Bayesian Inference Model for Performing Same/Different Task

David Krueger, Wei Ji Ma, Krešimir Josić

June 22, 2009

1 Introduction

Perception is often described as inferring the state of the world from noisy sensory data. Human perception has been shown to do this inference Bayes-Optimally for several simple perception tasks, such as cue combination (Whitley and Sahani 2008). Recent work has started to test the optimality of human perception in more complex tasks, involving multiple stimuli and structural inference. We attempt to test optimality of human performance in a same/different task using oriented Gabors (Task 2). While same/different tasks are relatively simple, optimal performance nonetheless involves structural inference, multiple stimuli, and a non-trivial decision rule.

We derived a Optimal Bayesian Observer (OBO) model for this task and ran 10 subjects. The OBO uses a Maximizing strategy; it maximizes the expected percent correct given the available information. Our model has several parameters, one of which is the observer's internal measure of uncertainty, σ, which should reflect the noisiness of observations. We designed two more tasks to try to fit this parameter for our model in the second part, Task 1A and Task 1B. Each subject ran only one Task 1.

Our first group of 7 subjects ran Task 1A. Task 1A had the same stimuli presentation as Task 2. But the task was to judge which Gabor (left or right) was oriented further clockwise. This was unambiguous because we instructed subjects that judgments were to be based on the smaller angle, and our the difference in orientation ranged from $-16.5°$ to $16.5°$. But our results indicated that it may have been too confusing or too difficult for most subjects anyhow. Task 1B was intended to be an easier task, but still similar enough to Task 2 that the value of σ would be comparable. Task 1B had only one Gabor presented at a time and the task was to judge whether the Gabor was tilted to the left (counterclockwise) or the right (clockwise) of a reference line. A group of 3 subjects ran Task 1B.

In our data analysis, we noticed several trends in the data which were in contrast with the OBO model. This lead us to include a Sampling model, with a non-deterministic decision rule. Comparing the results of both models, most subjects appeared to be following a strategy which was some combination of or compromise between the two models. So we created a mixed model which is a compromise between Sampling and Maximizing. We used Bayesian Model Comparison to compare our different models of the same/different Task.

1.1 Notation and Definitions

We say $P(x) = \mathcal{N}(x \,|\mu, \sigma^2)$ when x is normally distributed with mean μ and variance σ^2.

In both parts, we call the orientation of the left Gabor s_1 and the right s_2. In our models, we assume the observer forms noisy internal representations of the Gabors and call these orientations x_1 and x_2.

We use the following random variables:

$$\Delta_s = s_2 - s_1$$
$$\Delta_x = x_2 - x_1$$

We let x_1 and x_2 be normally distributed about s_1 and s_2 with variance σ_1^2 and σ_2^2, and let $\sigma = \sqrt{\sigma_1^2 + \sigma_2^2}$ so that:

$$\left\{\begin{array}{l} P(x_1 \,|\, s_1) = \mathcal{N}(x_1 \,|\, s_1, \sigma_1^2) \\ P(x_2 \,|\, s_2) = \mathcal{N}(x_2 \,|\, s_2, \sigma_2^2) \end{array}\right\} \Longrightarrow P(\Delta_x \,|\, \Delta_s) = \mathcal{N}(\Delta_x \,|\, \Delta_s, \sigma^2).$$

We also assume that $\sigma_1 = \sigma_2$, so then $\sigma = \sigma_1\sqrt{2}$.

In Task 1A/1B, the random variable C_1 is defined to equal 1 when "left" is the correct response, and 2 when "right" is the correct response, while R_1 is the subject or observer's response. In Task 2, the same/different task, we have $C_2 = 1$ for same trials, $C_2 = 2$ for different trials and $R_2 = 1$ when the subject or observer's response is "same" and $R_2 = 2$ when the response if "different".

2 Methods

2.1 Experiment Design

We ran two different versions of the experiment. First, we ran 7 subjects on a version of the experiment consisting of Task 1A followed by Task 2. Also, for this group of subjects, there was a bug in the code which caused all the "random" numbers to be the same for each subject. We then fixed the bug and ran 3 subjects on a version of the experiment consisting of Task 1B followed by Task 2. Subjects ran Task 1A or 1B first, and then Task 2.

2.2 Instructions to Subjects

There was a demo of two trials at the start of Task 1A/1B. The demo trials were exactly like practice trials, except that the Gabors were present until the subject responded, and then the message saying whether their response was correct or not remained on the screen until they pressed the space bar. An experimenter went through the demo with subjects. The demo was intended to show subjects the presentation and the experimenter pointed out the line patches and fixation cross on the screen when giving instructions to avoid confusion. The demo was also intended to ensure that subjects understood the task, since making judgments about which stimulus was further clockwise seemed ambiguous to some subjects.

All of the subjects were instructed:

1) That they would be asked to make judgments about the orientations of line patches.
2) That the position of the patches was irrelevant
3) To try to keep their vision focused on the fixation cross
4) That the experiment would last 2 hours including a 30 minute mandatory break

Subjects were then instructed about Task 1A/1B.
The task was explained and subjects were told that there would be 20 practice trials followed by 120 real trials, and that the practice trials were different in that they would receive feedback after each trial, while they would only know how well they did in the real trials upon completing all 120. They were also told that in the real trials, the patches would be on the screen for a much shorter time.
Subjects were asked to remain focused and do as well as they could. They were told they could take breaks any time if they felt like they were losing focus. Finally, they were told to see the experimenter for instructions on Task 2 after completing the first task and asked if they had any questions about anything relating to the experiment or task.

Before doing Task 2, subjects were instructed that their task was to report if the orientations of the patches were the same of different and reminded that position was irrelevant. They were told that Task 2 had 6 blocks, each with 40 practice trials followed by 200 real trials, and each block had a different difficulty level. They were told they could take breaks in between blocks or anytime they felt they were losing focus. They were instructed to take a 30 minute break after 3 blocks. They were reminded of the differences between practice and real trials, and reminded to stay focused, keep their vision focused on the fixation cross, and to do as well as they could.

Subjects doing Task 1A were also told that for about half the trials of Task 2, the orientations would be different and about half would be the same. Subjects doing Task 1B were not instructed of this.

Most of the instructions were also displayed on the screen before subjects performed the tasks. Reminder instructions were printed to the screen between blocks of Task 2. After 3 blocks of Experiment 2, subjects were asked to take a break, and the screen turned blue/green so that they would recognize and remember to do so (otherwise, they might have ignored the reminder, thinking it was the same as the reminder message between each block). For details see the experiment code (the version with Task 1A is experimentoneI.mat, the version with Task 1B is expetimentoneK.mat).

2.3 Experiment Design

Everything we told subjects was true. For Tasks 1A/1B, there was a demo with 2 trials, followed by 20 practice trials and 120 real trials. There were 6 blocks of trials for task 2, each with 40 practice trials (20 same / 20 different, randomly permuted) and 200 real trials (100 same / 100 different, randomly permuted) with breaks between blocks. In the practice trials for all three tasks, the Gabors were presented for TIME ms as opposed to TIME ms in the real trials. The value of σ_Δ was different for each block. The order of the blocks was randomized.

Subjects were allowed to take breaks between blocks, or any time they felt they were losing focus, and we forced them to take a 30 minute break after completing the third block of Task 2.

We used Matlab and Psychtoolbox to create the display and record subjects responses. There is a fixation cross in the center of the screen the subject is asked to focus on; the gabors appear on either side of the cross(see figure 1, below).

FIGURE

The subject responds by pressing left and right arrow keys in Tasks 1A/1B, corresponding to which Gabor they believe is further clockwise, and by pressing "s" and "d" keys in Task 2 for same and different, respectively.

For Tasks 1A and 2, each trial consisted of the following:

1) The fixation cross is present for TIME ms before the Gabors appear.
2) The Gabors are presented for TIME ms, the fixation cross stays on the screen.
3) The Gabors disappear, prompting the subject to respond.

4) The subject enters their response and the fixation cross flashes off and on (TIME ms), so the subject knows their response is registered.

For Task 1B, each trial consisted of the following:

1) The fixation cross and reference lines are present for TIME ms before the Gabor appears.
2) The Gabor is presented for TIME ms, the fixation cross and reference lines stays on the screen.
3) The Gabor disappears, prompting the subject to respond.
4) The subject enters their response and the fixation cross flashes off and on (TIME ms), so the subject knows their response is registered.

The orientations of the left Gabor in both parts is chosen randomly from a uniform distribution with all orientations equally likely (i.e. $\alpha = \frac{1}{n\pi}$ for some integer n):

$$P(s_1) = \alpha.$$

We used a Bayesian Adaptive Staircase method (Kontsevich and Tyler 1998) to choose the orientation of the second Gabor in Tasks 1A/1B. In Task 2, for "different" trials, it is normally distributed around the first with variance σ_Δ^2:

$$P(s_2 \,|\, s_1, C_2 = 2) = \mathcal{N}(s_2 \,|\, s_1, \sigma_\Delta^2)$$
$$P(s_2 \,|\, s_1, C_2 = 1) = \delta(s_2 - s_1).$$

3 The Model

3.1 Task 1A

We used the Bayesian Adaptive Staircase method described in Kontsevich and Tyler's paper to choose values of Δ_s, the intensity of the stimuli for both Task 1A and 1B, to maximize the expected gain in information about the subject's psychometric curve for the Gabor orientations. We use a psychometric curve which has three parameters, slope(σ), threshold (ψ), and the probability of an attention lapse (δ). We define $\lambda_t = (\sigma_t, \psi_t, \delta_t)$. For this method, we create a sample space of possible values for these parameters and update their probability distributions each trial by using Bayesian inference to condition on the result of that trial. The parameter δ is intended to model the possibility of subjects' missing the stimulus presentation and guessing Left or Right with equal probability. Although we were only interested in the parameter σ, including the parameters ψ and δ was found to provide better fits of σ to simulated data.

An unbiased observer's decision rule for this task is simple: respond s_1 if $x_1 > x_2$, and s_2 if $x_2 > x_1$ ($P(x_1 = x_2) = 0$). We can use this decision criterion to derive the form of an observer's psychometric function:

$$\begin{aligned}P(R_1 = 1 \mid \Delta_s) &= P(x_1 > x_2 \mid \Delta_s)\\ &= P(\Delta_x > 0 \mid \Delta_s).\end{aligned}$$

introducing bias, the expression becomes: $P(\Delta_x > \psi \mid \Delta_s)$.

Now, $P(\Delta_x \mid \Delta_s) = \mathcal{N}(\Delta_x \mid \Delta_s, \sigma^2)$, so we integrate and get an expression involving the error function, $\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$:

$$\begin{aligned}P(\Delta_x > \psi \mid \Delta_s) &= \int_{-\infty}^{\psi} \mathcal{N}(\Delta_x \mid \Delta_s, \sigma^2) d\Delta_x\\ &= \frac{1}{2}\left(1 + \mathrm{erf}\left(\frac{\psi - \Delta_s}{\sigma\sqrt{2}}\right)\right).\end{aligned}$$

Finally, we incorporate δ, the lapse probability, and the expression becomes:

$$P(R_1 = 1 \mid \Delta_s) = \frac{\delta}{2} + (1-\delta)\frac{1}{2}\left(1 + \mathrm{erf}\left(\frac{\psi - \Delta_s}{\sigma\sqrt{2}}\right)\right).$$

3.2 Task 1B

The model for this task is the same as for Task 1A, except that instead of comparing the orientations of two Gabors, the observer compares the orientation of a single Gabor to that of the reference line. Since the reference lines are fixed in position for all trials, we assume that the observer will learn this orientation during practice, so the internal representations of the reference lines are not noisy. So the model is identical, except only the noisiness of one Gabor is involved. Thus:

$$P(R_1 = 1 \mid \Delta_s) = \frac{\delta}{2} + (1-\delta)\frac{1}{2}\left(1 + \mathrm{erf}\left(\frac{\psi - \Delta_s}{\sigma_1\sqrt{2}}\right)\right).$$

3.3 Same/Different Task (Task 2)

Now we develop a theoretical model of the same/different task, with the assumption that the observer believes $P(C_2 = 1) = P(C_2 = 2) = 1/2$. The Optimal Bayesian Observer (OBO) model is a special case. In this model, the observer's response is based on the following criterion:

$$\mid \Delta_x \mid < f(\sigma_\Delta).$$

We consider several special cases:

1) A constant threshold model: $f(\sigma_\Delta) = k,$ a constant.

3) An Optimal Bayesian Observer model, which predicts the optimal strategy, given our assumptions, using Bayesian inference. We go over the Bayesian model below and show that it's decision rule fits this form.

We assume the OBO learns the distribution of Δ, including the parameter σ_Δ, and makes optimal use of this information in this task.
The OBO answers same iff

$$P(C_2 = 1 \,| x_1, x_2) > P(C_2 = 2 \,| x_1, x_2).$$

or, equivilently,

$$\log \frac{P(C_2 = 1 \,| x_1, x_2)}{P(C_2 = 2 \,| x_1, x_2)} = d > 0.$$

we call this value d the decision variable.

3.4 Deriving the OBO Decision Criterion

$$\begin{aligned}
d &= \log \frac{P(C_2 = 1 \,| x_1, x_2)}{P(C_2 = 2 \,| x_1, x_2)} \\
&= \log \left(\frac{\frac{P(x_1, x_2 \,| C_2 = 1) P(C_2 = 1)}{P(x_1, x_2)}}{\frac{P(x_1, x_2 \,| C_2 = 2) P(C_2 = 2)}{P(x_1, x_2)}} \right) \\
&= \log \frac{P(x_1, x_2 \,| C_2 = 1)}{P(x_1, x_2 \,| C_2 = 2)} + \log \frac{P(C_2 = 1)}{P(C_2 = 2)}
\end{aligned}$$

Now, we'll look at the numerator and denominator separately.

$$
\begin{aligned}
P(x_1, x_2 \,|\, C_2 = 1) &= \int\int P(x_1, x_2 \,|\, s_1, s_2, C_2 = 1) P(s_1) P(s_2) ds_1 ds_2 \\
&= \int\int P(x_1 \,|\, s_1) P(x_2 \,|\, s_2) P(s_1) P(s_2 \,|\, s_1, C_2 = 1) ds_1 ds_2 \\
&= \int\int P(x_1 \,|\, s_1) P(x_2 \,|\, s_2) P(s_1) \delta(s_1 - s_2) ds_1 ds_2 \\
&= \alpha \int P(x_1 \,|s_1) P(x_2 \,|\, s_2 = s_1) ds_1 \\
&= \alpha \int \mathcal{N}(x_1 \,|\, s, \sigma_1^2) \mathcal{N}(x_2 \,|\, s, \sigma_2^2) ds \\
&= \alpha * \mathcal{N}(x_1 \,|\, x_2, \sigma_1^2 + \sigma_2^2) \\
&= \alpha * \mathcal{N}(x_1 - x_2 \,|\, 0, \sigma_1^2 + \sigma_2^2)
\end{aligned}
$$

$$
\begin{aligned}
P(x_1, x_2 \,|\, C_2 = 2) &= \int\int P(x_1, x_2 \,|\, s_1, s_2, C_2 = 2) P(s_1, s_2 \,|\, C_2 = 2) ds_1 ds_2 \\
&= \int\int\int P(x_1 \,|\, s_1) P(x_2 \,|\, s_2) P(s_1) P(s_2 \,|\, s_1, C_2 = 2, \Delta_s) P(\Delta_s \,|\, C_2 = 2) ds_1 ds_2 d\Delta_s \\
&= \int\int\int P(x_1 \,|\, s_1) P(x_2 \,|\, s_2) P(s_1) \delta(s_2 - s_1 - \Delta_s) P(\Delta_s \,|\, C_2 = 2) ds_1 ds_2 d\Delta_s \\
&= \alpha \int\int P(x_1 \,|\, s_1) P(x_2 \,|\, s_2 = s1 + \Delta_s) P(\Delta_s \,|\, C_2 = 2) ds_1 d\Delta_s \\
&= \alpha \int\int \mathcal{N}(x_1 \,|\, s_1, \sigma_1^2) \mathcal{N}(x_2 \,|\, s_1 + \Delta_s, \sigma_2^2) \mathcal{N}(\Delta_s \,|\, \mu_\Delta, \sigma_\Delta^2) ds_1 d\Delta_s \\
&= \alpha \int\int \mathcal{N}(s_1 \,|\, x_1, \sigma_1^2) \mathcal{N}(s_1 \,|\, x_2 - \Delta_s, \sigma_2^2) \mathcal{N}(\Delta_s \,|\, \mu_\Delta, \sigma_\Delta^2) ds_1 d\Delta_s \\
&= \alpha \int \mathcal{N}(x_1 \,|\, x_2 - \Delta_s, \sigma_1^2 + \sigma_2^2) \mathcal{N}(\Delta_s \,|\, \mu_\Delta, \sigma_\Delta^2) d\Delta_s \\
&= \alpha \int \mathcal{N}(\Delta_s \,|\, x_2 - x_1, \sigma^2) \mathcal{N}(\Delta_s \,|\, \mu_\Delta, \sigma_\Delta^2) d\Delta_s \\
&= \alpha * \mathcal{N}(x_1 - x_2 \,|\, \mu_\Delta, \sigma^2 + \sigma_\Delta^2) \\
&= \alpha * \mathcal{N}(\Delta_x \,|\, 0, \sigma^2 + \sigma_\Delta^2)
\end{aligned}
$$

So now, substituting, we get:

$$
\begin{aligned}
d &= \log \frac{\alpha * \mathcal{N}(\Delta_x \,|\, 0, \sigma^2)}{\alpha * \mathcal{N}(\Delta_x \,|\, 0, \sigma^2 + \sigma_\Delta^2)} + \log \frac{P(C_2 = 1)}{P(C_2 = 2)} \\
&= \log \left(\frac{\frac{1}{\sqrt{2\pi\sigma^2}} * exp[-\frac{1}{2}\frac{(\Delta_x)^2}{\sigma^2}]}{\frac{1}{\sqrt{2\pi(\sigma^2+\sigma_\Delta^2)}} * exp[-\frac{1}{2}\frac{(\Delta_x)^2}{\sigma^2+\sigma_\Delta^2}]} \right) + \log \frac{P(C_2 = 1)}{P(C_2 = 2)} \\
&= \frac{1}{2} \log \left(\frac{\sigma^2 + \sigma_\Delta^2}{\sigma^2} \right) - \frac{1}{2} * \left(\frac{(\Delta_x)^2}{\sigma^2} - \frac{(\Delta_x)^2}{\sigma^2 + \sigma_\Delta^2} \right) + \log \frac{P(C_2 = 1)}{P(C_2 = 2)} \\
&= \frac{1}{2} \log \left(1 + \frac{\sigma_\Delta^2}{\sigma^2} \right) - \frac{\sigma_\Delta^2 (\Delta_x)^2}{2\sigma^2(\sigma^2 + \sigma_\Delta^2)} + \log \frac{P(C_2 = 1)}{P(C_2 = 2)}
\end{aligned}
$$

Now we can change our decision criterion, $d > 0$ into a condition on $|\,\Delta_x\,|$:

$$
d = \frac{1}{2} \log \left(1 + \frac{\sigma_\Delta^2}{\sigma^2} \right) - \frac{\sigma_\Delta^2 (x_1 - x_2)^2}{2\sigma^2(\sigma^2 + \sigma_\Delta^2)} + \log \frac{P(C_2 = 1)}{P(C_2 = 2)} > 0
$$

$$
\Longleftrightarrow
$$

$$
|\,\Delta_x\,| < \sigma \sqrt{\left(1 + \frac{\sigma^2}{\sigma_\Delta^2} \right) \left(\log \left(1 + \frac{\sigma_\Delta^2}{\sigma^2} \right) + 2 * \log \left(\frac{P(C_2 = 1)}{P(C_2 = 2)} \right) \right)} = B
$$

3.5 Overall Response Frequency Equals Prior Distribution For Sampling Strategy

Here is the general case for estimating a stimuli from an encoding. s is the true stimulus, $\vec{r}$ is the neural population response vector encoding the stimulus, this is decoded to create an estimate of the stimuli, $\hat{s}$. The overall response frequency is the distribution $P(\hat{s})$, and the prior distribution is $P(s)$. The Sampling decoding strategy is defined by by $P(\hat{s}\,|\,\vec{r}) = P(s\,|\,\vec{r})$. In contrast, the Maximizing strategy uses this decoding rule: $P(\hat{s}\,|\,\vec{r}) = \text{argmax}_s P(s\,|\,\vec{r})$.

$$
\begin{aligned}
P(\hat{s}) &= \int P(\hat{s} \,|\, s) P(s) ds \\
&= \int \int P(\hat{s} \,|\, \vec{r}) P(\vec{r} \,|\, s) P(s) ds d\vec{r} \\
&= \int P(\hat{s} \,|\, \vec{r}) P(\vec{r}) d\vec{r} \\
&= \int P(s \,|\, \vec{r}) P(\vec{r}) d\vec{r} \\
&= P(s).
\end{aligned}
$$

For our task, $s = C_2$, $\vec{r} = \Delta_x$ and $\hat{s} = R_2$. So the integral over $\hat{s}$ becomes a sum over C.

3.6 Further Predictions of the Model

If we condition the response (R_2) on the type of trial (C_2) we can predict the rate of True Positives (TP) and False Positives (FP). Here, we consider "sameness" a positive signal.

TP and FP are given by:

$$
\begin{aligned}
TP &= P(R_2 = 1 \,|\, C_2 = 1) \\
&= P(|\,\Delta_x\,| < B \mid C_2 = 1) \\
&= \int P(|\,\Delta_x\,| < B \mid \Delta_s) P(\Delta_s \,|\, C_2 = 1) d\Delta_s \\
&= \int P(|\,\Delta_x\,| < B \mid \Delta_s) \delta(\Delta_s) d\Delta_s \\
&= \int_{-B}^{B} \mathcal{N}(\Delta_x \,|\, 0, \sigma^2) d\Delta_x \\
&= \int_{-B}^{B} \frac{\exp\left(\frac{-\Delta_x^2}{2\sigma^2}\right)}{\sigma\sqrt{2\pi}} \, d\Delta_x \\
&= \operatorname{erf}\left(\frac{B}{\sigma\sqrt{2}}\right).
\end{aligned}
$$

$$
\begin{aligned}
FP &= P(R_2 = 1 \,|\, C_2 = 2) \\
&= P(|\,\Delta_x\,| < B \mid C_2 = 2) \\
&= \int P(|\,\Delta_x\,| < B \mid \Delta_s) P(\Delta_s \,|\, C_2 = 2) d\Delta_s \\
&= \int P(|\,\Delta_x\,| < B \mid \Delta_s) \mathcal{N}(\Delta_s \,|\, 0, \sigma_\Delta^2) d\Delta_s \\
&= \int_{-B}^{B} \mathcal{N}(\Delta_x \,|\, 0, \sigma^2 + \sigma_\Delta^2) d\Delta_x \\
&= \int_{-B}^{B} \frac{\exp\left(\frac{-\Delta_x^2}{2(\sigma^2+\sigma_\Delta^2)}\right)}{\sqrt{2\pi(\sigma^2 + \sigma_\Delta^2)}} \, d\Delta_x \\
&= \operatorname{erf}\left(\frac{B}{\sqrt{2(\sigma^2 + \sigma_\Delta^2)}}\right).
\end{aligned}
$$

FIGURE

Based on these values, we can also calculate the response frequencies, that is, $P(R_2 = 1)$ and $P(R_2 = 2)$. $P(R_2 = 1) = TP + FP$, and similarly, $P(R_2 = 2) = TN + FN$, where $TN = 1 - FP$ and $FN = 1 - TP$.

For the Bayesian Model Comparison, we also must know the probability of a certain response given a value of Δ_s. This will be the probability that $|\Delta_x|$ is less than B given Δ_s:

$$\begin{aligned} P(|\Delta_x| < B \,|\, \Delta_s) &= \int_{-B}^{B} \mathcal{N}(\Delta_x \,|\, \Delta_s, \sigma^2) d\Delta_x \\ &= \frac{1}{2}\left(\operatorname{erf}\left(\frac{B-\Delta_s}{\sigma\sqrt{2}}\right) - \operatorname{erf}\left(\frac{-B-\Delta_s}{\sigma\sqrt{2}}\right)\right) \end{aligned}$$

3.7 Limits as σ_Δ Goes to 0

$$\lim TP_{\sigma_\Delta \to 0} = TP\sigma_\Delta \to 0 \quad \lim TP_{\sigma_\Delta \to 0} = TP\sigma_\Delta \to 0$$

4 Results

FIGURES GALORE!

5 Conclusions

We didn't have time to run very many subjects and still have some outstanding experiment design issues to deal with. Also, in our model comparison, different models were the most likely for different subjects' data. So we cannot currently make the kind of broad statement about human perception which we intended to demonstrate through this experiment.

However, from our limited data set, it seems that humans do not, in general, perform this same/different task in a way that is Bayes-Optimal. Model 10 was the most likely model across subjects. This model used the Sampling strategy and had σ as a free parameter. We believe this strategy is more intuitive than the Maximizing (Optimal) model for this task.

But supposing this strategy is more intuitive, this still would not explain why Sampling occurs. We propose two possible adaptive benefits of this model. Since an observer must infer the structure of the task and the relevant variables, in this case, σ_Δ and $P(C_2)$, the observer does not have access to perfect information about these parameters. While it is possible to make reasonable estimates of these values given enough information, in many real world situations, it may be that these parameters are not fixed, but are changing in time. Even if they seem constant over a short time period, in a real world scenario, there would be numerous variables, changing with time, influencing these parameters. Furthermore, a change in the values of either parameter would lead to a change in the decision rule. There may be a trade-off between exploitation of known information and exploration, or checking for possible changes in the values of these parameters. Sampling may also be cheaper for the nervous system in terms of time and/or computational effort.

6 Extensions and Further Questions

There are many possible alterations to the tasks and models which would be interesting to try, but which we didn't have time for. There are also some interesting questions which remain unanswered and can and in some cases should be looked in to. Here are some of them.

As mentioned, Task 1A/1B did not provide a very good fit of σ for Task 2. Also, while some subjects values of σ from Task 1A/1B and Task 2 appeared reasonably close, for half of our subjects, the values of σ from Task 1A/1B were much larger than the σ values we fitted to the Task 2 data (SEE FIGURE). This seems to indicate that Task 1A/1B were significantly different from Task 2 in a way which we have not been able to explain intuitively, and that this is much more the case for some subjects than others. Indeed, the points clearly form separate clusters in THE FIGURE. We'd like to know why this is. We'd also like to find another Task which would provide a better estimate of σ for Task 2. If we can find answers to both of these problems, our conclusions would be strengthened, since we could then argue that each subject has a certain constant (over short periods of time) amount of noise in their observations of the stimuli, and that they take the degree of noisiness into account in their decision making process.

In all of our current models, we have assumed that the observer knows the values of σ_Δ, but we have no demonstration that this is the case. While the practice trials should facilitate learning of σ_Δ and $P(C_2)$, we'd like to have some idea of the reasonableness of this assumption. It should be possible, using Bayesian methods, to model the distributions of an observer's belief distributions over these parameters. This could affect our models of the task, or our experiment design. For instance, if it is found that an observer would not be able to make a "reasonably good" estimate of these parameters after 40 practice trials, we might consider having more practice trials, or modifying the models in some way to try to account for this. It should be noted, however, that the proposed analysis here is unverifiable and assumes that the observer is using Bayesian methods to estimate the parameters, and thus should probably only be used as a heuristic. Also, since the practice trials are longer than the real trials, the observer would probably have less noisy observations than during the real trials, and hence would also be able to extract more information about the parameters σ_Δ and $P(C_2)$ on each trial than during the real trials.

We noted that one reason we think humans might be Sampling in Task 2 is that the predictions of this strategy seem more intuitive, because the overall response frequency of R_2 matches the posterior distribution of C_2. But there are certainly other decision rules for which this is true. We'd like to derive an optimal decision rule for this task with this criteria as a constraint, or perhaps some looser constraint which would penalize models for disparities between response frequency and the prior, and find the model which is optimal in the sense that it minimizes some cost function which involves both this penalty and, of course, reward for greater expected percent correct.

It is known that certain orientations (vertical and horizontal, in particular) are more easily recognized than other. We assume independence of orientation throughout our models. This assumption is clearly a simplification, then, which may or may not be significant. In particular, in Task 1B, since the orientation of the reference line is constant, this may have an effect (for instance, if it was vertical, we would expect a lower estimate of σ here than in Task 2). Shuang examined trends in responses and correctness vs. orientation in some of our data and found that there was a dependence on orientation. This could affect models (introducing dependence on orientation) or experiment design (finding ways to control for this factor).

Professor Josić and Shuang developed a model of the task which included correlation between the orientations of the Gabors. While we haven't looked at correlation effects, and the Gabors should be placed far enough apart that they are not significant, this could be looked in to if there were a reason to suspect that correlation was playing a significant role.

Fatigue, learning, and other time dependencies may be at play. We have already made efforts to address fatigue in experiment design (breaks between blocks, mandatory 1/2 hour break), and we haven't uncovered any clear trends depending on time in the data, but this is certainly a factor to keep in mind.

We manipulated σ_Δ and tried to predict the response patterns for different values. But as experimenters, we also control the parameter $P(C_2)$. Our models also make different predictions when this value is changed. So we could run the same type of experiment but changing this parameter instead. It would be great to be able to run both experiments for the same subject on the same day. If a single model makes the best predictions in both cases, this would be much stronger evidence for that model. We could do Bayesian Model comparison using all the data from both versions.

We included a "lapse parameter", δ, in the model for Task 1A/1B. It seems reasonable that subjects would also have attention lapses in Task 2. This is already mentioned, but not implemented, in the Model Comparison Code. Models without a lapse parameter would make very poor predictions in some cases of attention lapse. For instance, if $\Delta_s = 50°$, and

$\sigma = 4°$, the probability of responding "same" is essentially zero in a model without δ. But in a model including the possibility of lapses, the probability of a response, regardless of stimuli intensity (Δ_s) never goes below $\delta/2$. So including the lapse parameter has the potential to improve fits.

As mentioned in the Experiment Design, the Bayesian Adaptive Staircase method of choosing stimuli intensities and estimating the psychometric curve tends to lead to large blocks of similar stimulus intensities, so that the correct response may be the same for many trials in a row. This may lead a subject to infer some dependency of the correct response on the previous correct responses. It may also make the task easier, if, for instance, $\Delta_s = 4.1°$, then on the next trial, $\Delta_s = 4.2°$. We would like to find a method of choosing stimuli intensity values which does as much as possible to maximize the expected gain in information (decrease in entropy) while creating the illusion of randomness. For example, one could predetermine what the correct response would be for each trial in advance, and then choose a stimulus intensity value which maximizes the expected gain in information under the constraint of the predetermined correct response. Or one could also begin constraining the possible intensities probabilistically (i.e. impose some penalty for choosing similar stimuli values twice in a row). There are numerous random number tests which are used to test psuedo-random number generators which might be of use in testing a proposed method, although humans are probably less sensitive than these tests.

We created code to test the effectiveness of the Bayesian Adaptive Staircase is estimating the true values of the parameters of a psychometric curve using simulated data and compared it to the more common Method of Constant Stimuli. So we can look at the probability distribution of different combinations of parameters after some number of trials. While the Bayesian Adaptive Staircase offers a clear advantage to the Method of Constant Stimuli, there is still considerable variance across the estimates from different simulations after 120 trials from different simulations, and for some simulations, the estimate was very different from the true value. So it may be that we should have more trials for Task 1A/1B to improve the estimate. It might also be interesting to compare the results of the method to the Cramer-Rao Lower Bound.

In the Bayesian Model Comparison Code, the size of the arrays is a constraint for some of the models with more parameters. One could split the trial arrays by same/different. For the same trials of a given block, the $P(R_2)$ is the same for each trial.

7 References