UC Berkeley

UC Berkeley Previously Published Works

Title

Toddlers infer higher-order relational principles in causal learning.

Permalink

<https://escholarship.org/uc/item/68z4n01t>

Journal

Psychological science, 25(1)

ISSN

0956-7976

Authors

Walker, Caren M
Gopnik, Alison

Publication Date

2014

DOI

10.1177/0956797613502983

Peer reviewed

Toddlers Infer Higher-Order Relational Principles in Causal Learning

Caren M. Walker and Alison Gopnik

University of California, Berkeley

Psychological Science
2014, Vol. 25(1) 161–169
© The Author(s) 2013
Reprints and permissions:
sagepub.com/journalsPermissions.nav
DOI: 10.1177/0956797613502983
pss.sagepub.com

Abstract

Children make inductive inferences about the causal properties of individual objects from a very young age. When can they infer higher-order relational properties? In three experiments, we examined 18- to 30-month-olds' relational inferences in a causal task. Results suggest that at this age, children are able to infer a higher-order relational causal principle from just a few observations and use this inference to guide their own subsequent actions and bring about a novel causal outcome. Moreover, the children passed a revised version of the relational match-to-sample task that has proven very difficult for nonhuman primates. The findings are considered in light of their implications for understanding the nature of relational and causal reasoning, and their evolutionary origins.

Keywords

cognitive development, causality, inference, learning, relational reasoning

Received 12/6/12; Revision accepted 7/23/13

Learning about causal relationships is one of the most important and challenging problems young humans face. Causal knowledge allows you to act on the world: If you know that A causes B, you can act on A to bring about B. Studies show that children as young as 16 to 24 months of age can quickly learn causal properties of objects from patterns of statistical contingency and can act on that knowledge to bring about effects (e.g., Gweon & Schulz, 2011; Meltzoff, Waismeyer, & Gopnik, 2012; Sobel & Kirkham, 2006; for reviews, see Gopnik, 2012; Gopnik & Wellman, 2012).

Much of this research on early causal learning has used a *blicket-detector* paradigm (Gopnik & Sobel, 2000), in which children learn which objects activate a novel machine. Children's inferences in these tasks go beyond associative learning, revealing the distinctive profile of causal inference. For example, children will use these inferences to design novel interventions (patterns of action they have never observed), to construct counterfactuals, and to make explicit causal judgments, including judgments about unobserved features (e.g. Gopnik et al., 2004; Gopnik & Sobel, 2000; Schulz, Gopnik, & Glymour, 2007; Sobel, Yoachim, Gopnik, Meltzoff, & Blumenthal, 2007).

However, much less is known about the development of children's ability to infer higher-order relational causal

principles. According to theory theorists of cognitive development, children are learning not only particular causal relationships, but also higher-order generalizations about causal structure (e.g., Carey, 2009; Gopnik & Meltzoff, 1997; Wellman & Gelman, 1992). Recent computational work also suggests that higher-order generalizations can help children learn new specific relationships from perceptual data more quickly (e.g., Goodman, Ullman, & Tenenbaum, 2011; Griffiths & Tenenbaum, 2007; Kemp, Perfors, & Tenenbaum, 2007).

Causal inferences might be more or less abstract, higher-order, or relational in many different ways. Here we focus on just one contrast: between object properties, such as shape or color, and higher-order relations between those properties, such as whether they are the same or different. For example, very young children can learn that red blocks activate a toy. At what age can children learn that two blocks that are the same (regardless of their color) can do so?

Empirical research using looking-time measures suggests that human infants may be able to recognize

Corresponding Author:

Caren M. Walker, Department of Psychology, University of California, Berkeley, 3210 Tolman Hall, Berkeley, CA 94720
E-mail: caren.walker@berkeley.edu

patterns of data that involve these higher-order relations (Dewar & Xu, 2010; Ferry, Hespos, & Gentner, 2012; Tyrrell, Stauffer, & Snowman, 1991). However, there is no evidence to date that infants can use those patterns to make causal inferences or guide subsequent actions.

In fact, earlier studies indicated that even preschoolers had difficulty making inferences in higher-order relational reasoning tasks (e.g., Christie & Gentner, 2007, 2010; Gentner, 2010). Children succeeded only when given labels or linguistic scaffolding to point out the pattern of similarity. Indeed, even when explicitly instructed to compare objects, 3-year-olds performed significantly below chance when test items were presented sequentially rather than side by side (Christie & Gentner, 2010).

These findings might lead to the conclusion that learning higher-order relations and using them to guide actions depends on direct instruction, language, and cultural input (e.g., Christie & Gentner, 2007; Gentner, 2003, 2010; Gentner, Anggoro, & Klibanoff, 2011). However, the tasks often relied on verbal categorizations of complex, multidimensional stimuli (e.g., Christie & Gentner, 2010). One study by Smith (1984) provides a hint that children might do better in a more goal-directed task with simpler materials. In particular, 2½-year-olds showed some understanding of identity matching in a nonverbal game.

Higher-order relational reasoning has also been studied extensively in nonhuman animals. Chimpanzees, like young infants, are able to spontaneously detect a relational pattern in habituation tasks (Oden, Thompson, & Premack, 1990). However, they have more difficulty with a relational match-to-sample task (Oden, Premack, & Thompson, 1988; Premack, 1983). In this task, animals observe a relational pattern: AA', BB', and CC' all lead to a reward. Then they are given a choice between AB (object match) and DD' (relational match). Although A and B have each been associated with the reward, an animal who has inferred the higher-order relational pattern should choose DD'. Premack and his colleagues found that chimpanzees could not solve this relational task without hundreds of trials with feedback (Premack, 1988) or training to use linguistic symbols for "same" (Premack, 1976, 1983; Premack & Premack, 1983, 2003).

Additional comparative studies have confirmed that this task is especially difficult for nonhuman primates and other animals (see Penn, Holyoak, & Povinelli, 2008). Moreover, when nonhuman animals, such as baboons, *do* solve this task, they require extended training and thousands of trials, which may indicate the use of simpler perceptual strategies, such as minimizing entropy in a perceptual array (Fagot, Wasserman, & Young, 2001; Wasserman, Fagot, & Young, 2001).

Do human children always require linguistic cues or extensive training to solve relational tasks, like the

preschoolers and primates in earlier studies? We designed a nonverbal blinket-detector task to explore when children could use higher-order relations to make causal inferences. Unlike the causal effects in previous blinket-detector studies, the causal effect in this task depended on whether the objects were the same or different, rather than on properties of the objects themselves.

Experiment 1a

In Experiment 1a, 21- to 24-month-olds were introduced to a novel toy that played music and to three unique pairs of identical blocks: AA', BB', and CC'. The experimenter placed blocks on the toy, and the toy either was or was not activated. Although individual blocks by themselves failed to activate the toy, pairs of identical blocks produced the effect. Immediately after this brief training, we examined whether the children had learned that the novel relational property (i.e., "same") produced the effect by asking them to activate the toy.

Method

Participants. A total of twenty-three 21- to 24-month-old toddlers participated in Experiment 1a (mean age = 23.0 months, *SD* = 1.05 months, range = 20.9–25.0 months; 13 girls, 10 boys). Three additional children were tested but excluded for fussiness or for failing to respond. Children were recruited from day-care centers and museums, and a range of ethnicities resembling the diversity of the local population was represented.

Materials. The toy was a 10- × 6- × 4-in. opaque white cardboard box containing a wireless doorbell. When a pair of blocks "activated" the toy, the doorbell played a melody. In fact, the toy was surreptitiously activated by a remote control. Six painted wooden blocks in assorted colors and shapes (three unique pairs of two identical blocks) were placed on the toy during the training phase. Six additional blocks, two novel pairs of identical blocks and two novel individual blocks, were used during the test phase.

Procedure. The procedure for Experiment 1a is illustrated in Figure 1. Following a warm-up, the toy was placed on the table. The experimenter said, "This is my toy. Some things make my toy play music, and some things do not make my toy play music." Children then observed while the experimenter placed six blocks (three unique pairs of "same" objects: AA', BB', and CC') on the table in front of the toy. She said, "Let's try one," selected a block (A), and placed it on top of the toy. No effect was produced. After a pause, the experimenter again said, "Let's try one," selected the paired block (A'), and placed

Fig. 1. Schematic representation of the training and test trials in Experiment 1a. On each training trial, a single block was placed on the toy (no activation), and then an identical block was added, activating the toy. This was repeated for the remaining two training pairs. On each test trial, three test blocks (novel paired block, familiar block, and novel distractor block) were presented. The experimenter then placed the target block on the toy, which produced no effect. The child was asked to select one test block to activate the toy.

it next to the first block (A) on top of the toy. This pair of objects (AA') activated the toy. The experimenter smiled and said, "Music," removed the blocks, and returned them to the pile of six. This procedure was repeated with the two remaining pairs (BB' and CC'). The order in which the pairs were used was randomized. Following all three demonstrations, all blocks were removed from the table.

Next, the experimenter produced three test blocks—one *novel paired block* (D), one *familiar block* (A), and one *novel distractor block* (E)—and placed them in a row on the table. The order of presentation was randomized. She said, "Let's try one," produced the *target block* (D'), and placed it on top of the toy. No effect was produced. The experimenter then pushed the toy and all three test blocks toward the child and asked, "Can you pick one of these (pointing to the test blocks) to make my toy play music?"

The first test block that the child placed on the toy was recorded. The toy was activated if the child correctly selected the novel paired block. If the child selected the familiar block or the novel distractor block, the toy was not activated. After this feedback, this procedure was repeated in a second test trial with a new set of test blocks.

If the toddlers acted on the basis of the previous association between the block and effect, they should have chosen the familiar block. If they simply preferred to try novel blocks, they should have picked the novel distractor block as often as the novel paired block. However, if the toddlers were able to learn the higher-order relation, they should have selected the novel paired block.

Coding and reliability. Children received 1 point for selecting the novel paired block and 0 points for selecting either of the other blocks in each trial. Responses were recorded by a second researcher, and all sessions were recorded for independent coding by a third researcher, who was naive to the purpose of the experiment. Interrater reliability was very high; the two coders agreed on 96% of the children's responses. Two minor discrepancies were resolved by a third party.

Results and discussion

Across the two test trials, children inferred the relational property and selected the novel paired block more often than expected by chance ($M = 1.22$, $SD = 0.74$), $\chi^2(2, N = 23) = 20.04$, $p < .001$. A Fischer exact test revealed no order effects for Test Trials 1 and 2, $p = .39$. Children chose the novel paired block (61% of trials) significantly more often than the novel distractor block (20%), $\chi^2(2, N = 23) = 14.15$, $p < .001$, and significantly more often than the familiar block (15%), $\chi^2(2, N = 23) = 14.09$,

$p < .001$. A minority of children (4%) placed more than one block on the toy simultaneously; these responses were scored as incorrect.

Previous proposals have suggested that children are unable to reason relationally because they tend to focus on the identity of objects that have been previously associated with the outcome (e.g., Gentner, 2010). We found no evidence of this. In fact, only 39% of participants who answered incorrectly on a given trial selected the familiar block, and the familiar block was not selected more often than the novel distractor block on these trials, $\chi^2(2, N = 23) = 2.43$, $p = .30$. This is particularly surprising given that the familiar block had been associated with the effect during training.

The results suggest that by 21 to 24 months of age, toddlers are able to infer a relational principle—"same"—from just a few pieces of evidence and use this inference to bring about a novel causal outcome. However, the children might have succeeded on this task because they imitated the experimenter's selection or because they preferred to match the blocks, regardless of training. Experiment 1b was designed to address these alternative explanations.

Experiment 1b

The procedure for Experiment 1b was identical to that for Experiment 1a, but the second object in each pair was occluded during the training trials. Because children observed only the first item in each pair, they were given no evidence for the relational property. If they simply imitated the experimenter or had a preexisting preference for matching, their performance would not differ from children's performance in Experiment 1a.

Method

Participants. Twenty 21- to 24-month-olds participated (mean age = 22.4 months, $SD = 1.8$ months, range = 20.8–25.6 months; 8 girls, 12 boys). Two additional children were tested but excluded for failing to respond. Recruitment procedures and demographics were the same as in Experiment 1a.

Materials and procedure. The materials and procedure were identical to those in Experiment 1a. However, the children did not observe the second object during the training trials. Instead, the second object was occluded by a 4 × 4-in. piece of cardboard, which was held in front of the block by the experimenter. Additionally, only one test trial was administered in order to avoid providing feedback. Therefore, each child could receive either 0 or 1 point. Interrater reliability for Experiment 1b was 100%.

Results and discussion

In the absence of evidence for the relational principle, only 40% of participants selected the novel paired block (exact binomial test, $p = .65$). This percentage was significantly different from the percentage of children (63%) who chose that block on the first trial in Experiment 1a, $p < .05$ by a Fischer exact test. Children's selections were evenly distributed; 40% of children selected the novel paired block, 35% selected the familiar block, and 25% selected the novel distractor. These results show that the findings from Experiment 1a could not have been the result of imitation or a bias to match.

Experiment 2

In the earlier primate studies using canonical relational match-to-sample tasks, pairs were presented simultaneously during training (e.g., the relation "same" was taught using pairs AA', BB', and CC'), and the animals had to choose between test pairs illustrating an object match (AB) or the relational match (DD'). Chimpanzees did not succeed spontaneously on this task—and had great difficulty even after engaging in trial and error over hundreds of trials. However, chimpanzees *were* able to solve a simpler match-to-sample task. The animals were first taught to match a test object (A) to a target object (A') through multiple positive- and negative-reinforcement trials over several weeks. They then generalized this pattern to novel objects without additional training (Oden et al., 1988; Premack, 1976; Premack & Premack, 1983, 2003).

For the task in Experiment 1a, like the simple match-to-sample task used with primates, the training objects were presented on the toy sequentially, and this may have made the task easier. However, children's performance in Experiment 1a also differed in several ways from primates' performance in previous studies. Children learned by observation—unlike the primate subjects, they did not initially make the responses themselves—and they spontaneously chose the novel paired block after observing only three trials. Additionally, they never observed that the "different" blocks would not produce the effect (whereas the primates engaged in reinforcement learning), so the association between the incorrect familiar block and the effect should have continued to be high during the test trials.

In order to make the comparison with the primate task clearer, we designed a causal task that was more directly analogous to the primate task because both "same" and "different" objects were placed on the toy simultaneously in pairs. This task also allowed us to explore whether children would infer the "different" relation as well as the "same" relation. We tested toddlers from a broader age range to explore possible developmental differences, recruiting children ages 18 to 30 months.

Participants were randomly assigned to one of two conditions: *same* or *different*. In the *same* condition, children were given two pieces of evidence that pairs of "same" objects (AA', BB') simultaneously placed on the toy produced the effect. We also provided two pieces of evidence that pairs of "different" objects (DE, FG) failed to produce the effect. In the *different* condition, children were given the same four pieces of evidence, but "different" pairs (DE, FG) produced the effect, and "same" pairs (AA', BB') failed to do so.

Method

Participants. Thirty-eight 18- to 30-month-olds participated (mean age = 25.8 months, $SD = 3.8$ months, range = 18.0–30.6 months; 21 girls, 17 boys); 19 children were randomly assigned to each condition (*same* and *different*). Seven additional children were tested but excluded: 4 because of failure to complete the study and 3 because of experimenter error. Recruitment procedures and demographics were the same as in Experiments 1a and 1b.

Materials. The same toy from Experiments 1a and 1b was used. Eight painted wooden blocks in assorted colors and shapes (2 pairs of "same" blocks and 2 pairs of "different" blocks) were placed on the toy in pairs during training. The blocks in each "same" pair were identical in color and shape, and the blocks in each "different" pair were distinct in color and shape (see Fig. 2). Four additional blocks were used during the test phase: one novel pair of "same" blocks and one novel pair of "different" blocks. The pairs of test blocks were placed on 4- × 4-in. plastic trays.

Procedure. The procedure for Experiment 2 is illustrated in Figure 2. Following a warm-up, the toy was placed on the table. The experimenter said, "This is my toy. Some things make my toy play music, and some things do not make my toy play music." Children then observed while the experimenter placed all eight training blocks (A, A', B, B', C, D, E, F) in a random arrangement on the table in front of the toy. The experimenter said, "Look at these things! We will try them on my toy." Then, the experimenter removed all objects from view, selected a pair of "same" blocks (e.g., AA'), and placed the blocks simultaneously on the toy. Children in the *same* condition observed the pair of objects activate the toy. The experimenter smiled and said, "Music! Let's try that again," picked up the pair of blocks, and placed them back on the toy a second time; the children again observed the outcome. After this second demonstration, the experimenter removed the pair, selected another pair—a "different" pair (e.g., CD)—and placed it on the toy. This time, children in the *same* condition observed no effect. As with the first pair, this demonstration was repeated.

Fig. 2. Schematic representation of training and test trials in the *same* and *different* conditions in Experiment 2. On each training trial, a pair of blocks were placed on the toy. In the *same* condition, the pairs of identical objects activated the machine. In the *different* condition, the pairs of distinct objects activated the machine. Participants observed four training trials (two causal and two inert). On each test trial, a novel pair of “same” blocks and a novel pair of “different” blocks were presented. The child was asked to select the pair that would activate the toy.

This procedure was repeated for the remaining “same” pair and “different” pair. Thus, all pairs were placed on the toy twice, and the children observed a total of eight outcomes (four positive and four negative).

Children in the *different* condition observed the same set of evidence as children in the *same* condition, with

one critical change: The pairs of “different” objects (e.g., CD) caused the toy to play music, whereas the pairs of “same” objects (e.g., AA') failed to activate the toy. There were no other differences in procedure. Which particular objects were included in each training and test pair was randomized across children.

Following the training phase in both conditions, the experimenter said, “Now it is going to be *your* turn. I want you to help me pick the ones that will make my toy play music!” The experimenter produced two pairs of test blocks: one novel “same” pair (GG′) and one novel “different” pair (HI). In order to avoid a novelty preference, we used novel objects in both test pairs. The pairs were presented to the child on plastic trays. The experimenter held up the two trays, shook them to get the child’s attention, and asked, “Can *you* pick the ones that will make my toy play music?” She then placed the trays on opposite sides of the table in front of the child. The side on which the correct pair was placed was randomized across subjects. The first tray that the child selected was recorded. A selection was defined as pointing to the tray, reaching to the tray, or picking up the objects on the tray.

If they had learned the relational property, then children in the *same* condition should have correctly selected the tray with the novel “same” objects, whereas children in the *different* condition should have correctly selected the tray with the novel “different” objects. Correct selections were given a score of 1, and incorrect selections were given a score of 0. Coding and recording procedures were identical to those in Experiments 1a and 1b. Interrater reliability was very high; the two coders agreed on all but one response to the test question.

Results and discussion

Across the two conditions, the children inferred the relational property and selected the correct pair more often than expected by chance ($M = .79$, $SD = .41$; chance = .50), $p < .02$ by exact binomial test. In fact, performance was identical in the *same* and *different* conditions, with 15 out of 19 children in each condition selecting the test pair that corresponded with the relation learned during the training trials. Additionally, logistic regression revealed no significant developmental change in performance between 18 and 30 months of age, $\chi^2(1, N = 38) = 0.11$, $p = .74$. The fact that children responded differentially in the otherwise identical *same* and *different* conditions also allowed us to rule out superficial explanations for the results, such as imitation or a preference for “same” or “different” pairs: Each condition acted as a control for the other condition.

Experiment 2 indicates that toddlers are able to infer the relational causal principles “same” and “different” from just a few pieces of evidence and to use this inference to intervene to bring about a novel causal outcome.

General Discussion

These findings show that human toddlers as young as 18 months can succeed on a relational causal match-to-sample task after only a few trials and without explicit

linguistic cues, instruction, or reward. This study has implications for current understanding of both causal and relational reasoning. In this paradigm, toddlers are able to quickly learn higher-order relational causal principles and use them to guide their actions. This ability appears to be in place surprisingly early—only a few months after the first evidence of the ability to learn about specific causal properties from contingency—and it may be in place even earlier. This may help explain how young children acquire the impressive causal knowledge evident in their intuitive theories (Carey, 2009; Gopnik & Wellman, 2012).

These findings also contrast with the striking failure of nonhuman primates to solve similar tasks, even when the relation is associated with a strong pattern of positive and negative reinforcement, and even after a large number of trials. This finding suggests that an ability to quickly learn relational causal concepts might be a dimension on which humans differ from other primates. This human advantage might in turn reflect the broader evolution of higher-order relational cognition (Penn et al., 2008) or causal cognition in general (Buchsbaum, Bridgers, Weisberg, & Gopnik, 2012; Byrne, 1995; Heyes & Frith, 2012).

Several questions for further study remain. One is whether the causal nature of this task was critical, or whether other aspects of the task, such as the fact that it involved goal-directed actions, might have made it easier for the children than relational tasks in other studies. It is also possible that children could succeed on this particular task by basing their causal inference on the observed association between the higher-order relational features and the effects. In other blinket-detector studies, children’s inferences go beyond association, but those studies would have to be replicated with the current relational design in order to rule out this possibility.

Further, it is possible that the children’s success was due to a perceptual heuristic, as has been suggested for nonhuman primates (Fagot et al., 2001; Penn et al., 2008; Wasserman et al., 2001). According to this argument, it is possible to solve relational match-to-sample tasks using the perceptual cue of entropy (i.e., the Shannon entropy of AA′ is 0, whereas that of AB is 1). Several features of Experiment 2 weigh against this possibility: The children saw pairs of objects (rather than multi-element displays), they observed only two positive and two negative trials, they never acted on an object, and their behavior was never reinforced. Indeed, no other species has come close to demonstrating the first-trial performance of these human children after so few observations (see Penn et al., 2008). Additionally, although human participants have been shown to be sensitive to entropy, findings suggest that additional processes of categorization likely play a role in the human conceptualization of “same”/“different” relations (Fagot et al., 2001). Nevertheless, future research examining this possibility would be informative.

Finally, it will be important to replicate this particular task with nonhuman primates to determine if, like children, they show greater success than on previously used relational reasoning tasks, or continue to have difficulty. Our protocol did not require a verbal response, so it may be useful in examining reasoning capacities in both preverbal human infants and nonhuman animals.

In conclusion, the current study does suggest that the ability to infer causal higher-order relations, an ability that could play a crucial role in further learning, is in place in humans from a very early age and does not depend on explicit linguistic cues or cultural scaffolding.

Author Contributions

C. M. Walker and A. Gopnik developed the study concept. Both authors contributed to the study design. Testing and data collection were performed by C. M. Walker. C. M. Walker performed the data analysis and interpretation under the supervision of A. Gopnik. C. M. Walker drafted the manuscript, and A. Gopnik provided critical revisions. Both authors approved the final version of the manuscript for submission.

Acknowledgments

We thank the parents and children who participated and the University of California, Berkeley, Early Childhood Centers; the Lawrence Hall of Science; and Habitot. We are grateful to Fei Xu, Tania Lombrozo, Keith Holyoak, and anonymous reviewers for providing feedback and to Rosie Aboody, Anna Akullien, Ela Banerjee, Sophie Bridgers, Sierra Eisen, Brynna Ledford, and especially Samantha Hubachek for facilitating data collection.

Declaration of Conflicting Interests

The authors declared that they had no conflicts of interest with respect to their authorship or the publication of this article.

Funding

This research was funded by the James S. McDonnell Foundation and by a National Science Foundation grant (BCS-1023875) to A. Gopnik.

References

- Buchsbaum, D., Bridgers, S., Weisberg, D. S., & Gopnik, A. (2012). The power of possibility: Causal learning, counterfactual reasoning, and pretend play. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367, 2202–2212.
- Byrne, R. W. (1995). *The thinking ape: Evolutionary origins of intelligence*. New York, NY: Oxford University Press.
- Carey, S. (2009). *The origin of concepts*. Oxford, England: Oxford University Press.
- Christie, S., & Gentner, D. (2007). Relational similarity in identity relation: The role of language. In S. Vosniadou & D. Kayser (Eds.), *Proceedings of the Second European Cognitive Science Conference* (pp. 601–666). London, England: Taylor & Francis.
- Christie, S., & Gentner, D. (2010). Where hypotheses come from: Learning new relations by structural alignment. *Journal of Cognition and Development*, 11, 356–373.
- Dewar, K. M., & Xu, F. (2010). Induction, overhypothesis, and the origin of abstract knowledge: Evidence from 9-month-old infants. *Psychological Science*, 21, 1871–1877.
- Fagot, J., Wasserman, E. A., & Young, M. E. (2001). Discriminating the relation between relations: The role of entropy in abstract conceptualization by baboons (*Papio papio*) and humans (*Homo sapiens*). *Journal of Experimental Psychology: Animal Behavior Processes*, 27, 316–328.
- Ferry, S., Hespos, S., & Gentner, D. (2012, June). *Prelinguistic relational concepts: Investigating the origins of analogical reasoning in infants*. Poster presented at the 18th Biennial International Conference on Infant Studies, Minneapolis, MN.
- Gentner, D. (2003). Why we're so smart. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and thought* (pp. 195–235). Cambridge, MA: MIT Press.
- Gentner, D. (2010). Bootstrapping the mind: Analogical processes and symbol systems. *Cognitive Science*, 34, 752–775.
- Gentner, D., Anggoro, F. K., & Klibanoff, R. S. (2011). Structure mapping and relational language support children's learning of relational categories. *Child Development*, 82, 1173–1188.
- Goodman, N. D., Ullman, T. D., & Tenenbaum, J. B. (2011). Learning a theory of causality. *Psychological Review*, 118, 110–119.
- Gopnik, A. (2012). Scientific thinking in young children: Theoretical advances, empirical research, and policy implications. *Science*, 337, 1623–1627.
- Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, 111, 3–32.
- Gopnik, A., & Meltzoff, A. N. (1997). *Words, thoughts, and theories*. Cambridge, MA: MIT Press.
- Gopnik, A., & Sobel, D. (2000). Detectingblickets: How young children use information about novel causal powers in categorization and induction. *Child Development*, 71, 1205–1222.
- Gopnik, A., & Wellman, H. (2012). Reconstructing constructivism: Causal models, Bayesian learning mechanisms, and the theory theory. *Psychological Bulletin*, 138, 1085–1108.
- Griffiths, T., & Tenenbaum, J. B. (2007). Two proposals for causal grammars. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 323–345). New York, NY: Oxford University Press.
- Gweon, H., & Schulz, L. (2011). 16-month-olds rationally infer causes of failed actions. *Science*, 332, 1524.
- Heyes, C., & Frith, U. (Eds.). (2012). New thinking: The evolution of human cognition [Theme issue]. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1599).

- Kemp, C., Perfors, A., & Tenenbaum, J. B. (2007). Learning overhypotheses with hierarchical Bayesian models. *Developmental Science*, 10, 307–321.
- Meltzoff, A. N., Waismeyer, A., & Gopnik, A. (2012). Learning about causes from people: Observational causal learning in 24-month-old infants. *Developmental Psychology*, 48, 1215–1228.
- Oden, D. L., Premack, D., & Thompson, R. K. (1988). Spontaneous transfer of matching by infant chimpanzees (*Pan troglodytes*). *Journal of Experimental Psychology: Animal Behavior Processes*, 14, 140–145.
- Oden, D. L., Thompson, R. K., & Premack, D. (1990). Infant chimpanzees spontaneously perceive both concrete and abstract same/different relations. *Child Development*, 61, 621–631.
- Penn, D. C., Holyoak, K. J., & Povinelli, D. J. (2008). Darwin's mistake: Explaining the discontinuity between human and nonhuman minds [Target article and commentaries]. *Behavioral & Brain Sciences*, 31, 109–178.
- Premack, D. (1976). *Intelligence in ape and man*. Hillsdale, NJ: Erlbaum.
- Premack, D. (1983). The codes of man and beasts [Target article and commentaries]. *Behavioral & Brain Sciences*, 6, 125–167.
- Premack, D. (1988). Minds without language. In L. Weiskrantz (Ed.), *Thought without language* (pp. 46–65). Oxford, England: Clarendon Press.
- Premack, D., & Premack, A. J. (1983). *The mind of an ape*. New York, NY: W. W. Norton.
- Premack, D., & Premack, A. J. (2003). *Original intelligence: Unlocking the mystery of who we are*. New York, NY: McGraw-Hill.
- Schulz, L., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science*, 10, 322–332.
- Smith, L. B. (1984). Young children's understanding of attributes and dimensions: A comparison of conceptual and linguistic measures. *Child Development*, 55, 363–380.
- Sobel, D. M., & Kirkham, N. Z. (2006). Blickets and babies: The development of causal reasoning in toddlers and infants. *Developmental Psychology*, 42, 1103–1115.
- Sobel, D. M., Yoachim, C. M., Gopnik, A., Meltzoff, A. N., & Blumenthal, E. J. (2007). The blicket within: Preschoolers' inferences about insides and causes. *Journal of Cognition and Development*, 8, 159–182.
- Tyrrell, D. J., Stauffer, L. B., & Snowman, L. G. (1991). Perception of abstract identity/difference relationships by infants. *Infant Behavior & Development*, 14, 125–129.
- Wasserman, E. A., Fagot, F., & Young, M. E. (2001). Same-different conceptualizations by baboons (*Papio papio*): The role of entropy. *Journal of Comparative Psychology*, 115, 42–52.
- Wellman, H. M., & Gelman, S. A. (1992). Cognitive development: Foundational theories of core domains. *Annual Review of Psychology*, 43, 337–375.