The Effects of Domain and Type of Knowledge on Category-Based Inductive Reasoning

Aimée Kay Crisp-Bright (a.k.crisp@dur.ac.uk)
Department of Psychology, Science Site, South Road
Durham, DH1 3LE, UK

Aidan Feeney (a.feeney@qub.ac.uk)
School of Psychology, Queen's University Belfast
Belfast, BT7 1NN, UK

Abstract

Accounts of category-based inductive reasoning can be distinguished by the emphasis they place on structured versus unstructured knowledge. In addition, it has been claimed that certain domains of structured knowledge are more available than others. Using a speeded task paradigm, participants rated the strength of inductive arguments in which the categories were either strongly or weakly associated and shared a taxonomic or causal relation.. Strongly associated categories received higher inductive strength ratings than weakly associated category pairs, regardless of the domain by which the categories were related. Strength of association was highly predictive of inductive strength ratings, but more additional variance was accounted for by beliefs about taxonomic and causal relations when people were not under time pressure. This suggests that, regardless of knowledge domain, maximizing inductive potency relies on the use of both structured and unstructured knowledge, depending on available mental resources.

Keywords: Category-Based Induction; Knowledge; Categorical Inferences; Reasoning.

Knowledge and Category-Based Induction

Category-based generalizations cover a class of inferences in which an object's category membership supports people's inferences about properties shared with other category members. For example, classifying an animal as a rabbit allows us to infer that it probably lives in a burrow. Furthermore, if we observe that the animal we have classified as a rabbit eats carrots, we are likely to infer that other rabbits and, perhaps hares, also eat carrots.

In order to understand what determines the likelihood that a property will be generalized from a known to a novel instance, we need to identify which aspects of our background knowledge are central to the induction process. Whereas some approaches view category-based induction as driven solely by associative or unstructured knowledge, such as featural overlap (Sloman, 1993), perceptual similarity (Sloutsky & Fisher, 2004) or semantic associations (Rogers & McClelland, 2004), apparently contradictory approaches place theory-based or structured knowledge at the centre of the inductive process, such as knowledge about stable category-hierarchies (Osherson, et al., 1990) and causal relations between categories (Kemp

&Tenenbaum, 2009). These contrasting types of knowledge in turn possess unique processing characteristics which differentially affect the reasoning output.

Unstructured Knowledge and Induction

Unstructured knowledge cannot be described by a higher order structure, abstract interrelationships or theories. It can include relations between entities based on contiguity, co-occurrence, similarity or associations. Several studies suggest that early category formation and induction is driven by the statistical properties inherent in the environment, such as co-occurrence and statistical distribution of perceptual features. For example, Sloutsky and Fisher's (2004) model of Similarity, Induction and Categorization (SINC) assumes that children perform categorization and inductive reasoning on the basis of perceptual similarity, in which the category label is simply treated as another feature contributing to increased similarity between different instances. These researchers also claim that there is only a gradual and developmentally late transition from exclusive reliance on similarity to the use of category membership as a basis for induction. This transition is largely seen as the product of explicit instruction and learning about general characteristics of categories (Fisher & Sloutsky, 2005).

Some proponents of associative approaches to category-based induction advocate that adult categorization and induction is also heavily influenced by similarity (Sloman, 1993) and associations in semantic memory (Rogers & McClelland, 2004). For example, Sloman's (1993) feature-based model explains generalizations purely in associative terms as the degree to which the presentation of the premise instances activates overlapping features of the conclusion instance. Arguments in which premise and conclusion categories share more features are stronger than arguments with little featural overlap between premise and conclusion. Consequently, there is no need to assume a stable category hierarchy. Sloman (1998) does not preclude the possibility that assessment of similarity can at times reflect a more effortful process which draws on knowledge about stable category hierarchies. However, he does suggest that the default mode of category-based induction reflects a predominantly intuitive thought process, requiring no processing effort or reference to class inclusion relations,

especially when people lack relevant knowledge, are under time pressure or have not been explicitly instructed to carefully consider their responses.

Structured Knowledge in Induction

An opposing approach to explaining inductive reasoning focuses on the influence of structured knowledge. The justification for assuming that structured knowledge can play an important role in category-based induction arises from several reasoning phenomena that cannot be explained exclusively by the use of unstructured or associative knowledge.

Osherson et al's (1990) Similarity-Coverage Model posits knowledge about stable taxonomic structure as an important source of information that people rely on when evaluating categorical arguments. Inductive evaluations reflect the weighted sum of two primary parameters, similarity and coverage. Similarity refers to the maximum average similarity between the premise and conclusion categories. Coverage refers to the degree to which the premise categories cover the featural space of the inclusive superordinate category and thus, calculation of coverage requires structured knowledge in the form of a stable hierarchy of categories. The coverage component of the model gives rise to the diversity effect, whereby dissimilar premise categories act as stronger evidence than similar premise categories. Although this phenomenon can be explained by Sloman's model, the developmental trajectory of the diversity effect (Lopez, Gelman, Gutheil & Smith, 1992) is more compatible with the assumption that people draw on structured knowledge about stable category hierarchies. Similarly, if sensitivity to diversity was based exclusively on unstructured associative knowledge, it would not be related to general cognitive ability (Feeney, 2007).

Approaches emphasizing the importance of unstructured knowledge also have no means of explaining effects that arise from considering underlying higher-order interrelationships between categories. Tenenbaum and Kemp (2009) and Shafto et al. (2008) have demonstrated that inductive reasoning about causal transmission can be dissociated from inductive inferences about physiological properties. Such dissociations suggest that the context or property people are reasoning about prompts them to draw on different and most relevant sources of structured knowledge. Making use of this kind of structured knowledge also gives rise to phenomena such as the causal asymmetry effect, whereby inferences about the transmission of diseases are deemed stronger from prey to predator than from predator to prey (Medin, Coley, Storms & Hayes, 2003; Shafto, et al., 2008). Again, it is hard to see how approaches relying exclusively on nondirectional unstructured knowledge might cogently explain such effects.

Processing Differences

On the surface it appears that approaches placing divergent emphasis on different types of knowledge are incompatible. However, recent evidence suggests that both structured and unstructured types of knowledge play an important role in inductive reasoning, and that they may be a source of individual differences. One of the major distinguishing features appears to be the nature of the mental processes that mediate the use of these contrasting types of knowledge. For example, Rehder (2009) explicitly suggests that the use of structured knowledge relies on an elaborate, analytical thought processes, whereas associative knowledge influences inductive reasoning fairly automatically and without much cognitive effort. Rehder (2009) taught participants about the causal links between category features of artificial categories. In line with the assumption that people draw on extensive causal knowledge, he demonstrated various phenomena, such as a causal asymmetry effect. However, he also found that there was a substantial minority of people whose patterns of inductions did not adhere to those predicted by his causal-based generalization model. Instead, they seemed to rely more on nondirectional associations between the category features.

This suggests that selective inductive reasoning can either be driven by structured knowledge based on theoretical conceptions about relations between categories within a domain, or on unstructured knowledge based on temporal contiguity or degree of association between the categories.

Testing for Effects of Knowledge Type

To test our hypothesis that category-based induction might be driven by different types of knowledge we used a paradigm developed by Shafto, Coley & Baldwin (2007) who were interested in the effects of knowledge domain on induction. Shafto et al (2007) presented participants with arguments consisting of taxonomically or ecologically related categories and manipulated time to respond. To test our hypothesis about differential effects of knowledge type, we also included a manipulation of between-category association. As access to structured knowledge seems to require slower and more elaborate reasoning, we expected people to rely more on unstructured knowledge when under time pressure.

Our design also allowed us to attempt to replicate Shafto et al's finding that whereas people's inferences about taxonomically related categories were unaffected when under time pressure, they gave lower inductive strength ratings to ecologically related categories when they had to respond rapidly. Because Shafto et al. did not control for level of association between their category pairs, it will be of interest to examine whether processing differences between knowledge domains still emerge when degree of association is equated between domains.

Methods

Participants

40 participants took part in the study. They were volunteers from Durham University, who received course credit for their participation. Their mean age was 24.2 years (SD= 7.8 years).

Design

The experiment had a 2 (timing: speeded versus delayed) by 2 (property: cells or disease) by 2 (relation: taxonomic or causal) by 2 (level of association: high versus low) mixed design, with timing as the between-subjects variable.

Materials and Procedure

There were 20 reasoning items consisting of a base category, a causally related target category and a taxonomically related target category. Causally related pairs were always from different superordinate categories, for example, plants and animals, or mammals and reptiles. In contrast, taxonomically related pairs were always from the same superordinate taxonomic category

For each item, there was a causal problem and a matching taxonomic induction problem, resulting in a total of 40 problems.

In order to control for level of association between the base category and its two target categories, 18 Durham University students were asked to rate how strongly pairs of words were associated on a scale from 1 (unrelated) to 9 (very strong association). Whilst no specific examples were given, when generating each rating participants were instructed to consider all kinds of possible relations, such as causal, functional, taxonomic etc, and were asked to give the first answer that came to mind. We selected only those 20 items with a similar level of association between the base and its alternative causal and taxonomic target categories. We then also derived a more objective measure of co-occurrence against which to verify our notion of association. We calculated the frequency with which the two categories co-occurred within six words on the World Wide Web by using a Google proximity search and used a formula suggested by Heylighen (2001) to calculate the conditional probability of co-occurrence:

$$Aw_{1\&w_2}= P\ (w_1\mid w_2) = \frac{P(w_1 \& w_2)}{P(w_1)} = \frac{N(w_1 \& w_2)}{N(w_1)}$$

In this equation, $P(w_1 \& w_2)$ represents the probability that a text contains both words w_1 and w_2, $P(w_1)$ represents the probability that it contains w_1 on its own. To calculate the conditional probability, one can simply count the number of times w_1 and w_2 co-occur and divide this by the number of times w_1 occurs by chance in the same text sample. We then took the mean of these two conditional probabilities and correlated this with our association strength ratings. These two measures were significantly correlated (Spearman's rho= .56, $p<$ .01) supporting our contention that we are

indeed measuring a construct of associative strength in which the activation of one leads to activation of the other, irrespective of the nature of relation between the two categories.

To explore the role that level of association plays in the availability of knowledge from different domains, a median split based on level of association was carried out on the selected items. Thus, for 10 items the association between the base and its target categories was classed as strong and for the remaining 10 items this association was classed as weak. For half the strongly and weakly associated items participants generalized diseases. For the other half, people evaluated inductive conclusions about cells, so whilst property was manipulated within-subjects, content was counterbalanced across participants in a Latin-square design.

Participants learnt that the base category had either a blank disease, such as disease 9T4, or blank cells, such as cells Lo8. They then rated the likelihood that the target category shared the disease or cells on a 9-point scale. For example, participants might be presented with the following induction problems:

Carrots have disease 3dfT.
How likely is it that Rabbits have disease 3dfT?
(causal/disease)

Carrots have disease ww3T.
How likely is it that Radishes have disease ww3T?
(taxonomic/disease)

Acorns have cells T4H.
How likely is it that Squirrels have cells T4H?
(causal/cells)

Acorns have cells eR2.
How likely is it that Walnuts have cells eR2?
(taxonomic/cells)

The induction problems were presented on a laptop. The premise and conclusions were presented simultaneously and appeared in a red font. Participants could only enter their response once the font changed to green. In the speeded condition, the font changed from green to red after one second and participants were instructed to read the problem and respond as fast as possible without sacrificing accuracy. In the delayed condition, the font only changed colour after 10 seconds and participants were instructed to carefully consider their responses. They entered their response on the key board by giving a rating between 1 and 9.

Post-Test

The post-test assessed people's beliefs about taxonomic and causal relatedness. For each of the 40 category pairs, participants were asked two questions, resulting in a total of 80 questions. One question asked them whether they believed that the two categories were from the same biological class and the other asked whether the two

categories were part of the same food chain. Participants could respond with YES, NO or DON'T KNOW, but were instructed to use the third option sparingly, as the emphasis was on their intuitions and beliefs rather than on factual correctness. The mean proportion of positive responses to the two post-test questions about biological group membership and food chain relations across the two timing conditions did not correlate with our web-based measure of co-occurrence (Spearman rho correlation coefficients ranged from -.18 to .16, all p's > .27), nor did it correlate with our subjective measure of associative strength (Spearman rho correlation coefficients ranged from .1 to .2, all p's > .18) suggesting that these measures did not reflect associative strength but represents beliefs based on more structured knowledge.

Results

To facilitate an initial factorial analysis of the data, mean inductive strength scores were calculated for the 5 problems representing the unique property by association by relation combination, resulting in 8 means for each participant. These were subjected to a 2 (property: disease or cell) by 2 (relation: causal or taxonomic) by association (high versus low) by 2 (timing: delayed or speeded) mixed-design ANOVA, with timing as the between-subject variable. We predicted effects of degree of association in our results. However, if association does not play an important role, we would expect to observe an interaction between timing and relation, with timing affecting causal but not taxonomic inferences, thus replicating Shafto et al's (2007) findings.

Although the effects of relation, $F_{(1, 38)} = 3.39$, $p = .073$, effect size d = .66, and timing, $F_{(1, 38)} = 3.18$, $p = .082$, effect size d = .6, were approaching significance, timing did not interact with any of the other variables. Thus, when we control for degree of association we do not replicate Shafto et al's finding.

The only large and reliable significant main effect was strength of association, $F_{(1, 38)} = 28.82$, $p < .0001$, effect size d = 2.0. As expected, inferences about closely associated categories ($M = 4.52$, $SE = .14$) were rated stronger than inferences about weakly associated categories ($M = 3.98$, $SE = .14$).

The only significant two-way interaction was between property and relation, $F_{(1, 38)} = 25.68$, $p < .0001$, effect size d = 1.7, suggesting that people showed some context-sensitive reasoning. Bonferroni posthoc tests showed that when reasoning about cells, people rated taxonomic inferences ($M = 5.01$, $SE = .2$) significantly stronger than causal inferences ($M = 3.79$, $SE = .22$, $p < .0001$, effect size d = .9). When reasoning about diseases, people rated causal inferences slightly higher ($M = 4.32$, $SE = .26$) than taxonomic inferences ($M = 3.89$, $SE = .17$) although this difference was not significant ($p = .16$, effect size d = .3). This might suggest that whereas physiological inferences are predominantly supported by taxonomic relations between categories, inferences about diseases can be made on the

basis of external mechanisms, in this case causal transmission, but also on the basis of more internal mechanisms, in this case taxonomic links and thus genetic relatedness.

None of the other higher-order interactions were significant (all p's > .08)

Regression Analyses

To explore how structured and unstructured types of knowledge influence category-based inductions under different conditions, we calculated mean inductive strength ratings for each item separately for the two types of property and timing conditions, resulting in 4 inductive strength scores for each item. Similarly, for each item we calculated the mean proportion of positive responses to the two post-test questions about biological group membership and food chain relations across the two timing conditions.

Multiple regression analyses were carried out on the mean inductive strength scores. We make the theoretical assumption that people will be influenced by strength of association regardless of timing manipulations. Hence, we entered this variable in block 1. In a second block, we added proportion of positive responses to the biological group question and food chain question as the independent predictor variables. This enabled us to evaluate the degree to which adding variables reflecting structured knowledge accounted for additional variance above and beyond strength of association.

All four regression analyses were significant, but different relevant knowledge influenced inductive strength under different conditions. Overall, larger multiple correlation coefficients were observed in the delayed condition, suggesting that people used different types of knowledge to inform their inferences when they had time to do so, whereas under time pressure, the ability to recruit relevant knowledge seemed to be attenuated.

Inferences about Diseases

As Figure 1 shows, speeded inductive reasoning about diseases (R = .59) was significantly predicted by strength of association (beta = .45, t = 3.13, $p = .003$). In the second block, knowledge about relevant causal food chain relations was also a significant predictor (beta = .35, t = 2.04, $p = .05$), whereas taxonomic knowledge was not a significant predictor (beta = .08, t = .44, $p = .67$). Together, adding these two structured knowledge variables accounted for a nonsignificant amount of additional variance (R^2 Change: 9.6%, $F_{(2, 36)} = 2.64$, $p = .09$).

In contrast, reasoning about diseases under delayed conditions (R = .68) was no longer significantly predicted by association (beta = .24, t = 1.8, $p = .08$). However, inductive strength was strongly predicted by relevant knowledge about food chain relations (beta = .61, t = 4.34, $p < .001$), but also by beliefs about biological relatedness (beta = .34, t = 2.33, $p = .03$). Adding the structured knowledge predictors in a second block did account for significantly more variance in inductive strength ratings

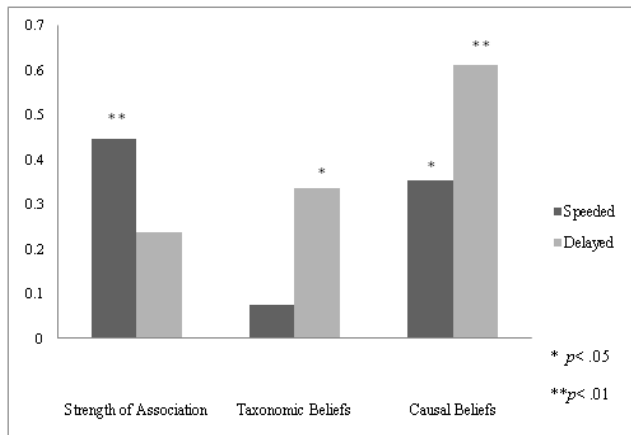than strength of association on its own (R^2 Change: 25.8%, $F_{(2, 36)} = 9.46$, $p < .001$).

Figure 1: Standardized Regression Coefficients for Predictive Relations between Taxonomic and Causal Beliefs, Strength of Association and Inductive Strength Ratings for Diseases

Inferences about Cells

Reasoning about cells showed a different pattern as shown in Figure 2. Under delayed conditions, strength of association was not a significant predictor of inductive strength (beta = .19, t = .15, $p = .14$). Inductive inferences were however predicted by beliefs about biological relatedness (R = .72) (beta = .48, t = 3.48, $p = .001$), and were negatively predicted by beliefs about causal relatedness (beta = -.31, t = -2.29, $p = .03$). Given that we had selected causal targets that were always from different superordinate categories, it is not surprising that causal beliefs were a negative predictor of inferences about cells.

As when reasoning about diseases, adding the structured knowledge predictors in a second block accounted for significantly more variance in inductive strength ratings than strength of association on its own when people were not under time pressure (R^2 Change: 44.2%, $F_{(2, 36)} = 16.33$, $p < .001$)

Speeded inductions about cells (R = .64) were predicted by strength of association (beta = .51, t = 3.76, $p = .001$) and were negatively predicted by beliefs about causal relatedness (beta = -.34, t = -2.5, $p = .05$). Taxonomic beliefs were not a significant predictor of speeded inductive strength ratings (beta = .11, t = .65, $p = .52$). However, adding the structured knowledge coefficients did explain some additional variance above strength of association on its own (R^2 Change: 16.7%, %, $F_{(2, 36)} = 5.09$, $p = .01$).
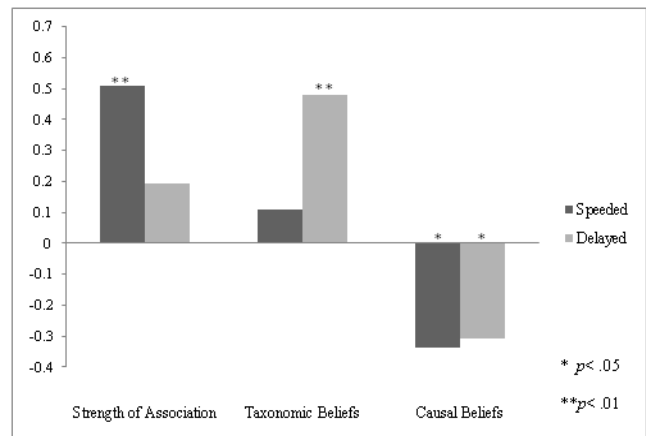

Figure 2: Standardized Regression Coefficients for Predictive Relations between Strength of Association, Taxonomic and Causal Beliefs and Inductive Strength Ratings for Cells

Discussion

Our main proposal was that knowledge effects in category-based induction can be distinguished with regards to two contrasting types of knowledge: effortlessly computable, unstructured knowledge such as strength of association (Rogers & McClelland, 2004) or similarity (Sloman 1993; Sloutsky & Fisher, 2004) on the one hand, and structured knowledge (Kemp & Tenenbaum, 2009, Shafto et al, 2008, Rehder, 2009), which requires more time and processing effort, on the other. Overall, our results strongly support this distinction between different types of knowledge that differ in their processing characteristics. The response timing paradigm used in the current experiment showed that strength of association was a stronger predictor of inductive strength ratings when people had to respond quickly. In contrast, structured causal and taxonomic knowledge became more important when people were forced to delay their response and hence had time to consider the nature of the relationship between the categories.

A secondary goal of this experiment was to explore whether differences in the accessibility of knowledge from different domains arises when level of association is controlled for. The results showed that once level of association was equated across causally and taxonomically related category pairs, the previously observed advantage for taxonomic knowledge (e.g. Shafto et al., 2007) was no longer observed. This suggests that no domain of knowledge is more privileged than any other.

With regards to our main proposal, there are several benefits of being able to draw on two types of knowledge that differ in their processing characteristics. The potency of inductive inferences can be maximized by recruiting structured knowledge, making inferences more sensitive to

contextual factors and relational constraints. It is difficult to see how connectionist models, whose hallmark processes are instantiated by nondirectional and automatic spreading activation, could explain how additional sources of knowledge, such as causal and taxonomic knowledge, selectively influence people's inferences about diseases when people have time but not when they have to respond rapidly.

However, people may not always have the time and available mental resources to try and draw on elaborate background knowledge. Thus, unstructured knowledge acquired through associations, temporal contiguity, or co-occurrence   provides a rich source of ecologically valid information at little or no processing cost (Evans, 2008; Smith & DeCoster, 2000). As demonstrated by Rogers & McClelland's (2004) PDP model, it is conceivable that frequently co-occurring categories would lead to a gradual adjustment of their semantic representations in memory, so that activation of one would either 'prime' or partially activate the representation of strongly associated categories.

Conclusion

We provide support for the claim that category-based inductive reasoning is influenced by two types of knowledge, structured and unstructured knowledge, which are mediated by two contrasting mental processes (Rehder, 2009). Use of unstructured knowledge, such as nondirectional associative strength (Sloman, 1993; Rogers & McClelland; 2004) seems to reflect a relatively effortless process, in which inductions are proportional to the degree to which activation of the premise and conclusion category representations in semantic memory overlap. However, this can be supplemented by the use of more elaborate structured knowledge (Kemp & Tenenbaum, 2009; Shafto et al., 2008). Structured knowledge encodes intuitive theories about the structural relationships between categories, such as knowledge about taxonomic connections or causal interactions. Use of this type of knowledge is constrained by cognitive resources but can maximize inductive potency of inferences beyond mere associative strength between categories.

References

Evans, J. S. B. T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology, 59*, 255-278.

Feeney, A. (2007). How many processes underline category-based induction? Effect of conclusion specificity and cognitive ability. *Memory & Cognition, 35*, 1830-1839.

Fisher, A. V., & Sloutsky, V. M. (2005). When induction meets memory: Evidence for gradual transition from similarity-based to category-based induction. *Child Development, 76*(3), 583-597.

Heylighen, F. (2001). *Mining Associative Meanings from the Web: from Word Disambiguation to the Global Brain*. Standaard Publishers.

Kemp, C., & Tenenbaum, J. B. (2009). Structured Statistical Models of Inductive Reasoning. *Psychological Review, 116*(1), 20-58.

Lopez, A., Gelman, S. A., Gutheil, G., & Smith, E. E. (1992). The Development of Category-Based Induction. *Child Development, 63*(5), 1070-1090.

Medin, D. L., Coley, J. D., Storms, G., & Hayes, B. K. (2003). A relevance theory of induction. *Psychonomic Bulletin & Review, 10*(3), 517-532

Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological Review, 97*(2), 185-200.

Rehder, B. (2009). Causal-Based Property Generalization. *Cognitive Science, 33*(3), 301-344

Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: a parallel distributed processing approach.* London: MIT.

Shafto, P., Coley, J. D., & Baldwin, D. (2007). Effects of time pressure on context-sensitive property induction. *Psychonomic Bulletin & Review, 14*(5), 890-894.

Shafto, P., Kemp, C., Bonawitz, E. B., Coley, J. D., & Tenenbaum, J. B. (2008). Inductive reasoning about causally transmitted properties. *Cognition, 109*(2), 175-192.

Sloman, S. A. (1993). Feature-Based Induction. *Cognitive Psychology, 25*(2), 231-280.

Sloman, S. A. (1998). Categorical inference is not a tree: The myth of inheritance hierarchies. *Cognitive Psychology, 35*(1), 1-33.

Sloutsky, V. M., & Fisher, A. V. (2004). Induction and categorization in young children: A similarity-based model. *Journal of Experimental Psychology-General, 133*(2), 166-188.

Smith, E. R., & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review, 4*(2), 108-131.