Is There a Default Similarity Distance for Categories?

Yaakov Kareev and Judith Avrahami
The Goldie Rotman Center for
Cognitive Science in Education
School of Education
The Hebrew University of Jerusalem

Abstract

How do people decide whether or not an item belongs to a new category, the variability of which they do not know? We postulate that people have a default similarity distance (DSD) which they use when no other information about the variability of a category is available. To test our claim, subjects were asked to tell how they would instruct a being from another world to distinguish members of a category, by showing pictures. The categories were from different levels thus differing in variability. For highly variable categories subjects tended to present multiple positive instances (thus indicating their extraordinary variability), whereas for narrow categories they tended to present negative instances (thus explicitly delimiting them). These results indicated that a norm, relative to which additional information is supplied, lay in between. Indeed, there was a level at which subjects apparently relied on DSD, finding it sufficient to show but a single exemplar of the category. This happened with basic-level categories for 8th graders and adults and with subordinate categories for 2nd graders, thus demonstrating a developmental trend in what is considered a normal standard category.

Dealing with the classification of new items into categories, researchers have focused mainly on what it is that a new item is compared with to decide whether it belongs to a particular category or not. That entity is variously claimed to be a prototype (Rosch, 1973; Rosch & Mervis, 1975), a set of all properties of members of the category (Hayes-Roth & Hayes-Roth, 1977), or a collection of all exemplars of the category hitherto encountered (Medin & Shaffer, 1978). All theories agree, however, that the new item need not be identical to any of the above; it just has to be sufficiently similar or, in Medin and Barsalou's (1987) terms "above a certain threshold" of similarity.

It is clear that there is no single value which defines "sufficient similarity," since lower similarity is allowed between items of a higher level category than between items of a lower level category. Thus, an object has to

be more similar to other bulldogs to be considered a bulldog than to other animals to still be considered an animal. Indeed, Fried and Holyoak (1984) suggest that the representation of a category includes both a mean value of the category and an indication of the density of its exemplars in a feature space. This density can be viewed as the variability or similarity distance allowed and expected between exemplars of the category.

To better understand classification we still need some explanation for situations in which classification is based on a single item or in which categories are formed with no feedback. We suggest that people have a notion of a "proper" distance or "plausible" similarity which serves as the threshold mentioned above. It is a default similarity distance (DSD), used when no information about the variability of the category or about its neighboring categories is available. We expect it to be some middle value, close to the mean similarity distance of categories known to the subject. Whatever the initial similarity distance assumed, it is continuously updated following subsequent encounters with exemplars which are known to belong to the category in question and ones that do not.

To test our hypothesis that people have such a DSD we had our subjects tell how they would teach a creature from another world to identify members of a certain category by showing it pictures. If there is no DSD one would expect subjects to provide for <u>any</u> category not only a representative member of it but also some indication of its variability, or allowed similarity distance. If, on the other hand, subjects take DSD into account and assume others to share it, they would see a lesser need to indicate the variability of categories for which DSD is more appropriate; for those categories, they may consider it sufficient to present but one typical exemplar.

The argument, then, goes as follows: for categories whose similarity distance is close to that denoted by DSD, subjects will be more likely to be

satisfied with a single representative exemplar of the category than for categories where DSD is inadequate. In the latter case subjects will be more likely to provide additional information to indicate the variability of the category.

How can the variability of a category be indicated? To indicate that the variability of a category is greater than the expected value, multiple, various exemplars belonging to that category can be used. When the variability of a category is narrower than the expected value, negative exemplars can be used - items which do not belong to the category in question but would belong if it were a category with the default distance. To insure a wide range of variabilities, we included categories of different levels: basic-level, superordinate, and subordinate.

To find out whether DSD changes with age, we employed subjects of different ages.

Method

Design. The study had a two-way factorial design with the variables of age and level of category.

Subjects. The sample consisted of 122 2nd graders (aged 7(3) to 8(2)) and 187 8th graders (aged 13(3) to 14(2)) from middle-class neighborhoods and 110 undergraduate students attending an introductory course in statistics at the Hebrew University.

Materials. The following items were used in the study: a) Superordinate categories: animal, plant, means of transportation; b) Basic-level categories: dog, bird, tree, mushroom, car, boat; c) Subordinate categories: bulldog, dachshund, wagtail, oal, cypress, fir, sedan-car, sports-car, sailing-boat, pedal-boat. All items are well known to Israeli children of the ages included in the study. The number of subjects at the three levels were: 2nd graders -

20, 37, 65; 8th graders - 27, 62, 98; college students - 22, 38, 50.

Procedure. The task was administered in groups. Each subject was handed a sheet of paper with the following instructions:

"Imagine that creatures from another world, which are very much like human beings, have landed on Earth. You have to instruct them, through the use of pictures, to identify things (such as a chair, a carrot, a tool). Imagine that you have at your disposal a collection of pictures which can include any picture you wish.

Which pictures would you choose to show them, so that when they encounter an object they will know whether or not it is a _____ ?"

The blank line was completed with one of the 19 items mentioned above.

Results and Discussion

Classification of answers. For the purposes of the present paper each answer list was characterized by its values on two dimensions:

a) Mentioning of multiple positive instances of the category: A list either contained multiple positive instances of the category (MP+) or did not (MP-).

b) Mentioning of negative instances of the category: A list either contained one or more negative instance (N+) or did not contain any negative instance (N-). A list was regarded as N+ if it contained at least one item from a category which shares the same immediate superordinate category with the item in question (for example, for bird "I'll show a picture of a butterfly and cross it out"; for bulldog, "I'll show a picture of a German shepherd").

The answers were independently evaluated by two judges. Agreement was very high (over 98%). The four resulting patterns of answers are discussed below.

a) (MP-/N-): Absence of multiple positive instances and absence of negative instances. This pattern was understood to mean that the subject relied on the DSD to indicate the variability of the category in question.

b) (MP-/N+): Absence of multiple positive instances and presence ofnegative instances. This pattern was understood to mean that the variability of the category in question was smaller than what the subject considered to be the norm, and had to be explicitly delimited by presenting negative instances.

c) (MP+/N-): Presence of multiple positive instances and absence of negative instances. This pattern was understood to mean that the variability of the category in question was larger than what the subject considered to be the norm, and had to be explicilty expanded by presenting more than one positive instance.

d) (MP+/N+): Presence of both multiple positive instances and negative instances. Here nothing can be inferred concerning the variability of the category in question relative to DSD.

Analysis. A corrolary of the claim that there exists a DSD for categories, is that category variability is explicitly indicated when DSD is inappropriate. Since the higher the level of a category the larger its variability, we expected the incidence of the MP+/N- pattern to increase with level. The opposite was expected for the MP-/N+ pattern: We expected it to be more prevalent the lower the category level. The relationship between category level and the incidence of the two patterns is depicted in Figure 1.

Since support for these two predictions is a prerequisite for any further analysis we first tested the main effect of level for the two patterns of responses in question. The analyses revealed a highly significant main effect of level both for the MP+/N- pattern of answer ($F(2,410) = 23.93$ $p < .001$) and for the MP-/N+ pattern of answer ($F(2,410) = 16.00$, $p < .001$)).

These two results establish that subjects performing the task were sensitive to category variability: The greater the variability the more likely they were to use multiple positive instances; the smaller the variability the more likely they were to use negative instances.

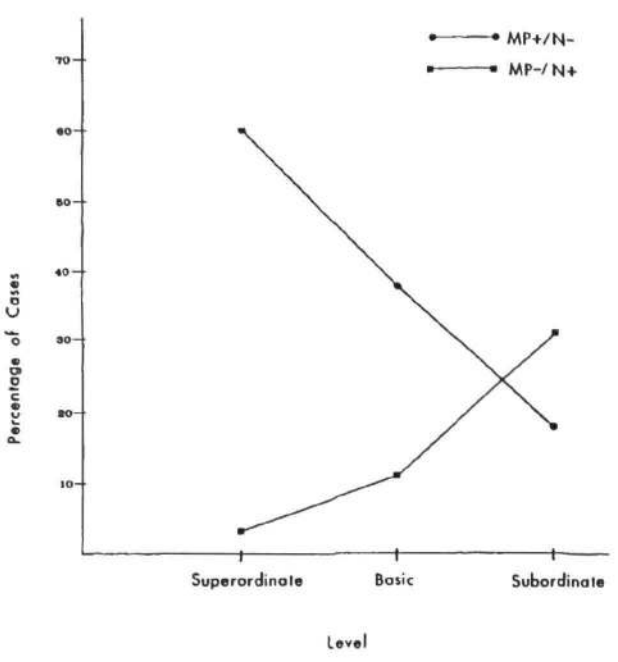

Figure 1: Incidence of MP+/N- and MP-/N+ Patterns of answers at the three levels.

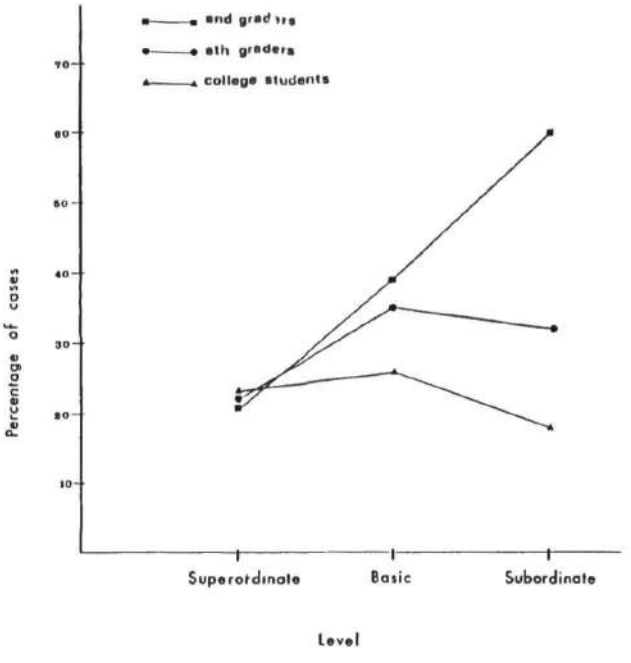

Figure 2: Incidence of the MP-/N- pattern of answers at the three category levels for each age group.

The results for the MP-/N- pattern of responses - the cases where subjects presumably relied on DSD - are presented in Figure 2. A two-way analysis of variance of the responses revealed significant effects of age ($F(2,410) = 3.79$, $p = .023$) and level ($F(2,410) = 3.19, p = .042$), as well as a significant interaction between the two variables ($F(4,410) = 2.42$, $p = .048$). The MP-/N- pattern was more prevalent the younger the subjects (its incidence was .39, .30, and .22 for the three age groups) and the lower the level (an incidence rate of .37, .33, and .22 for subordinate, basic-level and superordinate categories, respectively). Most interesting, the MP-/N- pattern was very common among 2nd graders teaching items from subordinate categories (60%), whereas for older subjects the mode was at the basic-level categories. This trend towards an increase in the size of DSD with age is also evidenced in the finding that for 8th graders the second most frequent MP-/N- cell was that of subordinate categories, while for the college students it was that of superordinate categories.

Our claim that people have a default similarity distance which they use and expect others to use when constructing a new category, provides a coherent and succinct explanation of the results. For each age group there was some category level for which a large proportion of the subjects found it unnecessary to indicate its variability. At the same time, subjects used multiple positive instances to indicate the greater than normal variability of more general categories and negative instances to indicate the relative narrowness of lower-level categories.

The results strongly imply that the size of DSD changes with age: For younger children the default variability was close to that of subordinate categories while for older children and adults the value was closer to the variability of basic-level categories. The evidence of the smaller size of children's default distance for categories is in line with findings of developmental studies of free categorization which indicate that younger children tend to create narrower categories than older ones (Nelson & Bonvillian, 1973, 1978; Saltz, Soller & Sigel, 1972).

The relationship between age and DSD has some implications for models of machine learning. Every model has to allow some similarily distance between items it classifies into the same category. Our findings imply that a model of human learning should initially use a steep generalization gradient and relax this requirement as more knowledge about the world accummulates. Thus, the present findings provide researchers in the field of machine learning with some idea of the size of the default similarity threshold to be installed for grouping stimuli into categories and the changes it should undergo with increased experience.

Bibliography

Fried, L. S. & Holyoak, K. J. (1984). Induction of category distributions: A framework for classification learning. Journal of Experimental Psychology: Learning, Memory and Cognition, 10, 234-257.

Hayes-Roth B. & Hayes-Roth, F. (1977). Concept learning and the recognition and classification of exemplars. Journal of Verbal Learning and Verbal Behavior, 16, 321-338.

Medin, D. L. & Barsalou, L. W. (1987). Categorization processes and categorical perception. In S. Harnad (Ed.) Categorical Perception. Cambridge, England: Cambridge University Press.

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

Nelson K. E. & Bonvillian J. D. (1973). Concepts and words in the 18-month-old: Acquiring concept names under controlled conditions. Cognition, 2, 435-450.

Nelson K. E. & Bonvillian J. D. (1978). Early language development: Conceptual growth and related processes between 2 and 4.5 years of age. In K. E. Nelson (Ed.) Children's language (Vol. 1). New York: Gardner.

Rosch, E. (1973). 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.

Rosch, E. & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. Cognitive Development, 7, 573-605.

Saltz E., Soller, E. & Sigel I. E. (1972). The development of natural language concepts. Child Development, 43, 1191-1202.