AVAILABILITY IN CATEGORY-BASED INDUCTION

Patrick Shafto

Massachussetts Institute of Techonology

John D. Coley & Anna Vitkin

Northeastern University

Induction can be supported by many kinds of knowledge. To be effective, reasoning should be context-sensitive; different kinds of knowledge should be selectively deployed in different situations. For example, in the domain of biology, when reasoning about the distribution of novel internal properties over species, taxonomic knowledge should be recruited since we know that taxonomic membership is not only related to perceptual similarity but is also a good predictor of shared unobservable anatomical features such as four-chambered hearts. However, when reasoning about the distribution of environmental toxins, ecological knowledge should be recruited since such a toxin would plausibly spread through an ecosystem. In this chapter, we address the factors that influence the recruitment of different kinds of knowledge in different contexts. We propose that different kinds of knowledge are differentially available across contexts. Using this concept of availability, we will address an array of experimental results, arguing for availability as a way to unite and explain a broad range of phenomena in category-based induction.

In a classic paper, Tversky and Kahneman (1973) discuss *availability* as a heuristic "by which people evaluate the frequency of classes or the likelihood of events"

(p. 207). This involves estimating frequency or probability "by the ease with which instances or associations are brought to mind" (p. 208). As such, availability on this view is essentially a metacognitive heuristic by which information is judged more likely or plausible based on an estimate of the effort involved in retrieving the information; indeed, Tversky and Kahneman argue that "[t]o assess availability it is not necessary to perform the actual operations of retrieval or construction. It suffices to assess the ease with which these operations could be performed" (p. 208). Our goal in this chapter is to use the spirit of this concept of availability—if not the letter—to unify a set of seemingly disparate findings on category-based inductive reasoning.

We see inductive inference as a process by which knowledge of relations among concepts is used to assess the likelihood that members of two categories will share a novel property. We will argue that the ease with which specific knowledge of conceptual relations comes to mind predicts the likelihood that such knowledge will be used to guide inductive inference. Below, we will present evidence that the availability of different kinds of knowledge to inform induction depends on two factors: acute (short-term) influences of context on availability and chronic (long-term) effects of experience, reflected in baseline differences in the availability of different knowledge structures. In this sense, availability can be thought of as a dynamic property of knowledge in memory that provides a ranking of the likelihood that different kinds of knowledge will be accessed in a given context (e.g. Horowitz et al, 1966, Tulving & Pearlstone, 1966).

Thus, like Tversky and Kahneman (1973), we argue that availability is a variable that mediates between prior knowledge and behavior via the relative accessibility of different kinds of knowledge in a particular situation. Here we part ways with Tversky and Kahneman because for our purposes an *estimate* of the ease with which information

might be accessed is not sufficient. Because specific knowledge must be brought to mind to guide or assess inductive inference, it *is* necessary to "perform the actual operations of retrieval or construction."

The structure of this chapter will be as follows. First, we will discuss the effects of two kinds of context on availability. We will then turn to how experience influences availability and interacts with context. Finally, we will relate this approach to other frameworks, and discuss novel predictions motivated by availability.

Availability in Category-Based Induction

A central problem in the psychology of induction is what kinds of knowledge we use to guide a particular inductive inference. Given all that we may know about a concept, which knowledge is relevant to a particular problem? This problem is beyond the scope of many previous models of inductive reasoning, which rely on a single kind of knowledge in a domain (though some recent progress has been made toward models supporting multiple kinds of knowledge, see Shafto, Kemp, et al, 2005; Shafto et al, in prep). We propose a framework for understanding reasoning in rich domains, explicitly linking the acute effects of immediate context and the chronic effects of long-term experience via the *availability* of different kinds of knowledge. We suggest that inductive generalizations are guided by the knowledge that is most available in a particular context, given past experience. Moreover, availability is dynamic in that changes in availability arise from one of two sources: changes in context or changes in underlying knowledge.

Context plays a crucial role in acute changes in availability. For any given inductive task, the presence of a biasing context can change the relative availability of different kinds of knowledge. For example, if you learn of a new *environmental toxin* affecting *frogs*, ecological knowledge about who may be exposed or causal knowledge

about how the toxin might be transmitted may become more available, leading you to expect that *herons* might be at risk (because they share a habitat with and eat frogs). In contrast, if you learn of a new *enzyme* discovered in frogs, you might expect the same enzyme to be present in *toads*, due to their close biological affinity, and never consider herons as a plausible candidate for the enzyme. Even if we know nothing about the nature of the property in question, any known examples that possess it may influence the availability of knowledge. For example, if told about a property that cows, chickens, and pigs have, knowledge about domesticated animals may become more available.

In the case where context does not provide a bias, availability reduces to a default way of thinking about a domain, which reflects the knowledge that has been rendered chronically most available by experience. Experience may lead to multiple kinds of knowledge all becoming equally available or a single kind of knowledge becoming more available over time. For example, the relative salience of the ecological versus taxonomic knowledge that inform the inferences about *frog* might differ for an ecologist working for Greenpeace, a molecular biologist, or a relatively uninformed undergraduate participant in a psychology experiment. The chronic effects of experience are manifest in baseline availability of default knowledge in a domain.

Importantly, acute and chronic changes in availability can interact, which can have important implications for reasoning in a domain. If, at baseline, multiple kinds of knowledge are equally available, reasoning will be highly sensitive to acute changes in availability as a result of context. However, if at baseline one kind of knowledge is a highly available default, much stronger context will be required to elicit reasoning based on the less-available knowledge. Thus, the interaction between acute and chronic changes

in availability has an important impact on the kinds of knowledge that are available to guide reasoning.

In the following sections, we will consider several studies that demonstrate that context and experience result in important changes in inductive reasoning. Most importantly, we will consider studies that show interactions between experience and context in inductive reasoning. We will argue that these effects are best understood in terms of acute and chronic changes in the availability of different kinds knowledge.

Context-based Changes in Availability

We suggest that the context provided by an inductive argument results in acute changes in the availability of different kinds of knowledge and different patterns of reasoning. In this section, we will consider a range of previous results investigating reasoning in knowledge-rich domains. We will re-consider this evidence in light of availability, particularly looking for qualitative changes in reasoning based on immediate inductive context. We consider two sources of context in these category-based induction tasks: the property in question, and the set of categories given in the problem.

Property Effects on Induction

One source of context in an inductive problem is the property to be generalized from premise to conclusion. Knowing something about the nature of this property can be informative as to what relations are recruited for making guesses about the distribution of that property. For example, if we are told that chickens have sesamoid bones we might conclude that other flightless birds might share a physiological, internal property such as bone structure. On the other hand, if we knew that chickens had the sesamoid flu we

might reconsider having chicken for dinner for fear of catching it ourselves. Although in each example the premise category is *chicken*, the conclusion categories we believe to share a property can change dramatically depending on what that property is. The systematic use of different kinds of knowledge to inform induction has been termed inductive selectivity. On our view, the immediate context provided by the property of an inductive argument can produce acute changes in the availability of different kinds of knowledge for inductive generalizations.

Heit and Rubinstein (1994 exp 1 & 2) proposed that different relations between premise and conclusion categories determined the strength of an inductive inference depending on the property a premise and conclusion were said to share. In their study they asked participants to make estimates of the probability that a pair of animal categories might share a property. Pairs of animals were chosen to be related by shared anatomy (e.g., *whale* and *bear*) or behavior (e.g., *whale* and *tuna*). Participants were asked about the likelihood of such pairs sharing an anatomical property like having two chambers in their liver that act as one, or a behavioral property such as traveling in a back-and-forth or zigzag trajectory. Heit and Rubinstein found that participants made the highest probability estimates for items where the relation between a pair of animals matched the kind of property they were asked to reason about (e.g., whales and bears have two chamber livers, whales and tuna travel in zigzag trajectories). These results suggest that the property of projection influenced the kind of knowledge that was recruited to support inductive inferences: anatomical properties made anatomical knowledge more available, whereas behavioral properties made behavioral knowledge more available.

Ross & Murphy (1999) have shown similar effects of property on the selective use of conceptual relations to guide inferences in the domain of food. They established that most participants cross-classified food into two major knowledge structures, taxonomic, based on shared features or composition, and script, based on what situations a food is consumed in. Participants in their study were asked to make biochemical or situational inferences about triplets of food. Participants were taught that the target food such as bagels had a biochemical property (enzyme) or situational property (eaten at an initiation ceremony). They were then asked to project the property to one of two alternatives, a taxonomic alternative such as crackers, or a script alternative such as eggs. Ross & Murphy found that participants made more taxonomic choices for biochemical properties and more script choices when considering situational properties. It seems that the nature of the property increased the availability of relevant knowledge about shared composition or situational appropriateness of a food, ultimately producing different patterns of induction for different kinds of properties.

Recent evidence also suggests that, like adults, children's inductive generalizations are sensitive to the property in question. Using a similar triad method to that of Ross & Murphy (1999), Nguyen & Murphy (2003) found that in the domain of food, seven-year old children (but not four-year-olds) made more taxonomic choices when reasoning about a biochemical property and more script choices when reasoning about a situational property. That is, they thought that a bagel would be more likely to share an enzyme with crackers but be more likely to be eaten at a ceremony with eggs. This suggests that the context provided by a property begins to mediate the differential availability of knowledge from an early age.

Further evidence of children's selective use of knowledge based on property comes from the domain of biology. Coley and colleagues (Coley, 2005; Coley, Vitkin, Seaton & Yopchick, 2005; Vitkin, Coley & Kane, 2005) asked school-aged children to consider triads of organisms with a target species, a taxonomic alternative (from the same superordinate class but ecologically unrelated), and an ecological alternative (from a different taxonomic class but related via habitat or predation). Children were taught, for example, that a banana tree had a property (either a disease or "stuff inside") and asked to choose if a calla lily (taxonomic match) or a monkey (ecological match) shared the property. Results clearly indicate that children were sensitive to the kind of property they were asked about, choosing taxonomic matches at above-chance levels when reasoning about insides, and choosing ecological matches at above-chance levels when reasoning about diseases that might plausibly spread through an ecosystem. In another task, children were asked to make open-ended predictions about what kinds of things might come to share properties (Vitkin, Coley & Hu, 2005). They were taught that a pair of animals either had a disease or "stuff inside", they were asked to project what other things might share this property and explain their reasoning. Children's responses were characterized as being based on similarity (taxonomy, shared features, etc) or interaction (contact through ecosystems). Children made more responses rooted in similarity when thinking about internal substances and more interaction-based responses when asked to consider diseases. Consistent with availability, this pattern of response demonstrated that children's generalizations were guided by different kinds of knowledge depending on what properties they were reasoning about.

Overall, we have seen that the property projected in an inductive task provides context for inferences, changing the kinds of information that are used make

generalizations. Specifically, the property being projected provides useful clues about what kinds of conceptual relations (disease transmitted through ecology, biochemical composition shared by taxonomic food classes, etc.) might support the generalization of such a property. This has been demonstrated in at least two different domains of knowledge, food and biology and in both children and adults. These results are consistent with the idea that the property being projected is one factor that influences the relative availability of different kinds of knowledge to support inductive reasoning.

Relations Among Premise and Conclusion Categories

Another line of evidence compatible with the proposal that availability mediates the basis of inductive reasoning examines the effects of relations among premise categories, or among premise and conclusion categories, as influencing the availability of different kinds of knowledge for guiding inductive inference. Work by Medin, Coley, Storms & Hayes (2003) identified several effects associated with the relationship between categories presented in the premises and conclusions, under the rubric of *relevance effects*. One key component of the relevance framework, as well as our availability framework, is the idea that salient relations among premise categories, or between premise and conclusion categories, may guide or constrain the inference supported by those premises. Medin et al. (2003) examine this idea with respect to two broad classes of phenomena, *causal relations* and *property reinforcement*. We concentrate on the latter. Medin et al. (2003) present several examples where increasing the salience of specific relations among categories in an inductive argument leads to violations of normative logic or of the predictions of similarity-based models of inductive reasoning.

The first of these is *non-diversity via property reinforcement*, which predicts that an argument with less diverse premises might be perceived to be stronger than an argument with more diverse premises if the premise categories of the more diverse argument reinforce a salient relation not shared by the conclusion category. For example, consider the following arguments:

A. Polar Bears have property X.

Antelopes have property X.

All Animals have property X.

B. Polar Bears have Property X.

Penguins have Property X.

All Animals have Property X.

From a strictly taxonomic point of view, *polar bears* (a mammal) and *penguins* (a bird) provide better coverage of the conclusion category *animal* than polar bears and antelopes, which are both mammals. Thus, a model based only on taxonomic knowledge must predict Argument B to be stronger. However, the salient property shared by polar bears and penguins—namely, adaptation to a cold climate—renders plausible the possibility that Property X is related to life below zero, and therefore might weaken the inference that all animals would share the property. Indeed, Medin at al. (2003) find that subjects in the U.S., Belgium and Australia on average rated arguments like A stronger than arguments like B[1]. This suggests that the salient property shared by the premise categories in B cancels out the greater coverage they provide of the conclusion category.

A second related phenomena discussed by Medin et al. (2003) is *conjunction fallacy via property reinforcement*, which predicts that arguments with a single conclusion

[1] Though recent evidence suggests that this effect may not be particularly robust (Heit & Feeney, 2005).

category might be perceived to be weaker than arguments with an additional conclusion category (a violation of normative logic because a conjunctive statement cannot be more probable than one of the component statements) if the second conclusion category reinforces a salient relation shared by all categories. For example, consider the following arguments:

C. Chickens have property X.

Cows have property X.

D. Chickens have Property X.

Cows have Property X.

Pigs have Property X.

Normative logic requires that "cows have property X" must be more likely than "cows and pigs have property X," but Medin et al. (2003) found that participants reliably rated arguments like D as more likely than arguments like C. The addition of *pigs* might serve to increase the availability of the knowledge about farm animals, and therefore strengthens Argument D relative to Argument C.

Finally, Medin et al. (2003) discuss *non-monotonicity via property reinforcement*. Monotonicity is the idea that all else being equal, adding premise categories that are proper members of the same superordinate as a conclusion category should strengthen the argument (see Osherson et al., 1990). Medin et al. (2003) predict that adding premises might weaken an argument if the added categories reinforce a relation shared by premise categories but not the conclusion category. For example, consider the following arguments:

E. Brown Bears have property X.

Buffalo have property X.

F. Brown Bears have Property X.

Polar Bears have Property X.

Black Bears have Property X.

Grizzly Bears have Property X.

Buffalo have Property X.

Monotonicity predicts that Argument F should at least as strong as Argument E because the premise categories in F necessarily cover the lowest inclusive category *mammals* at least as well as the premise of E, and are at least as similar to *buffalo*. However, Medin et al. (2003) find that—contrary to the principle of monotonicity—arguments like E are consistent given stronger ratings than arguments like F, presumably because the premises of F reinforce the relation of *being a bear*, and therefore make plausible the inference that Property X is particularly ursine in nature, and therefore unlikely to be true of buffalo.

Together, these results suggest that manipulation of relations among premise and/or conclusion categories can result in violations of normative logic (in the case of the conjunction fallacy) or violations of predictions derived from similarity-based models (in the case of non-diversity or non-monotonicity). We propose that these manipulations can be seen as manipulations of the availability of specific conceptual relations. By rendering specific relations such as *polar animals*, *farm animals* or *bears* highly available, these manipulations serve to overcome more general default approaches to evaluation of inductive arguments.

Using a more natural paradigm that allows participants to generate their own conclusions given a set of premises, Baker and Coley (2005; see also Coley, Baker & Kemp, 2004) have investigated whether spontaneous and relatively unconstrained inductive projections are sensitive to manipulations of relations among premise

categories. In this study, 30 college undergraduates were given pairs of premise species, taught a novel property said to be shared by the pair, and asked what other species might have the property, and why. Premise pairs were either drawn from the same superordinate category (taxonomically near pairs) or from different superordinates (taxonomically far pairs). Pairs were also ecologically related (via habitat, predator-prey relation, or ecological niche) or unrelated. See Table 1 for sample items.

Table 1. Sample items from Adult Open-Ended Projection Study.

Ecological Relation	Taxonomic Distance	
	Near	Far
Related	*Heron, Duck*	*Hawk, Field Mouse*
Unrelated	*Otter, Deer*	*Chipmunk, Bullfrog*

Reponses were coded according to the relationship between the given premise pair and the conclusion categories generated by the participants. The basis of an inference was judged to be *taxonomic* if participants' explanations emphasized that premise and conclusion categories belonged to the same class or kind, were similar in appearance, or similar in general. Responses were coded as *ecological* if participants' explanations relied on an interaction between premise and conclusion categories that highlighted a non-taxonomic relation such as a predator/prey relation, shared diet or habitat.

Results suggest that projections were sensitive to salient relations among premise categories. Specifically, taxonomic inferences were more frequent for taxonomically near pairs than for taxonomically far pairs, and were also more frequent for ecological unrelated than for ecologically related pairs. Likewise, ecological inferences were more frequent for ecologically related pairs than for unrelated pairs, although frequency of

ecological inferences was not affected by taxonomic distance among premise categories. In sum, relations among premise categories appear to influence the knowledge recruited to support generalizations.

These findings suggest that relations among premise categories may impact the availability of different kinds of knowledge for guiding spontaneous inferences. Premise categories that share salient taxonomic relations render such knowledge available, and thereby increase the likelihood of taxonomic generalizations. Likewise, premise categories that share salient ecological relations increase the availability and likelihood of ecological inferences.

Availability, Experience, and Default Domain Knowledge

Availability, as we have described it, depends on prior knowledge and context. In the previous section, we considered acute changes in availability due to the nature of the property being projected and relations among categories in an argument. In this section, we present evidence that increased knowledge and experience in a domain can lead to chronic changes in the relative availability of different kinds of knowledge for inductive reasoning. We will argue that experience-related changes in underlying knowledge—such as accrual of more facts, changes in the frequency with which different knowledge is used, and even fundamental changes in conceptual structure—are accompanied by changes in the baseline availability of different kinds of knowledge, and also by an increased sensitivity to context. In brief, we argue that the concept of availability is useful in explaining observed effects of experience on inductive reasoning.

Experience-related changes in availability of different kinds of knowledge

By definition, experts in a given domain know more than novices. However, in addition to the accrual of facts, expertise may also result in changes to the relative availability of different kinds of knowledge. We argue above that taxonomic knowledge may be a default guide for novices when reasoning about a given domain. However, several lines of evidence suggest that taxonomic relations consistently fail to predict expert reasoning to the same extent. For instance, López et al. (1997) show that when forced to choose the stronger of two inductive arguments about local mammals species, University of Michigan undergraduates' responses are almost unanimously in accord with the Osherson et al. (1990) taxonomically-based similarity coverage model. In contrast, responses of Itza' Maya participants—indigenous people of Guatamala who live in close contact with nature, and depend largely on hunting and swidden agriculture for subsistence, and therefore possess extensive knowledge of local flora and fauna—were at chance when evaluated against predictions of the similarity coverage model. Instead of relying on taxonomic relations, the Itza' recruited specific causal-ecological knowledge to evaluate the arguments. Follow up work by Proffitt, Coley and Medin (2000) revealed a pattern of induction among Chicago-area tree experts that was remarkably similar to that of the Itza', suggesting that domain-specific experience, rather than language or culture, is the driving factor. Indeed, this same general finding—increase in the relative salience of non-taxonomic relations for guiding induction—has been reported for commercial fishermen, professional musicians, and even undergraduates reasoning about alcohol (see Coley, Shafto, Stepanova & Baraff, 2005 for a review).

Not only does experience in general change the relative availability of different kinds of knowledge, but specific kinds of expertise also appear to lead to differential salience of knowledge among experts. For instance, Proffitt et al. (2000) report that of

three different groups of experts (taxonomists, landscapers, and maintenance workers), only taxonomists' patterns of induction relied on taxonomic knowledge. Taxonomists were also more likely to explain their inferences by referring to taxonomic factors such as typicality, family size, or diversity of premise species. Taxonomists, with their focus on knowledge about genealogical relations among species, tend to apply taxonomic knowledge even when reasoning about diseases among trees. However, landscapers and maintenance workers, who have other concerns, were more likely to apply other kinds of knowledge to the same task. Converging evidence comes from studies of categorization; Medin, Lynch, Coley & Atran (1997) found that these same groups of experts tended to sort tree species on the basis of relations specific to their area of expertise. For example, taxonomists tended to utilize botanical families, landscapers utilized categories like *specimen tree* that reflected a species' typical use in landscape design, and maintenance workers tended to use categories like *weed tree*, which reflected the ease of caring for different species. Thus, within a single domain, differential experience can render different knowledge available. These effects demonstrate how extensive experience can elicit chronic changes in availability.

Culture is another source of differential experience that could impact the availability of knowledge for induction. A striking example of culturally induced changes in availability can be found in the work of Choi, Nisbett, and Smith (1997). They investigated reasoning about biological and social categories among Americans and Koreans. Some have argued that western individualist cultures tend to process information analytically and generally attend to categories and rules, whereas eastern collectivist cultures tend to consider problems holistically, attending to categories in terms of interactions in a setting (see Nisbett, Peng, Choi & Norenzayan, 2001, for a

review). One interesting exception to this is in the domain of social categories and roles, which are argued to be particularly salient and important to eastern collectivist cultures. Based on these cultural differences, Choi et al (1997) predicted that manipulations of taxonomic category salience (specific versus general conclusions) would have a differential effect on American and Korean peoples' reasoning about biological and social categories. Specifically, they predicted that because taxonomic categories in general are less culturally salient to members of a collectivist culture, manipulations of salience should have a more pronounced effect on Koreans' biological inferences than on Americans', for whom taxonomic knowledge is already highly salient. In contrast, they predicted the opposite effect for social categories, where taxonomic knowledge is argued to be more salient for Koreans than for Americans. Consistent with these predictions, they found that when reasoning about biological categories, Koreans preferred taxonomic responses to arguments with general conclusions than to those with specific conclusions, whereas this manipulation had no effect on Americans' responses. Conversely, when reasoning about social categories, Americans preferred taxonomic responses to arguments with general conclusions than to those with specific conclusions, whereas this manipulation had no effect on Koreans' responses. These results are consistent with the claim that culture can be viewed as a kind of experience that may result in chronic changes in the availability of different kinds of knowledge in a domain.

In sum, experimental results consistently reveal differential use of taxonomic versus other knowledge in experienced versus novice populations. We propose that these differences can be understood as reflecting chronic changes in the relative availability of different kinds of knowledge that accompany the acquisition of expertise.

Experience-related changes in sensitivity to context

In this section, we focus on experience-based differences in the availability of knowledge to guide induction as a function of context. We argue that general changes in baseline availability of different kinds of knowledge also lead experts to be more sensitive to context. In other words, because more kinds of knowledge become available with experience, experts can draw on the most appropriate relation to guide a given inference. Below we present evidence that domain-specific experience is associated with increased inductive selectivity, and with increased sensitivity to relations among premise categories.

Experience-related changes in inductive selectivity. Recent evidence suggests that experience leads to an increase in inductive selectivity. For example, Shafto & Coley (2003) demonstrate experience-related changes in inductive selectivity by contrasting how experts (commercial fishermen) and novices (university undergraduates) reasoned about marine creatures. In this experiment, participants were given either a novel blank property ("has a property called sarca") or a novel disease ("has a disease called sarca") to reason about. Participants were given examples of creatures that had the property and then were asked to infer which other creatures (from a broad array of fish, sharks, whales, and crustaceans) would have the property. Results indicated marked differences in inductive selectivity between novices and experts. When told about blank properties, experts tended to generalize to taxonomically related creatures. However, when told about diseases, experts tended to generalize to creatures related in the food web, specifically by making directional inferences from prey to predators. In contrast, novices tended to generalize to taxonomically related creatures regardless of the property. This result is notable because although novices were unlikely to have the detailed knowledge

about marine food web relations upon which commercial fishermen depend for their livelihood, they undoubtedly possessed some rudimentary knowledge of marine food web relations (e.g., that sharks eat fish and not the other way around). This constitutes enough knowledge to make rough inferences based on food web information. Nevertheless, we observed no such reasoning among novices.

We interpret these results as suggesting that for novices, knowledge about food web relations is generally less available and that the context provided in this experiment (essentially, *disease*) did not create enough of a change in the availability of food web knowledge to overcome their taxonomic default. On the other hand, for experts who rely heavily on knowledge of food web relations among these creatures on a daily basis, taxonomic and food web knowledge have relatively similar baseline availabilities and the experimental context (*disease* versus *property*) was enough to manipulate the availability of the different kinds of knowledge.

Stepanova and Coley (2003, see also Coley, Shafto, Stepanova & Baraff, 2005; Stepanova, 2004) also contrasted reasoning by individuals with extensive or limited experience in a domain. However, rather than comparing experts and novices reasoning about a single domain, they compared a single population (US college undergraduates) reasoning about domain they have extensive experience with (alcoholic drinks) versus a domain with much less relevance to their daily lives (animals), on the assumption that for the typical college student, alcohol possesses greater relevance and cultural importance than animals, and undergraduates are likely to have more first-hand experience of, more frequent exposure to, and richer and more abundant folk theories about alcohol than animals. The task required participants to choose which of two pairs of premise categories provided better evidence for a generalization to *any alcohol* (or *any animal*).

Sets of premise categories were chosen so that one would clearly be stronger via diversity. In addition to being randomly assigned to the animal or alcohol conditions, participants were also randomly assigned to evaluate arguments about a chemical component or about getting sick. Results showed clear evidence for inductive selectivity when undergraduates were reasoning about alcohol. Specifically, participants reasoning about alcohol showed differential use of taxonomic knowledge as a function of property; these participants were more likely to make diversity-based inferences about getting sick than about a chemical component. Moreover, participants reasoning about alcohol also provided different explanations for their choices as a function of property; they were more likely to offer causal explanations for inferences about getting sick, but more likely to offer taxonomic explanations for inferences about a chemical component. In contrast, there was no evidence of inductive selectivity when participants were reasoning about animals; neither the relative frequency of diversity-based choices nor the type of explanations provided for those choices varied for inferences about a chemical versus getting sick. These results are in close accord with those of Shafto and Coley (2003) described above, and suggest that greater domain-specific experience may increase the potential availability of multiple conceptual relations, and therefore increase inductive selectivity as a function of property being projected.

Experience-related changes in sensitivity to premise relations. Recent developmental work also suggests that domain-specific experience may increase children's sensitivity to relations among premise categories as potential constraints on induction. In the Vitkin, Coley and Hu (2005) study described above, relations between premises as well as property effects were investigated using open-ended inductive projections among elementary school children from urban communities and rural communities. Relations

among animal pairs were manipulated in a two by two design such that each pair was either taxonomically close (from the same superordinate) or taxonomically far (from different superordinates) and either ecologically related (via habitat or predator-prey relations) or ecologically unrelated. Responses were coded as being based on similarity (taxonomy, shared features, etc) or interaction (contact through ecosystems). If experience increases sensitivity to context, then the greater opportunity for direct interaction with plants and animals in relatively intact ecosystems afforded by a rural environment may lead to increased sensitivity to relations among premise categories among rural children.

Indeed, urban and rural children differed strikingly with respect to their sensitivity to relations among premise categories. Rural children showed consistent sensitivity to differential relations among premise categories. Specifically, rural children made more similarity-based projections for taxonomically close pairs than for taxonomically far pairs. They also made more interaction-based projections for ecologically related pairs than for unrelated pairs, and for taxonomically far pairs than for taxonomically close pairs. None of these effects were evident for urban children. These results suggest that experience may mediate the availability of different conceptual relations for guiding children's spontaneous inferences. For biologically experienced rural children, relations among premise pairs was sufficient context to render taxonomic and ecological knowledge differentially available to inform inductive projections. In contrast, for biologically inexperienced urban children, this was not the case.

Taken together, these results suggest that the influence of experience on inductive reasoning can be thought of in terms of changes in the relative availability of different kinds of knowledge. Experience can be seen as increasing the relative availability of non-

default knowledge for guiding induction, as in the cases of fishermen reasoning about marine creatures and undergraduates reasoning about alcohol. Because multiple kinds of knowledge are equally available, expert reasoning is highly sensitive to context. In contrast, because one kind of knowledge is a highly available default for novices, much stronger context is required to elicit reasoning based on the less-available knowledge. Thus, changes in chronic availability results in increased sensitivity to context among experts relative to novices.

Summary: Availability in Category-Based Induction

Previous research has demonstrated two main factors that influence recruitment of different kinds of knowledge in category-based induction: experimental context and prior experience. We argue that these factors are best understood as manipulations of the *availability* of different kinds of knowledge.

We have reviewed previous work demonstrating that people's inductive generalizations are sensitive to the property to be projected. We find that different properties can lead to qualitatively different patterns of generalization, such as in the domain of foods where biochemical properties lead to taxonomic reasoning and situational properties lead to script-based reasoning. We propose that properties that are consistent with a particular kind of knowledge will increase the availability of that kind of knowledge. For example, reasoning about what foods would be eaten together at an initiation ceremony activates script knowledge about what foods typically co-occur, and therefore increases the likelihood of script-based inferences by rendering that knowledge temporarily available. Thus, the property of projection can lead to an acute change in availability of knowledge and characteristically different patterns of generalizations.

Similarly, we propose that relations between premise and conclusion categories are also able to elicit changes in availability. In cases such as *non-monotonicity via property reinforcement* as well as open-ended induction, the presence of known relations among the premises results in characteristically different patterns of generalization. For example, if you learn a new fact that is true of turtles and lizards, you might generalize it to other reptiles, whereas if you learn a new fact that is true of turtles and ducks, you might generalize it to pond creatures instead. We suggest that presenting information about categories that are united by a particular kind of knowledge leads to an acute increase in the availability of that kind of knowledge and an increased likelihood of an inference based on that knowledge.

More compelling evidence derives from the role of experience. In our framework, availability depends on context and prior knowledge accrued through experience. We have reviewed research with experts suggesting that experience can change the baseline availability of different kinds of knowledge. People tend to use knowledge that they have a deep understanding of and that has proven useful in the past. For example, evidence from tree experts suggests that a person's experiential background leads to chronic changes in the kinds of knowledge that are available for reasoning; by default, tree taxonomists tend to think about trees in terms of scientific taxonomies, whereas maintenance workers tend to think about trees in terms of maintenance concerns. This evidence suggests that experience can lead to chronic changes in the availability of different kinds of knowledge.

However, we think the most compelling evidence is manifest in interactions between experience and context. People are more likely to demonstrate robust inductive

selectivity in domains where they have extensive experience. For example, when reasoning about marine creatures, commercial fishermen demonstrated more inductive selectivity than university undergraduates; university undergraduates demonstrated more inductive selectivity when reasoning about alcohol than when reasoning about animals. We propose that experience facilitated inductive selectivity in these studies by causing an increase in the availability of non-default knowledge, essentially "leveling the playing field" and allowing for an increased sensitivity to context.

In sum, we see availability as a promising explanation for both context and experience effects in category-based induction. However, at this point, much of our case remains circumstantial. Though availability seems an apt explanation of previous results, one might ask, what do we gain by thinking about these effects in terms of availability?

Connections & Extensions

Traditional models of inductive reasoning have focused on a single kind of knowledge in a domain, eschewing the effects of context. For example, the similarity-coverage model (Osherson et al, 1990) focused on taxonomic knowledge in the domain of biology, using similarity between premises and conclusions as well as a taxonomic coverage to predict inductive generalizations. These models account for phenomena such as taxonomic diversity, the fact that people rate an anatomical property true of diverse premises such as *robins* and *ostriches* more likely to be shared by *all birds* than a property shared by *robins* and *doves*, and are generally in remarkably close accord with novices generalizations of anatomical properties[2]. Though able to capture reasoning

[2] The authors also propose that category members may be differentially available; for example, suggesting that robins are more available members of the category bird than turtledoves.

based on taxonomic knowledge, the similarity-coverage model does not account for reasoning based on other kinds of knowledge, and thus does not naturally extend to the kinds of reasoning we have been discussing here (see Smith et al, 1993 for an account of some property effects). Some context effects, such as inferences about behavioral and anatomical inferences in Heit & Rubinstein (1993) can be handled by allowing similarity to be defined context-sensitively; however, no similarity-based models can account for the kinds of asymmetric causal inferences in Shafto & Coley (2003).

Sloman's feature-based induction model (1993) also predicts inferences about taxonomic properties, but differs from the similarity-coverage model in not assuming a taxonomic structure over objects. Under (the geometric interpretation of) this model, prior knowledge is represented by an object-feature matrix, and inferences are generated by considering the proportion of shared properties between the premise and conclusion, relative to the total number of properties for the premise. There are many potential ways to extend the feature-based model to handle context effects, including adding context-specific feature weights. However, all extensions require the addition of abstract knowledge not present in the original model.

Rather than focusing on the fact that these models do not account for property effects, we think it worthwhile to emphasize commonalities between the similarity-coverage and feature-based models and our availability-based approach. The success of both of these models in predicting undergraduates' judgments suggests that knowledge about taxonomic relations is central to undergraduates' conceptualization of biological kinds. We have suggested that taxonomic knowledge is chronically more available to undergraduates, and therefore relied upon as a default strategy, while taxonomic and ecological knowledge are both chronically available to experts (such as commercial

fishermen). However, neither the similarity-based nor the feature-based model includes a means to explain the effects of properties on reasoning, or a natural way to explain how experience influences reasoning.

Some interesting parallels can be drawn between the availability framework and some previous work that has addressed aspects of induction discussed under our approach. Our notion of availability in inductive reasoning is analogous to Barsalou's (1982) distinction between *context-independent* and *context-dependent* properties of concepts[3]. Context-independent properties (e.g., basketballs *are round*) "are activated by the word for a concept on all occasions" (p. 82), whereas context-dependent properties (e.g., basketballs *float*) "are only activated by relevant contexts in which the word appears" (p. 82). Barsalou demonstrates that, for example, priming with relevant context facilitates the verification of context-dependent properties, but has no effect on the verification of context-independent properties. This distinction can be applied to our analysis of availability in category-based induction by granting that relations among concepts, as well as properties of individual concepts, may vary in their context-dependence. Thus, what we have called acute changes in availability correspond to priming a relevant context, whereas chronic changes in availability correspond to representational changes in the context-dependence of classes of relations among concepts (e.g., predator-prey relations). It remains to be seen whether any class of relations are truly context-independent, but taxonomic relations may be one candidate (e.g., Coley, Shafto, et al, 2005).

Our notion of availability as applied to category-based induction also fits nicely into Medin et al's (2003) Relevance framework. A central claim of the relevance

---

[3] We are grateful to Brett Hayes for pointing out this connection.

framework is that relations among premise categories, or between premise and conclusion categories, guide or constrain an inductive inference to the degree that such relations are deemed *relevant*. Relevance, in turn, is conceptualized in terms of *effort* and *effect*. The less effort needed to access or construct a relation, the more likely it is to be deemed relevant to a particular inference. Conversely, premises that have greater cognitive effect (in the sense of potentially leading to conceptual changes or allowing the derivation of novel conclusions) are more likely to be deemed relevant. On this view, our notion of availability can be seen as a more detailed way to think about the effort component of the Relevance framework. Specifically, the effort associated with a given conceptual relation reflects the availability of that knowledge; all else being equal, more available knowledge requires less effort to access and use. Thus, both the acute and chronic changes in availability reviewed above reflect—from the perspective of Relevance—acute and chronic changes in the effort required to access a given set of conceptual relations.

In the spirit of both Barsalou (1982) and Medin et al. (2003), availability makes two distinctions: chronic changes in availability of different kinds of knowledge grounded in experience, and acute changes in availability as a result of context. Chronic changes in availability account for why taxonomic knowledge is less effortful than ecological knowledge for biological novices and taxonomists but not for fishermen. Acute changes in availability reflect the fact that context manipulates effort to access different knowledge (cf. Heit & Bott, 2000). The interaction between chronic and acute changes in availability determines the degree to which people show inductive selectivity.

Though we believe availability provides a coherent framework uniting expertise and property differences in induction, we think that the true merit of thinking about

category-based induction in terms of availability will be in the guidance it provides in moving research forward. In the next section, we describe some recent work inspired by availability, and derive additional novel (but yet untested) predictions from this proposal.

Availability in action

We see two major challenges in the development of an availability-based framework. The first is to identify what kinds of predictions can be generated by understanding the relationship between knowledge and reasoning in terms of availability, and to begin to test those predictions. The second is to explain how chronic changes in availability arise with experience. In this section, we will outline some initial studies addressing the first challenge, and some preliminary ideas which address the second.

To be useful, any framework must generate new hypotheses as well as describe existing results. In a recent set of studies, we investigated availability as a possible explanation for the lack of inductive selectivity in novice populations (Shafto, Coley & Baldwin, submitted). To implicate availability, it is important to show that having knowledge is not sufficient for inductive selectivity. Previous research (Shafto & Coley, 2003) suggests that context (*novel diseases* and *novel properties*) did not elicit the selective use of taxonomic and ecological knowledge in biological novices. One reason novices may not have demonstrated inductive selectivity was a baseline difference in the availability of taxonomic and ecological knowledge. In a series of experiments (Shafto, Coley & Baldwin, 2005) we provided support for this claim by investigating the effects of context on novices' use of taxonomic and ecological knowledge, focusing on ecological relations that were familiar to novices. Pre-testing insured that the novices knew the taxonomic and ecological relations in question. However, despite demonstrated knowledge of ecological relations, participants consistently rated inductive

generalizations between taxonomically related species stronger than generalizations between ecologically related species. This was true regardless of whether they were reasoning about a blank property, disease, or toxin. In other words, possessing requisite knowledge of ecological relations was not sufficient for the selective use of that knowledge to guide induction. As Shafto and Coley (2003) found, the property manipulation did not render ecological knowledge sufficiently available to novices.

In two subsequent experiments, we provided evidence that taxonomic knowledge was more available than ecological knowledge in this population. First, we contrasted primed and unprimed similarity judgments for pairs of taxonomically or ecologically related species (following Ross & Murphy, 1999, Experiment 4). If ecological knowledge was chronically less available, then priming ecological categories should increase the availability of ecological knowledge, resulting in increased similarity ratings for ecological pairs in the primed versus the unprimed condition. In contrast, if taxonomic knowledge was already highly available, then taxonomic priming should elicit no change in similarity ratings. As predicted, priming was found to increase similarity ratings for ecologically related pairs but not taxonomically related pairs, consistent with the suggestion that ecological knowledge is less available than taxonomic knowledge.

A second experiment provided further evidence by contrasting inductive judgments with and without time pressure. We predicted that time pressure would decrease access to less available knowledge by curtailing memory search, but would not effect use of knowledge that was already highly available. In line with this prediction, likelihood ratings for ecological inferences decreased under time pressure relative to unspeeded judgments, whereas ratings for taxonomic inferences remained unchanged.

These results show that for novices in the domain of biology, taxonomic knowledge is chronically more available than ecological knowledge and suggest that differences in availability may impede inductive selectivity. More generally, this series of studies is one example of how the notion of availability can be used to generate novel predictions about the use of knowledge in category-based induction.

Apart from generating testable hypotheses about inductive reasoning, another important challenge for the availability framework is to elucidate the mechanisms by which experience elicits changes in availability of different kinds of knowledge. One route through which chronic changes may be elicited is the frequency with which knowledge is accessed. Naturally, we expect information that gets used frequently in a particular context to be more available (Horowitz et al, 1966). This assumption reflects the fact that our past experience should provide a useful reference in the future, a basic principle of memory (Anderson & Milson, 1989). Extending this idea to availability merely implements this assumption as a means to sort out what knowledge is deemed appropriate for a particular inference given our experience and the current context.

A second potential mechanism for eliciting chronic change in availability is representational efficiency. Availability should increase with increased representational efficiency. Here representational efficiency reflects a compromise between the accumulation of facts and a means to summarize the facts efficiently. For example, one reason that taxonomic knowledge may be highly available for biological reasoning is the fact that it provides a succinct summary of a large amount of factual knowledge. Taxonomic knowledge encompasses genetic information, anatomical information, information about shape, behavior, environment, etc. On the other hand, knowledge

valued by experts such as ecological knowledge does not have the same immediate payoff that taxonomic knowledge does. Taxonomic structures provide a simple way of encompassing all living kinds in a representational structure that is scale-invariant; all subsections of the taxonomy are branching trees. Perhaps an increase in availability of ecological knowledge for experts represents an increase in the efficiency with which ecological relations capture relevant knowledge. However, these proposals are speculation that will require extensive empirical research.

We suggest that the problem of reasoning in knowledge-rich domains is crucial to understanding human intelligence. We have focused on one aspect of this problem, how experience is brought into contact with context in informing inductive reasoning. We have argued that in any given context, different kinds of knowledge are more or less available, and that availability predicts how likely the knowledge is to be used in reasoning. Though much of our evidence at this stage is preliminary, the notion of availability unites existing work on knowledge and context-specific reasoning, and may provide a useful framework in which to investigate how knowledge is deployed in specific situations to guide category-based induction.

References

Anderson, J.R. & Milson, R. (1989). Human memory: An adaptive perspective. *Psychological Review*, 96, 703-719.

Baker, A. & Coley, J.D. (2005). *Taxonomic and ecological relations in open-ended induction*. Paper presented at the 27th Annual Conference of the Cognitive Science Society, Stresa, Italy.

Barsalou, L.W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition, 10*, 82-93.

Choi, I., Nisbett, R.E. & Smith, E.E. (1997). Culture, category salience, and inductive reasoning. *Cognition, 65,* 15-32.

Coley, J.D. (2005). *Relational Properties in Folk Biological Induction*. Paper presented at the Biennial Meetings of the Society for Research in Child Development, Atlanta, GA.

Coley, J.D., Baker, A, & Kemp, C. (2004). *Taxonomic and Ecological Relations in Open-Ended Induction*. Paper presented at the 45th Annual Meeting of the Psychonomic Society, Minneapolis, MN.

Coley, J.D., Shafto, P., Stepanova, O, & Barraff, E. (2005). Knowledge and category-based induction. In Ahn, W., Goldstone, R. L., Love, B. C., Markman, A. B., & Wolff, P. (Eds.), *Categorization inside and outside the laboratory: Essays in honor of Douglas L. Medin*. Washington, DC: American Psychological Association. pp. 69-86.

Coley, J.D., Vitkin, A.Z., Seaton, C.E. & Yopchick, J.E (2005). *Effects of Experience on Relational Inferences in Children: The Case of Folk Biology*. Paper presented at the 27th Annual Conference of the Cognitive Science Society, Stresa, Italy.

Heit, E. & Bott, L. (2000). Knowledge selection in category learning. In D. Medin (Ed.), *Psychology of Learning and Motivation, 39,* 163-199. Academic Press.

Heit, E. & Feeney, A. (2005). Relations between premise similarity and inductive strength. *Psychonomic Bulletin & Review, 12,* 340-344.

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

Horowitz, L.M., Norman, S.A., & Day, R.S. (1966). Availability and associative symmetry. *Psychological Review*. 73, 1-15.

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

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

Medin, D. L., Lynch, E. B., Coley, J. D, & Atran, S. (1997). Categorization and reasoning among tree experts: Do all roads lead to Rome? *Cognitive Psychology*, *32*, 49-96.

Nguyen, S.P. & Murphy, G.L. (2003). An apple is more than just a fruit: Cross-classification in children's concepts. *Child Development, 6,* 1783-1806.

Nisbett, R.E., Peng, K., Choi, I. & Norenzayan, A. (2001). Culture and systems of thought: Holistic versus analytic cognition. *Psychological Review, 108,* 291-310.

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

Pearlstone, Z. & Tulving, E. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*. 5, 381-391.

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

Ross, B.H. & Murphy, G.L. (1999). Food for thought: Cross-classification and category organization in a complex real-world domain. *Cognitive Psychology, 38,* 495-553.

Shafto, P. & Coley, J.D. (2003). Development of categorization and reasoning in the natural world: Novices to experts, naïve similarity to ecological knowledge. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *29*, 641-649

Shafto, P., Coley, J.D. & Baldwin, D. (2005). *Availability in Inductive Reasoning*. Paper presented at the 46th Annual Meeting of the Psychonomic Society, Toronto.

Shafto, P., Kemp, C., Baraff, E., Coley, J.D. & Tenenbaum, J.B. (2005). *Context-sensitive Induction*. Paper presented at the 27th Annual Conference of the Cognitive Science Society, Stresa, Italy.

Shafto, P., Kemp, C., Baraff, E., Coley, J.D. & Tenenbaum, J.B. (in prep). Inductive reasoning with causal knowledge.

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

Smith, E. E., Shafir, E., & Osherson, D. N. (1993). Similarity, plausibility, and judgments of probability. *Cognition*, *49*, 67-96.

Stepanova, O. & Coley, J.D. (2003). *Animals and Alcohol: The Role of Experience in Inductive Reasoning among College Students*. Paper presented at the 44th Annual Meeting of the Psychonomic Society, Vancouver, British Columbia.

Stepanova, O. (2004). *Vodka and vermin: Naïve reasoning about animals and alcohol.* Unpublished doctoral dissertation, Northeastern University.

Tversky, A. & Kahnemann, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5, 207-232.

Vitkin, A.Z., Coley, J.D. & Hu, R. (2005). *Children's Use of Relevance in Open-Ended Induction in the Domain of Biology*. Paper presented at the 27th Annual Conference of the Cognitive Science Society, Stresa, Italy.

Vitkin, A., Coley, J.D. & Kane, R. (2005). Salience of Taxonomic and Ecological Relations in Children's Biological Categorization. Paper presented at the Biennial Meetings of the Society for Research in Child Development, Atlanta, GA.