Cognitive Psychology 64 (2012) 35–73

Contents lists available at SciVerse ScienceDirect

Cognitive Psychology

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

An integrated account of generalization across objects and features ☆

Charles Kemp [a,*], Patrick Shafto [b], Joshua B. Tenenbaum [c]

[a] Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213, United States
[b] Department of Psychological and Brain Sciences, University of Louisville, Louisville, KY 40292, United States
[c] Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, United States

ARTICLE INFO

Article history:
Accepted 9 October 2011
Available online 15 November 2011

Keywords:
Generalization
Property induction
Causal reasoning
Knowledge integration

ABSTRACT

Humans routinely make inductive generalizations about unobserved features of objects. Previous accounts of inductive reasoning often focus on inferences about a single object or feature: accounts of causal reasoning often focus on a single object with one or more unobserved features, and accounts of property induction often focus on a single feature that is unobserved for one or more objects. We explore problems where people must make inferences about multiple objects and features, and propose that people solve these problems by integrating knowledge about features with knowledge about objects. We evaluate three computational methods for integrating multiple systems of knowledge: the output combination approach combines the outputs produced by these systems, the distribution combination approach combines the probability distributions captured by these systems, and the structure combination approach combines a graph structure over features with a graph structure over objects. Three experiments explore problems where participants make inferences that draw on causal relationships between features and taxonomic relationships between animals, and we find that the structure combination approach provides the best account of our data.

☆ An early version of this work was prepared in collaboration with Allison Berke and presented at the Neural Information Processing Systems conference in 2006. We thank Keith Holyoak and three anonymous reviewers for helpful comments on the manuscript, and Bob Rehder and Brian Milch for valuable discussions. This work was supported in part by NSF award CDI-0835797, the Pittsburgh Life Sciences Greenhouse Opportunity Fund, AFOSR MURI contract FA9550-05-1-0321, the William Asbjornsen Albert memorial fellowship (CK) and the Paul E. Newton Chair (JBT).

* Corresponding author. Address: Department of Psychology, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, United States.

E-mail addresses: ckemp@cmu.edu (C. Kemp), p.shafto@louisville.edu (P. Shafto), jbt@mit.edu (J.B. Tenenbaum).

doi:10.1016/j.cogpsych.2011.10.001

1. Introduction

Will that berry taste good? Is that table strong enough to sit on? Questions like these require a reasoner to predict whether an object has a feature that has not yet been observed. Two versions of this basic inductive challenge can be distinguished. *Across-object generalization* is a problem where a reasoner observes one or more objects that have a given feature (e.g. Tim has feature F) then decides whether other objects have the same feature (does Tim's twin brother Tom have feature F?). *Across-feature generalization* is a problem where a reasoner observes one or more features of a given object (e.g. Tim is obese) then makes inferences about other features of the same object (does Tim have diabetes?). These two generalization problems form a natural pair, and both can be viewed as inferences about the missing entries in a partially-observed object-feature matrix. Fig. 1 shows an example where the objects are animals of different species and the features are biological or behavioral attributes. Because the mouse and the rat are similar, observing that the mouse has gene X suggests that the rat is likely to carry the same gene (across-object generalization). If gene X causes enzyme Y to be expressed, then observing that the mouse has gene X suggests that the mouse is likely to express enzyme Y (across-feature generalization).

Across-object and across-feature generalization are typically studied in isolation but these two forms of generalization often interact. For example, given that Tim is obese, we might predict that Tim's twin brother Tom is more likely to have diabetes than an unrelated individual called Zach. This prediction appears to rely on across-object generalization (since Tim is obese, Tom is likely to be obese) and on across-feature generalization (if Tom is obese, then Tom is likely to have diabetes). Similarly, if we learn that the mouse in Fig. 1 carries gene X and that gene X causes enzyme Y to be expressed, we might predict that the rat is likely to express enzyme Y (Fig. 1c). Both of these predictions can be formulated as inferences about the missing entries in an object-feature matrix. We develop an account of generalization that handles inferences of this kind, and that includes both across-object and across-feature generalization as special cases.

Our approach is based on the idea of integrating multiple knowledge structures. An *object structure* can capture relationships among objects—for example, a structure defined over the three individuals previously introduced can indicate that Tim and Tom are more similar to each other than either is to Zach. A *feature structure* can capture relationships between features—for example, one feature struc-

(a) Across-object generalization

The mouse has gene X.

The rat has gene X.

(b) Across-feature generalization

The mouse has gene X.

The mouse has enzyme Y.

(c) Generalization across objects and features

The mouse has gene X.

The rat has enzyme Y.

(d)

	has sharp teeth	gnaws wood	climbs trees	is white	has gene X	has enzyme Y	...
mouse	1	1	0	1	1	?	
rat	1	1	0	0	?	?	
sheep	0	0	0	1	?	?	
squirrel	?	1	1	0	?	?	
⋮	?	?	?	?	?	?	

Fig. 1. Generalization problems involving a set of animals and their features. (a) Across-object generalization is a problem where a reasoner makes inferences about the distribution of a single feature—here "has gene X". The example shown is a one premise argument: given that the statement above the line is true, the reasoner must decide whether the statement below the line is likely to be true. (b) Across-feature generalization is a problem where a reasoner makes inferences about the features of a single object. The argument shown here is strong if gene X is known to cause enzyme Y to be expressed. (c) Generalization problems may require a reasoner to generalize across both objects and features. Here the reasoner is told that a given animal (the mouse) has a given feature (gene X), then asked to decide whether a different animal (the rat) has a different feature (enzyme Y). (d) Generalization can be formalized as the problem of filling in the missing entries in an object-feature matrix. The three problems in (a)–(c) are all special cases of this matrix completion problem.