The role of similarity in the development of categorization

Vladimir M. Sloutsky

Center for Cognitive Science, 208C Ohio Stadium East 1961, Tuttle Park Place, The Ohio State University, Columbus, OH 43210, USA

Early in development, humans exhibit the ability to form categories and overlook differences for the sake of generality. This ability poses several important questions: How does categorization arise? What processes underlie category formation? And how are categories mentally represented? We argue that the development of categorization is grounded in perceptual and attentional mechanisms capable of detecting multiple correspondences or similarities in the environment. We present evidence that: (a) similarity can drive categorization early in development; and (b) early in development, humans have powerful learning mechanisms that enable them to extract regularities in the environment. We conclude that, despite remaining challenges, the similarity-based approach offers a promising account of the development of categorization.

In the story *Funes, the Memorious*, Jorge Luis Borges described Funes – an individual with an amazing memory but no ability, or need, for abstraction. In particular, Funes could not understand that each entity, such as any specific number, is an instance of a more general category. For Funes, each number existed as a distinct, integral and primitive whole that could be labeled by a proper name. The rest of us, fortunately, understand that individuals (e.g. my dog) are members of more general categories (i.e. dogs), and this understanding is a crucially important component of cognition.

Categories are equivalence classes of different (i.e. discriminable) entities and categorization is the ability to form such categories and treat discriminable entities as members of an equivalence class. This ability to overlook differences for the sake of generality is important for several reasons. First, it is more resource efficient to incorporate a potentially infinite number of individual entities into a substantially smaller number of classes. Second, this ability supports organization of knowledge, such as enabling the creation of taxonomies by including smaller classes into larger ones (e.g. German shepherds → dogs → animals). Third, the ability to form categories supports induction, because members of the same class are assumed to share many unobserved properties (e.g. upon learning that that German shepherds use acid-based enzymes for digestion, one might generalize this knowledge to other dogs, or to other mammals). There are bi-directional relationships between categorization and knowledge: (a) knowing that an animal shares properties with members of the category 'cats' enables us to categorize this animal as a cat; and (b) knowing that an animal is a cat enables us to predict that the animal has properties of cats. The first ability could be defined as categorization and the second as inductive projection. We will refer to both abilities as inductive generalization.

While the importance of inductive generalization is widely accepted, several critical issues remain highly contested. How do people form categories and how does this ability arise? What processes underlie categorization? How are categories mentally represented, and how do these representations change in the course of cognitive development?

Theoretical approaches to categorization

It has been long believed that representations of categories develop from immature representations that are based on accidental features (or appearance similarities) to mature representations that are based on logical rules [1,2]. Under this view, mature representations include features that are individually necessary and jointly sufficient for a category membership [1–4]. For example, the concept 'prime number' includes two features: (a) an integer; and (b) divisible without remainder by one and by itself. Each feature is necessary and they are jointly sufficient to determine whether a number is a prime. According to this 'classical' view, conceptual development was viewed as a process of discovery of the necessary and sufficient features. However, in the 1980s these ideas came under severe attack: whereas some researchers have argued that many categories do not in principle conform to the 'classical' view because they simply do not have necessary and sufficient identifiable features, others have argued that children's representation of categories is more sophisticated than had been envisioned by the 'classical' view [5].

With the demise of the 'classical' view, an influential theoretical approach to conceptual development has emerged. This approach has been often referred to as the 'naïve-theory' or 'knowledge-based' position. The fundamental argument of the 'naïve theory' position is that mature categorization cannot develop from simpler perceptual components: even very early in development, people must have conceptual knowledge (i.e. naïve theories about the world) that would constrain which perceptual correspondences would be detected and which would be ignored (see [5,6] for reviews). There are several

Corresponding author: Vladimir M. Sloutsky (Sloutsky.1@osu.edu).

http://tics.trends.com 1364-6613/03/$ - see front matter doi:10.1016/S1364-6613(03)00109-8

more specific assumptions underlying this position. First, although there might be no truly 'essential' features defining a category, even young children believe that these features exist. Second, features differ in their conceptual centrality, which is determined by a causal status of the feature in question; even young children believe that essential features are the most central ones. Third, even for young children, perceptual features (e.g. appearances of entities) play a peripheral role in categorization compared with more conceptually central features. And fourth, even infants and young children believe that individual entities are members of more general categories, and that linguistic labels, presented as count nouns, denote categories.

Although this position has been quite influential, it has several serious limitations. First, the origins of conceptual knowledge are not specified. If this knowledge is innate, as it is often claimed, it should have developed from something. And if one accepts the possibility that this knowledge could have developed from simpler components in the course of evolution, this possibility weakens the argument that this knowledge could not in principle develop from simpler components. Successful computer simulations also weaken somewhat the assumption of innate conceptual knowledge: simple neural networks that have powerful learning mechanisms and little prior knowledge acquire categories [7]. Second, explanatory constructs (i.e. naïve theories) are as complex as the phenomena they purport to explain. Finally, the 'naïve theory' position is based on a large number of assumptions, and there is a growing body of evidence that some of these assumptions are either unwarranted or unnecessary. These are sufficiently strong reasons for exploring theoretical alternatives.

One theoretical alternative that has been prominent in the study of mature categorization is that inductive generalization is grounded in perceptual and attentional mechanisms capable of detecting multiple correspondences or similarities (see [5,7–10] for reviews). This position has recently gained strength in the study of conceptual development [7,11,12]. This position (we will refer to it as the 'similarity-based approach') challenges most of the assumptions of the theory-based approach. Most importantly, the similarity-based approach does not assume the existence of conceptual knowledge at the beginning of development; rather, it attempts to explain the development of this knowledge from simpler processes operating on simpler components. The central argument is that there are multiple correlations in the environment and that humans have perceptual and attentional mechanisms capable of extracting these regularities and establishing correspondences among correlated structures (see [7] for related arguments). Despite the remaining challenges, the similarity-based approach offers a promising account of the development of categorization and induction.

Because of space limitations, it is impossible do justice to both positions and this review will focus on the similarity-based position. However, to reflect theoretical contentions, the diverging points of view will be acknowledged wherever possible. I first review evidence challenging some assumptions of the 'naïve theory' approach, and then discuss powerful learning mechanisms, grounded in perception and attention, that enable humans to develop categories.

Attention, similarity and inductive generalization in young children

A major tenet of the 'naïve theory' approach is that even young children believe that some features of entities are conceptually central and that others are peripheral. In particular, linguistic labels presented as count nouns were claimed to be central properties, whereas appearances were claimed to be peripheral properties [13]. To demonstrate this point, researchers presented young children with a 'triad task' (Fig. 1).

However, there is mounting evidence to challenge the idea that labels are central and appearances are peripheral in young children's categorization. In particular, there is evidence that reliance on linguistic labels is not fixed, and that it can vary as a function of perceptual information. For example, children's reliance on linguistic labels in categorization and induction tasks differ for real 3-dimensional (3-D) objects and for line-drawing pictures: effects of labels are more pronounced for line-drawing pictures (such as those used in the original studies [14]) than for real 3-D objects [15]. Furthermore, under some conditions, perceptual similarity could be more important than matching labels: 3- and 4-year-olds are more likely to rely on similarity of motion than on the matching linguistic label [16].

This evidence seriously undermines the centrality assumption because properties of 'peripheral' information

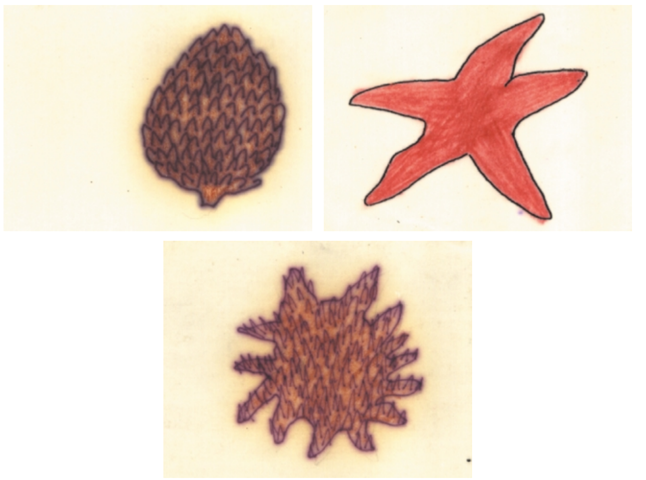

Fig. 1. An example of a triad task used in the experiments of Gelman and Markman [14], in which each triad included a target (the item below) and two test items. The target was introduced as a starfish. The upper-left test item was, according to the authors, perceptually similar to the target, and it was introduced as a pinecone. The upper-right test item was, according to the authors, perceptually dissimilar from the target, and it was introduced as a starfish. The child was told that the test starfish had an unobserved property (e.g. 'lives in water'), whereas the pinecone had another unobserved property (e.g. 'grows on trees'). When asked whether the target starfish lives in water or grows on trees, even 3-year-olds were more likely to rely on common label than on similar appearance [14]. It was concluded that (even for young children) linguistic labels relay conceptual information that is more central than perceptual similarity. (Picture courtesy of Susan A. Gelman).

should not affect the centrality of 'essential' features. It seems that labels gain their weight from an attentional mechanism (which presumes flexible attentional weights) rather than from the centrality assumption (which presumes fixed roles of central and peripheral features).

Another challenge to the idea of centrality of labels, as compared with perceptual similarity, comes from studies demonstrating that linguistic labels contribute to the overall similarity of compared entities: if two entities share a label, young children are more likely to say that these entities look alike [17]. Furthermore, this overall similarity – rather than linguistic labels alone – drives inductive generalizations [12]. Based on their model of similarity and the Luce's choice rule, these researchers predicted that the probability of inducing a property from the identically labeled test stimulus to the target varies as a function of ratio of similarities of each of the test stimulus to the target scaled by the weight of label match. To test these predictions, researchers systematically varied similarity of each of the test stimuli to the target by morphing pairs of animals into each other (Fig. 2). Young children's performance on induction tasks was well predicted by the estimated similarity ratios and weights of labels: more than 88% of variance for the induction task alone (Fig. 3) and more than 94% of the observed variance for induction, categorization and similarity judgment tasks was predicted by the model.

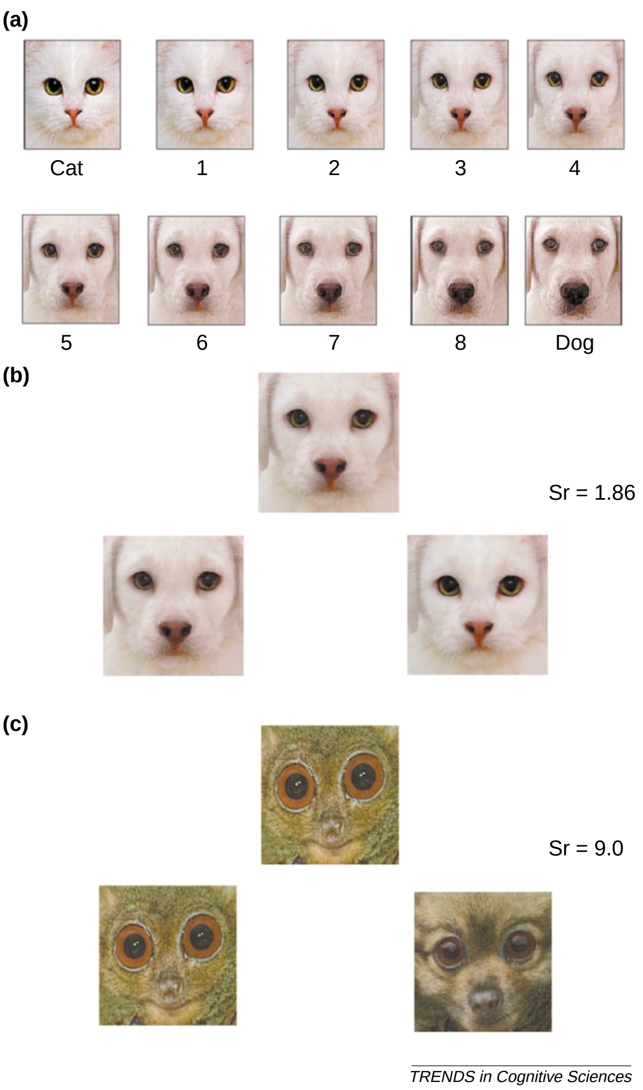

Fig. 2. To test predictions that inductive generalization is driven by the overall similarity, Sloutsky *et al.* [12] systematically varied the similarity of each of the test stimuli to the Target by morphing pairs of animals into each other (a). From these morphed sequences, they selected multiple triads with different similarity ratios. These similarity ratios (Sr) were estimated in separate experiments, in which participants were asked to select the test animal that looked more like the Target animal. In the triads illustrated here, where the Target animal is shown above two test pictures, (b) has a markedly lower Sr than (c).

Finally, there is evidence to suggest that the effects of labels on inductive generalization stem from attentional mechanisms rather than from the centrality assumption ([18], but see [19] for diverging arguments). In particular, infants [20,21] and young children [18] have attentional bias towards auditory stimuli (as compared with visual stimuli) and this bias disappears later in development. For example, when presented with a combination of a non-speech sound and visual arrangement, both infants and young children are more likely to attend to the auditory than to the visual stimuli, even though they have no difficulty attending to the visual stimuli when these are not accompanied by auditory stimuli (see Box 1 for details). In short, it seems that an attention-based mechanism of similarity computation can account for inductive generalization in young children.

The development of inductive generalization

The similarity-based account suggests that inductive generalizations are driven by similarity that is determined by automatically detected perceptual correspondences. However, the importance of the same correspondence might vary across contexts: for example, matching color is more important for categorizing something as a food item than it is for categorizing something as a car, whereas the reverse is true for shape. How do children know which correspondences are important, and for which categories?

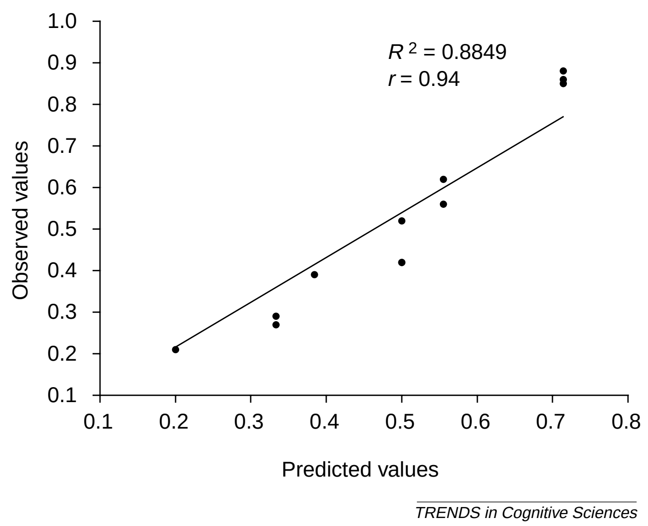

Fig. 3. Predicted and observed probabilities in induction experiments [12] in which 4- to 5-year-old children were presented with a triad task similar to that presented in Fig. 1. The values are predicted by a model of similarity [12,17] and observed values are also obtained. The results indicate that the model predicted more than 88% of the observed variance. (Data reproduced from Ref. [12] with permission from SRCD).

Box 1. Why do labels affect inductive generalization?

Two classes of explanation seem plausible: the prominence of labels could stem from language-specific factors, or it could stem from general auditory factors. One variant of the former explanation argues that young children assume that entities are members of categories and that labels presented as count nouns convey category membership [13]. Another variant of the former explanation argues that effects of labels might stem from special attention to human speech [32]. The general auditory explanation argues that the importance of labels might stem (at least partially) from this attentional bias, which was demonstrated in a series of studies [18] in which 4-year-olds and adults were presented with two stimuli sets, each consisting of a particular combination of visual and auditory patterns and were trained to select an arbitrarily designated 'target set' and to ignore the 'distracter set'. When training was complete, participants were tested. During testing the visual pattern from the target set was paired with a new auditory pattern, whereas the auditory pattern from the target set was paired with a new visual pattern (Table I). Results presented in Fig. I indicate that whereas adults almost invariably selected the visual match at testing, young children were more likely to select the auditory match. In control conditions, where either sound or picture was eliminated, both children and adults reliably selected the remaining trained stimulus.

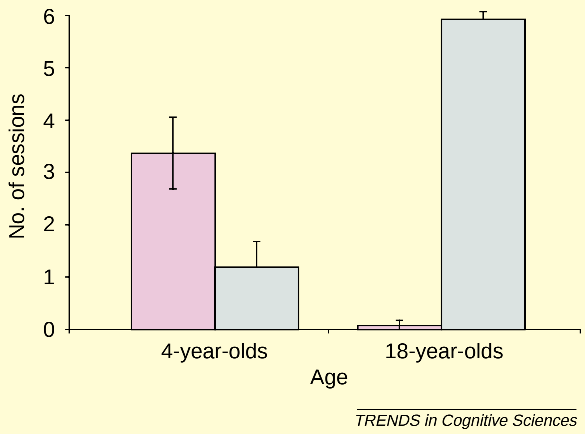

Fig. I. Number of test sessions (out of a maximum of 6) at above-chance performance by age group and stimulus modality (auditory, pink bars; visual, green bars) in an experiment by Sloutsky and Napolitano [18] (see text for details). Error bars represent standard errors of the mean. These results indicate that young children are more likely to focus on auditory stimuli, whereas adults are more likely to focus on visual stimuli. (Data reproduced from Ref. [18] with permission from SRCD).

Table 1. The overall structure of training and testing trials in Sloutsky and Napolitano [18]

Training session (n = 8 trials)		Testing session (n = 6 trials)	
VIS_1AUD_1 (trained set)	VIS_2AUD_2 (distracter set)	VIS_1AUD_{new} (test set A)	$VIS_{new}AUD_1$ (test set B)

AUD, auditory stimulus; VIS, visual stimulus.

We suggest that they do not have to know the importance of correspondences *a priori*, but that this knowledge could be a product of development and learning. We further argue that humans are endowed with powerful learning mechanisms that enable them to extract the importance of a particular match in a particular context. Some of these learning mechanisms are grounded in the ability to attend to and detect statistical regularities in the environment [7]. The remainder of this section reviews some of these mechanisms.

Perceptual learning

Perceptual learning is a process by which some features or stimulus dimensions become more distinct as a result of experience, whereas others become more equivalent. It is likely that statistical characteristics of stimuli underlie perceptual learning, although details of this process are not well understood. Perceptual learning has been shown to be an important mechanism for the development of categories during infancy: the perceptual system of infants can extract category-specific clusters of exemplars [22,23] (but see [24]).

Experiments using the preferential looking paradigm consisted of familiarization and test phases. During the familiarization phase, 3- to 4-month-old infants were presented with pictures of either dogs or cats. During test trials, a novel member of a familiarized category was paired with a member of a novel category. As a result of familiarization, infants learned to differentiate between members of the familiar and novel categories. Converging evidence coming from other infancy research paradigms suggests that experience with multiple exemplars might direct the perceptual system towards extracting important category-specific regularities [25].

Flexible adjustment of attentional weights

Some features are natural 'attention grabbers' because of their psychophysical characteristics (e.g. a sudden loud sound or a bright flash of light), whereas salience of others might change across various contexts. In particular, salience of the same feature might vary as a function of its statistical characteristics, such as diagnosticity or the ability to partition a set of entities [26]. For example, when there is a red triangle and a red circle, shape is more salient than color, whereas when there is a red triangle and a blue triangle, color is more salient than shape. Furthermore, if a feature is diagnostic, it might increase attentional weights of other correlated features. For example, when 2- to 3-year-olds were asked to extend a learned word to novel instances, they attended to both shape and texture if the entities presented had eyes (presence of eyes is diagnostic for the animate–inanimate distinction), whereas when the same entities were presented without eyes, participants attended only to shape [27] (but see [28]).

There is evidence that even young children can extract diagnosticity, and that diagnosticity might affect attentional weights of features. For example, researchers [29] presented 3-year-olds with novel objects, each accompanied with a novel label and information about the objects'

Box 2. Questions for Future Research

- How do mechanisms of inductive generalization change in the course of development?
- If early in development categorization is similarity-based, does it remain similarity-based later in development?
- If matching labels contribute to young children's similarity of labeled entities, does mere similarity of labels also contribute to similarity of labeled entities?
- What are the origins and the developmental course of abstract relational concepts?

functions. Each object had a distinct base shape and distinct parts. Participants were also presented with a non-member of the category that, by means of contrast, made either the base or the parts diagnostic and thus salient. Participants were then asked to extend the name to novel examples. Depending on the contrast condition, young children generalized either on the basis of shape (when shape was more diagnostic and thus salient) or parts (when parts were more diagnostic and thus salient). Hence the importance of features does not have to be known in advance – it can be 'created' by presenting supporting and contrasting examples.

Abstraction of detected regularities

Because many 'basic-level' categories (e.g. dogs or birds) have correlated structures (e.g. creatures who have feathers are likely to fly), the ability to detect specific and more abstract regularities might be an important learning mechanism supporting the development of categories. For example, suppose that in one's environment there are two correlated feature dimensions (e.g. color and shape), with each dimension having four values: most of the balls are red, most of the cars are yellow, most of the tables are white and most of the birds are green. The learning organism would first detect specific correlations between the red color and balls, and then a more abstract correlation between color and shape. This is because to detect correlations between the dimensions of color and shape, one needs a large body of evidence coming from the 16 cells of the 4-by-4 correlational matrix. However, a significantly reduced matrix is needed to detect a local contingency between a specific color and a specific shape. Therefore, when the number of instances is small, only local contingencies can be detected. As the number of instances grows, learners can detect more abstract correlations between dimensions and use these correlations in the development of categories.

Both the ability to extract regularities and the ability to use earlier detected specific regularities for the detection of more abstract regularities are present early in development [7,11,30,31]. In particular, in the course of language acquisition, babies first detect that a particular shape co-occurs with a particular label (e.g. 'cup-shaped' objects are called 'cups' whereas 'ball-shaped' objects are called 'balls'). As the corpus of evidence grows, they detect a more abstract regularity: similar-shaped entities have the same name, and this more abstract regularity, in turn, makes shape an important predictor of category membership [30].

Limits to the similarity-based approach and questions for future research

In summary, the similarity-based approach provides a plausible account of conceptual development. However, many questions remain unanswered, and they require further research. In addition, there are different types of category, and it is possible that perceptual similarity might ground some of these types, but not others. For example, natural kinds (e.g. *cat* or *bird*) are more likely to be grounded in perceptual similarity than abstract relational concepts (e.g. *fairness*, *proportionality*, *logical necessity* or *momentum*), and little is known about the development of abstract relational concepts. More research is needed to gain a fuller account of conceptual development (see Box 2. Questions for Future Research).

Acknowledgements

The writing of this paper was supported by a grant from the National Science Foundation (BCS # 0078945) to Vladimir M. Sloutsky.

References

1 Inhelder, B. and Piaget, J. (1964) *The Early Growth of Logic in the Child: Classification and Seriation*, Routledge Kegan Paul
2 Vygotsky, L.S. (1986) *Thought and Language*, MIT Press
3 Bruner, J.S. *et al.* (1956) *A Study of Thinking*, John Wiley & Sons
4 Smith, E.E. and Medin, D.L. (1981) *Categories and Concepts*, Harvard University Press
5 Murphy, G.L. (2002) *The Big Book of Concepts*, MIT Press
6 Keil, F.C. *et al.* (1998) Two dogmas of conceptual empiricism: implications for hybrid models of the structure of knowledge. *Cognition* 65, 103–135
7 McClelland, J.L. and Rogers, T.T. (in press) The parallel distributed processing approach to semantic cognition. *Nat. Rev. Neurosci.* 0, 000–000
8 Goldstone, R.L. (1994) The role of similarity in categorization: providing a groundwork. *Cognition* 52, 125–157
9 Hahn, U. and Ramscar, M. (2001) *Similarity and Categorization*, Oxford University Press
10 Sloman, S.A. and Rips, L.J. (1998) *Similarity and Symbols in Human Thinking*, MIT Press
11 Jones, S.S. and Smith, L.B. (2002) How children know the relevant properties for generalizing object names. *Dev. Sci.* 5, 219–232
12 Sloutsky, V.M. *et al.* (2001) How much does a shared name make things similar? Linguistic labels and the development of inductive inference. *Child Dev.* 72, 1695–1709
13 Gelman, S.A. and Coley, J.D. (1991) Language and categorization: the acquisition of natural kind terms. In *Perspectives on Language and Thought: Interrelations in Development* (Gelman, S.A. and Byrnes, J.P., eds) pp. 146–196, Cambridge University Press
14 Gelman, S.A. and Markman, E. (1986) Categories and induction in young children. *Cognition* 23, 183–209
15 Deak, G.O. and Bauer, P.J. (1996) The dynamics of preschoolers' categorization choices. *Child Dev.* 67, 740–767
16 Mak, B.S.K. and Vera, A.H. (1999) The role of motion in children's categorization of objects. *Cognition* 71, B11–B21
17 Sloutsky, V.M. and Lo, Y-F. (1999) How much does a shared name make things similar? Part 1: Linguistic labels and the development of similarity judgment. *Dev. Psychol.* 35, 1478–1492
18 Sloutsky, V.M. and Napolitano, A.C. (in press) Is a picture worth a thousand words? Preference for auditory modality in young children. *Child Dev.* 74, 000–000
19 Xu, F. (2002) The role of language in acquiring object kind concepts in infancy. *Cognition* 85, 223–250
20 Lewkowicz, D.J. (1988) Sensory dominance in infants: Part 1. Six-month-old infants' response to auditory–visual compounds. *Dev. Psychol.* 24, 155–171
21 Lewkowicz, D.J. (1988) Sensory dominance in infants: Part 2. Ten-month-old infants' response to auditory–visual compounds. *Dev. Psychol.* 24, 172–182

22 Quinn, P.C. *et al.* (1993) Evidence for representations of perceptually similar natural categories by 3-month-old and 4-month-old infants. *Perception* 22, 463–475
23 Quinn, P.C. and Eimas, P.D. (2000) The emergence of category representations during infancy: are separate and conceptual processes required? *J. Cogn. Dev.* 1, 55–61
24 Mandler, J.M. (2000) Perceptual and conceptual processes in infancy. *J. Cogn. Dev.* 1, 3–36
25 Mareschal, D. and Quinn, P.C. (2001) Categorization in infancy. *Trends Cogn. Sci.* 5, 443–450
26 Tversky, A. (1977) Features of similarity. *Psychol. Rev.* 84, 327–352
27 Jones, S.S. *et al.* (1991) Object properties and knowledge in early lexical learning. *Child Dev.* 62, 499–516
28 Booth, A.E. and Waxman, S.R. (2002) Word learning is 'smart': evidence that conceptual information affects preschoolers' extension of novel words. *Cognition* 84, B11–B22
29 Smith, L.B. *et al.* (1996) Naming in young children: a dumb attentional mechanism. *Cognition* 60, 143–171
30 Smith, L.B. *et al.* (2002) Object name learning provides on-the-job training for attention. *Psychol. Sci.* 13, 13–19
31 Younger, B.A. and Cohen, L.B. (1983) Infant perception of correlations among attributes. *Child Dev.* 54, 858–869
32 Balaban, M.T. and Waxman, S.R. (1997) Do words facilitate object categorization in 9-month-old infants? *J. Exp. Child Psychol.* 64, 3–26