Hampton, J. A. & Cannon, I. (2004). Category-based induction: An effect of conclusion typicality. Memory & Cognition, 32(2), 235 - 243. doi: 10.3758/BF03196855
<http://dx.doi.org/10.3758/BF03196855>

City Research Online

To appear in *Memory & Cognition 2004*

Category-based Induction - an effect of conclusion typicality

James A. Hampton

Iben Cannon

City University, London

Please address correspondence to:

Professor James A. Hampton

Psychology Department, City University

Northampton Square, London EC1V OHB, UK

hampton@city.ac.uk

Running HEAD: Category-Based Induction

Abstract

Category-based induction is the willingness of a thinker to project some newly learned property of one or more classes of objects to another class, on the basis of their shared membership in a common superordinate category. Previous research has established that the perceived strength of arguments of the form "Class A has property P, therefore class B has property P" is influenced by the similarity of A to B, and by the typicality or representativeness of A in a shared category, superordinate to both A and B. (The nature of P is also crucial, but we do not examine it in this study). There is however no prior evidence that the relation between B and the category is influential. Three experiments were designed to test whether the typicality of B in the superordinate category also has an effect on inductive argument strength. Using multiple regression (Experiment 1) and an experimental design (Experiment 3), an effect of conclusion typicality was found such that people are more willing to project properties to more typical conclusions. Experiment 2 ruled out conclusion familiarity as a potential confounding variable. The results are interpreted in the light of current models of category-based induction.

Category-based Induction - an effect of conclusion typicality

In a pioneering study, Rips (1975) posed problems of the following form to students:

> *If all <u>rabbits</u> on a particular island have a new type of contagious disease, then*
>
> *what proportion of <u>mice</u> would be expected to have the disease?*

This type of argument has been termed inductive in that it is presumably through this

form of reasoning that more general hypotheses are generated based on individual

observations. In particular, in contrast to deductive forms of reasoning such as

syllogisms or conditionals, there is no clear way of determining what the correct answer

in such a case may be. It is therefore of great interest to determine the factors within

such arguments that influence people when they come to make a judgment of

argument strength. This psychological question is, of course, largely independent of

the equally interesting epistemological question of the grounds that would <u>actually</u>

justify confidence in such an argument.

Rips (1975) worked with biological categories such as Birds. First, he took pair-

wise similarity judgments for a set of words from the same category, such as a sample

of different bird names, together with the category name "bird" itself. These similarity

judgments were then scaled with multi-dimensional scaling (MDS) to generate a two

dimensional similarity space. Effectively each concept name was placed on a two-

dimensional map, with the constraint that proximity on the map corresponded as

closely as possible to mean rated similarity for any pair of concepts. A second group of

participants were given a set of arguments of the form cited above, each argument

created by taking a different pair of concepts from the set of category members, and

arranging them in a particular order. This new group made judgments of the strength

of the inductive arguments. The analysis then used a regression to try to predict rated strength of an argument in terms of relative positions of the two concepts in the similarity space. For an argument of the form "If all A have a particular property, then what proportion of B will have it?", A is the premise term, B the conclusion term, and let us call C the most specific superordinate category that contains both A and B, (but which is not mentioned explicitly in the argument).   In Rips's study there were then three possible predictors of the argument strength (see Figure 1):

- the distance of the premise A from the category C

- the distance of the premise A from the conclusion B

- the distance of the conclusion B from the category C.

When these three variables were entered into the regression, Rips found that only the first two predictors were needed to account for variations in argument strength. The closer the premise was to the category (i.e. the more typical it was), and the closer the premise was to the conclusion (that is the more similar the premise and conclusion), then the stronger was the argument.   The typicality of the conclusion B (the proximity of B to C) was not a significant predictor of strength.   As a logical corollary of this result, Rips (1975) also showed that there is a premise-conclusion asymmetry effect in regard to typicality. When one of the category members is more typical than the other, then the strength of the argument TYPICAL $\Rightarrow$ ATYPICAL is stronger than the reverse argument. This asymmetry follows automatically from the finding that premise typicality affects argument strength, while conclusion typicality does not, since clearly the argument pair TYPICAL $\Rightarrow$ ATYPICAL will have a more typical premise than will the reversed pair ATYPICAL $\Rightarrow$ TYPICAL. (For the sake of

argument, it is assumed that similarity between the A and B terms is the same measured in each direction – as will necessarily be the case when taking proximity measures from MDS solutions. Evidence for this assumption is provided by Aguilar and Medin, 1999).

Subsequent work by Osherson et al. (1990) laid out evidence for a set of thirteen phenomena relating to category-based induction, using categories at different levels, and exploring arguments with either single or multiple premises. One methodological advance on Rips (1975) was their introduction of an improved format for the task. Asking about disease prevalence may bring to mind all kinds of different properties of the animals such as their diet and habitat which could influence the answer (see, for example, Coley et al. 1999). In order to keep the logical structure of the argument free from any specific effects of stored knowledge or background theory, other than knowledge of the taxonomic hierarchy of creatures categories, Osherson et al. (1990) adopted the following format:

Rabbits use serotonin as a neurotransmitter

therefore

Mice use serotonin as a neurotransmitter

Osherson et al. also tested for premise-conclusion asymmetry, but with equivocal results. For example, when asked to choose directly between the two arguments BAT $\Rightarrow$ MOUSE versus MOUSE $\Rightarrow$ BAT, responses were evenly divided. (Mouse would be a typical mammal, and bat an atypical mammal).   Likewise, when rather than making a forced choice, participants gave ratings of the conditional probability of the conclusion given the premise, there was no significant asymmetry

effect. The only condition where a reliable asymmetry effect <u>was</u> obtained was a condition where (a) the critical forced choice between the two arguments was placed in a context with other filler pairs, and (b) instructions were given that "there is always a difference in how much reason the facts of an argument give to believe its conclusion – however small this difference may be, we would like you to indicate for which argument the facts provide a better reason to believe the conclusion." Having noted this restriction on the generality of the asymmetry effect, one should also note that Osherson et al. (1990) do not state how many pairs of such arguments were tested per experiment, so their results (or lack of them) were quite possibly due to item-specific factors and a small sample size of items. It appears then that the evidence for premise-conclusion asymmetry may be fairly weak. This weakness could either reflect an absence of a premise typicality effect, or the presence of an equally strong conclusion typicality effect.

The lack of an effect of conclusion typicality reported by Rips (1975) has been further confirmed in several developmental studies, some of which also used adult groups as controls (Carey, 1985; Gelman & O'Reilly, 1988). In an extensive review of the literature, Heit (2000) concluded "there have been no reports to date of independent effects of the typicality of the conclusion category as opposed to the premise category". Perhaps the only result that suggests that there may be such an effect comes from a study by Sloman (1998) which took a category name as the <u>premise</u> of the argument. Arguments of the form PLANTS $\Rightarrow$ MOSSES should be perfectly convincing, provided that the participant agrees with the premise (such as *all plants have bryophytes*), and agrees that all mosses are indeed plants. However

Sloman showed that people still felt more confident in arguments that led to more typical conclusions – for example preferring PLANTS $\Rightarrow$ FLOWERS to PLANTS $\Rightarrow$ MOSSES. Sloman's result therefore sits rather uneasily with the apparent lack of a conclusion typicality effect when the premise is a category member rather than a category name. We therefore decided to investigate the issue of conclusion typicality further.

The aim of the current research was to set up an experimental test of the existence of a conclusion typicality effect. To this end, we considered carefully the best way in which to select materials for the test. It was not sufficient simply to select a premise and then two conclusions of differing typicality, all from the same category, since if typical items have a tendency to be more similar to one another than do typical to atypical items, an effect of conclusion typicality could simply reflect a difference in premise-conclusion similarity.

Sets of category members were therefore needed that were well balanced for premise typicality and premise-conclusion similarity but would vary maximally on conclusion typicality. One of the problems with earlier research is that if the premise concept is a underline{typical} member of the category, it follows that there will be a close correspondence between premise-conclusion similarity and conclusion typicality. As the points A and C in Figure 1 draw closer together, so the lengths of AB and CB are more and more constrained to be very similar. Given that premise-conclusion similarity (AB) is known to have the strongest effect on inductive argument strength, it is unlikely that there would be an independently detectable effect of conclusion typicality simply because of the multi-collinearity arising from using a very typical premise. Accordingly

we chose premises of intermediate typicality, which would enable us to choose conclusions of varied typicality but matched similarity to the premise.

Experiment 1

The first experiment adopted Rips' (1975) regression method. Independent measures were obtained of premise typicality, conclusion typicality, premise-conclusion similarity and the inductive strength of the argument from premise to conclusion. Inductive strength was then regressed on the other three variables.

Method

Participants. Forty-eight students at City University, London participated on a voluntary basis to provide the initial set of normative ratings, 25 for typicality and 23 for pair-wise similarities. An additional 19 students volunteered to provide a second set of similarity ratings for new pairings of the original items. The final balanced set of materials were then given to a third group of 36 students who rated the inductive strength of the premise-conclusion arguments, participating in return for course credit. There was no overlap between any of the groups.

Materials. Materials were generated in triplets within three biological categories – birds, mammals and insects. Pre-testing involved first obtaining ratings of typicality for a larger set of items in their respective categories, and obtaining similarity ratings for paired premise and conclusion concepts in each category. Order within the list was randomized within categories, except that for similarity ratings, pairs with the same premise (e.g. horse-cow and horse-bison) were kept maximally apart in the list. Order of premise and conclusion concepts within a pair was counter-balanced across materials. Each category was presented on a separate page of a three page booklet, and page order was balanced across participants for both tasks.

Analysis of the first set of ratings revealed that the initial construction of triplets had not succeeded in balancing the similarity of premise-conclusion pairs within each triplet. The same items were therefore re-paired into new triplets, and a second set of similarity ratings were obtained. It was then possible to construct a total of 22 triplets, 8 for mammals and 7 each for birds and insects.

Procedure. For the ratings of typicality and similarity, subjects were tested in a classroom setting, and completed the booklets in their own time without conferring. The typicality booklet contained 3 pages, one for each category, and used a 1-7 numbered scale, with 1 meaning very typical and 7 very atypical. Instructions emphasized the difference between familiarity and typicality (see Hampton and Gardiner, 1983). Similarity ratings were also blocked within category, in order to ensure that the basis used for similarity was relevant to the task. The category name was printed at the top of the page for both tasks. Similarity instructions asked participants to "rate each pair according to how similar or dissimilar the two instances appear to you within the category to which the pair belongs". A 1-7 numbered scale was again used, with 1 representing a very similar pair, and 7 a very dissimilar pair. In both tasks, if a participant did not know any of the items, they were instructed to underline the item and leave the scale blank. Seventeen participants (25%) were excluded from the analysis because they either marked more than 20% of items as unknown, or they failed to engage with the task (for example marking every item as very typical).

From an initial set of 30 a final set of 22 triplets were selected after re-pairing and re-measuring similarities. Within each triplet of premise A, typical conclusion B_{typ} and atypical conclusion B_{atyp}, the similarity of A to each conclusion was matched, and

9

the conclusions differed maximally in their typicality in the category. All items in the

final set were known to at least 75% of participants. (Some degree of unfamiliarity for

the atypical items was necessary in order to generate triplets that were otherwise

matched on similarity. Experiment 2 returns to the question of whether familiarity may

be confounding the effects of typicality.) A list of the triplets used, together with their

mean ratings is shown in Table 1.

INSERT TABLE 1 ABOUT HERE

Finally the 22 triplets were used to construct 44 inductive arguments, taking the

premise with either of the two conclusion terms. As in Osherson et al. (1990) 44

different biologically plausible blank predicates were used such as "needs Vitamin K for

liver function". Each argument was quantified with the word "All" in order to stress that

the premise and conclusion applied to whole kinds rather than to individuals. One

argument from each triplet was arranged in a random order in one booklet, while the

other was placed in a second booklet. Order was random, but blocked by category.

Each booklet contained equal numbers of typical and atypical arguments. For each

category two additional filler items were included in the first and sixth positions in the

list, in order to disguise the design. Two additional booklets were constructed by

reversing the order of arguments within each category list (while keeping the filler items

in the same place). Order of the three categories in the booklets was also

counterbalanced across participants. The booklets were distributed to students in a

classroom setting and were completed without time constraint or conferring.

Instructions were to consider the first statement as being a true fact, and then to judge

the probability of the conclusion being true in the light of the evidence provided by the

fact. Ratings were made on a 10 point scale, with 1 representing no faith in the

argument being true (very unlikely), and 10 representing strong belief that the argument could be true (very likely). As previously, participants were told to underline any items they did not know, and to leave the scale blank.

Results

Reliability of ratings. Across the 22 triplets, mean rated similarity for Premise-Typical conclusion and for Premise-Atypical conclusion pairs was identical at 4.43 (sd = .89 and 1.04 respectively). The average standard error for individual item pair similarity ratings was 0.37. Mean rated typicality was 2.30 (sd = .44) for the premise items, 1.72 (sd = .36) for the typical conclusion items, and 3.41 (sd = .61) for the atypical conclusion items. Across all premises the variance in typicality was 0.19, whereas across all conclusions the variance in typicality was 0.98. There was therefore 5 times as much variance in the conclusion typicalities as in the premise typicalities. As had been intended, the materials maximized the manipulation of conclusion typicality. A consequence of this procedure however was to produce greatly truncated variance in premise typicality. Reliability of the premise-conclusion similarity judgments across the 44 arguments was .84 (Spearman-Brown). Across the 66 typicality judgments, reliability was 0.86. This broke down into .89 for the 44 conclusion typicalities but only .52 for the 22 premise typicalities (because of the truncated variance.)

Mean inductive strength was calculated for each of the 44 arguments, and is shown in Table 1. Thirty-eight data points (less than 5%) were treated as missing because of unfamiliarity or failure to complete a rating. No participants or triplets needed to be omitted because of undue levels of missing data. Corrected split-half reliability for inductive strength was 0.82.

Regression analysis. All 44 arguments were entered into a multiple regression analysis to predict mean rated argument strength on the basis of the three predictors: premise-conclusion similarity, premise typicality and conclusion typicality. (Because of the direction of scoring of the scales, all predictors were expected to enter with negative coefficients.) As expected the regression showed a strong effect of premise-conclusion similarity (beta = -.733, p<.001). There was no reliable premise typicality effect (beta = .112, p>.10), probably because of the deliberately reduced variance and consequently low reliability for this measure. There was however a significant effect of conclusion typicality (beta = -.246, p = .029). Adjusted R square was .50. When dummy variables were included to take account of the factor of Category, the conclusion typicality effect was somewhat smaller, but remained significant (beta = .195, p =.047). Further tests showed that after entering premise-conclusion similarity as a first step, conclusion typicality would enter next with a significant beta (p = .035), whereas premise typicality would not (p = .46).

Lorch and Myers (1990) suggested that a more valid way to analyze regression data when using items as the random variable is to perform separate regression analyses on each individual participant, and then to combine the results. The set of ratings given by each of the 18 participants who rated inductive strength were taken individually, and 18 regression analyses were run to predict their individual ratings on the basis of the three predictors. Mean multiple R was 0.40, and mean R square was 0.18. The results confirmed the strong effect of premise-conclusion similarity (mean beta = -0.34, one sample t(17) = -8.00, p < .001), and the significant effect of conclusion typicality (mean beta = -0.12, one sample t(17) = -3.27, p < .005.) As before there was no significant effect of premise typicality. The beta coefficient was in

the predicted direction for 17 out of 18 participants for the similarity effect and for 14 out of 18 participants for the typicality conclusion effect, 5 of which were individually significant at .05, 1-tailed.

Discussion

The aim of Experiment 1 was to select a set of materials in such a way as to optimize the chances of obtaining an effect of conclusion typicality on the judged strength of a simple inductive argument. The normal correlation between premise-conclusion similarity and conclusion typicality was prevented from confounding the test of the hypothesis by deliberately choosing premises of intermediate typicality. With this correlation held artificially low (r = -.07, p >.5), variations in conclusion typicality were indeed found to predict inductive strength. More specifically, the higher the typicality of the conclusion concept in an inductive argument, the more strongly was the argument rated.

Experiment 2

One issue that remains a potential problem for interpreting the results of Experiment 1 concerns the differences in familiarity of the typical and atypical materials within the triplets. In order to construct well balanced triplets, it proved necessary to allow atypical conclusion concepts to be *prima facie* less familiar than typical conclusion concepts. Not all atypical items were unfamiliar – for example rat, and squirrel are probably more familiar than hyena or grizzly bear to a British sample. However should it turn out that familiarity of conclusion concepts predicts the strength of an inductive inference, our interpretation of the conclusion typicality effect is at risk. Indeed, some unpublished data from Collister and B.Tversky (personal communication)

suggests that people may prefer arguments with less familiar premises and more familiar conclusion terms.

In Experiment 2, therefore, a further sample of students from the same population were asked to provide ratings of familiarity for all 66 concepts used in Experiment 1. Mean familiarity ratings were calculated and correlated with the predictors and dependent variable from that Experiment.

Method.

Participants. 18 students at City University, London participated for course credit. None took part in any of the other experiments.

Procedure. Booklets were created in which the concepts from the triplets were listed in alphabetic order, blocked by category. A five point scale was provided, and participants rated the familiarity of each concept from 1 = very unfamiliar to 5 = very familiar.

Results

Mean familiarity ratings were calculated for each of the 66 concepts. Split half reliability was estimated at 0.92.   Mean familiarity across concepts ranged from 4.6 and above for highly familiar concepts such as dog, duck and ant, down to 1.9 for the less familiar bison. Mean familiarity was 3.8 (sd = 0.6) for the premises, 4.0 (sd = 0.5) for the typical conclusion concepts, and 3.2 (sd = 0.7) for the atypical conclusions. Within the triplets, 16 of the 22 had a more familiar typical than atypical conclusion term. However when compared to the results of Experiment 1, familiarity had no correlation with inductive strength of arguments.   Familiarity of the premise term correlated with inductive strength with r = 0.03, and familiarity of the conclusion term correlated with strength with r of -.07.   Furthermore, the familiarity variables failed to

enter any of the regression models examined in Experiment 1 with significant

coefficients.

There was therefore no evidence that the observed significant effect of

conclusion typicality was owing to differences in familiarity.

Experiment 3

To provide a further test of generality for the conclusion typicality effect,

Experiment 3 adopted a quasi-experimental design, so that generalization across both

items and subjects could be tested. Instead of making ratings of strength for individual

arguments, a forced-choice procedure was used in which a pair of arguments were

presented, and participants were asked to select the stronger. Osherson et al (1990)

used a similar procedure. The forced-choice procedure helps to focus attention on the

key factor of the conclusion item, since pairs of arguments had the same premise and

the same predicate but differed in their conclusions.

Method.

Participants.   Forty-eight students at City University London volunteered to

participate, either with no incentive or for course credit. None of the participants had

taken part in any of the earlier experiments.

Materials. The same set of 22 triplets were used as in Experiment 1. They

were organized into pairs with a common premise and two different conclusions, one

typical and one atypical. The same filler items were used within each category list to

disguise the manipulation of conclusion typicality.

Procedure. Booklets contained 28 argument pairs, comprising the 22 critical

test pairs blocked by category, and the 6 filler pairs. Participants completed the

booklets in a group setting, without time constraint. An example of an argument pair was as follows:

> *a) All Vultures have sesamoid bones, therefore all Robins have sesamoid bones*
>
> *b) All Vultures have sesamoid bones, therefore all Ostriches have sesamoid bones*

One argument was presented on the left of the page, and the other argument opposite it on the right of the page. Across booklets, typical and atypical arguments occurred equally often on the left or the right. Participants had to indicate which of the two they considered a stronger argument - which they felt was more likely to be true. Order of critical pairs was randomized within categories, with the filler arguments in positions 1 and 6. A second set of booklets used the reverse random order for critical pairs. Order of categories within the booklets was also counterbalanced, leading to 12 different booklets. Four participants completed each booklet. As before if a concept was unknown, participants were asked to underline it and move on to the next pair.

Results. Less than 6% of the data were omitted because of unknown items, and the data from all participants were used. The argument with the more typical conclusion was selected 58% of the time, where chance would be 50%. This result was significant across subjects, $t(47) = 3.89$, $p<.001$, two-tailed. Thirty-four out of 48 participants showed the effect. Similar t-tests were significant for each of the three categories individually, with p values between .05 and .003. Across items, 17 out of 22 paired arguments had a greater proportion of participants selecting the more typical conclusion, ($p<.01$ on a sign test). This result was also significant on a one-sample t-test across items ($t(21) = 2.37$, $p < .027$, two-tailed).   There was no significant correlation between the proportion of participants choosing the more typical conclusion

for an item, and the difference in the rated familiarity from Experiment 2 of the two

conclusions (r = 0.1).   There was therefore again no evidence that familiarity plays a

role in judging inductive strength.

Discussion

Experiment 3 confirmed the finding from Experiment 1 that the typicality of a

conclusion concept can affect the perceived strength of an inductive argument. When

given a choice between two arguments, one with a typical conclusion and one with an

atypical conclusion (with respect to the common superordinate category), participants

showed a significant tendency to prefer the typical one. The effect was relatively small

(as in Experiment 1) but was consistent, with 71% of participants and 77% of items

showing the predicted pattern.

General Discussion

The results presented here provide a clear demonstration of a robust effect of

the typicality of the conclusion term in category-based inductions. Before discussing

our own account of the data, we next discount three alternative accounts of the data,

and consider its implications for Sloman's (1993) and Osherson et al.'s (1990) models.

Choice of appropriate superordinate

If Osherson et al's model is correct, it could account for our data by supposing

that premises and typical conclusions were members of more narrowly defined

superordinates than the corresponding premises and atypical conclusions. For

example [duck + canary] may cue retrieval of Flying Birds, whereas [duck + ostrich]

would cue retrieval of Birds. Similarly [horse + cow] might cue Farm Animal, whereas

[horse + bison] would cue just Animal. The more narrowly defined the superordinate

category, then the more likely it is that the premise term will provide fuller coverage of the category – and hence the stronger will be the argument.

Basis of similarity

A related notion would be that the use of the biological blank predicate (e.g. uses serotonin as a neurotransmitter), could have triggered a different similarity metric from that considered when people just rated similarity alone (see Heit and Rubinstein, 1994). For example about half of the triples have an atypical conclusion term that comes from a different geographical region from the other two terms (e.g. Eagle – Crow – Cockatoo). Perhaps people might have used geographical region as an indicator of deeper biological similarity for the inductive judgment, but ignored it when making general similarity judgments in the pretest.

While these are both plausible explanations, there was no support for them in the data. To illustrate this claim, Table 1 lists the triples within each category in descending order of their effect size (argument strength for typical conclusion minus argument strength for atypical conclusion) in Experiment 1. The strongest effects of conclusion typicality in Experiment 1 were in triples where neither of these accounts apply (e.g. Centipede – Beetle – Silverfish, Koala – Tiger – Guinea pig). In addition the triple that went most strongly against the hypothesis was one to which both arguments do apply (Duck – Canary – Ostrich). Ostrich is from a different evolutionary niche and is not in the Flying Bird category, but (contrary to the overall trend) the argument from Duck to Ostrich was rated as stronger than from Duck to Canary. When coded as binary variables (using the first author's intuitions), neither of these hypotheses entered into the regressions for Experiment 1 with significant coefficients, whereas the effect of Conclusion typicality remained strong and significant. Of course,

a fully adequate test of these accounts would require further work, predicting argument strength from two new data sets. One would be data on the closest common superordinate that people associate with a pair of items, and the other would be similarity judgments taken from a biological perspective.

<u>Non-specific effects of typicality</u>

Our results depend critically on the preference of participants for arguments that have more typical conclusion terms.   An interesting third possibility is that there may be some overall non-specific bias towards accepting the truth of statements about typical concepts – regardless of any other considerations. The test for this would be to obtain estimates of the truth of the argument conclusions in the absence of the premises. It is implausible to expect such a bias, given that it would be reflecting typicality in relation to a category which is not mentioned in the task. Furthermore, if there were a bias, it would more plausibly relate to familiarity, and work the other way. The more ignorant we are of an item, then perhaps the more willing we are to accept that any given predicate should be true of it.   We therefore consider this account unlikely.

Given the demonstration that the typicality of the conclusion concept may affect inductive strength, our discussion now turns to how two particular current models may or may not be able to accommodate such an effect.

<u>Sloman's feature-based model</u>

The first model to be considered was proposed by Sloman (1993). In his feature-based model, when premises are introduced, the features possessed by those premise concepts are associated to the predicate. The strength of the argument associating the predicate to the conclusion is then determined by the feature overlap

between the conclusion term and the set of activated features, divided by a measure of the "magnitude" of the conclusion term. More specifically, argument strength will increase with premise-conclusion similarity, but will decrease the more features are possessed by the conclusion term.   Sloman's model is entirely feature based, and requires no consideration of set inclusion relations. It is therefore well suited to the approach proposed above. More radically, however, it does not require activation of any superordinate category in the case of single premise arguments such as those considered in this paper. All typicality effects are handled either by feature overlap, or by the degree of richness of the conclusion concept term. For example to account for the possible asymmetry between TYPICAL$\Rightarrow$ATYPICAL and ATYPICAL$\Rightarrow$TYPICAL arguments it is proposed that more typical conclusions will tend to have richer feature representations, and hence will have weaker argument strengths. Sloman (1993) discussed evidence based on six pairs of items for which typicality and richness of representation could be disconfounded. He concluded that the number of features possessed by the conclusion term is the critical factor in making weaker arguments, (see also Sloman and Wisnieski, 1992).

If this were the correct account of the asymmetry effect, and if typical concepts are distinguished by having richer feature representations, then holding the premise constant (as was done here) and varying typicality of the conclusion should show an <u>inverse</u> conclusion typicality effect; more typical conclusions should show weaker strength ratings, to the extent that they have richer feature representations. To explain the advantage for arguments with typical conclusions, Sloman would have to argue (in contradiction to his account of the asymmetry effect) that the typical conclusion items

used in these studies have <u>less rich</u> feature representations. The reader can check the

plausibility of this idea in Table 1, where the triples are listed in descending order of

effect size. To take the two triples with the strongest effect, it would imply that beetle

has a less rich representation than silverfish, and that tiger has fewer features

represented than guinea pig.

Note also that Experiment 2 failed to show any effects of familiarity of the

premise and conclusion concepts on inductive argument strength, whereas one might

reasonably expect concepts with richer feature representations to be judged as more

familiar.

<u>Osherson et al.</u>

The second model is Osherson et al.'s 1990 coverage model which suggests

that inductive strength in single premise arguments is determined by two things: the

similarity between premise and conclusion terms, and the average similarity between

the premise and other members of the lowest superordinate category that includes

both premise and conclusion. Their model therefore predicts that if premise typicality

and premise-conclusion similarity are held constant there will be no conclusion

typicality effect. Although their model is clearly successful at accounting for a wide

range of phenomena, nonetheless our data suggest that it is in need of modification.

We suggest, like Osherson et al, that there are two "routes" to argument

strength: one direct route through the similarity between the mental representations of

the premise and conclusion concepts, and a second indirect route via category

membership. The first route, via similarity, is presumably a very general route,

applying to any pair of items regardless of their category membership. It is the second

route involving membership in the common superordinate category that needs revision.

The indirect route involves two stages, both probably driven by a belief that members of biological categories share hidden biological properties. One simple modification of the Osherson et al. model would propose that when the superordinate category becomes implicated in the reasoning process, it is a prototype of the superordinate that is involved rather than the superordinate category considered as an equivalence set. Osherson et al.'s model proposes that once the premise or premises have activated the superordinate, then the resulting argument strength applies equally to all category members. The reasoning according to their model might thus be explicated as involving the following 2 stages:

1. *All vultures have sesamoid bones, therefore all birds have sesamoid bones,*

2. *All birds have sesamoid bones, all ostriches are birds, therefore all ostriches have sesamoid bones,*

The first step in the reasoning will have variable strength depending on premise typicality, but the second step is taken in their model to be always perfectly strong on logical grounds. (Step 2 has the classical form of a syllogism). There is evidence however from Sloman (1998) that people do not always respect the logic of step 2 (see also Hampton (1982) for failures in the transitivity of category superordination). As described in the introduction, Sloman showed that people still preferred arguments from a superordinate to a typical subset over arguments from the same superordinate to an atypical subset.

According to the proposed modification to the model, the reasoning would now proceed as follows:

1. *All vultures have sesamoid bones, therefore prototypical birds have sesamoid bones,*

2. *Prototypical birds have sesamoid bones, therefore all ostriches have sesamoid bones.*

Unlike the Osherson et al model, the middle term of the category-based induction in our model is the prototype for the category, and not the category as a whole class. The category is being represented as an intensional concept, and not as an extensional set of exemplars. (See Tversky and Kahneman (1982) and Hampton (1987) for additional evidence that through representing concepts in terms of their properties people are prone to ignore logical constraints on set membership.) By using intensional representation, the quantification of the middle argument "birds have sesamoid bones" is left vague, and so step 2 cannot be given maximum strength on logical grounds.

Step 1 is of variable strength and will depend as before on premise typicality. Step 2 however will also be of variable strength and will now depend on the conclusion typicality. Moreover these two steps rely on the same general principle as the direct route for assessing argument strength, namely the similarity between two concepts.

Osherson et al.'s model applies to a much wider range of inductive problems than those considered here. One phenomenon that would not be explained by a simple activation account like that offered here is the diversity effect. Arguments using two premise terms from the same category as the conclusion are considered less strong if the two premise terms are similar than if they are diverse. Activation of the prototype concept for the superordinate according to our proposal would therefore have to be greater in the case of diverse premises. Possible mechanisms to achieve this would to base prototype activation on a sum of similarity to category exemplars (as

in Osherson et al.'s model), or alternatively to base prototype activation on a measure of feature overlap between the disjunction of the features of the two premises and the prototype as in Sloman's (1993) model.

Asymmetry

As it stands, the proposed modification to the Osherson two-route model would predict that the final strength of the argument due to the indirect route would be the product of the strengths of the two links. However that would imply symmetry in the strength of arguments when the premise and conclusion terms are reversed. Some alternative account therefore remains to be given of the premise-conclusion asymmetry effect, assuming it to be reliable. One can look for an answer in two places – in the direct route and in the indirect route, either or both of which may introduce asymmetry.

Asymmetry in the direct route would arise if the underlying similarity between the concepts is itself asymmetric. There is evidence that similarity is greater for a pair of items when the more typical or salient term is the target. For example, Tversky (1977) reported that the similarity (as rated by US students) of Cuba to the USA is rated as greater than the similarity of the USA to Cuba. If premise-conclusion similarity were asymmetrical in a similar way, (and computed as the similarity of the conclusion term to the premise term), then that could account for the effect of reversing premise and conclusion within an argument. However, as noted in the introduction, Aguilar and Medin (1999) reported a failure to replicate asymmetries in similarity ratings.

A second locus for asymmetry would be in the indirect route. Note that to be consistent with asymmetry, our result implies that there must be both a premise and a conclusion typicality effect, but that the premise typicality effect be greater than the conclusion typicality effect. One way to achieve this would be for the generalization

gradient around a category prototype such as bird to be broader than that around a

subclass concept such as robin or ostrich. Having narrow generalization gradients,

individual concepts will only activate the general category name if they are close to it.

Typical premises such as robin will activate bird strongly, but atypical premises will not.

The effect of premise typicality is therefore very pronounced. Once activated, the

category name generalizes more broadly across all members of the category.

Although it will still activate typical conclusions more strongly than atypical conclusions,

the gradient of the typicality effect will be much shallower.

<u>Conclusion</u>

This paper has set out first to demonstrate the existence of a phenomenon

previously claimed not to exist – namely that the typicality of a conclusion category

affects the judgement of strength in category-based induction. The phenomenon was

demonstrated with two different procedures across three biological categories, and an

explanation in terms of familiarity was discounted. Finally, the implications of the result

were discussed. First they appear to directly contradict the account given by Sloman's

(1993) feature-based model for typicality effects, since (unless typical items should turn

out to have fewer features) the explanation he provides for the typicality asymmetry

effect predicts that less typical conclusions will have stronger argument strength.

Second, a modification of the Osherson et al. (1990) model was proposed in which the

strength of an argument is propagated via the superordinate category, but through

activation of the category prototype rather than the set of category members.

References

Aguilar, C.M., & Medin, D.L. (1999). Asymmetries of comparison. Psychonomic Bulletin & Review, 6, 328-337.

Carey,S. (1985). Conceptual Change in Childhood. Cambridge: MA: MIT Press.

Coley,J.D., Medin,D.L., Proffitt,J.B., Lynch,E., & Atran,S. (1999). Inductive reasoning in folkbiological thought. In D. L. Medin & S. Atran (Eds.), Folkbiology. (pp. 205-232). Cambridge, MA:

Gelman,S.A., & O'Reilly,A.W. (1988). Children's Inductive Inferences within Superordinate Categories: The Role of Language and Category Structure. Child Development, 59, 876-887.

Hampton,J.A. (1982). A demonstration of intransitivity in natural categories. Cognition, 12, 151-164.

Hampton,J.A. (1987). Inheritance of attributes in natural concept conjunctions. Memory & Cognition, 15, 55-71.

Hampton,J.A., & Gardiner,M.M. (1983). Measures of Internal Category Structure: a correlational analysis of normative data. British Journal of Psychology, 74, 491-516.

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.

Lorch,R.F., Jr., & Myers,J.L. (1990). Regression analyses of repeated measures data

in cognitive research. Journal of Experimental Psychology: Learning, Memory,

and Cognition, 16, 149-157.

Osherson,D.N., Smith,E.E., Wilkie,O., Lopez,A., & Shafir,E. (1990). Category-Based

Induction. Psychological Review, 97, 185-200.

Rips,L.J. (1975). Inductive Judgments about Natural Categories. Journal of Verbal

Learning and Verbal Behavior, 14, 665-681.

Sloman,S.A. (1993). Feature-based Induction. Cognitive Psychology, 25, 231-280.

Sloman,S.A., & Wisniewski,E. (1992). Extending the domain of a feature-based model

of property induction. In Proceedings of the Fourteenth Annual Conference of the

Cognitive Science Society, Hillsdale NJ: Erlbaum.

Sloman,S.A. (1998). Categorical inference is not a tree: The myth of inheritance

hierarchies. Cognitive Psychology, 35, 1-33.

Tversky,A. (1977). Features of Similarity. Psychological Review, 84, 327-352.

Tversky,A., & Kahneman,D. (1982). Judgments of and by representativeness. In

D.Kahneman, P.Slovic, & A.Tversky (Eds.), Judgment under uncertainty:

Heuristics and biases. (pp. 84-98). Cambridge: Cambridge University Press

Table 1. Triplets used in both experiments, together with mean ratings for premise typicality, similarity between the premise and each conclusion, and typicality and familiarity (from Experiment 2) for the conclusion items. Also shown is argument strength (Experiment 1).

Category = Mammals

Premise	Mean Premise Typicality	Mean Premise Familiarity	Typical Conclusion / Atypical Conclusion	Mean P-C Similarity	Mean Conclusion Typicality	Mean Conclusion Familiarity	Mean Argument Strength
Koala	2.44	3.39	Tiger	5.71	1.81	4.28	4.28
			Guinea Pig	5.47	3.13	4.28	2.78
Dog	1.25	4.89	Fox	2.53	1.81	4.61	6.00
			Coyote	2.00	3.87	2.78	4.83
Whale	2.56	4.00	Grizzly Bear	6.06	2.19	3.56	3.78
			Bat	6.31	3.56	3.94	2.78
Horse	1.75	4.72	Cow	3.76	1.44	4.67	5.89
			Bison	3.88	2.69	1.89	5.25
Boar	2.80	2.67	Deer	4.71	1.81	4.06	4.72
			Walrus	5.47	3.71	2.83	4.64
Zebra	2.31	3.89	Hippopotamus	5.59	2.13	3.78	3.28
			Squirrel	5.71	2.94	4.50	3.28
Hare	2.56	3.56	Goat	5.00	1.81	4.39	3.39
			Rat	4.35	3.06	4.50	4.06
Wolf	2.13	3.83	Hyena	3.00	2.63	3.67	5.89
			Dingo	2.65	4.00	2.33	6.56

Table 1 (continued)

Category = Birds

Premise	Mean Premise Typicality	Mean Premise Familiarity	Typical Conclusion Atypical Conclusion	Mean P-C Similarity	Mean Conclusion Typicality	Mean Conclusion Familiarity	Mean Argument Strength
Vulture	2.40	3.11	Sparrow	4.94	1.06	4.28	6.06
			Quail	4.69	3.54	2.72	4.62
Pelican	3.00	3.50	Parrot	3.18	1.38	4.17	6.47
			Toucan	3.13	3.42	2.56	5.19
Heron	2.31	2.94	Pheasant	4.50	2.19	3.61	5.07
			Emu	4.07	4.14	3.06	4.56
Eagle	1.31	3.94	Crow	3.82	1.19	4.17	5.28
			Cockatoo	5.00	2.38	3.33	4.89
Stork	2.93	3.33	Falcon	3.44	1.94	3.06	6.12
			Goose	3.19	3.00	3.94	6.28
Chicken	2.38	4.61	Magpie	4.81	1.57	3.61	5.17
			Flamingo	4.82	3.00	3.72	5.67
Duck	2.00	4.56	Canary	4.50	1.44	3.61	4.89
			Ostrich	4.12	3.33	3.50	6.83

Table 1 (continued)

Category = Insects

Premise	Mean Premise Typicality	Mean Premise Familiarity	Typical Conclusion Atypical Conclusion	Mean P-C Similarity	Mean Conclusion Typicality	Mean Conclusion Familiarity	Mean Argument Strength
Centipede	2.47	3.61	Beetle	4.35	1.73	4.33	5.81
			Silverfish	4.57	5.08	2.94	3.88
Butterfly	2.69	4.50	Ant	5.35	1.44	4.94	4.22
			Tick	5.50	3.29	2.89	3.35
Daddy-long-legs	2.19	4.06	Cockroach	4.41	1.63	4.06	4.56
			Praying Mantis	3.85	4.33	2.11	4.00
Dragonfly	2.38	3.72	Flea	4.53	1.94	3.39	3.83
			Maggot	5.06	3.53	3.28	3.56
Ladybird	2.13	4.28	Spider	4.65	1.50	4.78	3.94
			Termite	4.44	2.88	3.00	4.06
Wasp	2.06	4.39	Cricket	4.71	1.75	3.67	3.81
			Scorpion	4.94	3.19	3.50	4.11
Locust	2.53	3.06	Mosquito	3.94	1.50	4.11	4.39
			Earwig	4.23	3.00	2.94	5.14

Figure Captions

Figure 1

Representation of a category prototype C, with a typical instance A and an atypical instance B. As the distance AC gets shorter, so the correlation between AB and BC, as B takes different positions, becomes higher. AB represents premise-conclusion similarity, whereas AC and BC represent the premise and conclusion typicalities respectively.