fective in introducing information without calling attention to it. Obviously, there are many other ways to introduce new information. The experimental manipulation of subsequent information may constitute a useful technique for investigating the interaction of a person's specific experiences and subsequent knowledge related to those experiences.

REFERENCES

Anderson, J. R., & Bower, G. H. *Human Associative Memory*. Washington, DC: Winston, 1973.

Bartlett, F. C. *Remembering: A study in experimental and social psychology*. London: Cambridge, University Press, 1932.

Bransford, J. D. & Johnson, M. K. Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning and Verbal Behavior*, 1972, **11**, 717–726.

Clark, H. H., & Haviland, S. E. Psychological processes as linguistic explanation. In D. Cohen (Ed.), *The nature of explanation in linguistics*. Milwaukee: University of Wisconsin Press, in press.

Dooling, D. J., & Lachman, R. Effects of comprehension on retention of prose. *Journal of Experimental Psychology*, 1971, **88**, 216–222.

Harris, R. J. Answering questions containing marked and unmarked adjectives and adverbs. *Journal of Experimental Psychology*, 1973, **97**, 399–401.

Loftus, E. F. Reconstructing memory. The incredible eyewitness . *Psychology Today*, 1974, **8**, 116–119.

Loftus, E. F., & Palmer, J. C. Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior*, 1974, **13**, 585–589.

Loftus, E. F., & Zanni, G. Eyewitness testimony: The influence of the wording of a question. *Bulletin of the Psychonomic Society*, 1975, **5**, 86–88.

Rumelhart, D. E., Lindsay, P. H., & Norman, D. A. A process model of long-term memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory*. New York: Academic Press, 1972.

Rumelhart, D. E., & Norman, D. A. *Active semantic networks as a model of human memory*. Proceedings of the Third International Joint Conference on Artificial Intelligence, Stanford University, 1973.

Shepard, R. N. Learning and recall as organization and search. *Journal of Verbal Learning and Verbal Behavior*, 1966, **5**, 201–204.

Selfridge, O. G. & Neisser, U. Pattern recognition by machine. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought*. New York: McGraw Hill, 1963.

Tulving, E. & Thomson, D. M. Encoding specificity and retrieval processes in episodic memory. *Psychological Review*, 1973, **80**, 352–373.

Winograd, T. Understanding natural language. *Cognitive Psychology*, 1972, **3**, 1–191.

(Accepted May 6, 1975)

COGNITIVE PSYCHOLOGY **7**, 573–605 (1975)

Family Resemblances: Studies in the Internal Structure of Categories

ELEANOR ROSCH AND CAROLYN B. MERVIS

University of California, Berkeley

Six experiments explored the hypothesis that the members of categories which are considered most prototypical are those with most attributes in common with other members of the category and least attributes in common with other categories. In probabilistic terms, the hypothesis is that prototypicality is a function of the total cue validity of the attributes of items. In Experiments 1 and 3, subjects listed attributes for members of semantic categories which had been previously rated for degree of prototypicality. High positive correlations were obtained between those ratings and the extent of distribution of an item's attributes among the other items of the category. In Experiments 2 and 4, subjects listed superordinates of category members and listed attributes of members of contrasting categories. Negative correlations were obtained between prototypicality and superordinates other than the category in question and between prototypicality and an item's possession of attributes possessed by members of contrasting categories. Experiments 5 and 6 used artificial categories and showed that family resemblance within categories and lack of overlap of elements with contrasting categories were correlated with ease of learning, reaction time in identifying an item after learning, and rating of prototypicality of an item. It is argued that family resemblance offers an alternative to criterial features in defining categories.

As speakers of our language and members of our culture, we know that a chair is a more reasonable exemplar of the category *furniture* than a radio, and that some chairs fit our idea or image of a chair better than others. However, when describing categories analytically, most traditions of thought have treated category membership as a digital, all-or-none phenomenon. That is, much work in philosophy, psychology, linguistics, and anthropology assumes that categories are logical bounded entities, membership in which is defined by an item's posses-

This research was supported by grants to the first author (under her former name Eleanor Rosch Heider) by the National Science Foundation (GB-38245X), by the Grant Foundation, and by the National Institutes of Mental Health (1 R01 MH24316-01). We wish to thank David Johnson, Joseph Romeo, Ross Quigley, R. Scott Miller, Steve Frank, Alina Furnow, and Louise Jones for help with testing and analysis of the data. We also wish to thank Ed Smith, Ed Shoben, and Lance Rips for permission to refer to the multidimensional scaling study of superordinate categories which was performed jointly with them. Carolyn Mervis is now at Cornell Universtiy. She was supported by an NSF Predoctoral Fellowship during the research. Requests for reprints should be sent to Eleanor Rosch, Department of Psychology, University of California at Berkeley, Berkeley, CA 94720.

Copyright © 1975 by Academic Press, Inc.

sion of a simple set of criterial features, in which all instances possessing the criterial attributes have a full and equal degree of membership.

In contrast to such a view, it has been recently argued (see Lakoff, 1972; Rosch, 1973; Zadeh, 1965) that some natural categories are analog and must be represented logically in a manner which reflects their analog structure. Rosch (1973, 1975b) has further characterized some natural analog categories as internally structured into a prototype (clearest cases, best examples of the category) and nonprototype members, with nonprototype members tending toward an order from better to poorer examples. While the domain for which such a claim has been demonstrated most unequivocally is that of color (Berlin & Kay, 1969; Heider, 1971, 1972; Mervis, Catlin, & Rosch, 1975; Rosch, 1974, in press-c, in press-d), there is also considerable evidence that natural superordinate semantic categories have a prototype structure. Subjects can reliably rate the extent to which a member of a category fits their idea or image of the meaning of the category name (Rosch, 1973, 1975a), and such ratings predict performance in a number of tasks (Rips, Shoben & Smith, 1973; Rosch, 1973, 1975a, in press-c, 1975b; Smith, Rips, & Shoben, 1974; Smith, Shoben, & Rips, 1974).

However, there has, as yet, been little attention given to the problem of how internal structure arises. That is, what principles govern the formation of category prototypes and gradients of category membership? For some categories which probably have a physiological basis, such as colors, forms, and facial expressions of basic human emotions, prototypes may be stimuli which are salient prior to formation of the category, whose salience, at the outset, determines the categorical structuring of those domains (Ekman, 1971; McDaniel, Note 1; Rosch, 1974, 1975b). For the artificial categories which have been used in prototype research—such as families of dot patterns (Posner, 1973) and artificial faces (Reed, 1972)—the categories have been intentionally structured and/or the prototypes have been defined so that the prototypes were central tendencies of the categories. For most domains, however, prototypes do not appear to precede the category (Rosch, in press-a) and must be formed through principles of learning and information processing from the items given in the category. The present research was not intended to provide a processing model of the learning of categories or formation of prototypes; rather, our intention was to examine the stimulus relations which underlie such learning. That is, the purpose of the present research was to explore one of the major *structural* principles which, we believe, may govern the formation of the prototype structure of semantic categories.

This principle was first suggested in philosophy; Wittgenstein (1953) argued that the referents of a word need not have common elements in order for the word to be understood and used in the normal functioning

of language. He suggested that, rather, a family resemblance might be what linked the various referents of a word. A family resemblance relationship consists of a set of items of the form AB, BC, CD, DE. That is, each item has at least one, and probably several, elements in common with one or more other items, but no, or few, elements are common to all items. The existence of such relationships in actual natural language categories has not previously been investigated empirically.

In the present research, we viewed natural semantic categories as networks of overlapping attributes; the basic hypothesis was that members of a category come to be viewed as prototypical of the category as a whole in proportion to the extent to which they bear a family resemblance to (have attributes which overlap those of) other members of the category. Conversely, items viewed as most prototypical of one category will be those with least family resemblance to or membership in other categories. In natural categories of concrete objects, the two aspects of family resemblance should coincide rather than conflict since it is reasonable that categories tend to become organized in such a way that they reflect the correlational structure of the environment in a manner which renders them maximally discriminable from each other (Rosch, in press-a; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, in press).

The present structural hypothesis is closely related to a *cue validity* processing model of classification in which the validity of a cue is defined in terms of its total frequency within a category and its proportional frequency in that category relative to contrasting categories. Mathematically, cue validity has been defined as a conditional probability—specifically, the frequency of a cue being associated with the category in question divided by the total frequency of that cue over all relevant categories (Beach, 1964; Reed, 1972). Unfortunately, cue validity has been treated as a model in conflict with a prototype model of category processing where prototypes are operationally defined solely as attribute means (Reed, 1972). If prototypes are defined more broadly—for example, as the abstract representation of a category, or as those category members to which subjects compare items when judging category membership, or as the internal structure of the category defined by subjects' judgments of the degree to which members fit their "idea or image" of the category—then prototypes should coincide rather than conflict with cue validity. That is, if natural categories of concrete objects tend to become organized so as to render the categories maximally discriminable from each other, it follows that the maximum possible cue validity of items within each category will be attained (Rosch *et al.*, in press). The principle of family resemblance relationships can be restated in terms of cue validity since the attributes most distributed among members of a category and least distributed among members of contrasting categories are, by definition, the most valid cues to membership in the category in question. We use the term *family resemblance* rather than *cue validity* primarily to emphasize that we are dealing with a description of structural principles and not with a processing model. We believe that the principle of family resemblance relationships is a very general one and is applicable to categories regardless of whether or not they have features common to members of the category or formal criteria for category membership.

In all of the studies of the present research, family resemblances were defined in terms of discrete attributes such as *has legs, you drive it,* or *the letter B is a member*. These are the kinds of features of natural semantic categories which can be most readily reported and the features normally used in definitions of categories by means of lists of formal criteria. Insofar as the context in which an attribute occurs as part of a stimulus may always affect perception and understanding of the attribute, discrete attributes of this type may be an analytic myth. However, in one sense, the purpose of the present research was to show that it is not necessary to invoke attribute interactions or higher order gestalt properties of stimuli (such as those used by Posner, 1973; Reed, 1972; Rosch, Simpson, & Miller, Note 2) in order to analyze the prototype structure of categories. That is, even at the level of analysis of the type of discrete attributes normally used in definitions of categories by means of criterial features, we believe there is a principle of the structure of stimulus sets, family resemblances, which can be shown to underlie category prototype structure.

The present paper reports studies using three different types of category; superordinate semantic categories such as *furniture* and *vehicle,* basic level semantic categories such as *chair* and *car,* and artificial categories formed from sets of letter strings. For each type of stimulus, both aspects of the family resemblance hypothesis (that the most prototypical members of categories are those with most attributes in common with other members of that category and are those with least attributes in common with other categories) were tested.

Superordinate semantic categories are of particular interest because they are sufficiently abstract that they have few, if any, attributes common to all members (Rosch *et al.,* in press). Thus, such categories may consist almost entirely of items related to each other by means of family resemblances of overlapping attributes. In addition, superordinate categories have the advantage that their membership consists of a finite number of names of basic level categories which can be adequately sampled. Superordinate categories have the disadvantage that they do not have contrasting categories (operationally defined below); thus, the second half of the family resemblance hypothesis (that prototypical members of categories have least resemblance to other categories) had

to be tested indirectly by measuring membership in, rather than attributes in common with, other superordinate categories.

Basic level semantic categories are of great interest because they are the level of abstraction at which the basic category cuts in the world may be made (Rosch, in press-a; Rosch *et al.,* in press). However, basic level categories present a sampling problem since their membership consists of an infinite number of objects. On the positive side, basic level categories do form contrast sets, thus, making possible a direct test of the second part of the family resemblance hypothesis.

Artificial categories were needed because they made possible the study of prototype formation with adequate controls. In natural language domains of any type, categories have long since evolved in culture and been learned by subjects. Both prototypes and the attribute structure of categories are independent variables; we can only measure their correlations. Artificial categories are of use because attribute structures can be varied in a controlled manner and the development of prototypes studied as a dependent variable.

PART I: SUPERORDINATE SEMANTIC CATEGORIES

Experiment 1

Although it is always possible for an ingenious philosopher or psychologist to invent criterial attributes defining a category, earlier research has shown that actual subjects rate superordinate semantic categories as having few, if any, attributes common to all members (Rosch *et al.,* in press). Thus, if the "categorical" nature of these categories is to be explained, it appeared most likely to reside in family resemblances between members. Part of the purpose of the present experiment was to obtain portraits of the distribution of attributes of members of a number of superordinate natural language categories. Part of the hypothesis was that category members would prove to bear a family resemblance relationsip to each other. The major purpose of the experiment, however, was to observe the relation between degree of relatedness between members of the category and the rated prototypicality of those members. The specific hypothesis was that a measure of the degree to which an item bore a family resemblance to other members of the category would prove significantly correlated with previously obtained prototypicality ratings of the members of the category.

Method

Subjects. Subjects were 400 students in introductory psychology classes who received this 10 min task as part of their classroom work.

Stimuli. The categories used were the six most common categories of concrete nouns in English, determined by a measure of word frequency

(Kucera & Francis, 1967). All of the categories were ones for which norms for the prototypicality of items had already been obtained for 50–60 category members (Rosch, 1975a). These norms were derived from subjects' ratings of the extent to which each item fit their "idea or image" of the meaning of the category name. (The rating task and instructions were very similar to those used in Experiment 3 of the present research. A complete account of the methods for deriving the six superordinate categories and complete norms for all items of the six categories are provided in Rosch, in press-d.) The 20 items from each category used in the present experiment were chosen to represent the full range of goodness-of-example ranks. These items are listed, in their goodness-of-example order, in Table 1.

Procedure. Each of the 120 items shown in Table 1 was printed at the top of a page, and the pages assembled into packets consisting of six items, one from each superordinate category. Items were chosen randomly within a category such that each subject who received an item received it with different items from the other five categories and received the items representing each category in a different order. Each item was rated by 20 subjects. Each subject rated six items, one from each category.

Subjects were asked to list the attributes possessed by each item. Instructions were:

> This is a very simple experiment to find out the characteristics and attributes that people feel are common to and characteristic of different kinds of ordinary everyday objects. For example, for *bicycles* you might think of things they have in common like two wheels, pedals, handlebars, you ride on them, they don't use fuel, etc. For *dogs* you might think of things they have in common like having four legs, barking, having fur, etc.
>
> There are six pages following this one. At the top of each is listed the name of one common object. For each page, you'll have a minute and a half to write down all of the attributes of that object that you can think of. But try not to *just* free associate—for example, if bicycles just happen to remind you of your father, *don't* write down *father*.
>
> Okay—you'll have a minute and a half for each page. When I say turn to the next page, read the name of the object and write down the attributes or characteristics you think are characteristic of that object as fast as you can until you're told to turn the page again.

Measurement of family resemblance. To derive the basic measure of family resemblance, for each category, all attributes mentioned by subjects were listed and each item, for which an attribute had been listed, was credited with that attribute. Two judges reviewed the resulting table and indicated cases in which an attribute was clearly and obviously false. These attributes were deleted from the tabulation. The judges also indicated any attribute which had been listed for one or more items, but was clearly and obviously true of another item in the category for which

TABLE 1
SUPERORDINATE CATEGORIES AND ITEMS USED IN EXPERIMENTS 1 AND 2

	Category					
Item	Furniture	Vehicle	Fruit	Weapon	Vegetable	Clothing
1	Chair	Car	Orange	Gun	Peas	Pants
2	Sofa	Truck	Apple	Knife	Carrots	Shirt
3	Table	Bus	Banana	Sword	String beans	Dress
4	Dresser	Motorcycle	Peach	Bomb	Spinach	Skirt
5	Desk	Train	Pear	Hand grenade	Broccoli	Jacket
6	Bed	Trolley car	Apricot	Spear	Asparagus	Coat
7	Bookcase	Bicycle	Plum	Cannon	Corn	Sweater
8	Footstool	Airplane	Grapes	Bow and arrow	Cauliflower	Underpants
9	Lamp	Boat	Strawberry	Club	Brussel sprouts	Socks
10	Piano	Tractor	Grapefruit	Tank	Lettuce	Pajamas
11	Cushion	Cart	Pineapple	Teargas	Beets	Bathing suit
12	Mirror	Wheelchair	Blueberry	Whip	Tomato	Shoes
13	Rug	Tank	Lemon	Icepick	Lima beans	Vest
14	Radio	Raft	Watermelon	Fists	Eggplant	Tie
15	Stove	Sled	Honeydew	Rocket	Onion	Mittens
16	Clock	Horse	Pomegranate	Poison	Potato	Hat
17	Picture	Blimp	Date	Scissors	Yam	Apron
18	Closet	Skates	Coconut	Words	Mushroom	Purse
19	Vase	Wheelbarrow	Tomato	Foot	Pumpkin	Wristwatch
20	Telephone	Elevator	Olive	Screwdriver	Rice	Necklace