Available online at www.sciencedirect.com
ScienceDirect
Cognitive Psychology 55 (2007) 196–231
Cognitive Psychology
www.elsevier.com/locate/cogpsych

Beyond common features: The role of roles in determining similarity ☆

Matt Jones *, Bradley C. Love *

The University of Texas, Department of Psychology, 1 University Station, A8000, Austin, TX 78712, USA

Accepted 6 September 2006
Available online 13 November 2006

Abstract

Historically, accounts of object representation and perceived similarity have focused on intrinsic features. Although more recent accounts have explored how objects, scenes, and situations containing common relational structures come to be perceived as similar, less is known about how the perceived similarity of parts or objects embedded within these relational systems is affected. The current studies test the hypothesis that objects situated in common relational systems come to be perceived as more similar. Similarity increases most for objects playing the same role within a relation (e.g., predator), but also increases for objects playing different roles within the same relation (e.g., the predator or prey role in the hunts relation) regardless of whether the objects participate in the same instance of the relation. This pattern of results can be captured by extending existing models that extract meaning from text corpora so that they are sensitive to the verb-specific thematic roles that objects fill. Alternative explanations based on analogical and inferential processes are also considered, as well as the implications of the current findings to research in language processing, decision making, and category learning.

Keywords: Similarity; Relational similarity; Thematic similarity; Verb-specific thematic role; Language processing; Category learning; Role-governed categories; Latent semantic analysis; Incidental learning

☆ This work was supported by National Research Service Award F32-MH068965 from the National Institutes of Health to Matt Jones, and by Air Force Office of Scientific Research Grant FA9550-04-1-0226 and National Science Foundation CAREER Grant 0349101 to Bradley C. Love. The authors thank Art Markman, Ken McRae, and Bill Swann for helpful feedback on this research, and Brian Pomeroy, Momoko Sakamoto, and Elliot Saperstein for their assistance with stimulus design and data collection.
* Corresponding authors. Fax: +1 512 471 5935.
E-mail addresses: mattj@psy.utexas.edu (M. Jones), love@psy.utexas.edu (B.C. Love).

0010-0285/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.cogpsych.2006.09.004

1. Introduction

In nascent days of the Internet, the term "information superhighway" was used to help the uninitiated understand what the Internet was and what it did. Although the Internet and an actual highway system share few features (e.g., highways are made of concrete whereas the Internet is made of fiber cables), highlighting that the Internet and a highway both play the role of conduit allowed the general populace to appreciate the similarities between the two. More correctly, this metaphor or analogy highlighted the similarity of one transportation system, consisting of fiber cables and information, to another, consisting of highways and automobiles.

Research in analogical comparison has shown that people find systems to be more similar when they share common relational structure, such as in the above example (Gentner, 1983). However, an open question is whether the constituent objects embedded in these corresponding structures become more similar as well. That is, does similarity depend on relational information that is extrinsic to the items being compared? For example, can playing the same role (e.g., conduit) in different instances of a relational system make the two corresponding objects (fiber cables and highways) more similar? Furthermore, can such effects occur even in the absence of an overt metaphor or analogy? An additional question is whether other relational commonalities, beyond role matches, can affect perceived similarity as well. For example, do bread and butter become more similar over time because they participate in a common relation, albeit in different roles? If so, do the objects have to directly interact with one another or can perceived similarity increase simply by occurring in different instances of the same relational system (as with bread and icing, or butter and cake)? These are basic questions of critical importance to research in analogy, categorization, and semantic structure, yet have not been directly examined experimentally.

The question of whether concepts are represented (in part) in terms of the relational systems in which they participate also has important implications for language processing, in both comprehension and production. For example, a critical step in sentence comprehension is assignment of nouns or noun phrases to the thematic roles of verbs. This process is facilitated by semantic knowledge of which items provide better fits for each verb-specific role (McRae, Spivey-Knowlton, & Tanenhaus, 1998). In addition, studies of structural priming show that generation of syntactic structure during sentence production is guided by the thematic roles that items fill, beyond the influences of intrinsic properties such as animacy (Chang, Bock, & Goldberg, 2003). Knowledge of the verb-specific thematic roles that items tend to fill is closely related to knowledge of the relations in which they participate. Therefore a better understanding of the contribution of extrinsic relational information to semantic representations may help further our understanding of the interplay between semantic knowledge and processing of syntactic structure.

This article investigates the role of extrinsic relational information in semantic knowledge, by directly testing the ways in which relational commonalities can influence the perceived similarity of constituent objects. A computational model is forwarded to explain our empirical results and a number of alternative explanations are also considered. Finally, the connections to work in language, decision making, and categorization are explored.

2. Similarity and relational information

Early research on similarity assumed that similarity could be represented as a function of distance within a psychological stimulus space (Shepard, 1987). Research by Tversky (1977) showed this view to be insufficient, as the metric axioms of a distance-based representation are violated in many cases. Subsequent work has shown that the spatial representation is often adequate, provided it is augmented with a selective attention mechanism that effectively stretches the psychological space along attended dimensions (Nosofsky, 1986). Despite the variation in these theoretical approaches, they all share the critical assumption that objects are represented in terms of intrinsic features, and that similarity is determined solely by commonalities or overlap in these features.

More recently, researchers have looked at the effect of relations on similarity, but only of relations within the structures being compared (Gentner, Rattermann, & Forbus, 1993; Goldstone, Medin, & Gentner, 1991; Markman & Gentner, 1996; Medin, Goldstone, & Gentner, 1990). For example, Medin et al. (1990) found that the similarity between drawings of butterflies depended on the match in certain relations among each butterfly's parts (e.g., the relative size of the two wings). Although this study examined how relations affected the perceived similarity of the complete systems (i.e., the butterflies), there was no investigation of the effect of relations on the similarity between the individual components (the wings), nor of any possible effect on the similarity among butterflies due to interaction between them and other objects (e.g., plants or birds). Thus, although more contemporary work on perceived similarity considers a richer set of comparison processes to align representations, as opposed to relying on pre-established alignments of matching features or dimensions as in earlier accounts, these richer alignment accounts only consider the impact of intrinsic relations. In this regard, existing relational accounts of similarity have much in common with previous accounts of similarity that focus on intrinsic features. In both cases it is assumed that the similarity of two items can be determined by considering each in isolation, without regard for their relationships to other objects with a larger system.

The focus of the present study is on relational information that is extrinsic to the items being compared. Thus, we consider interactions and other relationships between objects or events that exist as separate entities, such as *the butterfly drinks nectar*. This type of relational information can be distinguished from relationships between categories themselves such as taxonomic relations (e.g., *lizards are reptiles*), componential relations (e.g., *the bicycle has wheels*), and relations that are internal to the object in question (e.g., *the house's walls hold up its roof*; this is an extrinsic property of *wall* and *roof*, but not of *house*). Examples of the types of relations we consider are spatial (*the computer is on the desk*), dynamic (*the cat stalks the mouse*), and causal (*cars create pollution*). Our main hypothesis is that objects and categories are often represented (in part) in terms of these extrinsic relational properties, and furthermore that such information plays a part in determining similarity, just as intrinsic featural or relational information does.

The notion of concepts being defined in terms of their relationships to other concepts has precedent in the philosophy of language, where it is known as conceptual role semantics (Block, 1987; Field, 1977). Conceptual role semantics is an outgrowth of use theories of meaning (e.g., Wittgenstein, 1953) that allows for use to include counterfactual or imagined actions in addition to communicative acts. Conceptual role semantics emphasizes the role of symbols within computational systems. Actual computational models have been built in which conceptual units derive meaning from their role within the larger system (e.g., Goldstone & Rogosky, 2002; Larkey & Love, 2003). Conceptual role semantics offers a response to the view that concepts must be innately specified because their definitions cannot be learned (e.g., Fodor, 1998). For instance, instead of having to learn the definition of a word, conceptual role semantics holds that one can learn a word by learning how to use it.

Conceptual role semantics posits that the meaning of a concept is determined by its functional role within a conceptual system (or semantic network). However, this notion of role includes not only relational information as presently defined, but also taxonomic and componential relations, which are often incorporated in existing models of semantic knowledge (e.g., Collins & Quillian, 1969). Here, we specifically focus on interaction-based associations, that is, those based on tokens of the concepts rather than on the concepts themselves (as is the case with taxonomic relations), and more specifically those based on extrinsic relations (e.g., *hammers hit nails*) rather than featural or componential relations (e.g., *birds have wings*).

The distinction we make here between features and internal relations on the one hand, and extrinsic relations on the other, is essentially the same as that between intrinsic and extrinsic properties as formulated by Barr and Caplan (1987). These researchers showed, using a feature-listing paradigm, that many natural categories are partially determined by extrinsic properties, and furthermore that the relative numerosity of such properties has important implications for category structure. The proposal that objects can be represented in terms of relational information is also related to Markman and Stilwell's (2001) theory of role-governed categories (see also Gentner & Kurtz, 2005; Markman, 2005). Markman and Stilwell note that many categories that cannot be easily defined in terms of intrinsic features, such as *thief* or *game* (cf. Wittgenstein, 1953) have simple definitions in terms of their members' roles in certain relational systems, such as *steals* (x, y, z) or *plays* (x, y). They further argue that role-governed categories are prevalent in our conceptual systems and that such categories are psychologically represented in terms of relational information. Representation of concepts in terms of their relationships to other concepts also captures many of the core principles of the theory-based approach to conceptual coherence, which posits that conceptual representations are largely determined by the naive theories that bind them together (Murphy & Medin, 1985).

Further evidence for the significance of relational properties comes from the feature norms collected by McRae, Cree, Seidenberg, and McNorgan (2005). To test the prevalence of extrinsic relational properties in these norms, one of the authors (MJ) and a naive coder separately categorized all 6485 properties listed for 541 categories (after excluding concept-level relations such as taxonomy and synonyms) into intrinsic and extrinsic (relational) properties. The result was that 39.9% of these properties were found to be extrinsic (relational) properties (agreement rate = 99.0%).

When considered as a property of a constituent object in a relational structure, relational information can be construed in a number of ways. Here, we describe three distinct ways in which constituent objects can relationally match, each of which could potentially affect their perceived similarity. The experiments presented below were designed to test the psychological reality of each of these three possibilities.

2.1. Roles

One possibility is that an object participating in a relation could be represented in terms of the role it plays in that relational system. For example, a hammer could have the property of playing the agent role in the *hit* relation. Thus, objects sharing such relation-specific roles could be perceived as more similar. For example, one reason *hammer* and *baseball bat* are similar could be that they both typically hit other objects (nails and baseballs). This notion of relational information closely corresponds to Markman and Stilwell's (2001) role-governed categories as well as to the majority of the extrinsic properties considered by Barr and Caplan (1987). Other evidence for the psychological significance of role-based properties comes from experiments demonstrating role-role priming (e.g., TORCH-CANDLE; Flores d'Arcais, Schreuder, & Glazenborg, 1985; Schreuder, Flores d'Arcais, & Glazenborg, 1984). However, there are difficulties interpreting these results due to methodological concerns, specifically confounding by featural commonalities (Pecher, Zeelenberg, & Raaijmakers, 1998; see also Hutchison, 2003, pp. 799–801). Thus the question of role-based priming is still unresolved. Here, we explore how role matches influence perceived similarity, using a counterbalanced experimental design that controls for featural commonalities.

2.2. Relations

There is the more general possibility that objects participating in common relations are perceived as more similar irrespective of the roles they play in those relations. For example, *hammer* and *baseball* might be similar due to their common association to the relation *hit*. This is the prediction made by certain models that extract word meaning based on occurrence statistics, including Latent Semantic Analysis (LSA; Landauer & Dumais, 1997) and Hyperspace Analogue to Language (HAL; Burgess, 1998; Lund & Burgess, 1996; see also McDonald & Lowe, 1998). These models work by tracking the contexts in which words appear and rating words with similar contexts as being semantically similar. Thus in particular if two nouns (denoting objects or categories) tend to co-occur with the same verb (denoting a relation), then they are predicted to be perceived as similar. (Further details of both models, as well as their predictions regarding role-based similarity, are discussed below in the modeling portion of this article.)

2.3. Scenarios

Finally, it is possible that objects become associated to the particular instances of relational systems in which they appear. This would cause direct interaction between two objects to make those objects similar to each other. For example, reading the sentence "The hammer hit the nail," which describes an instance of the relation *hit*, might lead to an increase in the similarity between *hammer* and *nail*. Such an effect might best be thought of as due not to commonalities between the objects but rather to direct associations. This type of common-scenario similarity effect, often termed thematic similarity, has been demonstrated (in a correlative design) by Wisniewski and Bassok (1999), who suggest that similarity can be affected by integrative processes (which produce direct associations) as well as by comparison processes (which lead to recognition of commonalities). Integration-based processes differ from comparison-based processes in that two objects are linked through a common relation rather than aligned. For example, Wisniewski and Bassok found that peanut butter is rated as similar to a knife, presumably not because of any common features but because of the thematic association derived from the use of the one to spread the other. Similarly, Lin and Murphy (2001) found that subjects instructed to group objects that "go together to form a category" were at least as likely to group thematically as taxonomically related items. For example, *chalk* was grouped with *blackboard* more often than with *marker* in a triad sorting task. In related work, Moss, Ostrin, Tyler, and Marslen-Wilson (1995) showed priming in a lexical decision task for instrument relations (e.g., hammer–nail) even in the absence of prior association (as measured by free association). A similar effect was found by Prior and Bentin (2003), who showed that co-occurrence of two words within a sentence facilitated learning of an association between those words in a subsequent paired-associates task.

2.4. Summary

Previous work in similarity has focused on the impact of features and relations intrinsic to objects and situations. For instance, work at the intersection of analogy and similarity has demonstrated that two situations containing corresponding relational structure are perceived as similar. In contrast, the effect of relations that are extrinsic to the items being compared has not been as extensively examined. We identify three ways in which objects can share extrinsic relational properties. These three forms of match are common role, relation, and scenario. Relation and scenario matches are integration-based in that they involve linking two objects through a common relation. Scenario matches are more specific than relation matches in that the objects appear in the same instance of the relation. Role matches can be viewed as comparison-based as they involve determining that two objects play analogous roles within different instances of the same relation. The following behavioral studies experimentally manipulate these types of match and evaluate their relative contributions to perceived similarity.

3. Experiment 1: Testing the contribution of relational information to similarity

Experiment 1 directly tests the influence of the three types of relational commonality (i.e., common roles, relations, and scenarios) on perceived similarity. The procedure for this and all following experiments involved an incidental learning task in which participants read about relations between common items. Each relation was named by a verb (e.g., *chase*) with two thematic roles-agent and patient. Each scenario was thus described by a single sentence in which one item filled each role (e.g., *The polar bear chases the seal*). Following the incidental task, participants were presented with a series of similarity comparisons among the same items, to evaluate how the relational information presented affected their perceived similarity. Each similarity comparison involved one base and two target objects; the participant's task was to select which target was most similar to the base.

The incidental task was structured such that objects varied between participants in terms of the relations in which they participated, their roles within those relations, and which objects they co-occurred with (i.e., in the same scenario). Pairs of objects differed in three principal ways in terms of the relational properties they had in common. First, two objects could participate in the same relation or in different relations. Second, objects participating in the same relation could play the same or opposite roles (agent vs. patient) within that relation. Third, objects participating in the same relation could occur in the same or different instances of that relation (i.e., they could directly interact or they could not).

Comparing the similarities of pairs of objects with different levels of relational commonality provides tests of the influence of different types of relational information on similarity. Four types of comparisons were used, as illustrated in Table 1. In each case, the two targets differed in the relational properties they shared with the base, so that one target (the primary target) was predicted to be selected more often than the other (alternative target). In comparisons evaluating the effect of common role, the targets differed in their role match to the base; that is, all three appeared with the same verb but the primary target filled the same thematic role as the base (agent or patient) whereas the alternative target did not. In tests of the effect of common relation, the targets differed in their relation match with the base; the primary target appeared in the same relation (but opposite role) as the base, and the alternative target appeared in a different relation. Comparisons evaluating the effect of common scenario were similar to those for common relation except that the primary target appeared in the same scenario as the base. In all three of these analyses, the targets were counterbalanced between two groups of subjects by varying the information in the incidental task, so that the primary target for one group was the alternative target for the other. Therefore the response rates of the two groups could be directly compared, and any difference ascribed to the manipulation of relational information rather than variation in baseline similarity prior to the experiment. In particular, this allowed us to control for effects of intrinsic (e.g., featural) properties, a critical confound in previous research (e.g., Wisniewski & Bassok, 1999). The final test, role vs. relation, directly contrasted the effects of common role and common relation. The items in these comparisons were structured as in the common role test for one group of participants and as in the common relation test for another. Comparing the two groups' strength of preference for the primary target provided a comparison between the boost in similarity from playing common roles (i.e., going from matching relation to matching relation and role) and that from appearing in the same relation (i.e., going from no match to matching relation). Again, comparing responses of two groups on the same similarity comparison ensures that any effects found are due to the manipulation of relational information.

Table 1
Logical structure of tests used in experiments

Test	Commonalities with base					
	Primary target			Alternative target		
	Relation	Role	Instance	Relation	Role	Instance
Common role	Same	Same	Different	Same	Different	Different
Common relation	Same	Different	Different	Different	—	—
Common scenario	Same	Different	Same	Different	—	—
Role vs. relation						
Role group	Same	Same	Different	Same	Different	Different
Relation group	Same	Different	Different	Different	—	—

Notes: Each test assesses the contribution of a specific type of relational information to similarity. Role refers to relation-specific role (e.g., chase-agent vs. chase-patient). Instance refers to specific occurrence of the relational system (i.e., whether the objects directly interacted). For each test it was predicted that participants would tend to select the primary target as more similar to the base than the alternative target. The role vs. relation test determines which of the common-role and common-relation effects is stronger, by comparing the strength of preference for the primary target between two groups.

3.1. Method

3.1.1. Participants

Participants were 40 undergraduate students at the University of Texas, Austin, who received partial credit in an introductory psychology course.

3.1.2. Stimuli

Stimuli in the first phase of the experiment were brief subject–verb–object (agent–verb–patient) sentences, using active verbs that describe physical events. Sentences were presented visually on a computer screen. These sentences described interactions among common objects (including people and animals), such as *The polar bear chases the seal*. These interactions were chosen to represent presumably familiar scenarios that are typical of the objects involved. In all cases the reversed interaction would be unrealistic (e.g., seals do not chase polar bears), so that participants would be unlikely to misremember or misencode the items' roles (Ferretti, McRae, & Hatherell, 2001). Stimuli in the second phase of the experiment were the same object names that had appeared in the first phase, presented in sets of three for similarity comparisons.

3.1.3. Design

The experiment involved four separate word groups. Each group consisted of seven nouns (objects) and two verbs (relations). One of these groups, along with the relevant

Table 2
Illustration of experimental design

Condition	Sentences (incidental learning)
1	The polar bear chases the seal
	The collie chases the cat
	The German shepherd herds the sheep
2	The polar bear chases the seal
	The German shepherd chases the cat
	The collie herds the sheep
3	The seal chases the fish
	The collie chases the cat
	The German shepherd herds the sheep
4	The seal chases the fish
	The German shepherd chases the cat
	The collie herds the sheep

Similarity comparisons

Base	Target 1	Target 2
Seal	Collie	German shepherd
Seal	Cat	Collie
Seal	Cat	German shepherd
Cat	Collie	German shepherd
Sheep	Collie	German shepherd

Notes: The full design consisted of four copies of the above structure, with four distinct word sets. The condition to which a participant was assigned for a given word set determined the sentences (s)he viewed, but all participants answered the same similarity comparisons. Thus each participant viewed 12 sentences (plus 4 catch trials) during incidental learning, and later answered 20 similarity comparisons. See the Appendix for words from the other three sets.

portion of the design, is shown in Table 2; the other three groups are listed in Appendix A. In the incidental learning phase, the words in each group were combined to create three sentences, with six of the nouns and one of the verbs appearing exactly once, and the other verb appearing twice. There were four ways in which these sentences, as a group, could be formed, corresponding to four between-subjects conditions for each word set. Assignment of participants to conditions in each word set was done by a Latin square, so that each participant was in each condition for exactly one of the four word sets. Thus, to determine the sentences for a given participant, first (s)he was assigned to a condition for each word set (e.g., condition 3 for set A, 4 for B, 1 for C, and 2 for D), and then the sentences corresponding to each set-condition pair were generated.

The similarity comparisons in the second phase of the experiment were the same for all participants. These consisted of 20 triples of items, five from each word set. (See Table 2 for examples.)

3.1.4. Procedure

At the start of the experiment, participants were informed that they were participating in an investigation of how they process simple sentences. They were told that they would view a series of sentences and that they were to envision each before answering a question about it. They were then presented with two blocks of the same 16 sentences, presented in a different random order each time (except for catch trials; see below). For the first block, they were asked to rate how realistic each sentence was; in the second block, they were asked to rate how interesting each sentence was. Ratings were given by pressing a key from "1" to "7" (with the direction of the scale randomized for each participant, but consistent between the two blocks). Each trial began with the instruction "Visualize the scene described in the following sentence" presented for 2 s. Following 1 s of blank screen, the sentence was presented for 2 s before the response prompt was presented ("How realistic is this scene?" or "How interesting is this scene?"). The sentence remained visible until the participant entered a response. The inter-trial interval was 1 s of blank screen.

Twelve of the 16 sentences were those determined by the experimental design, as described above. The other four sentences were used for catch trials to help ensure participants' active participation. These catch trials occurred in pairs, on trials 1, 2, 7, and 8 of the first block and 3, 4, 10, and 11 of the second block. Following the second trial in each pair, a "pop quiz" was given that queried one of the items from the previous sentence (e.g., *What broke the window?*). Participants selected their responses from among the correct item and three semantically similar lures (e.g., *stone*, *ball*, *brick*, *rock*). If incorrect, the participant heard a loud screeching sound and saw the message "Hey, wake up!" Objects and relations used in the catch-trial sentences were chosen to be dissimilar from all the items used in the experimental sentences.

For the second phase of the experiment, participants made 20 similarity comparisons among the items from the experimental sentences of the first phase. Each comparison comprised one base object and two target objects, with the task being to decide which target was more similar to the base. Each trial began with presentation of the question *Which is more similar to X?* where *X* represents the name of the base object. The base was presented in the center of the screen near the top, with the beginning of the question directly above it. After a 1 s delay, the target words were presented below the base at either side of the screen, along with response instructions to press the 'Q' key (which is on the left side of the keyboard) to select the item on the left and to press the 'P' key (which is on the right side of the keyboard) to select the item on the right. Left–right locations of the two targets were randomly assigned on each trial. After the participant made a response there was 1 s of blank screen before the next trial began.

Participants were seated approximately 24 in. from a 15-in. monitor. Primary stimuli were all presented in 30-point (5/12-in.) font, with instructions and response prompts in 24-point (1/3-in.) font. All text was white on a black background. All responses were made via a standard keyboard. The entire experiment lasted between 15 and 25 min.

3.2. Results

Preliminary analysis showed that all but three participants gave perfect responses on the four catch trials; of these three, two missed one question and the third missed two. Thus, it was decided to include all 40 participants in the analysis of similarity responses. (Ratings from the incidental tasks were not analyzed.)

The general strategy in analyzing the similarity data was to contrast response rates to specific comparisons between participants in different conditions. Because of the counterbalancing in the design, any reliable difference between two groups must be due to the differences in the interactions processed during the first phase of the experiment (as opposed to pre-existing knowledge), and in particular to the differences in relational information presented. The contrasts that were analyzed followed the logical structure of the four tests described above and summarized in Table 1. These contrasts were designed to evaluate the

Table 3
Illustration of contrasts

Test	Scenarios		Comparison		
	Group A	Group B	Base	Target 1	Target 2
Common role	The polar bear chases the *seal*	The *seal* chases the fish	Seal	Cat	Collie
	The *collie* chases the ***cat***	The ***collie*** chases the *cat*			
Common relation	The polar bear chases the *seal*	The polar bear chases the *seal*	Seal	Collie	German shepherd
	The ***collie*** chases the cat	The ***German shepherd*** chases the cat			
	The *German shepherd* herds the sheep	The *collie* herds the sheep			
Role vs. relation	The polar bear chases the *seal*	The *seal* chases the fish	Seal	Cat	Collie
	The *collie* chases the ***cat***	The German shepherd chases the ***cat***			
	The German shepherd herds the sheep	The *collie* herds the sheep			
Common scenario	The ***collie*** herds the *sheep*	The ***German shepherd*** herds the *sheep*	Sheep	Collie	German shepherd
	The *German shepherd* chases the cat	The *collie* chases the cat			

Notes: Items in bold indicate predicted responses for each group. Participants in group A for the role, relation, and scenario contrasts were predicted to choose the response listed as target 1, whereas group B participants would tend to choose target 2. For the role vs. relation contrast, participants in both groups were expected to choose target 1 most often. The hypothesis that common role has a stronger effect than common relation predicts this tendency to be greater for group A than group B; the alternative hypothesis makes the opposite prediction.

potential effects of the three types of relational information—common roles, relations, and scenarios—on similarity. Examples of the contrasts and their relation to the design of the implicit learning phase of the experiment are given in Table 3. To foreshadow the results, all three types of relational commonality increased the perceived similarity of matching objects. The effect size for role matches was greatest, whereas relation and scenario matches had roughly equal effects on perceived similarity. The overall results are shown in Fig. 1.

3.2.1. Common-role test

This analysis tested whether matching role increases the similarity between objects, beyond any effect of participating in the same relation. The analysis was based on similarity comparisons in which both targets were involved in the same relation as the base, but only one target matched the base's role (agent or patient within that relation). Response rates on each comparison were contrasted between two groups. Both groups had read a sentence in which target 1 was the patient and target 2 was the agent. Each group had also read a sentence in which the base appeared with the same verb, but the role that the base played differed between the two groups. An example is shown in Table 3. The prediction was that participants in each group would tend to select the target that had appeared in the same role as the base.

In the full design there were eight contrasts of this type, two from each word set. All eight showed the same effect, with participants tending to choose the target that matched the base's role as predicted. In other words, common role had a positive influence on similarity in every case. Overall, 69.4% of responses (111/160) conformed to this pattern. To evaluate the statistical significance of this effect, a logistic regression was performed on the data from all eight contrasts, with role as the focal factor and comparison included as a categorical variable to allow for varying base response rates to the eight comparisons (due to biases in similarity prior to the experiment). Each participant contributed four observations to this model, but these were always based on responses to unrelated sets of words and thus were assumed to be independent (e.g., a participant's baseline similarity between *seal* and *collie* was assumed to be independent of his or her baseline between *robot* and *carpenter*). The regression revealed a significant effect of role, $\chi^2(1) = 27.68$, $p < 10^{-6}$, indicating that similarity was reliably increased by common roles. The interaction of role with comparison was non-significant, $\chi^2(7) = 3.09$, $p > .5$, suggesting a uniform effect of common role regardless of the particular stimuli involved.

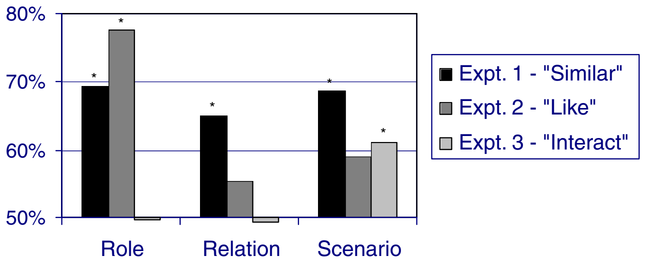

Fig. 1. Summary of results from all three experiments. Horizontal axis categories refer to common-role, common-relation, and common-scenario tests, respectively. Vertical axis gives magnitude of effect found in each experiment (chance is 50%). Entries marked with "*" are significant at $p < .05$.

3.2.2. Common-relation test

This analysis tested whether participation in matching relations alone is sufficient to influence similarity. The analysis was based on comparisons in which one target participated in the same relation as the base, but in the opposite role, and the other target was involved in a different relation (all three appeared in different scenarios). See Table 3 for an example. The two targets were counterbalanced between two groups of participants, whose response rates were contrasted. The prediction was that participants in each group would tend to select the target they had seen in the same relation as the base.

Overall there were four contrasts of this type, one from each word set. All four showed a positive effect of common relation on similarity, with a combined response rate of 65.0% (52/80) in favor of the relation-matched target. As before, a logistic regression model was fit to the combined data to test the significance of this effect. This model included two responses from each participant, but these could be assumed to be independent for the same reason stated above. The regression revealed a significant effect of relation, $\chi^2(1) = 7.91$, $p < .01$, indicating that similarity was reliably increased by common relations. The interaction of relation with comparison was non-significant, $\chi^2(3) = .80$, $p > .5$, suggesting a uniform effect of common relation regardless of the particular stimuli involved.

3.2.3. Role vs. relation test

This analysis directly compared the effects of common role and common relation. The analysis was based on comparisons in which one target always had the advantage of greater commonality with the base, but the nature of this advantage differed between two groups of participants. See Table 3 for an example. For one group, the advantaged target matched the base in both relation and role, whereas the other target matched in relation only (as in the common-role test). For the other group, the advantaged target matched the base's relation but not its role, and the other target mismatched entirely (as in the common-relation test). Thus the relative strength of the common-role and common-relation effects could be evaluated by determining which group responded more often with the advantaged item.

In the full design there were eight contrasts of this type, seven of which favored the effect of role over that of relation. Overall, 55.0% (45/80) of responses chose the advantaged target when the targets differed in relation match to the base, and 78.8% (63/80) chose the advantaged target when the two differed in role. The significance of this difference was tested with a logistic regression model having Information Type (role or relation) as one predictor and comparison as the other. Each participant contributed four observations to this model, which were assumed to be independent for the same reason as above. The model showed a significant main effect of Information Type, $\chi^2(1) = 21.04$, $p < 10^{-5}$, and a non-significant interaction, $\chi^2(7) = 6.10$, $p > .5$. Thus, the effect of common role was stronger than that of common relation, to a degree that appears to be independent of the stimuli used.

3.2.4. Common-scenario test

This analysis tested whether direct interaction (i.e., appearance in the same scenario) influences similarity. The analysis was based on comparisons in which one target appeared in the same scenario as the base (thus in the same relation but in the opposite role), and the other target appeared in a different scenario involving a different relation. See Table 3 for an example. The targets were counterbalanced between two groups of participants, whose response rates were contrasted. The prediction was that participants in each group would tend to select the target they had seen in the same scenario as the base.

The full design provided eight contrasts of this type. All eight contrasts were in the direction of a positive effect of co-occurrence on similarity. Overall, 68.8% of responses (220/320) followed this pattern. The significance of this effect was tested using logistic regression as before, with co-occurrence with the base and comparison as predictors. Each participant contributed four observations to this model, which could be assumed to be independent as before.[1] The model revealed a significant effect of co-occurrence, $\chi^2(1) = 24.45$, $p < 10^{-6}$. The interaction with comparison was non-significant, $\chi^2(7) = 4.06$, $p > .5$. Thus, results indicated a positive effect of direct interaction on similarity that was independent of the particular stimuli used.

3.2.5. Comparing common-relation and common-scenario effects

Because the information contributing to the co-occurrence effect includes the information contributing to the common-relation effect, it is important to ask whether the former effect is stronger. That is, there is the question of whether there is an effect of interaction per se, or whether the effect found above was simply a by-product of the common-relation effect.

The present design does not allow for a direct comparison of the magnitudes of these two effects, because their estimates are based on response rates to different sets of similarity comparisons. However, such a contrast is provisionally justified by the above results suggesting that both effects are independent of the stimuli used (i.e., by the null interaction effects found in both analyses). Therefore a logistic regression was performed combining data from the common-relation and co-occurrence analyses.[2] The effect of interest in this model was the interaction between Observation Type (i.e., whether the observation was taken from the relation test or the co-occurrence test) and relation/co-occurrence (the concatenation of the variables from the two separate analyses). This interaction was non-significant, $\chi^2(1) = .02$, $p > .5$. Thus, there is no evidence that the effect of co-occurrence was stronger than that of common relation.

3.3. Discussion

Experiment 1 demonstrates that one determinant of the perceived similarity between objects is the relational systems in which they participate. In all of the analyses described above, participants reliably tended to select the target object that shared relational properties with the base object. The experimental design allowed separate demonstrations of two different effects of this sort. First, similarity is increased when two entities are involved

[1] There were actually eight observations per participant relevant to this analysis, but only four were included in the model in order to satisfy the independence assumption. See Appendix A for details on how the data were randomly split.

[2] Half of the observations contributing to the co-occurrence analysis were potentially correlated with observations contributing to the common-relation analysis, in that the same participant gave two responses to overlapping sets of stimuli (one response contributing to each analysis). Therefore that half of the co-occurrence data was removed from the combined analysis. Importantly, though, the number of remaining co-occurrence observations equaled the total number of common-relation observations (as originally the former were twice as numerous). Furthermore, an analysis using all of the data yielded the same conclusion as the one presented here.

in the same type of relation (e.g., chasing), even when they appear in different events, and even when they play different roles within this relation (e.g., one is chasing and the other is being chased). Thus participation in common relations increases similarity. Second, similarity is further enhanced if both entities play the same role within a common relation, for example, when both chase other things or both are chased by other things. Thus correspondence of roles within a relation produces an additional increase in similarity. Analyses that directly contrasted these two effects showed that the common-role effect is stronger. Thus the structural sensitivity implied by this effect is quite significant and central to participants' object representations and similarity processes.

The third hypothesized effect, that of direct interaction, was found as well. This result provides experimental validation of the thematic similarity findings of Wisniewski and Bassok (1999). However, the effect was not significantly stronger than that of common relation. That is, entities appearing in the same event (e.g., one chasing the other) are not reliably more similar than entities appearing in different events involving the same relation. Because objects that co-occur within a given scenario are necessarily participating in the same relation, the present data suggest that this is what drives the co-occurrence effect. That is, there may be no effect of direct interaction per se, beyond that due to the common-relation effect. Thus previous evidence of thematic similarity might also be due to a common-relation effect, rather than to direct associations. Of course further research is warranted to answer this question more definitively. For now, it is also worth noting that the absence of a significantly stronger effect for cases of co-occurrence as opposed to common relations provides validation of the experimental paradigm used here. Specifically, it suggests that participants integrated information from the various scenarios presented, such that this information impacted their representations of the objects involved, rather than responding based directly on the surface properties of individual sentences (since the latter possibility would imply a much greater effect for the co-occurrence test).

4. Experiment 2

The design of Experiment 2 largely mirrors that of Experiment 1 with one key change in task instructions. The change in instructions is intended to test the robustness of Experiment 1's results and to explore possible contributions of multiple processes to judged similarity. Specifically, the term "similar" was replaced with the word "like." Thus in the second phase of Experiment 2 participants were asked to report which of the two target objects was "more like" the base object.

On first consideration, it might seem surprising to expect behavioral differences based on such a slight change of instruction. However, previous research indicates that thematic similarity is more evident under similarity rating task instructions than in closely related tasks. Therefore, one prediction is that the effects of thematic similarity (i.e., common relation and common scenario) will be attenuated in Experiment 2. In support of this claim, Markman and Hutchinson (1984) showed, in the same triad task as used here, that children aged two to five tend to choose a target that is thematically related to the base when asked to choose based on similarity, but will choose the featurally related target when asked to extend a novel verbal label that had been applied to the base. Gentner and Brem (1999) demonstrated an equivalent result with adults. In both studies, the relative influence of integration-based and comparison-based processes on perceived similarity varied through subtle changes in task instructions. If "like" encourages more of a comparison frame than "similar", we would expect to observe smaller effects for thematic (i.e., relation or scenario) concordance in Experiment 2 than in Experiment 1.

A second prediction is that the effect of role matches should increase under the "like" instructions. Although there is no preexisting research directly comparing "similar" and "like" task instructions, work in metaphor and simile processing does bear on this issue. Metaphors (*X is a Y*) and similes (*X is like a Y*) are both figurative statements, but metaphors are linguistically identical to literal categorization statements whereas similes are identical to literal comparison statements. Bowdle and Gentner (2005) found that this difference is psychologically relevant, with simile constructions better suited to nonconventional figurative statements (e.g., *The mind is [like] a garden*), and metaphors better suited to conventional figurative statements (e.g., *An opportunity is* [*like*] *a doorway*). Likewise, Gregory and Mergler (1990) found that similes are more successful at highlighting non-obvious similarities between *X* and *Y*. Bowdle and Gentner (2005) characterize the "is like a" simile construction as principally engaging analogical-comparison processes. In terms of the three types of relational match we explore in Experiment 2 (roles, relations, and scenarios), the aforementioned results suggest that "like" should serve to increase the effect of role matches as these matching objects play analogous roles within their respective instances of the relation.

In summary, Experiment 2 closely mirrors Experiment 1 and was predicted to largely replicate the basic results from that experiment. However, previous research suggests that the change to "like" instructions may serve to decrease the effect of thematic (i.e., relation and scenario) matches while increasing the effect of role matches, by encouraging more of a comparison or analogical frame as opposed to an integration frame.

4.1. Method

4.1.1. Participants

Participants were 28 undergraduate students at the University of Texas, Austin, who received partial credit in an introductory psychology course.

4.1.2. Stimuli, design, and procedure

Experiment 2 was identical to Experiment 1 except for the question presented on each trial of the similarity comparison phase. The wording used was "Which is more like ⟨*base object*⟩?"

4.2. Results

Data were analyzed identically to those of Experiment 1 (see above for details). Inspection of catch-trial responses showed that all but one of the participants performed perfectly on the four catch-trial questions; the final participant gave one incorrect response. Therefore all participants were included in later analyses. The overall results, along with comparisons to results from the other experiments, are given in Fig. 1.

The four contrasts composing the common-role test showed a combined rate of 77.7% (87/112) of responses in favor of the role-matched target. Logistic regression showed this effect to be significant, $\chi^2(1) = 42.86$, $p < 10^{-10}$. The interaction with comparison was nonsignificant, $\chi^2(7) = 9.91$, $p > .1$. In the data composing the common-relation test, 55.4% (31/56) of responses favored the relation-matched target. Both this effect and the interaction with comparison were non-significant ($\chi^2[1] = .78$ and $\chi^2[3] = 1.73$, respectively, both $ps > .5$). The test of role vs. relation showed 53.6% (30/56) of responses favoring the advantaged target when the targets differed in relation and 85.7% (47/56) when the targets differed in role. This difference was significant, $\chi^2(1) = 23.35$, $p < 10^{-5}$, whereas the interaction with comparison was non-significant, $\chi^2(7) = 9.69$, $p > .2$.

In the data for the common-scenario test, 58.9% (132/224) of responses corresponded to the target that had interacted with the base. This effect was marginally significant, $\chi^2(1) = 3.11$, $p = .078$, whereas the interaction with comparison was non-significant, $\chi^2(7) = 3.97$, $p > .5$. The regression model combining data from the common-relation and co-occurrence tests showed no significant differences between the two effects, $\chi^2(1) < .01$, $p > .5$.

5. Discussion

Experiment 2 largely replicated the results of Experiment 1, although the change in the similarity instruction did lead to some important differences that were anticipated. Specifically, the strong effect of common role was reproduced, and was in fact stronger than in Experiment 1 (see Fig. 1). However, this time the evidence for a common-relation effect was much weaker and was in fact not statistically significant. The co-occurrence (i.e., common scenario) effect was marginally significant, but, as in Experiment 1, was not significantly stronger than that of common relation, indicating once again that there is no effect of direct interaction per se. Thus, whether or not there is a non-zero effect of common relation under these instructions (i.e., whether clear evidence would be found with a greater number of participants), it is apparent that the influence of common role was stronger, and the influence of common relation weaker, for Experiment 2 as compared to Experiment 1.

The opposite effects of the change in task instruction on the common-role and common-relation effects suggests they are due to separate cognitive processes, which are activated to differing degrees by the terms "similar" and "like." We have motivated this division by appealing to the distinction the field makes between integration- and comparison-based processes and noting that scenario and relation matches are naturally classified as integration processes whereas role match can be derived through comparison processes.

Based on the above findings, we conjecture that the effect of common relations is a manifestation of thematic similarity, that is, similarity based on integrating thematic associations rather than aligning common properties (e.g., Bassok & Medin, 1997; Gentner & Brem, 1999; Wisniewski & Bassok, 1999). Because the present experiments showed no reliable effect of common scenario over common relation, it appears that direct association is not a prerequisite for thematic similarity. Rather, all that is necessary is participation in the same relation, or, more generally, association to the same schema (but in possibly different instances).

The distinction proposed here between integration and comparison-based similarity is consistent with, and in the same spirit as, previous work dissociating these processes (e.g., Gentner & Brem, 1999; Markman & Hutchinson, 1984). Gentner and Gunn (2001) showed that comparison (identifying commonalities) facilitates later listing of differences between the same items, whereas integration (identifying thematic relations) has the opposite effect. Smiley and Brown (1979) showed that although there exist age differences in preference between integration and comparison processes, individuals of all ages (from preschool through old age) have access to both processes. Wisniewski and Love (1998) have demonstrated that featural comparison and thematic integration processes can be selectively primed to influence the interpretation of ambiguous noun–noun combinations (e.g., an "elephant vet" is a large vet or a vet that treats elephants). Other work in noun–noun interpretation by Estes (2003) dissociates the effects of comparison and integration processes on the subsequent perceived similarity of the noun constituents—the constituents of noun–noun combinations interpreted through thematic integration are subsequently perceived as more similar whereas the opposite is true for constituents interpreted through comparison.

To sum up, comparison between the results of Experiments 1 and 2 suggests that the effect of role matches operates via a different process than that of relation and scenario matches. This distinction may correspond to the distinction in previous research between comparison and integration processes, and between featural and thematic similarity.

6. Experiment 3

One possible objection to the results of Experiments 1 and 2 is that they were due to unintended task demands. Specifically, it could be claimed that participants believed they should base their responses in phase 2 on the scenarios presented in phase 1, even though they would not act analogously in a naturalistic situation, simply because there was no other information present in the experiment. Our primary counter to such a claim would be that in fact there was a great deal of additional information available, namely the vast store of semantic knowledge about the objects in question that was present in participants' minds prior to the experiment. If participants used only the sentences they read to inform their similarity judgments, then there should be no variability in the base response rates to the various comparisons; response rates should depend only on the role of each comparison within the logical structure of the experimental design. However, this was not the case. Indeed, all five analysis types presented above show significant variability in base response rates to different similarity comparisons. That is, the main effect of comparison was significant in the logistic regression models (pooling data from both experiments) for the tests of common relation ($p < .01$), co-occurrence ($p < .05$), relation vs. co-occurrence ($p < .05$), common role ($p < .001$), and role vs. relation ($p < 10^{-6}$). Therefore responses were based on participants' prior knowledge of the objects in addition to the relational information from the incidental task.[3] This argues that the observed effects of relational information were mediated by changes in semantic knowledge.

Another way to determine whether participants were blindly using information directly from the scenarios is to test whether they still use that information when it is irrelevant to the task. Experiment 3 explores this possibility. This experiment mimicked the design and procedure of the first two, but instead of asking for similarity comparisons during the second phase, we asked participants which target was more likely to interact with the base. This question clearly favors responses in line with the co-occurrence effect found already (objects that were seen interacting should be perceived as likely to interact), but does not

[3] Variability in base response rates to different comparisons should not be confused with variability in the effect of relational information. The latter is measured by the interaction between comparison and role, relation, or co-occurrence. These interactions were all non-significant (as reported above), indicating that the effects of relational information generalize across stimuli.

have such clear implications for the common-role and common-relation effects. If anything, the role effect might be expected to reverse, because objects playing opposite roles might be more likely to interact (e.g., one animal chases and the other can be chased). However, the nature of the stimuli used makes many of them incompatible (e.g., a seal chasing a cat). Thus, the information presented in phase 1 is largely irrelevant to many of the comparisons in phase 2, and therefore by our account we predicted a null effect for both the common-role and common-relation tests. On the other hand, if the criticism mentioned above has merit, then participants should be expected to exhibit some manner of effect for both of these tests, due to their attempt to base responses directly on the scenarios they read.

6.1. Method

6.1.1. Participants

Participants were 88 undergraduate students at the University of Texas, Austin, who received partial credit in an introductory psychology course.

6.1.2. Stimuli, design, and procedure

Experiment 3 was identical to Experiments 1 and 2 except that in the triad task participants were instructed to choose the target object that is more likely to interact with the base object, rather than being asked to choose the target object that is more similar to (Experiment 1) or more like (Experiment 2) the base object.

6.2. Results

Data were analyzed identically to those of Experiments 1 and 2 (see above for details). Inspection of catch-trial responses showed that only one participant missed more than one of the four catch-trial questions (this participant missed two). Therefore all participants were included in later analyses. The primary results given in this section are also displayed in Fig. 1 for comparison to the first two experiments.

In the test for the effect of common role, the target matching the base's role was selected in 49.7% (175/352) of responses. Logistic regression showed this effect to be non-significant, $\chi^2(1) = .010$, $p > .5$, although there was a significant interaction with comparison, $\chi^2(7) = 16.32$, $p < .05$. Inspection of individual comparisons revealed that there was a strong positive effect of common role for the question of which was more likely to interact with a vein: a tunnel or railroad tracks. Because this comparison seems essentially meaningless, we conjecture that this was simply a Type I error. When the data from this comparison were removed from the analysis, the interaction disappeared ($\chi^2[6] = 8.23$, $p > .2$), while the main effect of role remained null.

In the test for the effect of common relations, the target matching the base's relation was selected in 49.4% (87/176) of responses. Logistic regression showed this effect to be non-significant, $\chi^2(1) = .027$, $p > .5$, with a marginally significant interaction with comparison, $\chi^2(3) = 6.54$, $p = .09$.

Finally, the co-occurrence test showed a 61.2% (431/704) rate of responding in favor of the target that had appeared in the same scenario as the base. This effect was significant, $\chi^2(1) = 17.04$, $p < 10^{-4}$, with a significant interaction with comparison, $\chi^2(7) = 22.41$, $p < .01$. Inspection of the data for individual comparisons showed that the effect of co-occurrence was restricted to the four comparisons based on the first two words sets (hence the interaction with comparison). The lack of an effect in the latter two words sets was likely due to ceiling effects in base response rates (e.g., 95% of participants in the analysis believed that a customer was more likely to interact with a cashier than with a stock-boy).

6.3. Discussion

The primary conclusion from Experiment 3 is that the effects of common roles and relations found in the first two experiments are specific to the comparison tasks used there. If these earlier results had been artifacts of the experimental design, due to participants' explicitly trying to base their responses on the patterns of sentences they had read, then similar effects should have occurred in Experiment 3. Instead, it was found that participants exhibited no effect of roles or relations, but only an effect of co-occurrence. These null effects were found despite the fact that the sample size was two to three times larger than in the previous two experiments. This shows that participants were using just the information that was relevant to the task at hand. When asked which objects are likely to interact, their responses depended only on which objects they had seen interacting; when asked which objects are similar (or alike), their responses depended on which objects had been involved in matching relations or had played identical roles. Thus, we can conclude that the pattern of effects in each case reflects the question asked, and in particular that effects of common roles and relations are present in people's subjective assessment of similarity.

7. Applying co-occurrence based models to relational learning

As mentioned above, the present results are partially anticipated by word learning models that exploit word occurrence statistics. For example, LSA (Landauer & Dumais, 1997) works by creating a matrix of occurrence counts that tracks the number of times each word appears in each context (where a context can be anything from a sentence to an entire document or discourse). This matrix is then subjected to a dimensional reduction via singular value decomposition. Words with similar context vectors in the reduced space are then rated by the model as being semantically similar. Thus, LSA predicts one major component of our results, namely the effect of common relation, because appearing with the same verb(s) implies appearing in similar contexts.

Despite its ability to exhibit a common-relation effect, LSA is unable to explain the effect found here for common roles, due to its complete insensitivity to sentence structure. More precisely, the construction of the occurrence matrix ignores grammatical constructions and the locations of words within a sentence. Thus, the input to the model does not carry information regarding whether a given noun was the agent or the patient of a given verb. It might be hoped that such structural information could be extracted from the corpus as a whole, in the same way that LSA extracts much other latent information, but we know of no specific mechanism by which this might happen. Furthermore, experimentation with the model shows that it consistently predicts greater similarity between thematic associates than between items that share features and roles. Some illustrative examples include: *cigarette* is rated as more similar to *smoker* than to *cigar*; *car* is closer to *driver* than to *truck* or *motorcycle*; *coffee* is closer to *cup* than to *tea*; *cow* is closer to *barn* than to *pig* or *horse*; and *fox* is closer to *rabbit* than to *wolf*.[4] Thus, although LSA provides an elegant explanation for how humans induce many aspects of word meaning from verbal input, the present results imply that the lack of structural sensitivity is a critical shortcoming of the model.

HAL's predictions regarding relational similarity are slightly different. HAL works by constructing a square matrix of word co-occurrence counts, so that the context of a word is given directly in terms of the words with which it appears (Lund & Burgess, 1996). This matrix stores separate counts for each word pair (in its lower and upper triangles) based on which word occurs first in the input stream. Thus the model is sensitive to word order. At first inspection it might appear that this would allow HAL to correctly predict the common-role effect, because a noun's position relative to a verb (before or after) is often correlated with its role in the relation specified by that verb (agent or patient). However, because HAL is only sensitive to word order and not to true syntactical structure, it is unable to recognize differences in sentence construction such as active vs. passive voicing. If trained on a corpus including both of these constructions, such that role and word order information were decorrelated, HAL would also fail to predict role-based similarity.

7.1. A structurally sensitive model of relational learning

The following demonstration shows how the principles underlying HAL and LSA can be extended to explain the full range of our results, including role-based similarity. Here we present a new model, Relations Offer Latently Extracted Similarities (ROLES), which combines the statistical approach used in previous models with sensitivity to sentence structure. ROLES is intended to showcase the untapped potential of extending corpus approaches so that they are sensitive to role information as demonstrated in Experiments 1 and 2. Illustrative simulations of ROLES are offered as a potential explanation of our experimental results.

The primary advance embodied by ROLES is that sentences are represented not as sets of words but rather as trees that encode the relational structure conveyed by each sentence. These trees are constructed in the same manner as in previously successful models of human analogy-making (Larkey & Love, 2003). Fig. 2A shows an example tree that applies to all of the sentences in the present study. The most important feature of this construction is that each verb is expanded into a subtree consisting of the verb itself attached to nodes representing each of its thematic roles; the objects filling these roles are attached to these latter nodes. Thus, the assumption of the model is that people extract verb-specific thematic roles and tag them during initial processing of the input. A further assumption is that structural elements such as the thematic roles of verbs are treated in the same way as regular words. As with HAL, the input to ROLES is a co-occurrence matrix among all words (and structural elements), with co-occurrence counts inversely weighted by the distance between words (as $weight = e^{-distance}$). However, rather than using distance in the input stream (i.e., number of intervening words), ROLES uses the graphwise distance obtained from its tree representation.

[4] These computations were performed using the One-to-Many Comparison utility available at http://lsa.colorado.edu, using the topic space General_Reading_up_to_1st_year_college (300 factors).

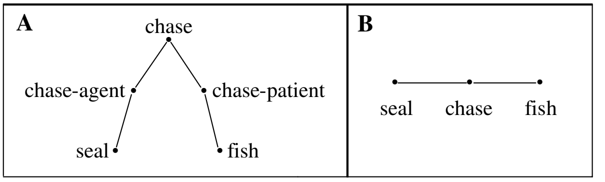

Fig. 2. Example graphs used to generate input for (A) ROLES and (B) the non-structural version of the model, for the sentence "The seal chases the fish."

Although Fig. 2A presents a simple graph containing one verb, the same approach applies to arbitrarily complex structures. For example, in the sentence *The captain sees the seal chase the fish*, the chase node shown in Fig. 2A would itself be bound to the verb-specific thematic role *see-patient*. Relations serving as arguments to other relations are accommodated by existing knowledge representation approaches (e.g., Larkey & Love, 2003) and the graphwise distance measure defined here holds for these complex structures as well.

Finally, as with LSA, ROLES' input matrix is diagonalized using singular value decomposition and its rank is then reduced by removing the smallest elements in this diagonal representation. The result is a matrix with lower dimensionality (i.e., rank) than the input, containing as rows a semantic vector for every word in the corpus.

In summary, the model assumes some syntactic preprocessing (not explicitly modeled here); semantic knowledge is then extracted from the resulting structured representation, rather than from an undifferentiated set or sequence of words. Although the present application does not involve other constructions such as the passive voice, it is assumed that the preprocessing produces a deep structure that is largely unaffected by such variations. That is, a sentence presented in passive voice would induce the same graph structure as in Fig. 2A, because the correspondence between items and thematic roles would be unchanged. Thus the model is sensitive to true semantic structure, not just word order (cf. Lund & Burgess, 1996).

Clearly, humans can extract thematic roles from events and text. Thus, including this information in the input is reasonable for the present use of ROLES as a psychological model. Still, one interesting question is how this requirement impinges on the practical applicability of ROLES. The need for ROLES's input to include thematic role information does not imply that ROLES could not be applied to corpora without human intervention. Existing systems can automatically extract thematic roles from text (Gildea & Jurafsky, 2002; Miller, Fox, Ramshaw, & Weischedel, 2000). The output of these systems could serve as the input to ROLES. These systems determine thematic roles by automatically extracting syntactic cues (e.g., word position, parse tree paths) that probabilistically predict the roles of verb constituents. After being trained on a relatively small hand-annotated corpus, these systems can achieve high levels of performance on novel corpora. A recent competition involving ten systems has demonstrated impressive performance improvements (Carreras & M'arques, 2004). The ten systems in the competition used very shallow parses that consisted only of chunks (e.g., noun phrases) of non-overlapping, non-embedded constituents. The sentences utilized in our experiment would be trivial for existing systems to tag.

As such systems improve, ROLES's performance should also improve. Nevertheless, it remains an open question whether ROLES could scale to large corpora in which numerous tagging errors would occur. Whether these errors would act as harmless noise or systematically corrupt the extracted meanings is a question to be resolved. Here, we focus on ROLES as a psychological explanation for the results of Experiments 1 and 2.

7.2. Simulation

To evaluate the ability of ROLES to explain the present empirical results, the model was applied to the set of twelve experimental sentences viewed by each participant (catch-trial sentences were omitted). Corpus approaches are typically applied to a large body of text. However, because of the experimental design of our studies, participants' ''internal corpus'' should only differ between groups in terms of their exposure to the study sentences. Implicit to the design and predictions of Experiments 1 and 2 is the assumption that more recent information will be disproportionately weighted (cf. Jones & Sieck, 2003; Stewart, Brown, & Chater, 2002). Thus, in order to evaluate ROLES's predictions for the present experiments it is sufficient to train the model on only the experimental sentences. Further research will determine whether ROLES scales up to explain psychological phenomena that depend on larger or more complex knowledge bases.

Training on the 12 study sentences involved converting each sentence into the corresponding graph (e.g., Fig. 2A), compiling the distance-weighted co-occurrence matrix for the corpus of sentences, and then reducing the dimensionality (rank) of this matrix through singular value decomposition. This resulted in a matrix whose rows represent semantic vectors for all of the words in the corpus. Similarity measures between pairs of words were then obtained as cosines between their respective semantic vectors. Response probabilities were generated by a standard Luce choice rule:

$$P(A) = \frac{e^{D^* sim(A,base)}}{e^{D^* sim(A,base)} + e^{D^* sim(B,base)}}$$

where A and B represent the targets and D is a scaling parameter. The scaling parameter and the number of factors used in the dimensional reduction were adjusted to provide the best least-squares fit to the response rates in the common-role, common-relation, and co-occurrence tests averaged over the first two experiments. The qualitative pattern, consisting of a positive effect for all three measures, a stronger effect of common role than common relation, and roughly equal effect of common relation and co-occurrence, was robust over a wide range of parameter settings. The best fit is shown in Table 4.

In order to evaluate the importance of the structured input assumption embodied in ROLES, a non-structural version of the model was also fit to the data. This model differed only in its input representation. Specifically, we removed the thematic role nodes and

Table 4
Summary of simulation results

Data source	Test			D	Factors	SSE
	Role	Relation	Scenario			
Humans	72.8%	61.0%	64.7%	—	—	—
ROLES	76.4	57.9	57.9	1.49	18	0.69
Non-structural model	50.0	55.8	66.7	1.97	32	5.51

Notes: Human data are averaged over Experiments 1 and 2. D represents response-scaling parameter. Factors column shows number of dimensions used in semantic representation. SSE is sum of squared error.

treated each sentence as a simple string of words (see Fig. 2B). Word order was still preserved. However, most of the grammatical structure was removed, and in particular agents and patients could no longer be distinguished (because the graphs were symmetric). For all parameter settings, this model produced no effect of role; that is, response rates for the relevant similarity comparisons were all exactly 50% (see Table 4 for the best fit). Therefore the structural representation is critical to explaining the present results, in particular participants' sensitivity to objects' roles. Conversely, inclusion of this richer input representation is sufficient to explain the structural sensitivity exhibited by participants in our experiments.

8. General discussion

Historically, accounts of object representation and perceived similarity have focused on isolated features (e.g., Tversky, 1977). Although more recent accounts have explored how objects, scenes, or situations containing common relational structures come to be perceived as similar (Medin et al., 1990), little is known about how the perceived similarity of parts or objects embedded within these relational systems is affected. Given that a large proportion of properties listed by subjects for common objects are extrinsic (McRae, Cree, et al., 2005), it seems likely that such properties affect similarity as well. The current studies and simulations directly address this critical, but under-explored, issue, and test the hypothesis that objects situated in common relational systems come to be perceived as more similar. Taken together, the results of Experiments 1 and 2 show that being situated within common relational systems increases similarity through at least two separate processes. First, objects that are involved in the same relation are more similar when they play the same role in that relation (e.g., *chaser*). To a lesser extent, involvement in the same relation also provides a boost in similarity even when the objects play different roles (e.g., agent and patient of the *chase* relation). These two types of similarity are in accord with the field's distinction between integration- and comparison-based processes.

In considering the first effect of common role, a natural question to ask is what precisely defines a role. At one extreme, roles can be situation-specific (e.g., *Vice President for Southwest Regional Sales during 2001*). At the other extreme, roles such as *agent* span across many predicates (Dowty, 1991). We propose that people operate at an intermediate level, specifically that the level of significant generalization is that of a frame (Baker, Fillmore, & Lowe, 1998). A frame is a schematic representation that specifies the interactions of participants and props (Fillmore, 1976). For example, the verbs *buy* and *sell* share a common frame, and thus roles defined with regard to one would generalize to those based on the other. (Of course, a verb with multiple senses will be associated with multiple frames; see Hare, McRae, & Elman, 2003, 2004.) Related approaches that decompose relations and roles into constituent features allow for generalization across related roles through feature overlap (Hummel & Holyoak, 1997). However, further research is needed to investigate and compare generalization at different levels, thus addressing the question of how specifically relational information is represented.

The second effect of relational information on similarity is non-role-specific: involvement in the same relation provides a boost in similarity even when the objects play different roles. Thus, the effect of common role mentioned above occurs over and above this effect of common relation. The effect of common relation seems to be a manifestation of thematic or integrative similarity (Bassok & Medin, 1997; Wisniewski & Bassok, 1999). However, we found no appreciable effect of direct interaction (appearance in the same scenario) beyond the effect of common relation. Thus, contrary to earlier views, enhanced similarity of objects participating in common scenarios may only depend on association to a common schema. As with the common-role effect, we predict that the relevant level of description for such schemas is not as specific as that determined by a particular verb, but instead is at the level of a frame.

It is important to note that the effects found here were likely not based on participants' learning new relational properties about objects. This is because the scenarios presented were chosen so as to be common occurrences for the objects involved. Thus, differences in participants' similarity responses probably arose from the increased salience of pre-existing relational properties, as a result of recent experience with these interactions. That similarity can be temporarily affected in such a way should not be surprising, since it is well known that similarity differentially depends on properties according to their current salience or relevance (e.g., Nosofsky, 1986; Tversky, 1977). Thus, similarity depends both on learned properties and on the present salience of these properties. Of course, the relational properties used in the present experiments must have been learned previously; that is, even though our results were likely not due to participants' acquisition of new relational knowledge, they do imply that such acquisition has taken place in the past. Presumably, similar results could be obtained in a task requiring actual learning (i.e., by using blank predicates or novel objects), although such a demonstration may require more exposure than was given here.

One domain in which our results have particular significance is theories of word learning that extract meaning from occurrence statistics, such as LSA (Landauer & Dumais, 1997) and HAL (Burgess, 1998; Lund & Burgess, 1996). These models are based on the same principle we advocate here, namely that concepts are represented largely in terms of their relationships to other concepts. In these models, the concepts and relationships in question are words and their co-occurrences. As explained earlier, the patterns of co-occurrence between nouns and verbs naturally account for the effect of common relations found here. However, the common-role effect found in the present experiments cannot be explained by mere co-occurrence. This implies that purely statistical techniques are insufficient to capture human cognition, because they do not adequately take account of sentence structure. Conversely, it was shown here that integrating co-occurrence statistics with structural information, specifically the verb-specific thematic roles filled by nouns, is sufficient to explain the full range of relational effects on similarity found in our studies. The model presented, ROLES, incorporates the core principles of past models but additionally assumes that verbal input is initially processed to extract grammatical structure and thematic roles. A semantic representation is then derived based on co-occurrence statistics among both true words and more grammatical elements such as verb-specific thematic roles. The model was fit here to the averaged data from Experiments 1 and 2, but can also account for the separable effects of common role and common relation, by differentially weighting the context-vector components based on different item types (verbs vs. their thematic roles) when computing similarity.

Much of the power of LSA comes from the dimensional reduction process, which produces indirect inferences about word meaning as a result of simultaneously accommodating large numbers of constraints. Landauer and Dumais (1997) argue that the power of this mechanism is only fully apparent when applied to large corpora and high-dimensional semantic spaces. The present simulations used only a small set of sentences and a correspondingly low-dimensional semantic representation. Still, the model was able to extract meaning based on relational information, even for words that had not directly co-occurred. It remains an open question how ROLES will behave on larger corpora, for which indirect inference should play an even greater role, but as with LSA we expect that its learning power would only increase.

One interesting prediction that follows from this approach to learning is that the model develops explicit representations for verb-specific thematic roles, such as the agent and patient in $hit(x, y)$. These roles will develop strong associations to the nouns that regularly fill them. This prediction is borne out by studies of verb-specific thematic roles by McRae and colleagues. First, McRae, Ferretti, and Amyote (1997) found that people can report typical features associated with such roles (i.e., the typical features of objects that fill the role), that these role/feature typicality ratings predict role/filler typicality ratings, and that compatibility between an object's features and roles affects reading times for ambiguous sentences. Second, Ferretti et al. (2001) found that verbs prime nouns that typically fill their thematic roles as well as common features of those objects, and that sentence fragments (e.g., "He arrested the" or "He was arrested by the") selectively prime nouns that fit the unfilled role (*criminal* or *cop*, respectively). These results suggest that people have learned concepts corresponding to the class of objects that fit each role of a given verb; such concepts would correspond to the representations of thematic elements extracted by ROLES.

8.1. Possible mechanisms of relational influences on similarity

The primary goal of these studies was to empirically investigate how perceived similarity is affected by embedding of objects within relational systems. The results can be interpreted from within a number of different theoretical frameworks and are potentially due to any of several possible mechanisms. Here, we discuss what we see as three of the most promising explanations along with some of their implications and predictions. It will be an important task for future research to determine their relative merit.

8.1.1. Roles as properties

The most straightforward explanation of the present results is that the roles played by objects, and more generally the relations in which they engage, are stored as semantic properties of those objects, just as intrinsic features are. Similarity between objects then reflects some measure of commonality taken over all properties, both relational and featural. This is essentially the stance taken by ROLES: co-occurrence with featural descriptors (e.g., adjectives) indicates intrinsic features of objects whereas co-occurrence with relational descriptors (verbs and their thematic roles) indicates relational properties, but both are treated in the same way when deriving semantic representations.

If the roles-as-properties interpretation is correct, then models of semantic knowledge (e.g., Collins & Loftus, 1975; Masson, 1995) need to be extended to include relational information (but see Gentner, 1975; Rumelhart & Levin, 1975). In the case of spreading activation models, this implies the inclusion of nodes representing the separate roles of a relation, with connections from these nodes to the relation itself as well as to the objects that typically fill those roles. Such an approach is also supported by evidence that verbs prime nouns that tend to fill their thematic roles (Ferretti et al., 2001) and vice versa

(McRae, Hare, Elman, & Ferretti, 2005). Spellman, Holyoak, and Morrison (2001) also present evidence, based on analogical priming, that semantic activation depends on the specific relations linking objects and argue that access to semantic knowledge is based on analogical mapping rather than simple spreading activation.

A related possibility that should be explored is that similarity depends not just on explicit first-order relations among objects, but also on higher-order relations. Bassok and Medin's (1997) findings support the notion that people are sensitive to implied higher-order relations within systems. For instance, they found that the sentence *The carpenter fixes the chair* is judged as more similar to the sentence *The electrician fixes the radio* than to *The electrician fixes the chair*. Apparently the match in the abstract concept of a person doing his or her job, which is a second-order relation between the relation *fix* and its arguments *carpenter* and *chair*, is more important to similarity than the featural match of the chair. In this case the relations in question were all intrinsic to the items being compared, but it is possible that such effects could also be found for extrinsic relational properties. For example, it would be interesting to see whether *carpenter* and *plumber* are more likely to be rated as similar by someone having read *The carpenter fixes the chair* and *The plumber installs the sink* than by someone who had read *The carpenter fixes the chair* and *The plumber fixes the radio*. Such an experiment would contrast the effects of higher-order relations against the effects of first-order relations found in the present study. If an effect of higher-order relations were found (as we anticipate it would be), then the preprocessing assumed by ROLES would need to be expanded to include not just grammatical tagging but also inferences of a more semantic nature. This would allow the input to the model to include higher-order relations that are not explicitly present but are inferable based on prior knowledge.

8.1.2. Similarity through analogy

The effect of common role may also be explainable as the result of participants' forming analogies between scenarios involving the same relation. Objects playing the same roles would then correspond to each other within these analogies, which could in turn lead to an increase in similarity. For example, the process of analogy formation involves putting objects into correspondence, which in turn involves comparing them (Gentner, 1983). It is known that the process of comparison by itself tends to increase the similarity between the items being compared (Boroditsky, in press), so the comparisons involved in analogy formation could increase the similarity between corresponding elements (in this case, objects playing the same role within the same relation). If this explanation is correct, then it implies a complex bi-directional relationship between the processes of analogy and similarity, in that analogy is driven by similarity between potentially corresponding elements (e.g., Gentner & Toupin, 1986; Goldstone, 1996) but the formation of an analogy also influences that same similarity. It also implies that analogy likely plays an important role in naturalistic category learning, where relational information about stimuli is abundant (as contrasted with most category learning experiments).

It is interesting to note that the analogy-based explanation is essentially an exemplar explanation, whereas the roles-as-properties proposal is essentially a prototype explanation. That is, under the analogy explanation, people compare two objects by recalling specific instances of their occurrence, and comparing the contexts of those occurrences. Under the roles-as-properties explanation, people rely on an abstracted representation of each object, which has properties that are determined by the contexts of individual occurrences but that are no longer necessarily associated to those specific instances. Just as it has been very difficult to distinguish between these competing explanations in the category learning literature (e.g., Nosofsky & Zaki, 2002; Smith & Minda, 2002), much work on the relationship between analogy, similarity, and category learning may be needed before the question can be adequately addressed in the present context. Furthermore, as in category learning, it may turn out that the best explanation lies somewhere in between exemplars and prototypes (Anderson, 1991) with situational factors determining the balance (Love, Medin, & Gureckis, 2004).

8.1.3. Mediation by feature salience

Finally, there is the possibility that appearance in a particular role could highlight certain intrinsic features of an object (cf. Barsalou, 1987). For example, reading about a polar bear chasing a seal may lead to inferences about the vulnerability and small size of the seal. Objects playing similar roles (e.g., a cat being chased by a dog) could thus have similar (temporarily) salient features, which could then lead to increased similarity via a features-only process. Slightly more general (non-role-specific) features could be responsible for the common-relation effect. As mentioned above, it is known that similarity is affected by the variable salience of features, so it seems plausible that a process of this sort could have occurred in the present experiments. Even if this is the case, it is important to point out that relational knowledge is still playing a critical role, in that it is driving the inference process. That is, in order for participants to infer features of an object based on its role, they must have pre-existing associations between the role (or relation) and the features, and in particular they must have learned concepts for individual roles and relations. Just how these roles and relations are represented, and how the knowledge they embody drives the inference process, are important questions for future research.

8.2. Implications beyond similarity

Issues of semantic representation have broad relevance. The general idea that participating in common relational structures can increase the perceived similarity of objects embedded in those structures has currency beyond predicting similarity judgments. The relevance of our studies to statistical approaches to word learning has already been considered in our discussion of the ROLES model. Additional theoretical connections are considered below.

8.2.1. Language learning and processing

A critical step in language comprehension involves assigning nouns to the thematic roles of verbs (Carlson & Tanenhaus, 1988; Fillmore, 1968). Our results suggest that semantic knowledge about objects may facilitate this process, as this knowledge incorporates information about verb-specific thematic roles that each object is likely to fill. This finding complements the work of McRae and colleagues (McRae et al., 1997, 1998; Ferretti et al., 2001), who have shown that the meaning of a verb incorporates information about objects that are likely to fill each of its thematic roles. Knowledge about the compatibility of nouns and verb-specific thematic roles plays a critical role in online sentence comprehension, by aiding in rapid resolution of structural ambiguities in garden path sentences (McRae et al., 1998) and by selectively priming nouns that provide good fits for unfilled thematic roles (Ferretti et al., 2001). Thus, relational knowledge may play a critical role in the interplay between language structure and meaning.

A related question is how acquisition of relational knowledge interacts with learning of syntax early in development. Role-specific learning should be difficult or impossible prior to knowledge of the language's syntax (because syntax is needed in general to assign nouns to thematic roles), but on the other hand it would be difficult to extract general syntactic rules without knowledge of which nouns are likely to fill which roles. A possible answer to this problem comes from Tomasello's verb-island hypothesis (Akhtar & Tomasello, 1997; Diessel & Tomasello, 2001; Tomasello, 1992), which states that children initially learn isolated syntactical regularities for individual verbs, based on the semantics of those verbs. This fits well with ROLES' prediction that separate concepts are learned for the thematic roles of individual verbs. Once such role-specific concepts are learned for a particular verb, based on an understanding of that verb's meaning and repeated (verbal and non-verbal) experience with that type of interaction, the language learner can recognize patterns in the order in which the nouns corresponding to the different roles tend to be uttered. This would allow initial formation of verb islands, which should eventually be combinable to induce the full syntax of the language (Tomasello, 1992).

The question of whether concepts are represented by intrinsic properties vs. extrinsic relations also relates to the question in language production of whether generation of syntactic structure is guided by the features of concepts (or lexical entries) vs. their thematic roles. This question has been investigated using the phenomenon of structural priming (Bock, 1986), by which people tend to reuse syntactic constructions from one sentence to the next, even in the absence of overlap in content. Bock, Loebell, and Morey (1992) show that the strength of structural priming can depend on features of concepts, such as animacy. Given a choice between active and passive constructions (which allows freedom in the assignment of agent and patient to subject and object), participants tend to produce the option that matches the prime in animacy of the subject. However, Chang et al. (2003) show that thematic roles also influence structural priming. Stimuli in their studies were based on the locative alternation (e.g., *The man sprayed the car with wax*/*The man sprayed wax on the car*), which varies the relative positioning of the location (*car*) and theme (*wax*). Animacy of the role fillers was controlled for to minimize possible featural effects. Chang et al. found that participants were biased towards the construction that matched the prime in order of thematic roles (location–theme vs. theme–location), thus demonstrating that thematic role information influences sentence structure. This result parallels the present findings, by showing that relational information plays a role similar to that of featural information in cognitive representations (in this case for the purposes of language production).[5]

Turning back to the present findings, an important question is whether the observed effects of relational information were driven by syntax or semantics. We have assumed that the common role effect is a semantic effect due to objects playing the same thematic role (e.g., agent), but it is also possible that the effect is due to both objects filling the same syntactic role (e.g., subject). These two interpretations could be discriminated by using passive constructions. For example, if participants were presented with *The polar bear chases the*

[5] Chang et al.'s (2003) findings differ from ours in that they found effects of verb-general thematic roles, whereas the present experiments investigated the effects of roles within specific relations. However, our findings do not rule out verb-general similarity effects (e.g., two objects becoming more similar due to acting as the agent in different relations), albeit likely weaker than those found here. Furthermore, Pickering and Branigan (1998) show that structural priming is stronger when the prime and target sentences contain the same verb.

seal and *The cat is chased by the collie*, the semantic hypothesis would predict the seal to be rated as more similar to the cat, because both are patients of the chase relation, whereas the syntactic hypothesis would predict the seal to be rated as similar to the collie, because both are objects. The semantic prediction seems more likely, because similarity is generally assumed to be a semantic construct, but further research along these lines is needed before the conclusion can be made definitively.

8.2.2. Decision making

Normatively inspired theories of decision making, based on expected utility theory (von Neumann & Morgenstern, 1947), assume that decisions are based on intrinsic values, or utilities, of outcomes (e.g., Kahneman & Tversky, 1979; Quiggin, 1993). One critical assumption of utility-based approaches is that the utility of an option is independent of the other options in the choice set. However, there is much evidence showing that this assumption is untenable, in that preferences between certain options can depend on which other options are present (Huber, Payne, & Puto, 1982; Simonson & Tversky, 1992; Stewart, Chater, Stott, & Reimers, 2003). For example, in the attraction (or asymmetric dominance) effect, preference for one object relative to another can be increased if there is a third option present that the first dominates on all dimensions (Huber et al., 1982). The attraction effect and other violations of independence are potentially explainable in the present framework by assuming that relations such as dominance are encoded as role-based properties of objects and that these (extrinsic) properties affect subjective valuations. Indeed, recent theories that have been developed to explain such phenomena all rely on relative evaluation of alternatives (Choplin & Hummel, 2002; González-Vallejo, 2002; Roe, Busemeyer, & Townsend, 2001; Stewart, Chater, & Brown, 2006; Tversky & Simonson, 1993). Therefore the effects of relational information found in the present study may share a common mechanism with processes underlying decision making.

8.2.3. Category learning

Perhaps the most important implications of the present study are those relevant to the study of concepts and categories. Past research in category learning has almost uniformly assumed that objects and categories are represented solely in terms of intrinsic features, without regard for how separate objects might interact. This assumption is implicit in the abundance of laboratory studies using isolated, non-interacting objects, such as geometrical figures or dot patterns (e.g., Posner & Keele, 1968; Shepard, Hovland, & Jenkins, 1961), as well as in more naturalistic studies of category representation (Rosch & Mervis, 1975). The same assumption is made in essentially every formal theory of category learning, including exemplar (Medin & Schaffer, 1978; Nosofsky, 1986), prototype (Reed, 1972; Smith & Minda, 2002), clustering (Anderson, 1991; Love et al., 2004), and rule-based models (Collins & Quillian, 1969; Nosofsky, Palmeri, & McKinley, 1994). However, recent work has brought into question the adequacy of the intrinsic feature assumption even for laboratory studies of category learning in restricted domains (Love & Markman, 2003).

In contrast to the majority of work in category learning, the results presented here demonstrate that relational information plays an important role in category representations. This suggests that relations among stimuli, when present, may also play an important role in the learning of novel categories. Evidence for this claim comes from recent work by Larkey, Narvaez, and Markman (2004), who show that under certain conditions participants prefer to categorize objects based on their roles rather than on their features, and by

Rehder and Ross (2001), who show that relational information can be used to learn categories even in the absence of any regularities in individual features. Taken together with these findings, our results imply that theories and empirical studies of categorization both need to be extended to consider the impact of relational information on learning. Extending current models may be as simple as allowing input representations to include relational properties, but may also require more complex and interactive processing of exemplars, including determination of analogical relationships among them (Tomlinson & Love, 2006). Deciding among the competing explanations offered here for our results will be a critical starting point in addressing this question.

Our results also have important implications for Markman and Stilwell's (2001) theory of role-governed categories. On the one hand, they provide direct empirical evidence that category representations include relational information. However, the work on role-governed categories has been primarily cast as an investigation into different category types. For example, Markman and Stilwell's dichotomy assumes that "feature-based categories…only have feature-based representations" (p. 341). The experiments here provide a quite different interpretation. In particular, the stimuli used all corresponded to concrete categories of the sort that would normally be referred to as "feature-based". Therefore our results are not limited to a subclass of concepts such as role-governed categories. Instead, they support the claim that essentially every category has some relational component. Thus, instead of separate category types, we are led to a view of different types of information, featural and relational, that both contribute to all categories albeit to varying degrees (see Wisniewski, Clancy, & Tillman, 2005, for a similar argument). This suggests that future work should focus on the implications of differential reliance on different types of information (cf. Barr & Caplan, 1987), rather than on strict taxonomies of category types.

Appendix A. Details of design and analyses

The full set of objects and relations used in all three experiments (and, in abstract form, in the simulations) is presented in Table A1. These items are grouped into four word sets, each containing seven object names (labeled A through G) and two relation names (X and Y). The four word sets all played identical roles in the design. For a given participant, each word set contributed three sentences to the incidental learning phase. The nature of these sentences was determined by the participant's assigned condition for that word set, according to the design given in Table A2. For example, a participant who was in

Table A1
Complete set of stimuli for Experiments 1–3

Word set	Object							Relation	
	A	B	C	D	E	F	G	X	Y
1	Polar bear	Seal	Fish	Collie	German shepherd	Cat	Sheep	Chase	Herd
2	Engineer	Robot	Car	Construction worker	Carpenter	Building	Foreman	Build	Listen to
3	Blood	Vein	Arm	Road	Railroad tracks	Tunnel	Mountain	Run through	Go over
4	Private	Sergeant	General	Stock-boy	Cashier	Manager	Customer	Take orders from	Assist

condition 2 for word set 3 viewed the sentences *The blood runs through the vein* (A X B), *The railroad tracks run through the tunnel* (E X F), and *The road goes over the mountain* (D Y G). The four conditions form a 2×2 design that varies the role of B (agent vs. patient) and the relations for D and E. Note that replacement of the entries in Table A2 with the items from word set 1 reproduces the top half of Table 2 in the main text.

The similarity comparisons (interaction comparisons for Experiment 3) used in all experiments and simulations are displayed in Table A3. These comparisons were replicated once for each word set, for a total of 20. Every participant responded to this same set of comparisons. All of the primary analyses involved contrasting response rates to particular comparisons across certain groups, as detailed in Table A4. For example, the test of the common-relation effect involved comparing response rates to the comparison $B \sim D/E$ (base B, targets D and E) between participants in conditions 1 and 2 for each word set. Condition 1 participants had seen B and D in the same relation whereas condition 2 participants had seen B and E in the same relation. Thus, a positive effect of common relation corresponds to the former group selecting the first target (D) more often (hence condition 1 is listed first in this line of the table). As with the previous two tables, each line in Table A4 corresponds to four contrasts, one from each word set. Thus the common-relation

Table A2
Design of incidental learning phase for experiments and simulations

Condition	Phase 1 sentences		
	Agent	Relation	Patient
1	A	X	B
	D	X	F
	E	Y	G
2	A	X	B
	E	X	F
	D	Y	G
3	B	X	C
	D	X	F
	E	Y	G
4	B	X	C
	E	X	F
	D	Y	G

Notes: Each participant viewed three sentences from each word set. These sentences were determined by the participant's assigned condition for that set, according to this table. See Table A1 for the actual words used.

Table A3
Comparisons presented in all experiments and simulations

Base	Target 1	Target 2
B	D	E
B	D	F
B	E	F
F	D	E
G	D	E

Note: Each row represents four comparisons—one for each word set (see Table A1). Thus a total of 20 comparisons were presented.

Table A4
Contrasts used in primary analyses in all experiments and simulations

Test	Comparison	Conditions
Role	B ~ D/F	3 vs. 1
	B ~ E/F	4 vs. 2
Relation	B ~ D/E	1 vs. 2
Role vs. relation	B ~ F/D	1 vs. 4
	B ~ F/E	2 vs. 3
Scenario	F ~ D/E	1,3 vs. 2,4
	G ~ D/E	2,4 vs. 1,3

Notes: Comparisons are written as base ~ target 1/target 2. Each row represents four comparisons—one for each word set—that all contributed to the particular test. For all contrasts, a positive effect (or greater effect of role than relation) corresponds to participants in the first listed condition(s) choosing target 1 more often than participants in the second condition(s).

analysis comprised data from four contrasts, and the other three analyses each comprised eight.

As stated in Footnote 1, data in the co-occurrence analysis could not be assumed to be independent, because each participant contributed responses to both F ~ D/E and G ~ D/E from each word set (regardless of the participant's condition for that set; see Table A4). Because such a pair of comparisons involves overlapping words, responses from a given participant could be correlated. Therefore only one observation from each such pair was included in the regression model. (Because estimates of means are not affected by dependence among variables, the average response rates reported for all three experiments were based on the full dataset.) Selection of which data to include was done randomly, subject to the constraint that, for each condition in each word set, the F ~ D/E data were retained for half of the participants in that condition and the G ~ D/E data were retained for the other half. In Experiment 2, each condition for each word set contained an odd number of participants (seven); therefore we retained three participants for F ~ D/E and four for G ~ D/E from each of conditions 1 and 2, while we retained four for F ~ D/E and three for G ~ D/E from each of conditions 3 and 4. This division ensured that all eight contrasts in the model had the same number of observations and that each was balanced.

As stated in Footnote 2, analyses comparing the effect of co-occurrence to that of common relation could only include half of the co-occurrence data, due to independence considerations. Specifically, the co-occurrence observations from participants in conditions 1 and 2 for each word set had to be excluded. This was done after the division of co-occurrence data described in the previous paragraph. The nature of that division ensured that even after this further reduction each of the eight co-occurrence contrasts remained balanced. (For Experiment 2, each F ~ D/E contrast now had slightly more total observations than each G ~ D/E contrast, but this presents no problems of interpretation for the model.)

References

Akhtar, N., & Tomasello, M. (1997). Young children's productivity with word order and verb morphology. *Developmental Psychology, 33*, 952–965.

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

Baker, C. F., Fillmore, C. J., & Lowe, J. B. (1998). The Berkeley FrameNet project. *Proceedings of COLING/ACL*, 86–90.

Barr, R. A., & Caplan, L. J. (1987). Category representations and their implications for category structure. *Memory and Cognition, 15*, 397–418.

Barsalou, L. W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 101–140). New York: Cambridge University Press.

Bassok, M., & Medin, D. L. (1997). Birds of a feather flock together: similarity judgments with semantically rich stimuli. *Journal of Memory and Language, 36*, 311–336.

Block, N. (1987). Functional role and truth conditions. *Proceedings of the Aristotelian Society LXI*, 157–181.

Bock, J. K. (1986). Syntactic persistence in language production. *Cognitive Psychology, 18*, 355–387.

Bock, J. K., Loebell, H., & Morey, R. (1992). From conceptual roles to structural relations: Bridging the syntactic cleft. *Psychological Review, 99*, 150–171.

Boroditsky, L. (in press). Comparison and the development of knowledge. *Cognition*.

Bowdle, B. F., & Gentner, D. (2005). The career of metaphor. *Psychological Review, 112*, 193–216.

Burgess, C. (1998). From simple associations to the building blocks of language: modeling meaning in memory with the HAL model. *Behavior Research Methods, Instruments, and Computers, 30*, 188–198.

Carlson, G. N., & Tanenhaus, M. K. (1988). Thematic roles and language comprehension. In W. Wilkins (Ed.), *Thematic relations* (pp. 263–288). New York: Academic Press.

Carreras, X., & M'arques, L. (2004). Introduction to the CoNLL-2004 shared task: Semantic role labeling. In H.T. Ng & E. Riloff (Eds.), Proceedings of CoNLL-2004. Boston.

Chang, F., Bock, K., & Goldberg, A. E. (2003). Can thematic roles leave traces of their places?. *Cognition 90*, 29–49.

Choplin, J. M., & Hummel, J. E. (2002). Magnitude comparisons distort mental representations of magnitude. *Journal of Experimental Psychology: General, 131*, 270–286.

Collins, A. M., & Loftus, E. E. (1975). A spreading activation theory of semantic processing. *Psychological Review, 82*, 407–428.

Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior, 8*, 240–247.

Diessel, H., & Tomasello, M. (2001). The acquisition of finite complement clauses in English: a corpus-based analysis. *Cognitive Linguistics, 12*, 97–141.

Dowty, D. R. (1991). Thematic proto-roles and argument selection. *Language, 67*, 547–619.

Estes, Z. (2003). A tale of two similarities: Comparison and integration in conceptual combination. *Cognitive Science, 27*, 911–921.

Ferretti, T. R., McRae, K., & Hatherell, A. (2001). Integrating verbs, situation schemas, and thematic role concepts. *Journal of Memory and Language, 44*, 516–547.

Field, H. (1977). Logic, meaning and conceptual role. *Journal of Philosophy, 69*, 379–408.

Fillmore, C. (1968). The case for case. In E. Bach & R. T. Harms (Eds.), *Universals in linguistic theory* (pp. 1–90). New York: Holt, Rinehart, & Winston.

Fillmore, C. J. (1976). Frame semantics and the nature of language. *Annals of the New York Academy of Sciences: Conference on the Origin and Development of Language and Speech, 280*, 20–32.

Flores d'Arcais, G. B., Schreuder, R., & Glazenborg, G. (1985). Semantic activation during recognition of referential words. *Psychological Research, 47*, 39–49.

Fodor, J. A. (1998). *Concepts: Where cognitive science went wrong*. New York: Oxford University Press.

Gentner, D. (1975). Evidence for the psychological reality of semantic components: The verbs of possession. In D. A. Norman & D. E. Rumelhart (Eds.), *Explorations in Cognition* (pp. 211–246). San Francisco: W.H. Freeman.

Gentner, D. (1983). Structure-mapping: a theoretical framework for analogy. *Cognitive Science, 7*, 155–170.

Gentner, D., & Brem, S. K. (1999). Is snow really like a shovel? Distinguishing similarity from thematic relatedness. In M. Hahn & S. C. Sontess (Eds.), *Proceedings of the Twenty-first Annual Meeting of the Cognitive Science Society* (pp. 179–184). Mahwah, NJ: Lawrence Erlbaum Associates.

Gentner, D., & Gunn, V. (2001). Structural alignment facilitates the noticing of differences. *Memory and Cognition, 29*, 565–577.

Gentner, D., & Kurtz, K. J. (2005). Learning and using relational categories. In W. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. Wolff (Eds.), *Categorization inside and outside of the lab: Festschrift in Honor of Douglas L. Medin*. Washington, DC: American Psychological Association.

Gentner, D., Rattermann, M. J., & Forbus, K. D. (1993). The roles of similarity in transfer: separating retrievability from inferential soundness. *Cognitive Psychology, 25*, 524–575.

Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science, 10*, 277–300.

Gildea, D., & Jurafsky, D. (2002). Automatic labeling of semantic roles. *Computational Linguistics, 28*(3), 245–288.

Gregory, M. E., & Mergler, N. L. (1990). Metaphor comprehension: In search of literal truth, possible truth, and metaphoricity. *Metaphor and Symbolic Activity, 5*, 151–173.

Goldstone, R. L. (1996). Alignment-based nonmonotonicities in similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22*, 988–1001.

Goldstone, R. L., Medin, D. L., & Gentner, D. (1991). Relational similarity and the nonindependence of features in similarity judgments. *Cognitive Psychology, 23*, 222–262.

Goldstone, R. L., & Rogosky, B. J. (2002). Using relations within conceptual systems to translate across conceptual systems. *Cognition, 84*, 295–320.

González-Vallejo, C. (2002). Making trade-offs: a probabilistic and context-sensitive model of choice behavior. *Psychological Review, 109*, 137–154.

Hare, M., McRae, K., & Elman, J. L. (2003). Sense and structure: meaning as a determinant of verb subcategorization preferences. *Journal of Memory and Language, 48*, 281–303.

Hare, M., McRae, K., & Elman, J. L. (2004). Admitting that admitting verb sense into corpus analyses makes sense. *Language and Cognitive Processes, 19*, 181–224.

Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives: violations of regularity and the similarity heuristic. *Journal of Consumer Research, 9*, 90–98.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: a theory of analogical access and mapping. *Psychological Review, 104*, 427–466.

Hutchison, K. A. (2003). Is semantic priming due to association strength or feature overlap?. A microanalytic review. *Psychonomic Bulletin and Review, 10*, 785–813.

Jones, M., & Sieck, W. R. (2003). Learning myopia: an adaptive recency effect in category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29*, 626–640.

Kahneman, D., & Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica, 47*, 263–291.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: the latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review, 104*, 211–240.

Larkey, L. B., & Love, B. C. (2003). CAB: connectionist analogy builder. *Cognitive Science, 27*, 781–794.

Larkey, L. B., Narvaez, L., & Markman, A. B. (2004). Categories among relations. In K. D. Forbus, D. Gentner, & T. Regier (Eds.), *Proceedings of the twenty-sixth annual meeting of the cognitive science society* (p. 1634). Hillsdale, NJ: Lawrence Erlbaum.

Lin, E. L., & Murphy, G. L. (2001). Thematic relations in adults' concepts. *Journal of Experimental Psychology: General, 130*, 3–28.

Love, B. C., & Markman, A. B. (2003). The nonindependence of stimulus properties in human category learning. *Memory and Cognition, 31*, 790–799.

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

Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments and Computers, 28*, 203–208.

Markman, A. B. (2005). What are categories, and why are they coherent? In W. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. Wolff (Eds.), *Categorization inside and outside of the lab: Festschrift in Honor of Douglas L. Medin*. Washington, DC: American Psychological Association.

Markman, A. B., & Gentner, D. (1996). Commonalities and differences in similarity comparisons. *Memory and Cognition, 24*, 235–249.

Markman, A. B., & Stilwell, C. H. (2001). Role-governed categories. *Journal of Experimental and Theoretical Artificial Intelligence, 13*, 329–358.

Markman, E. M., & Hutchinson, J. E. (1984). Children's sensitivity to constraints on word meaning: taxonomic versus thematic relations. *Cognitive Psychology, 16*(1), 1–27.

Masson, M. E. J. (1995). A distributed memory model of semantic priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 3–23.

McDonald, S., & Lowe, W. (1998). Modelling functional priming and the associative boost. In Proceedings of the 20th Annual Conference of the Cognitive Science Society, pp. 675–680. Mahwah, NJ: Erlbaum.

McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavioral Research Methods, Instruments, and Computers, 37*, 547–559.

McRae, K., Ferretti, T. R., & Amyote, L. (1997). Thematic roles as verb-specific concepts. *Language and Cognitive Processes, 12*, 137–176.

McRae, K., Hare, M., Elman, J. L., & Ferretti, T. R. (2005). A basis for generating expectancies for verbs from nouns. *Memory and Cognition, 33*, 1174–1184.

McRae, K., Spivey-Knowlton, M. J., & Tanenhaus, M. K. (1998). Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language, 38*, 283–312.

Medin, D. L., Goldstone, R. L., & Gentner, D. (1990). Similarity involving attributes and relations: judgments of similarity and difference are not inverses. *Psychological Science, 1*, 64–69.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85*, 207–238.

Miller, S., Fox, H., Ramshaw, L., & Weischedel, R. (2000). A novel use of statistical parsing to extract information from text. In Proceedings of the 1st Annual Meeting of the North American Chapter of the ACL (NAACL), 226–233.

Moss, H. E., Ostrin, R. K., Tyler, L. K., & Marslen-Wilson, W. D. (1995). Accessing different types of lexical semantic information: Evidence from priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 863–883.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92*, 289–316.

von Neumann, M., & Morgenstern, O. (1947). *Theory of games and economic behavior* (2nd ed.). Princeton, NJ: Princeton University Press.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115*, 39–57.

Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review, 101*, 53–79.

Nosofsky, R. M., & Zaki, S. R. (2002). Exemplar and prototype models revisited: response strategies, selective attention, and stimulus generalization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*, 924–940.

Pecher, D., Zeelenberg, R., & Raaijmakers, J. G. W. (1998). Does pizza prime coin? Perceptual priming in lexical decision and pro-nunciation. *Journal of Memory and Language, 38*, 401–418.

Pickering, M. J., & Branigan, H. P. (1998). The representation of verbs: evidence from syntactic priming in language production. *Journal of Memory and Language, 39*, 633–651.

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology, 77*, 353–363.

Prior, A., & Bentin, S. (2003). Incidental formation of episodic associations: the importance of sentential context. *Memory and Cognition, 31*, 306–316.

Quiggin, J. (1993). *Generalized expected utility theory: The rank-dependent model*. Boston: Kluwer Academic.

Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology, 3*, 382–407.

Rehder, B., & Ross, B. H. (2001). Abstract coherent categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*, 1261–1275.

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: a dynamic connectionist model of decision making. *Psychological Review, 108*, 370–392.

Rosch, E., & Mervis, C. B. (1975). Family resemblances: studies in the internal structure of categories. *Cognitive Psychology, 7*, 573–605.

Rumelhart, D. E., & Levin, J. A. (1975). A language comprehension system. In D. A. Norman & D. E. Rumelhart (Eds.), *Explorations in cognition* (pp. 179–208). San Francisco: W.H. Freeman.

Schreuder, R., Flores d'Arcais, G. B., & Glazenborg, G. (1984). Effects of perceptual and conceptual similarity in semantic priming. *Psychological Research, 45*, 339–354.

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science, 237*, 1317–1323.

Shepard, R., Hovland, C. L., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs: General and Applied, 75*, 1–42.

Simonson, I., & Tversky, A. (1992). Choice in context: tradeoff contrast and extremeness aversion. *Journal of Marketing Research, 29*, 281–295.

Smiley, S. S., & Brown, A. L. (1979). Conceptual preference for thematic or taxonomic relations: a nonmonotonic age trend from preschool to old age. *Journal of Experimental Child Psychology, 28*(2), 249–257.

Smith, J. D., & Minda, J. P. (2002). Distinguishing prototype-based and exemplar-based processes in dot-pattern category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*, 800–811.

Spellman, B. A., Holyoak, K. J., & Morrison, R. G. (2001). Analogical priming via semantic relations. *Memory and Cognition, 29*, 383–393.

Stewart, N., Brown, G. D. A., & Chater, N. (2002). Sequence effects in categorization of simple perceptual stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*, 3–11.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology, 53*, 1–26.

Stewart, N., Chater, N., Stott, H. P., & Reimers, S. (2003). Prospect relativity: how choice options influence decision under risk. *Journal of Experimental Psychology: General, 132*, 23–46.

Tomasello, M. (1992). *First verbs: A case study of early grammatical development*. Cambridge: Cambridge University Press.

Tomlinson, M., & Love, B. C. (2006). From pigeons to humans: Grounding relational learning in concrete examples. *Proceedings of the Twenty-first National Conference on Artificial Intelligence (AAAI-06)*. pp. 199–204.

Tversky, A. (1977). Features of similarity. *Psychological Review, 84*, 327–352.

Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management Science, 39*, 1179–1189.

Wisniewski, E. J., & Bassok, M. (1999). What makes a man similar to a tie? Stimulus compatibility with comparison and integration. *Cognitive Psychology, 39*, 208–238.

Wisniewski, E. J., Clancy, E. J., & Tillman, R. N. (2005). On different types of categories. In W. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. Wolff (Eds.), *Categorization inside and outside of the lab: Festschrift in Honor of Douglas L. Medin*. Washington, DC: American Psychological Association.

Wisniewski, E. J., & Love, B. C. (1998). Relations versus properties in conceptual combination. *Journal of Memory and Language, 38*, 117–202.

Wittgenstein, L. (1953). *Philosophical Investigations (Trans. G.E.M. Anscombe)*. New York: Macmillan.