J. Expt. Theor. Artif. Intell. 13(2001)329–358

Role-governed categories

ARTHUR B. MARKMAN and C. HUNT STILWELL

Department of Psychology, University of Texas, Austin, TX 78712, USA
e-mail: markman@psy.utexas.edu

Abstract. Theories of categorization have typically focused on the internal structure of categories. This paper is concerned with the external structure of categories. In particular, it is suggested that many categories specify the relational role that is played by category members. To support this claim, the paper distinguishes between traditional feature-based categories, relational categories (which specify a relational structure) and role-governed categories (which specify that an item plays a particular role within a relational structure). After discussing the relationship among these types of categories, the implications of this view for the study of category learning and category use are discussed.

Keywords: categories, relations, representations

Received September 2001; accepted October 2001

1. Introduction

Much research has been devoted to the way that category information is represented. This work has looked both at natural categories and artificial categories learned in the lab. A variety of proposals have been presented and defended including the idea that concepts are represented by prototypes (e.g. Reed 1972, Rosch and Mervis 1975), by collections of exemplars (e.g. Medin and Schaffer 1978, Nosofsky 1986), and by theory-based information that may explain how items fit into a category (e.g. Murphy and Medin 1985).

Despite this variety of approaches, all of these views share the assumption that category representations consist of some collection of properties. Where theories of categorization differ is in their hypotheses about the kinds of properties that are used to represent a concept and the structure of that information. That is, theories of categorization in cognitive psychology almost universally assume that there are *feature-based categories*. In a feature-based category, the category label refers to a set of properties denoted by the label. It is important to note that the term feature-based is being used to emphasize that category representations involve properties. Those properties might be features (i.e. discrete symbols that denote properties), but they might also be mental spaces or relational structures (Markman 1999).

In this paper, we suggest that the feature-based view is too restrictive, because it describes only one type of concept that people have.[1] A critical type of concept that is ignored by the focus on feature-based categories is categories that denote the role

Journal of Experimental & Theoretical Artificial Intelligence ISSN 0952–813X print/ISSN 1362–3079 online © 2001
Taylor & Francis Ltd http://www.tandf.co.uk/journals
DOI: 10.1080/09528130110100252

that an item plays in a more global relational structure. We call these categories *role-governed categories*. For example, a private in the army is a category defined by the relative position of a particular soldier within a military hierarchy. Privates are at the bottom of this hierarchy, and thus must take orders from soldiers of a variety of other ranks. A private may be recognized on the basis of features such as having a single stripe on his or her uniform sleeve. These features are useful for recognition and for inference. They are not critical for category membership, however. What defines a private as a private is his or her position within this social structure. Thus, private is a role-governed category.

Many concepts, beyond just those that make social distinctions, are actually defined by their position within a relational structure. For example, many artefact categories and superordinate categories are role-governed categories. In order to provide a structure for future research, we distinguish between four types of concepts. To provide a groundwork for this discussion, we begin with a brief description of the feature-based view of category representation. Then, we discuss the nature of role-governed category structures in detail. We also discuss two additional types of categories—relational categories and transformational categories—that must also be considered in any discussion of role-governed categories. Finally, we relate role-governed category structures to other proposals about category learning and category use.

2. The feature-based view of categories

Most general reviews of categorization assume some form of category representation in which the features of categories are subordinate to the category label (Smith and Medin 1981, Medin and Heit in press). For example, a featural view might assume that each category is represented by a prototype consisting of a set of features typically associated with a particular category (Smith *et al.* 1974, Tversky 1977). The category representation might also contain information about the degree to which particular features are criterial for a given category. There is nothing important about the assumption that a category is represented by a prototype here (i.e. it is possible for featural views of categorization to posit storage of exemplars or of multiple prototypes for each category). What is critical is that the featural information is subordinate to the category label. By 'subordinate', we mean that these features are contained within the category representation, and the label refers to this category representation.

Spatial theories of category representation also assume that the category representation is subordinate to the category label. On this view, concepts are points or vectors in a multidimensional space (Rips *et al.* 1973, Nosofsky 1986), and the category representation consists of values for a given concept along each dimension that defines the space.

Even theory-based views of categorization often assume that the representational information about the category is subordinate to the category label (Murphy and Medin 1985, Gopnik and Meltzoff 1997). Some variations of theory theories suggest that a part of the representation of a category consists of causal (or theory) information about the domain. For example, Murphy and Medin (1985) point out that a person who dives into a pool with their clothes on might be classified as drunk. This person does not have specific features that label them as drunk, but instead, a theory of drunken behaviour suggests that a person might do outlandish things.

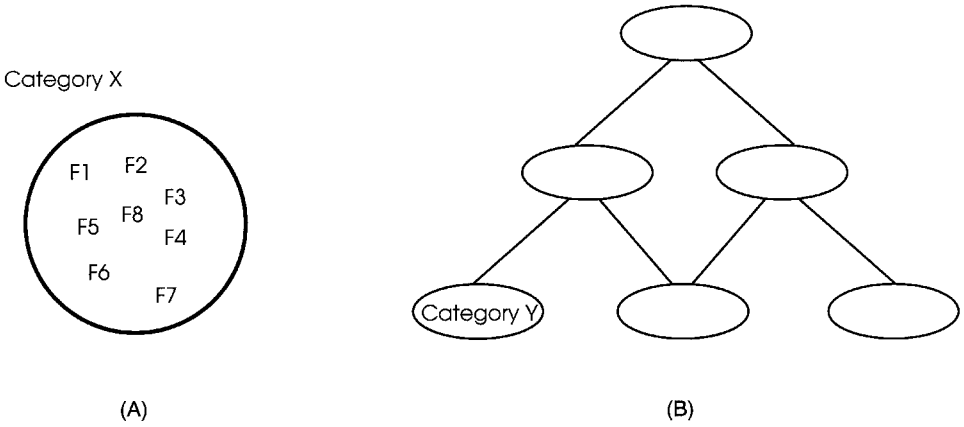

Figure 1. (a) Schematic of the kind of category structure typically assumed by theories; (b) schematic of an inverted category structure.

Diving into a pool with one's clothes on is outlandish, and so this person might be classified as being drunk on the basis of possessing a feature that is consistent with a theory of drunken behaviour.

Finally, proposals in psychology about the nature of category essences in artefacts and natural kinds assume that there is some property (which may be unknown) that makes an item a member of a particular category (e.g. Bloom 1998, 2000, Medin and Ortony 1989). Figure 1a shows a schematic for feature-based categories, in which the featural information is subordinate to the category label.

Given category representations of this type, theories of classification tend to focus on similarity relationships between the new exemplar and some stored category representation (i.e. a prototype or exemplars, Reed 1972, Nosofsky 1986). Thus, there is a graded degree to which an exemplar belongs to a category that decreases with the similarity of the new exemplar to a stored category representation (Rosch and Mervis 1975). This graded degree of membership is called *typicality*.

There is an interesting asymmetry, however, between category features and category labels. People are willing to accept an object as a member of a category on the basis of its overall similarity to other category members. Once an object is known to be a member of that category, however, people tend to be more rigid in their beliefs about the features the object is likely to have. For example, Yamauchi and Markman (2000) found that people were reluctant to infer that a member of a category had a feature value that was atypical for that category. Instead, their inferences tended to be of the most typical value for the category. Similarly, Gelman and her colleagues have found that children given a category label for an object tend to infer properties of that object consistent with the label rather than based on the perceptual similarity of the object to other items (e.g. Gelman and Markman 1986). Finally, there is evidence that children assume personality traits described as properties (e.g. 'John eats carrots') are less enduring than personality traits described as labels for the person (e.g. 'John is a carrot-eater'; Gelman and Heyman 1999).

These results suggest that there is a discontinuity between possession of a feature and membership within a category. This discontinuity flies in the face of the assumption of many models of classification (e.g. Nosofsky 1986, Anderson 1990,

Kruschke 1992) that category labels are just another feature dimension of categories (though perhaps a feature dimension that is weighted more strongly than other dimensions). What is it about category membership that makes it different from mere possession of a feature? To answer this question, we return to the notion of a role-governed category structure that we highlighted at the beginning of this paper.

3. Role-governed category structures

As shown in figure 1a, the feature-based view assumes that category labels refer to a category representation that consists of some collection of features. There are other ways that categories could be organized. In particular, many proposals about mental representations suggest that relational structures are important for understanding the world (e.g. Schank 1972, Norman and Rumelhart 1975, Schank and Abelson 1977). These relations may be used to represent spatial relations (Biederman 1987, Landau and Jackendoff 1993), causal information (Gentner and Stevens 1983, Ahn *et al.* 2000), events (Schank 1982, Schank and Abelson 1977) and verbs (Gentner 1978, 1981, Au 1986).

Many of these relations are socially, perceptually or culturally important, and thus recur often in a person's experience. Indeed, as we discuss below, there may be categories that refer to these relations. In addition, there are many categories that refer to items that play particular roles within a given relational structure. These categories, which we call *role-governed categories*, are the focus of this paper. For example, classifying a person as a professor signals that they play a particular role within a university. This view is schematized in figure 1b. The category label refers to items that play a specific role within a relational structure. In some sense, role-governed categories invert the traditional way of looking at category representation, because they suggest that categories are defined by their position within a more global relational structure that is superordinate to the category label.

Some proposals for theory-based categories appear to refer to categories that are defined by the relational structure in which they are embedded. For example, Carey (1985) draws parallels between conceptual change in children and in scientific theories. Presumably, elements within scientific theories derive their meaning from the roles they play within a theory, and thus are role-governed categories.

In this section, we address four issues critical to the role-governed category view. First, we take up the issue of category definitions. Second, we examine the relationship between role-governed categories and the feature-based view of categorization. Third, we explore the kinds of things that are likely to be represented by role-governed categories. Finally, we distinguish role-governed categories from related ideas like the concept of a relational category.

3.1. *Category definitions*

Category definitions have come to have a bad name in psychology (e.g. Smith and Medin 1981, also see Armstrong *et al.* 1983). Category definitions are usually described with respect to the feature-based view of categorization (this approach is often called the *classical view* of categorization). In this approach, a category definition is a set of necessary and sufficient conditions (e.g. features) that specify when some item is a member of a category. For example, geometry provides strict definitions for categories of geometric figures. A triangle is a three-sided closed

figure. Each of these features is necessary (e.g. no non-three-sided item is a triangle) and they are jointly sufficient to specify the category (i.e. no object with these properties is not a triangle).

The poor reputation of the classical view arises from two sources: influential philosophical arguments and the manifest difficulty of providing a good definition for categories. Wittgenstein (1968) provided the most widely cited argument against classical definitions. In a famous passage, he discussed the concept *game*, providing counter-examples to a variety of possible classical definitions, and concluding more generally that these definitions do not exist. Many psychologists have pointed out, in like style, that even simple concepts that ought to be the best candidates for having a classical definition are not well described by sets of necessary and sufficient features. For example, a bachelor seems straightforwardly defined as an unmarried adult male, but many counter-examples to this definition exist. For example, priests and gay men in long-term partnerships are, at best, poor examples of bachelors.

Many theorists have pointed out that the absence of classical definitions flies in the face of people's intuitions that categories do in fact have defining characteristics. Medin and Ortony (1989) suggested that people act as if categories have some *essence* that makes the categories what they are (Putnam 1988). People may not know what the essence is, but they act as if the category has one, leading to the intuition that there are some necessary and sufficient conditions that make an object a member of a category. Brooks (1999) suggests that the intuition that categories have definitions comes from the fact that people tend to learn categories holistically, and thus they never analyse their categories sufficiently to realize that the categories do not have a definition. Even for categories that people have analysed, they may assume there is an essence, but that they do not know precisely what it is (although some expert might know it; Putnam 1988).

Role-governed categories may provide another way out of this impasse for some categories. On this view, when an exemplar is classified as a member of a role-governed category, the item is being placed into a particular role within a more global relational structure. This position within the relational structure defines the category. Items that are permitted to play this role may differ in their properties. Thus, category definitions are not sets of necessary and sufficient features as the classical view of category representation posits. Because the classical view focuses on properties that are subordinate to the category label, it misses the potential defining power of the relational structure superordinate to the category.

On this view, Wittgenstein's argument about the category 'game' conflated two different types of representations. A game bears a particular relationship to people who play it. As we will see, properties typically associated with games (like being fun) are useful for recognizing that something is a game, for judging the typicality of games and for making predictions of features of games, but they are not defining properties of games. Wittgenstein selected examples of events that could play the role of a game, and then suggested that these items did not have some key property that ought to be in the definition of a game (or he selected items that would not play the role of a game and suggested that it would also be included by a particular definition).

So, what is the definition of a game? As discussed above, a role-governed category is one that names a specific relational role. For example, consider the simple relation play(x,y). This relation relates two arguments (x and y), each of which specifies a relational role. In the relation play(x,y), the first argument (x) specifies a 'player' (i.e.

one who plays) and the second argument specifies a *game* (i.e. something that is played). This analysis suggests that a game is defined as an object that plays the role of the second argument of the relation play(x,y). Similarly, the category 'player' is defined as something that plays the role of the first argument in the relation play(x,y).

In this example, the relational system that is superordinate to the role-governed category 'game' has a lexical entry ('play'). Not all role-governed categories participate in relational systems that have names. For example, a 'private' is a military rank; it is defined with respect to the social system of relations in the army. This system of relations does not have a specific name.

Notice that the definition of game does not involve any of the features that are typical of games (such as having rules or being fun). These features are important for the representation of games, because they can be used to make inferences or to communicate about games. These properties do not define something as a game, however. That is done by virtue of the position game within a simple relational structure. In the following section, we will discuss the role of category features in more detail.

One advantage of defining a category in terms of the relation it plays is that it enables the same item to be seen as a member of different categories. For example, a basketball fan views basketball as a game, because the fan sees basketball as something played by the participants. A professional basketball player may see basketball as a job and not a game, because it is work, and 'job' is the role-governed category that specifies the second argument to the relation work(x,y). Finally, an executive working for a professional basketball team may view basketball as a 'product' and not a 'game', because corporations may be seen as manufacturing something, and basketball may be seen as the product of this manufacturing process. None of these views of basketball are mistaken, nor are they metaphorical uses of categories. For an executive, basketball really is a product, because it plays that role within a relational system in which he or she is engaged.

Finally, we can see the use of role-governed categories in putative metaphorical uses of categories as when a parent says to a child, 'Be serious, this is not a game'. The parent is communicating that the offending event in question should not stand in the same relationship to the child as a game does. There are two important parts to this example. First, a role-governed category communicates that an item plays a particular relational role. Thus, role-governed categories are associated with a particular relation or relational structure.

Second, this view provides a way of comprehending negation. There are a myriad of ways that an object could fail to be a member of some category, yet people immediately understand negations like 'This is not a game'. The role-governed view suggests that negation blocks the possibility that the object fills that relational role. The probabilistic feature-based view of categorization has difficulty with these kinds of negations. On this view, if an item is specified as not being a game, it is not clear what is being communicated. There are a variety of features that games might have, and there are no good principles for specifying which ones are being negated.

A philosophical account of role-governed categories might specify the types of relations that can serve as the basis of these categories. However, we take a pragmatic view of categories in general. Category representations are formed when the category occurs sufficiently often in human experience to warrant a stored

representation. Thus, categories that might seem bizarre to us would come to be stored if they were important for subsequent cognitive processing. This aspect of categories makes them difficult to study, because there will be no logical principles that specify that a particular item will or will not be stored as a category.[2]

To summarize so far, we have suggested that many categories can be construed as naming items that play a particular relational role. We have contrasted game (which can be defined as being the second argument to the relation play(x,y)), and job (which can be defined as being the second argument to the relation work(x,y)). In the next section, we discuss the relationship between role-governed and feature-based categories. Then, we deal with a number of complexities raised by this view.

3.2. *Role-governed and feature-based categories*

The fact that role-governed categories are defined by their position within a more global relational structure does not mean that they are independent of the feature-based view of categories. Games, for example, still have features associated with them, such as being fun or involving scoring. Indeed, the probabilistic view of categorization is likely to be a reasonably accurate description of the kinds of information people store about categories (though for present purposes we remain agnostic about whether people are storing prototypes, some set of exemplars, theory-based information or some combination of the three).

These two types of representations must work in concert in order for the cognitive system to use categories effectively. Classifying a new instance as a member of a category requires making use of available features. The features of a newly presented exemplar that are likely to be available are perceptual properties, or perhaps information that might appear in a discourse context. Much of this information may be unrelated to the relational role that defines the category. Nonetheless, these properties are often useful for recognizing that an item is an instance of a category. For example, seeing a board, dice and pieces may reliably signal that a particular object is a game, even though these properties are not related directly to the notion of playing. This classification can take place using the kinds of mechanisms suggested by probabilistic exemplar and prototype models (Reed 1972, Nosofsky 1986, Kruschke 1992).

Once an item is known to be a member of a particular category, processing can go beyond the specific features that allowed it to be classified and to make use of this relational information. Thus, processes that operate on items that have already been classified, like inference and communication, will make greater use of this relational information. For example, this relational information may explain the finding described above that category labels tend to signal more enduring properties of categories than do features.

The view that there is a distinction between what a category 'is' and its features is similar to a suggestion by Fodor (1998). He argues that features of concepts are involved in category uses like recognition and inference, but that the category itself is not a mere amalgam of the features of category members. Where his approach differs from the one espoused here is that he suggests concepts are punctate, and as such are not related to each other in an explicit way. In contrast, we suggest that role-governed concepts name specific relational roles and, hence, are connected to relational concepts and to other role-governed concepts.

Now that we have provided an introduction to role-governed categories and their relationship to the feature-based view of categories, we take up two potential

objections. First, it might seem that this relational role information is just another feature of a category (albeit an important one) and so this proposal is just another variant of the feature-based view. It would be easier to support this argument if all role-governed categories involved the same kind of property. For example, if the relational role always specified the function of a category relative to people, then one could argue that what we are calling role-governed categories were just categories for which functional information was particularly important.

In contrast to this suggestion, the key defining properties of role-governed categories are determined by the particular role the objects play within a relational system. There are many possible relational roles that an object might fill. Though one could define the function of an item as the relational role that the item fills, this definition would miss the important variety of relational roles that an item might play. For example, 'games' are events, because of the role they play in the relation play(x,y), but 'players' are sentient beings. In addition, diverse categories like 'carnivore', 'private', 'professor' and 'oven' are all good candidates to be role-governed categories, but it is hard to see that there is a single meaning of *function* that would capture the defining aspects of these categories. Thus, we conclude that the defining properties are determined by their relational role rather than by some property of the objects.

Second, there is a potential circularity in the proposal about role-governed categories. If games are just things that are arguments to the relation play(x,y), what is to keep this relation from being defined as something that relates 'players' to 'games'? A critical part of this proposal is that not all concepts are role-governed concepts. In particular, some concepts are *relational* concepts. We will discuss relational concepts in more detail in the following section, but we give a few examples here. As one example example, many verbs (such as 'play') describe relations among a set of actors rather than relational roles (Gentner 1981, Tomasello and Merriman 1995). Prepositions also describe relations, often spatial and temporal relations (Herskovits 1986, Landau and Jackendoff 1993).

The meaning of relational concepts goes beyond just the arguments they relate. For example, Norman and Rumelhart (1975) present an extensive discussion of possible decompositional representations for relational concepts. For example, Gentner (1975) examined the difference between verbs of transfer like give, trade and sell. Verbs like give have simple relational structures involving only a simple transfer. For example,

$$\text{cause[do(x,a), change[possess(x,z), possess(y,z)]]} \quad (1)$$

In this case, x (the *giver*) has performed some action that caused some object z (the item given) possessed by the giver to be possessed by y (the recipient). A verb like 'give' can be related to other verbs like 'trade', because they involve some similar relational structure, except that trade also involves a reciprocal transfer of a second item from the second person to the first.

While there is admittedly no complete theory of relational meaning (as there is no complete theory of any aspect of semantics), this discussion suggests that relational meaning will not need to reduce to just some function of its arguments. Thus, it is not circular to define role-governed categories in terms of the relations in which they participate.

To summarize, role-governed categories and feature-based categories work in concert. The relational role played by a concept determines what it 'is', but it does

not provide information about the features of category members. Thus, tasks that rely on category membership will be guided by the relational role information. Tasks that often occur before category membership is known will be strongly influenced by the featural information. In the next section, we explore the kinds of categories that are likely to be role-governed.

3.3. *When are role-governed categories used?*

Not all categories are role-governed categories. In this section, we discus three types of categories: feature-based categories, role-governed categories, and relational categories. We outline the types of categories most likely to fall in each type and discuss the relationships among them. In addition, we speculate on the existence of a fourth type of category—*transformational categories*—at the end of the section.

We begin with a disclaimer. As discussed in the previous section, all role-governed categories have associated feature-based representations that contain information about characteristic properties of category members. Thus, when we discuss categories that are most likely to be role-governed categories, we are assuming that these categories also have feature-based representations.

3.3.1. *Relational categories*. In order for there to be role-governed categories, there must also be relational categories. Relational categories are those that pick out relations among items in the environment. There are (at least) four kinds of relational categories: verb categories, event scripts, preposition categories and comparatives (see Gentner and Boroditsky 2001 for a discussion of types of relational categories). We refer to some of these categories with the name of their grammatical class, because these categories are typically denoted in language with particular grammatical classes.

Much work has been done on verbs. A key aspect of verbs is that they tend to point outward to other concepts rather than having connections primarily to other representational elements that are part of the same relational structure (Gentner 1981, 1982a, Choi and Bowerman 1991). For example, we discussed the concept 'give' in the previous section. 'Give' is a verb that relates a giver, a recipient and an item given. A simple potential representation for 'give' based on the scheme developed by Norman and Rumelhart (1975) was presented above. The verb names that relational structure.

Many verbs describe actions. Events, which contain many actions, are also relational categories. For example, a person's script for going to the doctor's office contains information about relations among the subcomponents of the visit (like waiting in a waiting room and being examined), as well as causal relationships among events (Schank and Abelson 1977, Schank 1982). Thus, an event representation—like a verb representation—contains many links to concepts outside the event. For example, a doctor's appointment involves a particular doctor, patient and office.

Prepositions are also relational categories (Herskovits 1986, Cienki 1989, Landau and Jackendoff 1993, Regier 1996). Like verbs, they point outward, linking a small number of objects or events (usually two). Prepositions describe a number of different kinds of relationships. For example, prepositions are often used to help specify the arguments to the relational categories denoted by verbs. A key function of prepositions is to describe spatial or temporal relationships. In this way, prepositions may name categories of spatial relations. For example, 'on' denotes a support relationship in which one object rests on a second.[3]

Comparative categories are those that state possible relationships among quantities. For example, the concepts 'same', 'equal', 'greater' and 'less' are relational concepts that specify a particular relationship among a set of items. These terms can be quite difficult to learn. Smith (1989) suggests that children start with an ability to detect global similarities (e.g. that a pair of objects is identical in all respects) and only gradually do they become able to detect similarities along particular dimensions (e.g. that a pair of objects is similar only in color). Thompson and Oden (2000) suggest that primates do not display an understanding of the concept 'identical' (e.g. that two identical squares are in the same relation as two identical triangles) unless they have been given training with symbols (see Premack 1983 for an extensive discussion of this issue).

One thing that seems to be true across the types of relational categories is that it is difficult to verbalize the content of relational concepts (e.g. Sieck *et al.* 1999). This difficulty may be related to the fact that relational concepts must bind together a set of item concepts (often labelled in language as nouns). These noun concepts may have properties (particularly perceptual properties) that are more accessible than those of relational categories (Barsalou 1999). This difficulty talking about the content of relational categories may also contribute to problems people have with giving definitions for concepts. To the extent that role-governed categories are defined by their position within a relational structure, people may have little verbal access to the content of the relational structure beyond the label for a relational category that denotes that relational structure.

3.3.2. *Role-governed categories*. The reason for presenting relational categories first is that role-governed categories name particular roles within a relational structure defined by a relational category. When we introduced the notion of a role-governed category, we focused on 'games', which were an argument to the relation 'play(x,y)' (which is labelled by a verb). We suggest that role-governed categories are specific to particular verb and event categories. Preposition and comparative concepts are too general to provide the basis for role-governed concepts. For example, the argument to a spatial preposition would just be an object, because it is something that can have a spatial extent.

Role-governed categories are most often derived from verb and event categories. There has been substantial research on the types of arguments that verbs and events may have (e.g. Fillmore 1968, Levin 1993, Pustejovsky 1995). Role-governed categories reflect this diversity of argument types. First, using play(x,y) as an example, 'player' is a role-governed category for the agent of this relation and 'game' is a role-governed category for the object of this relation. In these examples, role-governed categories may be types of people, or they may be types of objects.

This proposal is related to Barsalou's (1983) work on ad hoc categories. Barsalou pointed out that ad hoc categories like 'things to wear on a first date' have many characteristics that make them like taxonomic categories. For example, just as taxonomic categories have a graded typicality structure (e.g. robins are more typical birds than are chickens), so too do ad hoc categories (e.g. ties are more typical attire for a first date than are Speedo bathing suits).

From the standpoint of feature-based views of categorization, similarities between ad hoc categories and taxonomic categories are a puzzle, because taxonomic categories are assumed to represent stable sets of features that describe category members. Ad hoc category members do not share a stable set of featural similarities.

From the perspective of role-governed categories, however, ad hoc categories are role-governed categories that involve unfamiliar relations. Ad hoc categories are defined by their position within a relational structure, though that relational structure may have to be constructed on-the-fly. Thus, while ad hoc categories are unlikely to have been stored in advance, they can be computed using the same processes that serve other role-governed categories. An ad hoc category that was used frequently would be stored in memory.

In the role-governed categories discussed so far (as well as with ad hoc categories), the category may bind together a heterogeneous set of individuals. For example, there is significant diversity in types of games. This diversity was the basis of Wittgenstein's discussion of the role of necessary and sufficient features in category representations. However, role-governed categories need not have this diversity. For example, an 'iron' is the implement used to press clothes; most irons used in people's homes have a characteristic shape. Nonetheless, classifying something as an iron suggests that it is meant to play a role in the relational structure of pressing clothes. An object with a novel shape and strange components may still be an iron if it plays the right relational role.

Much research has focused on the role of function and intended function on the categorization of artefacts, particularly in children (Keil 1989, Malt and Johnson 1992, Bloom 1996, 1998, Matan and Carey 2000). There are three basic results of this research. First, if an artefact is substantially transformed and then used for a new function, then people generally judge this object to be a member of the new category. For example, if an electric clothes dryer is dismantled, its heating element placed on a metal plate cut from the top of the dryer and the door handle placed on the top of the new contraption so that it can be used to press clothes, people judge that the new item is no longer a clothes dryer, and will judge that it has become an iron if the added features are consistent with those of known irons (Malt and Johnson 1992).[4]

Second, if an artefact is not substantially transformed, but it is used for a different purpose, it is judged to be a member of the category based on its original (intended) function. For example, if a person had a clothes dryer that they only used to put in dry wrinkled shirts to get them smooth, people would typically judge this to be a dryer (albeit one used for a strange purpose).

Finally, function is neither necessary nor sufficient for category membership. As Malt and Johnson (1992) point out, people's judgments of the likelihood category membership of an object with normal function, but bizarre features are significantly lower than their judgments for items with normal function and no information about features, suggesting that the presence of bizarre features decreased their belief that the object fit in the category (see Gentner 1982b and Kemler-Nelson *et al.* 2000 for related findings with children).

Role-governed categories are obviously compatible with the first set of findings. If an object becomes one that is used only for a single purpose, then its category membership should change to that of the new role that it plays. The second finding is more problematic from the standpoint of role-governed categories. A clothes dryer that is only used to smooth wrinkled shirts is still judged to be a dryer. Why does this object not change its category membership?

It is important to bear in mind that category structures are enduring representational entities in the cognitive system. Thus, they are meant to be useful for long-term interactions with objects in the world. Many artefacts can be used for functions other than the ones for which they were developed. If artefacts were recategorized

every time they were used for a different function, then the category system would be too labile to be useful. Thus, artefacts are recognized as members of a particular category using their feature-based representations. Still, there is a sense in which the defining nature of relational roles can be seen. Smith *et al.* (1974 following a suggestion by Lakoff) propose that the hedges that are reasonable to use with a category may suggest the status of its features. In particular, the hedge 'technically speaking' can be used with objects that have defining properties of a category, but not its characteristic features. Consistent with this observation, in the example above, we could say that 'technically speaking' the object in question is an iron.

So what use is the idea of a role-governed category if its feature-based representation is going to drive categorization in the situation just described? The dryer/iron example is generally farfetched (as are many examples that appear in psychology experiments). Most artefacts are created for a particular purpose and they retain that purpose. In fact, it is more often the case that there is variation in critical features and stability at the level of function. For example a 'dishwasher' may be a small front-loading unit with two racks. It may also be a large metal object in which racks are passed through it on a conveyor belt (in large cafeterias). Indeed, one of the authors heard the sentence 'We also have a dishwasher at home... his name is Art' spoken while he was in graduate school. What makes it appropriate to categorize each of these items as a dishwasher is that they play a common role in a relational structure, despite their manifest differences in perceptual features and internal mechanisms.

Role-governed categories may also be used to explain why putative essential properties of objects do not predict how an object will be classified. In one set of studies, Malt (1994) examined the concept 'water'. Most people would agree that if 'water' has an essence, it would be the chemical structure H_2O, yet the amount of H_2O in a variety of liquids was not a good predictor of the label that it would be given. For example, lemonade is not called water, yet it has a high percentage of H_2O. Ocean water may have a lower percentage of H_2O than lemonade, but it is called water.

We argue that there are two different water categories, one of which is a role-governed category. The role-governed version of water refers to standing bodies of clear liquid. For example, lakes, oceans, pools and puddles contain water, even though ocean water is mixed with salts, and pool water is mixed with a variety of chemicals like chlorine. The feature-based version of water refers to a clear odourless tasteless liquid that can be drunk. Any potable liquid with flavouring falls into a different category. This category is feature-based because it relies on a number of features beyond its relational role. It may seem unparsimonious to posit two separate categories that are labelled water, but, in fact, the term 'water' appears to be used in quite different contexts. Furthermore, as Klein and Murphy (2001) demonstrate, the many meanings of polysemous words appear to be treated as psychologically distinct just as are the meanings of homonyms.

As this example demonstrates, sub-types of role-governed categories are often differentiated by the particular features they have and hence are feature-based categories. A 'beverage' is any item that serves as the second argument to the relation drink(x,y), but the subtypes of beverages are distinguished by having particular properties. Water is one kind of beverage, and others (like lemonade or iced tea) are distinguished from water on the basis of their features.

To summarize, role-governed categories are often related to relational categories that specify the overarching relational structure in which they are embedded. Many role-governed categories specify participants or objects in verbs and events. Many general artefact categories are also role-governed categories. These role-governed categories tend to be the most general category that could satisfy a role (e.g. 'beverage'). The various subtypes of role-governed category are feature-based categories.

3.3.3. *Feature-based categories*. We will have somewhat less to say about feature-based categories, because they are the type most frequently studied by psychologists (and so quite a bit has been said already). All categories have feature-based representation, which provide information about properties that are characteristically associated with category members. What separates feature-based categories from relational and role-governed categories is that they only have feature-based representations.

Proposals for feature-based representations have suggested that some features that describe one concept may refer to other categories (or express relations to other categories). Barr and Caplan (1987) distinguished between intrinsic features, which refer only to a property of an object and extrinsic features, which require knowing about other objects to be interpreted. For example, for the concept 'hammer', having a handle would be an intrinsic feature, but being used to pound nails would be an extrinsic feature (because it expresses a relation between a hammer and nails). Goldstone (1996) also explored the acquisition of categories whose features are interrelated in laboratory studies. This work suggests that feature-based concepts need not be independent, as the properties of one concept may entail knowing about others.

In the previous section, we discussed that one meaning of the word water refers to a feature-based category consisting of a clear, tasteless liquid that is often used for cooking. This feature-based category is a subcategory of the role-governed category 'beverage' (which names the second argument to the relation drink(x,y)). Other contrasting feature-based categories of this type are 'lemonade' and 'tea'.

We suggest that it is characteristic of role-governed categories that subcategories of them are feature-based categories. These feature-based categories do not have a strict definition. For example, lemonade is generally yellow or pink, tastes like lemons and is both sweet and bitter. However, any of these properties could be missing, and the drink might still be classified as lemonade.

The categorization literature contains many models of feature-based category representations including prototypes, exemplars and theory-based information. The proposal for role-governed categories is independent of these distinctions. Thus, any type of feature-based category representation would be compatible with the presence of role-governed categories.

3.4. *Transformational categories*

So far, we have described what we believe are the three main types of categories: relational, role governed and feature based. In addition, we have taken a strong stance on role-governed categories by arguing that they are defined by the relational structure in which they take part (rather than simply involving the relational role). For example, we have examined the role-governed category 'game', which is defined

as the second argument to the relation play(x,y). We further suggested that 'player' is a role-governed category defined as the first argument to this relation.

A potential problem with this view is that there are often cases where more than one category of items may fill the same relational role. If a definition is supposed to specify all and only those members of a particular category, then having more than one category of items that can play a relational role is a serious problem.

In fact, some features do make their way into role-governed categories in the form of *selection restrictions* on the arguments to relations. For example, the first argument to the relation play(x,y) must be filled with a sentient being (such as a computer or a person). Many models of verbs and other relational systems assume that there are some restrictions on the types of fillers that a particular role may have. These restrictions are part of the semantics of the relational category. These selection restrictions become part of the definition of a role-governed category.

Of interest for the present analysis is that when more than one category of items can fill a particular relational role, it is often the case that the other categories (beyond the role-governed category) violate one of the selection restrictions for a relational argument. For example, a 'team' may also play the role of the first argument in the relation play(x,y). A team is not a single sentient being, but rather a sentient collection of beings that are unified into an individual. We suggest that concepts like 'team' are *transformational categories*. A transformational category is one that specifies a change in a selection restriction for a relation. For example, 'team' is used to transform a group into an individual so that it can be used as an argument that requires an individual.

Transformational categories differ from role-governed categories in that they do not specify a particular relational role. Instead, they are designed to avoid particular selection restrictions of relational roles. One way to see this distinction is that role-governed categories are specific to particular relations. 'Player' is related to play(x,y) and 'worker' is related to work(x,y). 'Team', however, can be used with both of these relations. A team can be a group that plays a game together or works on a job together.

A second way to see the role of transformational categories is to examine the effects of negation. As discussed above, negating a role-governed category suggests that a particular relational role is not being filled by that object and hence the overarching relational structure may be the wrong one (e.g. 'this is not a game'). Negating a transformational category negates the relationship between the transformational category and the role transformed. For example, saying 'they are not a team' means that the group is not functioning as an individual. Unlike negations of role-governed categories, negations of transformational categories do not negate the whole relational structure.

Finally, just as ad hoc categories are role-governed categories involving relational systems that did not previously have a role-governed category, metonymy may be viewed as a linguistic device for creating novel transformational categories. In a metonymy, one property of an individual is taken to stand for the whole. For example, in the sentence 'the newsletter called for a rally at city hall', the category 'newsletter' is used to refer to a transformational category in which a group (the people who wrote a newsletter) is treated as an individual for the purposes of participating in a particular relation (calling for a rally). Further research must explore the relationships among transformational categories and metonymy.

3.5. *Hierarchical category structure*

An important observation about natural categories (particularly for object categories) is that they are often organized into taxonomic hierarchies (e.g. Rosch *et al.* 1976). For example, pink lemonade is a type of lemonade, which is a type of beverage. Most research in this area focuses on issues like the level of categorization that people are likely to use or to learn in a given situation.

One aspect of hierarchical categories that has not been the subject of significant research is that natural categories vary in the strength of the hierarchical relations among categories. For example, the categories of animals and vehicles and their subordinates are strongly hierarchical (and, not surprisingly, they form the basis for most examples used in papers on hierarchical structure in natural categories). Many artefact categories are less strongly hierarchical. For example, ovens and refrigerators can be construed as types of appliances or perhaps kitchen appliances, but they are not as obviously connected to their superordinate as are subcategories of animals or vehicles.

Furthermore, the hierarchical structure of other types of categories is even weaker (or perhaps absent). For example, there has been some research on the hierarchical structure of event categories. Morris and Murphy (1990) found some evidence that there is a basic level for event categories, but the types of events they used (e.g. shopping) did not seem to have the same strong class-inclusion relations among categories that many object categories have. Rips and Conrad (1989) suggested that events may be organized instead into groups such that if one event is a kind of a second event, then the second event is a part of the first event. For example, if pondering is a kind of thinking, then thinking is seen as a part of pondering.

Finally, verbs do not seem to have strong hierarchies at all (e.g. Huttenlocher and Lui 1979. Instead, there appear to be clusters of related verbs. For example, Pavlicic and Markman (1997) observed that people found it easier to compare verbs from a single cluster (e.g. to give and to trade) than to compare verbs from different clusters (e.g. to give and to break). However, there was no evidence that view some verbs as being clearly subordinate to others as they do for nouns.

The differences in the strength of hierarchical category structure may be explicable given the types of categories described in this paper. A central function of hierarchical organization is to support property inference. Knowing that a poodle is a type of dog licenses a number of inferences about its behaviour, how to care for it and what it will eat. Property inheritance is a particularly important aspect of feature-based categories, where the central aspect of the representation consists of features of the objects. A superordinate category like 'vehicle' is a role-governed category, but its subordinate categories are feature-based. Thus, it is not surprising that there is a strong hierarchical structure for the subcategories of vehicle. Subordinates of vehicle may inherit the relational role they typically fill from vehicle.

In contrast, role-governed categories are defined by their relational role. Thus, while their properties may be important for making inferences and for recognizing category members, their defining property links the category to a relation. Thus, even if one role-governed category is known to be a subtype of another other category, that other category will specify another relation, and so the hierarchical link between one role-governed category and another will not be strong. For example, as discussed above, 'oven' is a role-governed category that is part of a relational structure describing the baking of substances. The best superordinate for

oven (e.g. appliance) does not specify the same relation, and hence the hierarchical structure for this category is not particularly useful.

More generally, this proposal may explain why people's concept hierarchies are not very deep (Randall 1976). Indeed, in most studies of taxonomic hierarchies, the hierarchical structure does not appear to get more abstract than the superordinate category (e.g. 'animal' is not obviously connected to a more superordinate category). We suggest that most superordinate categories (e.g. vehicle or beverage) are role-governed categories. There is little utility to having class-inclusion relations above these role-governed categories and so they are not stored. Thus, while it is logically possible to have concept hierarchies that extend all the way to some very abstract concept (e.g. 'thing', Lenat and Guha 1990), people's concept hierarchies do not extend beyond superordinate categories.

Finally, verbs and many events are relational categories. Relational categories may be similar to other categories, but there is no basis for property inheritance from one relational category to another, and hence there is no need for strict class inclusion relations among relational categories.

3.6. *Summary of the initial category types*

Figure 2 summarizes the four types of categories that we have proposed in this section, as well as the relations among them. Relational categories specify a relation or a relational structure. Role-governed categories name particular relational roles within a relational structure. Both relational categories and role-governed categories are associated with feature-based representations that provide information about properties that are typical of category members. These features may be used to recognize items and as the basis of predictions when category membership is used to make inferences about other properties.

Transformational categories can be connected to role-governed categories. These categories provide extensions of the relational role filled by a role-governed category. Finally, some categories have only feature-based representations. These feature-based categories are the type that have been studied most extensively in cognitive psychology. Often, these categories are subtypes of role-governed categories (as in the example of water, lemonade and tea described above). However, it is possible for there to be feature-based categories that are independent of role-governed categories. This analysis further suggests that category hierarchies are most prominent for feature-based categories and may be less useful for role-governed and relational categories.

Finally, we admit that we have taken a rather strong stand on role-governed categories in this paper by assuming that they are defined by their position within a relational structure. Many cognitive psychologists are (justifiably) uncomfortable with the notion of a definition. Another way to cast this proposal is to consider two sources of meaning in mental representation. A mental representation may come to have meaning either by virtue of corresponding to some entity in a represented world or because of its connections to other representational elements (see Dietrich and Markman in press, Markman 1999 and Markman and Dietrich 2000). The first source of meaning can be called *correspondence*, and the second can be called *functional role*.

The feature-based view often focuses on correspondence as a source of information about concepts. The properties in the concept representation are all supposed to correspond to some property either in the world or in some other representational

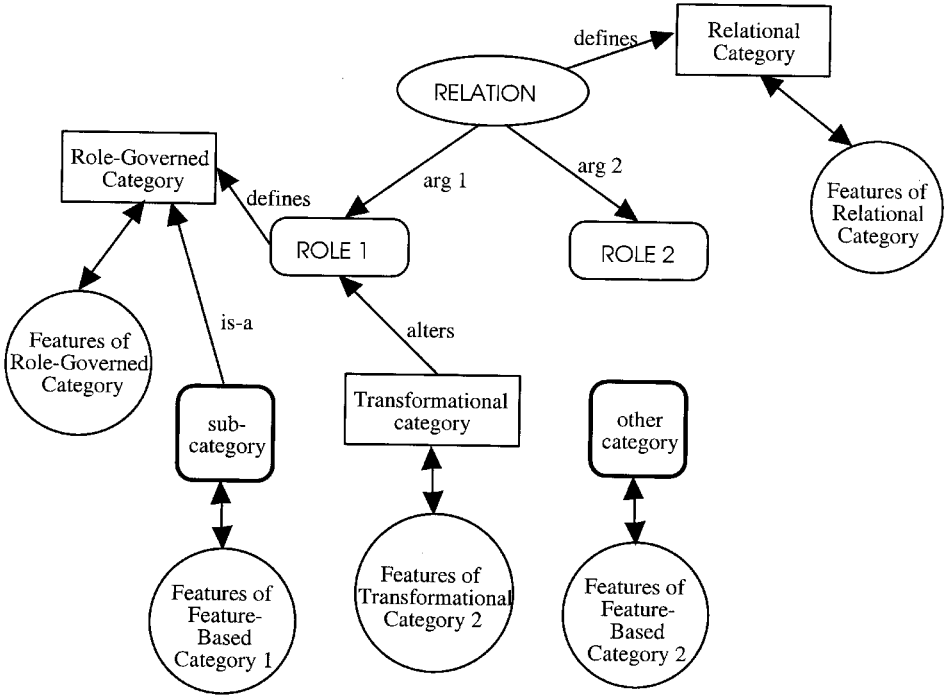

Figure 2. Summary of the types of categories described in this paper. Relational categories refer to relational structures. Role-governed categories refer to particular roles within a given relational structure. Transformational categories may be used to alter the properties of a particular relational role. Feature-based categories are structured only by the set of features they possess.

system. Thus, feature-based categories derive most of their meaning from the correspondence component of mental representation. In contrast, role-governed categories derive their meaning as a result of their (relational) connections to other representational elements. Thus, role-governed, relational and transformational categories derive most of their meaning through the functional role component of mental representation. Readers who are uncomfortable with categories having definitions may substitute this distinction without altering the core of the proposals in this paper.[5]

4. Learning and use of role-governed categories

In order to complete our discussion of role-governed categories, we must now focus on how role-governed categories are learned and used. Category learning has long been a central focus of research (e.g. Shepard *et al.* 1961), although most of this research has focused on how people learn to classify new items. There has also been research on category acquisition in children, though most of this work has focused on how children learn new labels (often count nouns). Finally, there has been an

increase in attention to category uses. Significant work has explored induction. Research has also been directed at other functions of categories such as communication and metaphor. We will examine uses of role-governed categories at the end of this section.

4.1. *Category learning*

Much work with children and adults examines category learning. In this section, we examine the implications of the existence of role-governed categories for this research. Most research on category acquisition has focused on feature-based categories. There has been a small amount of work on the acquisition of relational concepts (notably in the verb learning literature). There has been almost no research on learning of role-governed categories. Furthermore, as we will see in our discussion of category acquisition in adults, there are aspects of the methodologies used to study category acquisition in adults that are incompatible with role-governed categories.

4.1.1. *Category acquisition in adults*. Our proposal about category representation suggests a potential danger with the focus of much research on categorization in psychology. The typical study employs a variant of a classification method. For example, people may be shown an exemplar and be asked to classify it into one of a small number of groups. On each trial, feedback is given and learning continues until the items can be successfully classified. Another variant of this task is sorting, where people take a set of items and sort them into groups. On our view of categorization, this task overemphasizes the importance of features subordinate to the category. As discussed above, the kinds of feature structures posited by the feature-based view of categorization are important for classifying instances.

Two important things are missing from classification tasks. First, the category label is not used for anything in these tasks, and so there is no way to assess whether people have acquired information superordinate to the category. Second, the categories themselves are typically independent of other information, and so there are no relations that the items participate in.[6] As a result, the categories learned using the standard laboratory methods are not role-governed categories. Both of these factors bias classification studies toward feature-based categories.

There has been an increasing interest in the influence of category use on category learning (Markman *et al.* 1997, Ross 1996, 1997, 1999, 2000, Yamauchi and Markman 1998). This work has begun to explore tasks that may be more amenable to the exploration of role-governed categories. For example, Ross (1999) has examined medical diagnosis and treatment situations. He finds that features of diseases that are relevant to treatment decisions become more important in the category representation than features of the disease that are equally predictive of the disease, but are not relevant for treatment. Treatment decisions involve considering relationships among categories (e.g. a particular drug may be used to treat a particular disease). Thus, tasks like this are more likely to permit an exploration of role-governed categories than are the classification tasks more typical of laboratory research on category learning.

One line of research on category acquisition has begun to explore the role of general relational information on category acquisition. Rehder and Ross (2001) use the term *abstract coherent categories* to refer to categories whose members have features that are consistent with some general relational information. For example,

people found it easier to learn about items that could all be used for a particular purpose, even if the specific features of each exemplar were different. This research suggests that people might treat these categories as objects that play a particular role in a relational system rather than as collections of features. Of importance, people in these studies were asked to carry out a variety of tasks with these categories including making predictive inferences.

There has also been research on the influence of communication on category acquisition. It has long been known that people in a group communication setting will settle on a common way of viewing an ambiguous situation (Sherif 1936). Furthermore, pairs of people working together in referential communication tasks with novel objects tend to settle on a vocabulary for talking about these items (Krauss and Weinheimer 1964, Clark and Wilkes-Gibbs 1986, Garrod and Anderson 1987, Garrod and Doherty 1994). If the items are discussed for an extended period of time, the people involved in the discussion tend to become more similar in the way they classify the items (Markman and Makin 1998).

This effect of communication may help to explain a dissociation between similarity and category labels observed by Malt *et al.* (1999). They found systematic differences among the labels given to a set of pictures of containers (e.g. bottles, jars and boxes) by speakers of different languages. These differences occurred despite significant similarities in the similarity judgements given to these items by native speakers of these languages. Malt *et al.* suggested that communicative factors may influence the labels used in this task.

It is interesting to note that the objects used in the study by Malt *et al.* (1999) are primarily members of role-governed categories. For example, what makes something a 'jar' in English is not a particular set of features, but rather that it is expected to participate in a relational structure in which the object contains a substance that must be scooped out of it. Jars may be distinguished from bottles, which participate in a different set of relations (having the substance contained in them poured out rather than scooped out). While the perceptual features of an object may stay constant, the particular relation that someone expects an object to participate in (as well as the relations that someone finds salient) may vary substantially across cultures, leading to the observed variability among speakers of different languages.[7]

A third area of research that may help us better understand role-governed categories is research on event categories (Morris and Murphy 1990, Kersten and Billman 1997, Zacks *et al.* 2001). Events are relational categories that capture the relations among a set of actors. For example, shopping is an event category that describes a trip to the store. The participants in a shopping event, however, are role-governed categories. For example, a cashier is the person who determines how much someone should pay and carries out the monetary transaction within a shopping event. Exploring the way event categories are learned and structured should be useful for understanding role-governed categories.

Current research tends to treat event categories as ordinary feature-based categories. For example, Kersten and Billman (1997) were interested in the features of events that people tended to attach to an event category. In their studies, the agents (i.e. the potential role-governed categories) were simply geometric shapes that remained constant during learning. This tendency is also reflected in studies of verb acquisition in children (e.g. Behrend 1990, Naigles 1990), which tend to show an event to children along with a verb and to examine the features of events that must be present in order for the same verb to be applied later.

One interesting aspect of verb learning that is relevant to role-governed categories is that there is evidence that both children and adults have a bias to name actions using instrumental verbs (Behrend 1990). An instrumental verb is one that names the instrument used to perform an action as part of the verb. For example, 'to saw' incorporates the instrument (a saw) as part of the verb. Instrumental verbs also appear to be relatively easy for children to learn (Huttenlocher *et al.* 1983, Behrend 1990). This ease may reflect that instrumental verbs are strongly associated with particular role-governed categories. Learning these instrumental verbs may also be a vehicle that allows children to transition from having a feature-based category of an artefact to having a role-governed category (Gentner 1978, 1981).

In addition, there is evidence from research on the relationship between nouns and verbs that adults have some role-governed categories that are associated with particular verbs (McRae *et al.* 1997, Ferretti *et al.* 2001). For example, the noun *cop* is associated with the verb *arrest*, because police tend to be the agents who arrest criminals. Consistent with this relationship, Ferretti *et al.* (2001) found that verbs reliably prime nouns that name categories that are typically associated with them as agents and instruments of an action (though typical locations are not primed).

As these examples demonstrate, research on category acquisition in the lab has typically biased against role-governed categories. However, the development of new techniques that focus on other uses of categories is likely to provide evidence that can be used to constrain theories of how role-governed categories are learned. In particular, research that focuses on category uses that involve relations among categories is likely to be fruitful for shaping ideas about role-governed categories.

4.1.2. *Category acquisition in children.* Laboratory studies of category acquisition in preschool children tend to use more naturalistic tasks and materials than studies of adults. Many experiments have looked at how preschool children learn new labels (primarily nouns). There has also been significant discussion of the order of acquisition of nouns and verbs. As we discuss in this section, both of these issues bear on the idea of role-governed categories.

Studies of the acquisition of nouns have used a fairly naturalistic task. Preschool children are shown an object and it is given a novel label, and they are shown some additional objects and are asked to select the item that also shares the same label (e.g. 'this is a *dax*, can you show me another dax?'). This task is similar to the way many parents introduce new labels to their children.

One proposal for the way children learn new labels is that children assume novel labels refer to taxonomic categories (Markman and Hutchinson 1984, Markman 1989). The basis of taxonomic categories has not been well-defined, but this term refers to sets of objects that tend to be featurally similar in some way. More recent work has suggested that novel labels tend to be extended on the basis of specific types of properties. For example, novel count nouns tend to be applied to objects with a common shape (Imai *et al.* 1994, Landau *et al.* 1988). These proposals about word learning (and presumably the categories underlying these words) are feature-based.

Gentner (1981 1982a) proposed that nouns tend to be learned before verbs and other relational categories. On this view, verbs are relational categories that bind together a number of actors and objects (as described above). In order to learn a relational category that binds together a set of items, it is logically necessary to have some items that can be related. Thus, children must learn at least some nouns before they can acquire verbs. While there has been some dispute in the literature about the

degree to which noun learning precedes verb learning, there is much evidence that supports this proposal (though see Gopnik *et al.* 1996, Tardif 1996).

This developmental account suggests that role-governed categories will be learned late. In particular, children may first acquire feature-based noun categories. Once a few noun categories have been established, children may then learn relational categories like verbs and events. These relational categories bind together feature-based categories. Once relational categories have been acquired, role-governed categories may then be learned. That is, from a logical standpoint, it is not possible to learn a category that specifies a particular relational role until the relation has first been learned.

It is possible that role-governed categories are important in children's early word learning. In particular, many researchers have noted that after about 18 months children begin to learn words at a rapid rate. This phenomenon has sometimes been called the naming explosion (e.g. Markman 1989, though see Bloom 2000). One reason for this increase in the rate of word learning may be that as children begin to learn relational concepts, they develop the relational structure necessary to learn role-governed categories. Thus, only after some feature-based and relational categories have been acquired can children also acquire role-governed categories.

This section has been somewhat speculative. The bulk of the research on concept acquisition in children has focused on feature-based categories. There has been some attention to relational categories, particularly in the verb learning literature. By making a finer set of distinctions about types of categories, it may be possible to explore concept acquisition in children in a more fine-grained manner. In particular, while research has distinguished between syntactic categories of nouns (e.g. mass nouns vs. count nouns), this research has not distinguished among types of noun categories. Most of the nouns that have been explored have been ones that are likely to be feature-based categories like 'cow'. Much less attention has been paid to role-governed categories like 'game'.

4.1.3. *Summary*. Research on category acquisition in adults and children has focused primarily on feature-based categories with some work on relational categories. Role-governed categories have not been studied. One important reason for this omission is that the methods employed in most studies of category acquisition are incompatible with role-governed categories. The artificial stimuli used in studies typically create a small number of classes that differ from each other along a limited set of dimensions. The categories embody a feature-based view, because only feature information can be used to distinguish among categories. The categories are not embedded in a more global relational structure, and thus, neither relational nor role-governed categories are relevant in these studies.

In order to explore the acquisition of role-governed categories in the lab, methods must be developed that promotion the acquisition of relational information and relational role information. This will not be an easy task. In particular, in order to learn a role-governed category, people must first learn relational categories. Most existing relational categories already have role-governed categories associated with them, and so it is difficult to create novel role-governed categories that take advantage of people's background knowledge. Instead, more elaborate experimental situations must be developed in which a relational system is learned so that role-governed categories can be embedded within that structure. Given that the bulk of

psychological studies with adults are done with college undergraduates who participate in a study for about an hour, it is not surprising that most research has focused on feature-based categories, which can be learned independent of other knowledge.

4.2. *Category use*

Role-governed categories may also help us to understand other category uses besides classification. In particular, the feature-based view has tended to focus researchers on theories that posit that similarity relationships among categories influence the way categories are used. These theories have also focused on how features of one category might be carried over to another. If theories of categorization take into account the relational structures in which categories are embedded, new explanatory constructs for these category uses may become available. We illustrate this point with examples from category-based induction, metaphor and conceptual combination.

4.2.1. *Induction*. Induction is a critical function of categories. Knowing that an object is a member of a category allows us to predict a number of other properties that it may have. These predictive inferences are so central to category use that Anderson (1990) based his rational theory of categorization around the use of categories to make feature predictions.

Because of this importance, a number of theories have been developed about how people use their categories to make predictive inferences. In a classic paper, Osherson *et al.* (1990) examined factors that influence the strength of inductive arguments. For example, they found that people were more likely to posit that all birds have some property X, if all robins have it than if all penguins have it. In order to explain how people judge the strength of inductive arguments of this type, Osherson *et al.* posited the similarity-coverage model. This model is feature-based and relies on featural similarity among categories as the basis of predictive inferences. Essentially, the more (featurally) similar the categories in the premise are to the categories in the conclusion of an inductive argument, the stronger the argument is judged to be. Furthermore, the more that the categories in the premise span the range of items within a category specified by the conclusion of an argument, the stronger an argument is judged to be. Sloman (1993) developed an alternative feature-based approach to explain patterns of category-based induction.

Heit and Rubinstein (1994) extended these feature-based accounts by suggesting that people consider only features of particular types when evaluating feature predictions. For example, when evaluating a new behavioural property, only features relevant to behaviour might be examined, whereas when evaluating a new morphological property, only features relevant to morphology might be considered.

A number of findings suggest that these featural accounts may need to be extended. As one example, when experts make predictive inferences about new features of categories, they do not rely on similarity the way novices do. (Coley *et al.* 1999, Medin *et al.* 1997). Instead, causal or ecological factors may influence their judgments. For example, Coley *et al.* (1999) found that when tree experts reasoned about whether a particular tree would get a disease given that some other tree had that disease, they used information about how widely planted the trees were and how disease-resistant the trees normally are to make their decisions (see also Proffitt *et al.*

2000). Feature-based information was not used as strongly by the experts as by novices making predictions about novel properties.

Understanding patterns of expert prediction may require attention to the relational structures that are critical for role-governed categories. For example, Heit and Rubinstein (1994) demonstrated that when a new property being inferred about an animal is related to whether that animal is a predator, people focus selectively on properties related to predatory behaviour. The role-governed category view suggests that 'predator' is a role-governed category. The relational structure associated with predators thus becomes salient when this category is used. This relational structure may then be used to reason about potential behaviours of an animal. Predators need not share specific features with each other. Instead, they must have properties that allow them to successfully hunt other creatures. The particular features necessary to hunt successfully will depend more on the environment in which the animal lives and on the methods it uses to hunt (e.g. individually or in a group) than on its category membership. Thus, by using a role-governed category, the induction process can be focused on relevant relational information rather than on specific features.

Consistent with this suggestion, there has been some research suggesting that relations (within the feature-based portion of a category representation) are important determinants of category-based induction. Lassaline (1996) demonstrated that shared features between a pair of categories increased their perception of the similarity of the items but not their rated strength of an inductive argument. In contrast, a relation that connected a shared feature to the feature to be inferred increased the strength of inductive arguments. A similar finding was obtained by Wu and Gentner (1998).

The effects of expertise on category-based induction are also compatible with role-governed categories. For example, Proffitt *et al.* (2000) suggest that tree maintenance workers will classify some trees as 'weed trees'. Weeds in general (and weed as a subtype of weeds) are role-governed categories. They are plants that are growing in an area where they are not wanted. Reasoning about the growth of weed trees and how to deal with them requires having relational knowledge about trees. Experts have this type of relational knowledge that they use to reason in a domain.

This discussion touches on an important reason why the use of role-governed categories has been minimized in many previous studies of categorization. Most laboratory studies focus on categories that are unfamiliar to the people learning and using them. This focus on novel categories is based on two sources. First, in order to reduce variability across participants, category learning studies focus on materials that are not connected to other things that research participants might know about. In this way, categorization be studied independent of other knowledge effects.

Second, category learning research has its roots in work on animal behaviour. Early classification studies (e.g. Shepard *et al.* 1961), whose methods have shaped the kinds of studies that are still performed, used methods that were similar to those used to train pigeons and rats. In these studies, a stimulus is presented, a response is made and feedback is given. No other attempt is made to communicate about the categories with the participants. This method deviates significantly from the conditions under which people generally learn categories (Schank *et al.* 1986), but it is comparable with the kinds of methods used to study animal behaviour.

By minimizing the effects of prior knowledge and of social interactions in category learning, research has also focused psychology on feature-based categories to the exclusion of relational and role-governed categories. In order to study role-governed

categories, it will be necessary to perform more studies in which the categories are embedded in relational structures that are either learned during the experimental session or are within the expertise of the participants in the study.

4.2.2. *Metaphor*. Role-governed categories also provide an interesting perspective on category-based metaphors. Glucksberg and Keysar (1990) suggested that metaphorical statements may be processed as class-inclusion statements. For example, the metaphor 'my job is a jail' may be interpreted as meaning that one domain (e.g. my job) is being classified in a category exemplified by the other domain (e.g. a jail).

There is some debate in the metaphor literature over whether metaphors should be considered as class-inclusion statements or instead as comparisons (Wolff and Gentner 1992, 2000, Glucksberg *et al.* 1997, Bowdle and Gentner 1999, Gentner and Wolff 2000). This work suggests that when a metaphor is novel, it may be processed more like a comparison, but that with increasing conventionality it may be processed more like a class-inclusion statement. This debate is not relevant for the current discussion, as both of these views suggest that eventually some metaphors are processed as class-inclusion statements.

At the stage where metaphors are processed as class-inclusion statements, a feature-based view of categorization still has the problem of determining what features are used as the general category exemplified by an object. For example, why is a jail a prototypical example of a confining location as opposed to a house for criminals? The role-governed category view suggests that a jail is defined with respect to a relational structure having to do with law and criminals. Thus, a jail is something that confines by definition.

On this view, there is little distinction between metaphoric and literal use of role-governed categories. Because relations can easily take a variety of objects as arguments, placing the target of a metaphor into a new relational system is essentially a literal use of that system. For example, placing 'my job' into the relational role that a jail normally holds (and substituting other elements from this domain like bosses and workloads) could be done using the same comprehension processes needed to understand a literal story about jails. This discussion suggests that the failure to find evidence for a processing difference between literal and figurative language may occur because the same comprehension processes are involved in both (Glucksberg *et al.* 1982, Gibbs 1994, Gentner and Wolff 2000).

4.2.3. *Conceptual combination*. A third aspect of category processing that we will consider in relation to role-governed categories is conceptual combination. Conceptual combination involves the interpretation of phrases that involve two or more concepts. A conceptual combination may involve an adjective and a noun (e.g. 'the brown apple') or two noun concepts (e.g. 'the cup apple'). Both types of combinations have been studied (e.g. Smith and Osherson 1984, Smith *et al.* 1988), but we will focus this discussion on noun-noun combinations.

Research on conceptual combination has described three main strategies that people use to interpret novel noun-noun combinations (e.g. 'a zebra horse'). The first is called *relational* interpretation, and involves positing a relation between the nouns in the combination (e.g. a zebra horse is a horse that *lives with* zebras; Murphy 1988, Gagne and Shoben 1997, Gagne 2000). The second is *property* combinations, which involve carrying a property from the modifier noun (e.g. zebra) to the head noun

(e.g. horse as in 'a zebra horse is a horse with stripes; Wisniewski and Gentner 1991, Wisniewski 1997). Finally, there are *hybrid* combinations, in which the interpretation is a hybrid of the two concepts (e.g. a zebra horse is a combined horse and zebra).

There are many different processes that have been proposed as the basis for conceptual combination (Wisniewski 1997, Wisniewski and Love 1998, Bock and Clifton 2000, Costello and Keane 2000, Gagne 2000). A review of these proposals would be beyond the scope of this paper. The role-governed category view suggests that using role-governed categories in conceptual combinations should lead to relational interpretations involving the relation that defines the role-governed category. For example, as we discussed above, an 'iron' is an object used to press clothes. A 'carpet iron' might be an iron that is used to press carpeting. In this case, the role-governed category activates the relation that helps to define it, and suggests that the other argument to that relation can be filled by the modifier noun. Current research in conceptual combination does not distinguish among types of concepts, and thus may be missing a source of constraints on the novel noun-noun combinations are interpreted.

5. Conclusions

The feature-based view of categories dominates research and theory in psychology. Nearly all proposals for category representation as well as computational and mathematical models based on these proposals assume that categories are represented by collections of features (perhaps structured by relations). These proposals lead research to focus on the internal structure of category representations.

In this paper, we have distinguished among three primary types of categories: feature-based, relational and role-governed categories. We have also identified a fourth potential category type—transformational categories—that are used to allow an item to satisfy the selection restriction on a relational role when it would not do so ordinarily. As summarized in figure 2, category knowledge is embedded in relational structures. These relational structures may themselves be identified as relational categories. Furthermore, particular roles within these relational structures may define role-governed categories. Both relational and role-governed categories may have associated with them featural information that describes properties typical of category members. These features are important for category use.

By distinguishing among these types of categories, we aim to provide a basis for future research on category learning and category use. On the category learning front, we have noted that most current laboratory studies are biased to focus on feature-based categories. This bias arises from three sources. First, the materials in the studies are not embedded in a more global relational structure. Second, most category learning studies use a classification methodology. Classification may be more heavily influenced by the features typically associated with category members than are other functions of categories. Finally, the studies must be completed in a short period of time, which may not be sufficient for the acquisition of a novel relational structure.

Most theories of category use also assume feature-based categories. Often, relational information is needed to make inferences about categories and to understand metaphoric statements involving categories. Role-governed categories suggest which relations are likely to be important for these category uses. Further research

will have to explore the relationship between these category uses and what is learned about category members in the process of using them.

Acknowledgments

This work was supported by NSF grant SBR-9905013. The authors would like to thank Dedre Gentner, Brad Love, Todd Maddox and Douglas Medin for helpful comments on earlier drafts of this paper. Thanks also to Les Cohen, Randy Diehl, Eric Dietrich, Bill Geisler, Phil Gough, David Gilden, Ryan Gossen and Ross Speir for helpful discussions during the evolution of this project. Special thanks to Scott Stevens for great defence.

Notes

1. Many treatments of categorization use the term *category* to refer to a collection of entities in the world and the term *concept* to refer to a mental representation (Murphy and Medin 1985). In this paper, we will use these terms interchangeably. When we are interested in talking about a mental representation, we will talk about representations explicitly.
2. A further complication of this view is that the representation of a new concept is partly contingent on the representations of existing concepts. Thus, if a person has not yet learned about a particular relation, they may represent a new object category as a feature-based category. Another person who has already learned about a particular representation might represent that same category as a role-governed category.
3. It can be quite difficult to state the spatial relationship denoted by a preposition verbally (Miller and Johnson-Laird 1976). Landau and Jackendoff (1993) suggest that the kinds of relations specified by prepositions are consistent with theories of high-level visual representation like the one proposed by Marr (1982).
4. This pattern contrasts with what is observed for natural kinds, where transformations are typically judged not to change the category membership of the object.
5. One potential problem with functional role theories of semantics is that they lead to the problem of holism (Fodor and Lepore 1992), in which the meaning of one representational element depends in part on the meanings of all other representations to which it is connected. In practice, this problem is not likely to influence cognitive processing, as the limited capacity of working memory will generally focus people on information that is likely to be useful in the current context (Dietrich and Markman in press).
6. As discussed above, Goldstone (1996) examined the acquisition of categories that were interdependent, but these categories were not connected by relations.
7. Malt *et al.* (1999) found that the way objects were labelled could not be predicted on the basis of people's ratings of 'functional' similarity. To the extent that an artefact's function reflects its position within a relational structure, that would seem to contradict the present account. A measure of 'functional similarity' may be too gross to tease apart the issues discussed here. First, a relational structure in which contents are poured out and one in which they are scooped out are quite similar, and so there will be a compression in the range of similarity judgments. In addition, people may take other features into account when rating similarity such as the similarity of the context of use of a pair of items. These properties are assumed to be irrelevant to the category membership of an item.

References

Ahn, W. K., Kim, N. S., Lassaline, M. E., and Dennis, M. J., 2000, Causal status as a determinant of feature centrality. *Cognitive Psychology*, **41**: 361–416.

Anderson, J. R., 1990, *The Adaptive Character of Thought* (Hillsdale, NJ: Lawrence Erlbaum Associates).

Armstrong, S. L., Gleitman, L. R., and Gleitman, H., 1983, What some concepts might not be. *Cognition*, **13**: 263–308.

Au, T. K., 1986, A verb is worth a thousand words: The causes and consequences of interpersonal events implicit in language. *Journal of Memory and Language*, **25**: 104–122.

Barr, R. A., and Caplan, L. J., 1987, Category representations and their implications for category structure. *Memory and Cognition*, **15**(5): 397–418.

Barsalou, L. W., 1983, Ad hoc categories. *Memory and Cognition*, **11**: 211–227.

Barsalou, L. W., 1999, Perceptual symbol systems. *Behavioral and Brain Sciences*, **22**(4): 577–660.

Behrend, D. A., 1990, The development of verb concepts: Children's use of verbs to label familiar and novel events. *Child Development*, **61**: 681–696.

Biederman, I., 1987, Recognition-by-components: a theory of human image understanding. *Psychological Review*, **94**(2): 115–147.
Bloom, P., 1996, Intention, history, and artefact concepts. *Cognition*, **60**: 1–29.
Bloom, P., 1998, Theories of artefact categorization. *Cognition*, **66**: 87–93.
Bloom, P., 2000, *How children learn the meanings of words* (Cambridge, MA: The MIT Press).
Bock, J. S., and Clifton, C., 2000, The role of salience in conceptual combination. *Memory and Cognition*, **28**(8): 1378–1386.
Bowdle, B., and Gentner, D., 1999, Metaphor comprehension: From comparison to categorization. Paper presented at the Proceedings of the 21st annual conference of the Cognitive Science Society, Vancouver, BC.
Brooks, L. R., 1999, Orientation to meaning suppresses explicit knowledge of surface cues in category learning. *Abstracts of the Psychonomic Society*, **4**: 51.
Carey, S., 1985, *Conceptual change in childhood.* (Cambridge, MA: The MIT Press).
Choi, S., and Bowerman, M., 1991, Learning to express motion events in English and Korean: The influence of language-specific lexicalization patterns. *Cognition*, **41**: 83–121.
Cienki, A. J., 1989, *Spatial cognition and the semantics of prepositions in English, Polish and Russian.* (Munich: Verlag Otto Sagner).
Clark, H. H., and Wilkes-Gibbs, D., 1986, Referring as a collaborative process. *Cognition*, **22**: 1–39.
Coley, J. D., Medin, D. L., Proffitt, J. B., Lynch, E., and Atran, S., 1999, Inductive reasoning in folkbiological thought. In D. L. Medin and S. Atran (eds.) *Folkbiology* (Cambridge, MA: The MIT Press), pp. 205–232.
Costello, F. J., and Keane, M. T., 2000, Efficient creativity: constraint guided conceptual combination. *Cognitive Science*, **24**(2): 299–349.
Dietrich, E., and Markman, A. B., in press. Discrete thoughts: why cognition must use discrete representations.
Ferretti, T. R., McRae, K., and Hatherell, A., 2001, Integrating verbs, situation schemas, and thematic role concepts. *Journal of Memory and Language*, **44**: 516–547
Fillmore, C. J., 1968, The case for case. In E. Bach and R. T. Harms (eds) *Universals in linguistic theory* (New York: Holt, Rinehart, and Winston, Inc.), pp. 1–90.
Fodor, J. A., 1998, *Concepts: Where cognitive science went wrong* (New York: Oxford University Press).
Fodor, J. A., and Lepore, E., 1992, *Holism: A shoppers guide* (Cambridge, MA: Blackwell).
Gagne, C. L., 2000, Relation-based combinations versus property-based combinations: a test of the CARIN theory and the dual-process theory of conceptual combination. *Journal of Memory and Language*, **42**: 365–389.
Gagne, C. L., and Shoben, E. J., 1997, Influence of thematic relations on the comprehension of modifier-noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **23**(1): 71–87.
Garrod, S., and Anderson, A., 1987, Saying what you mean in dialogue: A study in conceptual and semantic co-ordination. *Cognition*, **27**: 181–218.
Garrod, S., and Doherty, G., 1994, Conversation, co-ordination and convention: An empirical investigation of how groups establish linguistic conventions. *Cognition*, **53**: 181–215.
Gelman, S. A., and Heyman, G. D., 1999, Carrot-eaters and creature-believers: The effects of lexicalization on children's inferences about social categories. *Psychological Science*, **10**(6): 489–493.
Gelman, S. A., and Markman, E. M., 1986, Categories and induction in young children. *Cognition*, **23**: 183–209.
Gentner, D., 1975, Evidence for the psychological reality of semantic components: the verbs of possession. In D. A. Norman and D. E. Rumelhart (eds) *Explorations in Cognition* (San Francisco: W. H. Freeman).
Gentner, D., 1978, On relational meaning: the acquisition of verb meaning. *Child Development*, **49**: 988–998.
Gentner, D., 1981, Some interesting differences between nouns and verbs. *Cognition and Brain Theory*, **4**(2), 161–178.
Gentner, D., (1982a), Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In S. A. Kuczaj (ed.), *Language development, Vol. 2: Language, thought and culture.* (Hillsdale, NJ: Lawrence Erlbaum Associates).
Gentner, D. (1982b). A study of early word meaning using artificial objects: What looks like a jiggy, but acts like a zimbo? In J. Gardner (ed.) *Readings in Developmental Psychology* (Boston: Little, Brown, and Company).
Gentner, D., and Boroditsky, L., 2001, Individuation, relativity, and early word learning. In M. Bowerman and S. C. Levinson (eds) *Language acquisition and conceptual development* (New York: Cambridge University Press), pp. 215–256.
Gentner, D., and Stevens, A. L. (1983). *Mental Models.* Hillsdale, NJ: Lawrence Erlbaum Associates.
Gentner, D., and Wolff, P., 2000, Metaphor and knowledge change. In E. Dietrich and A. B. Markman (eds) *Cognitive dynamics* (Mahwah, NJ: Lawrence Erlbaum Associates), pp. 295–342.

Gibbs, R. W., 1994, *The Poetics of Mind* (New York: Cambridge University Press).
Glucksberg, S., Gildea, P., and Bookin, H. B., 1982, On understanding nonliteral speech: can people ignore metaphors? *Journal of Verbal Learning and Verbal Behavior*, **21**(1): 85–98.
Glucksberg, S., and Keysar, B., 1990, Understanding metaphorical comparisons: Beyond similarity. *Psychological Review*, **97**(1): 3–18.
Glucksberg, S., McGlone, M. S., and Manfredi, D., 1997, Property attribution in metaphor comprehension. *Journal of Memory and Language*, **36**(1): 50–67.
Goldstone, R. L., 1996, Isolated and interrelated concepts. *Memory and Cognition*, **24**(5): 608–628.
Gopnik, A., Choi, S., and Baumberger, T., 1996, Cross-linguistic differences in early semantic and cognitive development. *Cognitive Development*, **11**: 197–227.
Gopnik, A., and Meltzoff, A. N., 1997, *Words, thoughts and theories*. (Cambridge, MA: The MIT Press).
Heit, E., and Rubinstein, J., 1994, Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **20**(2): 411–422.
Herskovits, A., 1986, *Language and spatial cognition: An interdisciplinary study of the prepositions in English* (New York: Cambridge University Press).
Huttenlocher, J., and Lui, F., 1979, The semantic organization of some simple nouns and verbs. *Journal of Verbal Learning and Verbal Behavior*, **18**: 141–162.
Huttenlocher, J., Smiley, P., and Charney, R., 1983, Emergence of action categories in the child: evidence from word meanings. *Psychological Review*, **90**(1): 72–93.
Imai, M., Gentner, D., and Uchida, N., 1994, Children's theories of word meaning: The role of shape similarity in early acquisition. *Cognitive Development*, **9**: 45–75.
Keil, F. C., 1989, *Concepts, Kinds and Cognitive Development* (Cambridge, MA: The MIT Press).
Kemler-Nelson, D. G., Russell, R., Duke, N., and Jones, K., 2000, Two-year-olds will name artefacts by their functions. *Child Development*, **71**(5): 1271–1288.
Kersten, A. W., and Billman, D., 1997, Event category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **23**(3): 638–658.
Klein, D. E., and Murphy, G. L., 2001. The representation of polysemous words. *Journal of Memory and Language*, **45**: 259–282.
Krauss, R. M., and Weinheimer, S., 1964, Changes in the length of reference phrases as a function of social interaction: a preliminary study. *Psychonomic Science*, **1**: 113–114.
Kruschke, J. K., 1992, ALCOVE: an exemplar-based connectionist model of category learning. *Psychological Review*, **99**: 22–44.
Landau, B., and Jackendoff, R., 1993, 'What' and 'where' in spatial language and spatial cognition. *Behavioral and Brain Sciences*, **16**(2): 217–266.
Landau, B., Smith, L. B., and Jones, S. S., 1988, The importance of shape in early lexical learning. *Cognitive Development*, **3**: 299–321.
Lassaline, M. E., 1996, Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **22**: 754–770.
Lenat, D. G., and Guha, R. V. (1990). Building Large Knowledge-based Systems. San Francisco: Addison-Wesley Publishers, Inc.
Levin, B., 1993, *English verb classes and alternations* (Chicago: University of Chicago Press).
Malt, B. C., 1994, Water is not H20. *Cognitive Psychology*, **27**(1): 41–70.
Malt, B. C., and Johnson, E. J., 1992, Do artefact concepts have cores. *Journal of Memory and Language*, **31**: 195–217.
Malt, B. C., Sloman, S. A., Gennari, S., Shi, M., and Wang, Y., 1999, Knowing versus naming: Similarity of the linguistic categorization of artefacts. *Journal of Memory and Language*, **40**: 230–262.
Markman, A. B., 1999, *Knowledge representation* (Mahwah, NJ: Lawrence Erlbaum Associates).
Markman, A. B., and Dietrich, E., 2000, In defense of representation. *Cognitive Psychology*, **40**: 138–171.
Markman, A. B., and Makin, V. S., 1998, Referential communication and category acquisition. *Journal of Experimental Psychology: General*, **127**(4): 331–354.
Markman, A. B., Yamauchi, T., and Makin, V. S., 1997, The creation of new concepts: a multifaceted approach to category learning. In T. B. Ward, S. M. Smith and J. Vaid (eds) *Creative Thought: An Investigation of Conceptual Structures and Processes* (Washington, DC: American Psychological Association), pp. 179–208.
Markman, E. M., 1989, *Categorization and Naming in Children* (Cambridge, MA: The MIT Press).
Markman, E. M., and Hutchinson, J. E., 1984, Children's sensitivity to constraints on word meaning: taxonomic vs. thematic relations. *Cognitive Psychology*, **16**(1): 1–27.
Marr, D., 1982, *Vision* (New York: W.H. Freeman and Company).
Matan, A., and Carey, S., 2000, Developmental changes within the core of artefact concepts. *Cognition*, **78**: 1–26.
McRae, K., Ferretti, T. R., and Amyote, L., 1997, Thematic roles as verb-specific concepts. *Language and Cognitive Processes*, **12**: 137–176.

Medin, D. L., and Heit, E., in press, Categorization. In D. E. Rumelhart and B. O. Martin (eds) *Handbook of Cognition and Perception: Cognitive Science* (San Diego: Academic Press).
Medin, D. L., and Ortony, A., 1989, Psychological essentialism. In S. Vosniadou and A. Ortony (eds) *Similarity and Analogical Reasoning* (New York: Cambridge University Press), pp. 179–195.
Medin, D. L., and Schaffer, M. M., 1978, Context theory of classification. *Psychological Review*, **85**: 207–238.
Medin, D. L., Lynch, E. B., Coley, J. D., and Atran, S., 1997, Categorization and reasoning among tree experts: do all roads lead to Rome? *Cognitive Psychology*, **32**(1): 49–96.
Miller, G. A., and Johnson-Laird, P. N., 1976, *Language and perception* (Cambridge, MA: Harvard University Press).
Morris, M. W., and Murphy, G. L., 1990, Converging operations on a basic level in event taxonomies. *Memory and Cognition*, **18**(4): 407–418.
Murphy, G. L., 1988, Comprehending complex concepts. *Cognitive Science*, **12**: 529–562.
Murphy, G. L., and Medin, D. L., 1985, The role of theories in conceptual coherence. *Psychological Review*, **92**(3): 289–315.
Naigles, L., 1990, Children use syntax to learn verb meanings. *Journal of Child Language*, **17**: 357–374.
Norman, D. A., and Rumelhart, D. E., 1975, *Explorations in Cognition* (San Francisco: W. H. Freeman).
Nosofsky, R. M., 1986, Attention, similarity and the identification-categorization relationship. *Journal of Experimental Psychology: General*, **115**(1): 39–57.
Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., and Shafir, E., 1990, Category based induction. *Psychological Review*, **97**(2): 185–200.
Pavlicic, T., and Markman, A. B. (1997). The structure of the verb lexicon: Evidence from a structural alignment approach to similarity. In *The Proceedings of the 19th Annual Meeting of the Cognitive Science Society*. Stanford, CA: Lawrence Erlbaum Associates.
Premack, D., 1983, Animal cognition. *Annual Review of Psychology*, **34**: 351–362.
Proffitt, J. B., Coley, J. D., and Medin, D. L., 2000, Expertise and category-based induction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **26**(4): 811–828.
Pustejovsky, J., 1995, *The Generative Lexicon* (Cambridge, MA: The MIT Press).
Putnam, H., 1988, *Representation and Reality* (Cambridge, MA: The MIT Press).
Randall, R. A., 1976, How tall is a taxonomic tree? Some evidence for dwarfism. *American Ethnologist*, **3**: 543–553.
Reed, S. K., 1972, Pattern recognition and categorization. *Cognitive Psychology*, **3**: 382–407.
Regier, T., 1996, *The human semantic potential* (Cambridge, MA: The MIT Press).
Rehder, B., and Ross, B. H., 2001, Abstract coherent categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **27**: 1261–1275.
Rips, L. J., and Conrad, F. G., 1989, Folk psychology of mental activities. *Psychological Review*, **96**(2): 187–207.
Rips, L. J., Shoben, E. J., and Smith, E. E., 1973, Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior*, **12**: 1–20.
Rosch, E., and Mervis, C. B., 1975, Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, **7**: 573–605.
Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., and Boyes-Braem, P., 1976, Basic objects in natural categories. *Cognitive Psychology*, **8**: 382–439.
Ross, B. H., 1996, Category representations and the effects of interacting with instances. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **22**(5): 1249–1265.
Ross, B. H., 1997, The use of categories affects classification. *Journal of Memory and Language*, **37**: 240–267.
Ross, B. H., 1999, Post-classification 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., 2000, The effects of category use on learned categories. *Memory and Cognition*, **28**(1): 51–63.
Schank, R. C., 1972, Conceptual dependency: a theory of natural language understanding. *Cognitive Psychology*, **3**: 552–631.
Schank, R. C., 1982, *Dynamic Memory* (New York: Cambridge University Press).
Schank, R. C., and Abelson, R., 1977, *Scripts, Plans, Goals and Understanding* (Hillsdale, NJ: Lawrence Erlbaum Associates).
Schank, R. C., Collins, G. C., and Hunter, L. E., 1986, Transcending inductive category formation in learning. *Behavioral and Brain Sciences*, **9**: 639–686.
Shepard, R. N., Hovland, C. L., and Jenkins, H. M., 1961, Learning and memorization of classifications. *Psychological Monographs*, **75**.
Sherif, M., 1936, *The psychology of social norms* (New York: Harper and Row Publishers).
Sieck, W. R., Quinn, C. N., and Schooler, J. W., 1999, Justification effects on the judgment of analogy. *Memory and Cognition*, **27**(5): 844–855.

Sloman, S. A., 1993, Feature-based induction. *Cognitive Psychology*, **25**(2): 231–280.
Smith, E. E., and Medin, D. L., 1981, *Categories and Concepts* (Cambridge, MA: Harvard University Press).
Smith, E. E., and Osherson, D. N., 1984, Conceptual combination with prototype concepts. *Cognitive Science*, **8**: 337–361.
Smith, E. E., Osherson, D. N., Rips, L. J., and Keane, M., 1988, Combining prototypes: A selective modification model. *Cognitive Science*, **12**: 485–527.
Smith, E. E., Shoben, E. J., and Rips, L. J., 1974, Structure and process in semantic memory: a featural model for semantic decisions. *Psychological Review*, **81**: 214–241.
Smith, L. B., 1989, From global similarities to kinds of similarities: The construction of dimensions in development. In S. Vosniadou and A. Ortony (eds) *Similarity and Analogical Reasoning* (New York: Cambridge University Press), pp. 146–178.
Tardif, T., 1996, Nouns are not always learned before verbs: evidence from Mandarin. *Developmental Psychology*, **32**: 492–504.
Thompson, R. K. R., and Oden, D. L., 2000, Categorical perception and conceptual judgments by nonhuman primates: the paleological monkey and the analogical ape. *Cognitive Science*, **24**(3): 363–398.
Tomasello, M., and Merriman, W. E. (eds) 1995, *Beyond names for things* (Hillsdale, NJ: Lawrence Erlbaum Associates).
Tversky, A., 1977, Features of similarity. *Psychological Review*, **84**(4): 327–352.
Wisniewski, E. J., 1997, When concepts combine. *Psychonomic Bulletin and Review*, **4**(2): 167–183.
Wisniewski, E. J., and Gentner, D., 1991, On the combinatorial semantics of noun pairs: Minor and major adjustments to meaning. In G. B. Simpson (ed.) *Understanding Word and Sentence* (Amsterdam: Elsevier Science Publishers).
Wisniewski, E. J., and Love, B. C., 1998, Relations versus properties in conceptual combination. *Journal of Memory and Language*, **38**: 177–202.
Wittgenstein, L., 1968, *Philosophical investigations*, trans. G.E.M. Anscombe (New York: MacMillan).
Wolff, P., and Gentner, D., 1992, The time course of metaphor comprehension. Paper presented at the Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society, Bloomington, IN.
Wolff, P., and Gentner, D., 2000, Evidence for role-neutral initial processing of metaphors. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **26**(2): 512–528.
Wu, M. L., and Gentner, D., 1998, Structure in category-based induction. In *The Proceedings of the Twentieth Annual Conference of the Cognitive Science Society*, (Madison, WI: Lawrence Erlbaum Associates), pp. 1154–1158.
Yamauchi, T., and Markman, A. B., 1998, Category learning by inference and classification. *Journal of Memory and Language*, **39**(1): 124–148.
Yamauchi, T., and Markman, A. B., 2000, Inference using categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **26**(3): 776–795.
Zacks, J. M., Tversky, B., and Iyer, G., 2001, Perceiving, remembering, and communicating structure in events. *Journal of Experimental Psychology: General*, **130**(1): 29–58.