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

Concepts Do More Than Categorize

Karen Olseth Solomon, Douglas L. Medin, and Elizabeth Lynch

Northwestern University

Address correspondence to:
Karen O. Solomon
Department of Psychology
Northwestern University
2029 Sheridan Road
Evanston, IL 60208
Telephone: (847) 467-6513
Fax: (847) 491-7859
k-solomon@nwu.edu

Abstract

Concepts underlie all higher-level cognitive processes. Until recently, the study of concepts has largely been the study of categorization. But categorization is only one conceptual function among several. We argue that concepts can not be understood sufficiently through the study of categorization, or any other function, in isolation, for two important reasons. First, concepts serve multiple functions which interact to affect conceptual structure and processing. Given the interrelated nature of these functions, one can not hope to understand conceptual structure or processing by studying any one function. Second, studying a single function in isolation encourages one to see cognitive processes which are distinct to each function, but discourages the discovery of processes that are common to multiple functions. For these two reasons, we suggest that concepts should instead be studied in the context of a system of interrelated functions.

Concepts are the building blocks of thought. How concepts are formed, used, and updated are therefore, central questions in cognitive science. The literature on concepts is vast (see Komatsu[1], Goldstone[2], Medin and Coley[3] for reviews and Nakamura, Taraaban, and Medin[4], van Mechelen et al.[5], Hirschfeld and Gelman[6], and Lamberts and Shanks[7] for recent edited volumes) and impossible to summarize in a short review. For this reason, we focus on a single point—namely, the idea that concepts have multiple functions which interact to affect conceptual structure and processing. Until recently, the study of concepts has largely been the study of categorization. But categorization is only one function among several. Concepts serve multiple functions, and, as we will see, these functions are not independent of one another; rather, they interact with and influence each other. We believe that these interactions undermine the popular strategy of studying categorization, or any other function, in isolation. Moreover, these interactions suggest that concepts should be studied in the context of a system of interrelated functions.

The remainder of this article is organized as follows. First, we outline some common functions of concepts. Then, we turn to recent research which demonstrates that these functions indeed do interact with one another. Finally, we conclude with an example of how an integrated approach can encourage the discovery of cognitive processes which extend across multiple functions..

Functions of Concepts.

A concept is notoriously hard to define. However, in this paper, we will refer to a *concept* as a mental representation that is used to meet a variety of cognitive functions. The most commonly studied function has been categorization, a process by which mental representations (or concepts) are used to classify entities. Recently, there has been a trend toward the study of two other conceptual functions, inference, and conceptual combination. But as we shall see, concepts are used to perform many cognitive tasks, and the boundaries between conceptual functions blur as conceptual functions interact and influence each other. Thus, in this section we will briefly illustrate a number of ways that concepts are used in everyday life, and in the following section discuss interrelations among these functions. Most researchers assume that conceptual representations include procedures for identifying whether an entity is a member of a category, a process often referred to as categorization. Categorization is not an end in itself, but rather it serves to connect old to new. Categorizing novel entities allows the cognitive system to bring relevant knowledge to bear in the service of understanding. Recognizing some unusual shape as a *toothbrush* allows one to understand its parts and their functions. A related function is inference. Knowledge of category membership supports predictions about behavior. For example, medical diagnostic categories allow physicians to predict what sorts of treatments will be effective. Concepts are also crucial in explanation and reasoning. Having categorized a young man as a *football fan*, one might be able to explain why he is walking down the street bare-chested with blue and yellow paint on his face bellowing out the Michigan rouser. Categories are also used to instantiate goals in planning[8]. For example, our football fan may create an ad hoc category of things to bring to a football game (e.g. beer, binoculars, seat cushion, and transistor radio).

Not only are new entities understood in terms of old, but new entities also modify and update concepts. That is, concepts support learning. This broad function is associated with a variety of questions such as how much weight to give new versus old information, when to set up subcategories, and so on. Moreover, given that instances can belong to multiple overlapping categories there is the issue of whether new information should modify all possible categories, a relevant subset, or only the category most associated with the new information. For example, if one learns that Oprah Winfrey is a multimillionaire, which of the many possible categories (e.g., African American, woman, talk show host, actress, etc.) should be updated with the information?

Concepts are also centrally involved in communication. The inter-personal aspect of concepts places strong constraints on categorization. Category membership is, in important respects, a social construction and whether something is a *bottle* or a *carton* may be a matter of convention[9,10] or context[11]. Finally, by combining concepts we can use a limited number of known concepts to create an unlimited number of new concepts. For example, we can use the concepts, *chocolate* and *rash* to create the new concept, *chocolate rash*, which presumably means a rash produced by eating chocolate.

These various functions place competing demands on how conceptual knowledge is organized. For example, ad hoc categories may encourage an organization scheme which is different from one that maximizes inference potential[12], but maximizing inferences as the sole consideration might lead to such narrow categories that novel entities may often fall into no pre-existing category. Hence, even a broad computational level analysis suggests that no single representation and processing scheme will maximize every conceptual function. If conceptual representations are a compromise between functions having somewhat different requirements, then one can't hope to understand representations by looking only at a single function (such as categorization). We turn now to some recent findings that reinforce the idea that concepts should be studied in the context of a system of interrelated functions.

Interplay Across Functions (Going Beyond Categorization).

As discussed above, people use concepts in a wide variety of ways. Yet, research on concepts has focused primarily on the process of categorization (see Medin and Coley[3] for a recent summary of this research). In this tradition, models of conceptual structure have addressed how people classify objects into categories. For instance, the prototype view assumes that objects are categorized by comparing them to summary representations or prototypes. Since a summary representation constitutes something of a mental average, a standard assumption of prototype models is that conceptual representations are modified whenever a new entity is categorized (exemplar models make the same assumption). A growing body of research, however, suggests that conceptual representations are also modified when they are used for other conceptual tasks besides categorization.

For example, consider the findings from a study recently reported by Ross[13]. In his experiment, participants diagnosed fictitious diseases according to sets of symptoms. After diagnosing a disease, participants decided on a treatment for the disease. Each treatment was based only on a subset of the possible symptoms of the disease. On a later categorization task, symptoms that were relevant to the treatment led to more accurate diagnoses of the disease than symptoms that were irrelevant to the treatment. Moreover, treatment relevant symptoms were also generated faster from the name of the disease, and judged to occur more frequently. Thus, treatment relevant symptoms became more important to the concept of the disease than treatment irrelevant symptoms. These results suggest that interacting with instances of a category can change the structure of the conceptual representation, which in turn affect other conceptual tasks.

Communicating about an entity also appears to affect conceptual structure and categorization. As one example, Markman and Makin[10] investigated the effect of communication on categorization. In their procedure, participants either built a LEGO model with another person or they built it alone. Participants who built the model with another person, and therefore were forced to communicate about the LEGO pieces, were more likely to later sort the LEGO pieces in a similar way than those subjects who built the model alone. Thus, communicating about a given entity can influence how it is later categorized. In a similar vein, Barsalou and Sewell[14] showed that a key component of communication, establishing a shared reference, had a large effect on standard phenomena such as typicality judgments. For example, when Emory University students were asked to take the perspective of a person in China their typicality judgments of examples of the category, *bird*, changed dramatically.

Barbara Malt[9] has shown that communication and linguistic convention are crucial to what people consider to be members of natural object categories. For example, people consider aspirin bottles, square hand lotion bottles, and baby bottles to all be members of the category *bottle* even though these members differ in form and function. Malt's research suggests a concept's structure may reflect a chaining process (see Lakoff[15] for a discussion of this), where entities are considered to be members of a category because of their connection to other category members . For example, even though a square hand lotion bottle has a different shape and function from most prototypical bottles, it may be considered a bottle because of its connection to other hand lotion bottles. Category members may also be linked through history or convention. For example, even though a Model T and a new car have very different shapes, they may both be thought of and called *automobiles* because they are linked historically. A child's *juice box*, however, may be considered a *bo*x rather than a *carton* simply because of the name the manufacturer gave it (see Malt et al.[16] for an argument that the categorization processes used for naming and the categorization processes used for object recognition may be distinct).

We have just seen that using a concept for a specific function alters the conceptual structure, which in turn, changes how the concept is later used for other functions. In each of the studies described above, we have assumed that participants used one conceptual structure across multiple functions. But sometimes one conceptual structure may not be sufficient to meet the competing demands of multiple functions. For example, consider how a particular group of tree experts (landscapers) organize and use their knowledge of trees. In a recent paper, Medin, Lynch, and Coley[17] found that landscapers used one conceptual structure for categorization, and another for reasoning. Specifically, the landscapers used a conceptual organization that reflected their goals as landscapers (e.g., shade trees, weed trees, and ornamental trees) to sort trees into categories (a categorization task), whereas they used a taxonomic organization to make biological inferences from one category to another (a reasoning task). The landscapers probably relied on a taxonomic organization for the reasoning task because they thought that their goal-centered organization would not support biological inferences (see Box 1 for a further discussion of this goal-centered organization). Their responses may have been different if the landscapers had been asked about inferences that were based on their goals as landscapers (e.g., if a Sugar Maple shades a particular house, do you think a Red Oak would also shade this house?). In any case, these results show that the landscapers had access to both conceptual organizations, and that they relied on the structure which could best serve the conceptual function at hand. (See Box 2 for a more in depth discussion of the nature of expertise or lack thereof.)

In the study described above, Medin et al. examined a domain of knowledge that was regularly used by their participants to serve a wide variety of conceptual functions (e.g., planning, communication, inference, etc.). If these different functions affect structure then by looking at the functions together it might be possible to make better predictions about conceptual structure. A study by Ross and Murphy[18] supports this hypothesis. They examined the complex domain of food because most people regularly categorize, communicate, make plans, solve problems, and make inferences about food. They found that indeed people do have multiple conceptual organizations for different kinds of food. Their participants classified food according to standard taxonomic categories (e.g., fruits, vegetable, and meats), but they also classified food according to situational categories (e.g., snacks, breakfast foods, and health foods) that cut across standard taxonomies. The situational categories were typically used to meet goals associated with a particular context. For example, someone might use her category of snack food to decide what to eat during an afternoon break. Unlike many goal-derived categories, however, these situation categories seem to be well-established in memory. Interestingly, participants used different types of categories to make different types of inferences; taxonomic categories were used to make inferences about biochemical properties (e.g., gives one energy, is sweet), whereas situational categories were used to make inferences about situational properties (e.g., costs less, is eaten in the morning). (See also Heit and Rubinstein[19]).

The two studies described above suggest that multiple conceptual organizations are common in knowledge rich domains. For this reason, conceptual models should be modified to include theories of how multiple structures are coordinated and updated. Such theories should address whether multiple organizations are available simultaneously or whether they compete with one another. Although this issue has not been addressed by cognitive psychologists, research from the area of social cognition suggests that multiple structures are not simultaneously available and may even inhibit one another. For example, Macrae, Bodenhausen, and Milne[20] found that after viewing an Asian woman eating noodles, stereotypical traits of Asians were activated and stereotypical traits of women were inhibited. In contrast, after viewing the same Asian woman putting on makeup, stereotypical traits of women were activated, and stereotypical traits of Asians were inhibited. These results obviously need to be replicated but they certainly are intriguing.

The issue of whether multiple category structures are simultaneously activated has important implications for how categories are updated. If only one category is activated at a time, then perhaps only the activated category gets updated with new information. For example, in Macrae et al.'s study, if someone observes an Asian woman putting on makeup, he may add this information to his category of women but not his category of Asians. Similarly, if the same person now observes the Asian woman eating noodles, he may now only update his category of Asians and not his category of women. But this kind of updating may reinforce stereotypical beliefs (e.g., women put on make-up, Asians are eat noodles), and suppress other possible beliefs (e.g, Asians put on make-up, women eat noodles). Hence, by biasing the information that is added to a particular category, multiple conceptual organizations may reinforce and even encourage stereotypes. In any event we are just beginning to understand how multiple category schemes combine, compete, and support reasoning.

Extending Processes Across Functions

As discussed above, conceptual functions have often been studied in isolation. One repercussion of this approach is that distinct conceptual functions are usually associated with distinct conceptual processes. Thus, even though researchers generally assume that a single conceptual representation can support multiple functions, they also assume that the processes that operate on this representation are different for each function. For instance, the cognitive processes used to categorize an object are assumed to be different than the processes used to make inferences about this object.

Taking a multi-functional approach to concepts, however, encourages one to see processes that are common to multiple functions. Consider Ed Wisniewski's recent work on conceptual combination[21,22]. By investigating how people interpret novel conceptual combinations, Wisniewski has discovered two distinct types of processing, *integration* and *comparison,* that all appear central to other conceptual functions. Broadly defined, integration refers to the process of finding a relation which meaningfully links two concepts together, whereas comparison refers to the process of aligning two concepts along common dimensions and then analyzing the similarities and differences between the concepts with respect to these dimensions.

According to Wisniewski[21], people use a comparison process to understand noun-noun conceptual combinations when the two nouns are easily aligned. For example, Wisniewski argues that people interpret noun-noun combinations with alignable nouns, such as *zebra horse*, by first aligning the two concepts, and then extending a salient property of the modifying noun to the head noun. For example, *zebra horse* might be interpreted as a horse with stripes, and a *barn house* might be interpreted as a house with a barn shaped roof. In contrast, when the two concepts are not easily aligned an integration process is more likely to be evoked. For example, people are likely to interpret a conceptual combination such as *stereo money*, by finding a relation which meaningfully integrates the two concepts. For example, *stereo money* might be money that was used to buy a stereo.

Interestingly, comparison and integration processes extend beyond conceptual combination to other conceptual tasks. Thus, it is not the type of task, or function that determines whether integration or comparison is used, but rather the relationship among the stimuli. Across different tasks, alignable stimuli (e.g., zebra, horse) tend to elicit comparison processes while non-alignable stimuli (e.g., zebra, grass) tend to elicit integration processes.

For instance, past theories have assumed that people assess the similarity of two concepts through a featural comparison process only. Integration, however, is also important. Wisniewski and Bassok[23], report a study in which participants judged both taxonomically-related concepts (e.g., cow, horse) and functionally-related concepts (e.g., cow, milk) to be similar. A standard theory of similarity would predict that taxonomically related stimuli would be judged to be similar, but the similarity of functionally related concepts is a finding not predicted by most similarity theories. Participants' justifications for their judgments suggested that they were assessing the similarity of alignable concepts (e.g., cow, horse) by comparing them, but they were assessing the similarity of functionally-related concepts (e.g., cow, milk) by finding a relation that integrated the two concepts. Thus, participants might have judged *cow* and *milk* to be similar because they could find a relation which meaningfully integrated the two concepts (e.g., a cow produces milk).

Recent evidence suggests that comparison and integration are also central to reasoning. Most models of inductive reasoning suggest that people draw inferences across categories by first comparing the two categories, and if they are judged to be sufficiently similar, extending a property from one category to another. For example, a person might decide that a property of *gerbils* is also true of *hamsters* because *hamsters* and *gerbils* are quite similar. Interestingly, integration also appears central to inductive reasoning. Lin[24] found that for certain examples, people were more likely to extend a property to a functionally-related category than to a taxonomically-related category. For example, participants who were told that *plumbers* had a certain bacteria were much more likely to infer that *pipes* also had the bacteria than *carpenters* had the bacteria, probably because they believed that the bacteria could be transmitted by physical contact. In other words, participants used the functional relation between *pipes* and *plumbers* to infer the category that was most likely to have the bacteria. Thus, because the integration process is clearly used in some cases, current theories of induction need to incorporate it. d the integration process and generally assume that inferences are made through a comparison process only. (See also Box 3).

Sloman[25] provides another counterexample to the assumption that the similarity of categories is the main determinant of when a property can be extended from one category to another. Sloman shows that the relation that links the initial category and property can determine whether someone will extend the property to a new category. For example, being told that ex-cons are hired as bodyguards may increase the likelihood that war veterans are also hired as bodyguards, however knowing that many ex-cons are unemployed may not increase the likelihood that many war veterans are unemployed. In short, these examples suggest that theories of inductive reasoning need to be modified to include the process of integration.

Conclusion.

We have discussed several lines of research which show that concepts can not be studied through categorization alone because conceptual functions often interact and influence one another. Some of the research showed that conceptual representations are often a compromise between the competing demands of multiple functions. Other research showed that when these competing demands are too great to be met by one conceptual representation, alternative representations emerge. Finally, while differences in conceptual functions generally correspond to differences in processing, an integrated approach can reveal types of processes which extend across functional boundaries. We hope that this review has demonstrated that a multi-functional approach to concepts is both promising and necessary, and while it may pose challenges for researchers, work conducted from this perspective is yielding intriguing findings with clear implications for models of conceptual behavior.

BOX 1. Knowledge affects internal structure of categories

All members of a category are not equal. The internal structure of categories is graded, such that some members are better representatives of the category than others. Graded structure plays a role in many cognitive tasks, like category verification, category listing, category learning, and induction. A popular way of explaining the representativeness of category members is by appeal to overall similarity. That is, a category member is more representative of a category to the extent that it is similar to other category members. Thus, it has long been assumed that the most representative category member is the one that is most similar overall to the remaining category members.

This explanation of representativeness fits well with the idea that the function of categories is to capture the overall similarity structure of the environment. We know, however, that the internal structure of categories is not always organized around overall similarity. Barsalou[a] showed that while the internal structure of taxonomic categories was organized around overall similarity, the internal structure of goal-derived categories like clothes to wear in the snow, was determined by ideals. The best example of this category (down jacket) is not the exemplar that is most like other category members; rather, it is the exemplar with the most extreme value on the goal-related dimension. The better the exemplar serves the goal of the category, the more typical is the exemplar of the category. Trait categories have also been shown to be organized around ideals[b,c,d]

So far the explanation seems to be that what determines the graded structure of a category depends upon the kind of category in question. Perhaps the graded structure of taxonomic categories is determined by overall similarity (central tendency) and the graded structure of goal-derived categories is organized around ideals. This explanation makes good sense, especially given that goal derived categories do not have coherent overall similarity structure. However, recent evidence suggests that internal structure is not as tightly linked to category type as we might think.

Lynch, Coley, and Medin[e] found that the internal structure of the category tree was organized around the ideals of height and weediness for a group tree experts. This study is the first to demonstrate that ideals can serve as the primary determinants of the graded structure of a natural kind taxonomic category. The Lynch, et al. results indicate that for experts, ideals are consistently the strongest predictor of GOE ratings for trees. That is, the most representative trees--those which are the best examples of the category tree--are not trees of average height, but rather trees of extraordinary height. Likewise, weediness served as a negative ideal for the experts. The more often a tree is considered a weed, the less representative that tree is of the category. Thus, for the experts in our sample, tree appears to be structured around positive and negative ideal dimensions rather than the average of several relevant dimensions. Furthermore, undergraduates with little experience with trees did not use height as a criterion for judging representativeness, but rather representativeness was organized around familiarity.

Atran[f] reports a similar finding. Atran found that Itzaj Maya living in the rainforest of Guatemala and University of Michigan undergraduates use a similar metric to judge the similarity of birds (in a free sorting task). However, they use different criteria to rate the representativeness of the birds. The University of Michigan undergraduates determine representativeness on the basis of overall similarity. Thus, the birds that are most similar overall to other birds (the passerines, or small songbirds) are also the most representative. For the Maya, birds are representative to the extent that they are culturally salient. For the Maya the ‘truest’ bird is the wild turkey, because it has tasty meat and is beautiful.

Rather than being determined solely on the basis of the content of the category (category type), these recent studies suggest that internal structure may be affected in large part by experience with a category. In the case of tree, the ideal of weediness appears to be related to the goals of the tree experts. The factor of meat tastiness in the case of the wild turkey also relates to goals. Lynch et al. do not think that height is a goal-related dimension. Perhaps it is related to aesthetics as in the case of the beauty of the wild turkey. We can not be sure what about experience alters the internal structure of these conceptual representations. Nonetheless, the two examples are fascinating exceptions to the generally accepted rule that overall similarity determines the internal structure of taxonomic categories.

BOX 2. A case of disexpertise?

It is difficult to escape defining expertise in relative terms---an expert is someone who has knowledge and/or skills that others lack. Suppose that a group of experts succeeded in passing their knowledge and skills on to the rest of us. In that event would we now say that everyone was an expert or that the experts had ceased to exist? Or conversely suppose that knowledge deteriorated across generations such that the most skilled of the current generation would have been only average in an earlier era. Would such people be experts or should the average person of the current generation be seen as manifesting a deficient set of skills?

The above questions are not entirely hypothetical. First a little more backgound. Many entities belong to categories that are hierarchically organized. For example, a fish could be categorized as a Rainbow Trout, a trout, a fish, a vertebrate, an animal, a living thing and so on. There is strong evidence from both cognitive psychology (Rosch et al.[a]) and anthropology (Berlin, Breedlove, and Raven[b]) pointing to the fact that one of these levels, referred to by Rosch et al. as the basic level, is particularly salient or psychologically privileged. The basic level is the level people prefer to use in naming, first learned by children, and the level at which people can categorize most rapidly. In the domain of plants and animals Berlin[c] refers to basic level categories as "beacons on the landscape of biological reality."

Although cognitive psychology and anthropology have converged on this notion of a privileged level the surprising thing is that they don't agree on what level is basic! In the case of biological categories Rosch et al. found that the level corresponding to bird, fish, and tree was basic whereas the level most salient according to Berlin et al. is more specific, corresponding to robin, trout, and oak. Why this difference? One possibility (see Coley , Medin, and Atran[d] for others) is that the populations studied by Berlin et al. (people in traditional, subsistence cultures) were more expert than the population studied by Rosch et al. (Berkeley undergraduates) and that the basic level changes as a function of expertise. There is some suggestive evidence that the basic level does become more specific with expertise (e.g. Tanaka and Taylor[e], Johnson and Mervis[f]). So the difference between Berlin et al. and Rosch et al. may be attributable to the average Maya or Jivaroo farmer being an expert relative to the average American college student.

An interesting converging observation supports the hypothesis that technologically-oriented cultures have undergone a decline or devolution in biological knowledge. Wolf, Medin, and Pankratz[g] examined the large data base of written material from the 16th to the 20th century contained in the Oxford English Dictionary (available on-line) for trends in the frequency and specificity of references to biological kinds. Specifically, they looked at tree terms at different levels of organization (e.g. tree, oak, white oak). After attending to a variety of potential artifacts (see their paper for details), they uncovered two clear trends. First, there was a reliable increase in the use of tree terms between the 6th and 19th centuries followed by a precipitous decline after the 19th century. Second, the decline in use of tree terms was much steeper for the generic level (e.g. oak) than for the more abstract term, tree. These trends are consistent with the idea that people from traditional societies are experts relative to American college students primarily because of a deterioration (historically speaking) of American college students' knowledge of biological kinds. This conclusion underlines the uncomfortable quality of relative definitions. But is it possible to develop an absolute definition of expertise?

Box 3. Use of categories in reasoning.

Suppose that there is some disease that we know Rats and Pocket Mice get and another disease that Tapirs and Squirrels get. Which disease is more likely to affect all mammals? Even if you don't have a good idea what a Tapir looks like you would probably guess that the disease that effects Tapirs and Squirrels is more likely to affect all mammals. Your rationale might be that Rats and Pocket Mice are pretty similar and that their disease might be specific to rodents. In reaching this decision you would be following a proscription from taxonomic analysis (e.g. Warburton[a]) that conclusions based on diverse sources of evidence are more likely to hold than conclusions based on less (taxonomically) diverse sources of evidence. On hypothetical reasoning problems like the one above, Lopez et al.[b] found that University of Michigan undergraduates overwhelmingly picked the more diverse pair of premises (e.g. Tapir, Squirrel rather than Rat, Pocket Mouse). Although this illustration seems fairly obvious, it shows that categories may serve the important role of guiding reasoning. You may be surprised to find that diversity-based reasoning is far from universal. Lopez et al. gave this same task to Itzaj Maya of Peten, Guatemala (who are mainly subsistence farmers). Although the organization of the category mammal is similar across cultures and departs from scientific taxonomy in similar ways the Maya participants either showed no diversity-based responding or even below chance diversity responses.

The absence of diversity-based responding among the Itzaj Maya is a puzzle. Is there something about how they interact with mammals or use the category, mammal, that leads them to ignore taxonomic relatedness? Is the category too abstract to be meaningful? Do people in traditional societies have difficulty with hypothetical questions? Some of the Itzaj justifications for their answers provide some hints. One person who selected the Rat, Pocket Mouse pair over the Tapir, Squirrel pair reasoned as follows: “The only way that a Tapir and Squirrel could catch the same disease would be if a bat bit them and, in that case, they would not spread the disease. The Rat and the Pocket Mouse are companions and disease could spread between them and then affect other mammals.” Another person selected the Tapir, Squirrel pair and explained his choice in terms of diversity of range and habitat, not taxonomic diversity. That is, the Itzaj appeared to justifying their choices in terms of ecological or causal reasoning.

Some further light on diversity-based reasoning comes from studies of category-based reasoning among American tree experts using arguments involving tree diseases[c]. Although the pattern of results varied somewhat as a function of type of expertise (landscaper, taxonomist, parks maintenance), parks maintenance personnel showed below chance diversity responding much like the Itzaj. All three types of experts frequently gave ecological/causal justifications. In short, it appears that taxonomic reasoning is not the only game in town and that knowledgeable people (Maya, tree experts) often use ecological reasoning. And undergraduates might too--suppose that Grass has some property x, is it more likely that Oak Trees also have this property or that Cows have this property. You might answer Cows in the assumption that the property might be something that Cows acquire by eating Grass.

Box 4. Outstanding questions.

When is one conceptual structure sufficient for multiple functions, and when do alternative organizations arise?

How are multiple conceptual structures accessed and updated?

When do more general processing strategies, such as integration and comparison, correspond with or diverge from task-specific processing?

Which conceptual functions are likely to place competing demands on conceptual structure and why?

Are there principles of categorization which hold across different cultures and different levels of expertise?

References

1. Komatsu, L. K. (1992) Recent views of conceptual structure. *Psychological Bulletin, 112*, 500-526

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

3. Medin, D.L., & Coley, J.D. (1998) Concepts and categorization. In Hochberg, J. & Cutting, J.E. (Eds.), *Handbook of perception and Cognition, Perception and cognition at century's end: History, philosophy, theory*. Academic Press

4. Nakamura, G. V., Taraban, R., & Medin, D. L. (Eds.). (1993) *The Psychology of Learning and Motivation: Categorization by Humans and Machines*, Vol. 28. San Diego, CA: Academic Press

5. VanMechelen, I. et al. (Eds.) (1993) *Categories and Concepts: Theoretical Views and Inductive Data Analysis*. London: Academic Press

6. Hirschfeld, L. A., & Gelman, S. A. (Eds.) (1994) *Mapping the Mind: Domain specificity in cognition and culture*. New York: Cambridge University Press

7. Lamberts, K. & Shanks. D. (1998) *Knowledge, Concepts, and Categories*. Cambridge: MIT Press

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

9. Malt, B. C. (1993) Concept Structure and Category Boundaries. In Nakamura, G.V., Medin, D.L., & Taraban, R. (Eds.). *Categorization by Humans and Machines*. Academic Press

10. Markman, A.B. & Makin, V.S. (in press) Referential communication and category acquisition. *Journal of Experimental Psychology: General*.

11. Brown, R. (1958) *Words and things: An introduction to language*. New York: MacMillan

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

13. Ross, B. (1997) The use of categories affects categorization. *J. Mem. Lang., 37*, 240-267

14. Barsalou, L.W., & Sewell, D.R. (1984) Constructing representations of categories from different points of view. Emory Cognition Project Technical

Report #2, Emory University

15. Lakoff, G. (1987) *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago: University of Chicago Press

16. Malt, B.C. et al. (in press) Knowing versus Naming: Similarity and the Linguistic Categorization of Artifacts. *J. Mem. Lang*.

17. Medin, D. L. et al. (1997) Categorization and reasoning among tree experts: Do all roads lead to Rome? *Cognit. Psychol., 32*, 49-96

18. Ross, B.H. & Murphy, G.L. (in press) Food for thought: Cross-classification and category organization in a complex real-world domain. *Cognit. Psychol.*.

19. Heit, E. & Rubinstein, J. (1994) Similarity and property effects in inductive reasoning. *J. Exper. Psychol.: LMC*, 20, 411-422

20. Macrae, C. N., Bodenhausen, G. V. & Milne, A. B. (1995) The dissection of selection in person perception: Inhibitory processes in social stereotyping. *Journal of Personality and Social Psychology*

21. Wisniewski, E. J. (1996) Construal and similarity in conceptual combination. *J. Mem. Lang., 35*, 1-20

22. Wisniewski, E. J. (1997) When concepts combine. *Psychonomic Bullentin & Review, 4*, 167-183

23. Wisniewski, E.J. & Bassok, M. (1996) On putting milk in coffee: The effect of thematic relations on similarity judgments. *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society*. (pp. 464-468). Hillsdale, NJ: Erlbaum

24. Lin, E.L. (1996) Thematic relations in adultís concepts and categorization. Unpublished doctoral dissertation, University of Illinois

25. Sloman, S. (1994) When explanations compete: the role of explanatory coherence on judgements of likelihood. *Cognition, 52*, 1-21

References for Box 1.

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

b. Bornkeneau, P. (1990) Traits as ideal-based and goal-derived social categories. *Journal of Personality and Social Psychology, 58*, 381-396

c. Read, S.J., Jones, D.K, & Miller, L.C. (1990) Traits as goal-based categories: The importance of goals in the coherence of dispositional categories. *Journal of Personality and Social Psychology, 58*, 048-1061

d. Chaplin, W.G., John, O.P., & Goldberg, L.R. (1988) Conceptions of states and traits: Dimensional attributes with ideals as prototypes. *Journal of Personality and Social Psychology, 54*, 541-557

e. Lynch, E., Coley, J. D., Medin, D. L. (in press) Tall is typical: Central tendency, ideal dimensions and graded category structure among tree experts and novices. *Memory & Cognition*.

f. Atran, S. (In press) Itzaj Maya folkbiological taxonomy. In D. L. Medin & S. Atran (Eds.), *Folkbiology*. Cambridge, MA: MIT Press

References for Box 2.

a. Rosch, E. et al. (1976) Basic objects in natural categories. *Cognit. Psychol., 8*, 573-605

b. Berlin, B., Breedlove, D., & Raven, P. (1973) General principles of classification and nomenclature in folk biology. *American Anthropologist, 74*, 214-242

c. Berlin, B. (1992) *Ethnobiological classification: Principles of categorization of plants and animals in traditional societies*. Princeton, NJ: Princeton University Press

d. Coley, J. D., Medin, D. L., Atran, S. (1997) Does rank have its privilege? Inductive inferences within folkbiological taxonomies. *Cognition, 64*, 73-112

e. Tanaka, J. W., & Taylor, M. (1991) Object categories and expertise: Is the basic level in the eye of the beholder? *Cognit. Psychol., 23*, 457-482

f. Johnson, K. E., & Mervis, C. B. (1998) Impact of intuitive theories on feature recruitment thoughout the continuum of expertise. *Memory & Cognition, 26*, 382-401

g. Wolf, P., Medin, D., & Pankratz, C. (in press) Evolution and devolution of folkbiological knowledge. *Cognition*.

References for Box 3.

a. Warburton, F. (1967) The purposes of classification. *Systematic Zoology, 16*, 241-245

b. Lopez, A. et al. (1997) The tree of life: Universals of folkbiological taxonomies and inductions. *Cognit. Psychol., 32*, 251-295

c. Coley, J. D. et al. (in press) Inductive Reasoning in Folkbiological Thought. In D. L. Medin & S. Atran (Eds.) *Folkbiology*. Cambridge, MA: MIT Press