

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Developmental Differences in the Status of Category Exceptions

Permalink

<https://escholarship.org/uc/item/5mp2p1m9>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 40(0)

Authors

Savic, Olivera

Blanco, Nathaniel

Sloutsky, Vladimir

Publication Date

2018

Developmental Differences in the Status of Category Exceptions

Olivera Savic (savic.7@osu.edu)

Department of Psychology, 1835 Neil Avenue,
Columbus, OH 43210, United States

Nathaniel Blanco (blanco.53@osu.edu)

Department of Psychology, 1835 Neil Avenue,
Columbus, OH 43210, United States

Vladimir Sloutsky (sloutsky.1@osu.edu)

Department of Psychology, 1835 Neil Avenue,
Columbus, OH 43210, United States

Abstract

In this paper we explored how people represent categories that include exceptions by examining contributions that features of regular and exception items make to determining category membership. We examined performance of 4-year-old children and adults and found significant developmental differences. While for 4-year-olds, deterministic features of regular items and exceptions contributed comparably to determining category membership, an asymmetry was found in adults. For adults, deterministic features of regular items contributed more to determining category membership than features of exceptions. The results are discussed in relation to the SUSTAIN clustering model of category learning (Love, Medin, & Gureckis, 2004).

Keywords: category learning; exceptions; SUSTAIN

Introduction

The primary goal of category learning is to allow for generalization of acquired knowledge to new instances. If we know that something belongs to the category of *birds*, we could assume that it has feathers, it lays eggs and it has wings that it uses to fly. However, in order for categories to be useful and adequately govern predictions and behavior, the learner needs to discover and encode the key attributes of the members of the category. If a learner represents birds as flying creatures, she will misclassify numerous flightless birds. Thus, learner's representation of birds would need to be updated so it includes flightless birds. Although it is hardly controversial that even young children can learn that penguin is a bird although it does not fly, there is little evidence on how such information affects boundaries of previously learned categories.

It is reasonable to expect that the contribution that features of penguins have in determining membership in the category of birds may depend on the nature of the representation of regular members of the category of birds (flying birds). Previous studies have shown that while older children and adults tend to form rule-based representations, focusing primarily on the most relevant features of the category, younger children tend to represent categories in a similarity-based manner, relying on both relevant and irrelevant

category features (see Sloutsky, 2010; Deng & Sloutsky, 2015, 2016). Since the way we represent regular members of categories changes through development, the status of exceptions and their role in categorization may also change.

The dominant view in categorization literature is that items that violate category expectations need to be represented separately (Palmeri & Nosofsky, 1995; Sakamoto & Love, 2004; Davis, Love & Preston, 2012). The special status of exceptions is based on the assumption that in the case of regular category members, learners optimize and represent only the key characteristics of the category – those that are highly predictive of category membership. Since for exceptions this kind of optimization may not be possible or needed (because they are rare), all of their features could be equally well represented, which could result in better memory for exceptions (Palmeri & Nosofsky, 1995). Based on this approach, regular and exception items are represented differently and thus their features may contribute differently to category membership decisions. While key features of regular items should be highly predictive of category membership, most of the features of exceptions may be comparably predictive or only predictive when considered together.

With developmental differences in mind, the special status of exceptions as described in previous studies (Palmeri & Nosofsky, 1995; Sakamoto & Love, 2004; Davis, Love & Preston, 2012) and a need for a separate representation could be predicted for adults but not for younger children. Since young children tend to rely on both relevant and irrelevant features when categorizing regular members of a category (Sloutsky, 2010; Deng & Sloutsky, 2015, 2016), they may represent both regular and exception items in the same way and thus the features of both item types may have comparable contributions in determining category membership.

Current study

Here we present a study whose main goal was to determine unique contributions that features of regular and exception items make to determining category membership. We further investigated how this may change through development.

Four-year-old children and adults were trained on two categories, each including an exception item. Following training, categorization performance was tested on new (generalization) items.

Previous studies (Palmeri & Nosofsky, 1995; Sakamoto & Love, 2004; Davis, Love & Preston, 2012) typically focused on a particular kind of exceptions – items that violate rule-defined structure by following the contrasting category rule. While this is an interesting case for models of category learning to study, this category structure is highly constraining. Further, it is unclear whether insights based on this peculiar structure could be generalized to other categories with exceptions – categories that are not necessarily represented by rules and exception items that do not follow the contrasting category rule.

In contrast to the majority of previous studies, we used a category structure of regular items that could be successfully learned both by forming a rule and by relying on the overall similarity between category members. Accordingly, exceptions from these categories could be represented based on overall appearance or individual features, since all their features were unique, different from the features of the members of their own category and contrasting category members.

Method

Participants

Participants were 28 four-year-old children (Mage = 54.1 months, range 47.2–59.6 months, 20 girls) and 47 adults (Mage = 231.3 months, range 216.1–289.7 months, 31 female).

Four-year-olds were recruited from preschools and childcare centers located in middle-class suburbs of Columbus. They were tested during their regular school hours in a quiet room in their preschool or childcare center. Adult participants were The Ohio State University undergraduate students. They were tested in a quiet room in the laboratory located on campus and they received course credits for their participation.

Materials

Stimuli were two categories of artificial creatures (Figure 1) that were created using Spore Creature Creator (Electronic Arts, 2009) and GIMP (GNU Image Manipulation Program; version 2.6.11; The GIMP Team). The two categories' creatures were accompanied by two novel labels: Momo and Lulu.

The Category Structure Participants were trained on two 6-dimensional categories (see Table 1). Both categories had five Regular members and one Exception item. Regular members had one deterministic and five probabilistic feature dimensions. The deterministic dimension was fully predictive, while probabilistic dimensions varied both between categories and within-category and were predictive only when taken together. Exceptions had all features unique

(different from within category members and different from the other Exception). Thus, all dimensions of Exceptions were fully predictive.

Despite the differences in their feature structure, we will refer to the features of both Item Types (i.e. Regulars and Exceptions) as probabilistic and deterministic in accordance to the structure of the Regular category members.

Table 1 shows kinds of items used during training and test phase. During training and Categorization test, participants were presented with Regular and Exception items which had all of their 6 features presented.

In the Generalization test there were eight kinds of items (see Table 1). These new items were based on Regular and Exception items participants saw during the training phase, but had either some features covered or in addition to some features being covered also had the deterministic feature incongruent with either type or category of the probabilistic features (see Figure 1).

The item structure presented in Table 1, allowed us to test participants categorization performance when they had to rely on 3 probabilistic features only (R-3P, E-3P; where R = Regulars, E = Exception, P = Probabilistic features) and to compare it to the performance when they could also (in addition to 3 probabilistic features) rely on deterministic feature that was either congruent (R-3P+D, E-3P+D; where D = Deterministic feature) or incongruent, i.e. belonged to the contrasting category (R-3P+Dc, E-3P+Dc; where Dc = Deterministic feature of the contrasting Category) or contrasting Item Type (R-3P+Dt, E-3P+Dt; where Dt = Deterministic feature of the contrasting Item Type). This design allowed us to estimate unique contributions of different types of deterministic features in making decision on category membership.

Procedure

All instructions and questions were phrased in the same way for both age groups. For adult participants they were presented on the screen, while for 4-year-olds they were read by a trained experimenter. Adult participants responded by pressing designated keys on a computer keyboard, while verbal responses of 4-year-olds were collected by a trained experimenter using a computer keyboard.

Experiment had four phases: instructions, training, Categorization test and Generalization test.

Instructions Following the cover story, participants were presented with the prototypes of the two categories. The prototypes were presented together, on the same screen, and participants were told that that is how members of the two categories usually look. Further, each of the six items' features was introduced. Probabilistic features were introduced first using the following sentence frame: "Momos usually have antennas like these, and Lulus usually have antennas like these". Deterministic feature was introduced last using the sentence: "All Momos have this kind of hands, and all Lulus have this kind of hands".

Training Participants were presented with the exemplars of the two categories and asked to classify them. Each item was presented individually. Participants were provided with corrective feedback after each response. Feedback was different for the two Items Types. For Regulars feedback emphasized the deterministic feature (“That’s right!/Oops! That’s a Momo (Lulu), because it has Momo’s (Lulu’s) hands!”). For Exceptions feedback emphasized their exceptions’ status (“That’s right!/Oops! That’s a Momo (Lulu), but it is a tricky one!”).

Participants were first presented with 10 Regular category members, after which additional 20 Regular and 12 Exception items were presented in random order.

Categorization test Participants were presented with 8 Regular and 8 Exception items that they saw during the training phase and they were asked to classify them. There was no feedback provided. The main purpose of the Categorization test was to provide an estimate of level of learning.

Generalization test Participants’ task was the same as in the Categorization test, but they were presented with new items (See Table 1). They were not provided with feedback for their response. The total number of trials was 64, 8 trials of each item type.

Table 1: The abstract category structure.

	Category A		Category B	
	Probabilistic	Rule	Probabilistic	Rule
Training (Categorization test)				
Regular	1	0	0	1
	0	1	0	1
	0	0	1	0
	0	0	0	1
	0	0	0	1
Exception	2	2	2	2
Generalization test				
R-3P	-	-	0	0
R-3P+D	-	-	0	0
R-3P+Dc	-	-	0	0
R-3P+Dt	-	-	0	0
E-3P	-	-	2	2
E-3P+D	-	-	2	2
E-3P+Dc	-	-	2	2
E-3P+Dt	-	-	2	2

*R = Regulars, E = Exception, P = Probabilistic features, D = Deterministic feature, “-“ = bubble (feature is covered), Dt = Deterministic feature of the contrasting Item Type, Dc = Deterministic feature of the contrasting Category.

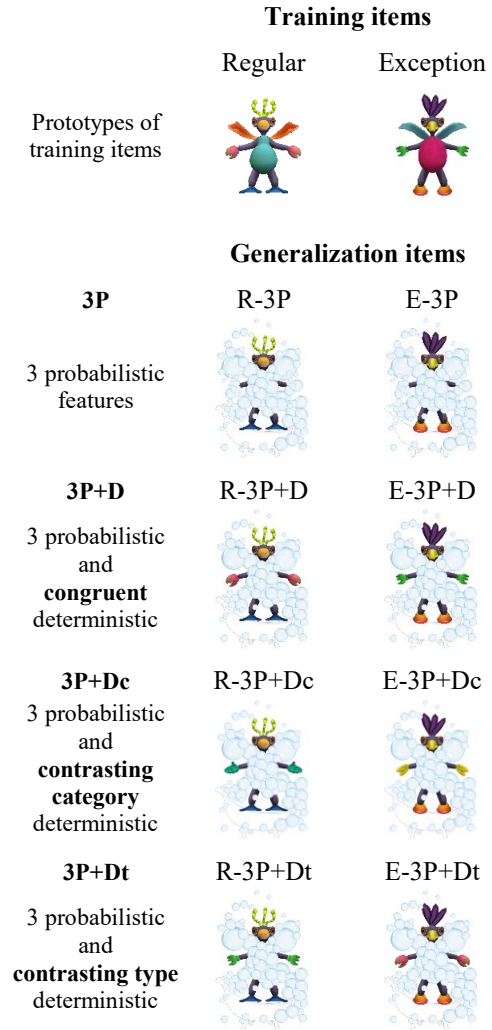


Figure 1: Examples of stimuli used in the study. Hands are the rule feature.

Results and Discussion

Preliminary Analyses: Categorization

The preliminary analyses focused on participants’ performance on items they were trained on. Table 2 shows the average proportion of accurately classified items during the Categorization test.

Both adults’ and 4-year-olds’ overall accuracy was above chance for both Item Types (one-sample t s > 6.24 , p s $< .001$, two-tailed). Differences in categorization accuracy were tested in a 2 (4-year-olds vs. Adults) \times 2 (Regulars vs. Exceptions) mixed ANOVA, with Age as a between-subjects factor and Item Type as a within-subjects factor. Adults performed better than 4-year-olds ($F(1, 73) = 13.91$, $MSE = .044$, $p < .001$, $\eta_p^2 = .160$), but both age groups performed equally well on both Item Types (Item Type: $F(1, 73) = .04$, $MSE = .029$, $p > .05$, $\eta_p^2 = .001$; Item Type \times Age: $F(1, 73) = 3.04$, $MSE = .029$, $p > .05$, $\eta_p = .040$).

Table 2: Mean accuracy (standard deviations) for performance at Categorization test

Age	Regular items	Exceptions
4-year-olds	0.73 (0.20)	0.77 (0.23)
Adults	0.91 (0.13)	0.85 (0.21)

Based on their performance on the Categorization test participants were selected for the main analyses. Six participants (3 from each age group) were excluded due to extremely poor performance (chance performance on both Item Types). In order to account for potential effects of categorization accuracy on generalization patterns, the remaining participants were classified as *good performers* (performed above .75 (6 out of 8) on both Regulars and Exceptions) or *poor performers* (performed below .75 on one of the Item Types). 36 (out of 45) adults and 11 (out of 25) 4-year-olds were classified as *good performers*. We performed all the analyses both on the subsample of good performers and on the whole sample. Since there were no differences in patterns of the results, we will report analyses performed on the whole sample only.

Generalization

We analyzed how congruent and incongruent deterministic features influence category decisions. Participants' performance based on probabilistic features only (baseline: 3P) was compared to their performance when in addition to probabilistic features they could also rely on deterministic feature. The deterministic feature was either congruent with probabilistic features (congruent: 3P+D), or incongruent in type (incongruent type: 3P+Dt) or category (incongruent type: 3P+Dc). To remind, in the incongruent category trials, a D feature from Category B was added to Category A or a D feature from Category A was added to Category B. In the incongruent type trials, a D feature from a regular item was added to an exception item or a D feature from an exception item was added to a regular item (with both replacements made within a category).

We estimated unique contributions of different types of deterministic features by subtracting categorization accuracy on items containing deterministic feature (3P+D; 3P+Dt; 3P+Dc) from categorization accuracy on 3P items. Estimates of contribution of a) congruent deterministic feature (3P – 3P+D), b) deterministic feature of contrasting Item Type (3P – 3P+Dt) and c) deterministic feature of contrasting category (3P – 3P+Dc) were subjected to 3 separate mixed ANOVAs with Item Type (Regulars vs. Exceptions) as a within-subjects factor and Age (4-year-olds vs. Adults) as a between-subjects factor.

Contribution of Congruent Deterministic Feature. The contribution of the congruent deterministic feature (see Figure 2A) was larger for Regular than for Exception items ($F(1, 67) = 8.98, MSE = .046, p < .01, \eta_p^2 = .118$). The main effect of Age and the effect of interaction were not significant

(both $ps > .10$). Advantage for Regulars' deterministic feature is expected and it shows that participants learned the category structure.

Contribution of Deterministic Feature of Contrasting Type These data are presented in Figure 2B. The analysis revealed main effect of Item Type ($F(1, 67) = 12.90, MSE = .048, p = .001, \eta_p^2 = .16$) and interaction between Item Type and Age ($F(1, 67) = 5.33, MSE = .048, p < .05, \eta_p^2 = .074$).

For 4-year-olds adding deterministic feature of the other Item Type (but same category) significantly increased performance for both Regular and Exception items. However, for Adults, only Exception items benefited from the deterministic feature of the Regular items, whereas features of Exceptions did not affect performance on Regular items. This resulted in asymmetry that can be observed in Figure 2 (panel b).

Contribution of Deterministic Feature of Contrasting Category

These data are presented in Figure 2C. The analyses revealed significant main effect of Item Type ($F(1, 67) = 20.99, MSE = .060, p < .001, \eta_p^2 = .239$) and significant Age by Item Type interaction ($F(1, 67) = 8.79, MSE = .060, p < .001, \eta_p^2 = .116$). The main effect of Age was not significant.

For both Age groups and Item Types, adding the contrasting category deterministic features moved participants membership decisions towards the category determined by the deterministic feature. Adults' performance was again significantly stronger affected by the deterministic feature of Regular items than Exceptions', while for children difference in the contributions of the two Item Types were small. This pattern is in accordance with the effect observed in the analyses presented above.

To sum up, we found significant developmental differences in the contribution of features of Regular and Exception items in determining category membership.

For younger participants (4-year-olds) features of both Item Types contributed comparably in determining category membership. Additionally, contributions were the same when deterministic features were congruent with probabilistic features and when they were added to the probabilistic features of the contrasting Category or contrasting Item Type.

Patterns of responses of Adults have revealed asymmetry in contributions of the two Item Types. Deterministic features of Regular items contributed significantly more than ones of Exceptions, to both Item Types. Note that adult participants did learn deterministic features of Exceptions, which can be seen both in small contribution they have when contributing to the congruent probabilistic features (Figure 2, panel a) and in a more significant effect they have when added to the probabilistic features of the contrasting category (Figure 2, panel c). However, adults did not rely on Exceptions' features when determining category membership of Regular items. On the other side, deterministic features of Regular items contributed equally when added to the probabilistic features of Regulars or when added to the probabilistic features of

Adults. Exceptions were classified more successfully when they had the deterministic feature of Regular items, than when they had type congruent deterministic feature.

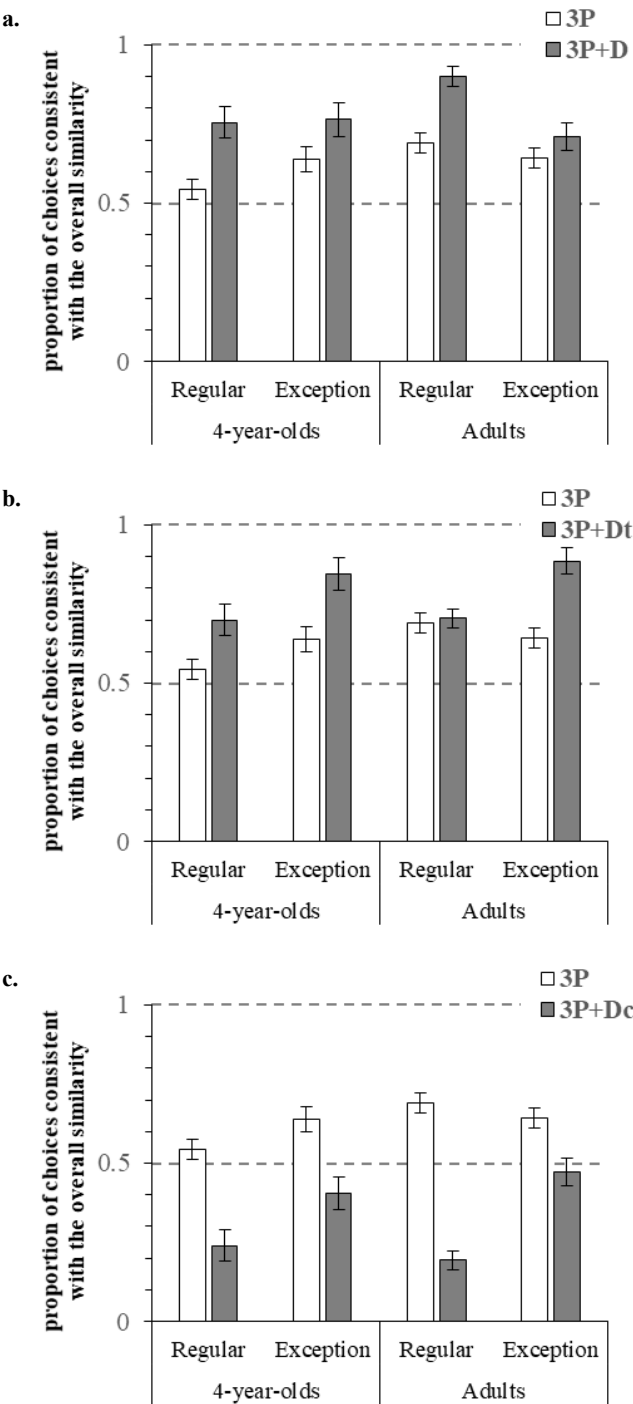


Figure 2: Proportion of choices consistent with the overall similarity as a function of Item Type and Age. Error bars represent standard errors of mean.

SUSTAIN simulation analyses

In order to get a better understanding of the mechanisms that may drive asymmetry effects in adults, we performed preliminary modeling analyses using SUSTAIN (Love, Medin, & Gureckis, 2004). After determining sets of parameters that predicted the pattern of adult’s categorization responses, we looked more closely at the representations formed by SUSTAIN. For each simulation, the model was given training analogous to that of human participants. We then examined the particular clustering solution formed by the model to represent the categories.

One particularly surprising finding was that SUSTAIN frequently formed a three-cluster representation, wherein for one category Regulars and Exceptions were represented separately, but they were represented jointly by a single hybrid cluster for the other category. Additionally, this three-cluster solution can perform surprisingly very well at categorization for the training items. A slightly more complex, but similar, five-cluster solution was also common. Together these two types of representations comprised the majority of the simulations. The frequency of these representations seemed to be the main factor driving the asymmetry effect in the model simulations, as they produced a large asymmetry in the contributions of Regulars’ and Exceptions’ deterministic features, while other cluster solutions produced much smaller asymmetries. This appears to be because one category’s Exceptions are subsumed in the same representation as its Regulars, and the hybrid cluster is more closely associated with the Regulars’ deterministic feature due to Regulars being more frequent. In other words, the Regulars’ features overshadow the Exceptions’ for the category with the joint representation.

Importantly this type of representation structure also makes the prediction that categorization accuracy should be unbalanced—for example, that accuracy should be better for members of Category A than for Category B (or vice versa). To further assess whether this three-cluster type of representation accurately characterizes the human data, we investigated the extent to which participants’ responses were unbalanced in the predicted way. We classified individual participants as balanced or unbalanced based on their categorization responses, and found that a large proportion of adults were unbalanced. We then examined the pattern of classification in balanced and unbalanced participants, and found that the asymmetry effect was much larger in unbalanced participants. So, the human data appear to confirm the predictions of SUSTAIN, that unbalanced representations between the two categories drives the asymmetry effect. The modeling results are interesting in that, contrary to the notion that exceptions have a special status in memory, they suggest Regulars and Exceptions may sometimes be represented together, leading to unexpected asymmetries in how individual features of Regulars and Exceptions contribute to categorization responses.

General discussion

In early childhood, category membership may be comparably affected by regular category members and exceptions. In our study, 4-year-old children relied on both features of regulars and features of exceptions when categorizing regular-like and exception-like items. Not only that both types of features were used when making category decisions, but deterministic features of regular items were equally useful for deciding on category membership of items that had probabilistic features of regulars and for items that had probabilistic features of exceptions. The same was true for the contribution of features of exceptions.

On the other hand, regular and exception items may have a different status in adulthood. Adult patterns showed asymmetry in contributions that features of regular and exception items have in categorization. Deterministic features of regular items made a large contribution to both categorizing regular-like and exception-like items, while features of exceptions had small contribution for exception-like items and did not affect categorization of regular-like items.

This pattern is in accordance with our predictions based on developmental differences in category learning (Sloutsky, 2010; Deng & Sloutsky, 2015, 2016). Based on previous findings (Deng & Sloutsky, 2015, 2016), adults, but not young children, are expected to optimize their attention towards the most relevant category features, which may result in asymmetry in contribution of deterministic features of regular and exception items.

The preliminary modeling analyses using SUSTAIN (Love, Medin, & Gureckis, 2004) gave us further insights into potential factors that may be driving the asymmetry effects in adults. The model predicted asymmetry effect by forming representations wherein regulars and the exception of one category were represented together, by one hybrid cluster, while for the other category they were represented by separate clusters. Since regular items were more frequent, submersing the exception into the same representation with the regular items made the deterministic feature of regular items more closely associated with the category. Importantly, human data appear to confirm the predictions of SUSTAIN. We found that a significant proportion of adult participants formed some kind of unbalanced representations – e.g. performing significantly better for one category than another. The finding that regular and exception items may sometimes be represented together is novel and it contrasts the dominant view of a need for a special status of category members that violate category expectations.

Acknowledgments

This research was supported by NIH grants R01HD078545 and P01HD080679 to VMS.

References

Davis, T., Love, B.C., & Preston, A.R. (2012). Learning the Exception to the Rule: Model-Based fMRI Reveals

Specialized Representations for Surprising Category Members. *Cerebral Cortex*, 22, 260-273.

Deng, W., & Sloutsky, V. M. (2015). The development of categorization: Effects of classification and inference training on category representation. *Developmental Psychology*, 51, 392–405.

Deng, W., & Sloutsky, V. M. (2016). Selective attention, diffused attention, and the development of categorization. *Cognitive Psychology*, 91, 24-62.

Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A network model of human category learning. *Psychological Review*, 111, 309–332.

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.

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

Sloutsky, V. M. (2010). From perceptual categories to concepts: What develops? *Cognitive Science*, 34, 1244-1286.