ELSEVIER

Cognitive Science 28 (2004) 45–74

COGNITIVE SCIENCE

http://www.elsevier.com/locate/cogsci

Feature centrality and property induction

Constantinos Hadjichristidis[a,*], Steven Sloman[b],
Rosemary Stevenson[c], David Over[d]

[a]*Department of Psychology, University of Plymouth, UK*
[b]*Department of Cognitive & Linguistic Sciences, Brown University, USA*
[c]*University of Durham, Durham, UK*
[d]*Department of Philosophy, University of Sunderland, Sunderland, UK*

Received 11 June 2001; received in revised form 27 August 2003; accepted 1 September 2003

Abstract

A feature is central to a concept to the extent that other features depend on it. Four studies tested the hypothesis that people will project a feature from a base concept to a target concept to the extent that they believe the feature is central to the two concepts. This centrality hypothesis implies that feature projection is guided by a principle that aims to maximize the structural commonality between base and target concepts. Participants were told that a category has two or three novel features. One feature was the most central in that more properties depended on it. The extent to which the target shared the feature's dependencies was manipulated by varying the similarity of category pairs. Participants' ratings of the likelihood that each feature would hold in the target category support the centrality hypothesis with both natural kind and artifact categories and with both well-specified and vague dependency structures.

Keywords: Induction; Categorical inference; Conceptual centrality; Projectibility

1. Introduction

Suppose that you read in a newspaper that a new species of swan has been discovered deep in the Canadian wilderness. The news story has little information in it, and you wonder about the probable features of this new species. You try to predict features of the new species from what you know about the features of familiar species of swan. In such prediction, you are said, technically, to project features from the familiar *base* species to the new *target* species.

* Corresponding author. Tel.: +44-1752-233147; fax: +44-1752-233362.
E-mail address: dinos@plymouth.ac.uk (C. Hadjichristidis).

0364-0213/$ – see front matter © 2003 Cognitive Science Society, Inc. All rights reserved.
doi:10.1016/j.cogsci.2003.09.001

You might hesitate to project that the new swans are white in color, but you will confidently project that they have hearts. Why? Why does our willingness to project features or predicates depend on the feature being projected? This question is about inductive reasoning, our capacity to generate new knowledge in the face of uncertainty, and has puzzled philosophers and psychologists alike.

1.1. Determinants of property projection

Goodman (1955/1983) argued that the *projectibility* of a predicate—people's willingness to believe it true of a general class from observations of particular cases—is determined by the extent of its so-called *entrenchment*. Familiar predicates (e.g., has a heart) gain entrenchment by appearing in successfully projected hypotheses (e.g., All swans have a heart; All bears have a heart).[1] Unfamiliar predicates (e.g., likes to eat alfalfa) may inherit entrenchment from their parent predicates (e.g., has a characteristic diet).

More recently, some psychologists have tried to explain projectibility, at least partially, in terms of similarity.[2] If projectibility is determined by the similarity between a base and target category, only a flexible notion of similarity will suffice (Heit & Rubinstein, 1994; Lassaline, 1996; Sloman, 1994; Smith, Shafir, & Osherson, 1993). Candidate predicates pick out certain properties of the base. A projection is strong to the extent that the target is believed to share these properties. Heit and Rubinstein (1994), for instance, found that for the anatomical predicate has a liver with two chambers inferences were stronger from chickens to hawks than from tigers to hawks, whereas for the behavioral predicate prefers to feed at night, the order of the preference was reversed. They argued for the existence of two types of similarity: anatomical and behavioral. An inference was strong to the extent that the type of candidate property matched the relevant type of similarity between the premise and conclusion categories.

Does either entrenchment or flexible similarity suffice to explain the projectibility of a predicate? We are not optimistic. Goodman's account attempts to explain induction in terms of linguistic practices: A hypothesis is projectible depending on the frequency and success of prior projections. However, infants and animals without language are capable of successful induction. Contrary to Goodman's suggestion, it seems that inductive practices drive linguistic practices and not the other way around (see Sloman & Lagnado, in press). Furthermore, frequency of successful projection seems incapable of capturing certain cases of comparative projectibility. Hypotheses regarding central properties seem to be projected less frequently than hypotheses regarding homogeneous diagnostic properties, which are often noncentral. For example, having a heart seems to be projected less frequently than being black amongst ravens. However, the former predicate seems to be more projectible than the latter.

As for flexible similarity, which properties do candidate predicates pick out (cf. Heit, 2000)? Evidence suggests that predicates select those properties that help explain their presence in the base category (see Lassaline, 1996; Sloman, 1994; Smith et al., 1993). Our aim in this paper is to propose and evaluate another account that emerges from people's knowledge about the internal composition and mechanisms that govern instances of a category. One implication of this knowledge is that all properties are not created equal; some play a central role in the activity of an object, others are more peripheral (Sloman, Love, & Ahn, 1998). Having a heart is more central to swans than being white because more of their other properties depend upon having a heart than on swans' color. The proposal that we pursue is that feature centrality is a determinant of projectibility. People will be willing to project central over less central properties to the extent that the categories involved share relevant internal structure. In an effort to find support for this proposal, experiments are reported that operationalize centrality in two ways. In each case, we manipulate the dependency structure of target properties, thereby making these properties more or less central.

1.2. Feature centrality and mutability

Immutability is an aspect of features that reflects the extent to which they resist mental transformation. Sloman et al. (1998) showed that people reliably order conceptual features along an immutability scale by using tasks that asked people to consider an object that is missing a feature but is otherwise intact. For example, their participants reliably said that an object that has all the characteristics of a robin except that it does not eat is harder to imagine than an object that has all the characteristics of a robin except that it does not chirp. Factor and experimental analyses were used to show that feature immutability cannot be reduced to known measures of category structure like variability. The curved shape of bananas is homogeneous across bananas, yet is mutable (see Medin & Shoben, 1988). Immutability is also not salience. The stripes of a zebra are salient, but nevertheless relatively mutable. Finally, immutability is not diagnosticity. Human fingerprints are diagnostic, but not immutable.

Sloman et al. (1998) suggested that one aspect determining the extent to which a feature resists mental transformation is its degree of conceptual centrality. They viewed concepts as reducible to features bound by dependency relations and represented them by dependency graphs. A recursive definition of centrality was supported according to which a feature is conceptually central to the extent that it has many dependents and to the extent that those dependents are themselves central by virtue of having many dependents, and so on. This notion of centrality is driven by an economy principle. People prefer small changes to large ones, and therefore prefer changing features whose effects would not ripple through the conceptual network. The assumption is that transforming a feature affects its dependents, which in turn affect their dependents, etc. Therefore, a feature is conceptually central to the extent that changing it would cause other features to change. Conceptual centrality and immutability are distinct. A feature may resist mental transformation, be immutable, just because it is hard to reach or heavy or for some other reason independent of dependency relations.

To illustrate what we mean by feature centrality, let $X \rightarrow Y \rightarrow Z$ represent the causal knowledge that people have about features X, Y, Z of a concept, where $\rightarrow$ means "causes." All else being equal, people should be least willing to transform property X mentally because doing so would also transform properties Y and Z. People should be relatively willing to transform property Z mentally because doing so would leave properties X and Y unaffected. In support of their account, Sloman et al. (1998) showed that immutability correlated more with the extent to which properties depended on a feature, than with the extent to which a feature depended on other properties (see Study 2 and Appendix B). Note that X is the most conceptually central property but Y occupies the central position in the causal chain. Conceptual centrality reflects the power that a property has to influence other properties of a concept.

The dependency relations that govern conceptual centrality are causal or mostly causal, including direct and indirect causation (e.g., the heart causes blood to pump through the body), enablement (e.g., veins enable blood to reach each cell), and prevention (e.g., the immune system prevents infections). However, some dependency relations might not be causal but rather temporal (e.g., blood arrives at the hips before the feet) or vague (see Section 6.4). Hence, we use the generic term "dependency relations" to describe them. If dependency relations are seen as explanations, the model suggests that a feature is central to the extent that it helps explain the existence of other features. This line of reasoning is similar to a proposal made by Quine (1977) who suggested that the most central predicates are those that support other predicates, in the same way that the axioms of a logical system are its most central statements. Our aim, in this article, is to examine the idea that feature centrality, as just defined, influences feature induction across concepts with similar structure.

1.3. The centrality hypothesis and dependency structure

Our proposal is that people project features from base to target concepts to the extent that they believe the candidate features to be central in the target. However, if the existence of the property in the target is unknown, then its centrality must be too. Therefore, centrality itself must be estimated from the base and projected to the target concept. To a first approximation, centrality can be projected directly from base to target if the base and target are believed to share dependency structure. That is, the more dependency structure two categories share, the more likely they are to share similar immutable elements. Therefore, centrality is only useful for induction in proportion to categories' dependency match. If more than just the centrality of a property is known, if the specific dependency structure of the property in the base is accessible, then centrality can be estimated from knowledge of the presence of the specific dependency structure in the target. That is, when knowledge about a property is specific and not vague, then these specifics can be used to estimate centrality, a more precise indicator of projectibility than similarity.

Thus, our *centrality hypothesis* has two parts:

1. *Centrality claim*: All else being equal, features that are central for a concept are more projectible to other concepts than features that are less central.
2. *Dependency match claim*: The preference to project central over less central features across concepts diminishes as the concepts share less dependency structure. When the candidate feature's dependencies are specified, the preference to project central over less central features diminishes as the concepts share fewer of the specified dependencies.

The centrality hypothesis is relevant when the target is not already known to have the precise dependencies in question. When the precise dependencies are known, induction can take place by reasoning about them directly.

The centrality claim predicts, for example, that a hormone upon which many of a dolphin's functions depend (a central property for dolphins) should be more projectible to seals than a hormone upon which only few of a dolphin's functions depend (a less central property for dolphins). The dependency match claim adds that the influence of centrality is proportional to the extent to which the target is expected to share the candidate feature's dependencies. That is because the more the target is expected to share the candidate feature's dependencies, the more likely the candidate feature is to retain its centrality status in the target. For example, a central feature of a dolphin is more likely to be central in a porpoise than in seaweed. Therefore, a central feature of a dolphin, like a heart, is more likely to be projected to a porpoise than to seaweed. As the extent of shared dependencies decreases, the expected difference in the centrality status of the central and the less central features should decrease as well. For very dissimilar category pairs, like dolphins and seaweed, the centrality effect may level off or even reverse. For example, people might be more willing to project a less central predicate, e.g., affected by an accidental oil spill, from a dolphin to seaweed than a central one, e.g., has a heart.

Centrality is likely to influence feature induction across concepts for two reasons. First, centrality is a proxy for homogeneity. Properties that are least variable are most projectible by definition; a property that does not vary must hold across categories. However, people often do not have direct access to variability. The number of available instances of a category may be insufficient to make a reliable judgment of variability. The property under consideration, e.g., an internal organ, might also not be observable. Central properties should be perceived to be the least variable because people find them the hardest to modify across categories. Consider again the abstract $X \rightarrow Y \rightarrow Z$ chain, describing a causal relation between three features of a concept. Set the variability of feature X at an arbitrary value reflecting noise. We expect that variability of its causal consequent, Y, to be a function of the variability of feature X plus noise, and so on for property Z. Noise, and therefore variability, should build up as one passes along the chain.

A second motivation for centrality is that an important goal of induction is to increase conceptual coherence (Thagard, 1989). Coherence implies understanding, and understanding how objects work is perhaps the major goal of induction because it promotes successful prediction, generalization, and creativity (see Rozenblit & Keil, 2002). Induction of a property increases understanding to the degree that the property explains other properties. Central properties provide more explanatory value than noncentral properties, if dependencies are understood as explanations. Central properties come earlier in the explanatory chains that describe the internal structure of a category.

1.4. Tests of the centrality hypothesis

We report four studies that examined the centrality hypothesis using single-premise categorical arguments, in which a predicate is projected from a base to a target concept. For example, consider the argument:

Eagles have an ulnar artery.
Therefore, falcons have an ulnar artery.

in which having an ulnar artery is projected from eagles to falcons. The statement above the line is the premise of the argument and is assumed to be true. The statement below is the conclusion. An argument is psychologically strong to the extent that its conclusion is judged to follow from its premise. When there is no independent reason to believe the conclusion, the strength of an argument is the judged probability that the conclusion holds given the premise. The strength of the argument reflects the projectibility of the feature from the base to the target concept.

1.4.1. Unspecified dependencies

People's knowledge of dependency structure is often vague and ill formed (Keil, 1995). One of the virtues of the centrality hypothesis is that it does not require well-specified dependency knowledge. Even vague knowledge is sufficient as long as it determines a centrality ordering. Experiments 1, 2, and 4 test this proposition by defining centrality using unspecified dependency relations (i.e., "depends on"). Our prediction was that centrality defined from such relations would influence their feature projections.

1.4.2. Feature centrality manipulations

Feature centrality is operationalized in our studies in two ways. The first is by multiple dependencies, and the second is by a single dependency chain. Experiments 1, 2, and 4 operationalized relative feature centrality in the first way, from multiple dependencies. Many of the base category's functions depend on central features, and few of its functions depend on less central features. Experiment 3 operationalized relative feature centrality in the second way, by a single causal chain. We used a single chain of the form $X \rightarrow Y \rightarrow Z$. Both operationalizations derive from our centrality definition: Features are central to the degree that other features depend on them.

1.4.3. Dependency match manipulations

According to our dependency match claim, the influence of centrality on projectibility is proportional to the degree to which the target concept is believed to share the feature's dependencies, which are specified for the base concept. In Experiments 1, 2, and 4 the feature's dependencies were left unspecified. Thus, the extent to which the target shared these dependencies was not known. The question of what measure of similarity best captures the relevant dependency match between base and target was addressed empirically. Experiments 1 and 2 used animal categories and manipulated biological similarity, a measure that could be characterized as generic in that domain. To the extent that biological similarity captures the relevant dependency match in the animal domain, our centrality hypothesis predicts a tendency to project central over less central features, which should decrease from the high to the low similarity conditions. Experiment 4 used artifact categories and the feature's dependencies were functions. It examined two measures of similarity: functional similarity and surface similarity. Experiment 4 allowed us to examine whether a more specific measure of similarity, functional similarity, is needed to capture the relevant dependency match between the target and the base.

In Experiment 3, we examined the dependency match claim in a context where the dependencies were well-specified and the relevant dependency structure in the base and target concepts—animals—was obvious (e.g., relating to metabolism). All target concepts shared the relevant dependency structure irrespective of their level of overall similarity. Therefore, unlike in the other experiments, in Experiment 3 we predicted a preference for projecting central features over less central features at all levels of similarity.

1.4.4. Domain-general or domain-specific?

Centrality is a structural aspect of representations and so should be a domain-general constraint. Artifacts, for instance, are also understood via intuitive theories, which can be represented by asymmetric dependency links. Thus, centrality should be an aspect of artifact representations as well. People know, for instance, that the fire-fighting equipment of fire engines helps them perform their designed function, which is to extinguish fires. Hence, the fire-fighting equipment of fire engines is one of their relatively central features. The present studies addressed this issue by using both natural kind categories (Experiments 1–3) and artifact categories (Experiment 4).

2. Experiment 1: centrality as the number of dependent properties

Experiment 1 examined the dependency match hypothesis in a context where the target category is not known to share the features' dependencies in the base category. In these circumstances, we expected that participants would estimate the degree to which the target shares the feature's dependencies, by appealing to the degree to which the target shares known dependencies with the base. The greater the target's similarity to the base, the greater the likelihood that the target shares the feature's dependencies. When dependencies do not match, the central feature need not have any inductive priority over the noncentral feature. In short, the preference for projecting central over less central features should be proportional to the similarity of the two categories.

2.1. Method

2.1.1. Participants

Twenty-four University of Durham students volunteered.

2.1.2. Design and materials

Experiment 1 crossed Centrality (central versus less central features) with Similarity (high versus medium versus low biologically similar animal pairs) in a repeated measures design.

Participants were informed that an animal had two features. One feature was the central feature: many of the animal's functions depended on it. The other feature was less central: few of the animal's functions depended on it. Participants were asked to estimate the likelihood that another animal had each of these features. High similarity items paired two similar mammals (e.g., rhinos–hippos). Medium similarity items paired two dissimilar mammals (e.g., beavers–hippos). Low similarity items paired a bird and a mammal (e.g., falcons–hippos). A check on the similarity manipulation was carried out by an independent group of 11 University of Durham students. These students were asked to rate the biological similarity of the 18 animal pairs of Experiment 1 on a 0–10 scale, where 0 was labeled as "highly dissimilar" and 10 as "highly similar." The mean (*SE*) similarity ratings were 7.86 (.30) for the high similarity condition, 3.26 (.61) for the medium similarity condition, and 1.76 (.56) for the low similarity condition validating the similarity manipulation.

Table 1
The 18 premise-conclusion categories of Experiments 1 and 3

Triple	Similarity			Conclusion
	High	Medium	Low	
1	Rhino	Beaver	Falcon	Hippo
2	Squirrel	Bear	Eagle	Mouse
3	Cow	Ferret	Robin	Horse
4	Gorilla	Zebra	Blackbird	Chimp
5	Lion	Raccoon	Sparrow	Tiger
6	Seal	Deer	Swallow	Dolphin

The 18 premise categories were organized into six sets of triples. All three categories of a premise triple were paired with a single conclusion category to form three pairs, one in each of the three similarity conditions. Thus, there were six category pairs in each of the conditions. Table 1 presents the six premise triples and their conclusion categories. Each item in each similarity condition contained both the central and the less central property. Table 2 presents an example from each Similarity by Centrality condition.

Three types of biological features were used: enzymes, hormones, and neurotransmitters. The central and less central counterparts of an item used features from the same biological class (e.g., enzymes). Three lists of materials were constructed from the original list of 18 items in which feature-types were assigned to arguments according to a Latin square design. This ensured that each feature-type occurred equally in each similarity condition and that, across the three lists, each feature-type appeared once with each argument within a triple. Each of these lists had a counterpart with the names of the central and less central features reversed, making six lists in all. With three similarity conditions in each of the six triples, there were 18 items in each list.

Table 2
Sample items from each similarity by centrality condition of Experiment 1

High similarity item	
Fact: Many of a squirrel's physiological functions depend on the enzyme amylase, but only a few depend on the enzyme streptokinase. Please rate the likelihood of the following statements.	
Central	A. Mice have amylase. _ %
Less central	B. Mice have streptokinase. _ %
Medium similarity item	
Fact: Many of a bear's physiological functions depend on the hormone ACTH, but only a few depend on the hormone LH. Please rate the likelihood of the following statements.	
Central	A. Mice have ACTH. _ %
Less central	B. Mice have LH. _ %
Low similarity item	
Fact: Many of an eagle's physiological functions depend on the neurotransmitter GABA, but only a few depend on the neurotransmitter glycine. Please rate the likelihood of the following statements.	
Central	A. Mice have GABA._ %
Less central	B. Mice have glycine._ %

2.1.3. Procedure

Participants were presented with a booklet, containing one of the six lists of materials. They were asked to imagine that they had recently obtained an authoritative book on animal biology, only to find out that some of the pages were missing or were torn apart. They were left with some excerpts from the book stating that a given animal has a biological feature. Their task was to rate the likelihood that another animal had the feature, on a scale ranging from 0 to 100%. Participants were asked to treat each test example separately. They worked through examples before starting the experiment. Time was not limited but participants were encouraged to work quickly.

2.2. Results

Table 3 summarizes the results. Central features were more projectible than less central ones and this difference in projectibility was more pronounced in the high similarity condition than in the other similarity conditions.

Two 2 (Centrality) $\times$ 3 (Similarity) ANOVA were carried out. In the F_1 analysis both factors were repeated measures, in the F_2 analysis only Centrality was repeated measures. The analyses showed a main effect of Centrality ($F_1(1, 23) = 11.51$, $MSE = 165.01$, $p < .005$; $F_2(1, 15) = 24.23$, $MSE = 20.59$, $p < .001$) and a main effect of Similarity ($F_1(1.19, 27.40) = 16.24$,[3] $MSE = 499.75$, $p < .001$; $F_2(2, 15) = 50.84$, $MSE = 24.06$, $p < .001$). The Similarity by Centrality interaction was also significant ($F_1(1.34, 30.70) = 8.85$, $MSE = 287.68$, $p < .005$; $F_2(2, 15) = 20.75$, $MSE = 20.59$, $p < .001$).

To examine the dependency match claim, according to which the centrality effect is proportional to similarity, a Page's L test for ordered alternatives was carried out. The similarity conditions were placed in increasing order and the difference between the central and the less central projectibility score for each subject in each similarity condition was calculated. The resulting "centrality effect" scores were cast in a two-way table with 24 rows and three columns and scores in each row were ranked from 1 to 3. Higher scores were assigned higher ranks. The null hypothesis is that average ranks across similarity conditions are equal. The alternative hypothesis is that at least one of the average ranks gets larger as similarity increases. The average rank scores were 2.41 for the high, 1.92 for the medium, and 1.67 for the low similarity conditions, $L(N = 24, k = 3) = 306$; $z = 2.60$, $p < .005$. Subsequently, the difference in the rank totals between each pair of similarity conditions were calculated. The rank total for the high similarity condition (58) was

Table 3
Mean (*SE*) percentage likelihood estimates as a function of centrality and similarity for Experiment 1

Similarity	Centrality	
	Central	Less central
High	74.88 (3.05)	54.83 (5.53)
Medium	52.79 (3.87)	47.54 (4.23)
Low	43.92 (4.12)	47.42 (4.60)

significantly higher than the rank total for the medium (46) and low (40) similarity conditions. The rank total for the medium similarity condition was marginally higher than the rank total for the low similarity condition. The observed difference was 6.00, and the critical difference at .05 level was 8.29. Overall, the results support the dependency match claim.

Correlation analyses comparing mean similarity ratings to corresponding mean likelihood ratings across items were also computed. Our dependency match claim predicts a high correlation between similarity ratings and the difference in likelihood between the central and the less central features, that is, the centrality effect. Supporting our claim, this correlation was very high (.90, $p < .01$). The correlation between similarity and likelihood ratings for central properties was .98, $p < .01$, and between similarity and likelihood ratings for the less central properties was .61, $p < .05$. The difference between the two correlations was found to be significant by Williams' (1959) t test for non-independent r's, $t(15) = 5.02$, $p < .01$.

2.3. Discussion

The main purpose of Experiment 1 was to assess the centrality hypothesis with centrality defined vaguely in terms of the number of properties that depend upon a feature. In this context, we predicted that the effect of centrality on projectibility should be proportional to similarity. It was. Across participants, the preference to project central over less central features was more pronounced in the high similarity condition. The difference between medium and low similarity conditions was marginal. Across items, correlation analyses showed a very high correlation between similarity ratings and the centrality effect.

3. Experiment 2: blank predicates

One objective of Experiment 2 was to replicate the results of Experiment 1. A second objective was to test an alternative hypothesis: that the results of Experiment 1 are due to a similarity heuristic whose use is invited only when features have many dependent properties, like our central features. This interpretation is supported by the higher correlation between similarity and likelihood ratings for the central property than for the less central property (.98 and .61, respectively). In contrast, we hold that centrality has an effect beyond mere similarity. Central features are more projectible to similar targets than less central features. This happens, not just because such targets share more of their dependent properties, but because the projection of a central feature provides support for these properties.

We compared the hypotheses by including a condition in which no information was given about the features. The features were unfamiliar and their dependency relations were unspecified. These no-information-given features, e.g., <u>has a left aortic arch</u>, are so-called *blank predicates*. Induction based on blank predicates is greatly influenced by base-target similarity (e.g., Rips, 1975). So, if similarity was the only operative factor, then central properties should not be any more projectible than blank predicates. If, as we propose, centrality has an effect beyond mere similarity, then we should observe more willingness to project central features over blank predicates. As a second means to decide between these hypotheses, we computed the correlation between similarity and likelihood estimates for central features, less central features, and blank predicates. To the extent that our hypothesis is correct, these correlations should not differ.

3.1. Method

3.1.1. Participants

Thirty-four University of Durham first-year undergraduates in psychology volunteered to participate.

3.1.2. Design and materials

Centrality (central versus less central versus blanks predicates) was crossed with Similarity (high versus medium versus low biologically similar animal pairs) in a repeated measures design.

Three triples from Experiment 1 were used, the first three triples of Table 1, with each member of a triple assigned to a different similarity condition. As in Experiment 1, three types of features were used: enzymes, hormones, and neurotransmitters. For a given item, all features were of the same type. Three lists of materials were constructed that counterbalanced the assignment of feature-type to animal pairs. Each of these lists had a counterpart where the names of the central and the less central features were reversed, making six lists in all.

3.1.3. Procedure

Each participant had to rate nine items, three in each similarity condition. Each item asked for three likelihood estimates, one in each centrality condition. Participants, therefore, had to provide 27 likelihood estimates in all. Table 4 presents a sample item from each Similarity by Centrality condition constructed using the second triple of Table 1. Participants were given a booklet containing one of the six lists of materials. The instructions were similar to those of Experiment 1. Participants worked through examples before starting the experiment.

3.2. Results

The results are summarized in Table 5. Central features were more projectible than both less central and blank features. The difference in projectibility between central and either less central or blank features was proportional to similarity.

3.2.1. Participant analyses

A 3 (Centrality) × 3 (Similarity) ANOVA was conducted across participants with repeated measures on both factors. F_2 statistics were not calculated because there were only three items per similarity condition. A main effect of Centrality ($F_1(1.49, 49.17) = 15.11$, $MSE = 504.02$, $p < .001$) was detected, and a main effect of Similarity ($F_1(1.46, 48.30) = 21.09$, $MSE = 440.05$, $p < .001$). The interaction was also significant ($F_1(2.26, 74.65) = 8.05$, $MSE = 137.56$, $p < .001$).

Table 4
Sample items from each similarity by centrality condition of Experiment 2

High similarity item

Facts: Squirrels have the neurotransmitter taurine upon which lots of their physiological functions depend, the neurotransmitter glycine upon which few of their physiological functions depend, and the neurotransmitter tyrosine about which you have no information. Please rate the likelihood of the following statements.

Central	Mice have taurine. _ %
Less central	Mice have glycine. _ %
Blank	Mice have tyrosine. _ %

Medium similarity item

Facts: Bears have enzyme aliesterase upon which lots of their physiological functions depend, enzyme streptokinase upon which few of their physiological functions depend, and enzyme elastase about which you have no information. Please rate the likelihood of the following statements.

Central	Mice have aliesterase. _ %
Less central	Mice have streptokinase. _ %
Blank	Mice have elastase. _ %

Low similarity item

Facts: Eagles have hormone MSH upon which lots of their physiological functions depend, hormone ACTH upon which few of their physiological functions depend, and hormone TSH about which you have no information. Please rate the likelihood of the following statements.

Central	Mice have MSH. _ %
Less central	Mice have ACTH. _ %
Blank	Mice have TSH. _ %

Central properties were more projectible than less central properties ($F_1(1, 33) = 14.72$, $MSE = 370.55$, $p < .001$) and this effect depended on level of similarity ($F_1(1.36, 45.02) = 9.19$, $MSE = 144.88$, $p < .005$). To examine the predicted pattern of interaction, a Page's L test was carried out similar to that of Experiment 1. The average rank scores were 2.35 for the high, 2.06 for the medium, and 1.59 for the low similarity conditions. The results supported our prediction: $L(N = 34, k = 3) = 434$; $z = 3.15$, $p < .001$. The rank total for the high similarity condition (80) was greater than the rank total for the medium similarity condition (70), which was in turn greater than the rank total for the low similarity condition (54) (p's $< .05$). The results support the dependency match claim.

Table 5
Mean (*SE*) percentage likelihood estimates as a function of centrality and similarity for Experiment 2

Similarity	Centrality		
	Central	Less central	Blank
High	80.43 (2.65)	62.00 (4.14)	57.74 (4.19)
Medium	60.98 (4.50)	52.30 (3.85)	47.92 (3.90)
Low	55.76 (4.45)	51.64 (4.10)	47.83 (4.28)

As predicted, central properties were also more projectible than blank properties ($F_1(1, 33) = 18.82$, $MSE = 568.15$, $p < .001$) and that effect also depended on level of similarity ($F_1(1.34, 44.09) = 11.01$, $MSE = 131.59$, $p = .001$). A Page's L test produced very similar results to the one comparing central versus less central properties. The average rank scores were 2.40 for the high, 2.04 for the medium, and 1.56 for the low similarity conditions: $L(N = 34, k = 3) = 436.5$; $z = 3.46$, $p < .001$. The rank total for the high similarity condition (81.5) was greater than the rank total for the medium similarity condition (69.5), which was in turn greater than the rank total for the low similarity condition (53) (p's $< .05$). Less central properties were more projectible than blank properties ($F_1(1, 33) = 4.65$, $MSE = 872.96$, $p < .05$) but that effect was independent of level of similarity ($F_1 < 1$).

3.2.2. Item analyses

Correlation analyses were carried out across items comparing similarity estimates to likelihood estimates. The respective correlation coefficients for the central, less central, and blank predicates were .98, .93, and .86 (p's $< .01$). Thus, contrary to the mere similarity interpretation, all types of features invited the use of similarity to estimate projectibility. Furthermore, the influence of centrality on projectibility was proportional to similarity irrespective of whether it was calculated as the difference between the central and the less central properties ($r = .94$, $p < .01$) or as the difference between the central and blank properties ($r = .88$, $p < .01$).

3.3. Discussion

One objective of Experiment 2 was to replicate the results of Experiment 1. It did. Across participants, the preference to project central over less central features was significantly more pronounced in the high followed by the medium followed by the low similarity condition. Across items, there was a very high correlation between similarity and the centrality effect, regardless of whether the centrality effect was calculated using the less central properties or the blank properties.

A second objective of Experiment 2 was to show that feature centrality has an effect beyond mere similarity. One way we did so was by comparing the projectibility of central features to that of blank predicates. We reasoned that if centrality has an effect beyond similarity, then central features should be more projectible than blank predicates. This was the result obtained. Central features were more projectible than both blank predicates and less central features. Another way we tested the mere similarity claim was by comparing the correlation coefficients between similarity and likelihood estimates for the three types of features; they were all very high. Taken together, the results strongly suggest, as our centrality hypothesis proposes, that the effect of centrality is mediated by but not reducible to similarity.

A rather counterintuitive result of Experiment 2 was that less central properties were more projectible than blank predicates. Since no information was given about the blank predicates, one might have expected them to have an intermediate level of centrality. Perhaps putting all types of features side-by-side for evaluation made participants interpret blank predicates as the least central.[4]

4. Experiment 3: centrality from a feature's position in a dependency chain

One objective of Experiment 3 was to examine the hypothesis that a feature is more projectible than its dependents. This hypothesis derives from the proposition that centrality can be inferred from local pairwise dependencies. Relative feature centrality was operationalized using a dependency chain that in the abstract had the form $X \rightarrow Y \rightarrow Z$, where $\rightarrow$ means "causes." Feature X is more central than feature Z because two features depend on X but none on Z. This hypothesis gains support from categorization studies (e.g., Ahn & Dennis, 1997; Ahn & Lassaline, 1995) that show that causal features influence category decisions more than effect features (*the causal status hypothesis*). Because causal relations are a special case of dependency relations, these studies suggest that a feature is weighted more strongly in categorization decisions than its dependents. Inference is a major function of categorical knowledge in that knowing that an object has a novel property suggests that other similar objects will also have the property. To the extent therefore that the status of a feature in a dependency chain influences categorization, it should also influence inference.

A second objective of Experiment 3 was to show that, when people have knowledge about the specific properties that depend on a candidate feature, they use this knowledge to estimate projectibility. To that end, Experiment 3 used a similar design as previous experiments except that dependent properties were specific and familiar (e.g., relating to metabolic rate). All target categories—animals—were known to possess these properties. We predicted that the preference to project central over less central features would be unaffected by similarity. That is, we predicted main effects of centrality and similarity but no interaction.

4.1. Method

4.1.1. Participants

Thirty-three University of Durham undergraduates volunteered to participate.

4.1.2. Design and materials

Experiment 3 crossed Centrality (central versus less central features) with Similarity (high versus medium versus low biologically similar animal pairs) in a repeated measures design.

Participants were informed that an animal had two features: one upon which a specific function depended (the central feature) and another that depended upon that same function (the less central feature). Participants were asked to estimate the likelihood that another animal had each of these features. The design was identical to that of Experiment 1 with the exception that the central and less central counterparts of an item were split into two separate items making 36 items overall. Table 6 presents an example from each Similarity by Centrality condition.

4.1.3. Procedure

The procedure used was identical to Experiments 1 and 2.

4.2. Results

The results are summarized in Table 7. Central features were more projectible than less central features and this effect was not influenced by similarity.

Table 6
Sample items from each similarity by centrality condition of Experiment 3

Central	Less central
High similarity items	
Fact: Rhinos have the enzyme lipase. For rhinos, the enzyme lipase regulates metabolism. Rate the likelihood of the following statement. Hippos have lipase. _ %	Fact: Rhinos have the enzyme protease. For rhinos, the enzyme protease is regulated by metabolic rate. Rate the likelihood of the following statement. Hippos have protease. _ %
Medium similarity items	
Fact: Beavers have the hormone prolactin. For beavers, the hormone prolactin regulates blood flow. Rate the likelihood of the following statement Hippos have prolactin. _ %	Fact: Beavers have the hormone renin. For beavers, the hormone renin is regulated by blood flow. Rate the likelihood of the following statement Hippos have renin. _ %
Low similarity items	
Fact: Falcons have the neurotransmitter acetylcholine. For falcons, the neurotransmitter acetylcholine helps detect predators. Rate the likelihood of the following statement. Hippos have acetylcholine. _ %	Fact: Falcons have the neurotransmitter noradrenalin. For falcons, the levels of the neurotransmitter noradrenalin increase after seeing a predator. Rate the likelihood of the following statement. Hippos have noradrenalin. _ %

Two 2 (Centrality) × 3 (Similarity) ANOVA were carried out. In the participants analysis (F_1) both factors were repeated, in the items analysis (F_2) only Centrality was a repeated factor. The analyses showed a main effect of Centrality ($F_1(1, 32) = 4.75$, $MSE = 214.47$, $p < .05$; $F_2(1, 15) = 19.55$, $MSE = 9.81$, $p < .001$), a main effect Similarity ($F_1(1.34, 43.00) = 68.48$, $MSE = 347.09$, $p < .001$; $F_2(2, 15) = 42.16$, $MSE = 69.28$, $p < .001$), and no interaction (both F's < 1).

Across items, correlation analyses comparing mean similarity ratings to corresponding mean likelihood ratings were carried out. The correlation between similarity and argument strength ratings for central properties was .91, $p < .01$, and between similarity and argument strength ratings for the less central properties was .93, $p < .01$. No association was detected between similarity ratings and the difference in argument strength ratings between the central and the less central features, i.e., the centrality effect, $r = -.02$.

Table 7
Mean (*SE*) percentage likelihood estimates as a function of centrality and similarity in Experiment 3

Similarity	Centrality	
	Central	Less central
High	72.26 (2.87)	67.09 (3.56)
Medium	47.46 (4.05)	44.11 (3.97)
Low	43.02 (3.74)	37.93 (3.60)

4.3. Discussion

One purpose of Experiment 3 was to assess whether a feature is more projectible than its dependents. It was. This result supports the primary claim of the centrality hypothesis and suggests that people are sensitive to the directionality of dependency relations when projecting properties. A second purpose was to show that when the dependency structure in the base is accessible, centrality is estimated directly from the knowledge of the specific dependency structure in the target. In Experiment 3, all target categories were known to share the relevant dependency structure because that structure was obvious and familiar (e.g., relating to blood flow). So central properties promised to increase the targets' coherence by the same amount, irrespective of the targets' overall similarity with the base. Supporting this claim, there was a main effect of centrality but no similarity by centrality interaction. Note that Experiment 3 failed to show an interaction with a design that was powerful enough to show such an interaction in Experiments 1 and 2. This suggests that, when dependency structure is available, inductive inference is not just a matter of assessing similarity.

In Experiment 3 the effect of centrality was small. This might be due to background knowledge effects. We might have inadvertently contradicted general knowledge by using the same type of properties to construct the central and the less central counterparts of an item. For example, although we stated that an enzyme was controlled by metabolic rate, participants may still have thought that the enzyme regulated metabolism. Such background knowledge or belief may have scaled down or, in some cases, even reversed the intended asymmetry.

5. Experiment 4: extending the centrality hypothesis to artifacts

The present account assumes a domain-general inference mechanism. Many recent studies have demonstrated effects of domain-specific knowledge in concept learning and use (e.g., Carey, 1985; Murphy & Medin, 1985; Pazzani, 1991; Wattenmaker, 1995; for a review see Heit, 1997). Such evidence appears to challenge the possibility that a domain-general inference process underlies categorical induction (cf. Hirschfield & Gelman, 1994).

5.1. Conceptual centrality can explain apparent domain differences

Functional features (e.g., <u>is used to pound nails</u>) are generally important for identifying artifacts, whereas compositional features (e.g., <u>has a left aortic arch</u>) are important for natural kinds (see Barton & Komatsu, 1989). Exceptions to this rule have also been found (Malt & Johnson, 1992; Sloman, Malt, & Fridman, 2001). Using the materials from both Barton and Komatsu (1989) and Malt and Johnson (1992), Ahn (1998) showed that in both cases causal centrality was correlated with judged category importance. Ahn's last two studies examined artificial natural kind and artifact categories while directly manipulating the causal status of a feature. The results showed that high-causal status features were judged as more important in categorical decisions than low-causal status features across both artifacts and natural kind categories. Taken together, Ahn's studies strongly suggest that a feature's causal status constrains categorization. Causal status is a special case of conceptual centrality, therefore her studies corroborate the idea that conceptual centrality constrains categorical inference.

Like knowledge about natural kinds, knowledge about artifacts can be represented by asymmetric dependency links connecting their properties (cf. Sloman & Malt, 2003). Thus, the features of an artifact should also differ in terms of conceptual centrality. Being able to freeze stuff, for instance, is a conceptually central feature of refrigerators; being white is not. As with animal categories, we expect that the more central a feature for a concept the higher its projectibility to other concepts that share its attributes and dependencies.

Experiment 4 examined the centrality hypothesis with artifacts. Like Experiments 1 and 2, relative feature centrality was manipulated via the number of dependencies: Many unspecified functions depended upon the central feature but few upon the less central feature. Unlike the previous experiments, Experiment 4 manipulated two types of similarity: functional similarity and surface similarity. We examined whether the relevant dependency match between the target and the base categories is best captured by a more specific measure of similarity, in the present case functional. We used three similarity conditions reflecting degree of functional similarity: high, medium, and low. The medium similarity condition was lower in terms of surface similarity than the low similarity condition. The medium and low similarity conditions were matched in terms of overall similarity (the average of functional and surface similarity). If functional similarity captures the relevant dependency match between target and base, then functional similarity should exert the most influence on the preference to project central over less central features. If a generic similarity measure captures the relevant dependency match between target and base, then overall similarity should exert the most influence on the centrality effect.

Our experiments so far have tried to minimize participants' background knowledge of the candidate features. Animal categories, being highly complex, are well suited to study our hypothesis in part because they possess many properties that people do not know about. Few people outside biological sciences know the names or precise functions of biological substances. For artifacts, and especially for simple ones like chairs, credible novel features are harder to generate.[5] For this reason, the candidate features in the present experiment were named in Greek (using Latin characters). Centrality was defined by the number of dependent properties. Participants were instructed that they would learn two facts about a category, and that they would be asked to generalize these to another category. The task was presented in the context of a TV show that promised a large sum of money for the best answers.

5.2. Method

5.2.1. Participants

The participants were 22 first-year undergraduates of the University of Durham participating in a tutorial.

5.2.2. Design

This experiment crossed Centrality with Functional similarity (high versus medium versus low functionally similar pairs of artifacts) in a repeated measures design. The medium similarity condition was higher in terms of physical features than the low similarity condition. The

Table 8
Premise triples along with the target categories for the items of Experiment 4

Premise triples			Target
High functional high surface	Medium functional low surface	Low functional high surface	
Lorry	Cargo boat	Toy-truck	Truck
Trumpet	Music box	Plastic saxophone	Saxophone
Washing machine	Car-wash	Refrigerator	Dishwasher
Clock	Hourglass	Compass	Watch
Mac	Calculator	TV	IBM-PC
Microwave	Camp stove	Freezer	Oven
Boeing 747	Rocket ship	Remote-control plane	Concorde
Bungalow	Tent	Barbie's house	House

medium and low similarity conditions were matched in terms of average functional and surface similarity.

5.2.3. Materials

High similarity items were constructed by pairing two categories that came from the same superordinate (e.g., lorry–truck), medium similarity items by pairing two functionally similar but physically dissimilar categories (e.g., cargo-boat–truck), and low similarity items by pairing two functionally dissimilar but physically similar categories (e.g., toy-truck–truck). The materials were organized in eight triples in the same way as in the previous experiments. One member of each triple appeared in the high similarity condition, one in the medium similarity condition, and a third in the low similarity condition. The difference between similarity conditions was the category term in the premise of the argument. Table 8 presents the 24 categories used and Table 9 a sample item from each Similarity by Centrality condition. Like in Experiments 1 and 2, relative centrality was defined by manipulating the number of properties depending on a feature.

5.2.4. Selection of base-target pairs

The authors constructed a series of triples with a single target that were expected to differ in their surface and functional similarity. The list was then further refined to produce a list of eight triples, with one member of a triple in each Functional similarity condition (see Table 9). This list was then presented to a separate group of 12 participants who were asked to rate each item for its functional and surface similarity, on a scale ranging from 0 to 100. For the functional similarity estimates, participants were asked; "Please rate how similar are the following pairs of objects in terms of *function*. The higher the rating the more *functionally similar* you think that the pair of objects are." For the surface similarity estimates, participants were instructed as follows: "Please rate how similar are the following pairs of objects in terms of their *surface properties* (e.g., how do they look). The higher the rating, the more *similar looking* you think that the pair of objects are." The results are shown in Table 10. The assignment of category pairs to similarity conditions was validated by participants' judgments.[6,7]

Table 9
Sample items from each similarity by centrality condition of Experiment 4

High similarity item	
Fact: Lots of a lorry's functions depend on *michani* but only a few depend on *kitrino*. Please rate the likelihood of the following statements.	
Central	A. Trucks have *michani*. _ %
Less central	B. Tracks have *kitrino*. _ %
Medium similarity item	
Fact: Lots of a cargo-boat's functions depend on *propela*, but only a few depend on *kathisma*. Please rate the likelihood of the following statements.	
Central	A. Trucks have *propela*. _ %
Less central	B. Trucks have *kathisma*. _ %
Low similarity item	
Fact: Lots of a toy-truck's functions depend on *rodes*, but only a few depend on *autokolita*. Please rate the likelihood of the following statements.	
Central	A. Trucks have *rodes*. _ %
Less central	B. Trucks have *autokolita*. _ %

In each questionnaire, half of the questions in each similarity condition asked for the likelihood that the target has the central feature first and the less central feature second. For the other half, this order was reversed. The order of evaluation of the central and the less central features was counterbalanced across participants. Two presentation orders were used, each order presented to half the participants. In the present experiment a single feature was assigned to each base-target pair. We felt it unnecessary to counterbalance the assignment of features to premise triples because the features were in Greek.

5.2.5. Procedure

The items were presented in booklets. Participants were asked to imagine that they were finalists in a TV quiz show a few questions away from winning the grand prize: £100,000. The presenter unveiled the last task named "It's all Greek to me!" They were told that they would be informed that an object had two features. Their task was to rate the likelihood of another object having each of these features, on a scale ranging from 0 to 100%. The catch was that the names of the features were given in Greek. Participants worked through an example before proceeding with the test items.

Table 10
Mean (*SE*) functional similarity and surface similarity ratings for the three similarity conditions of Experiment 4

Similarity	Similarity type		Mean
	Functional	Surface	
High	77.72 (5.97)	67.37 (4.67)	72.55 (4.21)
Medium	43.66 (4.55)	20.90 (5.08)	32.28 (3.29)
Low	17.55 (2.65)	55.97 (4.60)	36.77 (2.65)

Table 11
Mean (*SE*) percentage likelihood estimates as a function of centrality and functional similarity for Experiment 4

Similarity	Centrality	
	Central	Less central
High	61.91 (2.28)	42.88 (3.42)
Medium	35.37 (2.66)	30.71 (3.30)
Low	33.72 (2.80)	30.76 (3.11)

5.3. Results

Table 11 summarizes the results. The size of the centrality effect for the high similarity condition was larger than that for the other conditions. The centrality effects for the medium and low similarity conditions were of equal size.

Two 2 (Centrality) $\times$ 3 (Functional similarity) analyses of variance were carried out. In the F_1 analysis both factors were repeated, in the F_2 only Centrality was repeated. There was a main effect of Centrality ($F_1(1, 21) = 12.91$, $MSE = 200.25$, $p < .001$; $F_2(1, 21) = 34.83$, $MSE = 27.10$, $p < .001$), and a main effect of Functional similarity ($F_1(1.82, 38.11) = 81.23$, $MSE = 78.11$, $p < .001$; $F_2(2, 21) = 17.56$, $MSE = 119.20$, $p < .001$). The interaction was also significant ($F_1(1.26, 26.41) = 13.37$, $MSE = 100.63$, $p = .001$; $F_2(2, 21) = 11.43$, $MSE = 27.10$, $p < .001$).

To examine the predicted pattern of interaction, a Page's L test was carried out. The average rank scores were 2.64 for the high, 1.66 for the medium, and 1.70 for the low similarity conditions; $L(N = 22, k = 3) = 285.5$; $z = 3.24$, $p < .001$. The locus of this effect was a significant difference between the rank total for the high similarity condition (58) and the rank total for the medium (36.5) and low (37.5) similarity conditions. No significant difference was observed between the medium and low conditions.

For items, we also performed correlation analyses comparing mean similarity ratings from each similarity measure (functional, surface, and average) to corresponding mean likelihood ratings. The results are presented in Table 12. The centrality effect was most influenced by overall similarity. In fact, overall similarity produced the highest correlations across all comparisons. The correlation between overall similarity and likelihood ratings for central properties

Table 12
Correlations across items for Experiment 4

	Similarity		
	Surface	Functional	Overall
Correlation between similarity estimates and likelihood estimates			
Central	.56**	.73**	.82**
Less central	.55**	.66**	.77**
Central–less central	.51*	.69**	.77**

Note. Overall similarity: average surface and functional similarity.

* Correlation is significant beyond the .05 level (two-tailed).

** Correlation is significant beyond the .01 level (two-tailed).

was .82, $p < .01$, between overall similarity and likelihood ratings for less central properties was .77, $p < .01$, and between overall similarity ratings and likelihood differences (reflecting the centrality effect) was .77, $p < .01$. As in Experiments 1 and 2, these results suggest that the relevant dependency structure is best captured by a generic similarity measure. Participants apparently considered both functional and physical features in projecting properties from base to target concepts.

Assuming a correspondence between the present high, medium, and low similarity levels and those of the natural kind experiments, likelihood judgments in Experiment 4 were about 20 points lower than corresponding judgments of Experiment 2 (compare the first two columns of Table 5 to those of Table 11). The lower likelihood judgments of the present experiment might reflect the uncertainty caused by the features being named in Greek.

5.4. Discussion

The primary aim of Experiment 4 was to extend our centrality hypothesis to artifacts. The results are supportive. Central features were more projectible than less central ones as a function of similarity. A secondary aim was to detect whether the centrality effect would be most influenced by a specific and relevant similarity measure, in this case functional, or by a generic similarity measure. The centrality effect was most influenced by overall similarity. Overall similarity correlated the highest with likelihood ratings for central features, less central features, and their difference. Furthermore, the participant analyses showed no difference in the centrality effect between the medium and low similarity conditions, which were matched for overall similarity. In sum, in Experiment 4 overall similarity was the best surrogate for relevant dependency match.

Although Experiment 4 showed an influence of centrality on categorical inference with artifacts, questions remain open about how judgments were made. One suggestion might be that participants interpreted the Greek as naming familiar features of the premise categories. If this had occurred, however, then the central and the less central features should have been equally projectible. They were not, so we reject this possibility. At the other extreme, participants may have used a general centrality heuristic, something like "project the central feature no matter what." This second possibility must also be discarded because it fails to explain the centrality by similarity interaction. A third possibility is that people interpreted the features as familiar features that fit the bill, features that agreed with the centrality descriptions. For instance, they might have interpreted the central feature *michani* as engine (its real meaning), and the less central feature *kitrino* as yellow-colored (its real meaning). That is, the novel central features might have been interpreted as generalizable familiar properties. Notice that the only information participants were supplied with was vague information about a feature's centrality. So the possibility that novel central features cue familiar generalizable features suggests that feature centrality influences generalizability.

We do not doubt that artifact and natural kind concepts differ. Such differences, however, might concern the ways in which dependency patterns in each domain cluster and converge (Keil, 1995). The features of an artifact are generally seen to depend on its designed function, e.g., much about a chair depends on its relational function as something to sit on. The features of a natural kind are generally seen to depend on its underlying intrinsic properties, e.g., the features of an element depend on its atomic number, and many features of a species depend on the processes that keep it alive. Such clustering may reflect the belief that natural kinds have essences or vital forces, whereas artifacts do not (Atran, 1990).

6. General discussion

6.1. Feature centrality constrains inference

The main objective of the present studies was to test the two claims of our centrality hypothesis. The first claim is that the more central a feature is in a base concept, the higher its projectibility to a target concept. To test this claim, we manipulated feature centrality. Experiments 1, 2, and 4 operationalized centrality using the number of a feature's dependent properties, whereas Experiment 3 operationalized centrality using a single dependency chain. That multiple definitions of centrality produced centrality effects is evidence that conceptual centrality influences feature induction.

The second claim is that the tendency to project central over less central features is proportional to the extent that concepts' share a dependency structure that is relevant for inference. To test this claim, we manipulated the extent to which two concepts shared a dependency structure. In Experiment 3, all category pairs had the same relevant structure, only one dependency was relevant and it was specified. As expected, similarity and centrality did not interact in this experiment; the centrality effect was the same at different levels of similarity. In Experiments 1 and 2 the category pairs differed in the amount of relevant shared structure. In these experiments, as expected, the effect of centrality was proportional to similarity.

In our artifact study, Experiment 4, we manipulated functional and surface similarity in an effort to examine what type of similarity best captures the relevant dependency match between base and target. The features' dependencies were functional. We reasoned that if a specific similarity measure was used as a surrogate for the relevant dependency match, then it should be functional similarity that mediates the centrality effect. Alternatively, if a generic similarity measure was used to estimate the relevant dependency match, then the centrality effect should be proportional to overall similarity. Overall similarity was defined as average functional and surface similarity. The latter hypothesis was supported. Overall similarity exerted the strongest influence. Taken together, Experiments 1, 2, and 4 suggest that a generic measure of similarity seems to be used to estimate the extent to which the target shares the feature's dependencies. These experiments also imply that people's categorical inductions show strong centrality effects when the category pairs are highly similar to each other. For example, one should expect these effects for pairs like eagles–falcons in the animal domain, and trumpet–saxophone in the artifact domain.

Our account posits that categorical induction is influenced by feature centrality, a domain-general aspect of representations. It, therefore, predicts that the centrality hypothesis applies across domains. Our hypothesis was supported with both natural kinds and artifacts. Taken together, the findings suggest that conceptual centrality provides a domain-independent constraint on feature projection. Finally, to assess the claim that people are sensitive to vague dependency relations, Experiments 1, 2, and 4 left the dependency relations unspecified. All these experiments showed a centrality effect, thus supporting the claim.

6.2. The explanatory power of centrality

6.2.1. Perceived homogeneity effects

The centrality hypothesis can explain phenomena that have been previously attributed to perceived homogeneity or stability (e.g., Gelman, 1988; Nisbett, Krantz, Jepson, & Kunda, 1983). Nisbett et al. (1983, Study 1), for instance, informed participants that one instance of a new bird, the shreeble, had been observed to be blue and to nest in a eucalyptus tree. Participants were asked to estimate the percentage of all shreebles that manifested each characteristic. Participants judged that a higher percentage of shreebles would nest in eucalyptus trees than would be blue. The researchers attributed this finding to the former property being perceived as more homogeneous than the latter.

The centrality hypothesis suggests that some beliefs about homogeneity derive from beliefs about feature centrality. Centrality influences perceived homogeneity because it suggests causal importance, and causal importance suggests invariance. In support of this claim, Sloman et al. (1998, Study 5) showed that manipulating the centrality of a feature influences subsequent judgments of its frequency: The greater the judged centrality of a feature, the greater its judged frequency. Applying the centrality hypothesis to Nisbett et al.'s (1983) study, the suggestion is that participants preferred to project nesting habits over color to other shreebles because they believed that nesting habits would support more of the shreebles' other properties.

6.2.2. Category-feature interactions

The centrality hypothesis can also account for category by feature effects in property induction. It can explain, for instance, Heit and Rubinstein's (1994) finding that the projection of anatomical features is influenced by anatomical similarity, whereas the projection of behavioral features is influenced by both anatomical and behavioral similarity. Anatomical properties depend on anatomical features, whereas behavioral properties depend on both behavioral and anatomical features. Like in Experiment 4, this is another case where a more generic measure of similarity is the best surrogate of the extent to which the target shares the candidate feature's dependencies. In contrast to Heit and Rubinstein (1994), we claim that centrality has an effect beyond mere context-dependent similarity. An anatomical feature that many properties depend on should be more projectible than an anatomical feature that no information is given about. This prediction was supported in Experiment 2.

6.3. Projectibility and explanatory coherence

By appealing to centrality, people can maximize the coherence of their inductive inferences. By appealing to features that can explain as many of the target's properties as possible, their explanations are more likely to be consistent with one another than explanations generated separately. Thagard (1989) proposed that hypotheses cohere not only with propositions that they help explain, as our centrality hypothesis suggests, but also with propositions that help to explain them. Applied to categorical induction, the implication is that features will be projected to a target to the extent that the features can be explained by properties of the target. Support for the claim can be found in Lassaline (1996) and Sloman (1994, 1997). Sloman (1997), for instance, has shown that statements lend support to one another to the extent that they share an explanation. Many computer programmers have bad backs; therefore, many secretaries have bad backs, for example, was judged stronger than Many successful members of minority groups have a hard time financing a house; therefore, many secretaries have a hard time financing a house, because in the first argument the target shares the feature that explains the predicate in the premise, whereas explanations differ in the second argument.

Our centrality claim is consistent with such findings, if dependencies are interpreted as explanations. If a feature is central, then the property needed to explain it is even more central. This is a consequence of our centrality definition. Furthermore, similar categories are more likely to share central properties because central properties are the least likely to be transformed across similar categories. To illustrate, consider the following plausible explanatory chain for the concept sparrow: having wings $\rightarrow$ ability to fly $\rightarrow$ ability to nest in trees, where $\rightarrow$ means "enables." Ability to nest in trees should be less projectible than the ability to fly because explaining the former calls for a less central property (e.g., ability to fly) than explaining the latter (e.g., having wings).

6.4. Vague dependencies or explanations?

The proposition that a property is projectible to the extent that it increases the explanatory coherence of the target category, either by explaining its properties or by being explained by them, can account for many effects of property induction. However, there are reasons for not limiting the scope of dependencies to explanatory dependencies. Some effects cannot be captured by explanatory connections but could be captured by generic dependencies. This may be seen in a series of studies showing that categorization is constrained by general beliefs which exist in the absence of specific causal knowledge (Keil, 1995; Keil, Smith, Simons, & Levin, 1998; Simons & Keil, 1995). Simons and Keil (1995) presented children with a target object, an animal or a machine, together with a set of potential insides: the insides of an animal, a machine, a pile of rocks, or a pile of blocks. When children were asked to match the target objects with the correct insides, even preschool children expected the insides of machines and animals to differ. They systematically picked different insides but sometimes they picked the wrong ones. Children's decisions were therefore driven by vague notions of centrality, not by specific causal knowledge. On the basis of such findings, Keil et al. (1998, p. 42) argue that "basic notions of causal centrality may emerge early ... the ability to perceive and learn causal patterns may be just as fundamental as the ability to learn typicality and frequency distributions." The centrality hypothesis implies that children projected different insides to animals and machines because they believed that different insides were needed to account for their different appearances and behaviors.

6.5. Implications for models of categorical inference

For models of categorical inference to account for centrality effects, they must represent features, relations between features, directionality, and they must weight features in proportion to their dependencies. Current categorical inference models cannot account for centrality effects because they lack on one or more of these characteristics.

6.5.1. Comparison models

Certain models of induction (the regression model of Rips, 1975, the *similarity-coverage* model of Osherson, Smith, Wilkie, López, & Shafir, 1990, and the *feature-coverage* model of Sloman, 1993) attempt to explain argument strength using fixed indices of categories' relatedness. For example, Rips' model represents categories as solutions in a multidimensional space and predicts argument strength from distance between solutions. Sloman's model represents categories as vectors of values over a set of features, and predicts argument strength from the projection of premise vectors onto the conclusion vector. These models are insensitive to the particularities of the predicate being projected, including its centrality.

In an attempt to capture how properties influence induction, Smith et al. (1993) proposed the GAP model. The GAP model deals with non-blank predicates, such as <u>can bite through barbed wire</u>, which call on prior beliefs. Such predicates are assumed to invite an examination of the plausibility of the argument's premise. Consider an argument with premise <u>Poodles can bite through barbed wire</u>, and conclusion <u>German shepherds can bite through barbed wire</u>. In order to assume that the premise is true in this case, we must significantly change our beliefs about the strength of poodles or of barbed wire. Once we have made that change, we judge that the argument is strong. However, if <u>German shepherds can bite through barbed wire</u> were the premise and <u>Poodles can bite through barbed wire</u> the conclusion, we need not change our beliefs about German shepherds or barbed wire. In this case, we would judge the argument to be weak. Broadly, the more implausible the premise is judged on the basis of prior beliefs, the more these beliefs are revised when the premise is assumed. The more our beliefs are revised, the wider the basis we have for inferring the conclusion. However, the GAP model cannot account for centrality effects because it does not represent relations between features.

6.5.2. Bayesian models

Heit (1998) proposed a Bayesian model of categorical inference in which the goal of induction is to estimate the range of a property. The model states that people estimate the range of a property based on a prior hypothesis about that range. Premises serve as new evidence that people use to revise their confidence in the hypothesis. For this revision, they rely on Bayes's rule. Heit's model does not, however, explain why there are differences in projectibility. It can only represent such differences by choosing appropriate degrees of belief in prior hypotheses. Sanjana and Tenenbaum (2003) also offer a Bayesian model of categorical inference, but one that is equally incapable of explaining centrality effects. The model derives probabilities from category clusters and also has no representation of relations amongst properties.

6.5.3. Structural-alignment models

Unlike the models considered so far, structural-alignment models (e.g., Gentner, 1983, 1989) represent features, relations between features, and directionality. Specifically, they represent categories as features embedded in hierarchical systems of relations. A situation where a man fixes a robot may be represented by Fix(Man, Robot), and a situation where a robot fixes a car by Fix(Robot, Car). These models assume that object comparison involves aligning the objects' relations and properties. Elements connected to the resulting system of relations that are present in the base but absent in the target constitute potential inferences. Alignment is assumed to follow *the principle of systematicity* according to which a match of higher order relations is preferred over a match of lower order relations or properties. In our example, the prediction is that alignment will be based on matching the relation Fix (Fix–Fix, Man–Robot, Robot–Car), rather than the attribute Robot (Robot–Robot, Man–Car). Evidence corroborates that projectibility is determined by the principle of systematicity (see Clement & Gentner, 1991; Wu & Gentner, 1998). However, structural-alignment models cannot account for centrality effects because they give no inductive preference to central properties over their dependents. Consider an argument with premise Robins have wings that enable them to fly. Structural-alignment models do not predict a preference to project wings over flying to another bird. Such a preference was shown in Experiment 3. In sum the finding that centrality influences induction cannot be satisfactorily accounted by current models of categorical inference.

6.6. *Towards a unified account*

What is the relation between the various models of categorical inference? As Gentner and Medina (1998) propose, the relation may be that in the absence of sufficient knowledge (such as in the case of blank predictates), people fall back on default methods such as similarity-based strategies. Some form of similarity- or feature-based model might be adequate in such cases for representing judgments of projectibility. When information is available that relates the candidate feature to other properties of a concept, as in the present experiments, people will make use of such information. In such cases, judgments of projectibiltiy will be constrained by structural aspects like the centrality status of the candidate feature. In cases where such knowledge is vague, a vague centrality heuristic seems to be used. Such cases are common. We often make inferences under conditions of ignorance such as when we lack relevant causal knowledge or when we are operating in a novel context. In cases where the plausibility of premises is a question, the GAP model may prove the best. Notice though that the GAP model is built on the assumption that explainable predicates pick out dimensions of the premise category that help explain the predicate. For an argument with the premise Poodles can bite through barbed wire, the model assumes that dimensions like strength and ferocity get potentiated both of which help explain the ability of dogs to bite through barbed wire. The point is that a notion of coherence based on pairwise relations might be at the core of a unifying model of property inference. Finally, in cases where people can draw an inference directly from causal knowledge they will do so. For example, the argument Lakes are contaminated with toxic waste; therefore, fish in the lakes are contaminated with toxic waste should be judged strong because of a strong causal link between the argument's premise and conclusion (see Medin, Coley, Storms, & Hayes, 2003).

Corroborating evidence that inferential strategies are knowledge-dependent comes from literature on the *diversity phenomenon*, the finding that arguments with two premises are preferred to the extent their premises are dissimilar to one another. Evidence suggests that the diversity phenomenon, a similarity-based phenomenon, is more widespread among individuals who are naive in a domain. Western adults who have scarce knowledge of folkbiology exhibit more diversity than Western experts (e.g., Bailenson, Shum, Atran, Medin, & Coley, 2002; Osherson et al., 1990; Proffitt, Coley, & Medin, 2000). Itza Mayas exhibit diversity when reasoning about abstract logic problems, but not when reasoning about categories they are familiar with like trees or plants (López, Atran, Coley, Medin, & Smith, 1997).

6.7. Conclusions

We have attempted to reduce much of the problem of feature induction to the problem of determining a feature's centrality in a network of dependency relations. Feature projection depends on the specific categories used. The same feature may lend great support to certain targets but little to others. Being round is a projectible feature for wheels but not for oranges, because roundness is central in understanding wheels but not oranges. There are other ways people use knowledge to estimate projectibility. Categorical arguments of the form Lakes are contaminated with toxic waste; therefore, fish in the lakes are contaminated with toxic waste are strong because of a strong causal link between the argument's premise and conclusion (see Medin et al., 2003). Arguments of the form Poodles can bite through barbed wire; therefore, German shepherds can bite through barbed wire seem to invite strategies that take the relative degree of surprise of statements into account. The centrality heuristic is one of many ways people have of bringing causal and other sorts of knowledge to bear on the problem of inductive inference. In the information-rich real world environments where people frequently have vague knowledge of dependence between properties, centrality might prove to be a common inferential heuristic.

Notes

1. Predicates, hypotheses, and arguments are underlined.
2. Although Goodman (1955/1983) attempted to solve the problem of confirmation, that is, the problem of which hypotheses are confirmed from any given evidence, his discussion was limited to within-category projection (e.g., projecting green from a sample of emeralds to all emeralds). The present article examines between-category projection (e.g., projecting the enzyme Dihedron from rabbits to dogs), which is a species of the same problem. Rips (1975; see also Gelman, 1988; Heit, 2000; Shipley, 1993) has extended Goodman's concepts of entrenchment and projectibility to discuss between-category projection. We do the same here.
3. Whenever a repeated factor had three levels, we reduced the degrees of freedom of the accompanying F tests by multiplying their normal values by the corresponding Greenhouse–Geisser epsilon. These adjustments aimed to counteract departures from the sphericity assumption on the variance–covariance matrix (see Howell, 1997, pp. 464–466).
4. We thank an anonymous reviewer for this suggestion.
5. The claim is not that complexity completely differentiates natural kinds from artifacts. Complexity varies within a domain, e.g., water is less complex than zebra; pencil is less complex than microwave oven.
6. Statistical analyses support the assignment of category pairs to similarity conditions. The mean functional and surface similarity ratings for the high similarity condition was higher than that for the medium or low similarity conditions; the mean ratings for the low and medium similarity conditions did not differ. Low similarity pairs were more similar in terms of surface features than medium similarity pairs, whereas medium

similarity pairs were more similar in terms of functional features than low similarity pairs.

7. Rob Goldstone pointed out that our participants' folk psychological notion of superficial similarity may not correspond to the theoretic, psychological notion. Some of our medium similarity pairs, like cargo-boats and trucks, may be in fact more similar on the surface than some of our low similarity pairs, like toy-trucks and trucks. For instance, cargo-boats and trucks may be more confusable than toy-trucks and trucks. Peoples' estimates of surface similarity may be driven by their notion that toy-trucks are designed to resemble trucks. "Surface similarity" is defined with respect to our participants' ratings.

Acknowledgments

Rosemary, in her gracious way, made many contributions to this paper. Sadly, she did not live to see its completion. This research was supported by a fellowship from the Psychology Department of the University of Durham, and the writing of this article was supported in part by ESRC grant R000239074 and NASA grant NCC2-1217. This article is based on a doctoral dissertation. We thank Colleen Seifert, Rob Goldstone, and three anonymous reviewers for their comments on earlier drafts.

References

Ahn, W. (1998). Why are different features central for natural kinds and artifacts?: The role of causal status in determining feature centrality. *Cognition*, *69*, 135–178.

Ahn, W., & Dennis, M. J. (1997). Dissociation between categorization and similarity judgment: Differential effect of causal status on feature weights. In M. Ramscar, U. Hahn, E. Cambouropoulos, & H. Pain (Eds.), *Proceedings of the interdisciplinary workshop on similarity and categorisation* (pp. 1–7). Edinburgh, Scotland: Department of Artificial Intelligence, University of Edinburgh.

Ahn, W., & Lassaline, M. E. (1995). Causal structure in categorization. In J. D. Moore & J. F. Lehman (Eds.), *Proceedings of the 17th annual conference of the cognitive science society* (pp. 521–526). Mahwah, NJ: Erlbaum.

Atran, S. (1990). *Cognitive foundations of natural history: Towards an anthropology of science*. New York: Cambridge University Press.

Bailenson, J. N., Shum, M., Atran, S., Medin, D. L., & Coley, J. (2002). A birds eye view: triangulating biological categorization and reasoning within and across cultures. *Cognition*, *84*, 1–53.

Barton, M. E., & Komatsu, L. K. (1989). Defining features of natural kinds and artifacts. *Journal of Psycholinguistic Research*, *18*, 433–447.

Carey, S. (1985). *Conceptual change in childhood*. Cambridge, MA: Bradford Books.

Clement, J., & Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. *Cognitive Science*, *15*, 89–132.

Gelman, S. A. (1988). The development of induction with natural kind and artifact categories. *Cognitive Psychology*, *20*, 65–95.

Gentner, D. (1983). Structure mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 155–170.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199–241). New York: Cambridge University Press.

Gentner, D., Medina, J. (1998). Similarity and the development of rules. In S. A. Sloman & L. J. Rips (Eds.), *Similarity and symbols in human thinking* (pp. 177–211). Cambridge, MA: MIT Press.

Goodman, N. (1955/1983). *Fact, fiction, and forecast* (3rd ed.). Cambridge, MA: Harvard University Press.

Heit, E. (1997). Knowledge and concept learning. In K. Lamberts & D. Shanks (Eds.), *Knowledge, concepts, and categories* (pp. 7–41). London: Psychology Press.

Heit, E. (1998). A Bayesian analysis of some forms of inductive reasoning. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition* (pp. 248–274). Oxford: Oxford University Press.

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

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

Hirschfield, L. A., & Gelman, S. A. (Eds.). (1994). *Mapping the mind*. Cambridge: Cambridge University Press.

Howell, D. C. (1997). *Statistical methods for psychology* (4th ed.). USA: Duxbury Press.

Keil, F. C. (1995). The growth of causal understanding of natural kinds. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary approach* (pp. 234–262). New York: Oxford University Press.

Keil, F. C., Smith, C. W., Simons, D. J., & Levin, D. T. (1998). Two dogmas of conceptual empiricism: Implications for hybrid models of the structure of knowledge. In S. A. Sloman & L. J. Rips (Eds.), *Similarity and symbols in human thinking* (pp. 17–49). Cambridge, MA: MIT Press.

Lassaline, M. E. (1996). Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning Memory, and Cognition*, *22*, 754–770.

López, A., Atran, S., Coley, J. D., Medin, D. L., & Smith, E. E. (1997). The tree of life: Universal and cultural features of folkbiological taxonomies and inductions. *Cognitive Psychology*, *32*, 251–295.

Malt, B. C., & Johnson, E. C. (1992). Do artifact concepts have cores? *Journal of Memory and Language*, *31*, 195–217.

Medin, D. L., Coley, J. D., Storms, G., & Hayes, B. K. (2003). A relevance theory of induction. *Psychonomic Bulletin & Review*, *10*, 517–532.

Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology*, *20*, 158–190.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*, 289–316.

Nisbett, R. E., Krantz, D. H., Jepson, C., & Kunda, Z. (1983). The use of statistical heuristics in everyday inductive reasoning. *Psychological Review*, *90*, 339–363.

Osherson, D. N., Smith, E. E., Wilkie, O., López, A., & Shafir, E. (1990). Category-based induction. *Psychological Review*, *97*, 185–200.

Pazzani, M. J. (1991). Influence of prior knowledge on concept acquisition: Experimental and computational results. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 416–432.

Proffitt, J. B., Coley, J. L., & Medin, D. L. (2000). Expertise and category-based induction. *Journal of Experimental Psychology: Learning Memory, and Cognition*, *26*, 811–828.

Quine, W. V. O. (1977). Natural kinds. In S. P. Schwartz (Ed.), *Naming, necessity, and natural kinds* (pp. 155–175). Ithaka, NY: Cornell University Press.

Rips, L. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning, and Verbal Behavior*, *14*, 665–681.

Rozenblit, L., & Keil, F. C. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science*, *26*, 521–562.

Sanjana, N., & Tenenbaum, J. B. (2003). Bayesian models of inductive generalization. In: S. Becker, S. Thrun & K. Obermayer (Eds.), *Advances in Neural Information Processing Systems*, *15*. MIT Press.

Simons, D. J., & Keil, F. C. (1995). An abstract to concrete shift in the development of biological thought: The insides story. *Cognition*, *56*, 129–163.

Sloman, S. A. (1993). Feature based induction. *Cognitive Psychology*, *25*, 231–280.

Sloman, S. A. (1994). When explanations compete: The role of explanatory coherence on judgments of likelihood. *Cognition*, *52*, 1–21.

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

Sloman, S. A., & Lagnado, D. A. (in press). The problem of induction. In: R. Morrison & K. Holyoak (Eds.), Cambridge, *Handbook of Thinking and Reasoning*.

Sloman, S. A., Love, B. C., & Ahn, W. (1998). Feature centrality and conceptual coherence. *Cognitive Science*, *22*, 189–228.

Sloman, S. A., & Malt, B. (2003). Artifacts are not ascribed essences, nor are they treated as belonging to kinds. *Language and Cognitive Processes*, *18*, 563–582.

Sloman, S. A., Malt, B. C., & Fridman, A. (2001). Categorization versus similarity: The case of container names. In U. Hahn & M. J. A. Ramscar (Eds.), *Similarity and categorization*. Oxford: Oxford University Press.

Smith, E. E., Shafir, E., & Osherson, D. (1993). Similarity, plausibility, and judgments of probability. In P. N. Johnson-Laird & E. Shafir (Eds.), *Reasoning and decision making* (pp. 67–96). Cambridge, MA and Oxford, UK: Blackwell.

Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences*, *12*, 435–502.

Wattenmaker, W. D. (1995). Knowledge structures and linear separability: Integrating information in object and social categorization. *Cognitive Psychology*, *28*, 274–328.

Wu, M., & Gentner, D. (1998). Structure in category based induction. In M. A. Gernsbacher & S. J. Derry (Eds.), *Proceedings of the 20th annual conference of the cognitive science society* (pp. 1154–1158). Mahwah, NJ: Erlbaum.