

Correlated properties in artifact and natural kind concepts

Ken McRae

Department of Psychology
Meliora Hall
University of Rochester
Rochester, NY 14627
kenm@psych.rochester.edu

Abstract¹

Property intercorrelations are viewed as central to the representation and processing of real-world object concepts. In contrast, prior research into real-world object concepts has incorporated the assumption that properties are independent and additive. In two studies, the role of correlated properties was explored. Property norms had been collected for 190 natural kinds and artifacts. In Experiment 1, property intercorrelations influenced performance in a property verification task. In Experiment 2, concept similarity, as measured by overlap of *independent* properties, predicted short interval priming latency for artifacts. In contrast, concept similarity, as measured by overlap of *correlated* property pairs, predicted short interval priming for natural kinds. The influence of property intercorrelations was stronger for natural kinds because they tended to contain a higher proportion of correlated properties. It was concluded that people encode knowledge about independent and correlated properties of real-world objects. Presently, a Hopfield network is being implemented to explore implications of allowing a system to encode property intercorrelations. Finally, results suggest that semantic relatedness can be defined in terms of property overlap between concepts.

Property-based representations of concepts were used to investigate aspects of the semantic representation of simple real-world objects, such as DOG and COUCH. Models based on conceptual primitives or properties have been used in a number of areas of cognitive science, including vision (Biederman, 1987), lexical representation (e.g., McClelland & Rumelhart, 1981), and concepts and categorization

(e.g., Gluck & Bower, 1988; Hintzman, 1986; Kruschke, 1992; Rosch & Mervis, 1975; Malt & Smith, 1984).

In previous investigations of real-world object concepts, the simplifying assumption has been made that concepts are composed of independent and additive properties (e.g., Rosch & Mervis, 1975). In direct contrast, a number of people have claimed that property co-occurrences are an important organizing principle for objects in the world (e.g., Boyd, 1984; Keil, 1989; Malt & Smith, 1984; Medin, 1989; Rosch et al., 1976). In fact, it has been claimed that concepts are organized around clusters of intercorrelated properties. For example, the concept BIRD is assumed to be organized around a cluster of intercorrelated properties that includes: HAS WINGS, HAS FEATHERS, FLIES, and HAS A BEAK.

Evidence for encoding and use of property co-occurrence information has been found in **artificial** concept experiments by Medin and Schwanenflugel (1981), Medin et al. (1982), and Younger and Cohen (1983). Due to these findings, recent computational models of categorization have incorporated mechanisms capable of simulating effects of correlated properties (e.g., Gluck, Bower, & Hee, 1989; Hintzman, 1986; Kruschke, 1992; Nosofsky, 1986). However, only one study has investigated the influence of correlated properties in **real-world** concepts. Malt and Smith (1984) found **no** evidence to support the hypothesis that correlated properties predict typicality judgements above and beyond the predictive power of independent properties. Therefore, at this point in time, despite the fact that property intercorrelations are viewed as central to concepts and categorization, no empirical evidence exists to suggest that they are encoded in conceptual representations of real-world objects. The purpose of this paper was to explore potential effects of property intercorrelations in semantic tasks involving real-world object concepts.

¹ This research was partially supported by an FCAR (Quebec) graduate fellowship and an NSERC (Canada) postdoctoral fellowship.

Target property: HUNTED BY PEOPLE

<u>Concept</u>	<i>DEER</i>	<i>DUCK</i>
<u>Correlated Property</u>	<u>Shared variance (%)</u>	<u>Correlated Property</u> <u>Shared variance (%)</u>
1) is herbivorous	71.9	1) an animal 22.0
2) has antlers	68.3	2) lives in water 14.9
3) lives in the woods	49.2	3) migrates 12.9
4) lives in the wild	43.3	4) swims 10.9
5) a mammal	23.5	
6) an animal	22.0	
7) is brown	14.8	
8) has hooves	10.2	
9) has 4 legs	10.0	
10) has fur	6.8	
11) has legs	6.6	

strength (DEER) = 71.9+68.3+49.2+43.3+23.5+22.0+14.8+10.2+10.0+6.8+6.6 = 326.6

strength (DUCK) = 22.0+14.9+12.9+10.9 = 60.7

Note: shared variance refers to percentage of shared variance between HUNTED BY PEOPLE and the listed property.

Figure 1 - Computing strength of intercorrelation

Semantic property norms were collected for 19 exemplars from 10 categories, 190 in total. There were 4 natural kind categories: BIRDS, MAMMALS, FRUITS, and VEGETABLES; and 6 artifact categories: CLOTHING, FURNITURE, KITCHEN ITEMS, TOOLS, VEHICLES, and WEAPONS. Three hundred subjects were divided into 10 groups. Each group was asked to produce lists of properties for 19 of the 190 exemplars. The exemplars and property lists were used to create stimuli and representations for the following experiments.

Experiment 1 - Property Verification

I tested the hypothesis that intercorrelational strength among properties within a concept influences performance in a property verification task. In this task, a concept name (e.g., DEER) was presented, followed by a property name (e.g., HAS ANTLERS). A subject was asked to indicate whether or not the property was reasonably true of the object to which the concept name referred. For example, the correct response was "yes" to DEER-HAS ANTLERS, but "no" to HORSE-HAS SCALES.

To select stimuli, a target property was paired with two concepts: one in which the correlations between the target property and other properties within it were as strong as possible, and a second concept in which the correlations between the target property and other properties within it were as weak

as possible. It was hypothesized that a property that was strongly intercorrelated with other properties within a concept would be more easily verified than a property that was weakly intercorrelated.

Method

Subjects. Twenty McGill university undergraduates were paid \$2 each to participate.

Materials. The stimuli were 37 pairs of concepts, each sharing a target property. The intercorrelations between the target property and other properties in one concept were as strong as possible, but were as weak as possible in the matched concept. For example, as shown in Figure 1, HUNTED BY PEOPLE is more strongly intercorrelated with properties of DEER than properties of DUCK. The stimuli were chosen using the following algorithm. Properties that were part of fewer than three concepts were removed, then proportion of shared variance was computed between each property pair. Proportion of shared variance was computed between properties represented as 190-element vectors of production frequencies (1 element per concept). Strength of intercorrelation of a property within a concept was measured as the sum of the proportion of shared variance between the target property and all properties within the concept with which it was significantly correlated. This computation is illustrated in Figure 1.

Mean intercorrelational strength differed significantly between groups (strong: 175; weak: 20). Because a critical property appeared in both groups, all variables associated with it, such as reading time, were equated on an item by item basis. The groups were also equated for production frequency of a target property, concept familiarity, and number of properties produced for a concept. Two lists were constructed so that a subject saw each target property only once. Further methodological detail can be found in McRae (1991).

Procedure. Each trial proceeded as follows: a 1500 ms intertrial interval; an asterisk in the centre of the screen for 500 ms; a blank screen for 100 ms; the concept for 400 ms; and the property until the subject responded. Subjects were instructed to press the "yes" key as quickly and accurately as possible if the property was reasonably true of the concept, or press the "no" key otherwise.

Results

One-tailed paired samples t-tests were used to determine the effect of intercorrelational strength on decision latency and (square root) errors. Subjects were faster to judge that a property was part of a concept if there were strong intercorrelations (820 ms) than if there were weak or no intercorrelations (912 ms), $t(19) = 4.22, p < .0003$ by subjects, $t(36) = 4.23, p < .0002$ by items. Furthermore, fewer errors were committed for items in the strong condition (4.3%) than in the weak condition (13.0%), $t(19) = 4.34, p < .0003$ by subjects, $t(36) = 3.24, p < .002$ by items.

Discussion

A clear effect of property intercorrelations was found on the ease with which properties were verified. Property verification necessarily involves computation of an exemplar concept and computation of the meaning of a property. The effect of property intercorrelations was attributed to computing the meaning of the target property, rather than computing the exemplar concept. Given that production frequency is assumed to indicate the strength with which a property is activated when a concept name is read, there should be no difference between groups due to computing exemplar concepts because production frequency was equated across groups. In contrast, if computing the meaning of a property is viewed as initiating a pattern completion process, its effect on verification is apparent. If co-occurrence information is encoded in memory, then when a property's representation is computed, properties correlated with it become activated. Because the item pairs were

constructed so that the target property was more highly intercorrelated with the properties of one concept, the pattern for that concept would have been completed to a greater extent by the target property, resulting in an easier verification decision. For example, when the meaning of HUNTED BY PEOPLE was computed, it more strongly activated properties associated with wild mammals that have antlers and live in the woods, than properties associated with birds. In summary, it can be concluded that knowledge about independent and correlated properties is encoded in memory for the meaning of words.

Experiment 2 - Short Interval Semantic Priming

If property intercorrelations are encoded in semantic memory, then concepts must be represented in terms of independent and correlated properties. Therefore, similarity between concepts may depend on shared correlated property pairs, as well as independent properties. I investigated this possibility with a short interval semantic priming task. Each target appeared with a similar prime (e.g., EAGLE-HAWK) and a dissimilar prime (e.g., SANDALS-HAWK). Short interval semantic priming should be a sensitive measure of concept similarity for two main reasons. First, a number of studies strongly suggest that short interval priming is sensitive to effects associated with computing the meaning of the prime and target (e.g., de Groot, 1984; den Heyer, Briand, & Dannenbring, 1983; Neely, 1977). Second, because each target appears with a similar and a dissimilar prime, it serves as its own control. Therefore, variance due to the decision, as well as variance due to nuisance variables such as target length and frequency, can be factored out in the analyses. It was hypothesized that the magnitude of priming effects can be predicted by concept similarity in terms of independent and correlated properties.

Method

Subjects. Forty-eight McGill university undergraduates were paid \$3 each to participate.

Materials. Ninety similar prime-target pairs served as stimuli, 9 similar prime-target pairs from each of the 10 categories (e.g., TRUCK-VAN, EAGLE-HAWK). In every case, the prime and target had been normed by different subjects. In each of the 4 tasks, a subject made a semantic decision to the target, the decision being: "is it animate?", "is it an object?", "is it made by humans?", or "Does it grow?". Two lists were constructed so that no subject saw a target twice.

Further methodological detail can be found in McRae (1991).

Procedure. Each subject performed all 4 tasks, which were blocked. Each trial proceeded as follows: a 1500 ms intertrial interval; an asterisk in the centre of the screen for 250 ms; a blank screen for 250 ms; the prime for 200 ms; a mask, &&&&&&&&&&, for 50 ms; and the target until the subject responded. Subjects were instructed to silently read the prime, then respond as quickly and accurately as possible to the target. Examples were given to portray what was meant by: an animate thing (something that is alive, a cockroach is but a keyboard is not); an object (something that is tangible or concrete, a pair of scissors is but the sky is not); something made by humans (something that is manufactured by people, a razor is but a butterfly is not); and something that grows (something that grows on its own, a spider grows but a door does not).

Results

The frequency of errors for each task was less than 4% and were not further analyzed. Two items, EMU and STARLING, were removed because it was apparent from this and other studies that their referents were unknown to many subjects.

Subjects were faster to make decisions to targets preceded by similar primes in 3 of the 4 tasks: "is it animate?" (63 ms priming effect), $F(1, 188) = 13.10$, $p < .01$ by subjects, and $F(1, 96) = 12.46$, $p < .01$ by items; "is it an object?" (50 ms), $F(1, 188) = 9.98$, $p < .01$ by subjects, and $F(1, 96) = 15.70$, $p < .01$ by items; and "is it made by humans?" (53 ms), $F(1, 188) = 6.93$, $p < .01$ by subjects, and $F(1, 96) = 8.82$, $p < .01$ by items. The priming effect in the "does it grow?" task was not significant (23 ms), $F(1, 188) = 3.40$, $p > .5$ by subjects, and $F(1, 96) = 1.66$, $p > .1$ by items.

Regression analyses. In the independent properties representation, each concept was represented as a list of individual properties weighted by production frequency. Properties that were produced for an exemplar by a minimum of one-sixth of the subjects were included, resulting in 1,242 properties to describe the 190 exemplars. Concept similarity was computed by measuring the cosine of the angle between concept vectors in 1,242-dimensional property space. To construct a correlated properties representation, the set of 1,242 properties was first reduced to the 240 that were part of three or more concepts. Proportion of shared variance between each pair of properties was computed; it measured the tendency for two properties to occur together in the world (as described by the norms). The correlated properties representation contained one to ten vectors for each pair of significantly correlated

properties ($p < .01$), contingent upon amount of shared variance. The number of units assigned to each property pair is shown in parentheses: (10) $R^2 > .964$; (9) $R^2 > .864$; (8) $R^2 > .764$; (7) $R^2 > .664$; (6) $R^2 > .564$; (5) $R^2 > .464$; (4) $R^2 > .364$; (3) $R^2 > .264$; (2) $R^2 > .164$; (1) $R^2 > .064$.

With this algorithm, properties that participated in a greater number of correlated property pairs received greater representation, as did properties that participated in highly correlated pairs. A concept was represented as a pattern across the resulting 2,630 vectors, including 1,190 unique property pairs. A value was given to each unit on the basis of a three-part function. If a concept had neither property from the correlated property pair, the unit was set to 0. If a concept possessed both properties, the unit was set to the sum of the production frequencies. If a concept contained one of the properties, then it violated the correlation and the unit's value was the negated production frequency of the property that was possessed by the concept. Therefore, the more strongly a concept possessed one property of the pair without the other, the greater the violation. For example, the correlated property pair, FLIES-HAS FEATHERS, shared 42.5% of their variance. The pair was represented by 4 units. According to the norms, a CARROT neither FLIES nor HAS FEATHERS (activation = 0); an EAGLE both FLIES and HAS FEATHERS (activation = 13 + 16 = +29), and an OSTRICH HAS FEATHERS, but never FLIES (activation = -22). Concept similarity was measured by the cosine of the angle between concept vectors in the 2,630-dimensional space defined by these pairs.

In the regression analyses, prime-target concept similarity was used to predict mean response latency for similar prime-target pairs (LAMP-CHANDELIER) after mean response latency for dissimilar prime-target pairs (GOOSE-CHANDELIER) had been entered. Similarity for dissimilar prime-target pairs was excluded from the regression analyses because it carried little information. Regression analyses were conducted separately for the complete set of 88 stimulus pairs, the 54 artifact pairs, and the 34 natural kind pairs. Similarity, as measured by overlap of independent properties, was used to predict amount of priming. Because any existing theory assumes that independent properties are encoded in semantic memory, the burden of proof was placed on demonstrating an effect of correlated properties; the predictive ability of correlated properties was tested with concept similarity in terms of independent properties already forced into the regression equation. The independent properties predicted a significant proportion of variance in priming effects for the 88 prime-target pairs ($R^2 = .125$, $p < .001$), but the correlated

properties did not ($R^2 = .004$, $F < 1$). More importantly, overlap of independent properties accounted for a significant proportion of variance of priming effects for artifact prime-target pairs ($R^2 = .150$, $p < .005$), but not for natural kinds ($R^2 = .040$, $p > .2$). Conversely, for natural kinds ($R^2 = .211$, $p < .009$), but not for artifacts ($R^2 = .003$, $F < 1$), overlap of property intercorrelations accounted for a significant proportion of residual variance of priming effects.

Discussion

A clear influence of property intercorrelations on concept similarity was found. Artifact similarity is captured by independent properties, and natural kind similarity by the correlations among them. These results are consistent with those found by McRae (1991) in a same/different category decision task.

Predictions for natural kinds were more dependent on property intercorrelations for two reasons. A greater proportion of natural kind property pairs (10.9%) than artifact pairs (5.8%) were significantly correlated ($p < .01$). Natural kinds also had much richer representations over property pairs. On average, they possessed a significantly greater number of correlated property pairs than did artifacts (83 versus 36). In contrast, there was no difference in terms of individual properties (natural kinds: 17; artifacts: 15). Furthermore, on average, for the stimuli of Experiment 2, natural kind prime-target pairs shared a significantly greater number of correlated property pairs (55) than did artifact prime-target pairs (22).

Why did successful prediction by correlated properties preclude prediction by independent properties? Concept similarity in terms of independent properties increases monotonically with number of shared properties. In contrast, it was possible for an additional shared property to decrease concept similarity in terms of correlated property pairs. If concept_a possessed both properties of a correlated pair, then concept_b was more similar to concept_a if it possessed neither property than if it possessed one property of the pair (i.e., violated the correlation). Predictions based on the independent properties representation may have been affected to the extent that these cases were present. Across the 88 similar prime-target pairs used in Experiment 2, the mean number of cases where one concept possessed both properties of a pair and the other possessed only one was significantly greater for natural kinds (43) than for artifacts (17). Therefore, concept similarity in terms of independent properties failed to make reliable predictions about natural kinds

because of its insensitivity to shared and violated correlated property pairs.

In summary, because natural kinds possess a rich structure of property intercorrelations, similarity among natural kind concepts is largely determined by overlap of correlated property pairs. In contrast, similarity among artifacts is largely determined by independent properties; because few property intercorrelations exist for artifacts, they carry little weight in determining concept similarity. These results provide support for an hypothesis in the philosophical literature concerning differences between natural kinds and artifacts. Boyd (1984) has argued that natural kinds are natural precisely because they are structured around sets of "contingently clustered" properties. That is, Boyd has claimed that the network of causally interconnected properties is much richer for natural kinds than for artifacts, confirmed empirically in the analyses reported above. Similarly, Keil (1989) has claimed that natural kinds have an "essence", which is the underlying source of a dense network of property intercorrelations. In contrast, artifact properties have weaker causal interconnections because they cohere around intended function of the creator.

Semantic relatedness. Experiment 2 is also important from the perspective of investigating the source of short interval semantic priming effects. To date, these effects have been attributed to "semantic relatedness", but the source of semantic relatedness has not been clearly defined. For example, in a recent study by Hodgson (1991), he compared priming effects for six types of semantic relatedness. He found that no type of relatedness produced consistently greater priming effects. I have shown that short interval semantic priming effects can be predicted by concept similarity as measured by independent properties and correlated property pairs. In other words, semantic relatedness can be defined, at least in part, by property-based concept similarity.

Conclusions

Property intercorrelations are viewed as an important element of human knowledge. Although it has been found that property intercorrelations affect performance in artificial concept studies, and recent models of categorization have incorporated mechanisms that can account for these effects, an effect of correlated properties on semantic tasks involving real-world object concepts had not been demonstrated. The experiments reported here provide strong evidence that property intercorrelations are encoded in conceptual representations of real-world objects and affect performance on semantic tasks. Therefore, any reasonable model of the representation and use of concepts must incorporate a mechanism by

which co-occurrences among properties are encoded. Currently, Virginia de Sa and I are exploring implications of using a Hopfield (1982) network to learn and process property-based conceptual representations.

References

- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*, 409-429.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*, 115-147.
- Boyd, R. (1984). Natural kinds, homeostasis, and the limits of essentialism. Paper presented at Cornell University.
- de Groot, A. M. B. (1984). Primed lexical decision: Combined effects of the proportion of related prime-target pairs and the stimulus-onset asynchrony of prime and target. *The Quarterly Journal of Experimental Psychology*, *36A*, 253-280.
- den Heyer, K., Briand, K., & Dannenbring, G. L. (1983). Strategic factors in a lexical-decision task: Evidence for automatic and attention-driven processes. *Memory & Cognition*, *11*, 374-381.
- Gluck, M. A., Bower, G. H., & Hee, M. R. (1989). A configural-cue network model of animal and human associative learning. In *Proceedings of the Eleventh annual Conference of the Cognitive Science Society*, Ann Arbor, Michigan. August 16-19, 1989 (pp. 323-332). Hillsdale, NJ: Erlbaum.
- Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, *93*, 411-428.
- Hodgson, J. M. (1991). Informational constraints on pre-lexical priming. *Language and Cognitive Processes*, *6*, 169-264.
- Hopfield, J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Science, USA*, *79*, 2254-2258.
- Keil, F. C. (1989). *Concepts, kinds, and cognitive Development*. Cambridge, MA: MIT Press.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, *99*, 22-44.
- Malt, B. C., & Smith, E. E. (1984). Correlated properties in natural categories. *Journal of Verbal Learning and Verbal Behavior*, *23*, 250-269.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part I. An account of basic findings. *Psychological Review*, *88*, 375-407.
- McRae, K. (1991). *Independent and correlated properties in artifact and natural kind concepts*. Unpublished doctoral dissertation, McGill University, Montreal.
- Medin, D. L. (1989). Concepts and conceptual structure. *American Psychologist*, *44*, 1469-1481.
- Medin, D. L., Altom, M. W., Edelson, S. M., & Freko, D. (1982). Correlated symptoms and simulated medical classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *8*, 37-50.
- Medin, D. L., & Schwanenflugel, P. J. (1981). Linear separability in classification learning. *Journal of Experimental Psychology: Learning, Memory, and Memory*, *7*, 355-368.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*, 289-316.
- Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*, *106*, 22-254.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39-57.
- Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch and B. B. Lloyd (Eds). *Cognition and categorization* (pp. 259-303). Hillsdale, NJ: Lawrence Erlbaum.
- Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, *104*, 192-233.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, *7*, 573-605.
- Younger, B. A., & Cohen, L. B. (1983). Infant perception of correlations among attributes. *Child Development*, *54*, 858-867.