Psychological Bulletin
1992, Vol. 111, No. 2, 291–303

0033-2909/92/$3.00

Explaining Basic Categories: Feature Predictability and Information

James E. Corter
Teachers College
Columbia University

Mark A. Gluck
Center for Molecular and Behavioral Neuroscience
Rutgers University

The category utility hypothesis holds that categories are useful because they can be used to predict the features of instances and that the categories that tend to survive and become preferred in a culture (*basic-level* categories) are those that best improve the category users' ability to perform this function. Starting from this hypothesis, a quantitative measure of the utility of a category is derived. Application to the special case of substitutive attributes is described. The measure is used successfully to predict the basic level in applications to data from hierarchies of natural categories and from hierarchies of artificial categories used in category-learning experiments. The relationship of the measure to previously proposed indicators of the basic level is discussed, as is its relation to certain concepts from information theory.

Categorization is one of the most basic cognitive functions. Why is the ability to categorize events or objects important to an organism? An obvious answer to this question is that categories are important because they often have functional significance for the organism. Another familiar answer is that grouping objects into categories allows for efficient storage of information about these groups of objects. One purpose of this article is to explore connections between these two answers regarding the utility of categories.

The idea that categories serve certain functions for the organism raises the possibility that some categories fulfill these functions better than others. The clearest evidence that certain natural categories are "better" than others stems from the work on "basic-level" categories (Mervis & Rosch, 1981; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). A basic-level category is one that is preferred by people over its superordinate and subordinate categories. For example, when shown a picture of a particular object, most people will identify it as a chair rather than as furniture or a kitchen chair. From this and other evidence, *chair* is considered to be a basic-level category for most people.

A variety of empirical phenomena demonstrates the superiority of basic-level categories (Mervis & Rosch, 1981). As suggested above, when people are shown an object, they tend to name it at the basic level (Rosch et al., 1976). In recognition tasks, people recognize basic-level objects faster than either subordinates or superordinates (Jolicoeur, Gluck, & Kosslyn, 1984; Rosch et al., 1976). Basic-level names generally have arisen earlier in the development of languages (Berlin, Breedlove, & Raven, 1973), and basic categories are used earlier in the naming and other behavior of young children (Anglin, 1977; Brown, 1958; Horton & Markman, 1980; Mervis & Crisafi, 1982). Furthermore, these basic-level names tend to be shorter and more frequent words in English than names of superordinate or subordinate categories.

All these types of evidence tend to pick out the same level of category (e.g. *chair*) as optimal in hierarchies of natural categories. Thus it is difficult to establish whether the causes of basic-level effects are linguistic (e.g., word length and frequency), developmental (basic-level terms are learned earlier by children), or "structural" (because of the particular features associated with a category). However, a number of experimental studies have used hierarchies of "artificial" categories to demonstrate that structural factors are sufficient to induce basic-level-type effects (Gluck, Corter, & Bower, 1992; Hoffmann & Ziessler, 1983; Mervis & Crisafi, 1982; Murphy & Smith, 1982). In a typical result from these studies, categories were defined at three levels of generality, with the feature–category associations defined in such a way that the middle level was expected to be "basic." These predictions were confirmed by showing that the middle-level categories were learned most quickly or could be named most quickly after they were learned.

These experimental results demonstrate that feature structure (i.e., feature–category association) is sufficient to induce the sort of category-learning phenomena that are associated with basic levels. The results do not mean that developmental and linguistic factors, such as age of acquisition and word frequency, do not influence how easily people can use category names. But we believe that these developmental and linguistic factors arise as effects of the optimality of certain categories, rather than as causes. This view follows from a theory of why basic-level categories arise, which we describe in this article. The theory is structural, in that it explains the utility of categories in terms of the features associated with the categories or category names. Furthermore, the theory suggests a potential quantitative measure of the "goodness" of a category. The theory is normative, in that it assumes that the best categories are those that serve the twin purposes of maximizing feature predictability and optimizing information transfer in communication between people (Gluck & Corter, 1985).

Many people have provided us with suggestions, criticisms, and comments on this work over the last 6 years. However, we must especially thank Pat Langley and Doug Fisher for their early and continuing encouragement. Doug Fisher's own research has contributed to our understanding of many of the issues discussed here. We also thank three anonymous reviewers for their insightful and detailed feedback.

Correspondence concerning this article should be addressed to James E. Corter, Box 41, Teachers College, Columbia University, New York, New York 10027.

Previously Proposed Structural Measures of Category Goodness

Rosch and co-workers (Rosch et al., 1976) introduced the idea that basic levels might arise as a consequence of natural patterns of co-occurrence of features in the world. They termed their approach a structuralist one, in that they did not propose specific process models for the learning of categories, but rather sought to describe the relations between stimuli and their features that underlie such learning.

Rosch et al. (1976) suggested several structural measures as candidates for picking out basic-level categories. One of these suggestions was that basic-level categories may be those for which the average *cue validity* is maximal. Cue validity is defined as the conditional probability of an object being in a category c, given that it possesses some feature f, $P(c|f)$. For example, the validity of the cue *wings* for the category *bird* is the probability that an object is a bird, given that it has wings. However, as Murphy (1982) pointed out, cue validity will always be maximal for the most general or inclusive level of a hierarchy. The probability that an object with wings is an animal is logically at least as great as the probability that it is a member of a subset of animals (e.g., a bird). However, because the basic level is rarely if ever the most inclusive level, cue validity cannot be a valid indicator of the basic level.

An alternate hypothesis is that basic-level categories are those for which *category validity* is maximal, where category validity is defined as the conditional probability of an object having some feature f given that it is in some category c, $P(f|c)$. Categories with high category validity would be those categories for which inferences about features can be made with confidence. But, as Medin (1983) noted, category validity suffers from the opposite problem as does cue validity: It generally selects the most specific categories because they tend to have the least variability in features. For example, the probability that something can fly given that it is a robin is higher than the probability that something can fly given that it is a bird. Thus, category validity is also an unsatisfactory predictor of the basic level.

Cue validity and category validity both reflect desirable patterns of association between categories and their features: Cue validity reflects the extent to which a feature can be used to uniquely identify a category, whereas category validity reflects the extent to which category membership can be used to make inferences about the presence of a feature. Thus basic-level categories may be ideal categories because they represent a tradeoff between these two factors. Furthermore, some function combining these two measures might be maximal for the intermediate-level categories that are often found to be basic (Medin, 1983). Jones (1983) suggested that basic-level categories might be those for which the product of cue validity and category validity, $P(f|c)P(c|f)$, is maximal. He termed this the feature *collocation* measure.

Jones (1983) suggested applying the collocation measure in the following way to predict which level of a hierarchy should be basic. Consider a single feature f, denoting the presence of an attribute or quality (e.g., *can fly*). For each category in a nested set (i.e., a hierarchy) such as *robin-bird-animal*, the product $P(f|c)P(c|f)$ is computed. Feature f is assigned to the category for which this measure is maximal. For example, if the product $P(\textit{can fly}|c)P(c|\textit{can fly})$ is greater for *bird* than for *robin* and *animal*, this feature will be assigned to the category *bird*. Across all features, a summary score for each category is defined as the number of features assigned to that category. In all subsequent discussion, we refer to this summary score as the *feature-possession* score. In the *robin–bird–animal* hierarchy, *bird* would be identified as the basic level if its feature-possession score exceeded that of *robin* and *animal*.

Jones's (1983) measure seems empirically promising in that (unlike cue validity and category validity) it can be maximal for intermediate-level categories. Furthermore, a normative justification can be constructed for the measure on the basis of the usefulness of attempting to simultaneously maximize cue validity and category validity. However, this normative argument does not necessarily lead to the multiplicative combination rule for combining the two criteria. That is, why should it be the product of cue and category validity that is maximized rather than some other function? Furthermore, the intermediate step of calculating feature-possession scores seems questionable. Why should a feature be assigned to only a single category for which the collocation is maximal? This implies that a feature contributes evidence for the "basicness" of only a single category. These details of the measure and its suggested method of application seem in need of better normative or empirical justification.

Other measures have also been proposed that might distinguish basic-level categories. Another such proposal is that basic-level categories might be those that maximize within-category similarity and minimize between-category similarity (Medin, 1983; Mervis & Rosch, 1981). Because within-category similarity will in general be maximal for the most specific categories, and between-categories similarity minimal for the most general categories (cf. Medin, 1983; Tversky, 1977), the two measures are at odds with each other (as for cue and category validity), and actually some joint function combining the two measures is presumably maximal for basic-level categories. Although this suggestion seems promising, it has not been explored in detail. Empirical evidence that seems to support such a view was provided by Mervis and Crisafi (1982) and by Murphy and Brownell (1985), who showed that category differentiation (presumably related to some joint function of within- and between-categories similarity) predicted speed of identification of line drawings in terms of the categories. Finally, one of the several structural characteristics proposed by Rosch and Mervis (Mervis & Rosch, 1981; Rosch et al., 1976) to be associated with basic categories was that basic-level categories are those that carry the most information about attributes. Although this suggestion was not fleshed out by Rosch and her co-workers, it foreshadows the ideas developed in the present article.

To summarize, either Jones's (1983) collocation measure or some measure combining within- and between-categories similarity seems potentially capable of identifying basic-level categories. However, to judge the relative merits of these and related suggestions, either empirical data or normative considerations ought to be brought to bear. In this article, we describe a normative justification for the existence of basic levels, on the basis of an account of the usefulness of categories to the categorizer.

We then show that the theory leads to a quantitative measure of category utility that, besides having an explicit normative justification, successfully predicts the basic level in several applications with data from natural and artificial concept hierarchies. We also discuss the relationship of the measure to the notion of informativeness and to Jones's collocation measure.

Utility of Categories: Feature Predictability

We propose that a category is useful to the extent that it can be expected to improve our ability to (a) accurately predict the values of features for members of that category and (b) efficiently communicate information to others about the features of instances of that category. We show that under plausible assumptions about how a person makes predictions or guesses about the values of features, a category that is optimal for one of these purposes also tends to be optimal for the other. We then discuss the implications of this view for explaining basic-level phenomena and identifying basic-level categories.

Our basic assumption is that there is functional value for a person to have accurate information about the features of things. For example, an organism searching for food needs to know whether a particular plant part is poisonous, nutritive, sweet, tough, and so on. Some features of instances may be useful only indirectly—for example, to generate tests to confirm tentative identifications. However, because a person will experience a variety of need states and goals across time, generally, there is value for the person to have accurate information about all or virtually all the features of instances.

We term this the *category utility hypothesis:* A category is useful to the extent that it can be expected to improve the ability of a person to accurately predict the features of instances of that category. We formalize this idea as follows. Let category c be defined as a finite set of instances, $c = \{o_1, o_2, \ldots, o_n\}$. We assume that an instance o_i can be described by a finite set of discrete features, $F = \{f_1, f_2, \ldots, f_m\}$. Consider a hypothetical situation in which the only information a person R has about instance o_i is the fact that it is a member of category c. What value does the knowledge that o_i is a member of c impart? We propose that one measure of this value is the increase in R's ability to correctly guess the features of o_i that results when R is given the information that o_i is a member of c.

With no knowledge about whether o_i is a member of c, R's best guess about the probability of o_i possessing feature f_k would be based on $P(f_k)$, the base-rate probability of f_k. How well can R be expected to do in guessing whether instance o_i possesses feature f_k in this situation? To derive R's expected score, we need to make an assumption regarding the type of guessing strategy R adopts. Evidence from the literature on probability learning suggests that there are strong general tendencies for people (a) to correctly induce the relative frequencies of events, given a large enough sample of trials (instances) and (b) to adopt a probability-matching strategy for guessing outcomes in such tasks (Estes, 1972). Assuming that R has correctly induced $P(f_k)$ from past experience and that R adopts a probability-matching strategy, R should guess that instance o_i possesses feature f_k with probability $P(f_k)$ when given no information about the category membership of o_i. Because such a guess issued randomly will be correct with probability $P(f_k)$ (the base-rate relative frequency of f_k), the overall probability of R's correctly guessing that o_i possesses feature f_k is given by $P(f_k)P(f_k) = P(f_k)^2$. This follows because in the absence of any information about o_i, R's guess is independent of the actual occurrence or nonoccurrence of f, thus the probabilities of o_i actually possessing feature f and of R guessing that it does can be multiplied to obtain the probability of the joint event.

Now suppose that R has some criterial test that informs him whenever an instance is a member of c. This test could be a message from another person, or it could be some feature-based identification procedure. Again assume that R adopts a probability-matching strategy and that R has correctly induced $P(f_k|c)$ from past experience. Then whenever R is informed that o_i is an instance of c, R should guess that o_i has feature f_k with probability $P(f_k|c)$. Because such a guess will be correct in this conditional distribution with probability $P(f_k|c)$, R's expected probability of "success" is $P(f_k|c)^2$ (following the same logic as for the no-information case). Thus the identification of o_i as a member of c will change R's probability of correctly guessing that o_i possesses feature f_k (i.e., R's expected score on a single trial) from $P(f_k)^2$ to $P(f_k|c)^2$. Finally, note that the message that o_i is a member of c will occur with probability $P(c)$ (assuming perfectly reliable identification of c), but with probability $1 - P(c)$, no such message occurs. We can thus show that the expected increase in R's ability to predict whether o_i possesses feature f_k, given reliable information as to when an instance is a member of c, is given by

$$P(c)[P(f_k\,|\,c)^2 - P(f_k)^2] + (1 - P(c))[P(f_k)^2 - P(f_k)^2] = P(c)[P(f_k\,|\,c)^2 - P(f_k)^2].$$

Summing this measure across all m features possessed by instances of c, we have

$$CU(c, F) = P(c) \sum_{k=1}^{m} [P(f_k\,|\,c)^2 - P(f_k)^2]. \quad (1)$$

We term this measure the *category utility* of c. Note that although it is assumed that R has correctly induced both $P(f_k)$ and $P(f_k|c)$, it is not necessary to assume that R can give an accurate assessment of $P(c)$. If R's test for the presence of c is perfectly reliable, or (put another way) if whenever an instance is a member of c, R is always so informed, then this term of the model can be interpreted as the real-world relative frequency of category c. In the present article, we assume that this is the case, thus in the following discussion we refer to $P(c)$ as the relative frequency of c. However, if R were not always informed that o_i is an instance of c, then $P(c)$ should be interpreted as the probability with which R is informed that c is present, rather than as the actual base-rate probability of c.

As an example, consider the population of newspapers and the subcategory of newspapers referred to as *tabloids* (category t). All tabloids are printed in black ink (feature f_1), and 90% of them feature pictures on the cover (f_2). Thus $P(f_1|t) = 1.0$ and $P(f_2|t) = .9$. For the population of newspapers as a whole, $P(f_1) = 1.0$ and $P(f_2) = .6$. If 33% of all newspapers are tabloids, $P(t) = .33$), then the category utility of the category t (using only these two features) is equal to $P(t)\,[(P(f_1|t)^2 - P(f_1)^2) + (P(f_2|t)^2 - P(f_2)^2)] = (.33)\,[(1.0 - 1.0) + (.81 - .36)] = (.33)\,[.0 + .45] = .15$. The interpretation is that being able to identify a newspaper as a tabloid leads to an expected increase in the ability to predict one of its features, namely the feature of having a picture on the front cover. Note that the feature of being printed in black ink, although it has a high category validity, $P(f_1|t) = 1$, does not contribute to the category utility score, because this feature can be predicted just as well for any newspaper.

Application to Substitutive Features

So far we have been assuming that instances can be described by the presence or absence of certain features or properties (e.g., a face either does or does not have a mustache). Tversky (1977) termed such features *additive,* and pointed out that, in contrast, many objects are better described by sets of mutually exclusive and exhaustive properties, which he termed *substitutive* features. For example, the attribute eye color might have several possible values, say brown, blue, or green. Thus the features brown eyes, blue eyes, green eyes are said to be substitutive. Garner (1978) made a similar distinction between *features* (denoting presence or absence of a property) and *dimensions* (which have mutually exclusive levels). We prefer to use the term *attribute* instead of *dimension,* which has connotations of continuous variation. Hence we will refer to the attribute of eye color, for example, with attribute values of brown, green, blue.

The category utility measure applies quite naturally to this situation. Assume that instances are described by a value on each of m attributes, where the jth attribute F_j has n_j mutually exclusive values, $f_{j1}, f_{j2}, \ldots, f_{jn}$. Then the probabilities of the n_j values of attribute F_j must sum to 1, because they are mutually exclusive and exhaustive. However, the category measure still applies, providing a measure of R's expected increase in ability to predict the value of an instance on attribute F_j, given information on whether the instance is a member of c:

$$CU(c, F_j) = P(c) \sum_{k=1}^{n_j} [P(f_{jk}|c)^2 - P(f_{jk})^2]. \quad (2)$$

Note that in this situation, R's feature-prediction score will be lowest when all of the values of attribute F are equally likely and highest when one value of F has probability 1. In fact, prediction will be perfect in the latter case. This is shown in Figure 1a, which presents a plot of R's expected score in predicting the two values of an attribute $F_j = \{f_1, f_2\}$ as a function of $P(f_1)$. Category c will be valuable to the extent that many attributes show large variability in the population (i.e., the no-information condition) and little or no variability within the category. The measure given in Equation 2 can be summed across the set of m attributes to obtain an overall measure of the expected increase in ability to predict the values of all the attributes of an instance.

Relation of Feature Predictability to Informativeness

As mentioned previously, one of the proposed characteristics of basic-level categories is that they are the most informative in terms of features (e.g. Mervis & Rosch, 1981). In this section, we show that there is a close relationship between the predictability of features as measured by category utility and the informativeness of a category in terms of its features, as measured by standard concepts of information theory (Shannon & Weaver, 1949).

Assume that the information available to R (for receiver) concerning the membership or nonmembership of o_i in category c is a message from another person T (for transmitter). From the perspective of information theory, it is natural also to consider as a message the fact that the feature f_k is present. The information that would be conveyed by transmission of this message is defined to be $-\log P(f_k)$. Furthermore, this message has probability $P(f_k)$ of occurring, therefore the expected information associated with this message is given by $-P(f_k)\log P(f_k)$. Similarly, when o_i is known to be a member of c, the information associated with message f_k is given by $-\log P(f_k|c)$, and its expected information is $-P(f_k|c)\log P(f_k|c)$.

When instances are described in terms of substitutive attributes, the n_j mutually exclusive values of attribute F_j constitute a message set, $F_j = \{f_{j1}, f_{j2}, \ldots, f_{jn}\}$. The expected information carried by a such a set of messages is termed the *uncertainty* of the set,

$$U(F_j) = -\sum_{k=1}^{n_j} P(f_{jk}) \log P(f_{jk}).$$

Uncertainty is maximal when all messages in the set are equiprobable and minimal when one message has probability 1 and the other messages have probability 0. Thus the uncertainty of a message set (i.e., a set of attribute values) is inversely related to the expected score that would be obtained in trying to predict the attribute values using a feature-matching strategy. This can be seen clearly in Figure 1, which contains a graph of the negative uncertainty of a message set $F_j = \{f_1, f_2\}$ as a function of $P(f_1)$ (Figure 1b) along with the expected score function (Figure 1a). The two curves are virtually identical, with each function attaining a minimum value when both features have equal probabilities, $P(f_1) = P(f_2) = .5$, and maximum values when one feature has probability 1 and the other messages have probability 0.

With these information-theoretic concepts, we can define the informativeness of category c as the expected reduction in uncertainty about the values of attribute F_j, given access to reliable messages concerning the membership of instances in category c. This expected reduction is equal to

$$P(c)[U(F_j) - U(F_j|c)]$$
$$= P(c)[-\sum_{k=1}^{n_j} P(f_{jk}) \log P(f_{jk}) + \sum_{k=1}^{n_j} P(f_{jk}|c) \log P(f_{jk}|c)]$$
$$= P(c) \sum_{k=1}^{n_j} [P(f_{jk}|c) \log P(f_{jk}|c) - P(f_{jk}) \log P(f_{jk})]. \quad (3)$$

This expression is obviously closely related to the category utility measure defined for a substitutive attribute (2). The measures differ only in the substitution of a log probability in Equation 3 for the actual probability in each of the terms of Equation 2. Because these log probabilities are monotonic with the actual probabilities, the resulting measures give virtually identical results in assessing the relative goodness of categories. This demonstrates that the category utility measure provides one possible quantitative instantiation of the informal notion that

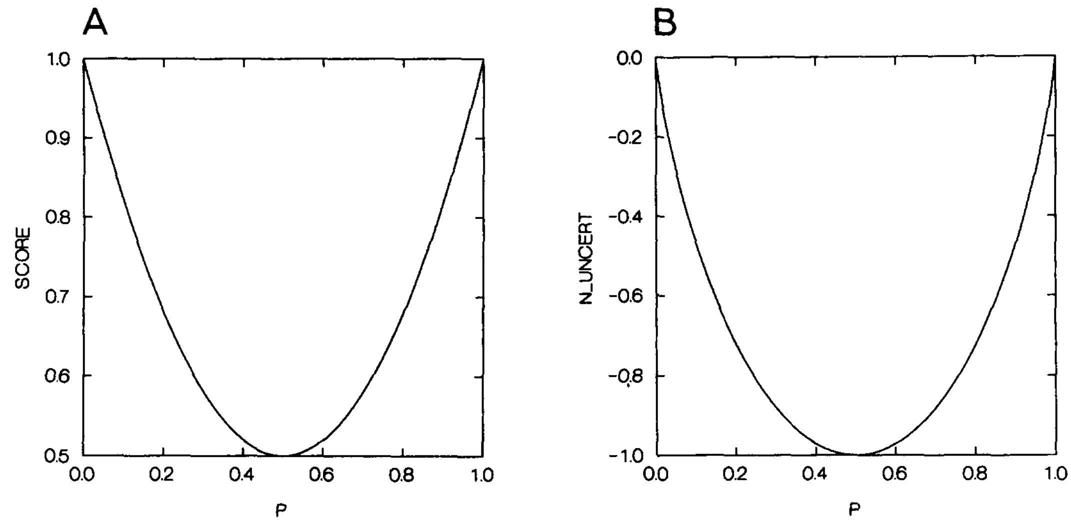

Figure 1. (A) Plot of expected score, $P(f_1)^2 + P(f_2)^2$, as a function of $P(f_1)$. (B) Plot of negative uncertainty, $P(f_1) \log P(f_1) + P(f_2) \log P(f_2)$, as a function of $P(f_1)$.

basic-level categories are the most informative in terms of features (Mervis & Rosch, 1981; Rosch et al., 1976).

Relation to Previously Proposed Structural Measures

The category utility model, expressed most simply in Equation 1, represents the goodness of a category in terms of three basic factors: the probability or relative frequency of the category c, the category validity of c for feature f, $P(f|c)$, and the overall base rate of feature f, $P(f)$. The category utility of c is defined as the difference between squared category validity and the squared base rate, weighted by the relative frequency of category c. Thus the model can be seen as a generalization of a simple category validity model. For example, the fact that most members of Alpha Omega fraternity like rock music does not contribute to make *members of Alpha Omega* a useful subcategory of the population of college students, because most college students like rock music. And if there are only one or two members of Alpha Omega, it again will not be a particularly useful category.

The category utility measure can also be related to Jones's (1983) collocation measure, $P(f|c)P(c|f)$. The major difference between the measures is in the role ascribed by the category utility model to $P(f)$, the base-rate probability or relative frequency of feature f. If we neglect the $P(f)^2$ term of Equation 1, we reduce the category utility measure to $P(c)P(f|c)^2$. By substituting $P(f \cap c)/P(c)$ for one of the $P(f|c)$ terms in this expression, we obtain $P(c)P(f|c)\ [P(f \cap c)/P(c)]$, which reduces to $P(f|c)P(f \cap c)$. Substituting $P(c|f)P(f)$ for $P(f \cap c)$, we have $P(f|c)P(c|f)P(f)$. Thus the first term of the category utility measure can be seen as a weighted function of the collocation measure, in which the collocation measure for each feature (or attribute value) is weighted by the relative frequency of that feature.[1] The second term of the category utility measure also involves a presumed influence of the base-rate probability of feature f.

In applications to identifying the basic level in an actual category hierarchy, the two measures may give quite different predictions as to category optimality. These differences may result not only from the role of the base rates of features but also from the intermediate steps involved in applying the measure in the manner suggested by Jones (1983). In these steps, a feature is assigned only to that category in a hierarchical set such as *robin–bird–animal* for which the collocation measure is maximal. For each category, a feature-possession score is computed as the number of features assigned to that category (rather than to superordinate or subordinate categories). It is these feature-possession scores, rather than averaged values of the collocation measure itself, that are used to identify the basic level. In the next section, we report applications of the category utility measure to identify the basic level in several category hierarchies used in previous studies of basic levels and present comparative data on the performance of the collocation measure. We apply the collocation measure both with and without the intermediate step of calculating the feature-possession scores, to assess what effect this has on the performance of the measure.

To summarize, the category utility hypothesis and the category utility measure derived from it suggest several characterizations of the type of category that will be most useful in a given context. The value or utility of a category c is assumed to depend on three factors: the relative frequency of the category, the relative frequency of each feature within the category, and the base rate of each feature in the context population.

Applications to Predicting Basic Levels

In this section, we examine the adequacy of the category utility measure to predict the basic level in hierarchies of natu-

[1] This relation between the two measures was pointed out by Doug Fisher (1987).

ral categories and in the artificial stimulus hierarchies used in several experimental studies of basic-level learning. We calculated the mean value of the category utility measure for categories at each level of the hierarchy to see if basic-level categories have a higher category utility than either subordinate or superordinate categories.

Natural Taxonomic Categories

Rosch et al. (1976) had subjects list the features of nine hierarchical sets of natural categories, including both biological categories (e.g., *trees*) and nonbiological ones (e.g., *vehicles*). These feature lists were then amended by judges to check for consistency and eliminate obvious errors. Although Rosch et al. did not report their actual data, Tversky and Hemenway (1984) presented the actual judge-amended feature lists for three of the hierarchies: fruit, furniture, and musical instruments.

Each of these hierarchies consists of 10 categories: 1 superordinate category (e.g. *fruit*), 3 basic-level categories (*apple, peach, grapes*), and 6 subordinate categories, 2 for each basic-level category (*delicious apple, macintosh apple; freestone peach, cling peach; concord grapes, green seedless grapes*). For the musical instruments hierarchy, the basic-level categories used were *guitar, piano,* and *drum;* for furniture, basic categories were *table, lamp,* and *chair.* The data consist of a list of features for each category. For example, *fruit* is associated with the features *seeds, sweet, you eat it,* and *apple* is associated with *stem, core, skin, juicy, round, grows on trees,* and so forth.

For analysis, the feature list for each hierarchy was represented as an $6 \times m$ matrix with binary entries, where m is the total number of features and 6 is the number of subordinate-level categories included in each hierarchy. Each entry x_{ik} was coded as a 1 if category i possessed the kth feature and as a 0 otherwise. Features were assumed to be "inherited"; that is, any features of a category were assumed also to be features of the categories below it in the hierarchy. For example, *you eat it* was listed by subjects as a feature of the superordinate category *fruit;* therefore, it was assumed also to be a feature of the categories *apple, grapes, Macintosh apple,* and so on, and each of these categories received a score of 1 for this feature. The feature structure for the fruit hierarchy is given in Table 1. For each category, category utility was computed using the relative frequency within the superordinate category as the base rate of that feature. The values of the measure for each category were averaged across categories within each level. For comparison, Jones's (1983) collocation measure was also computed for each category. This was done both with and without the additional step of calculating feature-possession scores (as described in the previous section). Results of these analyses are reported in Table 2. The row labeled Collocation contains the mean value of the statistic $P(f_j|c)P(c|f_j)$ summed across features; the row labeled Feature possession contains the mean values of the feature-possession scores (i.e., the mean number of features for which that level had the highest value of the collocation measure).

Values reported are the mean values of the measures for categories at each level (superordinate, basic, subordinate). For the category utility measure, the value for the superordinate category is 0.000 in every case because the superordinate category serves as the context population for this analysis. Thus the increase in predictability that is due to identifying an object as an instance of the superordinate category is equal to 0. More interesting are the values of the measures for the intermediate-level categories, which are the basic-level categories. In each hierarchy, the mean values of category utility are highest for this level. Furthermore, when comparisons of each basic-level category are made with its superordinate and with its subordinate categories, in all nine cases (three in each hierarchy), the ordering of the levels corresponds to that of the mean values in Table 2. This provides support for the category utility hypothesis. Collocation, on the other hand, fails to identify the basic level. In all three hierarchies, the mean values of the collocation measure incorrectly pick out the superordinate level as best. When the data are examined separately for each of the nine basic categories, the superordinate categories are incorrectly identified as best in seven out of the nine cases. The exceptions are the basic categories *guitar* and *chair,* which are picked out as superior both to their superordinates and to their subordinates. The feature-possession scores defined using the collocation values incorrectly pick out the superordinate level as best in every case.

Given the theoretical plausibility of the collocation measure, this empirical inadequacy is surprising. Comparing the calculation of this measure for the superordinate category and the first basic category yields some insight into the nature of the superordinate advantage. The collocation measure is calculated as the sum across all features of $P(f|c)P(c|f)$. For the superordinate category, the term $P(c|f)$ is equal to 1 for all 25 features in Table 1, thus the measure reduces to simply the sum of the $P(f|c)$ terms: $1 + 1 + 1 + .33 + .33 + .33 + .83 + .33 + .67 + .17 + .17 + .17 + .17 + .33 + .33 + .33 + .33 + .33 + .17 + .33 + .33 + .33 + .17 + .17 + .17 = 9.83$. For the first basic category, (consisting of the first two rows of Table 1), the $P(c|f)$ terms vary; many

Table 1
Feature Structure for the Fruit Hierarchy

Category	Features																								
Delicious apple	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0
Macintosh apple	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Freestone peach	1	1	1	0	0	0	0	0	1	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0	0
Cling peach	1	1	1	0	0	0	1	0	1	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0
Concord grapes	1	1	1	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	0	0
Green seedless grapes	1	1	1	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	0	1	1

Table 2
Mean Values of Category Utility, Collocation, and Feature-Possession Score Measures for Hierarchies of Natural Categories

Measure	Subordinate	Basic	Superordinate
	Level		
Musical instrument			
Category utility	0.71	1.37	0
Collocation	3.10	6.67	8.67
Feature possession	0.33	5.00	15.67
Fruit			
Category utility	0.69	1.13	0
Collocation	4.17	6.97	9.83
Feature possession	1.33	5.33	18.33
Furniture			
Category utility	0.84	1.54	0
Collocation	4.00	7.28	8.17
Feature possession	0.75	6.67	16.58

of them are considerably less than 1. For example, $P(c|f) = 0$ for the last 12 features in the table, which are features associated with other basic or subordinate categories. Thus it is not surprising that the mean value of the collocation measure should be greatest for the superordinate category.

Using the superordinate category itself as the context population may provide an unfair advantage for the category utility measure in these comparative evaluations, because no increase in feature predictability can result from identifying something as an instance of the superordinate category. However, there is no accompanying data describing related superordinate categories or relevant base rates for features. One possibility for simulating a plausible context for the given superordinate categories is to define a new hypothetical contrasting superordinate category in each case. This was done by simply replicating each given superordinate category to create a contrasting superordinate category. That is, from a given superordinate category described by 6 subordinate-level categories and k features, a new superhierarchy was created with 12 categories and $2k$ features. For the new contrast superordinate category, the original k features all have values of 0. Conversely, for the original superordinate category, the features from $(k+1)$ to $2k$ have all 0s. The extended hierarchies were analyzed as above, and the mean values of the three measures were calculated for each level.

The values for the collocation measure and the feature-possession scores for the original six categories were not affected by the expansion of the hierarchy. Thus, both measures still incorrectly identified the superordinate category as best in each of the three hierarchies. However, the values of category utility were affected by this change of context. For the furniture hierarchy, the basic results did not change; that is, all three basic-level categories had higher category utility than both their superordinate category and their subordinate categories. For the fruits hierarchy, however, the expansion of the context changed the relative ordering of the levels. Each basic-level category here had category utility higher than its subordinates but lower than its superordinate. For the musical instruments hierarchy, the results were mixed: Two basic categories (*guitar* and *piano*) were indicated as optimal, but the third (*drum*) was indicated as inferior to its superordinate (but still better than its subordinates).

Thus both versions of the collocation measure perform poorly with these data, incorrectly picking out the superordinate level as basic in nearly every case. Results for the category utility measure seem more promising, because it picks out the correct basic category in nine out of nine cases in the analyses using a restricted (single superordinate) context and in five out of nine cases with the expanded contexts. Note, too, that these expanded contexts probably exaggerate the distinctiveness of the superordinate categories, because they assume no feature overlap between the superordinate category and its new contrast category. For example, in the expanded fruit hierarchy, the superordinate level has three features that are perfectly predictive of the category, making the superordinate categories quite distinctive. It may be more realistic to assume that at least some of these features (e.g., *sweet* or *you eat it*) are shared with salient contrast categories. Redefining the expanded contexts to reflect this fact would reduce the distinctiveness, hence the utility, of the superordinates, possibly to a greater degree than for the basic or subordinate categories. To check this hypothesis, further analyses were conducted in which the expanded context was altered so that the first two features were assumed to be shared between the contrasting superordinate categories. For this modified context, the category utility measure picked out two out of three basic categories (*apple* and *peach*) as optimal in the fruit hierarchy. The collocation measure still picked the superordinate category as better than each of the basic categories. These exploratory analyses demonstrate the sensitivity of the category utility measure to changes in the context population of categories and features and also show the consistency of the collocation measure's tendency to select the superordinate level as best.

Studies Using Artificial Category Hierarchies

We have argued that basic-level categories are those that have evolved to take maximum advantage of the distribution of features in the environment, so that information about attribute values can be efficiently summarized and transmitted. People's individual cognitive mechanisms and learning strategies also may have evolved in such a way as to take maximal advantage of such structure (cf. Anderson, 1988). If this is so, then category utility might also serve as a predictor of individual performance in category-learning experiments. To test this, we applied the utility measure to data from two experiments that found basic-level learning effects using artificial stimuli.

The Murphy and Smith (1982) Experiment

The first set of data analyzed was from Experiment 1 of Murphy and Smith (1982). Subjects in this experiment were taught names for categories at all three levels of a hierarchy of artificial categories. In a later testing phase, subjects were shown a picture of a stimulus item along with a category name and were asked to verify whether the stimulus was a member of the named category. The stimulus materials (described below) were designed so that the middle level of the hierarchy was expected to be basic. The verification reaction times confirmed this prediction, with categories being verified fastest at the middle level and slowest at the superordinate level.

The stimuli consisted of 16 line drawings identified to the subjects as examples of fictitious tools. Figure 2 contains four examples of these stimuli, one selected from each of the four basic-level categories. The hierarchy consisted of two superordinate categories, identified only by function (*used for pounding* vs. *used for cutting*). Each of these superordinate categories was itself divided into two intermediate- (basic-) level categories. Each of the four intermediate categories was subdivided into two subordinate-level categories by varying one of the attributes of the intermediate-level category (e.g., the type of handle). Finally, there were two exemplars of each of the subordinate categories, a large drawing and a small one.

Because the stimuli were described as tools, they can be decomposed into parts in a straightforward manner. Each tool can be described (and is differentiated from other tools) by the shapes of its handle, shaft, and head. This suggests a natural representation of the perceptual aspects of the stimuli in terms of four multivalued attributes: three representing the shapes of the handle, shaft, and head and one additional size attribute with two values (large and small). The actual coding used is given in Table 3. Note that the names given in Table 3, for example, pounder and hammer, are used only for convenient description of the stimuli. They were not available to subjects.

For each category at each level, we computed the value of category utility and both versions of the collocation measure. One difficulty with comparing the utility measures to Jones's (1983) collocation measure is that the latter measure is defined in terms of binary features, denoting presence or absence of a property, rather than in terms of multivalued attributes. Therefore, to apply this measure to the current stimuli, we recoded the four multivalued attributes as 18 binary features, each corresponding to a single value of one of the attributes. The results of the analyses are given in Table 4, along with the mean verification reaction times from Murphy and Smith (1982).

The category utility measure correctly identifies the intermediate (basic) level as best. However, the collocation measure and the feature-possession scores incorrectly identify the superordinate level as best. It is interesting to examine why the collocation measure incorrectly rates the superordinate level as superior to the basic level. Let HN1 denote the first values of the *handle* attribute, and let SH1, HD1, and SZ1 denote the first values of the *shaft, head,* and *size* attributes, respectively (Table 2). Then the collocation value for the first basic-level category b1 is calculated as the value of $P(f|\text{b1})P(\text{b1}|f)$ summed over feature values HN1, SH1, HD1, HD2, SZ1, and SZ2. This sum is equal to $(1)(1) + (1)(1) + (.5)(1) + (.5)(1) + (.5)(.25) + (.5)(.25) = 3.25$. For the first superordinate category, sp1, the product $P(f|\text{sp1})P(\text{sp1}|f)$ is summed over the features HN1, HN2, HN3, SH1, SH2, HD1, HD2, HD3, SZ1, and SZ2. This sum is equal to $(.5)(1) + (.25)(1) + (.25)(1) + (.5)(1) + (.5)(1) + (.25)(1) + (.25)(1) + (.5)(1) + (.5)(.5) + (.5)(.5) = 3.50$. For the first three attributes, the relevant feature values are nested within the category, thus the $P(c|f)$ term is always 1, and the $P(f|c)$ terms always sum to 1 for a given attribute such as *handle.* For example, $P(\text{SH1}|\text{sp1})P(\text{sp1}|\text{SH1}) + P(\text{SH2}|\text{sp1})P(\text{sp1}|\text{SH2}) = (.5)(1) + (.5)(1) = 1$. The last attribute, *size,* breaks the tie and indicates

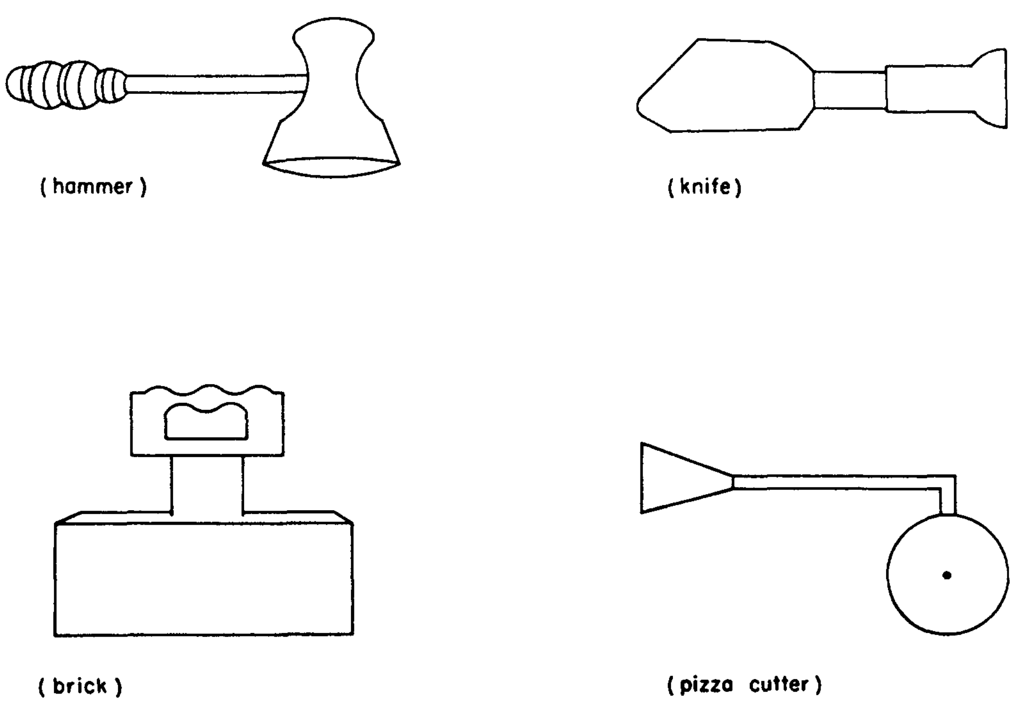

Figure 2. Examples of stimuli from each of four basic-level categories of tools. From "Basic Level Superiority in Picture Categorization" by G. L. Murphy and E. E. Smith, 1982, *Journal of Verbal Learning and Verbal Behavior, 21,* p. 3. Copyright 1982 by Academic Press. Reprinted by permission.

Table 3
Coding of the Murphy & Smith (1982) Stimuli in Terms of 4 Substitutive Attributes

Item	Category			Attribute			
	Superord	Basic	Subord	Handle	Shaft	Head	Size
1	Pounder	Hammer	Hammer 1	1	1	1	1
2				1	1	1	2
3			Hammer 2	1	1	2	1
4				1	1	2	2
5		Brick	Brick 1	2	2	3	1
6				2	2	3	2
7			Brick 2	3	2	3	1
8				3	2	3	2
9	Cutter	Knife	Knife 1	4	3	4	1
10				4	3	4	2
11			Knife 2	4	3	5	1
12				4	3	5	2
13		Pizza cutter	Pizza 1	5	4	6	1
14				5	4	6	2
15			Pizza 2	5	5	6	1
16				5	5	6	2

Superord = superordinate, subord = subordinate.

the superordinate category as superior. For this dimension, the $P(f|c)$ terms are equal to .5 for both bl and spl, but the $P(c|f)$ term is equal to .5 for Category spl and to .25 for Category bl. One can thus induce from these findings and similar examples that the collocation measure indicates that having multiple values on some attribute within a category is no different than having a single value for all members of the category, as long as all these feature values are nested within the category.

The Hoffmann and Ziessler (1983) Experiment

Hoffmann and Ziessler (1983) conducted an experiment similar to that of Murphy and Smith (1982). They divided subjects into three groups, each of which learned a different category hierarchy defined on eight schematic drawings. Each of the category hierarchies had three levels, with two top-level, four middle-level, and eight bottom-level categories. Schematic descriptions of the three hierarchies, taken from Hoffmann and Ziessler, are shown in Figure 3.

The three stimulus sets are identified as Begriffshierarchie I, II, and III. For each set, the eight stimuli (corresponding to the eight bottom-level categories) are shown along the bottom of the chart along with their exemplar names. The four middle-level and two top-level categories in each set are represented schematically by just those features that are common to all members of these categories. For example, in Begriffshierarchie I, the stimuli in the first top-level category, *ril,* all share a common jagged shape. The rils are further subdivided into two middle-level categories, *kas* and *jad,* which share common interior symbols. Finally, these middle-level categories each have two exemplars (bottom-level categories), which are differentiated by their bottom edges. Thus the three hierarchies differ in the degree to which exemplars of categories at different levels share attributes. Our representation of these stimuli was based on a straightforward coding of the drawings using three attributes: *shape, interior,* and *bottom edge,* with two, four, and four possible values, respectively.

Results of the analyses are shown in Table 5. For Hierarchy 1, the value of category utility was tied for the top (basic) and middle levels. Collocation and the feature-possession scores, on the other hand, correctly identified the top level as best. For Hierarchy 2, in which the middle level was basic, category utility correctly identified the middle level as optimal, and the collocation measure was tied for the top and middle levels. The feature-possession scores incorrectly identified the top level as basic. For Hierarchy 3, category utility correctly picked out the bottom level as basic, but collocation and the feature-possession scores incorrectly identified the top level as best.

Table 4
Mean Values of Category Utility, Collocation, and Feature-Possession Score Measures (With Mean "True" Verification Reaction Times) for Categories Used by Murphy & Smith (1982)

Measure	Level		
	Subordinate	Basic	Superordinate
Category utility	0.30	0.47	0.31
Collocation	2.13	3.25	3.50
Feature possession	3.67	5.67	8.67
Reaction time (in ms)	723	678	879

Summary of Applications

In applications to data from three hierarchies of natural categories, category utility showed promise as a valid indicator of the basic level, correctly picking out nine out of nine basic categories as superior to their superordinate and subordinate categories. In additional analyses using a plausible expanded context population of categories, category utility picked out the basic category in a majority of cases. The collocation measure

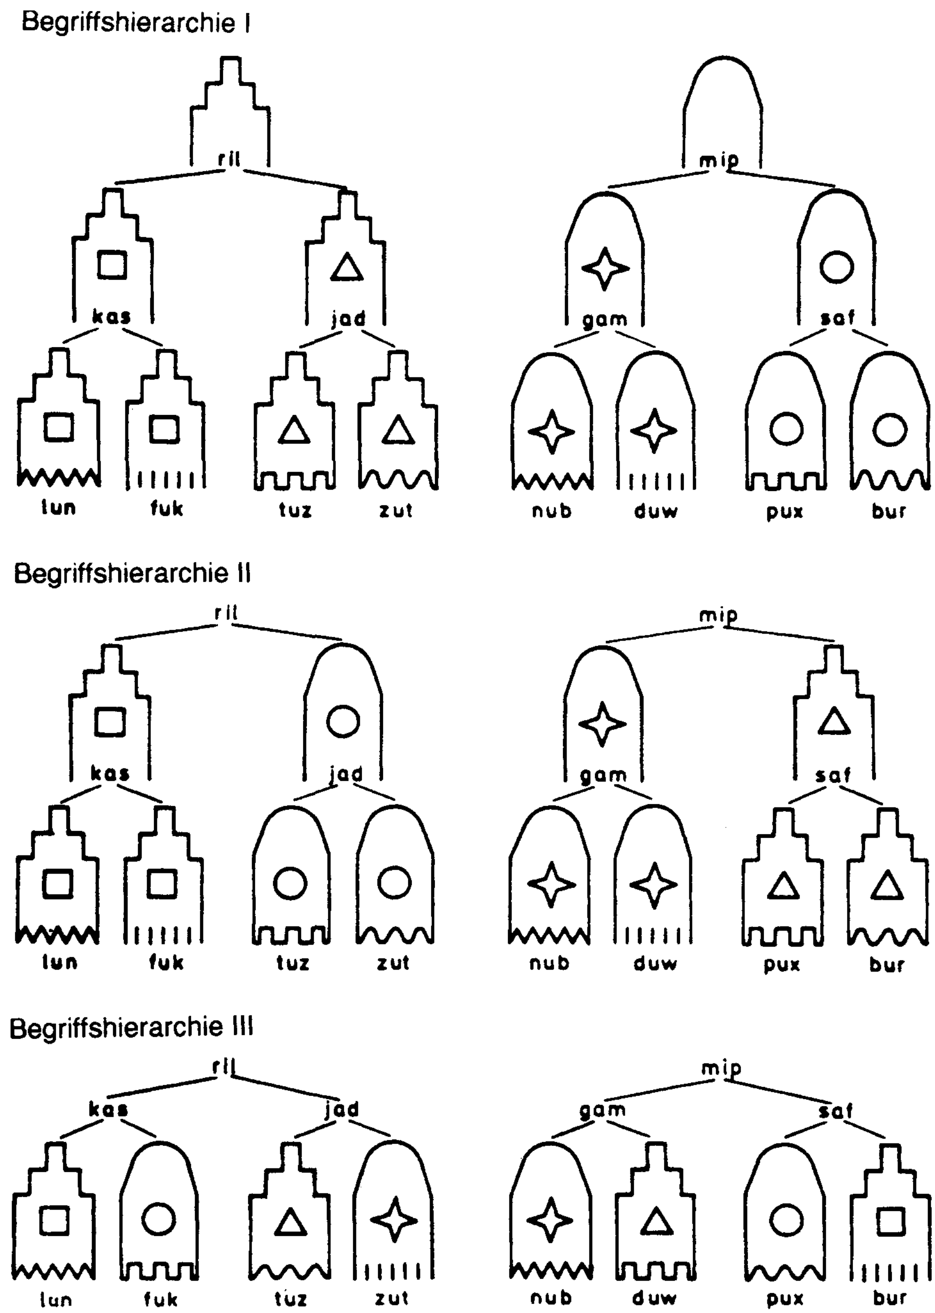

Figure 3. Schematic descriptions of three stimulus hierarchies. From "Objectidentifikation in kunstlichen Begriffshierarchien" by J. Hoffmann and C. Ziessler, 1983, *Zeitschrift fur Psychologie, 194,* p. 140. Copyright 1983 by Johann Ambrosius Barth. Reprinted by permission.

and the feature-possession scores proposed by Jones (1983) incorrectly picked out the superordinate level as best in seven out of nine and nine out of nine cases, respectively, regardless of context. For several hierarchies of artificial stimuli from laboratory studies of basic-level learning, category utility correctly identified the basic level in four out of four cases (although in one hierarchy the measure was tied for two levels). Jones's collocation measure did not perform as well, picking the correct level in only two out of four cases (again, in one of these successful cases the measure was tied for the basic and another level). The feature-possession scores invariably selected the superordinate level as basic, thus these scores identified the correct level in only one out of the four artificial stimulus hierarchies.

We believe that these results provide evidence that the category utility measure offers a useful metric for evaluating the goodness of categories and for identifying the basic level in hierarchies of natural and artificial categories. The collocation measure, although potentially adequate for this purpose on the-

Table 5
Mean Values of Category Utility, Collocation, and Feature-Possession Score Measures for the 3 Stimulus Hierarchies Used by Hoffmann & Ziessler (1983)

Measure	Level		
	Bottom	Middle	Top
Hierarchy 1			
Category utility	0.25	0.38	0.38
Collocation	1.25	2.00	2.50
Feature possession	2.00	3.00	5.00
Hierarchy 2			
Category utility	0.25	0.38	0.13
Collocation	1.25	2.00	2.00
Feature possession	1.67	3.67	4.67
Hierarchy 3			
Category utility	0.25	0.13	0
Collocation	1.25	1.25	1.50
Feature possession	2.50	2.00	5.50

Hierarchy 1 = top-level basic, Hierarchy 2 = middle-level basic, Hierarchy 3 = bottom-level basic.

oretical grounds, does not perform as well in identifying the basic level. Jones's (1983) suggested method for applying the collocation measure, involving comparing values of collocation "vertically" in a hierarchy for each feature and computing feature-possession scores for each category, seems to result in even poorer performance in identifying the basic level.

Discussion

One of Rosch et al.'s (1976) hypotheses about the structural characteristics of basic levels was that the basic level is "the most inclusive level at which the objects of a category possess numbers of attributes in common" (p. 392). If this is taken to mean that all or nearly all of the instances of a category should possess a given feature, then Rosch et al. seem to be suggesting that two factors influence the basicness of a category: the relative frequency (i.e., inclusiveness) of the category and some function of the category validities of features (i.e., the probability of a feature or attribute value given the category; cf. Medin, 1983; Tversky, 1977). The category utility measure is a quantitative measure that directly instantiates this idea, in that it defines the relative goodness of a category in terms of these two factors, plus a third: the base-rate distribution of attribute values in the context population. Inclusion of this third factor allows for effects of such factors as the distinctiveness of contrast categories and the nature of the context population of instances.

Previous structural analyses of categories and their properties have assumed that objects can be analyzed in terms of additive binary features, which represent the presence or absence of a property. Such additive features are independent in the sense that a given instance may (at least in theory) have any combination of these features. In certain of the analyses presented in this article, we have used an alternate representation scheme: substitutive attributes comprised of several possible attribute values. It is of course possible to represent such a substitutive attribute by a set of binary features in which each feature corresponds to a single value of the substitutive attribute. However, this representation ignores the fact that values of a substitutive attribute are mutually exclusive. Because an analysis that is based on substitutive attributes takes account of the logical relationships between the attribute values, we believe that this type of representation is most appropriate in many applications. Another reason why analysis in terms of substitutive attributes is useful is because such coding seems particularly appropriate for data structures (i.e., frames) commonly used in machine learning and other artificial intelligence work.

Applications in Machine Learning

Category utility provides a quantitative measure of the goodness of a category for summarizing and transmitting information. Given the normative arguments for why it is desirable to define categories that maximize this informational value, as well as the empirical evidence offered above indicating that the measure is highly correlated with human learning of artificial categories, category utility seems particularly well suited to use as a criterion for machine learning algorithms. In fact, Fisher (1987) designed an incremental "conceptual clustering" program, dubbed COBWEB, that incorporates a criterion measure that is based on category utility, in response to early reports of the category utility hypothesis (Corter & Gluck, 1985; Gluck & Corter, 1985). Specifically, COBWEB seeks partitions of the objects that maximize the mean category utility for the partition's constitutent categories. This system and its descendents have been influential in the machine-learning research community (e.g., Gennari, Langley, & Fisher, 1989).

Other Normative Accounts of Cognitive Processes

We have argued that the best categories are those that best facilitate inferences about features and communication about the features of objects. We have suggested that something analogous to evolutionary pressures have operated to select these maximally informative categories as the categories that will tend to be used frequently and survive in a language. Categories that are highly informative (as measured by category utility) will tend to be most useful to people in their individual struggle to adapt and will tend to be most useful to social groups in their attempts to develop an efficient and informative language. Basic-level categories are those categories for which this informativeness and facilitation of feature prediction is maximal, compared with superordinate and subordinate categories.

This evolutionary argument resembles other normative arguments concerning adaptive or evolutionary influences on the concepts and categories used by humans. Freyd (1983) suggested that one type of constraint on the concepts used by a society of cognitive beings is that concepts must be shareable. That is, concepts will tend to survive in a community of cognitive beings to the extent that they can be easily described and communicated among people without distortion or loss of information. In our analysis, we have emphasized the value of communicating information rather than the linguistic and other constraints that impede such communication. Nevertheless, we share Freyd's assumption that communication among people can select (and perhaps even shape) certain concepts to be preserved and used by a society.

More recently, Anderson (1988, 1990) has also taken a normative approach, analyzing several cognitive mechanisms from the standpoint of what type of performance is optimal or "rational." His analysis also begins by assuming that the main benefit of categories is to aid in prediction of feature values, although his specific formulation differs from that of Gluck and Corter (1985) in terms of what prior information the categorizer is assumed to have available.

Finally, we wish to point out that although we refer to our analysis as a normative account of the use of categories, at certain points we have assumed nonoptimal strategies on the part of the categorizer, for example in assuming that the categorizer follows a probability-matching strategy in guessing feature values. We have done this only where overwhelming experimental evidence exists to show that such deviations from "optimal" performance are in fact the rule. Thus the category utility measure is meant to provide a realistic measure of the value of a category to humans. On the other hand, it has been argued (Fisher, 1987) that in certain contexts, use of a probability-matching strategy (which generates a representative sample of predicted feature values) might actually be optimal.

Potential Extensions and Limitations

We believe that the category utility hypothesis, and its instantiation in the category utility measure, might fruitfully be applied to analyze phenomena other than the existence of basic levels. Certainly the measure seems to offer a general metric for the utility of categories, whether or not the categories in question form a hierarchy or nested set.

However, some caution should be used in attempting to apply the measure to explain the process of category learning. We have suggested that the set of categories that are most used and tend to survive in any natural language will be those that summarize relatively large amounts of information (as measured by category utility; Corter & Gluck, 1985). The idea that individual learning mechanisms might have adapted so as to tend to learn such "best" categories is a separate hypothesis, which should be examined as such. Still, a variety of factors operate in determining how people learn categories in real-life situations, and people seem to have great flexibility in the use of strategies to assist their learning. Thus the view that category utility (or any other measure) could serve to predict the order or ease of learning of categories in all situations seems naive.

In particular, two factors that have been shown to play a role in the actual learning of categories are not directly reflected in the category utility measure. First, attributes of objects differ in salience or prominence. The category utility measure as formulated above does not provide for differential salience or importance of features, although it would be straightforward to extend the model by allowing differential weights on features. Second, the model does not explicitly take account of correlations among features, whereas data from category-learning experiments (e.g. Medin, Altom, Edelson, & Freko, 1982) show that people are sensitive to correlations among features. The category utility hypothesis suggests that it is valuable to be able to predict the value of a feature and that it is reasonable to add the measure across features, regardless of any correlations among features. However, models of unsupervised category learning incorporating this or any other type of additive-cue criterion (Medin & Schaffer, 1978) will still tend to find categories described by sets of correlated features in many applications (Anderson, 1990; Fisher & Langley, 1990).

We do not view these possible shortcomings as fatal, because we do not advance the category utility hypothesis as a theory of category learning. Rather we are taking what Rosch termed a *structuralist* approach and attempting to delineate the category–feature associations that underlie the motivations for learning and use of categories.

Finally, note that general criticisms have been raised against purely structural explanations of basic levels. For example, Murphy and Medin (1985) have suggested that the coherence of categories depends in part on whether people can construct theories about why the category should exist or why certain features should covary. But this objection does not address the fact that people can learn feature–category associations even in the absence of obvious causal connections, as in experiments using artificial stimuli or in everyday superstitious behavior. Thus, although structural explanations may not account for all aspects of people's categorization performance, some account must be made of people's ability to notice and take advantage of feature–category associations. The contrast between a structuralist approach and one that is based on causal coherence parallels the distinction made in the machine-learning literature between techniques for similarity-based and explanation-based generalization (e.g. Dietterich, 1986; Lebowitz, 1986; Mitchell, Keller, & Kedar-Cabelli, 1986). A reviewer has pointed out that one of the uses of theories may be to help specify the space of relevant features. Once the space of features is determined with the aid of domain theories, then similarity-based learning and categorization processes may possibly be used. Thus both types of inference processes may be needed to adequately characterize human learning and reasoning.

References

Anderson, J. (1988, November). *A rational analysis of categorization.* Paper presented at the Twenty-Ninth Annual Meeting of the Psychonomic Society, Chicago.

Anderson, J. R. (1990). *The adaptive character of thought.* Hillsdale, NJ: Erlbaum.

Anglin (1977). *Word, object, and conceptual development.* New York: Norton.

Berlin, B., Breedlove, D. E., & Raven, P. H. (1973). General principles of classification and nomenclature in folk biology. *American Anthropology, 75,* 214–242.

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

Corter, J. E., & Gluck, M. A. (1985, June). *The pragmatics of semantics: Information, communication, and levels of abstraction.* Paper presented at the meeting of the Society for Philosophy and Psychology, University of Toronto, Toronto, Ontario, Canada.

Dietterich, T. G. (1986). Learning at the knowledge level. *Machine Learning, 1,* 287–316.

Estes, W. (1972). Research and theory on the learning of probability. *Journal of the American Statistical Association, 67,* 81–102.

Fisher, D. H. (1987). Knowledge acquisition via incremental conceptual clustering. *Machine Learning, 2,* 139–172.

Fisher, D. H., & Langley, P. (1990). The structure and function of natu-

ral categories. In G. Bower (Ed.), *The psychology of learning and motivation* (Vol. 26, pp. 241–284). San Diego, CA: Academic Press.
Freyd, J. J. (1983). Shareability: The social psychology of epistemology. *Cognitive Science, 7,* 191–210.
Garner, W. R. (1978). Selective attention to attributes and to stimuli. *Journal of Experimental Psychology: General, 107,* 287–308.
Gennari, J. H., Langley, P., & Fisher, D. H. (1989). Models of incremental concept formation. *Artificial Intelligence, 40,* 11–62.
Gluck, M. A., & Corter, J. E. (1985). Information, uncertainty, and the utility of categories. In *Proceedings of the Seventh Annual Conference of the Cognitive Science Society* (pp. 283–288). Hillsdale, NJ: Erlbaum.
Gluck, M. A., Corter, J. E., & Bower, G. H. (1992). *Basic levels in the learning of category hierarchies: An adaptive network model.* Manuscript in preparation.
Hoffmann, J., & Ziessler, C. (1983). Objectidentifikation in kunstlichen Begriffshierarchien [Object identification in artificial concept hierarchies]. *Zeitschrift fur Psychologie, 194,* 135–167.
Horton, M. S., & Markman, E. M. (1980). Developmental differences in the acquisition of basic and superordinate categories. *Child Development, 51,* 708–719.
Jolicoeur, P., Gluck, M., & Kosslyn, S. (1984). Pictures and names: Making the connection. *Cognitive Psychology, 16,* 243–275.
Jones, G. V. (1983). Identifying basic categories. *Psychological Bulletin, 94,* 423–428.
Lebowitz, M. (1986). Not the path to perdition: The utility of similarity-based learning. In *Proceedings of the Fifth National Conference of Artificial Intelligence* (pp. 533–537).
Medin, D. L. (1983). Structural principles of categorization. In B. Shepp & T. Tighe (Eds.), *Interaction: Preception, development and cognition* (pp. 203–230). Hillsdale, NJ: Erlbaum.
Medin, D. L., Altom, M. W., Edelson, S. M., & Freko, D. (1982). Correlated symptoms and simulated medical classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8,* 37–50.
Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85,* 207–238.
Mervis, C. B., & Crisafi, M. A. (1982). Order of acquisition of subordinate-, basic-, and superordinate-level categories. *Child Development, 53,* 258–266.
Mervis, C. B., & Rosch, E. (1981). Categorization of natural objects. *Annual Review of Psychology, 32,* 89–115.
Mitchell, T., Keller, R., & Kedar-Cabelli, S. (1986). Explanation-based learning: A unifying view. *Machine Learning, 1,* 47–80.
Murphy, G. L. (1982). Cue validity and levels of categorization. *Psychological Bulletin, 91,* 174–177.
Murphy, G. L., & Brownell, H. H. (1985). Category differentiation in object recognition: Typicality constraints on the basic category advantage. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11,* 70–84.
Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92,* 293–316.
Murphy, G. L., & Smith, E. E. (1982). Basic level superiority in picture categorization. *Journal of Verbal Learning and Verbal Behavior, 21,* 1–20.
Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8,* 382–439.
Shannon, C., & Weaver, W. (1949). *The mathematical theory of communication.* Chicago: University of Illinois Press.
Tversky, A. (1977). Features of similarity. *Psychological Review, 84,* 327–352.
Tversky, B., & Hemenway, K. (1984). Objects, parts, and categories. *Journal of Experimental Psychology: General, 113,* 169–193.

Received July 27, 1990
Revision received April 9, 1991
Accepted May 10, 1991 ■

Zahn-Waxler Appointed New Editor, 1993–1998

The Publications and Communications Board of the American Psychological Association announces the appointment of Carolyn Zahn-Waxler as editor of *Developmental Psychology.* Zahn-Waxler is associated with the National Institute of Mental Health. As of January 1, 1992, manuscripts should be directed to

Carolyn Zahn-Waxler
4305 Dresden Street
Kensington, Maryland 20895

Manuscript submission patterns make the precise date of completion of the 1992 volume uncertain. The current editor will receive and consider manuscripts through December 1991. Should the 1992 volume be completed before that date, manuscripts will be redirected to the incoming editor for consideration in the 1993 volume.