration for temporal judgements. *PLOS Computational Biology* 12(4):e1004865. Available at: http://dx.plos.org/10.1371/journal.pcbi.1004865. [aDR]

Adams W. J., Graf E. W. & Ernst M. O. (2004) Experience can change the "light-from-above" prior. *Nature Neuroscience* 7:1057–58. Available at: https://doi.org/10.1038/nn1312. [MN]

Adelson E. H. (1993) Perceptual organization and the judgment of brightness. *Science* 262(5142):2042–44. [aDR]

Adler W. T. & Ma W. J. (2018a) Comparing Bayesian and non-Bayesian accounts of human confidence reports. *PLoS Computational Biology.* https://doi.org/10.1371/journal.pcbi.1006572. [aDR]

Adler W. T. & Ma W. J. (2018b) Limitations of proposed signatures of Bayesian confidence. *Neural Computation* 30(12):3327–54. https://www.mitpressjournals.org/doi/abs/10.1162/neco_a_01141. [aDR]

Ais J., Zylberberg A., Barttfeld P. & Sigman M. (2015) Individual consistency in the accuracy and distribution of confidence judgments. *Cognition* 146:377–86. Available at: http://www.ncbi.nlm.nih.gov/pubmed/26513356. [aDR]

Aitchison L., Bang D., Bahrami B. & Latham P. E. (2015) Doubly Bayesian analysis of confidence in perceptual decision-making. *PLoS Computational Biology* 11(10):e1004519. [aDR]

Akaike H. (1974) A new look at the statistical model identification. *IEEE Transactions on Automatic Control* 19(6):716–23. [WJM]

Alais D. & Burr D. (2004) The ventriloquist effect results from near-optimal bimodal integration. *Current Biology* 14(3):257–62. doi:10.1016/j.cub.2004.01.029. [GMC, J-PN, aDR]

Allen M., Frank D., Schwarzkopf D. S., Fardo F., Winston J. S., Hauser T. U. & Rees G. (2016) Unexpected arousal modulates the influence of sensory noise on confidence. *eLife* 5:e18103. Available at: http://elifesciences.org/lookup/doi/10.7554/eLife.18103. [aDR]

Anderson B. L. (2015) Can computational goals inform theories of vision? *Topics in Cognitive Science* 7:274–86. [HCB]

Anderson B. L., O'Vari J. & Barth H. (2011) Non-Bayesian contour synthesis. *Current Biology* 21(6):492–96. Available at: http://linkinghub.elsevier.com/retrieve/pii/S0960982211001746. [HCB, aDR]

Anderson J. R. (1991) The adaptive nature of human categorization. *Psychological Review* 98:409–29. [BMT]

Anderson M. L. (2014) *After phrenology: Neural reuse and the interactive brain.* MIT Press. [RW]

Anton-Erxleben K., Henrich C. & Treue S. (2007) Attention changes perceived size of moving visual patterns. *Journal of Vision* 7(11):5.1–9. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=17997660&retmode=ref&cmd=prlinks. [aDR]

Anton-Erxleben K., Herrmann K. & Carrasco M. (2013) Independent effects of adaptation and attention on perceived speed. *Psychological Science* 24(2):150–59. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=23241456&retmode=ref&cmd=prlinks. [aDR]

Appelle S. (1972) Perception and discrimination as function of stimulus orientation. *Psychological Bulletin* 78:266–78. [CS]

Arrow K. J. (1958) Utilities, attitudes, choices: A review note. *Econometrica* 26:1–23. [JM]

Aru J. & Bachmann T. (2017) Expectation creates something out of nothing: The role of attention in iconic memory reconsidered. *Consciousness and Cognition* 53:203–10. [TB]

Balcetis E. (2016) Approach and avoidance as organizing structures for motivated distance perception. *Emotion Review* 8(2):115–28. Available at: https://doi.org/10.1177/1754073915586225. [aDR]

Balcı F., Freestone D. & Gallistel C. R. (2009) Risk assessment in man and mouse. *Proceedings of the National Academy of Sciences of the United States of America* 106(7):2459–63. Available at: http://www.pnas.org/content/106/7/2459.full. [PS]

Balcı F., Freestone D., Simen P., deSouza L., Cohen J. D. & Holmes P. (2011a) Optimal temporal risk assessment. *Frontiers in Integrative Neuroscience* 5:56. Available at: https://www.frontiersin.org/articles/10.3389/fnint.2011.00056/full. [PS]

Balcı F., Simen P., Niyogi R., Saxe A., Hughes J. A., Holmes P. & Cohen J. D. (2011b) Acquisition of decision making criteria: Reward rate ultimately beats accuracy. *Attention, Perception & Psychophysics* 73(2):640–57. Retrieved September 11, 2015. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3383845&tool=pmcentrez&rendertype=abstract. [arDR, PS]

Bang J. W. & Rahnev D. (2017) Stimulus expectation alters decision criterion but not sensory signal in perceptual decision making. *Scientific Reports* 7:17072. Available at: http://www.nature.com/articles/s41598-017-16885-2. [aDR]

Bang J. W., Shekhar M. & Rahnev D. (in press) Sensory noise increases metacognitive efficiency. *Journal of Experimental Psychology*. http://dx.doi.org/10.1037/xge0000511. [aDR]

Baranski J. V. & Petrusic W. M. (1994) The calibration and resolution of confidence in perceptual judgments. *Perception & Psychophysics* 55(4):412–28. Available at: http://www.ncbi.nlm.nih.gov/pubmed/8036121. [aDR]

Baranski J. V. & Petrusic W. M. (1995) On the calibration of knowledge and perception. *Canadian Journal of Experimental Psychology* 49(3):397–407. Available at: http://www.ncbi.nlm.nih.gov/pubmed/9183984. [aDR]

Baranski J. V. & Petrusic W. M. (1999) Realism of confidence in sensory discrimination. *Perception & Psychophysics* 61(7):1369–83. Available at: http://www.ncbi.nlm.nih.gov/pubmed/10572465. [aDR]

Barlow H. B. (1961) Possible principles underlying the transformation of sensory messages. In: *Sensory communication*, ed. W. A. Rosenblith, pp. 217–34. MIT Press. [aDR, AAS]

Barlow H. B. (1990) A theory about the functional role and synaptic mechanism of visual after-effects. In: *Vision: Coding and efficiency*, ed. C. Blakemore, pp. 363–75. Cambridge University Press. Available at: http://books.google.com/books?hl=en&lr=&id=xGJ_DxN3eygC&oi=fnd&pg=PA363&dq=a+theory+about+the+functional+role+and+synaptic+mechanism+of+visual+after+effects&ots=VsSUzK0vpB&sig=lZX28LU68XpGk9T8zoLwY8WOJBs. [aDR]

Barth H., Lesser E., Taggart J. & Slusser E. (2015) Spatial estimation: A non-Bayesian alternative. *Developmental Science* 18:853–62. [HCB]

Battaglia P. W., Jacobs R. A. & Aslin R. N. (2003) Bayesian integration of visual and auditory signals for spatial localization. *Journal of the Optical Society of America A, Optics and Image Science* 20(7):1391–97. [aDR]

Battaglia P. W., Kersten D. & Schrater P. R. (2011) How haptic size sensations improve distance perception. *PLoS Computational Biology* 7(6):e1002080. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=21738457&retmode=ref&cmd=prlinks. [aDR]

Bays P. M. & Dowding B. A. (2017) Fidelity of the representation of value in decision-making. *PLoS Computational Biology* 13(3):e1005405. Available at: http://www.ncbi.nlm.nih.gov/pubmed/28248958. [aDR]

Beauchamp M. S. (2005) Statistical criteria in fMRI studies of multisensory integration. *Neuroinformatics* 3(2):93–113. Available at: http://doi.org/10.1385/NI. [J-PN]

Bechtel W. (2009) Looking down, around, and up: Mechanistic explanation in psychology. *Philosophical Psychology* 22:543–64. [CZ]

Bechtel W. & Shagrir O. (2015) The non-redundant contributions of Marr's three levels of analysis for explaining information-processing mechanisms. *Topics in Cognitive Science* 7(2):312–22. [CZ]

Beck J. M., Ma W. J., Pitkow X., Latham P. E. & Pouget A. (2012) Not noisy, just wrong: The role of suboptimal inference in behavioral variability. *Neuron* 74(1):30–39. Available at: https://doi.org/10.1016/j.neuron.2012.03.016. [CC, aDR]

Behrmann M. & Avidan G. (2005) Congenital prosopagnosia: Face-blind from birth. *Trends in Cognitive Sciences* 9(4):180–87. Available at: http://doi.org/10.1016/j.tics.2005.02.011. [JS]

Beierholm U., Shams L., Ma W. J. & Koerding K. P. (2008) Comparing Bayesian models for multisensory cue combination without mandatory integration. Paper presented at the Advances in Neural Information Processing Systems 20 (NIPS 2007). [GMC]

Berger J. O. (1985) *Statistical decision theory and Bayesian analysis*. Springer. [aDR]

Berliner J. E. & Durlach N. I. (1973) Intensity perception. IV. Resolution in roving-level discrimination. *Journal of the Acoustical Society of America* 53(5):1270–87. Available at: http://www.ncbi.nlm.nih.gov/pubmed/4712555. [aDR]

Björkman M., Juslin P. & Winman A. (1993) Realism of confidence in sensory discrimination: The underconfidence phenomenon. *Perception & Psychophysics* 54(1):75–81. Available at: http://www.ncbi.nlm.nih.gov/pubmed/8351190. [aDR]

Blakemore C. & Van Sluyters R. C. (1975) Innate and environmental factors in the development of the kitten's visual cortex. *Journal of Physiology* 248(3):663–716. Available at: https://doi.org/10.1113/jphysiol.1975.sp010995. [MN]

Bogacz R. (2007) Optimal decision-making theories: Linking neurobiology with behaviour. *Trends in Cognitive Sciences* 11(3):118–25. Available at: http://www.sciencedirect.com/science/article/pii/S1364661307000290. [AM, aDR]

Bogacz R., Brown E., Moehlis J., Holmes P. & Cohen J. D. (2006) The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review* 113(4):700–65. Available at: http://doi.apa.org/getdoi.cfm?doi=10.1037/0033-295X.113.4.700. [aDR, PS]

Bogacz R., Hu P. T., Holmes P. J. & Cohen J. D. (2010) Do humans produce the speed-accuracy trade-off that maximizes reward rate? *Quarterly Journal of Experimental Psychology* 63(5):863–91. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2908414&tool=pmcentrez&rendertype=abstract. [AM, aDR]

Bohil C. J. & Maddox W. T. (2001) Category discriminability, base-rate, and payoff effects in perceptual categorization. *Perception & Psychophysics* 63(2):361–76. [aDR]

Bohil C. J. & Maddox W. T. (2003a) On the generality of optimal versus objective classifier feedback effects on decision criterion learning in perceptual categorization. *Memory & Cognition* 31(2):181–98. [aDR]

Bohil C. J. & Maddox W. T. (2003b) A test of the optimal classifier's independence assumption in perceptual categorization. *Perception & Psychophysics* 65(3): 478–93. [aDR]

Booth D. A. (2015) Scientific measurement of sensory preferences using stimulus tetrads. *Journal of Sensory Studies* 30:108–27. doi:10.1111/joss.12143. [DAB]

Booth D. A. (2017) How a mind works. Contrasts with twentieth century psychology. Working paper, University of Sussex. doi:10.13140/RG.2.2.21854.74564. [DAB]

Booth D. A. & Freeman R. P. J. (1993) Discriminative feature integration by individuals. *Acta Psychologica* 84:1–16. [DAB]

Booth D. A., Freeman R. P. J., Konle M., Wainwright C. J. & Sharpe O. (2011a) Perception as interacting psychophysical functions. Could the configuring of features replace a specialised receptor? *Perception* 40:509–29. [DAB]

Booth D. A., Higgs S., Schneider J. & Klinkenberg I. (2010) Learned liking versus inborn delight. Can sweetness give sensual pleasure or is it just motivating? *Psychological Science* 21:1656–63. doi:10.1177/0956797610385356. [DAB]

Booth D. A., Sharpe O., Freeman R. P. J. & Conner M. T. (2011b) Insight into sight, touch, taste and smell by multiple discriminations from norm. *Seeing and Perceiving* 24:485–511. doi:10.1163/187847511X588773. [DAB]

Bossaerts P. & Murawski C. (2017) Computational complexity and human decision-making. *Trends in Cognitive Sciences* 21(12):917–29. Available at: http://dx.doi.org/10.1016/j.tics.2017.09.005. [aDR]

Bowers J. S. & Davis C. J. (2012a) Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin* 138(3):389–414.Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=22545686&retmode=ref&cmd=prlinks. [HCB, AH, arDR, MZ]

Bowers J. S. & Davis C. J. (2012b) Is that what Bayesians believe? Reply to Griffiths, Chater, Norris, and Pouget (2012) *Psychological Bulletin* 138(3):423–26. Available at: http://doi.apa.org/getdoi.cfm?doi=10.1037/a0027750. [aDR]

Box G. E. P. (1979) Robustness in scientific model building. In: *Robustness in statistics*, ed. R. L. Launer & G. N. Wilkinson, pp. 201–36. Academic Press. [rDR]

Brainard D. H., Longère P., Delahunt P. B., Freeman W. T., Kraft J. M. & Xiao B. (2006) Bayesian model of human color constancy. *Journal of Vision* 6(11):1267–81. [aDR]

Brayanov J. B. & Smith M. A. (2010) Bayesian and 'anti-Bayesian' biases in sensory integration for action and perception in the size-weight illusion. *Journal of Neurophysiology* 103 (3):1518–31. Available at: http://jn.physiology.org/cgi/doi/10.1152/jn.00814.2009. [aDR]

Brenner N., Bialek W. & de Ruyter van Steveninck R. (2000) Adaptive rescaling maximizes information transmission. *Neuron* 26(3):695–702. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=10896164&retmode=ref&cmd=prlinks. [aDR]

Briggs R. A. (2017) Normative theories of rational choice: Expected utility. In: *The Stanford encyclopedia of philosophy*. Springer. Available at: https://plato.stanford.edu/archives/spr2017/entries/rationality-normative-utility/. [rDR]

Bronfman Z. Z., Brezis N., Moran R., Tsetsos K., Donner T. & Usher M. (2015) Decisions reduce sensitivity to subsequent information. *Proceedings of the Royal Society B: Biological Sciences* 282(1810):20150228. Available at: http://www.ncbi.nlm.nih.gov/pubmed/26108628. [aDR]

Brooke J. B. & MacRae A. W. (1977) Error patterns in the judgment and production of numerical proportions. *Perception & Psychophysics* 21(4):336–40. Available at: http://www.springerlink.com/index/10.3758/BF03199483. [aDR]

Bülthoff H. H. & Mallot H. A. (1988) Integration of depth modules: Stereo and shading. *Journal of the Optical Society of America A, Optics and Image Science* 5(10):1749–58. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=3204438&retmode=ref&cmd=prlinks. [aDR]

Burr D., Banks M. S. & Morrone M. C. (2009) Auditory dominance over vision in the perception of interval duration. *Experimental Brain Research* 198(1):49–57. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=19597804&retmode=ref&cmd=prlinks. [aDR]

Burr D. & Cicchini G. M. (2014) Vision: efficient adaptive coding. *Current Biology* **24** (22):R1096–98. doi:10.1016/j.cub.2014.10.002. [GMC]

Busemeyer J. R. & Myung I. J. (1992) An adaptive approach to human decision making: Learning theory, decision theory, and human performance. *Journal of Experimental Psychology: General* **121**(2):177–94. Available at: http://psycnet.apa.org/journals/xge/121/2/177.html. [aDR]

Busse L., Ayaz A., Dhruv N. T., Katzner S., Saleem A. B., Schölvinck M. L., Zaharia A. D. & Carandini M. (2011) The detection of visual contrast in the behaving mouse. *Journal of Neuroscience* **31**(31):11351–61. Available at: http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.6689-10.2011. [aDR]

Cangelosi A., Schlesinger M. & Smith L. B. (2015) *Developmental robotics: From babies to robots*. MIT Press. [MN]

Carandini M. & Heeger D. J. (2012) Normalization as a canonical neural computation. *Nature Reviews Neuroscience* **13**(1):51–62. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3273486&tool=pmcentrez&rendertype=abstract. [aDR]

Carrasco M. (2011) Visual attention: The past 25 years. *Vision Research* **51**(13):1484–525. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3390154&tool=pmcentrez&rendertype=abstract. [aDR]

Carrasco M., Ling S. & Read S. (2004) Attention alters appearance. *Nature Neuroscience* **7**(3):308–13. Available at: http://www.ncbi.nlm.nih.gov/pubmed/14966522. [aDR]

Çavdaroglu B., Zeki M. & Balcı F. (2014) Time-based reward maximization. *Philosophical Transactions of the Royal Society B: Biological Sciences* **369**(1637):20120461. Available at: http://rstb.royalsocietypublishing.org/content/369/1637/20120461. [PS]

Chambers C., Fernandes H. & Kording K. (2017a) Policies or knowledge: Priors differ between perceptual and sensorimotor tasks. *bioRxiv* 132829. Available at: https://doi.org/10.1101/132829. [CC]

Chambers C., Sokhey T., Gaebler-Spira D. & Kording K. (2017b) The development of Bayesian integration in sensorimotor estimation. *bioRxiv* 136267. Available at: https://doi.org/10.1101/136267. [CC]

Charles L., Van Opstal F., Marti S. & Dehaene S. (2013) Distinct brain mechanisms for conscious versus subliminal error detection. *NeuroImage* **73**:80–94. Available at: http://www.ncbi.nlm.nih.gov/pubmed/23380166. [aDR]

Cheadle S., Wyart V., Tsetsos K., Myers N., de Gardelle V., Castañón S. H. & Summerfield C. (2014) Adaptive gain control during human perceptual choice. *Neuron* **81**(6):1429–41. Available at: http://www.cell.com/article/S0896627314000518/fulltext. [aDR]

Chemero A. (2009) *Radical embodied cognitive science*. MIT Press. [RW]

Chen C.-C. & Tyler C. W. (2015) Shading beats binocular disparity in depth from luminance gradients: Evidence against a maximum likelihood principle for cue combination. *PLoS ONE* **10**(8):e0132658. Available at: http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0132658. [aDR]

Chen X. L., McNamara T. P., Kelly J. W. & Wolbers T. (2017) Cue combination in human spatial navigation. *Cognitive Psychology* **95**:105–44. [MZ]

Cheng K., Shettleworth S. J., Huttenlocher J. & Rieser J. J. (2007) Bayesian integration of spatial information. *Psychological Bulletin* **133**(4):625–37. [MZ]

Chiang T.-C., Lu R.-B., Hsieh S., Chang Y.-H. & Yang Y.-K. (2014) Stimulation in the dorsolateral prefrontal cortex changes subjective evaluation of percepts. *PLoS ONE* **9**(9):e106943. Available at: http://dx.plos.org/10.1371/journal.pone.0106943. [aDR]

Cicchini G. M., Anobile G. & Burr D. C. (2014) Compressive mapping of number to space reflects dynamic encoding mechanisms, not static logarithmic transform. *Proceedings of the National Academy of Sciences of the United States of America* **111**(21):7867–72. doi:10.1073/pnas.1402785111. [GMC]

Cicchini G. M., Arrighi R., Cecchetti L., Giusti M. & Burr D. C. (2012) Optimal encoding of interval timing in expert percussionists. *Journal of Neuroscience* **32**(3):1056–60. doi:10.1523/JNEUROSCI.3411-11.2012. [HCB, GMC]

Cicchini G. M. & Kristjánsson, Á. (2015) Guest editorial: On the possibility of a unifying framework for serial dependencies. *i-Perception* **6**(6). doi:10.1177/2041669515614148. [GMC]

Cicchini G. M., Mikellidou K. & Burr D. (2017) Serial dependencies act directly on perception. *Journal of Vision* **17**(14):6. doi:10.1167/17.14.6. [GMC]

Cicchini G. M., Mikellidou K. & Burr D. C. (2018) The functional role of serial dependencies. *Proceedings of the Royal Society B: Biological Sciences* **285**:20181722. doi:10.1098/rspb.2018.1722. [GMC]

Clark J. J. & Yullie A. L. (1990) *Data fusion for sensory information processing*. Kluwer Academic. [aDR]

Collins A. G. E., Albrecht M. A., Waltz J. A., Gold J. M. & Frank M. J. (2017) Interactions among working memory, reinforcement learning, and effort in value-based choice: A new paradigm and selective deficits in schizophrenia. *Biological Psychiatry* **82**(6):431–39. Available at: http://doi.org/10.1016/j.biopsych.2017.05.017. [JS]

Colombo M. & Hartmann S. (2017) Bayesian cognitive science, unification, and explanation. *British Journal for the Philosophy of Science* **68**(2):451–84. [CZ]

Conner M. T., Haddon A. V., Pickering E. S. & Booth D. A. (1988) Sweet tooth demonstrated: Individual differences in preference for both sweet foods and foods highly sweetened. *Journal of Applied Psychology* **73**:275–80. [DAB]

Conner M. T., Land D. G. & Booth D. A. (1987) Effects of stimulus range on judgments of sweetness intensity in a lime drink. *British Journal of Psychology* **78**:357–64. [DAB]

Cooper G. F. (1990) The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence* **42**(2–3):393–405. [aDR]

Cowan N. (2005) *Working memory capacity*. Psychology Press. [aDR]

Craver C. F. (2013) Functions and mechanisms: A perspectivalist view. In: *Functions: Selection and mechanisms*, ed. P. Huneman, pp. 133–58. Springer. [CZ]

Crawford L. E. & Duffy S. (2010) Sequence effects in estimating spatial location. *Psychonomic Bulletin & Review* **17**:725–30. [HCB]

Creelman C. D. & Macmillan N. A. (1979) Auditory phase and frequency discrimination: A comparison of nine procedures. *Journal of Experimental Psychology: Human Perception and Performance* **5**(1):146–56. Available at: http://www.ncbi.nlm.nih.gov/pubmed/528924. [aDR]

Dakin S. C. (2001) Information limit on the spatial integration of local orientation signals. *Journal of the Optical Society of America A, Optics and Image Science* **18**(5):1016–26. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=11336204&retmode=ref&cmd=prlinks. [aDR]

Danks D. (2008) Rational analyses, instrumentalism, and implementations. In: *The probabilistic mind: Prospects for rational models of cognition*, ed. N. Chater & M. Oaksford, pp. 59–75. Oxford University Press. [CZ]

Darwin C. (1859/1985) *On the origin of species*. Penguin. [RW]

Davis-Stober C. P., Park S., Brown N. & Regenwetter M. (2016) Reported violations of rationality may be aggregation artifacts. *Proceedings of the National Academy of Sciences of the United States of America* **113**(33):E4761–63. Available at: http://www.ncbi.nlm.nih.gov/pubmed/27462103. [aDR]

Dawes R. M. (1980) Confidence in intellectual vs. confidence in perceptual judgments. In: *Similarity and choice: Papers in honor of Clyde Coombs*, ed. E. D. Lantermann & H. Feger, pp. 327–45. Han Huber. [aDR]

Dawkins R. (1982) *The extended phenotype: The gene as the unit of selection*. Freeman. [RW]

Dayan P. (2014) Rationalizable irrationalities of choice. *Topics in Cognitive Science* **6**(2):204–28. Available at: http://doi.wiley.com/10.1111/tops.12082. [aDR]

de Gardelle V., Kouider S. & Sackur J. (2010) An oblique illusion modulated by visibility: Non-monotonic sensory integration in orientation processing. *Journal of Vision* **10**(10):6. [AAS]

de Gardelle V. & Mamassian P. (2015) Weighting mean and variability during confidence judgments. *PLoS ONE* **10**(3):e0120870. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=4368758&tool=pmcentrez&rendertype=abstract. [aDR]

de Gardelle V. & Summerfield C. (2011) Robust averaging during perceptual judgment. *Proceedings of the National Academy of Sciences of the United States of America* **108**(32):13341–46. doi:10.1073/pnas.1104517108. [aDR, CS]

Dekker T. M., Ban H., van der Velde B., Sereno M. I., Welchman A. E. & Nardini M. (2015) Late development of cue integration is linked to sensory fusion in cortex. *Current Biology* **25**(21): 2856–61. Available at: https://doi.org/10.1016/j.cub.2015.09.043. [MN, aDR]

Dekker T. M. & Nardini M. (2016) Risky visuomotor choices during rapid reaching in childhood. *Developmental Science* **19**(3):427–39. Available at: https://doi.org/10.1111/desc.12322. [MN]

de Lange F. P., Rahnev D., Donner T. H. & Lau H. (2013) Prestimulus oscillatory activity over motor cortex reflects perceptual expectations. *Journal of Neuroscience* **33**(4):1400–10. Available at: http://www.ncbi.nlm.nih.gov/pubmed/23345216. [aDR]

Del Cul A., Dehaene S., Reyes P., Bravo E. & Slachevsky A. (2009) Causal role of prefrontal cortex in the threshold for access to consciousness. *Brain* **132**(Pt. 9):2531–40. Available at: http://www.ncbi.nlm.nih.gov/pubmed/19433438. [aDR]

Denison R. N., Adler W. T., Carrasco M. & Ma W. J. (2018) Humans incorporate attention-dependent uncertainty into perceptual decisions and confidence. *Proceedings of the National Academy of Sciences of the United States of America* **115**(43):11090–95. doi: 10.1073/pnas.1717720115. [aDR]

de Wit, M M., van der Kamp J. & Withagen R. (2015) Visual illusions and direct perception: Elaborating on Gibson's insights. *New Ideas in Psychology* **36**:1–9. [RW]

Diedrichsen J., Shadmehr R. & Ivry R. B. (2010) The coordination of movement: Optimal feedback control and beyond. *Trends in Cognitive Sciences* **14**:31–39. [YPS]

Ding L. & Hikosaka O. (2006) Comparison of reward modulation in the frontal eye field and caudate of the macaque. *Journal of Neuroscience* **26**:6695–703. [ES]

Dobs K., Bülthoff I., Breidt M., Vuong Q. C., Curio C. & Schultz J. (2014) Quantifying human sensitivity to spatio-temporal information in dynamic faces. *Vision Research* **100**:78–87. Available at: http://doi.org/10.1016/j.visres.2014.04.009. [JS]

Dobs K., Bülthoff I. & Schultz J. (2016) Identity information content depends on the type of facial movement. *Scientific Reports* **6**:34301. Available at: http://doi.org/10.1038/srep34301. [JS]

Dobs K., Ma W. J. & Reddy L. (2017) Near-optimal integration of facial form and motion. *Scientific Reports* **7**:11002. Available at: http://doi.org/10.1038/s41598-017-10885-y. [JS]

Dounskaia N. & Shimansky Y. (2016) Strategy of arm movement control is determined by minimization of neural effort for joint coordination. *Experimental Brain Research* **234**:1335–50. [YPS]

Dreyfus H. & Taylor C. (2015) *Retrieving realism.* Harvard University Press. [RW]

Drugowitsch J., DeAngelis G. C., Angelaki D. E. & Pouget A. (2015) Tuning the speed-accuracy trade-off to maximize reward rate in multisensory decision-making. *eLife* 4: e06678. Available at: http://www.ncbi.nlm.nih.gov/pubmed/26090907. [aDR]

Drugowitsch J., DeAngelis G. C., Klier E. M., Angelaki D. E. & Pouget A. (2014a) Optimal multisensory decision-making in a reaction-time task. *eLife* 3:e03005. Available at: http://elifesciences.org/content/early/2014/06/14/eLife.03005.abstract. [aDR]

Drugowitsch J., Moreno-Bote R., Churchland A. K., Shadlen M. N. & Pouget A. (2012) The cost of accumulating evidence in perceptual decision making. *Journal of Neuroscience* 32(11):3612–28. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3329788&tool=pmcentrez&rendertype=abstract. [aDR]

Drugowitsch J., Moreno-Bote R. & Pouget A. (2014b) Relation between belief and performance in perceptual decision making. *PLoS ONE* 9(5):e96511. Available at: http://dx.plos.org/10.1371/journal.pone.0096511. [aDR]

Drugowitsch J. & Pouget A. (2012) Probabilistic vs. non-probabilistic approaches to the neurobiology of perceptual decision-making. *Current Opinion in Neurobiology* 22(6):963–69. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3513621&tool=pmcentrez&rendertype=abstract. [aDR]

Drugowitsch J., Wyart V., Devauchelle A.-D. & Koechlin E. (2016) Computational precision of mental inference as critical source of human choice suboptimality. *Neuron* 92 (6):1398–411. Available at: http://dx.doi.org/10.1016/j.neuron.2016.11.005. [RM, aDR, CS, VW]

Duffy S., Huttenlocher J. & Crawford L. E. (2006) Children use categories to maximize estimation. *Developmental Science* 9:597–603. [HCB]

Duffy S. & Smith J. (2017) Category effects on stimulus estimation: Shifting and skewed frequency distributions – A reexamination. *Psychonomic Bulletin & Review* 25 (5):1740–50. Available at: https://doi.org/10.3758/s13423-017-1392-7. [HCB]

Eberhardt F. & Danks D. (2011) Confirmation in the cognitive sciences: The problematic case of Bayesian models. *Minds and Machines* 21(3):389–410. Available at: http://link.springer.com/10.1007/s11023-011-9241-3. [aDR]

Eckstein M. P. (2011) Visual search: A retrospective. *Journal of Vision* 11(5):14. Available at: http://www.journalofvision.org/content/11/5/14.abstract. [aDR]

Endler J. A. (1986) *Natural selection in the wild.* Princeton University Press. [RW]

Ernst M. O. & Banks M. S. (2002) Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415(6870):429–33. Available at: http://dx.doi.org/10.1038/415429a. [GMC, J-PN, aDR, AAS]

Esins J., Schultz J., Stemper C., Kennerknecht I. & Bülthoff I. (2016) Face perception and test reliabilities in congenital prosopagnosia in seven tests. *i-Perception* 7(1). Available at: http://doi.org/10.1177/2041669515625797. [JS]

Esins J., Schultz J., Wallraven C. & Bülthoff I. (2014) Do congenital prosopagnosia and the other-race effect affect the same face recognition mechanisms? *Frontiers in Human Neuroscience* 8:759. Available at: https://doi.org/10.3389/fnhum.2014.00759. [JS]

Etienne A. S. & Jeffery K. J. (2004) Path integration in mammals. *Hippocampus* 14:180–92. [MZ]

Evans K. K., Birdwell R. L. & Wolfe J. M. (2013) If you don't find it often, you often don't find it: Why some cancers are missed in breast cancer screening. *PLoS ONE* 8(5): e64366. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3667799&tool=pmcentrez&rendertype=abstract. [aDR]

Evans K. K., Tambouret R. H., Evered A., Wilbur D. C. & Wolfe J. M. (2011) Prevalence of abnormalities influences cytologists' error rates in screening for cervical cancer. *Archives of Pathology & Laboratory Medicine* 135(12):1557–60. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3966132&tool=pmcentrez&rendertype=abstract. [aDR]

Evans N. J. & Brown S. D. (2017) People adopt optimal policies in simple decision-making, after practice and guidance. *Psychonomic Bulletin & Review* 24:597–606. Available at: https://link.springer.com/article/10.3758%2Fs13423-016-1135-1. [PS]

Fawcett T. W., Fallenstein B., Higginson A. D., Houston A. I., Mallpress D. E., Trimmer P. C. & McNamara J. M. (2014) The evolution of decision rules in complex environments. *Trends in Cognitive Sciences* 18(3):153–61. [RM]

Fechner G. T. (1860) *Elemente der psychophysik.* Breitkopf und Härtel. [aDR]

Fechner G. T. (1882) Revision der Hauptpunkte der Psychophysik: Sektion 21 (pp. 290–99). Breitkopf und Hertel. Reprinted and translated in *Psychological Research* 49(1987):209–12. [TB]

Feng S., Holmes P., Rorie A. & Newsome W. T. (2009) Can monkeys choose optimally when faced with noisy stimuli and unequal rewards? *PLoS Computational Biology* 5 (2):e1000284. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2631644&tool=pmcentrez&rendertype=abstract. [aDR, ES]

Fetsch C. R., Pouget A., Deangelis G. C. & Angelaki D. E. (2012) Neural correlates of reliability-based cue weighting during multisensory integration. *Nature Neuroscience* 15(1):146–54. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=22101645&retmode=ref&cmd=prlinks. [aDR]

Fetsch C. R., Turner A. H., DeAngelis G. C. & Angelaki D. E. (2009) Dynamic reweighting of visual and vestibular cues during self-motion perception. *Journal of Neuroscience* 29:15601–12. [J-PN]

Firestone C. & Scholl B. J. (2016) Cognition does not affect perception: Evaluating the evidence for "top-down" effects. *Behavioral and Brain Sciences* 39:e229. Available at: http://www.ncbi.nlm.nih.gov/pubmed/26189677. [aDR]

Fischer J. & Whitney D. (2014) Serial dependence in visual perception. *Nature Neuroscience* 17(5):738–43. Available at: http://dx.doi.org/10.1038/nn.3689. [aDR, GMC]

Fiser J., Berkes P., Orbán G. & Lengyel M. (2010) Statistically optimal perception and learning: From behavior to neural representations. *Trends in Cognitive Sciences* 14 (3):119–30. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2939867&tool=pmcentrez&rendertype=abstract. [aDR]

Fisher R. A. (1926) The arrangement of field experiments. *Journal of the Ministry of Agriculture of Great Britain* 33:503–13. [WJM]

Fitts P. M. (1966) Cognitive aspects of information processing: III. Set for speed versus accuracy. *Journal of Experimental Psychology* 71(6):849–57. [aDR]

Fleming R. W. (2011) Visual perception: Bizarre contours go against the odds. *Current Biology* 21:R259–61. [HCB]

Fleming S. M. & Daw N. D. (2017) Self-evaluation of decision-making: A general Bayesian framework for metacognitive computation. *Psychological Review* 124(1):91–114. http://dx.doi.org/10.1037/rev0000045. [aDR]

Fleming S. M. & Lau H. (2014) How to measure metacognition. *Frontiers in Human Neuroscience* 8:443. Available at: http://journal.frontiersin.org/Journal/10.3389/fnhum.2014.00443/abstract. [aDR]

Fleming S. M., Maloney L. T. & Daw N. D. (2013) The irrationality of categorical perception. *Journal of Neuroscience* 33(49):19060–70. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=24305804&retmode=ref&cmd=prlinks. [aDR]

Fleming S. M., Maniscalco B. & Ko Y. (2015) Action-specific disruption of perceptual confidence. *Psychological Science* 26(1):89–98. Available at: http://www.ncbi.nlm.nih.gov/pubmed/25425059. [aDR]

Fleming S. M., Massoni S., Gajdos T. & Vergnaud J.-C. (2016) Metacognition about the past and future: Quantifying common and distinct influences on prospective and retrospective judgments of self-performance. *Neuroscience of Consciousness* 2016(1):niw018. Available at: https://academic.oup.com/nc/article-lookup/doi/10.1093/nc/niw018. [aDR]

Fleming S. M., Ryu J., Golfinos J. G. & Blackmon K. E. (2014) Domain-specific impairment in metacognitive accuracy following anterior prefrontal lesions. *Brain* 137(10):2811–22. Available at: http://brain.oxfordjournals.org/content/early/2014/08/06/brain.awu221.long. [aDR]

Foo P., Warren W. H., Duchon A. & Tarr M. J. (2005) Do humans integrate routes into a cognitive map? Map- versus landmark-based navigation of novel shortcuts. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 31:195–215. [MZ]

Forstmann B. U., Ratcliff R. & Wagenmakers E.-J. (2016) Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions. *Annual Review of Psychology* 67:641–66. Available at: http://www.annualreviews.org/eprint/2stAyEdsCkSk9MpsHMDV/full/10.1146/annurev-psych-122414-033645. [aDR]

Frens M. A. & Van Opstal A. J. (1998) Visual-auditory interactions modulate saccade-related activity in monkey superior colliculus. *Brain Research Bulletin* 46:211–24. [J-PN]

Fritsche M., Mostert P. & de Lange F. P. (2017) Opposite effects of recent history on perception and decision. *Current Biology* 27(4):590–95. Available at: http://dx.doi.org/10.1016/j.cub.2017.01.006. [aDR]

Frund I., Wichmann F. A. & Macke J. H. (2014) Quantifying the effect of intertrial dependence on perceptual decisions. *Journal of Vision* 14(7):9. doi:10.1167/14.7.9. [GMC, aDR]

Fuller S., Park Y. & Carrasco M. (2009) Cue contrast modulates the effects of exogenous attention on appearance. *Vision Research* 49(14):1825–37. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=19393260&retmode=ref&cmd=prlinks. [aDR]

Gallistel C. R. (2005) Deconstructing the law of effect. *Games and Economic Behavior* 52 (2):410–23. [JM]

Ganmor E., Landy M. S. & Simoncelli E. P. (2015) Near-optimal integration of orientation information across saccades. *Journal of Vision* 15(16):8. Available at: http://jov.arvojournals.org/article.aspx?doi=10.1167/15.16.8. [aDR]

Garcia S. E., Jones P. R., Reeve E. I., Michaelides M., Rubin G. S. & Nardini M. (2017) Multisensory cue combination after sensory loss: Audio-visual localization in patients with progressive retinal disease. *Journal of Experimental Psychology: Human Perception and Performance* 43(4):729–40. Available at: http://doi.apa.org/getdoi.cfm?doi=10.1037/xhp0000344. [aDR]

García-Pérez M. A. & Alcalá-Quintana R. (2010) The difference model with guessing explains interval bias in two-alternative forced-choice detection procedures. *Journal of Sensory Studies* 25(6):876–98. Available at: http://doi.wiley.com/10.1111/j.1745-459X.2010.00310.x. [aDR]

García-Pérez M. A. & Alcalá-Quintana R. (2011) Interval bias in 2AFC detection tasks: Sorting out the artifacts. *Attention, Perception, & Psychophysics* 73(7):2332–52. Available at: http://www.springerlink.com/index/10.3758/s13414-011-0167-x. [aDR]

Geisler W. S. (2011) Contributions of ideal observer theory to vision research. *Vision Research* 51(7):771–81. [arDR]

Geisler W. S. & Najemnik J. (2013) Optimal and non-optimal fixation selection in visual search. *Perception ECVP Abstract* 42:226. Available at: http://www.perceptionweb.com/abstract.cgi?id=v130805. [aDR]

Gekas N., Chalk M., Seitz A. R. & Series P. (2013) Complexity and specificity of experimentally-induced expectations in motion perception. *Journal of Vision* 13 (4):8. Available at: http://jov.arvojournals.org/article.aspx?articleid=2121832. [aDR]

Gepshtein S., Burge J., Ernst M. O. & Banks M. S. (2005) The combination of vision and touch depends on spatial proximity. *Journal of Vision* 5(11):1013–23. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=16441199&retmode=ref&cmd=prlinks. [aDR]

Gershman S. J., Horvitz E. J. & Tenenbaum J. B. (2015) Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science* 349 (6245):273–78. Available at: http://www.ncbi.nlm.nih.gov/pubmed/26185246. [aDR]

Gibson J. J. (1966) *The senses considered as perceptual systems*. Houghton Mifflin. [RW]

Gibson J. J. (1979/1986) *The ecological approach to visual perception*. Houghton Mifflin. [RW]

Gibson J. J. & Radner M. (1937) Adaptation, after-effect and contrast in the perception of tilted lines. I. Quantitative studies. *Journal of Experimental Psychology* 20(5):453–67. Available at: http://doi.apa.org/getdoi.cfm?doi=10.1037/h0059826. [arDR]

Gigerenzer G. (1991) From tools to theories: A heuristic of discovery in cognitive psychology. *Psychological Review* 98(2):254–67. [CZ]

Gigerenzer G. (2004) Striking a blow for sanity in theories of rationality. In: *Models of a man: Essays in the memory of Herbert A. Simon*, ed. M. Augier & J. G. March, pp. 389–409. MIT Press. [AM]

Gigerenzer G. & Brighton H. (2009) Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science* 1(1):107–43. [aDR]

Gigerenzer G. & Gaissmaier W. (2011) Heuristic decision making. *Annual Review of Psychology* 62:451–82. Available at: https://doi.org/10.1146/annurev-psych-120709-145346. [CC]

Gigerenzer G., Hoffrage U. & Kleinbölting H. (1991) Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review* 98(4):506–28. Available at: http://www.ncbi.nlm.nih.gov/pubmed/1961771. [aDR]

Gigerenzer G. & Selten R., eds. (2001) *Bounded rationality: The adaptive toolbox*. MIT Press. [AM]

Gigerenzer G. & Selten R. (2002) *Bounded rationality*. MIT Press. [aDR]

Gigerenzer G., Todd P. M. & the ABC Research Group. (1999) *Simple heuristics that make us smart*. Oxford University Press. [AM]

Girshick A. R., Landy M. S. & Simoncelli E. P. (2011) Cardinal rules: Visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience* 14 (7):926–32. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3125404&tool=pmcentrez&rendertype=abstract. [LSG, aDR, CS, VW]

Glennerster A., Tcheang L., Gilson S. J., Fitzgibbon A. W. & Parker A. J. (2006) Humans ignore motion and stereo cues in favor of a fictional stable world. *Current Biology* 16 (4):428–32. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=16488879&retmode=ref&cmd=prlinks. [aDR]

Gobell J. & Carrasco M. (2005) Attention alters the appearance of spatial frequency and gap size. *Psychological Science* 16(8):644–51. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=16102068&retmode=ref&cmd=prlinks. [aDR]

Goble J., Zhang Y., Shimansky Y., Sharma S. & Dounskaia N. (2007) Directional biases reveal utilization of arm's biomechanical properties for optimization of motor behavior. *Journal of Neurophysiology* 98:1240–52. [YPS]

Gold J. M., Murray R. F., Bennett P. J. & Sekuler A. B. (2000) Deriving behavioural receptive fields for visually completed contours. *Current Biology* 10(11):663–66. Available at: http://www.sciencedirect.com/science/article/pii/S0960982200005236. [aDR]

Goodman N. D., Frank M. C. & Griffiths T. L., Tenenbaum J. B., Battaglia P. W. & Hamrick J. B. (2015) Relevant and robust: A response to Marcus and Davis (2013) *Psychological Science* 26(4):539–41. Available at: http://pss.sagepub.com/lookup/doi/10.1177/0956797614559544. [aDR]

Gorea A., Caetta F. & Sagi D. (2005) Criteria interactions across visual attributes. *Vision Research* 45(19):2523–32. Available at: http://www.ncbi.nlm.nih.gov/pubmed/15950255. [aDR]

Gorea A. & Sagi D. (2000) Failure to handle more than one internal representation in visual detection tasks. *Proceedings of the National Academy of Sciences of the United States of America* 97(22):12380–84. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=17350&tool=pmcentrez&rendertype=abstract. [aDR]

Gorea A. & Sagi D. (2001) Disentangling signal from noise in visual contrast discrimination. *Nature Neuroscience* 4(11):1146–50. Available at: http://www.ncbi.nlm.nih.gov/pubmed/11687818. [aDR]

Gorea A. & Sagi D. (2002) Natural extinction: A criterion shift phenomenon. *Visual Cognition* 9(8):913–36. [aDR]

Gori M., Del Viva M., Sandini G. & Burr D. C. (2008) Young children do not integrate visual and haptic form information. *Current Biology* 18(9):694–98. Available at: https://doi.org/10.1016/j.cub.2008.04.036. [MN, aDR]

Gould S. J. & Lewontin R. (1978) The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist program. *Proceedings of the Royal Society B: Biological Sciences* 205:581–98. [RW]

Green D. M. & Swets J. A. (1966) *Signal detection theory and psychophysics*. John Wiley & Sons. [aDR, VW]

Griffin D. & Tversky A. (1992) The weighing of evidence and the determinants of confidence. *Cognitive Psychology* 24(3):411–35. Available at: http://www.sciencedirect.com/science/article/pii/001002859290013R. [aDR]

Griffiths T. L., Chater N., Norris D. & Pouget A. (2012) How the Bayesians got their beliefs (and what those beliefs actually are): Comment on Bowers and Davis (2012) *Psychological Bulletin* 138(3):415–22. Available at: https://doi.org/10.1037/a0026884. [CC, aDR]

Griffiths T. L., Lieder F. & Goodman N. D. (2015) Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science* 7 (2):217–29. Available at: http://doi.wiley.com/10.1111/tops.12142. [aDR]

Grzeczkowski L., Clarke A. M., Francis G., Mast F. W. & Herzog M. H. (2017) About individual differences in vision. *Vision Research* 141:282–92. Available at: http://doi.org/10.1016/j.visres.2016.10.006. [JS]

Grzywacz N. M. & Balboa R. M. (2002) A Bayesian framework for sensory adaptation. *Neural Computation* 14(3):543–59. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=11860682&retmode=ref&cmd=prlinks. [aDR]

Gu Y., Angelaki D. E. & DeAngelis G. C. (2008) Neural correlates of multisensory cue integration in macaque MSTd. *Nature Neuroscience* 11(10):1201–10. Available at: http://www.nature.com/doifinder/10.1038/nn.2191. [aDR]

Guastella A. J., Mitchell P. B. & Dadds M. R. (2008) Oxytocin increases gaze to the eye region of human faces. *Biological Psychiatry* 63(1):3–5. Available at: http://doi.org/10.1016/j.biopsych.2007.06.026. [JS]

Hadad B. S., Maurer D. & Lewis T. L. (2011) Long trajectory for the development of sensitivity to global and biological motion. *Developmental Science* 14:1330–39. Available at: https://doi.org/10.1111/j.1467-7687.2011.01078.x. [MN]

Hammett S. T., Champion R. A., Thompson P. G. & Morland A. B. (2007) Perceptual distortions of speed at low luminance: Evidence inconsistent with a Bayesian account of speed encoding. *Vision Research* 47(4):564–68. Available at: http://www.ncbi.nlm.nih.gov/pubmed/17011014. [aDR]

Hands D. W. (2014) Normative ecological rationality: Normative rationality in the fast-and-frugal-heuristics research program. *Journal of Economic Methodology* 21 (4):396–410. [AM]

Hanks T. D., Mazurek M. E., Kiani R., Hopp E. & Shadlen M. N. (2011) Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *Journal of Neuroscience* 31(17):6339–52. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3356114&tool=pmcentrez&rendertype=abstract. [aDR]

Hanks T. D. & Summerfield C. (2017) Perceptual decision making in rodents, monkeys, and humans. *Neuron* 93(1):15–31. Available at: http://dx.doi.org/10.1016/j.neuron.2016.12.003. [aDR]

Harvey N. (1997) Confidence in judgment. *Trends in Cognitive Sciences* 1(2):78–82. Available at: http://www.ncbi.nlm.nih.gov/pubmed/21223868. [aDR]

Hassan O. & Hammett S. T. (2015) Perceptual biases are inconsistent with Bayesian encoding of speed in the human visual system. *Journal of Vision* 15(2):9. Available at: http://jov.arvojournals.org/article.aspx?articleid=2213273. [aDR]

Hauser C. K., Zhu D., Stanford T. R. & Salinas E. (2018) Motor selection dynamics in FEF explain the reaction time variance of saccades to single targets. *eLife* 7:e33456. [ES]

Hawkins G. E., Forstmann B. U., Wagenmakers E.-J., Ratcliff R. & Brown S. D. (2015) Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *Journal of Neuroscience* 35(6):2476–84. Available at: http://www.jneurosci.org/content/35/6/2476.full. [aDR]

Healy A. F. & Kubovy M. (1981) Probability matching and the formation of conservative decision rules in a numerical analog of signal detection. *Journal of Experimental Psychology: Human Learning and Memory* 7(5):344–54. [aDR]

Heft H. (2007) The social constitution of perceiver–environment reciprocity. *Ecological Psychology* 19:85–105. [RW]

Heitz R. P. (2014) The speed-accuracy tradeoff: History, physiology, methodology, and behavior. *Frontiers in Neuroscience* 8:150. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=4052662&tool=pmcentrez&rendertype=abstract. [aDR]

Helmholtz H. (1867) *Handbuch der Physiologischen Optik*. Allg. Enzyklopadie der Physik, 9 Bd. Voss. [AAS]

Helmholtz H. L. F. (1856) *Treatise on physiological optics*. Thoemmes Continuum. [aDR]

Henriques J. B., Glowacki J. M. & Davidson R. J. (1994) Reward fails to alter response bias in depression. *Journal of Abnormal Psychology* 103(3):460–66. Available at: http://www.ncbi.nlm.nih.gov/pubmed/7930045. [aDR]

Hikosaka O., Nakamura K. & Nakahara H. (2006) Basal ganglia orient eyes to reward. *Journal of Neurophysiology* 95:567–84. [ES]

Hillis J. M., Ernst M. O., Banks M. S. & Landy M. S. (2002) Combining sensory information: Mandatory fusion within, but not between, senses. *Science* 298(5598):1627–30. Available at: http://www.sciencemag.org/cgi/doi/10.1126/science.1075396. [aDR]

Hohwy J., Roepstorff A. & Friston K. (2008) Predictive coding explains binocular rivalry: An epistemological review. *Cognition* **108**(3):687–701. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=18649876&retmode=ref&cmd=prlinks. [aDR]

Holmes P. & Cohen J. D. (2014) Optimality and some of its discontents: Successes and shortcomings of existing models for binary decisions. *Topics in Cognitive Science* **6**(2):258–78. Available at: http://www.ncbi.nlm.nih.gov/pubmed/24648411. [aDR]

Hotaling J. M., Cohen A. L., Shiffrin R. M. & Busemeyer J. R. (2015) The dilution effect and information integration in perceptual decision making. *PLoS ONE* **10**(9): e0138481. Available at: https://doi.org/10.1371/journal.pone.0138481. [BMT]

Howes A., Lewis R. L. & Vera A. (2009) Rational adaptation under task and processing constraints: implications for testing theories of cognition and action. *Psychological Review* **116**(4):717–51. [AH]

Howes A., Warren P., Farmer G., El-Deredy W. & Lewis R. L. (2016) Why contextual preference reversals maximize expected value. *Psychological Review* **123**(4):368–91. [AH]

Hubel D. H. & Wiesel T. N. (1970) The period of susceptibility to the physiological effects of unilateral eye closure in kittens. *Journal of Physiology* **206**(2):419–36. Available at: http://www.ncbi.nlm.nih.gov/pubmed/5498493. [rDR]

Hurlemann R. (2017) Oxytocin-augmented psychotherapy: Beware of context. *Neuropsychopharmacology* **42**(1):377. Available at: http://doi.org/10.1038/npp.2016.188. [JS]

Hurlemann R., Scheele D., Maier W. & Schultz J. (2017) Oxytocin drives prosocial biases in favor of attractive people. *Behavioral and Brain Sciences* **40**:e30. Available at: http://doi.org/10.1017/S0140525X16000510. [JS]

Huttenlocher J., Hedges L. V. & Vevea J. L. (2000) Why do categories affect stimulus judgment? *Journal of Experimental Psychology: General* **129**:220–41. [HCB]

Huys Q. J. M., Daw N. D. & Dayan P. (2015) Depression: A decision-theoretic analysis. *Annual Review of Neuroscience* **38**(1):1–23. Available at: http://doi.org/10.1146/annurev-neuro-071714-033928. [JS]

Ikeda T. & Hikosaka O. (2003) Reward-dependent gain and bias of visual responses in primate superior colliculus. *Neuron* **39**:693–700. [ES]

Insel T., Cuthbert B., Garvey M., Heinssen R., Pine D. S., Quinn K., Sanislow C. & Wang P. (2010) Research domain criteria (RDoC): Toward a new classification framework for research on mental disorders. *American Journal of Psychiatry* **167**(7):748–51. Available at: http://doi.org/10.1176/appi.ajp.2010.09091379. [JS]

Isoda M. & Hikosaka O. (2008) A neural correlate of motivational conflict in the superior colliculus of the macaque. *Journal of Neurophysiology* **100**:1332–42. [ES]

Jack C. E. & Thurlow W. R. (1973) Effects of degree of visual association and angle of displacement on the "ventriloquism" effect. *Perceptual and Motor Skills* **37**(3):967–79. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=4764534&retmode=ref&cmd=prlinks. [aDR]

Jacobs D. M. & Michaels C. F. (2007) Direct learning. *Ecological Psychology* **19**:321–49. [RW]

Jacobs R. A. (1999) Optimal integration of texture and motion cues to depth. *Vision Research* **39**(21):3621–29. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=10746132&retmode=ref&cmd=prlinks. [aDR]

Jannati A. & Di Lollo V. (2012) Relative blindsight arises from a criterion confound in metacontrast masking: Implications for theories of consciousness. *Consciousness and Cognition* **21**:307–14. [TB]

Jardri R., Duverne S., Litvinova A. S. & Deneve S. (2017) Experimental evidence for circular inference in schizophrenia. *Nature Communications* **8**:14218. Available at: http://doi.org/10.1038/ncomms14218. [JS]

Jastrow J. (1892) Studies from the University of Wisconsin: On the judgment of angles and positions of lines. *American Journal of Psychology* **5**(2):214–48. Available at: http://www.jstor.org/stable/1410867?origin=crossref. [aDR]

Jaynes E. (1957/2003) *Probability theory: The logic of science.* (Original lectures published 1957). Available at: http://www.med.mcgill.ca/epidemiology/hanley/bios601/GaussianModel/JaynesProbabilityTheory.pdf. Cambridge University Press. [arDR, AAS]

Jazayeri M. & Movshon J. A. (2007) A new perceptual illusion reveals mechanisms of sensory decoding. *Nature* **446**(7138):912–15. Available at: http://www.nature.com/doifinder/10.1038/nature05739. [aDR]

Jazayeri M. & Shadlen M. N. (2010) Temporal context calibrates interval timing. *Nature Neuroscience* **13**(8):1020–26. doi:10.1038/nn.2590. [GMC]

Jesteadt W. (1974) Intensity and frequency discrimination in one- and two-interval paradigms. *Journal of the Acoustical Society of America* **55**(6):1266–76. Available at: http://scitation.aip.org/content/asa/journal/jasa/55/6/10.1121/1.1914696. [aDR]

Jolij J. & Lamme V. A. F. (2005) Repression of unconscious information by conscious processing: Evidence from affective blindsight induced by transcranial magnetic stimulation. *Proceedings of the National Academy of Sciences of the United States of America* **102**(30):10747–51. Available at: http://www.pnas.org/content/102/30/10747.abstract. [aDR]

Jones M. & Love B. C. (2011) Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition.

Behavioral and Brain Sciences **34**(4):169–88. Available at: http://www.journals.cambridge.org/abstract_S0140525X10003134. [BCL, aDR, CZ, MZ]

Jones P. R. & Dekker T. M. (2017) The development of perceptual averaging: Learning what to do, not just how to do it. *Developmental Science* **21**:e12584. Available at: https://doi.org/10.1111/desc.12584. [MN]

Juslin P., Nilsson H. & Winman A. (2009) Probability theory, not the very guide of life. *Psychological Review* **116**(4):856–74. Available at: http://psycnet.apa.org/record/2009-18254-007 [aDR]

Kahneman D. (1968) Method, findings, and theory in studies of visual masking. *Psychological Bulletin* **70**:404–25. [TB]

Kahneman D., Slovic P. & Tversky A. (1982a) Judgment under uncertainty. *Science* **185**(4157):1124–31. Available at: https://doi.org/10.1093/oxfordhb/9780195376746.013.0038. [CC]

Kahneman D., Slovic P. & Tversky A. (1982b) *Judgment under uncertainty: Heuristics and biases.* Cambridge University Press. [MZ]

Kahneman D. & Tversky A. (1979) Prospect theory: An analysis of decision under risk. *Econometrica* **47**(2):263–92. Retrieved March 11, 2017. Available at: http://www.jstor.org/stable/1914185?origin=crossref. [AM, arDR]

Kalenscher T., Tobler P. N., Huijbers W., Daselaar S. M. & Pennartz C. (2010) Neural signatures of intransitive preferences. *Frontiers in Human Neuroscience* **4**:49. Available at: http://journal.frontiersin.org/article/10.3389/fnhum.2010.00049/abstract. [aDR]

Kaneko Y. & Sakai K. (2015) Dissociation in decision bias mechanism between probabilistic information and previous decision. *Frontiers in Human Neuroscience* **9**:261. Available at: http://journal.frontiersin.org/article/10.3389/fnhum.2015.00261/abstract. [aDR]

Karşılar H., Simen P., Papadakis S. & Balcı F. (2014) Speed accuracy trade-off under response deadlines. *Frontiers in Neuroscience* **8**:248. Available at: https://www.frontiersin.org/articles/10.3389/fnins.2014.00248/full. [PS]

Keren G. (1988) On the ability of monitoring non-veridical perceptions and uncertain knowledge: Some calibration studies. *Acta Psychologica* **67**(2):95–119. Available at: http://www.sciencedirect.com/science/article/pii/0001691888900078. [aDR]

Kersten D., Mamassian P. & Yuille A. (2004) Object perception as Bayesian inference. *Annual Review of Psychology* **55**:271–304. [CC]

Keshvari S., van den Berg R. & Ma W. J. (2012) Probabilistic computation in human perception under variability in encoding precision. *PLoS ONE* **7**(6):e40216. [WJM]

Kiani R., Corthell L. & Shadlen M. N. (2014) Choice certainty is informed by both evidence and decision time. *Neuron* **84**(6):1329–42. Available at: http://www.sciencedirect.com/science/article/pii/S0896627314010964. [aDR]

Kiani R. & Shadlen M. N. (2009) Representation of confidence associated with a decision by neurons in the parietal cortex. *Science* **324**(5928):759–64. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2738936&tool=pmcentrez&rendertype=abstract. [aDR]

Kinchla R. A. & Smyzer F. (1967) A diffusion model of perceptual memory. *Perception & Psychophysics* **2**(6):219–29. Available at: http://www.springerlink.com/index/10.3758/BF03212471. [aDR]

Knight R., Piette C. E., Page H., Walters D., Marozzi E., Nardini M., Stringer S. & Jeffery K. J. (2014) Weighted cue integration in the rodent head direction system. *Philosophical Transactions of the Royal Society B: Biological Sciences* **369**(1635):20120512. [MZ]

Knill D. C. & Pouget A. (2004) The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences* **27**(12):712–19. Available at: http://www.ncbi.nlm.nih.gov/pubmed/15541511. [RB, aDR, MZ]

Knill D. C. & Richards W., eds. (1996) *Perception as Bayesian inference.* Cambridge University Press. [AAS]

Knill D. C. & Saunders J. A. (2003) Do humans optimally integrate stereo and texture information for judgments of surface slant? *Vision Research* **43**(24):2539–58. [aDR]

Koizumi A., Maniscalco B. & Lau H. (2015) Does perceptual confidence facilitate cognitive control? *Attention, Perception & Psychophysics* **77**(4):1295–306. Available at: http://www.ncbi.nlm.nih.gov/pubmed/25737256. [aDR]

Körding K. P., Beierholm U., Ma W. J., Quartz S., Tenenbaum J. B. & Shams L. (2007) Causal inference in multisensory perception. *PLoS ONE* **2**(9):e943. Available at: http://dx.plos.org/10.1371/journal.pone.0000943. [aDR]

Körding K. P. & Wolpert D. M. (2004) Bayesian integration in sensorimotor learning. *Nature* **427**(6971):244–47. Available at: http://dx.doi.org/10.1038/nature02169 . [aDR, AAS]

Körding K. P. & Wolpert D. M. (2006) Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences* **10**(7):319–26. Available at: https://doi.org/10.1016/j.tics.2006.05.003. [RB, CC, aDR]

Koriat A. (2011) Subjective confidence in perceptual judgments: A test of the self-consistency model. *Journal of Experimental Psychology: General* **140**(1):117–39. Available at: http://www.ncbi.nlm.nih.gov/pubmed/2129932. [aDR]

Kreuder A.-K., Scheele D., Wassermann L., Wollseifer M., Stoffel-Wagner B., Lee M. R., Hennig J., Maier W. & Hurlemann R. (2017) How the brain codes intimacy: The

neurobiological substrates of romantic touch. *Human Brain Mapping* **38**(9):4525–34. Available at: http://doi.org/10.1002/hbm.23679.   [JS]

Landy M. S., Banks M. S. & Knill D. C. (2011) Ideal-observer models of cue integration. In: *Sensory cue integration*, ed. J. Trommershäuser, K. P. Körding & M. S. Landy, pp. 5–29. Oxford University Press.   [aDR]

Landy M. S., Goutcher R., Trommershäuser J. & Mamassian P. (2007) Visual estimation under risk. *Journal of Vision* **7**(6):4. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2638507&tool=pmcentrez&rendertype=abstract.   [aDR]

Landy M. S. & Kojima H. (2001) Ideal cue combination for localizing texture-defined edges. *Journal of the Optical Society of America A, Optics and Image Science* **18**(9):2307–20. Available at: http://www.cns.nyu.edu/~msl/papers/landykojima01.pdf.   [aDR]

Landy M. S., Maloney L., Johnston E. B. & Young M. (1995) Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Research* **35**(3):389–412.   [aDR]

Langer M. S. & Bülthoff H. H. (2001) A prior for global convexity in local shape-from-shading. *Perception* **30**(4):403–10. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils.elink.fcgi?dbfrom=pubmed&id=11383189&retmode=ref&cmd=prlinks.   [aDR]

Lau H. & Passingham R. E. (2006) Relative blindsight in normal observers and the neural correlate of visual consciousness. *Proceedings of the National Academy of Sciences of the United States of America* **103**(49):18763–68. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1693736&tool=pmcentrez&rendertype=abstract.   [aDR]

Lau H. & Rosenthal D. (2011) Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences* **15**(8):365–73. Available at: http://www.ncbi.nlm.nih.gov/pubmed/21737339.   [aDR]

Lauwereyns J., Watanabe K., Coe B. & Hikosaka O. (2002) A neural correlate of response bias in monkey caudate nucleus. *Nature* **418**:413–17.   [ES]

Lawson J. L. & Uhlenbeck G. E., eds. (1950) *Threshold signals*. McGraw-Hill. Retrieved August 7, 2018. Available at: https://trove.nla.gov.au/work/13687839?q&versionId=41165553.   [rDR]

Lawson R. P., Mathys C. & Rees G. (2017) Adults with autism overestimate the volatility of the sensory environment. *Nature Neuroscience* **20**(9):1293–99. Available at: https://doi.org/10.1038/nn.4615.   [MN]

Leekam S. R., Hunnisett E. & Moore C. (1998) Targets and cues: Gaze-following in children with autism. *Journal of Child Psychology and Psychiatry* **39**(7):951–62. Available at: http://www.ncbi.nlm.nih.gov/pubmed/9804028.   [JS]

Lennie P. (2003) The cost of cortical computation. *Current Biology* **13**(6): 493–97. Available at: https://www.sciencedirect.com/science/article/pii/S0960982 203001350.   [aDR]

Leshowitz B. (1969) Comparison of ROC curves from one- and two-interval rating-scale procedures. *Journal of the Acoustical Society of America* **46**(2B): 399–402. Available at: http://scitation.aip.org/content/asa/journal/jasa/46/2B/10.1121/1.1911703.   [aDR]

Lewis R. L., Howes A. & Singh S. (2014) Computational rationality: Linking mechanism and behavior through bounded utility maximization. *Topics in Cognitive Science* **6** (2):279–311.   [AH]

Lewontin R. C. (1970) The units of selection. *Annual Review of Ecology and Systematics* **1**:1–18.   [RW]

Li V., Herce Castanon S., Solomon J. A., Vandormael H. & Summerfield C. (2017) Robust averaging protects decisions from noise in neural computations. *PLoS Computational Biology* **13**(8):e1005723. doi:10.1371/journal.pcbi.1005723.   [CS]

Liberman A., Fischer J. & Whitney D. (2014) Serial dependence in the perception of faces. *Current Biology* **24**(21):2569–74. doi:10.1016/j.cub.2014.09.025.   [GMC, aDR]

Ling S. & Carrasco M. (2006) When sustained attention impairs perception. *Nature Neuroscience* **9**(10):1243–45. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils.elink.fcgi?dbfrom=pubmed&id=16964254&retmode=ref&cmd=prlinks.   [aDR]

Liu T., Abrams J. & Carrasco M. (2009) Voluntary attention enhances contrast appearance. *Psychological Science* **20**(3):354–62. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils.elink.fcgi?dbfrom=pubmed&id=19254239&retmode=ref&cmd=prlinks.   [aDR]

Love B. C., Medin D. L. & Gureckis T. M. (2004) SUSTAIN: A network model of category learning. *Psychological Review* **111**:309–32.   [BCL]

Luce R. D. (2013) Analogs in Luce's global psychophysical theory of Stevens's psychophysical regression effect. *American Journal of Psychology* **126**:47–52.   [DAB]

Lupyan G. (2012) Linguistically modulated perception and cognition: The label-feedback hypothesis. *Frontiers in Psychology* **3**:54. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils.elink.fcgi?dbfrom=pubmed&id=22408629&retmode=ref&cmd=prlinks.   [aDR]

Lupyan G. (2017) The paradox of the universal triangle: Concepts, language, and prototypes. *Quarterly Journal of Experimental Psychology* **70**(3):389–412. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils.elink.fcgi?dbfrom=pubmed&id=26731302&retmode=ref&cmd=prlinks.   [aDR]

Luu L. & Stocker A. A. (2016) Choice-induced biases in perception. *bioRxiv* 043224. Available at: http://biorxiv.org/content/early/2016/04/01/043224.abstract.   [aDR]

Ma W. J. (2010) Signal detection theory, uncertainty, and Poisson-like population codes. *Vision Research* **50**(22):2308–19. Available at: http://www.sciencedirect.com/science/article/pii/S004269891000430X.   [aDR]

Ma W. J. (2012) Organizing probabilistic models of perception. *Trends in Cognitive Sciences* **16**(10):511–18. Retrieved March 2, 2013. Available at: http://www.ncbi.nlm.nih.gov/pubmed/22981359.   [rDR]

Ma W. J., Beck J. M., Latham P. E. & Pouget A. (2006) Bayesian inference with probabilistic population codes. *Nature Neuroscience* **9**(11):1432–38. Available at: http://www.ncbi.nlm.nih.gov/pubmed/17057707.   [aDR]

Mack A., Erol M., Clarke J. & Bert J. (2016) No iconic memory without attention. *Consciousness and Cognition* **40**:1–8.   [TB]

Mack M. L., Preston A. R. & Love B. C. (2013) Decoding the brain's algorithm for categorization from its neural implementation. *Current Biology* **23**:2023–27.   [BCL]

Macmillan N. A. & Creelman C. D. (2005) *Detection theory: A user's guide*. 2nd edition. Erlbaum.   [aDR]

Maddox W. T. (1995) Base-rate effects in multidimensional perceptual categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **21** (2):288–301. Available at: http://www.ncbi.nlm.nih.gov/pubmed/7738501.   [aDR]

Maddox W. T. (2002) Toward a unified theory of decision criterion learning in perceptual categorization. *Journal of the Experimental Analysis of Behavior* **78**(3):567–95. Available at: http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1284916/.   [aDR]

Maddox W. T. & Bohil C. J. (1998a) Base-rate and payoff effects in multidimensional perceptual categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **24**(6):1459–82.   [aDR]

Maddox W. T. & Bohil C. J. (1998b) Overestimation of base-rate differences in complex perceptual categories. *Perception & Psychophysics* **60**(4):575–92.   [aDR]

Maddox W. T. & Bohil C. J. (2000) Costs and benefits in perceptual categorization. *Memory & Cognition* **28**(4):597–615.   [aDR]

Maddox W. T. & Bohil C. J. (2001) Feedback effects on cost-benefit learning in perceptual categorization. *Memory & Cognition* **29**(4):598–615.   [aDR]

Maddox W. T. & Bohil C. J. (2003) A theoretical framework for understanding the effects of simultaneous base-rate and payoff manipulations on decision criterion learning in perceptual categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **29**(2):307–20.   [aDR]

Maddox W. T. & Bohil C. J. (2004) Probability matching, accuracy maximization, and a test of the optimal classifier's independence assumption in perceptual categorization. *Perception & Psychophysics* **66**(1):104–18.   [aDR]

Maddox W. T. & Bohil C. J. (2005) Optimal classifier feedback improves cost-benefit but not base-rate decision criterion learning in perceptual categorization. *Memory & Cognition* **33**(2):303–19.   [aDR]

Maddox W. T., Bohil C. J. & Dodd J. L. (2003) Linear transformations of the payoff matrix and decision criterion learning in perceptual categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **29**(6):1174–93.   [aDR]

Maddox W. T. & Dodd J. L. (2001) On the relation between base-rate and cost-benefit learning in simulated medical diagnosis. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **27**(6):1367–84. Available at: http://www.ncbi.nlm.nih.gov/pubmed/11713873.   [aDR]

Maia T. V., Huys Q. J. M. & Frank M. J. (2017) Theory-based computational psychiatry. *Biological Psychiatry* **82**(6):382–84. Available at: http://doi.org/10.1016/j.biopsych.2017.07.016.   [JS]

Maier A., Scheele D., Spengler F. B., Menba T., Mohr F., Güntürkün O., Stoffel-Wagner B., Kinfe T. M., Maier W., Khalsa S. S. & Hurlemann R. (2018) Oxytocin reduces a chemosensory-induced stress bias in social perception. *Neuropsychopharmacology*. Available at: http://doi.org/10.1038/s41386-018-0063-3.   [JS]

Maiworm M. & Röder B. (2011) Suboptimal auditory dominance in audiovisual integration of temporal cues. *Tsinghua Science & Technology* **16**(2):121–32.   [aDR]

Maksimov M., Vaht M., Harro J. & Bachmann T. (2013) Can common functional gene variants affect visual discrimination in metacontrast masking? *PLoS ONE* **8**(1): e55287. doi:10.1371/journal.pone.0055287.   [TB]

Maksimov M., Vaht M., Harro J. & Bachmann T. (2015a) Single 5HTR2A-1438 A/G nucleotide polymorphism affects performance in a metacontrast masking task: Implications for vulnerability testing and neuromodulation of pyramidal cells. *Neuroscience Letters* **584**:129–34.   [TB]

Maksimov M., Vaht M., Murd C., Harro J. & Bachmann T. (2015b) Brain dopaminergic system related genetic variability interacts with target/mask timing in metacontrast masking. *Neuropsychologia* **71**:112–18.   [TB]

Maloney L. T. & Landy M. S. (1989) A statistical framework for robust fusion of depth information. In: *Proceedings of Society of Photo-Optical Instrumentation Engineers (SPIE) 1119, Visual Communications and Image Processing IV*, ed. W. A. Pearlman, pp. 1154–63. SPIE. Available at: http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=1262206.   [aDR]

Maloney L. T. & Mamassian P. (2009) Bayesian decision theory as a model of human visual perception: Testing Bayesian transfer. *Visual Neuroscience* **26**(1):147–55. Available at: https://doi.org/10.1017/S0952523808080905.   [CC, aDR]

Maloney L. T. & Zhang H. (2010) Decision-theoretic models of visual perception and action. *Vision Research* **50**(23):2362–74. Available at: http://www.ncbi.nlm.nih.gov/pubmed/20932856.   [aDR]

Maniscalco B. & Lau H. (2010) Comparing signal detection models of perceptual decision confidence. *Journal of Vision* **10**(7):213. Available at: http://jov.arvojournals.org/article.aspx?articleid=2138292.   [aDR]

Maniscalco B. & Lau H. (2012) A signal detection theoretic approach for estimating meta-cognitive sensitivity from confidence ratings. *Consciousness and Cognition* 21(1):422–30. Available at: http://www.ncbi.nlm.nih.gov/pubmed/22071269. [aDR]

Maniscalco B. & Lau H. (2015) Manipulation of working memory contents selectively impairs metacognitive sensitivity in a concurrent visual discrimination task. *Neuroscience of Consciousness* 2015(1):niv002. Available at: http://nc.oxfordjournals.org/content/2015/1/niv002.abstract. [aDR]

Maniscalco B. & Lau H. (2016) The signal processing architecture underlying subjective reports of sensory awareness. *Neuroscience of Consciousness* 2016(1): niw002. [aDR]

Maniscalco B., Peters M. A. K. & Lau H. (2016) Heuristic use of perceptual evidence leads to dissociation between performance and metacognitive sensitivity. *Attention, Perception & Psychophysics* 78(3):923–37. Available at: http://www.ncbi.nlm.nih.gov/pubmed/26791233. [aDR]

Manning C., Dakin S. C., Tibber M. S. & Pellicano E. (2014) Averaging, not internal noise, limits the development of coherent motion processing. *Developmental Cognitive Neuroscience* 10:44–56. Available at: https://doi.org/10.1016/j.dcn.2014.07.004. [MN]

Marcus G. F. & Davis E. (2013) How robust are probabilistic models of higher-level cognition? *Psychological Science* 24(12):2351–60. Available at: http://pss.sagepub.com/content/24/12/2351.abstract?ijkey=42fdf6a62d20a7c5e573d149a973e121f7ae2626&keytype2=tf_ipsecsha. [aDR]

Marcus G. F. & Davis E. (2015) Still searching for principles: A response to Goodman et al. (2015) *Psychological Science* 26(4):542–44. Available at: http://pss.sagepub.com/lookup/doi/10.1177/0956797614568433. [aDR]

Markman A. B., Baldwin G. C. & Maddox W. T. (2005) The interaction of payoff structure and regulatory focus in classification. *Psychological Science* 16(11):852–55. Available at: http://www.ncbi.nlm.nih.gov/pubmed/16262768. [aDR]

Markowitz J. & Swets J. A. (1967) Factors affecting the slope of empirical ROC curves: Comparison of binary and rating responses. *Perception & Psychophysics* 2(3):91–100. Available at: http://www.springerlink.com/index/10.3758/BF03210301. [aDR]

Marlin B. J., Mitre M., D'amour J. A., Chao M. V. & Froemke R. C. (2015) Oxytocin enables maternal behaviour by balancing cortical inhibition. *Nature* 520(7548):499–504. Available at: http://doi.org/10.1038/nature14402. [JS]

Marr D. (1982) *Vision: A computational investigation into the human representation and processing of visual information*. W. H. Freeman. [LSG, VW, CZ]

Martignon L. & Hoffrage U. (2002) Fast, frugal, and fit: Simple heuristics for paired comparison. *Theory and Decision* 52(1):29–71. [AM]

Massoni S. (2014) Emotion as a boost to metacognition: How worry enhances the quality of confidence. *Consciousness and Cognition* 29:189–98. Available at: http://www.ncbi.nlm.nih.gov/pubmed/25286128. [aDR]

Massoni S., Gajdos T. & Vergnaud J.-C. (2014) Confidence measurement in the light of signal detection theory. *Frontiers in Psychology* 5:1455. Available at: http://www.ncbi.nlm.nih.gov/pubmed/25566135. [aDR]

Maunsell J. H. (2004) Neuronal representations of cognitive state: Reward or attention? *Trends in Cognitive Sciences* 8:261–65. [ES]

Mayr E. (2002) *What evolution is*. Phoenix. [RW]

Mayr E. (2004) *What makes biology unique?* Cambridge University Press. [RW]

McClelland J. L. (2013) Integrating probabilistic models of perception and interactive neural networks: A historical and tutorial review. *Frontiers in Psychology* 4:503. [BMT]

McConachie H. R. (1976) Developmental prosopagnosia. A single case report. *Cortex* 12 (1):76–82. Available at: http://psycnet.apa.org/record/1976-20939-001. [JS]

McCurdy L. Y., Maniscalco B., Metcalfe J., Liu K. Y., de Lange F. P. & Lau H. (2013) Anatomical coupling between distinct metacognitive systems for memory and visual perception. *Journal of Neuroscience* 33(5):1897–906. Available at: http://www.ncbi.nlm.nih.gov/pubmed/23365229. [aDR]

McKoon G. & Ratcliff R. (2016) Adults with poor reading skills: How lexical knowledge interacts with scores on standardized reading comprehension tests. *Cognition* 146:453–69. Available at: http://doi.org/10.1016/j.cognition.2015.10.009. [JS]

McMillen T. & Holmes P. (2006) The dynamics of choice among multiple alternatives. *Journal of Mathematical Psychology* 50:30–57. [PS]

Metcalfe J. & Shimamura A. P. (1994) *Metacognition: Knowing about knowing*. MIT Press. [aDR]

Meyer D. E. & Kieras D. E. (1997) A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic mechanisms. *Psychological Review* 104(1):3–65. [AH]

Meyer J., Wiczorek R. & Günzler T. (2014) Measures of reliance and compliance in aided visual scanning. *Human Factors* 56(5):840–49. [JM]

Michael E., de Gardelle V., Nevado-Holgado A. & Summerfield C. (2015) Unreliable evidence: 2 Sources of uncertainty during perceptual choice. *Cerebral Cortex* 25 (4):937–47. Available at: http://www.ncbi.nlm.nih.gov/pubmed/24122138. [aDR]

Michael E., de Gardelle V. & Summerfield C. (2014) Priming by the variability of visual information. *Proceedings of the National Academy of Sciences of the United States of America* 111(21):7873–78. Available at: http://www.ncbi.nlm.nih.gov/pubmed/24821803. [aDR]

Mondloch C. J., Le Grand R. & Maurer D. (2002) Configural face processing develops more slowly than featural face processing. *Perception* 31:553–66. Available at: https://doi.org/10.1068/p3339. [MN]

Morales J., Solovey G., Maniscalco B., Rahnev D., de Lange F. P. & Lau H. (2015) Low attention impairs optimal incorporation of prior knowledge in perceptual decisions. *Attention, Perception & Psychophysics* 77(6):2021–36. Available at: http://www.ncbi.nlm.nih.gov/pubmed/25836765. [aDR]

Moran R. (2015) Optimal decision making in heterogeneous and biased environment. *Psychonomic Bulletin & Review* 22:38–53. [RM]

Moritz S. & Woodward T. S. (2007) Metacognitive training in schizophrenia: From basic research to knowledge translation and intervention. *Current Opinion in Psychiatry* 20 (6):619–25. Available at: http://doi.org/10.1097/YCO.0b013e3282f0b8ed. [JS]

Mou W. & Zhang L. (2014) Dissociating position and heading estimations: Rotated visual orientation cues perceived after walking reset headings but not positions. *Cognition* 133:553–71. [MZ]

Moustafa A. A., Kéri S., Somlai Z., Balsdon T., Frydecka D., Misiak B. & White C. (2015) Drift diffusion model of reward and punishment learning in schizophrenia: Modeling and experimental data. *Behavioural Brain Research* 291:147–54. Available at: http://doi.org/10.1016/j.bbr.2015.05.024. [JS]

Mozer M. C., Pashler H. & Homaei H. (2008) Optimal predictions in everyday cognition: The wisdom of individuals or crowds? *Cognitive Science* 32(7):1133–47. [aDR]

Mueller S. T. & Weidemann C. T. (2008) Decision noise: An explanation for observed violations of signal detection theory. *Psychonomic Bulletin & Review* 15(3):465–94. Available at: http://www.springerlink.com/index/10.3758/PBR.15.3.465. [aDR]

Mulder M. J., Bos D., Weusten J. M. H., van Belle J., van Dijk S. C., Simen P., van Engeland H. & Durston S. (2010) Basic impairments in regulating the speed-accuracy tradeoff predict symptoms of attention-deficit/hyperactivity disorder. *Biological Psychiatry* 68(12):1114–19. Available at: http://doi.org/10.1016/j.biopsych.2010.07.031. [JS]

Müller J. (1837–1840/1938) *The elements of physiology*. Murray. [RW]

Müller-Lyer F. C. (1889) Optische Urteilstäuschungen. *Archiv für Anatomie und Physiologie, Physiologische Abteilung* 2:263–70. [rDR]

Myers C. W., Lewis R. L. & Howes A. (2013) Bounded optimal state estimation and control in visual search: Explaining distractor ratio effects. *Proceedings of the Annual Meeting of the Cognitive Science Society* 35:1049–54. [AH]

Nardini M., Bedford R. & Mareschal D. (2010) Fusion of visual cues is not mandatory in children. *Proceedings of the National Academy of Sciences of the United States of America* 107(39):17041–46. Available at: https://doi.org/10.1073/pnas.1001699107. [MN, aDR]

Nardini M., Jones P., Bedford R. & Braddick O. (2008) Development of cue integration in human navigation. *Current Biology* 18(9):689–93. Available at: https://doi.org/10.1016/j.cub.2008.04.021. [MN, aDR, MZ]

Navajas J., Hindocha C., Foda H., Keramati M., Latham P. E. & Bahrami B. (2017) The idiosyncratic nature of confidence. *Nature Human Behaviour* 1(11):810–18. Available at: http://www.nature.com/articles/s41562-017-0215-1. [aDR]

Navajas J., Sigman M. & Kamienkowski J. E. (2014) Dynamics of visibility, confidence, and choice during eye movements. *Journal of Experimental Psychology: Human Perception and Performance* 40(3):1213–27. Available at: http://www.ncbi.nlm.nih.gov/pubmed/24730743. [aDR]

Norton E. H., Fleming S. M., Daw N. D. & Landy M. S. (2017) Suboptimal criterion learning in static and dynamic environments. *PLoS Computational Biology* 13(1): e1005304. [aDR]

Odegaard B., Wozny D. R. & Shams L. (2015) Biases in visual, auditory, and audiovisual perception of space. *PLoS Computational Biology* 11(12):e1004649. Available at: http://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1004649. [aDR]

Ohshiro T., Angelaki D. E. & DeAngelis G. C. (2011) A normalization model of multisensory integration. *Nature Neuroscience* 14(6):775–82. Available at: https://doi.org/10.1038/nn.2815. [MN]

Olzak L. A. (1985) Interactions between spatially tuned mechanisms: Converging evidence. *Journal of the Optical Society of America A, Optics and Image Science* 2 (9):1551–59. [aDR]

Oruç I., Maloney L. T. & Landy M. S. (2003) Weighted linear cue combination with possibly correlated error. *Vision Research* 43(23):2451–68. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=12972395&retmode=ref&cmd=prlinks. [aDR]

Osgood C. E. (1953) *Method and theory in experimental psychology*. Oxford University Press. [aDR]

Osmmy O., Moran R., Pfeffer T., Tsetsos K., Usher M. & Donner T. (2013) The time scale of perceptual evidence integration can be adapted to the environment. *Current Biology* 23:981–86. [RM]

Oud B., Krajbich I., Miller K., Cheong J. H., Botvinick M. & Fehr E. (2016) Irrational time allocation in decision-making. *Proceedings of the Royal Society B: Biological Sciences* 283(1822):20151439. Available at: http://www.ncbi.nlm.nih.gov/pubmed/26763695. [aDR]

Palmer C. J. & Clifford C. W. G. (2017) Functional mechanisms encoding others' direction of gaze in the human nervous system. *Journal of Cognitive Neuroscience* 29 (10):1725–38. Available at: http://doi.org/10.1162/jocn_a_01150. [JS]

Palmer C. J., Lawson R. P., Shankar S., Clifford C. W. G. & Rees G. (2018) Autistic adults show preserved normalisation of sensory responses in gaze processing. *Cortex* **103**:13–23. Available at: http://doi.org/10.1016/j.cortex.2018.02.005.   [JS]

Palminteri S., Wyart V. & Koechlin E. (2017) The importance of falsification in computational cognitive modeling. *Trends in Cognitive Sciences* **21**(6):425–33.   [VW]

Partos T. R., Cropper S. J. & Rawlings D. (2016) You don't see what I see: Individual differences in the perception of meaning from visual stimuli. *PLoS ONE* **11**(3):e0150615. Available at: http://doi.org/10.1371/journal.pone.0150615.   [JS]

Peck C. J., Jangraw D. C., Suzuki M., Efem R. & Gottlieb J. (2009) Reward modulates attention independently of action value in posterior parietal cortex. *Journal of Neuroscience* **29**:11182–91.   [ES]

Peters M. A. K., Ma W. J. & Shams L. (2016) The size-weight illusion is not anti-Bayesian after all: A unifying Bayesian account. *PeerJ* **4**:e2124. Available at: http://www.ncbi. nlm.nih.gov/pubmed/27350899.   [aDR]

Petrini K., Remark A., Smith L. & Nardini M. (2014) When vision is not an option: Children's integration of auditory and haptic information is suboptimal. *Developmental Science* **17**(3):376–87. Available at: http://onlinelibrary.wiley.com/doi/ 10.1111/desc.12127/full.   [aDR]

Petzschner F. H. & Glasauer S. (2011) Iterative Bayesian estimation as an explanation for range and regression effects: A study on human path integration. *Journal of Neuroscience* **31**(47):17220–29. Available at: http://www.jneurosci.org/content/31/47/ 17220.   [aDR]

Petzschner F. H., Glasauer S. & Stephan K. E. (2015) A Bayesian perspective on magnitude estimation. *Trends in Cognitive Sciences* **19**:285–93.   [HCB]

Pinto N., Doukhan D., DiCarlo J. J. & Cox D. D. (2009) A high-throughput screening approach to discovering good forms of biologically inspired visual representation. *PLoS Computational Biology* **5**(11):e1000579.   [WJM]

Plaisier M. A., van Dam L. C. J., Glowania C. & Ernst M. O. (2014) Exploration mode affects visuohaptic integration of surface orientation. *Journal of Vision* **14**(13):22. Available at: http://www.ncbi.nlm.nih.gov/pubmed/25413627.   [aDR]

Pleskac T. J. & Busemeyer J. R. (2010) Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review* **117**(3):864–901. Available at: http://www.ncbi.nlm.nih.gov/pubmed/20658856.   [aDR]

Poulton E. C. (1989) *Bias in quantifying judgements*. Taylor & Francis.   [DAB]

Powers R., Mathys C. & Corlett P. R. (2017) Pavlovian conditioning–induced hallucinations result from overweighting of perceptual priors. *Science* **357**:596–600.   [TB]

Preciado D., Munneke J. & Theeuwes J. (2017) Mixed signals: The effect of conflicting reward- and goal-driven biases on selective attention. *Attention Perception & Psychophysics* **79**:1297–310.   [ES]

Prsa M., Gale S. & Blanke O. (2012) Self-motion leads to mandatory cue fusion across sensory modalities. *Journal of Neurophysiology* **108**(8):2282–91. Available at: http:// eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=22832567&retmode =ref&cmd=prlinks.   [aDR]

Pynn C. T. (1972) Intensity perception. III. Resolution in small-range identification. *Journal of the Acoustical Society of America* **51**(2B):559–66. Available at: http:// scitation.aip.org/content/asa/journal/jasa/51/2B/10.1121/1.1912878.   [aDR]

Rahnev D., Bahdo L., de Lange F. P. & Lau H. (2012a) Prestimulus hemodynamic activity in dorsal attention network is negatively associated with decision confidence in visual perception. *Journal of Neurophysiology* **108**(5):1529–36. Available at: http:// www.ncbi.nlm.nih.gov/pubmed/22723670.   [aDR]

Rahnev D., Koizumi A., McCurdy L. Y., D'Esposito M. & Lau H. (2015) Confidence leak in perceptual decision making. *Psychological Science* **26**(11):1664–80. Available at: http://pss.sagepub.com/lookup/doi/10.1177/0956797615595037.   [aDR]

Rahnev D., Kok P., Munneke M., Bahdo L., de Lange F. P. & Lau H. (2013) Continuous theta burst transcranial magnetic stimulation reduces resting state connectivity between visual areas. *Journal of Neurophysiology* **110**(8):1811–21. Available at: http://www.ncbi.nlm.nih.gov/pubmed/23883858.   [aDR]

Rahnev D., Lau H. & de Lange F. P. (2011a) Prior expectation modulates the interaction between sensory and prefrontal regions in the human brain. *Journal of Neuroscience* **31**(29):10741–48.   [aDR]

Rahnev D., Maniscalco B., Graves T., Huang E., de Lange F. P. & Lau H. (2011b) Attention induces conservative subjective biases in visual perception. *Nature Neuroscience* **14**(12):1513–15. Available at: http://www.ncbi.nlm.nih.gov/pubmed/ 22019729.   [aDR]

Rahnev D., Maniscalco B., Luber B., Lau H. & Lisanby S. H. (2012b) Direct injection of noise to the visual cortex decreases accuracy but increases decision confidence. *Journal of Neurophysiology* **107**(6):1556–63. Available at: http://www.ncbi.nlm.nih.gov/ pubmed/22170965.   [aDR]

Rahnev D., Nee D. E., Riddle J., Larson A. S. & D'Esposito M. (2016) Causal evidence for frontal cortex organization for perceptual decision making. *Proceedings of the National Academy of Sciences of the United States of America* **113**(20):6059–64. Available at: http://www.pnas.org/content/early/2016/05/04/1522551113.full?tab= metrics.   [aDR]

Ramachandran V. (1990) Interactions between motion, depth, color and form: The utilitarian theory of perception. In: *Vision: Coding and efficiency*, ed. C. Blakemore, pp. 346–60. Cambridge University Press.   [CC, aDR]

Rand M. K. & Shimansky Y. P. (2013) Two-phase strategy of neural control for planar reaching movements: I. XY coordination variability and its relation to endpoint variability. *Experimental Brain Research* **225**:55–73.   [YPS]

Ratcliff R. (1978) A theory of memory retrieval. *Psychological Review* **85**(2):59–108. Available at: http://doi.org/10.1037/0033-295X.85.2.59.   [JS]

Ratcliff R. & McKoon G. (2008) The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation* **20**(4):873–922.   [AM]

Ratcliff R., Perea M., Colangelo A. & Buchanan L. (2004) A diffusion model account of normal and impaired readers. *Brain and Cognition* **55**(2):374–82. Available at: http:// doi.org/10.1016/j.bandc.2004.02.051.   [JS]

Ratcliff R., Smith P. L., Brown S. D. & McKoon G. (2016) Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences* **20**(4):260–81. Available at: http://doi.org/10.1016/j.tics.2016.01.007.   [JS]

Ratcliff R. & Starns J. J. (2009) Modeling confidence and response time in recognition memory. *Psychological Review* **116**(1):59–83. Available at: http://www. pubmedcentral.nih.gov/articlerender.fcgi?artid=2693899&tool=pmcentrez&rendertype =abstract.   [aDR]

Ratcliff R., Thapar A. & McKoon G. (2010) Individual differences, aging, and IQ in two-choice tasks. *Cognitive Psychology* **60**(3):127–57. Available at: http://doi.org/10. 1016/j.cogpsych.2009.09.001.   [JS]

Ratcliff R., Thapar A. & McKoon G. (2011) Effects of aging and IQ on item and associative memory. *Journal of Experimental Psychology: General* **140**(3):464–87. Available at: http://doi.org/10.1037/a0023810.   [JS]

Rauber H. J. & Treue S. (1998) Reference repulsion when judging the direction of visual motion. *Perception* **27**(4):393–402. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/ eutils/elink.fcgi?dbfrom=pubmed&id=9797918&retmode=ref&cmd=prlinks.   [aDR]

Raviv O., Ahissar M. & Loewenstein Y. (2012) How recent history affects perception: The normative approach and its heuristic approximation. *PLoS Computational Biology* **8** (10):e1002731. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi? dbfrom=pubmed&id=23133343&retmode=ref&cmd=prlinks.   [aDR]

Reckless G. E., Bolstad I., Nakstad P. H., Andreassen O. A. & Jensen J. (2013) Motivation alters response bias and neural activation patterns in a perceptual decision-making task. *Neuroscience* **238**:135–47. Available at: http://www.ncbi.nlm.nih.gov/ pubmed/23428623.   [aDR]

Reckless G. E., Ousdal O. T., Server A., Walter H., Andreassen O. A. & Jensen J. (2014) The left inferior frontal gyrus is involved in adjusting response bias during a perceptual decision-making task. *Brain and Behavior* **4**(3):398–407. Available at: http://www. pubmedcentral.nih.gov/articlerender.fcgi?artid=4055190&tool=pmcentrez&rendertype =abstract.   [aDR]

Reed E. S. (1982) The corporeal ideas hypothesis and the origin of scientific psychology. *Review of Metaphysics* **35**:731–52.   [RW]

Regenwetter M., Cavagnaro D. R., Popova A., Guo Y., Zwilling C., Lim S. H. & Stevens J. R. (2017) Heterogeneity and parsimony in intertemporal choice. *Decision* **5**(2):63–94. Available at: http://doi.apa.org/getdoi.cfm?doi=10.1037/dec0000069.   [aDR]

Regenwetter M., Dana J. & Davis-Stober C. P. (2010) Testing transitivity of preferences on two-alternative forced choice data. *Frontiers in Psychology* **1**:148. Available at: http:// journal.frontiersin.org/article/10.3389/fpsyg.2010.00148/abstract.   [aDR]

Regenwetter M., Dana J., Davis-Stober C. P. & Guo Y. (2011) Parsimonious testing of transitive or intransitive preferences: Reply to Birnbaum (2011) *Psychological Review* **118**(4):684–88. Available at: http://doi.apa.org/getdoi.cfm?doi=10.1037/a0025291. [aDR]

Renart A. & Machens C. K. (2014) Variability in neural activity and behavior. *Current Opinion in Neurobiology* **25**:211–20. Available at: http://dx.doi.org/10.1016/j.conb. 2014.02.013.   [aDR]

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.   [RB]

Roach N. W., Heron J. & McGraw P. V. (2006) Resolving multisensory conflict: A strategy for balancing the costs and benefits of audio-visual integration. *Proceedings of the Royal Society B: Biological Sciences* **273**(1598):2159–68. doi:10.1098/rspb.2006.3578. [GMC, aDR]

Rosas P., Wagemans J., Ernst M. O. & Wichmann F. A. (2005) Texture and haptic cues in slant discrimination: Reliability-based cue weighting without statistically optimal cue combination. *Journal of the Optical Society of America A, Optics and Image Science* **22**(5):801–809.   [aDR]

Rosas P. & Wichmann F. A. (2011) Cue combination: Beyond optimality. In: *Sensory cue integration*, ed. J. Trommershäuser, K. P. Körding & M. S. Landy, pp. 144–52. Oxford University Press.   [aDR]

Rosas P., Wichmann F. A. & Wagemans J. (2007) Texture and object motion in slant discrimination: Failure of reliability-based weighting of cues may be evidence for strong fusion. *Journal of Vision* **7**(6):3. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink. fcgi?dbfrom=pubmed&id=17685786&retmode=ref&cmd=prlinks.   [aDR]

Rosenberg A. & McShea D. W. (2008) *Philosophy of biology: A contemporary introduction*. Routledge.   [rDR]

Rosenberg A., Patterson J. S. & Angelaki D. E. (2015) A computational perspective on autism. *Proceedings of the National Academy of Sciences of the United States of*

America **112**(30):9158–65. Available at: https://doi.org/10.1073/pnas.1510583112. [MN, JS]

Rumsfeld D. (2011) *Known and unknown: A memoir.* Penguin. [RB]

Russell R., Duchaine B. & Nakayama K. (2009) Super-recognizers: People with extraordinary face recognition ability. *Psychonomic Bulletin and Review* **16**(2):252–57. Available at: http://doi.org/10.3758/PBR.16.2.252. [JS]

Russell S. J. & Subramanian D. (1995) Provably bounded-optimal agents. *Journal of Artificial Intelligence Research* **2**:575–609. [AH]

Saarela T. P. & Landy M. S. (2015) Integration trumps selection in object recognition. *Current Biology* **25**(7):920–27. Available at: http://www.ncbi.nlm.nih.gov/pubmed/25802154. [aDR]

Sabra A. I. (1989) *The optics of Ibn Al-Haytham, Books I–III: On direct vision.* Warburg Institute. [aDR]

Samaha J., Barrett J. J., Sheldon A. D., LaRocque J. J. & Postle B. R. (2016) Dissociating perceptual confidence from discrimination accuracy reveals no influence of metacognitive awareness on working memory. *Frontiers in Psychology* **7**:851. Available at: http://journal.frontiersin.org/Article/10.3389/fpsyg.2016.00851/abstract. [aDR]

Sanborn A. N., Griffiths T. L. & Navarro D. J. (2010) Rational approximations to rational models: Alternative algorithms for category learning. *Psychological Review* **4**:1144–67. [BMT, CZ]

Sandberg K., Timmermans B., Overgaard M. & Cleeremans A. (2010) Measuring consciousness: Is one measure better than the other? *Consciousness and Cognition* **19** (4):1069–78. [TB]

Sanders J. I., Hangya B. & Kepecs A. (2016) Signatures of a statistical computation in the human sense of confidence. *Neuron* **90**(3):499–506. Available at: http://www.cell.com/article/S0896627316300162/fulltext. [aDR]

Scheele D., Kendrick K. M., Khouri C., Kretzer E., Schläpfer T. E., Stoffel-Wagner B., Güntürkün O., Maier W. & Hurlemann R. (2014) An oxytocin-induced facilitation of neural and emotional responses to social touch correlates inversely with autism traits. *Neuropsychopharmacology* **39**(9):2078–85. Available at: http://doi.org/10.1038/npp.2014.78. [JS]

Scheele D., Schwering C., Elison J. T., Spunt R., Maier W. & Hurlemann R. (2015) A human tendency to anthropomorphize is enhanced by oxytocin. *European Neuropsychopharmacology* **25**(10):1817–23. Available at: http://doi.org/10.1016/j.euroneuro.2015.05.009. [JS]

Scheele D., Wille A., Kendrick K. M., Stoffel-Wagner B., Becker B., Güntürkün O., Maier W. & Hurlemann R. (2013) Oxytocin enhances brain reward system responses in men viewing the face of their female partner. *Proceedings of the National Academy of Sciences of the United States of America* **110**(50):20308–13. Available at: http://doi.org/10.1073/pnas.1314190110. [JS]

Schmiedek F., Oberauer K., Wilhelm O., Süss H.-M. & Wittmann W. W. (2007) Individual differences in components of reaction time distributions and their relations to working memory and intelligence. *Journal of Experimental Psychology: General* **136** (3):414–29. Available at: http://doi.org/10.1037/0096-3445.136.3.414. [JS]

Schroeder C. E., Wilson D. A., Radman T., Scharfman H. & Lakatos P. (2010) Dynamics of active sensing and perceptual selection. *Current Opinion in Neurobiology* **20**:172–76. Available at: https://doi.org/10.1016/j.conb.2010.02.010 [J-PN]

Schulman A. I. & Mitchell R. R. (1966) Operating characteristics from yes-no and forced-choice procedures. *Journal of the Acoustical Society of America* **40**(2):473–77. Available at: http://www.ncbi.nlm.nih.gov/pubmed/5911357. [aDR]

Schultz J. & Bülthoff H. H. (2013) Parametric animacy percept evoked by a single moving dot mimicking natural stimuli. *Journal of Vision* **13**(4):1–19. Available at: http://doi.org/10.1167/13.4.15. [JS]

Schurger A., Kim M.-S. & Cohen J. D. (2015) Paradoxical interaction between ocular activity, perception, and decision confidence at the threshold of vision. *PLoS ONE* **10**(5):e0125278. Available at: http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0125278. [aDR]

Schwiedrzik C. M., Ruff C. C., Lazar A., Leitner F. C., Singer W. & Melloni L. (2014) Untangling perceptual memory: Hysteresis and adaptation map into separate cortical networks. *Cerebral Cortex* **24**(5):1152–64. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=23236204&retmode=ref&cmd=prlinks. [aDR]

Sciutti A., Burr D., Saracco A., Sandini G. & Gori M. (2015) Development of context-dependency in human space perception. *Experimental Brain Research* **232**:3965–76. [HCB]

See J. E., Warm J. S., Dember W. N. & Howe S. R. (1997) Vigilance and signal detection theory: An empirical evaluation of five measures of response bias. *Human Factors* **39** (1):14–29. Available at: http://hfs.sagepub.com/cgi/doi/10.1518/001872097778940704. [aDR]

Senju A., Southgate V., White S. & Frith U. (2009) Mindblind eyes: An absence of spontaneous theory of mind in Asperger syndrome. *Science* **325**(5942):883–85. Available at: http://doi.org/10.1126/science.1176170. [JS]

Seriès P. & Seitz A. R. (2013) Learning what to expect (in visual perception). *Frontiers in Human Neuroscience* **7**:668. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=24187536&retmode=ref&cmd=prlinks. [aDR]

Servan-Schreiber D., Printz H. & Cohen J. D. (1990) A network model of catecholamine effects: Gain, signal-to-noise ratio, and behavior. *Science* **249**(4971):892–95. [RM]

Shagrir O. (2010) Marr on computational-level theories. *Philosophy of Science* **77**(4):477–500. [CZ]

Shalom S. & Zaidel A. (2018) Better than optimal. *Neuron* **97**(3):484–87. Available at: https://doi.org/10.1016/j.neuron.2018.01.041. [J-PN]

Shannon C. E. & Weaver W. (1949) *The mathematical theory of information.* University of Illinois Press. [DAB]

Shen S. & Ma W. J. (2016) A detailed comparison of optimality and simplicity in perceptual decision making. *Psychological Review* **123**(4):452–80. Available at: http://www.ncbi.nlm.nih.gov/pubmed/27177259. [WJM, aDR]

Shen S. & Ma W. J. (in press) Variable precision in visual perception. Available at: http://psycnet.apa.org/doiLanding?doi=10.1037%2Frev0000128. *Psychological Review.* [WJM]

Sherman M. T., Seth A. K., Barrett A. B. & Kanai R. (2015) Prior expectations facilitate metacognition for perceptual decision. *Consciousness and Cognition* **35**:53–65. Available at: http://www.sciencedirect.com/science/article/pii/S1053810015000926. [aDR]

Shi L., Griffiths T. L., Feldman N. H. & Sanborn A. N. (2010) Exemplar models as a mechanism for performing Bayesian inference. *Psychonomic Bulletin and Review* **17**:443–64. [BMT]

Shimansky Y. P. (2011) State estimation bias induced by optimization under uncertainty and error cost asymmetry is likely reflected in perception. *Biological Cybernetics* **104**:225–33. [YPS]

Shimansky Y. P., Kang T. & He J. (2004) A novel model of motor learning capable of developing an optimal movement control law online from scratch. *Biological Cybernetics* **90**:133–45. [YPS]

Shimansky Y. P. & Rand M. K. (2013) Two-phase strategy of controlling motor coordination determined by task performance optimality. *Biological Cybernetics* **107**:107–29. [YPS]

Simen P., Balcı F., deSouza L., Holmes P. & Cohen J. D. (2011) A model of interval timing by neural integration. *Journal of Neuroscience* **31**(25):9238–53. Available at: https://doi.org/10.1523/JNEUROSCI.3121-10.2011. [PS]

Simen P., Cohen J. D. & Holmes P. (2006) Rapid decision threshold modulation by reward rate in a neural network. *Neural Networks* **19**:1013–26. Available at: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1808344/. [PS]

Simen P., Contreras D., Buck C., Hu P., Holmes P. & Cohen J. D. (2009) Reward-rate optimization in two-alternative decision making: Empirical tests of theoretical predictions. *Journal of Experimental Psychology: Human Perception and Performance* **35**:1865–94. Available at: http://dx.doi.org/10.1037/a0016926. [arDR, PS]

Simon H. A. (1955) A behavioral model of rational choice. *Quarterly Journal of Economics* **69**(1):99–118. [AM]

Simon H. A. (1956) Rational choice and the structure of the environment. *Psychological Review* **63**(2):129–38. [AM, aDR]

Simon H. A. (1957) A behavioral model of rational choice. In: *Models of man, social and rational: Mathematical essays on rational human behavior in a social setting*, pp. 99–118. Wiley. [aDR]

Simon H. A. (1996) *The sciences of the artificial*, 3rd edition. MIT Press. [AM]

Simoncelli E. P. (2003) Vision and the statistics of the visual environment. *Current Opinion in Neurobiology* **13**(2):144–49. [CS]

Snyder J. S., Schwiedrzik C. M., Vitela A. D. & Melloni L. (2015) How previous experience shapes perception in different sensory modalities. *Frontiers in Human Neuroscience* **9**:594. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom= pubmed&id=26582982&retmode=ref&cmd=prlinks. [aDR]

Solovey G., Graney G. G. & Lau H. (2015) A decisional account of subjective inflation of visual perception at the periphery. *Attention, Perception & Psychophysics* **77**(1):258–71. Available at: http://www.ncbi.nlm.nih.gov/pubmed/25248620. [aDR]

Soltani A., De Martino B. & Camerer C. (2012) A range-normalization model of context-dependent choice: A new model and evidence. *PLoS Computational Biology* **8**(7): e1002607. [RM]

Song C., Koizumi A. & Lau H. (2015) A behavioral method to manipulate metacognitive awareness independent of stimulus awareness. In: *Behavioral methods in consciousness research*, ed. M. Overgaard, pp. 77–85. Oxford University Press. [aDR]

Song C., Kanai R., Fleming S. M., Weil R. S., Schwarzkopf D. S. & Rees G. (2011) Relating inter-individual differences in metacognitive performance on different perceptual tasks. *Consciousness and Cognition* **20**(4):1787–92. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3203218&tool=pmcentrez&rendertype=abstract. [aDR]

Spence M. L., Dux P. E. & Arnold D. H. (2016) Computations underlying confidence in visual perception. *Journal of Experimental Psychology: Human Perception and Performance* **42**(5):671–82. Available at: http://www.ncbi.nlm.nih.gov/pubmed/26594876. [aDR]

Spengler F. B., Scheele D., Marsh N., Kofferath C., Flach A., Schwarz S., Stoffel-Wagner B., Maier W. & Hurlemann R. (2017a) Oxytocin facilitates reciprocity in social communication. *Social Cognitive and Affective Neuroscience* **12**(8):1325–33. Available at: http://doi.org/10.1093/scan/nsx061. [JS]

Spengler F. B., Schultz J., Scheele D., Essel M., Maier W., Heinrichs M. & Hurlemann R. (2017b) Kinetics and dose dependency of intranasal oxytocin effects on amygdala reactivity. *Biological Psychiatry* 82:885–94. Available at: http://doi.org/10.1016/j.biopsych.2017.04.015. [JS]

Spitzer B., Waschke L. & Summerfield C. (2017) Selective overweighting of larger magnitudes during numerical comparison. *Nature Human Behaviour* 1:0145. doi:10.1038/s41562-017-0145. [CS]

Stanford T. R., Shankar S., Massoglia D. P., Costello M. G. & Salinas E. (2010) Perceptual decision making in less than 30 milliseconds. *Nature Neuroscience* 13:379–85. [ES]

Starns J. J. & Ratcliff R. (2010) The effects of aging on the speed–accuracy compromise: Boundary optimality in the diffusion model. *Psychology and Aging* 25(2):377–90. Available at: http://dx.doi.org/10.1037/a0018022. [aDR, PS]

Starns J. J. & Ratcliff R. (2012) Age-related differences in diffusion model boundary optimality with both trial-limited and time-limited tasks. *Psychonomic Bulletin & Review* 19(1):139–45. Available at: http://www.ncbi.nlm.nih.gov/pubmed/22144142. [aDR]

Stein B. E. & Meredith M. A. (1993) *The merging of the senses*. MIT Press. [J-PN]

Stewart N., Brown G. D. A. & Chater N. (2005) Absolute identification by relative judgment. *Psychological Review* 112:881–911. [DAB]

Stocker A. A. & Simoncelli E. P. (2006a) Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience* 9(4):578–85. Available at: http://dx.doi.org/10.1038/nn1669. [aDR, AAS]

Stocker A. A. & Simoncelli E. P. (2006b) Sensory adaptation within a Bayesian framework for perception. In: *Advances in neural information processing systems 18* (proceedings from the conference, Neural Information Processing Systems 2005), ed. Y. Weiss & B. Schölkopf & J. C. Platt. Available at: https://papers.nips.cc/book/advances-in-neural-information-processing-systems-18-2005. [aDR]

Stocker A. A. & Simoncelli E. P. (2008) A Bayesian model of conditioned perception. In: *Advances in neural information processing systems 20* (proceedings from the conference, Neural Information Processing Systems 2007), ed. J. C. Platt, D. Koller, Y. Singer & S. Roweis. Available at: https://papers.nips.cc/paper/3369-a-bayesian-model-of-conditioned-perception. [aDR]

Stone J. V. (2011) Footprints sticking out of the sand. Part 2: Children's Bayesian priors for shape and lighting direction. *Perception* 40(2):175–90. Available at: https://doi.org/10.1068/p6776. [MN]

Stone L. S. & Thompson P. (1992) Human speed perception is contrast dependent. *Vision Research* 32(8):1535–49. Available at: http://www.ncbi.nlm.nih.gov/pubmed/1455726. [aDR]

Störmer V. S., Mcdonald J. J. & Hillyard S. A. (2009) Cross-modal cueing of attention alters appearance and early cortical processing of visual stimuli. *Proceedings of the National Academy of Sciences of the United States of America* 106(52):22456–61. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=20007778&retmode=ref&cmd=prlinks. [aDR]

Striepens N., Matusch A., Kendrick K. M., Mihov Y., Elmenhorst D., Becker B., Lang M., Coenen H. H., Maier W., Hurlemann R. & Bauer A. (2014) Oxytocin enhances attractiveness of unfamiliar female faces independent of the dopamine reward system. *Psychoneuroendocrinology* 39:74–87. Available at: http://dx.doi.org/10.1016/j.psyneuen.2013.09.026. [JS]

Stüttgen M. O., Schwarz C. & Jäkel F. (2011) Mapping spikes to sensations. *Frontiers in Neuroscience* 5:125. [CZ]

Summerfield C. & Koechlin E. (2010) Economic value biases uncertain perceptual choices in the parietal and prefrontal cortices. *Frontiers in Human Neuroscience* 4:208. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3024559&tool=pmcentrez&rendertype=abstract. [aDR]

Summerfield C. & Tsetsos K. (2015) Do humans make good decisions? *Trends in Cognitive Sciences* 19(1):27–34. Available at: https://doi.org/10.1007/s11103-011-9767-z. [CC, aDR]

Sun J. & Perona P. (1997) Shading and stereo in early perception of shape and reflectance. *Perception* 26(4):519–29. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=9404497&retmode=ref&cmd=prlinks. [aDR]

Susilo T. & Duchaine B. (2013) Advances in developmental prosopagnosia research. *Current Opinion in Neurobiology* 23(3):423–29. Available at: http://doi.org/10.1016/j.conb.2012.12.011. [JS]

Sweeny T. D., Wurnitsch N., Gopnik A. & Whitney D. (2015) Ensemble perception of size in 4–5-year-old children. *Developmental Science* 18(4):556–68. Available at: https://doi.org/10.1111/desc.12239. [MN]

Swets J. A. & Green D. M. (1961) Sequential observations by human observers of signals in noise. In: *Information theory: Proceedings of the fourth London symposium*, ed. C. Cherry, pp. 177–95. Butterworth. [aDR]

Swets J. A., Tanner W. P. & Birdsall T. G. (1961) Decision processes in perception. *Psychological Review* 68(5):301–40. Available at: http://www.ncbi.nlm.nih.gov/pubmed/13774292. [aDR, CZ]

Tachibana Y. & Hikosaka O. (2012) The primate ventral pallidum encodes expected reward value and regulates motor action. *Neuron* 76:826–37. [ES]

Takikawa Y., Kawagoe R. & Hikosaka O. (2004) A possible role of midbrain dopamine neurons in short- and long-term adaptation of saccades to position-reward mapping. *Journal of Neurophysiology* 92:2520–29. [ES]

Takikawa Y., Kawagoe R., Itoh H., Nakahara H. & Hikosaka O. (2002) Modulation of saccadic eye movements by predicted reward outcome. *Experimental Brain Research* 142:284–91. [ES]

Tanner T. A., Haller R. W. & Atkinson R. C. (1967) Signal recognition as influenced by presentation schedules. *Perception & Psychophysics* 2(8):349–58. Available at: http://www.springerlink.com/index/10.3758/BF03210070. [aDR]

Tanner W. P. (1956) Theory of recognition. *Journal of the Acoustical Society of America* 28:882–88. [aDR]

Tanner W. P. (1961) Physiological implications of psychophysical data. *Annals of the New York Academy of Sciences* 89:752–65. Available at: http://www.ncbi.nlm.nih.gov/pubmed/13775211. [aDR]

Tauber S., Navarro D. J., Perfors A. & Steyvers M. (2017) Bayesian models of cognition revisited: Setting optimality aside and letting data drive psychological theory. *Psychological Review* 124(4):410–41. [aDR]

Taubert J., Alais D. & Burr D. (2016) Different coding strategies for the perception of stable and changeable facial attributes. *Scientific Reports* 6:32239. doi:10.1038/srep32239. [GMC]

Taylor S. F., Welsh R. C., Wagner T. D., Phan K. L., Fitzgerald K. D. & Gehring W. J. (2004) A functional neuroimaging study of motivation and executive function. *NeuroImage* 21(3):1045–54. Available at: http://www.ncbi.nlm.nih.gov/pubmed/15006672. [aDR]

Tcheang L., Bulthoff H. H. & Burgess N. (2011) Visual influence on path integration in darkness indicates a multimodal representation of large-scale space. *Proceedings of the National Academy of Sciences of the United States of America* 108:1152–57. [MZ]

Tenenbaum J. B. & Griffiths T. L. (2006) Optimal predictions in everyday cognition. *Psychological Science* 17(9):767–73. [aDR]

Tenenbaum J. B., Kemp C., Griffiths T. L. & Goodman N. D. (2011) How to grow a mind: Statistics, structure, and abstraction. *Science* 331(6022):1279–85. Available at: http://www.ncbi.nlm.nih.gov/pubmed/21393536. [aDR]

Teufel C., Subramaniam N., Dobler V., Perez J., Finnemann J., Mehta P. R., Goodyer I. M. & Fletcher P. C. (2015) Shift toward prior knowledge confers a perceptual advantage in early psychosis and psychosis-prone healthy individuals. *Proceedings of the National Academy of Sciences of the United States of America* 112(43):13401–406. Available at: http://doi.org/10.1073/pnas.1503916112. [JS]

Thomas R., Nardini M. & Mareschal D. (2010) Interactions between "light-from-above" and convexity priors in visual development. *Journal of Vision* 10:6. Available at: https://doi.org/10.1167/10.8.6. [MN]

Thompson P. (1982) Perceived rate of movement depends on contrast. *Vision Research* 22(3):377–80. Available at: http://www.ncbi.nlm.nih.gov/pubmed/7090191. [aDR]

Thompson P., Brooks K. & Hammett S. T. (2006) Speed can go up as well as down at low contrast: Implications for models of motion perception. *Vision Research* 46(6–7):782–86. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=16171842&retmode=ref&cmd=prlinks. [aDR]

Thura D., Beauregard-Racine J., Fradet C.-W. & Cisek P. (2012) Decision making by urgency gating: Theory and experimental support. *Journal of Neurophysiology* 108(11):2912–30. Available at: http://www.ncbi.nlm.nih.gov/pubmed/22993260. [aDR]

Todd P. M. & Gigerenzer G. (2003) Bounding rationality to the world. *Journal of Economic Psychology* 24(2):143–65. [AM]

Todorov E. (2004) Optimality principles in sensorimotor control. *Nature Neuroscience* 7:907–915. [YPS]

Tomassini A., Morgan M. J. & Solomon J. A. (2010) Orientation uncertainty reduces perceived obliquity. *Vision Research* 50:541–47. [LSG, AAS]

Treisman M. & Faulkner A. (1984) The setting and maintenance of criteria representing levels of confidence. *Journal of Experimental Psychology: Human Perception and Performance* 10(1):119–39. Available at: http://discovery.ucl.ac.uk/20033/. [aDR]

Trommershäuser J. (2009) Biases and optimality of sensory-motor and cognitive decisions. *Progress in Brain Research* 174:267–78. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=19477345&retmode=ref&cmd=prlinks. [aDR]

Trommershäuser J., Körding K. P. & Landy M. S., eds. (2011) *Sensory cue integration*. Oxford University Press. [aDR]

Tse P. U. (2005) Voluntary attention modulates the brightness of overlapping transparent surfaces. *Vision Research* 45(9):1095–98. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=15707917&retmode=ref&cmd=prlinks. [aDR]

Tsetsos K., Moran R., Moreland J., Chater N., Usher M. & Summerfield C. (2016a) Economic irrationality is optimal during noisy decision making. *Proceedings of the National Academy of Sciences of the United States of America* 113(11):3102–107. Available at: http://www.pnas.org/content/early/2016/02/24/1519157113.long. [RM, aDR, CS]

Tsetsos K., Moran R., Moreland J., Chater N., Usher M. & Summerfield C. (2016b) Reply to Davis-Stober et al.: Violations of rationality in Tsetsos et al. (2016) are not

aggregation artifacts. *Proceedings of the National Academy of Sciences of the United States of America* **113**(33):pE4764–66. [RM]

Tsetsos K., Pfeffer T., Jentgens P. & Donner T. H. (2015) Action planning and the time-scale of evidence accumulation. *PLoS ONE* **10**(6):e0129473. [aDR]

Tsotsos J. K. (1993) The role of computational complexity in perceptual theory. *Advances in Psychology* **99**:261–96. [aDR]

Turatto M., Vescovi M. & Valsecchi M. (2007) Attention makes moving objects be perceived to move faster. *Vision Research* **47**(2):166–78. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=17116314&retmode=ref&cmd=prlinks. [aDR]

Turnbull W. H. (1961) *The correspondence of Isaac Newton. Vol. 3, 1688–1694.* Cambridge University Press. [aDR]

Turner B. M. (under review) Toward a common representational framework for adaptation. [BMT]

Turner B. M., Gao J., Koenig S., Palfy D. & McClelland J. L. (2017) The dynamics of multimodal integration: The averaging diffusion model. *Psychonomic Bulletin and Review* **24**:1819–43. [BMT]

Turner B. M. & Van Zandt T. (2014) Hierarchical approximate Bayesian computation. *Psychometrika* **79**:185–209. [BMT]

Turner B. M., Van Zandt T. & Brown S. (2011) A dynamic, stimulus-driven model of signal detection. *Psychological Review* **118**:583–613. [BMT]

Ulehla Z. J. (1966) Optimality of perceptual decision criteria. *Journal of Experimental Psychology* **71**(4):564–69. Available at: http://www.ncbi.nlm.nih.gov/pubmed/5909083. [aDR]

van Beers R. J., Sittig A. C. & Denier van der Gon J. J. (1999) Integration of proprioceptive and visual position-information: An experimentally supported model. *Journal of Neurophysiology* **81**:1355–64. [J-PN]

van Beers R. J., Sittig A. C. & van der Gon Denier J. J. (1996) How humans combine simultaneous proprioceptive and visual position information. *Experimental Brain Research* **111**(2):253–61. [aDR]

van Bergen R. S., Ji Ma W., Pratte M. S. & Jehee J. F. M. (2015) Sensory uncertainty decoded from visual cortex predicts behavior. *Nature Neuroscience* **18**(12):1728–30. Available at: https://doi.org/10.1038/nn.4150. [MN]

van Boxtel J. J. A., Peng Y., Su J. & Lu H. (2017) Individual differences in high-level biological motion tasks correlate with autistic traits. *Vision Research* **141**:136–44. Available at: http://doi.org/10.1016/j.visres.2016.11.005. [JS]

van den Berg R., Awh E. & Ma W. J. (2014) Factorial comparison of working memory models. *Psychological Review* **121**(1):124–49. [WJM]

van den Berg R., Yoo A. H. & Ma W. J. (2017) Fechner's law in metacognition: A quantitative model of visual working memory confidence. *Psychological Review* **124**(2):197–214. [aDR]

vandormael H., Castañón S. H., Balaguer J., Li V. & Summerfield C. (2017) Robust sampling of decision information during perceptual choice. *Proceedings of the National Academy of Sciences of the United States of America* **114**(10):2771–76. Available at: http://www.pnas.org/lookup/doi/10.1073/pnas.1613950114. [aDR]

van Rooij I. (2008) The tractable cognition thesis. *Cognitive Science* **32**(6):939–84. Available at: http://doi.wiley.com/10.1080/03640210801897856. [aDR]

van Wert M. J., Horowitz T. S. & Wolfe J. M. (2009) *Attention, Perception & Psychophysics* **71**(3):541–53. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2701252&tool=pmcentrez&rendertype=abstract. [aDR]

Varey C. A., Mellers B. A. & Birnbaum M. H. (1990) Judgments of proportions. *Journal of Experimental Psychology: Human Perception and Performance* **16**(3):613–25. Available at: http://www.ncbi.nlm.nih.gov/pubmed/2144575. [aDR]

Vaziri-Pashkam M. & Cavanagh P. (2008) Apparent speed increases at low luminance. *Journal of Vision* **8**(16):9. Available at: http://www.ncbi.nlm.nih.gov/pubmed/19146275. [aDR]

Vehtari A., Gelman A. & Gabry J. (2017) Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing* **27**(5):1413–32. [WJM]

Vickers D. (1979) *Decision processes in visual perception.* Academic Press. [aDR]

Vickers D. & Packer J. (1982) Effects of alternating set for speed or accuracy on response time, accuracy and confidence in a unidimensional discrimination task. *Acta Psychologica* **50**(2):179–97. [aDR]

Viemeister N. F. (1970) Intensity discrimination: Performance in three paradigms. *Perception & Psychophysics* **8**(6):417–19. Available at: http://www.springerlink.com/index/10.3758/BF03207037. [aDR]

Vilares I. & Kording K. (2011) Bayesian models: The structure of the world, uncertainty, behaviour, and then brain. *Annals of the New York Academy of Sciences* **1224**(1):22–39. Available at: https://doi.org/10.1111/j.1749-6632.2011.05965.x. [CC]

Vincent B. (2011) Covert visual search: Prior beliefs are optimally combined with sensory evidence. *Journal of Vision* **11**(13):25. [aDR]

Vintch B. & Gardner J. L. (2014) Cortical correlates of human motion perception biases. *Journal of Neuroscience* **34**(7):2592–604. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=24523549&retmode=ref&cmd=prlinks. [aDR]

Vlassova A., Donkin C. & Pearson J. (2014) Unconscious information changes decision accuracy but not confidence. *Proceedings of the National Academy of Sciences of the United States of America* **111**(45):16214–18. Available at: http://www.pnas.org/content/early/2014/10/24/1403619111.short. [aDR]

von Winterfeldt D. & Edwards W. (1982) Costs and payoffs in perceptual research. *Psychological Bulletin* **91**(3):609–22. [aDR]

Vul E., Goodman N., Griffiths T. L. & Tenenbaum J. B. (2014) One and done? Optimal decisions from very few samples. *Cognitive Science* **38**(4):599–637. Available at: http://www.ncbi.nlm.nih.gov/pubmed/24467492. [aDR]

Wagenmakers E.-J. & Farrell S. (2004) AIC model selection using Akaike weights. *Psychonomic Bulletin & Review* **11**(1):192–96. [WJM]

Wainwright M. J. (1999) Visual adaptation as optimal information transmission. *Vision Research* **39**(23):3960–74. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=10748928&retmode=ref&cmd=prlinks. [aDR]

Wallace S., Coleman M., Pascalis O. & Bailey A. (2006) A study of impaired judgment of eye-gaze direction and related face-processing deficits in autism spectrum disorders. *Perception* **35**(12):1651–64. Available at: http://doi.org/10.1068/p5442. [JS]

Wang J. X., Kurth-Nelson Z., Tirumala D., Soyer H., Leibo J. Z., Munos R., Blundell C., Kumaran D. & Botivnick M. (2017) Learning to reinforcement learn. *ArXiv* 1611.05763. Available at: https://arxiv.org/abs/1611.05763. [MN]

Ward L. M. & Lockhead G. R. (1970) Sequential effects and memory in category judgments. *Journal of Experimental Psychology* **84**(1):27–34. Available at: https://scholars.duke.edu/display/pub651252. [aDR]

Wark B., Lundstrom B. N. & Fairhall A. (2007) Sensory adaptation. *Current Opinion in Neurobiology* **17**(4):423–29. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=17714934&retmode=ref&cmd=prlinks. [aDR]

Warren D. H. & Cleaves W. T. (1971) Visual-proprioceptive interaction under large amounts of conflict. *Journal of Experimental Psychology* **90**(2):206–14. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=5134326&retmode=ref&cmd=prlinks. [aDR]

Watanabe M., Cromwell H. C., Tremblay L., Hollerman J. R., Hikosaka K. & Schultz W. (2001) Behavioral reactions reflecting differential reward expectations in monkeys. *Experimental Brain Research* **140**:511–18. [ES]

Watson C. S., Kellogg S. C., Kawanishi D. T. & Lucas P. A. (1973) The uncertain response in detection-oriented psychophysics. *Journal of Experimental Psychology* **99**(2):180–85. [aDR]

Webster M. A. (2015) Visual adaptation. *Annual Review of Vision Science* **1**:547–67. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=26858985&retmode=ref&cmd=prlinks. [aDR]

Webster M. A., Kaping D., Mizokami Y. & Duhamel P. (2004) Adaptation to natural facial categories. *Nature* **428**(6982):557–61. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=15058304&retmode=ref&cmd= prlinks. [aDR]

Webster M. A. & MacLeod D. I. A. (2011) Visual adaptation and face perception. *Philosophical Transactions of the Royal Society B: Biological Sciences* **366** (1571):1702–25. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=21536555&retmode=ref&cmd=prlinks. [aDR]

Wei K. & Körding K. P. (2011) Causal inference in sensorimotor learning and control. In: *Sensory cue integration,* ed. J. Trommershäuser, K. Körding & M. S. Landy, pp. 30–45. Oxford University Press. [aDR]

Wei X.-X. & Stocker A. A. (2013) Efficient coding provides a direct link between prior and likelihood in perceptual Bayesian inference. In: *Advances in neural information processing systems 25* (proceedings from the conference, Neural Information Processing Systems 2012), ed. F. Pereira, C. J. C. Burges, L. Bottou & K. Q. Weinberger. Available at: https://papers.nips.cc/paper/4489-efficient-coding-provides-a-direct-link-between-prior-and-likelihood-in-perceptual-bayesian-inference. [aDR]

Wei X.-X. & Stocker A. A. (2015) A Bayesian observer model constrained by efficient coding can explain "anti-Bayesian" percepts. *Nature Neuroscience* **18**:1509–17. Available at: http://dx.doi.org/10.1038/nn.4105. [aDR, AAS, CS, VW]

Wei X. X. & Stocker A. A. (2017) Lawful relation between perceptual bias and discriminability. *Proceedings of the National Academy of Sciences of the United States of America* **114**(38):10244–49. [AAS, CS]

Weil L. G., Fleming S. M., Dumontheil I., Kilford E. J., Weil R. S., Rees G., Dolan R. J., Blakemore S.-J. (2013) The development of metacognitive ability in adolescence. *Consciousness and Cognition* **22**(1):264–71. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3719211&tool=pmcentrez&rendertype=abstract. [aDR]

Weiskrantz L. (1996) Blindsight revisited. *Current Opinion in Neurobiology* **6**(2):215–20. Available at: http://www.ncbi.nlm.nih.gov/pubmed/8725963. [aDR]

Weiss Y., Simoncelli E. P. & Adelson E. H. (2002) Motion illusions as optimal percepts. *Nature Neuroscience* **5**(6):598–604. Available at: http://www.nature.com/neuro/journal/v5/n6/full/nn858.html. [aDR]

White C., Ratcliff R., Vasey M. & McKoon G. (2009) Dysphoria and memory for emotional material: A diffusion-model analysis. *Cognition & Emotion* **23**(1), 181–205. Available at: http://doi.org/10.1080/02699930801976770. [JS]

Whiteley L. & Sahani M. (2008) Implicit knowledge of visual uncertainty guides decisions with asymmetric outcomes. *Journal of Vision* **8**(3):2.1–15. Available at: http://www.journalofvision.org/content/8/3/2. [aDR]

Whiteley L. & Sahani M. (2012) Attention in a Bayesian framework. *Frontiers in Human Neuroscience* 6:100. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=22712010&retmode=ref&cmd=prlinks%5Cnpapers3://publication/doi/10.3389/fnhum.2012.00100. [aDR]

Wiecki T. V., Poland J. & Frank M. J. (2014) Model-based cognitive neuroscience approaches to computational psychiatry. *Clinical Psychological Science* 3(3):378–99. Available at: http://doi.org/10.1177/2167702614565359. [JS]

Wilimzig C., Tsuchiya N., Fahle M., Einhäuser W. & Koch C. (2008) Spatial attention increases performance but not subjective confidence in a discrimination task. *Journal of Vision* 8(5):1–10. Available at: http://www.ncbi.nlm.nih.gov/pubmed/18842078. [aDR]

Wilmer J. B., Germine L., Chabris C. F., Chatterjee G., Gerbasi M. & Nakayama K. (2012) Capturing specific abilities as a window into human individuality: The example of face recognition. *Cognitive Neuropsychology* 29(5–6):360–92. Available at: http://doi.org/10.1080/02643294.2012.753433. [JS]

Winman A. & Juslin P. (1993) Calibration of sensory and cognitive judgments: Two different accounts. *Scandinavian Journal of Psychology* 34(2):135–48. Available at: http://doi.wiley.com/10.1111/j.1467-9450.1993.tb01109.x. [aDR]

Withagen R. (2004) The pickup of nonspecifying variables does not entail indirect perception. *Ecological Psychology* 16:237–53. [RW]

Withagen R. & Chemero A. (2009) Naturalizing perception: Developing the Gibsonian approach to perception along evolutionary lines. *Theory & Psychology* 19:363–89. [RW]

Witt J. K. (2011) Action's effect on perception. *Current Directions in Psychological Science* 20(3):201–206. Available at: http://cdp.sagepub.com/content/20/3/201.short. [aDR]

Witt J. K., Proffitt D. R. & Epstein W. (2005) Tool use affects perceived distance, but only when you intend to use it. *Journal of Experimental Psychology: Human Perception and Performance* 31(5):880–88. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=16262485&retmode=ref&cmd=prlinks. [aDR]

Wohlgemuth A. (1911) *On the after-effect of seen movement.* Cambridge University Press. Available at: https://books.google.com/books?id=Z6AhAQAAIAAJ. [aDR]

Wolfe J. M., Brunelli D. N., Rubinstein J. & Horowitz T. S. (2013) Prevalence effects in newly trained airport checkpoint screeners: Trained observers miss rare targets, too. *Journal of Vision* 13(3):33. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3848386&tool=pmcentrez&rendertype=abstract. [aDR]

Wolfe J. M., Horowitz T. S. & Kenner N. M. (2005) Cognitive psychology: Rare items often missed in visual searches. *Nature* 435(7041):439–40. Available at: http://dx.doi.org/10.1038/435439a. [aDR]

Wolfe J. M., Horowitz T. S., Van Wert M. J., Kenner N. M., Place S. S. & Kibbi N. (2007) Low target prevalence is a stubborn source of errors in visual search tasks. *Journal of Experimental Psychology: General* 136(4):623–38. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2662480&tool=pmcentrez&rendertype=abstract. [aDR]

Wolfe J. M. & Van Wert M. J. (2010) Varying target prevalence reveals two dissociable decision criteria in visual search. *Current Biology* 20(2):121–24. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2818748&tool=pmcentrez&rendertype=abstract. [aDR]

Wu S.-W., Delgado M. R. & Maloney L. T. (2011) The neural correlates of subjective utility of monetary outcome and probability weight in economic and in motor decision under risk. *Journal of Neuroscience* 31(24): 8822–31. Available at: https://doi.org/10.1523/JNEUROSCI.0540-11.2011. [MN]

Wyart V. & Koechlin E. (2016) Choice variability and suboptimality in uncertain environments. *Current Opinion in Behavioral Sciences* 11:109–15. Available at: http://dx.doi.org/10.1016/j.cobeha.2016.07.003. [aDR, VW]

Wyart V., Myers N. E. & Summerfield C. (2015) Neural mechanisms of human perceptual choice under focused and divided attention. *Journal of Neuroscience* 35(8):3485–98. Available at: http://www.jneurosci.org/content/35/8/3485.abstract?etoc. [aDR]

Xu Y., Regier T. & Newcombe N. S. (2017) An adaptive cue combination model of human spatial reorientation. *Cognition* 163:56–66. [MZ]

Yamins D. L. K., Hong H., Cadieu C. F., Solomon E. A., Seibert D. & DiCarlo J. J. (2014) Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences of the United States of America* 111(23):8619–24. Available at: http://www.ncbi.nlm.nih.gov/pubmed/24812127. [aDR]

Yasuda M. & Hikosaka O. (2017) To wait or not to wait—separate mechanisms in the oculomotor circuit of basal ganglia. *Frontiers in Neuroanatomy* 11:35. [ES]

Yeshurun Y. & Carrasco M. (1998) Attention improves or impairs visual performance by enhancing spatial resolution. *Nature* 396(6706):72–75. Available at: http://eutils.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=9817201&retmode=ref&cmd=prlinks. [aDR]

Yeshurun Y., Carrasco M. & Maloney L. T. (2008) Bias and sensitivity in two-interval forced choice procedures: Tests of the difference model. *Vision Research* 48(17):1837–51. Available at: http://www.sciencedirect.com/science/article/pii/S0042698908002599. [aDR]

Yeung N. & Summerfield C. (2012) Metacognition in human decision-making: Confidence and error monitoring. *Philosophical Transactions of the Royal Society of London: Series B, Biological Sciences* 367(1594):1310–21. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3318764&tool=pmcentrez&rendertype= abstract. [aDR]

Yoder R. M., Clark B. J. & Taube J. S. (2011) Origins of landmark encoding in the brain. *Trends in Neurosciences* 34(11):561–71. [MZ]

Yu A. J. & Cohen J. D. (2009) Sequential effects: Superstition or rational behavior? In: *Advances in neural information processing systems 21* (proceedings from the conference, Neural Information Processing Systems 2008), ed. D. Koller, D. Schuurmans, Y. Bengio & L. Bottou. Available at: https://papers.nips.cc/book/advances-in-neural-information-processing-systems-21-2008. [aDR]

Zacksenhouse M., Bogacz R. & Holmes P. (2010) Robust versus optimal strategies for two-alternative forced choice tasks. *Journal of Mathematical Psychology* 54(2):230–46. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3505075&tool=pmcentrez&rendertype=abstract. [aDR]

Zak I., Katkov M., Gorea A. & Sagi D. (2012) Decision criteria in dual discrimination tasks estimated using external-noise methods. *Attention, Perception & Psychophysics* 74(5):1042–55. Available at: http://www.ncbi.nlm.nih.gov/pubmed/22351481. [aDR]

Zamboni E., Ledgeway T., McGraw P. V. & Schluppeck D. (2016) Do perceptual biases emerge early or late in visual processing? Decision-biases in motion perception. *Proceedings of the Royal Society B: Biological Sciences* 283(1833):20160263. Available at: http://rspb.royalsocietypublishing.org/content/283/1833/20160263. [aDR]

Zednik C. (2017) Mechanisms in cognitive science. In: *The Routledge handbook of mechanisms and mechanical philosophy*, ed. S. Glennan & P. Illari, pp. 389–400. Routledge. [CZ]

Zednik C. & Jäkel F. (2014) How does Bayesian reverse-engineering work? In: *Proceedings of the 36th Annual Conference of the Cognitive Science Society*, ed. P. Bello, M. Guarini, M. McShane & B. Scassellati, pp. 666–71. Cognitive Science Society. [CZ]

Zednik C. & Jäkel F. (2016) Bayesian reverse-engineering considered as a research strategy for cognitive science. *Synthese* 193:3951–85. [CZ]

Zhang H. & Maloney L. T. (2012) Ubiquitous log odds: A common representation of probability and frequency distortion in perception, action, and cognition. *Frontiers in Neuroscience* 6:1. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3261445&tool=pmcentrez&rendertype=abstract. [aDR]

Zhang H., Morvan C. & Maloney L. T. (2010) Gambling in the visual periphery: A conjoint-measurement analysis of human ability to judge visual uncertainty. *PLoS Computational Biology* 6(12):1001023. Available at: http://dx.plos.org/10.1371/journal.pcbi.1001023. [aDR]

Zhang L. & Mou W. (2017) Piloting systems reset path integration systems during position estimation. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 43:472–91. [MZ]

Zhao M. & Warren W. H. (2015a) Environmental stability modulates the role of path integration in human navigation. *Cognition* 142:96–109. [MZ]

Zhao M. & Warren W. H. (2015b) How you get there from here: Interaction of visual landmarks and path integration in human navigation. *Psychological Science* 26:915–24. [MZ]

Zylberberg A., Barttfeld P. & Sigman M. (2012) The construction of confidence in a perceptual decision. *Frontiers in Integrative Neuroscience* 6:79. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3448113&tool=pmcentrez&rendertype=abstract. [aDR]

Zylberberg A., Roelfsema P. R. & Sigman M. (2014) Variance misperception explains illusions of confidence in simple perceptual decisions. *Consciousness and Cognition* 27:246–53. Available at: http://www.sciencedirect.com/science/article/pii/S1053810014000865. [aDR]