Reconciling Categorization and Memory via Environmental Statistics

Arjun Devraj¹, Thomas L. Griffiths^{1,2}, Qiong Zhang^{3,4,5}

Computer Science Department¹, Psychology Department², Princeton University;
Psychology Department³, Computer Science Department⁴, Center for Cognitive Science⁵,
Rutgers University – New Brunswick

Author Note

Experiment 1 appeared in a paper in the non-archival conference proceedings of the Annual Meeting of the Cognitive Science Society in 2021. A pre-print of the manuscript is available at psyarxiv.com. This study was preregistered. This work was supported by a CV Starr Fellowship awarded to Q.Z. and AFOSR FA9550-18-1-0077 awarded to T.L.G..

Correspondence concerning this article should be addressed to Qiong Zhang (qiong.z@rutgers.edu).

Abstract

How people represent categories and how those representations change over time is a basic question about human cognition. Previous research has demonstrated that people categorize objects by comparing them to category prototypes in early stages of learning but consider the individual exemplars within each category in later stages. However, these results do not seem consistent with findings in the memory literature showing that it becomes increasingly easier to access representations of general knowledge than representations of specific items over time. Why would one rely more on exemplar-based representations in later stages of categorization when it is more difficult to access these exemplars in memory? To reconcile these incongruities, our study proposed that previous findings on categorization are a result of human participants adapting to a specific experimental environment, in which the probability of encountering an object stays uniform over time. In a more realistic environment, however, one would be less likely to encounter the same object if a long time has passed. Confirming our hypothesis, we demonstrated that under environment statistics identical to typical categorization experiments the advantage of exemplar-based categorization over prototype-based categorization increases over time, replicating previous research in categorization. In contrast, under realistic environmental statistics simulated by our experiments the advantage of exemplar-based categorization over prototype-based categorization decreases over time. A second set of experiments replicated our results, while additionally demonstrating that human categorization is sensitive to the category structure presented to the participants. These results provide converging evidence that human categorization adapts appropriately to environmental statistics.

Keywords: category learning; prototype model; exemplar model; memory retrieval; environmental statistics; rational analysis

Reconciling Categorization and Memory via Environmental Statistics

Introduction

Computational models of categorization have played a key role in understanding the representations people use when learning categories. Two prominent models of categorization include the prototype model, which posits that a stimulus is categorized by comparing it to solely the “mean” object, or prototype, of each category (Posner and Keele, 1968; Posner and Keele, 1970; Reed, 1972; Mervis and Rosch, 1981; Homa et al., 1981), and the exemplar model, which asserts that a stimulus is categorized by comparing it with all of the objects for each category (Medin and Schaffer, 1978; Nosofsky, 1986). Earlier findings strongly favored the exemplar model (McKinley and Nosofsky, 1996; Medin and Schaffer, 1978; Shin and Nosofsky, 1992), but these studies focused on the final stage of learning and did not consider the progression of category learning over time. Smith and Minda (1998) studied the fits of the prototype and exemplar models of categorization to human performance on a word categorization task over the course of an experiment and found that although the exemplar model dominates in the end, the prototype model has a strong advantage in early stages of learning. This result is consistent with the proposal that individuals categorize using a simple rule-based approach in early stages of category learning and more exemplar-specific strategies in later stages (Nosofsky, Palmeri, et al., 1994a).

While prototype-based categorization and exemplar-based categorization have been extensively studied in the literature, less attention has been paid to the constraints of memory on categorization; namely, how much one is able to use prototype-based or exemplar-based categorization is constrained by how accessible these representations are in memory. Prototype-based categorization requires retrieving generalized knowledge or summary statistics about the category, whereas exemplar-based categorization requires retrieving memory about specific instances. Categorizing using generalized knowledge in early stages and switching to an exemplar-based model in later stages, as documented by

Smith and Minda (1998), does not seem consistent with findings in the memory literature showing that it becomes increasingly easier to access representations of general knowledge than representations of specific items over time—i.e., memory of specific items decays much faster than memory of the gist, or general knowledge (Posner and Keele, 1970; Zeng et al., 2021). Why would one rely more on exemplar-based representations in late stages of categorization when it is more difficult to access these exemplars in memory?

Memory of specific items decays quickly for a reason. From a rational perspective, there is no need to retain information in memory if it is no longer needed, as determined by the statistics of the environment (Anderson and Schooler, 1991). Therefore, memory decay obeys a power-law function, defined as

$$f(x) = ax^k, \tag{1}$$

because the relationship between stimulus recency and the need odds, the odds that the stimulus will be required in the future, often obeys the power law in realistic environments, as examined from environmental sources such as *The New York Times*, parental speech, and electronic mail (Anderson and Schooler, 1991; Figure 1a). In other words, we remember what happened five minutes ago better than what happened five days ago *because* what happened five minutes ago is more likely to be relevant to the current moment than what happened five days ago. In contrast, categorization experiments, including Smith and Minda (1998), usually use stimuli that appear with the same frequency regardless of the last time the stimulus was seen (Figure 1b). This creates unrealistic environmental statistics where what happened five minutes ago is equally likely to be relevant to the current moment compared with what happened five days ago. To adapt to this new environment, human participants would over-represent stimuli in their memory that were last seen a long time ago.

To systematically examine the effect of environmental statistics on categorization behavior, we create an experimental condition in Experiment 1 that better reflects the power-law relationship between stimulus recency and the need odds as seen from

Figure 1. The environmental recency function, displaying the relationship between stimulus recency and its need odds, calculated as $P(\text{stim})/(1 - P(\text{stim}))$ where $P(\text{stim})$ is the probability of seeing the stimulus. (A) is based on New York Times headlines (1986 and 1987), reproduced from Figure 4a in Schooler & Anderson (1991). (B) is calculated based on the experimental design of Smith and Minda (1998). (C) and (D) display average data from the 60 sequences of stimuli generated for the experimental condition in Experiment 1. In (C) and (D), $P(\text{stim})$ was calculated using a window size of 15 trials. The linear log-log relationship in (D) indicates that (C) is a power-law function, as taking the log of both sides of Equation 1 gives the linear function $\log(f(x)) = k \log(ax)$.

environmental statistics encountered in the real world (Figures 1c, 1d). We also create a control condition in which stimuli are presented with uniform frequency over time regardless of the last time a stimulus was seen, identical to Figures 1b, following closely the setup in Smith and Minda (1998). We hypothesize that setting up the categorization experiment with realistic environmental statistics in the experimental condition will reduce the accessibility of exemplar-based representations in memory over time (assuming that human memory rationally adapts to the environment), therefore reducing or inverting the trend previously observed in favoring prototype-based representations early on and exemplar-based representations at late stages.

The goal of the current work is to demonstrate how human categorization behavior is a result of rationally adapting to the structure of the environment—different trajectories of prototype-based versus exemplar-based representations will emerge under different environmental statistics. Our efforts in providing a rational account of categorization regarding the environmental structure are in the same spirit to previous work that provides a rational account of categorization regarding the category structure: specifically, when a learner should choose to use prototype-based or exemplar-based representations is determined by the category structure presented to the learner (Griffiths et al., 2007; Briscoe and Feldman, 2011). Prototype-based or exemplar-based representations correspond to different strategies of categorization, both being specific cases of a unifying model that is capable of optimally representing a given category structure (Griffiths et al., 2007). Similarly, it has been suggested that the literature that favors exemplar-based representations typically involve experiments that use poorly differentiated category structures (Medin et al., 1983; Medin and Schaffer, 1978; Medin and Schwanenflugel, 1981; Medin and Smith, 1981; Nosofsky, Gluck, et al., 1994; Nosofsky, Palmeri, et al., 1994b); those category structures can weaken the urge to form prototype-based clusters of exemplars (Smith and Minda, 1998). Given the important role of category structure during categorization, in Experiment 2, we explore whether results obtained in Experiment 1 can

be generalized to different category structures, by changing the number of category exceptions. We expect that when there are more category exceptions (non-linearly separable category structure), it is rational for participants to rely more on exemplar-based representations to encode these exceptions; whereas when there are no category exceptions (linearly separable category structure), prototype-based representations and exemplar-based representations are equally useful in supporting categorization performance. Of particular interest is how participants’ categorization strategies over different stages of learning, as examined in Experiment 1, interact with different category structures.

The plan of the paper is as follows. We first explain the formulations of the prototype and exemplar models of categorization. Next, we delineate how we constructed the presentation of stimuli in the experimental condition in order to reflect the environmental statistics of the real world. Finally, we discuss the behavioral and model results from both our replication of Smith and Minda (1998) in the control condition and our manipulations of the environmental statistics in the experimental condition. Lastly, we conduct a second set of experiments to examine how our findings regarding the effect of the environmental statistics generalize to different category structures. These results provide converging evidence that human categorization behavior is a result of rationally adapting to the structure of the environment.

Background

Prototype Model. Various versions of the prototype model exist (Medin and Schaffer, 1978; Reed, 1972); we use here the formulation specified by Smith and Minda (1998, 2011). The prototype model compares the observed stimulus to the prototype for each category. First, the distance d_{i,P_k} between stimulus i and the prototype for category $k \in \{1, 2\}$, P_k , is computed as

$$d_{i,P_k} = \sum_{j=1}^m w_j |i_j - P_{k,j}|, \quad (2)$$

where i_j is the stimulus' value for dimension j , and $P_{k,j}$ is the value of the prototype of category k for dimension j , and w_j is the attentional weight assigned to dimension j . Each attentional weight w_j is constrained to take on a value in $[0, 1]$, and the weights together sum to 1 ($\sum_{j=1}^m w_j = 1$).

Once the raw distance d_{i,P_k} has been calculated, stimulus i 's similarity, η_{i,P_k} , to the category k prototype P_k is computed as

$$\eta_{i,P_k} = e^{-c*d_{i,P_k}}, \quad (3)$$

where c is a sensitivity parameter constrained to $[0, 20]$. The sensitivity parameter has the effect of amplifying or shrinking psychological space (Smith and Minda, 1998). Finally, the probability that stimulus i (S_i) will be categorized into category 1 (R_1) is given by

$$P(R_1|S_i) = \frac{\eta_{i,P_1}}{\eta_{i,P_1} + \eta_{i,P_2}}. \quad (4)$$

Exemplar Model. The exemplar model has been developed and generalized from the original context model (Medin, 1975; Nosofsky, 1984; Nosofsky, 1986; Nosofsky, 1987; Nosofsky, 1988; Palmeri and Nosofsky, 1995; McKinley and Nosofsky, 1995). This formulation of the exemplar model was used and thoroughly described in Smith and Minda (1998). The exemplar model compares the observed stimulus to all of the previously seen exemplars in each category in order to generate a category prediction for the observed stimulus. The distance $d_{i,x}$ and similarity measure $\eta_{i,x}$ between stimulus i and exemplar x are calculated identically as in the prototype model. However, the exemplar model considers all exemplars in order to generate a prediction, so the probability that stimulus i (S_i) will be categorized into category 1 (R_1) is

$$P(R_1|S_i) = \frac{\sum_{x \in C_1} \eta_{i,x}}{\sum_{x \in C_1} \eta_{i,x} + \sum_{x \in C_2} \eta_{i,x}}, \quad (5)$$

where C_1 is the set of exemplars for category 1 and C_2 is the set of exemplars for category 2.

Because our experiment used stimuli with six dimensions, both models had a total of seven parameters: the six attentional weights ($w_1, \dots, w_6$) and the sensitivity parameter (c).

The two models were then fit to participant data from the control and experimental conditions of the experiment.

Experiment 1

Experiment 1 explores the possibility that participants' categorization strategies over different stages of learning depend on the environmental statistics. To this end, participants' performance is modeled using the basic and most prominent prototype and exemplar models discussed earlier. In the control condition, set up similarly to Experiment 2 of Smith and Minda (1998), we expect to replicate the previous results in which the prototype-based model would fit human data better early on and the exemplar-based model would fit human data better later in learning; in the experimental condition with realistic environmental statistics, we predict that the trend observed in the control condition would be reduced or inverted. Code and data for the experiment can be accessed through this link: TBD. The hypothesis, design, sample size, and analysis of Experiment 1 have been preregistered (<https://aspredicted.org/yq984.pdf>).

Methods

Procedure. The task required participants to repeatedly categorize 14 stimuli into two categories. The stimuli were six-letter nonsensical words taken from Appendix A of Smith and Minda (1998). Each stimulus can be thought of as a six digit string of bits, such that the prototype for category 1 was 000000 and the prototype for category 2 was 111111. Seven stimuli belonged to category 1 [000000, 100000, 010000, 001000, 000010, 000001, 111101], and the other seven stimuli belonged to category 2 [111111, 011111, 101111, 110111, 111011, 111110, 000100]. Each digit and position in the binary string corresponds to a unique letter; for example, the actual stimulus corresponding to 000000 that participants see is *gafuzi*, and the stimulus corresponding to 010000 is *gyfuzi*. In each trial, the participant was shown one of the 14 stimuli and had unlimited time to select a

response of 1 for category 1 and 2 for category 2. For exactly one second after each trial, the participant was shown *Correct* or *Incorrect* accordingly.

Figure 2. The stimulus sequences shown to participants were generated algorithmically from power-law functions. (A) Each curve corresponds to the power-law function for a particular stimulus. (B) Since data is aggregated, each curve corresponds to the i -th stimulus seen across all participants. For example, the first curve illustrates the average proportion of each bin allocated to the first stimulus (regardless of its identity) shown to the 60 participants. (C) Each color corresponds to a unique stimulus $i \in \{1, 14\}$, and the figure highlights the significant temporal clustering and subsequent decay caused by our stimulus generation algorithm based on power-law functions. Parameter values of $a = 1$, $k = -0.3$, $r = 1000$, $n = 350$, $s = 10$, and $\theta = a * (\frac{n}{r})^k$ were chosen to adequately match the power-law function in Figure 5a.

Control Condition. The control condition, similar to Smith and Minda’s (1998) experiment (the only difference being 616 trials in our experiment vs. 560 trials in the

original experiment), consisted of 616 trials divided into 44 blocks, each consisting of 14 trials. Each block consisted of a random permutation of the 14 stimuli such that each stimulus was shown exactly once per block.

Experimental Condition. To create the environmental statistics in Figure 1a, we developed an algorithm to generate a sequence of stimuli for the experiment. The specific pattern of stimulus presentation should, in turn, influence the need odds and hence the participant's retention of objects in memory. The experimental design first involved assigning each of the fourteen stimuli its own power-law function. We added several new parameters to the power-law function to generate the necessary experimental conditions:

$$f(t) = \begin{cases} 0 & t < t_0 \\ s & t = t_0 \\ a\left(\frac{t-t_0}{r}\right)^k - \theta & t_0 < t \leq t_0 + n \\ 0 & t > t_0 + n \end{cases} \quad (6)$$

where t is the trial number, t_0 is the start trial for the stimulus to be introduced, r is the range, n is the number of trials the object's power-law function should be sampled from, s is the starting value for the power-law function, and θ is a calibration parameter.

Power-law functions for consecutive stimuli were introduced in intervals of 35 trials. As in the control condition, each participant participated in a total of 616 trials. Figure 2a displays the resulting power-law functions for each stimulus over the entirety of the experiment. The order of the stimuli was randomly assigned for each participant.

To assign a stimulus to each trial based on the continuous power-law functions, we binned trials into discrete bins of 35 trials, generated a probability distribution for the bin based on the relative values of the various stimulus power-law functions over the trials in the bin, and generated the stimuli for trials in the bin by sampling from this distribution. Figure 2b shows how the resulting sequences of stimuli adequately reflected the original power-law functions. Figure 2c shows an example stimulus sequence for a single

participant. Finally, we plotted how the need odds vary as a function of stimulus recency given the sequences of stimuli generated by this method (Figures 1c, 1d), and verified that they closely reflect the environmental statistics in the real world such as those derived from headlines of *The New York Times* (Figure 1a).

Participants. We collected data from 60 participants for each condition (Control, Experimental). For reference, Smith and Minda (1998) had 32 participants for each experiment. Participants were recruited from the 18-23 age range and English-speaking population using Prolific.

Model Fitting. Trials were split and grouped into 11 trial segments. For the control condition, following the procedure in Smith and Minda (1998), the first trial segment corresponds to trials at the beginning of the experiment, and the last trial segment corresponds to trials towards the end of the experiment. For the experimental condition, the first trial segment corresponds to the earliest stage when a stimulus has just been introduced (i.e., at the beginning of the power-law functions in Figure 2a), and the last trial segment corresponds to the latest stage when a stimulus has been introduced for a long time (i.e., towards the end of the power-law functions in Figure 2a).

For the comparison of the control condition to Smith and Minda’s (1998) results, we used the sum of squared errors (SSE) between the observed and predicted probabilities as the measure of fit, which we sought to minimize:

$$SSE = \sum_{i=1}^{14} (P(R_1|S_i) - \hat{p}_{1,S_i})^2 \quad (7)$$

where $P(R_1|S_i)$ is the model-generated probability that stimulus i belongs to category 1 based on an entire trial segment (56 trials) of data, and $p_{1,i}$ is the proportion of trials in the trial segment (out of those in which stimulus i was seen) in which the participant actually categorized stimulus i in category 1.

We ran Scipy’s Sequential Least Squares Programming (SLSQP) method with 10 initial random configurations to obtain the best-fit parameters. The model was fit for each participant receiving the control condition over each trial segment, as in Smith and Minda

(1998). The resulting best-fit parameters were used to calculate the SSE that was ultimately used as the measure of fit in Figure 4. For the experimental condition, because trials in the same trial segments did not represent chronological groupings of trials in absolute time, the model was fit for each participant over the entire experiment. This meant that we used the mean-squared error (MSE) over all trials in the trial segment, with residuals computed on a trial-by-trial basis, as the measure fit:

$$MSE = \frac{1}{56} \sum_{j=1}^{56} (P(R_1|T[j]) - \hat{p}_{1,T[j]})^2 \quad (8)$$

where j is the trial number, $T[j]$ corresponds to the stimulus S_i that was seen on trial j of the trial segment, and $\hat{p}_{1,T[j]}$ is either 0 or 1 depending on the participant's observed prediction on trial j . We ensured that the exemplar model only considered exemplars seen *thus far* since not all exemplars were seen from the beginning in the experimental condition.

Results

Figure 3. Average categorization accuracy over trial segments. Error bars denote the standard error of the mean.

Behavioral Results. The categorization accuracies per trial segment in each condition were averaged over all participants and graphed over time, resulting in Figure 3.

In the control condition, significant learning occurred between the first ($M = .62$, $SD = .11$) and last ($M = .73$, $SD = .18$) trial segment, $t(59) = -5.08$, $p < .001$. In the experimental condition, under realistic environment statistics where a stimulus is presented less and less frequently over trial segments, categorization accuracy decreases from the second ($M = .92$, $SD = .08$) to last ($M = .81$, $SD = .15$) trial segment, $t(59) = 7.99$, $p < .001$. The second trial segment was compared in the experimental condition because this was the point that best reflected the beginning of the power-law function—when the participant had repeatedly seen the stimulus and was past the very initial stage of learning.

Figure 4. Comparison of model fits of prototype-based versus exemplar-based categorization for (A) the original experiment, reproduced from Smith and Minda (1998), (B) the control condition, and (C) the experimental condition after introducing realistic environmental statistics. Models were fit for each participant’s performance over a trial segment (a,b) or over the entire experiment (c) and then averaged over all participants for the condition. Error bars denote the standard error of the mean.

Model Results. The model results over trial segments of the control condition in Figure 4b are compared with the results from Smith and Minda (1998) in Figure 4a. We qualitatively replicate the key results from Smith and Minda (1998), where the advantage of the exemplar model over the prototype model is increasing over time. To formally test if the difference in model fits between the prototype and exemplar models changed over the course of the experiment, we ran a two-way ANOVA on the model fits with the trial

segment (first, last) and model type (prototype, exemplar) as within-subject variables. There is significant interaction between model type and trial segment ($F(1, 59) = 19.72$, $p < .001$). The model results over trial segments of the experimental condition (Figure 4c) are compared with the results from the control condition, when also fit over the entire experiment and using MSE for fit. In contrast to the control condition, under environmental statistics that reflect real-world environments, the advantage of the exemplar model over the prototype model is decreasing over trial segments in the experimental condition ($F(1, 59) = 19.72$, $p < .001$).

Experiment 2

In addition to environmental statistics, Experiment 2 also assessed the effect of category structure on the fits of the prototype and exemplar models, by altering the number of category exceptions (i.e., 0, 2, 4). Similarly to Experiment 1, participants' performance is modeled using the basic and most prominent prototype and exemplar models. Since the prototype model is equivalent to assuming a linear boundary between categories (Ashby and Maddox, 1993), we expect that when there are more category exceptions (corresponding to a less linearly separable category structure), it is rational for participants to rely more on exemplar-based representations to encode these exceptions; whereas when there are no category exceptions (resulting in a linearly separable category structure), participants would rely similarly on both prototype-based or exemplar-based representations. Of particular interest is how participants' category strategies over different stages of learning depends jointly on the category structure and environmental structure. Since Experiment 1 was conducted using a category structure with 2 exceptions, we additionally expect to replicate the results in Experiment 1 in the conditions with 2 exceptions in Experiment 2. Code and data for the experiment can be accessed through the same link: TBD. The hypothesis, design, sample size, and analysis of Experiment 2 have also been preregistered (<https://aspredicted.org/9me7s.pdf>).

Methods

Procedure. Experiment 2 has a total of six conditions. We examine category structures with 0 exceptions, 2 exceptions (as in Experiment 1), and 4 exceptions, with a control and experimental condition for each level of exceptions. The sequences of stimuli for each of the experimental and control conditions were generated identically as in Experiment 1, using the same binning and sampling method with the piece-wise power law function in Equation 6 for the experimental condition and repeated blocks of random permutations of the 14 stimuli for the control condition; in comparison to Experiment 1, only the actual stimuli changed, depending on the number of exceptions. Stimuli were still six-letter nonsensical words modeled by six-bit binary strings, with the category prototypes of 000000 for category 1 and 111111 for category 2. For the conditions with 0 exceptions, the seven stimuli in category 1 were 000000, 100000, 010000, 001000, 000010, 000001, and 000100, and the seven stimuli in category 2 were 111111, 011111, 101111, 110111, 111011, 111110, and 111101. For the conditions with 2 exceptions, the category structure was identical to that described in Experiment 1: 000000, 100000, 010000, 001000, 000010, 000001, and 111101 for category 1 and 111111, 011111, 101111, 110111, 111011, 111110, 000100 for category 2. For the conditions with 4 exceptions, the seven stimuli in category 1 were 000000, 100000, 010000, 000010, 000001, 111101, and 101111, and the seven stimuli in category 2 were 111111, 011111, 110111, 111011, 111110, 000100, and 001000. Each digit and position in the binary string corresponds to a unique letter; for example, the actual stimulus corresponding to 000000 that participants see is **gafuzi**, and the stimulus corresponding to 010000 is **gyfuzi**.

Participants. This experiment was run on Prolific with 60 participants for each of the six conditions. Same as in Experiment 1, participants were recruited from the 18-23 age range and English-speaking population.

Model Fitting. The model-fitting procedure for both control and experimental conditions was identical to that in Experiment 1.

Results

Figure 5. Average categorization accuracy over trial segments for all three control conditions and all three experimental conditions, by number of exceptions. Error bars denote the standard error of the mean.

Behavioral Results. We replicated key behavioral results from Experiment 1: in control conditions (0 exceptions; 2 exceptions; 4 exceptions), participants improve their categorization accuracy over trial segments, whereas in experimental conditions, participants' categorization accuracy decreases over trial segments. Additionally, we examined the effect of exceptions on categorization accuracy. As the number of exceptions increases, categorization accuracy declines for both the control and experimental conditions. To formally test these effects, we conducted a two-way mixed ANOVA on the participants' categorization accuracies with trial segment (first/second, last) as a within-subject variable and the number of category exceptions (0, 2, 4) as a between-subjects variable for both the control and experimental conditions. There were significant effects on categorization accuracy by both the number of exceptions (control condition: $F(2, 177) = 91.78$, $p < .001$; experimental condition: $F(2, 177) = 22.61$, $p < .001$), and trial segment (control condition: $F(1, 177) = 173.53$, $p < .001$; experimental

condition: $F(1, 177) = 14.76, p < .001$).

Model Results. There are two goals for the modeling analyses. The first goal is to examine the effect of category structure (i.e., number of exceptions) on the fits of the prototype and exemplar models. Figure 6 shows that as the number of exceptions increases, the exemplar model’s advantage over the prototype model also increases. We conducted a two-way ANOVA that analyzed the effects of model type (prototype, exemplar), a within-subject variable, and number of exceptions (0, 2, 4), a between-subjects variable, on the model fits in the middle (fifth) trial segment. The interaction between model type and number of exceptions was significant for both the control conditions, $F(2, 177) = 3.82, p = .024$, and experimental conditions, $F(2, 177) = 160.05, p < .001$.

The second goal of the modeling analyses is to examine whether results obtained in Experiment 1 can be generalized to different category structures (linearly-separable: 0 exceptions; non-linearly separable: 2 exceptions, 4 exceptions). Consistent with the observations in Experiment 1, for conditions with two or four exceptions, the exemplar model’s advantage over the prototype model decreases over time in the experimental condition but increases over time in the control condition, as shown in Figure 6. To formally test these effects, we ran a two-way ANOVA to analyze the effects of model type (prototype, exemplar) and trial segment (first/second, last) on the model fits. For the conditions with two exceptions (i.e., replication of Experiment 1), the interaction between model type and trial segment was significant for both the control condition, $F(1, 59) = 27.06, p < .001$, and the experimental condition, $F(1, 59) = 68.59, p < .001$. Similarly, with four exceptions, the interaction was significant for both the control condition, $F(1, 59) = 14.03, p < .001$, and experimental condition, $F(1, 59) = 93.08, p < .001$. However, with zero exceptions, model fits for both the prototype and exemplar models follow nearly identical trends over time and do not demonstrate the effects previously observed in Experiment 1.

Figure 6. Comparison of model fits of prototype-based versus exemplar-based categorization for all six conditions, each defined by (1) control or experimental version of the experiment and (2) the number of exceptions in the category structure. Models were fit for each participant’s performance over a trial segment (A-C) or over the entire experiment (D-F) and then averaged over all participants for the condition. Error bars denote the standard error of the mean.

General Discussion

While prototype-based and exemplar-based categorization have been extensively studied in the literature, less attention has been paid to the constraints of memory on categorization under different environmental statistics. In our results, we demonstrate that human categorization rationally adapts to the environmental statistics. In Experiment 1, our findings from the control condition qualitatively replicated Smith and Minda’s (1998) findings that when stimuli are presented with uniform frequency over time, there is an opportunity to over-learn the stimuli, which leads to an increased advantage of the

exemplar model over the prototype model over time. The results from the experimental condition confirm our hypothesis: when stimulus presentation more accurately reflects realistic environmental statistics, there is reduced need to retrieve the same exemplars as time goes by, leading to a decreased advantage of the exemplar model over the prototype model. Experiment 2 replicated our results, while additionally demonstrating that human categorization is sensitive to the category structure presented to participants: the trend observed in adapting to different environmental statistics only extends to category structures with exceptions present. We now turn to the broader implications of these results.

It has been pointed out that memory and categorization rest on similar representations, but differ in whether identification is made at the instance/individual level or the category level (Anderson, 1991). By adjusting the level of identification, rational models of categorization can be used to explain results of fan effects in a memory task (Anderson, 1991). In spite of the theoretical link on how categorization relates to memory, there has been a lack of empirical exploration in this direction, as measures that are of interest to memory researchers (e.g., separate analyses on instance-level representations) have not been directly addressed by category learning paradigms (De Brigard et al., 2017). However, some work in the literature has started to draw parallels between category learning and memory studies. In several studies, following standard categorization experiments where participants made judgement at the category level, participants completed an additional recognition test that allows for probing the representations at the instance level. These studies identified a number of parallels between category learning and memory studies on the effects of schemas (Palmeri and Nosofsky, 1995; Rojahn and Pettigrew, 1992; Sakamoto and Love, 2004).

The current paper is another step towards linking the two pillars of cognitive psychology: memory and categorization. Our findings demonstrate that what can account for rational memory behavior can also be used to explain categorization behavior. Since

categorization requires not only retrieving category representations—which characterizes memory behavior—but also requires a step to use these representations to make an identification on the category, it is expected that what influences retrieval may also have downstream effects on the categorization behavior. Our results are consistent with the rational analysis of memory by Anderson (1990) in deciding whether one should retrieve a particular memory representation. A rationally designed information-retrieval system would retrieve memory structures ordered by their need probability p , and stop retrieving when $pG < C$, where G is the reward associated with retrieving a target memory and C is the cost in considering the memory. Under realistic environmental sources where p decreases as time goes by (Anderson and Schooler, 1991; Figure 1a), this retrieval strategy predicts that information that has occurred more recently is more likely to be retrieved by the memory system (Anderson, 1990). Similarly, in our categorization experiments, under realistic environmental statistics where p decreases as time goes by (i.e., experimental condition), stimuli that have occurred further in the past are less likely to show up again and therefore less likely to be retrieved during categorization, leading to decreased access to exemplar-based representations over time; this contrasts with environmental statistics typically used in the categorization literature (i.e., control condition), where p stays constant as time goes by. The latter case provides an opportunity for over-learning individual stimuli, leading to increased access to exemplar-based representations over time regardless of when the stimuli were last seen. To summarize, our results from the control and experimental conditions in Experiments 1 and 2 are consistent with predictions from a rational-retrieval strategy when p is determined by the environmental statistics.

Furthermore, our results from Experiment 2 on the effect of category structure are consistent with predictions from a rational-retrieval strategy when the reward G is determined by the category structure. Specifically, when there are no exceptions, prototype-based representations or exemplar-based representations are equally useful in supporting categorization performance, leading to the same G values. When there are

category exceptions present, there is larger G associated with exemplar-based representations, as exemplar-based representations are capable of encoding both exceptions and non-exceptions whereas prototype-based representations provide a summary statistics representative of only non-exceptions. As a result, we observe that the advantage of exemplar-based representations over prototype-based representations increases as the number of exceptions increases, and that when there are no exceptions, there is no change in the difference of exemplar-based and prototype-based model fits over different stages of category learning.

While the memory literature shows that it becomes increasingly easier to access representations of abstract knowledge than representations of specific items over time (Posner and Keele, 1970; Zeng et al., 2021), categorization has been shown to rely on abstract knowledge in early stages and switch to an exemplar-based model in later stages (Smith and Minda, 1998). Our findings help to reconcile the incongruities between these results by demonstrating that different categorization behavior arises as a result of rationally adapting to different environmental statistics. In both experiments, human participants rationally adjust to the need probability, relative costs, and rewards associated with each object. Future work will explore more complex environmental statistics and cost-reward structures to better understand how memory and categorization work in tandem to enable efficient category learning in the real world.

References

- Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, 2, 396–408.
- Anderson, J. R. (1990). *The adaptive character of thought*. Psychology Press.
- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98(3), 409–429.
- Ashby, F. G., & Maddox, W. T. (1993). Relations between prototype, exemplar, and decision bound models of categorization. *Journal of Mathematical Psychology*, 37, 372–400.
- Briscoe, E., & Feldman, J. (2011). Conceptual complexity and the bias/variance tradeoff. *Cognition*, 118(1), 2–16.
- De Brigard, F., Brady, T. F., Ruzic, L., & Schacter, D. L. (2017). Tracking the emergence of memories: A category-learning paradigm to explore schema-driven recognition. *Memory & Cognition*, 45(1), 105–120.
- Griffiths, T., Canini, K., Sanborn, A., & Navarro, D. (2007). Unifying rational models of categorization via the hierarchical Dirichlet process. *Proceedings of the 29th Annual Conference of the Cognitive Science Society*, 323–328.
- Homa, D., Sterling, S., & Trepel, L. (1981). Limitations of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Human Learning and Memory*, 7(6), 418–439.
- McKinley, S. C., & Nosofsky, R. M. (1995). Investigations of exemplar and decision bound models in large, ill-defined category structures. *Journal of Experimental Psychology: Human Perception and Performance*, 21, 128–148.
- McKinley, S. C., & Nosofsky, R. M. (1996). Selective attention and the formation of linear decision boundaries. *Journal of Experimental Psychology: Human Perception and Performance*, 22, 294–317.

- Medin, D. L. (1975). A theory of context in discrimination learning. In G. H. Bower (Ed.), *The psychology of learning and motivation*. Academic Press.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85, 207–238.
- Medin, D. L., & Smith, E. E. (1981). Strategies and classification learning. *Journal of Experimental Psychology: Human Learning and Memory*, 7, 241–253.
- Medin, D. L., Dewey, G. I., & Murphy, T. D. (1983). Relationships between item and category learning: Evidence that abstraction is not automatic. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9(4), 607–625.
- Medin, D. L., & Schwanenflugel, P. J. (1981). Linear separability in classification learning. *Journal of Experimental Psychology: Human Learning and Memory*, 7(5), 355.
- Mervis, C. B., & Rosch, E. (1981). Categorization of natural objects. *Annual Review of Psychology*, 32(1), 89–115.
- Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 104–114.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115, 39–57.
- Nosofsky, R. M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 87–109.
- Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 700–708.
- Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994a). Rule-plus-exception model of classification learning. *Psychological Review*, 101, 53–79.

Nosofsky, R. M., Gluck, M. A., Palmeri, T. J., McKinley, S. C., & Glauthier, P. (1994).

Comparing modes of rule-based classification learning: A replication and extension of Shepard, Hovland, and Jenkins (1961). *Memory & Cognition*, 22(3), 352–369.

Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994b). Rule-plus-exception model of classification learning. *Psychological review*, 101(1), 53–79.

Palmeri, T. J., & Nosofsky, R. M. (1995). Recognition memory for exceptions to the category rule. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 548–568.

Posner, M. I., & Keele, S. W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology*, 83, 304–308.

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77(3, Pt.1), 353–363.

Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, 3, 382–407.

Rojahn, K., & Pettigrew, T. F. (1992). Memory for schema-relevant information: A meta-analytic resolution. *British Journal of Social Psychology*, 31(2), 81–109.

Sakamoto, Y., & Love, B. C. (2004). Schematic influences on category learning and recognition memory. *Journal of Experimental Psychology: General*, 133(4), 534–553.

Shin, H. J., & Nosofsky, R. M. (1992). Similarity-scaling studies of dot-pattern classification and recognition. *Journal of Experimental Psychology: General*, 121, 278–304.

Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology*, 24, 1411–1436.

Zeng, T., Tompary, A., Schapiro, A. C., & Thompson-Schill, S. L. (2021). Tracking the relation between gist and item memory over the course of long-term memory consolidation [preprint]. *bioRxiv*, 304–308.

<https://www.biorxiv.org/content/10.1101/2021.01.05.425378v1>