Journal of Experimental Child Psychology 141 (2016) 1–22

Contents lists available at ScienceDirect

Journal of Experimental Child Psychology

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

Children's use of interventions to learn causal structure

Teresa McCormack [a,*], Neil Bramley [b], Caren Frosch [c], Fiona Patrick [b], David Lagnado [b]

[a] School of Psychology, Queen's University Belfast, Belfast, Northern Ireland BT7 1NN, UK
[b] Division of Psychology and Language Sciences, University College London, London WC1E 6BT, UK
[c] School of Psychology, University of Leicester, Leicester LE1 9HN, UK

ARTICLE INFO

Article history:
Received 20 October 2014
Revised 29 April 2015
Available online 22 August 2015

Keywords:
Causal learning
Bayesian modelling
Scientific learning
Causal structure
Active learning
Self directed learning

ABSTRACT

Children between 5 and 8 years of age freely intervened on a three-variable causal system, with their task being to discover whether it was a common cause structure or one of two causal chains. From 6 or 7 years of age, children were able to use information from their interventions to correctly disambiguate the structure of a causal chain. We used a Bayesian model to examine children's interventions on the system; this showed that with development children became more efficient in producing the interventions needed to disambiguate the causal structure and that the quality of interventions, as measured by their informativeness, improved developmentally. The latter measure was a significant predictor of children's correct inferences about the causal structure. A second experiment showed that levels of performance were not reduced in a task where children did not select and carry out interventions themselves, indicating no advantage for self-directed learning. However, children's performance was not related to intervention quality in these circumstances, suggesting that children learn in a different way when they carry out interventions themselves.

* Corresponding author.
E-mail address: t.mccormack@qub.ac.uk (T. McCormack).

http://dx.doi.org/10.1016/j.jecp.2015.06.017
0022-0965/

Introduction

Most developmental studies of causal learning (e.g., Bullock, Gellman, & Baillargeon, 1982; Gopnik, Sobel, Schulz, & Glymour, 2001; Shultz, 1982) have required children to judge whether an event is causally efficacious. However, in learning about the world, what is at issue is often not just whether a specific variable has a particular causal power but also the structure of the causal relations between a set of variables. You might observe that Events A, B, and C tend to co-occur (e.g., that when you feel stressed you are likely to drink more heavily and also that your blood pressure is raised). This co-occurrence is consistent with a variety of different causal structures; for example, the structure may be a common cause in which A independently causes both B and C, B ← A → C (stress causes heavier drinking and also independently causes raised blood pressure), or a causal chain in which A causes B which causes C, A → B → C (stress causes heavier drinking which raises blood pressure). How do we distinguish between these possibilities? Intervening on a causal system potentially provides very important information (Hagmayer, Sloman, Lagnado, & Waldmann, 2007; Sloman & Lagnado, 2005; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003). Observing what happens if we intervened on B only would allow us to distinguish between the two suggested causal structures; if, assuming no background causes, we make B occur on its own (engage in heavy drinking when not stressed), and C does not occur (blood pressure is not elevated), we can rule out the A → B → C causal chain.

A variety of studies have required adults to infer the structure of the relations between sets of variables (e.g., Fernbach & Sloman, 2009; Kushnir, Gopnik, Lucas, & Schulz, 2010; Lagnado & Sloman, 2004, 2006; Sobel & Kushnir, 2006; Steyvers et al., 2003). In several of these studies, participants learned the causal structure by deciding what interventions to make on elements in the system, carrying out these interventions and observing their effects (Bramley, Lagnado, & Speekenbrink, 2015; Lagnado & Sloman, 2004, 2006; Sobel & Kushnir, 2006; Steyvers et al., 2003). Such interventions are assumed to reveal conditional dependencies or independencies between variables, and adults' success on these tasks has been interpreted as being consistent with the causal Bayes net approach to causal learning (e.g., Glymour, 2001; Gopnik et al., 2004) that captures causal learning in terms of the construction of causal models based on conditional probability information. A major advantage of this approach is that it specifies how and why interventions on a system yield richer information about the causal (in)dependencies between variables than that which is available through observation of patterns of covariation (Hagmayer et al., 2007; Steyvers et al., 2003; Waldmann & Hagmayer, 2005).

The causal Bayes net approach has been extensively adopted by developmental psychologists interested in explaining children's learning about causation (Gopnik, 2012; Gopnik & Wellman, 2012; Gopnik et al., 2004). The majority of studies in this tradition have involved children learning whether an object possesses a particular causal power, usually on the basis of observing the experimenter's actions (e.g., whether an object makes a box light up and play a tune; Gopnik & Sobel, 2000; Kushnir & Gopnik, 2007; Sobel, Tenenbaum, & Gopnik, 2004). Relatively few studies have used tasks in which children themselves decide which interventions to carry out in order to discover the causal structure of a system (e.g., whether it has a causal chain or common cause structure). Such studies are particularly important because they can be used to assess young children's effectiveness in generating and testing hypotheses about the causal relations between sets of variables. Moreover, a key advantage of the causal Bayes net approach over most other accounts of causal learning is that it can capture this more complex type of learning, distinguishing between different causal paths as well as identifying variables' ultimate effects.

One study that did examine children's ability to learn causal structure by means of making interventions on a system is that of Schulz, Gopnik, and Glymour (2007, Experiment 3), in which 4- and 5-year-olds intervened on a causal system involving a box with two gears. Children needed to decide whether each gear moved by itself or whether one of the gears caused the other one to move. Children could remove each gear in turn from the box to examine whether the other gear worked on its own when the box itself was switched on. They gave their answers about the relations between the gears by selecting from a set of anthropomorphized pictures of the two gears that depicted different possible relations between them. Performance on this task was mixed. Children did not all reliably generate the