The Basic Level and Privilege in Relation to Goals, Theories, and Similarity

Douglas L. Medin, Elizabeth B. Lynch, and John D. Coley

Northwestern University
Department of Psychology
2029 Sheridan Road
Evanston, IL 60208-2710

Scott Atran

CNRS-CREA, Ecole Polytechnique
and University of Michigan

Abstract

It has been argued that some categories are *privileged* in that the correlational structure of features in the environment creates natural chunks that are compelling to learners. If so, then differences in goals should not affect the organization and use of such so-called privileged categories. This paper reports research where the amount and type of expertise is varied for the same domain. Goals exert a major effect on category organization as well as on category-based induction. Even for the case where some categories seemed to be relatively privileged, the basis for this privilege may derive from linguistic cues rather than strictly empirical structure.

Introduction

What are the fundamental principles of category organization and why are some categories more coherent, easy to learn, and useful than other categories? These issues continue to be central to our understanding of concepts and conceptual behavior. In the present paper we focus on two broad questions: 1. How do similarity and theory structure categories? and 2. What makes some categories privileged relative to others? We begin by placing our work in the context of earlier research, next describe some new psychological experiments, and then bring out the implications of our findings for similarity theories and notions of privilege.

A. Similarity versus Theory. Questions about the respective roles of similarity versus theory or explanation serve to organize much of the ongoing work on categorization (Komatsu, 1992; Medin and Heit, in press; Murphy, 1991, 1993a, b; Smith and Heise, 1992 for reviews and Burns, 1992; Fisher, Pazzani, and Langley, 1991; VanMechelen, Hampton, Michalski, and Theuns, 1993; Nakamura, Taraban, and Medin, 1993 for relevant edited volumes). To make similarity meaningful, there must be constraints on what features or aspects enter into a comparison and the weights assigned to them. We are seeing both more sophisticated models of similarity and models of categorization capable of addressing the flexible and dynamic aspects of similarity (e.g. Anderson, 1991; Barsalou, 1991, 1992, 1993; Billman, 1992; Clapper and Bower, 1994; Estes, 1994; Gentner and Markman, 1994; Medin, Goldstone, and Gentner, 1993; Myers, Lohnmeller, and Well, 1994; Nosofsky, Gluck, Palmeri, McKinley, and Glanthier, 1994).

Meanwhile, there has been corresponding progress in understanding the role of knowledge and theories in organizing concepts (Gelman and Wellman, 1991; Keil, 1989; Malt, 1995; Michalski, 1989; Murphy, 1993a,b; Pazzani, 1991; Rips and Collins, 1993; Waldman and Holyoak, 1992; Walker, 1992, and edited volumes by Hirschfeld and Gelman, 1994; Michalski and Tecuci, 1994). Although one might view knowledge-driven learning as an alternative to similarity-based approaches, the scholarly consensus has been in the direction of models which combine similarity-based and explanation-based categorization (to give but a few examples, see Billman and Knutson, 1996; Choi, McDaniel, and Busemeyer, 1993; Fisher, Pazzani, and Langley, 1991; Fisher, and Yoo, 1993; Heit, 1993, in press; Mooney, 1993; Wisniewski and Medin, 1994, a,b).

Although empirical work on similarity-based and knowledge-driven learning has not restricted itself to artificial materials (e.g. Brooks, Norman, and Allen, 1991) neither has it made systematic contact with the kinds of categories associated with theories which emphasize domain-specific categorization principles (e.g. Hirschfeld and Gelman, 1994). The focus of the present project is on one such domain, biological categorization. We shall not argue that categorization processes are necessarily domain-specific, but rather we focus on the point that there are a number of research issues that become more coherent when analyzed within the context of particular domains. Biological reasoning is of particular interest because, in this domain, theorists have laid out distinct positions on the role of mind and world in categorization. Furthermore, the domain of biological kinds seems particularly attractive from the perspective of similarity-based categorization -- surely our perceptual system is an adaptation to the natural world and if similarity models are going to succeed anywhere, it should be here.

B. Mind and world. It is straightforward to discern points on an underlying dimension that represents the extent to which goals and belief systems versus world structure are said to influence organization and reasoning with biological categories. At one end of the continuum are those who view taxonomies as direct reflections of the structure of the natural world. This view is associated primarily, but not exclusively, with anthropology. On this account folkbiological categories are essentially given by encounters with the environment and virtually "crying out to be named" (Berlin, 1992). The idea is that people are innately prepared by their perceptual apparatus to form categories and little conceptual effort is needed to produce folkbiological classification because in Boster's (1991) terms, "the source of structure in biological similarity judgments is in the world, not in the brain."

A diametrically opposite position is that goals and beliefs determine categorization and reasoning. This view argues that empirical observations underdetermine categories and that the very notion of similarity is theory-laden (Murphy and Medin, 1985; Ellen, 1993). In the words of Gelman and Coley (1991), "to calculate similarity, it is necessary to have a theory of which features are important and how they should be weighted." From this perspective, one might expect differences in biological categorization and reasoning to the extent that associated goals, theories, and belief systems differ.

Although the above two views represent end points on a continuum, they may not be so different as first appears. To begin with, advocates of the "structure in the world" position would concede that structure is relative to a perceptual (and conceptual) apparatus capable of apprehending it. In the same vein, human goals, theories, and interests may have universal components that are not readily separated from world structure. For example, Anderson's (1990) rational model of categorization is premised on a (universal) goal of understanding the world and making predictions about it. Goals and theories must make contact with the world and do not have the luxury of ignoring structure.

In the research motivating the present paper, goals and interests are varied by virtue of the occupational demands and activities associated with different types of expertise for a common biological category, trees. Our experts include taxonomists, landscapers, and parks personnel. Before describing these differences in greater detail, we turn to another dimension relevant to the contributions of mind and world to categories, hierarchical level.

C. The basic level and other privileged categories. Another perspective on the integration of similarity and theory is the view that their relative contributions vary as a function of types or levels of categorization. For example, in a now classic paper Rosch, Mervis, and their associates (Rosch, Mervis, Gray, Johnson, and Boyes-Braem, 1976) argued that "in the real world information-rich bundles of perceptual and functional attributes occur that form natural discontinuities and that basic cuts in categorization are made at these discontinuities" (p. 385). That is, the correlational structure of features of entities in the environment creates natural clusters or <u>basic level</u> categories which play a central role in many cognitive processes associated with categorization. Basic level categories such as <u>Chair</u>, <u>Hammer</u>, and <u>Dog</u> may be contrasted with more general superordinate categories (<u>Furniture</u>, <u>Tool</u>, and <u>Animal</u>) and more specific subordinate concepts (e.g. <u>Recliner</u>, <u>Ballpeen hammer</u>, <u>Labrador retriever</u>). Rosch et al (1976) evaluated a number of criteria associated with category use and they showed a remarkable convergence on the basic level of categorization. Basic level categories are the most inclusive categories that <u>1</u> possess numerous common attributes <u>2</u> people interact with using similar motor programs <u>3</u> have similar shapes and <u>4</u> can be identified from averaged shapes of members of the class. Furthermore, basic level categories are preferred in adult naming, first learned by children, and most easily identified (relative to subordinate and superordinate categories). It is as if basic level categories are out in the environment waiting for people to apprehend them.

Although researchers have not tended to employ the variety of converging measures educed by Rosch, the basic level superiority has proven to be quite robust (e.g. see Lassaline, Wisniewski, and Medin, 1992 for a review). There is however, at least one serious puzzle, a puzzle that has recently motivated a significant portion of our research. Rosch et al. (1976) made initial guesses about which level in a hierarchy would prove to be basic and these guesses were generally correct. For biological taxonomies, however, the initial guesses were based on anthropological observations (Berlin, Breedlove, and Raven, 1966, 1973) and in each case these guesses turned out to be wrong. For example, instead of <u>Oak</u>, <u>Trout</u>, and <u>Eagle</u> being basic, <u>Tree</u>, <u>Fish</u>, and <u>Bird</u> met the Rosch et al criteria. The puzzle concerns why the anthropological and psychological criteria did not converge. Berlin (1992) argues that there is one level of categorization of plants and animals that is especially salient but he believes that this privileged level corresponds to the genus level in scientific taxonomy (e.g. Oak). The level that Rosch et al found to be basic corresponds to a much higher level in a scientific taxonomy, <u>Class</u> or even <u>Division</u>. Why the difference?

1. Expertise? One way of reconciling Rosch with Berlin is to suggest that the basic level changes with expertise (Rosch et al., 1976 discussed this possibility; see also Mervis, 1987). People living in urban, high technology settings may be relative novices with respect to living kinds. Surprisingly, there has been only a handful of studies of expertise and categorization (Chi, 1983; Chi, Hutchinson, and Robin, 1989; Chi, Feltovich, and Glaser, 1981; Gobbo and Chi, 1986; Johnson and Mervis, 1994; see also Honeck, Firment, and Case, 1987) and even less work on changes in the basic level with expertise (Tanaka and Taylor, 1991; Johnson, 1992; Mervis, Johnson, and Scott, 1993; Palmer, Jones, Hennessy, Unze, and Dick, 1989). These latter studies reveal changes in naming practices and feature listings with expertise as well as some evidence that subordinate level categorization may come to be as rapid as basic level categorization.

2. Multiple privileged levels? A complementary hypothesis derives from the original Berlin et al folkbiological data (see also Berlin, 1992). Berlin suggested that five or even six levels of taxonomic structure were privileged in the sense of being stable across cultures: 1 unique beginner (plants versus animals), 2 life form (e.g. tree, grass, bush, bird, fish), 3 intermediate (e.g. hooved mammals), 4 folk-generic (usually corresponding to genus), 5 folk-specific (corresponding to species), and 6 varietal (e.g. kinds of corn). Not all of these levels are equally salient. For example, intermediates are less common and the folk-specific and varietal levels are much more likely to be seen in agricultural societies than in foraging cultures. Berlin's analysis raises the possibility that there may be multiple privileged levels in a taxonomic hierarchy. Note that the biological category level that Rosch et al (1976) argued was basic (e.g. bird, fish) corresponds to Berlin's life form level (see also Mandler, Bauer, and McDonough, 1991). It is possible that both the life form and folk-generic levels are psychologically salient but that the subjects tested by Rosch et al were unfamiliar with the particular genera used. The bird watchers studied by Tanaka and Taylor (1991) would have had the requisite experience and their rapid subordinate level categorization is consistent with the folk-generic level being salient (for their dog experts, the subordinate level would correspond to folk-specific or varietal).

Figure 1 shows one way of conceptualizing multiple privileged levels of categorization. There are many candidate metrics for ease of categorization (e.g. Corter and Gluck, 1992; Jones, 1983; Murphy, 1982) but suppose, for illustrative purposes , that categories are salient if the ratio of within category similarity to between category similarity exceeds some criterion. Then (A,B) would be salient because of high within category similarity, (A,B,C,D,E,F,G,H) salient because of low between category similarity, and (A,B,C,D) would be salient for intermediate reasons. (This analysis assumes that similarity is fixed; it is certainly possible that experience and expertise leads to increased sensitivity or sensitivity to new features or dimensions). All three categories should be more salient than the category (A,B,C,D,E,F).

3. Comparability of measures? Rosch et al (1976) pinpointed the basic level with a range of converging measures not seen before or since. Therefore, apparent cross-cultural differences or changes with expertise may reflect changes in some measures but not others. The pattern of convergences and divergences as a function of knowledge may prove to be particularly informative with respect to categorization theories.

D. Goal and theories. Of course, interest in expertise goes considerably beyond possible changes within the basic level. Knowledge and expertise may lead to changes not just in category salience but also in category organization and reasoning with categories (Boster and Johnson, 1989; Carey, 1985; Carey and Spelke, 1994; Johnson, Mervis, and Boster, 1992; Keil, 1994). There is at least suggestive evidence that goals and functions associated with expertise may influence category organization. For example, Boster and Johnson (1989) found that expert fishermen judged similarities among fish on both functional and morphological criteria, while novices used morphological criteria alone. As a result, novice judgments corresponded better to scientific taxonomy than the judgments of experts. Experts were also more variable in their judgments, raising the possibility that they had multiple classification schemes. Boster and Johnson (1989) did not look at hierarchical levels. One implication of the basic level (as given by the environment) perspective is that functions and theories may influence categorization more at the superordinate and subordinate level than at the basic level.

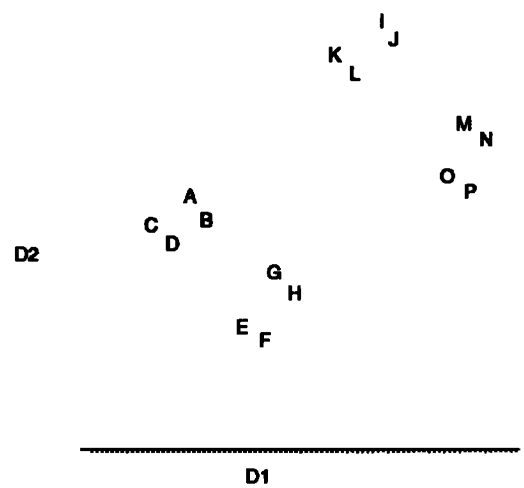

Figure 1. Spatial representation of similarity as the inverse of distance.

E. Overview of present research. Recently, we have been analyzing levels of categorization, category organization, and category-based reasoning as a function of expertise. In one main project, we have been assessing the categorization of and reasoning about trees by taxonomists, landscapers, and parks personnel. As may be obvious from these occupations, a shared ability to identify trees may be coupled with distinctly different goals and interests. Parks personnel spend considerable time planting new trees and culling and pruning diseased, damaged and aging trees. Landscapers focus on aesthetic functions and are only secondarily concerned with maintenance. Taxonomists are involved in a variety of actions ranging from systematics to education and consulting. A key question is the extent to which these subgroups differ in their conceptual behavior. Finally, we include one cross-cultural measure of basicness, comparing Northwestern University undergraduates with a Maya Indian group in Guatemala. In addition to a host of other differences, the Maya are expert folkbiologists, especially in comparison with undergraduates.

Studies of Biological categorization and Reasoning

Our primary aim is to examine category organization as a function of knowledge and type of expertise. How do different goals, theories, and functions affect category organization and do they affect all levels of a taxonomic hierarchy or are some levels (e.g. the basic level) immune from such factors?

A. Expert sorting. In the first experiment, we examine the groupings of tree species, and explanations for those groupings, provided by three types of tree experts. Are there differences in spontaneously generated taxonomies? If there are systematic differences, the precise nature of them might specify the ways in which human interests and goals impact folkbiological taxonomic systems. There are several distinct theoretical positions with respect to differences in categorization as a function of type of expertise. We distinguish four.

1. Predictions concerning expertise. One possibility is that, so to speak, all roads lead to Rome--differing goals and concerns may not penetrate the compelling spontaneous categorization that biological kinds afford. But one does not have to be a structure-in-the-world theorist to anticipate the absence of group differences. Goals and associated differences in patterns of interaction may not affect biological classification because experts interact with biological kinds in so many different ways that a general purpose organizational scheme has the greatest utility. In either case we would expect a high correspondence between folk and scientific taxonomy, justifications that are uniform and predominantly taxonomic, and similarly structured taxonomic hierarchies across groups.

Alternatively it may be that goals and activities do affect biological classification but that their influence on categorization is difficult to detect because biological objects have highly correlated features. To the extent that features are correlated within categories, it will increase the likelihood that classifiers with different orientations will create the same categories for different reasons. For example, perhaps feathers are important to Bob because he uses feathers to make pillows. Bob has a category of all the objects that he can use to procure feathers. Alternatively, perhaps wings are important to Jane because she studies flight of vertebrates. Thus Jane has the category "winged vertebrates." Bob and Jane have almost identical categories, but their rationale for placing the objects into these categories differs dramatically. In short, similar categories in terms of content may mask differences in the features relevant to those categories in the minds of the categorizers (see also Boster & D'Andrade, 1989).

Of course there is also psychological research that would, by extension, support more substantial differences as a function of type of expertise. Different goals, activities, and knowledge may well lead to distinct hierarchical structures—fundamentally different ways of organizing nature. This possibility represents the strongest and most pervasive way in which human goals might shape folkbiological taxonomy. This leads to two further possibilities.

Models of categorization allow for selecting weighting of features or dimensions (e.g. Medin and Schaffer, 1978; Nosofsky, 1992; and see Wisniewski and Medin, 1994 for the more radical possibility that theories may influence the construction and construal of features). Distinct types of expertise may lead to different patterns of attention to and weighting of morphological properties of trees. One could think of these differences as yielding a variable stretching and shrinking of dimensions of some multidimensional similarity space (e.g. Boster and D'Andrade, 1989; Nosofsky, 1992). On this account the same general similarity space would describe all experts and any differences might be attributed to variable weighting of a set of common dimensions. For example, taxonomists may differ from landscape and parks maintenance groups in giving extra weight to features associated with reproduction. If differences in the weighting of dimensions were strong enough to overcome the influence of correlated dimensions, one might observe group differences in sorting as well as differences in the justifications.

The final position we shall examine is motivated by Barsalou's analyses of ad hoc or goal-derived categories (Barsalou, 1982, 1983, 1985, 1991). He argues that categories constructed in the service of goals (e.g. things to take on a camping trip) violate the correlational structure of the environment. In addition, he contrasts context-independent properties of concepts (e.g. basketballs are round) and context-dependent properties (basketballs can float) and notes that goals often serve to activate context-dependent properties. Especially important for our purposes, Barsalou suggests that frequently used goal-derived categories may develop well-established category representations much like those of common taxonomic categories (see Barsalou and Ross, 1986).

Barsalou's analysis suggests that extensive use of goal-derived categories may lead to the following patterns of spontaneous sorting: 1. the correspondence between sorting and scientific taxonomic distance should be low, 2. the justifications for sortings should be utilitarian (linked to goals), 3. the folk taxonomic structure may look more like a series of goal-relevant cross-cuts than a true hierarchical taxonomy. Finally, goals do not necessarily partition the full set of entities in a domain. This raises the possibility that sorting may reflect a mixture of goal-derived and taxonomic categories.

The job descriptions of our three sets of experts suggest that landscape personnel are most likely to frequently use goal-derived categories. Because much of their work involves placing the right tree in the appropriate context, utilitarian functions (providing shade, aesthetic qualities, etc.) are salient to landscapers. On the other hand, the day to day activities of parks maintenance personnel involve caring for an existing population of trees, rather than using trees to achieve particular goals. Finally, we assume that the taxonomists will use scientific, rather than goal-derived categories because most of them frequently use the scientific taxonomy in their work.

In summary, there are theoretical motivations to expect anything from no effect of type of expertise on sorting to a radical restructuring in the service of goals. The justifications for sorts provide data bearing on the idea that the same categories may be created for different reasons. The degree to which expert groups differ in their clustering of tree species should indicate the degree to which human goals and interests influence folkbiological classification.

2. Results. Space does not permit a full analysis of the procedures and results (see Medin, Lynch, Coley, and Atran, in press for details) and we are limited to a summary. Participants were asked to sort 48 tree names into as many categories they wished to create clusters that go together "by nature." In addition the experts were asked to justify or describe the basis for each category. This procedure was repeated to successively lump or split categories to produce a hierarchical taxonomy (the results would be essentially the same if we only considered the initial sorting.

The sorting of the taxonomist group provides a useful comparative standard. Individual taxonomies were combined to produce a consensual taxonomy (again for procedural details and mathematical justification, see Medin et al., in press). For ranks ranging from genus to family up to division, as taxonomic distance increased, mean distance in their consensual clustering increased. Furthermore, the taxonomist's sorting never broke up the genus level. Finally, justifications were almost exclusively in terms of taxonomy and morphology.

The Parks personnel were similar in certain key respects to the taxonomists. Sorting and justifications were primarily in terms of taxonomic and morphological properties. The correlation between scientific taxonomic distance and folk distance was lower, in part because in several cases the morphological properties used in sorting are not necessarily properties given high weight in scientific taxonomy. For example, consensual categories based on type of fruit or nut partially crosscut scientific taxonomy. In general, the greatest correspondence with scientific taxonomy was observed at the genus and family levels. The parks group also differs from the taxonomist group in its inclusion of a clear utilitarian category, weed trees (fast-growing weak-wooded trees that create maintenance problems).

Where does the parks group fall with respect to the question of whether the genus level is privileged in sorting? The picture is mixed but perhaps the best summary statement is that genus is relatively but not absolutely privileged. It is relatively privileged in that the genus level was more likely to be preserved and less likely to be broken up than the family level. In addition, there was a much larger jump in folk distance between the genus level and the family level than between the family level and higher taxonomic ranks. But genus is not absolutely privileged. Even in the most favorable condition where folk terms and science converge on the genus level, the sorting revealed a corresponding category only about half the time.

The pattern of sorting and justification for the landscape group showed the largest influence of utilitarian factors and the largest departure from scientific taxonomy. The correlation of landscapers' folk distance with scientific distance was low and genus-level categories were broken up over half the time. The greatest correspondence between scientific and folk taxonomy appeared to be at the family level though even here it was not strong. The justifications of landscape group members also reflect utilitarian concerns. Weed status, landscape utility, size, and aesthetic value were mentioned more frequently as justifications than morphological properties. It is hard to escape the impression that the landscape sorting is a utilitarian structure imposed on nature rather than a reflection on categories salient in the environment.

Overall, the sorting and justification profile for the landscape group corresponds remarkably well with Barsalou's (1983) analysis of goal-derived categories. The sortings crosscut the correlational structure of the environment (as embodied in scientific taxonomy) were justified in terms of multiple utilitarian factors, and created a hierarchy based on combinations of goal-relevant factors. The modest agreement with science and with the other two groups of experts appears to derive from a combination of goals incompletely partitioning the set of trees and a convergence of morphological and utilitarian criteria on the same categories. Apparently years of experience lead the goal-derived categories of landscape workers to become well-established in memory. The next experiment addresses the question of whether these categories are also used in reasoning.

Summary. Taken together, these comparisons among reveal specific group commonalities and differences. Although there were some clusters of trees that either are similar enough, or convergent enough across goals to be put together regardless of type of expertise, group differences stand out as salient. Most strikingly, what were organized as taxonomic categories by parks and taxonomist groups appeared as goal-derived categories for the landscape group. This does not show that this is the only organization available for the landscape group (it surely is not) but it underlines the accessibility of a utilitarian organization. It remains to be seen whether the nontaxonomic bases for categorization carry over into other conceptual tasks.

B. Category-Based Reasoning.

1. Rationale. The second experiment provides a potentially converging measure of category organization in the form of a category-based reasoning task (Rips, 1975). In this experiment, we examine competing bases for extending novel properties from one tree species to another. Experts are told that a tree has some novel

property and then asked which of two alternative trees is most likely to have that property. Alternatives vary in their scientific (and folk) distance from the target. The question of interest is the degree to which different kinds of categories promote inferences about physiological properties. For example, we might say that a new disease has been discovered that Boxelders can get and ask whether it is more likely that a Sugar Maple or a Cottonwood can also get the disease. The Boxelder is a maple (its other common name is Ashleaf Maple) and, therefore, a taxonomically motivated answer would be Sugar Maple. Boxelders and Cottonwoods are both frequently nominated as garbage trees so a functionally motivated answer would be Cottonwood. Alternatively, if such categories only represent a few relevant shared features but fail to capture important underlying similarities, they might not promote inferences as well as scientific categories do; i.e., the expert might still prefer to draw an inference from Boxelder to Sugar maple.

We used three distinct types of so-called "blank" properties (a novel disease, novel physiological property or novel property associated with cross-fertilization), to see if different kinds of properties evoked different forms of category organization. If results converge with those of Experiment 1, taxonomists' responses should closely follow scientific relations, whereas those of landscape and parks personnel should not necessarily do so.

2. Results. Again, taxonomists provide a standard that can be contrasted with the other groups. Taxonomists pick the taxonomically (and presumably morphologically) more similar tree over the alternative more than 90 percent of the time. Landscapers agree more with scientific taxonomy than with their own consensual sorting (83 percent versus 63 percent) whereas parks workers show the opposite pattern (87 percent agreement with their consensual sorting versus 77 percent agreement with taxonomy). At a finer level of analysis, landscapers pick the taxonomic choice 93 percent of the time when the standard and the taxonomic alternative are of the same genus. At levels above the genus the taxonomic alternative is picked only 54 percent of the time. When taxonomy and consensual sorting are in agreement at higher levels (for Landscapers) that alternative is selected 88 percent of the time. Overall, our observations suggest that the genus level is privileged in the reasoning of landscapers.

It is harder to separate scientific from folk taxonomy for the parks personnel but, relative to landscapers, they show a smaller influence of shared genus. For example, on the triad where Boxelder is the base and the alternatives are Cottonwood and Sugar Maple, parks personnel consistently select Cottonwood (a fellow garbage tree) rather than Sugar Maple (a fellow maple). It does seem that the parks subgroup privileges the genus level when it is marked in language (an observation consistent with this is that a number of parks participants grouped the American Mountain Ash with the Green Ash and White Ash in sorting, even though the American Mountain Ash belongs to a different family and order taxonomically). In short, the folk generic level may be privileged for parks personnel but the reason for this may have as much to do with language as with perceptual salience.

For both the parks and taxonomist subgroups one could say that the same category organization revealed by sorting carries over into reasoning. For the taxonomists this is essentially equivalent to scientific taxonomy. Some of their justifications were in terms of known patterns of disease but these are highly correlated with taxonomic distance. For the parks subgroup, the taxonomy guiding reasoning appears to be their consensual folk taxonomy. When folk and scientific taxonomy made different predictions, they were more likely than landscaper personnel to reason in accord with their own consensual taxonomy. When folk and scientific taxonomy conflicted and the scientific relation was above the genus level, parks subgroup responses were consistent with consensual folk taxonomy as a basis for inferences. The correlation of morphological characters with utilitarian function makes it difficult to pinpoint the exact basis for reasoning but the data strongly suggest that at least weed status, and perhaps landscape utility, influenced choices.

In sharp contrast to the other two groups, the landscape subgroup used a reasoning strategy that departed substantially from their consensual sorting. In sorting, landscapers were most likely to cross-classify genus-level groups of trees; in reasoning, they almost never drew inferences on the basis of folk categories when the scientific choice was at the genus level. As mentioned above, when folk and scientific distance conflicted, landscapers were more likely than the parks subgroup to base inferences on science. We hasten to point out that the landscape reasoning data agreed more with parks consensual sorting than with either their own consensual sorting or with science (though in the latter case the differences were small). The fact that parks subgroup sorting and science predicted the landscape subgroup reasoning about equally well makes the basis for landscape reasoning ambiguous. Folk nomenclature had no reliable effect on patterns of reasoning for landscapers. The landscape folk categories <u>weed</u> and <u>ornamental</u> had some effect on landscapers' reasoning, but only at levels above the genus.

Overall, results support the claim that genus-level categories are inductively privileged. For the landscape subgroup this privilege appears for scientific genera regardless of whether folkterminology marks it; for the parks subgroup, privilege extends most clearly to folkgenerics marked by common terms. It remains to be seen whether this difference can be explained in terms of amount of formal education. The next experiment extends our analysis of rank and privilege to other populations.

C. The basic level and Inductive strength: Experiment 3.

1. Rationale. One important function of categories is to promote inferences; if we know a property is true of members of a given category, we can make an educated guess as to whether that property is true of other related categories. This inductive potential should be maximized

at the privileged level of taxonomy. Maximizing induction potential as a single criterion would dictate the most specific categories in that any property true of higher level categories must be true of subordinate categories. Whatever is true of all fish will be true of all trout and all rainbow trout. On the other hand, people encounter more fish than trout and more trout than rainbow trout and induction potential may trade off number inferred properties and confidence in inductions with frequency. In the limiting case of specificity one would not generalize beyond an individual case. The need to draw inferences and make predictions in the face of uncertainty, however, may pick an intermediate level of categorization as maximizing the tradeoff between frequency and accuracy.

In all of the Rosch et al experiments, the logic of locating the basic level was the same; the basic level was the level above which much information was lost, and below which little information was gained. For instance, in a feature-listing task, subjects listed a mean of 3 common features for the superordinate category furniture, a mean of 9 features for basic level furniture categories (e.g., table) and an average of 10.3 features for subordinate furniture categories (e.g., kitchen table). There was a large gain in information when going from the superordinate to the basic level (6 new common features are added, in this example), and only a slight gain going from the basic to the subordinate (1.3 features). The third experiment extends this logic to inductive reasoning; inductive inferences to a privileged category should be significantly stronger than inferences to more general categories and not significantly weaker than inferences to more specific categories. If table (as opposed to kitchen table or furniture) is the basic level, then inferences about tables should be roughly as strong as inferences about kitchen tables, and much stronger than inferences about furniture. Inductive strength should show a sharp drop as one moves above the basic level.

2. Design. Recall that anthropologists studying traditional cultures find that the folk-generic level (e.g., sparrow, trout, oak) is privileged; psychologists, studying mostly urbanized populations, find that the "basic" level for living things maps onto the rank of life-form (e.g., bird, fish, tree). To help shed light on this puzzle, we undertook a series of studies using category-based induction to investigate privileged taxonomic levels. As in the second experiment participants were told that an unfamiliar property was true of all members of a category, but then were asked how likely it was to be true of the members of a more general category. Using Rosch's logic for diagnosing the basic level, inferences to a privileged level should be judged more likely than those to more general levels, and no less likely than those to more specific levels. If Berlin is correct about the primacy of folk-generic categories, this pattern should center on taxa of the folk-generic rank. Specifically, we would expect inferences to folk-generic taxa to be relatively strong, inferences to taxa of lower order only marginally stronger, and inferences to taxa of higher rank to be relatively weak. That is, there should be a breakpoint or elbow between the folk-generic level and higher levels.

The research participants in this study were both Northwestern University undergraduates and Maya from the Peten region of Guatemala (see Coley, Medin, and Atran , submitted, for details). If expertise or contact with a rich natural world is the critical factor in induction, then undergraduates may privilege the life-form level and Maya the folkgeneric level in induction.

3. Results and Discussion. Overall, inferences to folk-generic categories were consistently stronger than inferences to more general categories, and no weaker than inferences to more specific categories (for details, see Coley, et al., submitted). This result held for both the undergraduate and Maya populations. The gain in inductive strength was greatest moving from life-form to folk-generic categories, suggesting that folk-generic categories are psychologically "basic" with respect to induction. For American students, this result is surprising because Rosch's work suggests that "basic object categories" for living things fall at a more general level.

Our results leave us with the two questions about the relations between folkbiological taxonomy and induction. First, why did the privileged level for Americans with respect to induction fail to correspond to Rosch's "basic" level? Second, why were categories at the same level of specificity inductively privileged for both Americans, with little first-hand experience of the natural world, and Maya, who have a great deal?

One possible account for the discrepancy between results of the present experiments and those of Rosch et al. is to propose different functions for within- and between-category similarity in the sets of tasks. On this account, proposed by many researchers since Rosch et al. (see Lassaline et al., 1992, for a review), the inductive strength of a category is driven by within-category similarity alone, whereas status as a "basic object category" results from maximum within-category similarity relative to between-category similarity. For example, trout might be privileged for induction over fish because people perceive all trout to be more alike--and therefore likely to share properties--than all fish; i.e., within-category similarity is higher for trout than fish. In contrast, fish rather than trout might be diagnosed as a "basic level category" because although all trout are a lot alike, they are also fairly similar to other fish. All fish are somewhat similar--less so than trout--but are very different from other animals. So, fish is "basic" because the ratio of within-category similarity to between-category similarity is higher for fish than for trout. On this account, life-form categories (e.g., fish, tree) may be privileged on Roschian tasks because they maximize within-category similarity relative to between-category similarity, and folk-generic categories (e.g., trout, oak) may be privileged in induction because they maximize within-category similarity alone.

Figure 2a illustrates one way in which, on this account, the inductively privileged level and the "basic" level need not coincide. In Figure 2a, the largest drop in within-category similarity occurs between the folk-generic and the life-form level, predicting an inductive advantage for folk-generic categories. However, the criterion for basic level status--within-category similarity relative to between-category similarity--is maximized at the life-form level (as indicated in Figure2a by the distance between the two lines). Thus, it is possible to conjecture patterns of within- and between-category similarity that yield folk-generic categories privileged for induction but life-form categories privileged for perceptual tasks.

Unfortunately, this account fails when one considers our results, and logical constraints on the relation between within- and between-category similarity. First, according to the pure similarity account, inductive strength should reflect within-category similarity only. We consistently found that the most pronounced dropoff in inductive strength was between the folk-generic and life-form level. This suggests that the largest drop in within-category similarity also occurs between the folk-generic and life-form level. Second, within-category similarity at a given level is logically equivalent to between-category similarity at the immediately subordinate level. For instance, a judgment of within-category similarity for a life-form (e.g., how similar are all fish) is equivalent to aggregated judgments of between-category similarity for folk-generic subordinates of that life-form (e.g., how similar is trout to bass, shark to goldfish, etc.). Therefore, the major breakpoints in within- and between-category similarity cannot both occur between the folk-generic and life-form levels, as required in Figure 2a. On the contrary, within-category similarity at the life-form level must correspond to between-category similarity at the folk-generic level. To suppose otherwise would literally require two different answers to precisely the same question.

Figure 2b is a revision of Figure 2a based on these observations. In accordance with (1), ratings of inductive strength collapsed across studies are substituted in 2b for the hypothetical within-category similarity ratings in 2a. In accordance with (2), between-category similarity at each level in Figure 2b is estimated by projecting the approximations of within-category similarity (i.e., aggregate ratings of inductive strength) to the immediately subordinate level. As Figure 2b indicates, if the major break in within-category similarity occurs between the folk-generic and life-form levels, then the ratio of within- to between-category similarity is maximized at the folk-generic level as well. If these observations are correct, it is impossible to simultaneously observe (1) the largest drop in within-category similarity between folk-generic and life-form categories, and (2) the ratio of within- to between-category similarity maximized at the life-form level. Therefore, a purely similarity-based account such as described above cannot explain both the present results favoring the folk-generic level for induction and Rosch et al.'s findings favoring the life-form level for categorization.

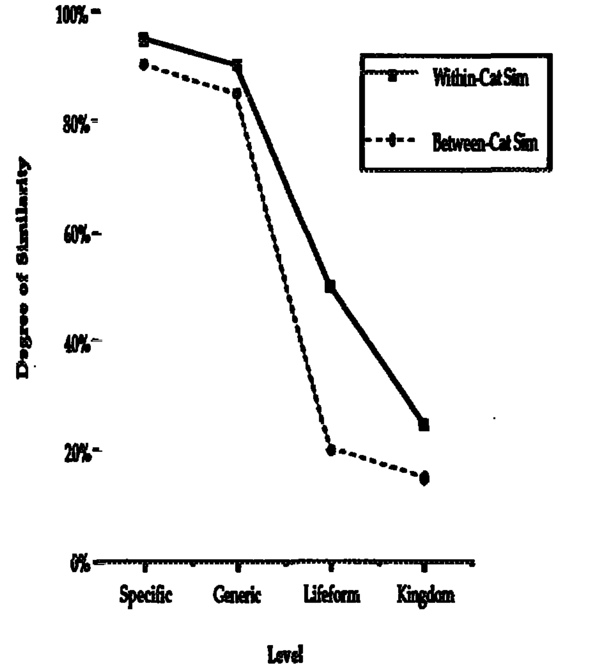

Figure 2a

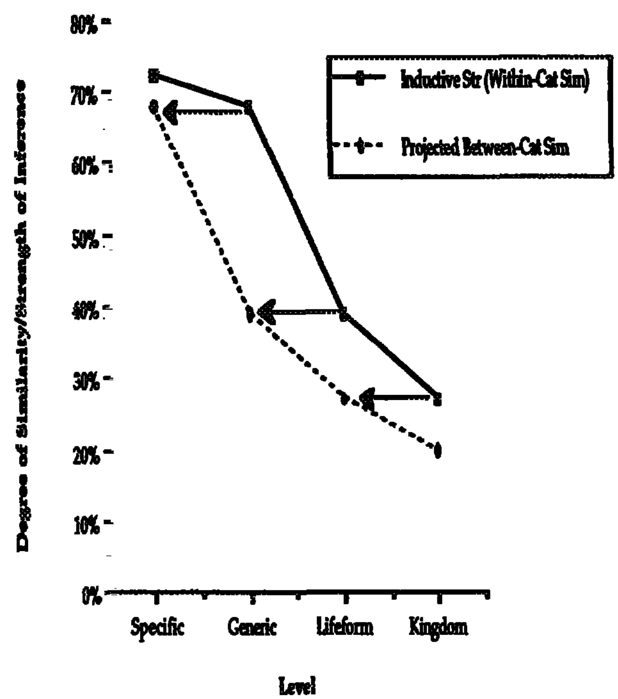

Figure 2b

Although the similarity-based account cannot explain the discrepancy, it could nevertheless be the case that Roschian tasks measuring the "basic level" and our induction task are tapping different competencies. Rosch's results show that (1) the folk-generic level is not the locus of knowledge for urban Americans, and (2) the folk-generic is not the most differentiated level perceptually for Americans. In contrast to Rosch's measures, ours are neither perceptual nor knowledge-dependent. Rather, they are linguistic and expectation-dependent. Rather than having participants list features that they knew to be true of categories, we asked them to project properties chosen expressly for the purpose of being unknown. Instead of testing participants' knowledge, we have tested their expectations. Our results show that despite lack of knowledge and perceptual differentiation, Americans expect that the folk-generic level is most useful for inductive inference.

Why would the most inductively useful level be differ from the "basic" level as indicated by knowledge and perception? Implicitly learning the logic of nomenclatural patterns of inclusion (a red oak is a kind of oak), and explicitly learning inclusion relations of folk-generics under life-forms (an oak is a kind of tree) may be enough to set up a semblance of a folk-biological taxonomy and to flag the folk-generic level despite little experience with members of the categories. But even very limited experience with a subset of biological kinds may be sufficient to fix the locus at the folk-generic level. This may be analogous to what Shipley (1993) refers to an overhypothesis (see also Goodman, 1955). For example, people may know that members of familiar folk-generics are alike in predictable ways, and different from other familiar folk-generics in predictable ways. This knowledge may set up an assumption about novel kinds of animals. For example, if people are told about some unfamiliar mammal, the quagga, they are likely to expect that quaggas will differ from other mammals and be similar to each other in external appearance, in some internal properties, and in a variety of other ways.

The American college students in our study presumably have little first-hand experience of or dependence on the natural world, but they know that elms, maples and oaks are kinds of trees, and expect (perhaps on the basis of nomenclature alone) that red oaks are kinds of oaks, and that northern red oaks are kinds of red oaks. But they may not even know enough to tell the difference between an oak and an elm. There may be universally fixed ranks in folkbiological classification (Berlin, 1992), and the folk-generic level may carry with it the expectation of an inherent physical nature, or underlying essence (Atran, 1987; Medin and Ortony, 1989). We propose that this relatively sophisticated knowledge of nomenclatural patterns in folkbiological taxonomy, coupled with an expectation that (folk-generic?) categories have an essence, would lead to our observed results. Americans may know more features and perceptual affinities at the more abstract level of tree and fish. Nevertheless, they may expect the strongest clusters of properties to cohere at the folk-generic level.

The second question raised by our findings is why we found the same pattern among folkbiologically naive Americans and folkbiologically sophisticated Maya? We have argued that, for Americans, linguistic cues may target an expectation of essence at the folk-generic rank for biological kinds despite lack of familiarity with the target. In situations where folk are likely to be well-acquainted with local living kinds, as in "traditional" societies like the Maya, the weight of ethnobiological evidence argues that perceptual cues converge on folk-generics as being psychologically basic. In fact, the coincidence of domain-general heuristics with domain- specific presumptions may well represent the default case for human understanding of living kinds under normal (evolutionarily-attuned) environmental conditions. This speculation remains to be tested and it certainly is possible that people from traditional societies will perform the same as undergraduates on the sorts of perceptual tasks used by Rosch and her associates. But on the other hand, there is some evidence that the basic level may change with expertise (Tanaka and Taylor, 1991) and people in traditional societies may be experts relative to American undergraduates.

Summary and Conclusions

To what degree do folkbiological conceptual systems reflect universal patterns of feature covariation in the world, or universal habits of mind, and to what degree do they reflect specific goals and needs of the categorizer? In this paper we have focused on the questions of how similarity and theory structure categories and the basis of category privilege. Let's first review differences associated with type of expertise.

Differences in taxonomies. One difference between the groups was in the structure of and justifications for their taxonomies, a difference that may reflect the different purposes for which each group classifies trees. Taxonomists basically reproduced scientific taxonomy. Parks subgroup taxonomies tended to be justified on the basis of morphological and sometimes utilitarian concerns but overall correlated moderately well with science. Finally, the landscapers' used the widest variety of justifications, including many conjunctive justifications ("large native specimen trees"). These findings are consistent with a "taxonomy" based on a number of semi-independent utilitarian dimensions being imposed over the set of trees.

Overall, parks experts can be characterized in terms of heavy emphasis on morphological properties seamlessly integrated with moderate attention to goal-related utilitarian factors and reflecting an influence of common names. The same features, weighted differently (e.g., parks personnel tend to weight features like wood strength and leaf shape more heavily than taxonomists), appear to drive the classification systems of both the taxonomist and parks subgroups. Likewise, both groups "stuck to their (folk)taxonomic guns" in reasoning, drawing inferences which, for the most part, agreed with their sortings. The landscapers, in contrast, violated the correlational structure of the stimulus set to such a degree that structure-based accounts would be severely strained. No stretching or shrinking of a common space would transform their consensual clustering to conform to that of either of the other two groups (save for the degenerate case of a set of dimensions with zero weighting for various subgroups). Barsalou's analysis of goal-derived categories and his conjecture that frequently used goal-derived categories may become well-established in long-term memory (Barsalou, 1982, 1983, 1991), predicts the pattern of landscape sorting in striking detail.

Indeed, the mismatch of landscapers' sorting and reasoning data provides further support for the above analysis. Assume that the landscapers' consensual taxonomy reflects goal-derived, special-purpose categories rather than general-purpose categories expected to maximally capture similarities in the world. If so, there is no reason to expect these goal-derived categories to support inferences about any properties except those related to the goals. The reproductive, disease, and physiological properties used in Experiment 2 have little relation to the functional, utilitarian goals that seemed to drive the landscapers' sortings from Experiment 1. Therefore, landscapers sensibly abandoned their salient but inductively less useful (in this case, at least) consensual taxonomy in favor of a general-purpose taxonomy (either science or an organization approximating that of the parks workers).

In sum, the first experiment revealed an influence of goals on the structuring of categories and no absolute privilege to the genus level. The second experiment was more supportive of the genus level as privileged but, in the case of the parks subgroup, only if the genus level was marked by language. This suggests that perceived structure is not simply determined by how things are in the world but rather that perceived structure is guided by language. Categorization models can and should be applied to these observations.

The third experiment demonstrated a remarkable convergence across cultures with respect to rank and privilege. Specifically, the folkgeneric (roughly genus) level appeared to be psychologically privileged for both knowledgeable Maya informants and undergraduates who have much less direct contact with biological kinds. Note that these findings also converge with our observations on tree experts whose reasoning was strongly influenced by genus-level matches, particularly when shared genus was marked by language. Although on the surface these results appear to support a "structure in the world" framework, we believe that they are most consistent with the idea that they are driven by language and expectations. As we have seen, the idea of perceptual similarity does not seem capable of simultaneously explaining Rosch's categorization results and our reasoning results. But it's equally clear that ideas about presumption of essence and expectations remain speculative in the absence of a body of evidence associated with more precise predictions. We also need cross-cultural measures of privilege that are more perceptually-oriented.

Although findings that hold across tree experts, undergraduates and Guatemala Maya may have the air of universality, that inductive leap may be premature. For example, we are currently studying ecologists and the structure of biological communities seems to have a lot of action at the species level. For example, the Burr Oak plays a very different role from that of other oaks found in the Savannah. It may be that ecologists will privilege the species level rather than genus. One should also note that horticulturists focus attention on the varietal level and perhaps their inductions would be even more narrowly constrained.

At this point one might even worry that the key notion of a basic, privileged level is suspect. It is not a large step from the idea of multiple privileged levels or different levels being privileged for different conceptual functions to the conclusion that there the very idea of basic or privileged level is not useful. We believe, however, that despair should not be the response to the first indication that categorization principles might not be as parsimonious as one might wish. Indeed there are fascinating questions that remain to be asked and answered concerning expertise, development, and cultural similarities and differences. For example, there is clear evidence that infants learning names for things have clear expectations concerning the extensions of category labels (see Waxman, 1994 for a review). But the question remains whether American infants are surprised to find that things as different as robins and Sparrows are both typically called "birds" or are Maya infants surprised that things as similar as Robins and Sparrows have different names? Or are infants from both cultures ready for either possibility? When we are able to answer questions such as this we will be in a much better position to understand the roles of similarity and theories in the development and expression of conceptual behavior.

References

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

Anderson, J. R. 1991. The adaptive nature of human categorization. *Psychological Review 98*: 409-429.

Atran, S. 1987. Ordinary constraints on the semantics of living kinds: A commonsense alternative to recent treatments of natural-object terms. *Mind and Language 2*: 27-63.

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

Barsalou, L. W. 1983. Adhoc categories. *Memory & Cognition 11*: 211-227.

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

Barsalou, L. W. 1991. Deriving categories to achieve goals. In G. H. Bower (Ed.), *The Psychology of Learning and Motivation: Advances in research and theory*, vol. 27. New York: Academic Press.

Barsalou, L. W. 1992. Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A. C. Collins, S. E. Gathercole, M. A. Conway, & P. E. M. Morris (Eds.), *Theories of memories*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Barsalou, L. W. 1993. Structure, flexibility, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A. C. Collins, S. E. Gathercole, & M. A. Conway (Eds.), *Theories of Memory*. London: Erlbaum.

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

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

Berlin, B., Breedlove, D. E. & Raven, P. H. 1966. Folk taxonomies and biological classification. *Science 154*: 273-275.

Berlin, B., Breedlove, D. E. & Raven, P. H. 1973. General principles of classification and nomenclature in folk biology. *American Anthropologist 75*: 214-242.

Billman, D. 1992. Modeling category learning and use: Representation and processing. In B. Burns ed. *Percepts, concepts, and categories: The representation and processing of information*. New York: Elsevier.

Billman, D., & Knutson, J. F. 1996. Unsupervised concept learning and value systematicity: A complex whole aids learning the parts. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22: 458-475.

Boster, J. S. 1991. The information economy model applied to biological similarity judgment. In L. Resnik, J. Levine, & S. Teasley eds. *Perspectives on socially shared cognition*. Washington, DC: American Psychological Association.

Boster, J. S., & D'Andrade, R. 1989. Natural and human sources of cross-cultural agreement in ornithological classification. *American Anthropologist 88*:569-583.

Boster, J. S., & Johnson, J. C. 1989. Form or function: A comparison of expert and novice judgments of similarity among fish. *American Anthropologist 91*: 866-889.

Brooks, L. R., Norman, G. R., & Allen, S. W. 1991. Role of specific similarity in a medical diagnosis task. *Journal of Experimental Psychology: General 120*: 278-287.

Burns, B. ed. 1992. *Percepts, concepts, and categories: The representation and processing of information*. New York: Elsevier.

Carey, S. 1985. *Conceptual change in childhood*. Cambridge, MA: Bradford Books.

Carey, S., & Spelke, E. 1994. Domain-specific knowledge and conceptual change. In L. A. Hirschfeld & S. A. Gelman eds. *Mapping the mind*. Cambridge: Cambridge University Press.

Chi, M. T. H. 1983. Knowledge-derived categorization in young children. In D. R. Rogers & J. A. Sloboda eds. *The acquisition of symbolic skill*. New York: Plenum Press.

Chi, M. T. H., Feltovich, P., & Glaser, R. 1981. Categorization and representation of physics problems by experts and novices. *Cognitive Science 5*: 121-152.

Chi, M. T. H., Hutchinson, J. E., & Robin, A. F. 1989. How inferences about novel domain-related concepts can be constrained by structured knowledge. *Merrill-Palmer Quarterly 35(1)*: 27-62.

Choi, S., McDaniel, M. A., & Busemeyer, J. R. 1993. Incorporating prior biases in network models of conceptual rule learning. *Memory & Cognition 21*: 413-423.

Clapper, J. P., & Bower, G. H. 1994. Category invention in unsupervised learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition20*: 443-460.

Coley, J. D., Medin, D. L., Atran, S., & Lynch, E. (submitted for publication). Does privilege have its rank? Folkbiological taxonomy and induction in two cultures.

Corter, J. E., & Gluck, M. A. 1992. Explaining basic categories: Feature predictability and information. *Psychological Bulletin 111*: 291-303.

Ellen, R. 1993. *The cultural relations of classification*. Cambridge, MA: Cambridge University Press.

Estes, W. K. 1994. *Classification and Cognition* 161-163. Oxford: Oxford University Press.

Fisher, D. H., Pazzani, M. J., & Langley, P. eds. 1991. *Concept formation: Knowledge and experience in unsupervised learning*. San Mateo, CA: Morgan Kaufman.

Fisher, D., & Yoo, J. P. 1993. Categorization, concept learning, and problem solving: A unifying view. In G. V. Nakamura, R. Taraban, & D. L. Medin eds. *The Psychology of Learning and Motivation: Categorization by* *humans and machines*, vol. 29. San Diego: Academic Press, Inc.

Gelman, S. A., & Coley, J. D. 1991. Language and categorization: The acquisition of natural kind terms. In J. P. Byrnes & S. A. Gelman eds. *Perspectives on language and thought: Interrelations in development* (pp. 146-196). Cambridge: Cambridge University Press.

Gelman, S. A., & Wellman, H. M. 1991. Insides and essences: Early understandings of the nonobvious. *Cognition 38*: 213-244.

Gentner, D., & Markman, A. B. 1994. Structural alignment in comparison: No difference without similarity. *Psychology Science 5*: 152-158.

Gobbo, C. & Chi, M. T. H. 1986. How knowledge is structured and used by expert and novice children. *Cognitive Development 1*: 221-237.

Goodman, N. 1955. *Fact, fiction and forecast*. Cambridge, MA: Harvard University Press.

Heit, E. 1993. Modeling the effects of expectations on recognition memory. *Psychological Science 4*: 244-252.

Heit, E. (in press). Models of the effects of prior knowledge on category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.

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

Honeck, R. P., Firment, M., & Case, T. J. S. 1987. Expertise and categorization. *Bulletin of the Psychonomic Society 25*: 431-434.

Johnson, K. E. 1992. The effect of expertise on hierarchical systems of categorization. Unpublished doctoral dissertation, Emory University, Atlanta, GA.

Johnson, K. E., & Mervis, C. B. 1994. Microgenetic analysis of first steps in children's acquisition of expertise on shorebirds. *Developmental Psychology 30*: 418-435.

Johnson, K. E., Mervis, C. B., & Boster, J. S. 1992. Developmental changes within the structure of the mammal domain. *Developmental Psychology 28*: 74-83.

Jones, G. V. 1983. Identifying basic categories. *Psychological Bulletin 94*: 423-428.

Keil, F. C. 1989. *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.

Keil, F. C. 1994. The birth and nurturance of concepts by domains: The origins of concepts of living things. In L. A. Hirschfeld & S. A. Gelman eds., *Mapping the mind*. Cambridge: Cambridge University Press.

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

Lassaline, M. E., Wisniewski, E. J., & Medin, D. L. 1992. Basic levels in artificial and natural categories: Are all basic levels created equal? In B. Burns ed. *Percepts, concepts, and categories: The representation and processing of information*. New York: Elsevier.

Malt, B. C. 1995. Category coherence in cross-cultural perspective. *Cognitive Psychology 29*: 85-148.

Mandler, J. M., Bauer, P. J., & McDonough, L. 1991. Separating the sheep from the goats: Differentiating global categories. *Cognitive Psychology 23*: 263-298.

Medin, D. L., & Heit, E. J. (in press). Categorization. In D. Rumelhart & B. Martin eds. *Handbook of cognition and perception*. Hillsdale, NJ: Erlbaum.

Medin, D. L., & Ortony, A. 1989. Psychological essentialism. In S. Vosniadou & A. Ortony eds. *Similarity and Analogical Reasoning*.

Medin, D. L. & Schaffer, M. M. 1978. Context theory of classification learning. *Psychological Review 85*: 207-238.

Medin, D. L., Goldstone, R. L., & Gentner, D. 1993. Respects for similarity. *Psychological Review 100*: 254-278.

Medin, D. L., Lynch, E. B., Coley, J. D., Atran, S. (in press). Categorization and reasoning among tree experts: Do all roads lead to Rome? *Cognitive Psychology*.

Mervis, C. B., Johnson, K. E., & Scott, P. 1993. Perceptual knowledge, conceptual knowledge, and expertise: *Comment on Jones and Smith. Cognitive Development 8*: 149-155.

Michalski, R. S. 1989. Two-tiered concept meaning, inferential matching, and conceptual cohesiveness. In S. Vosniadou & A. Ortony eds. *Similarity and analogical meaning*. New York: Cambridge University press.

Michalski, R. S., & Tecuci, G. Eds. 1994. *Machine learning: A multistrategy approach*, vol. 4. San Francisco: Morgan Kaufman.

Mooney, R. J. 1993. Integrating theory and data in category learning. In G. V. Nakamura, R. Taraban, & D. L. Medin eds. *The Psychology of Learning and Motivation: Categorization by humans and machines*, vol. 29. San Diego: Academic Press, Inc.

Murphy, G. L. 1982. Cue validity and levels of categorization. *Psychological Bulletin 91*: 174-177.

Murphy, G. L. 1991. Parts in object concepts: Experiments with artificial categories. *Memory & Cognition, 19*: 423-438.

Murphy, G. L. 1993a. Theories and concept formation. In I. V. Mechelen, J. Hampton, R. Michalski, & P. Theuns eds. *Categories and concepts: Theoretical views and inductive data analysis*. London: Academic Press.

Murphy, G. L. 1993b. A rational theory of concepts. In G. V. Nakamura, R. Taraban, & D. L. Medin eds. *The Psychology of Learning and Motivation: Categorization by humans and machines*, vol. 29. San Diego: Academic Press, Inc.

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

Myers, J. L., Lohmeier, J. H., & Well, A. D. 1994. Modeling probabilistic categorization data: Exemplar meaning and connectionist nets. *Psychological Science 5:* 83-89.

Nakamura, G. V., Taraban, R., & Medin, D. L. eds. 1993 *The Psychology of Learning and Motivation: Categorization by humans and machines*, vol. 29. San Diego: Academic Press, Inc.

Nosofsky, R. M. 1992. Exemplar-based approach to relating categorization, identification, and recognition. In F. G. Ashby ed. *Multidimensional models of perception and cognition*. Hillsdale, NJ: Erlbaum.

Nosofsky, R. M., Gluck, M. A., Palmeri, T. S., McKinley, S. C., & Glauthier, P. 1994. Comparing models of rule-based classification learning: A replication and extension of Shepard, Hovland, and Jenkins 1961. *Memory & Cognition 22*: 352-369.

Palmer, C. F., Jones, R. K., Hennessy, B. L., Unze, M. G., & Pick, A. D. 1989. How is a trumpet known? The "basic object level" concept and perception of musical instruments. *American Journal of Psychology 102*: 17-37.

Pazzani, M. J. 1991. Influence of prior knowledge on concept acquisition: Experimental and computational results. *Journal of Experimental Psychology: Learning, Memory, and Cognition 17*, 416-432.

Rips, L. J. 1975. Inductive judgements about natural categories. *Journal of Verbal Learning and Verbal Behavior 14*: 665-681.

Rips, L. J., & Collins, A. 1993. Categories and resemblance. *Journal of Experimental Psychology: General 122*: 468-486.

Rosch, E., Mervis, C. B., Gray, W., Johnson, D., & Boyes-Braem, P. 1976. Basic objects in natural categories. *Cognitive Psychology 8*: 573-605.

Shipley, E. F. 1993. Categories, hierarchies, and induction. In D. L. Medin ed. *The Psychology of Learning and Motivation, Vol. 30*. New York: Academic Press.

Smith, L. B., & Heise, D. 1992. Perceptual similarity and conceptual structure. In B. Burns ed *Percepts, concepts, and categories: The representation and processing of information*. New York: Elsevier.

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

VanMechelen, I., Hampton, J., Michalski, R., & Theuns, P. eds. 1993. *Categories and Concepts: Theoretical Views and Inductive Data Analysis*. London: Academic Press.

Waldman, M. R., & Holyoak, K. J. 1992. Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General 121*: 222-236.

Walker, S. J. 1992. Supernatural beliefs, natural kinds, and conceptual structure. *Memory & Cognition 20*: 655-662.

Waxman, S. R. 1994. The development of an appreciation of specific linkages between linguistic and conceptual organization. *Lingua 92*:229-257.

Wisniewski, E. J., & Medin, D. L. 1994. On the interaction of theory and data in concept learning. *Cognitive Science 18*:221-281.

Wisniewski, E. J., & Medin, D. L. 1994a The fiction and nonfiction of features. In R. S. Michalski & G. Tecuci eds. *Machine Learning Vol. 4*, San Mateo, CA: Morgan Kaufmann.

Wisniewski, E. J., & Medin, D. L. 1994b. On the interaction of theory and data in concept learning. *Cognitive Science 18*: 221-281.