Distinguishing Name Centrality From Conceptual Centrality

Woo-kyoung Ahn
Department of Psychology
Yale University, P. O. Box 208205
New Haven, CT 06520-8205
ahn@pantheon.yale.edu

Steven Sloman
Department of Cognitive & Linguistic Sciences
Brown University, Box 1978
Providence, RI 02912
Steven_Sloman@brown.edu

Abstract

The features of a concept differ in their centrality. Having a seat is more central to the concept of chair than is having arms. This paper claims that centrality is not a homogeneous phenomenon in that it has at least two aspects, conceptual and naming. We propose that a feature is central to naming in proportion to the feature's category validity, the probability of the feature given the category. In contrast, a feature is conceptually central (immutable) to the extent the feature is depended on by other features. We predict that conceptual and naming centrality diverge as categories become more specific. An experiment is reported that provides corroborating evidence. Increasing the specificity of object categories increased the judged mutability of representative features without affecting their judged appropriateness for determining names.

Introduction

Nobody can say clearly and definitely what a game is, or what a chair is, or even what a bachelor is. Wittgenstein (1953) was surely right that few, if any, categories can be identified with a set of defining features. Of course, most of us know a game or a chair when we see one, so our concepts must cohere in some sense, even if they do not consist of a set of necessary and sufficient conditions.

Even if features are rarely defining, the coherence of a concept does seem to rely more on some features than others. Not all games are competitive (consider ring around-the-rosy), but being competitive is clearly more important to being a game than is (say) being outdoors. "Being competitive" is more important to the concept of *game* in many senses: Something is more likely to be a game if it is competitive than if it is outdoors (many events that are not games are outdoors); imagining a game that is not competitive is harder than imagining one that is not outdoors; competition comes to mind more rapidly when one thinks about games in general than does the outdoors; a game that is not competitive is a more atypical game than one that is not outdoors. In all these senses, and others, competition is more central to "gamehood" than is being outdoors. In general, concepts are comprised of aspects or features that differ in their centrality. Smith, Shoben, and Rips (1973) argued that some features are "core" and others merely "characteristic." But they themselves appreciated the implausibility of this position. Core is merely another term for defining, so if concepts don't have defining features, they don't have core features either. Features differ in their centrality, but their differences are graded; they are neither dichotomous nor trichotomous.

How many kinds of centrality are there? The notion of centrality would be vacuous if there were too many. If every question about feature differences could be answered only by proffering a new, unique kind of centrality, then the notion of centrality would offer nothing to the answer. Sloman, Love, and Ahn (submitted) provide evidence that this is not the case; notions of centrality converge. They found that several measures of centrality were highly correlated. Nevertheless, in this paper we discuss the need for a distinction between two notions of centrality, one pertaining to conceptual structure, the other to naming. We show that the distinction is more apparent at more specific levels of categorization (e.g., green seedless grapes) and less apparent at more abstract levels (e.g., grapes). We first argue for the distinction and its meaningfulness and then report an experiment that attempts to dissociate measures of the different kinds of centrality.

Conceptual centrality: mutability

One sense of centrality is conceptual. A feature is central to an object because the feature binds the mental representation of the object. Having internal organs is central to our concept of "mammal" because we use internal organs and their relations to other mammalian features to make sense of the life-cycle, appearance, and activity of mammals (Keil, 1989). We use the term *mutability* to refer to conceptual centrality. Features differ by degrees of mutability, the most central being the most immutable. A feature is mutable of an object to the extent that the feature can be mentally transformed without altering the object's representation. Tails are mutable of cats because we can easily mentally transform our notion of "cat" to something that is the same in all respects to "cat" except that it does not have a tail while still thinking about cats.

We define mutability operationally, using a variety of convergent measures. A feature is immutable with respect to an object to the extent that people i. claim that they would be surprised to encounter the object without the feature; ii. claim that imagining the object without the feature is difficult; iii. rate the object without the feature as atypical of the object category iv. require time to mentally transform the object to one that does not have the feature; and v. rate the object without the feature as dissimilar to an ideal instance of the object category. Sloman et al. (submitted) and Love (1996) demonstrated convergence amongst these measures by showing that they were all highly correlated. For instance, people's ratings on the following questions about an apple tended to converge; i. How surprised would you be to encounter an apple that did not grow on trees? ii. How easily can you imagine a real apple that does not grow on trees? iii. How good an example of an apple would you consider an apple that does not ever grow on trees?

Moreover, the measures did not correlate with measures of other dimensions on which features differ. Ratings of centrality did not correlate with ratings of feature salience, such as a feature's degree of prominence (How prominent in your conception of an apple is that it grows on trees?), nor with ratings of feature diagnosticity, like cue validity -- the probability of the category given the feature (Of all things that grow on trees, what percentage are apples?).

Our earlier work (Sloman & Love, 1995; Sloman et al., submitted) has explored the hypothesis that mutability is a function of a concept's dependency structure. Specifically, we have collected evidence to support the hypothesis that a feature is immutable to the extent that other immutable features depend on it. For instance, for the concept "chair," the feature "has a seat" was judged to have many features depending on it, including "is comfortable," "you sit on it," and "holds people." However, people judged that few features depended on "is made of wood." Measures of this degree of dependency correlated with immutability judgments in that people judged that imagining a chair without a seat is more difficult than imagining one that is not made of wood.

Some of the evidence is in the form of close fits between a simple model of concepts' dependency structure and our mutability judgment data. The model is a constraint satisfaction network that uses asymmetric connections between features to represent empirically-derived dependency relations. A feature is conceptually central, according to our hypothesis, to the extent that transforming the feature would necessitate the transformation of other features; i.e., to the extent that the feature helps to maintain the current state of the network. People's judgments on various measures of mutability were highly correlated with the predictions made by this simple network.

Name centrality

Features can be central in a different sense as well, a sense that may be more common. A feature can be central to an object by virtue of the feature's power to determine the appropriateness of a label for the object. Having light-colored hair determines whether someone can appropriately be called "blond," and having four doors is central to whether a car can appropriately be called a "four-door sedan." Exceptions can be imagined, a natural blond might dye their hair, but the general tendency stands: Cooperative conversation and reference demands that an object referred to by the term "blond" have, or appear to have, or once had, or have under some special conditions, light-colored hair.

This notion of centrality is not the same as the conceptual one. To see this, note that having four doors is not immutable; it is not central to the *concept* of four-door sedan. A sedan without four doors is assuredly not, strictly speaking, a four-door sedan. But one's notion of four-door sedan in the vast majority of contexts would be much more concerned with attributes like having a motor, having seats, safety, comfort, and reliability than it is with four doors. So having four doors is central to naming without being central to conceiving. Consider another example. A person is appropriately labeled an American if they hold United States citizenship. But holding U. S. citizenship hardly describes what it means to be an American. Much more central to the concept "American" is language, political principles, and cultural icons. Indeed, many Canadians who grew up with stories about George Washington and prime-time TV might constructively be *thought* of as "American."

The naming centrality of a feature is closer to the notion of definingness than conceptual centrality is. A feature is central in naming to the extent that the feature obtains when the relevant object does. "Four door" is central for an object to be a four-door sedan because it is necessary. If **x** is a four-door sedan, then **x** has four doors (strictly speaking). If **x** is an American, then **x** has U. S. citizenship. However, name centrality is not dichotomous; it comes in degrees. Some features are *fairly* central because they *usually* obtain when the object does (e.g., tents are usually, but not always, enclosed on all sides). Hence, we take it that a feature is central in category naming if the probability of the feature given the category is high. This probability measure is often referred to as category validity.

Mutability versus name centrality

In most cases, the two types of centrality are highly correlated. After all, a category label should, to some extent, relate to the category's most stable, conceptually central elements. The label "bird" suggests something that has wings and can fly and, indeed, having wings and flying are relatively immutable properties of birds. If the label did not help to identify such properties, then the label would not effectively distinguish birds from categories with different immutable features. Furthermore, immutable properties are likely to be attributes that we want to refer to in everyday discourse. In discussing the symptoms, etiology, or care of a disease, for example, successful communication requires that we make reference to whatever we know about the conceptually central aspects of the disease, like the virus that

causes it. So immutable features can bind the category that we need a label for and thus be necessary. Whenever conceptual structure orders features in line with category validity, mutability and name centrality will be correlated.

Nevertheless, mutability should not be confused with name centrality. The two types of centrality are dissociable because when naming, sometimes we wish to discriminate an object from others that share immutable properties with it. In particular, features are more likely to be necessary but nevertheless mutable when discourse concerns specific -- as opposed to abstract -- categories. At lower levels of abstraction, categories tend to share their dependency structure with other categories. One kind of sedan has much the same dependency structure as another kind of sedan. Ford sedans differ in important ways from Nissan sedans, but the differences are not in their most immutable properties. Both have engines, seats, steering wheels, etc. The features that are critical for determining whether an object is a Ford sedan or a Nissan sedan include some that are necessary without being conceptually central. These include the label that appear on the object's name plate, the details of the object's shape, its gas mileage, color, etc. Specific categories are likely to include features that dissociate name centrality from mutability because, by virtue of sharing their dependency structure with other categories, they will contain features that are distinct from those of other categories and yet are not the most conceptually central.

This leads directly to the hypothesis of this study. We propose that the name centrality and mutability of a feature are more likely to dissociate at lower levels of category abstraction. Therefore, we varied the level of abstraction of a variety of categories and took measurements of name centrality and of mutability of selected features at each level. We predict that necessary features will be more mutable at the specific level than at the abstract level. To measure name centrality, we asked subjects to judge the appropriateness of a label for an object missing the relevant feature. To measure mutability, we asked subjects to rate the ease of a mental transformation involving the feature.

Experiment

Method

Participants

Seventeen undergraduate students at Yale University participated in this experiment. Five participated in the experiment as partial fulfillment of requirement for an introductory psychology course. Twelve received $7.00 for their participation in this and other short experiments following this one.

Materials

Nine taxonomies with two levels of categories each were selected from Rosch, Mervis, Gray, Johnson, & Boyes-Braem (1976a). The full list of the 18 categories is provided in Table 1. Among these categories, 9 were specific or what Rosch et al. (1976a) called "subordinate-level" categories. Categories one level higher than each of these 9 specific categories constituted the abstract categories in the experiment. These categories are all at what Rosch et al. (1976a) called the "basic level". Five of these taxonomies involved artifacts and four involved biological kinds.

For each category at each level, one feature was selected from Rosch, Mervis, Gray, Johnson, & Boyes-Braem's (1976b) full set of attributes. This set was compiled by asking undergraduates to write down attributes that they could think of for each category and by eliminating features that were not commonly listed. The following criteria were used to select features from this full set. The feature associated with a specific category should be one of the features that distinguishes it from other specific categories sharing a superordinate. For instance, "being large" distinguishes a grand piano from an upright piano, and therefore was selected as a feature for *grand piano*. The feature associated with an abstract category was one of the features that generally hold in all of its subordinates. For instance, "makes music" holds for both upright and grand pianos, and therefore was selected as a feature for *piano*. The features used in the experiment appear in Table 1 in parentheses next to their corresponding category.

Table 1. Categories and features used in the experiment.

Specific category and its feature	Abstract category and its feature
grand piano (is large)	piano (makes music)
green seedless grapes (is green)	grapes (is juicy)
Phillips head screwdriver (has a cross-shaped tip)	screwdriver (has a handle)
Levis (is blue)	pants (have legs)
living room chair (has a cushion)	chair (has a seat)
sports car (is small)	car (has an engine)
birch tree (has white bark)	tree (has a trunk)
seabass (is large)	fish (swims)
cardinal (is red)	bird (has wings)

Based on the category-feature list in Table 1, two sets of questions were developed, one for name centrality judgments and the other for mutability judgments. Questions involving the name centrality judgment were in the form, "Suppose an object is in all ways like X except it does not have feature Y. How appropriate would it be to call this

object X?" where X was one of the categories in Table 1 and Y was its corresponding feature. Following each question, a 9-point scale was presented where "9" indicated "very appropriate" and "1" indicated "very inappropriate." Questions involving the mutability judgment were in the form, "Imagine an X that has all the usual characteristics and properties of X. Now, change this image of X so that it is in all ways like X except it does not have feature Y. Rate the ease of this transformation." Following each question, a 9-point scale was presented where "9" indicated "very difficult" and "1" indicated "very easy."

Procedure

For the name centrality judgment questions, participants were instructed "to evaluate the appropriateness of a label." They were further told, "Sometimes, objects seem to require certain attributes to warrant a specific label. It would be inappropriate, for example, to call a man a 'bachelor' if he was known to have a spouse. We'd like to know what properties you think are necessary to apply a label to an object. Your task is to rate how appropriate a label is for an object that is missing a specified property." The 9-point scale was explained by showing participants example questions; a cloud not floating in the sky would receive "2" if they thought that something would not normally be called a cloud if it did not float in the sky. They were also told that a car which did not have a horn might receive "8" if they thought that something would probably still be called a car even if it did not have a horn.

For the mutability judgment questions, participants were instructed "to evaluate the ease with which you can transform an image of an object." They were further instructed, "We'll ask you to imagine an ideal object, and then change some specified part or aspect it. For example, imagine a door, then transform it in your mind into a door without a doorknob. We would like to know how easily you can complete this transformation. Your task, then, is to rate the ease of the transformation required to get from the original to the mutated form." After being instructed on the 9-point scale, they received examples such as "Imagine a textbook that has all the normal characteristics and properties of a textbook. Now, change this image of the textbook so that it is in all ways like a textbook except that it does not have pages. Rate the ease of this transformation." Then for this example, they were told that "'8' might be your response if you think that it would be difficult to transform your image of a textbook to one that lacks pages." They also received an example which might receive a lower rating: "Imagine a coat that has all the normal characteristics and properties of a coat. Now, change this image of the coat so that it is in all ways like a coat except that it does not have a hood."

Each question was displayed on a computer screen one at a time. Participants entered their responses on the keyboard. The responses were displayed on the computer screen and the participants could correct any mistake before they proceeded to the next question. The entire experiment was self-paced. Eight participants received the name centrality judgment questions first, followed by the mutability judgment questions. Nine received the mutability judgment questions first, followed by the name centrality judgments. For each type of judgment, the order of the 18 questions was randomized for each participant. The experiment was programmed using Psyscope 1.1 (Cohen, MacWhinney, Flatt, & Provost, 1993).

Results

The order in which the participants received the two types of questions did not affect the responses. Therefore, all data will be collapsed across the two orders in the following analyses.

The name centrality and the mutability scales used in the experiments are in the opposite direction in that 1 in the name centrality scale (i.e., very inappropriate to call it X without feature Y) means the feature is very central whereas 1 in the mutability scale (i.e., very easy to transform feature Y in X) means the feature is peripheral. To facilitate presentation of the results, the name centrality scale was reversed so that the two scales had the same directionality.

As shown in Table 2, the name centrality ratings were essentially the same for the abstract and the specific level categories, indicating that the features selected in this experiment did not vary in centrality substantially with respect to naming. However, the mutability ratings show that the features unique to the specific level were easier to mentally transform than those associated with the abstract level. A two-way analysis of variance was conducted with the type of task and the level of categories as two within-subject variables. As predicted, there was a reliable interaction between the two factors, $F(1,15) = 4.68$, $MSe = 0.83$, $p < .05$. In addition, regardless of the type of task, the features selected for the abstract categories were judged to be reliably more central (mean of 5.10) than those selected for the specific categories (mean of 4.11), $F(1,15) = 35.96$, $MSe = .42$, $p < .01$. This main effect is mainly due to the low mutability ratings at the specific level. Finally, regardless of the level of categories, the ratings on the name centrality judgment task were reliably higher (mean of 5.66) than those on the mutability judgment task (mean of 3.54), $F(1,15) = 42.86$, $MSe = 1.68$, $p < .01$.

Table 2. Mean mutability and name centrality ratings. On both scales, 1 means less central and 9 means more central.[SS1]

	Mutability Judgment	Name centrality Judgment
Abstract	4.28	5.90
Specific	2.81	5.42

The pattern of results was generally consistent across the nine items. For 8 of 9 categories, the mutability judgment for the feature at the specific level was lower than for the feature at the abstract level (the exception was "car"). In

contrast, the name centrality judgment was lower at the specific than at the abstract level for only 4 categories.

Discussion

In summary, the experiment shows a dissociation between mutability and name centrality. Changing the specificity of the object category increased the average mutability of representative features without affecting their average name centrality. Hence, centrality is not a homogeneous phenomenon; it has at least two aspects, conceptual and naming. We propose that a feature is central for naming in proportion to the probability of the feature given the category. In contrast, we propose that a feature is conceptually central to the extent that the feature is depended on by others.

Implications

A debate has raged in philosophy (e.g., Austin, 1961; Putnam, 1975; Wittgenstein, 1953) and in psychology (e.g., Armstrong, Gleitman, & Gleitman, 1983; Malt, 1994; Smith, Shoben, & Rips, 1973) about whether concepts include defining features. Some have argued that concepts have "theoretical" cores (Keil, 1989; Murphy & Medin, 1985), and others that concepts have unresolved "essences" (Medin & Ortony, 1989). Our work has two implications relevant to this debate. First, features can vary in their centrality without being dichotomous -- i.e., defining versus non-defining. Indeed, in earlier work we have established that people's ratings of centrality are graded, not dichotomous (Sloman et al., submitted). Second, feature centrality is relative to the function being served by the feature. Concepts have multiple facets. The importance of a feature depends not only on the identity of the feature and its relation to other conceptual features, but also on the goal of the agent using the concept. In particular, using a concept to name an object requires different information about the internal structure of the concept than does transforming the concept in the service of conceptual tasks like imagination, typicality judgment, and problem-solving.

The necessity of distinguishing conceptual and name centrality points to a gap between the concepts employed by thought and those employed by language (Brown, 1958). Language makes its own particular demands on conceptual structure in the service of communication. Specifically, labeling depends more on knowledge about the relative frequency of features in a category, and less on knowledge about the relations between features. This may be so because relative frequency information generally can be obtained with more confidence and more reliability. This is not to say that language and thought make use of mutually exclusive sets of concepts, only that conceptual structure must be viewed from one perspective to account for linguistic usage and a different one to account for thought.

In addition, the fact that the dissociation between conceptual and naming centrality occurred at the specific level of categories suggests one possible reason why "basic level" categories are psychologically more basic (Rosch et al., 1976a). This level might be the one at which instances share the most dependency structure. Because conceptual centrality is determined by dependency structure, conceptual centrality is identical to naming centrality at the basic level. We speculate that the relative ease of processing at the basic level is due to this convergence. When they diverge, as they do at the subordinate level, language and conception are no longer entirely mutually supportive. This might make processing difficult. Referring to a "1956 dime" rather than just a "dime" is confusing if I am discussing a financial transaction involving 10 cents.

Finally, the need to distinguish these two forms of categorization suggests that other forms may require their own special analyses. For example, object recognition almost surely comprises a different set of processes than does classification in abstract domains like law and football strategy.

Acknowledgments

We would like to thank two anonymous reviewers for their helpful suggestions on the earlier draft and Peter Jaros and Helen Sullivan for running the experiment reported in this paper. This project was supported by a National Science Foundation Grant (NSF-SBR 9515085) to Woo-kyoung Ahn and a Salomon Faculty Research Award to Steven Sloman.

References

Armstrong, S. L., Gleitman, L., & Gleitman, H. (1983). What some concepts might not be. *Cognition, 13*, 263-308.

Austin, J. L. (1961). *Philosophical Papers*. Oxford: Oxford University Press.

Brown, R. (1958). How shall a thing be called? *Psychological Review, 65*, 14-21.

Cohen J. D., MacWhinney B., Flatt M. & Provost J. (1993). Psyscope: A new graphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments & Computers, 25*, 257-271.

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

Love, B. C. (1996). Mutability, conceptual transformation, and context. *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society*, San Diego, CA. Hillsdale: Erlbaum.

Love, B. C. & Sloman, S. A. (1995). Mutability and the determinants of conceptual transformability. *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*, Pittsburgh, PA. Hillsdale: Erlbaum.

Malt, B. C. (1994). Water is not H_2O. *Cognitive Psychology, 27*, 41-70.

Medin, D. L., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.),

Similarity and analogical reasoning . New York, NY: Cambridge University Press.

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

Putnam, H. (1975). The meaning of 'meaning.' In H. Putnam, *Mind, language, and reality: Philosophical papers, vol. 2*. Cambridge: Cambridge University Press.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976a). Basic objects in natural categories. *Cognitive Psychology, 8*, 382-439.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976b). Basic objects in natural categories. *Working Paper #43.* The Language Behavior Research Laboratory, UC Berkeley.

Sloman, S. A., Love, B. C., & Ahn, W (submitted). Feature centrality and conceptual coherence.

Wittgenstein, L. (1953). *Philosophical Investigations.* New York: MacMillan.