Feature–feature causal relations and statistical co-occurrences in object concepts

CHRIS McNORGAN, RACHEL A. KOTACK, DEBORAH C. MEEHAN, AND KEN McRAE
University of Western Ontario, London, Ontario, Canada

Influences of feature–feature statistical co-occurrences and causal relations have been found in some circumstances, but not others. We hypothesized that detecting an influence of these knowledge types hinges crucially on the congruence between the task and type of knowledge. We show that both knowledge types influence tasks that tap feature relatedness. Detailed descriptions of causal theories were collected, and co-occurrence statistics were based on feature production norms. Regression analyses tested the influences of these knowledge types in untimed relatedness ratings and speeded relatedness decisions for 65 feature pairs spanning a range of correlational strength. Both knowledge types influenced both tasks, demonstrating that causal theories and statistical co-occurrences between features influence conceptual computations.

Lexical concepts are ubiquitous in human cognition, and language succinctly transmits a great deal of information about them via the computation of word meaning. Concepts are often described in terms of semantic features, and people readily provide features for object concepts in production tasks (McRae, Cree, Seidenberg, & McNorgan, 2005; Rosch & Mervis, 1975). For example, for the word *robin*, a participant might list *flies*, *has feathers*, *has wings*, *has a beak*, *lays eggs*, *builds nests*, and *eats worms*. Of course, people know much more about concepts than simply lists of features. In particular, they have extensive knowledge of various distributional statistics regarding environmental structure (Cree & McRae, 2003) as well as background and theory-based knowledge that influences people's use of concepts (Murphy & Medin, 1985).

An important confluence of distributional statistics and theory-based knowledge concerns the manner in which features co-occur. Features do not occur independently of one another; there is statistical structure in the patterns of feature co-occurrence across concepts. That is, there is a continuum of variation in the degree to which the presence of one feature signals the presence of another. For example, *has feathers* and *has a beak* are highly correlated because the various types of birds that have feathers are very likely to have a beak as well. In contrast, *has a tail* and *has hooves* are relatively weakly correlated because things in the world that have hooves always have a tail, but there are many types of animals that have a tail but do not have hooves. Moreover, for some correlated feature pairs, people possess a theory for why they are correlated, such as the fact that *has wings* is causally related to *flies*. It has been shown multiple times that people are aware (and can learn) that certain features co-occur within concepts (Chin-

Parker & Ross, 2002; Malt & Smith, 1984). Furthermore, knowledge of feature co-occurrences and causal relations has been shown to influence performance in a range of tasks, from offline feature inferences (Gelman, 2003) to online feature verification (McRae, Cree, Westmacott, & de Sa, 1999).

In fact, the tasks used to tap knowledge of feature co-occurrences and the manner in which people learn about concepts are both important factors in showing influences of feature co-occurrences (Chin-Parker & Ross, 2002). Although many studies have focused on categorization per se and the intentional category learning paradigm, the utility of various aspects of conceptual representation and computation is much broader than solely categorization or laboratory category learning tasks. Concepts are used for many purposes, including language comprehension and production, object recognition, conceptual combination, and making various sorts of inferences. The present research investigates the manner in which knowledge of feature–feature causal relations and statistical co-occurrences forms part of our conceptual representations and influences conceptual tasks. In the present article, we use people's judgments regarding how features are related to one another to provide insight into these central aspects of concepts.

Causal Relations Between Features

In this article, the term *causal relations* is used to refer to explicit beliefs about causal connections between features, in concert with definitions used previously (Ahn, Marsh, Luhmann, & Lee, 2002; Murphy & Wisniewski, 1989; although *causal theories* often refers to a wider range of causally based knowledge). A causal theory between two features explains how one feature, such as *has*

K. McRae, mcrac@uwo.ca

a blade, is causally responsible for a second, such as *used for cutting* (Murphy & Medin, 1985). Theories of causal relations among features are central to concepts and categorization. For example, they allow one to infer properties of category members (Gelman, 2003). Evidence also suggests that categories may cohere around causal relations. For example, Rehder and Hastie (2004) found that when causal relations within categories were violated, feature induction and category classification ratings were lower, suggesting that these relations play a role in determining category membership.

Statistical Co-Occurrences Between Features

In this article, knowledge of *statistical co-occurrences* is used to refer to implicit knowledge of the statistical tendency for two features to be correlated. In the present research, two features were considered correlated if they tended to appear in the same basic-level concepts. For example, although someone may have never explicitly considered that things that grow on trees are often used in pies, they may have frequently encountered pies made of apples, peaches, or cherries. Furthermore, if asked, people typically are able to verify these co-occurrences (i.e., they will likely concur that things that grow on trees are often used in pies). However, it appears that for most feature pairs of this sort, people rarely think explicitly about their co-occurrence, nor do they necessarily possess a theory about why these features tend to go together. Nonetheless, some argue that people possess this knowledge and that it influences behavior (Chin-Parker & Ross, 2002; McRae, de Sa, & Seidenberg, 1997).

We focus on feature co-occurrences across basic-level concepts for two main reasons. First, the vast majority of previous research on feature correlations in natural concepts has focused on the basic level, and feature co-occurrence at this level provides insights into conceptual representations and processing (Ahn et al., 2002; Malt & Smith, 1984). More pragmatically, we possess feature production norms for 541 living and nonliving things, which allows the computation of statistical co-occurrences for a large number of feature pairs across a sizeable sample of concepts. Experiment 1 includes an explanation of how these correlations were calculated. It is, of course, possible to study feature correlations at other levels, such as within basic-level categories, such as *dog*. Although the present studies do not directly address co-occurrences at the subordinate level, they presumably underlie people's basic-level concepts and thus influence the features that they produce when given basic-level concept names, such as *dog*.

Evidence for Causal Relations

Task type can play a crucial role in determining whether a researcher finds evidence for the computation and use of a specific type of knowledge (Jones & Smith, 1993). Tasks such as speeded feature verification are associated with rapid, relatively automatic processes that emphasize lower level or implicit knowledge. In contrast, tasks such as untimed typicality ratings allow participants to engage in higher level reasoning and to integrate information from a number of sources. Thus, because knowledge

that is closer to the implicit end of the continuum (such as statistical co-occurrences) is expected to differ from more explicit knowledge (such as causal relations) with respect to the time at which the information is available and is used, speeded and untimed tasks may not be equally sensitive to these knowledge types (McRae et al., 1997; Sloman, Love, & Ahn, 1998).

Causal relations have been shown to influence a number of untimed tasks. For example, Malt and Smith (1984) examined how typicality judgments are affected by broken correlations, which occur when only one member of a correlated feature pair appears in a concept (such as having wings, but not being able to fly). Not all correlated feature pairs had the same effect on typicality ratings, leading Malt and Smith to speculate that knowledge of statistical co-occurrences may be limited to only certain salient relations.

Ahn et al. (2002) hypothesized that the feature pairs in Malt and Smith (1984) that influenced typicality ratings were those for which participants possess a theory-based relation. Participants rated the correlational magnitude of Malt and Smith's feature pairs, and these ratings were used to distinguish between explicit (high magnitude) and implicit (low magnitude) pairs. In a subsequent experiment, participants indicated the nature and direction of the relationship with a labeled arrow and numerically rated its strength. Because a greater number of participants provided labeled relations for explicitly related pairs and these relationships were judged to be stronger, Ahn et al. (2002) concluded that explicitly related features were theory based. Finally, in a typicality rating experiment using artificial category exemplars, broken correlations influenced typicality ratings only when they involved explicit relations. Ahn et al. (2002) concluded that "the correlations that are present in our representations of real-world concepts are the ones with explanatory relations" (p. 115).

Evidence that causal relations are used in speeded tasks has not been demonstrated as clearly. Lin and Murphy (1997) presented participants with one of two cover stories explaining the form-function relations of parts of objects depicted in line drawings, and these cover stories emphasized contrasting relations between features. Background causal relations influenced response latencies in a speeded picture categorization task. However, these results' relevance to the use of background knowledge in the online computation of concepts is limited in two ways. First, response latencies averaged over 2,100 msec, which is quite long in comparison with the sub-1,000-msec response latencies found in speeded experiments demonstrating the effects of statistical co-occurrences (McRae et al., 1999). Second, the authors argued that background knowledge influenced the salience of critical object parts established during concept acquisition, rather than the computations required to perform the categorization task.

In Palmeri and Blalock (2000), participants learned to classify drawings that ostensibly were created by two groups of children. Group labels were either meaningful (drawn by "creative" vs. "noncreative" children) or meaningless ("Group 1" vs. "Group 2"). Background

knowledge regarding creativity influenced speeded categorization decisions in which participants were given a 300-msec response deadline. However, Palmeri and Blalock concluded that this resulted from a holistic response strategy. They argued that participants in the creative–noncreative condition responded on the basis of overall image complexity (creative children’s drawings tended to contain more detail); thus, these results do not clearly show that background knowledge is computed and used quickly in speeded tasks.

In summary, causal relations correspond to explicit knowledge that appears to take some time to compute. Accordingly, causal relations influence untimed tasks that are sensitive to them, but it is unclear whether they directly influence speeded versions of these same tasks.

Evidence for Statistical Co-Occurrences

In statistical learning approaches, such as attractor networks, statistically based knowledge of feature co-occurrences is highlighted. In most studies of feature correlations in artificial concepts, researchers have used perfectly correlated feature pairs (Chin-Parker & Ross, 2002). In the real world, however, correlation is a matter of degree; some pairs of features are much more strongly correlated than others. Attractor networks encode the degree to which features co-occur across concepts, and do so as part of learning what features belong to what concepts (see McRae, 2004, for a detailed discussion). A network’s knowledge of feature correlations that is encoded in feature–feature weights, for example, influences its settling dynamics and is the primary determinant of its semantically based pattern completion abilities (Cree, McRae, & McNorgan, 1999; McRae et al., 1999). Thus, because the rate of feature activation depends on the number and magnitude of correlations in which features participate, these networks predict that the influence of statistical co-occurrences should be most evident in speeded tasks that are sensitive to the temporal dynamics of conceptual computations.

McRae et al. (1997) computed statistical co-occurrences using a large set of feature production norms. They contrasted speeded feature verification with untimed feature typicality ratings (“How typical is *hunted by people* for *deer*?”). The degree to which a concept’s features were correlated with one another predicted feature verification latencies, but not feature typicality ratings. Similar effects were found with a speeded semantic similarity priming task versus an untimed similarity rating task. McRae et al. (1999) reported further evidence for the role of feature correlations in conceptual computations. In simulations of feature verification, their Hopfield (1982, 1984) network predicted an interaction between SOA (300 or 2,000 msec) and the degree to which a concept’s features are correlated with one another. In addition, the predicted interaction was in opposing directions depending on whether the concept name or feature name was presented first. These predictions were borne out by the human experiments and appear to be extremely difficult to account for using causal relations.

This discussion highlights the importance of congruency between knowledge type and task type: Speeded and

untimed tasks appear to be differentially sensitive to statistical co-occurrences and causal relations, thus making it difficult to directly compare their influences. However, congruency is not based solely on temporal parameters; the degree to which a task directly taps the type of knowledge in question also matters (Chin-Parker & Ross, 2002). Therefore, it is possible that if the correspondence between knowledge type and task type is sufficiently direct, an influence of both types of knowledge will be found even when the temporal parameters are less than optimal.

The present experiments address these issues. Experiment 1 provided detailed descriptive statistics on fully articulated causal theories between correlated feature pairs. Experiment 2 investigated the individual contributions of causal relations and statistical co-occurrences in an untimed relatedness rating task designed to directly tap people’s knowledge of how features co-occur across basic-level living and nonliving things. Finally, Experiment 3 used a binary feature relatedness decision task in a speeded analogue of Experiment 2. These experiments show that both causal relations and statistical co-occurrences influence both tasks.

EXPERIMENT 1

Although considerable research has shed light on the role of background knowledge and theories, the content of people’s theories remains unclear. Theories constructed on the basis of experimenter intuition (Murphy & Wisniewski, 1989) or explicitly stated relationships (Ahn et al., 2002) do not necessarily reflect people’s beliefs. When causal knowledge has been measured empirically, it has been described along the dimensions of direction (Ahn et al., 2002; Sloman et al., 1998), strength (Ahn et al., 2002), and depth (Ahn, Kim, Lassaline, & Dennis, 2000; Sloman et al., 1998). Causal direction is intuitively important because it determines when the presence of one feature can be inferred from the presence of another. For example, most people would agree that being *worn on the feet* generally causes objects to *come in pairs*; thus, this causal relationship is strong. Conversely, because *comes in pairs* often does not cause an object to be *worn on the feet*, the causal relation in this direction is presumably weaker or nonexistent. Causal strength depends not only on direction, but also on the degree to which the features covary. For example, there is low causal strength in either direction between the color of an appliance and its function because these features vary independently. Lastly, causal relations are transitive; that is, if A causes B, and B causes C, then A is an indirect cause of C. Causal depth expresses the role that a feature plays in such a chain; A has greater causal depth than does B, which in turn has greater causal depth than does C. Causal depth is important because relations involving deeper features may be more important for category membership (Ahn et al., 2000).

There are, however, other properties of causal theories that may be important to understanding how causal knowledge is learned, computed, and used. First, a theory’s explanatory content is not fully captured by the dimensions

of depth and direction; a complete description of causal knowledge should identify the actors and their relationships in a causal chain. Second, because a major issue concerns how causal theories and statistical co-occurrences influence one another during learning, it is important to distinguish between prior held—or *a priori*—theories versus *ad hoc* theories that may be generated during the course of an experiment. Because an effect cannot precede its cause, only *a priori* theories influence prior learning. In Experiment 1, we used open-ended interviews to collect fully elucidated causal explanations for a set of feature pairs that were correlated to varying degrees across basic-level concepts according to a large set of semantic feature production norms.

Method

Participants. Twenty-one University of Western Ontario students were paid \$15 for their participation. Because the interview required a relatively sophisticated understanding of the distinction between correlation and causation, only graduate and upper year undergraduate students were interviewed. The responses of 1 participant were dropped because only a single causal theory was provided, signifying a lack of effort. In all experiments reported herein, all participants were native English speakers and had either normal or corrected-to-normal visual acuity. No participant served in more than one experiment.

Materials. Feature–feature statistical co-occurrence was estimated using semantic feature production norms for 541 basic-level concrete concepts (McRae et al., 2005). In the norming task, participants listed features when given the names of approximately 20 concepts (30 participants listed features for each concept). Care was taken to ensure that no participant listed features for sets of similar concepts. For the present purposes, this means that participants did not produce sets of correlated features for similar concepts (i.e., the concepts most likely to give rise to such correlations). In the norming task, no mention was made of feature correlations; participants simply listed features for a set of rather dissimilar concepts. Thus, this avoided any circularity in using feature correlations calculated from the norms to predict feature relatedness ratings or binary feature relatedness decisions. Using the entire set of norms, which included responses from approximately 750 participants in total, all features were retained that were provided by at least 5 of 30 participants for a specific concept. A 541 concepts $\times$ 2,526 features matrix was then constructed, where each matrix element corresponded to the number of participants listing a specific feature for a specific concept. Thus, each feature was represented by a 541-element vector so that a Pearson correlation could be computed between each feature pair. For the present study, we attempted to avoid spurious correlations by

considering only pairs involving the 340 features that were listed for more than 3 concepts. The Pearson correlation between each feature pair was squared to obtain shared variance between features. The significance threshold was somewhat arbitrarily set to $r^2 = .05$ ($p < .0001$). A total of 5,949 feature pairs met this criterion, representing 10.4% of the 57,360 possible correlated pairs.

We calculated feature correlations across all 541 concepts, which differed from other research in which correlations were calculated separately for each set of concepts within a superordinate category. We believe that this is the most psychologically valid way of computing them. The major difficulties in computing correlations within superordinate categories concern determining category membership. According to the taxonomic features provided by participants in the norms, many basic-level exemplars were part of multiple superordinates, such as *knife* (utensil and weapon) or *eagle* (bird, animal, predator, carnivore). Others were not part of any clear category (e.g., *napkin*, *garage*, *mirror*). Therefore, it was not clear how to restrict the computation of statistical co-occurrence within categories, what it would mean if the same features were correlated to different degrees in overlapping categories, and why people would not learn such distributional statistics across concepts that do not fit into any superordinate category.

The stimuli for all three experiments consisted of 65 feature pairs that were pseudorandomly selected from the set of 5,949 correlated pairs (see Appendix A). Pairs in the upper range of shared variance were somewhat overrepresented because we expected that the number of theories provided would be positively correlated with shared variance (see Figure 1). The distribution of correlational strengths in the experimental pairs otherwise approximated the distribution within the entire set. Note that we purposely used a somewhat low bound on degree of shared variance between features (a minimum of 5%) because the major analyses in the present experiments were correlational, and it was therefore important to have a suitable range.

Five additional lead-in pairs that spanned the range of shared variance were selected as practice items. For example, *is ugly–lives in deserts* was selected as a weakly correlated pair ($r^2 = .05$), and *eats grass–eaten as meat* was selected as a strongly correlated pair ($r^2 = .51$).

Feature pairs were ordered randomly to create a master list. Because 23 features appeared in two pairs, we ensured that no participant was presented with the same feature in the same ordinal position twice in order to avoid cuing directional effects. For example, *has a tail* was paired with both *is furry* and *has hooves*, and was thus listed as the first feature in one pair and as the second feature in the other. There were four lists. In two of the lists, the features within a pair were presented in one order (*has wings–flies*) and in the other order in the remaining two lists (*flies–has wings*). In addition, the 65 items were presented in one order in the forward lists and in the reverse order in the backward lists in order to alleviate effects of practice or fatigue.

Figure 1. Distribution of correlated pairs at each level of shared variance.

The five lead-in pairs appeared in the same order at the beginning of all lists. Five interviewees were randomly assigned to each list.

Procedure. Interviews were recorded in a quiet room. Interviewees were read a set of instructions to inform them of the distinction between correlation and causation and to provide an example of how theories might be stated (Appendix B). For each feature pair, the interviewee was asked a series of questions in a predetermined order. A flowchart of the interview questions appears in Appendix C. Explanations were repeated back verbatim to participants, allowing them to revise or elaborate on their responses, which promoted theory production. No feedback was given regarding theory content, even if directly requested. Interviewees sometimes were encouraged to provide theories, particularly when they were taking some time thinking about it. Interviews took about 50 min.

The causal theories and answers to the battery of questions were transcribed from their original digital source and were then coded independently into symbolic logic notation by two judges, one being the first author and the other being a research assistant who was trained to derive symbolic logic notation from the interviews. The symbolic logic notation used the well-known set of logical operators, representing AND, OR, and IMPLIES (or CAUSES), which took propositional tokens as operands. A propositional token is a phrase (or proposition) for which a truth value can be assigned. Take, as an example, one of the causal theories between the features *worn on the feet* and *comes in pairs*, “Most people have two feet. So if you get something to wear on the feet, you are going to need a pair of them.” This theory contains three propositional tokens: (A) $\exists x|x$ has two feet, (B) $\exists y|x$ wears y on the foot, and (C) y comes in pairs. These tokens were connected with logical operators to arrive at the symbolic logic representation for the theory $(A \cdot B) \rightarrow C$. Translated back to English, this notation states that the conjunction of the facts that people have two feet and that some item is worn on the foot (of a person) causes that item to come in pairs. The final symbolic logic representations used in the analyses were arrived at by consensus between the two judges. Note that—as in this example—interviewees were not required to state precisely the terms *causes* or *implies* for an explanation to be counted as a causal theory.

Results and Discussion

Theories that included causal chains from either A to B or B to A were classified as causal relations. Invalid theories were those wherein a spurious third variable, C, causes both A and B, and thus no causal relations were strictly described between A and B. For example, 1 participant gave the following theory for *grows on trees—is juicy*: “Having something be juicy may cause it to be grown on trees just because lots of fruit is grown on trees, and fruit is generally juicy.” This participant initially stated that being juicy causes something to grow on trees, but her elaboration indicates that the property of being a fruit causes something to grow on trees and to be juicy. In other words, a third proposition—(C) x is a fruit—is the cause of both (A) x grows on trees and (B) x is juicy. Thus, although this may be legitimately held knowledge, it cannot be construed as a causal relationship between A and B. Rather, it exemplifies the interviewee’s awareness that both features frequently co-occur in fruit.

In addition to providing statistics concerning fully articulated causal theories, Experiment 1 makes the important distinction between a priori and ad hoc theories. A priori theories are those for which participants indicated that they had thought about the causal relationship prior to the interview. The intention is to measure the pervasiveness of causal theories that may have influenced prior learning. In

contrast, ad hoc theories are those that causally relate two features, but interviewees indicated that they had never before considered this relationship. Invalid theories are neither a priori nor ad hoc, because they do not causally relate two features. We believe that using self-report to distinguish between a priori and ad hoc theories is justifiable for three reasons. First, no other method is described in the literature, and self-report was most suited to the interview format. Second, if theories are explicitly available and used, participants should reasonably have awareness of their use. Third, it seems equally reasonable to accept self-report of theory use as to accept self-report of theory content, as has been done in other studies.

Theory pervasiveness. Two measures of theory pervasiveness were used. In addition to the a priori measure, the *liberal* measure was defined as the total number of attempted theory productions, including a priori theories, ad hoc theories, and invalid theories. The intention is to measure people’s willingness to present causal explanations without regard to content. Interviewees generally agreed on what pairs had a priori and ad hoc theories; interrater reliability across the 20 participants—estimated using the Spearman–Brown correction formula—was .81 for a priori theories and .81 for the liberal theory count. The liberal theory count per feature pair ($M = 7.4$ of a possible 20, $SE = 0.8$) was significantly greater than the number of a priori theories ($M = 2.2$, $SE = 0.4$) [$t(64) = 10.38$, $p < .001$]. In addition, 64% of all theories were ad hoc. That is, participants provided numerous valid theories during the experiment that did not reflect prior-held causal knowledge. Therefore, if all participants’ attempts at explaining a relationship were counted as a causal theory, actual theory pervasiveness would be greatly overestimated.

Figure 2 shows a clear discrepancy between the liberal and a priori estimates of theory pervasiveness. For 38 of 65 feature pairs, either 0 or 1 participant provided an a priori theory, and the frequency of feature pairs as a function of the number of participants providing a priori theories exponentially decreases. For no pair did more than 11 of 20 participants provide an a priori theory. In contrast, for the liberal theory count, the distribution of pairs is approximately equal across intervals.

These findings seem particularly important for researchers who employ empirically derived theories. If one assumes that an effect does not precede its cause, then new theories generated during an experiment cannot have influenced prior learning. The graduate and upper-level undergraduate participants had little difficulty providing ad hoc theories on demand. Thus, studies of causal relations should distinguish between a priori and ad hoc theories. Although we believe that the a priori theory count is more appropriate than the liberal theory count, both variables were used as predictors in Experiments 2 and 3 to address any concerns regarding the subjective identification and relatively small proportion of a priori theories. Note that although the absolute numbers differ substantially for the two measures, they are highly correlated [$r(63) = .86$, $p < .001$].

Participants were more likely to provide theories for feature pairs that were more strongly correlated. Shared

Figure 2. Distribution of feature pairs at each causal theory count.

variance between feature pairs correlated with both the number of a priori theories [$r(63) = .28, p < .05$] and the liberal theory count [$r(63) = .34, p < .01$].

Noticing statistical co-occurrences without theories. Some have argued that explanatory theories are required for learning feature correlations, so that only causally related features are perceived as being statistically related (Ahn et al., 2002; Murphy & Medin, 1985). If this is so, then the expected proportion of co-occurrences noticed without a theory is zero. A single-sample t test showed that the mean proportion of correlations noticed in the absence of an a priori or an ad hoc theory ($M = .49, SE = .03$) was significantly greater than zero [$t(64) = 17.54, p < .001$]. That is, in almost 50% of the cases, statistical correlations were noticed even though participants did not produce a causal theory.

Without employing causal theories, participants may have learned many statistical co-occurrences through simple observation. Shared variance between features correlated with the number of participants who reported noticing that the features co-occur [$r(63) = .42, p < .001$]. This result is consistent with the statistical learning explanation because highly correlated feature pairs have a higher probability of being attended together, thus providing increased opportunity to learn the correlation through observation.

Noticing co-occurrences with theories. Although the previous analyses suggest that causal theories are not required to learn statistical co-occurrences, theories presumably assist in learning them (and vice versa). Murphy and Medin (1985), for example, suggested that theories may help categorize objects in situations where justification or explanation is required. Accordingly, the probability of noticing a feature co-occurrence correlated with both the number of a priori theories [$r(63) = .55, p < .001$] and the liberal theory count [$r(63) = .61, p < .001$].

Theory complexity. It seems reasonable that people are less likely to possess and/or generate theories that describe complex relationships. Although it might be obvious that a pair of features co-occur, the causal mechanism behind the co-occurrence may be obscure. Theory complexity is defined as the number of logical operators in the logical notation. The simplest causal relations have a complexity

of 1 because the proposition *A causes B* can be expressed using only the IMPLIES operator. More complex theories involve longer causal chains (e.g., *A causes C, which causes B*) or contingencies (e.g., *A and C together cause B*), using additional operators. Mean complexity was computed for each feature pair and included all ad hoc and a priori theories. Complexity negatively correlated with the number of a priori theories [$r(55) = -.42, p < .001$] and the liberal theory count [$r(55) = -.37, p < .01$]. Thus, more complex explanatory theories are less commonly held, or are at least less commonly generated. Keil (2003) suggested that a number of converging factors limit the depth of one's explanatory knowledge in many domains while simultaneously masking the inadequacy of one's explanatory knowledge. He attributed this illusion of explanatory depth to factors that promote the sort of shallow explanatory theories provided by participants in Experiment 1.

Summary. Experiment 1 provides information regarding the content and pervasiveness of the types of causal theories that people possess. Participants provided causal explanations—even for feature co-occurrences they had never before noticed—while at the same time noticing many feature co-occurrences without also possessing a theory. These analyses suggest that a person's ability to provide a causal explanation on demand is not a reliable indicator of the set of causal theories of which he or she makes use and that the probability of noticing feature co-occurrences depends on both distributional statistics and causal knowledge.

EXPERIMENT 2

To test for influences of causal relations and statistical co-occurrences, we used an untimed feature relatedness rating task ("On a scale of 1 to 9, to what extent do these two features tend to occur together in living and/or nonliving things in the world?"). Untimed tasks have been used to test for influences of causal theories in part because it has been assumed that only slower tasks make use of them (Sloman et al., 1998). Therefore, feature relatedness rating should be sensitive to causal relations. In contrast, untimed tasks generally lack sensitivity to statistical co-occurrence knowledge (Ahn et al., 2002; McRae et al.,

Figure 3. Mean relatedness ratings as a function of percentage of shared variance.

1997). However, because feature relatedness ratings directly tap knowledge of feature–feature correspondences, they may be sensitive to statistical co-occurrences as well. Correlational analyses were used to test the influence of the two knowledge types on relatedness ratings.

Method

Participants. Forty-nine University of Western Ontario undergraduates participated for either course credit or monetary reimbursement. Three participants were excluded because they left more than two responses blank, and 6 participants were excluded because they used only a small range of responses, indicating a lack of effort.

Materials. The stimuli consisted of the 65 correlated feature pairs and five lead-in items that were used in Experiment 1. The feature pairs were listed one per line along with a 9-point rating scale, as shown in this example:

		Not at all related			Moderately related			Extremely highly related		
		1	2	3	4	5	6	7	8	9
flies	has a beak									

As occurred in Experiment 1, crossing the pairwise (AB and BA) and listwise (forward and reverse) ordering created four lists.

Procedure. Participants circled a number on the scale beside each feature pair to indicate the extent to which the two features tend to occur together in the same living and/or nonliving things in the world. Three sample ratings were provided as examples of high, low, and intermediate ratings, and participants were instructed to try to use the entire range of the scale.

Results and Discussion

The Spearman–Brown correction formula estimated the mean interrater reliability to be .96. Mean relatedness ratings across the 65 feature pairs varied from 2.5 (*has pockets—is blue*) to 6.2 (*has fins—has gills*), with an overall mean of 4.5 and a standard error of 0.1. Figure 3 shows that shared variance between features correlates with the mean relatedness rating [$r(63) = .42, p < .001$]. Thus, although Ahn et al. (2002) and Malt and Smith (1984) failed to find an effect of statistical correlations in typicality ratings of artificial exemplars, the present results suggest that a sensitive untimed task is influenced by people's knowledge of statistical co-occurrences.

Figure 4 shows that the number of a priori theories correlates with mean relatedness rating [$r(63) = .48, p < .001$]. Furthermore, Figure 5 shows that the liberal theory count does as well [$r(63) = .49, p < .001$]. Thus, relatedness judgments are higher for feature pairs for which theories are more commonly held and generated. Consistent with Ahn et al. (2002), this is further evidence that causal relations play a role in explicit judgments of feature relatedness.

Causal relations and statistical co-occurrences influence an untimed task designed to directly probe knowledge of feature relatedness. The causal relation results cohere with previous findings regarding untimed conceptual tasks (Ahn et al., 2002; Malt & Smith, 1984). In terms of statistical co-occurrences, these results provide further insight into the conditions under which people's knowledge of feature correlations influences conceptual tasks. They are in accordance with Chin-Parker and Ross (2002), who showed that people are sensitive to within-category feature correlations in novel concepts when tested in a feature prediction task or a double feature classification task, but not in a typicality rating task. They also provide further evidence that the failure to find an influence of statistical co-occurrences in untimed categorization tasks (Murphy & Wisniewski, 1989), typicality ratings (Ahn et al., 2002; Malt & Smith, 1984), and similarity judgments and feature typicality ratings (McRae et al., 1997) is due to knowledge–task congruency rather than a lack of feature correlation knowledge.

EXPERIMENT 3

Speeded binary feature relatedness judgments require direct judgments of how features are related, and should therefore be sensitive to feature co-occurrence knowledge. We hypothesized that consistent with McRae et al. (1999; McRae et al., 1997), relatedness decision latencies should be predicted by knowledge of statistical co-occurrences. Because of the equivocal nature of the results of Lin and Murphy (1997) and Palmeri and Blalock (2000), it is unclear whether causal relations can be accessed sufficiently

Figure 4. Mean relatedness ratings as a function of number of a priori theories.

Figure 5. Mean relatedness ratings as a function of the liberal theory count.

rapidly to influence speeded relatedness decisions. In Experiment 3 we examined the influence of these knowledge types on speeded relatedness decisions at increasing SOAs, which allowed us to investigate whether additional time to process the features influences the availability or use of each knowledge type. We hypothesized that increasing the SOA between features might increase the probability that causal relations would influence decision latency. To some extent, this hypothesis was made on the basis of the work of Sloman et al. (1998), who stated that “We propose that early access to concepts ignores the content of relations; only slower processing makes use of it” (p. 204). It was unclear whether the influence of statistical co-occurrences would change over SOA.

Method

Participants. A total of 188 University of Western Ontario undergraduates participated for course credit. We randomly assigned 20, 18, 19, 18, and 19 participants to each list for the 400-, 800-, 1,200-, 1,600-, and 2,000-msec SOAs, respectively.

Materials. The 65 feature pairs from Experiments 1 and 2 were used. An additional 65 unrelated pairs (e.g., *covered with needles* and *spouts water*) were included so that 50% of the pairs required a “yes” response and 50% required a “no” response. No feature was used in both related and unrelated feature pairs, although the 23 features appearing in 2 related feature pairs were listed such that they did not appear in the same pairwise position twice, as was the case in Experiments 1 and 2. There were two lists containing all pairs, each of which presented items in a different pairwise order.

To avoid cuing relatedness, features appearing in related and unrelated pairs were matched along a number of dimensions, including the number of characters per feature name (related, $M = 12.6$, $SE = 0.4$; unrelated, $M = 13.1$, $SE = 0.3$) [$t_2(63) = 0.91$, $p > .3$], number of words per feature name (related, $M = 2.6$, $SE = 0.1$; unrelated, $M = 2.6$, $SE = 0.1$) [$t_2(63) = 0.09$, $p > .9$], and the number of concepts associated with each feature (related, $M = 14.8$, $SE = 1.2$; unrelated, $M = 11.0$, $SE = 1.8$) [$t_2(63) = 1.68$, $p > .05$]. Features were also matched with respect to their distribution across nine feature types (Cree & McRae, 2003) because feature type typically is cued linguistically; for example, functional features often take the form *used for . . .*, whereas visual form and surface features often take the form *has a . . .* or *made of . . .*. The distribution of features across nine feature types did not differ for related versus unrelated pairs [$\chi^2(8) = 0.88$, $p > .90$]. There were 30 practice pairs (15 related and 15 unrelated). No feature was used in both the practice set and the experimental trial set.

Procedure. Participants were tested individually using PsyScope (Cohen, MacWhinney, Flatt, & Provost, 1993) on a Macintosh G3 computer equipped with a 17-in. color Sony Trinitron monitor. Decision latencies were recorded with millisecond accuracy using a CMU button box that measured the time between the onset of the second feature and the button press. Participants used the index finger of their dominant hand for a “yes” response and the index finger of their nondominant hand for a “no” response.

Each trial began with a fixation point (+) in the center of the screen for 500 msec. For each of the five SOA conditions, the fixation point was replaced with the first feature for 400, 800, 1,200, 1,600, or 2,000 msec; then, the second feature appeared one line below it. Both features remained on the screen until the participant responded. Trials were presented in random order. Participants were instructed to read the feature names silently and to indicate as quickly and accurately as possible whether the two features go together in living and/or nonliving things. Participants completed 30 practice trials and then 130 experimental trials. Each session took less than 30 min to complete.

Design. ANOVAs were conducted using participants (F_1) and items (F_2) as random variables. The dependent variables were decision latency and the square root of the number of errors (Myers, 1979). The independent variables were relatedness (related vs. unrelated) and SOA (400, 800, 1,200, 1,600, or 2,000 msec). Relatedness was within participants but between items, whereas SOA was between participants but within items. Planned comparisons were conducted on decision latencies to test for the effects of relatedness at each SOA. Regression analyses were conducted at each SOA to test the influences of statistical co-occurrences and causal theories on decision latencies.

Results and Discussion

Decision latency analyses excluded incorrect responses. Decision latencies greater than three standard deviations above the grand mean were replaced by the cutoff value (less than 2% of the scores). Nine participants (2 from the 800-msec SOA, 3 from 1,600 msec, and 4 from 2,000 msec) had mean decision latencies three standard deviations above the grand mean and were excluded. Eight participants (2 from the 400-msec SOA, 2 from 800 msec, 3 from 1,600 msec, and 1 from 2,000 msec) had error rates exceeding three standard deviations above the grand mean for each condition and were excluded.

Table 1 presents the mean decision latencies and proportions of errors at each SOA. Overall, related feature pairs ($M = 1,031$ msec, $SE = 18$ msec) were responded to more quickly than unrelated pairs ($M = 1,315$ msec, $SE = 31$ msec) [$F_1(1,180) = 275.94$, $p < .001$; $F_2(1,128) = 82.04$, $p < .001$]. Relatedness interacted with SOA because the relatedness effect ranged from 391 msec at the 400-msec SOA to 225 msec at the 1,200-msec SOA [$F_1(4,180) = 3.35$, $p < .05$; $F_2(4,512) = 18.98$, $p < .001$]. The relatedness effect did not vary monotonically over SOA, since it increased to 304 msec at an SOA of 1,600 msec. Planned comparisons showed that related feature pairs were responded to more quickly than unrelated feature pairs at all five SOAs: 400 msec [$F_1(1,180) = 114.55$, $p < .001$; $F_2(1,182) = 129.33$, $p < .001$], 800 msec [$F_1(1,180) = 37.45$, $p < .001$; $F_2(1,182) = 50.52$, $p < .001$], 1,200 msec [$F_1(1,180) = 36.03$, $p < .001$; $F_2(1,182) = 41.50$, $p < .001$], 1,600 msec [$F_1(1,180) = 62.32$, $p < .001$; $F_2(1,182) = 81.78$, $p < .001$], and 2,000 msec [$F_1(1,180) = 42.82$, $p < .001$; $F_2(1,182) = 54.91$, $p < .001$]. Finally, mean decision

Table 1
Mean Decision Latencies (in Milliseconds) and Error Rates for
Experiment 3 Related and Unrelated Pairs

SOA	Decision Latency					Proportion of Errors				
	Unrelated		Related		Difference	Unrelated		Related		Difference
	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>		<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	
400	1,498	84	1,107	52	391	.08	.01	.06	.01	.02
800	1,212	64	973	40	239	.05	.01	.07	.01	-.02
1,200	1,214	62	989	36	225	.05	.01	.05	.01	0
1,600	1,346	73	1,042	36	304	.07	.01	.07	.01	0
2,000	1,291	42	1,039	27	252	.05	.01	.06	.01	-.01

latencies varied across SOAs (400 msec, $M = 1,302$ msec, $SE = 66$ msec; 800 msec, $M = 1,092$ msec, $SE = 51$ msec; 1,200 msec, $M = 1,101$ msec, $SE = 48$ msec; 1,600 msec, $M = 1,194$ msec, $SE = 53$ msec; 2,000 msec, $M = 1,165$ msec, $SE = 32$ msec) [$F_1(4,180) = 2.78$, $p < .05$; $F_2(4,512) = 122.77$, $p < .001$].

There were no significant effects for error rates, mainly because—importantly—participants made few errors, being able to reliably distinguish between related versus unrelated feature pairs [for relatedness, both $F_s < 1$; for SOA, $F_1(4,180) = 1.70$, $p > .2$; $F_2(4,512) = 5.72$, $p < .001$; for the interaction, $F_1(4,180) = 1.13$, $p > .3$; $F_2(4,512) = 1.26$, $p > .2$].

In summary, decision latencies were shorter for related than for unrelated feature pairs. The relatedness effect should be interpreted with caution, however, because it involves comparing what participants might have considered a “yes” (related) versus a “no” (unrelated) judgment. Finally, although feature pairs that were correlated as low as at 5% shared variance were included in the stimuli, participants had little trouble discriminating between related and unrelated pairs.

Speeded effects of relations and correlations. Of primary interest were the influences of causal relations and statistical co-occurrences on decision latencies. Therefore, regression analyses were performed at each SOA to predict mean decision latencies using shared variance between features, the number of a priori theories, and the liberal theory count. Because the features must be read before responding, a number of reading time variables were first entered into the regression equation, including the number of content words in a feature name, the frequency of content words, and the length in characters of all words. Word frequency was calculated as the average of the natural logarithm of word frequency counts taken from the British National Corpus (Burnard, 2000). Feature length in characters included spaces between words. Experimental list was also included because different participants were presented with the two lists, and decision latencies differed by list.

Table 2 shows the partial correlations after the reading time variables had been entered into each regression equation. All three variables predicted significant variance in decision latencies at all five SOAs. Although the partial correlations are not huge—ranging from $-.22$ to $-.40$, with the majority around $-.30$ —the fact that they are significant at all SOAs provides five replications for

each variable. In addition, the predictive ability of the three variables did not systematically change with SOA. We had hypothesized that the influence of causal relations might increase with additional time between features, but this was not the case.

The independent variables for a priori and liberal theories differed from those in Experiment 2. In that experiment, we used the total number of theories across the two feature orders because participants had time to look back and forth between features. In Experiment 3, because participants were presented with the first feature for a period ranging from 400 to 2,000 msec prior to the second, we used directionally specific theory counts. Using the total number of theories as predictors produced similar results (a priori, partial correlations of $-.33$, $-.40$, $-.30$, $-.31$, and $-.35$ across the five SOAs; liberal, partial correlations of $-.34$, $-.42$, $-.30$, $-.31$, and $-.31$; all $p_s < .001$).

From a statistical learning point of view, better predictions for statistical co-occurrences would be expected for pairs of features in which both are relatively directly observable. To test this idea, we conducted regression analyses using only pairs in which both features referred to parts, surface features, or functional features. This resulted in 74 items (37 pairs in each order). The partial correlations for statistical co-occurrences were noticeably stronger (partial correlations of $-.36$, $-.46$, $-.34$, $-.36$, and $-.45$ across the five SOAs, all $p_s < .01$). In contrast, the partial correlations for the number of a priori theories ($r = -.22$, $p > .07$; $r = -.32$, $p < .01$; $r = -.16$, $p > .1$; $r = -.22$, $p > .07$; $r = -.31$, $p < .01$, across the five SOAs) and the liberal theory count ($r = -.30$, $p < .05$; $r = -.35$, $p < .01$; $r = -.19$, $p > .1$; $r = -.23$, $p > .06$; $r = -.28$, $p < .05$) decreased somewhat. In addition, when the Experiment 2 relatedness ratings were analyzed using the same 37 feature pairs, feature

Table 2
Partial Correlations From Regressions
on Relatedness Decision Latencies ($N = 130$)

SOA	Predictor Variable		
	Statistical Co-Occurrences	A Priori Theories	Liberal Theory Count
400	-.23*	-.33***	-.35***
800	-.28**	-.40***	-.40***
1,200	-.22*	-.28**	-.27**
1,600	-.26**	-.30**	-.31**
2,000	-.30***	-.32***	-.27**

* $p < .05$. ** $p < .01$. *** $p < .001$.

correlations were a somewhat better predictor ($r = .54$, $p < .01$), with the a priori ($r = .39$, $p < .05$) and liberal ($r = .42$, $p < .01$) predictions decreasing slightly.

Summary. Both theory-based and statistically based knowledge predicted decision latencies across all five SOAs, suggesting that they play a role in speeded judgments of feature relatedness. The results are in concert with other influences of feature correlations on speeded tasks (McRae et al., 1999). Furthermore, the demonstration of an influence of theory-based knowledge supports an interpretation of the results of Lin and Murphy (1997) and Palmeri and Blalock (2000) that background knowledge can influence speeded tasks.

GENERAL DISCUSSION

This research made three main contributions. Experiment 1 advanced the investigation of theory-based knowledge by collecting descriptions of causal relations between features in a detailed manner. It elucidated aspects of the relationships among causal theories, feature correlations, noticing correlations with and without prior theories, and theory complexity. Experiment 2 used a relatedness rating task to provide the first evidence of an influence of feature–feature statistical co-occurrences across natural basic-level concepts in an untimed task. Experiment 3 introduced a speeded relatedness decision task that showed an influence of causal relations in a speeded task. Thus, congruency between knowledge type and task type depends not only on matching the temporal parameters of mental computations and tasks, but also on the directness with which a task taps the relevant knowledge type, a conclusion that echoes that of Chin-Parker and Ross (2002).

Causal Theories

As was detailed by Murphy (2002) and Rehder and Murphy (2003), causal theories and background knowledge have been found to influence a plethora of types of concept learning and conceptual processing tasks. The present experiments add further evidence. In addition, we differentiated between a priori and ad hoc theories in real-world concepts and used both measures of causal relations to account for judgments of feature relatedness. Because the two measures showed similar predictive strength, one might be tempted to conclude that there is no advantage to distinguishing between these variables. However, although the two measures were highly correlated, the prevalence of a priori theories differed markedly from that of making any attempt at a theory at all. Therefore, the distinction does have theoretical consequences in terms of the degree to which causal theories influence conceptual processing.

One important aspect of theory-based theories is that individual concepts cohere around causal theories (Sloman et al., 1998). Consequently, much of the literature on the role of causal relations in categorization calculates correlations among features within a category, focusing on causal theories that bind features within a concept (Ahn et al., 2002; Malt & Smith, 1984; Murphy & Wisniewski,

1989). Given the theoretical importance of causal relationships to category coherence, it is possible that we may have found a greater influence of causal theories in our experiments if reference categories had been provided to participants. Thus, the present methodology may have underestimated the role and prevalence of causal theories. In addition, these factors may have been even stronger if, for example, social concepts had been used rather than concrete object concepts.

Statistical Co-Occurrences

Although there seems to be little debate regarding the influence of causal relations, the same cannot be said of the influence of feature correlations. However, there presently exists a body of research that—taken as a whole—provides what we believe is a clear picture regarding the facts that people learn feature correlations and that such correlations influence conceptual computations. This research includes investigating real-world lexical concepts and concept learning paradigms using novel categories and exemplars. Four factors have been identified that determine whether feature correlations are learned: the type or goal of learning, the quantity of learning experiences provided, the overall structure in the stimuli, and the tasks used to test people's knowledge.

A number of studies with adults have shown a lack of sensitivity to feature correlations when intentional category learning is used (Chin-Parker & Ross, 2002; Murphy & Wisniewski, 1989). It now is clear that learning to categorize exemplars into two categories on the basis of feedback promotes learning cues that enable successful classification, but impinges on learning other types of environmental structure. In contrast, tasks such as incidental category or inference learning typically lead to learning feature correlations (Wattenmaker, 1991, 1993). Presumably, concept learning in the real world includes varied experiences and learning situations.

In terms of quantity of learning, Tangen and Allan (2004) showed that sensitivity to causal structure depends on the number of learning trials. Whereas small numbers of learning trials promote causal reasoning, large numbers of trials promote statistical learning. Most category learning experiments use relatively few learning experiences. In contrast, undergraduate students presumably have vast experience with the sorts of object concepts typically used in studies such as the present one.

Billman and Knutson (1996), on the basis of Billman and Heit's (1988) modeling, showed that people find it easier to learn correlations that are part of a coherent system of correlations, as is the case in natural lexical concepts but is not the case in most artificial concept-learning experiments. Given this fact, it is not surprising that experiments in which participants learn about novel stimuli that instantiate a sparse statistical structure often fail to demonstrate an influence of statistical co-occurrences.

These three factors combine to influence the probability that people will learn statistically based co-occurrences. In addition, as was already discussed, the task used to test for this knowledge is critical for demonstrating such effects. That people naturally learn and use feature correlations is

or is not surprising, depending on one's theoretical perspective. If one assumes that people must explicitly notice that two features co-occur and must construct a link between them in a prototype list or a semantic network, then this appears to be a computationally intractable problem, given the huge number of possible feature pairs in the world. In contrast, for those researchers who view the brain as being designed for correlational learning and thus emphasize statistical learning as a building block of knowledge acquisition, people cannot help but learn feature correlations. Research to this point has supported the latter view.

The Interplay Between Them

Actual co-occurrences between features do not depend on people's explanations. For example, things with wings fly independently of whether people have an explanation for this relationship. Thus, feature correlations obtained from feature production norms presumably reflect real-world statistical structure, although their magnitudes may also be influenced in part by causal knowledge. The features used in the present experiments were drawn from McRae et al.'s (2005) semantic feature norms; therefore, the statistical correlations reflect tendencies for people to list these features together for specific concepts. Murphy and Medin (1985) stated—we believe correctly—that causal relationships might influence what features are listed in such norms and in this way might affect calculations of statistical correlations between features.

The mechanism by which causal relations can influence statistical co-occurrence knowledge is the same as that behind all statistically based learning. Some have proposed that explanatory theories focus attention on certain feature co-occurrences (Murphy & Medin, 1985; Murphy & Wisniewski, 1989). Consequently, these co-occurrences are effectively oversampled, causing them to become particularly salient. For example, to explain why whales are mammals rather than fish, a student may be presented with comparisons between whales and other mammals that highlight various common mammalian properties, causing the student to give an unusual amount of consideration to these salient relationships. The more frequently these causal relations are considered, the more frequently the underlying featural representations are simultaneously active, thereby strengthening the connections between relevant features.

However, there is likely a reciprocal influence of statistical co-occurrences on causal relations. People do not possess all possible causal theories; rather, causal theories explain noticed co-occurrences. For example, in the case of accidental discoveries such as that of penicillin, noticed co-occurrences (e.g., petri dishes that contain penicillin mold are bacteria free) are used to generate causal hypotheses, which are subsequently refined. Therefore, just as theories can influence statistical learning, real-world co-occurrence statistics constrain the theories that people generate.

This research helps to clarify issues concerning the roles of causal relations and statistical co-occurrences in the representation, computation, and use of concepts. People are capable of learning both statistically based and causally based relations, and when adults are tested using

sensitive tasks, it is apparent that both are important aspects of conceptual knowledge.

AUTHOR NOTE

This work was supported by a Natural Sciences and Engineering Research Council Doctoral Scholarship and a SHARCNET Graduate Research Fellowship to C.M., and by Natural Sciences and Engineering Research Council Grant OGP0155704 and National Institutes of Health Grants R01-DC0418 and R01-MH6051701 to K.M. Part of the manuscript formed C.M.'s University of Western Ontario master's thesis. Correspondence concerning this article should be addressed to C. McNorgan or K. McRae, Department of Psychology, University of Western Ontario, London, ON, N6A 5C2 Canada (e-mail: cmcnorg4@uwo.ca or mcrae@uwo.ca).

REFERENCES

- AHN, W.-K., KIM, N. S., LASSALINE, M. E., & DENNIS, M. J. (2000). Causal status as a determinant of feature centrality. *Cognitive Psychology*, **41**, 361-416.
- AHN, W.-K., MARSH, J. K., LUHMANN, C. C., & LEE, K. (2002). Effect of theory-based feature correlations on typicality judgments. *Memory & Cognition*, **30**, 107-118.
- BILLMAN, D., & HEIT, E. (1988). Observational learning from internal feedback: A simulation of an adaptive learning method. *Cognitive Science*, **12**, 587-625.
- BILLMAN, D., & KNUTSON, J. (1996). Unsupervised concept learning and value systematicity: A complex whole aids learning the parts. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **22**, 458-475.
- BURNARD, L. (2000). British National Corpus user reference guide version 2.0. Retrieved August 25, 2004, from www.hcu.ox.ac.uk/BNC/World/HTML/urg.html.
- CHIN-PARKER, S., & ROSS, B. H. (2002). The effect of category learning on sensitivity to within-category correlations. *Memory & Cognition*, **30**, 353-362.
- COHEN, J. D., MACWHINNEY, B., FLATT, M., & PROVOST, J. (1993). PsyScope: An interactive graphic system for designing and controlling experiments in the psychology laboratory using Macintosh computers. *Behavior Research Methods, Instruments, & Computers*, **25**, 257-271.
- CREE, G. S., & McRAE, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of *chipmunk*, *cherry*, *chisel*, *cheese*, and *cello* (and many other such concrete nouns). *Journal of Experimental Psychology: General*, **132**, 163-201.
- CREE, G. S., McRAE, K., & McNORGAN, C. (1999). An attractor model of lexical conceptual processing: Simulating semantic priming. *Cognitive Science*, **23**, 371-414.
- GELMAN, S. A. (2003). *The essential child: Origins of essentialism in everyday thought*. Oxford: Oxford University Press.
- HOPFIELD, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, **79**, 2554-2558.
- HOPFIELD, J. J. (1984). Neurons with graded response have collective computational features like those of two-state neurons. *Proceedings of the National Academy of Sciences*, **81**, 3088-3092.
- JONES, S. S., & SMITH, L. B. (1993). The place of perception in children's concepts. *Cognitive Development*, **8**, 113-139.
- KEIL, F. C. (2003). Categorisation, causation, and the limits of understanding. *Language & Cognitive Processes*, **18**, 663-692.
- LIN, E. L., & MURPHY, G. L. (1997). Effects of background knowledge on object categorization and part detection. *Journal of Experimental Psychology: Human Perception & Performance*, **23**, 1153-1169.
- MALT, B. C., & SMITH, E. E. (1984). Correlated properties in natural categories. *Journal of Verbal Learning & Verbal Behavior*, **23**, 250-269.
- McRAE, K. (2004). Semantic memory: Some insights from feature-based connectionist attractor networks. In B. H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 45, pp. 41-86). San Diego: Academic Press.

- MCRAB, K., CREE, G. S., SEIDENBERG, M. S., & McNORGAN, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods*, **37**, 547-559.
- MCRAB, K., CREE, G. S., WESTMACOTT, R., & DE SA, V. R. (1999). Further evidence for feature correlations in semantic memory. *Canadian Journal of Experimental Psychology*, **53**, 360-373.
- MCRAB, K., DE SA, V. R., & SEIDENBERG, M. S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, **126**, 99-130.
- 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.
- MURPHY, G. L., & WISNIEWSKI, E. J. (1989). Feature correlations in conceptual representations. In G. Tiberghien (Ed.), *Advances in cognitive science: Vol. 2. Theory and applications* (pp. 23-45). Chichester, U.K.: Ellis Horwood.
- MYERS, J. L. (1979). *Fundamentals of experimental design*. Boston: Allyn & Bacon.
- PALMERI, T. J., & BLALOCK, C. (2000). The role of background knowledge in speeded perceptual categorization. *Cognition*, **77**, B45-B57.
- REHDER, B., & HASTIE, R. (2004). Category coherence and category-based property induction. *Cognition*, **91**, 113-153.
- REHDER, B., & MURPHY, G. L. (2003). A knowledge-resonance (KRES) model of category learning. *Psychonomic Bulletin & Review*, **10**, 759-784.
- ROSCH, E., & MERVIS, C. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, **7**, 573-605.
- SLOMAN, S. A., LOVE, B. C., & AHN, W. (1998). Feature centrality and conceptual coherence. *Cognitive Science*, **22**, 189-228.
- TANGEN, J. M., & ALLAN, L. G. (2004). Cue interaction and judgments of causality: Contributions of causal and associative processes. *Memory & Cognition*, **32**, 107-124.
- WATTENMAKER, W. D. (1991). Learning modes, feature correlations, and memory-based categorization. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **17**, 908-923.
- WATTENMAKER, W. D. (1993). Incidental concept learning, feature frequency, and correlated properties. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **19**, 203-222.

APPENDIX A
Correlated Feature Pairs Used in Experiments 1, 2, and 3

Feature A	Feature B	% Shared Variance	Feature A	Feature B	% Shared Variance
has fins	has gills	81.3	has pockets	worn by men	14.5
comes in pairs	worn on feet	74.7	is crunchy	is nutritious	14.1
made of wool	is knitted	69.5	is breakable	has a lid	13.2
swims	lives in water	67.2	builds nests	has feathers	12.6
worn for warmth	made of wool	65.2	is yellow	has a peel	12.5
tastes sour	is citrus	62.1	has fur	lives in the wilderness	11.7
lives in zoos	lives in Africa	61.1	has long legs	is pink	11.5
is sharp	used for cutting	58.5	made of porcelain	is breakable	11.2
has a shell	is slow	52.8	has a tail	has hooves	11.0
flies	has a beak	50.9	used by children	has wheels	10.1
eaten in summer	has a pit	47.2	is furry	has a tail	8.6
made of cotton	has sleeves	43.9	has strings	used in orchestras	8.0
has scales	has fins	41.6	used by blowing air through	is loud	7.8
has a lid	used for holding things	36.7	is rectangular	is flat	7.3
found in bathrooms	used for holding water	35.9	used for war	is dangerous	7.2
grows on trees	is juicy	34.1	has buttons	worn for covering	7.2
used in bands	made of brass	33.9	is green	is edible	7.1
made of brick	has a roof	33.4	has wheels	used for moving things	6.7
floats	used on water	32.7	has feathers	has a neck	6.3
eaten by cooking	grows in the ground	32.5	is dangerous	lives in jungles	6.2
lives in the jungle	lives in zoos	30.0	is red	grows in gardens	6.1
has an engine	used for transportation	27.8	travels in herds	has horns	6.1
eaten in salads	is crunchy	27.2	lives in the wilderness	used in circuses	5.9
used in orchestras	used by blowing air through	26.2	is colorful	is pretty	5.9
has seeds	is juicy	24.5	has a trigger	is loud	5.8
is fun	used for playing	23.3	lays eggs	has webbed feet	5.6
has feet	migrates	21.5	has whiskers	has fur	5.5
used for hunting	has a trigger	21.1	is blue	has pockets	5.5
is comfortable	used by sitting on it	20.8	has legs	lives on farms	5.4
worn by women	made of silk	18.5	has a pit	tastes good	5.3
used for passengers	has an engine	17.8	is brown	has 4 legs	5.1
has scales	swims	17.1	is thin	has a pointed end	5.1
flies	sings	16.9			

APPENDIX B
Sample Theories Given in Experiment 1

Straightforward Theory

is hard—made of metal: Most metals are hard, and so if something is made of metal, then this causes the object to also be hard.

Roundabout Theory

has a wooden handle—used for construction: If you are using something for construction, it's probably going to be some kind of tool, and you might have to hold on to it, which will require it to have a handle. The handle will have to be durable, but lightweight, so a wooden handle would probably be a good type of handle to have. Thus, *used for construction* causes the object to have a wooden handle.

No Causal Theory

has wings—chirps: Obviously these two features are related because they are both properties of many birds. However, having wings doesn't cause an object to chirp, nor does the ability to chirp cause something to have wings, and so these features are only accidentally related.

APPENDIX C
Interview Flowchart for Experiment 1

(Manuscript received May 13, 2005;
 revision accepted for publication December 22, 2005.)