

This excerpt from

The Motion Aftereffect.

George Mather, Frans Verstraten and Stuart Anstis, editors.

© 1998 The MIT Press.

is provided in screen-viewable form for personal use only by members of MIT CogNet.

Unauthorized use or dissemination of this information is expressly forbidden.

If you have any questions about this material, please contact cognetadmin@cognet.mit.edu.

Chapter 10

Principles of Categorization

Eleanor Rosch

The following is a taxonomy of the animal kingdom. It has been attributed to an ancient Chinese encyclopedia entitled the *Celestial Emporium of Benevolent Knowledge*:

On those remote pages it is written that animals are divided into (a) those that belong to the Emperor, (b) embalmed ones, (c) those that are trained, (d) suckling pigs, (e) mermaids, (f) fabulous ones, (g) stray dogs, (h) those that are included in this classification, (i) those that tremble as if they were mad, (j) innumerable ones, (k) those drawn with a very fine camel's hair brush, (l) others, (m) those that have just broken a flower vase, (n) those that resemble flies from a distance. (Borges 1966, p. 108)

Conceptually, the most interesting aspect of this classification system is that it does not exist. Certain types of categorizations may appear in the imagination of poets, but they are never found in the practical or linguistic classes of organisms or of man-made objects used by any of the cultures of the world. For some years, I have argued that human categorization should not be considered the arbitrary product of historical accident or of whimsy but rather the result of psychological principles of categorization, which are subject to investigation. This chapter is a summary and discussion of those principles.

The chapter is divided into five parts. The first part presents the two general principles that are proposed to underlie categorization systems. The second part shows the way in which these principles appear to result in a basic and primary level of categorization in the levels of abstraction in a taxonomy. It is essentially a summary of the research already reported on basic level objects (Rosch et al., 1976). Thus the second section may be omitted by the reader already sufficiently familiar with that material. The third part relates the principles of categorization to the formation of prototypes in those categories that are at the same level of abstraction in a taxonomy. In particular, this section attempts to clarify the operational concept of prototypicality and to separate that concept from claims concerning the role of prototypes in cognitive processing, representation, and learning for which there is little evidence. The fourth part presents two issues that are problematical for the abstract principles of categorization stated in the first part: (1) the relation of context to basic level objects and prototypes; and (2) assumptions about the nature of the attributes of real-world objects that underlie the claim that there is structure in the world.

From chapter 8 in *Concepts: Core Readings*, ed. E. Margolis and S. Laurence (Cambridge, MA: MIT Press, 1978/1999), 189–206. Reprinted with permission.

The fifth part is a report of initial attempts to base an analysis of the attributes, functions, and contexts of objects on a consideration of objects as props in culturally defined events.

It should be noted that the issues in categorization with which we are primarily concerned have to do with explaining the categories found in a culture and coded by the language of that culture at a particular point in time. When we speak of the formation of categories, we mean their formation in the culture. This point is often misunderstood. The principles of categorization proposed are not as such intended to constitute a theory of the development of categories in children born into a culture nor to constitute a model of how categories are processed (how categorizations are made) in the minds of adult speakers of a language.

The Principles

Two general and basic principles are proposed for the formation of categories: The first has to do with the function of category systems and asserts that the task of category systems is to provide maximum information with the least cognitive effort; the second has to do with the structure of the information so provided and asserts that the perceived world comes as structured information rather than as arbitrary or unpredictable attributes. Thus maximum information with least cognitive effort is achieved if categories map the perceived world structure as closely as possible. This condition can be achieved either by the mapping of categories to given attribute structures or by the definition or redefinition of attributes to render a given set of categories appropriately structured. These principles are elaborated in the following.

Cognitive Economy

The first principle contains the almost common-sense notion that, as an organism, what one wishes to gain from one's categories is a great deal of information about the environment while conserving finite resources as much as possible. To categorize a stimulus means to consider it, for purposes of that categorization, not only equivalent to other stimuli in the same category but also different from stimuli not in that category. On the one hand, it would appear to the organism's advantage to have as many properties as possible predictable from knowing any one property, a principle that would lead to formation of large numbers of categories with as fine discriminations between categories as possible. On the other hand, one purpose of categorization is to reduce the infinite differences among stimuli to behaviorally and cognitively usable proportions. It is to the organism's advantage not to differentiate one stimulus from others when that differentiation is irrelevant to the purposes at hand.

Perceived World Structure

The second principle of categorization asserts that unlike the sets of stimuli used in traditional laboratory-concept attainment tasks, the perceived world is not an unstructured total set of equiprobable co-occurring attributes. Rather, the material objects of the world are perceived to possess (in Garner's, 1974,

sense) high correlational structure. That is, given a knower who perceives the complex attributes of feathers, fur, and wings, it is an empirical fact provided by the perceived world that wings co-occur with feathers more than with fur. And given an actor with the motor programs for sitting, it is a fact of the perceived world that objects with the perceptual attributes of chairs are more likely to have functional sit-on-able-ness than objects with the appearance of cats. In short, combinations of what we perceive as the attributes of real objects do not occur uniformly. Some pairs, triples, etc., are quite probable, appearing in combination sometimes with one, sometimes another attribute; others are rare; others logically cannot or empirically do not occur.

It should be emphasized that we are talking about the perceived world and not a metaphysical world without a knower. What kinds of attributes *can* be perceived are, of course, species-specific. A dog's sense of smell is more highly differentiated than a human's, and the structure of the world for a dog must surely include attributes of smell that we, as a species, are incapable of perceiving. Furthermore, because a dog's body is constructed differently from a human's, its motor interactions with objects are necessarily differently structured. The "out there" of a bat, a frog, or a bee is surely more different still from that of a human. What attributes *will* be perceived given the ability to perceive them is undoubtedly determined by many factors having to do with the functional needs of the knower interacting with the physical and social environment. One influence on how attributes will be defined by humans is clearly the category system already existent in the culture at a given time. Thus, our segmentation of a bird's body such that there is an attribute called "wings" may be influenced not only by perceptual factors such as the gestalt laws of form that would lead us to consider the wings as a separate part (Palmer 1977) but also by the fact that at present we already have a cultural and linguistic category called "birds." Viewing attributes as, at least in part, constructs of the perceiver does not negate the higher-order structural fact about attributes at issue, namely that the attributes of wings and that of feathers do co-occur in the perceived world.

These two basic principles of categorization, a drive toward cognitive economy combined with structure in the perceived world, have implications both for the level of abstraction of categories formed in a culture and for the internal structure of those categories once formed.

For purposes of explication, we may conceive of category systems as having both a vertical and horizontal dimension. The vertical dimension concerns the level of inclusiveness of the category—the dimension along which the terms collie, dog, mammal, animal, and living thing vary. The horizontal dimension concerns the segmentation of categories at the same level of inclusiveness—the dimension on which dog, cat, car, bus, chair, and sofa vary. The implication of the two principles of categorization for the vertical dimension is that not all possible levels of categorization are equally good or useful; rather, the most basic level of categorization will be the most inclusive (abstract) level at which the categories can mirror the structure of attributes perceived in the world. The implication of the principles of categorization for the horizontal dimension is that to increase the distinctiveness and flexibility of categories, categories tend to become defined in terms of prototypes or prototypical instances that contain

the attributes most representative of items inside and least representative of items outside the category.

The Vertical Dimension of Categories: Basic-Level Objects

In a programmatic series of experiments, we have attempted to argue that categories within taxonomies of concrete objects are structured such that there is generally one level of abstraction at which the most basic category cuts can be made (Rosch et al. 1976a). By *category* is meant a number of objects that are considered equivalent. Categories are generally designated by names (e.g., *dog*, *animal*). A *taxonomy* is a system by which categories are related to one another by means of class inclusion. The greater the inclusiveness of a category within a taxonomy, the higher the level of abstraction. Each category within a taxonomy is entirely included within one other category (unless it is the highest level category) but is not exhaustive of that more inclusive category (see Kay 1971). Thus the term *level of abstraction* within a taxonomy refers to a particular level of inclusiveness. A familiar taxonomy is the Linnean system for the classification of animals.

Our claims concerning a basic level of abstraction can be formalized in terms of cue validity (Rosch et al. 1976a) or in terms of the set theoretic representation of similarity provided by Tversky (1977, and Tversky and Gati 1978). Cue validity is a probabilistic concept; the validity of a given cue x as a predictor of a given category y (the conditional probability of y/x) increases as the frequency with which cue x is associated with category y increases and decreases as the frequency with which cue x is associated with categories other than y increases (Beach 1964a, 1964b; Reed 1972). The cue validity of an entire category may be defined as the summation of the cue validities for that category of each of the attributes of the category. A category with high cue validity is, by definition, more differentiated from other categories than one of lower cue validity. The elegant formulization that Tversky (1978) provides is in terms of the variable "category resemblance," which is defined as the weighted sum of the measures of all of the common features within a category minus the sum of the measures of all of the distinctive features. Distinctive features include those that belong to only some members of a given category as well as those belonging to contrasting categories. Thus Tversky's formalization does not weight the effect of contrast categories as much as does the cue validity formulation. Tversky suggests that two disjoint classes tend to be combined whenever the weight of the added common features exceeds the weight of the distinctive features.

A working assumption of the research on basic objects is that (1) in the perceived world, information-rich bundles of perceptual and functional attributes occur that form natural discontinuities, and that (2) basic cuts in categorization are made at these discontinuities. Suppose that basic objects (e.g., chair, car) are at the most inclusive level at which there are attributes common to all or most members of the category. Then both total cue validities and category resemblance are maximized at that level of abstraction at which basic objects are categorized. This is, categories one level more abstract will be superordinate categories (e.g., furniture, vehicle) whose members share only a few attributes

among each other. Categories below the basic level will be bundles of common and, thus, predictable attributes and functions but contain many attributes that overlap with other categories (for example, kitchen chair shares most of its attributes with other kinds of chairs).

Superordinate categories have lower total cue validity and lower category resemblance than do basic-level categories, because they have fewer common attributes; in fact, the category resemblance measure of items within the superordinate can even be negative due to the high ratio of distinctive to common features. Subordinate categories have lower total cue validity than do basic categories, because they also share most attributes with contrasting subordinate categories; in Tversky's terms, they tend to be combined because the weight of the added common features tends to exceed the weight of the distinctive features. That basic objects are categories at the level of abstraction that maximizes cue validity and maximizes category resemblance is another way of asserting that basic objects are the categories that best mirror the correlational structure of the environment.

We chose to look at concrete objects because they appeared to be a domain that was at once an indisputable aspect of complex natural language classifications yet at the same time was amenable to methods of empirical analysis. In our investigations of basic categories, the correlational structure of concrete objects was considered to consist of a number of inseparable aspects of form and function, any one of which could serve as the starting point for analysis. Four investigations provided converging operational definitions of the basic level of abstraction: attributes in common, motor movements in common, objective similarity in shape, and identifiability of averaged shapes.

Common Attributes

Ethnobiologists had suggested on the basis of linguistic criteria and field observation that the folk genus was the level of classification at which organisms had bundles of attributes in common and maximum discontinuity between classes (see Berlin 1978). The purpose of our research was to provide a systematic empirical study of the co-occurrence of attributes in the most common taxonomies of biological and man-made objects in our own culture.

The hypothesis that basic level objects are the most inclusive level of classification at which objects have numbers of attributes in common was tested for categories at three levels of abstraction for nine taxonomies: tree, bird, fish, fruit, musical instruments, tool, clothing, furniture, and vehicle. Examples of the three levels for one biological and one nonbiological taxonomy are shown in table 10.1. Criteria for choice of these specific items were that the taxonomies contain the most common (defined by word frequency) categories of concrete nouns in English, that the levels of abstraction bear simple class-inclusion relations to each other, and that those class-inclusion relations be generally known to our subjects (be agreed upon by a sample of native English speakers). The middle level of abstraction was the hypothesized basic level: For nonbiological taxonomies, this corresponded to the intuition of the experimenters (which also turned out to be consistent with Berlin's linguistic criteria); for biological categories, we assumed that the basic level would be the level of the folk generic.

Table 10.1
Examples of taxonomies used in basic object research

Superordinate	Basic level	Subordinate
Furniture	Chair	Kitchen chair
		Living-room chair
	Table	Kitchen table
		Dining-room table
	Lamp	Floor lamp
		Desk lamp
Tree	Oak	White oak
		Red oak
	Maple	Silver maple
		Sugar maple
	Birch	River birch
		White birch

Subjects received sets of words taken from these nine taxonomies; the subject's task was to list all of the attributes he could think of that were true of the items included in the class of things designated by each object name. Thus, for purposes of this study, attributes were defined operationally as whatever subjects agreed them to be with no implications for whether such analysis of an object could or could not be perceptually considered prior to knowledge of the object itself. Results of the study were as predicted: Very few attributes were listed for the superordinate categories, a significantly greater number listed for the supposed basic-level objects, and not significantly more attributes listed for subordinate-level objects than for basic-level. An additional study showed essentially the same attributes listed for visually present objects as for the object names. The single unpredicted result was that for the three biological taxonomies, the basic level, as defined by numbers of attributes in common, did not occur at the level of the folk generic but appeared at the level we had originally expected to be superordinate (e.g., *tree* rather than *oak*).

Motor Movements

Inseparable from the perceived attributes of objects are the ways in which humans habitually use or interact with those objects. For concrete objects, such interactions take the form of motor movements. For example, when performing the action of sitting down on a chair, a sequence of body and muscle movements are typically made that are inseparable from the nature of the attributes of chairs—legs, seat, back, etc. This aspect of objects is particularly important in light of the role that sensory-motor interaction with the world appears to play in the development of thought (Bruner, Olver, and Greenfield 1966; Nelson 1974; Piaget 1952).

In our study of motor movements, each of the sets of words used in the previous experiment was administered to new subjects. A subject was asked to describe, in as much finely analyzed detail as possible, the sequences of motor movements he made when using or interacting with the object. Tallies of

agreed upon listings of the same movements of the same body part in the same part of the movement sequence formed the unit of analysis. Results were identical to those of the attribute listings; basic objects were the most general classes to have motor sequences in common. For example, there are few motor programs we carry out to items of furniture in general and several specific motor programs carried out in regard to sitting down on chairs, but we sit on kitchen and living-room chairs using essentially the same motor programs.

Similarity in Shapes

Another aspect of the meaning of a class of objects is the appearance of the objects in the class. In order to be able to analyze correlational structures by different but converging methods, it was necessary to find a method of analyzing similarity in the visual aspects of the objects that was not dependent on subjects' descriptions, that was free from effects of the object's name (which would not have been the case for subjects' ratings of similarity), and that went beyond similarity of analyzable, listable attributes that had already been used in the first study described. For this purpose, outlines of the shape of two-dimensional representations of objects were used, an integral aspect of natural forms. Similarity in shape was measured by the amount of overlap of the two outlines when the outlines (normalized for size and orientation) were juxtaposed.

Results showed that the ratio of overlapped to nonoverlapped area when two objects from the same basic-level category (e.g., two cars) were superimposed was far greater than when two objects from the same superordinate category were superimposed (e.g., a car and a motorcycle). Although some gain in ratio of overlap to nonoverlap also occurred for subordinate category objects (e.g., two sports cars), the gain obtained by shifting from basic-level to subordinate objects was significantly less than the gain obtained by shifting from superordinate to basic-level objects.

Identifiability of Averaged Shapes

If the basic level is the most inclusive level at which shapes of objects of a class are similar, a possible result of such similarity may be that the basic level is also the most inclusive level at which an averaged shape of an object can be recognized. To test this hypothesis, the same normalized superimposed shapes used in the previous experiment were used to draw an averaged outline of the overlapped figures. Subjects were then asked to identify both the superordinate category and the specific object depicted. Results showed that basic objects were the most general and inclusive categories at which the objects depicted could be identified. Furthermore, overlaps of subordinate objects were no more identifiable than objects at the basic level.

In summary, our four converging operational definitions of basic objects all indicated the same level of abstraction to be basic in our taxonomies. Admittedly, the basic level for biological objects was not that predicted by the folk genus; however, this fact appeared to be simply accounted for by our subjects' lack of knowledge of the additional depth of real-world attribute structure available at the level of the folk generic (see Rosch et al. 1976a).

Implications for Other Fields

The foregoing theory of categorization and basic objects has implications for several traditional areas of study in psychology; some of these have been tested.

Imagery

The fact that basic-level objects were the most inclusive categories at which an averaged member of the category could be identified suggested that basic objects might be the most inclusive categories for which it was possible to form a mental image isomorphic to the appearance of members of the class as a whole. Experiments using a signal-detection paradigm and a priming paradigm, both of which have been previously argued to be measures of imagery (Peterson and Graham 1974; Rosch 1975c), verified that, in so far as it was meaningful to use the term *imagery*, basic objects appeared to be the most abstract categories for which an image could be reasonably representative of the class as a whole.

Perception

From all that has been said of the nature of basic classifications, it would hardly be reasonable to suppose that in perception of the world, objects were first categorized either at the most abstract or at the most concrete level possible. Two separate studies of picture verification (Rosch et al. 1976a; Smith, Balzano, and Walker 1978) indicate that, in fact, objects may be first seen or recognized as members of their basic category, and that only with the aid of additional processing can they be identified as members of their superordinate or subordinate category.

Development

We have argued that classification into categories at the basic level is overdetermined because perception, motor movements, functions, and iconic images would all lead to the same level of categorization. Thus basic objects should be the first categorizations of concrete objects made by children. In fact, for our nine taxonomies, the basic level was the first named. And even when naming was controlled, pictures of several basic-level objects were sorted into groups "because they were the same type of thing" long before such a technique of sorting has become general in children.

Language

From all that has been said, we would expect the most useful and, thus, most used name for an item to be the basic-level name. In fact, we found that adults almost invariably named pictures of the subordinate items of the nine taxonomies at the basic level, although they knew the correct superordinate and subordinate names for the objects. On a more speculative level, in the evolution of languages, one would expect names to evolve first for basic-level objects, spreading both upward and downward as taxonomies increased in depth. Of great relevance for this hypothesis are Berlin's (1972) claims for such a pattern for the evolution of plant names, and our own (Rosch et al. 1976a) and Newport and Bellugi's (1978) finding for American Sign Language of the Deaf, that

it was the basic-level categories that were most often coded by single signs and super- and subordinate categories that were likely to be missing. Thus a wide range of converging operations verify as basic the same levels of abstraction.

The Horizontal Dimension: Internal Structure of Categories: Prototypes

Most, if not all, categories do not have clear-cut boundaries. To argue that basic object categories follow clusters of perceived attributes is not to say that such attribute clusters are necessarily discontinuous.

In terms of the principles of categorization proposed earlier, cognitive economy dictates that categories tend to be viewed as being as separate from each other and as clear-cut as possible. One way to achieve this is by means of formal, necessary and sufficient criteria for category membership. The attempt to impose such criteria on categories marks virtually all definitions in the tradition of Western reason. The psychological treatment of categories in the standard concept-identification paradigm lies within this tradition. Another way to achieve separateness and clarity of actually continuous categories is by conceiving of each category in terms of its clear cases rather than its boundaries. As Wittgenstein (1953) has pointed out, categorical judgments become a problem only if one is concerned with boundaries—in the normal course of life, two neighbors know on whose property they are standing without exact demarcation of the boundary line. Categories can be viewed in terms of their clear cases if the perceiver places emphasis on the correlational structure of perceived attributes such that the categories are represented by their most structured portions.

By prototypes of categories we have generally meant the clearest cases of category membership defined operationally by people's judgments of goodness of membership in the category. A great deal of confusion in the discussion of prototypes has arisen from two sources. First, the notion of prototypes has tended to become reified as though it meant a specific category member or mental structure. Questions are then asked in an either-or fashion about whether something is or is not the prototype or part of the prototype in exactly the same way in which the question would previously have been asked about the category boundary. Such thinking precisely violates the Wittgensteinian insight that we can judge how clear a case something is and deal with categories on the basis of clear cases in the total absence of information about boundaries. Second, the empirical findings about prototypicality have been confused with theories of processing—that is, there has been a failure to distinguish the structure of categories from theories concerning the use of that structure in processing. Therefore, let us first attempt to look at prototypes in as purely structural a fashion as possible. We will focus on what may be said about prototypes based on operational definitions and empirical findings alone without the addition of processing assumptions.

Perception of typicality differences is, in the first place, an empirical fact of people's judgments about category membership. It is by now a well-documented finding that subjects overwhelmingly agree in their judgments of how good an example or clear a case members are of a category, even for categories about whose boundaries they disagree (Rosch 1974, 1975b). Such

judgments are reliable even under changes of instructions and items (Rips, Shoben, and Smith 1973; Rosch 1975b, 1975c; Rosch and Mervis 1975). Were such agreement and reliability in judgment not to have been obtained, there would be no further point in discussion or investigation of the issue. However, given the empirical verification of degree of prototypicality, we can proceed to ask what principles determine which items will be judged the more prototypical and what other variables might be affected by prototypicality.

In terms of the basic principles of category formation, the formation of category prototypes should, like basic levels of abstraction, be determinate and be closely related to the initial formation of categories. For categories of concrete objects (which do not have a physiological basis, as categories such as colors and forms apparently do—Rosch 1974), a reasonable hypothesis is that prototypes develop through the same principles such as maximization of cue validity and maximization of category resemblance¹ as those principles governing the formation of the categories themselves.

In support of such a hypothesis, Rosch and Mervis (1975) have shown that the more prototypical of a category a member is rated, the more attributes it has in common with other members of the category and the fewer attributes in common with members of the contrasting categories. This finding was demonstrated for natural language superordinate categories, for natural language basic-level categories, and for artificial categories in which the definition of attributes and the amount of experience with items was completely specified and controlled. The same basic principles can be represented in ways other than through attributes in common. Because the present theory is a structural theory, one aspect of it is that centrality shares the mathematical notions inherent in measures like the mean and mode. Prototypical category members have been found to represent the means of attributes that have a metric, such as size (Reed 1972; Rosch, Simpson, and Miller 1976).

In short, prototypes appear to be just those members of a category that most reflect the redundancy structure of the category as a whole. That is, if categories form to maximize the information-rich cluster of attributes in the environment and, thus, the cue validity or category resemblance of the attributes of categories, prototypes of categories appear to form in such a manner as to maximize such clusters and such cue validity still further within categories.

It is important to note that for natural language categories both at the superordinate and basic levels, the extent to which items have attributes common to the category was highly negatively correlated with the extent to which they have attributes belonging to members of contrast categories. This appears to be part of the structure of real-world categories. It may be that such structure is given by the correlated clusters of attributes of the real world. Or such structure, may be a result of the human tendency once a contrast exists to define attributes for contrasting categories so that the categories will be maximally distinctive. In either case, it is a fact that both representativeness within a category and distinctiveness from contrast categories are correlated with prototypicality in real categories. For artificial categories, either principle alone will produce prototype effects (Rosch et al. 1976b; Smith and Balzano, personal communication) depending on the structure of the stimulus set. Thus to perform experiments to try to distinguish which principle is the *one* that deter-

mines prototype formation and category processing appears to be an artificial exercise.

Effects of Prototypicality on Psychological Dependent Variables

The fact that prototypicality is reliably rated and is correlated with category structure does not have clear implications for particular processing models nor for a theory of cognitive representations of categories (see the introduction to Part III of Rosch and Lloyd 1978 and Palmer 1978). What is very clear from the extant research is that the prototypicality of items within a category can be shown to affect virtually all of the major dependent variables used as measures in psychological research.

Speed of Processing: Reaction Time

The speed with which subjects can judge statements about category membership is one of the most widely used measures of processing in semantic memory research within the human information-processing framework. Subjects typically are required to respond true or false to statements of the form: *X* item is a member of *Y* category, where the dependent variable of interest is reaction time. In such tasks, for natural language categories, responses of true are invariably faster for the items that have been rated more prototypical. Furthermore, Rosch et al. (1976b) had subjects learn artificial categories where prototypicality was defined structurally for some subjects in terms of distance of a gestalt configuration from a prototype, for others in terms of means of attributes, and for still others in terms of family resemblance between attributes. Factors other than the structure of the category, such as frequency, were controlled. After learning was completed, reaction time in a category membership verification task proved to be a function of structural prototypicality.

Speed of Learning of Artificial Categories (Errors) and Order of Development in Children

Rate of learning of new material and the naturally obtainable measure of learning (combined with maturation) reflected in developmental order are two of the most pervasive dependent variables in psychological research. In the artificial categories used by Rosch et al. (1976b), prototypicality for all three types of stimulus material predicted speed of learning of the categories. Developmentally, Anglin (1976) obtained evidence that young children learn category membership of good examples of categories before that of poor examples. Using a category-membership verification technique, Rosch (1973) found that the differences in reaction time to verify good and poor members were far more extreme for 10-year-old children than for adults, indicating that the children had learned the category membership of the prototypical members earlier than that of other members.

Order and Probability of Item Output

Item output is normally taken to reflect some aspect of storage, retrieval, or category search. Battig and Montague (1969) provided a normative study of the probability with which college students listed instances of superordinate

semantic categories. The order is correlated with prototypicality ratings (Rosch 1975b). Furthermore, using the artificial categories in which frequency of experience with all items was controlled, Rosch et al. (1976b) demonstrated that the most prototypical items were the first and most frequently produced items when subjects were asked to list the members of the category.

Effects of Advance Information on Performance: Set, Priming

For colors (Rosch 1975c), for natural superordinate semantic categories (Rosch 1975b), and for artificial categories (Rosch et al. 1976b), it has been shown that degree of prototypicality determines whether advance information about the category name facilitates or inhibits responses in a matching task.

The Logic of Natural Language Use of Category Terms: Hedges, Substitutability into Sentences, Superordination in ASL

Although logic may treat categories as though membership is all or none, natural languages themselves possess linguistic mechanisms for coding and coping with gradients of category membership.

1. *Hedges.* In English there are qualifying terms such as “almost” and “virtually,” which Lakoff (1972) calls “hedges.” Even those who insist that statements such as “A robin is a bird” and “A penguin is a bird” are equally true, have to admit different hedges applicable to statements of category membership. Thus it is correct to say that a penguin is technically a bird but not that a robin is technically a bird, because a robin is more than just technically a bird; it is a real bird, a bird par excellence. Rosch (1975a) showed that when subjects were given sentence frames such as “X is virtually Y,” they reliably placed the more prototypical member of a pair of items into the referent slot, a finding which is isomorphic to Tversky’s work on asymmetry of similarity relations (Tversky & Gati 1978).

2. *Substitutability into sentences.* The meaning of words is intimately tied to their use in sentences. Rosch (1977) has shown that prototypicality ratings for members of superordinate categories predict the extent to which the member term is substitutable for the superordinate word in sentences. Thus, in the sentence “Twenty or so birds often perch on the telephone wires outside my window and twitter in the morning,” the term “sparrow” may readily be substituted for “bird” but the result turns ludicrous by substitution of “turkey,” an effect which is not simply a matter of frequency (Rosch 1975d).

3. *Productive superordinates in ASL.* Newport and Bellugi (1978) demonstrate that when superordinates in ASL are generated by means of a partial fixed list of category members, those members are the more prototypical items in the category.

In summary, evidence has been presented that prototypes of categories are related to the major dependent variables with which psychological processes are typically measured. What the work summarized does not tell us, however, is considerably more than it tells us. The pervasiveness of prototypes in real-

world categories and of prototypicality as a variable indicates that prototypes must have some place in psychological theories of representation, processing, and learning. However, prototypes themselves do not constitute any particular model of processes, representations, or learning. This point is so often misunderstood that it requires discussion:

1. To speak of *a prototype* at all is simply a convenient grammatical fiction; what is really referred to are judgments of degree of prototypicality. Only in some artificial categories is there by definition a literal single prototype (for example, Posner, Goldsmith, and Welton 1967; Reed 1972; Rosch et al. 1976b). For natural-language categories, to speak of a single entity that is the prototype is either a gross misunderstanding of the empirical data or a covert theory of mental representation.
2. Prototypes do not constitute any particular processing model for categories. For example, in pattern recognition, as Palmer (1978) points out, a prototype can be described as well by feature lists or structural descriptions as by templates. And many different types of matching operations can be conceived for matching to a prototype given any of these three modes of representation of the prototypes. Other cognitive processes performed on categories such as verifying the membership of an instance in a category, searching the exemplars of a category for the member with a particular attribute, or understanding the meaning of a paragraph containing the category name are not bound to any single process model by the fact that we may acknowledge prototypes. What the facts about prototypicality do contribute to processing notions is a constraint—process models should not be inconsistent with the known facts about prototypes. For example, a model should not be such as to predict equal verification times for good and bad examples of categories nor predict completely random search through a category.
3. Prototypes do not constitute a theory of representation of categories. Although we have suggested elsewhere that it would be reasonable in light of the basic principles of categorization, if categories were represented by prototypes that were most representative of the items in the category and least representative of items outside the category (Rosch and Mervis 1975; Rosch 1977), such a statement remains an unspecified formula until it is made concrete by inclusion in some specific theory of representation. For example, different theories of semantic memory can contain the notion of prototypes in different fashions (Smith, 1978). Prototypes can be represented either by propositional or image systems (see Kosslyn 1978 and Palmer 1978). As with processing models, the facts about prototypes can only constrain, but do not determine, models of representation. A representation of categories in terms of conjoined necessary and sufficient attributes alone would probably be incapable of handling all of the presently known facts, but there are many representations other than necessary and sufficient attributes that are possible.
4. Although prototypes must be learned, they do not constitute any particular theory of category learning. For example, learning of prototypicality

in the types of categories examined in Rosch and Mervis (1975) could be represented in terms of counting attribute frequency (as in Neuman 1974), in terms of storage of a set of exemplars to which one later matched the input (see Shepp 1978 and the introduction to Part II of Rosch and Lloyd 1978), or in terms of explicit teaching of the prototypes once prototypicality within a category is established in a culture (e.g., “Now that’s a *real* coat.”)

In short, prototypes only constrain but do not specify representation and process models. In addition, such models further constrain each other. For example, one could not argue for a frequency count of attributes in children’s learning of prototypes of categories if one had reason to believe that children’s representation of attributes did not allow for separability and selective attention to each attribute (see Garner 1978 and the introduction to Part II of Rosch and Lloyd 1978).

Two Problematical Issues

The Nature of Perceived Attributes

The derivations of basic objects and of prototypes from the basic principles of categorization have depended on the notion of a structure in the perceived world—bundles of perceived world attributes that formed natural discontinuities. When the research on basic objects and their prototypes was initially conceived (Rosch et al. 1976a), I thought of such attributes as inherent in the real world. Thus, given an organism that had sensory equipment capable of perceiving attributes such as wings and feathers, it was a fact in the real world that wings and feathers co-occurred. The state of knowledge of a person might be ignorant of (or indifferent or inattentive to) the attributes or might know of the attributes but be ignorant concerning their correlation. Conversely, a person might know of the attributes and their correlational structure but exaggerate that structure, turning partial into complete correlations (as when attributes true only of many members of a category are thought of as true of all members). However, the environment was thought to constrain categorizations in that human knowledge could not provide correlational structure where there was none at all. For purposes of the basic object experiments, perceived attributes were operationally defined as those attributes listed by our subjects. Shape was defined as measured by our computer programs. We thus seemed to have our system grounded comfortably in the real world.

On contemplation of the nature of many of the attributes listed by our subjects, however, it appeared that three types of attributes presented a problem for such a realistic view: (1) some attributes, such as “seat” for the object “chair,” appeared to have names that showed them not to be meaningful prior to knowledge of the object as chair; (2) some attributes such as “large” for the object “piano” seemed to have meaning only in relation to categorization of the object in terms of a superordinate category—piano is large for furniture but small for other kinds of objects such as buildings; (3) some attributes such as “you eat on it” for the object “table” were functional attributes that seemed to require knowledge about humans, their activities, and the real world in order

to be understood (see Miller 1978). That is, it appeared that the analysis of objects into attributes was a rather sophisticated activity that our subjects (and indeed a system of cultural knowledge) might well be considered to be able to impose only *after* the development of the category system.

In fact, the same laws of cognitive economy leading to the push toward basic-level categories and prototypes might also lead to the definition of attributes of categories such that the categories once given would appear maximally distinctive from one another and such that the more prototypical items would appear even more representative of their own and less representative of contrastive categories. Actually, in the evolution of the meaning of terms in languages, probably both the constraint of real-world factors and the construction and reconstruction of attributes are continually present. Thus, given a particular category system, attributes are defined such as to make the system appear as logical and economical as possible. However, if such a system becomes markedly out of phase with real-world constraints, it will probably tend to evolve to be more in line with those constraints—with redefinition of attributes ensuing if necessary. Unfortunately, to state the matter in such a way is to provide no clear place at which we can enter the system as analytical scientists. What is the unit with which to start our analysis? Partly in order to find a more basic real-world unit for analysis than attributes, we have turned our attention to the contexts in which objects occur—that is, to the culturally defined events in which objects serve as props.

The Role of Context in Basic-Level Objects and Prototypes

It is obvious, even in the absence of controlled experimentation, that a man about to buy a chair who is standing in a furniture store surrounded by different chairs among which he must choose will think and speak about chairs at other than the basic level of “chair.” Similarly, in regard to prototypes, it is obvious that if asked for the most typical African animal, people of any age will not name the same animal as when asked for the most typical American pet animal. Because interest in context is only beginning, it is not yet clear just what experimentally defined contexts will affect what dependent variables for what categories. But it is predetermined that there will be context effects for both the level of abstraction at which an object is considered and for which items are named, learned, listed, or expected in a category. Does this mean that our findings in regard to basic levels and prototypes are relevant only to the artificial situation of the laboratory in which a context is not specified?

Actually, both basic levels and prototypes are, in a sense, theories about context itself. The basic level of abstraction is that level of abstraction that is appropriate for using, thinking about, or naming an object in most situations in which the object occurs (Rosch et al. 1976a). And when a context is not specified in an experiment, people must contribute their own context. Presumably, they do not do so randomly. Indeed, it seems likely that, in the absence of a specified context, subjects assume what they consider the normal context or situation for occurrence of that object. To make such claims about categories appears to demand an analysis of the actual events in daily life in which objects occur.

The Role of Objects in Events

The attempt we have made to answer the issues of the origin of attributes and the role of context has been in terms of the use of objects in the events of daily human life. The study of events grew out of an interest in categorizations of the flow of experience. That is, our initial interest was in the question of whether any of the principles of categorization we had found useful for understanding concrete objects appeared to apply to the cutting up of the continuity of experience into the discrete bounded temporal units that we call *events*.

Previously, events have been studied primarily from two perspectives in psychology. Within ecological and social psychology, an observer records and attempts to segment the stream of another person's behavior into event sequences (for example, Barker and Wright 1955; Newton 1976). And within the artificial intelligence tradition, Story Understanders are being constructed that can "comprehend," by means of event scripts, statements about simple, culturally predictable sequences such as going to a restaurant (Shank 1975).

The unit of the event would appear to be a particularly important unit for analysis. Events stand at the interface between an analysis of social structure and culture and an analysis of individual psychology. It may be useful to think of scripts for events as the level of theory at which we can specify how culture and social structure enter the individual mind. Could we use events as the basic unit from which to derive an understanding of objects? Could we view objects as props for the carrying out of events and have the functions, perceptual attributes, and levels of abstraction of objects fall out of their role in such events?

Our research to date has been a study rather than an experiment and more like a pilot study at that. Events were defined neither by observation of others nor by a priori units for scripts but introspectively in the following fashion. Students in a seminar on events were asked to choose a particular evening on which to list the events that they remembered of that day—e.g., to answer the question what did I do? (or what happened to me?) that day by means of a list of the names of the events. They were to begin in the morning. The students were aware of the nature of the inquiry and that the focus of interest was on the units that they would perceive as the appropriate units into which to chunk the days' happenings. After completing the list for that day, they were to do the same sort of lists for events remembered from the previous day, and thus to continue backwards to preceding days until they could remember no more day's events. They also listed events for units smaller and larger than a day: for example, the hour immediately preceding writing and the previous school quarter.

The results were somewhat encouraging concerning the tractability of such a means of study. There was considerable agreement on the kinds of units into which a day should be broken—units such as making coffee, taking a shower, and going to statistics class. No one used much smaller units: That is, units such as picking up the toothpaste tube, squeezing toothpaste onto the brush, etc., never occurred. Nor did people use larger units such as "got myself out of the house in the morning" or "went to all my afternoon classes." Furthermore,

the units that were listed did not change in size or type with their recency or remoteness in time to the writing. Thus, for the time unit of the hour preceding writing, components of events were not listed. Nor were larger units of time given for a day a week past than for the day on which the list was composed. Indeed, it was dramatic how, as days further and further in the past appeared, fewer and fewer events were remembered although the type of unit for those that were remembered remained the same. That is, for a day a week past, a student would not say that he now only remembered getting himself out of the house in the morning (though such “summarizing” events could be inferred); rather he either did or did not remember feeding the cat that day (an occurrence that could also be inferred but for which inference and memory were introspectively clearly distinguishable). Indeed, it appeared that events such as “all the morning chores” as a whole do not have a memory representation separate from memory of doing the individual chores—perhaps in the way that superordinate categories, such as furniture, do not appear to be imageable *per se* apart from imaging individual items in the category. It should be noted that event boundaries appeared to be marked in a reasonable way by factors such as changes of the actors participating with ego, changes in the objects ego interacts with, changes in place, and changes in the type or rate of activity with an object, and by notable gaps in time between two reported events.

A good candidate for the basic level of abstraction for events is the type of unit into which the students broke their days. The events they listed were just those kinds of events for which Shank (1975) has provided scripts. Scripts of events analyze the event into individual units of action; these typically occur in a predictable order. For example, the script for going to a restaurant contains script elements such as entering, going to a table, ordering, eating, and paying. Some recent research has provided evidence for the psychological reality of scripts and their elements (Bower 1976).

Our present concern is with the role of concrete objects in events. What categories of objects are required to serve as props for events at the level of abstraction of those listed by the students? In general, we found that the event name itself combined most readily with superordinate noun categories; thus, one gets dressed with clothes and needs various kitchen utensils to make breakfast. When such activities were analyzed into their script elements, the basic level appeared as the level of abstraction of objects necessary to script the events; e.g., in getting dressed, one puts on pants, sweater, and shoes, and in making breakfast, one cooks eggs in a frying pan.

With respect to prototypes, it appears to be those category members judged the more prototypical that have attributes that enable them to fit into the typical and agreed upon script elements. We are presently collecting normative data on the intersection of common events, the objects associated with those events and the other sets of events associated with those objects.² In addition, object names for eliciting events are varied in level of abstraction and in known prototypicality in given categories. Initial results show a similar pattern to that obtained in the earlier research in which it was found that the more typical members of superordinate categories could replace the superordinate in sentence frames generated by subjects told to “make up a sentence” that used the

superordinate (Rosch 1977). That is, the task of using a given concrete noun in a sentence appears to be an indirect method of eliciting a statement about the events in which objects play a part; that indirect method showed clearly that prototypical category members are those that can play the role in events expected of members of that category.

The use of deviant forms of object names in narratives accounts for several recently explored effects in the psychological literature. Substituting object names at other than the basic level within scripts results in obviously deviant descriptions. Substitution of superordinates produces just those types of narrative that Bransford and Johnson (1973) have claimed are not comprehended; for example, "The procedure is actually quite simple. First you arrange things into different groups. Of course, one pile may be sufficient [p. 400]." It should be noted in the present context that what Bransford and Johnson call context cues are actually names of basic-level events (e.g., washing clothes) and that one function of hearing the event name is to enable the reader to translate the superordinate terms into basic-level objects and actions. Such a translation appears to be a necessary aspect of our ability to match linguistic descriptions to world knowledge in a way that produces the "click of comprehension."

On the other hand, substitution of subordinate terms for basic-level object names in scripts gives the effect of satire or snobbery. For example, a review (Garis 1975) of a pretentious novel accused of actually being about nothing more than brand-name snobbery concludes, "And so, after putting away my 10-year-old Royal 470 manual and lining up my Mongol number 3 pencils on my Goldsmith Brothers Formica imitation-wood desk, I slide into my over-size squirrel-skin L. L. Bean slippers and shuffle off to the kitchen. There, holding *Decades* in my trembling right hand, I drop it, *plunk*, into my new Sears 20-gallon, celadon-green Permanex trash can [p. 48]."

Analysis of events is still in its initial stages. It is hoped that further understanding of the functions and attributes of objects can be derived from such an analysis.

Summary

The first part of this chapter showed how the same principles of categorization could account for the taxonomic structure of a category system organized around a basic level and also for the formation of the categories that occur within this basic level. Thus the principles described accounted for both the vertical and horizontal structure of category systems. Four converging operations were employed to establish the claim that the basic level provides the cornerstone of a taxonomy. The section on prototypes distinguished the empirical evidence for prototypes as structural facts about categories from the possible role of prototypes in cognitive processing, representation, and learning. Then we considered assumptions about the nature of the attributes of real-world objects and assumptions about context—insofar as attributes and contexts underlie the claim that there is structure in the world. Finally, a highly tentative pilot study of attributes and functions of objects as props in culturally defined events was presented.

Notes

1. Tversky formalizes prototypicality as the member or members of the category with the highest summed similarity to all members of the category. This measure, although formally more tractable than that of cue validity, does not take account, as cue validity does, of an item's dissimilarity to contrast categories. This issue is discussed further later.
2. This work is being done by Elizabeth Kreusi.

References

- Anglin, J. Les premiers termes de référence de l'enfant. In S. Ehrlich and E. Tulving (Eds.), *La mémoire sémantique*. Paris: Bulletin de Psychologie, 1976.
- Barker, R., and Wright, H. *Midwest and its children*. Evanston, Ill.: Row-Peterson, 1955.
- Battig, W. F., and Montague, W. E. Category norms for verbal items in 56 categories: A replication and extension of the Connecticut category norms. *Journal of Experimental Psychology Monograph*, 1969, 80 (3, Pt. 2).
- Beach, L. R. Cue probabilism and inference behavior. *Psychological Monographs*. 1964, 78 (Whole No. 582). (a)
- Beach, L. R. Recognition, assimilation, and identification of objects. *Psychological Monographs*, 1964, 78 (Whole No. 583). (b)
- Berlin, B. Speculations on the growth of ethnobotanical nomenclature. *Language in Society*, 1972, 1, 51–86.
- Berlin, B. Ethnobiological classification. In E. Rosch and B. B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1978.
- Borges, J. L. *Other inquiries 1937–1952*. New York: Washington Square Press, 1966.
- Bower, G. *Comprehending and recalling stories*. Paper presented as Division 3 presidential address to the American Psychological Association, Washington, D.C., September 1976.
- Bransford, J. D., and Johnson, M. K. Considerations of some problems of comprehension. In W. Chase (Ed.), *Visual information processing*. New York: Academic Press, 1973.
- Bruner, J. S., Olver, R. R., and Greenfield, P. M. *Studies in cognitive growth*. New York: Wiley, 1966.
- Garis, L. The Margaret Mead of Madison Avenue. *Ms.*, March 1975, pp. 47–48.
- Garner, W. R. *The processing of information and structure*. New York: Wiley, 1974.
- Garner, W. R. Aspects of a stimulus: Features, dimensions, and configurations. In E. Rosch and B. B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1978.
- Kay, P. Taxonomy and semantic contrast. *Language*, 1971, 47, 866–887.
- Kosslyn, S. M. Imagery and internal representation. In E. Rosch and B. B. Lloyd (Eds.), *Cognition and Categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1978.
- Lakoff, G. Hedges: A study in meaning criteria and the logic of fuzzy concepts. *Papers from the eighth regional meeting, Chicago Linguistics Society*. Chicago: University of Chicago Linguistics Department, 1972.
- Miller, G. A. Practical and lexical knowledge. In E. Rosch and B. B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1978.
- Nelson, K. Concept, word and sentence: Interrelations in acquisition and development. *Psychological Review*, 1974, 81, 267–285.
- Neuman, P. G. An attribute frequency model for the abstraction of prototypes. *Memory and Cognition*, 1974, 2, 241–248.
- Newport, E. L., and Bellugi, U. Linguistic expression of category levels in a visual-gestural language: A flower is a flower is a flower. In E. Rosch and B. B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1978.
- Newton, D. Foundations of attribution: The perception of ongoing behavior. In J. Harvey, W. Ickes, and R. Kidd (Eds.), *New directions in attribution research*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1976.
- Palmer, S. Hierarchical structure in perceptual representation. *Cognitive Psychology*, 1977, 9, 441–474.
- Palmer, S. E. Fundamental aspects of cognitive representation. In E. Rosch and B. B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1978.

- Peterson, M. J., and Graham, S. E. Visual detection and visual imagery. *Journal of Experimental Psychology*, 1974, 103, 509–514.
- Piaget, J. *The origins of intelligence in children*. New York: International Universities Press, 1952.
- Posner, M. I., Goldsmith, R., and Welton, K. E. Perceived distance and the classification of distorted patterns. *Journal of Experimental Psychology*, 1967, 73, 28–38.
- Reed, S. K. Pattern recognition and categorization. *Cognitive Psychology*, 1972, 3, 382–407.
- Rips, L. J., Shoben, E. J., and Smith, E. E. Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior*, 1973, 12, 1–20.
- Rosch, E. On the internal structure of perceptual and semantic categories. In T. E. Moore (Ed.), *Cognitive development and the acquisition of language*. New York: Academic Press, 1973.
- Rosch, E. Linguistic relativity. In A. Silverstein (Ed.), *Human communication: Theoretical perspectives*. New York: Halsted Press, 1974.
- Rosch, E. Cognitive reference points. *Cognitive Psychology*, 1975, 7, 532–547. (a)
- Rosch, E. Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 1975, 104, 192–233. (b)
- Rosch, E. The nature of mental codes for color categories. *Journal of Experimental Psychology: Human Perception and Performance*, 1975, 1, 303–322. (c)
- Rosch, E. Universals and cultural specifics in human categorization. In R. Brislin, S. Bochner, and W. Lonner (Eds.), *Cross-cultural perspectives on learning*. New York: Halsted Press, 1975. (d)
- Rosch, E. Human categorization. In N. Warren (Ed.), *Advances in cross-cultural psychology* (Vol. 1). London: Academic Press, 1977.
- Rosch, E., and Lloyd, B. B. *Cognition and categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1978.
- Rosch, E., and Mervis, C. B. Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 1975, 7, 573–605.
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., and Boyes-Braem, P. Basic objects in natural categories. *Cognitive Psychology*, 1976, 8, 382–439. (a)
- Rosch, E., Simpson, C., and Miller, R. S. Structural bases of typicality effects. *Journal of Experimental Psychology: Human Perception and Performance*. 1976, 2, 491–502. (b)
- Shank, R. C. The structure of episodes in memory. In D. G. Bobrow and A. Collins (Eds.), *Representation and understanding: Studies in cognitive science*. New York: Academic Press, 1975.
- Shepp, B. E. From perceived similarity to dimensional structure: A new hypothesis about perspective development. In E. Rosch and B. B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1978.
- Smith, E. E. Theories of semantic memory. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes* (Vol. 5). Hillsdale, N.J.: Lawrence Erlbaum Associates, 1978.
- Smith, E. E., and Balzano, G. J. Personal communication, April 1977.
- Smith, E. E., Balzano, G. J., and Walker, J. H. Nominal, perceptual, and semantic codes in picture categorization. In J. Cotton and R. Klatzky (Eds.), *Semantic factors in cognition*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1978.
- Tversky, A. Features of similarity. *Psychological Review*, 1977, 84, 327–352.
- Tversky, A., and Gati, I. Studies of similarity. In E. Rosch and B. B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1978.
- Wittgenstein, L. *Philosophical investigations*. New York: Macmillan, 1953.

This excerpt from

The Motion Aftereffect.

George Mather, Frans Verstraten and Stuart Anstis, editors.

© 1998 The MIT Press.

is provided in screen-viewable form for personal use only by members of MIT CogNet.

Unauthorized use or dissemination of this information is expressly forbidden.

If you have any questions about this material, please contact cognetadmin@cognet.mit.edu.