Evidence for Interrelated and Isolated Concepts from Prototype and Caricature Classifications

Robert L. Goldstone
Indiana University at Bloomington
rgoldsto@ucs.indiana.edu

Abstract

Previous research (Goldstone, 1991) has suggested that concepts differ in their degree of dependency on other concepts. While some concepts' characterizations depend on other simultaneously acquired concepts, other concepts are relatively isolated. The current experiments provide a new measure of a concept's interrelatedness/isolation. It is assumed that if the prototype of a concept is classified with greater accuracy than a caricature, then the concept is relatively independent of the influences of other concepts. If a caricature is more easily categorized than the prototype, then the concept is relatively dependent on other concepts. If these assumptions are made, then the current experiments provide converging support for a interrelated/isolated distinction. Instructing subjects to form images of the concepts to be acquired, or infrequently alternating categories during presentation, yields relatively isolated concepts. Instructing subjects to try to discriminate between concepts, or frequently alternating categories, yields relatively interrelated concepts.

The Interrelated/Isolated Distinction

In a recent paper, Goldstone (1991) argued for a distinction between two methods for representing concepts. One type of concept [originally called a "negatively-defined concept," a term borrowed from Saussure (1915/1959)] is characterized in terms of, or depends upon, other concepts at the same level of abstraction. The other type of concept (called a "positively-defined concept") does not depend on other such concepts for its representation. The distinction was partly motivated by the intuition that concepts seem to be directly accessed during object recognition, but also seem to be intricately connected to each other. Concepts are both isolated from, and interrelated to, other concepts.

Several frameworks for interrelated concepts exist. Semantic network models (Collin & Quillian, 1969) represent concepts by labelled links to other concepts. For example, the concept *canary* might have a *kind-of* link to *bird*, a *has-a* link to *beak*, and a *lives-in* link to *tree*. Barr and Caplan (1987) have discussed extinsically defined concepts that are defined by features that relate the concept to another concept. The feature "used to hit nails" is an extrinsic property of *hammer* because it relates *hammer* to the separate concept of *nail*. Winston, Chaffin, and Herrmann (1987) present a taxonomy of the various types of relations between concepts.

Perhaps the clearest examples of isolated concepts are feature detectors and templates. Idealized feature detectors are typically understood as responding to particular stimulus aspects regardless of the states of other feature detectors (Hubel & Wiesel, 1968). A feature detector that distinguishes red objects from non-red objects can still function properly even if other feature detectors are eliminated. Templates are photograph-like representations. A concept's template is typically based only on the visual characteristics of the concept's instances, and not on other concepts in the system.

Empirical Tools for Determining Concept Structure

The overall strategy for this research is to use converging operations (Garner, Hake, & Erikson, 1956) to identify relatively interrelated and isolated concepts. Experimental manipulations are developed that are expected to alter the representations of the concepts to be learned. At the same time, empirical indicators for isolated and interrelated concepts are developed. Ideally, the task manipulations that are intended to yield

isolated concepts will yield isolated concepts as measured by the indicators. The reason for testing multiple manipulations and indicators is that each manipulation/indicator, by itself, provides only an indirect and perhaps fallible lens onto category structure. Confidence that a particular manipulation or indicator is related to category structure increases if it provides results that are consistent with other methods.

Goldstone (1991) used the influence of nondiagnostic, relative to diagnostic, features as an indicator of conceptual representation. A nondiagnostic feature is a feature that does not, by itself, provide any basis for predicting one category rather than another. If two categories both have feature X present in 80% of their instances, then feature X does not provide information in favor of one of the categories over the other. However, in some circumstances, categorization accuracy is greater when feature X is present. If concepts are isolated, then any feature that tends to be present in the concept's instances will be part of the concept's representation, and may be used for categorization decisions. The distinction between a nondiagnostic and a diagnostic feature is only relevant when a concept is defined in relation to another concept. For interrelated concepts, diagnostic features will be more important than nondiagnostic features because they allow comparisons of the form "Concept Y has more of property P, relative to Concept X." As predicted by these assumptions, diagnostic features were particularly important relative to nondiagnostic features in tasks that biased subjects to create interrelated concepts.

The current experiments provide a converging measure of conceptual representation by assessing the relative ease of caricature and prototype categorizations. The prototype of a concept is its central tendency. It is the item that is closest on average to other category members. Caricatures of a concept assume dimension values that depart from the central tendency in the opposite direction from the central tendency of other concepts to be simultaneously acquired. Figure 1 shows members of two categories that are defined by their size and color. Five cm is the central tendency of concept B because it is the average size of B items. Six cm is a caricature of concept B because it is more than B's average of 5 cm and A's average of 2 cm is less than B's average. Just as a caricature of a politician in a cartoon exaggerates certain distinguishing features of the politician, so a caricature of a concept exaggerates a dimension value that distinguishes the concept from other concepts.

Is the prototype or caricature of category B more easily categorized as belonging to category B? If a concept is completely isolated, an advantage for the concept's prototype is expected. If a concept is defined without recourse to other concepts, then the representation that best exemplifies a concept will be its prototype. For completely isolated concepts, there is no difference between a caricature and other distortions of the concept that are equally distant from the central tendency. However, if a concept is defined relative to other concepts, then rules of the sort "Concept B items are larger than Concept A items" or "Concept B is large,

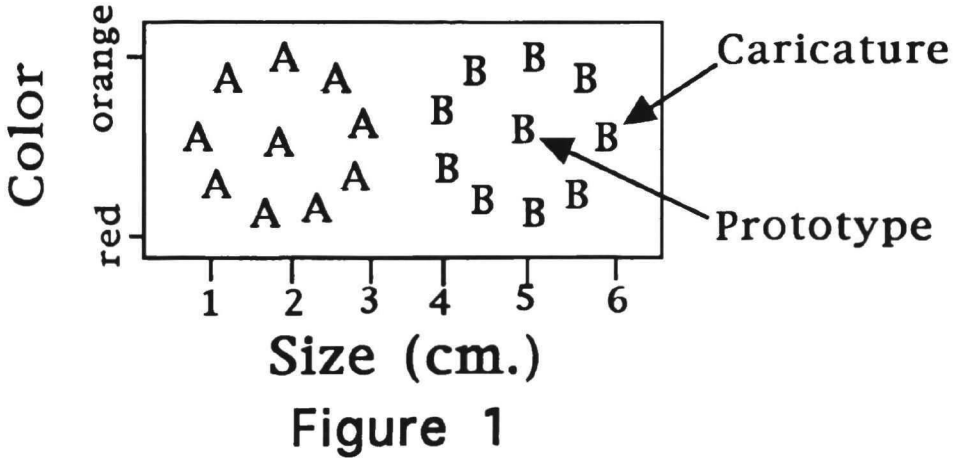

Figure 1

relative to Concept A" will be likely to develop. Caricatures better fit these relational rules than do prototypes. While items of 5 or 6 cm both satisfy the rule "larger than Concept A items," the item of 6 cm satisfies this rule more obviously.

Experiment 1

Goldstone (1991) found that instructions that were expected to bias subjects to create isolated concepts resulted in a relatively large influence of nondiagnostic features on categorization accuracy. Experiment 1 is an attempt to provide converging support for this influence of instructions using the categorization of caricatures and prototypes as indicators. One group of subjects is told to create images for the concepts that are learned. These Imagery instructions are predicted to promote isolated concepts. If concepts are represented as images, then they can be possessed without making reference to other concepts. A second group of subjects is instructed to look for stimulus aspects that serve to distinguish one category from another. A concept's distinguishing features are only diagnostic relative to another concept. While an image can be generated for a concept without any knowledge of the other concepts being acquired, the selection of distinguishing/diagnostic features for a concept requires knowledge of the other candidate concepts.

Method

Materials. Sample materials are shown in Figure 2. Stimuli consisted of seven vertical bars joined together to form a histogram-like shape. Each bar assumed one of 6 different values (1.0= shortest, 6.0=tallest, each unit corresponding to 1.5 cm). The histograms belonged to one of four categories. Each category was associated with one long bar, and three bars with heights with values of 1.0. These four bars are diagnostic because, if they have a height value above 1 then they provide information in favor of one and only one of the categories. The lengthened diagnostic bar for a category has a value of 2.0 on one quarter of the trials, 3.0 on one half of the trials, and 5.0 on one quarter of the trials. Thus, a value of 3.0 for a lengthened diagnostic bar is considered the prototypical value, because it occurs most frequently, and because its value is intermediate to the other possible values. A value of 5.0 for a diagnostic bar is considered a caricature because it exaggerates the length of the bar that is particularly long for the category, and it exaggerates the length in the direction opposite to the other category's bars.

The other three bars are nondiagnostic. At the beginning of an experiment, three randomly determined bar heights are generated. These are the standard values for the nondiagnostic bars. The bars are nondiagnostic because all four category items share the same standard set of values for these bars. On one half of the trials, a histogram is presented with the standard values for the nondiagnostic bars. On the other half of the trials, new random heights are generated. Thus, subjects get a great deal of experience with the set of standard nondiagnostic bar heights, and get less experience with all other nondiagnostic bar heights.

Procedure. Twenty-six undergraduates from Indiana University were divided into two instruction groups. One group of subjects (the Image group) was given the instructions, "While you are learning the two categories, you should try to form an image of what each category looks like. Use these images to help you categorize the pictures that you see." The other group of subjects (the Discriminate group) was given the instructions, "While you are learning the two categories, you should try to find particular features in the pictures that help you distinguish between the two categories."

Subjects were presented with 384 trials in all. On each trial, a histogram appeared and subjects pressed one of four keys to indicate their proposed categorization for the object. Subjects then received feedback indicating the correct response.

Results and Discussion

As Figure 3A shows, there is a general speed advantage for categorizing caricatures (height = 5) relative to prototypes (height = 3). In addition, subjects are faster to categorize instances when the nondiagnostic bars are in their standard configuration. Finally, there is a main effect of instructions, with Discriminate subjects categorizing more quickly than Image subjects.

Figure 3A also shows two interactions of interest - between height of lengthened diagnostic bar and instructions, $F(2, 12) = 6.5$, $p < .05$, and between height and configuration of nondiagnostic bars, $F(2, 12) = 6.2$, $p < .05$. Both of these interactions are predicted by the isolated/interrelated framework. According to the first interaction, the speed advantage of caricatures over prototypes is particularly pronounced when subjects are given Discriminate rather than Image instructions. Given that the magnitude of the caricature advantage is assumed to be directly related to the degree of dependency between concepts, it is appropriate that the caricature advantage is greatest when subjects are told to search for features that serve to distinguish between the categories. According to the second interaction, the caricature advantage is greatest when items are presented with a random configuration of nondiagnostic features. One interpretation of this result is that when the overall stimulus configuration is similar to a familiar item, then a process similar to template-matching transpires. Conversely, when the item, because of unfamiliar values for nondiagnostic values, is not similar to many previous items, then rules of the form "third bar is relatively long" may be used.

Experiment 2

Experiment 1 showed that the categorization advantage of caricatures over prototypes is amplified in the Discriminate condition, a condition that is hypothesized to produce relatively interrelated concepts. In Experiment 2, the interrelatedness of concepts is manipulated by alternating the presented categories either frequently or infrequently. If categories are alternated rarely, then subjects will see long series of items that belong to the same category. If categories are alternated frequently, then a picture from one category will most often be followed by a picture from a different category.

The *apriori* assumption for this experiment is that frequent alternation of categories will yield concepts that are relatively interrelated, and infrequent alternation will yield relatively isolated concepts. If several instances of

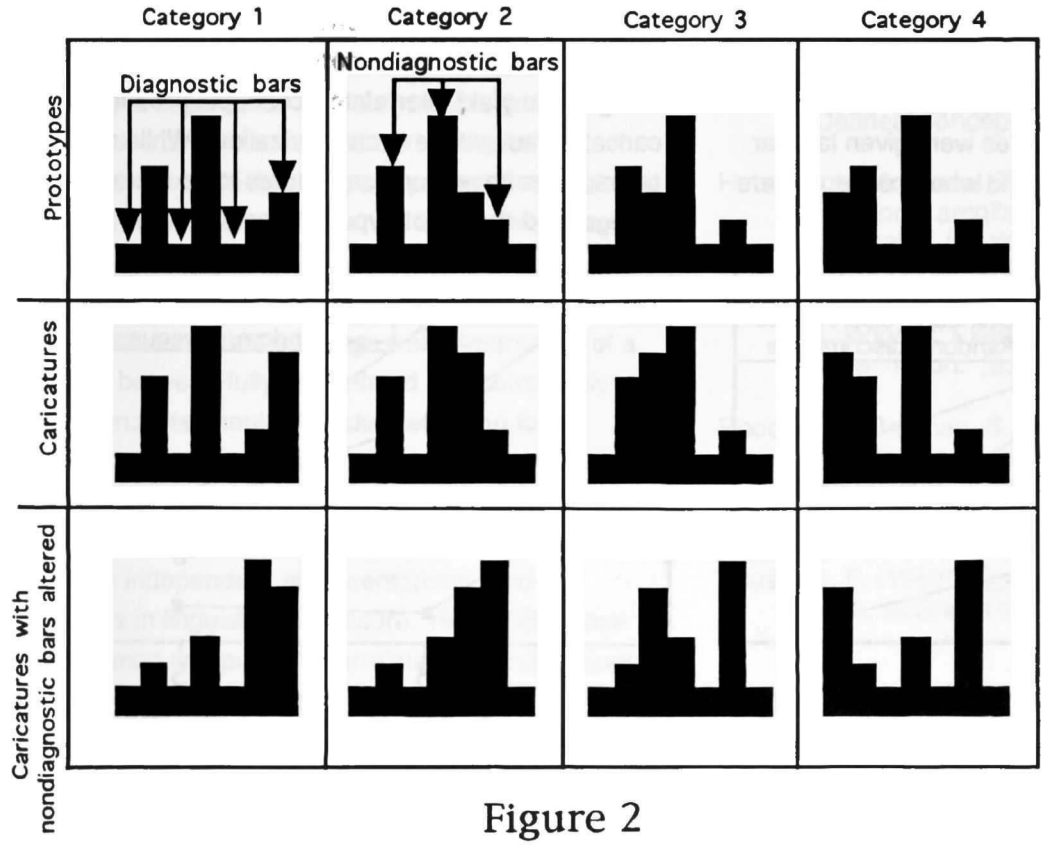

Figure 2

the same concept are presented in a series, then the concept's representation will most likely be based on the common characteristics of the instances. Influences from the other concept will be relatively modest. If categories are alternated frequently, then there will be more opportunity for interplay between the developing concepts.

Method

The materials and general procedure from Experiment 1 were used. Twenty-four Indiana University undergraduates were split into two groups. For one group of subjects, the frequent alternation group, the probability of presenting an item from the same category as the previous item was 7%, and the probability of presenting an item from one of the three other categories was 31%. For the infrequent alternation group, the probability of presenting an item from the same category as the previous item was 31% and the probability of presenting an item from one of the three other categories was 23%.

Results and Discussion

The results for Experiment 2 are shown in Figure 3B. As with Experiment 4, caricatures were generally categorized more quickly than prototypes. Subjects were faster when nondiagnostic features were given familiar values (as with Experiment 1), and when concepts were alternated infrequently.

There are also two interactions involving the height of the lengthened diagnostic bar. First, the speed advantage of caricature over prototype categorization is greater when concepts were alternated frequently rather than infrequently, $F(2, 12) = 5.6$, $p < .05$. Second, the speed advantage of caricature over prototype categorization is greater when nondiagnostic features are given familiar rather than unfamiliar values, $F(2, 12) = 4.04$, $p < .05$. These results are consistent with the predictions of the interrelated/isolated framework. Frequent alternation of categories is expected to yield interrelated concepts, because instances from different concepts will more likely be compared to each other. The presence of familiar nondiagnostic bar values may yield relatively isolated concepts because it facilitates an overall template match rather than a selective use of discriminating features.

Conclusions

Both of the experiments provide converging support to previous results (Goldstone, 1991) that have suggested an isolated/interrelated distinction among concepts. Experimental manipulations that were designed to yield interrelated concepts produced a strong caricature advantage in categorization. While other researchers have found caricatures to be more easily categorized than prototypes (Nosofsky, 1991; Rhodes,

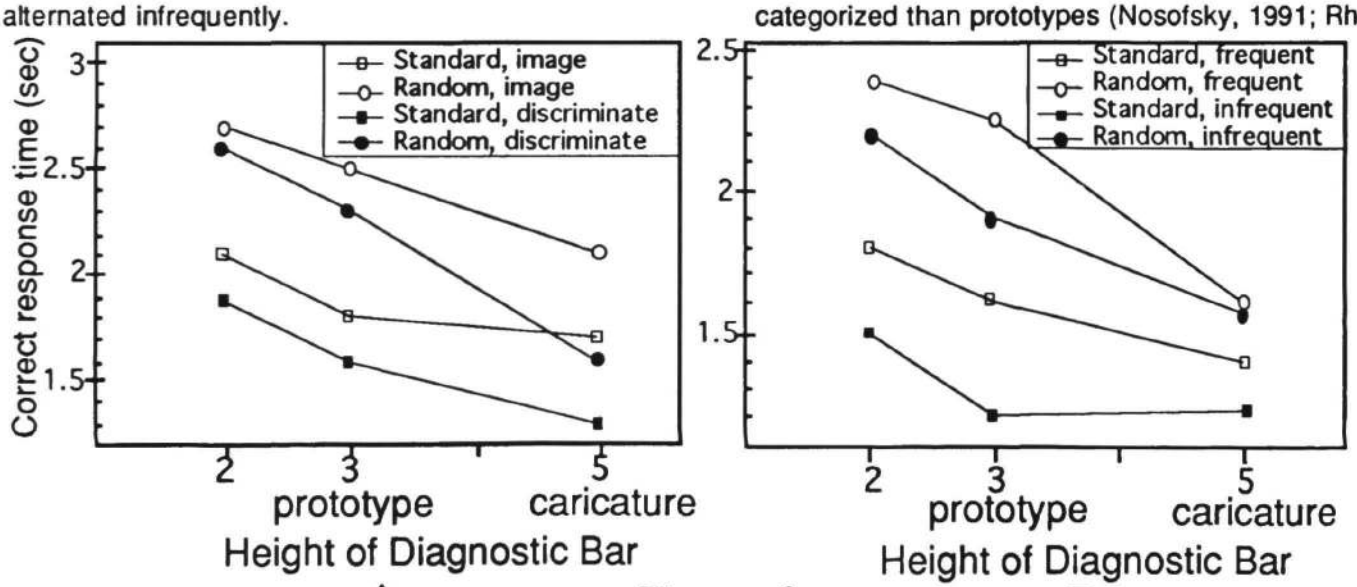

Figure 3

Brennan, & Carey, 1987), the current experiments augment this work by showing that the relative advantage of caricature categorization depends on task factors. Factors that promote interrelated concepts (Discriminate instructions and frequent alternation between categories) increase the caricature advantage; factors that promote isolated concepts (Image instructions and infrequent category alternation) decrease the caricature advantage.

Recently, a connectionist simulation, RECON, has been developed (Goldstone, 1993) that models varying degrees of concept isolation/interrelation by varying the influence of concept nodes on each other. Activation recurrently flows not only from input features to concept nodes, but also from concept nodes to input features, and between concept nodes. The influence of concepts upon each other is increased by increasing the flow of activation between concept nodes or by increasing the learning rate of connections between concepts. When interconceptual influences increase, caricatures become increasingly well categorized relative to prototypes, and nondiagnostic features become less influential.

A full account of concept learning should explain how concepts seem to be intricately related to each other, yet also directly accessed during object recognition. Towards this end, this paper has developed a distinction between isolated and interrelated concepts. Virtually all concepts fall somewhere between these two extremes. Concepts can be located on the continuum by empirical measures that converge well. Awareness of a continuum between fully interrelated and completely isolated concepts should provide motivation for researchers to expand their models to incorporate both varieties of concepts. While researchers in pattern recognition have usually hypothesized concepts with completely independent representations, and researchers in linguistics (Saussure, 1915/1959) have often assumed completely interrelated representations, the majority of concepts may require both representations. Detailing the interactions between independent perceptual representations and interconceptual relations will very likely prove to be an important part of the agenda of cognitive science.

Acknowledgments

This paper has greatly benefitted from helpful suggestions from John Kruschke, Douglas Medin, Robert Nosofksy, and Edward Smith.

Bibliography

Barr, R. A., & Caplan, L. J. (1987). *Category representations and their implications for category structure.* Memory & Cognition, 15, 397-418.

Collins, A. M., & Quillian, M. R. (1969). *Retrieval time from semantic memory.* Journal of Verbal Learning and Verbal Behavior, 8, 240-247.

Garner, W. R., Hake, H. W., & Eriksen, C. W. (1956). Operationism and the concept of perception. Psychological Review, 63, 149-159.

Goldstone, R. L. (1991). Feature diagnosticity as a tool for investigating positively and negatively defined concepts. Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society. (pp. 263-268). Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Goldstone, R. L. (1993). Positively and Negatively Defined Concepts. Unpublished manuscript.

Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. Journal of Neurophysiology, 195, 215-243.

Nosofsky, R. M. (1991). Typicality in logicall defined categories: Exemplar-similarity versus rule instantiation. Memory & Cognition, 19, 131-150.

Rhodes, G., Brennan, S., & Carey, S. (1987). Identification and ratings of caricatures: Implications for mental representations of faces. Cognitive Psychology, 19, 473-497.

Saussure, F. (1959). Course in general linguistics. New York: McGraw-Hill.