

## Words can slow down category learning

Chandra L. Brojde · Chelsea Porter · Eliana Colunga

© Psychonomic Society, Inc. 2011

The uniquely human ability to use language not only allows for effective communication but may also provide necessary grounding for certain cognitive activities (e.g., Carruthers, 2011; Clark, 2006; Condry & Spelke, 2008; Jackendoff, 1996). Researchers have suggested that words may make abstract notions more tangible, allowing for the formation of concepts that are easier to interpret and manipulate (Clark, 2006), or that words provide additional processing power while task switching (Emerson & Miyake, 2003), reasoning about false beliefs (Newton & de Villiers, 2007), and learning new categories (Lupyan, Rakison, & McClelland, 2007).

Lupyan et al. (2007) found that named categories are easier to learn than unnamed categories, even when the names are redundant. In their study, participants were shown a series of novel objects (“aliens”), asked to respond to them (either to approach or escape), and then given auditory feedback (a buzz or chime). In the word condition, the feedback was followed by the presentation of a redundant category name. Participants in the word condition outperformed those in the no-word condition. This study demonstrated that words can enhance category learning, but the question of just what words do remains.

Two accounts have been proposed of how words may influence categorization. First, words may enhance a mental representation, providing an additional robust, perhaps symbolic, representation for the whole category of objects. For example, the word “cup” may allow us to think about the category *cup* more efficiently than the phrase “concave

object with a handle” and more abstractly than a mental picture of a cup. In this sense, words provide “material symbols” that augment category representations (Clark, 2006; Lupyan et al., 2007).

An alternative account is that the effects of words are attentional. Rather than augmenting representations, words may shift attention to the perceptual dimension relevant for the task at hand. In a categorization task, words may act as contextual cues, highlighting the features of an object that have proven to be predictive of category membership. In general, cues that have been probabilistically associated with some stimulus in the past enhance the detection, processing, and learning about that stimulus (Brady & Chun, 2007; Chun & Jiang, 1998), and attentional shifts to linguistic cues have been well documented (e.g., Christie & Klein, 1995; Robinson & Sloutsky, 2007; Slobin, 2003; Talmy, 2000). In addition, category learning has been shown to shift attention to diagnostic dimensions in the short term (e.g., Goldstone, Lippa, & Shiffrin, 2001), and possibly in the long term (e.g., Winawer et al., 2007).

In the context of object categorization, shape is a property that is typically predictive of category membership; every theory of object recognition deals with shape features (e.g., Biederman, 1987; Ullman, 1996). Moreover, basic-level categories are generally characterized by a prototypical shape (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976), and when forming novel lexical categories, people attend to shape over other features, such as material or color (e.g., Imai & Gentner, 1997). In short, these findings suggest that words can act as a contextual cue that shifts attention to predictive features and that shape is likely to be a historically predictive dimension for category membership.

---

C. L. Brojde (✉) · C. Porter · E. Colunga  
Department of Psychology and Neuroscience,  
University of Colorado at Boulder,  
Boulder, CO 80309, USA  
e-mail: chandrab@colorado.edu

The following two experiments were designed to distinguish between these two accounts—that words enhance mental representations or that they shift attention to relevant properties. If words act as “material symbols,” they should aid category learning, no matter which property is relevant to categorization. However, if words shift attention to a historically predictive property, such as shape, they should *hurt* category learning when the relevant feature does not have a history of being predictive. Thus, in both experiments, participants learned to categorize a set of novel stimuli by one of two different dimensions, with and without words.

## Experiment 1

In **Experiment 1**, participants learned to categorize novel stimuli by shape or by texture, with or without redundant category names. If words generally augment representations, performance should improve in the word condition, regardless of whether the categories to be learned are shape or texture based. However, if words shift attention to typically predictive dimensions—shape, in this case—the effect of having a word should be different depending on the categorizing dimension, and the addition of a word should be detrimental when learning texture-based categories.

### Method

#### Participants

A group of 80 undergraduates participated, with 20 randomly assigned to each of the four conditions in a 2

(relevant dimension: shape, texture)  $\times$  2 (label: word, no word) design.

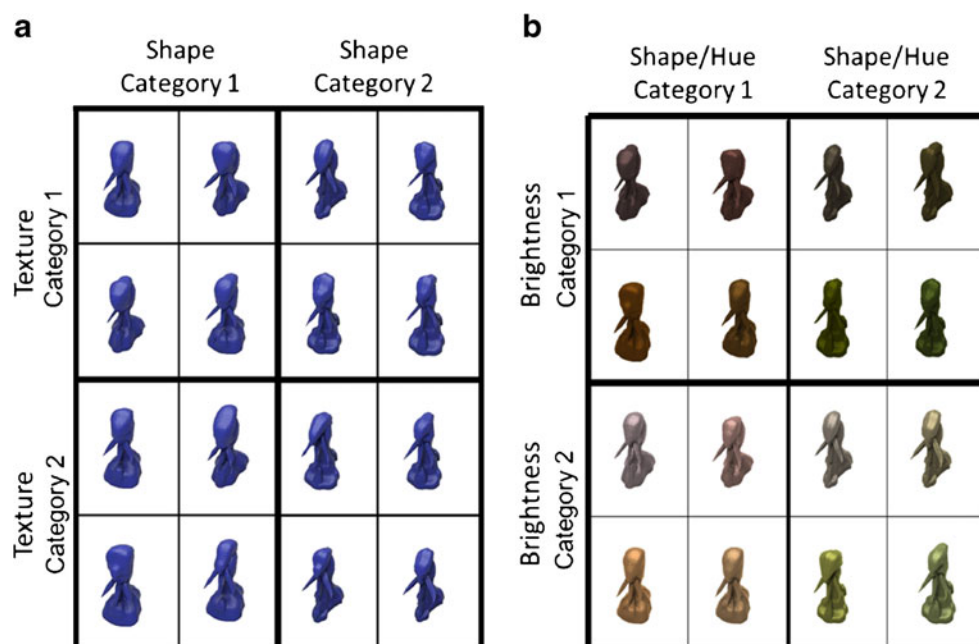
### Materials

The stimuli consisted of the same subset of 16 YUFO “alien” stimuli (Gauthier, James, Curby, & Tarr, 2003) used in Lupyan et al. (2007), but here the stimuli were texturized to add an orthogonal categorizing dimension. The two shape-based categories (Fig. 1a) were characterized by either a larger head ridge with a skinny base for one group, or a smaller head ridge with a larger base for the other (cf. Gauthier et al., 2003). The two texture-based categories were characterized by the amount of scaling and relief added in Photoshop—denser, deeper textures for one group, and sparser, shallower textures for the other.

In a preliminary experiment, 13 additional participants were shown all possible pairs for all of the shapes used (without textures) and all of the textures used (on squares) and were asked to provide similarity judgments on a 1–7 scale. A 2 (dimension: shape, texture)  $\times$  2 (pair type: within categories, between categories) repeated measures ANOVA on their similarity ratings showed that items from the same category were judged to be more similar to each other than items from different categories,  $F(1, 12) = 143.63$ ,  $\eta^2 = .92$ ,  $p < .001$ , although texture-based categories were more discriminable than the shape-based categories,  $F(1, 12) = 6.95$ ,  $\eta^2 = .37$ ,  $p < .05$ .

The final stimulus set consisted of the 16 textures layered onto the 16 shapes, so that the resulting 16 stimuli could be categorized in two orthogonal ways, by shape or by texture (Fig. 1a). For half of the participants, shape was

**Fig. 1** Final stimuli used in **Experiment 1** and **2**. The dark vertical lines represent the two categories of aliens separated (a) by shape for **Experiment 1** and (b) by shape/hue for **Experiment 2**. The dark horizontal lines represent the two categories separated (a) by texture for **Experiment 1** and (b) by brightness for **Experiment 2**



the relevant dimension; for the other half, texture was the relevant dimension. Additionally, in the word condition, “grecious” and “leebish” were used to label the two different categories, counterbalanced across participants. In the no-word condition, participants heard no words.

### Procedure

**Training trials** As in Lupyan et al. (2007), participants read directions asking them to imagine that they were exploring another planet and that on this planet they would encounter alien life forms. For each alien, participants were to choose whether to approach it or move away from it, based on its appearance. For each trial, a cartoon explorer appeared at the top, bottom, left, or right of the screen. After 500 ms, one of the 16 aliens appeared in the middle of the screen. Participants used the arrow keys to indicate which direction the explorer should move, and the explorer moved accordingly. So, if the explorer appeared above the alien, the “down” key indicated a choice to approach and the “up” key indicated a choice to escape, but if the explorer appeared below the alien, the same keys would indicate the opposite responses, thereby eliminating procedural learning.

A total of 200 ms after participants made their choice and the explorer had stopped moving, the participants received auditory feedback: a chime if correct, and a buzz if incorrect. In the word condition, 300 ms after the chime/buzz, participants heard a female voice say either “grecious” or “leebish,” depending on the category the alien belonged to, approachable or not. In the no-word condition, participants heard only the feedback sounds, but both the total trial durations and the exposure to the stimuli were equal across the two label conditions. Participants saw nine blocks of stimuli with 16 trials per block—that is, 1 trial for each of the 16 alien stimuli.

**Verification trials** As in Lupyan et al. (2007), we measured whether the participants had learned the association between the words and the stimuli. Verification trials were included after a random 10% of training trials in the word condition. In these trials, a random picture of one of the aliens appeared with the following question: “Is this one leebish [grecious]? yes/no.” No feedback was given for verification trials.

## Results

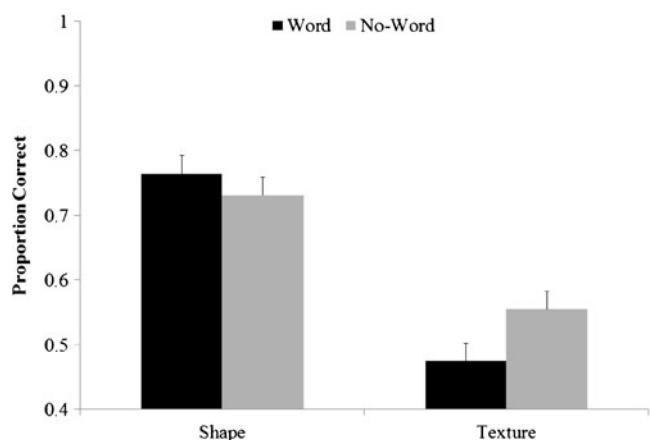
### Training trials

To investigate the effect of words on learning shape-based versus texture-based categories, a Block x Relevant Dimension (shape, texture) x Label (word, no word)

ANOVA was conducted on the average proportions correct. There was no main effect of label,  $F(1, 76) = 0.69, p > .4$ , so words did not help categorization all around. Instead, and consistent with the attentional shift account, there was a significant interaction between label and relevant dimension,  $F(1, 76) = 4.07, \eta^2 = .05, p < .05$  (Fig. 2). The effect of having redundant category labels depended on whether the relevant dimension was shape or texture. Planned comparisons revealed that, when categorizing by texture, participants performed better in the no-word condition than in the word condition,  $t(38) = 2.01, p = .02$ . In contrast, when categorizing by shape, there was no effect of label,  $t(38) = 0.84, n.s.$  There were also main effects of block,  $F(8, 608) = 16.08, \eta^2 = .18, p < .001$ , and relevant dimension,  $F(1, 76) = 70.10, \eta^2 = .48, p < .001$ , and a significant Block x Relevant Dimension interaction,  $F(8, 608) = 5.55, \eta^2 = .07, p < .001$ ; participants were more accurate over time and more accurate when learning to categorize by shape than by texture. In general, participants performed better than chance in the shape condition (in all blocks in the word condition, or after the second block in the no-word condition), but participants in the texture condition never performed above chance. The overall poor performance in the texture condition may seem puzzling, given the results of the discriminability pretest, but recall that those textures were presented on squares; overlaying textures on YUFOs might have changed their relative discriminability.

### Verification trials

A one-sample  $t$  test showed that participants learned the words “leebish” and “grecious” at above chance levels when categorizing by shape,  $t(19) = 11.55, p < .001$ , but not when categorizing by texture,  $t(19) = 1.66, n.s.$ , and better



**Fig. 2** Results of the Experiment 1 training trials for proportions correct

in the shape condition than in the texture condition,  $t(39) = 6.21, p < .001$ . In addition, there was a significant Pearson's correlation between accuracy in the verification and the training trials for the shape condition,  $r(18) = .54, p < .01$ , but not for the texture condition,  $r(18) = .20, n.s.$

## Discussion

**Experiment 1** shows that words hinder learning for texture-based categories. This result is consistent with the idea that words shift attention to shape, making it harder to learn to categorize by other dimensions. However, the poor performance in the texture condition leaves open the possibility that the deleterious effect of words when categorizing by texture might be the result of making an already hard task even harder. Moreover, the effect of words in the shape condition did not reach significance, possibly because participants reach ceiling very quickly when categorizing by shape. Another demonstration of the attentional effect of words would be to show a shift in attention away from an alternative dimension—that is, other than shape—that is equally valid as a basis for categorization. For example, if both shape and hue are valid covarying diagnostic dimensions, would participants be more likely to learn to categorize by shape alone than by both shape and hue in the word than in the no-word condition?

## Experiment 2

To test whether the detrimental effect of the redundant category label would extend to other dimensions that are as readily learnable as shape, we manipulated brightness instead of texture. To test whether words would shift attention to shape selectively even when there was a second diagnostic dimension, we included hue as a second relevant dimension in addition to shape; shape and hue covaried, and each of them independently determined category membership. Note that in this shape/hue condition, correct performance could reflect attention to shape, hue, or both. Thus, in addition to measuring accuracy, we collected reports on categorization strategy use after the experiment was completed. If, as suggested by **Experiment 1**, words shift attention to shape and away from other dimensions, participants should report using shape alone more in the word condition than in the no-word condition, even though they could have successfully used hue, shape, or both.

## Method

### Participants

A group of 133 undergraduate students were assigned to one of four conditions in a  $2 \times 2$  design with relevant

dimension (shape/hue, brightness) and label (word, no word) as between-subject variables. The final 53 participants (40%) were asked about their strategy use.

### Materials

To create the color set, we first chose 8 colors in the Munsell color space, starting from a random color; each color was one hue step apart from the previous one. These same 8 hues were selected at two different brightness levels, four steps apart from each other, to make the final 16 colors used. Thus, the 16 colors used could be separated, orthogonally, into two subsets, either by hue (8 greenish, 8 reddish) or by brightness (8 light, 8 dark). For the final stimuli, shown in Fig. 1b, the 16 colors were layered onto the 16 YUFO shapes from **Experiment 1** such that the 8 greenish colors (4 light and 4 dark) were on the aliens with larger head ridges and narrower bases, and the reddish colors (4 light and 4 dark) were on the aliens with smaller head ridges and wider bases. The resulting 16 stimuli were then divisible, orthogonally, in two ways: either by shape and hue or by brightness.

Again, 12 additional participants rated the similarity among pairs of aliens on a 1–7 scale, and a 2 (relevant dimension: shape/hue, brightness)  $\times$  2 (pair type: within categories, between categories) repeated measures ANOVA showed that they were discriminable,  $F(1, 11) = 76.23, \eta^2 = .88, p < .001$ . Unlike in **Experiment 1**, this time there was no significant interaction between relevant dimension and pair type.

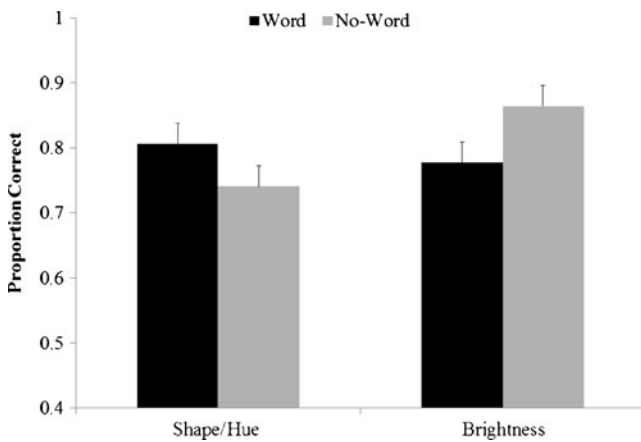
### Procedure

Participants followed the same procedure as in **Experiment 1** with the new stimuli. After the training portion of the experiment, the participants were asked, “What strategies did you use to help distinguish the aliens?”

## Results

### Training trials

As in **Experiment 1**, a Block  $\times$  Relevant Dimension (shape/hue, brightness)  $\times$  Label (word, no word) ANOVA on the average proportions correct revealed the predicted significant interaction between label and relevant dimension,  $F(1, 129) = 8.21, \eta^2 = .06, p < .01$  (Fig. 3). When categorizing by brightness, participants performed better without words,  $t(65) = 2.65, p = .01$ . In contrast, when categorizing by shape/hue, there was no effect of label,  $t(64) = 1.26, n.s.$  Unlike in **Experiment 1**, however, there was no main effect of relevant dimension grouping,  $F(1, 129) = 0.44, n.s.$ ; both categorizations were equally readily learned. As in **Experiment 1**, participants' perfor-



**Fig. 3** Results of the Experiment 2 training trials for proportions correct

mance improved across blocks,  $F(8, 1032) = 110.93$ ,  $\eta^2 = .46$ ,  $p < .001$ . However, a three-way interaction between block, label, and relevant dimension showed that the negative effect of words when categorizing by brightness held for each of the first five blocks, all  $ps < .05$ , but disappeared afterward, all  $ps > .05$ . In contrast, the positive effect of words when categorizing by shape/hue was present only in the first block,  $p < .05$ .

#### Reported strategies

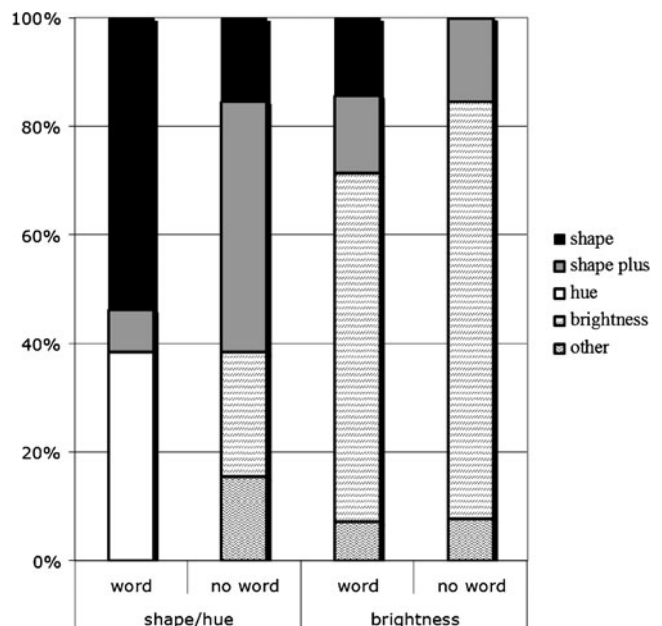
Results from participants' self-reported strategies offer a complementary view of the processes underlying the effects of words on category learning. Figure 4 shows the strategies reported, classified into one of five categories, based on the nouns and adjectives used in response to the question "What strategies did you use to help distinguish the aliens?": (1) shape (e.g., round, pointy, the word "shape"), (2) hue (e.g., green, purple, "color"), (3) brightness (e.g., light, pastel, brighter), (4) shape-plus, a combination of shape and brightness or hue, and (5) don't know or no strategy used. As predicted, in the shape/hue condition, participants in the word condition were more likely to report using shape exclusively and less likely to report using both shape and hue than participants in the no-word condition, Fisher's exact  $p = .04$ . Although the same number of participants (8/13) reported using a shape-based strategy in both label conditions, in the word condition 7 participants used shape alone, as compared to only 2 in the no-word condition. There was no effect of label on reported strategies when learning to categorize by brightness (Fig. 4, last two columns),  $\chi^2(3) = 2.02$ , n.s., with most participants in both label conditions using brightness alone (word, 9/14; no word, 11/13). This is not surprising, since brightness was the only diagnostic dimension in this condition, and all participants figured this out by the end of the experiment. The key finding, however, was that in the condition in

which participants could have attended to shape, hue, or both, participants were more likely to exclusively report attending to shape and to ignore hue when they heard words than when they did not.

Overall, participants who reported using a correct strategy performed significantly better than those who did not,  $t(51) = 4.20$ ,  $p < .001$ , suggesting that they were not solving the task through implicit learning. This was true in both the brightness condition,  $t(25) = 2.94$ ,  $p < .001$ , and the shape/hue-no-word condition,  $t(11) = 3.44$ ,  $p < .05$ . (We did not test the shape/hue-word condition, since no participants reported incorrect strategies there.) Furthermore, in the shape/hue condition, participants who reported using both the shape and hue dimensions performed better than those using only one of the dimensions,  $t(19) = 2.03$ ,  $p = .057$ , indicating that attending to both dimensions rather than to just one might be advantageous.

#### Verification trials

Participants learned the words "leebish" and "grecious" at above-chance levels when categorizing by either shape/hue,  $t(34) = 7.44$ ,  $p < .001$ , or brightness,  $t(33) = 14.38$ ,  $p < .001$ , and in contrast with Experiment 1, learned the words equally well in both conditions,  $t(67) = 1.85$ , n.s. In addition, there was a significant Pearson's correlation between accuracy on verification trials and throughout training for the shape/hue condition,  $r(31) = .65$ ,  $p < .01$ , and for the brightness condition,  $r(32) = .73$ ,  $p < .01$ , suggesting that participants learned the categories better as they learned the words, regardless of categorizing dimension.



**Fig. 4** Self-reported strategies used by participants in Experiment 2 to categorize aliens

## General discussion

Considerable past research on the influence of words on cognition has focused on their advantageous effects. The present results suggest that words direct attention to some stimulus dimensions over others in ways that can sometimes be detrimental. Thus, words do not always help in learning categories; they help when the task-relevant dimension in the current task aligns with relevant dimensions in previous similar tasks, and they hurt when the dimensions do not align. This detrimental effect of redundant words was observed for two nonshape dimensions, texture and brightness, and the attentional shift toward shape for named categories was confirmed in the strategies used for distinguishing categories as reported by participants.

Intriguingly, aside from shifting attention to shape, words may foster selective attention in general. The strategy reports in [Experiment 2](#) suggest that people are more likely to focus attention on *one* dimension when given a word: When participants could use both shape and hue as a basis for categorization, in the word condition they tended to use one or the other dimension, but in the no-word condition they tended to use both. This result complements Lupyan's (2009) finding that verbal interference disrupts the ability to selectively attend to task-relevant dimensions. It further suggests that words may encourage selective attention even when such attention is neither necessary nor advantageous.

Previous research showing negative effects of words, such as verbal overshadowing, has involved words with already known meanings. For example, generating verbal descriptions or category names for stimuli decreases memory performance (Lupyan, 2008; Melcher & Schooler, 1996). These studies suggest that known category labels interfere with the processing of perceptual information, making the details of complex stimuli hard to remember. This same blurring-of-details effect may be behind some of the positive effects of novel words when the task involves not memory for individual instances, but categorization over instances. The present study finds a negative effect of *novel* words, but only when learning categories that are not based on shape, suggesting that shifting attention to shape may be a default for novel words that do not yet have meanings.

Of course, different sorts of words can be used to label different aspects of an object—object properties, actions, and even abstract notions—so why would novel words, as category names, direct attention to shape as a default? Love and Markman (2003) suggested that shape is integral to categories and that other properties, such as color, are represented in relation to shape; in addition, developmental studies have noted the centrality of shape in determining

category membership as a cue to function or kind (Gelman & Bloom, 2000; Prasada, Ferenz, & Haskell, 2002). The origin of the link between shape and novel words, however, may be developmental and rooted in the processes of attentional learning (e.g., Robinson & Sloutsky, 2007; Yoshida & Smith, 2003). It seems possible that words shift attention to shape because in English and many other languages, object names are systematically correlated with object categories (Colunga & Smith, 2002, 2005), and more specifically with shape-based categories (Samuelson & Smith, 1999). Over time, we may learn that things that have the same name are likely to have similar shapes, and a category name becomes a contextual cue that shifts attention away from nonshape dimensions. If this were the case, we might expect a different pattern of results in speakers of languages such as Yucatec Mayan, with its material-based nouns (Lucy & Gaskins, 2001). In other words, words are what they do; they are tools. We use them, and they become specialized as that tool, directing attention to historically task-relevant dimensions. It is not representation, but usage, that matters.

**Author Note** The authors thank Gary Lupyan for providing the original materials used in Lupyan et al. (2007). We also thank A. Miyake and Y. Munakata for helpful comments.

## References

- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*, 115–147.
- Brady, T. F., & Chun, M. M. (2007). Spatial constraints on learning in visual search: Modeling contextual cuing. *Journal of Experimental Psychology: Human Perception and Performance*, *33*, 798–815.
- Carruthers, P. (2011). Language in cognition. In E. Margolis, R. Samuels, & S. Stich (Eds.), *The Oxford handbook of philosophy of cognitive science*. Oxford: Oxford University Press.
- Christie, J., & Klein, R. (1995). Familiarity and attention: Does what we know affect what we notice? *Memory & Cognition*, *23*, 547–550.
- Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, *36*, 28–71.
- Clark, A. (2006). Language, embodiment, and the cognitive niche. *Trends in Cognitive Sciences*, *10*, 370–374.
- Colunga, E., & Smith, L. B. (2002). What makes a word? In W. D. Gray & C. Schunn (Eds.), *Proceedings of the Twenty-Fourth Annual Conference of the Cognitive Science Society* (pp. 214–219). Mahwah: Erlbaum.
- Colunga, E., & Smith, L. B. (2005). From the lexicon to expectations about kinds: A role for associative learning. *Psychological Review*, *112*, 347–382.
- Condry, K. F., & Spelke, E. (2008). The development of language and abstract concepts: The case of natural number. *Journal of Experimental Psychology: General*, *137*, 22–38.
- Emerson, M. J., & Miyake, A. (2003). The role of inner speech in task switching: A dual-task investigation. *Journal of Memory and Language*, *48*, 148–168.

- Gauthier, I., James, T. W., Curby, K. M., & Tarr, M. J. (2003). The influence of conceptual knowledge on visual discrimination. *Cognitive Neuropsychology*, *20*, 507–523.
- Gelman, S. A., & Bloom, P. (2000). Young children are sensitive to how an object was created when deciding what to name it. *Cognition*, *76*, 91–103.
- Goldstone, R. L., Lippa, Y., & Shiffrin, R. M. (2001). Altering object representations through category learning. *Cognition*, *78*, 27–43.
- Imai, M., & Gentner, D. (1997). A cross-linguistic study of early word meaning: Universal ontology and linguistic influence. *Cognition*, *62*, 169–200.
- Jackendoff, R. (1996). How language helps us think. *Pragmatics and Cognition*, *4*, 1–34.
- Love, B. C., & Markman, A. B. (2003). The non-independence of stimulus properties in category learning. *Memory & Cognition*, *31*, 790–799.
- Lucy, J. A., & Gaskins, S. (2001). Grammatical categories and the development of classification preferences: A comparative approach. In M. Bowerman & S. C. Levinson (Eds.), *Language acquisition and conceptual development* (pp. 257–283). Cambridge: Cambridge University Press.
- Lupyan, G. (2008). From chair to “chair”: A representational shift account of object labeling effects on memory. *Journal of Experimental Psychology: General*, *137*, 348–369.
- Lupyan, G. (2009). Extracommunicative functions of language: Verbal interference causes selective categorization impairments. *Psychonomic Bulletin & Review*, *16*, 711–718.
- Lupyan, G., Rakison, D. H., & McClelland, J. L. (2007). Language is not just for talking: Redundant labels facilitate learning of novel categories. *Psychological Science*, *18*, 1077–1083.
- Melcher, J. M., & Schooler, J. W. (1996). The misremembrance of wines past: Verbal and perceptual expertise differentially mediate verbal overshadowing of taste memory. *Journal of Memory and Language*, *35*, 231–245.
- Newton, A. M., & de Villiers, J. G. (2007). Thinking while talking. *Psychological Science*, *18*, 574–579.
- Prasada, S., Ferenz, K., & Haskell, T. (2002). Conceiving of entities as objects and as stuff. *Cognition*, *83*, 141–165.
- Robinson, C. W., & Sloutsky, V. M. (2007). Visual processing speed: Effects of auditory input on visual processing. *Developmental Science*, *10*, 734–740.
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, *8*, 382–439.
- Samuelson, L. K., & Smith, L. B. (1999). Early noun vocabularies: Do ontology, category structure and syntax correspond? *Cognition*, *73*, 1–33.
- Slobin, D. I. (2003). Language and thought online: Cognitive consequences of linguistic relativity. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and thought* (pp. 157–191). Cambridge: MIT Press.
- Talmy, L. (2000). *Toward a cognitive semantics*. Cambridge: MIT Press.
- Ullman, S. (1996). *High-level vision*. Cambridge: MIT Press.
- Winawer, J., Witthoft, N., Frank, M. C., Wu, L., Wade, A. R., & Boroditsky, L. (2007). Russian blues reveal effects of language on color discrimination. *Proceedings of the National Academy of Sciences*, *104*, 7780–7785.
- Yoshida, H., & Smith, L. B. (2003). Known and novel noun extensions: Attention at two levels of abstraction. *Child Development*, *74*, 564–577.