

# Is Snow Really Like a Shovel?

## Distinguishing Similarity from Thematic Relatedness

Dedre Gentner (gentner@nwu.edu)  
Northwestern University  
Department of Psychology  
2029 Sheridan Rd.  
Evanston, IL 60208

Sarah K. Brem (sbrem@soe.berkeley.edu)  
University of California, Berkeley  
Graduate School of Education  
4533 Tolman Hall, #1670  
Berkeley, CA 94720-1670

### Abstract

Traditionally, thematic relatedness (*chicken and egg*) and similarity (*chicken and turkey*) have been thought of as distinct phenomena, the former the result of associative processes, and the latter reflecting comparison processes. However, recent studies (Bassok & Medin, 1996; Wisniewski & Bassok, 1996) suggest that similarity is a result of both association and comparison. This could call for a radical redefinition of similarity as inherently fused with association. We term this view the *integration* account. We consider an alternative, the *confusability* account, under which thematic influences intrude upon assessments of similarity but are not an essential part of the similarity process. We present two experiments supporting the confusability account. The first indicates that comparison and association are independent processes. The second shows that thematic influences rise with increased cognitive load. We believe that while a redefinition of similarity is not warranted, similarity is more vulnerable to error and intrusion than is generally thought.

Similarity is central to cognitive science. It plays a role in psychological and computational models of analogical reasoning (Forbus, Gentner & Law, 1995; Gentner, 1988; Holyoak & Thagard, 1989), problem solving (Ross, 1989; Novick, 1988), categorization (Markman & Wisniewski, 1997; Medin & Heit, in press; Murphy & Allopenna, 1994) and inference (Clement & Gentner, 1991; Lassaline, 1996). Much research in cognitive development has been predicated on this distinction: e.g., research that seeks the causes of the developmental shift from early reliance on thematic associations to later reliance on common-category relations (e.g., Markman, 1989; Vygotsky, 1978). Even models that question the degree to which similarity plays a role in cognitive processes base their criticisms on the assumption that similarity is a well-defined phenomenon (e.g., Murphy & Medin, 1985; Sloman & Rips, 1998). Therefore, when the possibility arises that our basic conceptualization of similarity is in need of radical restructuring, the implications are considerable. Recently, two studies have suggested just that.

Traditionally, the similarity between two items is thought to be a function of their distance in mental space, the overlap of their features (Tversky, 1977) or of their shared structure (Gentner, 1983, 1988; Gentner & Markman, 1997; Medin, Goldstone, & Gentner, 1993). All of these approaches assume that similarity results from a process of comparison, while thematic relatedness stems from association. The difference between these two types of

relationships is clear when one examines their predicate-argument structures. For example, according to Gentner (1988), a *shovel* and a *spade* are similar because they share **relations** (e.g., MOVE [shovel, here, there, stuff]/MOVE [spade, here, there, stuff]) and **attributes** (e.g., HANDLE [shovel]/HANDLE [spade]). In contrast, thematic associates like *snow* and *shovel* generally do not share attributes, and take different roles in relational predicates (e.g., MOVE [**shovel**, here, there, **snow**]). Because *snow* and *shovel* do not share features of appearance or use, both Gentner's and Tversky's models predict low similarity. Snow and shovel are related, though; experience informs people that snow and shovels interact, appear in the same propositions, and commonly co-occur. But this hinges upon noting relations that *associate* snow and shovel, rather than on *comparing* the two and noting common structure.

Perhaps because the distinction between thematic association and similarity seems obvious in many contexts its validity has not been the focus of much research. It seems clear that to base categorizations, word extensions, and inferences on thematic associations would frequently cause error. For example, whales are thematically related to the plankton they eat and to the harpoons with which whalers hunt them. However, to conclude that plankton or harpoons are warm-blooded like whales would be incorrect. Likewise, thematic associations play roles that similarity cannot. For example, associations allow us to detect and make predictions on the basis of covariation (Kelley, 1973;

Novick & Cheng, 1990). The association between smoke and fire leads us to sound the alarm when we see dark clouds billowing from an apartment window. But, looking for smoke when we see fog because they are perceptually similar would be a mistake. Differences such as these have provided a strong argument for treating similarity and thematic association as separate processes.

Beyond these intuitions, much of cognitive research depends on the theoretical assumption that thematic association and similarity relationships are distinct. Therefore, the finding that the two are not psychologically separable would radically alter our understanding of cognition. Yet, some recent theories suggest a strong link between similarity and association. Sloman (1996) proposes that similarity and association are processed by a single system, with abstract rules reposing in a separate system. Further, Bassok and Medin (1996) have recently found that when judging the similarity of two sentences, participants are influenced by thematic relationships between the sentences. For example, participants judged the sentences *The carpenter fixed the chair* and *The carpenter sat on the chair* to be similar because "the carpenter sat on the chair to see if the repair would hold." Likewise, Wisniewski and Bassok (1996) compared the similarity ratings assigned to similar pairs (*milk-lemonade*), thematic pairs (*milk-cow*) and pairs sharing both similarity and thematic relationships (*milk-coffee*). Thematic relationships significantly increased similarity ratings: e.g., *milk* and *coffee* were rated as more similar than *milk* and *lemonade*. Although previous research has shown that similarity and thematic relationships are confusable by children (e.g., Bauer & Mandler, 1989), these experiments go further in suggesting that these confusions occur for adults as well.

These results force us to consider the possibility that traditional models oversimplify the concept of similarity. Under Medin and Bassok's (1996) view, similarity is an integration of independent processes of comparison and association. However, an even more radical position can be taken, consistent with Sloman's position. Under this view, which we term the *integration account*, thematic associations and comparisons are the result of a unitary process. Similarity is thus inherently intermixed with thematic relatedness.

Another possibility is that the comparison process is distinct from association, and thematic influences arise as the result of thematic intrusions that interfere with similarity judgments. According to this *confusability account*, similarity is the result of the comparison process, but this process can be derailed by other factors. Such derailment might occur if participants have difficulty distinguishing between the mental output that arises from accessing associations as opposed to the output of a separate, independent comparison process.

Several lines of reasoning lead us to entertain the confusability account. First, Markman (1989) argues that the developmental shift towards a preference for taxonomic groupings does not indicate a loss of salience for thematic relationships, but a rise in the salience of taxonomic relationships. Likewise, Smiley and Brown (1979) argue that the shift is one of preference, rather than a radical

restructuring of knowledge. If adults remain sensitive to thematic associations, thematic relationships could interfere with other processes.

A second line of reasoning concerns individual differences. We have found evidence of substantial variation in people's ability to distinguish and identify similarity relations. In a screening task (further described below), participants were given a standard (e.g., **dart**) and had to choose which alternative was most similar: **bullseye** or **javelin**. Eleven percent of the participants were consistently unable to distinguish similarity from thematic association (Undifferentiating). Another 41% showed varying degrees of differentiation. If such confusion were universal, it would suggest a basic fusion of the processes of comparison and association, supporting the integration account. However, 48% of participants show no confusion of similarity and thematic association (Differentiating).

Our proposal is that similarity and thematic relatedness result from two separate processes. However, the output of these processes is sometimes difficult to differentiate through direct introspection, particularly when similarity is very low (the conditions in which thematic intrusions have been found in similarity tasks). We conjecture that the ability to make the distinction reliably and explicitly is learned as part of the development of metacognitive skills. There is precedent for the view that internal cognitive distinctions must be learned. Markman (1979) found that 8-11 year-old children failed to report any comprehension difficulties when reading stories that contained blatant inconsistencies (e.g., that ants rely on smell, but have no noses and cannot smell). Nevertheless, the children were slower to read sentences that led to strong inconsistencies. This suggests that they engaged in some kind of inconsistency processing but could not *label* it as such. Elsewhere, we see evidence for people's inability to accurately distinguish between the cognitive and emotional states produced by real and implanted memories (e.g., Loftus, 1997), and an inability to correctly identify what they know or do not know (e.g., Koriat, 1993). Additionally, there is evidence of developmental shifts in the ability to reflect upon the products of cognition (e.g., Flavell, Green & Flavell, 1990; Kitchener & King, 1994). Going further, there is evidence for historical shifts in the use of comparison and association. Medieval alchemy treated similarity and thematic relatedness as interchangeable to a greater degree than in modern science (Gentner & Jeziorski, 1993). In short, we suggest that reliably distinguishing the sensation of relatedness via commonalities from the sensation of relatedness via thematic associations may require a degree of metacognitive sophistication.

On the confusability account, the findings of Bassok & Medin (1996) and Wisniewski & Bassok (1996) result from people's failure to explicitly distinguish the results of an associative retrieval process and the results of a comparison process. But if these processes are indeed separable, it should be possible to find a task that draws only on comparison. In Experiment 1, we do this by utilizing a word learning task.

## Experiment 1

To find out whether people can focus purely on similarity, we needed a task that naturally promotes a strong focus on commonalities. We drew on findings from cognitive development. Across many tasks researchers have noted a thematic preference in young children (Gelman & Baillargeon, 1983), and a shift towards taxonomic preferences with age and experience (Smiley & Brown, 1979). However, one task in particular has been found to elicit a reliable shift towards commonality, even among preschool children: namely, the task of extending a novel word to new items. In the classic word extension task, a child is shown a **dog** as standard, together with a **cat** (perceptually and taxonomically similar), and a **bone** (thematically related). If the child is simply asked to choose one that 'goes with' the dog, or even to choose 'another one' s/he will typically choose the bone. But if the child is told that dogs are called "feps" in a foreign language and asked to find the other "fep," s/he chooses the cat (Markman & Hutchinson, 1984; Waxman & Gelman, 1986).

The word learning task appears to invite a focus on common structure. Thus it can provide a test of whether people possess an internally separable comparison process. In this experiment we ask whether adults will show a strong focus on commonality in a word-extension task. If adults are like preschoolers in that they show only a weak preference for common structure, the results would be consistent with the integration account. To the extent that adults show a strong preference for common structure in the word extension condition, this would support the separability of comparison and association.

We showed participants triads of objects of the form:

**dog**  
**bone**                      **cat**

So as not to bias the experiment in favor of independent processes, we designed the triads to have strong thematic associates and similarity choices that were clear but not extremely close to the standard (e.g., not in the same basic level category). In the baseline condition we asked subjects to choose the item that "goes with" the standard. In the word condition, we told them the name of the standard (e.g., 'blicket') in a foreign language, and asked them to say which alternative was also a 'blicket'. In the baseline task, since both similar and thematic items are related to the base, either could in principle be an appropriate choice. However, because we chose highly salient thematic relations, we expect to see an advantage for the thematic alternative.

The key prediction concerns the word task. If comparison is indeed a separable process, then there should be a strong similarity focus in the word task, even if baseline responding is strongly thematic. In contrast, an integration account that postulates a single process combining comparison and association, would fail to predict a difference in the mental output generated by the two tasks.

### Method

**Materials.** Participants saw triads consisting of a standard and three alternatives: a comparison alternative, an

association alternative, and an unrelated foil. There were 12 such item sets. (See Table 1 for examples).

**Procedure & Design.** There were two levels of Task Type (between subjects): Word extension and Baseline. In the Word condition, subjects saw a sentence of the following format: "In a foreign country, 'X,' the people call spoons 'blicks.' Which of these is also a blick?" In the Baseline condition, subjects saw a sentence of the following format: "Which of the following goes with spoon?" In both tasks, the three alternatives were presented below the sentence, and subjects circled their choice.

**Participants.** Twenty-two Northwestern undergraduates participated, 11 in each condition.

Table 1. Examples of items used in Experiments 1 and 2

Base item	Comparison alternative	Association alternative	Unrelated foil
spoon	ladle	cereal	shirt
rocket	missile	astronaut	belt
garlic	onion	vampire	cement

### Results & Discussion

The results of Experiment 1 are presented in Table 2. As predicted, participants' choices depended overwhelmingly on Task Type. In the Baseline task, participants chose the similar item in 2.2% of the cases. In the Word task, they chose the similar item in 96.9% of the cases. This difference is significant (by participants,  $t(11) = 33.41$ ,  $p < 0.0001$ ; by items,  $t(11) = 45.51$ ,  $p < 0.0001$ ).

Table 2. Experiment 1: Percentage of similarity responses as a function of Task Type.

Task Type	% Similarity Responses
<b>Baseline</b>	2.2%
<b>Word extension</b>	96.9%

These results suggest that similarity and thematic relatedness are distinct, separable processes for adults. The strong form of the integration account is thus seriously undermined by these findings, as they cannot be accounted for by a unitary process fusing comparison and association.

Interestingly, we found no individual differences in Experiment 1. In the Baseline task, only two participants chose any similar alternatives, doing so 8% and 16% of the time. In the Word task, only two participants chose thematic alternatives, doing so 8% and 25% of the time. This striking contrast with the screening results described above suggests that the two processes may be implicitly called forth by different cognitive tasks. Just as task support helps children focus on different types of relationships, task support may aid adults in distinguishing similarity from association.

The results of Experiment 1 support the confusability account, and are not compatible with conceptualizing comparison and association as the result of a single, undifferentiating process. Participants can focus on

similarity relations even in the fact of very salient thematic relations when a naturalistic task requires them to do so.

## Experiment 2

While the results of Experiment 1 support the proposal that comparison and association are independent processes producing distinct output, we now need to address the confusion of these outputs. As noted above, these confusions occurred not only in Bassok & Medin (1996) and Wisniewski & Bassok (1996), but in our own screening studies. Under the confusability account, these intrusions of thematic relatedness into similarity judgments arise because when participants experience difficulty in introspectively separating the outputs of comparison and association. This account implies that the more difficult it becomes to engage in introspection, the greater should be the confusion between similarity and thematic relatedness.

In Experiment 2, we asked subjects to make a similarity choice under time pressure, and varied whether there was a competing thematic relationship present. We used triads of words containing a standard item, an item similar to the standard, and either a thematic associate of the standard, or an unrelated foil.

Participants were to ignore thematic associations, and choose only items that were similar to the standard. We provided examples and substantial practice with feedback. To manipulate cognitive load, and thus vary participants' difficulty in examining the output of their mental processes, we employed two deadlines, one at 1000 ms and one at 2000 ms.

Any account would predict more errors at the shorter deadline. But the integration and confusability accounts make very different predictions regarding errors at the longer deadline. The confusability account predicts that shorter deadlines will differentially increase the error rate in the presence of thematic distractors. Thus there will be an interaction between triad type and deadline. In contrast, in the integration account, the inability to set aside thematic influences arises from fusion at the process level. Such fusion should compromise the separation of comparison and association at both the shorter and the longer deadlines. Thus, both accounts predict more errors at the shorter deadline, but only the 'separate but confusable processes' account predicts an interaction between deadline and the type of triad.

### Method

**Screening task.** We tested 702 participants as part of a group testing session, measuring their ability to distinguish similarity from thematic association using a triads similarity task. Given a standard (e.g., **dart**), they chose which alternative was most similar: e.g., **bullseye** or **javelin**. We selected two extreme groups of participants. The Undifferentiating group (11% of the participants) were unable to distinguish similarity from association in over 90% of cases. The Differentiating group (48% of participants) made the distinction correctly in over 90% of cases. The remaining 41%, who were intermediate in their performance, were omitted.

The screening task was completed an average of six

weeks prior to participation in this experiment, and no less than three weeks prior to participation. No mention of the screening task was made at the time of the actual experiment.

**Participants.** Eighty-four Northwestern undergraduates participated. From the initial screening group, 54 Differentiating and 30 Undifferentiating participants were drawn. All were fluent speakers of English.

**Materials.** Participants saw 64 triads consisting of a standard, a similar item, and either a thematically related alternative or an unrelated alternative. (See Table 1.)

**Design and Procedure.** We manipulated two variables within-subjects: Deadline (1000ms vs. 2000ms) and Alternative Type (Thematic vs. Unrelated). Deadline was blocked and counterbalanced across participants. Within each block, both Thematic and Unrelated triads were seen. Across participants, all items appeared in all conditions. We also included the categorical variable, Group, with two levels (Differentiated and Undifferentiated).

Participants were instructed to choose the similar alternative in each triad. The two types of relationships were described, and examples were given as illustrations. To allow participants to become accustomed to the procedure, they saw 24 practice items at the new deadline at the beginning of each block. They received feedback as to the accuracy of their responses during this practice session, but not during the test blocks.

The stimuli were presented by computer. For each item, participants saw lines of asterisks, which were replaced by the three words composing the triad. The triad remained on screen until the participants pressed the left or right cursor key to choose the left or right alternative. If participants failed to respond within the deadline, a buzzer sounded and the words "Too slow" flashed on the screen. Once participants made their choice, the next item was presented after a 1000 ms delay.

## Results & Discussion

All responses exceeding the deadline were discarded. The percentage of similarity responses is given in Table 3.

Table 3: Percentage of Similar responses in Experiment 2.

Differentiating Participants		
	1000ms	2000ms
<b>Alternative Type</b>		
<b>Thematic</b>	72.9%	86.9%
<b>Unrelated</b>	88.3%	96.8%
<b>Undifferentiating Participants</b>		
	1000ms	2000ms
<b>Alternative Type</b>		
<b>Thematic</b>	68.4%	84.6%
<b>Unrelated</b>	84.0%	95.1%

There is a main effect of Alternative Type; this is consistent with both accounts. Participants' performance when Similar alternatives were paired with Unrelated foils

was superior to their performance when Thematic alternatives were present (by item,  $F(1,63) = 49.24$ ,  $p < 0.001$ ; by subject,  $F(1, 82) = 72.78$ ,  $p < 0.001$ ). Also, as would be expected under either account, overall performance was better at the 2000ms deadline than at the 1000ms deadline (by item,  $F(1, 63) = 69.68$ ,  $p < 0.001$ ; by subject,  $F(1,82) = 44.36$ ,  $p < 0.001$ ).

However, consistent only with the confusability account, there was an interaction between Deadline and Alternative Type. Participants' ability to identify the Similar alternative was more vulnerable to time pressure in the presence of an Thematic distractor than in the presence of an Unrelated foil. The interaction is significant by subject, though marginal by item (by subject,  $F(1, 82) = 6.83$ ,  $p < 0.01$ ; by item,  $F(1,63) = 3.16$ ,  $p = 0.08$ ). The integration account predicts no such improvement in performance at the longer deadline.

With respect to individual differences, the performance of the Undifferentiating group was slightly lower in all conditions (by item,  $F(1, 63) = 4.80$ ,  $p < 0.05$ ; by subject,  $F(1,82) = 4.37$ ,  $p < 0.05$ ). The interaction of Deadline by Group was not significant (by item,  $F(1,63) = 1.98$ ,  $p > 0.10$ , by subject,  $F(1,82) = 1.03$ ,  $p > 0.10$ ), nor was any other interaction involving Group (all  $F_s < 1.1$ ). It appears that both Undifferentiating and Differentiating participants are challenged by shorter deadlines. This is consistent with a metacognitive process distinguishing comparison from association that requires some time to complete. Individual differences could arise because of differences in the speed and reliability of this process. Consistent with this possibility, the group differences were not specific to the similarity-thematic distinction. Rather, there was simply an overall reduction in accuracy for the Undifferentiating group.

It is also noteworthy that the Undifferentiating group did far better in this task than in the screening task. They made over 90% thematic false alarms on the screening task, but in Experiment 2 they were correct (i.e., able to ignore the association) 68% of the time at the shorter deadline, and 84% of the time at the longer deadline. This is a striking improvement. We conjecture that the initial practice session may have sharpened their ability to discern the thematic-similarity distinction. If so, this would be further evidence for separable processes with output that requires experiential practice to label reliably.

### General Discussion

The results presented here are most consistent with the confusability account: thematic association and similarity are distinct processes, and thematic associations intrude upon similarity judgments. Experiment 1 provides evidence that association and similarity are distinct processes, in that participants could set aside association based on the nature of the task, choosing similar alternatives in the Word task despite the dominance of thematic alternatives in the Baseline task. However, Experiment 2 suggests that even people who make this distinction easily and consistently under normal conditions can falter under strict deadlines. This is in line with the confusability account; it is harder for

participants to set aside thematic associations when they have to make quick decisions.

Although the strong form of the integration account can probably be dismissed, intermediate positions remain viable. For example, perhaps comparison and association are independent processes that are often (or perhaps typically) combined in judgments of similarity (Medin & Bassok, 1996). We also note that very common similarity relationships may be cached, rather than being recomputed each time they are encountered. Under these circumstances, similarity and thematic relatedness would often coincide.

Our conjecture is that the major culprit in the thematic intrusion phenomenon is a lack of introspective awareness of cognitive states. This is evidenced by the fact that even people who easily distinguished between similarity and association in the screening task experienced difficulty with increased cognitive load. As we noted in the introduction, there is a considerable body of evidence suggesting that people are fallible in their ability to reflect accurately upon their own mental states. For example, Goldstone (1994) comments on the fact that people can be highly variable in their interpretation of a similarity rating task, sometimes equating the task of judging similarity with that of judging purely perceptual similarity. (See Goldstone, Medin, & Gentner, 1993, for other instances of variability in rated similarity). The results of Experiment 2 as well as the screening task suggest that people experience difficulties in distinguishing and internally labeling the mental states produced by association on the one hand and comparison on the other.

Such confusion could significantly affect reasoning tasks, as well as a person's insight into their own responses. For example, in a speeded version of the word-learning task used in Experiment 1, would participants make errors of the sort seen in Experiment 2? Another intriguing question is whether the contextual tasks that support people's ability to focus on similarity also allow them to introspectively access the reason for their actions (see Markman's (1979) studies of inconsistency detection). Answers to these questions will not only help us to understand how similarity and association are cognitively represented and used, but may also shed light on individual differences in reasoning and problem solving. In the end, we believe that these new possibilities will lead to constructive debates over the concept of similarity, a concept so pervasive in cognitive science research that it is often taken for granted.

### Acknowledgments

The authors thank Philip Wolff, Andrea Boyes, and an anonymous reviewer for comments on this version and previous drafts. We also thank Kathleen Braun for her help in analyzing Experiment 1. This research was supported in part by the National Science Foundation grant SBR-9511757 awarded to the first author.

### References

- Bassok, M. & Medin, D. L. (1996). Birds of a feather flock together: Similarity judgments with semantically rich stimuli. *Journal of Memory and Language*, 36, 311-336.

- Bauer, P. J., & Mandler, J. M. (1989). Taxonomies and triads: Conceptual organization in one-to-two-year-olds. Cognitive Psychology, *21*, 156-184.
- Cheng, P. W. & Novick, L. R. (1990). A probabilistic contrast model of causal induction. Journal of Personality & Social Psychology, *58*, 545-567.
- Clement, C. A. & Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. Cognitive Science, *15*, 89-132.
- Flavell, J. H., Green, F. L., & Flavell, E. R. (1990). Developmental changes in young children's knowledge about the mind. Cognitive Development, *5*, 1-27.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. Cognitive Science, *7*, 155-170.
- Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou and A. Ortony (Eds.) Similarity, analogy, and thought (pp. 189-241). Cambridge, England: Cambridge University Press.
- Gentner, D. & Jeziorski, M. (1993). The shift from metaphor to analogy in Western science. In A. Ortony (Ed.) Metaphor and thought. (pp. 447-480). New York: Cambridge University Press.
- Gentner, D. & Markman, A. B. (1997). Structure mapping in analogy and similarity. American Psychologist, *52*, 45-56.
- Goldstone, R. L. (1994). The role of similarity in categorization: Providing a groundwork. Cognition *52*(2), 125-157.
- Holyoak, K. J. & Thagard, P. (1989). Analogical mapping by constraint satisfaction. Cognitive Science, *13*, 295-355.
- Kelley, H. H. (1967). Attribution theory in social psychology. Nebraska Symposium on Motivation, *15*, 192-238.
- King, P. M. & Kitchener, K. S. (1994). Developing Reflective Judgment : Understanding and Promoting Intellectual Growth and Critical Thinking in Adolescents and Adults. San Francisco: Jossey-Bass.
- Koriat, A. (1993). How do we know what we know? The accessibility model of the feeling of knowing. Psychological Review, *100*, 609-139.
- Loftus, E. F. (1997). Creating childhood memories. Applied Cognitive Psychology, *11*, S75-S86.
- Markman, A. B. (1997). Constraints on analogical inference. Cognitive Science, *21*(4), 373-418.
- Markman, A. B. & Wisniewski, E. (1997). Similar and different: The differentiation of basic-level categories. Journal of Experimental Psychology: Learning, Memory, and Cognition, *23*(1), 54-70.
- Markman, E. M. (1979). Realizing that you don't understand: Elementary school children's awareness of inconsistencies. Child Development, *50*, 643-655.
- Markman, E. M. (1989). Categorization and naming in children. Cambridge, MA: MIT Press.
- Markman, E. M. & Hutchinson, J. E. (1984). Children's sensitivity to constraints on word meaning: Taxonomic versus thematic relations. Cognitive Psychology, *16*, 1-27.
- Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. Psychological Review, *100*, 254-278.
- Medin, D. L. & Heit, E. J. (in press). Categorization. In D. Rumelhart & B. Martin (Eds.), Handbook of cognition and perception. Hillsdale, NJ: Erlbaum.
- Murphy, G. L. & Allopenna, P. D. (1994). The locus of knowledge effects in concept learning. Journal of Experimental Psychology: Learning, Memory and Cognition, *20*, 904-919.
- Murphy, G. L. & Medin, D. L. (1985). The role of theories in conceptual coherence. Psychological Review, *92*, 289-316.
- Slooman, S. A. (1996). The empirical case for two systems of reasoning. Psychological Bulletin, *119* (1), 3-22.
- Tversky, A. (1977). Features of similarity. Psychological Review, *84*, 327-352.
- Slooman, S. A. & Rips, L. J. (1998). Similarity as an explanatory construct. Cognition, *65*, 87-101.
- Vygotsky, L. S. (1978). Mind in society: The development of higher psychological processes. M. Cole, V. John-Steiner, S. Scribner, & E. Souberman (Eds.), Cambridge, MA: Harvard University Press.
- Waxman, S. R. & Gelman, R. (1986). Preschoolers' use of superordinate relations in classification and language. Cognitive Development, *1*, 139-156.
- Wisniewski, E. J. & Bassok, M. (1996). On putting milk in coffee: The effect of thematic relations on similarity judgments. Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society. (pp. 464-468). Hillsdale, NJ: Erlbaum.