

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Categorization, Information Selection and Stimulus Uncertainty

Permalink

<https://escholarship.org/uc/item/20w5v933>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 39(0)

Authors

Halpern, David J.

Gureckis, Todd M.

Publication Date

2017

Peer reviewed

Categorization, Information Selection and Stimulus Uncertainty

David J. Halpern (david.halpern@nyu.edu)

Todd M. Gureckis (todd.gureckis@nyu.edu)

Department of Psychology, NYU, 6 Washington Place
New York, NY 10003 USA

Abstract

Although a common assumption in models of perceptual discrimination, most models of categorization do not explicitly account for uncertainty in stimulus measurement. Such uncertainty may arise from inherent perceptual noise or external measurement noise (e.g., a medical test that gives variable results). In this paper we explore how people decide to gather information from various stimulus properties when each sample or measurement is noisy. The participant's goal is to correctly classify the given item. Across two experiments we find support for the idea that people take category structure into account when selecting information for a classification decision. In addition, we find some evidence that people are also sensitive to their own perceptual uncertainty when selecting information.

Keywords: attention, categorization, information sampling

Categorizing objects or situations into meaningful groups is a critical cognitive ability. Many existing theories of categorization share an unrealistic assumption that information about the features of a to-be-categorized object can be directly and precisely observed (Medin & Schaffer, 1978; Nosofsky, 1986; Smith & Minda, 1998; Love, Medin, & Gureckis, 2004). But this is often not the case; for example, doctors do not always have full access to all the information about patient symptoms (i.e., stimulus features) but instead can only rely on a patient's self-report and medical tests which are subject to selective reporting and noise (perceptual or otherwise). If a doctor orders a cholesterol test, they must take into account that the patient's true levels are somewhat different than those reported by the test due to error or noise. There is often uncertainty not only about the *category* of an object but also about its specific dimension or feature values.

If noise and uncertainty are the affliction, then the antidote is the fact that categorizers can often select how feature or stimulus measurements are made. In many cases, a good strategy for selecting measurements (e.g., repeating the same measurement when it is known to be noisy) can significantly improve categorization performance. Here we ask what strategies people use to select stimulus measurements in the service of categorization. While previous approaches to this question tend to focus on how people select stimulus features to view without noise before making a categorization decision (e.g., Nelson, McKenzie, Cottrell, & Sejnowski, 2010 or in eye and mouse tracking studies, Rehder & Hoffman, 2005; Matsuka & Corter, 2008; Blair, Watson, Walshe, & Maj, 2009) here we explore the effect of measurement or perceptual noise on information sampling strategies.

Categorizing in a Noisy World

We begin by presenting an Ideal Actor analysis of categorization under measurement noise which extends the General

Recognition Theory (Ashby & Townsend, 1986). We then present a series of experiments exploring how people sample information about category features under conditions of stimulus noise.

Categorizing without perceptual noise The standard model of a probabilistic binary categorization task (Medin & Schaffer, 1978; Nosofsky, 1986; Smith & Minda, 1998) assumes that on a given trial t , a category $C_t \in \{A, B\}$ is drawn randomly and a stimulus \mathbf{s}_t is generated from the distribution associated with C_t (see Figure 1a). Subjects are assumed to either learn the parameters of these distributions through experience or description. Based on the information in \mathbf{s}_t , subjects guess which category $C_t \in \{A, B\}$ it was generated from. For example, in our first experiment, the stimuli have two dimensions, color and orientation, i.e. $\mathbf{s}_t = [s_{t_{\text{orientation}}}, s_{t_{\text{color}}}]$ and the category distributions are bivariate Gaussian distributions $\mathcal{N}(\boldsymbol{\mu}_{C_t}, \boldsymbol{\Sigma})$ where the mean is dependent on the category and we can decompose the covariance matrix as $\boldsymbol{\Sigma} = \begin{bmatrix} \sigma_{\text{orientation}}^2 & \rho \sigma_{\text{orientation}} \sigma_{\text{color}} \\ \rho \sigma_{\text{orientation}} \sigma_{\text{color}} & \sigma_{\text{color}}^2 \end{bmatrix}$. The Ideal Observer decision rule, assuming knowledge of the category distributions, is to use the log-likelihood ratio

$$l(\mathbf{s}_t) = \frac{\log P(\mathbf{s} = \mathbf{s}_t | C_t = A)}{\log P(\mathbf{s} = \mathbf{s}_t | C_t = B)} \quad (1)$$

This rule responds A when $l(\mathbf{s}_t) > 0$ and B otherwise (Ashby & Gott, 1988).

Categorizing with perceptual noise In the 1980s, Ashby and colleagues developed General Recognition Theory (GRT), a family of models of multidimensional classification that assumed perceptual noise, in contrast to the models above (Ashby & Townsend, 1986). Here we discuss a special case of GRT that uses the Ideal Observer decision rule and makes two critical assumptions about perceptual noise: that perceptual noise has a normal distribution and that the noise in perceiving each stimulus dimension is independent. In this model, on trial t the subject perceives a percept $\mathbf{p}_t = [p_{t_{\text{color}}}, p_{t_{\text{orientation}}}]$ with probability $\mathcal{N}(\mathbf{p}_t; \mathbf{s}_t, \boldsymbol{\Sigma}_p)$ where $\boldsymbol{\Sigma}_p = \begin{bmatrix} \sigma_{\text{orientation}}^2 & 0 \\ 0 & \sigma_{\text{color}}^2 \end{bmatrix}$. $\boldsymbol{\Sigma}_p$ is diagonal due to the independent noise assumption and $\sigma_{\text{orientation}}^2$ and σ_{color}^2 are a subject's perceptual noise for orientation and color respectively. When making their categorization decision, subjects only have access to \mathbf{p}_t rather than \mathbf{s}_t as before (see Figure 1b). Therefore, the Ideal Observer's decision rule is now based on the log likelihood ratio $l(\mathbf{p}_t)$ where \mathbf{p}_t is distributed under each category as $\mathcal{N}(\mathbf{p}_t; \boldsymbol{\mu}_{C_t}, \boldsymbol{\Sigma} + \boldsymbol{\Sigma}_p)$

Reducing uncertainty by making measurements One strategy for dealing with uncertainty in the percept \mathbf{p}_t is to exert control over the amount of information collected about each stimulus dimension. In our task, the categorizer is given a fixed budget of κ noisy measurements \mathbf{m}_t of a stimulus made up of two dimensions. These measurements can be allocated so that there are $\lambda_{orientation}$ measurements of the orientation dimension and $\kappa - \lambda_{orientation}$ measurements of color with $\lambda_{orientation} \in [0, \kappa]$. Conceptually, this is like a doctor ordering multiple runs of medical tests and collect more or fewer runs of some tests over others. As in our experiment, orientation measurements are each independently distributed with probability $\mathcal{N}(s_t, \sigma_{m_{orientation}})$ and similarly for color. Because these measurements are all normally distributed, the mean measurement $\bar{\mathbf{m}}_t$ is distributed as $\mathcal{N}(s_t, \Sigma_m)$, where

$$\Sigma_m = \begin{bmatrix} \frac{\sigma_{m_{orientation}}^2}{\lambda_{orientation}} & 0 \\ 0 & \frac{\sigma_{m_{color}}^2}{\kappa - \lambda_{orientation}} \end{bmatrix}. \text{ We now assume that subjects}$$

estimate the mean of several colors or oriented lines with a constant variance, largely supported by the literature on ensemble perception (e.g., Maule & Franklin, 2016; Dakin & Watt, 1997). We assume the same perceptual distribution as above except $\bar{\mathbf{p}}_t$ represents a subjects estimate of the mean measurement and Σ_p represents their noise in estimating that mean. Let $\Sigma_{mp} = \Sigma + \Sigma_m + \Sigma_p$. Therefore,

$$P(\bar{\mathbf{p}} = \bar{\mathbf{p}}_t | C_t) = \mathcal{N}(\bar{\mathbf{p}}_t; \mu_{C_t}, \Sigma_{mp}) \quad (2)$$

The optimal decision rule depends on the log likelihood ratio $l(\bar{\mathbf{p}}_t)$ as above.

Optimizing stimulus measurement In our most interesting setting (see Figure 1c), we allow subjects to choose λ in order to maximize their own categorization performance. This is akin to the doctor trying to optimize the probability of a correct diagnosis while trying to keep the cost of running the medical tests below a budget. To make the dependence on λ explicit, we can rewrite Σ_{mp} as $\Sigma_{mp}(\lambda_{orientation}) =$

$$\begin{bmatrix} (\sigma_{orientation}^2 + \frac{\sigma_{orientation}^2}{\lambda_{orientation}} + \sigma_{p_{orientation}}^2) & \rho \sigma_{orientation} \sigma_{color} \\ \rho \sigma_{orientation} \sigma_{color} & (\sigma_{color}^2 + \frac{\sigma_{color}^2}{\kappa - \lambda_{orientation}} + \sigma_{p_{color}}^2) \end{bmatrix}. \text{ The}$$

expected percent correct of an Ideal Observer with a given $\lambda_{orientation}$ should be (Anderson, 1958):

$$EC(\lambda_{orientation}) = \frac{1}{2} \operatorname{erfc} \left(-\frac{\sqrt{(\mu_A - \mu_B)^T \Sigma_{mp}(\lambda_{orientation})^{-1} (\mu_A - \mu_B)}}{2\sqrt{2}} \right) \quad (3)$$

The Ideal Actor should then set $\lambda_{orientation}$ to $\lambda_{orientation}^* = \arg \max_{\lambda_{orientation}} EC(\lambda_{orientation})$.

Theoretical Predictions

In the following experiment, subjects perform a categorization task with stimuli that vary on color and orientation dimensions. Subjects only receive κ noisy signals of what the color and orientation of the stimulus are but choose a number $\lambda_{orientation}$ of signals to receive of the orientation dimension. By varying the type of category structure (the




Figure 1: (a) Graphical model of the standard categorization model where stimuli s depend on the category C . (b) a model analogous to General Recognition Theory where stimuli s depend on the category C but are never directly observable. Instead, the observer has access to p which itself depends on s but is corrupted with noise. (c) An active learning model of categorization under uncertainty. Here λ reflects the number of measurements m made of the stimulus s . The categorizer has access to p which depends on the category, stimulus, and the information sampling strategy (m and λ).

two $P(s|C)$ for category A and B) experimentally and using natural variation in perceptual noise, we can test two predictions about how subjects should select signals based on the above Ideal Actor model. We first divide category structures into three groups: 1D-color structures where $P(\text{orientation}|A) = P(\text{orientation}|B)$, 1D-orientation structures where $P(\text{color}|A) = P(\text{color}|B)$ and 2D structures otherwise. We predict:

1. The rank ordering of subjects choice of $\lambda_{orientation}$ should be $1\text{D-color} \leq 2\text{D} \leq 1\text{D-orientation}$
2. Within 2D categories, subjects choice of $\lambda_{orientation}$ will be modulated by the subject's measured perceptual noise.

Intuition for the two hypotheses can be seen in the Ideal Actor predictions in Figure 2 which plots expected Ideal Observer percent correct as a function of strategy and category structure in (a) and perceptual noise in (b). In the 1D-color situation accuracy is optimized by collecting zero orientation samples. In the 1D-orientation, accuracy is highest with 10 orientation samples. In the 2-D case, accuracy is optimized with 5 samples (Figure 2a). Furthermore the peak of these functions changes with perceptual noise (Figure 2b).

Experiment 1

In order to test these predictions, we used a task divided into six phases. In the first two phases, we estimated the free parameters of the Ideal Actor, $\sigma_{p_{orientation}}^2$ and $\sigma_{p_{color}}^2$, the subjects' noise in estimating the mean of a number of colored dots and oriented lines using a 2AFC task. In the next three phases,




Figure 2: (a) Ideal Observer expected percent correct as a function of condition (described in Table 1) and strategy. The rank order of the max of these curves is 1D-color \leq 2D \leq 1D-orientation. (b) Ideal Observer expected percent correct if subjects had no perceptual noise (in red). In our experiment, we estimated noise parameters ($\sigma_{orientation}$ and σ_{color}) using a 2AFC task. Shown in blue and purple are Ideal Observers with estimated noise parameters from two different subjects. Perceptual noise has a strong influence on optimal strategy which is indicated by the arrow

subjects underwent extensive training to learn the category and measurement distribution parameters (assumed known by the Ideal Actor). In the final test phase, subjects performed a categorization task where they could decide how allocate a fixed budget of samples to the two stimulus dimensions. All phases were 100 trials except the first category learning phase which was 200.

Participants Sixty one participants were recruited through Amazon Mechanical Turk. Participants received \$8 for participating in the experiment with a performance based bonus of up to \$10. Ten trials were selected at random from the entire experiment and participants were awarded a bonus of \$0.25 for each trial correct. Participants were randomly assigned to the eight conditions described in Table 1.

Stimuli and Procedure All stimuli in the experiment were generated randomly by drawing samples from the generative model. To generate the stimuli, each sample corresponded to either the angle of an oriented line relative to the circle or the color of a dot where the number was the angle on a circle of radius 60 in CIE 1976 (L^* , a^* , b^*) color space. The locations of the colored dots on the screen were determined by force layout, an algorithm within the d3 javascript visualization library (Bostock, Ogievetsky, & Heer, 2011). Examples of these stimuli can be seen in (Figure 3). Throughout the experiment, the “measurement noise” $\sigma_m = \sigma_{m_{orientation}} = \sigma_{m_{color}} = .6$.

Perceptual Noise Estimation Phases. We adapted a 2AFC task from Jogan and Stocker (2014) designed to estimate subjects’ noise in estimating a property of a stimulus. We conducted this task in two phases, one for orientation and one for color (with order counterbalanced across subjects). On each trial of the task, three stimuli $s \in \{\text{test}, \text{reference}_1, \text{reference}_2\}$ were presented. The subject was asked which of the two reference stimuli was closer in terms of the property of interest to the test stimulus and responded by pressing the appropriate computer key. The specific properties of interest here were the average orientation of a set of several lines or the aver-

age color of a set of several dots. The stimuli during these phases looked the same as the stimuli in the later categorization phase (Figure 3b) but with just one feature present. On a given trial t , stimuli were generated by drawing n_t samples from $\mathcal{N}(\mu_{st}, \sigma_m^2)$ with $n_t \in [1, 10]$ to keep the range and set sizes of the stimuli the same as in the later categorization experiment. μ_{st} was selected on each trial by a Bayesian adaptive procedure (Kontsevich & Tyler, 1999). Using an Ideal Observer analysis detailed in Jogan and Stocker (2014), we can estimate the perceptual noise parameters for identifying the mean of the stimuli based on subjects’ performance in this task.

Category structures/conditions. Throughout the four categorization phases, stimuli were generated from a single pair of categories, chosen from a set of eight categories described in Table 1. For each subject, each dimension was shifted by a random amount chosen from a uniform on $[0, 2\pi]$ to wash out any effect of a specific stimulus range.

Condition	μ_A	μ_B	Σ
1D-orientation	[0, 0]	[1, 0]	$\begin{bmatrix} .2 & 0 \\ 0 & .2 \end{bmatrix}$
1D-color	[0, 1]	[0, 0]	$\begin{bmatrix} .2 & 0 \\ 0 & .2 \end{bmatrix}$
2D ₁	[0, 0]	[.24, 62]	$\begin{bmatrix} .22 & -.2 \\ .2 & .22 \end{bmatrix}$
2D ₂	[.24, 62]	[0, 0]	$\begin{bmatrix} .22 & -.2 \\ .2 & .22 \end{bmatrix}$

Table 1: Category parameters for Experiment 1. There were eight categories total, with the other 4 being the same but with the entries of μ_A and μ_B swapped. For the main analysis, we collapse the conditions that share the parameters.

Category Learning Phase. In the category learning phase, a category was drawn from a uniform distribution and a bivariate sample was drawn from that category’s associated distribution. This sample was converted to the color and orientation stimulus space using the procedure described above (see Figure 3a for an example). Subjects responded by hitting the “Q” key if they thought the stimulus was in category A and the “P” key if the stimulus was in category B. Subjects then received feedback on whether they were correct or incorrect depending on the category structures defined above.

Measurement Noise Learning Phases The measurement noise learning phases were meant to acclimatize the subject to the effects of measurement noise on categorization. The stimuli were created by sampling from the full generative model for the task as described in the theory section and converting samples to the stimulus space as described above. During this phase, subjects did not get to choose the number of measurements of each dimension. Instead, stimuli in the first phase included ten measurements of each dimension ($\kappa = 20$ and $\lambda_{orientation} = 10$) and stimuli in the second included a total of ten measurements with a random number of them allocated to orientation ($\kappa = 10$ and $\lambda_{orientation} \in [0, 10]$) See Figure 3b for an example of the stimuli in this phase. Note that there

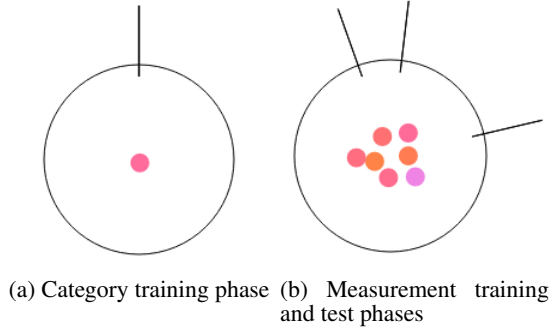


Figure 3: Example categorization trials

are multiple lines/dots in this stimulus reflecting the multiple noisy “measurements” made of each dimension. In order to gain an intuitive understanding of the measurement procedure, subjects were told “we showed the color of the stimulus to 10 people and the location of the stimulus to 10 different people. Later each of them had to re-create what they saw from memory. Your task will be to take their recreations and try to guess what category you think the original stimulus belonged to.” After every trial, subjects would receive feedback on their categorization judgement as well as feedback about what the true stimulus s_t had been on that trial.

Test Phase During this phase, subjects chose on a slider how many measurements of each dimension they would see on each trial ($\lambda_{orientation}$). Stimuli were then generated in the same manner as in the measurement noise learning phases. Subjects then performed the classification task as in the previous training sections.

Results

For each subject, we computed the posterior over the $\sigma_{orientation}^2$ and σ_{color}^2 parameters using the analysis described in Jogan and Stocker (2014). In order to check that the Bayesian adaptive procedure converged towards the correct estimate, Jogan and Stocker use a diagnostic called the Boundary Index (BI) a measure of the number of trials that were chosen to be at the boundary of the space. All of our subjects were below the recommended threshold of 0.9.

Using the perceptual noise posterior, we compute a posterior over Ideal Actor strategies for every subject in our experiment. Since each subject only experienced a single set of categories in the test phase, the Ideal Actor only uses a single $\lambda_{orientation}$ parameter for the entire experiment. However, subjects were able to change their choice of $\lambda_{orientation}$ parameter on every trial and relatively few subjects used just a single value. In order to compare subjects to the model, we took an average each subjects’ setting of $\lambda_{orientation}$. We chose a priori to average only the second half of test trials to ensure that subjects had stabilized after having experience with using the slider. The results turn out to be unchanged even if we use all of the data from the test phase.

In order to test our first hypothesis we used Kendall’s τ , a




Figure 4: (a) Scatter plot of subjects strategies vs. the Ideal Actor (IA) strategy posterior mean. X error bars are ± 1 SD of the IA posterior and Y-error bars are standard error of mean subject strategy. (b) Expected number additional incorrect Ideal Observer (IO) trials relative to the Ideal Actor (IA). We do not include error bars on the actual performance since the sampling distribution is over trial sequences while the sampling distributions on the Ideal Observers are over measurement selections averaged over trial sequences.

common non-parametric rank-correlation method. We found a significant monotonic relationship between category structure and subjects’ $\lambda_{orientation}$ choice with 1D-color < 2D < 1D-orientation (Kendall $\tau(59) = 0.52$, $p < 1e-8$).

We also found a significant linear relationship between subjects’ $\lambda_{orientation}$ parameters and the posterior mean Ideal Actor (Pearson $r(59) = 0.65$, $p < 1e-8$). This result was still significant using just the data from the 2D structure where differences in Ideal Actor strategies are only due to differences in estimated perceptual noise (Pearson $r(30) = 0.38$, $p = 0.03$). This provides weak evidence for our second hypothesis.

Given that many subjects did not use just a single $\lambda_{orientation}$ throughout the whole experiment, what was the cost of their suboptimal choice? Did they know to avoid choices that would lead to significantly worse performance? To answer this, we compare the theoretical performance of the Ideal Actor to what we call the *subject Ideal Observer*, the theoretical performance of an Ideal Observer who chose $\lambda_{orientation}$ on every trial as the subject did. The subject Ideal Observer performance is of interest because it isolates the expected decrease in performance solely due to choice of $\lambda_{orientation}$. In contrast, differences between subjects’ actual performance and the Ideal Actor may be for several reasons unrelated to the information selection strategy. In Figure 4b, we compare the Ideal Actor to subject Ideal Observer performance (in blue), subjects’ actual performance (in black) and a baseline where the Ideal Observer who chose $\lambda_{orientation}$ randomly (in purple). Only seven out of sixty-one subjects had subject Ideal Observers that did not perform significantly better than the baseline suggesting that most subjects were sensitive to the costs of choosing $\lambda_{orientation}$ incorrectly.

Discussion

We found some preliminary evidence suggesting that people take category structure and perceptual noise into account. While the correlation between the Ideal Actor and subjects strategies was significant, subjects deviated from the Ideal Actor in other significant ways. In particular, subjects of-

ten used multiple $\lambda_{orientation}$'s throughout the experiment and actual categorization performance did not match the subject Ideal Observer – two substantive suboptimalities. There may be several reasons for this including that subjects might not use the Ideal Observer rule for categorization or subjects did not learn the exact category parameters in the time allotted. It is difficult to assess subject knowledge in this task since the average subject only had 75% agreement with the Ideal Observer during the last 10 trials of the category learning phase. Also, a MANOVA found that subjects' disagreement with the Ideal Observer and suboptimality was significantly different across category types (Wilk's $\Lambda = .62$, $F(2, 58) = 7.6$, $p = 1e-4$). This suggests that subjects might have significantly different knowledge about the category across conditions. Finally our estimates of $\sigma_{orientation}^2$ and σ_{color}^2 may not be perfect which might bias our Ideal Actor model.

Not learning the category parameters is probably the most serious issue since the Ideal Actor strategy depends heavily on these parameters. In order to address this, we conducted a second study involving only binary-valued stimulus features. With only a finite number of possible stimuli, we can easily check whether subjects have "learned" the category in the sense of having a high agreement with the ideal observer when selecting the category for a given stimulus.

Experiment 2

Participants Thirty three participants who did not participate in the previous experiment were recruited through Amazon Mechanical Turk. Payment was the same as in Experiment 1.

Categorization Task Subjects in this task were instructed that they needed to help a doctor discover how to categorize patients presenting certain symptoms. Subjects would see the outcome of two medical tests represented as the color of horizontal and vertical lines (blue if positive and red if negative). All four of the possible outcomes can be seen in Figure 5 (a). Based on the stimulus, subjects would have to determine which of two diseases (A or B) the patient had. These diseases (or categories) were defined as bivariate Bernoulli distributions over possible test outcomes. Let 1 denote a positive test and 0 denote a negative test. Then let $P(\cdot|D)$ be a matrix where each entry with index $[v, h]$ indicates the probability of the vertical test taking on value $v \in (0, 1)$ and the horizontal test taking on value $h \in (0, 1)$ given that the patient has disease $D \in A, B$. We used category conditions described in Table 2.

In the first phase of the experiment, subjects simply saw the stimuli in Figure 5 with the above probabilities and subjects were told that the tests were performed with no measurement noise. In the later phases, subjects were told that they now would see κ tests on every trial with $\lambda_{horizontal}$ measurements of the horizontal tests. These tests had Bernoulli measurement noise, i.e. the probability of the horizontal test outputting k tests with the true value on each trial was $\binom{\lambda_{horizontal}}{k} p^k (1-p)^{(\lambda_{horizontal}-k)}$. In this experiment, we used

Condition	$P(A)$	$P(B)$
1D-horizontal	$\begin{bmatrix} .5 & 0 \\ .5 & 0 \end{bmatrix}$	$\begin{bmatrix} 0 & .5 \\ 0 & .5 \end{bmatrix}$
1D-vertical	$\begin{bmatrix} .5 & .5 \\ 0 & 0 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 \\ .5 & .5 \end{bmatrix}$
2D	$\begin{bmatrix} .5 & 0 \\ 0 & .5 \end{bmatrix}$	$\begin{bmatrix} 0 & .5 \\ .5 & 0 \end{bmatrix}$

Table 2: Category parameters for Experiment 2. There were six conditions total, with the other three being the same but with the category labels swapped. For the main analysis, we collapse the conditions that share the parameters.

a p of .8. Figure 5 shows an example of a noisy stimulus (i.e., multiple measurements of the horizontal or vertical line segment) with its true value below. The choice of κ and $\lambda_{horizontal}$ in each phase was exactly the same as in Experiment 1 with subjects having a choice of $\lambda_{horizontal}$ in the last phase as before. The first three phases of this experiment each consisted of 200 trials and the last phase had 100. While we could not derive a general analytic solution to the Ideal Actor in this case, we can easily compute the strategy by enumerating all of the potential observed measurements. We also assume that in this case the effects of perceptual noise on performance are minimal.




Figure 5: Example experiment 2 stimuli

Results

According to hypothesis 1 above, the rank order of the subject $\lambda_{horizontal}$ in each category structure should be $[1D\text{-vertical} \leq 2D \leq 1D\text{-horizontal}]$. We again found a significant monotonic relationship between category structure and subjects' $\lambda_{horizontal}$ choice in the direction we hypothesized (Kendall $\tau(31) = 0.747$, $p < 1e-9$). In addition, we found a significant linear relationship (Pearson $r(31) = 0.855$, $p < 1e-9$) meaning that the exact number of samples subjects chose were proportionally similar to the Ideal Actor. One interesting feature of the data was that most of the errors in subjects responses were in the one-dimensional categories, which may be due to a general hesitancy to only sample information about one feature.

We can also perform the same cost analysis as for the previous experiment. Figure 6b shows that only one subject did not select measurements significantly better than the random




Figure 6: (a) Violin plot of subjects strategies as a function of category structure. (b) Expected number of additional incorrect trials relative to the Ideal Actor

baseline. We can also check whether subjects truly learned the categories: on average subjects had a 96% agreement with the Ideal Observer in the second half of the category learning phase and only 1 subject was below 90%. In addition, based on a MANOVA, there was no effect of category structure on suboptimal measurement selection or agreement with Ideal Observer (Wilk's $\Lambda = .88$, $F(2, 30) = .96$, $p = .43$) suggesting that none of the conditions were more difficult than the others.

Conclusion

We developed a new categorization paradigm in order to study people's strategies for information selection. These tasks allowed us to study human information selection in categorization tasks with measurement and perceptual noise, which we argue is the typical situation in everyday categorization. We analytically derived an Ideal Actor model of this task and from that derived two qualitative predictions for human behavior: 1) that subjects would be sensitive to the category structure and 2) their own perceptual noise. In Experiment 1, the predictions for perceptual noise were not fit to the selection task but estimated in a separate psychophysics task. Across two experiments, we demonstrated that most subjects take into account the category structure. The first experiment provided some evidence that subjects take into account perceptual noise as well although the evidence is somewhat weaker.

In order to get a better understanding of people's strategies, future work could address several additional questions including whether people are sensitive to the costs of information collection (see Meder and Nelson (2012) for some evidence that they do not) of different costs for correct or incorrect answers. Another direction might be whether people may be more sensitive to certain features of the categories (such as differences in the mean) than others (like the feature covariance). Finally, information selection has been proposed to be important in several other domains. Feature-based perceptual attention (Scolari, Edward, & Serences, 2014) can be thought of as a type of information selection and our model has parallels with some existing models in that literature (Palmer, 1990). A future experiment could use our model to investigate how people allocate perceptual resources during categorization. Many economists have recently investigated limited

information as an explanation for many economic phenomena (Caplin, 2015) but have often assumed that people collect information optimally. Using this model and measurement selection task could allow assessment of how people actually select information in choice situations.

Acknowledgments This research was supported by NSF grant BCS-1255538 and a John S. McDonnell Foundation Scholar Award to TMG.

References

- Anderson, T. W. (1958). *An introduction to multivariate statistical analysis* (1st ed.). John Wiley and Sons, Inc.
- Ashby, F. G., & Gott, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of experimental psychology. Learning, memory, and cognition*, *14*(1), 33–53.
- Ashby, F. G., & Townsend, J. T. (1986). Varieties of perceptual independence. *Psychological review*, *93*(2), 154–179.
- Blair, M. R., Watson, M. R., Walshe, R. C., & Maj, F. (2009). Extremely selective attention: eye-tracking studies of the dynamic allocation of attention to stimulus features in categorization. *Journal of experimental psychology. Learning, memory, and cognition*, *35*(5), 1196–206.
- Bostock, M., Ogievetsky, V., & Heer, J. (2011). D3 data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, *17*(12), 2301–2309.
- Caplin, A. (2015). Measuring and Modeling Attention. *Annual Review of Economics*, *8*(1), 1–43.
- Dakin, S. C., & Watt, R. J. (1997). The computation of orientation statistics from visual texture. *Vision Research*, *37*(22), 3181–3192.
- Jogan, M., & Stocker, A. (2014). A new two-alternative forced choice method for the unbiased characterization of perceptual bias and discriminability. *Journal of Vision*, *14*(3:20), 1–18.
- Kontsevich, L. L., & Tyler, C. W. (1999). Bayesian adaptive estimation of psychometric slope and threshold. *Vision Research*, *39*(16), 2729–2737.
- Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: a network model of category learning. *Psychological review*, *111*(2), 309–332.
- Matsuka, T., & Corter, J. E. (2008). Observed attention allocation processes in category learning. *Quarterly Journal of Experimental Psychology*, *61*(7), 1067–1097.
- Maule, J., & Franklin, A. (2016, Mar). Accurate rapid averaging of multihue ensembles is due to a limited capacity subsampling mechanism. *J. Opt. Soc. Am. A*, *33*(3), A22–A29.
- Meder, B., & Nelson, J. D. (2012). Information search with situation-specific reward functions. *Judgment and Decision Making*, *7*(2), 119–148.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*(3), 207–238.
- Nelson, J. D., McKenzie, C. R. M., Cottrell, G. W., & Sejnowski, T. J. (2010). Experience matters: information acquisition optimizes probability gain. *Psychological science*, *21*(7), 960–9.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of experimental psychology. General*, *115*(1), 39–61.
- Palmer, J. (1990). Attentional limits on the perception and memory of visual information. *Journal of experimental psychology. Human perception and performance*, *16*(2), 332–350.
- Rehder, B., & Hoffman, A. B. (2005). Eyetracking and selective attention in category learning. *Cognitive Psychology*, *51*(1), 1–41.
- Scolari, M., Edward, F., & Serences, J. T. (2014). Feature- and Object-Based Attentional Modulation in the Human Visual System. In *Oxford handbook of attention* (pp. 573–600).
- Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*(6), 1411–1436.