Journal of Experimental Child Psychology 152 (2016) 335–342

Contents lists available at ScienceDirect

Journal of Experimental Child Psychology

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

Brief Report

Learning from non-representative instances: Children's sample and population predictions

Jordan T. Thevenow-Harrison *, Charles W. Kalish

Department of Educational Psychology, University of Wisconsin–Madison, Madison, WI 53706, USA

ARTICLE INFO

Article history:
Available online 27 August 2016

Keywords:
Probability learning
Inductive inference
Cognitive development

ABSTRACT

What do children learn from biased samples? Most samples people encounter are biased in some way, and responses to bias can distinguish among different theories of inductive inference. A sample of 67 4- to 8-year-old children learned to make conditional predictions about a set of sample items. They then made predictions about the properties of new instances or old instances from the training set. The experiment compared unbiased and biased sampling. Given unbiased samples, participants used what they learned to make predictions about population and sample instances. With biased samples, children were less accurate/confident about inferences about the population than about the sample. Children used information in a biased sample to make predictions about items in that sample, but they were less likely to generalize to new items than when samples were unbiased.

Introduction

Learning from examples is both basic and central to forming knowledge of the world. There is significant debate about how best to characterize the psychological processes of learning from examples and how to understand potential developmental changes in such processes (Colunga & Smith, 2008; Griffiths, Chater, Kemp, Perfors, & Tenenbaum, 2010; Xu, 2007). Some argue that learning is a process of forming associations (McClelland et al., 2010), whereas others argue that learning is best understood as evidential inference (Griffiths et al., 2010; Kalish & Thevenow-Harrison, 2014). Both positions

* Corresponding author.
E-mail address: jtth@education.wisc.edu (J.T. Thevenow-Harrison).

http://dx.doi.org/10.1016/j.jecp.2016.07.013
0022-0965/

are well characterized as "statistical learning." Associative accounts focus on learning descriptive statistics; how do people notice patterns of co-occurrence in distributions of examples? Evidential accounts focus on learning inferential statistics; how do people use examples as evidence to draw conclusions about generative processes or populations (see Kalish, Kim, & Young, 2012)? Both associative and evidential accounts imply that people are sensitive to distributional features in samples. The accounts differ, however, in their treatment of the differences between samples and populations.

Consider the following situation. A fisherwoman catches a collection of fish that vary in size (big and small) and color (green and red). As she catches fish, she notices that all the big ones are red and all the small ones are green. How she notices and mentally represents the patterns (e.g., memory for specific examples, prototypes, conditional probabilities) are questions about the descriptive statistics she learns. Distributional features of the caught fish such as the co-occurrences of sizes and colors, the number of examples, and the salience of the different attributes will affect the patterns she notices. Associative and evidential accounts agree here, but they diverge in characterizing what happens when the fisherwoman tries to use the patterns she noticed to make predictions. Prediction is a matching process in associative accounts; if she sees a big fish and wants to predict its color, she can match the target fish to the fish in her creel and return a judgment (e.g., it is more like the red fish than the green fish). In evidential accounts, prediction is an inference to a generative process; if she sees a big fish, she considers what her experience has led her to conclude about the population of big fish. Critically, her conclusion may depend on two inferential features of the problem. The first inferential feature is the prediction target; is she making a prediction about a fish from the creel, about another fish from the pond she was fishing in, or about a fish from some other pond? The second feature is the representativeness of her sampling strategy; was she using a particular kind of bait or fishing from a location that might have biased her sample? Whether or how such inferential features affect predictions is where associative and evidential accounts diverge.

The specific evidential prediction is an interaction between prediction target and representativeness of sampling strategy. Strategy should matter for inference about new items from the population but not for inferences about old items from the sample. If the fisherwoman is making a prediction about a fish from her creel, the sampling strategy that produced the fish (e.g., the bait, the location) is irrelevant. One can guess that a big fish selected from the creel will be red without worrying about how that fish ended up in the creel in the first place. In contrast, predicting that a new big fish from the pond will be red does depend on the representativeness of the original sample. If the sampling strategy is biased, one cannot (easily) use the sample of caught fish to make a prediction about a new fish. This interactive prediction was the focus of the empirical study reported below; will children use information about samples to make predictions about old (sample) examples but not new (population) examples when samples are non-representative but use sample information for both old and new items when samples are representative?

Previous research suggests that young children are sensitive to both prediction target and sampling strategy. Several studies have found that young children and infants expect populations to resemble samples when the sampling strategy is representative (e.g., blindfolded selection) but not when it is not (e.g., deliberate selection) (Xu & Denison, 2009; Xu & Tenenbaum, 2007). Young children are able to use a mismatch between sample and population to infer non-representative sampling (Kushnir, Xu, & Wellman, 2010). Kalish and colleagues (2012) found that young school-aged children were better able to make predictions based on two distinct samples (e.g., fish from two creels) when the prediction targets came from the combined sample (from the creels) than when the targets were new instances (from the pond). Younger children did not distinguish between predictions about old versus new targets because they had difficulty in integrating information from the two samples (i.e., tracking the descriptive statistics). The researchers interpreted this result as suggesting that the presence of two samples with different distributions left children confused or suspicious of the sampling strategy.

In this experiment, participants encountered examples with a correlation between two binary features. The process of selecting the examples was either representative (e.g., random) or not (e.g., predetermined). Participants made a series of conditional predictions about individual examples drawn either from the sample or from the large population. Two groups of children participated in the study: 4- and 5-year-olds and 7- and 8-year-olds. Kalish and colleagues (2012) provided evidence that young school-age children were sensitive to inferential relations of samples to population. Preschool-aged