



# Journal of Experimental Psychology: Learning, Memory, and Cognition

VOL. 11, No. 4

OCTOBER 1985

## Ideals, Central Tendency, and Frequency of Instantiation as Determinants of Graded Structure in Categories

Lawrence W. Barsalou  
Emory University

Three possible determinants of graded structure (typicality) were observed in *common taxonomic* categories and *goal-derived* categories: (1) an exemplar's similarity to *ideals* associated with goals its category serves; (2) an exemplar's similarity to the *central tendency* of its category (family resemblance); and (3) an exemplar's *frequency of instantiation* (people's subjective estimates of how often it is encountered as a category member). Experiment 1 found that central tendency did not predict graded structure in goal-derived categories, although it did predict graded structure in common taxonomic categories. Ideals and frequency of instantiation predicted graded structure in both category types to sizeable and equal extents. A fourth possible determinant—*familiarity*—did not predict typicality in either common taxonomic or goal-derived categories. Experiment 2 demonstrated that both central tendency and ideals causally determine graded structure, and work showing that frequency causally determines graded structure is discussed. Experiment 2 also demonstrated that the determinants of a particular category's graded structure can change with context. Whereas ideals may determine a category's graded structure in one context, central tendency may determine a different graded structure in another. It is proposed that graded structures do not reflect invariant structures associated with categories but instead reflect people's dynamic ability to construct concepts.

A central theme in categorization research for the last decade has been that categories possess graded structure. Instead of being

equivalent, the members of a category vary in how good an example (or how typical) they are of their category (Rips, Shoben, & Smith, 1973; Rosch, 1973; 1975; Smith, Shoben, & Rips, 1974). In *birds*, for example, American college students agree that *robin* is very typical, *pigeon* is moderately typical, and *ostrich* is atypical. In addition, nonmembers of a category vary in how good a nonmember they are of the category (Barsalou, 1983). For example, *chair* is a better nonmember of *birds* than is *butterfly*. *Graded structure* refers to this continuum of category representativeness, beginning with the most representative members of

The research in this article was supported by Grant MH 13905 from the National Institute of Mental Health to Gordon Bower, a National Science Foundation Graduate Fellowship to the author, and Experimental Training Grant MH 15757-02 from the National Institute of Mental Health to Stanford University. Parts of this article were drawn from a doctoral dissertation submitted to Stanford University (Barsalou, 1981) and an Emory Cognition Project Technical Report (Barsalou, 1984).

I am grateful to Gordon Bower for supporting this project and to Gordon Bower, Herbert Clark, Ellen Markman, and David Rumelhart for direction during its initial development. I am also grateful to Douglas Medin for his comments on every written form of this research and for his guidance in developing it. In addition, this article has benefited from the comments of Jeffrey Farrar, James Hampton, Barbara Malt, Gregory Murphy, Brian Ross, Daniel Sewell, and several anonymous reviewers.

Request for reprints should be sent to Lawrence W. Barsalou, Department of Psychology, Emory University, Atlanta, Georgia 30322.

a category and continuing through its atypical members to those nonmembers least similar to category members. No other variable is as important as graded structure in predicting performance on a wide range of categorization tasks (e.g., category acquisition, exemplar production, category verification). In addition, graded structure occurs in a diverse range of categories, suggesting that it may be a universal property of categories. The large body of work addressing graded structure is reviewed in Mervis and Rosch (1981), Smith and Medin (1981), and Medin and Smith (1984).

This article addresses the issue of what determines graded structure. Why are some exemplars of a category more typical than others? Two experiments examine three possible determinants of graded structure: *central tendency*, *ideals*, and *frequency of instantiation*. A fourth possible determinant—*familiarity*—is also briefly considered.

### Central Tendency

Following the work of Rosch and Mervis (1975), there has been widespread acceptance that an exemplar's typicality depends on its *family resemblance*, where family resemblance is defined as an exemplar's average similarity to other category members and its average dissimilarity to members of contrast categories. The more similar an exemplar is to other category members and the less similar it is to members of contrast categories, the higher its family resemblance, and the more typical it is of its category. *Dog*, for example, is very similar to other members of *mammals* and not very similar to members of contrast categories (e.g., *fish*, *birds*). In contrast, *whale* is not as similar to other *mammals* and is highly similar to the members of a contrast category (i.e., *fish*). Consequently *dog* is more typical of *mammals* than is *whale*.

Another way to view an exemplar's family resemblance is as its similarity to central tendency (Hampton, 1979; Smith et al., 1974), where central tendency refers to any kind of central tendency information about a category's exemplars (e.g., average, median, or modal values on dimensions, highly probable properties, etc.). As just discussed, an exemplar's family resemblance is defined in part as its average similarity to other category members.

However its average similarity to other category members must be at least roughly the same as its similarity to their central tendency (Barsalou, 1983). This is analogous to the average difference between a number and several other numbers being the same as the difference between the first number and the average of the others. In a related manner, an exemplar's average dissimilarity to the members of contrast categories must be at least roughly the same as its dissimilarity to their central tendencies. Consequently an exemplar's family resemblance can be specified either as its average similarity and dissimilarity to category members and nonmembers, or as its similarity and dissimilarity to their central tendencies.

Although people could determine family resemblance in either of these two ways, determining family resemblance through comparisons to central tendencies may be more psychologically plausible—comparing an exemplar to central tendencies requires much fewer comparisons than comparing an exemplar to members and nonmembers. Regardless of how people actually derive family resemblance, however, similarity to central tendencies and similarity to members and nonmembers are functionally equivalent at the level of predicting typicality. Because this article primarily addresses functional relations between typicality and other variables, family resemblance and similarity to central tendencies will be assumed to be equivalent.

### Ideals

Ideals, which provide another possible determinant of graded structure, are characteristics that exemplars should have if they are to best serve a goal associated with their category. For example, an ideal for *foods to eat on a diet* is *zero calories*. The fewer calories an exemplar has, the better it serves the goal associated with its category, namely, *lose weight*. This ideal appears to determine graded structure in that exemplars with decreasing numbers of calories become increasingly good exemplars of the category. Similarly for *things to take from one's home during a fire*, finding exemplars near the ideal of *highest possible value* is relevant to the goal of *minimizing loss*; therefore this property appears to determine the category's graded structure.

Most categories have one ideal. For example, *eat at* may have *lowest cost, highest possible proximity*. On a given occasion, a person is pursuing a goal, and the most memorable exemplar is the one that is most important for achieving that goal. For example, when a person is pursuing a quick meal, the most memorable exemplar is the one that is closest to the goal.

Ideals differ from central tendencies in two ways. First, an exemplar can be the central tendency of a category (although they may not be), but it cannot be the central tendency of a contrast category. For example, *eat on a diet* is the central tendency of *foods to eat on a diet*, but it is not the central tendency of *foods to eat*. Instead of lying at the extreme end of a dimension, an ideal is a few category members away from the central tendency. Instead of lying at the extreme end of a dimension, an ideal is a few category members away from the central tendency.

Central tendencies are determined by the origin of the category. Central tendencies are determined by the origin of the category. Central tendencies are determined by the origin of the category. Central tendencies are determined by the origin of the category. Central tendencies are determined by the origin of the category.

### Frequency of Instantiation

Rosch, Smith, and Mervis (1976) found that the frequency of instantiation of an exemplar is a very sensitive predictor of its typicality. Rosch, Smith, and Mervis (1976) found that the frequency of instantiation of an exemplar is a very sensitive predictor of its typicality. Rosch, Smith, and Mervis (1976) found that the frequency of instantiation of an exemplar is a very sensitive predictor of its typicality.

Most categories probably have more than one ideal. For example, *possible restaurants to eat at* may have the ideals of *lowest possible cost*, *highest possible quality*, and *closest possible proximity*. The most important ideal(s) on a given occasion may depend on the goal a person is pursuing. If the goal is to have a memorable experience, then *high quality* may be most important. But if the goal is to have a quick meal, then *high quality* may succumb to *close proximity* and *low cost*.

Ideals differ from central tendency in at least two ways. First, ideals generally do not appear to be the central tendencies of their categories (although they may occasionally be). *Zero calories*, for example, is certainly not the central tendency with respect to *calories* for *things to eat on a diet*; nor is *closest possible proximity* the central tendency with respect to *distance* for *possible places to eat at*. Ideals tend to be extreme values that are either true of only a few category members or true of none at all. Instead of lying at the center of categories (as does central tendency), they generally lie at the periphery.<sup>1</sup>

Central tendency and ideals also differ in origin. Central tendency depends directly on the exemplars of a category, and more specifically, on the particular exemplars a person has experienced. Although people may form impressions of a category's central tendency through hearsay, they may generally acquire such information through experience with exemplars. In contrast, ideals may often be determined independently of exemplars, being acquired through the process of planning how to achieve goals before exemplars are ever encountered.

#### Frequency of Instantiation and Familiarity

Rosch, Simpson, and Miller (1976) and Mervis, Catlin, and Rosch (1976) argued that frequency does not determine graded structure, although their tests of frequency were not very sensitive. Rosch, Simpson, et al. (1976) pitted family resemblance against frequency and found that only family resemblance predicted typicality. However, their design was not capable of detecting simultaneous effects of family resemblance and frequency. Consequently frequency could have had an effect, but was not detected because it was the weaker of the two factors. Mervis et al. (1976) found

that an exemplar's word frequency in Kučera and Francis's (1967) analysis did not predict typicality. However it is by no means clear that word frequency is a good measure of how often people encounter exemplars in their everyday routines. Other measures of frequency may be better predictors of typicality.

More recent work has contradicted these initial reports, finding that familiar exemplars are perceived as more typical than unfamiliar exemplars (Ashcraft, 1978; Glass & Meany, 1978; Hampton & Gardiner, 1983; Malt & Smith, 1982).<sup>2</sup> Familiarity can be defined as someone's subjective estimate of how often they have experienced an entity across *all* contexts. However an alternative form of frequency that could determine graded structure is frequency of instantiation, which can be defined as someone's subjective estimate of how often they have experienced an entity as a member of a *particular category*. Whereas familiarity is a *category-independent* measure of frequency, frequency of instantiation is a *category-specific* measure of frequency. For example, people generally appear more familiar with *chair* than with *log*, having experienced *chair* more often across all contexts. However people have probably experienced *log* more often as an instantiation of *firewood*. Increases in familiarity and frequency of instantiation could both be associated with increasing typicality. Although both possible determinants receive attention here, the focus will be on frequency of instantiation, because initial inspection of categories suggested it as the more important factor.

In summary, a number of factors could determine graded structure, including central tendency, ideals, frequency of instantiation, and familiarity. Because previous work has not observed ideals and frequency of instantiation, and because previous work has generally not

<sup>1</sup> Ideals are not always the most extreme values possible on a dimension. Exemplars of *clothes to wear in the snow*, for example, vary along the dimension of *how warm they keep people*, with the ideal not being *as warm as possible* (which could be fatal) but being instead *as much warmth as is necessary for survival and comfort*. This ideal, however, is probably not the central tendency of the category.

<sup>2</sup> McCloskey (1980) also reported effects of familiarity on conceptual processing, but his work primarily addressed the role of a category term's familiarity instead of an exemplar term's familiarity, which is of interest here.

performed comprehensive tests of possible determinants, one of the purposes of this project was to observe all four of these possible determinants simultaneously. The focus on these factors is not meant to imply that they are the only possible determinants of graded structure. Instead it is highly likely that other factors also determine typicality. For example, Hampton and Gardiner (1983) review findings that address whether the number of properties associated with an exemplar determines its typicality. In addition, Lakoff (in press) presents a number of other possible determinants, which in at least some cases appear to be composites of the factors examined here.

#### Common Taxonomic and Goal-Derived Categories

Previous work showing that central tendency and familiarity determine graded structure has focused on typicality in one particular kind of category, namely, *common taxonomic* categories (e.g., *birds, furniture, fruit*). However, because Hampton (1981) found that graded structure is not well-predicted by central tendency in some abstract categories, there is reason to believe that the generality of the previous studies is limited. The factors that determine graded structure may vary widely across categories. Consequently a second purpose of this project was to observe typicality in another kind of category, what will be referred to as *goal-derived* categories (e.g., *things not to eat on a diet, things to take from one's home during a fire, birthday presents*).

It should be noted that Barsalou (1983) found graded structure in *ad hoc* categories, which are those goal-derived categories that have been constructed to achieve a novel goal and that therefore are not well-established in memory. Once an ad hoc category is frequently used and becomes well-established in memory, however, it is no longer ad hoc by this definition (see Barsalou, 1983, pp. 224-225). Consequently goal-derived categories include both ad hoc categories and better established categories that were once ad hoc.

Of course it would be ideal to distinguish common taxonomic and goal-derived categories in terms of simple definitions. Unfortunately such definitions have not as yet been forthcoming, although it is as least possible to

provide characteristic properties for each category type. One way common taxonomic and goal-derived categories generally appear to differ has to do with the "correlational structure of the environment." Correlational structure refers to the fact that properties in the physical environment are not independent; that is, a given property generally co-occurs with certain other properties but not with others. *Feathers*, for example, typically co-occurs with *wings* and *beak*, but not with *tires* and *engine*. As discussed by Rosch and Mervis (1975) and Rosch, Mervis, Gray, Johnson, and Boyes-Braem (1976), common taxonomic categories appear to circumscribe sets of things in the environment that share these clusters of co-occurring properties. Consequently these categories reflect the correlational structure of the environment. Many exemplars of *birds*, for example, share co-occurring properties that rarely occur outside the category, thereby making exemplars of this category very similar to each other and very dissimilar to nonmembers.

In contrast, goal-derived categories generally appear to violate the correlational structure of the environment. Many goal-derived categories include some members from each of several common taxonomic categories, but never all the members from a given one. *Things to take on a camping trip*, for example, includes members of *food, clothing, tools*, and so on, but it does not include all members. Because the members of these goal-derived categories are often quite dissimilar to each other and very similar to many nonmembers, they do not maximize the correlational structure of the environment. Other goal-derived categories contain subsets of one particular common taxonomic category. For example, someone with a back problem might be interested in *chairs that provide good back support*. In these cases, goal-derived categories do not maximize correlational structure because many noncategory members are highly similar to category members (e.g., *chairs that provide good back support* are very similar to *chairs that do not*). In general, because goal-derived categories do not maximize the correlational structure of the environment, they are not very salient and do not stand out as natural groups. Instead they appear to only become salient when relevant to currently pursued goals.

Another and goal-differ has taxonomic sification, often used entities in appear to More spec level categ taxonomic Kosslyn, Mervis, et 1984). In use goal-schema v achieve th ample, a in schema as where t what to ta egories of of transpo case facili tions for 1 fully achiev schema v priate in t egories pr which ins

Although rived cate extent to structure distinction For categories niture are may often goal-deriv for classifi are not d properties derived ca

Dete  
ir

These g taxonomic that differ graded str

Another way in which common taxonomic and goal-derived categories generally seem to differ has to do with category use. Common taxonomic categories are often used for classification, whereas goal-derived categories are often used for instantiation. When classifying entities in the environment, people primarily appear to use common taxonomic categories. More specifically, people generally prefer basic level categories, which are a subset of common taxonomic categories (Jolicoeur, Gluck, & Kosslyn, 1984; Murphy & Smith, 1982; Rosch, Mervis, et al., 1976; B. Tversky & Hemenway, 1984). In contrast, people primarily appear to use goal-derived categories for instantiating schema variables while achieving goals. To achieve the goal of *taking a vacation*, for example, a planner has to instantiate variables in schematic knowledge about *vacations*, such as *where to go*, *who to go with*, *how to get there*, *what to take*, and so on. The goal-derived categories of *places to go*, *people to go with*, *types of transportation*, and *things to pack in a suitcase* facilitate locating and selecting instantiations for these variables. In general, successfully achieving a goal requires that people bind schema variables with instantiations appropriate in the current setting. Goal-derived categories provide pools of instantiations from which instantiations can be chosen.

Although common taxonomic and goal-derived categories generally seem to differ in the extent to which they reflect correlational structure and in the way they are used, these distinctions are by no means clear-cut or defining. For example, some common taxonomic categories such as *vehicles*, *clothing*, and *furniture* are highly related to people's goals and may often be used for instantiation. Conversely goal-derived categories may at times be used for classification. Although these distinctions are not defining, they provide characteristic properties of common taxonomic and goal-derived categories.

#### Determinants of Graded Structure in Common Taxonomic and Goal-Derived Categories

These general differences between common taxonomic and goal-derived categories suggest that different factors may determine their graded structures. To begin with, central ten-

dency may be highly salient in people's representations of common taxonomic categories and thereby become the standard by which typicality is judged. Central tendency may be salient in common taxonomic categories for the following two reasons. First, because these categories generally reflect correlational structure, people may use them as a means of representing the structure of the environment. If so, then acquiring central tendency information for these categories provides *representative information* about the kinds of entities the environment contains. Central tendency information is clearly more representative than ideal information, because the former has a much higher likelihood of occurring for a category's exemplars than the latter.

A second reason central tendency information may be salient in common taxonomic categories has to do with their use. Because these categories are often used for classification, their representations may be designed to maximize classification performance. It is well-known in the category verification literature that classifying an entity proceeds faster to the extent the entity is similar to the category standard (Smith, 1978). Basing classification standards on central tendency information minimizes the average distance of category members to category standards (e.g., it is a statistical fact that the average absolute distance from all points in a set to one particular point is minimized when that point is the median; Hayes, 1973, p. 223). Therefore the average difficulty of performing classifications is minimized when the category standard contains central tendency information (as opposed to ideals). Because a primary use of common taxonomic categories is to serve classification, it would not be surprising if central tendency information were central to their representations.

In contrast, ideals may become highly salient in people's representations of goal-derived categories and thereby become the standards by which typicality is judged. Because goal-derived categories generally serve goals, their representations may contain ideals in order to maximize goal achievement. As people consider possible instantiations of a category, they can compare them to the category's ideals and thereby pick the exemplar or exemplars that will result in maximal goal satisfaction. For

