

A Detailed Comparison of Optimality and Simplicity in Perceptual Decision Making

Shan Shen
Baylor College of Medicine

Wei Ji Ma
New York University

Two prominent ideas in the study of decision making have been that organisms behave near-optimally, and that they use simple heuristic rules. These principles might be operating in different types of tasks, but this possibility cannot be fully investigated without a direct, rigorous comparison within a single task. Such a comparison was lacking in most previous studies, because (a) the optimal decision rule was simple, (b) no simple suboptimal rules were considered, (c) it was unclear what was optimal, or (d) a simple rule could closely approximate the optimal rule. Here, we used a perceptual decision-making task in which the optimal decision rule is well-defined and complex, and makes qualitatively distinct predictions from many simple suboptimal rules. We find that all simple rules tested fail to describe human behavior, that the optimal rule accounts well for the data, and that several complex suboptimal rules are indistinguishable from the optimal one. Moreover, we found evidence that the optimal model is close to the true model: First, the better the trial-to-trial predictions of a suboptimal model agree with those of the optimal model, the better that suboptimal model fits; second, our estimate of the Kullback–Leibler divergence between the optimal model and the true model is not significantly different from zero. When observers receive no feedback, the optimal model still describes behavior best, suggesting that sensory uncertainty is implicitly represented and taken into account. Beyond the task and models studied here, our results have implications for best practices of model comparison.

Keywords: optimality, perception, visual search, ideal observer, model comparison

Many forms of human perception seem close to the ideal set by Bayesian optimality (Geisler, 2011; Körding et al., 2007), according to which the brain maximizes performance when given noisy and ambiguous sensory input. In most of these cases, such as in cue combination (e.g., Alais & Burr, 2004; Ernst & Banks, 2002; Gu, Angelaki, & Deangelis, 2008; for more examples, see Trommershäuser, Körding, & Landy, 2011), the optimal decision rule is simple. It has been argued that when the optimal decision rule is complex, the brain has an incentive to use a computationally simple and reasonably effective, though strictly suboptimal, decision rule—also called a heuristic (Gigerenzer & Gaissmaier, 2011; Simon, 1956).

Evidence for this proposal has been mixed. On the one hand, there is no strong evidence for optimality in complex tasks. In some studies that claim that people follow a complex optimal rule,

simple suboptimal rules were not considered (e.g., Geisler & Perry, 2009; see also Bowers & Davis, 2012). In other studies, the optimal rule fitted about equally well as a simple rule (Mazyar, van den Berg, Seilheimer, & Ma, 2013; Palmer, Verghese, & Pavel, 2000; Qamar et al., 2013). In some perceptual experiments, the optimal rule outperformed simple rules, but only a few simple rules were tested (Ma, Navalpakkam, Beck, van den Berg, & Pouget, 2011; Ma, Shen, Dziugaite, & van den Berg, 2015). Finally, it has been argued that in reports of near-optimality in cognitive tasks (Chater, Tenenbaum, & Yuille, 2006; Norris, 2006), the optimal model was given excessive flexibility in order to fit the data (Bowers & Davis, 2012; Gigerenzer, 2004; Jones & Love, 2011), making the optimality label suspect.

On the other hand, some claims of suboptimality might be premature. For ball catching, a suboptimal but simple “gaze heuristic” has been proposed (Dienes & McLeod, 1996; McLeod, Reed, & Dienes, 2003). However, in such complex sensorimotor tasks, the definition of optimality depends on largely unknown costs, making it difficult to conclusively claim that behavior is suboptimal; in fact, in a simplified setting, ball catching was found to be near-optimal (Faisal & Wolpert, 2009; López-Moliner, Field, & Wann, 2007).

Thus, optimality and simplicity have, so far, not been directly compared in a strongly distinguishing paradigm. In addition, none of the studies mentioned in the previous two paragraphs made an effort to establish how much room there is for an untested model to fit better than the study’s favored model, whether optimal or simple. To address both issues, we used a visual categorization task that did not suffer from the shortcomings we mentioned in the

Shan Shen, Department of Neuroscience, Baylor College of Medicine; Wei Ji Ma, Center for Neural Science and Department of Psychology, New York University.

We thank Sebastiaan van Opheusden for suggesting the analysis in the “How good are the best models?” section. Wei Ji Ma is supported by award number R01EY020958 from the National Eye Institute, and award number W911NF-12-1-0262 from the Army Research Office.

Correspondence concerning this article should be addressed to Wei Ji Ma, Center for Neural Science and Department of Psychology, 4 Washington Place, New York University, New York, NY 10003. E-mail: weijima@nyu.edu

previous two paragraphs. Optimality was not ambiguous, and there are many plausible simple rules that make substantially different predictions from the optimal rule. Moreover, in analyzing the data from this task, we attempt to establish how close our best-fitting model is to the unknown true model.

We want to emphasize that in this article, *simplicity* refers to the number of operations in the observer's decision rule. In previous literature, simplicity has been used in at least two other meanings. The first is the simplicity of the observer's interpretation of a visual scene. The "simplicity principle," initiated by Wertheimer and other Gestalt psychologists, states that the observer reports the simplest interpretation of a visual scene that is consistent with the sensory input, for example, leading to a preference for perceiving continuity (Chater, 1996). In our task, the two hypotheses about the state of the world, left and right, are equally complex in this sense. The second meaning is simplicity in terms of the number of parameters (Spiegelhalter, Best, Carlin, & Van Der Linde, 2002). In our case, all the models we will discuss below have two parameters, and the only thing that differs is the observer's decision rule. Thus, we are not talking about simplicity from the point of view of the experimenter comparing models, but about simplicity of a decision strategy for the observer doing the task.

Experiment 1

We designed our experiment such that the statistical structure of the task was easy for subjects to learn, but optimal inference was hard. The task was a visual target categorization task on four oriented stimuli. On each trial, we presented observers with a target and three identical distractors (Figure 1A); the target orientation and the (common) distractor orientation were drawn independently from the same Gaussian distribution with a mean of vertical and a standard deviation of $\sigma_s = 9.06^\circ$ (Figure 1B). Subjects reported whether the target stimulus (the unique stimulus)

was tilted to the right or to the left of vertical. Correctness feedback was provided on each trial.

Apparatus and Stimuli

Subjects were seated at a viewing distance of approximately 60 cm. All stimuli were displayed on a 21-in. LCD monitor with a refresh rate of 60 Hz and a resolution of $1,280 \times 1024$ pixels. The background luminance was approximately 29.3 cd/m^2 . Each stimulus display contained four stimuli, placed on an invisible circle centered at the center of the screen and with a radius of 5° of visual angle. The angular positions of the stimuli were -135° , -45° , 45° , and 135° relative to the positive horizontal axis. Each stimulus was a Gabor patch with a peak luminance of approximately 35.2 cd/m^2 , a spatial frequency of 3.13 cycles per degree, a standard deviation of 8.18 pixels, and a phase of 0. On each trial, three of the stimuli were identical; these were distractors. The fourth stimulus, whose location was chosen from the four possible locations with equal probabilities, was the target. Target and distractor orientations were independently drawn from a Gaussian distribution with a mean of 0° (vertical) and a standard deviation of 9.06° .

Experimental Procedures

Each trial started with a fixation dot on a blank screen (500 ms), followed by a stimulus display (50 ms). Then, a blank screen was shown until the subject made a response. Subjects pressed a button to report whether the target was tilted to the right or to the left relative to vertical. After the response, correctness feedback was given by coloring the fixation dot red or green (500 ms; Figure 1A).

The experiment consisted of three sessions on different days. Each session consisted of five blocks, and each block contained

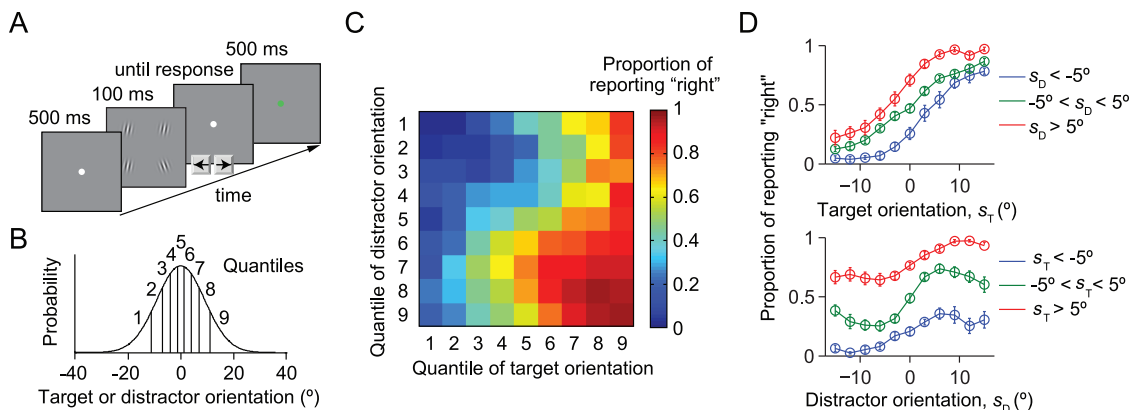


Figure 1. Task and data. (A) Trial procedure. Each display contains four items, three of which have a common orientation; these are the distractors. Subjects report whether the fourth item (the target) is tilted to the left or to the right with respect to vertical. The target location is randomly chosen on every trial. (B) The target orientation and the common distractor orientation are independently drawn from the same Gaussian distribution with a mean of 0° (vertical) and a standard deviation of 9.06° . For plotting purposes, we divided orientation space into 9 quantiles. (C) Proportion of reporting "right" (color, or gray scale) as a function of target and distractor orientation quantiles. (D) Proportion of reporting "right" as a function of target orientation s_T (top) and distractor orientation s_D (bottom). Error bars are standard errors of the mean. The bottom curves are not expected to be monotonic (see text). See the online article for the color version of this figure.

200 trials, for a total of $3 \times 5 \times 200 = 3,000$ trials per subject. To avoid the learning effect, data from the first session (1,000 trials) were excluded in the analysis. Nine subjects participated in this study (seven female; age range = 23 to 30 years), all of them scientifically trained but naïve to the task. At the beginning of the first session, subjects were orally briefed about all details of the experimental design. We showed them 10 samples drawn from the Gaussian stimulus distribution, along with a graph of that distribution. We explained the distribution as follows: “The target orientation and the common distractor orientation are drawn from the same bell-shaped distribution; the orientations are most often near-vertical, and the width of the distribution is about 10° .”

Experiment 1 Data

Human behavior in this task exhibits several interesting patterns (Figure 1C). The proportion of “right” responses increases monotonically with target orientation (Figure 1D, top), but also depends on distractor orientation. More surprising is that the proportion of “right” responses does not increase monotonically with distractor orientation (Figure 1D, bottom). On second thought, this is intuitive: Suppose the target is vertical and the distractors are slightly tilted to the right. Then, the distractors can easily be confused with the target, and therefore the subject will often report “right.” By contrast, when the distractors are strongly tilted to the right, they are easily identified as distractors and a sophisticated subject will only use the item that is most likely to be the target in their response (in this example, make a random guess).

Models

We study four categories of models: the optimal model, simple heuristic models, two-step models, and generalized sum models.

The optimal model. The optimal observer has learned the generative model of the task and incorporates this knowledge during decision making. We first discuss the generative model (Figure 2A). We denote the binary variable, C , the direction of tilt relative to vertical (-1 for left, $+1$ for right). A distractor orientation s_D is drawn from a Gaussian distribution with a mean of 0 and a standard deviation of σ_s . A target orientation s_T is drawn from that same Gaussian distribution but truncated to only negative values (when $C = -1$) or only positive values (when $C = 1$); thus, the conditional distribution of s_T is a half-Gaussian. Stimuli appear at four fixed locations. The target location, L , is chosen with equal probability for each of the four possibilities. The orientation at that location is s_T , and the orientations at the three other locations are all s_D . Finally, we assume that the observer makes a noisy measurement x_i of each orientation. We assume that each x_i is independently drawn from a Gaussian distribution whose mean is the true stimulus (s_T or s_D) and whose standard deviation is σ .

We are now ready to describe the optimal observer’s inference process. Given a set of measurements $\mathbf{x} = (x_1, x_2, x_3, x_4)$, the optimal observer considers each item a potential target and computes the weighted average over locations L of the probability that an assumed target at L was tilted right, with weight given by the probability that the target was at L . The posterior over C ($C = 1$ or -1) is

$$p(C|\mathbf{x}) = \sum_{L=1}^N p(C|\mathbf{x}, L)p(L|\mathbf{x}), \quad (1)$$

where N denotes the set size (in our case, $N = 4$). The weighted average over L is an example of the Bayesian operation of marginalization. We can simplify Equation 1 to

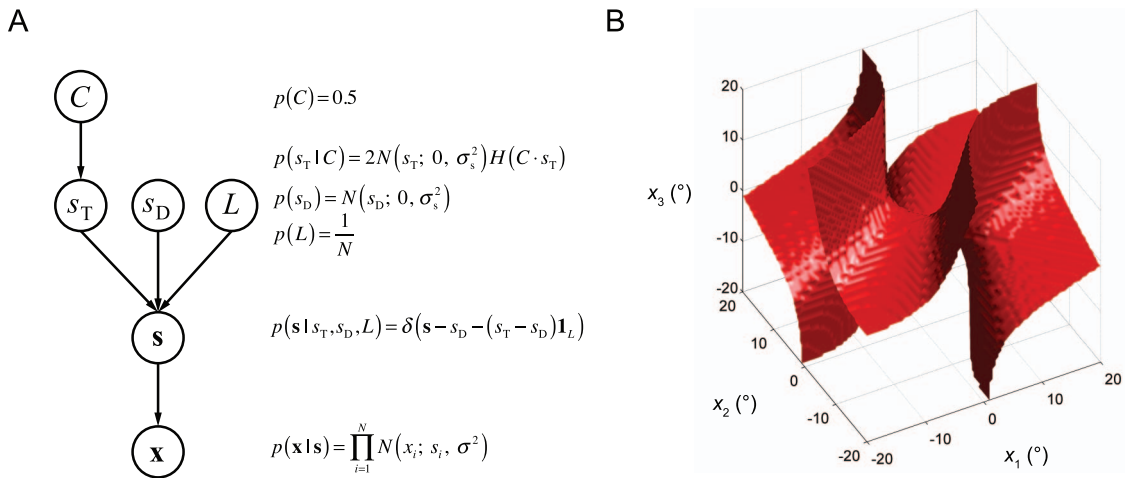


Figure 2. Generative model. (A) Each node represents a random variable, each arrow a conditional probability distribution. Distributions are shown in the equations on the side. $N(x; 0, \sigma^2)$ denotes a normal distribution with a mean of 0 and a variance of σ^2 . $H(x)$ denotes the Heaviside function. $\mathbf{1}_L$ denotes a vector in which the L^{th} entry equals 1 and all others equal 0. $\delta(x)$ is the Dirac delta function. This diagram specifies the distribution of the measurements, \mathbf{x} . The optimal observer inverts the generative model and computes the conditional probability of C given \mathbf{x} . (B) Decision boundary of the optimal decision rule if the set size N were equal to 3. Each point in the three-dimensional space represents a set of measurements $\mathbf{x} = (x_1, x_2, x_3)$. On one side of the boundary (the side that includes the all-positive octant), the optimal observer reports “right,” on the other side, “left.” See the online article for the color version of this figure.

$$\begin{aligned}
p(C|\mathbf{x}) &= \sum_{L=1}^N \frac{p(\mathbf{x}|C, L)}{p(\mathbf{x}|L)} p(C) p(\mathbf{x}|L) \frac{p(L)}{p(\mathbf{x})} \\
&= \frac{p(C)}{p(\mathbf{x})} \sum_{L=1}^N p(\mathbf{x}|C, L) p(L).
\end{aligned} \tag{2}$$

Given that the priors over C and L are uniform, and $p(\mathbf{x})$ is just a normalization, we only need to compute $p(\mathbf{x}|C, L)$, the likelihoods of $C = \pm 1$ given L . To compute these, the optimal observer marginalizes over all unknown variables other than C and L . These variables are the stimulus orientations, \mathbf{s} , the values of target and distractor orientations, s_T and s_D :

$$\begin{aligned}
p(\mathbf{x}|C, L) &= \int p(\mathbf{x}|\mathbf{s}) p(\mathbf{s}|C, L) ds \\
&= \int \int \int p(\mathbf{x}|\mathbf{s}) p(\mathbf{s}|s_T, s_D, L) p(s_T, s_D|C) ds ds_T ds_D \\
&= \int \int \int \left(\prod_{i=1}^N p(x_i|s_i) \right) \delta(\mathbf{s} - s_D - (s_T - s_D)\mathbf{1}_L) \\
&\quad \times p(s_T|C) p(s_D) ds ds_T ds_D \\
&= \left(\int p(x_L|s_T) p(s_T|C) ds_T \right) \left(\int \left(\prod_{i \neq L} p(x_i|s_D) \right) p(s_D) ds_D \right),
\end{aligned} \tag{3}$$

where $\mathbf{1}_L$ is a vector of zeroes except for a 1 in the L^{th} dimension and $\delta(\mathbf{x})$ is Dirac delta function.

The optimal observer reports ‘‘right’’ when the posterior probability of a right tilt, $p(C = 1|\mathbf{x})$, exceeds the posterior probability of a left tilt, $p(C = -1|\mathbf{x})$. Starting from Equation 2, we see that this is the case when

$$\sum_{L=1}^N p(\mathbf{x}|C = 1, L) > \sum_{L=1}^N p(\mathbf{x}|C = -1, L). \tag{4}$$

Evaluating Equation 3 and then substituting in Equation 4, we find that the optimal decision rule (Opt rule) is to report ‘‘right’’ when

$$\sum_{L=1}^N e^{\text{Weight}_L(\mathbf{x})} \text{erf} \frac{\frac{x_L}{\sigma^2}}{\sqrt{2\left(\frac{1}{\sigma^2} + \frac{1}{\sigma_s^2}\right)}} > 0, \tag{5}$$

where $\text{Weight}_L(\mathbf{x})$ reflects the strength of the evidence that the target is at L :

$$\text{Weight}_L(\mathbf{x}) = -\frac{x_L^2}{2(\sigma_s^2 + \sigma^2)} - \frac{\bar{x}_{\setminus L}^2}{2\left(\sigma_s^2 + \frac{\sigma^2}{N-1}\right)} - \frac{N-1}{2\sigma^2} \text{Var } \mathbf{x}_{\setminus L}. \tag{6}$$

Here, $\bar{x}_{\setminus L}$ and $\text{Var } \mathbf{x}_{\setminus L}$ are the sample mean and variance of \mathbf{x} with the L^{th} element left out, respectively. The first term in the expression for $\text{Weight}_L(\mathbf{x})$ comes from $\int p(x_L|s_T) p(s_T|C) ds_T$, and the other two terms come from $\int \left(\prod_{i \neq L} p(x_i|s_D) \right) p(s_D) ds_D$ in Equation 3. To aid our intuition, we visualize the optimal decision rule for the hypothetical scenario that set size $N = 3$, in Figure 2B.

Each of the components of the weight term has an intuitive meaning: They can be interpreted as being associated with the target, the distractor mean, and the distractor variance, respectively. Because both the target orientation and the common distractor orientation are drawn from a Gaussian distribution with a

mean of 0, the L^{th} item is more likely to be the target when the L^{th} measurement is closer to 0 (target term), and the mean of the measurements at the other locations is closer to 0 (distractor mean term). Moreover, the subject knows that the three distractor orientations are the same, so an item is more likely to be the target when the variance of the measurements at the other three locations is smaller (distractor variance term). In the following, we will examine whether human subjects take all these aspects of the task statistics into account.

Simple heuristic models. Several simple heuristic decision rules are plausible for this task. These rules, some of which have been widely used in previous studies, postulate how the observer integrates information across locations. According to the (signed) Max rule, the observer reports the direction of tilt of the measurement that is most tilted in either direction (Baldassi & Verghese, 2002; Eckstein, 1998; Green & Swets, 1966; Nolte & Jaarsma, 1967; Palmer et al., 2000). According to the Sum rule, the observer reports the direction of tilt of the sum of the four measurements (Baldassi & Burr, 2000; Green & Swets, 1966; Kramer, Graham, & Yager, 1985; Palmer et al., 2000). We also conceived three new simple heuristics. According to the Min rule, the observer reports the direction of tilt of the measurement that is *least* tilted in either direction. According to the Var rule, the observer reports the direction of tilt of the measurement, which, when left out, leaves the smallest variance of the remaining three measurements. Finally, according to the Sign rule, the observer reports the common direction of tilt when all measurements are tilted in the same direction, and the least frequent direction of tilt otherwise (guessing in case of a tie). All of these rules perform above chance, but they make very different predictions from the Opt rule.

Two-step models. We then considered ‘‘two-step models,’’ in which the observer does not marginalize over target location but instead follows a simple strategy of first deciding which item is the target, then reporting its direction of tilt. This suboptimal decision process is similar to the Take The Best heuristic algorithm (Gigerenzer & Goldstein, 1996), in which decisions are based solely on the cue that best discriminates between options, and to some other two-step models in perception (Fleming, Maloney, & Daw, 2013; Jazayeri & Movshon, 2007).

Different two-step models can be constructed based on how the target is selected in the first step. One way is to transform the Opt model into a two-step model by replacing the (optimal) marginalization over L in Equation 1 with (suboptimal) maximization: The first step is to infer the target location, \hat{L} , and the second step is to infer the target tilt, \hat{C} , given that location:

$$\begin{aligned}
\hat{L} &= \underset{L}{\text{argmax}} p(L|\mathbf{x}), \\
\hat{C} &= \underset{C}{\text{argmax}} p(C|\mathbf{x}, \hat{L}).
\end{aligned} \tag{7}$$

When we evaluate the first line of Equation 7 and make use of the fact that $p(L)$ is uniform, we get

$$\begin{aligned}
\hat{L} &= \underset{L}{\text{argmax}} p(\mathbf{x}|L) \frac{p(L)}{p(\mathbf{x})} \\
&= \underset{L}{\text{argmax}} p(\mathbf{x}|L) \\
&= \underset{L}{\text{argmax}} \left(\int \left(\prod_{i \neq L} p(x_i|s_D) \right) p(s_D) ds_D \right).
\end{aligned} \tag{8}$$

Retracing the derivation from Equation 3 to Equation 5 shows that finding \hat{L} amounts to maximizing $\text{Weight}_L(\mathbf{x})$ in Equation 6:

$$\hat{L} = \underset{L}{\operatorname{argmax}} \text{Weight}_L(\mathbf{x}). \quad (9)$$

The second step is inferring the target tilt if we assume that the target location is \hat{L} :

$$\begin{aligned} \hat{C} &= \underset{C}{\operatorname{argmax}} \frac{p(x|C, \hat{L})}{p(x|\hat{L})} p(C) \\ &= \underset{C}{\operatorname{argmax}} p(x|C, \hat{L}) \\ &= \underset{C}{\operatorname{argmax}} \int p(x_L | s_T) p(s_T | C) ds_T \\ &= \operatorname{sgn} x_L. \end{aligned}$$

The simplicity of this expression is easily understood: If the observer assumes that the target is in the L^{th} location, then only the measurement at that location should be used for the report of target tilt.

Whereas Equation 9 uses the entire $\text{Weight}_L(\mathbf{x})$ expression to decide on a target location, we will also consider the possibility that the target location is decided based on maximizing a subset of the terms in that expression. This would correspond to taking into account only part of the statistics of the task. In total, we end up with seven two-step models. Two of these we discussed already: the Min model (target term only) and the Var model (distractor variance term only). The remaining five models we name MaxT2, MaxT12, MaxT13, MaxT23, and MaxT123 (T refers to *term*), for which the numbers refer to the terms kept in the expression.

Generalized sum models. As a third category of suboptimal models, we considered “generalized sum models”, in which the observer does marginalize over potential target locations. We construct this category by systematically perturbing the optimal decision rule. The perturbation consisted of leaving out a subset of the terms in the weight term (Equation 6), for a total of seven models (excluding the null perturbation), which we refer to as SumErf models. We use the nomenclature SumErfT2, SumErfT12 etc. in parallel to the two-step models, except that SumErf without any suffixes is the model in which the weight term is replaced by 0.

A further approximation, $\operatorname{erf}(x) \approx x$, results in a total of eight new models (eight because the null perturbation with this approximation is a new model), which we call SumX models. For example, SumXT12 replaces “erf” in Equation 5 by the identity and keeps only the first and second terms in Equation 6. Since the SumX model in which none of the terms in Equation 6 is kept is identical to our original Sum model, the SumX category contains seven new models.

Modeling Methods

Quantifying decision rule complexity. To quantify the computational complexity of the decision rules, we first counted the number of arithmetic operations. We treated functions of the measurements \mathbf{x} as arguments, as they change from trial to trial. The other parameters (σ and σ_s) were taken as constants. We defined four types of operations:

1. Linear operations: addition and subtraction between arguments and constants, or between arguments; multiplication of an argument by a constant.
2. Quadratic operations: multiplication of arguments.

3. Nonlinear operations: nonlinear operations (exponentials, reciprocals, square root, error functions) of arguments.
4. Sorting operations: taking the maximum or minimum, or comparing with zero.

We combined products of exponentials of polynomials in \mathbf{x} into a single exponential of a polynomial, and simplified each polynomial expression to have no parentheses. This definition of complexity is designed only for comparing observer decision rules in a perceptual task.

To evaluate decision rule complexity in a more biological manner, we used a neural encoding model. The observations on a given trial do not consist of a scalar measurement, x , and noise level, σ , but of a vector of spike counts of a group of orientation-tuned neurons, \mathbf{r} . For spike count variability belonging to the “Poisson-like” family of distributions (exponential family with linear sufficient statistics), x and σ^2 can be identified with $\frac{\mathbf{a} \cdot \mathbf{r}}{\mathbf{b} \cdot \mathbf{r}}$ and $\frac{1}{\mathbf{b} \cdot \mathbf{r}}$, respectively, where \mathbf{a} and \mathbf{b} are constant vectors (Ma, 2010). Then, we expressed the weight term into a rational polynomial function and counted the number of occurrences of $\mathbf{a} \cdot \mathbf{r}$ and $\mathbf{b} \cdot \mathbf{r}$. We counted the other operations in the same way as above.

The results of both quantifications are shown in Table A1 of the Appendix. According to either measure, the optimal decision rule is much more complex than any of these simple heuristic rules, and more complex than some of the two-step rules. The generalized sum models tend to be similarly complex to the Opt model. We realize that our ways of quantifying decision rule complexity are somewhat arbitrary, but this problem is shared by any study that claims that people use “simple” decision rules.

Model predictions. Each model had the same two parameters: sensory precision, $J = \frac{1}{\sigma^2}$ (where σ is the standard deviation of the sensory noise), and lapse rate, λ . We tested combinations of parameter values on grid. The grid for J consisted of 31 equally spaced values between 0.001 and 0.3. The grid for λ consisted of 51 equally spaced values between 0 and 1. For each of the 17 models, each of the 31×51 parameter combinations, each of the nine subjects, and each of their 2,000 trials, we computed the probability of reporting “right” ($\hat{C} = 1$) given the target and distractor orientations, s_T and s_D , on that trial, $p(\hat{C} = 1 | s_T, s_D, M, J, \lambda)$, where M denotes a certain model. Because this probability could not be computed analytically, we used Monte Carlo simulations with 1,000 sampled measurement vectors (x_1, x_2, x_3, x_4); for each, we applied the model’s decision rule and counted the proportion of “right” reports. This served as an estimate of the probability of reporting “right” in the absence of lapses $p(\hat{C} = 1 | s_T, s_D, M, J, \lambda = 0)$. The probability of reporting “right” in the presence of lapses was then

$$\begin{aligned} p(\hat{C} = 1 | s_T, s_D, M, J, \lambda) \\ = (1 - \lambda)p(\hat{C} = 1 | s_T, s_D, M, J, \lambda = 0) + 0.5\lambda. \end{aligned}$$

Model fitting. To fit each model for a given subject, we used its model predictions obtained above to compute the log likelihood of the parameter combination (J, λ) , which is the logarithm of the probability of all of the subject’s responses given the model and each parameter combination:

$$\begin{aligned}
\log L_M(J, \lambda) &\equiv \log p(\text{data} | M, J, \lambda) \\
&= \log \prod_i^{N_{\text{trials}}} p(\hat{C}_i | s_{T,i}, s_{D,i}, M, J, \lambda) \\
&= \sum_{i=1}^{N_{\text{trials}}} \log p(\hat{C}_i | s_{T,i}, s_{D,i}, M, J, \lambda)
\end{aligned} \tag{10}$$

where i is the trial index, N_{trials} is the number of trials, and we have assumed that there are no sequential dependencies between trials. We found the values of the parameters on the grid that maximized $\log L_M(J, \lambda)$. We verified that, in no case, the maximum was on an edge of the grid. Using a different method for finding the maximum of the likelihood function, namely, a custom-built evolutionary algorithm, gave approximately the same fits and qualitatively identical model comparison results.

Model comparison. For Bayesian model comparison, we estimated the marginal likelihood of each model for each subject, $p(\text{data} | M)$. We obtained the marginal likelihood of a model by integrating the parameter likelihood over the parameters,

$$\begin{aligned}
p(\text{data} | M) &= \iint p(\text{data} | M, J, \lambda) p(J, \lambda) dJ d\lambda \\
&= \iint L_M(J, \lambda) p(J, \lambda) dJ d\lambda.
\end{aligned}$$

For the parameter prior $p(J, \lambda)$, we assumed a uniform distribution on the rectangle defined by the ranges mentioned above. For numerical convenience, we took the logarithm, giving the log marginal likelihood (LML):

$$\begin{aligned}
\text{LML}(M) &= \log \iint L_M(J, \lambda) p(J, \lambda) dJ d\lambda \\
&= \text{LL}_{\max}(M) + \log \iint e^{\log L_M(J, \lambda) - \text{LL}_{\max}(M)} p(J, \lambda) dJ d\lambda,
\end{aligned} \tag{11}$$

where $\text{LL}_{\max}(M) = \max_{J, \lambda} L_M(J, \lambda)$. This form prevents numerical problems associated with the integrand becoming too small. We used a Riemann sum over the previously defined parameter grid to approximate the integral.

As an alternative to Bayesian model comparison, we could use one of several well-known information criteria: the Akaike information criterion (AIC), a small-sample variant of it (AICc), and the Bayesian information criterion (BIC). All of these penalize the maximum log likelihood, $\text{LL}_{\max}(M)$, by a term increasing with the number of free parameters. However, because all models have two parameters, all penalty terms cancel when taking the difference between two models, and the difference reduces to the difference between the models' $\text{LL}_{\max}(M)$ values. All information criteria therefore yield the same differences, and we will refer to them as *IC.

Model recovery. To validate our methods, we performed a model recovery analysis. We generated nine synthetic data sets of

3,000 trials from each one of the 25 models. Each data set corresponded to an actual subject, in the following sense. To make our synthetic data realistic, we chose each parameter value to be the weighted average of the maximum-likelihood estimates of that parameter obtained from the subject's data under each of the 25 models, weighted by the posterior probability of the model. (Many other ways of choosing parameters for synthetic data will also work.) We then fitted each of the 25 models to each of the 225 synthetic data sets. We found that the correct model had the highest LML and the lowest *IC in the same 137 out of 225 cases. Most of the confusions (79/88) arose within one of two sets of models. One set contains the Opt, SunErfT3, SumErfT13, SumErfT23, SumXT3, SumXT13, SumXT23, and SumXT123 models. The other set contains Sum, MaxT2, SumErf, SumErfT1, SumErfT2, SumErfT12, SunXT1, SumXT2, and SumXT12. Within each of these two sets, we cannot distinguish the models. When we averaged over different data sets generated from the same model, the correct model won in all cases (25/25; for the confusion matrix of mean LML, see Figure A1 of the Appendix).

No data set generated from a simple heuristic or a two-step model was fitted better by the Opt model than by the true model. Moreover, the LML and *IC differences were large (see Table 1), indicating that one is very unlikely to mistake data from those models for being generated from Opt. In other words, the Opt model is not able to mimic the simple models.

Modeling Results

Comparison between the optimal model and the simple heuristic models. Figure 3 shows the fits of the Opt model and the simple heuristic models. The Opt model fits the data best (Figures 3A to D). In particular, the Opt model accurately accounts for the counterintuitive nonmonotonicity of proportion correct as a function of distractor orientation (Figure 1D, bottom; Figure 3D). All simple heuristic models exhibited systematic, clearly visible deviations from the data (Figures 3B to D). The LML of the Opt model exceeded that of the Sum, Max, Min, Var, and Sign models by 244 ± 48 , 161 ± 43 , 282 ± 30 , 99 ± 28 , and 189 ± 32 , respectively (paired t tests, $p \leq .01$). The Opt model was most likely for each of the nine subjects individually. We obtained nearly identical results using the *IC information criteria (Figure A2). This suggests that people prioritize optimality over simplicity in this visual search paradigm.

Comparison between the optimal model and the two-step models. None of the two-step models fits the data well (Figure 4A to C). Bayesian model comparison confirms this: The mean LML of the Opt model exceeds that of the two-step models by 253 ± 48 , 256 ± 47 , 130 ± 34 , 38 ± 10 , and 126 ± 34 ($p < .01$).

Table 1
Test for False Alarms in Model Recovery

Opt tested on	LML # wins	LML difference	*IC #wins	*IC difference
Simple heuristic	0/45	[−348, −42] (mean −173)	0/45	[42, 347] (mean 173)
Two-step	0/45	[−382, −23] (mean −161)	0/45	[23, 381] (mean 161)
Generalized sum	12/70	[−436, 0] (mean −205)	12/70	[0, 435] (mean 204)

Note. Comparison between the Opt model and other models on synthetic data sets generated with suboptimal models. LML: log marginal likelihood, *IC: information criterion. #wins indicates the number of data sets for which the optimal model fits best. For the LML and *IC differences, we are giving the range and the mean.

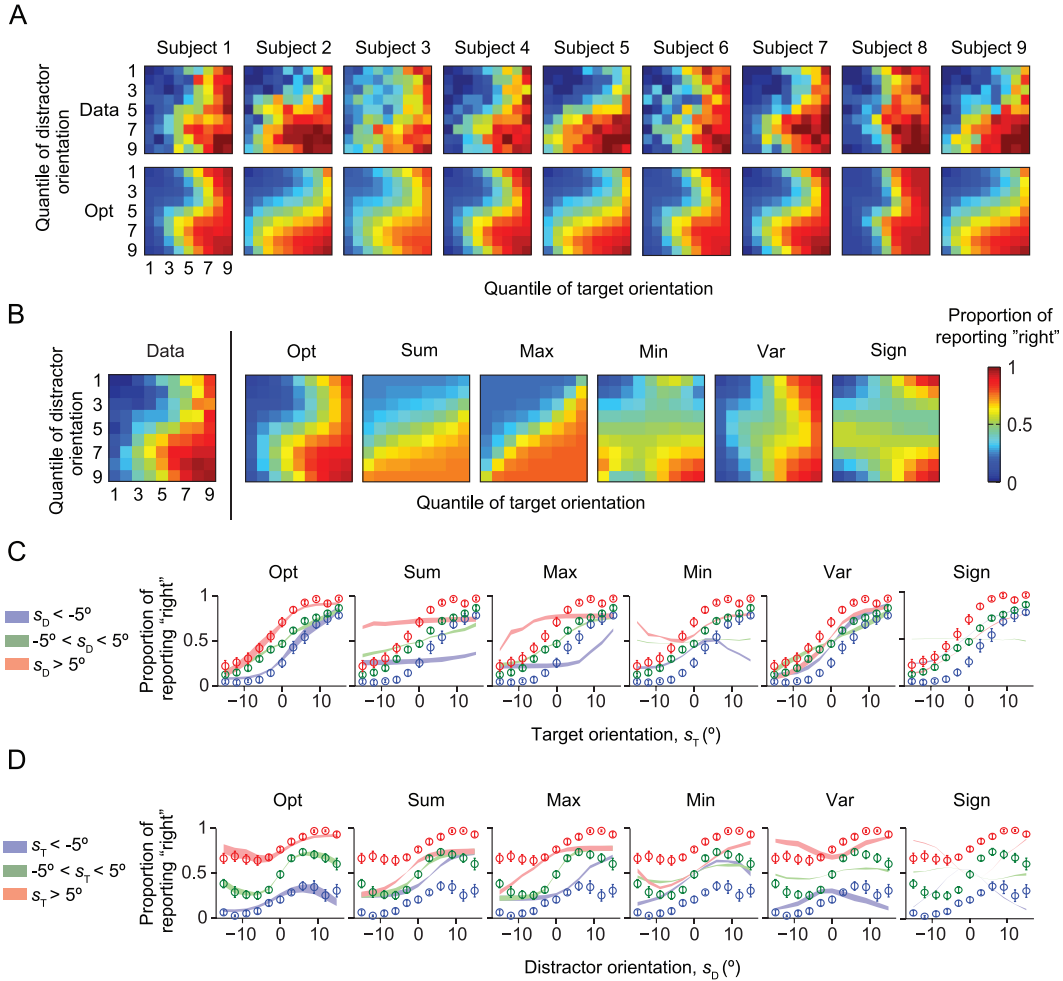


Figure 3. Model fits of the Opt model and the simple heuristic models. The Opt model fits better than the heuristic models. (A) Proportion of reporting “right” (color, or gray scale) as a function of target and distractor orientation quantiles, for individual subjects. The top plot shows the data, the bottom the fits of the Opt model. (B) As (A), averaged over subjects. The leftmost plot shows the data from Fig. 1C, the other plots the model fits. (C) Proportion of reporting “right” as a function of target orientation s_T . Circles and error bars: data; shaded areas: model fits. (D) Proportion of reporting “right” as a function of distractor orientation s_D . See the online article for the color version of this figure.

The Opt model is more likely than any of the two-step models for all nine subjects. We obtained nearly identical results using the *IC information criteria (Figure A2). This suggests that in perceptual decision making, people do not follow the heuristic of only relying on the most discriminative cue.

Comparison between the optimal model and the generalized sum models. So far, none of the simple models tested can account for human behavior. Among the “generalized sum models,” those without the distractor variance term (Term 3) fitted poorly to human data (Figures 5 and 6). The LML of the Opt model was higher than the log likelihoods of the models without Term 3 by 256 ± 47 (SumErf), 265 ± 46 (SumErfT1), 256 ± 47 (SumErfT2), 266 ± 46 (SumErfT12), 266 ± 46 (SumXT1), 249 ± 48 (SumXT2), and 265 ± 46 (SumXT12; $p < .001$; Figure 7A, B). However, models that included Term 3 were not distinguishable from the Opt model: The LML differences were -1.9 ± 1.3

(SumErfT3, $p = .22$), 21.7 ± 9.1 (SumErfT13, $p = .05$), -8.2 ± 4.1 (SumErfT23, $p = .10$), -4.8 ± 3.3 (SumXT3, $p = .20$), 12.2 ± 5.7 (SumXT13, $p = .08$), -8.9 ± 5.8 (SumXT23, $p = .19$), and -3.1 ± 2.1 (SumXT123, $p = .20$). We obtained nearly identical results using the *IC information criteria (Figure A2). Thus, like in the synthetic data used for model recovery, seven suboptimal models (SumErfT3, SumErfT13, SumErfT23, SumXT3, SumXT13, SumXT23, and SumXT123) fitted the data about as well as the Opt model. However, these are by no means simple models (see Table 1). Our results suggest that the distractor variance term (Term 3) is a crucial component of the decision rule, implying that subjects used the knowledge that the three distractors have the same orientation in their decision.

Model agreement. We found that seven (of 24) suboptimal models can describe human behavior equally well as the Opt model. This might seem unsatisfactory, because it seems that no

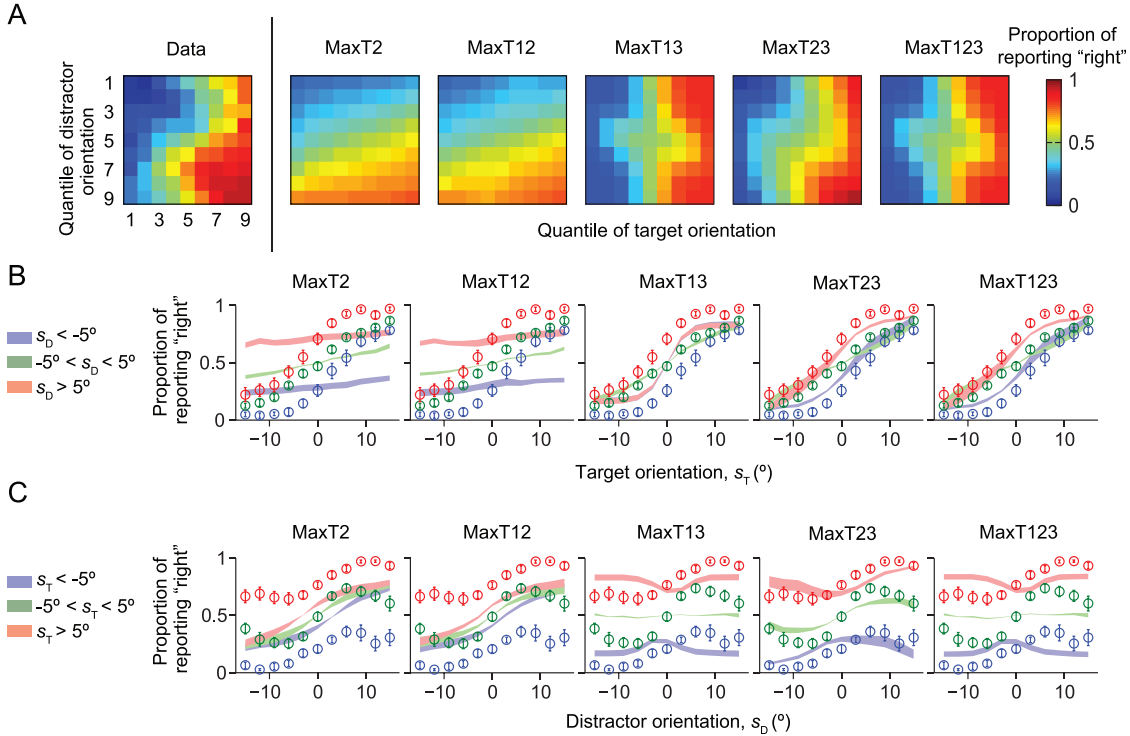


Figure 4. Model fits of the two-step models. The Opt model fits better than the two-step models. (A) Proportion of reporting “right” (color, or gray scale) as a function of target and distractor orientation quantiles, averaged over subjects. The leftmost plot shows the data from Fig. 1C, the other plots the model fits. (B) Proportion of reporting “right” as a function of target orientation s_T . Circles and error bars: data; shaded areas: model fits. (C) Proportion of reporting “right” as a function of distractor orientation s_D . See the online article for the color version of this figure.

definite conclusion can be drawn about whether people are optimal or not. Instead, we argue that it is flawed to think of the distinction between optimal and suboptimal models as categorical, and that the seven well-fitting suboptimal models are, for all practical purposes, identical to the Opt model. A first hint of this was provided by the earlier model recovery analysis, in which we found that models within this group of eight models are indistinguishable (Figure A1). However, we can quantify similarity between models in more detail. We will do this in two ways.

One way to think of model similarity is in terms of predicted accuracy. We computed the proportion correct predicted by different models given the same set of simulated stimulus vectors (consisting of 100,000 trials) across a range of values of the noise parameter σ (0.01 to 10, with a step of 0.01). We found that the accuracy levels predicted by these eight models are nearly identical (Figure 8A), which means that observers using any of these eight decision rules would all obtain near-maximal accuracy.

This analysis still leaves two possibilities: The seven alternative models do well because they are nearly identical to the Opt model in their trial-to-trial predictions, or because they account better than the Opt model for some trials and worse for others. In other words, we are not just interested in whether an observer uses a rule that allows near-maximal accuracy for a given noise level but also in whether they use the specific decision rule in Equation 5.

To distinguish between these possibilities, we examined the trial-to-trial agreement between the model predictions of the Opt

model and of each alternative model, under a lapse rate of zero. (A nonzero lapse rate would simply replace some responses by coin flips and would not change any of the following results.) At each value of the noise level σ , we simulated 100,000 measurement vectors drawn from the generative model. For each measurement vector, the Opt and the alternative model make a deterministic prediction for the binary response. Comparing these predictions across all measurement vectors, we obtain a “proportion agreement.” We then plotted this quantity as a function of the noise level σ (Figure 8B). The seven models that were indistinguishable from the Opt model made trial-to-trial predictions that agreed with those of the Opt model on more than 95% of trials. This shows that the models that fit about as well as the Opt model do so *because* the predictions of these models and the Opt model strongly agree from trial to trial.

A structure on the space of decision rules? Although we tested many more models than is common in psychophysics studies, the set of all possible models is obviously infinite, which make it hard to infer whether an untested model would fit better. However, we are in the special circumstance that our models only differ in their decision rules. Here, we make use of this property to explore the structure of the decision rule space, with the goal of making inferences beyond the decision rules that we tested.

In our task, a decision rule is a mapping from a measurement vector $\mathbf{x} \in \mathbb{R}^d$ and the parameter σ to a binary response $\hat{C} = \pm 1$. There are infinitely many such mappings; when we fix σ , one can

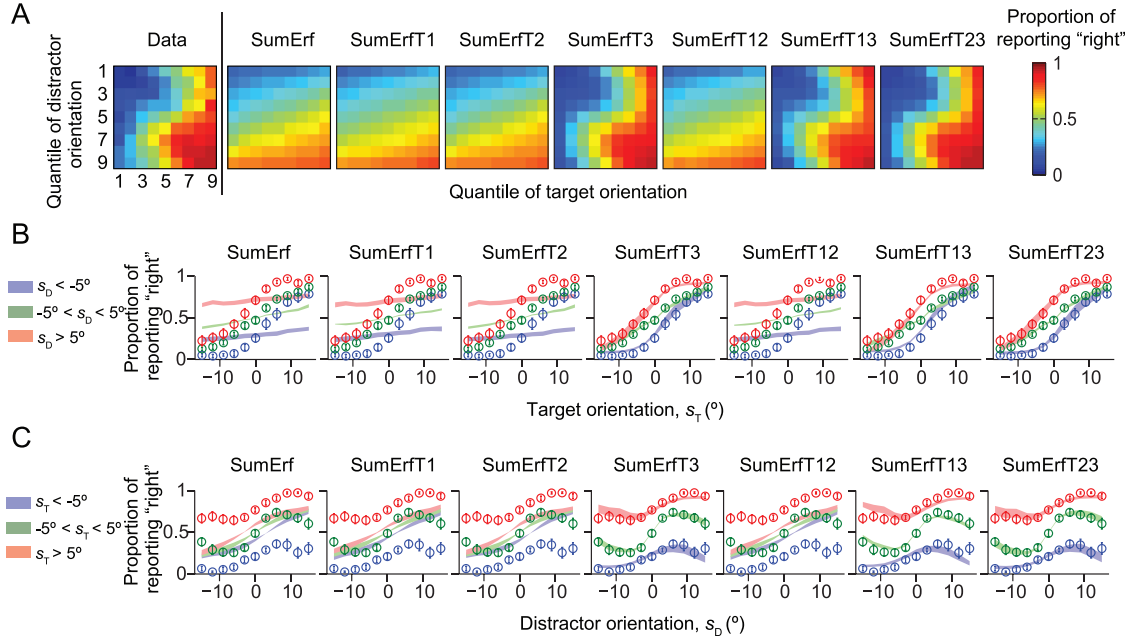


Figure 5. Model fits of the generalized sum models of the SumErfT* type. Models contain term 3 (SumErfT3, SumErfT13, and SumErfT23) fit about as well as the Opt model. (A) Proportion of reporting “right” (color, or gray scale) as a function of target and distractor orientation quantiles, averaged over subjects. The leftmost plot shows the data from Fig. 1C, the other plots the model fits. (B) Proportion of reporting “right” as a function of target orientation s_T . Circles and error bars: data; shaded areas: model fits. (C) Proportion of reporting “right” as a function of distractor orientation s_D . See the online article for the color version of this figure.

think of them as different ways of dividing four-dimensional space into black and white regions (for a three-dimensional analog showing the decision boundary, see Figure 2B). One could shrink the space of decision rules by imposing that d must be antisymmetric under the sign flip $\mathbf{x} \mapsto -\mathbf{x}$ and invariant under permutations of the elements of \mathbf{x} . In addition, one could impose smoothness constraints or neural constraints. Even so, d -space will still be infinite-dimensional and will not be a vector space in an obvious way. However, it is possible to equip it with a metric structure, which we will describe now.

In the previous section, we introduced the proportion of trial-to-trial agreement between two decision rules. There, it was a function of σ , but we can reduce it to a single number by averaging over σ :

$$\text{Agreement}(d_1, d_2) = \langle \delta_{d_1(\mathbf{x}, \sigma), d_2(\mathbf{x}, \sigma)} \rangle_{\mathbf{x}, \sigma}. \quad (12)$$

In this equation, δ denotes the Kronecker delta function, which equals 1 when its two subscripted integer arguments are equal to each other and 0 otherwise. The average $\langle \cdot \rangle$ is over everything: noise realizations, stimuli, and the parameter σ (for which we again use the range $[0, 10]$). Agreement is linearly related to the Hamming distance (Hamming, 1950) computed across the two binary strings of \hat{C} values obtained by evaluating both decision rules on randomly sampled (σ, \mathbf{x}) pairs. This means that Agreement equips the space of decision rules with a metric structure.

The Agreement metric is less general than Kullback–Leibler (KL) or Jensen-Shannon (JS) divergence: Agreement is designed for comparing two deterministic decision rules acting on the same internal representation, whereas the divergences can be used to characterize

similarity between any two models. However, Agreement has the advantage of capturing the *trial-to-trial* similarity between model predictions: Whereas the divergences only compare the predicted *distributions* of responses conditioned on the stimuli, \mathbf{s} , Agreement compares the predicted *individual responses* conditioned on the internal representation, \mathbf{x} . Consider an example in which a specific stimulus \mathbf{s} is repeated many times, and two Models A and B both predict the subject to report “right” on 70% of these trials. Then, KL/JS divergence between A and B will be 0. However, Agreement between A and B might differ greatly. One extreme possibility is that A and B make the exact same prediction for every \mathbf{x} that is generated from \mathbf{s} , that is, Agreement is 100%. The other extreme possibility is that A and B make maximally distinct predictions; in this case, Agreement is only 40%.

Relation between agreement with Opt and goodness of fit.

Now that we have used Agreement to define a structure on model space, we can examine how goodness of fit (LML) depends on Agreement. As a first step, we visualized the space of decision rules using multidimensional scaling. This method converts the matrix of Agreement values between pairs of decision rules ($25 \times 24/2 = 300$ values in total) into distances in a low-dimensional space (Borg & Groenen, 2005). We found that the models that lie farther away from the Opt model in this low-dimensional space tend to fit worse (Figure 8C). This suggests that Agreement is informative about goodness of fit. We now examine this suggestion in more detail.

Given a data set, an ideal conclusion to draw would be that one particular decision rule, say d^* , has the highest LML in the space

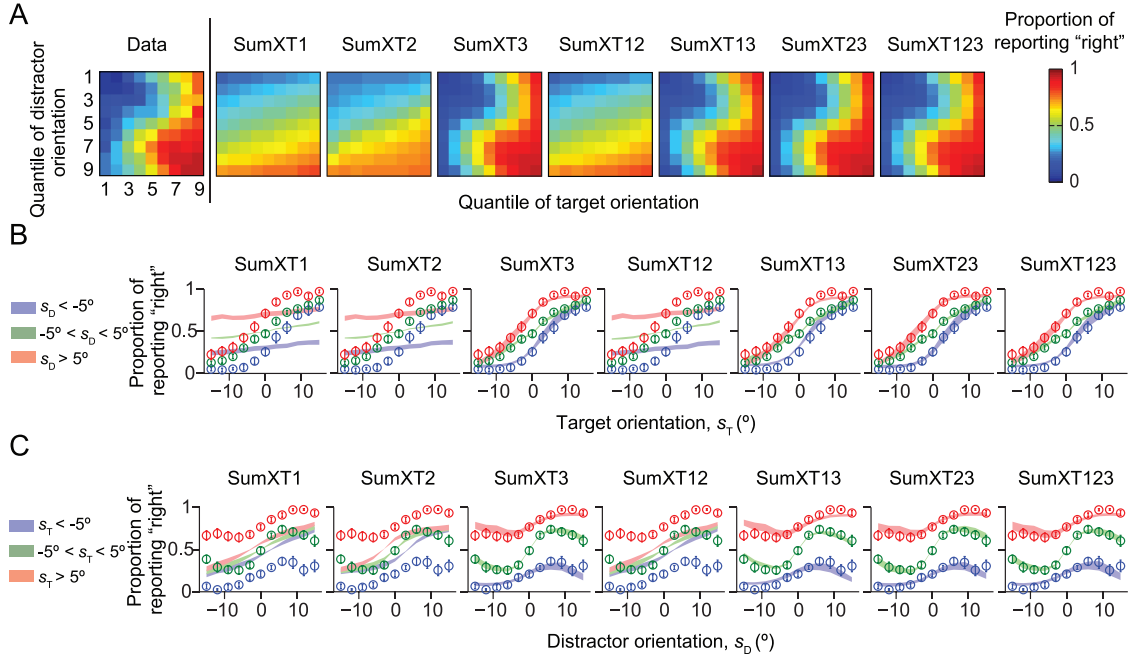


Figure 6. Model fits of generalized sum models of the SumXT* type. Models contain term 3 (SumXT3, SumXT13, SumXT123, and SumXT123) fit about as well as the Opt model. (A) Proportion of reporting “right” (color, or gray scale) as a function of target and distractor orientation quantiles, averaged over subjects. The leftmost plot shows the data from Fig. 1C, the other plots the model fits. (B) Proportion of reporting “right” as a function of target orientation s_T . Circles and error bars: data; shaded areas: model fits. (C) Proportion of reporting “right” as a function of distractor orientation s_D . See the online article for the color version of this figure.

of all plausible mappings d . In practice, we can only test a small number of d , so it is helpful to know whether there exists a monotonic relationship between LML and Agreement with d^* . In fact, a sufficient condition for d^* being the only maximum in all of d -space is that LML is lower whenever Agreement to d^* is lower. The LML landscape does not necessarily obey this property globally, and is in fact unlikely to; for example, there could be local maxima (multimodality) or ridges. However, the property might still hold in most of the d -space, and therefore we think that the correlation between LML and Agreement (CLA) is informative.

Across the 25 models we tested, we find that a model’s LML is strongly correlated with its Agreement with the Opt model, with a correlation of 0.99 (Figure 8D). The correlation was much lower when we correlated LML with the Agreement with a model other than Opt, except for the seven other best models (Figure 9A, Figure A3). These observations suggest that the more similar a model’s lapse-free trial-to-trial predictions are to the eight best models, the better the model (with a lapse rate) fits the human data.

In synthetic data generated using parameters fitted to subject data (see the “Model recovery” section), CLA with the Opt model as the reference model was only high when the data were generated from the seven other best models (Figure 9B, Figure A4). Moreover, the CLA with a specific model as the reference model was high when the synthetic data were generated using that model (Figure 9C, Figure A5). Therefore, the observed high correlation between a model’s LML on subject data and its Agreement with one of the eight best models is consistent with the eight best models being close to the true model underlying the data.

However, this argument relies on the small, rather arbitrary, and possibly biased set of models that we tested here. In particular, the high correlation would be unsurprising if the models tested all reside in a local neighborhood of the Opt model. In addition, a drawback of the analysis above is that Agreement was defined on models without a lapse rate, whereas LML was computed on the same models with a lapse rate. To further investigate how close the eight best models are to the true model underlying the data, we now introduce an independent approach.

How Good Are the Best Models?

Here we use an information-theoretical method to determine how well a model fits in an “absolute” sense, specifically, how much room there is for an untested model to fit better than the eight best models. The basic idea can be illustrated with a simple example: If a biased coin has a probability p of coming up heads, and we try to account for the outcomes of N independent tosses of that coin, then the best we can do is to state for each toss that the probability of heads is p . When N is large, the log likelihood of this model will be the sum of $Np \log p$, from the heads outcomes, and $N(1-p) \log (1-p)$, from the tails outcomes, which is the negative entropy of a sequence of this coin’s toss outcomes. No model can have a higher log likelihood: The log likelihood of any model is bounded from above by the pure stochasticity of the data. The argument below amounts to estimating how close the log likelihood of our best models is to the upper bound given by the negative entropy.

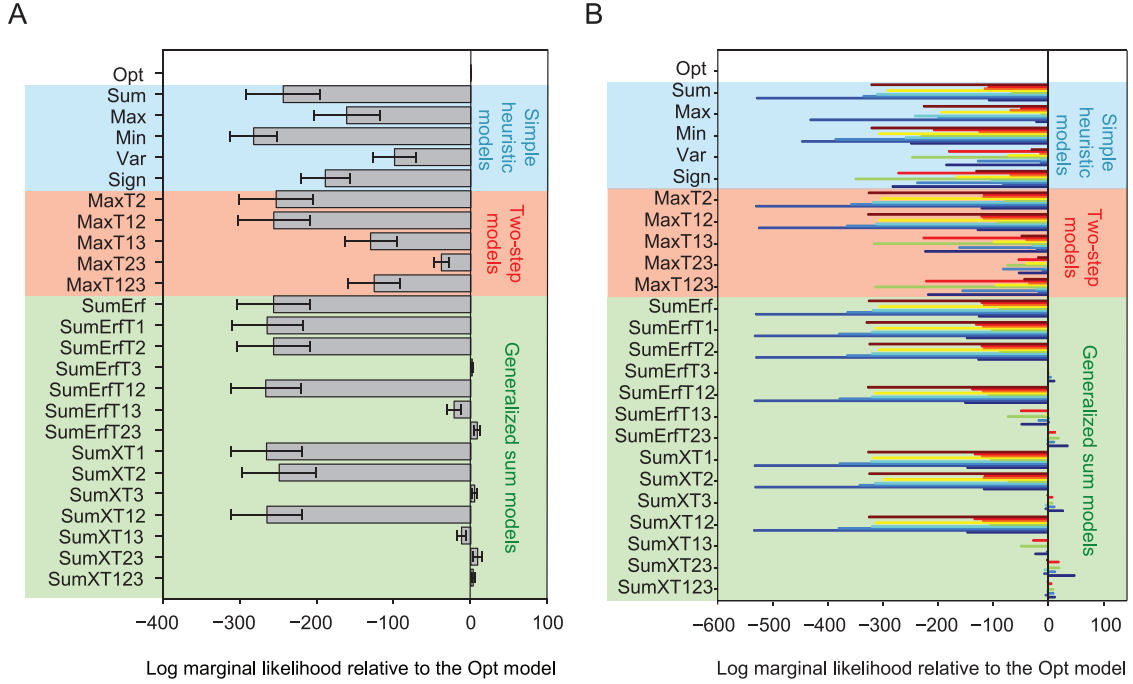


Figure 7. Model comparison. (A) Mean and standard error of the mean across subjects of the difference in log marginal likelihood between each model and the Opt model. (B) Difference in log marginal likelihood between each model and the Opt model for individual subjects; bars of different colors (or gray scales) represent different subjects. See the online article for the color version of this figure.

KL divergence. The data D produced by a model M follow a probability distribution $p(D|M)$. If we assume that there is a true model M_{true} , then its data distribution is $p(D|M_{\text{true}})$. A principled measure of the distance between a model M and the true model, M_{true} , is the KL divergence between these two data distributions (Cover & Thomas, 2005):

$$D_{\text{KL}}(p(D|M_{\text{true}})||p(D|M)) \equiv \sum_D p(D|M_{\text{true}}) \log \frac{p(D|M_{\text{true}})}{p(D|M)}.$$

This quantity is always non-negative and can be evaluated as

$$D_{\text{KL}}(p(D|M_{\text{true}})||p(D|M)) = -H(p(D|M_{\text{true}})) + H(p(D|M_{\text{true}}), p(D|M)), \quad (13)$$

where

$$H(p(D|M_{\text{true}})) = -\sum_D p(D|M_{\text{true}}) \log p(D|M_{\text{true}}) \quad (14)$$

is the entropy of the data distribution $p(D|M_{\text{true}})$, and

$$H(p(D|M_{\text{true}}), p(D|M)) = -\sum_D p(D|M_{\text{true}}) \log p(D|M) \quad (15)$$

is the cross-entropy between $p(D|M_{\text{true}})$ and $p(D|M)$, and, as we will see below, closely related to model log likelihood. Because D_{KL} is non-negative, the negative entropy is an upper bound on the negative cross-entropy. The KL divergence corresponds to unexplained variation, the entropy to subject stochasticity (“unexplainable variation”), and the negative cross-entropy to the goodness of

fit (“explained variation”). Thus, goodness of fit, KL divergence, and unexplained variation sum up to a perfect fit (zero). No model can fit the data better than allowed by the stochasticity of subject responses.

Both terms in Equation 13 involve a sum over all possible data sets, D s, generated from model M_{true} , but we have only a single one available, namely, the subject data. Therefore, both terms have to be estimated based on that one data set, which we denote by “data,” as before. Before we resolve this issue, we will first simplify Equations 14 and 15 by assuming independence of trials and discretizing the stimuli.

Negative entropy term. A possible data set consists of a binary vector of length N_{trials} , consisting of $+1$ and -1 responses: $D = \hat{C} \in \{\pm 1\}^{N_{\text{trials}}}$. We now assume that the trials are conditionally independent of each other. We can then evaluate the negative entropy of $p(D|M_{\text{true}})$, starting from Equation 14:

$$\begin{aligned} -H(p(\hat{C}|M_{\text{true}})) &= \sum_{\hat{C} \in \{\pm 1\}^{N_{\text{trials}}}} p(\hat{C}|M_{\text{true}}) \log p(\hat{C}|M_{\text{true}}) \\ &= \sum_{\hat{C}_1 = \pm 1} \dots \sum_{\hat{C}_{N_{\text{trials}}} = \pm 1} \left(\prod_{i=1}^{N_{\text{trials}}} p(\hat{C}_i|M_{\text{true}}) \right) \log \left(\prod_{i=1}^{N_{\text{trials}}} p(\hat{C}_i|M_{\text{true}}) \right), \\ &= \sum_{i=1}^{N_{\text{trials}}} \sum_{\hat{C}_i = \pm 1} p(\hat{C}_i|M_{\text{true}}) \log p(\hat{C}_i|M_{\text{true}}) \end{aligned}$$

where \hat{C}_i is the subject’s response on the i^{th} trial.

To make further progress, we need to define unique stimulus conditions that have sufficiently many trials. To this end, we binned the data as in Figure 1B: nine quantiles for s_T crossed,

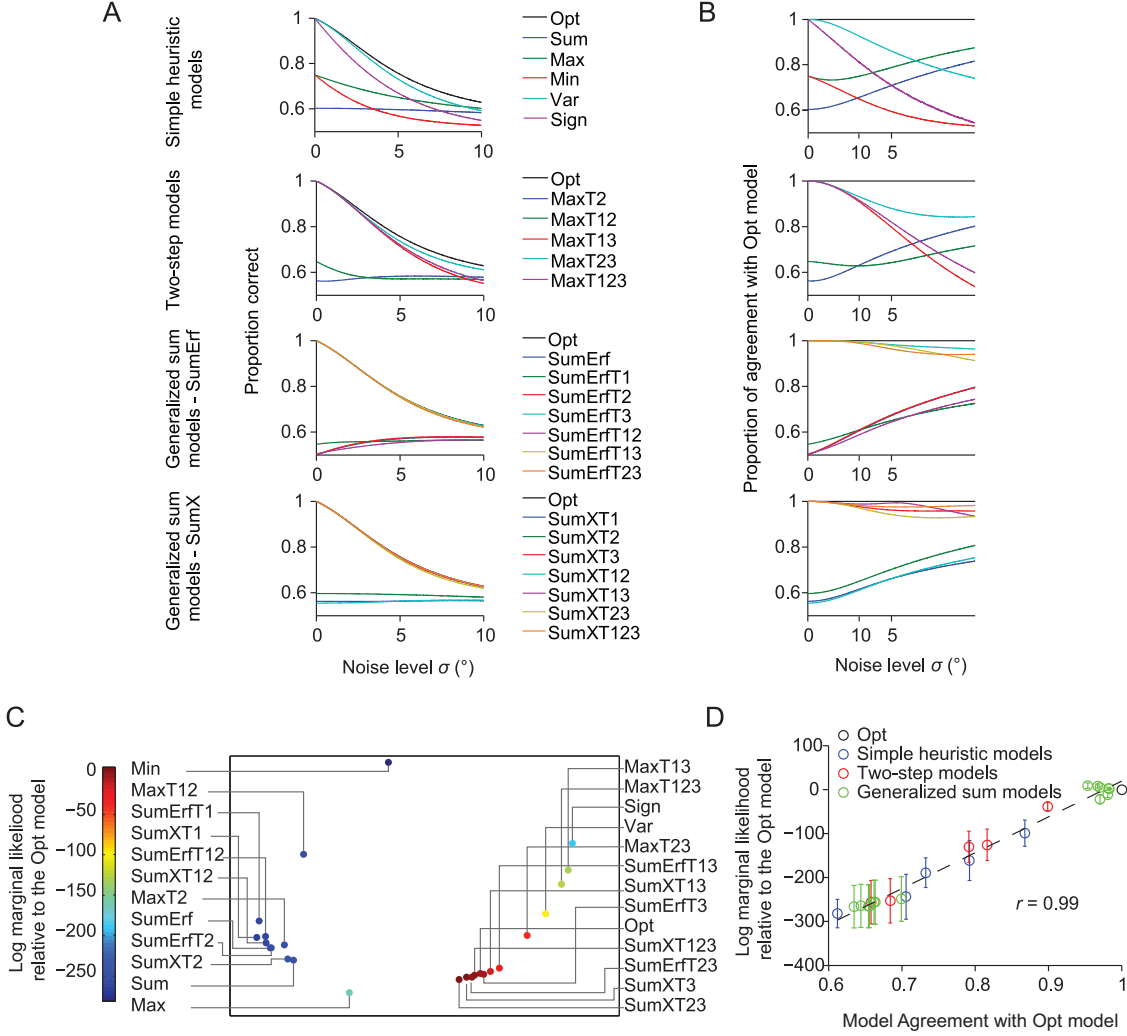


Figure 8. Model similarity and goodness of fit. (A) Proportion correct as a function of the noise level σ for all models. (B) Proportion of trials for which a model makes the same prediction as the Opt model, as a function of σ . (C) Averaged prediction agreement (“Agreement”) visualized using multi-dimensional scaling. Each dot represents a model, and the distance between two models represents the disagreement between those models. The color (or gray scale) of a dot represents its log marginal likelihood. Models that agree more with the Opt model tend to have a higher log marginal likelihood. (D) Mean (open circle) and standard error of the mean (error bar) across subjects of a model’s log marginal likelihood as a function of its Agreement with the Opt model. Each dot indicates a model. The solid line represents the best linear fit. r is the Pearson correlation. See the online article for the color version of this figure.

with nine quantiles for s_D . We denote the number of trials when the stimuli are in the j^{th} stimulus bin ($j = 1, \dots, 81$) by N_j , and, among those, the number of trials when the subject responded “right,” by n_j . Thus, we reduce each subject’s data to 81 counts. We verified that the number of bins did not meaningfully affect our results.

Using this discretization, the distribution $p(\hat{C}_j | M_{\text{true}})$ is the same for all trials i in the same stimulus condition. Therefore, we can group the trials by stimulus bin j , and negative entropy becomes

$$-H(p(\hat{C} | M_{\text{true}})) = \sum_{j=1}^{81} N_j \sum_{\hat{C}_j=\pm 1} p(\hat{C}_j | M_{\text{true}}) \log p(\hat{C}_j | M_{\text{true}}).$$

Defining $\pi_j \equiv p(\hat{C}_j = 1 | M_{\text{true}})$, the probability of reporting “right” under the unknown true model, we have

$$-H(p(\hat{C} | M_{\text{true}})) = \sum_{j=1}^{81} N_j (\pi_j \log \pi_j + (1 - \pi_j) \log (1 - \pi_j)).$$

(16)

Negative cross-entropy term. We now turn to the cross-entropy of $p(D | M_{\text{true}})$ and $p(D | M)$, as given by Equation 15. This term is difficult to estimate, because responses in different stimulus bins are independent only conditional on the parameters. Unfortunately, the parameters θ (J and λ) are unknown and have to be estimated from the data or marginalized over. Here,

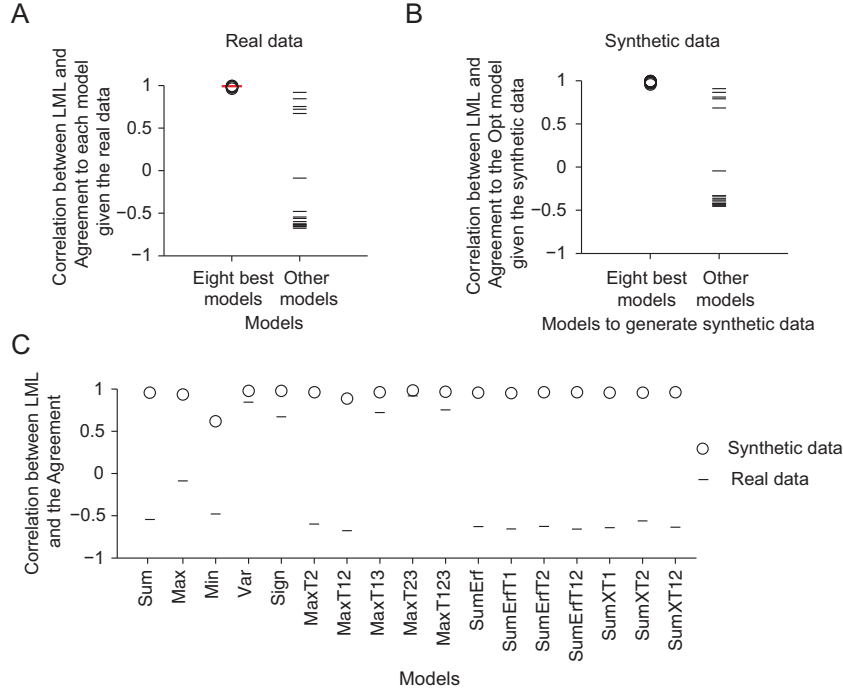


Figure 9. Correlation between log marginal likelihood and Agreement (CLA) as a potential measure of the global maximum of goodness of fit in model space. (A) Agreement, and therefore CLA, is computed relative to a reference model. CLA is high when the reference model is the Opt model (red horizontal line, see also Fig. 8D) or one of the seven other best models (see also Fig. A3). CLA is significantly lower when the reference model is a different model (Wilcoxon rank-sum test, $p = 8.4 \times 10^{-5}$). (B) Given synthetic data generated from one of the eight best models, CLA with the Opt model as the reference is high. Given synthetic data generated from a model outside of the eight best, CLA with the Opt model as the reference is significantly lower (Wilcoxon rank-sum test, $p = 8.4 \times 10^{-5}$, see also Fig. A4). This serves as a negative control for the high CLA with the Opt model as a reference (red horizontal line in (A) and Fig. 8D). (C) Given synthetic data generated from any one model, the CLA with that model as the reference is high (> 0.9). Moreover, given synthetic data generated from a model outside of the eight best, the CLA with that model as a reference is significantly higher than given the real data (Wilcoxon signed-rank test, $p = 2.9 \times 10^{-4}$, see also Fig. A5). This serves as a positive control for the low CLAs with the models outside of the eight best as reference models (circles in (A)). See the online article for the color version of this figure.

the former is easier. For each model, we used every other trial to obtain maximum likelihood estimates $\hat{\theta}$ of the parameters (using the same parameter grid as before), and then used these values to evaluate model predictions on the other half of the trials. Then, we can use a trial factorization and stimulus binning analogous to the ones done above for the entropy, to arrive at

$$-H(p(\hat{C}|M_{\text{true}}), p(\hat{C}|M, \hat{\theta})) = \sum_{j=1}^{81} N_j (\pi_j \log p(\hat{C}_j = 1 | M, \hat{\theta}) + (1 - \pi_j) \log p(\hat{C}_j = -1 | M, \hat{\theta})) \quad (17)$$

where $p(\hat{C}_j | M, \hat{\theta})$ denotes the probability of response \hat{C}_j under model M on a trial on which the stimuli are in the j^{th} bin. In practice, we computed the latter values as a weighted average of the proportions of \hat{C}_j across a fine grid of stimulus combinations in the j^{th} bin, with weights given by the probabilities of those stimulus combinations.

Estimating the terms: Deviance approach. We now have expressions for the two terms in the KL divergence under the trial

factorization and stimulus binning, Equations 16 and 17. The KL divergence is the difference between these two quantities. In computing this, we face a problem: $\{\pi_j\}$, the predicted proportions of “right” responses under the true model, are unknown because the true model is unknown. We will first describe the standard way to deal with this problem, and then our way.

The standard way to deal with the problem that $\{\pi_j\}$ are unknown is to estimate them as the *empirical* proportions of $\hat{C}_j = 1$ responses, in other words, as $\hat{\pi}_j = \frac{n_j}{N_j}$. Then, the estimator of the negative entropy becomes

$$-\hat{H}(p(\hat{C}|M_{\text{true}})) = \sum_{j=1}^{81} N_j (\hat{\pi}_j \log \hat{\pi}_j + (1 - \hat{\pi}_j) \log (1 - \hat{\pi}_j)) \quad (18)$$

and the estimator of the negative cross-entropy becomes

$$-\hat{H}(p(\hat{C}|M_{\text{true}}), p(\hat{C}|M, \hat{\theta})) = \text{LL}_{\text{cv}}(M), \quad (19)$$

where

$$\begin{aligned} \text{LL}_{\text{cv}}(M) \equiv & \sum_{j=1}^{81} (n_j \log p(\hat{C}_j = 1 | M, \hat{\theta}) \\ & + (N_j - n_j) \log p(\hat{C}_j = -1 | M, \hat{\theta})) \end{aligned} \quad (20)$$

is the cross-validated log likelihood of model M (compare Equation 10). The difference between these two terms is then the estimator of KL divergence:

$$\begin{aligned} \hat{D}_{\text{KL}}(p(D | M_{\text{true}}) \| p(D | M)) \\ = \sum_{j=1}^{81} N_j \left(\hat{\pi}_j \log \frac{\hat{\pi}_j}{p(\hat{C}_j = 1 | M, \hat{\theta})} + (1 - \hat{\pi}_j) \log \frac{1 - \hat{\pi}_j}{p(\hat{C}_j = -1 | M, \hat{\theta})} \right). \end{aligned}$$

Up to an irrelevant factor of 2, this is also known as the *deviance* (Wichmann & Hill, 2001). It is common to perform a statistical test on the deviance to determine whether it is significantly different from 0. If it is not, then the model is statistically “as good as possible.”

Although such an analysis of deviance is widespread, it suffers from a fundamental problem: $\hat{\pi}_j = \frac{n_j}{N_j}$ is an unbiased estimator of π_j , but $\hat{\pi}_j \log \hat{\pi}_j$ is not an unbiased estimator of $\pi_j \log \pi_j$. In fact, the bias in estimating entropy has been characterized in detail (Grassberger, 1988, 2003). Therefore, we explore a different solution here.

Estimating the terms: New approach.

Negative entropy term. To estimate the negative entropy, we use the Grassberger estimator (Grassberger, 2003), evaluated on the same half of the data as used to estimate the cross-entropy:

$$-\hat{H}_G(p(D | M_{\text{true}})) = -\sum_{j=1}^{81} N_j \left(G_{N_j} - \frac{1}{N_j} (n_j G_{n_j} + (N_j - n_j) G_{N_j - n_j}) \right), \quad (21)$$

where the numbers G_n are obtained through $G_0 = 0$, $G_1 = -\gamma - \log 2$ (where $\gamma \approx 0.577215$ is Euler’s constant), $G_2 = 2 - \gamma - \log 2$, and for $n \geq 1$, $G_{2n+1} = G_{2n}$ and $G_{2n+2} = G_{2n} + \frac{2}{2n+1}$. Thus, $G_{2n} = -\gamma - \log 2 + \frac{2}{1} + \frac{2}{3} + \frac{2}{5} + \frac{2}{2n+1}$.

Negative cross-entropy term. The negative cross-entropy term is linear in π_j and therefore does not suffer from the same biased estimation problem. Therefore, we use the estimator of Equation 19, $-\hat{H}(p(\hat{C} | M_{\text{true}}), p(\hat{C} | M, \hat{\theta})) = \text{LL}_{\text{cv}}(M)$.

KL divergence. Our estimator of the KL divergence is then the difference of Equations 21 and 19:

$$\hat{D}_{\text{KL}}(p(D | M_{\text{true}}) \| p(D | M)) = -\hat{H}_G(p(D | M_{\text{true}})) - \text{LL}_{\text{cv}}(M). \quad (22)$$

Significance testing. To test whether D_{KL} is significantly greater than 0, we can no longer assume a chi-squared distribution, as is common for establishing significance of deviance (Collett, 2002, Section 3.8; Wichmann & Hill, 2001). In fact, we do not know how to compute a confidence interval on the expression in Equation 22. Therefore, we make a further approximation by regarding our Grassberger estimate of negative entropy as the truth, and only computing a confidence interval on the cross-entropy term. This can lead to false alarms (a model is falsely declared as being substantially different from the truth) but not to misses; thus, it is a conservative approach if we aim to test whether our best models are indistinguishable from the truth.

To compute a confidence interval for the cross-entropy, we use a Bayesian approach (i.e., we compute a credible interval), that is, we compute the posterior probability distribution over Equation 17. First, we compute the posterior over π_j assuming a flat prior; a Jeffreys prior would not substantially change our results. Then, the posterior over π_j is a beta distribution with mean $\frac{n_j}{N_j}$ and variance $\frac{(n_j+1)(N_j-n_j+1)}{(n_j+2)^2(N_j+3)}$. The next step is to approximate the beta distribution with a normal distribution with the same mean and variance, and make use of the independence of the π_j . Then, the posterior over $-\hat{H}(p(\hat{C} | M_{\text{true}}), p(\hat{C} | M, \hat{\theta}))$ has mean LL_{cv} as before (Equation 20), and variance

$$\sum_{j=1}^{81} \left(N_j \log \frac{p_j(\hat{C} = 1 | M, \hat{\theta})}{p_j(\hat{C} = -1 | M, \hat{\theta})} \right)^2 \frac{(n_j+1)(N_j-n_j+1)}{(N_j+2)^2(N_j+3)}. \quad (23)$$

This yields a 95% credible interval of the estimate of the KL divergence:

$$\left[\begin{aligned} & -\hat{H}(p(D | M_{\text{true}})) - \text{LL}_{\text{cv}}(M) - 1.96 \cdot \sqrt{\text{Var}(\text{LL}_{\text{max}}(M))}, \\ & -\hat{H}(p(D | M_{\text{true}})) - \text{LL}_{\text{cv}}(M) + 1.96 \cdot \sqrt{\text{Var}(\text{LL}_{\text{max}}(M))} \end{aligned} \right]. \quad (24)$$

For every subject, we computed the LL_{max} of a random coin-flip model ($-N_{\text{trials}} \cdot 0.5 \log 2$), the Grassberger estimate of negative entropy (Equation 21), $\text{LL}_{\text{cv}}(M)$ for all models, and the 95% credible interval of the estimate of the negative cross-entropy for the Opt model (Equation 24; Figure 10, Table 2). For most subjects, the estimate of negative cross-entropy is lower than the estimate of negative entropy; for two subjects, they are reversed. Because it is mathematically impossible for negative cross-entropy to exceed the negative entropy, this is an indication of estimation error. For eight of nine subjects, the negative entropy was not significantly higher than the negative cross-entropy of the Opt model, indicating that the Opt model fits the data very well in an “absolute” sense (see Table 2). The same applies to the seven models that are indistinguishable from the Opt model (see Table A2).

Across subjects, the mean KL divergence between the Opt model and the true model is not significantly greater than 0 ($p = .15$, one-sided Wilcoxon signed-rank test), confirming that the Opt model explains most of the explainable variation. The same is true for the seven other best models (SumErfT3, $p = .25$; SumErfT13, $p = .13$; SumErfT23, $p = .25$; SumXT3, $p = .21$; SumXT13, $p = .13$; SumXT23, $p = .18$; SumXT123, $p = .21$).

We can restate the comparison between negative cross-entropy and negative entropy in perhaps more intuitive terms. The negative cross-entropy, when divided by the number of trials and then exponentiated, represents the geometric mean probability of the model correctly predicting the subject’s response on a given trial. Similarly, the negative entropy divided by the number of trials and exponentiated represents the geometric mean probability of correctly predicting the subject’s response on a given trial, given the empirical response frequencies. For example, for the Opt model, the prediction accuracy values are 0.63 ± 0.01 and 0.62 ± 0.01 , respectively (for individual subjects, see Table 2). This means that the Opt model predicts subject responses about as well as possible given random variability in the data. Taken together, these results

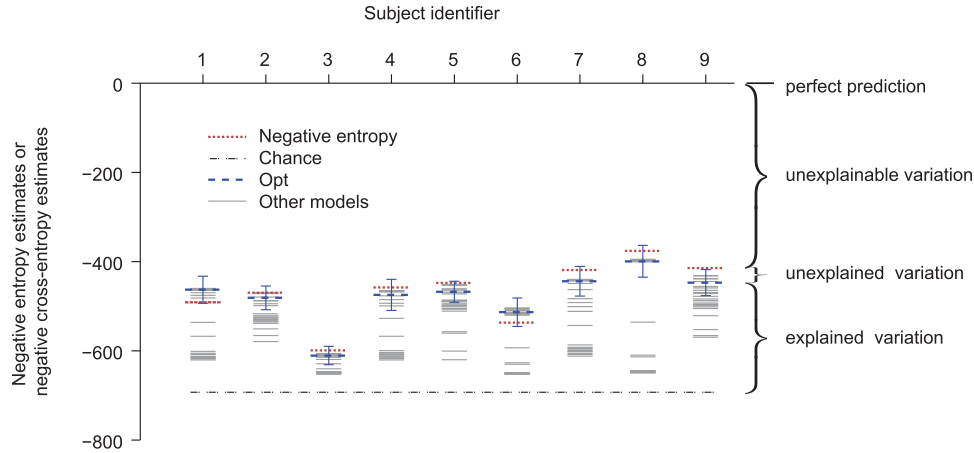


Figure 10. Information-theoretical estimate of how good the eight best models are. Each column represents a subject. For each subject, the red dense dashed line represents an estimate of the negative entropy of the data, the dot dashed black line the negative cross-entropy between a coin-flip model and the true model, the blue sparse dashed line an estimate of the negative cross-entropy between the Opt model and the true model, and the gray lines estimates of the negative cross-entropies between other models and the true model. The error bar represents an estimate of the 95% credible interval of the negative cross-entropy between the Opt model and the true model. The estimate of the negative cross-entropy between the Opt model and the true model is not significantly different from the estimate of the negative entropy of the data (one-sided Wilcoxon signed-rank test, $p = 0.15$), suggesting that the Opt model explains most of the explainable variation. The same holds for the seven other best models (see main text). See the online article for the color version of this figure.

show that there is relatively little room for an untested model to fit better than our eight best models.

Experiment 2

Although the Opt decision rule was derived using the principles of Bayesian inference, our finding that the rule describes human behavior well does not imply that people reason with probabilities. In particular, an observer can, through the trial-to-trial feedback provided in the experiment, learn the Opt rule as a policy or look-up table, without ever representing probabilities in their brain. In earlier work, we therefore distinguished optimal computation from probabilistic computation (Ma, 2012). In the latter, the observer uses an implicit representation of sensory uncertainty, or even of an entire probability distribution over a sensory stimulus, in downstream computation.

There is substantial evidence that people do this (for a review, see Ma & Jazayeri, 2014), but not much from tasks in which the decision rules are as complicated as here.

To vary uncertainty, one can vary stimulus reliability. In Experiment 2, we determined whether people are optimal when we train them at high stimulus reliability and test at low stimulus reliability. This test is a special case of “Bayesian transfer,” a term coined by Maloney and Mamassian (2009) to indicate that a truly probabilistic Bayesian observer should maintain priors, likelihoods, and cost functions as building blocks that can be mixed and matched according to the task demands. The experimental prediction is that people should generalize nearly immediately to a new prior, likelihood, or cost function. In our case, we generalize from a narrow likelihood function (high stimulus reliability) to a wide one (low stimulus reliability).

Table 2

For Most Subjects, KL Divergence Is Not Significantly Different From 0

ID	Negative entropy	Prediction accuracy upper bound	95% CI of negative cross-entropy of Opt	95% CI of prediction accuracy of Opt	95% CI of D_{KL}	p value
1	−492	.61	[−494, −433]	[.61, .65]	[−59, −2]	.97
2	−470	.63	[−508, −455]	[.60, .63]	[−15, 38]	.20
3	−599	.55	[−632, −590]	[.53, .55]	[−9, 33]	.14
4	−458	.63	[−510, −440]	[.60, .64]	[−18, 52]	.17
5	−448	.64	[−491, −444]	[.61, .64]	[−4, 43]	.05
6	−537	.58	[−545, −482]	[.58, .62]	[−55, 8]	.92
7	−419	.66	[−478, −411]	[.62, .66]	[−8, 59]	.07
8	−376	.69	[−435, −364]	[.64, .69]	[−8, 59]	.10
9	−414	.66	[−476, −418]	[.62, .66]	[4, 64]	.02

Note. The p value is the estimated probability that the negative cross-entropy is equal to or higher than the negative entropy (with higher being theoretically impossible). CI = confidence interval.

Method

Experiment 2 was identical to Experiment 1 except for the following differences. The experiment consisted of three sessions, held on different days. Each session contained six blocks of 150 trials each. During the instruction phase on the first day, subjects viewed a plot of the Gaussian orientation distribution $p(s)$ for a single item; the meaning of this plot was explained to them using vocabulary that matched their educational background. Subjects also viewed 30 stimuli whose orientations were drawn from $p(s)$.

Blocks 1 and 4 of each session were training blocks, in which the stimuli were presented at high contrast (peak luminance of the Gabor patch = 120 cd/m²), and correctness feedback was provided after each trial. The other four blocks were testing blocks. The stimuli in these blocks had lower contrast (peak luminance = 56 cd/m²) and importantly, no feedback was provided. Five subjects (three female) participated; all were naïve to the goals of experiment.

The instruction and training were intended to make the subjects learn the stimulus distributions and correct any category (left/right) biases. However, the likelihood functions on the testing trials could not be learned from the training trials, because the stimuli had different reliability.

Results

We performed the same analyses as in Experiment 1 on the testing trials of Experiment 2. We found that the Opt model is still one of the best-fitting models (Figure 11A and B), and Bayesian model comparison revealed model rankings consistent with Experiment 1 (Figure 11C, D). The Opt model and the seven other best models from Experiment 1 have higher LML than the other models.

Moreover, the Agreement and information-theoretic analyses yield the same conclusions as in Experiment 1. First, a model's LML is strongly correlated with its Agreement with the Opt model ($r = .99$, Figure 11F), and much less with its Agreement with a model outside of the best eight (Figure 11G). Second, our estimate of KL divergence between any of the eight best models and the true model is not significantly different from 0 (Figure 11H; e.g., Opt: $p = .31$, one-sided Wilcoxon signed-ranks test). These results suggest that sensory uncertainty is internally represented and that people combine this information with the task demands to achieve near-optimal performance, instead of using a fixed policy or look-up table.

Discussion

In a relatively complicated perceptual decision-making task involving multiple stimuli, we tested the optimal decision rule against a series of suboptimal rules, including many simple heuristic rules. We found that the optimal rule describes the rich patterns in human behavior extremely well, and better than the heuristic rules that we tested.

It should be kept in mind that our model observer is not optimal in an absolute sense (Ma, 2012) because measurement noise is not zero. In fact, this noise itself might reflect suboptimality in earlier stages of processing (Pouget, Beck, Ma, & Latham, 2013), or reflect attentional limitations (Mazyar et al., 2013). Our notion of optimality pertains solely to the decision rule applied to noisy sensory information.

We found that seven suboptimal models fit the data as well as the Opt model. However, those models are indistinguishable from the Opt model both in a model recovery test and in trial-to-trial model agreement. This argues for dropping the hard distinction between optimal and suboptimal, and instead talking about “models indistinguishable from optimal.”

Although we tested a relatively large number of models, model space contains infinitely many more models. Therefore, we introduced and computed two quantities to estimate how well the Opt model and the seven other best models fit relative to untested models: (a) the correlation across models between LML and trial-to-trial agreement with one of the best models, and (b) an estimate of the KL divergence between the eight best models and the true model. The first approach defines a structure on the space of decision rules, but it only applies when decision rules are deterministic and act on the same internal representation. (However, it can likely be generalized by using KL divergence instead of Agreement.) The second approach is much more general but involves many approximations. Taken together, however, we believe that we have provided evidence that it is difficult to find a rule that fits the data better than the Opt rule (and rules indistinguishable from it).

Even without trial-to-trial feedback, the Opt model still provides the best fit, and the Opt model is still close to the true model in an absolute sense, according to both the Agreement analysis and the KL divergence estimates, suggesting that near-optimal behavior is not a result of establishing a look-up table through feedback, but is a result of probabilistic computation.

Figure 11 (opposite). Results of Experiment 2, in which feedback was withheld. (A) Proportion of reporting “right” (color, or gray scale) as a function of each combination of target and distractor orientation quantiles (1 to 9), averaged over all 5 subjects. The left plot shows the data, the right the fits of the Opt model. (B) Proportion of reporting “right” as a function of target orientation. Circles and error bars: data; shaded areas: Opt model fits. (C) Proportion of reporting “right” as a function of distractor orientation. (D) Mean and standard error of the mean across subjects of the log marginal likelihood of each model relative to the Opt model. Shades of different colors (or gray scale) indicate the category of a model. (E) Log marginal likelihood of each model minus that of the Opt model, for individual subjects. In the bar plots, each color (or gray scale) represents a different subject. (F) As Fig. 8D, for Experiment 2. (G) As Fig. 9A, for Experiment 2. CLA is high when the reference model is the Opt model (red horizontal line) or one of the seven other best models, and significantly lower otherwise (Wilcoxon rank-sum test, $p = 8.4 \times 10^{-5}$). (H) As Fig. 10, for Experiment 2. The estimate of the negative cross-entropy between the Opt model and the true model is not significantly different from the estimate of the negative entropy of the data (one-sided Wilcoxon signed-rank test, $p = 0.31$), suggesting that the Opt model explains most of the explainable variation. The same conclusion holds for the seven other best models. See the online article for the color version of this figure.

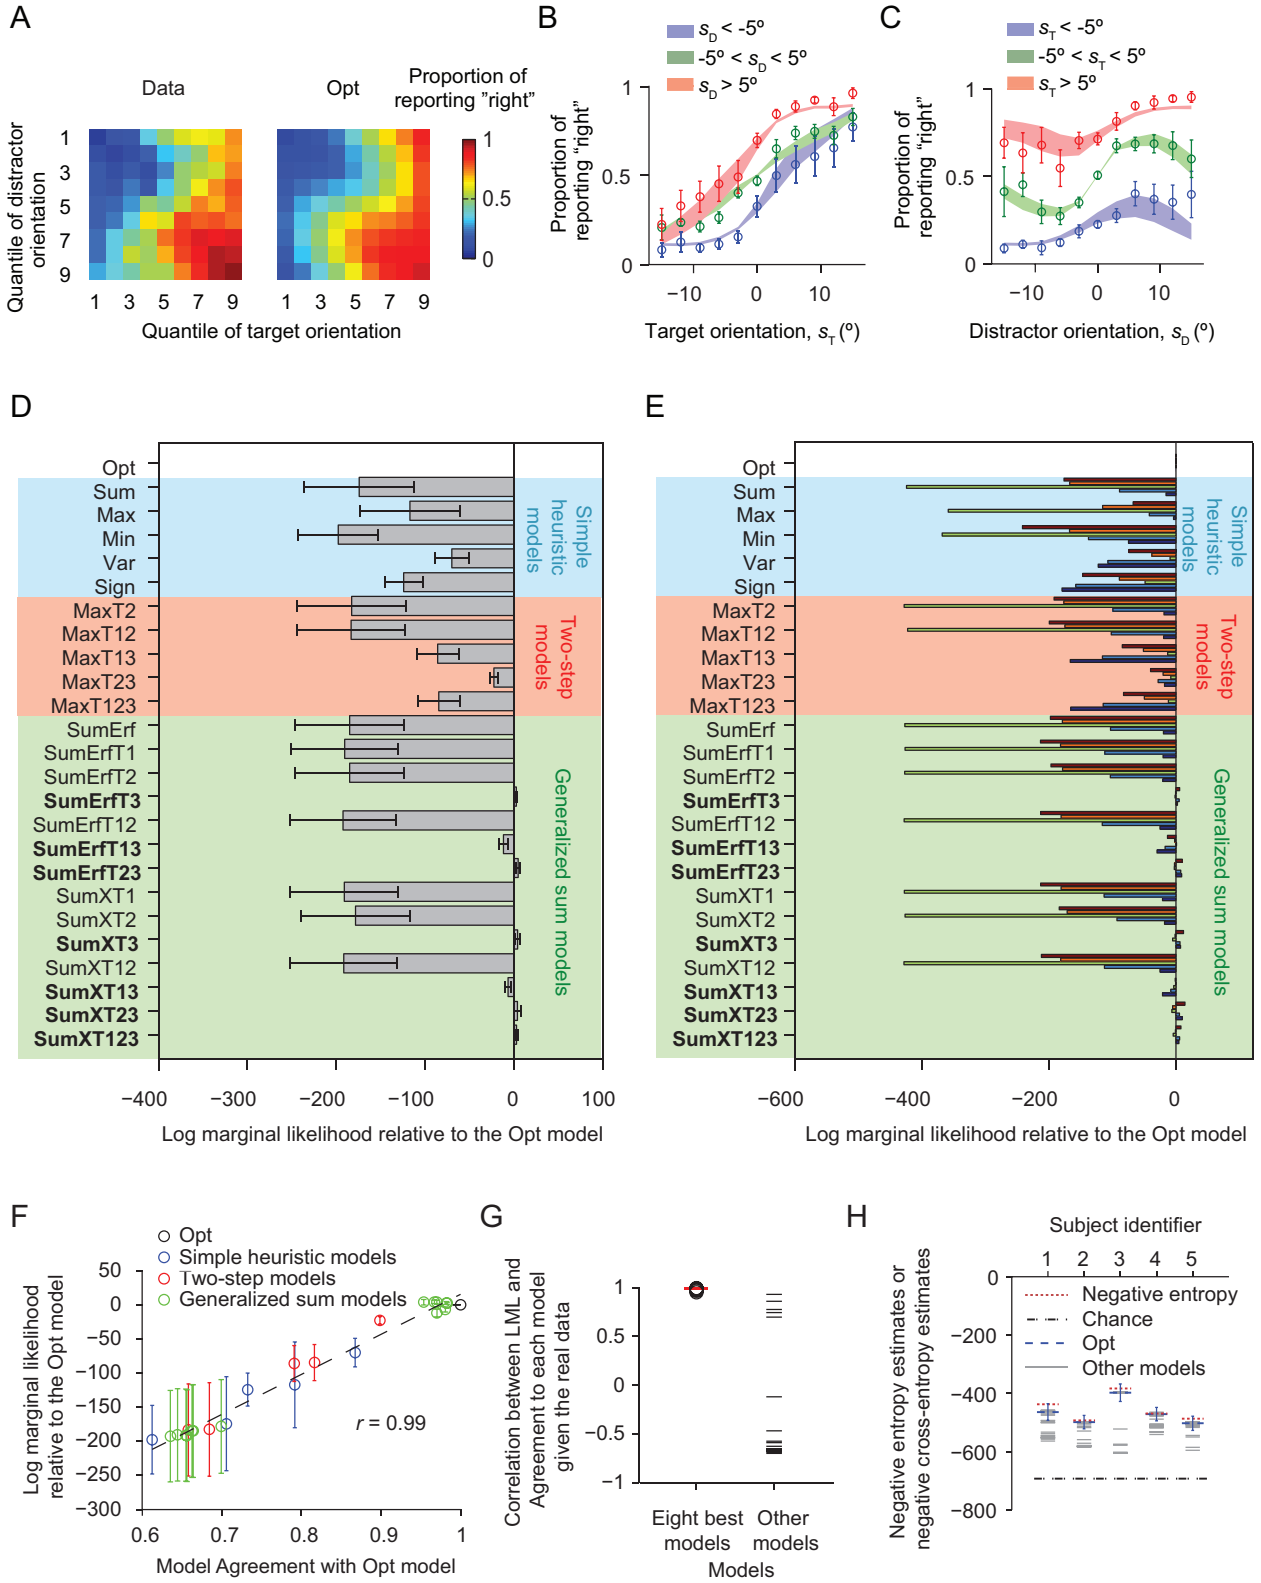


Figure 11 (opposite).

Although we were able to define “complexity” for our mathematically specified decision rules, the way in which we did so was rather arbitrary. Coming up with a general definition of the complexity of a mapping between internal representation and decision might be even more difficult, for several reasons. First, it is unclear how to define complexity when the decision rule does not admit a neat mathematical description, as is the case for MAP estimation when the task distributions are empirical (Griffiths & Tenenbaum, 2006). Second, a mathematically complex computational-level rule might be simple for neurons to implement—or conversely. For example, some complex optimization problems can be solved using “simple” network operations (Deneve, 2008; Hopfield & Tank, 1985; Hornik, Stinchcombe, & White, 1989; Nessler, Pfeiffer, Buesing, & Maass, 2013; Nessler, Pfeiffer, & Maass, 2008). These ambiguities in the meaning of the complexity or simplicity of a decision process pose as much of a challenge to those proposing simple heuristics as an organizing principle as they do to us. In fact, such ambiguities may make any debate between optimality and simplicity ill-defined. A better-defined broad question is in what task domains human behavior is indistinguishable from optimal.

Although it seems safe to claim that subjects are not always near-optimal in perceptual tasks, the evidence for that statement is surprisingly weak. First, apparent suboptimality can arise when there are many parameters to learn, the distributions are complex, the parameters vary greatly across trials (Landy, Goutcher, Trommershäuser, & Mamassian, 2007), or the subjects are not fully attending to a stimulus (Morales et al., 2015). However, when given a large number of training trials (Körding & Wolpert, 2004) or a real-life backstory (Seydell, McCann, Trommershäuser, & Knill, 2008), people exhibit near-optimal behavior even under those conditions. When subjects do not perfectly learn the parameters, the optimal model should take into account evolving posteriors over those parameters; subjects might then still be optimal, given the limited information about the parameters they have available (Fiser, Berkes, Orbán, & Lengyel, 2010). Second, unexplained biases that seem suboptimal have been reported when people estimate the direction of moving dots (Jazayeri & Movshon, 2007; Rauber & Treue, 1998; Szpiro, Spering, & Carrasco, 2014); however, such biases might be caused by the inhomogeneous nature of the tuning curves in the population encoding the stimulus (Wei & Stocker, 2015). Third, it has been suggested that people use a probability matching strategy, which is suboptimal (Wozny, Beierholm, & Shams, 2010). However, evidence in that study was mixed and somewhat indirect; in a direct comparison of probability matching and optimal estimation, the latter described human data better (Acuna, Berniker, Fernandes, & Körding, 2015). In our data, when we tested a probability matching version of the Opt model on our data, we found that its log likelihood is lower than that of the original Opt model by 73 ± 23 (see Figure 12), and that the original Opt model is more likely for all nine subjects. Fourth, in some domains of perception, heuristic models have historically been popular. An example is the maximum-of-outputs model in visual search (Nolte & Jaarsma, 1967). However, a review of studies in which that model did well showed that the optimal model described human data as well or better (Ma et al., 2015). Thus, when we restrict ourselves to perceptual studies in which the generative model is well characterized, observers have presumably fully learned the generative




Figure 12. Comparison between probability matching version of the Opt model and the Opt model. Difference in log marginal likelihood between the probability matching model and the Opt model for individual subjects. The last column shows the mean and standard error of the mean of this difference.

model, and proper model comparison has been performed, strong evidence for suboptimality in human behavior seems to be absent. Once recent challenge to this claim involved a task combining an uncertain perceptual judgment with a speeded reaching movement; the claim was that subjects’ behavior obeyed a two-step model (Fleming et al., 2013). This needs to be explored further.

What does our work imply for cognitive decision-making tasks? In some cases, such as predicting the weather, choosing a job, or playing chess, it might forever remain unknown whether people use near-optimal or simple rules, because both optimality and simplicity are hard to define, and people’s prior beliefs, computational constraints, and utility functions are unknown and very difficult to estimate. However, there is a rich arena of cognitive tasks that are restricted, parameterized, and allow for quantitative modeling. Examples include the learning of category boundaries (Ashby & Maddox, 2005), strategies for information gathering (Coenen, Rehder, & Gureckis, 2015; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003; Wason, 1960), estimation of everyday quantities (Griffiths, Chater, Kemp, Perfors, & Tenenbaum, 2010), category-based induction (Osherson, Smith, Wilkie, López, & Shafir, 1990), and intuitive physics (Battaglia, Hamrick, & Tenenbaum, 2013). In all these realms, it is possible and important to develop large numbers of plausible, suboptimal, simple models and to perform analyses similar to the ones we did here. This is, however, rarely done (Bowers & Davis, 2012), placing claims of near-optimality and Bayesian reasoning in such tasks on shaky ground. On the other hand, there is also no evidence in any of these domains that people use simple heuristics as proposed by Gigerenzer. It might be that people are near-optimal in tasks in which probabilities have to be manipulated implicitly rather than explic-

itly (Chen, Ross, & Murphy, 2014; Maloney, Trommershäuser, & Landy, 2007).

Beyond the task or models studied here, our approach might help to establish criteria that must be satisfied before it can be claimed that any one model describes reality: (a) The model should be compared with, and outperform, a large number of alternative models; (b) The better the trial-to-trial predictions of the model agree with that of an alternative model, the better that alternative model should fit; and (c) The KL divergence between the model and the underlying true model should not be estimated to be significantly different from zero.

References

- Acuna, D. E., Berniker, M., Fernandes, H. L., & Kording, K. P. (2015). Using psychophysics to ask if the brain samples or maximizes. *Journal of Vision*, *15*(3), 7. <http://dx.doi.org/10.1167/15.3.7>
- Alais, D., & Burr, D. (2004). The ventriloquist effect results from near-optimal bimodal integration. *Current Biology*, *14*, 257–262. <http://dx.doi.org/10.1016/j.cub.2004.01.029>
- Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology*, *56*, 149–178. <http://dx.doi.org/10.1146/annurev.psych.56.091103.070217>
- Baldassi, S., & Burr, D. C. (2000). Feature-based integration of orientation signals in visual search. *Vision Research*, *40*, 1293–1300. [http://dx.doi.org/10.1016/S0042-6989\(00\)00029-8](http://dx.doi.org/10.1016/S0042-6989(00)00029-8)
- Baldassi, S., & Verghese, P. (2002). Comparing integration rules in visual search. *Journal of Vision*, *2*, 559–570. <http://dx.doi.org/10.1167/2.8.3>
- Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, *110*, 18327–18332. <http://dx.doi.org/10.1073/pnas.1306572110>
- Borg, I., & Groenen, P. J. F. (2005). *Modern multidimensional scaling* (Vol. 94). New York: Springer Series in Statistics.
- Bowers, J. S., & Davis, C. J. (2012). Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin*, *138*, 389–414. <http://dx.doi.org/10.1037/a0026450>
- Chater, N. (1996). Reconciling simplicity and likelihood principles in perceptual organization. *Psychological Review*, *103*, 566–581. <http://dx.doi.org/10.1037/0033-295X.103.3.566>
- Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences*, *10*, 287–291. <http://dx.doi.org/10.1016/j.tics.2006.05.007>
- Chen, S. Y., Ross, B. H., & Murphy, G. L. (2014). Implicit and explicit processes in category-based induction: Is induction best when we don't think? *Journal of Experimental Psychology: General*, *143*, 227–246. <http://dx.doi.org/10.1037/a0032064>
- Coenen, A., Rehder, B., & Gureckis, T. (2015). Strategies to intervene on causal systems are adaptively selected. *Cognitive Psychology*, *79*, 102–133.
- Collett, D. (2002). *Modelling binary data* (2nd ed.). Boca Raton, London, New York, Washington, DC: CRC Press.
- Cover, T. M., & Thomas, J. A. (2005). *Elements of information theory*. Los Altos, CA: John Wiley & Sons.
- Deneve, S. (2008). Bayesian spiking neurons II: Learning. *Neural Computation*, *20*, 118–145. <http://dx.doi.org/10.1162/neco.2008.20.1.118>
- Dienes, P. Z., & McLeod, P. (1996). Do fielders know where to go to catch the ball, or only how to get there? *Journal of Experimental Psychology: Human Perception and Performance*, *22*, 531–543.
- Eckstein, M. P. (1998). The lower visual search efficiency for conjunctions is due to noise and not serial attentional processing. *Psychological Science*, *9*, 111–118. <http://dx.doi.org/10.1111/1467-9280.00020>
- Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*, 429–433. <http://dx.doi.org/10.1038/415429a>
- Faisal, A. A., & Wolpert, D. M. (2009). Near optimal combination of sensory and motor uncertainty in time during a naturalistic perception-action task. *Journal of Neurophysiology*, *101*, 1901–1912. <http://dx.doi.org/10.1152/jn.90974.2008>
- Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: From behavior to neural representations. *Trends in Cognitive Sciences*, *14*, 119–130. <http://dx.doi.org/10.1016/j.tics.2010.01.003>
- Fleming, S. M., Maloney, L. T., & Daw, N. D. (2013). The irrationality of categorical perception. *The Journal of Neuroscience*, *33*, 19060–19070. <http://dx.doi.org/10.1523/JNEUROSCI.1263-13.2013>
- Geisler, W. S. (2011). Contributions of ideal observer theory to vision research. *Vision Research*, *51*, 771–781. <http://dx.doi.org/10.1016/j.visres.2010.09.027>
- Geisler, W. S., & Perry, J. S. (2009). Contour statistics in natural images: Grouping across occlusions. *Visual Neuroscience*, *26*, 109–121. <http://dx.doi.org/10.1017/S0952523808080875>
- Gigerenzer, G. (2004). Striking a blow for sanity in theories of rationality. In M. Augier & J. G. March (Eds.), *Models of a man: Essays in memory of Herbert A. Simon* (pp. 389–409). Cambridge, MA: MIT Press.
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, *62*, 451–482. <http://dx.doi.org/10.1146/annurev-psych-120709-145346>
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650–669. <http://dx.doi.org/10.1037/0033-295X.103.4.650>
- Grassberger, P. (1988). Finite sample corrections to entropy and dimension estimates. *Physics Letters, A*, *128*, 369–373. [http://dx.doi.org/10.1016/0375-9601\(88\)90193-4](http://dx.doi.org/10.1016/0375-9601(88)90193-4)
- Grassberger, P. (2003). Entropy estimates from insufficient samplings. <http://arXiv:physics/0307138>
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. Los Altos, CA: John Wiley & Sons.
- Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in Cognitive Sciences*, *14*, 357–364. <http://dx.doi.org/10.1016/j.tics.2010.05.004>
- Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science*, *17*, 767–773. <http://dx.doi.org/10.1111/j.1467-9280.2006.01780.x>
- Gu, Y., Angelaki, D. E., & DeAngelis, G. C. (2008). Neural correlates of multisensory cue integration in macaque MSTd. *Nature Neuroscience*, *11*, 1201–1210. <http://dx.doi.org/10.1038/nn.2191>
- Hamming, R. W. (1950). Error detecting and error correcting codes. *The Bell System Technical Journal*, *29*, 147–160. <http://dx.doi.org/10.1002/j.1538-7305.1950.tb00463.x>
- Hopfield, J. J., & Tank, D. W. (1985). “Neural” computation of decisions in optimization problems. *Biological Cybernetics*, *52*, 141–152.
- Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, *2*, 359–366. [http://dx.doi.org/10.1016/0893-6080\(89\)90020-8](http://dx.doi.org/10.1016/0893-6080(89)90020-8)
- Jazayeri, M., & Movshon, J. A. (2007). A new perceptual illusion reveals mechanisms of sensory decoding. *Nature*, *446*, 912–915. <http://dx.doi.org/10.1038/nature05739>
- Jones, M., & Love, B. C. (2011). Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *The Behavioral and Brain Sciences*, *34*, 169–188; discussion 188–231.
- 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. <http://dx.doi.org/10.1371/journal.pone.0000943>

- Körding, K. P., & Wolpert, D. M. (2004). The loss function of sensorimotor learning. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, *101*, 9839–9842. <http://dx.doi.org/10.1073/pnas.0308394101>
- Kramer, P., Graham, N., & Yager, D. (1985). Simultaneous measurement of spatial-frequency summation and uncertainty effects. *Journal of the Optical Society of America, A, Optics, Image & Science*, *2*, 1533–1542. <http://dx.doi.org/10.1364/JOSAA.2.001533>
- Landy, M. S., Goutcher, R., Trommershäuser, J., & Mamassian, P. (2007). Visual estimation under risk. *Journal of Vision*, *7*(6), 4. <http://dx.doi.org/10.1167/7.6.4>
- López-Moliner, J., Field, D. T., & Wann, J. P. (2007). Interceptive timing: Prior knowledge matters. *Journal of Vision*, *7*, 11.1–11.8.
- Ma, W. J. (2010). Signal detection theory, uncertainty, and Poisson-like population codes. *Vision Research*, *50*, 2308–2319. <http://dx.doi.org/10.1016/j.visres.2010.08.035>
- Ma, W. J. (2012). Organizing probabilistic models of perception. *Trends in Cognitive Sciences*, *16*, 511–518. <http://dx.doi.org/10.1016/j.tics.2012.08.010>
- Ma, W. J., & Jazayeri, M. (2014). Neural coding of uncertainty and probability. *Annual Review of Neuroscience*, *37*, 205–220. <http://dx.doi.org/10.1146/annurev-neuro-071013-014017>
- Ma, W. J., Navalpakkam, V., Beck, J. M., van den Berg, R., & Pouget, A. (2011). Behavior and neural basis of near-optimal visual search. *Nature Neuroscience*, *14*, 783–790. <http://dx.doi.org/10.1038/nn.2814>
- Ma, W. J., Shen, S., Dziugaite, G., & van den Berg, R. (2015). Requiem for the max rule? *Vision Research*, *116*, 179–193. <http://dx.doi.org/10.1016/j.visres.2014.12.019>
- Maloney, L. T., & Mamassian, P. (2009). Bayesian decision theory as a model of human visual perception: Testing Bayesian transfer. *Visual Neuroscience*, *26*, 147–155. <http://dx.doi.org/10.1017/S0952523808080905>
- Maloney, L. T., Trommershäuser, J., & Landy, M. S. (2007). Questions without words: A comparison between decision making under risk and movement planning under risk. In W. D. Gray (Ed.), *Integrated models of cognitive systems* (pp. 297–314). New York: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780195189193.003.0021>
- Mazyar, H., van den Berg, R., Seilheimer, R. L., & Ma, W. J. (2013). Independence is elusive: Set size effects on encoding precision in visual search. *Journal of Vision*, *13*(5), 8. <http://dx.doi.org/10.1167/13.5.8>
- McLeod, P., Reed, N., & Dienes, Z. (2003). Psychophysics: How fielders arrive in time to catch the ball. *Nature*, *426*, 244–245. <http://dx.doi.org/10.1038/426244a>
- 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*, 2021–2036. <http://dx.doi.org/10.3758/s13414-015-0897-2>
- Nessler, B., Pfeiffer, M., Buesing, L., & Maass, W. (2013). Bayesian computation emerges in generic cortical microcircuits through spike-timing-dependent plasticity. *PLoS Computational Biology*, *9*(4), e1003037. <http://dx.doi.org/10.1371/journal.pcbi.1003037>
- Nessler, B., Pfeiffer, M., & Maass, W. (2008). Hebbian learning of Bayes optimal decisions. *Proceedings of the 21st Conference on Advances in Neural Information Processing Systems; December 2008*. Vancouver, Canada: NIPS 2008.
- Nolte, L. W., & Jaarsma, D. (1967). More on the detection of one of M orthogonal signals. *The Journal of the Acoustical Society of America*, *41*, 497. <http://dx.doi.org/10.1121/1.1910360>
- Norris, D. (2006). The Bayesian reader: Explaining word recognition as an optimal Bayesian decision process. *Psychological Review*, *113*, 327–357. <http://dx.doi.org/10.1037/0033-295X.113.2.327>
- Osherson, D. N., Smith, E. E., Wilkie, O., López, A., & Shafir, E. (1990). Category-based induction. *Psychological Review*, *97*, 185–200. <http://dx.doi.org/10.1037/0033-295X.97.2.185>
- Palmer, J., Verghese, P., & Pavel, M. (2000). The psychophysics of visual search. *Vision Research*, *40*, 1227–1268. [http://dx.doi.org/10.1016/S0042-6989\(99\)00244-8](http://dx.doi.org/10.1016/S0042-6989(99)00244-8)
- Pouget, A., Beck, J. M., Ma, W. J., & Latham, P. E. (2013). Probabilistic brains: Knowns and unknowns. *Nature Neuroscience*, *16*, 1170–1178. <http://dx.doi.org/10.1038/nn.3495>
- Qamar, A. T., Cotton, R. J., George, R. G., Beck, J. M., & Prezhdo, E., Laudano, A., . . . Ma, W. J. (2013). Trial-to-trial, uncertainty-based adjustment of decision boundaries in visual categorization. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, *110*, 20332–20337.
- Rauber, H. J., & Treue, S. (1998). Reference repulsion when judging the direction of visual motion. *Perception*, *27*, 393–402. <http://dx.doi.org/10.1068/p270393>
- Seydell, A., McCann, B. C., Trommershäuser, J., & Knill, D. C. (2008). Learning stochastic reward distributions in a speeded pointing task. *The Journal of Neuroscience*, *28*, 4356–4367. <http://dx.doi.org/10.1523/JNEUROSCI.0647-08.2008>
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*, 129–138. <http://dx.doi.org/10.1037/h0042769>
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society Series B, Statistical Methodology*, *64*, 583–639. <http://dx.doi.org/10.1111/1467-9868.00353>
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E. J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*, 453–489. http://dx.doi.org/10.1207/s15516709cog2703_6
- Szpiro, S. F., Spering, M., & Carrasco, M. (2014). Perceptual learning modifies untrained pursuit eye movements. *Journal of Vision*, *14*, 8. <http://dx.doi.org/10.1167/14.8.8>
- Trommershäuser, J., Körding, K. P., & Landy, M. S. (2011). *Sensory cue integration*. New York, NY: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780195387247.001.0001>
- Wason, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *The Quarterly Journal of Experimental Psychology*, *12*, 129–140. <http://dx.doi.org/10.1080/17470216008416717>
- Wei, X.-X., & Stocker, A. A. (2015). A Bayesian observer model constrained by efficient coding can explain “anti-Bayesian” percepts. *Nature Neuroscience*, *18*, 1509–1517. <http://dx.doi.org/10.1038/nn.4105>
- Wichmann, F. A., & Hill, N. J. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception & Psychophysics*, *63*, 1293–1313. <http://dx.doi.org/10.3758/BF03194544>
- Wozny, D. R., Beierholm, U. R., & Shams, L. (2010). Probability matching as a computational strategy used in perception. *PLoS Computational Biology*, *6*(8), e1000871. <http://dx.doi.org/10.1371/journal.pcbi.1000871>

(Appendix follows)

Appendix

Decision Rules of All Models

Here, we give for every model the condition for which the observer reports “rightward.”

Optimal Model

$$\text{Opt: } \sum_{L=1}^N \text{erf} \frac{\frac{x_L}{\sigma^2}}{\sqrt{2\left(\frac{1}{\sigma^2} + \frac{1}{\sigma_s^2}\right)}} \exp\left(-\frac{x_L^2}{2(\sigma_s^2 + \sigma^2)} - \frac{\bar{x}_L^2}{2\left(\sigma_s^2 + \frac{\sigma^2}{N-1}\right)} - \frac{N-1}{2\sigma^2} \text{Var } \mathbf{x}_L\right) > 0,$$

Simple Heuristic Models

$$\text{Sum: } \sum_{L=1}^N x_L > 0$$

$$\text{Max: } x_{\text{argmax}_L |x_L|} > 0$$

$$\text{Min: } x_{\text{argmin}_L |x_L|} > 0$$

$$\text{Var: } x_{\text{argmax}_L \text{Var } \mathbf{x}_L} > 0$$

Two-Step Models

$$\text{MaxT2: } x_{\text{argmax}_L (-\bar{x}_L^2)} > 0$$

$$\text{MaxT12: } x_{\text{argmax}_L \left(-\frac{x_L^2}{2(\sigma_s^2 + \sigma^2)} - \frac{\bar{x}_L^2}{2\left(\sigma_s^2 + \frac{\sigma^2}{N-1}\right)}\right)} > 0$$

$$\text{MaxT13: } x_{\text{argmax}_L \left(-\frac{x_L^2}{2(\sigma_s^2 + \sigma^2)} - \frac{N-1}{2\sigma^2} \text{Var } \mathbf{x}_L\right)} > 0$$

$$\text{MaxT23: } x_{\text{argmax}_L \left(-\frac{\bar{x}_L^2}{2\left(\sigma_s^2 + \frac{\sigma^2}{N-1}\right)} - \frac{N-1}{2\sigma^2} \text{Var } \mathbf{x}_L\right)} > 0$$

$$\text{MaxT123: } x_{\text{argmax}_L \left(-\frac{x_L^2}{2(\sigma_s^2 + \sigma^2)} - \frac{\bar{x}_L^2}{2\left(\sigma_s^2 + \frac{\sigma^2}{N-1}\right)} - \frac{N-1}{2\sigma^2} \text{Var } \mathbf{x}_L\right)} > 0$$

Generalized Sum Models

$$\text{SumErf: } \sum_{L=1}^N \text{erf} \frac{\frac{x_L}{\sigma^2}}{\sqrt{2\left(\frac{1}{\sigma^2} + \frac{1}{\sigma_s^2}\right)}} > 0$$

$$\text{SumErfT1: } \sum_{L=1}^N \text{erf} \frac{\frac{x_L}{\sigma^2}}{\sqrt{2\left(\frac{1}{\sigma^2} + \frac{1}{\sigma_s^2}\right)}} \exp\left(-\frac{x_L^2}{2(\sigma_s^2 + \sigma^2)}\right) > 0,$$

(Appendix continues)

$$\text{SumErfT2: } \sum_{L=1}^N \operatorname{erf} \frac{x_L}{\sigma^2} \exp \left(-\frac{\bar{x}_L^2}{2(\sigma_s^2 + \frac{\sigma^2}{N-1})} \right) > 0,$$

$$\text{SumErfT3: } \sum_{L=1}^N \operatorname{erf} \frac{x_L}{\sigma^2} \exp \left(-\frac{N-1}{2\sigma^2} \operatorname{Var} \mathbf{x}_L \right) > 0,$$

$$\text{SumErfT12: } \sum_{L=1}^N \operatorname{erf} \frac{x_L}{\sigma^2} \exp \left(-\frac{x_L^2}{2(\sigma_s^2 + \sigma^2)} - \frac{\bar{x}_L^2}{2(\sigma_s^2 + \frac{\sigma^2}{N-1})} \right) > 0,$$

$$\text{SumErfT13: } \sum_{L=1}^N \operatorname{erf} \frac{x_L}{\sigma^2} \exp \left(-\frac{x_L^2}{2(\sigma_s^2 + \sigma^2)} - \frac{N-1}{2\sigma^2} \operatorname{Var} \mathbf{x}_L \right) > 0,$$

$$\text{SumErfT23: } \sum_{L=1}^N \operatorname{erf} \frac{x_L}{\sigma^2} \exp \left(-\frac{\bar{x}_L^2}{2(\sigma_s^2 + \frac{\sigma^2}{N-1})} - \frac{N-1}{2\sigma^2} \operatorname{Var} \mathbf{x}_L \right) > 0,$$

$$\text{SumXT1: } \sum_{L=1}^N x_L \exp \left(-\frac{x_L^2}{2(\sigma_s^2 + \sigma^2)} \right) > 0,$$

$$\text{SumXT2: } \sum_{L=1}^N x_L \exp \left(-\frac{\bar{x}_L^2}{2(\sigma_s^2 + \frac{\sigma^2}{N-1})} \right) > 0,$$

$$\text{SumXT3: } \sum_{L=1}^N x_L \exp \left(-\frac{N-1}{2\sigma^2} \operatorname{Var} \mathbf{x}_L \right) > 0,$$

$$\text{SumXT12: } \sum_{L=1}^N x_L \exp \left(-\frac{x_L^2}{2(\sigma_s^2 + \sigma^2)} - \frac{\bar{x}_L^2}{2(\sigma_s^2 + \frac{\sigma^2}{N-1})} \right) > 0,$$

$$\text{SumXT13: } \sum_{L=1}^N x_L \exp \left(-\frac{\bar{x}_L^2}{2(\sigma_s^2 + \frac{\sigma^2}{N-1})} - \frac{N-1}{2\sigma^2} \operatorname{Var} \mathbf{x}_L \right) > 0,$$

$$\text{SumXT23: } \sum_{L=1}^N x_L \exp \left(-\frac{\bar{x}_L^2}{2(\sigma_s^2 + \frac{\sigma^2}{N-1})} - \frac{N-1}{2\sigma^2} \operatorname{Var} \mathbf{x}_L \right) > 0,$$

$$\text{SumXT123: } \sum_{L=1}^N x_L \exp \left(-\frac{x_L^2}{2(\sigma_s^2 + \sigma^2)} - \frac{\bar{x}_L^2}{2(\sigma_s^2 + \frac{\sigma^2}{N-1})} - \frac{N-1}{2\sigma^2} \operatorname{Var} \mathbf{x}_L \right) > 0,$$

(Appendix continues)

Table A1
Complexity of Observer Decision Rules

Model	Arithmetic operations				Neural operations					
	Q	L	NL	Sort	$\mathbf{a} \cdot \mathbf{r}$	$\mathbf{b} \cdot \mathbf{r}$	Q	L	NL	Sort
Opt	32	59	8	1	164	656	724	183	20	1
Sum	0	3	0	1	4	16	15	3	1	1
Max	0	0	4	2	4	4	0	0	4	2
Min	0	0	4	2	4	4	0	0	4	2
Var	28	40	0	2	48	144	164	44	4	2
Sign	0	0	0	4	48	172	188	52	4	2
MaxT2	28	40	0	2	112	412	456	124	4	2
MaxT12	32	48	0	2	104	396	440	108	4	2
MaxT13	32	48	0	2	80	316	356	84	4	2
MaxT23	28	40	0	2	160	652	720	172	4	2
MaxT123	32	48	0	2	4	4	0	11	12	1
SumErf	0	7	4	1	12	16	12	19	20	1
SumErfT1	8	11	8	1	52	176	192	63	20	1
SumErfT2	28	51	8	1	52	148	168	55	20	1
SumErfT3	28	51	8	1	116	416	460	135	20	1
SumErfT12	32	59	8	1	108	400	444	119	20	1
SumErfT13	32	59	8	1	84	328	360	95	20	1
SumErfT23	28	51	8	1	12	16	12	11	12	1
SumXT1	8	7	4	1	52	176	192	55	12	1
SumXT2	28	47	4	1	52	148	162	47	12	1
SumXT3	28	47	4	1	116	416	460	127	12	1
SumXT12	32	55	4	1	108	400	444	111	12	1
SumXT13	32	55	4	1	84	320	360	87	12	1
SumXT23	28	47	4	1	164	656	724	175	12	1
SumXT123	32	55	4	1	164	656	724	183	20	1

Note. The left part of the table shows the number of arithmetic operations of different types in each decision rule: linear operations (L), quadratic operations (Q), other nonlinear operations (NL), and sorting operations (Sort). The right part of the table shows the numbers of neural operations if the decision variable were directly transformed into a neural quantity using the theory of probabilistic population codes. All models have only two free parameters; only the complexity of the observer's computation differs.

(Appendix continues)

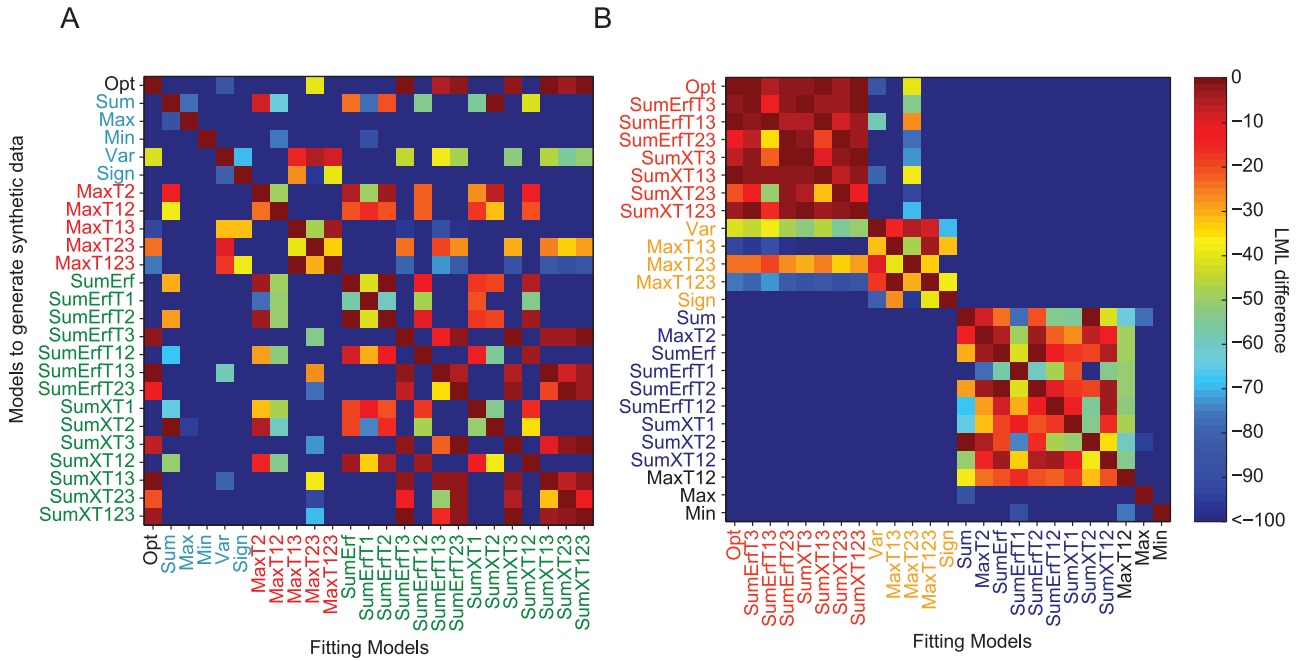


Figure A1. Model recovery analysis. We tested how well synthetic data sets generated from each model (rows) were fitted by each model (columns). (A) Model confusion matrix. The color in a cell represents the difference in log marginal likelihood between a model and the winning model for the corresponding data set. Dark red on the diagonal means that the model used to generate the data was found to be most likely. (B) As in (A), but with models clustered by Agreement. Models in red are models with high Agreement to the Opt model. Models in blue are from a different model set with low Agreement to the Opt model, but similar to each other. Models in orange have higher Agreement to the Opt than to the blues models, but they are still well distinguishable from the Opt model. Also refer to Fig. 8C. See the online article for the color version of this figure.

(Appendix continues)

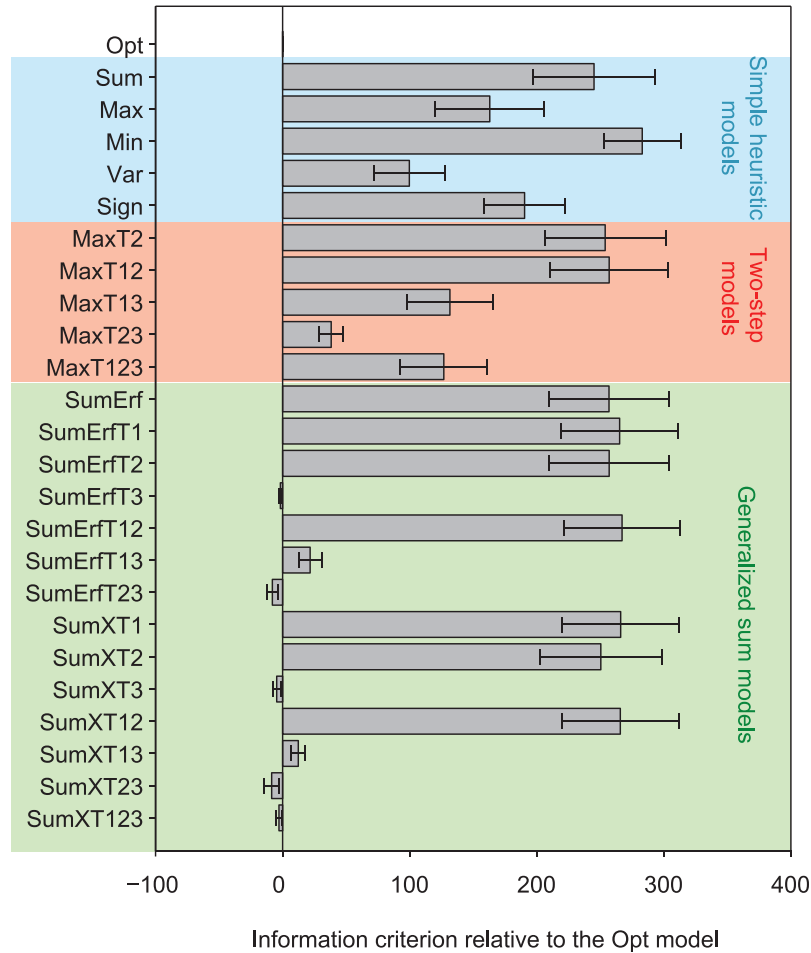


Figure A2. Model comparison based on information criteria. Mean and standard error of the mean across subjects of the difference in information criterion (a small-sample variant of the Akaike Information Criterion or the Bayesian Information Criterion) between each model and the Opt model. Note that all models have two parameters, therefore all information criteria yield the same differences between models. See the online article for the color version of this figure.

(Appendix continues)




Figure A3. Correlation between log marginal likelihood and Agreement with any one model given the real data. Related to Fig. 8D and Fig. 9A. Each plot shows, given the real data, the mean (open circle) and standard error of the mean (error bar) across subjects of a model's average log marginal likelihood per trial as a function of its Agreement with a reference model; the reference model differs between plots. The dashed line represents the best linear fit. r is the Pearson correlation. The names of eight best models are in boldface. See the online article for the color version of this figure.

(Appendix continues)




Figure A4. Correlation between log marginal likelihood and Agreement with the Opt model given synthetic data generated from any one model. Related to Fig. 9B. For each plot, we generated 9 synthetic data sets from a different generating model. The plot shows the mean (open circle) and standard error of the mean (error bar) of a model's average log marginal likelihood per trial as a function of its Agreement with the Opt model based on those data sets. The dashed line represents the best linear fit. r is the Pearson correlation. The names of eight best models are in boldface. Given synthetic data generated from one of the eight best models, the correlation is high. Given synthetic data generated from a model outside of the eight best, the correlation is low. See the online article for the color version of this figure.

(Appendix continues)




Figure A5. Correlation between log marginal likelihood and Agreement with a reference model given synthetic data generated from that reference model. Related to Fig. 9C. For each plot, we generated 9 synthetic data sets from a different generating model. The plot shows the mean (open circle) and standard error of the mean (error bar) of a model’s average log marginal likelihood per trial as a function of its Agreement with the generating model. The dashed line represents the best linear fit. r is the Pearson correlation. Given synthetic data generated from any one model, the correlation is high. See the online article for the color version of this figure.

(Appendix continues)

Table A2

For All Models and All Subjects, the Probability That the Negative Cross-Entropy is Equal to or Higher Than the Negative Entropy

Model	1	2	3	4	5	6	7	8	9
Opt ^a	.97*	.2*	.14*	.17*	.05*	.92*	.071*	.099*	.015
Sum ^b	<10 ⁻¹⁶	<10 ⁻³	<10 ⁻¹¹	<10 ⁻¹⁶	.0045	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻³
Max ^b	<10 ⁻⁹	.034	<10 ⁻⁵	<10 ⁻¹⁶	.069*	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	.01
Min ^b	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶
Var ^b	.89*	<10 ⁻⁷	.22*	.043	<10 ⁻¹⁶	.98*	<10 ⁻⁴	.15*	<10 ⁻¹⁶
Sign ^b	<10 ⁻³	.062*	.0011	<10 ⁻⁵	.15*	<10 ⁻⁶	<10 ⁻⁵	<10 ⁻¹⁶	.005
MaxT2 ^b	<10 ⁻¹⁶	<10 ⁻⁴	<10 ⁻¹³	<10 ⁻¹⁶	<10 ⁻³	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻⁶
MaxT12 ^b	<10 ⁻¹⁶	<10 ⁻⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻⁴	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻⁸
MaxT13 ^b	.73*	<10 ⁻¹⁴	.023	.0053	<10 ⁻¹⁶	.96*	<10 ⁻⁹	.12*	<10 ⁻¹⁶
MaxT23 ^b	.94*	.073*	.28*	.12*	<10 ⁻⁴	.97*	<10 ⁻³	.15*	<10 ⁻⁷
MaxT123 ^b	.82*	<10 ⁻¹⁰	.067*	.015	<10 ⁻¹⁶	.98*	<10 ⁻⁷	.12*	<10 ⁻¹⁶
SumErf ^b	<10 ⁻¹⁶	<10 ⁻⁴	<10 ⁻¹¹	<10 ⁻¹⁶	<10 ⁻³	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻⁵
SumErfT1	<10 ⁻¹⁶	<10 ⁻⁴	<10 ⁻¹⁶	<10 ⁻¹⁶	.0075	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻⁷
SumErfT2	<10 ⁻¹⁶	<10 ⁻⁴	<10 ⁻¹²	<10 ⁻¹⁶	<10 ⁻³	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻⁵
SumErfT3 ^a	.96*	.47*	.16*	.33*	.068*	.93*	.13*	.15*	.024*
SumErfT12	<10 ⁻¹⁶	<10 ⁻¹⁰	<10 ⁻¹⁴	<10 ⁻¹⁶	<10 ⁻⁴	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻⁷
SumErfT13 ^a	.96*	.096*	.18*	.29*	<10 ⁻⁷	.96*	.027	.19*	<10 ⁻⁵
SumErfT23 ^a	.95*	.5*	.15*	.33*	.37*	.91*	.13*	.15*	.16*
SumXT1	<10 ⁻¹⁶	<10 ⁻⁴	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻³	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻⁶
SumXT2	<10 ⁻¹⁶	<10 ⁻³	<10 ⁻¹²	<10 ⁻¹⁶	.0016	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻⁴
SumXT3 ^a	.95*	.52*	.15*	.32*	.18*	.89*	.15*	.13*	.072*
SumXT12	<10 ⁻¹⁶	<10 ⁻⁵	<10 ⁻¹⁴	<10 ⁻¹⁶	<10 ⁻⁴	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻¹⁶	<10 ⁻⁷
SumXT13 ^a	.96*	.25*	.17*	.33*	.0022*	.92*	.073*	.15*	.0016
SumXT23 ^a	.93*	.46*	.12*	.3*	.41*	.86*	.1*	.12*	<10 ⁻³
SumXT123 ^a	.95*	.5*	.15*	.33*	.19*	.9*	.15*	.13*	.065*

Note. Each column corresponds to a subject. For the eight best models, most subjects have $p > .05$, indicating these models are good in an "absolute" sense.

^a The eight best models. ^b Simple heuristic models and two-step models.

* $p > .05$.

Received March 9, 2015
 Revision received January 4, 2016
 Accepted January 6, 2016 ■