DRAFT – Comments welcome

Genter, D & Kurtz, K. J. (in press). Relational Categories. To appear in Ahn, W.K., Goldstone, R.L., Love, B.C., Markman, A.B., & Wolff, P.W. (Eds.), *Categorization inside and outside the lab*. Washington, DC: American Psychological Association.

Relational Categories

Dedre Gentner

Department of Psychology

Northwestern University

Kenneth J. Kurtz

Department of Psychology

State University of New York at Binghamton

address correspondence to:

Dedre Gentner
Department of Psychology
Northwestern University
2029 Sheridan Road
Evanston, IL 60208
e-mail: gentner@northwestern.edu

Abstract

This work is concerned with the acquisition and use of *relational categories*—categories such as *barrier*, *gift*, and *enemy*, whose membership is determined by the relations they state. We suggest that these kinds of categories, though relatively neglected in the field, are many in number, frequent in use, and important in cognitive processing. We review their patterns of acquisition and use and show that they differ markedly from standard entity categories in a number of ways. Finally, we argue that the distinction between relational categories and entity categories is a continuum, not a dichotomy, and that the framework developed for relational categories applies to standard categories as well.

Learning and Using Relational Categories

This work is concerned with the acquisition and use of relational categories. By *relational category* we mean a category whose membership is determined by a common relational structure, rather than by common properties. For instance, for X to be a *bridge*, X must connect two other entities or points; for X to be a *carnivore*, X must eat animals. Relational categories contrast with entity categories like *tulip* or *camel*, whose members share many intrinsic properties. Relational categories cohere on the basis of a core relationship fulfilled by all members. This relation may be situation-specific (e.g., *passenger* or *accident*) or enduring (e.g., *carnivore* or *ratio*). Relational categories abound in ordinary language. Some are restricted in their arguments: for example, carnivores are animals who eat other animals. But for many relational categories, the arguments can range widely: for example, a *bridge* can connect two concrete locations, or two generations, or two abstract ideas. As with *bridge*, the instances of a relational category can have few or no intrinsic properties in common with one another.

Research on categories has mostly ignored relational categories, focusing instead on entity categories—categories that can be characterized in terms of intrinsic similarity among members, like those shown in Figure 1. Further, as Kloos and Sloutsky (2004) point out, theories of categorization have often operated under the assumption that all concepts are fundamentally alike. However, as Medin and his colleagues (Medin, Lynch, & Coley, 1997; Medin, Lynch, & Solomon, 2000) have argued, categories are not uniform in character, and the variations support a range of different functions. In this paper we contrast relational categories—categories whose members satisfy a specified relational structure—with entity categories—categories whose members have highly overlapping intrinsic features and feature correlations.

------------Figure 1 about here.-------------------

The exceedingly minor role of relational categories in the empirical literature might suggest that they are rare or unimportant in human psychology, but just the opposite is true. Informal ratings of the 100 highest frequency nouns in the British National Corpus revealed that about half are relational nouns (Asmuth & Gentner, in preparation). We suspect that for most of us, sentences like (2) (containing relational nouns) occur as or more often than sentences like (1) (containing entity nouns):

(1) The dog chased a rat into the shed.

(2) The investigation showed that the plan was a mistake from the start.

A simple exercise shows the importance of relational categories in communication. Try, without using any relational nouns, to paraphrase the following sentence (or sentence (2) above):

(3) His allies vowed payback for the betrayal.

Even if we allow verbs and prepositions (themselves, of course, relational terms) any such attempt is at best unwieldy, if not impossible.

Relational categories can be divided into *relational role categories* (or *role categories*) and *relational schema categories* (or *schema categories*). *Schema categories* are names of relational systems: e.g., *robbery*, *betrayal*, *slavery* and *sister*. *Role categories* are categories whose members all play the same role in a relational schema.[1] For example, *robbery* is a relational schema category with three arguments, each of which is a relational role category:

robbery (*thief, goods, victim*).

The three relational role categories are *thief* (agent who steals), *goods* (the things transferred), and *victim* (the one stolen from). Not all the relational roles have to be specified on a given occasion. For example, in the sentence “Bush stole the election” the thief and goods-stolen categories are given, but the victim is left unmentioned.

We have noted that entity categories are characterized by high intrinsic similarity among the members, while relational categories are characterized by fit to a core relational structure. Along with this goes an important distinction between dense (or rich) and sparse representations (Gentner, 1981; Kloos & Sloutsky, 2004). The members of a given entity category are characterized by richly interconnected feature structure. The high intrinsic similarity among members is a natural consequence of the fact that they share many features and feature correlations. In contrast, because the members of a given relational category share only a sparse relational structure, there may be no obvious intrinsic similarities among members[2]. Entity categories can be thought of as first-order partitions of the world (Gentner, 1982) and relational categories as second-order ways of organizing and linking those first-order partitions.

To summarize, e*ntity categories*—such as *cow, tulip, and radish*—are characterized by rich sets of intrinsic features and feature correlations. Relational categories are characterized by sparse, rule-like relational structures. *Relational role categories* —such as *carnivore* and *robber*—are characterized by their extrinsic relation to a schema; *relational schema categories* —such as *bridge* and *robbery*—are characterized by internal relational structures that take external arguments.

Nouns, Verbs, Entities and Relations

One way to begin exploring the contrast between entity categories and relational categories is to consider some known contrasts between nouns and verbs. There are some obvious parallels between relational noun categories and the categories named by verbs. Relational nouns, like verbs and prepositions, have meanings that are centered around relations with other concepts. Relational nouns are also similar to verbs in that they take arguments (they are semantically unsaturated). For example, the relational schema noun *robbery* denotes a system of relations among the arguments (*thief*, *goods*, *victim*), much as does the verb *to rob.* There is a clear conceptual affinity (and often a morphological relation as well) between such relational nouns and their related verbs. All this suggests a speculative analogy:

Relational nouns : Entity nouns :: Verbs : Nouns

This analogy is of course only partial, because the contrast between nouns and verbs involves a shift in syntactic as well as semantic category. Also, it applies more naturally to a relational schema category like *robbery* (which, verblike, denotes a system of relations) than to a relational role category like *thief*[3] (which invites a relational schema of which it is one argument).

Given these parallels, it is worth exploring whether noun-verb contrasts can illuminate relation-entity contrasts. Gentner (1981, 1982) laid out a set of phenomena that differentiate verbs from concrete nouns, summarized in Table 1. Verbs are slower to be acquired than nouns (Caselli et al., 1995; Gentner, 1982; Gentner and Boroditsky, 2001); poorer in memory than nouns, both in recognition and in recall (e.g., Kersten & Earles, in press); more mutable in meaning under semantic strain (Gentner & France, 1988); less prone to be borrowed in language contact (e.g., Sobin, 1982); and less stable in translation between languages than nouns (Gentner, 1981). Verbs are also more polysemous than nouns at a given word frequency.

For example, Webster's Dictionary (1961) shows a mean of 7.3 word senses for the twenty highest-frequency nouns and a mean of 12.4 word senses for the twenty highest-frequency verbs.

----Table 1 about here ------------

This comparison leads us to something of a paradox. The above set of contrasts leads to a general portrait of verbs as poor in memory and variable in meaning relative to nouns. Yet verbs are considered syntactically and semantically central in sentence structure. The verb is often described as the core of a sentence, in that it conveys the central set of events and relations in which the nouns participate. Chafe (1970, pp. 97-98) analogized the verb and nouns to the sun and planets, respectively: "anything which happens to the sun affects the entire solar system," whereas "a noun is like a planet whose internal modifications affect it alone, and not the solar system as a whole." How can verbs be so central and yet so elusive?

Part of the answer lies in the syntactic role of the verb, but another part lies in semantic structure. Gentner (1981) proposed that a distinction between dense (or rich) and sparse representations underlies many of the differences between concrete nouns and verbs. Gentner proposed that concrete object-concepts have greater conceptual density (that is, more links between their internal components) than do similarly basic relational concepts.. The greater conceptual density of entity concepts contributes to their perceptual availability and thus to their early acquisition. Because of their interconnected, tightly-packed structure, such entities are naturally individuated via experience with the world (the Natural Partitions hypothesis) (Gentner, 1982; Gentner & Boroditsky, 2001). A corollary assumption is that the ratio of internal links to external links is higher for object concepts than for relational concepts, leading to greater stability for object concepts in conceptual combination (Gentner & France, 1988).

One corollary of the fact that verbs have sparse representations is that the meanings of verbs—even “concrete” verbs like motion verbs—vary more cross-linguistically do the meanings of concrete nouns (Gentner’s (1982) Relational Relativity hypothesis). Verb meanings include only part of the available relational information, and just *which* information is selected varies across languages (e.g., Bowerman & Choi, 2003; Levinson, 1996; Slobin, 1996). For example, Talmy’s (1975) seminal research showed that languages differ in which semantic elements are incorporated into motion verbs: the path of the moving figure (as in Spanish); the manner of its motion (as in English) or the shape of the figure (as in Atsugewi).

This line of theorizing predicts that names for entities (especially animate beings, which are especially individuable) should be learned earlier than relational terms cross-linguistically (Gentner, 1982; Gentner & Boroditsky, 2001). The prelinguistic infant has already individuated many entities, and has only to attach the word to the referent. But relational terms such as verbs and prepositions pose a greater challenge. Their referents are not simply “out there” in the experiential world; they are linguistically selected. To learn what a verb means, the child must discover which aspects of the situation enter into its meaning in her language.

Consistent with this prediction, children readily take novel words as names for whole objects (Markman, 1989), even as early as 13 months of age (Waxman & Markow, 1995). Cross-linguistically, nouns predominate over verbs in children’s early production and comprehension[4] as evidenced by findings from English, Italian, Japanese, Korean, Mandarin, and Navajo (Gentner, 1982; Gentner & Boroditsky, 2001, in preparation; Goldin-Meadow, Seligman, & Gelman, 1976; Imai, Haryu, & Okada, in preparation; Tardif, Gelman & Xu, 1999). Research by Gillette, Gleitman, Gleitman and Lederer (1999) demonstrates the greater ease of learning concrete noun referents from observation. They showed adults silent videos of mothers talking to young children, with beeps marking the instances of a particular noun or verb. Their task was to guess the word the mother had uttered at the beeps. After six different instances of a given word, they were able to guess correctly 45% of the time for nouns, but only 15% of the time for verbs. Thus pure observational learning was far more effective for deriving noun meanings than verb meanings.

With this background, we now return to our analogy: *relational categories : entity categories :: verbs : concrete nouns.* Table 2 lays out the implications of this analogy. Although our purpose here is to stimulate further research, some parts of the table do have confirming evidence. Preliminary research with Jennifer Asmuth suggests that relational categories are more mutable and less well remembered than entity categories (Asmuth & Gentner, in preparation). Further, as we will discuss, relational categories are slower to be acquired than entity categories.

---------------------Table 2 about here ---------------------

Prior Research on Relational Categories

The contrast between entity categories and relational categories has some precedents. Prior work has distinguished between perceptual features and functional features (Bruner, Greenfield, & Olver, 1966; Miller & Johnson-Laird, 1976) and between concrete and abstract features (e.g., Paivio, 1971). Barr and Caplan (1987) explored the distinction between intrinsically represented and extrinsically represented categories. They defined an *intrinsic feature* as one that is true of an entity considered in isolation (such as “has wings” for a bird) and an *extrinsic feature* as one expressing the relationship between two or more entities (for example, “used to work with” for a hammer). Using a variety of converging measures, they divided their 13 categories into an intrinsically based set (in our terms, entity categories)—

mammals, birds, flowers, fruit, vegetables and trees—and an extrinsically based set (relational categories)—weapons, vehicles, furniture, toys, tools, and sports—with one intermediate category (clothing). Barr and Caplan found that relational categories showed less agreement across categories on membership judgments, a finding echoed in the Kurtz and Gentner (2001) study described later.

Goldstone and colleagues (1996; Goldstone, Steyvers, & Rogosky, 2003) proposed a continuum from categories understood well in isolation (*entity categories*, in our terminology) to those that are highly dependent on other concepts for their meaning (*relational categories*, in our terms). Two signatures of relational concepts are relative insensitivity to nondiagnostic category features (consistent with the rule-like flavor of relational categories) and better categorization performance for caricatured instances that emphasize differences between neighboring categories. Rehder & Ross (2001) investigated the relational basis of what they termed *abstract coherent categories* and found that such categories can be acquired based on relationships that are orthogonal to the specific attributes, as long as the relationships arise in a manner consistent with prior expectations. Such categories lack feature overlap, but take their structure from systems of features that are understood to support a common abstract relationship.

Kloos and Sloutsky (2004) made a related distinction between *natural kind* concepts like *bird*, which have rich correlational structure, and *nominal kind* concepts like *acceleration*, which are based on a sparse rule structure, drawing on Gentner's (1981) distinction between dense and sparse concepts. Kloos & Sloutsky investigated the learning of entity and relational categories using an artificial world. Members of entity categories shared a set of correlated features, whereas members of relational categories shared a single relation. The task was to learn a category either by observation (i.e., by being presented with many instances of the category) or through a rule-like definition (e.g., “All members of the category have relation X”). The results suggest a dissociation: entity categories were learned better by observation than were relational categories (which were extremely poorly learned in this condition), whereas relational categories were learned better than entity categories when an explicit rule was given. These results support the hypothesis that a difference in representational density may underlie some of the psychological distinctions between the two kinds of categories.

Another important precedent is Markman and Stilwell’s (2001) paper in which they distinguish four kinds of categories: feature-based categories, relational categories, role-governed categories, and transformational categories. *Feature-based* categories correspond to our entity categories, and *role-governed* categories to our relational role categories. Markman and Stilwell define *relational categories* (including verb categories, event scripts, preposition categories and comparatives) as those that pick out relations in the environment (Gentner & Boroditsky 2001). For example, the verb *plays (x, y)* is a relational category with two associated relational role categories: *player* and *game.* Finally, they define a *transformational category* as one that specifies a change in the selectional restriction for a relation: for example, *team* is used to transform a group to an individual so that it can serve as an argument to relations that require an individual.

We agree with Markman and Stilwell’s conceptual distinction between relational role categories and relational schema categories. However, we see the transformational criterion, whereby naming a complex system allows it to serve as the argument of other predicates, as an important characteristic of relational categories in general. This property—that relational categories allow relations to serve as arguments to other predicates (as in "damage prevention" or "reversal of fortune")—is a source of great expressive power in human language.

Another important line of related work is Barsalou's (1983, 1985) investigation of goal-derived categories, such as *things to take out of the house in case of fire*. We see these as relational role categories. Their members typically lack intrinsic similarity; for example, the above category can include a pet cat, family photos, a checkbook, a computer, etc. The only commonality the members need have is that to serve as "something you value highly." Barsalou found that goal-derived categories—such as *foods not to eat when on a diet*—show graded structure around ideals (properties that optimally promote goal resolution) rather than around central tendency. Again, the centrality of a specific goal is consistent with the relatively sparse, rule-like nature of relational category representations.

Another class of ad hoc categories in Barsalou's work is *thematic groupings*—items that coalesce around some common idea or activity. Because thematic groupings can be confused with relational categories, we take a moment to clarify. An example of a thematic grouping is *things associated with going to a movie*: e.g., ticket, popcorn, program, a date, and so on. These things all play different roles in the movie-going schema. They coalesce by virtue of covering all or most of the roles in the schema. In contrast, in a true relational role category such as *ticket*, the members always play the same role. Whether a ticket is for a movie, a concert, or a bus trip, it serves as a token of entitlement to attend the event. Just as the members of intrinsic categories share properties, the members of relational categories share specific relations. For relational role categories like *ticket*, the shared relation is participation in the same role in a schema. For relational schema categories like *going to a movie*, the shared information is a common set of interconnected relations. But for thematic categories like *things associated with* *going to a movie*, there is no commonality beyond participating somehow in the movie schema. Because thematic groupings are associational, and not characterized by commonalities, we do not include them among relational categories.

Entity and Relational Categories in Adult Processing

Kurtz and Gentner (2001; in preparation) compared traditional entity categories with relational categories using an exemplar generation task[5]. Our goals were to seek evidence for the intrinsic-extrinsic distinction and to discover signatures of underlying organization among different kinds of categories. Participants were asked to generate as many examples as they could in a four-minute interval for each category cue. The relational categories were *trap*, *weapon*, *guide*, *signal*, *barrier*, *tool*, *filter*, and *shield*. The entity categories were *animal*, *plant*, *vegetable*, *fruit*, *vehicle*, *household appliance*, *type of dwelling*, and *musical instrument*.[6] We expected relational categories to be less fluent, less generative, and less consistent than taxonomic categories. As noted above, the instances of a relational category may have little or no intrinsic similarity to each other. Extrapolating from the general finding that purely relational similarity is more difficult to access in memory than rich object similarity, we expected far less inter-item reminding for relational categories than for entity categories (Gentner et al, 1993; Holyoak & Koh, 1995).

Table 3 shows some example generation lists. As predicted, the mean number of responses produced was much higher for entity categories (*M*=23.2) than for relational categories (*M*=14.2). People were also much faster to generate members of entity categories than of relational categories: the mean interresponse time was 6.4 sec. for entity categories and 10.8 sec. for relational categories. There was also more agreement among participants on the exemplars of entity categories. Finally, independent ratings of inter-item similarity showed higher item similarity within entity categories than within relational categories. This pattern is consistent with our claim that entity categories, but not relational categories, are based on intrinsic similarity.

-------Table 3 about here ----------

Not only were people less fluent at generating members of relational categories, they were also less able to guess the category given the examples. When new participants were given sets of high-consensus category exemplars and asked to state what the items had in common, they were much more likely to arrive at the initial category (or a synonymous description) for entity categories than for relational categories. In sum, entity categories showed greater fluency in exemplar-generation, greater interterm similarity, and greater category transparency given the exemplars than did relational categories.

At this point, it may appear that relational categories are the lame duck of the category world. Indeed, in the next section we will add to the list of their handicaps when we discuss their acquisition by children. However, two further observations from the Kurtz and Gentner study sounds a different note. First, we observed informally that there were many more creative entries for the relational categories than for the entity categories. For example, after listing standard exemplars like *wall* and *fence* for the category *barrier*, people went on to list *poverty* and *lack of education*. These analogical extensions are reminiscent of the patterns of metaphorical extension seen for relational terms in natural language. Second, in an additional pilot study, we asked participants to list the categories to which a single example (*dog*) could belong (See Table 4). Here the fecundity of entity and relational categories is reversed: whereas a given exemplar might belong to six or seven entity categories (themselves hierarchically related), it could easily belong to dozens of relational categories. Further, these relational categories were not typically hierarchically related, nor were they generally mutually exclusive; they were simply applicable relational structures.

-------Table 4 about here ----------

The pattern shown in Table 4 points up another contrast in conceptual structure between relational categories and entity categories. Relational categories, like verbs, appear to exist in multiply branching bushy thickets, rather than in strict taxonomies. Table 4 also underscores the large number of relational categories to which even a single exemplar can belong. We noted earlier that there are about as many relational categories as entity categories in everyday language. The large number of relational categories in ready use suggests that despite their lack of transparency they have considerable utility in thought and language.

Relational categories are in many ways at the opposite pole from entity categories. Entity categories (especially basic-level categories such as familiar animals) sort experience into mutually exclusive categories in clear vertical taxonomies. Just the opposite is true for relational categories. You can't be both a wolf and a sheep, but you can easily be both predator and prey. Again, whereas the *entity* superordinates for *wolf* are clearly delimited and vertically arranged (*canine*, *mammal*, *animal*, *life form*), its *relational* superordinates are abundant in number and cross-cutting in structure (*carnivore, predator*, *prey of cougar*, *eater of sheep*, *bane of ranchers*, *endangered species*, *member of pack*, *cousin of dogs*, and so on). Entity categories afford fast, fluent categorization of their members (especially basic-level categories). As we have seen, relational categories do not; even when given a *set* of members, people find it difficult to identify the relational category they belong to. Entity categories partition the experiential world; relational categories provide relations among these partitions.

Relational Categories in Acquisition

Research on language acquisition bears on relational categories in two ways: first, as noted earlier, it shows that verbs and prepositions are learned more slowly than concrete nouns (Gentner, 1982); second, a very small set of studies has compared the acquisition of entity and relational (although, as in studies of adult categorization, this research has centered on entity categories). The available evidence paints a consistent picture: relational concepts are acquired later than entity concepts (Gentner & Rattermann, 1991; Gentner & Medina, 1988). First, there is considerable evidence that infants readily become sensitive to basic-level entity categories (Jusczyk, 1997; Quinn et al., 2002). Second, as noted above, children's early word learning appears geared towards object-based categories (e.g., Gentner, 1982; Golinkoff & Hirsh-Pasek, 1990; Imai, et al., in preparation; Markman, 1989; Waxman & Markow, 1995). Young children tend to take a new word as the name of an object, and to extend the term to similar entities, often to items with intrinsic perceptual commonalities such as common shape (Baldwin, 1989; Gentner, 1977; Imai, Gentner, & Uchida, 1994; Smith, Landau & Jones, 1992). A third indication that entity nouns enter the vocabulary before relational nouns comes from examining the MacArthur Communicative Developmental Inventory, which serves as a reasonable upper-bound estimate of what children might know at a given age. For 8-16-month-olds the MCDI has 296 nouns, of which 93% are entity nouns (objects, animals and people) and 7% are mixed entity-relational nouns (there are no purely relational nouns). For 17-30-month-olds there are 411 nouns, of which 79% are entity nouns, 13% are mixed, and 8% are relational nouns.

Another indication of the greater learnability of entity categories is that when children do learn relational terms, they often initially treat them as entity terms (Gentner & Rattermann, 1991). For example, a *brother* may be described as a boy about 12 years old, rather than as any male (however young or old) who is someone's sibling, or an *uncle* as a nice man with a pipe rather than any male in a sibling relationship with one's mother or father (E. Clark, 1993; Keil & Batterman, 1984). Likewise, Keil and Batterman found that 4-year-olds conceive of an *island* as a place with sand and palms. Only later do they learn the relational descriptors—e.g., that an *island* is *any* subcontinental body of land surrounded by water. Hall and Waxman (1993) found that 3 ½ -year-olds had difficulty learning novel relational nouns denoting concepts like *passenger*. Even when they were explicitly told (for example) "This one is a *blicket* BECAUSE IT IS RIDING IN A CAR," children tended to interpret the novel noun as referring to the object category and to extend it to a similar-looking doll rather than to another doll riding in a vehicle.

Learning Relational Categories

Given all these difficulties, one might wonder how relational categories are ever acquired. The answer lies partly in the process of analogy. There is a substantial body of research showing that structure-mapping processes act to highlight common relational systems—*systems of interconnected knowledge* linked by higher-order causal, mathematical, or perceptual relations (Clement & Gentner, 1991; Loewenstein, Thompson, & Gentner, 1999) Analogical processes—i.e., structural alignment and mapping—are crucial in the learning of relational categories (Gentner, 1983; Gentner, 2003). But if comparison is a key process in learning the meanings of relational terms, then the next question is how do such comparisons come about? One way is by progressive alignment: by experiencing juxtapositions that progress from highly similar pairs to pairs with less surface similarity, and eventually to pairs that differ strongly in their entity characteristics but share relational patterns.

A study by Kotovsky and Gentner (1996) shows how such comparison processes can promote the learning of relational categories. The matches were based on higher-order perceptual relations such as *symmetry* or *monotonic increase* (Chipman & Mendelson, 1979; Halford, 1987). Children were given a perceptual matching task—e.g., matching a standard oOo with alternatives xXx vs. Xxx. Four year-olds could readily choose the similar alternative when the dimensions matched—that is, when both the standard and the alternatives varied on either size or shading—but they were at chance on cross-dimensional triads. They apparently saw no likeness between *little-big-little* and *light- dark-light*. Then a new group of 4-year-olds was given the same sixteen triads (also in a free similarity task without feedback) with one crucial difference: these children received all the within-dimension triads first (grouped so that similar triads were juxtaposed) followed by the cross-dimension triads. The results were dramatically different. Children who first received the relatively concrete same-dimension matches were subsequently able to see the purely relational match between the cross-dimensional figures—e.g., *little-big-little* with *light-dark-light.* Kotovsky and Gentner suggested that blocking the within-dimension pairs provided concentrated experience with close (literally similar) matches. On these matches, children experienced virtually foolproof close alignments that repeatedly yielded the same relational structure. This repeated experience of the relational structure (symmetry or monotonicity)—a kind of *progressive alignment* experience—prepared children to see the common higher-order relational pattern when it appeared across dimensions.

This brings us to a key aspect of the process: re-representation. We suggest that children initially represented the relations in a dimensionally specific fashion: that is, the difference in magnitude was conflated with the dimension of difference. Their representations were roughly *bigger* (x, y) and *darker* (a, b), and these of course do not match. We suggest that the within-dimension comparisons made the higher-order pattern of *symmetry* sufficiently salient to be noticed in a cross-dimensional triad. Comparison is a structurally greedy process;

one match invites another. In SME, when a relation matches, the process tries to put its arguments into correspondence (Falkenhainer, Forbus, & Gentner, 1989). Thus experiencing a higher-order match invites an alignment among the lower-order relations. If they don't match, this is an invitation to try re-representing the lower-order relations: for example, to separate the *change-of-magnitude* relation from the specific dimension of change. In this way the nonmatching relations (*bigger* (x, y) and *darker* (a, b)) are reconstrued as partially matching relations—*greater* (*size*[x], *size*[y]) and *greater* (*shading*[a], *shading*[b]). The (unlike) dimensions can then be put into correspondence by virtue of their like relational roles (See Forbus et al. ,1989; Yan, Forbus, & Gentner, 2003 for a more detailed discussion.)

This change in representation allows the child to notice common relational schemas across different dimensions. This kind of extraction of magnitude relations from the specific dimensions in which they are instantiated is one of the great disembeddings that occur during development. It is what permits generative understanding of cross-dimensional metaphors like "I'm feeling lower than ever today; nothing seems to raise my spirits" (Lakoff & Johnson, 1980; Nagy, 1974). It also underlies scientific analogies such as the electricity/waterflow analogy, in which *current moving through a voltage differential* is seen as analogous to *water moving through a pressure differential* (Gentner & Gentner, 1983). Carried far enough, this process of abstracting relations from their specific dimensional contexts results in general relational structures that can be applied across semantic domains.

The reminding bottleneck. If comparison is at the heart of relational learning, where do the appropriate pairs come from? Some of them arise through tutorial planning, as in the study just described. Other comparisons are brought about by the learner's own curiosity, when (for example) an infant repeatedly puts in and dumps outs an object from a cup. But this kind of self-constructed comparison can happen only in certain domains. A third route to comparison is via remindings of prior experience. But this route is limited by of the perverse nature of reminding processes. If human learners were reliably reminded of relationally similar instances, then alignments between current situations and remindings would lead to relational abstractions. But a vast amount of research shows that similarity-based reminding from long-term memory is strongly responsive to overall similarity, including surface similarity, and relatively insensitive to relational similarity (Brooks, Norman, & Allan, 1991; Forbus, Gentner, & Law, 1995; Gentner, Rattermann & Forbus, 1993; Holyoak & Koh, 1987; Reeves & Weisberg, 1994; Ross, 1987). Of course, surface-based remindings are not all bad. Very often, things that look alike are alike relationally as well (Regehr & Brooks, 1993; Medin & Ortony, 1989). Gentner (1989) referred to this as the *kind world* hypothesis: what looks like a tiger generally is a tiger. In some domains, experience naturally presents exemplars that allow remindings and comparisons to progress from near to far. In such arenas, relational categories can be learned via natural experience.

But these three routes—chance experiential juxtapositions, matches constructed by the learner, and reminding-based comparisons—are not enough to account for human relational learning. Many important relational structures are neither instantiated thickly enough to afford spontaneous comparison nor readily constructed by the child, and for these, experiential comparisons are not a reliable route to learning. (We have only to look at human history to establish that many important relational concepts—*momentum, planetary orbits, the derivative of a quantity, regression to the mean, sensitivity versus bias*—are not inevitably derived from human experience.) A further way to arrive at relational categories is through cultural guidance: in particular, through linguistic labels.

The role of language. As noted above, there is ample evidence that noun labels invites children to focus on category-level commonalities, overriding competing associations. Might common labels invite noticing relational commonalities as well as common intrinsic properties? This question is largely unexplored. In contrast to the vast amount of research investigating the acquisition of nouns denoting object categories, very little work has looked at the acquisition of relational nouns.

To begin the exploration, Gentner and Klibanoff (in preparation) investigated the acquisition of novel relational nouns. Because we assumed this would be challenging for young children we developed a methodology in which parallel relational situations were used to illustrate the new term. The idea was to use comparison across instances as a way to facilitate noticing common abstractions, as discussed above (e.g., Gentner & Namy, 1999; Jameson & Gentner 2003; Loewenstein & Gentner, 2001). Three-, four- and six-year-olds were shown picture cards that depicted the same relation across two contexts, and were then tested on a third context. In the *relational language* condition, a novel relational noun was used to name the relation: e.g., "The knife is the *blick* for the watermelon, and the ax is the *blick* for the tree." Then they were shown a new picture (e.g., a piece of paper) along with three alternatives: a pair of scissors (*same-relation*, correct), a pencil (*thematic*), and another piece of paper (*object/taxonomic*). They were asked "What would be the *blick* for the paper?" The No-language group saw the same materials without the new word: “The knife goes with the watermelon, and the ax goes with the tree in the same way. What would go with the paper in the same way?” We reasoned that if children could interpret the new word as naming a relational category, then performance would be better with the word than without it, because the word could serve to draw attention to the common relation. But if children at this stage take the term to be an object name, then performance would be worse with the relational language than without it. A third possibility, of course, is that children would ignore the term, yielding no effect of relational language.

The results showed a strong developmental effect. The 3-year-olds performed at chance in both versions of the task; they showed no insight into the common relation[7] and chose randomly among the three alternatives. In contrast, for the older groups, children performed well above chance in both versions of the task. Further, these children performed better with the novel relational noun than without it. It appears that by the age of four, the meanings of novel relational nouns can be learned in supportive contexts that permit abstracting common relations across situations. Further, by this stage, words can serve as invitations to *relational* categories as well as entity categories.

Children's initial representations are often described as highly conservative. Their knowledge is described as "concrete," "situated" or "contextually embedded." We suggest that comparison processes foster allow learners to see that the same relational patterns may apply across specific situations. In this way comparison promotes the abstraction or disembedding of relations from their initial rich contexts. This gradual abstraction of initially conservative, context-specific representations can give rise to relational categories. Such a view is consistent with Medin and Ross's (1989) proposal that sophisticated learning can arise from comparison across highly specific instances, and with Forbus and Gentner's (1986) proposal of initially conservative representations that give rise to abstractions via through comparison processes.

The findings of Gentner and Klibanoff further suggest that this abstraction process can be promoted by relational labels (see also Gentner & Rattermann, 1991). So the capacity to see consistent mappings between structures across different dimensions is promoted both by direct experiential comparisons and by learning relational language that invites comparison of situations that share the relational label. Gentner (2003) termed this “symbolic juxtaposition.” This idea is an extension of the principle that “words are invitations to form categories” (Brown, 1958; Waxman & Markow, 1995). We suggest that words are invitations to *compare exemplars* (Gentner & Namy, 1999). By giving two things the same name, we invite children to compare them regardless of their surface similarity. This is important in learning relational categories, whose members are superficially unlike.

We hypothesize that relational categories may be particularly likely to be led by language, because their common core tends to be unsupported by object commonalities and is thus less obvious on first inspection[8]. In such cases, the likelihood of spontaneous alignment of exemplars simply owing to perceived similarity would be low. Learning words for relations may have ramifications beyond the immediate category. For example, knowing a relational term may increase the likelihood that the learner will be reminded of prior instances of the relation when a new instance occurs (Gentner, 2003). This in turn will increase the likelihood of encoding relations in the same way across different situations Such representational uniformity is important in achieving the ability to perceive relational patterns across different contexts (Gentner, Forbus & Law, 1995).

Further Implications and Conclusions

Relational structure in ordinary categories. Although we have treated the entity-relational distinction, it is better seen as a continuum. Ordinary basic-level object categories also have relational information as part of their representations. For example, a cow is understood to have both characteristic internal relations—such as that its food passes through multiple stomachs—and characteristic external relations: *eats grass, lives in barn, provides milk to farmers,* and so on. These are sometimes listed among the properties of cows, but they are clearly relations between cows and other entities. There are several lines of evidence supporting the claim that standard entity categories include relational information. Ross and Murphy (1999) investigated the conceptual organization of foods and showed that items in the domain are cross-classified in terms of multiple categories, including those grounded in the events in which the food items participate. Participants naturally and efficiently applied situational categories such as *main courses* as well as taxonomic categories such as *meats* for sorting and inferential purposes. Wisniewski and Medin (1994) found that people's categorizations of children's drawings were influenced not just by the direct similarity among the drawings but by their beliefs about the groups who had drawn them (farm children vs. city children, gifted children vs. normal children, etc.).

There is mountingevidence that causal structure is important in the representation and use of ordinary entity categories. Ahn and colleagues (Ahn, 1998; Ahn et al, 2000) find that features that cause other features take on greater explanatory importance and greater weight in membership judgments, along the lines suggested in the theory view of categorization (Murphy & Medin, 1985). Rehder (2003) proposes that a feature's importance may be derived from its degree of causal connectivity and suggests that category membership is evaluated relative to a theory-like causal model that is part of the category representation. Sloman, Love & Ahn (1998) have shown that the features of a category can be reliably differentiated in terms of the extent to which they participate in dependency relations that determine conceptual coherence.

Some kinds of categories are heavily influenced by culturally specific relations. For example, Sloman and Malt (2003) argue that artifact categories are not defined by essences, and suggest that their meanings are partly based on functional relations. Cross-cultural studies of folk biology by Medin and his colleagues (Lynch, Coley & Medin, 2000; Medin, Lynch, &

Solomon, 2000: Solomon, Medin & Lynch, 1999) have found that, in addition to biological taxonomic categories, relational categories such as *ornamental tree* and *weed* exist in many systems. Indeed, landscape experts often rely more heavily on these relational categories than on standard taxonomic categories.

Another source of evidence that general category usage involves relationally structured representations comes from studies of category-based induction (Blok & Gentner, 2000; Lassaline, 1996; Wu & Gentner, 1998). In Lassaline's studies—which utilized standard basic-level animal categories—people were given the opportunity to infer a new property of a novel category, based on its similarity to an existing category. She found that people preferred to carry over properties that had a relational connection to shared features, rather than properties not connected to the matching facts. Similar results were obtained by Wu and Gentner. This propensity to make inferences based on mapping connected relational systems is a hallmark of analogical mapping (Clement & Gentner, 1991; Markman, 1997). Further, Heit and Rubenstein (1994) showed that people make stronger inferences when the inferred property is linked to the common structure shared by the categories: e.g., *tuna*→*whale* yields stronger behavioral inferences, but *bear*→*whale* yields stronger anatomical inferences. Overall, these findings suggest that category-based induction involves the alignment of relational structures, just as in analogy.

What we are suggesting here is that the entity-relational contrast is not a dichotomy between representations based on perceptual attributes and those based on relational representations. Rather, it is a continuum from representations based on *both* relational structure and rich perceptual information to ones based *only* on relational structure. Consistent with this view, categories have often been described as having two sides—characteristic vs.

defining features (Keil & Batterman, 1984; Rips, Shoben & Smith, 1973); core features vs. identificatory procedures (Smith & Medin, 1981); or sense vs. reference (Frege, 1892). Artifact categories, such a *hammer* or *knife*, are clear examples of categories that have both relational and entity sides; we know how to identify them perceptually, but we also know their causal affordances.

Similarity and theory in conceptual structure. In current research on conceptual structure, there is a deep divide between similarity-based and theory-based accounts. The traditional version of the similarity-based view holds that the correlational structure of features in the environment gives rise to categories characterized by systematic patterns of within-category similarity and between-category difference (Rosch & Mervis, 1975). This view that similarity is the basis for category structure has been highly influential. However, Murphy and Medin (1985; Medin, 1989) and Rips (1989) among others have challenged the idea that similarity explains category coherence. They note that the internal structure of a category is richer than a list of features and must include relationships within and between category examples. Further, phenomena have emerged that challenge the feature similarity account, such as Medin & Shoben's (1988) demonstration that gray clouds are considered more similar to black clouds than to white clouds, but gray hair is more similar to white hair than to black hair. The tension inherent in the need for a constrained, yet rich basis for category coherence poses a major challenge to theorists (Goldstone, 1994).

Our focus on the role of relations can go a considerable way toward addressing the major challenges suggested by the theory view of concepts (Gentner & Medina, 1998; Medin, Goldstone & Gentner, 1993). For example, in the Medin and Shoben example above, gray hair and white hair share a causal relation to aging in humans, whereas gray clouds and black clouds share the relation of predicting storms. When relational structure is considered as a key component of category representation, these effects can be derived from relational commonalities. Again, when relational similarity is taken into account, there is often no quarrel between similarity-based and theory-based categorization.

Rules and relational categories. In many ways, relational categories are closer to classical categories than are entity categories; as discussed above, their sparse, rule-like representations often yield clear criteria for membership. Yet, paradoxically, the membership of relational categories is more open than that of entity categories. For example, recall the informal observation from the Kurtz and Gentner generation study that relational categories were far more likely to be metaphorically extended to include abstract exemplars (such as *poverty* as a kind of *barrier)* than were entity categories. Barr and Caplan (1987) also found more gradedness in membership judgments for relational categories than for entity categories. They noted that when a relational category can be widely applied, its felicity of application seems to decrease as the arguments depart further from the canonical set. Indeed, in the Kurtz and Gentner study, the abstract exemplars typically occurred late in the generation lists, after a large number of concrete exemplars.

Relational categories are both more clear in their intensions and more open in their extensions than entity categories. And on further examination, this ceases to be a paradox. For densely represented concepts like *raccoon*, the entry requirements for membership are thickly interconnected; this makes it hard to state the intension, but easy to determine concrete membership. But once outside the sharp membership boundary, it is difficult to decide what else might count as a legitimate extension. In contrast, the clarity of relational intensions allows them to have broad extensions.

Historically, many relational categories may arise from metaphorical extension of entity categories, especially for categories in which causal structure is especially prominent, such as artifact categories: e.g., *bridge, sanctuary, trap* (Bowdle & Gentner, in press; Sweetser, 1990). For example, Zharikov and Gentner (2002) noted that *sanctuary* once meant a place of worship. Through analogical extensions (we hypothesize) it was gradually stripped of its entity properties, so that it can now be used to mean "a safe place." Goldstone (1995) makes a useful distinction between *default* and *directed* similarity, a contrast akin to Gentner's (1983) distinction between literal similarity and analogy. The former underlies the graded structure and broad inferential power of entity categories, while the latter describes the focal, context-specific sense of similarity that applies for goal-derived categories or analogical relationships. When directed analogical similarity processes are applied to entity categories, the result is often a metaphorical extension like those above—another indication that entity categories include relational structure.

Coherence and correspondence. The contrast between entity concepts and relational concepts maps onto the distinction between correspondence-based concepts, whose meanings can be conceived of in terms of reference to entities or sets in the world, and coherence-based concepts, whose meanings are derived from their relations with other concepts. These two kinds of concepts have very different psychological patterns and theoretical affinities. Entity concepts—categories of concrete objects and animate beings—are correspondence-driven: they derive meaning by pointing to referents in the world. Relational concepts are coherence-driven; they take on meaning through their relations with other concepts. Their sparse representations can be linked to other concepts in orderly ways that contribute to fixing their meanings. For example, consider the fact that many verbs have antonyms (*come/go, give/ take, find/lose, remember/forget*) as do many relational nouns (a*bsence/presence, creation/destruction, friend/enemy*. But concrete entity nouns rarely do.[9] If we consider that antonyms are terms that match except for a single polar or binary alignable difference, then it follows that sparsely represented concepts (like relational concepts) afford more antonymic possibilities.

Correspondence-based concepts link naturally to referential theories of meaning and to psychological theories that emphasize the role of perception and interactive experience between the thinker and the environment. Coherence-based concepts are emphasized in accounts that postulate internal structure, such as schemas (Rumelhart & Ortony, 1977), semantic nets (Collins & Quillian, 1969), qualitative process theory (Forbus, 1984), and other accounts of knowledge representation (See Markman, 1999, for a review.) The two views implicitly assume different relationships between language and thought. According to the correspondence view, concepts arise naturally from the way our perceptual capacities operate on the experiential world. The role of language is simply to name those concepts according to their kind: action, object, spatial relation, etc. The mapping is from world to word. In contrast, the coherence-driven view assumes a large role for the semantics of one's language and culture in determining the way information is structured. The extreme of the coherence-driven view was expressed in Saussure's writings on language—for example, in the idea that "each linguistic term derives its value from its opposition to all other terms" (Saussure 1916/1966: 88) and that "Language is a system of interdependent terms in which the value of each term results solely from the simultaneous presence of the others ..." (Saussure 1916/1966: 114). If we substitute "concept" for "linguistic term" we get what could serve as the strong coherence view of concepts.

Goldstone and Rogosky (2002) tested the sufficiency of coherence and the interaction between coherence and correspondence in an ingenious computational study. The intuitive question they addressed is whether having the same set of internal coherence relations is enough to allow two people to communicate. They designed a system, ABSURDIST, whose job was to put two representations in correspondence using the structural consistency constraint from structure-mapping. In the initial studies, the two representations each consisted of a set of nodes with simple relations. The relational vocabulary was deliberately impoverished, consisting solely of numerical ratings of similarity between each pair of concepts within each system. To put this in perspective, imagine a representation of a *robbery* that consisted solely of the similarity ratings between the *thief* and the *stolen goods*, the *victim*, the *robbery*, and so on. This of course was not meant to be an adequate representation of meaning (hence the name "ABSURDIST"). Indeed, it would seem impossible to derive any sensible correspondences between two such representations. Yet even with all these handicaps, ABSURDIST was able to derive correspondences fairly accurately for small sets (between 3 and 15 items) provided the amount of noise was low. There are two other interesting results. First, as long as noise was low, the system did *better* with more nodes (not worse, as might have been predicted). Second, when even one external referential correspondence was added (by requiring two nodes to match, as would happen if they referred to the same entity in the world), the results for larger sets of nodes improved markedly.

These results show that pure coherence—even over an absurdly limited internal relational vocabulary with only one representational currency, degree of similarity—can go a long way towards achieving translatability between speakers. Overall, ABSURDIST demonstrates two points: First, it shows the power of coherence relations—of a concept's relations to other concepts in the mind—in establishing meaning. Second, it shows that adding even a few correspondence pointers to the external world greatly improves the ability to align two conceptual systems. If we extrapolate to realistic human knowledge representations, it is clear that we will need many distinct kinds of relations. A more powerful structure-mapping process, such as SME, is required to carry out alignment and inference across such relationally delineated representations (Falkenhainer, Forbus & Gentner 1989; Hummel & Holyoak, 1997; Larkey & Love, 2003; Ramscar & Yarlett, 2003).

But despite the importance of aligning internal relational representations, the size and intricacy of human representations demands that there also be some external pointers to ground the alignment, and this is where entity categories enter in. It is doubtful that relational terms could even be learned without some entity terms to anchor them. This may one reason that names for entities are so prominent in children's language. They serve to permit alignment with adult minds, and they facilitate arriving at common relational information. With increasing knowledge, it becomes possible to find more complex alignments with other minds—alignments based on relational categories and relational systems.

Acknowledgments

This research was supported by NSF grant SBR-95-11757 and ONR contract N00014-92-J-1098 to Dedre Gentner, and by NSF-ROLE award 21002/REC-0087516. We thank Jennifer Asmuth, Jeff Loewenstein, Art Markman, and Phillip Wolff for discussions of these issues, and Jason Jameson, Brad Love and Vladimir Sloutsky for insightful comments on this paper.

References

Ahn, W. K. (1998). Why are different features central for natural kinds and artifacts? The role of causal status in determining feature centrality. *Cognition, 69* (2) Dec 1998, 135-178.

Ahn, W. K., Kim, N. S., Lassaline, M. E., & Dennis, M. J. (2000). Causal status as a determinant of feature centrality. *Cognitive Psychology, 41* (4), 361-416.

Asmuth, J. A., & Gentner, D. (in preparation). Context sensitivity of relational nouns.

Baldwin, D. A. (1989). Priorities in children's expectations about object label reference: Form over color. *Child Development, 60,* 1289-1306.

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

Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition, 11,* 211-227.

Barsalou, L. W. (1985). Ideas, central tendency, and frequency of instantiation as determinants of graded structures in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11* (4), 629-654.

Blok, S. V., & Gentner, D. (2000). Reasoning from shared structure. *Proceedings of the 22nd meeting of the Cognitive Science Society,* Philadelphia, PA.

Bowerman, M., & Choi, S. (2003). Space under construction: Language-specific spatial categorization in first language acquisition. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and cognition* (pp. 387-428). Cambridge, MA: MIT Press.

Brooks, L. R, Norman, G. R., & Allen, S. W. (1991). Role of specific similarity in a medical diagnostic task. *Journal of Experimental Psychology: General*. *120*(3), 278-287.

Brown, R. (1958). How shall a thing be called? *Psychological Review, 65*(1), 14-21.

Bruner, J. S., Greenfield, P., Olver, R. (1966). *Studies in cognitive growth.* New York: Wiley.

Caselli, M. C., Bates, E., Casadio, P., Fenson, J., Fenson, L., Sanderl, L., & Weir, J. (1995). A cross-linguistic study of early lexical development. *Cognitive Development, 10,* 159-199.

Chafe, W. L. (1970). Discourse structure and human knowledge. In J. B. Carroll & R. O. Freedle (Eds.), *Language comprehension and the acquisition of knowledge,* pp. 97-98. Washington, D. C.: V. H. Winston.

Chipman, S. F., & Mendelson, M. J. (1979). Influence of six types of visual structure on complexity judgments in children and adults. *Journal of Experimental Psychology: Human Perception and Performance, 5,* 365-378.

Clark, E. V. (1993). *The lexicon in acquisition.* Cambridge: Cambridge University Press.

Clement, C. A., & Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. *Cognitive Science, 15,* 89-132.

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

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence, 41,* 1-63.

Forbus, K. D. (1984). Qualitative process theory. *Journal of Artificial Intelligence, 24,* 85-168.

Forbus, K. D., & Gentner, D. (1986). Learning physical domains: Toward a theoretical framework. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 2, pp. 311-348). Los Altos, CA: Kaufmann.

Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science, 19,* 141-205.

Gentner, D. (1981). Some interesting differences between nouns and verbs. *Cognition and Brain Theory, 4,* 161-178.

Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In S. Kuczaj (Ed.), *Language development: Vol. 2. Language, thought and culture* (pp. 301-334). Hillsdale, NJ: Erlbaum.

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

Gentner, D. (1989). Mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199-241). London: Cambridge University Press.

Gentner, D., & Boroditsky, L. (in preparation). Early acquisition of nouns and verbs: Evidence from Navajo.

Gentner, D., & Boroditsky, L. (2001). Individuation, relativity and early word learning. In M. Bowerman & S. Levinson (Eds.), *Language acquisition and conceptual development* (pp. 215-256). Cambridge, UK: Cambridge University Press.

Gentner, D., & France, I. M. (1988). The verb mutability effect: Studies of the combinatorial semantics of nouns and verbs. In S. L. Small, G. W. Cottrell, & M. K. Tanenhaus (Eds.), *Lexical ambiguity resolution: Perspectives from psycholinguistics, neuropsychology, and artificial intelligence* (pp. 343-382). San Mateo, CA: Kaufmann.

Gentner, D., & Gentner, D. R. (1983). Flowing waters or teeming crowds: Mental models of electricity. In D. Gentner & A. L. Stevens (Eds.), *Mental models* (pp. 99-129). Hillsdale, NJ: Erlbaum.

Gentner, D. and Klibanoff, R. S. (in preparation). On acquiring *gift*: the acquisition of relational nouns.

Gentner, D., & Markman, A. B. (1997). Structure-mapping in analogy and similarity. *American Psychologist, 52*(1), 45-56.

Gentner, D., & Namy, L. (1999). Comparison in the development of categories. *Cognitive Development, 14*, 487-513.

Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In S. A. Gelman & J. P. Byrnes (Eds.), *Perspectives on language and thought: Interrelations in development* (pp. 225-277). London: Cambridge University Press.

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

Gillette, J., Gleitman, H., Gleitman, L., & Lederer, A. (1999). Human simulations of vocabulary learning. *Cognition, 73*(2), 135-176.

Goldin-Meadow, S., Seligman, M. E. P., & Gelman, R. (1976). Language in the two-year old. *Cognition, 4*, 189-202.

Goldstone, R. L. (1994). The role of similarity in categorization: Providing a groundwork. *Cognition, 52*, 125-157.

Goldstone, R. L. (1995). Mainstream and avant-garde similarity. *Psychologica Belgica, 35(2/3),* 145-165.

Goldstone, R. L. (1996). Isolated and interrelated concepts. *Memory & Cognition,* 31, 169-180.

Goldstone, R. L., Steyvers, M., & Rogosky, B. J. (2003). Conceptual interrelatedness and caricatures. *Memory & Cognition, 31,* 169-180.

Golinkoff, R. M., & Hirsh-Pasek, K. (1990). Let the mute speak: What infants can tell us about language acquisition. *Merrill-Palmer Quarterly, 361,* 67-91.

Halford, G. S. (1987). A structure-mapping approach to cognitive development. *International Journal of Psychology, 22*, 609-642.

Heit, E., & Rubinstein, J. (1994). Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*, 411-422.

Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory & Cognition, 15,* 332-340.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review, 104*(3), 427-466.

Imai, M., Gentner, D., & Uchida, N. (1994). Children's theories of word meaning: The role of shape similarity in early acquisition. *Cognitive Development*, *9*, 45-75.

Imai, M., Haryu, E., & Okada, H. (in preparation). Mapping novel nouns and verbs onto dynamic action events: Are verb meanings easier to learn than noun meanings for Japanese children?

Jameson, J., and Gentner, D. (2003). Mundane comparisons can facilitate relational understanding. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the Twenty-fifth Annual Meeting of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum.

Jusczyk, P. W. (1997). *The discovery of spoken language*. Cambridge: MIT Press.

Keil, F. C., & Batterman, N. (1984). A characteristic-to-defining shift in the development of word meaning. *Journal of Verbal Learning and Verbal Behavior, 23,* 221-236.

Kersten, A. W., & Earles, J. L. (in press), Semantic context influences memory for verbs more than memory for nouns. *Memory and Cognition.*

Kloos, H, & Sloutsky, V. (2004). Are natural kinds psychologically distinct from nominal kinds? Evidence from learning and development. *Proceedings of the Meeting of the Cognitive Science Society*, Chicago, IL.

Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development, 67,* 2797-2822.

Kuhn, T. S. (1962). *The structure of scientific revolutions.* Chicago: University of Chicago Press.

Kurtz, K. J., & Gentner, D. (2001). Kinds of kinds: Sources of category coherence. *Proceedings of the Twenty-third Annual Conference of the Cognitive Science Society,* 522-527.

Kurtz, K. J., & Gentner, D. (in preparation). Distinguishing kinds of categories.

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by* (pp. 3-34). Chicago, IL: University of Chicago Press.

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

Lassaline, M. E. (1996). Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22,* 754-770.

Loewenstein, J., & Gentner, D. (2001). Spatial mapping in preschoolers: Close comparisons facilitate far mappings. *Journal of Cognition and Development*, *2*(2), 189-219.

Loewenstein, J., Thompson, L., & Gentner, D. (1999). Analogical encoding facilitates knowledge transfer in negotiation. *Psychonomic Bulletin & Review, 6*(4), 586-597.

Lynch, E. B., Coley, J. D., & Medin, D. L. (2000). Tall is typical: central tendency, ideal dimensions and graded category structure among tree experts and novices. *Memory and Cognition, 28* (1), 41-50.

Markman, A. B. (1997). Constraints on analogical inference. *Cognitive Science, 21*(4), 373-418.

Markman, A. B. (1999). *Knowledge representation*. Mahwah, NJ: Lawrence Erlbaum Associates.

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

Markman, E. M. (1989). *Categorization in children: Problems of induction*. Cambridge, MA: MIT Press.

Medin, D. L. (1989). Concepts and conceptual structure. *American Psychologist, 44*, 1469-1481.

Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review, 100* (2), 254-278.

Medin, D. L., Lynch, E. B., & Coley, J. D. (1997). Categorization and reasoning among tree experts: Do all roads lead to Rome? *Cognitive Psychology, 32, 49-96.*

Medin, D. L., Lynch, E. B., & Solomon, K. O. (2000). Are there kinds of concepts? *Annual Review of Psychology, 51*, 121-147

Medin, D. L., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 179-195). New York: Cambridge University Press.

Medin, D. L., & Ross, B. H. (1989). The specific character of abstract thought: Categorization, problem-solving, and induction. In R. J. Sternberg (Ed.), *Advances in the psychology of human intelligence* (Vol. 5, pp. 189-223). Hillsdale, NJ: Erlbaum.

Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology, 20*, 158-190.

Miller, G. A., & Johnson-Laird, P. N. (1976). *Language and perception.* Cambridge, MA: Belknap Press.

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

Nagy, W. (1974). *Figurative patterns and the redundancy in lexicon.* Unpublished doctoral dissertation, University of California at San Diego.

Paivio, A. (1971). *Imagery and verbal processes.* New York: Holt, Rinehart, & Winston. (Reprinted by Erlbaum, 1979).

Quinn, P.C., Bhatt, R.S., Brush, D., Grimes, A., & Sharpnack, H. (2002) Development of form similarity as a gestalt grouping principle in infancy. *Psychological Science 13(4)* 320.

Ramscar, M. & Yarlett, D. (2003). Semantic grounding in models of analogy: An environmental approach. *Cognitive Science 27*: 41-71.

Rattermann, M. J., & Gentner, D. (1998). More evidence for a relational shift in the development of analogy: Children's performance on a causal-mapping task. *Cognitive Development, 13*, 453-478.

Reeves, L. M., & Weisberg, R. W. (1994). The role of content and abstract information in analogical transfer. *Psychological Bulletin, 115*(3), 381-400.

Regehr, G., & Brooks, L. R. (1993). Perceptual manifestations of an analytic structure: The priority of holistic individuation. *Journal of Experimental Psychology: General*, *122*(1), 92-114.

Rehder, B. (2003). Categorization as causal reasoning. *Cognitive Science*, 27, 709-748.

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

Rips, L. J. (1989). Similarity, typicality, and categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 21-59). New York: Cambridge University Press.

Rips, L. J., Shoben, E. J., & Smith, E. E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior, 12,* 1-20.

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

Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*(4) 629-639.

Ross, B. H. (1999). Postclassification category use: The effects of learning to use categories after learning to classify. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25*(3), 743-757.

Ross, B. H., & Murphy, G. L. (1999). Food for thought: Cross-classification and category organization in a complex real-world domain. *Cognitive Psychology, 38,* 495-553.

Rumelhart, D. E., & Ortony, A. (1977). Representation of knowledge in memory. In R. C. Anderson, R. J. Spiro, & W. E. Montague (Eds.), *Schooling and the acquisition of knowledge*. Hillsdale, NJ: Erlbaum.

Saussure, F. de, *Course in general linguistics* (C. Bally & A. Sechehaye, Eds., W. Baskin, Trans.). New York: McGraw Hill, 1966. (Original work published 1916)

Slobin, D. I. (1996). From “thought and language” to “thinking for speaking”. In J. J. Gumperz & S. C. Levinson (Eds.), *Rethinking linguistic relativity* (pp. 70-96). Cambridge, MA: Cambridge University Press.

Sloman, S. A., Love, B.C. & Ahn, W.K. (1998). Feature centrality and conceptual coherence. *Cognitive Science, 22*, 189-228.

Sloman, S. A., & Malt, B. C. (2003). Artifacts are not ascribed essences, nor are they treated as belonging to kinds. *Language and Cognitive Processes*, *18*, 563-582.

Smith, L. B., Jones, S. S., & Landau, B. (1992). Count nouns, adjectives, and perceptual properties in children's novel word interpretations. Developmental Psychology, 28, 273-286.

Sobin, N. (1982). Texas Spanish and lexical borrowing. In J. Amastae & L. Elias-Olivares (Eds.), *Spanish in the United States: Sociolinguistic aspects*. Cambridge: Cambridge University Press, pp. 166-181.

Solomon, K. O., Medin, D. L., & Lynch, E. B. (1999). Concepts do more than categorize. *Trends in Cognitive Science, 3*(3), 99-105.

Sweetser, E. (1990). *From etymology to pragmatics*. Cambridge: Cambridge University.

Talmy, L. (1975). Semantics and syntax of motion. In J. Kimball (Ed.), *Syntax and semantics* (Vol. 4, pp. 181-238). New York: Academic Press

Tardif, T., Gelman, S. A., & Xu, F. (1999). Putting the noun bias in context: A comparison of English and Mandarin. *Child Development, 70*(3), 620-635.

Waxman, S. R., & Hall, D. G. (1993). The development of a linkage between count nouns and object categories: Evidence from fifteen- to twenty-month-old infants. *Child Development, 64,* 1224-1241.

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

Wisniewski, E. J., & Medin, D. L. (1994). On the interaction of theory and data in concept learning. *Cognitive Science, 18,* 221-281.

Wu, M., & Gentner, D. (1998). Structure in category-based induction. *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society*, 1154-1158.

Yan, J. Forbus, K., & Gentner, D. (2003). A theory of rerepresentation in analogical matching. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the Twenty-fifth Annual Meeting of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum.

Zharikov, S., & Gentner, D. (2002). Why do metaphors seem deeper than similes? In W. D. Gray & C. D. Schunn (Eds.), *Proceedings of the Twenty-Fourth Annual Conference of the Cognitive Science Society* (pp. 976-981). Fairfax, VA: George Mason University.

Figure Captions

Figure 1. An example entity category

Figure 2. Sample materials used in Kotovsky and Gentner's (1996) progressive alignment study.

Table 1. Summary of the characteristics of verbs as compared to concrete nouns (adapted from Gentner, 1981, 1982; Gentner & Boroditsky, 2001).

Cross-linguistically variable

Hard to translate

Not borrowed in language contact

Hard to learn

Late in acquisition

Late in second-language learning

Context-sensitive in meaning

Polysemous

Poor in memory retrieval

Distributed at higher word frequencies

Impaired in Broca's aphasia

Table 2

Predicted contrasts between entity categories and relational categories extrapolated from the analogy with nouns and verbs

Entity Categories	*Relational Categories*
Object reference	Relational meaning
Concrete nouns, proper nouns	Verbs, prepositions, relational nouns
Perceptually given	Linguistically constructed
Locally pre-individuated through perception	Individuated relative to system of concepts
Stable across languages	Variable cross-linguistically
Acquired early	Acquired late
Correspondence-driven	Coherence-driven

Table 3
Sample generation lists for a relational superordinate category (barrier) and an entity superordinate category (vegetable) adapted from Kurtz and Gentner (2001, in preparation)

Relational Role Category *Barrier*		Entity Category *Vegetable*	
Protocol 1	Protocol 2	Protocol 1	Protocol 2
a wall	fence	carrots	carrots
an insult	a wall	zucchini	celery
a person	post-emotional trauma	asparagus	radishes
a door	barbed wire fence	broccoli	lettuce
a smell	gender	sprouts	string bean
a fear	ditch	potato	snow pea
a river	age	tomato	peas
a forest	moat	celery	corn
a tree	an ocean	artichoke	potato
a cliff	a mountain	water cress	onion
an ocean	land mines	beans	chive
a toxic fume	a fear	refried beans	cabbage
time	low self-concept	string bean	bean sprouts
lack of resources	morals	green beans	brussel sprouts
no money	a door	peas	yams
anger		lettuce	sweet potato
emotion		cabbage	broccoli
love		cauliflower	cauliflower
tiredness		cucumber	soy beans
a duty		pickles	garbanzo beans
a requirement		radishes	ginger
loss of power		squash	cucumber
loss of desire		corn	zucchini
a mountain		spinach	water chestnuts
trouble		onion	parsley
danger		olives	squash
loss of direction		beet	
loss of support		person in a coma	

Table 4. Pilot data from category generation task: Categories generated for the exemplar *dog* (adapted from Kurtz & Gentner, 2001; in preparation).

Taxonomic responses:

canine, animal, mammal, being, organism, living thing, physical object

Relational responses:

carnivore, pet, creature, guard, companion, friend, guide, hunter, racer, playmate, rescuer, fighter, showpiece, barrier, social parasite, threat, weapon, food, profit-maker, host for parasites, disease, carrier, cat chaser, swimmer, escapee, mess-maker, transportation

Figure 1. A standard category

Note. From "Les Vaches / Cows" by Atelier Nouvelles Images.

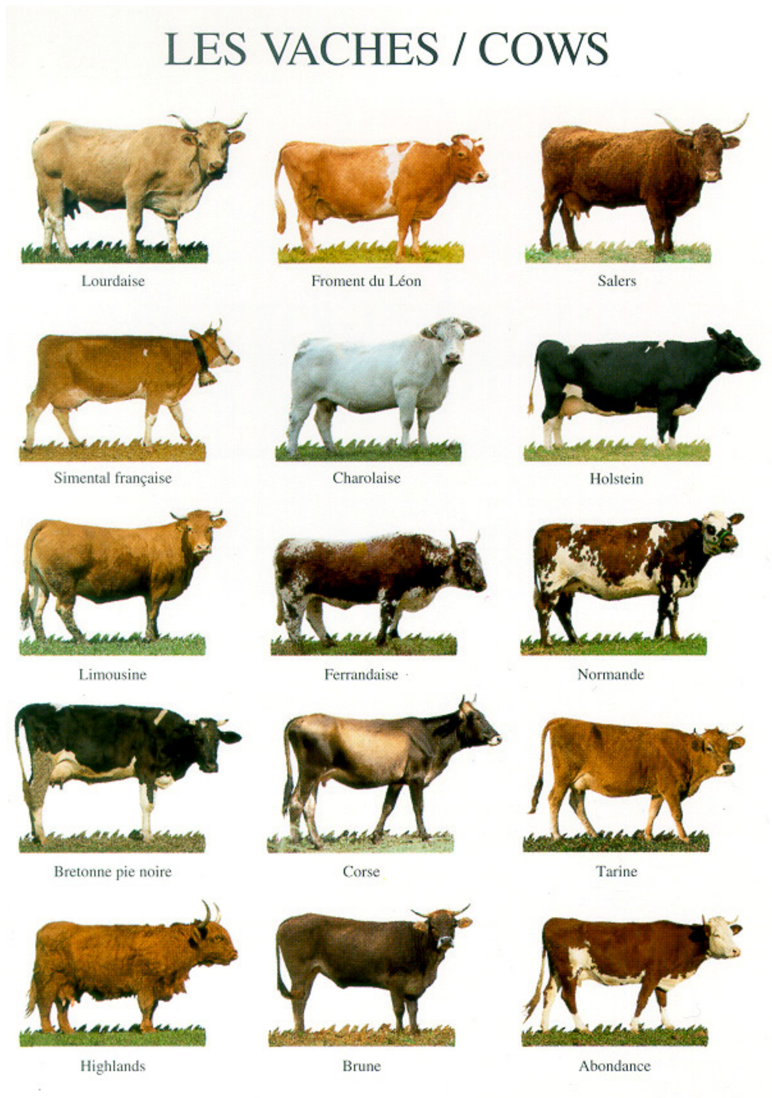

Figure 2. Sample Stimuli from Kotovsky & Gentner

Within-Dimension

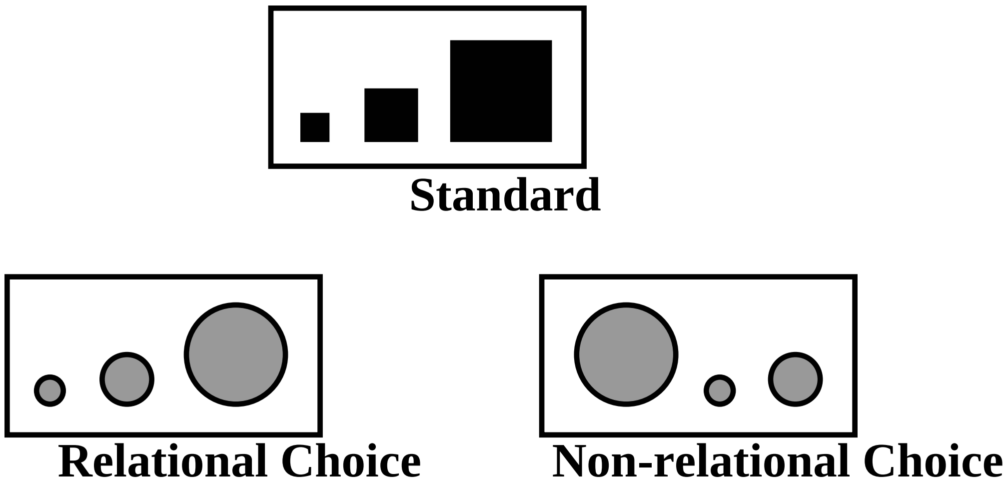

Across-Dimension

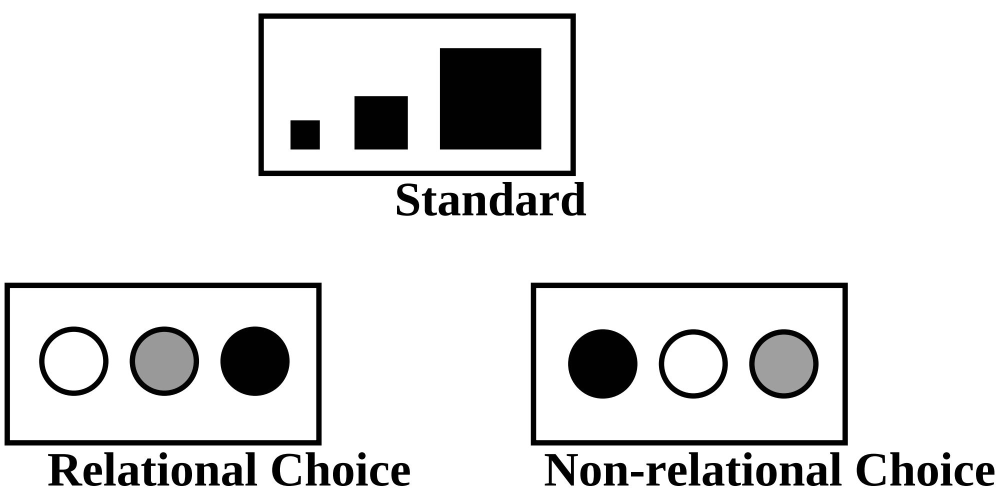

[1] Although this paper focuses chiefly on nominal relational categories, we note that relational schema categories include categories named by prepositions (e.g., *in*(x,y) and *between*(x,y,z)), comparative adjectives (e.g., *bigger*(x,y) and *meaner*(x,y)), and verbs (e.g., *approach*(x,y) and *give*(x,y,z)).

[2] However, although relational categories tend to be more abstract than entity categories, the distinction between them cannot be reduced to an abstract-concrete difference. For example, superordinate entity categories such as *object* and *animal* are abstract but not relational.

[3] However, note that there are often morphological/derivational relationships between relational role nouns and verbs (e.g., *victim/victimize*) as well as between relational role nouns and relational schema nouns (e.g., *thief/thievery*; *robber/robbery*).

[4] Although there has been dispute on this point, we believe the weight of the evidence favors Gentner's proposed semantic universal of early noun dominance in child vocabularies (see Gentner & Boroditsky, 2001, for a summary).

[5] The study included four kinds of categories: entity categories, relational categories, ad-hoc goal-derived categories (e.g., things to sell at a garage sale) and thematic categories (e.g., items associated with working at an office desk).

[6] The superordinates of entity categories are sometimes entity categories (like plant or animal) and sometimes abstract relational categories (see Markman & Stilwell, 2001).

[7] The poor performance of the 3-year-old group is consistent with the claim of a *relational shift* from an early focus on entity-level commonalities to a later focus on relational commonalities (Gentner & Rattermann, 1991; Rattermann & Gentner, 1998).

[8] In addition, the analogy with verbs suggest that relational categories may be cross-culturally variable; learning language may be instrumental in forming the category.

[9] An exception is that many biological entity categories have antonyms based on gender: e.g., *mare/stallion, ewe, ram.* Interestingly, kinship terms—a relational category system—are far richer in antonyms: *mother/father*, *sister/brother*, *father/son, son/daughter.*

Figure 1. A standard category

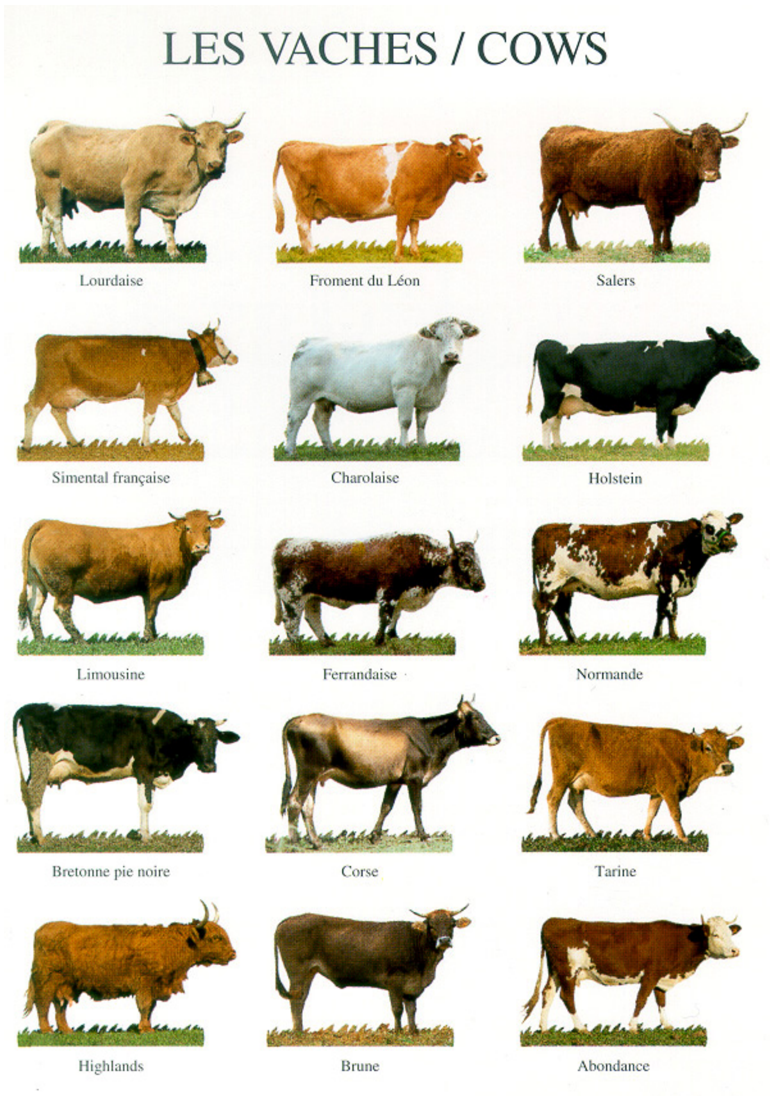

Figure 2. Sample Stimuli from Kotovsky & Gentner