

Structure in Category-Based Induction

Melissa Lin Wu (m-wu@nwu.edu)

Department of Psychology, Northwestern University
2029 Sheridan Rd., Evanston, IL 60208 USA

Dedre Gentner (gentner@ils.nwu.edu)

Department of Psychology, Northwestern University
2029 Sheridan Rd., Evanston, IL 60208 USA

Abstract

We investigated category-based inference tasks, contrasting the predictions of structural alignment theory as applied to categorization with those of feature-overlap models of similarity. We provide evidence for the differential level of importance of causal information in category-based inference tasks, as predicted by the systematicity principle (Gentner, 1983). Our basic paradigm consists of a task in which participants decide between inferences based on shared causal antecedents or shared attributes. Experiment 1 demonstrated a preference for the causal inference when the target animal shares one attribute with one of the base animals and one causal antecedent with the other base. In Experiment 2, we found that this preference holds even when the target animal shares greater attribute similarity with the noncausal base (i.e., the target shares *two* attributes with one base and one causal antecedent with the other). Experiment 2 also served to demonstrate that this result can indeed be attributed to the influence of causal structure, and not to surface stimulus properties, such as sentence length. Overall, the results agreed with the predictions of structural alignment theory and were inconsistent with a feature-overlap account.

Introduction

One important function of concepts is to support inferences --deriving new information from previous knowledge. For instance, given that you know that your dog has an appendix, you might infer this property for your cat. But you would probably not infer it for your vacuum cleaner. Given this, what are the constraints on the category-based inference process?

Two prominent theories of category-based induction posit that featural similarity is important for arguments like *Dogs have an appendix; therefore, cats have an appendix*. The similarity-coverage model (Osherson, Smith, Wilkie, Lopez, & Shafir, 1990; Osherson, Stern, Wilkie, Stob, & Smith, 1991) encodes categories as exemplars, which are ultimately comprised of features (e.g., "has fur," "has four legs"). The strength of the above argument (i.e., people's confidence in the truth of the conclusion given the truth of the premises) would be a function of the similarity between cat and dog, which is directly related to the number of shared features of their exemplars and inversely related to the number of distinctive features. The model also makes use of

hierarchical relations between categories. Exemplars of the lowest-level category that includes both dogs and cats, mammal, are compared to exemplars of the premise category, dog.

In contrast to the similarity-coverage model, the feature-coverage model (Sloman, 1993) does not assume hierarchical relations. Rather, categories at all levels (i.e., superordinate, basic, and subordinate) are represented by features. Argument strength is directly related to the featural overlap between premise and conclusion categories and inversely related to the familiarity and complexity of the conclusion category.

Although both of these models account for an impressive array of phenomena (see Osherson et al, 1990 and Sloman, 1993), there is room for an alternative. These models share the assumption that exemplars are represented as sets of features, which form the basis for the computation of feature-overlap or similarity. However, some recent evidence suggests that exemplar comparison is accomplished by alignment of structural representations (Clement & Gentner, 1991; Lassaline, 1996; Heit & Rubinstein, 1994). These results follow from an approach to similarity based on representations consisting of features and relations between them (Gentner, 1989; Medin, Goldstone, & Gentner, 1994). In *structural alignment*, parts of two structured representations are put into correspondence. Inferences are possible when predicates exist that are absent in one representation but connected to the common aligned system. According to the *systematicity principle*, systems of relations are preferred over individual features and higher-order relations (i.e., relations between relations), especially causal relations, are preferred over lower-order relations (Gentner, 1983). For instance, the fact that both the solar system and a hydrogen atom involve attraction and revolution is more important than the fact that both have spherical entities. Assuming the systematicity principle is at work in inference tasks, inferences connected to causal structure should be preferred over those unconnected.

In accord with this prediction, Clement and Gentner (1991) found that participants who read analogous stories preferred to make an inference from a shared causal antecedent, relative to an unconnected fact. They were also more confident in this inference.

Lassaline (1996) applied structural alignment directly to the issue of category-based inference. She demonstrated that adding a causal relation to an argument about animals increased its rated strength. For example, adding the relation *For Animal B, a weak immune system causes an acute sense of smell* increased the strength of the argument *Animal A has a weak immune system, skin that has no pigment, and dry flaky skin; Animal B has a weak immune system and an acute sense of smell; therefore, Animal A also has an acute sense of smell.*

In addition, Lassaline found that adding shared features (e.g., "has muscular forearms") to the premises also increased inductive strength. However, she did not address the question of which factor--relational or featural--is more important for inferential strength. If the systematicity principle holds for category-based inference tasks, causal relations should be more important, and we should see a preference for an inference connected to higher-order (causal) structure.

Experiments

Our experiments attempted to test whether category-based induction tasks are governed by the process of structural alignment.^{*} In particular, we tested the systematicity principle. We created a category-based induction task similar to Lassaline's (1996). For both experiments, we followed Lassaline in using fictional animals (e.g., Animal A). Stimulus items consisted of descriptions of one target animal and two base animals, as shown in Table 1. The "causal" base consisted of an unfamiliar but plausible causal relation and the "attribute" base consisted of a list of features. The target shared the causal antecedent of the causal animal and all but one attribute of the attribute animal. Participants were told to choose the stronger of two possible inferences that could be made about the target using the information about the bases. There were two obvious choices: the causal consequent from the causal base or the remaining attribute from the attribute base.

Attributes were grouped such that they were independent of each other, the causal antecedent, and the causal consequent. We varied which of the attribute animal's features were shared versus inferred. For instance, "has protruding canines" was a shared feature for some participants and a potential inference for others. The order of stimulus items in the packet was also varied between subjects.

^{*} In using the term "category-based induction," we do not necessarily assume that participants are basing inferences on knowledge of taxonomic relationships or on a categorical representation of "animal." Rather, the process of structural alignment may operate on instances, species or perhaps at an intermediate abstraction of a group of species (e.g., at a level akin to pragmatic reasoning schemas, Cheng & Holyoak, 1985).

Experiment 1

Our objective in this experiment was to test the prediction that participants would prefer the systematic (causal) inference. This prediction follows from a structural account of similarity. In this study, the target animal shared one of the two listed attributes with the attribute animal. According to a "flat" or featural account of similarity, there should be no preference for either the attribute or causal inference, since both base animals are equally similar to the target. In drawing predictions from the flat similarity account, we assumed a dot product similarity computation with equal weighting of features, over features explicitly present in the three exemplars only. The causal animal description consisted of a statement that a causal antecedent caused a consequent ($A \rightarrow C$). Although not explicitly stated, it seemed likely that participants would assume that the causal animal had the causal antecedent (A) and the consequent (C). Had they done so, then the causal and attribute animals would be equally similar to the target. However, one might argue that participants' representation of the causal animal differed from that assumed. For instance, they may have assumed the base contained only the causal relation ($A \rightarrow C$) and the antecedent (A), or that it contained only the causal relation and neither the antecedent nor the consequent. However, in none of these cases does the computed similarity of the causal animal to the target exceed that of the attribute animal to the target. This is true assuming either a dot product or Tverskian computation (Tversky, 1977).

Sixteen Northwestern University undergraduates participated in partial fulfillment of a course requirement. Materials consisted of a packet of 16 of the stimulus items previously described.

Table 1: Sample Item from Experiment 1

The following facts are true of Animal A [*the Attribute base*]:

- Has large oil glands
- Has protruding canines

The following facts are true of Animal B [*the Causal base*]:

- Adenoviruses cause an increased risk for tumors

The following facts are true of Animal C [*the Target*]:

- Has large oil glands
- Has an adenovirus

Please infer only one property for Animal C (circle one only):

- Has protruding canines
- Has an increased risk for tumors

Note: Information in brackets was not presented in experimental materials.

The results show a strong preference for causal inference. We assigned participants to the "predominantly causal" group if they chose the causal consequent on at least 13 of 16 trials, a significantly greater number that would be expected by chance, $p < .01$. Figure 1 shows the results of the sign test on participants. Thirteen out of 16 participants had predominantly causal inferences, a significantly greater number of participants than would be expected by chance, $p < .05$. No participant had predominantly attribute inferences (13 or more attribute inferences), and the remaining three participants demonstrated a pattern that did not differ from chance. The same procedure was done on items, with similar results.

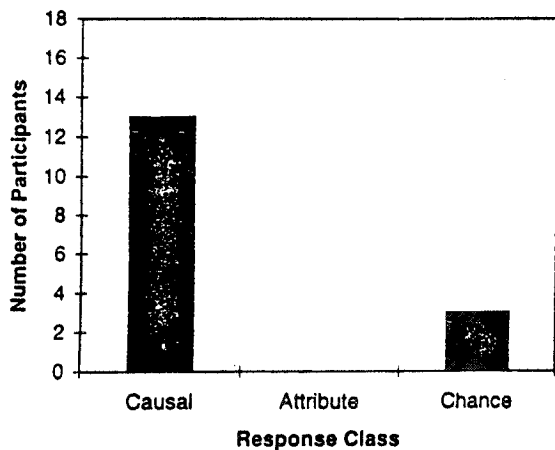


Figure 1: Number of participants in each response class, Experiment 1.

These results bear out the prediction of structural alignment theory: there appears to be a preference to infer from a base connected to causal structure, as predicted by the systematicity principle. The finding of a causal preference (Clement & Gentner, 1991) appears to apply within a category-based inference task (Lassaline, 1996), even when featural similarity is equalized. However, before drawing conclusions, we wished to address some possible concerns of this study. First, perhaps participants were initially attending to only the featural similarity between the target and each base, but since the bases did not differ in this regard, they felt compelled to use another strategy. Under this scenario, participants inferred from the causal base because it was more distinctive. A more definitive test for the preference for causal inference would involve constructing stimuli in which the attribute base is more featurally similar to the target than is the causal base.

Another concern is with the syntactic differences between the descriptions of the causal and attribute animals. The

causal animal description consisted of one long statement, while the attribute animal description consisted of facts on separate lines. Perhaps participants chose the causal consequent simply because the causal information was presented in a way that made it seem more important or salient.

A final concern is with generalizability. The same set of 16 causal consequents was used for all participants. This leaves open the possibility of stimulus effects. For example, if the causal choices happened to yield more applicable or enduring inference choices, then the preference for the causal choice could not be taken to imply a preference for higher-order structure per se. For instance, people may be more willing to infer "bears many offspring" than "lives in Madagascar" because the first fact is true of more animals than the second. Likewise, the first fact seems like a better inference than "is sleepy" because it is a less transient property. Although the differences in our stimuli were never this extreme, it may be the case that the facts used as causal consequents were slightly more applicable and enduring, and were favored by participants as inferences because of it.

Experiment 2

This experiment was designed to determine whether a preference for inferences from a base connected to causal structure holds when using a stronger test. We contrasted the causal animal with an animal with high attribute similarity to the target animal. As in Experiment 1, the target animal shared a causal antecedent with the causal animal. However, this time, the target shared *two* out of three independent features with the attribute animal, rather than one of two independent features, as in Experiment 1. In this design, the attribute base exceeds the causal base in its degree of feature overlap to the target. Thus, a flat similarity view predicts a preference to infer from the attribute base, whereas a structural similarity account predicts a preference to infer from the causal base.

A second objective was to rule out alternative interpretations of the data from the first experiment. To make the causal and attribute base descriptions more similar in appearance to each other, and thus eliminate possible differences in salience, we added the information that the causal base animal had the causal antecedent and consequent properties to the causal base description. We also wished to ensure that the causal consequent property was chosen as an inference because of the causal relation connecting it to the shared causal antecedent, and not because of any other conceptual differences between it and the second feature of the attribute base. To this end, we varied whether a given possible response was associated with the causal base or the attribute base. Specifically, for a given item, half of the subjects saw one of the inference choices (e.g., "high risk for strokes") as the causal consequent and the other choice (e.g., "underdeveloped visual system") as the second attribute feature, while for the other half, the correspondences were reversed.

Sixteen paid Northwestern University undergraduates participated. Materials consisted of a packet of 12 stimulus items like the sample shown in Table 2.

The data again demonstrate a preference for the causal inference. We assigned participants to the "predominantly causal" group if they chose the causal consequent on at least 11 of the 12 experimental trials, a significantly greater number than would be expected by chance, $p < .01$. As shown in Figure 2, 14 out of 16 participants had predominantly causal inferences, a significantly greater number of participants than would be expected by chance, $p < .01$. One participant had predominantly attribute inferences, and the remaining participant had a pattern of responses that did not differ from chance. The same analysis was done on items, with similar results.

Table 2: Sample Item from Experiment 2

The following facts are true of Animal A [*the Attribute base*]:

- Has muscular forearms
- Has sensitive gums
- Has high risk for strokes

The following facts are true of Animal B [*the Causal base*]:

- Has an overactive thyroid
- Has an underdeveloped visual system
- An overactive thyroid causes an underdeveloped visual system

The following facts are true of Animal C [*the Target*]:

- Has muscular forearms
- Has sensitive gums
- Has an overactive thyroid

Please infer only one property for Animal C (circle one only):

- Has high risk for strokes
- Has an underdeveloped visual system

These results rule out the alternative interpretations for our earlier findings and support the systematicity principle. If we count all shared assertions, then by flat similarity models, participants should prefer the inference from the attribute animal, because it shares two of three features with the target, whereas the causal animal shares only one of three features. Thus, in a situation where a flat similarity account would predict a clear preference for a particular attribute inference, participants overwhelmingly choose the causal consequent. These results are evidence for the role of structural similarity in category-based induction.

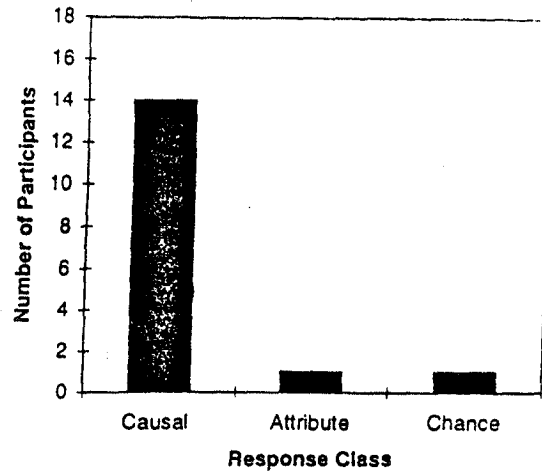


Figure 2: Number of participants in each response class. Experiment 2.

Discussion

These studies attempted to shed light on the issue of whether featural or structural similarity is more crucial in evaluation of category-based inferences. The two experiments reported in this paper suggest that connection to causal structure drives inference. When given two inference choices, people prefer to make the inference from the base that is connected to causal structure in the target, over the base that shares independent features. This generalization is true even when the alternative inference is from a target with greater attribute similarity with the base (i.e., with a larger set of shared independent features). Taken together, these findings would appear to contradict the predictions from feature-overlap versions of a similarity-based account of category-based induction (i.e., Sloman, 1993; Osherson et al., 1991).

Our findings add to the body of literature demonstrating the influence of systematicity and structural similarity on inference. Gentner, Rattermann, and Forbus (1993) asked participants to rate the inferential soundness, or degree to which one could make accurate inferences, of matched stories. Rated soundness was positively related to the degree of shared relational structure. Bowdle and Gentner (in press) found that people were more likely to make an inference from a more systematic story to a less systematic story than the other way around. Moreover, the inferences tended to be information connected to common causal structure.

Other evidence points to the role of systems of interrelated relations in category-based induction. Heit and Rubinstein (1994) found that people make stronger inferences in a category-based inference task when the kind of property to be inferred (anatomical or behavioral) matches the kind of similarity between the animal in the premise and the animal

in the conclusion (anatomical or behavioral). For instance, people judge the likelihood that whales travel in a zig-zag trajectory higher when told that tuna have that property, relative to when they are told that bears have that property. This is presumably because whales and tuna match behaviorally (they both swim) while whales and bears do not (they match anatomically--both are mammals).

Heit and Rubinstein's results can be seen as evidence for the importance of shared structure linked to the property to be inferred. Tuna and whales can be represented as structures containing features and relations connecting these features. Moreover, one might argue that participants assume that some (behavioral, and perhaps, anatomical) attributes and relations of tuna are causally related to the property of traveling in a zig-zag trajectory. In order to evaluate the argument, participants align their representations of tuna and whale. They then search for support for the inference that whales travel in a zig-zag trajectory. They examine attributes and relations that are causally connected to the property in tuna. If these are shared by the target, whale, then they contribute to the support measure. The more shared relevant attributes and relations that are meaningfully connected to the candidate inference, the greater the support for the candidate inference.

We have argued that similarity is indeed important for inference, but that similarity involves more than simple feature lists. Rather, representations encode features and background knowledge, theories, and causal relations.

Acknowledgments

We thank Ginger Gunn, Phil Wolff, Jeff Loewenstein, Shelli Feist, Ken Kurtz, and Ken Forbus for comments on this research. This research was funded by the National Science Foundation and the Office of Naval Research.

References

- Bowdle, B.F., & Gentner, D. (in press). Informativity and asymmetry in comparisons. *Cognitive Psychology*.
- Cheng, P.W., & Holyoak, K.J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology*, 17, 391-416.
- Clement, C.A., & Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. *Cognitive Science*, 15, 89-132.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D., Ratterman, M., & Forbus, K.D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology*, 25, 524-575.
- Heit, E., & Rubinstein, J. (1994). Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 411-422.
- Lassaline, M.E. (1996). Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 754-770.
- Medin, D.L., Goldstone, R.L., & Gentner, D. (1994). Respects for Similarity. *Psychological Review*, 100, 254-278.
- Osherson, D.N., Smith, E.E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological Review*, 97, 185-200.
- Osherson, D.N., Stern, J., Wilkie, O., Stob, M., & Smith, E.E. (1991). Default probability. *Cognitive Science*, 15, 251-269.
- Sloman, S.A. (1993). Feature-based induction. *Cognitive Psychology*, 25, 231-280.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84, 327-352.