The role of category coherence in experience-based prediction

ANDREA L. PATALANO
Wesleyan University, Middletown, Connecticut

AND

BRIAN H. ROSS
University of Illinois, Urbana-Champaign, Illinois

Both real-world category knowledge and instance-based sample data are often available as sources of inductive inference. In three experiments using natural social categories, we test the influence of general category knowledge on the use of category instances to make property inductions both to other category members and to others in the population. We find that a category's coherence—the extent to which its features are interrelated through prior knowledge (Murphy & Medin, 1985)—influences inductions positively to new category members and negatively to the population. This effect of coherence is strongest with small as compared with large samples of instances. The results are interpreted from both similarity and explanation-based perspectives.

A central way in which categories differ from arbitrary collections of entities is that categories support broad inductive reasoning. For example, that an entity is a dog is more informative about other properties of the entity, such as what it eats or where it lives, than that it is a brown thing. Category-based induction has been a major focus of categorization research (Nisbett, Krantz, Jepson, & Kunda, 1983; Wisniewski, 2002) both for the practical implications, as well as to help understand what makes a collection of entities exist as a psychologically meaningful whole. The strength of a category-based induction is influenced by many factors including prior knowledge about a category, recent experiences with category instances, and the nature of the property in question. In this article, we examine how a particular type of prior knowledge, *category coherence*, influences category-based induction from experience with category instances, using natural social categories.

Category coherence refers to the extent to which category features make sense or go together in light of prior knowledge (Murphy & Medin, 1985). A category of birds with wings that can fly, for instance, is more coherent than one with wings that cannot fly, assuming knowledge that wings support flying. Past research has shown that category coherence facilitates learning of categories (see Murphy, 2002, for review) and has begun to address the role of coherence in category-based induction (Patalano, Chin-Parker, & Ross, 2006; Rehder & Hastie, 2001, 2004). Most relevant to the present work, Patalano et al. found that given properties of natural job and hobby categories (e.g., "80% of yacht club members [or county

clerks] prefer Chinese to Mexican food") are more often extended to a new category member of a high coherence category (e.g., yacht club member) than to a member of a low coherence category (e.g., county clerk). Coherence was operationalized as the rated uniformity of category members, which was highly correlated with other coherence measures including category informativeness for inference and the presence of underlying features giving rise to surface ones (Haslam, Rothschild, & Ernst, 2000), as well as the number of underlying featured generated in a feature listing task. For example, for yacht club member, underlying features included: is educated, possesses great family wealth, lives an indulgent lifestyle, and is very sociable. In contrast, more superficial features were generated for county clerk including: has a secure government job with benefits, is skilled in sending bills, and is polite at the office. Although much past work had used novel experimental materials, Patalano et al. showed that real-world social categories, a rich set of categories about which inductions are often made, also vary in coherence.

The described research with natural social categories focused on the extension of a category property to an individual category member. However, many everyday inductions are at least initially motivated by an observation of a pattern across experienced instances (see Heit, 2000, for review). For example, it may be upon experiencing multiple county clerks who prefer Chinese food to Mexican food that one begins to extend this property to other county clerks. Although much research has examined people's willingness to extend properties from sample data to other category members, the role of category co-

A. L. Patalano, apatalano@wesleyan.edu

herence in such inductions has not been investigated. The most closely related work is by Nisbett et al. (1983) who examined perceived property variability and sample size. They found that properties with low perceived variability (e.g., skin color for an island tribe) are more likely to be extended than more variable ones (e.g., body weight), with the difference decreasing as experienced sample size increases. Like property variability, category coherence might be especially important in making inductions from instances. However, given the classic finding that sample information sometimes overwhelms prior knowledge (base rates; Kahneman & Tversky, 1973), it is important to investigate the role of coherence in this context.

OVERVIEW OF EXPERIMENTS 1A AND 1B

The present research examines the influence of category coherence on induction from sample data to cover two issues. First, we ask whether coherence influences the strength of the induction—do higher coherence categories lead to stronger inductions even when information about the property is learned from samples? Second, how does this influence change with sample size? Experiments 1A and B test the extent to which category coherence is related to the perceived strength of inductions from sample data to other category members versus to the population at large, while Experiment 2 extends the finding by introducing variations in sample size to assess the extent to which the coherence effect increases as sample size decreases. Through this work, we hope to develop a richer understanding of how coherence influences category-based induction, and to situate the influences of coherence in a broader system of sample-based probabilistic reasoning.

Experiment 1 contrasts two different ways in which category coherence may influence the strength of induction from category instances. We use natural social categories that have been previously rated on coherence. In Experiment 1A, we consider the basic premise that there is a greater willingness to transfer a novel property from some experienced category instances to another category member when the category is high rather than low in coherence. For example, suppose one observed that a number of category members prefer Chinese food to Mexican food despite an equal preference for these two cuisines in the population, and one then judges the probability that a new member of the category shares that preference. We predict that one should be more likely to extend this property for a highly coherent category, such as yacht club members, than for a less coherent category, such as county clerks.

In Experiment 1B, we test the nonobvious prediction that low coherence categories will lead to stronger inductions to noncategory members (i.e., members of the general population). High coherence categories tend to have members with common deep features and relations, making them both more similar to each other and less similar to the population, compared to less coherent categories. The more coherent a social category, the more its members may be viewed as different from the average person in both known and yet to be discovered ways. In other words,

although greater coherence promotes inductions to other category members, it might also lead to weaker inductions to randomly chosen people from the population.

EXPERIMENT 1A

Participants read scenarios involving 6 category members with a novel preference and were asked to assess the likelihood that the next member would have the same preference. The sample size was chosen with the intuition that it was large enough to motivate a category-based inference, but small enough to plausibly be due to something else, such as a chance run. If coherence influences judgment, we should see greater likelihood estimates for high as compared with low coherence categories. If, instead, likelihood judgments are based exclusively on the presence of strong sample data, responses should not differ systematically across categories.

Method

Participants. Twelve undergraduates (4 men and 8 women) at the University of Illinois were paid for their participation in this 30-min task.

Materials. Thirty-two categories were drawn from an existing database of job and hobby categories for which uniformity ratings ("On a scale of 1–9, how similar are members of this category to one another?"), and frequency estimates ("Estimate the number of people in the United States per 1,000 who belong to this category") were available. The 16 high coherence categories had uniformity ratings ≥ 5 ($M = 5.7$), while the 16 low coherence ones had ratings ≤ 4 ($M = 3.4$). The two sets were matched on estimated frequency of category members in the population at large (Ms = 10/1,000) and number of job and hobby categories to eliminate two possible confounds. See Table 1 for all categories and a subset of properties.

Procedure. Participants worked at their own pace on booklets of 32 problems (16 high and 16 low coherence), presented in either of 2 random orders, with 3 problems per page. A representative high coherence problem follows:

> Think about violinists and what they are like. What is one important characteristic associated with this category?

> Suppose that half of all people prefer broccoli and half prefer cauliflower. Imagine that the first 6 violinists you encounter prefer broccoli to cauliflower. How likely do you think it is that the next violinist you encounter will prefer broccoli to cauliflower as well? _____ (0%–100%)

The first question ensured cognitive processing of the category; responses were not analyzed. Properties (e.g., vegetable preference) were relatively blank; that is, though generally familiar, they were not associated with the categories, based on experimenter intuitions.

Results

The overall mean likelihood judgment was 63%. Across participants, the mean for high coherence problems (66%, $SE = 2.9$) was significantly higher than for low coherence ones (60%, $SE = 2.9$) [$t(11) = 5.66, p < .001$]. The same pattern was found across problems [$t(15) = 3.41, p = .004$]. A scatterplot of category coherence by mean likelihood judgment (Figure 1) illustrates the relationship between these variables [$r(31) = .37, p = .038$, with only the low coherence category of "novelist" visibly different from the overall pattern; $r(30) = .59, p < .001$, with this 1 of the 32 categories removed].

Table 1
Experiment 1B Choice Results and High and Low Coherence Categories Used Across Experiments

High Coherence			Low Coherence				Low-Coh
Category	Coh	Freq	Category	Coh	Freq	Property	Choice Results
military general	6.9	2	scuba diver	3.2	3	*VISA* to *MasterCard*	71%
brain surgeon	6.6	3	quilter	3.9	2	*wood floors* to *wall-to-wall carpeting*	29%
duck hunter	6.3	10	tour guide	3.4	11	*study of physics* to *study of chemistry*	86%
rare-sculpture collector	6.1	5	radio disc jockey	4.0	3	*Newsweek* to *Time magazine*	64%
tango dancer	5.7	4	novelist	1.8	2	*Labor Day* to *Memorial Day*	64%
nurse	5.7	18	book-club member	3.0	25	*going bowling* to *going biking*	38%
violinist	5.7	11	baker	3.8	14	*eating broccoli* to *eating cauliflower*	57%
computer programmer	5.6	10	world traveler	3.1	11	*watching NBC* to *watching ABC*	79%
mime	5.6	3	massage therapist	3.8	5	*risk of tornadoes* to *risk of hurricanes*	86%
pediatrician	5.5	11	camper	3.6	16	*brick houses* to *stone houses*	51%
children's librarian	5.4	4	bungee jumper	3.6	3	*morning exercise* to *evening exercise*	57%
accountant	5.3	20	star-gazer	2.6	30	*London* to *Rome*	64%
surfer	5.2	3	comedian	3.1	5	*espresso* to *cappuccino*	59%
gambler	5.2	41	apartment manager	3.3	15	*phone* to *email*	36%
skateboarder	5.1	12	taxi driver	3.8	15	*the color red* to *the color blue*	64%
lumberjack	5.1	5	model-train collector	4.0	4	*game shows* to *sitcoms*	57%

Note—Coherence (Coh) is on a 9-point scale from 1 (low) to 9 (high). Frequency (Freq) refers to number of people out of 1,000 (with individual responses log-transformed, averaged, and reverse-transformed to reduce influence of extreme scores). Properties are a subset from Experiments 1A and 2. Choice results are from Experiment 1B, in which more abstract blank properties (e.g., "Vegetable A") were used.

EXPERIMENT 1B

In addition to predicting that inductions from sample data to new category members should be greater for high coherence categories, we also predict that the inductions to the general population should be less for high coherence categories. We conducted a preliminary study with 16 participants, identical to Experiment 1A except that the induction question asked about an individual drawn randomly from the population rather than another category member. As predicted, this time we found a small, but significant, advantage for inference judgments for the *low* coherence categories, 49% versus 46% [$t(15) = 2.17, p = .046$]. A scatterplot of category coherence by mean likelihood judgment (Figure 1) illustrates the relationship between these variables [$r(30) = -.27, p = .142$, with the "novelist" category removed for consistency with Experiment 1A]. However, though the preliminary results complemented Experiment 1A, the small size of the difference led us to consider ways to address the issue without relying on absolute judgments, thereby motivating Experiment 1B.

In Experiment 1B, we extended the results to a task that directly set high and low coherence categories in opposition. Specifically, scenarios were presented in which 6 individuals from a high coherence category had one property and 6 individuals from a low coherence category had the complementary property. A novel individual was introduced who belonged to *neither* category and participants were instructed to record the individual's most likely preference. Abstract predicates (e.g., "Book A") were used to ensure that participants had no expectations about property distributions in the population. We expected to find that the novel individual would now most often be deemed to have the property associated with the low coherence rather than the high coherence category.

Method

Participants. Fourteen undergraduates (5 men and 9 women) at Wesleyan University participated in exchange for course credit in this 30-min task.

Materials. The materials were the same as those used in previous experiment, except for the following changes, as illustrated by the example below.

List one typical characteristic of bakers: _____
List one typical characteristic of violinists: _____

Suppose that Book N and Book O are two popular books. Imagine that:
The first 6 bakers you encounter prefer Book N to Book O.
The first 6 violinists you encounter prefer Book O to Book N.
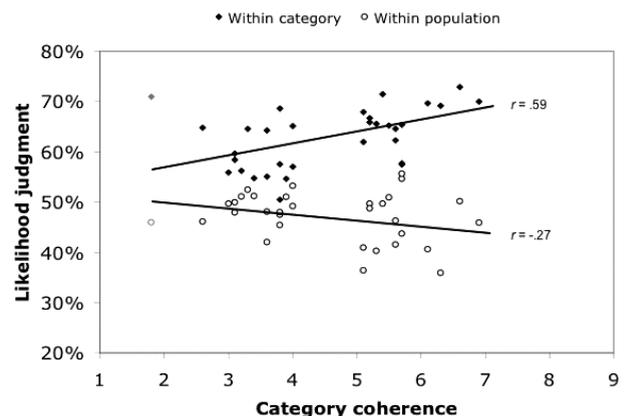
Figure 1. Relationship between category coherence and judged likelihood of property extension. The solid circles are the 32 data points from Experiment 1A in which the extension is to another category member. The unfilled circles are from the pilot study to Experiment 1B in which the extension is to a member of the general population.

You then encounter someone who is *neither* a baker nor a violinist. Which book do you think this person prefers, Book N or Book O? _____

Problems were constructed by using the 32 categories from Experiment 1A to make 16 new problems. Each problem paired a high and a low coherence category matched on frequency (see Table 1 for pairings). Two versions of the materials were made, each with different pairings of categories and properties and different random orders. Abstract blank predicates (e.g., "Book O") were used to ensure that participants had no expectations about distributions of properties in the population. Participants were instructed to choose between the properties associated with each category.

Procedure. The procedure was the same as Experiment 1A, except that the instructions were changed slightly to reflect the new task.

Results

No effect of materials was found ($p > .10$) so reported results are collapsed over materials. The mean percentage of times each response was associated with the low coherence category was 59% ($SE = 3.5$), which is reliably greater than chance (50%) [$t(13) = 2.57, p = .020$]. The same pattern was found across problems [$t(15) = 2.04, p = .057$], with 81% of problems in the expected direction, as shown in the last column of Table 1.

DISCUSSION OF EXPERIMENTS 1A AND 1B

Experiment 1A reveals greater willingness to extend a property from experienced category instances to a new category member when a category is high as compared with low coherence. This is complemented by Experiment 1B, which shows greater willingness to generalize properties of low as compared with high coherence categories to the population at large.

While the difference between high and low coherence problems was relatively small in Experiment 1A, the effect must be considered in light of the fact that judgments over 70% were rarely given despite the low probability (<5%) of encountering 6 category members in a row with the same preference by chance alone. This suggest poor understanding of the laws of probability or a belief that prior knowledge offers little basis for explaining connections between the categories and properties. Regardless, the results provide clear evidence of the influence of category coherence on category-based induction in the context of sample data.

EXPERIMENT 2

Now that we know coherence influences inductions from sample data, a basic question is whether this effect depends on the size of the sample—is there an interaction between coherence and sample size? In Experiment 2, we extend the results of Experiment 1A by considering the influence of category coherence across a range of sample sizes (e.g., experience with 3, 6, 9, or 12 yacht club members who prefer Chinese to Mexican food). One possibility is that coherence and sample size have additive influences, with the influence of coherence constant across levels of sample size. Another is that there is an interaction such

that coherence boosts perceived inductive strength only when the data do not independently provide strong evidence, such as when a sample size is small.

Method

Participants. Twenty-four undergraduates (12 men and 12 women) at the University of Illinois participated in exchange for course credit in this 30-min task.

Materials. The materials were the same as Experiment 1A except that 4 versions of the booklet were created. Within each version, each of 4 sample sizes (3, 6, 9, and 12) was used for 8 problems; these were counterbalanced across booklets.

Procedure. The procedure was the same as Experiment 1A except that problems were presented in a different random order for each participant (rather than in 2 random orders for all participants).

Results and Discussion

No effect of materials was found ($p > .10$) so reported results are collapsed over materials. The results are illustrated in Figure 2. The figure shows a main effect of coherence, replicating Experiment 1A, in that higher likelihood judgments were given for high coherence as compared with low coherence categories [$F(1,23) = 17.23, MS_e = 33.95, p < .001$]. There was also a main effect of sample size, in that likelihood judgments increased with sample size [$F(3,69) = 14.43, MS_e = 78.92, p < .001$, for basic effect; $F(1,23) = 19.30, MS_e = 172.52, p < .001$, for linear trend]. A linear interaction between coherence and sample size was also found. Linear increases in likelihood judgment in response to sample size were greater for low coherence as compared with high coherence categories [$F(3,69) = 2.19, MS_e = 46.14, p = .097$, for basic interaction; $F(1,23) = 6.31, MS_e = 45.32, p = .019$, for linear interaction].

An analysis by problem yielded similar results. There was again a linear effect of sample size [$F(1,30) = 14.49, MS_e = 66.30, p < .001$, for linear trend], and a marginally significant linear interaction between coherence and sample size [$F(1,23) = 3.89, MS_e = 66.30, p = .058$, for linear interaction]. The effect of coherence not did not reach significance [$F(1,30) = 1.93, MS_e = 103.71, p = .175$].

The results suggest that when the empirical evidence in the form of sample data is weak, coherence is recruited to assess the likelihood that a property extends to a category. When the empirical evidence is greater, the data
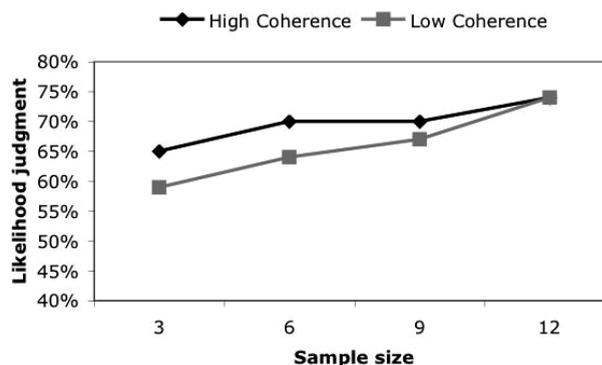

Figure 2. Mean likelihood judgment as a function of category coherence and sample size in Experiment 2.

provide strong independent evidence and coherence no longer matters much. An alternative explanation is that the two factors are additive but that there is a response ceiling of about 75% beyond which, for whatever reason, people rarely respond. We cannot rule out this ceiling explanation, however, such a ceiling is unlikely given that all participants had maximum judgments well above 75% (12/15 were ≥90% and the remaining 3 were 85%).

GENERAL DISCUSSION

These studies focus on the critical question of the interplay of category coherence and experience with instances, and extend our understanding in three ways. First, although Patalano et al. (2006) showed that category coherence predicts extension of general category knowledge to individual category members, the current findings show that the influence occurs even with instance data. Specifically, there is a greater willingness to extend properties across individual members of high as compared with low coherence categories. Second, the same coherence that makes for stronger inductions within a category makes for weaker inductions to the population because coherence serves to differentiate category members from the average member of the population. Third, the influence of coherence, however, does decrease with increasing instance experience. For small samples, patterns (e.g., 3 yacht club members who prefer Chinese to Mexican food) can be attributed to chance unless there is some reason, such as coherence, to believe that the category members might share some property. With larger samples, the data are overwhelming and presumably cannot be attributed to chance and so are attributed to category membership regardless of the category's coherence.

Accounting for the Findings

While the present studies show that coherence influences induction from sample data, they do not directly address the question of *how* coherence might be used in reasoning. We consider two possible accounts. First, coherence is highly correlated with rated similarity (our operationalization is a similarity measure, uniformity) which, unlike past accounts, is likely to include not just shared features, but also weighting of underlying features and interconnections among features (Rehder & Hastie, 2001). It may be that this similarity is used as a heuristic for estimating coherence, even without conscious assessment of the nature of the similarity. For example, an evidence accumulation process (similar to random walk models) might be used in which the depth or interconnectedness of activated features serve as evidence for the overall coherence of the category.

A second account of these results is that coherence supports explanation. Past research suggests that people often use explanation as evidence for inference (Heit & Rubinstein, 1994; Sloman, 1997). Sloman (1997), for example, found greater willingness to transfer the property "makes a good bodyguard" from ex-convicts to weight lifters than to bachelors, because the first two categories use the same explanation that members must be strong.

Because high coherence categories have more deep features and stronger interrelations among features, one would expect that they would promote more plausible explanations as well. In Experiments 1a and 2, when the samples could be attributed to chance, the higher coherence might be more likely to provide an explanation. However, with large runs, even the low coherence categories would be assumed to have some explanation behind the run. For Experiment 1B, the more plausible explanation for the high coherence category would tend to rely on the category and thus make it less likely to extend to the population. Such explanations, or even an estimation of the likelihood of being able to construct them (which could itself be based on similarity), might provide a means by which coherence and experiential information might influence inductive inference.

We have no evidence that would help decide between these explanations for the current data, though the estimation account would need a way to incorporate information about sample size to account for the findings of Experiment 2. More generally, our best guess is that each of these influences of coherence may occur under different circumstances. Speculatively, large number of decisions or the need for speeded responses might lead people to make more heuristic assessments, whereas important decisions or ones in which the property is related to the categories, might lead to more explanation-based responses.

Category-Based Induction With Real-World Samples

In these studies, the sample data that provided experiential evidence—a summary of how many category members had a particular property—admittedly did not represent a typical real-world unfolding of information. For instance, there was no category knowledge that was incrementally experienced over time, such as meeting and eating lunch with a series of yacht club members or county clerks, as would be more common in everyday life. There are a number of differences between the experimental task and many real-world situations including the simultaneous versus sequential experience of exemplars, the salience of properties in common across exemplars, and opportunities to obtain explanatory information. It is not clear how such situations would influence findings here. On the one hand, all everyday inferences are likely less strong when exemplars are experienced individually and over time because there is little reason to encode seemingly category-irrelevant properties in memory. On the other hand, the differences between more and less coherent categories might be magnified by the fact that fewer experiences are needed for transfer across high coherence categories. Future work is needed to explore these issues.

Conclusions

Both real-world category knowledge and instance-based sample data are often available as sources of inductive inference. Using natural social categories varying in coherence, we found that the overall coherence of a category influences property induction from instances.

Category coherence both increases the strength of within-category property inductions and decreases the strength of inductions to other members of the population. The work extends a growing literature on the importance of coherence for category learning and use in induction.

REFERENCES

HASLAM, N., ROTHSCHILD, L., & ERNST, D. (2000). Essentialist beliefs about social categories. *British Journal of Social Psychology*, **39**, 113-127.

HEIT, E. (2000). Properties of inductive reasoning. *Psychonomic Bulletin & Review*, **7**, 569-592.

HEIT, E., & RUBINSTEIN, J. (1994). Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **20**, 411-422.

KAHNEMAN D. & TVERSKY, A. (1973). On the psychology of prediction. *Psychological Review*, **80**, 237-251.

MURPHY, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.

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

NISBETT, R. E., KRANTZ, D. H., JEPSON, C., & KUNDA, Z. (1983). The use of statistical heuristics in everyday inductive reasoning. *Psychological Review*, **90**, 339-363.

PATALANO, A. L., CHIN-PARKER, S., & ROSS, B. H. (2006). The importance of being coherent: Category coherence, cross-classification, and reasoning. *Journal of Memory & Language*, **54**, 407-424.

REHDER, B., & HASTIE, R. (2001). Causal knowledge and categories: The effects of causal beliefs on categorization, induction, and similarity. *Journal of Experimental Psychology: General*, **130**, 323-360.

REHDER, B., & HASTIE, R. (2004). Category coherence and category-based property induction. *Cognition*, **91**, 113-153.

SLOMAN, S. A. (1997). Explanatory coherence and the induction of properties. *Thinking & Reasoning*, **3**, 81-110.

WISNIEWSKI, E. J. (2002). Concepts and categorization. In H. Pashler & D. Medin (Eds.), *Stevens' Handbook of experimental psychology: Vol. 2. Memory and cognitive processes* (3rd ed., pp. 467-531). Hoboken, NJ: Wiley.