Cognition 109 (2008) 175–192

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

Inductive reasoning about causally transmitted properties ☆

Patrick Shafto [a,*], Charles Kemp [b], Elizabeth Baraff Bonawitz [c], John D. Coley [d], Joshua B. Tenenbaum [c]

[a] Department of Psychological and Brain Sciences, University of Louisville, 317 Life Sciences, Louisville, KY, USA
[b] Department of Psychology, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, USA
[c] Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA, USA
[d] Department of Psychology, Northeastern University, 360 Huntington Ave, Boston, MA, USA

ARTICLE INFO

Article history:
Received 31 August 2007
Revised 16 May 2008
Accepted 12 July 2008

Keywords:
Property induction
Inductive reasoning

ABSTRACT

Different intuitive theories constrain and guide inferences in different contexts. Formalizing simple intuitive theories as probabilistic processes operating over structured representations, we present a new computational model of category-based induction about causally transmitted properties. A first experiment demonstrates undergraduates' context-sensitive use of taxonomic and food web knowledge to guide reasoning about causal transmission and shows good qualitative agreement between model predictions and human inferences. A second experiment demonstrates strong quantitative and qualitative fits to inferences about a more complex artificial food web. A third experiment investigates human reasoning about complex novel food webs where species have known taxonomic relations. Results demonstrate a double-dissociation between the predictions of our causal model and a related taxonomic model [Kemp, C., & Tenenbaum, J. B. (2003). Learning domain structures. In *Proceedings of the 25th annual conference of the cognitive science society*]: the causal model predicts human inferences about diseases but not genes, while the taxonomic model predicts human inferences about genes but not diseases. We contrast our framework with previous models of category-based induction and previous formal instantiations of intuitive theories, and outline challenges in developing a complete model of context-sensitive reasoning.

1. Introduction

Any familiar thing can be thought about in a multitude of ways. A cat is a creature that climbs trees, eats mice, has whiskers, belongs to the category of felines, and was revered by the ancient Egyptians. Knowledge of all of these kinds plays an important role in inductive inference. If we learn that cats suffer from a recently discovered disease, we might think that mice also have the disease – perhaps the cats picked-up the disease from something they ate. Yet if we learn that cats carry a recently discovered gene, lions and leopards seem more likely to carry the gene than mice. Flexible inferences like these are a hallmark of human reasoning, which is notable for the selective application of different kinds of knowledge to different kinds of problems.

Psychologists have confirmed experimentally that inductive inferences vary depending on the property involved. When adults are told about genes or other internal anatomical properties, they tend to generalize to taxonomically related categories (Osherson, Smith, Wilkie, L'opez, & Shafir, 1990). When told about novel diseases, however, adults may generalize to categories related by a causal mechanism of disease transmission, such as a food web

☆ Preliminary versions of this work were presented at the International Conference on Thinking 2004 and the 27th annual meeting of the Cognitive Science Society. This work was supported by the Albert Fellowship to C.K., NSF Grant BCS-0236338 to J.D.C., and the Paul E. Newton Chair to J.B.T. The authors thank Vikash Mansinghka for helpful comments and discussion.

* Corresponding author.
E-mail address: p.shafto@louisville.edu (P. Shafto).

0010-0277/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.cognition.2008.07.006

(Shafto & Coley, 2003). Across development, children demonstrate increasingly distinct patterns of inference for properties such as drinking versus riding (Mandler & McDonough, 1996, 1998a, 1998b), anatomic versus transient properties (Gelman & Markman, 1986), and anatomy versus beliefs (Springer, 1996; Solomon, Johnson, Zaitchik, & Carey, 1996). Psychologists have also suggested, at least in principle, how complex inferences like these might work. Flexible inductive inferences are supported by *intuitive theories* (Murphy & Medin, 1985; Carey, 1985; Keil, 1989), or "causal relations that collectively generate or explain the phenomena in a domain" (Murphy, 1993). In any given domain, more than one theory may apply, and different patterns of inference will be observed depending on which theory is triggered.

Although a theory-based approach is attractive in principle, formalizing the approach is a difficult challenge. Recent work by Kemp and Tenenbaum (2003) has proposed a model for taxonomic theories. Here we describe and test a Bayesian theory-based model of induction about causally transmitted properties. This new model is a rational analysis of reasoning about causal transmission in the sense of Anderson (1990). The model consists of two parts: a generative theory that defines prior beliefs, and Bayesian inferential machinery that generalizes novel concepts by combining observed examples with prior beliefs.

We begin by discussing the problem of context-sensitive induction, and explain why theories and causal knowledge are important to understanding context-sensitive induction. We then present our model of causal property induction and the Bayesian framework for theory-based inference. A first experiment investigates undergraduates' reasoning about species with familiar taxonomic and food web relations, demonstrating qualitative fits between model predictions and human inferences. A second experiment shows that the model predicts human inferences about the distribution of diseases over a more complex artificial food web. In a third experiment, we contrast the fits of causal and taxonomic models to human generalizations of diseases and genes over known species, showing that the causal model predicts inferences about diseases but not genes, and the taxonomic model predicts inferences about genes but not diseases. Finally, we discuss our contributions to understanding the relationship between prior knowledge and reasoning, and outline challenges in developing a full model of context-sensitive induction.

2. Context-sensitive induction

In category-based induction tasks (Rips, 1975), participants are given one or more examples of categories that have a novel property. For example, participants may be told, "Lions have gene XR-35", where the property is gene XR-35, and lions are one example of things that have the property. Participants are then asked to judge the probability that other categories have the property; for example, "How likely is it that tigers have gene XR-35, like lions?" Properties are chosen such that participants have no specific knowledge about which categories have the properties, and predictions must be generated based on prior knowledge about the kind of property and the categories in question. Many elements of the context may influence reasoning in these tasks. Several important sources of context are the property being generalized, the sampling of example categories, instructions, and general demand characteristics. In this paper, we are concerned with the effects of different kinds of properties on inductive generalization.

Research has confirmed that the properties used strongly influence the inductive inferences of both children and adults. For example, Gelman and Markman (1986) found that 4-year-old children generalize internal anatomical and behavioral properties ("has cold blood") but not idiosyncratic properties ("gets cold at night") between members of the same category. Working with adults, Heit and Rubinstein (1994) showed that inferences differ when reasoning about behavioral versus anatomical properties. For example, participants were more willing to generalize between taxonomically matched species such as bears and whales when reasoning about properties such as "has a liver with two chambers that act as one". However, when reasoning about a behavioral property such as "usually travels in a back-and-forth, or zig-zag, trajectory", participants were more willing to generalize between behaviorally matched species such as tuna and whales. More recent research has shown that fishermen generalize diseases, but not properties, over food web relations, with inferences being stronger from prey to predators than from predators to prey (Shafto & Coley, 2003). These experimental examples underscore the importance of properties in inductive reasoning.

Previous models of property induction have had difficulty explaining sensitivity to context. Consider first the similarity-coverage model (Osherson et al., 1990), the best known model of category-based inductive reasoning. It predicts inferences about novel properties based on similarities between pairs of categories and a hierarchy of taxonomic relations among categories. The model makes accurate predictions about human generalizations in default contexts, when people are reasoning about generic biological properties that seem to refer to anatomy or physiology. However, accounting for inferences about anatomical and behavioral properties such as those in Heit and Rubinstein (1994) would require extending the model to allow context-sensitive notions of similarity. Even if this amendment is allowed, similarity-based approaches cannot naturally account for the causal asymmetries demonstrated in Shafto and Coley (2003) because ratings of similarity between predators and prey do not show strong asymmetries (see also Medin, Coley, Storms, & Hayes, 2003). To be fair, the similarity-coverage model was not designed with multiple contexts in mind; nevertheless, any comprehensive model of category-based induction will have to deal with the general phenomenon of context-sensitive reasoning, and reasoning about causally transmitted properties in particular.

Sloman (1993) proposed a more flexible feature-based approach to modeling property induction. Instead of appealing to stable notions of similarity or taxonomy, Sloman posits that each category is represented by a large, potentially context-sensitive, set of features. The strength