Cognition 156 (2016) 30–40

Contents lists available at ScienceDirect

Cognition

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

Original Articles

The early emergence and puzzling decline of relational reasoning: Effects of knowledge and search on inferring abstract concepts

Caren M. Walker [a,*], Sophie Bridgers [b], Alison Gopnik [c]

[a] University of California, San Diego, Department of Psychology, 9500 Gillman, # 0109, La Jolla, CA 92093-0109, United States
[b] Stanford University, Department of Psychology, 450 Serra Mall, Jordan Hall, Building 420, Stanford, CA 94305, United States
[c] University of California, Berkeley, Department of Psychology, 3210 Tolman Hall, Room 1221, Berkeley, CA 94720, United States

ARTICLE INFO

Article history:
Received 20 December 2015
Revised 8 July 2016
Accepted 14 July 2016

Keywords:
Cognitive development
Causal learning
Relational reasoning
Explanation
Bayesian inference

ABSTRACT

We explore the developmental trajectory and underlying mechanisms of abstract relational reasoning. We describe a surprising developmental pattern: Younger learners are better than older ones at inferring abstract causal relations. Walker and Gopnik (2014) demonstrated that toddlers are able to infer that an effect was caused by a relation between two objects (whether they are the same or different), rather than by individual kinds of objects. While these findings are consistent with evidence that infants recognize *same-different* relations, they contrast with a large literature suggesting that older children tend to have difficulty inferring these relations. Why might this be? In Experiment 1a, we demonstrate that while younger children (18–30-month-olds) have no difficulty learning these relational concepts, older children (36–48-month-olds) fail to draw this abstract inference. Experiment 1b replicates the finding with 18–30-month-olds using a more demanding intervention task. Experiment 2 tests whether this difference in performance might be because older children have developed the general hypothesis that individual kinds of objects are causal – the high initial probability of this alternative hypothesis might override the data that favors the relational hypothesis. Providing additional information falsifying the alternative hypothesis improves older children's performance. Finally, Experiment 3 demonstrates that prompting for explanations during learning also improves performance, even without any additional information. These findings are discussed in light of recent computational and algorithmic theories of learning.

1. Introduction

A growing literature indicates that children as young as 16 months of age are able to learn specific causal properties from contingency information and can act on that knowledge to bring about novel effects in the world (see Gopnik & Wellman, 2012 for a review). But when and how can children learn more abstract causal principles? The ability to quickly learn abstract and specific relations in tandem might explain how children acquire the impressive amount of causal knowledge evident in their early intuitive theories about the world.

In the current paper, we examine children's developing ability to infer an abstract causal principle – a *relation* between objects that causes an effect (i.e., the relation "same" or "different") – from a limited set of observations. Walker and Gopnik (2014) recently demonstrated that toddlers (18–30-month-olds) are surprisingly adept at learning and using these relational concepts in a causal relational match-to-sample (RMTS) task. In this study, children were assigned to either a *same* or *different* condition, and observed as four pairs of objects (two "same" pairs and two "different" pairs) were placed on a toy that played music. In the *same* condition, pairs of identical objects activated the toy while pairs of different objects did not. This pattern of activation was reversed for the *different* condition. During test, children were given a choice between two novel pairs: one pair of *same* and one pair of *different* objects, and asked to select the pair that would activate the toy. Children overwhelmingly selected the pair that was consistent with their training. These results suggest that the ability to reason about abstract relations is in place very early – emerging spontaneously only a few months after the first evidence of children's ability to learn about the specific causal properties of individual objects.

Walker and Gopnik's (2014) results are consistent with some research demonstrating early competence in abstracting *same-different* relations in infancy. In particular, research relying on looking-time and visual search measures suggest that infants as young as 7- and 9-months-old may be able to recognize data that

* Corresponding author at: Department of Psychology, University of California, San Diego, 9500 Gillman Dr., La Jolla, CA 92093-0109, United States.
E-mail address: carenwalker@ucsd.edu (C.M. Walker).

http://dx.doi.org/10.1016/j.cognition.2016.07.008
0010-0277/

involve *same-different* relations in visual displays from very few trials (Dewar & Xu, 2010; Ferry, Hespos, & Gentner, 2015; Tyrrell, Stauffer, & Snowman, 1991; see also, Hochmann, Mody, & Carey, 2016).

Intuitively, it might seem plausible that more abstract hypotheses, such as *same* and *different*, would be acquired later than lower-level, concrete ones based on specific features of objects. However, theoretical advances drawing on Bayesian accounts of the "blessing of abstraction" (Goodman, Ullman, & Tenenbaum, 2011) combined with empirical research on early learning (Dewar & Xu, 2010; Schulz, Goodman, Tenenbaum, & Jenkins, 2008) suggest that children's ability to learn abstract principles need not progress in a bottom-up manner. Instead, Hierarchical Bayesian Models formalize how it may be possible to infer relations between objects and events among multiple levels of abstraction simultaneously (Griffiths & Tenenbaum, 2009; Tenenbaum, Griffiths, & Kemp, 2006).

In fact, there is experimental evidence supporting the claim that children are able to grasp certain abstract principles at the same time, or even *before* they learn the specific causal relations underlying them (Gelman & Gottfried, 1996; Kemp, Perfors, & Tenenbaum, 2007; Lehrer & Schauble, 1998; Mansinghka, Kemp, Tenenbaum, & Griffiths, 2006; Rozenblit & Keil, 2002; Schulz et al., 2008; Tenenbaum & Niyogi, 2003; Tenenbaum et al., 2006). For example, decades of evidence from developmental studies of psychological essentialism (e.g., Gelman, 2003; Keil, 1989) has demonstrated that children assume that animals from similar species are likely to share internal structures. Importantly, they can do this well before they can identify just what those internal structures actually are.

This account may help to explain the growing evidence that basic relational concepts are available much earlier than previously believed. On the other hand, these results contrast with a much larger body of research demonstrating that older, preschool-aged children consistently experience difficulty with relational matching (e.g., Christie & Gentner, 2007, 2010, 2014; Gentner, 2010). If relational learning is indeed a continuous process, as has been proposed (e.g., Gentner & Medina, 1998; Gentner & Namy, 1999; Mix, 2008; Richland, Morrison, & Holyoak, 2006), and *same-different* concepts are already available very early in development (Ferry et al., 2015; Smith, 1984; Tyrrell et al., 1991; Walker & Gopnik, 2014), why do older children often fail to demonstrate this knowledge? How might we interpret this apparent developmental reversal in which abstract reasoning seems to emerge in the first two years of life, but then decline in early childhood?

First, it is possible that older children failed to exhibit relational reasoning in previous studies because of methodological problems – the tasks were simply too difficult. The toddlers in Walker and Gopnik (2014) may have succeeded because the novel causal procedure simply made the task easier (see also, Smith, 1984). Similarly, there is a large literature indicating the dissociation between children's knowledge as measured in looking-time tasks and their ability to act on this knowledge across a variety of developmental domains (e.g., Hood, Cole-Davies, & Dias, 2003; Kirkham, Cruess, & Diamond, 2003; Zelazo, Frye, & Rapus, 1996). These possibilities may account for differences between younger (Ferry et al., 2015; Walker & Gopnik, 2014) and older (Christie & Gentner, 2014) children's performance on *same-different* relational reasoning tasks.

In Experiment 1a below, we therefore present participants with exactly the same reasoning task used in Walker and Gopnik (2014). After replicating this previous work with 18–30-month-olds, we also assess an additional group of 18–30-month-olds, using another test of toddlers' causal understanding of the relational concepts (Experiment 1b). In addition to coding which pair of blocks the children selected (by pointing) to activate the toy in the causal RMTS task, we also coded whether the children themselves put the correct novel pair of blocks on top. This ability to design a new intervention, and to act on a cause in order to produce its effect has been argued to be a particularly telling signature of true causal understanding (Pearl, 2000; Woodward, 2003).

Centrally, Experiment 1a also compares performance of 18–30-month-olds with that of older children (ranging from 30 to 48-month-olds) on exactly the same task. We include the full range of ages from 18 to 48 months to test if there is a continuous developmental trajectory. If the toddlers in Walker and Gopnik (2014) indeed succeeded because of the particular methodological features of the task, then we would expect that older children would succeed as well. If they fail, however, this decline cannot be explained as a result of the methodological differences between tasks assessing the presence of relational concepts in toddlers, and those assessing older children.

There is at least one reason why younger children might genuinely outperform older children in learning these causal relational concepts, independent of method. It may be that while 3-year-olds are able to reason on the basis of relations, they are less likely to infer relational causes because they have learned that the properties of individual objects are especially likely to have causal powers. This leads to a bias. When they see a block on the toy they assume that some feature of that individual object, its color or shape or weight, was responsible for the effect, rather than the relation between blocks. Indeed, preschool-aged children often demonstrate a bias to attend to individual object kinds, which has been proposed to interfere with relational processing (e.g., Christie & Gentner, 2007, 2010, 2014; Gentner, 1998; Gentner & Medina, 1998; Gentner & Rattermann, 1991). A parallel bias has been observed in a variety of causal learning tasks, in which preschool-aged children assume that causal powers are inherent to individual objects (e.g. Gopnik & Sobel, 2000).

Why would this bias affect older learners and not younger ones? In probabilistic model accounts, learners explain newly observed evidence by searching through a space of potential hypotheses and testing these hypotheses against the data (e.g., Gopnik & Wellman, 2012). To do this, learners combine two probabilities: the "prior" – the probability of a particular hypothesis being true before any data are observed, and the "likelihood" – the probability of the observed data given that a particular hypothesis is true. Combining these two probabilities with Bayes rule produces the "posterior" – the probability of the hypothesis being true given the observed data. A learner can then compare the posteriors of different hypotheses, settling on the ones with the highest probabilities.

These models predict that if the prior probability of one hypothesis is high, then it will take stronger data to overturn it in favor of another hypothesis. But in addition to formulating specific hypotheses, learners can also formulate "overhypotheses" or "framework principles" (Goodman, 1955; Goodman et al., 2011; Kemp et al., 2007). Having an overhypothesis leads the learner to assign a higher prior probability to certain types of hypotheses, and so constrains children's interpretation of new data (Kemp et al., 2007). As a result, in order for the learner to consider a hypothesis that is inconsistent with the overhypothesis, the learner would need more evidence supporting this competing hypothesis than if she began with no prior expectations and assigned all possible hypotheses an equal prior probability (i.e., a "flat" prior).

From a probabilistic models perspective, then, we might say that younger children have a "flatter" prior distribution: they are equally likely to entertain hypotheses about individual object properties and about relations. In the case of Walker and Gopnik's (2014) causal reasoning task, an abstract principle of simplicity, as proposed by Lombrozo (2007), might lead toddlers to initially prefer a relational hypothesis over an individual object hypothesis, since a relational hypothesis proposes fewer causes to account for the data. Indeed, previous work demonstrates that