How does the presence of a label affect attention to other features?

Amy Perfors (amy.perfors@adelaide.edu.au)
School of Psychology, University of Adelaide

Daniel J. Navarro (daniel.navarro@adelaide.edu.au)
School of Psychology, University of Adelaide

Abstract

Are labels cues to category membership or simply highly salient features? This question is difficult to answer definitively because of the challenge in identifying empirical predictions that would be distinct in each case: either way, one would expect labels to be highly interesting, easy to process, and preferentially used as the basis of generalization. Here we suggest that one difference should be in how the label directs (or fails to direct) attention to the other, less-salient features of the object. We perform a categorization experiment with complex objects containing many low-salience features, and find that labels affect attention to the other features in the same way that highly salient features such as color or sounds do (and unlike an explicit cue to category membership). This results in a diminished ability to use the less-salient features of the categories to generalize appropriately.
Keywords: labels; features; categorization; generalization

Introduction

Shakespeare once famously asked "What's in a name?" Over the past few decades, psychologists have studied the scientific version of this question: what is the role of labels in category learning? How do labels affect categorization: the categories people form, the inductions they license, and the generalizations they make? What assumptions about labels do people bring to the tasks of word and category learning? These questions have been of special interest in the study of language acquisition, because understanding the assumptions children bring to the problem of word learning is key to understanding their linguistic development.

Much evidence demonstrates that children assume that labels are special in some way. Infants familiarized to items from a novel category will treat it like a category if they hear a label attached to the items, but not if they hear a non-linguistic sound (Balaban & Waxman, 1997; Fulkerson & Waxman, 2007) or hear nothing at all (Waxman & Markow, 1995; Waxman & Braun, 2005). Moreover, infants use labels but not sounds for individuation (Xu, 2002) and as a basis for inductive inference (Gelman & Markman, 1987; Davidson & Gelman, 1990; Graham, Kilbreath, & Welder, 2004).

Why do labels have this special status? Although infants appear uniquely interested in speech (Vouloumanos & Werker, 2004), they are equally capable of learning mappings involving non-linguistic sounds as words (Roberts & Jacob, 1991; Woodward & Hoyne, 1999). This suggests that the "specialness" of labels is not solely due to increased attention or interest in speech in general (although it may be related to the fact that the input is auditory; see Robinson & Sloutsky, 2004, 2006). Furthermore, when labels are inconsistent with apparent category structure or similarity, infants and children are much more reluctant to form categories based on them (Davidson & Gelman, 1990; Waxman & Braun, 2005; Plunkett, Hu, & Cohen, 2008); this may suggest that words are important because they tend to pick out useful categories. Perhaps children make the assumption that labels map cleanly onto category structure because labels are referential: younger infants will categorize using symbolic forms other than words (e.g., gestures or pictograms) if they are used in a referential context (Namy, 2001; Campbell & Namy, 2003), and older infants will use labels to pick out global categories only if they are presented in person by an experimenter rather than a recording(Fulkerson & Haaf, 2003). Another possibility is that infants assume that words identify useful categories because they statistically tend to do so (Samuelson & Smith, 1999), and infants' statistical learning mechanisms are well-attuned for picking this sort of pattern up (Smith, Jones, & Landau, 1996).

As this discussion illustrates, there is some disagreement about how and why labels are special. It may be that labels are special because they are linguistic – referential and used for communication – and infants realize this (Balaban & Waxman, 1997; Namy, 2001; Xu, 2002; Fulkerson & Waxman, 2007). Alternatively, it may be that infants have learned to pay special attention to words because they are statistically likely to be useful indicators of category structure (Smith et al., 1996). The special status of labels may also be perceptual in origin: perhaps labels play a unique role in category formation because of their auditory properties (Robinson & Sloutsky, 2004, 2006, 2007).

This debate parallels a similar, but not identical, discussion in the adult literature – one focused on whether labels act as category indicators or just a highly salient feature. On one hand, labels certainly do appear to hold a privileged psychological status in some ways. When objects share a label, this is sufficient to increase their similarity (e.g., Goldstone, Lippa, & Shiffrin, 2001), and people often make inductive inferences based on an object's label rather than its features or overall similarity(e.g., Yamauchi & Markman, 2000; Johansen & Kruschke, 2005). On the other hand, formal models of categorization have often been remarkably successful at matching human performance simply by treating labels as another – possibly highly salient – feature of the stimulus (e.g., Anderson, 1991; Gliozzi, Mayor, Hu, & Plunkett, 2009).

One of the difficulties inherent in resolving this debate is that it is hard to identify characteristics that an indicator of category membership would have but a very salient feature would not. For instance, one might suggest that the differ-

ence might be that if something is an indicator of category membership, it should be used to pick out categories even when it seems to be inconsistent with the observed similarity or category structure. There is evidence that this is the case for labels when they are mildly inconsistent (Yamauchi & Markman, 2000), but not when they grow too inconsistent (Davidson & Gelman, 1990; Waxman & Braun, 2005). But does this mean that words are strong markers of category membership or salient object features? The problem is that the results make sense under either theory. On one hand, if labels are especially salient features then one would expect them to be followed even if other (less salient) features seemed to pick out a different category structure; on the other hand, if labels are treated as markers to category membership without being features themselves, they could still be such strong markers that they are nearly impossible to override.

More generally, both highly salient features and cues to category membership should share many other characteristics: easy to represent, quick to process, and preferentially used as a basis for generalization. What, then, is the difference between them? To address this question, it helps to consider the two possibilities individually.

- *What are the cognitive effects of a salient feature?* Much work suggests that salient features share two important characteristics. One is that they tend to be the features that people examine first when making choices (e.g., Tversky, 1972; Gigerenzer & Goldstein, 1996). The other is that if the feature is predictive and useful, it will become even more salient over the course of learning (Kruschke, 1992, 2003). As a consequence, if a feature is initially quite salient and later turns out to be predictive of category membership, even more attention will be devoted to it, and the attention devoted to the *other* features will decrease commensurately, particularly if they themselves are not salient or are difficult to process.

- *What are the cognitive effects of a cue to category membership?* Less research bears directly on this question, but we can begin by considering the case of something that is unequivocally a cue to category membership and also unequivocally *not* a feature: explicit instruction. Imagine telling someone that objects from category *A* were sorted into one box and objects from category *B* were sorted into another. Those boxes (along with the instructions) would be cues to category membership, but not features of the objects. How would this affect processing of the objects? Not surprisingly, providing this kind of structure in the visual presentation of stimuli tends to improve learning by calling attention to the relevant features and minimizing the processing load imposed on the learner (e.g., Bruner, Goodnow, & Austin, 1956, ch. 4). As a result the effect on attention is expected to be in the opposite direction: those object features that are less salient, will be processed much more than they otherwise might.

Do the cognitive effects of labelling look more like those of features, or of cues to category membership? We address this question by presenting participants with a simple categorization task involving objects with numerous non-salient and difficult-to-process features paired with a *category indicator* of some sort. In one condition, the category indicator is intended to be a strong cue to category membership: the objects are explicitly categorized by being sorted into boxes. In two other conditions, the category indicator is a highly salient non-linguistic feature (a color or a non-linguistic sound). In two final conditions, the category indicator is a label (either written or oral). After sorting the objects, participants are asked how they would classify new objects for which the category indicator is unknown. Importantly, because the category indicator is unknown and the other features so complex and low-salience, performance on the generalization task reflects how much people have attended to those features. If the category indicator acts like a cue to category membership by calling attention *to* the less-salient features, generalization on the basis of them should be improved when given the indicator; however, if the category indicator acts more like a salient feature by directing attention *away from* the less-salient features, generalization should be poor. Our results suggest that labels behave much like other extremely salient features in the way that they focus attention away from other features of an object.

Method

92 adult participants were recruited from the University of Adelaide and surrounding community and were paid $5 for their participation in the half-hour experiment. Two participants were excluded due to failure to understand the task, leaving 18 people in each of five possible conditions. Each participant saw a series of trials in which they were asked to sort novel objects into categories. They were then asked two generalization questions about how they would categorize additional objects without category indicators. Each of the objects has eight features, four of which vary coherently according to the category structure, and four of which are random. In half of the trials (the NO INDICATOR trials), participants were asked to sort these objects into clusters. In the other half (the INDICATOR trials) the task was the same except that the objects were also each associated with a category indicator, the nature of which varied by condition.

Items. Each item consists of a square with four symbolic characters (one in each quadrant) surrounded by circles (also containing symbolic characters) at each corner; we refer to each location as one of the eight low-salience *features* of the objects, and the particular character in that location as its *feature value*. Each feature can take on a value corresponding to one of ten specific characters, and there is no overlap of possible character sets (feature values) from feature to feature. For each participant and trial, features were generated independently, according to the following pattern: four features are randomly selected to be *dispersed*, meaning that they do not respect category structure because they are uniformly se-
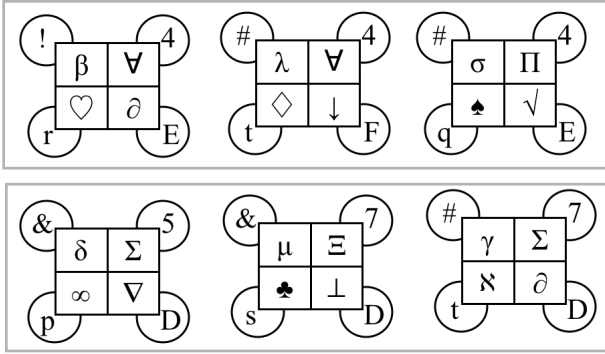
Figure 1: Example INDICATOR trial in the BOXES condition (for visual clarity, we show 6 objects rather than 8 or 16). In this condition, objects are presented already sorted into boxes corresponding to two categories. Here the four coherent features are the two upper circles, the upper right square, and the lower right circle. These features have a 75% coherence level: each of the four coherent features has 25% probability of being "flipped" from the value appropriate to its category.
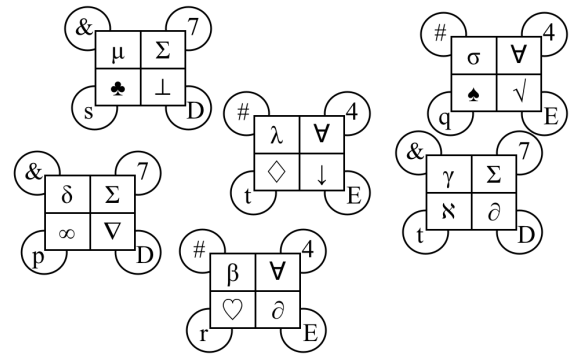

Figure 2: Example NO INDICATOR trial, which participants in all conditions were exposed to. In this sort of trial, people are told to sort the objects in whatever way appears sensible, and are not told in advance how many categories there are or what features are important or useful. In this trial the coherence level is 100%: each of the four coherent features (which are the same as in Figure 1) follow the category structure precisely.

lected from the possible set of values for that feature. The other four are *coherent*, meaning that they correspond to the underlying category structure: feature A corresponds to category structure if all members of category X share a the same feature value for A (say, all of them have a δ in the upper left corner of the square). We systematically varied the coherence[1] level of the four *coherent* features so that half of the trials involved items with a coherence level of 75%, and half involved a coherence of 100%. This mimics real-world categories, which have a probabilistic, graded structure.[2] It is possible to identify the correct categories on the basis of the coherent features, as people have succeeded in doing in other studies (Perfors & Tenenbaum, 2009). However, because these features are numerous, of low salience, and representationally complex, it can be difficult.

Sample objects as they appeared in the experiment are shown in Figures 1 and 2.

Trial structure. Each participant saw eight NO INDICATOR and eight INDICATOR trials. In order to ensure that participants were not relying on external knowledge about how many categories the correct sorting contained, trials varied in the number of items (8 or 16) and the number of categories (2 or 4). Since items varied also in coherence, this resulted in the following factorial design: 2 (INDICATOR or NO INDICATOR) x 2 (coherence level of 75% or 100%) x 2 (containing 8 or 16 items total) x 2 (categories made of 2 or 4 items). This resulted in 16 trials per participant. Due to a coding error, trials with 8 items and 2 categories were not properly counterbalanced according to category indicator, so all analyses excluded these trials and therefore consisted of 12 trials per participant. Figure 1 shows the sort of situation a partic-

---

[1] A coherence of c means that a feature value has a $(100 - c)\%$ chance of being randomly generated rather than following category structure.

[2] There were no interaction effects between coherence and any of the results of interest here, so all analyses combine coherence levels.

ipant might see on an INDICATOR trial in the BOXES condition, while Figure 2 shows a typical NO INDICATOR trial.

Conditions. The five conditions are defined by the nature of the category indicator involved in the INDICATOR trials. In the BOXES condition, participants saw the objects already pre-sorted into boxes; this is intended as an explicit cue to category membership, and was described to participants as such. In two of the other conditions, the objects in the INDICATOR trials were associated with a label. In the WRITTEN LABEL condition, participants were told that the label would be written above the object. To evaluate whether it mattered if the label was presented visually or orally, in the ORAL LABEL condition, the label was presented out loud (over headphones) whenever the participant clicked on the object. Since participants had to click on all objects in order to sort them, they ended up hearing the labels for every object at some point. The label conditions were compared to two conditions in which the category indicator was simply a highly salient feature. In the COLOR condition, objects were colored (unlike the objects in the NO INDICATOR trials and other conditions, which were always white). And in the SOUND condition, objects were associated with non-linguistic sounds (distinct buzzes, beeps, and tone sequences without semantic associations). As in the ORAL LABEL condition, these sounds were heard through headphones whenever the participant clicked on the object.

Procedure. Each trial consisted of two phases. The first was the "sorting" phase, in which participants were presented with all of the objects in the trial randomly scattered on the computer screen and asked to sort them in categories. (The exception is the INDICATOR trials in the BOXES condition, in which the objects appeared already sorted with square "boxes" drawn around each of the categories, as depicted in Figure 1). During the sorting phase, participants were allowed unlimited time in which to move the objects around on the screen by clicking and dragging them into clusters. They then drew boxes around the objects to indicate cate-
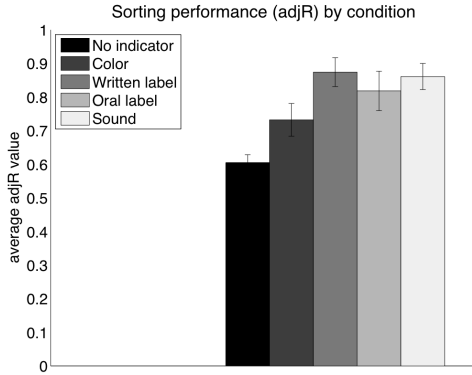
Figure 3: Performance in the sorting task. Subjects in the ORAL LABEL, WRITTEN LABEL, and SOUND conditions used the category indicators to sort at close to optimal levels. When there was no category indicator, people were able to use the less-salient features to sort, but were significantly worse than when there was one.

gories. People were told ahead of time that not all trials would have the same number of items or categories, and they should just sort in whatever way seemed sensible.

After the sorting task was completed, the items remained on the screen and participants were presented with two generalization questions in random order. In *first-order* generalization, participants were shown one of the items they had sorted (without category indicator) and asked which of two novel items would go in the same category as that one. The correct answer had the four coherent features in common with the first, and the incorrect answer had the four other features in common. The *second-order* generalization trials were identical, except that the item shown to the participants had specific feature values that had not been seen before: a person could only answer correctly if they realized that the coherent features (rather than specific values) were what mattered for category organization. As in Perfors and Tenenbaum (2009), our participants performed identically in the first- and second-order generalization, so all analyses collapse them together into one variable, *gen*.

Results

There are two natural questions to ask. First, does the nature of the category indicator affect people's sorting behavior? Second, does it affect how people pay attention to the other, less-salient features of the objects? We can address the second question by examining generalization performance in each condition, since our generalization tasks do not include the category indicator and therefore necessarily rely on the other features. The answer to the first is important for knowing how to interpret the answer to the second: for instance, if generalization performance is poorer because people cannot figure out the correct categories, that does not tell us anything about how people are attending to the less-salient features *given* those categories. We therefore begin with addressing how sorting performance depends on the nature of the category indicator.

Sorting performance

Sorting performance is evaluated using a standard measure for evaluating the similarity between two clusterings of items known as the adjusted Rand Index (*adjR*) of Hubert and Arabie (1985). In this case, we use *adjR* to measure the similarity between the correct category clustering and the category assignments made by the participants. An *adjR* of 1 indicates that the clusters are identical, while 0 is the score one would expect from two random clusterings; scores below 0 indicate that the clusters match less than one would expect by chance.

Figure 3 indicates that category indicator has a strong effect on sorting performance.[3] Participants in the ORAL LABEL, WRITTEN LABEL, and SOUND conditions sorted nearly optimally, which suggests that they used the category indicators to create their categories (since sorting according to category indicator is optimal sorting). Participants on the NO INDICATOR trials were able to use the less-salient features to sort at an above-chance level, but performed worse than when given a category indicator. Finally, people in the COLOR condition sorted halfway in-between, suggesting that color was a more salient feature than the symbolic characters, but not as salient as labels or sounds.

Generalization

Based on sorting performance it appears that participants generally created sensible categories. Were they able to form generalizations about category membership based on the less-salient features? We test this, as explained earlier, by presenting participants with additional items and asking how they would categorize a novel item they had not seen before. Figure 4 demonstrates that generalization in the BOXES condition was generally superior to generalization in the other conditions, all of which were similar to each other.[4] Since generalization depends on what the participant notices about the less-salient features *other* than the category indicator, this suggests that in the BOXES condition people were paying more attention to those features than in any of the other conditions.

These two results, taken together, drive the main conclusion of this paper: labels appear to act more like highly salient features than overt category indicators (boxes). Labels, like highly salient features, support accurate sorting, but are associated with poorer levels of generalization to new items. We have suggested that the reason for this may be because the labels and salient features are directing attention away from the non-salient features during the sorting task; this impairs generalization because attention to the non-salient features is

---

[3]A one-way Anova on *adjR* by condition was significant: $F(4, 158) = 9.77, p = 4.34e^{-7}$. Post-hoc comparisons using the Tukey-Kramer test indicated that the mean *adjR* in the NO INDICATOR condition was significantly different than mean *adjR* in the ORAL LABEL, WRITTEN LABEL, and SOUND conditions.

[4]A one-way Anova on generalization by condition was significant: $F(5, 176) = 2.91, p = 0.0149$. Post-hoc comparisons using the Tukey-Kramer test indicated that the generalization in the BOXES condition was significantly different from the NO INDICATOR and ORAL LABEL conditions, and nearly significantly different from the other three.
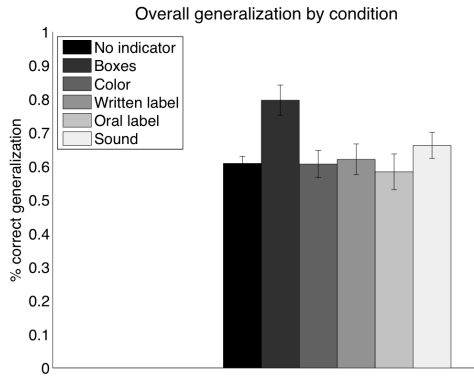
Overall generalization by condition

Figure 4: Generalization on the basis of the non-salient features in the BOXES condition was superior to generalization in the other conditions, suggesting that participants in the other conditions did not attend as much to the less-salient features when generalizing.

necessary for accurate generalization of novel items (which are not associated with a label or highly salient feature). This would explain why generalization in those conditions is lower than generalization in the BOXES condition.

However, one minor yet confusing aspect of these results remains: if the salient features are truly directing attention away from the non-salient features, why is generalization performance not poorer on the INDICATOR trials than the NO INDICATOR trials, at least in all conditions other than BOXES? After all, it might be assumed that people are *less* able to use the non-salient features when they have the distracting, highly-salient features around, especially since those features do a very good job at picking out the category members.

Relating sorting and generalization

We address this question by realizing that two factors drive generalization performance, which depends ultimately on knowing which of the less-salient features pick out which categories. It therefore requires not only being able to attend to and identify the less-salient features, but also knowing what the correct categories are. On the INDICATOR trials in the WRITTEN LABEL, ORAL LABEL, SOUND, and COLOR conditions, participants may be less able to attend to the non-salient features, but be better at identifying the categories in the first place. These factors may therefore be cancelling each other out, resulting in generalization that is very similar to the NO INDICATOR conditions.

This possibility yields a testable prediction, namely that in the NO INDICATOR trials sorting performance should be positively correlated with generalization, but in the INDICATOR trials it should be more irrelevant.[5] We would not expect it to be *entirely* irrelevant since, after all, one must be able to identify the categories in order to generalize correctly. However, the converse is not necessarily true: identifying the categories in the INDICATOR conditions would not imply that

one should be able to generalize correctly, since generalization requires attention to the less-salient features but categorization does not. We test this by calculating the correlation between sorting accuracy (*adjR*) and generalization (*gen*) for both the INDICATOR and NO INDICATOR trials. Although both are significant, the size of the effect on the INDICATOR trials is markedly weaker.[6] While not conclusive, this is consistent with our interpretation of the results: sorting is less predictive of generalization in the INDICATOR trials because sorting does not depend on the less-salient features in those trials, unlike in the NO INDICATOR situation.

Discussion

This research is motivated by the question of whether labels are cues to category membership or simply highly salient features. The question is difficult to answer in part because it is hard to predict what would be empirically different in each case: no matter what, one would expect labels to be highly interesting, easy to process, and preferentially used as the basis of generalization (but also to be ignorable if they were inconsistent with category structure). We suggest that one difference between cues to category membership and highly salient features is their effect on the processing of the other, less salient features of the objects: highly salient features should direct attention away from the less salient ones, while cues to category membership should direct attention toward them. We tested this by presenting participants with a sorting task involving objects with many complex, low-salience features, and then posing generalization questions that required attention to the less-salient features to answer correctly. Our main results, shown in Figures 3 and 4, suggest that labels act more like highly salient features than they act like boxes (an explicit external cue to category membership).

One might object that this result is not very surprising. After all, stimuli in the BOXES condition may be easier to process since they have one fewer feature – the cue to category membership is the box and the visual organization of the objects, not any features inherent to them. However, in a very real sense this is precisely our point: if something is acting as a cue to category membership, it *should* improve performance by reducing the load required to process the actual features of the objects. Labels, whether oral or written, did not do that in our study.

An important subtlety lies in how we define salience. In what way are the labels in our study really "highly salient"? All of them except for the written label are perceptually noticeable; is this what we mean? The written label was actually fairly small relative to the size of the entire object, so why do people treat it as highly salient? In answer, we note the importance of distinguishing *perceptual* salience from what we might call *conceptual* salience. A feature is perceptually salient because our basic perceptual mechanisms automatically notice and process it preferentially or more easily; this

---

[5]Note that when we refer to sorting in the INDICATOR trials, we are excluding the BOXES condition, since participants do not actually have to sort anything – the items are already placed into boxes. All of these analysis, therefore, excluded the BOXES condition from the INDICATOR trials.

[6]Spearman's: INDICATOR: $\rho = 0.192, p = 0.007$; NO INDICATOR: $\rho = 0.499, p < 0.0001$.

might be true of speech input (Vouloumanos & Werker, 2004) or auditory input in early childhood (Robinson & Sloutsky, 2004). By contrast, a feature may be *conceptually* salient if we have learned to attend to it preferentially for more abstract conceptual reasons – perhaps because it has proven useful in the past, or if it is easier to process because we have practiced processing it for many years. If the written labels are highly salient, this is probably the sense in which they are. The distinction between the two types of salience gets somewhat blurry at the edges, since many features may be both perceptually and cognitively salient, or change in salience over time. The important point, however, is that for our purposes something is salient if it invites preferential attention or is easier to process; that may be because of perceptual factors, learned conceptual factors, or some mixture of both, and we do not address that question in this work.

One limitation of our study is the fact that it was presented entirely on a computer using bizarre objects with many representationally complex features. The complexity of the features was intentional since we wanted to maximize our chances of creating a situation in which low attention to the features had a measurable effect on generalization; however, it is possible that, due to the unnaturalness of the situation, people adopted a strategy unlike that which they use in the real world. It is also possible that labels, since they are normally referential and communicative, might have a different effect when presented in a communicative, social context rather than on a computer. There is evidence that for children, labelling by a person results in different behavior to labelling by a recorder (Fulkerson & Haaf, 2003), and that non-labels can behave more like labels when presented in a referential context (Campbell & Namy, 2003). However, it is unclear how (or if) these findings will generalize to children, to more naturalistic stimuli, or to different contexts; future work is necessary.

Acknowledgments

References

Anderson, J. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*(3), 409–429.

Balaban, M., & Waxman, S. (1997). Do words facilitate object categorization in 9-month-old infants? *Journal of Experimental Child Psychology*, *64*, 3–26.

Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1956). *A study of thinking*. New York, NY: Wiley.

Campbell, A., & Namy, L. (2003). The role of social-referential context in verbal and non-verbal symbol learning. *Child Development*, *74*(2), 549–563.

Davidson, N., & Gelman, S. (1990). Inductions from novel categories: The role of language and conceptual structure. *Cognitive Development*, *5*, 151–176.

Fulkerson, A., & Haaf, R. (2003). The influence of labels, non-labeling sounds, and source of auditory input on 9- and 15-month-olds' object categorization. *Infancy*, *4*(3), 349–369.

Fulkerson, A., & Waxman, S. (2007). Words (but not tones) facil-itate object categorization: Evidence from 6- and 12-month-olds. *Cognition*, *105*(1), 218–228.

Gelman, S., & Markman, E. (1987). Young children's inductions from natural kinds: The role of categories and appearances. *Child Development*, *58*, 1532–1541.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650-669.

Gliozzi, V., Mayor, J., Hu, J.-F., & Plunkett, K. (2009). Labels as features (not names) for infant categorization: A neurocomputa-tional approach. *Cognitive Science*, *33*, 709–738.

Goldstone, R., Lippa, Y., & Shiffrin, R. (2001). Altering object representations through category learning. *Cognition*, *78*, 27–43.

Graham, S., Kilbreath, C., & Welder, A. (2004). Thirteen-month-olds rely on shared labels and shape similarity for inductive infer-ences. *Child Development*, *75*(2), 409–427.

Hubert, L., & Arabie, P. (1985). Comparing partitions. *Journal of Classification*, *2*, 193–218.

Johansen, M., & Kruschke, J. (2005). Category representation for classification and feature inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(6), 1433–1458.

Kruschke, J. K. (1992). Alcove: An exemplar-based connectionist model of category learning. *Psychological Review*, *99*(1), 22-44.

Kruschke, J. K. (2003). Attention in learning. *Current Directions in Psychological Science*, *12*, 171-175.

Namy, L. (2001). What's in a name when it isn't a word? 17-month-olds' mapping of nonverbal symbols to object categories. *Infancy*, *2*(1), 73–86.

Perfors, A., & Tenenbaum, J. (2009). Learning to learn categories. In *31st Annual Conference of the Cognitive Science Society*.

Plunkett, K., Hu, J.-F., & Cohen, L. (2008). Labels can override perceptual categories in early infancy. *Cognition*, *106*, 665–681.

Roberts, K., & Jacob, M. (1991). Linguistic versus attentional in-fluences on nonlinguistic categorization. *Cognitive Development*, *6*, 355–375.

Robinson, C., & Sloutsky, V. (2004). Auditory dominance and its change in the course of development. *Child Development*, *75*(5), 1387–1401.

Robinson, C., & Sloutsky, V. (2006). Auditory overshadowing and categorization: When decreased visual processing facilitates cat-egorization. In *28th Annual Conference of the Cognitive Science Society*.

Robinson, C., & Sloutsky, V. (2007). Linguistic labels and catego-rization in infancy: Do labels facilitate or hinder? *Infancy*, *11*(3), 233–253.

Samuelson, L., & Smith, L. (1999). Early noun vocabularies: Do ontology, category structure, and syntax correspond? *Cognition*, *73*, 1–33.

Smith, L., Jones, S., & Landau, B. (1996). Naming in young chil-dren: A dumb attentional mechanism? *Cognition*, *60*, 143–171.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, *79*, 281-299.

Vouloumanos, A., & Werker, J. (2004). Tuned to the signal: The privileged status of speech for young infants. *Developmental Science*, *7*(3), 270–276.

Waxman, S., & Braun, I. (2005). Consistent (but not variable) names as invitations to form object categories: New evidence from 12-month-old infants. *Cognition*, *95*, B59–B68.

Waxman, S., & Markow, D. (1995). Words as invitations to form categories: Evidence from 12- to 13-month-old infants. *Cognitive Psychology*, *29*, 257–302.

Woodward, A., & Hoyne, K. (1999). Infants' learning about words and sounds in relation to objects. *Child Development*, *70*(1), 65–77.

Xu, F. (2002). The role of language in acquiring object kind con-cepts in infancy. *Cognition*, *85*, 223–250.

Yamauchi, T., & Markman, A. (2000). Inference using categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*(3), 776–795.