Cahiers de Psychologie Cognitive, 1986, 6 (2), 187-202.

CONCEPTS: STATIC DEFINITIONS OR CONTEXT-DEPENDENT REPRESENTATIONS ?

Lawrence W. Barsalou* and Douglas L. Medin**

* Department of Psychology, Emory University, Atlanta, GA 30322, USA.

** Department of Psychology, 603 E. Daniel, University of Illinois, Champaign, IL 61820, USA.

Résumé :

Les théories relatives à la formation et à l'utilisation des concepts les ont traditionnellement décrits à l'aide de définitions statiques. Bien que ces théories maximisent l'exactitude des inférences dans les situations de classification, l'exclusion d'informations non-essentielles des concepts (ie., qui ne relèvent pas de la définition) donne lieu à une réduction importante des connaissances. L'information non-essentielle joue un rôle central dans les représentations individuelles des catégories parce qu'elle engendre des inférences utiles. De plus, lorsque les individus tiennent compte du contexte, les inférences basées sur l'information non-essentielle peuvent devenir quasi exactes. La littérature démontre l'importance de plusieurs types d'informations non-essentielles : prototypes, exemplaires, propriétés corrélatives et théories intuitives. L'information non-essentielle permet aux individus d'émettre des inférences catégorielles, basées sur leur expérience personnelle et sur le contexte momentané, qui ont de fortes chances d'être vraies. Nous argumentons également contre la théorie "essence-et-identification" proposée pour remplacer les théories traditionnelles. Cette théorie soutient que les définitions ne comprennent que l'essence des concepts, bien que d'autres types d'informations puissent affecter la performance. Nous suggérons que l'essence d'un concept reflète, non pas une définition statique, mais plutôt l'expérience et les théories intuitives des individus.

Key words : concepts, categories, classification, knowledge, representation.

Mots clés : concepts, catégories, classification, connaissance, représentation.

1. INTRODUCTION

One of the basic assumptions of Cognitive Science is that people have extensive world knowledge. When relevant to the current situation, knowledge from long term memory is utilized to comprehend what is occurring, to make predictions about what will happen, to guide action, and to structure other basic cognitive operations. Given the importance of knowledge in human behavior, many attempts have been made to characterize its structure and to identify the processes that act upon it. Implicit in many of these accounts has been the assumption that knowledge is comprised of relatively static and well-bounded packets of information (e.g., definitions, schemas, frames, scripts), each of which represents some particular kind of thing (e.g., *birds, weddings*). Most importantly, these accounts assume that the same packet of information is used to represent all instances of a category across a wide range of contexts.

The classical theory of categories.

The classical theory of categories provides the quintessential example of a view that postulates static and well-bounded packets of knowledge (e.g., Katz, 1972; Lyons, 1968). We focus on the classical view for this reason, although much of what we say also applies to other theories of representation as well. According to the classical view, categories are defined by sets of properties that are jointly sufficient and singly necessary for category membership (see Smith and Medin, 1981, for further discussion). As Lakoff (in press) points out, classical theorists often assume that such definitions exist by virtue of the objective world. Because the objective world contains various kinds of things that share defining properties, human knowledge comes to represent these properties.

Several important conclusions follow from this view. First, it follows that the representation of a category in memory (what we will refer to as a *concept*) is well-bounded, since some clearly specifiable set of properties comprises it. Second, different people should have the same concept, or at least highly similar concepts, for a particular category, assuming that everyone has reasonably good access to the objective world or to information about it. Third, concepts should be static within a particular individual, since the defining properties for a category should rarely if ever change, objectively speaking. Because what objectively defines a category never changes, a category should always be represented by the same static and well-bounded packet of information across different individuals and across different contexts.

Preview.

The remainder of this paper provides a brief overview of a fundamental shift in how theorists have come to view concepts over the last ten years or so. This shift not only reflects dissatisfaction with the classical theory of categories, but also with the more general

assumption that people's knowledge is neatly parsed into static and well-bounded packets of information. In section 2, we argue that people's need to generate useful expectations pushes concept representations away from the static packets of knowledge associated with the classical view. In sections 3 and 4, we briefly and selectively review recent research showing that people's representations of categories vary dynamically with their experience and current context. In section 5, we consider and reject the core-plus-identification defense of the classical view. And in section 6, we conclude that concepts, instead of being viewed as static definitions, should be viewed as dynamic, context-dependent representations.

2. THE TRADEOFF BETWEEN LOGICAL ADEQUACY AND USEFUL EXPECTATIONS

The classical theory of concepts places a premium on accurate inferences from category membership. Consider the category of *birds* from the classical perspective. Since all birds have feathers, this property can be assumed to be true for any entity classified as a bird. Similarly, *living thing* can also be assumed to be true, through the logical inheritance relation that exists between *birds* and *animals*. Since only definitional properties occur in concepts, according to the classical view, all the inferences that someone can draw about an entity from knowing its category membership will be accurate. Inaccurate inferences are impossible because every property that could be inferred from category representations necessarily belongs to all exemplars. In addition, this view has the virtue of limiting miscommunication, since the only inferences that a listener can draw from a speaker's use of a category name are those true by definition.

A problem with representing categories in this way is that a tremendous amount of potentially important information is discarded. For example, since not all birds fly or have hollow bones (e.g., penguins and ostriches), these properties would not be included in the concept for *bird*. Yet if some entity is found to be a bird, then there is a high probability that it flies and has hollow bones. Of further informative value are the related generalizations that birds with hollow bones fly and that birds without hollow bones do not. It may also be important in many cases to entertain inferences that do *not* have a high probability of being true. For example, in encountering a snake while rock climbing, it could be vitally important to consider the fact that some snakes are poisonous, even though most are not. Given the value of such expectations, the requirement that inferences be one-hundred percent reliable is clearly an unaffordable luxury.

In addition, taking advantage of information in particular contexts can improve predictions. For example, a mushroom found in the wild can be expected to have a much higher likelihood of being poisonous than one found in a grocery store. Although context-specificity improves predictions, it undermines static, context-independent definitions.

Given the advantages of being able to generate useful expectations, it follows that natural categorization systems may generally operate with less than perfect accuracy. Awareness of this fact has caused many theorists to shift their concern from the logical structure of concepts to relations between concepts and the world that enable the production of useful expectations (see Johnson-Laird, Herrmann, and Chaffin, 1984, for a different perspective on this point). The remainder of our paper documents this shift from viewing concepts as context-independent, definitional representations to viewing them as context-dependent, experience-based representations.

3. THE ROLE OF EXPERIENCE IN THE REPRESENTATION OF CATEGORIES

The findings in this section illustrate how people's long term context (i.e., their experience) results in category representations that differ substantially from static definitions.

Prototype theory.

Prototype theory represents a major departure from the classical view, since it rejects the idea that conceptual representations are necessarily definitional (Rosch and Mervis, 1975; Smith, Shoben, and Rips, 1974; Smith and Medin, 1981). First, it has been frequently argued that definitions do not exist for many categories (e.g., see Wittgensteins's, 1953, analysis of *games*). Second, even when definitions exist, representations of categories include more than just singly necessary and jointly sufficient properties. In particular, many theorists have suggested that people represent categories with *prototypes*, which contain properties *characteristic* of category members (e.g., Hampton, 1979; Rosch and Mervis, 1975; Smith et al., 1974). So although *made of wood* is not defining for *furniture*, it is often true or characteristic of category members and so may be included in the prototype for the category. According to prototype theory, the representation of a category contains properties that have a high, although not necessarily perfect, likelihood of occurring for its exemplars. To the extent a property is characteristic of category members, it becomes central in the prototype for the category.

Another basic assumption of prototype theory is that people abstract characteristic properties of a category through their experience with category instances. Although Rosch and Mervis (1975) argued that the prototypes people abstract reflect the objective structure of the environment, another way to view prototypes is as reflecting experience. It follows from this view that if different individuals experience different sets of category instances, they will abstract different prototypes. The prototypes two different people abstract for the same category may contain at least somewhat different characteristic properties that reflect their different experiences with the category. This assumption has been incorporated into theories proposed by Bourne, Ekstrand, Lovallo, Kellogg, and Hiew (1976) and Kellogg (1981) on concept learning.

Exemplar theory.

Along with classical definitions and prototypes, a third way people could represent categories is with exemplars (Brooks, 1978; Medin and Schaffer, 1978). Instead of inducing a definition or abstracting a prototype for a category, people could simply store memories of particular exemplars and later use specific exemplars or sets of exemplars to represent the category. When classifying a new stimulus, for example, people could search memory for the exemplar or exemplars most similar to the stimulus and classify the new stimulus by analogy with the category associated with the retrieved exemplar(s).

It follows from this view that people's representations of categories reflect their experience. To the extent people experience different distributions of exemplars for a category, their representations of it will vary.

Correlated properties.

Another way to view some exemplar theories is in terms of correlated properties. For example, the Medin and Schaffer (1978) context theory of classification is based on the idea that all the properties defining an exemplar are perfectly correlated with one another and in a sense provide context for each other. Much work has since indicated that new stimuli possessing correlations of properties encoded via previous exemplars are classified more easily than stimuli not possessing these correlated properties (e.g., Medin and Schaffer, 1978; Medin and Schwanenflugel, 1981; Medin, Altom, Edelson, and Freko, 1982). Informally, most people would agree that small birds are more likely to sing than large birds (Malt and Smith, 1984). Since such knowledge is not definitional, it could not be derived from the definition of *bird*. Instead awareness of such correlations reflects people's experience with particular category exemplars (1).

Intuitive theories.

Classical, prototype, and exemplar theories all assume that concepts are comprised of properties. However most such theories pay little if any attention to the origins of these properties or to relations between properties. However recent theorists have argued that

---

(1) It should be noted that sensitivity to correlated properties does not necessarily imply that subjects are using exemplars to perform classification. As noted by Medin and Schaffer (1978, p. 231-232), summarized representations that store frequencies for how often pairs of properties occur (and for higher-order combinations of properties) also exhibit sensitivity to correlated properties (e.g., Reitman and Bower, 1973).

people's intuitive theories about the world play central roles in conceptual processing (Murphy and Medin, 1985; Lakoff, in press; McCauley, in press; Neisser, in press; Schank, 1985; Schank, Collins, and Hunter, in press). Intuitive theories can be viewed as constructs that people use to *explain* a category's structure, behavior, origin, and so forth. These theories select, interpret, and organize properties in the process of forming category representations. An implicit personality theory, for example, may structure the category of *introverts* by specifying and interrelating the behaviors associated with it (e.g., *self-reflection, nervousness, shyness*).

The role of theories in categorization again demonstrates the importance of experience. To the extent individuals acquire different intuitive theories about the world, they may develop different representations of categories. For example, lay people and trained botanists may have different theories about plants and, as a result, may form different representations of them.

Instability of category representations.

Many of the findings in this section and the next have to do with *graded structure*. Graded structure has two components (Barsalou, 1983, 1985). First, exemplars within a category vary in how typical or in how good an example they are of the category. In *birds*, for example, *robin*, *pigeon*, and *ostrich* are generally perceived as decreasing in typicality. Second, entities that do not belong to a category vary in how typical they are of non-category members. With respect to non-members of *birds*, for example, *chair*, *airplane*, and *butterfly*, are generally perceived as decreasing in typicality. Although the phenomenon of graded structure has been one of the primary motivating factors responsible for the rejection of the classical view (Smith and Medin, 1981), our present concern is with the relation between graded structure and concept stability (2).

---

(2). The same general methodology was used for all the experiments discussed in this section. Subjects were presented with the name of a category (e.g., *birds*), followed by the names of exemplars belonging to that category arranged in a random order. In most experiments, half the categories were common taxonomic categories (e.g., *furniture*, *birds*) and half were goal-derived categories (e.g., *things to take on a camping trip*, *things to buy at the grocery store*); see Barsalou (1985) for further discussion of these category types. Subjects' task was to judge each exemplar (e.g., on a seven point scale) for how good an exemplar it is of the category. With respect to *groups* of subjects, agreement in graded structure reflected how well the *average* typicality judgments for exemplars across subjects correlated between the two groups. With respect to *individuals*, agreement reflected how well judgments correlated either between two individuals or within the same individual on two occasions.

– *Instability between populations.* Barsalou and Sewell (1984, Experiment 2) found that different populations of subjects perceive different graded structures in the same category. Barsalou and Sewell asked Emory University undergraduates and Emory University faculty to judge typicality in the same 20 categories and found very low correlations between the average graded structures of the two populations (= .23 on the average). Presumably, the substantial difference in how these two populations perceive the same categories reflects different experience with these categories.

– *Instability between individuals.* Barsalou (in press) reviews work on how well different people within the same population agree on the graded structure of the same category (also see Barsalou and Sewell, 1984). Statistically, between-subject agreement can be viewed as the average correlation between all possible pairs of subjects in terms of how they order exemplars by typicality. Across many different studies using different subjects, different categories, different exemplars, and different contexts, the average between-subject agreement was generally only around .50. That is, one subject's typicality judgments generally only correlated .50 with another subject's typicality judgments from the same population. Another way of viewing this is that one subject's typicality judgments generally accounted for only around 25% of the variance in another subject's typicality judgments. Presumably, the substantial difference in how different individuals perceive the same category reflects different experience with the category.

– *Instability within individuals.* Barsalou, Sewell, and Ballato (1986) explored how stable a given individual's graded structure for the same category is over a one month period. More specifically, they asked subjects to judge typicality in the same 40 categories on two days a month apart. The average agreement across subjects and categories for how well a subject's rankings on the first day correlated with their rankings for the same category on the second day was generally around .78. That is, a subject's rankings on the first day only accounted for around 61% of the variance in their rankings for the same category on the second day. Presumably, the substantial difference in how the same individual perceives the same category reflects the fact that their experience with the category over a one month period has changed how they view it. Because individuals show much higher stability after only an hour's delay (.92), the instability at longer delays does not simply reflect measurement error or some other source of random variability – instead it reflects change in the individual's representation of the category.

These studies also addressed the question of which exemplars *within* a category are primarily responsible for changes in a person's graded structure. Across a variety of task conditions, moderately typical exemplars were always least stable, whereas prototypical and atypical exemplars were always more stable and were generally equally stable to each other. Prototypical and atypical exemplars, however, also showed substantial instability. Because atypical exemplars showed more stability than moderately typical exemplars, these

results are inconsistent with the view that instability increases as exemplars approach category boundaries. These results are also inconsistent with the view that typicality and membership judgments reflect a unitary process, since Mc Closkey and Glucksberg (1978) found that the instability of membership decisions *did* increase as exemplars approached category boundaries.

4. THE ROLE OF CONTEXT IN THE REPRESENTATION OF CATEGORIES

The findings just reviewed illustrate how an individual's long term context or experience affects their category representations. These next findings illustrate how an individual's current context affects their category representations.

Linguistic context.

Roth and Shoben (1983) presented subjects with the name of a category embedded in a linguistic context followed by the names of exemplars. Subjects' task was to judge each exemplar for how good an example it was of its category. Roth and Shoben found that different linguistic contexts resulted in very different graded structures for the same category. For example, when *animal* was encountered in a sentence about milking, *cow* and *goat* were perceived as typical; whereas when *animal* was encountered in a sentence about riding, *horse* and *mule* were perceived as typical. In another experiment, Roth and Shoben found that linguistic context also affected how quickly subjects were able to access information about exemplars from memory.

These results demonstrate that the context in which people encounter a category can have a major influence on how they perceive it. There does not appear to be a static packet of information that represents a category in all contexts. Instead different information seems to represent the same category in different contexts.

Point of view.

Barsalou and Sewell (1984) had American undergraduates take different points of view while judging typicality. For example, subjects took various international points of view, such as for the average American, Chinese, or French citizen; or subjects took various domestic points of view, such as for the average housewife, businessman, or hippie. In general, subjects taking different points of view perceived different graded structures in the same categories. In many cases, the ways in which subjects taking different points of view ordered exemplars for the same category were uncorrelated or inversely related. Since subjects taking the same point of view agreed with each other (average between-subject agreement of around .45), this result does not reflect random responding. We assume

instead that this agreement reflects stereotypes about points of view that are shared by members of a population. Again it appears that the context in which a category is perceived -in this case a stereotype- results in different representations of the same category.

Ad hoc categories.

Barsalou (1983) observed that people frequently construct new categories as they become necessary to achieving a current but novel goal (i.e., ad hoc categories). For example, if someone is going camping for the first time, he or she may construct ad hoc categories for *places to go camping, times to go camping, things to take camping, things to beware of while camping*, and so forth. This ability to construct new categories further illustrates people's highly dynamic and flexible ability to construct representations that meet contextual constraints. Whereas the findings reported for linguistic context and point of view illustrate that contexts alter the representation of an existing category, the existence of ad hoc categories illustrates that contexts can engender completely new category representations.

5. IMPLICATIONS FOR THE CLASSICAL VIEW

There do not appear to be static packets of information that represent categories. In terms of the classical view, there generally do not appear to be definitions associated with categories that reflect the objective structure of the environment. Instead people's representations of categories appear to vary substantially with their long term context and with their current context.

The core-plus-identification defense of the classical view.

One way the classical view has attempted to handle these findings has been to propose that people's representations of categories are comprised of core properties and identification procedures (Armstrong, Gleitman, and Gleitman, 1983; Miller and Johnson-Laird, 1976; Osherson and Smith, 1981, 1982; Smith and Medin, 1981). Core properties comprise a static rule or definition that ultimately determines category membership, whereas identification procedures provide heuristic means of identifying category members. These identification procedures utilize characteristic or probable properties of exemplars and are not guaranteed to provide accurate classifications. With respect to *bachelor*, for example, the core might contain *unmarried, adult, male*, and *human;* whereas identification procedures might utilize properties such as *young, heterosexual, dates many women*, and so on.

To defend the existence of definitional cores, the core-plus-identification view explains effects of experience and context in terms of identification procedures. For example, this view assumes that the

prototypical properties, correlated properties, and exemplars that people acquire from experience simply reside in identification procedures instead of in cores. Similarly this view might argue that instability between populations, between individuals, and within individuals simply reflects the fact that different experience between individuals and within individuals over time results in different identification procedures. A similar argument can be made for context effects, namely, people use different identification procedures in different contexts. Most importantly, this view holds that, while identification procedures vary, definitional cores remain constant both between and within individuals. It follows from this position that there are static packets of information associated with categories and that it is these packets that are ultimately the most important parts of category representations.

Identifying the cores of concepts. TThere are a number of problems with this defense. First, if there are definitional cores for categories, then we should be able to identify them. To the extent we fail to identify such cores, we should be worried about their existence and, in turn, about there generally being static structures associated with categories. Yet Wittgenstein's (1953) challenge to provide definitions for categories like *games* has seldom been met. Even for scientific categories, where definitions should readily be found, it is often difficult to discover definitions. For example, it has been well-known in the science of biology for some time that classical definitions do not exist for many species (e.g., Sober, 1984; Stebbins, 1957). Moreover, even when technical definitions exist for certain categories, it is not at all clear that people know and use them.

In general, there is no reason to be confident about the existence of definitional cores, either in nature or in human knowledge. Although our inability to identify cores does not necessarily imply their lack of existence, our continued inability to identify cores decreases our confidence in their presence.

The instability of definitional cores. If there are definitional cores for categories, then different people should share the same core for a category, and a given person should use the same core across different situations. To test this, Bellezza (1984) asked subjects to provide definitions for concrete nouns (e.g., *car, cat, chair*), where each subject defined each noun twice on two days a week apart. If people have definitional cores for categories, then it seems reasonable that they should be able to access them when asked to define the words that refer to them. Surprisingly, when Bellezza broke the definitions provided by subjects down into propositions, he found substantial instability. More specifically, the average overlap *between* the definitions of different subjects for the same word was only .22. In addition, the average overlap *within* a subject for his or her two attepts to define the same word was only .48. These results certainly do not support the hypothesis that there are definitional cores for categories. Different people's definitions for the same noun barely resemble each other, and the same person's definition of a noun on two days a week apart shows substantial instability.

The instability of category membership decisions. In defense of the core-plus-identification view, one could argue that people really do have definitional cores but are unable to state them verbally. As a result, instability is observed in verbal definitions. To maintain this defense, however, it is necessary at some point to provide evidence for the existence of cores. If no such evidence can be obtained, then we have no reason to believe in their existence. One tack might be to argue that people can not directly report cores but that cores surface indirectly in tasks involving category membership. Even though people can not describe their definitions for categories, they utilize them when deciding whether stimuli belong to categories. It is difficult to think of any other categorization task for which cores should be more important.

However a study by McCloskey and Glucksberg (1978) strongly suggests that people do not utilize definitional cores when deciding whether stimuli belong to categories. Subjects in this study received the names of categories and a large number of possible members on two days a month apart. Subjects' task on each occasion was to decide whether each possible member belonged to a particular category. McCloskey and Glucksberg found substantial differences *between* subjects in the exemplars that they said belonged to a particular category. In addition, McCloskey and Glucksberg found that *individual* subjects frequently changed their mind over the one month period about whether a particular exemplar belonged to a category.

These classification data argue against the existence of definitional cores. Different subjects do not appear to have the same core for a category, and the same subject does not appear to have a stable core. If cores existed, classification performance should be much more stable. Not only are subjects' verbal definitions of categories unstable, there do not appear to be cores that determine classification judgments.

It is always possible to argue that there is yet some other function of concepts for which cores operate in their pristine form. As long as that function remains unspecified, however, there is the danger that cores are like unicorns – they can be conceived of but are rarely if ever perceived.

Are the cores that exist important? As mentioned earlier, it is obviously true that some categories have clear-cut definitions (e.g., *bachelors*). But in many such cases, the intuitive theories that are intertwined with concepts may relegate these definitions to auxiliary status. For example, *bachelor* seems closely intertwined with traditional notions of being marriageable. Consequently most people would not consider a priest or a homosexual male to be a bachelor, even though they are, technically speaking. This suggests that prototypes and exemplars constrained by intuitive theories play a central role in even those categories having definitions. Instead of being central to categorization, definitions, when they exist, may generally play minor roles and only come into play in highly specialized contexts. For example, definitions may be important in technical domains where precision is important (e.g., in mathematics) or in other domains where making correct classifications has high payoffs (e.g., being

sure in a tax audit that a particular man is really a bachelor and is not married).

Experientially-based cores.

Although concepts may generally not have definitional cores, they may typically have experientially-based cores. As found by Barsalou (1982), certain information is automatically activated every time the representation of a category is accessed, regardless of the current context. For example, *rattlesnake* automatically activates *poisonous* across contexts, *diamond* automatically activates *valuable*, and *skunk* automatically activates *unpleasant smell* (also see Barsalou and Ross, 1986). As suggested by Barsalou and Bower (1980), properties become core properties because they are processed in conjunction with a category on so many occasions that they become automatized.

Experientially-based cores differ from definitional cores in the following two ways. First, experientially-based cores are not definitional. There is no a priori reason why any core property should be either necessary or sufficient for category membership. Instead it is simply information that occurs frequently enough for a category in someone's experience to become automatized. As suggested by Barsalou and Bower (1980), information may be processed frequently because it is diagnostic for category membership or because it is relevant to a typical function the category serves. Information may also be processed frequently if it is highly relevant to people's intuitive theory about a category (Murphy and Medin, 1985).

A second way in which experientially-based and definitional cores differ is that only experientially-based cores vary between and within individuals. Between individuals, people with different experiences may frequently process different properties for the same category such that their core properties vary. Although there is no reason to expect that experientially-based cores should necessarily be the same for different people, common cultural experience and common intuitive theories may result in their being some core information that is shared. In fact, the moderate between-subject agreement that exists for graded structure (as discussed earlier) may be the result of culturally shared core information (Barsalou, in press). Within individuals, there is no reason to expect that a person's experientially-based core for a category will remain stable over time. Instead it could change, assuming that the information a person frequently processes for a category changes. Such changes may generally be slow and small, except when highly frequent experience with the category occurs during a short time period (as in job training or school).

Although experientially-based cores may exist for categories, there is no reason to believe that they solely comprise the representations of categories. Instead they may only comprise part of the representation of a category on a particular occasion, along with other information that is part of the representation because it is relevant in that context (Barsalou, 1982). Whereas core properties provide people with a means of producing expectations that are

likely to be true from their experience, context-dependent properties provide a means of producing expectations that are likely to be true in the current context.

6. THE MEANING OF "CONCEPT"

The term, "concept", has traditionally referred to the definitional information by which people determine category membership. As just discussed, this view proposes that typicality effects, exemplar effects, context effects, and so on reflect less important, non-definitional information associated with categories. For example, Armstrong et al. (1983) take this position in explaining the fact that formal categories such as *odd number* and *squares* have graded structure. They say (p. 284), "... [graded structure] must have been a fact about something other than the structure of concepts," thereby implying that definitions or cores lie at the heart of concepts and that typicality effects are incidental and immaterial phenomena.

Yet, as we have seen, there is no reason to believe that definitions regularly exist or, when they do, that they play a central role in how people use concepts. Instead the information that theorists refer to as peripheral to the structure of concepts is the information that appears to be doing most of the work. In support of this, recent views of the categorization literature (Mervis and Rosch, 1981; Medin and Smith, 1984) have generally found that no other variable is as important as typicality in predicting performance on basic categorization tasks such as acquisition, production, and classification.

Since definitions do not appear central to human categorization, "concept" may be better used in referring to the representations that people typically construct for categories. Although these representations may not serve to discriminate all possible category members from non-members, or to generate infallible inferences, they do serve people's representational purposes. More specifically, they provide people with a way of representing a category such that they can generate useful expectations that are likely to be true from their experience and that are likely to be true in the current context.

ACKNOWLEDGMENTS

We are grateful to Renée Baillargeon and Robert McCauley for their assistance with this paper and to two reviewers for their helpful comments. This work was supported by National Science Foundation grant IST 8308984 to Lawrence W. Barsalou and by U.S. Public Health Services grand MH 32370 to Douglas L. Medin.

ABSTRACT:

Theories of concepts have traditionally proposed that concepts are static definitions. Although such theories maximize the accuracy of inferences made during classification, their exclusion of non-defining information from concepts results in the loss of important sources of knowledge. Because non-defining information generates useful inferences, it is central to people's representations of categories. In addition, when context is taken into account, inferences based on non-defining information can become quite accurate. We review literature showing the importance of many kinds of non-defining information, including prototypes, exemplars, correlated properties, and intuitive theories. This information enables individuals to make inferences about a category that are likely to be true on the basis of their personal experience and likely to be true in the current context. We also present arguments against the "core-plus-identification" revision of traditional theories in which definitions comprise the cores of concepts but other kinds of information influence performance. We propose instead that the cores of concepts reflect experience and intuitive theories rather than static definitions.

REFERENCES:

Armstrong, S.L., Gleitman, L.R., and Gleitman, H. (1983). On what some concepts might not be. *Cognition, 13*, 263-308.

Barsalou, L.W. (1982). Context-independent and context-dependent information in concepts. *Memory and Cognition, 10*, 82-93.

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

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

Barsalou, L.W. (in press). The instability of graded structure: Implications for the nature of concepts. To appear in U. Neisser (Ed.), *Concepts reconsidered: The ecological and intellectual bases of categories*. Cambridge: Cambridge University Press.

Barsalou, L.W., and Bower, G.H. (1980). A priori determinants of a concept's highly accessible information. Paper presented at the Meeting of the American Psychological Association, Montreal, September 1980.

Barsalou, L.W., and Ross, B.H. (1986). The roles of automatic and strategic processing in sensitivity to superordinate and property frequency. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12*, 116-134.

Barsalou, L.W., and Sewell, D.R. (1984). Constructing representations of categories from different points of view. Emory Cognition Project Report #2, Emory University.

Barsalou, L.W., Sewell, D.R., and Ballato, S.M. (1986). The instability of graded structure. Work in progress, Emory University.

Bellezza, F.S. (1984). Reliability of retrieval from semantic memory: Noun meanings. *Bulletin of the Psychonomic Society, 22,* 377-380.

Bourne, L.E., Jr., Ekstrand, W.R., Lovallo, W.R., Kellogg, R.T., and Hiew, C.C. (1976). Frequency analysis of attribute identification. *Journal of Experimental Psychology: General, 105,* 294-312.

Brooks, L. (1978). Non-analytic concept formation and memory for instances. In E. Rosch and B.B. Lloyd (Eds.), *Cognition and categorization.* Hillsdale, NJ: Erlbaum.

Hampton, J.A. (1979). Polymorphous concepts in semantic memory. *Journal of Verbal Learning and Verbal Behavior, 18,* 441-461.

Johnson-Laird, P.N., Herrmann, D.J., and Chaffin, R. (1984). Only connections: A critique of semantic networks. *Psychological Bulletin, 96,* 292-315.

Katz, J.J. (1972). *Semantic theory.* New York: Harper & Row.

Kellogg, R.T. (1981). Feature frequency in concept learning: What is counted? *Memory and Cognition, 9,* 157-163.

Lakoff, G. (in press). *Women, fire, and dangerous things: What categories tell us about the nature of thought.* Chicago: University of Chicago Press.

Lyons, J. (1968). *Introduction to theoretical linguistics.* Cambridge: Cambridge University Press.

Malt, B.C., and Smith, E.E. (1984). Correlated properties in natural categories. *Journal of Verbal Learning and Verbal Behavior, 23,* 250-269.

McCauley, R. (in press). The role of theories in a theory of concepts. To appear in U. Neisser (Ed.), *Concepts reconsidered: The ecological and intellectual bases of categories.* Cambridge: Cambridge University Press.

McCloskey, M., and Glucksberg, S. (1978). Natural categories: Well-defined or fuzzy sets? *Memory and Cognition, 6,* 462-472.

Medin, D.L., Altom, M.W., Edelson, S.M., and Freko, D. (1982). Correlated symptoms and simulated medical classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8,* 37-50.

Medin, D.L., and Schaffer, M.M. (1978). A context theory of classification learning. *Psychological Review, 85,* 207-238.

Medin, D.L., and Schwanenflugel, P.J. (1981). Linear separability in classification learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 7,* 355-368.

Medin, D.L., and Smith, E.E. (1984). Concepts and concept formation. *Annual Review of Psychology, 35,* 113-138.

Mervis, C.B., and Rosch, E. (1981). Categorization of natural objects. *Annual Review of Psychology, 32,* 89-115.

Miller, G.A., and Johnson-Laird, P.N. (1976). *Language and Perception.* Cambridge, MA: Harvard University Press.

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

Neisser, U. (in press). From direct perception to conceptual structure. To appear in U. Neisser (Ed.), *Concepts reconsidered: The ecological and intellectual bases of categories.* Cambridge: Cambridge University Press.

Osherson, D.N., and Smith, E.E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition, 9,* 35-58.

Osherson, D.N., and Smith, E.E. (1982). Gradedness and conceptual combination. *Cognition, 12*, 299-318.
Reitman, J.S., and Bower, G.H. (1973). Storage and later recognition of exemplars. *Cognitive Psychology, 4*, 194-206.
Rosch, E.H., and Mervis, C.B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology, 7*, 573-605.
Roth, E.M., and Shoben, E.J. (1983). The effect of context on the structure of categories. *Cognitive Psychology, 15*, 346-378.
Schank, R.C. (1985). Looking at learning. In L. Steels and J.A. Campbell (Eds.), *Progress in artificial intelligence*. Chichester, UK: Ellis Horwood Ltd, 17-29.
Schank, R.C., Collins, G.C., and Hunter, L.E. (in press). Transcending inductive category formation in learning. *The Behavioral and Brain Sciences*.
Smith, E.E., and Medin, D.L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
Smith, E.E., Shoben, E.J., and Rips, L.J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review, 81*, 214-241.
Sober, E. (1984). *Conceptual issues in evolutionary biology*. Section 7. Cambridge, MA: MIT Press.
Stebbins, R. (1957). Intra-specific sympatry in the lungless salamander, *Ensatina Eschscholtzi*. *Evolution, 11*, 265-270.
Wittgenstein, L. (1953). *Philosophical investigations*. New York: Macmillan.