

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Evaluating Two Mechanisms of Flexible Induction: Selective Memory Retrieval and Evidence Explanation

Permalink

<https://escholarship.org/uc/item/4qw5c1zb>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 35(35)

ISSN

1069-7977

Authors

Vasilyeva, Nadya
Coley, John

Publication Date

2013

Peer reviewed

Evaluating Two Mechanisms of Flexible Induction: Selective Memory Retrieval and Evidence Explanation

Nadya Y. Vasilyeva (Nvasil@Brandeis.Edu)

Department of Psychology, Brandeis University
415 South Street, Waltham MA 02454 USA

John D. Coley (j.coley@neu.edu)

Department of Psychology, Northeastern University
360 Huntington Avenue, Boston MA 02115 USA

Abstract

We report three studies examining mechanism of property-sensitive induction. First, we demonstrate that, contrary to a common assumption, property does not influence retrieval of knowledge about premise categories. Second, we introduce property-driven explanations as a possible source of property effects and provide first evidence for this proposal.

Keywords: induction; property effects; retrieval; explanation.

Generating hypotheses about uncertain outcomes from limited evidence – inductive inference – is a pervasive cognitive activity. In order to be successful, inductive inference must be flexible. For example, if you learn that a new *influenza virus* has been discovered in chickens, you may reasonably get concerned about your own health; but if chickens were announced to carry a certain *defective gene*, you are much less likely to worry about catching one during your next meal. Indeed, a vast body of empirical evidence demonstrates that people make systematically different inferences when they project different properties (see Coley & Vasilyeva, 2010, for a review). Heit and Rubinstein (1994) proposed that property affects induction by indicating different subsets of features as relevant for evaluating premise-conclusion similarity. Goodman (1972) provided a logical argument for constrained recruitment of features: since any category has a potentially infinite set of features and can be infinitely similar to any other category, it is a necessary logical requirement for inductive inference to impose some initial constraints to limit a subset of relevant features.

Although it is generally agreed that induction requires constrained recruitment of prior knowledge, and there is evidence that projected property may provide one such constraint (Coley & Vasilyeva, 2010; Heit & Rubinstein, 1994), the mechanism of property-based constrained recruitment remains unclear. Existing models of induction either do not specify the psychological mechanism of property effects (McDonald, Samuels, & Rispoli, 1996; Medin, Coley, Storms, & Hayes, 2003; Rips, 1975; Osherson, Smith, Wilkie, Lopez, & Shafir, 1990; Sloman, 1993; Sloutsky & Fisher, 2004), or acknowledge the computational nature of their account that may not correspond to actual psychological processes involved in inductive inference (e.g. Tenenbaum, Kemp, & Shafto, 2007; Heit, 1998).

We report three studies that examine two candidate psychological mechanisms of property effects in induction: property-moderated retrieval of relevant knowledge about premise categories from long-term memory, and generating explanations of why premise categories might have the property to begin with. In contrast to the majority of studies

on induction that use argument evaluation task, we employed inference *generation* task: participants were given an inductive premise and asked to generate their own conclusions. Coley & Vasilyeva (2010) demonstrated that this task provides a particularly sensitive measure of participants' spontaneous use of different kinds of relevant knowledge, in the context of an ecologically valid inductive problem.

Property-Moderated Knowledge Retrieval as a Mechanism of Property Effects

Generation of an inductive hypothesis is inherently knowledge driven; when one learns that A has a novel property X, one uses what they know about A and its relations to other things to form guesses about what else is likely to have X. One source of input to inductive inference is *knowledge about premise categories*. When such categorical knowledge is accessed, a probabilistically determined subset of features and relations that comprise the representation of that concept becomes available as a raw material for the inference. For example, if A turned out to be a duck, such features as “is a bird”, “flies”, “lives in ponds”, “quacks” and “eaten by foxes” may come to mind. Although there are many different types of knowledge, knowledge about living things is commonly divided into two broad classes: *taxonomic* knowledge is based on relations of intrinsic similarity between members, whereas contextual, or *ecological* knowledge, is based on extrinsic relations between members and other entities. For example, ducks belong to the taxonomic category of *birds* and ecological categories of *aquatic animals* and *fox prey*. Each of these types of knowledge can serve as a basis for an inductive projection from ducks – to other birds, or other aquatic animals, or things that eat ducks.

In addition to the premise category, *knowledge about the property* can also serve as a source of information. If X in the example above is replaced with a more specific property, such as “carries a certain disease” or “has a certain gene”, new knowledge is brought to the table: independently of what we know about ducks, we also know something about diseases and genes: what they are, whether they can be transmitted via contact, etc. How can property influence what projections people end up making? One possibility is that property constrains what types of premise knowledge are used to produce an inductive hypothesis.

A premise category label, as any word, is connected to a vast amount of conceptual knowledge; this knowledge is unlikely to be retrieved in its entirety on any given occasion (McElree, Murphy & Ochoa, 2006). Rather, retrieval of conceptual information from long term memory is selective and depends on context (e.g. Barsalou, 1982; Swinney,

1979). Within an inductive problem, property may serve as context for the premise category(-ies), and it could affect what information about these categories is retrieved. For example, in response to a premise like “ducks have gene X”, one may be more likely to retrieve taxonomic knowledge about ducks (bird, have feathers whereas for (“ducks have parasite X”), one may be more likely to retrieve ecological knowledge about ducks (aquatic, prey to foxes). This mechanism is consistent with the flexible similarity proposal by Heit and Rubinstein (1994) and the Bayesian model of induction by Heit (1998). The advantage of this proposal is that it is more specific, it focuses on describing the underlying psychological process rather than on modeling outcomes of such a process, and it can be tested with behavioral data.

To evaluate this proposal, we conducted two studies. Study 1 examined knowledge *recruitment*: how knowledge about premise categories predicts outcome inferences about different properties. Study 2 examined knowledge *retrieval*: how knowledge about premise categories is activated in context of different properties during inference generation.

Study 1: Property-Specific Knowledge Recruitment

To begin to specify the role of property in inference generation, we examined how it affects *recruitment* of category knowledge, or the extent to which available knowledge about premise categories ends up being used in the outcome inferences. For example, if among many facts about ducks, one knows that they live in water, and one uses this knowledge to project a property from ducks to other aquatic animals, we can say that the knowledge has been recruited. The question is whether the nature of the property affects the likelihood of recruiting ecological versus taxonomic knowledge about premise categories.

To address this question, we measured available knowledge about a set of animal categories and examined the predictive relationship between this knowledge and inferences generated about the same set of animal categories (i.e. knowledge recruitment). Most critically for evaluating the first proposal, we manipulated the property in the inference generation task between ecologically-biasing, neutral, and taxonomically-biasing. If property moderates recruitment of knowledge about premise categories, the predictive relationship between category knowledge and inferences should vary with the property. Based on Coley & Vasilyeva (2010), we expected that property could facilitate recruitment of congruent knowledge, and/or inhibit recruitment of incongruent knowledge. For example, taxonomic properties should facilitate recruitment of taxonomic knowledge, and inhibit recruitment of ecological knowledge; ecological properties should show the converse pattern.

Method

Feature-Listing Task Twenty nine participants were given a list of 42¹ animal names and for each animal were asked to write down anything they could think of that was “generally true of that animal”.

Inference-Generation Task One hundred participants were given 42 open ended-induction questions about same

42 animals; each stated that a property was true of a single animal species, and asked what other species were likely to have the property. For example, “*GENE T5 is found in DUCKS. What else is likely to have gene T5? Why?*” Property was manipulated within subjects; participants saw two examples of three kinds of properties: ecological (*flu, parasite*), taxonomic (*gene, cell*) and neutral (*substance, property*).² Each participant was presented with seven questions about each property—each with a different animal premise—in random order. The dependent variables were the frequencies of taxonomic and ecological inferences.

Results

Data Coding Four or five trained coders coded features and inferences into two broad classes. *Taxonomic* (Tax) features and inferences invoked category membership, perceptual features, or non-interactive aspects of behaviors and physiology (e.g. Tax-feature: “bird”; Tax-inference: “other birds will have it”). In contrast, *ecological* (Eco) features and inferences invoked animals’ diet, habitat, or other interactions with entities in their environment (e.g. Eco-feature: “lives in water”; Eco-inference: “other animals that live in water”). Each feature and inference was coded as taxonomic, or ecological, or both, or neither³. For every animal, the mean counts of features coded as Tax or Eco were taken as the measures of the amount of salient taxonomic and ecological knowledge about that animal. To quantify inferences, relative frequencies of participants making Tax and Eco inferences about that animal were calculated separately for each property type resulting in 6 means per animal.

Property effects Results showed strong property effects. Eco-inferences were generated most frequently when the property was ecological, followed by neutral and taxonomic properties ($F(2,82)=95.05, p<.001, \eta^2_p=.70$); this pattern was reversed for Tax-inferences ($F(2,82)=64.64^4, p<.001, \eta^2_p=.61$; all planned pairwise comparisons p 's<.001).

Knowledge recruitment We examined relations between premise category knowledge and property in predicting inferences in 12 simple linear regressions. In one triplet of regressions, ecological premise category knowledge served as a predictor of Eco-inferences, separately for ecological, neutral and taxonomic property. The other triplets covered the three remaining combinations between knowledge type and inference type, broken down by the property. The standardized regression coefficients are shown in Figure 1.

Eco-inferences For Eco-inferences (Fig. 1a) the predictive power of knowledge varied with the property. Eco-knowledge was overall a positive, albeit non-significant, predictor of Eco-inferences when participants were reasoning about a neutral ($R^2=.057, \beta=.239, p=.132$) or taxonomic ($R^2=.034, \beta=.242, p=.128$) property, but its

² The 6 properties were selected from a larger pool of properties based on a norming study measuring participants’ beliefs about the distribution of properties across taxonomic and ecological categories.

³ Coding categories were not mutually exclusive; a given response could receive multiple codes if it unambiguously invoked multiple codable kinds of relations.

⁴ Data were scored and analyzed by item. All the analyses on proportions reported below were conducted on arcsine-transformed data, while the reported means are non-transformed and presented as percentages.

¹ Feature-listing data from 1 animal were lost.

contribution was stronger in the presence of an ecological property ($R^2=.124$, $\beta=.352$, $p=.024$). When we examined the contribution of Tax-knowledge to Eco-inferences, overall larger amounts of Tax-knowledge were associated with lower frequency of Eco-inferences (all β 's are negative), and this relationship again varied with property. Tax-knowledge inhibited Eco-inferences marginally when the property was neutral ($R^2=.076$, $\beta= -.276$, $p=.08$), and reliably so when it was reinforced by a taxonomic property ($R^2=.12$, $\beta= -.346$, $p=.027$). Relative to taxonomic and neutral properties, ecological property largely neutralized the inhibitory effect of Tax-knowledge on Eco-inferences ($R^2=.007$, $\beta= -.083$, $p=.605$).

Tax-inferences Although, as shown in Fig. 1b, the sign and ordering of predictors generally follow the predicted pattern of strengthening effects of congruent knowledge and weakening effects of incongruent knowledge (with two exceptions), Tax-inferences were not significantly predicted by category knowledge (all p 's > .121).

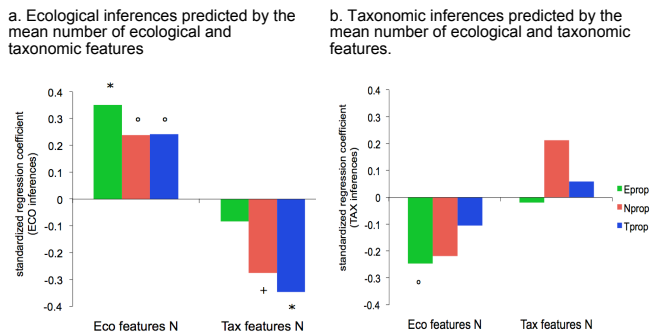


Figure 1: Eco- and Tax-knowledge about animals predicting relative frequency of Tax- and Eco-inferences about these animals, in the context of ecological, neutral and taxonomic properties. * $p < .05$, + $p < .1$, ° $p < .15$

Discussion

Overall, property had a profound effect on the inferences participants generated, and category knowledge predicted ecological, but not taxonomic inferences. Eco-inferences were facilitated by congruent (ecological) knowledge about premise categories, and inhibited by incongruent (taxonomic) knowledge. Most importantly, the relation between knowledge and Eco-inferences varied with the property: property strengthened effects of congruent knowledge, and weakened effects of incongruent knowledge. The overall pattern of congruent facilitation and incongruent inhibition held for Tax-inferences as well, although it was weaker and did not reach significance. Even though evidence of a relationship between knowledge and inferences was present for Eco-inferences but absent for Tax-inferences, it is sufficient in order to provide an “existence proof” for moderating effect of property on knowledge recruitment. These results are consistent with Heit and Rubinstein’s (1994) proposal about a flexible similarity metric, but they go beyond similarity relations and demonstrate flexible recruitment of ecological knowledge about contextual and interaction-based relations among animals.

Study 2: Property-Specific Knowledge Retrieval

Even though this demonstration of property effects in knowledge *recruitment* provides a useful constraint on the general underlying *retrieval* process, it does not specify it completely. Study 2 directly examines knowledge retrieval by measuring activation of premise category knowledge in real time, as it is accessed during inference generation. We borrowed a cross-modal priming paradigm from Swinney (1979). Participants were auditorily presented with a taxonomic, ecological or neutral property and an animal premise, and were asked to generate possible conclusions. In addition, upon hearing the property and animal, participants were presented with a lexical decision task involving targets related to salient taxonomic or ecological knowledge about the premise animal. For example, a participant might hear a property, *gene*, followed by the animal, *duck*, and, after a varying time interval, see on the screen a taxonomic target *bird*, or an ecological target *pond*, or an unrelated target *sofa*, or a non-word *soach*. The task was to decide whether the letter sequence was a word. The time to respond to the related targets was taken as a measure of activation of Tax- or Eco-knowledge about the premise animal. If knowledge about *duck* is activated, we expect decisions about related targets (*bird* and *pond*) to be faster than about unrelated targets (*sofa*). If property moderates knowledge retrieval, we would expect decisions about ecologically related targets (*pond*) to be faster in the presence of an ecological property and/or slower in the presence of a taxonomic property relative to a neutral, or non-biasing property context. Similarly, we would expect decisions about taxonomically related targets (*bird*) to be faster in the presence of a taxonomic and/or slower in the presence of an ecological property compared to neutral.

Method

Materials The stimuli for the induction task were 36 animal premises, each belonging to one salient taxonomic category (*mammal, bird, reptile, fish, insect*) and one salient habitat-based ecological category (*forest, desert, pond, ocean, savannah*). Each of the animals was presented in the context of an inductive problem about one of the six properties from Study 1 (*flu, gene, etc.*, presented with unique alphanumeric codes (X5, Z9)). All the animal names and properties were recorded in the voice of a female native speaker of English.

Thirty-six words and 36 pronounceable non-words were used as targets for the lexical decision task. The targets (derived from feature responses in Study 1) were taxonomically related, ecologically related, or unrelated to the animals used as premise categories in the induction task. The strength of association of the taxonomic, ecological, and unrelated targets to the corresponding animals, as well as lack of direct associations between properties and target words, were normed with another group of 18 native speakers of English.

Participants and Procedure One hundred eleven native speakers of English were tested individually, on a MacBook laptop running Superlab 4.0.4 software and set up with headphones and a microphone. The experiment consisted of an open-ended induction task with intervening lexical decision task. Participants were instructed that they would be listening to utterances that would introduce a property, followed by an animal that possesses that property, and their

task was to say out loud (at a cued moment) other species likely to share that property, along with a short justification. Participants were also informed that at “random” moments a sequence of characters would appear on the screen, and their task was to identify it as a word or a non-word as quickly as possible without sacrificing accuracy, using the response buttons. Each trial began with a 2000msec pause; then a participant heard the property to be projected (e.g., *flu M3*), followed by a pause of 1000msec, followed by the name of the premise animal (e.g., *bear*), followed by a pause and a signal to start speaking. At varying SOAs (stimulus onset asynchrony: 400, 900, or 1650msec from the onset of the animal name), a target for lexical decision appeared on the screen and stayed there until the participant responded or for 3500msec, whichever came first. No accuracy feedback was provided. After a 2000msec pause following the participant’s response or the end of lexical decision target presentation, a short beep signaled that the participant could start saying their inference. Participants had 15 seconds to say their response, after which the experiment automatically moved on to the next trial.

Design The main independent variable was property type (taxonomic: *gene, cell*; ecological: *flu, parasite*; neutral: *substance, property*). The second independent variable was the target word type (taxonomic vs. ecological vs. unrelated). Each non-filler animal was yoked to one target word type (taxonomic, ecological, or unrelated). We also varied SOA, but for the sake of brevity, this manipulation will not be discussed here. All the reported analyses were collapsed across SOA.

Results

Does property moderate retrieval of knowledge about premise animals? If so, property type should facilitate responses to property-congruent targets and/or interfere with responses for property-incongruent targets.

Accuracy A 3(target type: eco, tax, unrelated) x 3(property: eco, tax, neutral) repeated measures ANOVA on accuracy in lexical decision task showed no main effect of property ($F(2,220)=.145, p=.865$). The effect of target type was significant ($F(2,220)=4.33, p=.014$): eco- and tax-targets were verified more accurately than unrelated targets ($t(110)=2.63, p=.010$; $t(110)=2.58, p=.011$); the former two did not differ ($t(110)=.33, p=.744$). Most importantly, the effect of target type did not interact with property ($F(4,440)=.57, p=.683$). This suggests that participants were retrieving category-relevant knowledge, but that such retrieval was not moderated by property.

Reaction Time. Reaction time results were consistent with the accuracy analyses. A 3 (target type: eco, tax, unrelated) x 3 (property: eco, tax, neutral) repeated measures ANOVA on RT showed no main effect of property ($F(2,218)=.44, p=.656$) and a significant effect of target type ($F(2,218)=4.73, p=.010$): eco-targets were verified faster than unrelated targets ($t(110)=3.20, p=.002$); and tax-targets were verified marginally faster than unrelated targets ($t(110)=1.80, p=.074$); the former two did not differ ($t(110)=1.46, p=.148$). Again, most importantly, the effect of target type did not interact with the property ($F(4,436)=1.24, p=.293$), suggesting that property does not moderate knowledge retrieval.

Surprisingly, the speed of lexical decisions about filler items (non-words) was affected by the property ($F(2,220)=14.95, p<.001$): decisions were slower in the

presence of ecological (1343msec, $t(110)=4.88, p<.001$) and taxonomic (1320msec, $t(110)=3.99, p<.001$) than neutral property (1253msec); the former two did not differ ($t(110)=1.45, p=.151$).

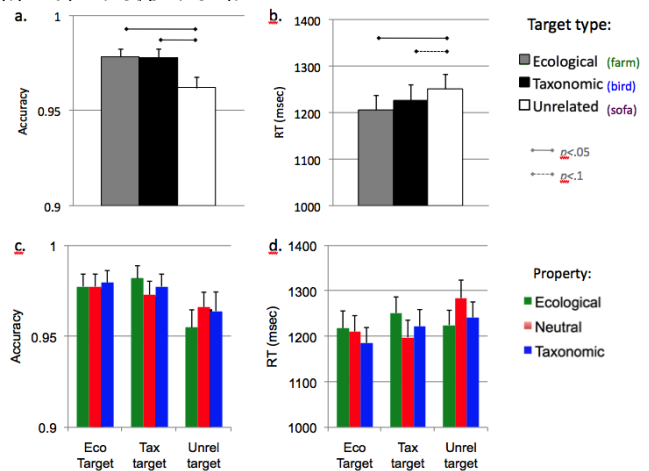


Figure 2: Effect of target type and property on lexical decision accuracy and reaction time. Verification of related targets (eco, tax) was more accurate (a) and faster (b) than unrelated targets. This effect did not depend on property (panels c and d). Error bars: 1 SEM.

Discussion

Based on the results of Study 1, we expected to find effects of property on knowledge retrieval: specifically, facilitation of property-congruent knowledge and inhibition of property-incongruent knowledge. However, we found no evidence of property moderating knowledge retrieval. Of course, we cannot completely rule out the possibility that the lack of property effects was caused by some procedural flaws of the study. However, because we *did* see property effects on some lexical decisions about filler words, we know that the method is in principle capable of detecting property effects. And because the premise category *did* differentially prime related vs. unrelated targets, we know that the method is capable of detecting differential priming. Therefore, it is likely that we failed to see property effects on retrieval because *property does not moderate retrieval of knowledge* about premise categories from long-term memory.

Study 3: Property-Moderated Explanation as a Mechanism of Property Effects

Study 1 demonstrated selective property-moderated *recruitment* of categorical knowledge to inform inferences. However, Study 2 found no moderating effects of property on knowledge *retrieval* in real time. If, as we argue, this finding reflects the actual absence of property effects on retrieval rather than an experimental failure, we need to look for another mechanism whereby property can guide selective recruitment of taxonomic and ecological information by inferences. The mechanism that we examine in Study 3 is based on *property-moderated explanation of evidence*.

As suggested by Sloman’s (1994) work, explanation of evidence may affect evaluation of inductive arguments. Even when the similarity between premise and conclusion is held constant, if both can be explained by reference to the

same principle, the perceived strength of an inductive argument can be higher than when premise and conclusion statements require different explanations.

Several features of explanation make it a good candidate mechanism for property effects in induction. For instance, explanation is flexible: there are multiple ways to explain any given observation. A *formal* explanation refers to categories or inherent properties; a *causal* explanation refers to the proximal mechanisms of change; and a *teleological* explanation refers to ends, goals or functions (Lombrozo, 2006). On the subsumption account proposed by Williams & Lombrozo (2010), explaining an observation involves identifying a larger pattern of which the observation is a part. In this sense, explanation identifies a relevant subset of knowledge about the observation that can serve as a basis for generalizing to other cases – thus satisfying the logical prerequisite for induction stipulated by Goodman (1972).

How might explanation provide a mechanism for property effects in induction? If different properties lend themselves to different explanations, and if explaining consists of identifying an observation as a part of a larger pattern or regularity, then different properties might determine whether a premise of an inductive argument is viewed as a part of one regularity or another (e.g. formal explanations might highlight taxonomic relations, whereas causal explanations might highlight ecological relations). Thus, construction of different explanations could engender differential *recruitment* of knowledge, and ultimately different hypotheses about how a property might generalize without the necessity of differential *retrieval* of knowledge.

In Experiment 1, although asked to explain why they generated particular conclusion categories, participants often spontaneously provided explanations for why a premise category exhibited a given property. To evaluate the explanation mechanism, we examined these spontaneous explanations to determine whether different properties were associated with different types of explanations. We expected taxonomic properties to provoke predominantly formal explanations referring to classes of objects (that would eventually translate into category-based, or taxonomic inferences) and ecological properties to lead to predominantly causal explanations, referring to interactions between animals and/or their environment (that would eventually translate into ecological inferences). We had no specific predictions about teleological explanations.

Method

Three trained coders independently re-coded all the inferences collected in Study 1 for the presence of formal, causal and teleological explanations. Twelve percent of inferences (467 out of 3920 codable responses) contained spontaneous explanations. These explanations were coded as *formal* (explanations that referred to kind membership, e.g. “mammal gene”, or “this is a bird flu”), *causal* (explanations describing a “story” of interactions between animals and other entities, or a sequence of events resulting in the premise category having the property, e.g. “vultures may get flu E5 from the dead and decaying animals they feed off of” or “[the gene will be found in] fish since pelicans eat them, the pelicans might develop that gene from the fish”), *teleological* (explanations referring to goals, functions or purposes of properties, e.g. “these cells are to protect them from the cold” or “perhaps B6-cells defend deer from particular viruses that they are exposed to”), or

other (idiosyncratic or vague explanations that could not be assigned to any of the three categories).

Results

Scoring. For each animal, we calculated the percentage of subjects who generated each type of explanation out of total number of participants who reasoned about that animal, separately for each property type. This yielded 12 percentages, or relative frequency scores, per animal (3 property types x 4 explanation types), that were arcsine-transformed for the analyses. Uncodable explanations were rare (less than 2% of participants per animal) and were excluded from the following analyses.

Analyses To provide support for the proposal that property affects inferences via explanation, we need to show that different properties are associated with different explanations, and that different explanations are associated with different inferences.

Relations between property and explanation. The main question, whether different properties trigger different types of explanations, was addressed by a 3 (property: eco, tax, neutral) x 3 (explanation: formal, causal, teleological) ANOVA on relative frequency of explanations. The overall likelihood of providing an explanation did not vary with the property ($F(2, 82)=.044, p=.957$). However, explanations differed in frequency ($F(2,76)=18.836, p<.001, \eta^2_{part}=.315$): causal explanations were more frequent (5.3%) than formal (2.8%, or teleological explanations (2.4%, $t(41)\geq 4.59, p<.001, d\geq .71$), which did not differ from each other.

Of most theoretical interest was the significant interaction between property and explanation type ($F(4,164)=34.442, p<.001, \eta^2_p=.457$, see Fig 3a). Explanations clearly varied with property: for ecological properties, causal explanations were more frequent than formal explanation, which were more frequent than teleological explanations ($t(41)\geq 4.52, p<.001, d\geq 0.70$). For neutral properties, causal explanations were also more frequent than formal or teleological explanations, which did not differ from each other ($t(41)\geq 3.00, p\leq .005, d\geq 0.45$). In contrast, for taxonomic properties, formal and teleological explanations were more frequent than causal explanations ($t(41)\geq 3.42, p\leq .001, d\geq 0.52$). These results demonstrate a link between property and explanation type.




Figure 3: a. Percentage of explained inferences involving causal, formal, and teleological explanations for ecological, neutral and taxonomic properties. b. Percentage of Tax- and Eco-inferences for responses with formal, teleological, and causal explanations.

Relations between explanation and inference To examine the link between explanation and inference, we again focused on the subset of responses from Study 1 that included spontaneous explanations. We calculated the percentage of Tax- and Eco-inferences that accompanied

each type of explanation. There was a clear association between explanation type and inference (Fig. 3b): responses that included formal or teleological explanations were much more likely to result in Tax- than Eco-inferences ($t(41) \geq 4.38$, $p < .001$, $d \geq 0.68$). In contrast, responses that included causal explanations were much more likely to include Eco- than Tax-inferences, $t(41) = 8.61$, $p < .001$, $d = 1.33$). This systematic relationship between explanation type and inferences, taken together with the evidence for the relationship between property and explanation type, is consistent with the proposal that explanations mediate the effect of property on inferences.

Discussion

To examine whether explanations might moderate property effects in induction, we asked whether different properties were associated with different explanations, and then whether different explanations were associated with different inferences. We have answered both questions in the affirmative. First, different properties triggered different types of explanations. When participants were reasoning about ecological properties, the majority of explanations they provided were causal, referring to a mechanism that could have endowed the animal with the property (e.g., “Owls eat mice and could contract the flu from the mice that it eats”). In contrast, when participants were reasoning about taxonomic properties, they were less likely to use causal explanations, preferring formal explanations (“this cell could be specific to jaguars”) or teleological explanations (“T5 is something to keep them warm”). Second, different explanations were associated with different inferences. Causal explanations were more likely to accompany ecological inferences, whereas formal and teleological explanations were more likely to accompany taxonomic inferences.

These findings are consistent with the idea that explanations serve as a mediator between properties and inferences. We acknowledge that these analyses are correlational, and therefore do not provide direct evidence that explanations play a causal role in property-specific inductive inference. Nevertheless, an informal comparison of effect sizes indicates that the mean effect of explanations on inferences ($d = 1.18$) is larger than the mean effect of properties on inferences ($d = 0.92$). This suggests that property-driven explanations are likely to affect inferences directly, rather than being a mere correlate of properties.

General Discussion

We provided evidence that property effects do not take place in retrieval. This questions the existing, but not tested, assumption in the field about the mechanism of property effects based on context-dependent retrieval of information from semantic memory (Heit & Rubinstein, 1994). We also provided some promising evidence that property effects may result from participants generating explanations for the presence of the property in the premise category. If this finding persists, it could strengthen connections between research on explanations and on induction. Most researchers agree that these are related, but very little supporting empirical work exists, although it is increasingly acknowledged that the presence of an available explanation can reduce reliance on overall similarity and override effects of similarity and diversity on induction (see Lombrozo, 2006, for a review). In this project we demonstrated that

explanations do not just “mess up” existing regularities in induction, but may in fact be an important part of the mechanism of one such established regularity – property effects in induction.

To sum up, this project makes a step towards specifying the mechanism of property effects in induction in two ways. First, it suggests that property effects do *not* work via property-based retrieval of knowledge about premise categories from memory. Second, it introduces property-driven explanations as a possible source of property effects. Of course, these proposals are not mutually exclusive, and our main suggestion for the further research would be not to abandon studying knowledge retrieval in induction, but to expand research on the mechanism of property effects to include explanations.

References

- Barsalou, L.W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition*, *65*, 15-32.
- Coley, J.D. & Vasilyeva, N.Y. (2010) Generating inductive inferences: premise relations and property effects. In Ross, B. (Ed.) *The Psychology of Learning and Motivation* (53). Burlington: Academic Press.
- Goodman, N. (1972). Seven structures on similarity. In *Problems and Projects*. New York, NY: Bobbs-Merrill.
- Heit, E. (1998). A Bayesian analysis of some forms of inductive reasoning. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition*. Oxford: Oxford University Press.
- Heit, E. & Rubinstein, J. (1994). Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 411-422.
- Lombrozo, T. (2006). The structure and function of explanations. *Trends in Cognitive Sciences*, *10*(10), 464-470.
- McDonald, J., Samuels, M., & Rispoli, J. (1996). A hypothesis assessment model of categorical argument strength. *Cognition*, *59*, 199-217.
- McElree, B., Murphy G.L., & Ochoa T. (2006). Time course of retrieving conceptual information: A speed-accuracy trade-off study. *Psychonomic Bulletin & Review*, *13*, 848-853.
- Medin, D. L., Coley, J. D., Storms, G. & Hayes, B. K. (2003) A relevance theory of induction. *Psychonomic Bulletin & Review*, *10*, 517-532.
- Osherson, D.N., Smith, E.E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological Review*, *97*, 185-200.
- Rips, L. J. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning & Verbal Behavior*, *14*, 665-681.
- Slovan, S.A. (1993). Feature-based induction. *Cognitive Psychology*, *25*, 213-280.
- Slovan, S.A. (1994). When explanations compete: the role of explanatory coherence on judgments of likelihood. *Cognition*, *52*, 1-21.
- Sloutsky, V.M., & Fisher, A.V. (2004). Induction and categorization in young children: A similarity-based model. *Journal of Experimental Psychology: General*, *133*, 166-188.
- Swinney, D. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning & Verbal Behavior*, *18*, 645-659.
- Tenenbaum, J. B., Kemp, C., and Shafto, P. (2007). Theory-based Bayesian models of inductive reasoning. In Feeney, A. & Heit, E. (Eds.), *Inductive reasoning*. Cambridge University Press.
- Williams, J.J. & Lombrozo, T. (2010). The role of explanation in discovery and generalization: evidence from category learning. *Cognitive Science*, *34*, 776-806.