

The Role of Correlational Structure in Learning Event Categories

Alan W. Kersten
Dorrit O. Billman

School of Psychology
Georgia Institute of Technology
Atlanta, GA 30332-0170
email: billman@pravda.gatech.edu

Abstract

How do people learn categories of simple, transitive events? We claim that people attempt to recover from input the predictive structure that is the basis of 'good', inferentially rich categories. Prior work with object categories found facilitation in learning a component relation (e.g. feathers covary with beak) when that correlation was embedded in a system of other, mutually relevant correlations. Little research has investigated event categories, but researchers have suggested that verb meanings (hence perhaps event categories) might be organized quite differently from noun meanings (and object categories). Thus it is far from clear whether the learning biases or procedures found for object categories will also appear for event learning. Two experiments investigated the effects of systematic correlational structure on learning the regularities comprising a set of event categories. Both found the same pattern of facilitation from correlational coherence as found earlier with object categories. We briefly discuss relations to 1) other constraints on concept learning that focus on the organization of the whole system of concepts and 2) learning paradigms that produce competition, not facilitation, between correlated cues.

Event Category Learning

What makes some event categories harder or easier to learn than others? By events we mean simple, "verb-sized" interactions between agents and patients. Categories are generalizations across multiple such events. In our experiments subjects view animations of simple, transitive events in an unsupervised learning paradigm and we assess what regularities they learn. Our work has two broad motivations: 1) to identify what makes some systems of categories natural and coherent but others

arbitrary and ad hoc and 2) to assess whether proposals developed for object categories apply to event categories.

We propose that 'natural', coherent categories 1) support useful predictions about new instances and 2) facilitate the learning of attribute relations within a category and relations to contrast categories. We look at ease of learning, particularly in unsupervised tasks, as an important index of category 'goodness'.

Our experiments ask how the organization of a system of categories affects learning components of that system. We believe that understanding coherence or ease of learning requires considering a system of categories, not each category in isolation. Investigating system-wide, structural constraints on learning has been a small but visible component of current research. Keil's (1979) concept of predicability, Markman's (1989) mutual exclusivity constraint, and perhaps the suggestions about the role of background theory in categorization (Murphy and Medin, 1985) are examples of work in this area.

We believe that the notion of correlational structure is critical for understanding category coherence. Rosch (1978) claimed that a good category captures rich correlational structure in the world and hence is inferentially rich; knowing something is a bird allows you to predict many of its properties. We extend Rosch's claim about structure to two learning principles. First, learners are biased to seek out categories with rich correlational structure so that learning any component correlation (say between feathers and flying) is facilitated if other attribute values also correlate (singing and having a beak). Thus we predict facilitation among correlated cues (Billman and Heit, 1988), not competition as in several other learning models (Rescorla & Wagner, 1972; Gluck & Bower, 1988). Second, learning benefits if correlations among the same attributes are present consistently across a system of contrast categories

(values of body covering and locomotion correlate for FISH as well as BIRD).

The current experiments test the first principle, asking if learning a correlation between two attributes is facilitated in a system where these attributes also covary with others. We compare learning the correlation in this context to learning the same correlation in differently organized systems. These experiments paralleled prior work on object categories which found support for this principle (Billman & Jeong, 1989; Billman & Knutson, 1990). To investigate events we had to identify attributes relevant to event categories; we did this by reference to work on perceptual properties of simple events (Michotte, 1946/1963) and to work identifying aspects of event meaning that are reliably identified in syntax and verb meaning.

Although event categories are widely acknowledged as important, concept research has overwhelmingly focused on object categories. Given the paucity of research on event categories, it was unclear whether the results for event categories would parallel those for objects. The closest research comes from study of verbs, not event categories per se. Gentner (1981) has hypothesized that relational concepts such as verbs have little correlational structure. Huttenlocher and Lui (1979) further claim that verbs, unlike nouns, possess a matrix-like organization, with little correlation between elements of meaning such as direction, instrument, intent, and manner. If real-world event categories indeed have little correlational structure, category learners would most likely not have a bias toward learning novel event categories with such structure. However, should correlational structure facilitate the learning of event as well as object categories, this would suggest that the presence of a coherent system of correlations plays a very pervasive role in learning. In addition, a commonality between object and event categorization would be identified, encouraging comparison between analyses of object categories (and nouns) and event categories (and verbs).

Experiment 1

Given the view that categories capture correlational structure, we investigated the ease of learning component correlations. The first experiment tested whether a correlation is more easily learned in isolation or when the attributes participating in that correlation also participate in other correlations. We predicted that correlations would be easier to learn in the presence of other correlations, due to a bias of category learners to seek out such correlational structure.

Method

Subjects. Thirty-six undergraduates at the Georgia Institute of Technology volunteered as subjects for class credit.

Stimuli.

Learning Phase. Every event consisted of the actions of a square character, the *agent*, and a circular character, the *patient*, on a varying background, the *environment*. The characters left their starting locations when the subject pressed the mouse button, with the agent always moving towards the patient. When the agent reached the patient, the *state change*, or change in appearance of the patient, took place. The patient then always moved away from the agent.

Each event varied on 7 attributes, each with 3 possible values. Three attributes specified static properties of the objects or their environment. Four specified dynamic properties which combined to produce collisions, chases etc. with different outcomes for the patient. The attributes were: (1) the agent color: red, green, or blue; (2) the patient color: purple, brown, or yellow; (3) the nature of the environment: a fine grain, squiggly lines, or small ovals; (4) the state change: blowing up, shrinking, or flashing; (5) the path of the agent after the state change: movement toward the patient, away from the patient, or remaining at the place where it met the patient; (6) the path of the patient before the state change; and (7) the agent's manner of motion: smooth, direct motion; oscillation perpendicular to the direction of motion; or surging forward in bursts.

For each subject, at least two of these attributes were correlated, such that the value of one attribute could be predicted given the value of the other. The correlation between these two attributes was designated the *target rule*. This was the only correlation present for subjects in the *isolating condition*. For subjects in the *structured condition*, these two attributes were correlated with two other attributes chosen at random. Three different target rules were used and we compared the ease of learning each target rule in the structured versus isolating conditions. Figure 1 shows an example of the correlations present in the structured and isolating conditions.

Test Phase. Each subject's knowledge of his or her target rule was tested. This was done by collecting ratings of events with correctly or incorrectly matched values of the attributes in the target rule. To test knowledge of the target rule in isolation, it was necessary to hide any attributes which were correlated with the attributes in that rule. If this were not done and one attribute had an incorrect value in the structured condition, multiple rules




Figure 1: Example schemas for two subjects in Experiment 1. Dark lines indicate correlations. Target rule between environment and state change is present in both conditions. For structured condition, these two attributes are correlated with two other attributes.

besides the target would also be incorrect. Thus, two correlated attributes were hidden from view on each test event for structured condition subjects. Two random attributes were hidden for each test of the target rule for isolating condition subjects. To make the test phase minimally instructive, two filler 'rules' were also tested. For structured condition subjects, these were two correlations which were present in the learning phase. For isolating condition subjects, these rules had not been present, so these subjects could not have had knowledge of the correct pairings on the filler rules. The target rule and each filler rule were each tested 18 times. On half of these trials, the values of the attributes in the rule were correctly matched, while on half they were matched differently than in the learning phase.

Procedure. The learning phase consisted of 120 animated events. Subjects were instructed that they would be seeing events on another planet, and were to learn the kinds of events which took place on this planet. Each event was initiated by the subject. There were five breaks during learning. After the learning phase, subjects were told that their knowledge of events would be assessed. Six correct, instruction displays familiarized subjects with how attributes would be hidden from view. Subjects were then told to rate each test display, based on the available attributes, for how good an example it was of possible events on planet Daysee. Subjects could repeat each of the 54 events they were to rate up to 3 times.

Design. The design consisted of two factors. One of these was correlational structure. The two levels

of this factor were structured and isolating. The other factor was the target rule. The three correlations comprising this factor were between patient color and agent path, state change and environment, and manner of motion and patient path. The dependent measure was rating accuracy on the 18 events testing the target rule.

Results

Accuracy scores were derived from rating scores by finding the difference between each subject's rating and the correct rating and then subtracting this from 2. Thus, a perfect rating for an event was awarded a score of 2, while the most incorrect rating received a -2.

Structured condition subjects had higher accuracy scores, averaging .89, compared to the average isolating condition score of .31. The correlation between state change and environment was rated most accurately, followed next by patient color and agent path, and finally by patient color and manner of motion (see Figure 2).

Due to significant heterogeneity of variance, a Brown-Forsythe test was performed instead of an ANOVA. This test revealed a significant effect of correlational structure, with an $F(1,15)$ of 8.82 ($p < .01$). The effect of rule was also found to be significant, producing an $F(2,15)$ of 15.83 ($p < .001$). The interaction was insignificant, with an $F(2,15)$ of 0.56.




Figure 2: Experiment 1 accuracy scores for the three rules. Rule 1: state change and environment. Rule 2: patient color and agent path. Rule 3: patient path and agent manner of motion.

Experiment 2

The results of Experiment 1 supported the notion that correlational structure, in the sense of multiple related correlations, facilitates the learning of individual correlations. However, there is an alternative interpretation of the data. Since four of the seven attributes were correlated for the structured condition, only three attributes varied randomly. In contrast, five of the seven attributes varied randomly for the isolating condition. Thus, the effects attributed to correlational structure in Experiment 1 could instead have been a result of differences in the amount of randomness in the two conditions.

Experiment 2 was an attempt to replicate Experiment 1 while controlling for the amount of randomness. Randomness was conceptualized as the number of possible events allowed within the specification of correlational structure for a given subject. For example, in the isolating condition of Experiment 1, the five randomly varying attributes were free to form any combination with one another and with the two correlated attributes, allowing 3^6 possible combinations. The structured condition only allowed 3^4 combinations (from A1-A2-A3-A4 all covarying with each other). In Experiment 2, the number of possible combinations in the comparison, or crossed, condition was reduced to 3^4 by introducing two correlations which were independent of the target rule (A1-A2, A3-A4, and A5-A6 pair-wise correlations). Thus, the number of possible events was equalized across conditions,

while the crossed condition was still low in systematic, correlational structure.

Method

Subjects. Thirty undergraduates at the Georgia Institute of Technology volunteered as subjects for class credit.

Stimuli.

Learning Phase. The learning phase differed from that of Experiment 1 in that subjects in the crossed condition saw correlations between three pairs of attributes. No attribute was correlated with more than one other attribute for these subjects. Every structured condition subject was presented with correlations among the same four attributes: agent path, environment, manner of motion, and state change (see Figure 3).

Test Phase. The test phase differed from that of Experiment 1 in that crossed condition subjects were tested on each of the three rules present in the learning phase 18 times. One attribute from each pair which was not being tested was hidden for each test phase event for these subjects. The test phase procedure was identical to that of Experiment 1 for structured subjects.

Procedure. Same as Experiment 1.

Design. The design of Experiment 2 differed from that of Experiment 1 in the rule factor. The three target rules for this experiment were between agent path and environment, agent path and manner of motion, and manner of motion and environment.

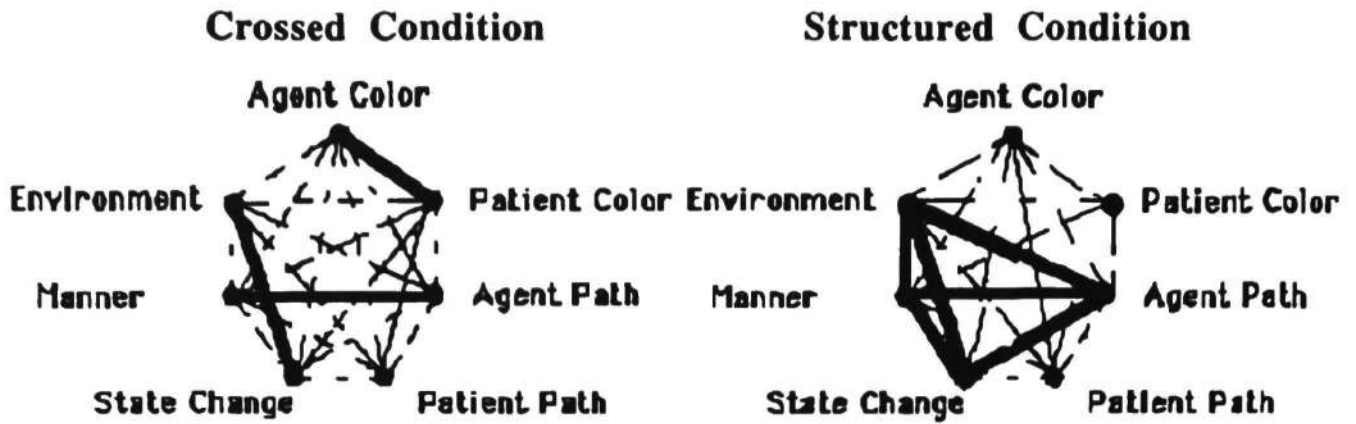


Figure 3: Schemas for two conditions in Experiment 2. Dark lines indicate correlations. Target rule between manner and agent path is present in both conditions. For structured condition, these two attributes are correlated with two other attributes.

Results

Structured condition subjects again scored higher in rating accuracy. These subjects averaged 1.08, compared to .29 for crossed subjects. The correlation between agent path and manner of

motion was easiest to learn, followed next by agent path and environment, and finally by manner of motion and environment (see Figure 4). An ANOVA revealed significant effects of correlational structure ($F = 8.18, p < .01$) and rule ($F = 3.49, p < .05$). The interaction was not significant ($F = .03$).

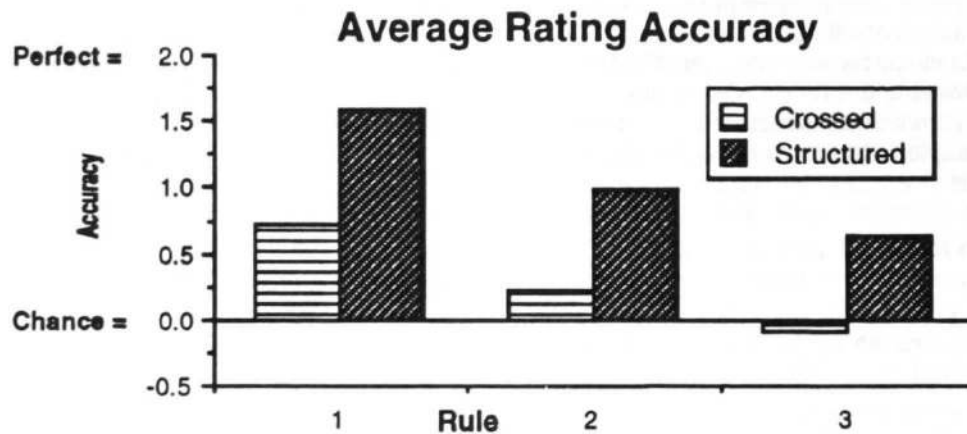


Figure 4: Experiment 2 accuracy scores for the three rules. Rule 1: agent path and agent manner of motion. Rule 2: agent path and environment. Rule 3: environment and agent manner of motion.

Discussion

Our experiments found that a systematic set of correlations determining event categories facilitated learning of a component relation. Experiment 1 showed that a given relation between two event attributes was more likely to be learned when that relation was part of a system of interpredictive attributes (Structured Condition) than in a 'simpler' system where the target rule was the only rule in the system (Isolating Condition). Experiment 2 showed facilitation of learning the target rule embedded in a more structured system over a system where there were a) the same number of possible events but b) differently organized correlations.

This research investigated correlational structure of events. We believe that other types of relations, most notably causal ones, are also important and are part of peoples' representations of event categories. However, we have focused on the 'data-driven' aspect of correlation; the learner can observe correlations in data, but can only infer cause. Clearly, the two are related. In particular, we believe that a learner biased to recover clusters of mutually relevant, correlated attributes will be finding just those correlations likely to reflect a common underlying cause. Thus, we investigated how properties of the data affect data-driven learning, but anticipate that the biases or principles investigated are just those that usefully interface with theory-driven processes.

Our findings with event categories parallel our earlier findings with object categories and suggest that the study of events can be a means of investigating the domain generality of principles initially proposed for object categories. Our findings also suggest that verb meanings may have more correlational structure than has been noted, if one takes the correspondence between verbs and events and between nouns and objects seriously. For example, verbs such as "gallop" and "read" suggest strong predictions about other components of the described event: gallop will probably have a horse for its agent; "read" will have an animate (and literate) agent and something written as patient. Event concepts and verb meaning may interact in interesting ways.

In both experiments we did not just ask if it is easier to learn about a category when there were many category predictors than when there were a few, but whether an individual, identical component pattern is learned faster. This identifies one important way in which the organization of a category system as a whole impacts learning its components. Further, the finding of facilitation among components identifies a useful bias for

unsupervised learning and suggests a quite different view than that of competitive cue models from highly supervised tasks (e.g. classical conditioning).

References

- Billman, D.O. & Heit, E. 1988. Observational learning from internal feedback: A simulation of an adaptive learning method. *Cognitive Science* 12:58-69.
- Billman, D., & Jeong, A. 1989. Systematic correlations facilitate learning component rules in spontaneous category formation. Paper presented at the Psychonomics Society, Atlanta, November, 1989.
- Billman, D., & Knutson, J. 1990. Systematic versus orthogonal sets of correlations benefit unsupervised concept learning. Paper presented at the Psychonomics Society, New Orleans, November, 1990.
- Gentner, D. 1981. Some interesting differences between verbs and nouns. *Cognition and Brain Theory* 4:161-178.
- Gluck, M.A., & Bower, G.H. 1988. Evaluating an adaptive network model of human learning. *Journal of Memory and Language* 27:166-195.
- Huttenlocher, J., & Lui, F. 1979. The semantic organization of some simple nouns and verbs. *Journal of Verbal Learning and Verbal Behavior* 18:141-162.
- Keil, F.C. 1979. *Semantic and Conceptual Development: An Ontological Perspective*. Cambridge, Mass.: Harvard University Press.
- Markman, E.M. 1989. *Categorization and naming in children: Problems of induction*. Cambridge, Mass.: MIT Press, Bradford Books.
- Michotte, A. 1963. *The perception of causality*. (T.R. Miles & Elaine Miles, Trans.). New York: Basic Books (Original work published 1946).
- Murphy, G.L. & Medin, D.L. 1985. The role of theories in conceptual coherence. *Psychological Review* 92:289-316.
- Rescorla, R.A., & Wagner, A.R. 1972. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A.H. Black, & W.F. Prokasy eds., *Classical conditioning: II. Current research and theory*. New York: Appleton-Century-Crofts.
- Rosch, E. 1978. Principles of categorization. In E. Rosch, & B.B. Lloyd, *Cognition and categorization*. Hillsdale, NJ: Erlbaum.