

The Design of the Learning Environment Shapes Preschoolers' Causal Inference

Alexandra Rett (arett@ucsd.edu)¹, Elizabeth Bonawitz (elizabeth.bonawitz@rutgers.edu)²,
& Caren Walker (carenwalker@ucsd.edu)¹

¹ Department of Psychology, UC San Diego, La Jolla, CA 92093

² Department of Psychology, Rutgers University, Newark, NJ 07102

Abstract

In the present study, we examine whether the design of the learning environment can impact causal inference in very young children. Specifically, we assess whether the physical features of a novel toy can facilitate children's recognition of an abstract, relational hypothesis (*same-different*) that they typically fail to discover. Three-year-olds were presented with an identical pattern of evidence that was consistent with a relational hypothesis (i.e., pairs of same or different blocks cause a toy to activate) using one of two causal toys. In the *standard* condition, blocks were placed in pairs on top of the toy, while in the *relational* condition, each block was placed inside one of two transparent openings on either side of the toy. The physical design of the latter toy was intended to highlight the relationship between pairs of blocks. Results suggest that even 3-year-olds' causal inferences are sensitive to design, with children in the *relational* condition more likely to infer the abstract relation than those in the *standard* case. These results provide strong evidence that design serves as a constraint on causal inference in early childhood. Findings are discussed in terms of their implications for creating intuitive learning environments for young children.

Keywords: cognitive development; causal inference; relational reasoning; learning environments; design

Introduction

When reasoning about novel causal relationships, learners must select the most likely hypothesis from a range of underdetermined possibilities. For example, to activate a novel appliance, you might consider several possible interventions: the 'on/off' switch might have to be flipped, the reset button on the circuit interrupter might have to be depressed, or perhaps both the switch and the button together activate the device. Depending on your prior belief in the likelihood of each candidate cause and your subsequent observations, you then select the most likely action. If, for example, the switch is in the 'on' position and the appliance does not activate, it provides evidence that it must be activated in conjunction with the interrupter reset. However, one could also imagine a seemingly infinite number of alternative ways the causal system may work. Perhaps the appliance is voice activated, or the buttons need to be pushed in a particular repeating order, or there is an additional hidden switch somewhere else on the device.

To solve the infinite hypothesis search problem, recent work emphasizing the psychological processes underlying inductive inference has proposed that learners likely "sample" from this vast space of hypotheses, based on prior knowledge (e.g., Bonawitz, Denison, Griffiths, & Gopnik, 2014; Ullman, Goodman, & Tenenbaum, 2012; Tenenbaum,

Griffiths & Kemp, 2006). Thus, instead of considering all possible hypotheses and weighing each against the observed evidence, learners may only generate a subset of the most likely candidates to evaluate (Bonawitz & Griffiths, 2010). Critically, the specific subset of hypotheses that is generated for a particular learning problem may depend on a variety of factors, including their prior probability, their relevance to the current problem, priming, and so forth (e.g., Dougherty & Hunter 2003; Flin, Slaven & Stewart, 1996; Klein, 1993; Weber et al., 1993; Schunn & Klahr, 1993; Koehler, 1994). In fact, even young children are sensitive to input that constrains the hypotheses they consider, including information about the problem they are trying to solve, how the data were sampled, who generated the evidence, and why (e.g., Buchsbaum, Gopnik, Griffiths, & Shafto, 2011; Butler & Markman, 2012; Gergely, Bekkering & Kiraly, 2002; Walker, Lombrozo, Legare, & Gopnik, 2014).

Accordingly, any input that changes a learner's prior expectations about the most likely causal structure can influence the hypotheses they privilege, and ultimately apply. Here, we consider a specific environmental cue that, to our knowledge, has not yet been examined: the visible *design* of the object itself. If children use information about an object's design to constrain the hypotheses they generate, changes in the *physical features* of the learning context might influence causal learning and discovery. That is, the design of a causal system may serve to increase or decrease the salience of some hypotheses over others.

Effects of Design on Behavior

Although object design has not been specifically examined in the context of causal learning, there are several reasons to expect that the physical features of the learning context may influence children's causal inference. Indeed, nearly all of the objects we interact with include some element of design, and we often use these cues to infer information about an object's function. For example, if a door has no handle, the only way to enter is to push. While this action seems intuitive, the design is intentional. The creator constructed the door so that its physical features would constrain the permissible actions. Norman (1988) includes such constraints as one of several principles of good design, recognizing that design impacts reasoning about object function. A large body of literature has also explored the ways in which subtle environmental influences, or "nudges," have disproportional effects on human choice (Thaler & Sunstein, 2008), impacting hygiene (Holland, Hendriks, & Aarts, 2005), energy use, (Allcott & Mullainathan, 2010), and health (Thorndike, Sonnenberg,

Riis, Barraclough, & Levy, 2012; van Nieuw-Amerongen, Kremers, De Vries, & Kok, 2011), among others.

Other applied research has also begun to examine whether environmental design can change the way we *learn* in select educational contexts. For example, museum designers have used exhibit access, visibility, and object affordances to encourage visitor exploration, engagement, and understanding (e.g., adding a knob to a display suggests that an object can be moved, adding a glass window on the side of a machine encourages visitors to view the internal mechanism; see Allen, 2004; Wineman & Peponis, 2010; Shin, Park & Kim, 2014). Here, we go beyond this past applied work to consider whether similar cues can influence the salience of certain concepts or reasoning strategies in the context of causal learning. That is, we test whether elements of design influence a learner's prior beliefs about the likelihood of a particular causal hypothesis, given some pattern of evidence.

To illustrate how the design of an object might impact a learner's beliefs about its function, we return to our novel appliance. If you are familiar with electronic machines, you might believe that before you can turn something on, you must connect its cord to a power source. Once learned, this general principle can be widely applied to novel cases, even before observing any evidence about how a particular appliance functions. However, now consider a situation in which you are confronted with an appliance that has *two* cords. In this case, your prior belief that a single power cord must be plugged in to turn on the machine seems less probable. You might instead form a hypothesis that the two cords must *both* be connected before the machine will turn on. This demonstrates an even more general assumption that the features of an object are relevant to its function. This sort of abstract causal principle, or "overhypothesis," is a belief about the *kinds* of hypotheses that are most likely to be true (Goodman, 1955; Kemp, Perfors, Tenenbaum, 2007). Based on the learner's prior experience, it might seem unlikely to observe a second power cord that is unnecessary for the machine's operation (without an alternative explanation for the presence of the second cord). In this way, the visible features of an object serve as critical design cues that constrain the hypotheses that are generated about its causal structure (Norman, 1988).

Some existing support for the proposal that an object's design serves to constrain inferences about causal structure can be found in the literature examining human reasoning about artifacts (i.e., human-made objects). That is, both children and adults view features of artifacts as reflective of that object's function and intended use (Keil, 1992; Keleman, 1999; Keleman, Seston, & Saint Georges, 2012). For example, Kelemen and colleagues (2012) showed preschool-aged children two objects that were equally optimal for performing a particular function (i.e., both objects featured a flat surface that could be used to crush popcorn), but one of them had additional salient features that suggested it could *also* be used for an additional purpose (i.e., spikes along the object's handle). When asked

which object was designed for the target purpose (crushing popcorn), 3- and 4-year-old children privileged the object with a more efficient design.

Magid and colleagues (2015) also provide evidence that children relate an object's design to its function. The authors argue that young learners represent the abstract criteria for solving a problem, before arriving at a precise solution. These criteria are based on how well a particular hypothesis matches the abstract "form" of the problem to be solved. Specifically, 4 and 5-year-olds mapped the *type* of effect produced (a discrete vs. continuous visual effect) to the *type* of mechanism that produced it (a binary "on/off" switch vs. a dial), providing evidence that children relate the physical structure of an object's causal mechanism to its effect. Additionally, 4- and 5-year-olds have also been shown to map the quantity and diversity of object functions (e.g., making cupcakes vs. making cupcakes *and* wrapping presents) to make inferences about the complexity of the design of its internal mechanism (Ahl & Keil, 2016).

The Current Approach

In the prior work reviewed above, learners made inferences about the design of objects, given information about possible functions. Here, we ask whether children can perform a more challenging task -- whether they will be more likely to generate a particular causal hypothesis, given the object's design. In particular, we present a conceptual case in which 3-year-olds typically fail to discover a relational hypothesis. We then assess whether the object's design influences learning by observing whether subtle changes to the physical structure of the causal system leads to the successful identification of the abstract relational cause.

Specifically, we present 3-year-olds with a relational reasoning problem that they systematically fail at this age (Walker, Bridgers & Gopnik, 2016). In this task, children are introduced to a novel toy that plays music for some objects and not for others (i.e., a "blicket detector," Gopnik & Sobel, 2000). They then observe pairs of blocks being placed on top of the toy. When 3-year-olds are provided with evidence that the toy's activation is caused by the relation between the two blocks in each pair (i.e., whether the blocks are the same or different), rather than by individual object kinds (i.e., blocks of a particular shape and color), they failed to make the correct causal inference at test (see Figure 1).

Notably, younger children (18 to 30-month-olds) successfully infer *same-different* relations in this task, suggesting that later failures are due to a difference in tendency, not a lack of relational competence (Walker et al, 2016; Walker & Gopnik, 2017; Walker, Walker, & Gopnik, under review). In other words, these developmental data provide evidence that older children are capable of inferring such relations, even if they do not spontaneously generate them in most learning scenarios. Critically, this proposal contrasts with decades of research suggesting that preschoolers were simply unable to reason on the basis of

these abstract relations (e.g., Christie & Gentner, 2010; 2014).

Based on these findings, it has been proposed that 3-year-olds' long-documented failure to infer *same-different* relations results from a learned bias in the form of an overhypothesis that privileges the role of individual objects over the relations between them (Walker et al., 2016; Carstensen & Walker, 2017). Walker and colleagues (2016, Exp 3) provide additional support for this idea, demonstrating that prompting children to explain during training trials significantly increases their tendency to endorse the relational hypothesis at test. The authors propose that explanation likely serves as an *internal* constraint on hypothesis search, leading learners to privilege more abstract solutions. This domain therefore provides a promising case study to explore the proposal that an *external* constraint, namely the design of an object, can also influence hypothesis generation in causal learning.

In order to assess whether the tendency to discover the relational hypothesis may be sensitive to constraints imposed by physical design, we made one small modification to the standard causal relational task: Rather than placing pairs of blocks on top of the toy on a single, large platform, the blocks were inserted into two transparent openings (see Figure 2). By adding these two intentionally designed openings, a learner who treats object design as relevant to their causal inferences might consider *why* the causal system included these features. These two openings therefore not only draw attention to the presence of two objects, but also suggest a particular affordance: that the machine activates by combining the two. As a result, this may raise the possibility that the *relation between* the blocks—rather than the identity of the blocks themselves—is relevant to the causal structure, leading to the discovery of the relational hypothesis. We return to consider the implications of this particular design choice in the discussion.

Alternatively, it may be the case that the design of the causal system has no effect on 3-year-olds' endorsement of the relational hypothesis. As noted, children at this age repeatedly fail to spontaneously privilege relational information (Christie & Gentner, 2010; 2014; Walker et al., 2016), suggesting a strong prior for hypotheses based on individual object kinds. To correctly infer the relational hypothesis in this case, children must integrate information about the object's design with their prior beliefs about likely causes, taking into account why object design is relevant, *and* weighing this information more heavily than their prior commitment to the object-based hypothesis. That said, if children's failure to infer abstract relations indeed results from a difference in *tendency*, rather than a lack of *competence* (as has been suggested), *and* they are sensitive to the design of the learning context, then we might reasonably expect them to successfully infer the abstract relational hypothesis, even following such a minor modification to the standard task.

Methods

Participants

A total of 152 3-year-olds participated in the study, with 76 children randomly assigned to either the *standard toy* ($M = 41.9$ months; 36 female) or *relational toy* ($M = 41.6$ months; 37 female) conditions. Within each condition, half of the children observed evidence consistent with the *same* relation and half observed evidence consistent with the *different* relation. Sample size satisfies a power analysis with power $> .8$, given an alpha of $.05$ and an effect size of $.3$ (medium). An additional 9 participants were excluded due to experimenter error (3), failure to complete the study (4), parent interference (1), or interference by another child (1). Children were recruited and tested in the lab, at preschools, and at museums. All participants were tested in a quiet, private room with the experimenter.

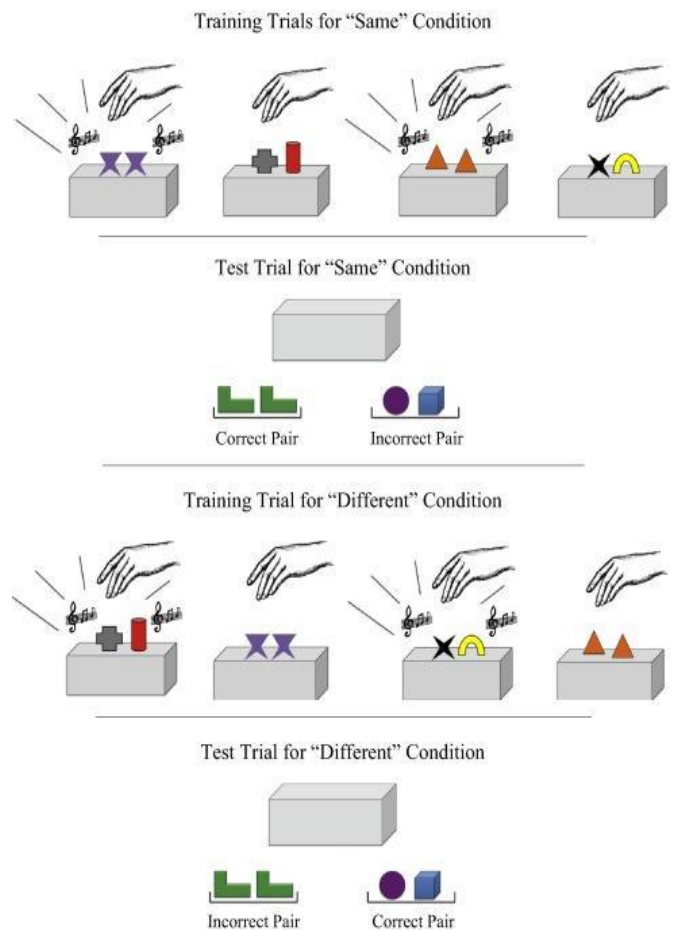


Figure 1. Schematic illustration of evidence presented during training and test trials in the *standard* condition (reprinted from Walker et al., 2016, Exp. 1). Identical pairs and outcomes were presented in the *relational* condition, using the relational toy (see Fig. 2).




Figure 2. Standard and relational toys

Materials and Procedure

The materials and procedure for the *standard toy* condition replicate those used by Walker et al. (2016, Exp. 1; see Fig. 1). Children were seated at a table across from the experimenter. The experimenter began by placing an opaque cardboard box on the table, saying “This is my toy! Sometimes when I put things on top, the toy will play music, and other times it does not. Should we try some and see how it works?” As in previous research, the toy appeared to activate and play a novel melody in response to certain combinations of blocks. In fact, the experimenter activated a wireless doorbell inside the box by surreptitiously pressing a button.

A total of 4 pairs of *same* and *different* painted wooden blocks (2 pairs of *same* and 2 pairs of *different*) were used during the training trials. After introducing the toy, the experimenter produced two blocks in either the *same* or *different* relation (depending upon the condition), and said, “Let’s try!,” and put both blocks on top of the toy, simultaneously. The toy played music and the experimenter said, “Music! My toy played music!” The experimenter then picked up the blocks and set them back on the toy, which again played music, saying “Music! These ones made my toy play music!” She then repeated this procedure with a new pair of blocks in the opposite relation. The new pair did not make the toy play music, and the experimenter responded to the first try with, “No music! Do you hear anything? I don’t hear anything,” and after the second try, said “No music. These ones did not make my toy play music.” This pattern was repeated with two additional pairs of blocks, one in each relation. The experimenter always began with a causal pair (identical blocks in the *same* condition and blocks of unique colors and shapes in the *different* condition), and then alternated inert, causal, inert, using novel blocks in each new pair, and randomizing the specific blocks between participants.

After the four training trials, the experimenter said “Now that you’ve seen how my toy works, I need your help finding the things that will make it play music. I have two choices for you.” The experimenter presented the child with

two new pairs composed of novel blocks, one “same” pair and one “different” pair. Each pair was presented on a plastic tray, which the experimenter held up, saying, “I have these, and I have these (directing the child’s attention to each pair). Only one of these trays has the things that will make my toy play music. Can you point to the tray that has the things that will make it play?” The trays were then placed out of the child’s reach, on either side of the toy, with each pair set an equal distance from the child. The order and side of presentation of the correct pair counterbalanced between participants. The experimenter recorded the child’s first point or reach, scoring the response as correct (1) if the child chose the test pair (same or different) that corresponded to her training, and incorrect (0) for the opposite pair.

The materials and procedures for the *relational toy* condition were identical to those in the *standard toy* condition with one critical difference: The design of the toy was modified to include two transparent openings located on either side (see Fig. 2). The openings were constructed using clear, 2” x 2” hard plastic boxes. When children observed each of the training trials described above, pairs of blocks were inserted into the two openings (one block on either side), rather than placed on top of the toy. This was the only difference between the two conditions.




Figure 3. Mean proportion of correct relations by condition. Error bars indicate +/- 1 SEM. Chance performance is indicated by the dotted line.

Results

Replicating previous work (Walker et al., 2016), 3-year-old children in the *standard toy* condition responded at chance (46%), $p = .57$ (two-tailed, exact binomial), with no difference in performance between *same* (53%) and *different* (40%) training trials, $p = .35$ (two-tailed, Fisher’s exact). In contrast, 3-year-olds in the *relational toy* condition succeeded in selecting the test pair that was consistent with their training (66%), $p = .008$ (two-tailed exact binomial), performing identically in *same* and *different* training trials ($p = 1$, two-tailed, Fisher’s exact). Comparing performance across conditions, children in the

relational toy condition significantly outperformed those in the *standard* condition ($p = .022$, two-tailed, Fisher's exact) in inferring *same-different* relations.

Discussion

In the current study, we present findings demonstrating that children are indeed sensitive to the physical design of the learning context when reasoning about causal relationships. Although 3-year-olds in the *standard toy* condition failed to recognize the relational hypothesis (replicating prior work), increasing the salience of this hypothesis through the application of a relatively subtle design cue significantly increased their tendency to engage in relational reasoning in this task. In addition to providing evidence for the role of design in constraining causal inference, these data provide additional support for the proposal that children's failure on relational reasoning tasks results from a difference in *tendency*, not a lack of competence (e.g., Walker & Gopnik, 2014; Walker et al., 2016; Carstensen & Walker, 2017; Walker & Gopnik, 2017).

These results are particularly striking given that 3-year-olds have repeatedly failed to spontaneously privilege relational information (Christie & Gentner, 2010; 2014), suggesting a very strong prior to prefer individual object kinds. In order to use the design of the learning context to override this tendency and privilege the relational hypothesis, these very young children had to make a particularly sophisticated inference: They must have noticed this subtle design cue, inferred its relevance to the system's causal structure (i.e., that an object's design is relevant for its function), and weighed this information more heavily than their (strong) prior commitment to the object-based hypothesis.

These surprising findings therefore suggest that relatively minor elements of design can radically change the distribution of a learner's prior expectations, constrain the type of hypotheses that are generated, influence learning outcomes, and even facilitate the early discovery of new causal beliefs. Our results join prior research suggesting that hypothesis generation can be influenced by a variety of cognitive factors (e.g., Dougherty & Hunter 2003; Flin, Slaven & Stewart, 1996; Klein, 1993; Weber et al., 1993; Schunn & Klahr, 1993; Koehler, 1994), prompts (Walker et al., 2014, 2016; Williams & Lombrozo, 2010) and social inferences (Butler & Markman, 2012; Buchsbaum et al., 2011; Gergely, Bekkering, & Kiraly, 2002), and extends this work to include the structure of the learning environment itself. Ongoing work examines whether and how design influences even more entrenched causal beliefs and biases (e.g., in adults; Walker, Rett, & Bonawitz, in prep), and considers how design may interact with other constraints, such as pedagogical cues or prompts to explain. For instance, in some contexts, children privilege an object's visible affordances over an actor's intentional behavior when reasoning about how an artifact is intended to be used (e.g., DiYanni & Keleman, 2008). Future work will explore to what extent learners may be reasoning about the

intentions of the designer (as a social agent) when making inferences based on these environmental cues.

There are also open questions surrounding how the particular design modifications used in this experiment influence children's reasoning. One possibility is that the addition of exactly two transparent openings on either side of the toy directly primed the relational hypothesis. Another possibility is that this design cue simply served to disrupt children's initial intuitions about the likely causal mechanism, leading them to consider alternatives more broadly. If so, this may have made it more likely for children to discover the relational hypothesis, albeit indirectly. Future work is needed to address these important questions.

Finally, these results have clear practical implications for early science education, and in particular, the design of formal and informal learning environments intended for children. Our findings dovetail with literature in education pointing to the importance of "mise en place" or setting the stage for learning (Weisberg, Hirsh-Pasek, Golinkoff, & McCandliss, 2014). As demonstrated here, children are sensitive to relatively subtle physical cues in the learning environment when they are engaged in causal reasoning. This simple manipulation led children to consider a relational hypothesis that they typically fail to spontaneously produce. Our findings therefore highlight the importance of careful design when aiming to teach children specific concepts, given that the visible features of objects may increase or decrease the salience of the available evidence, and change the learner's interpretation of their observations. It is impossible to create artifacts without also making specific design choices, so being aware of how these features might be used to facilitate reasoning can have major consequences for learning and instruction. These findings therefore open up new avenues for future work examining how the design of learning environments can be used to support belief revision and guide early learning and discovery.

Acknowledgements

We are grateful to Mike Frank for his early contributions in conceptualizing the methods for this study. We thank Alicia Lunardhi and Emily To for their efforts towards data collection and Nicky Sullivan for facilitating recruitment. We also thank the Fleet Science Center, the Birch Aquarium at Scripps Institute of Oceanography, and the New Children's Museum, as well as all of the participating preschools and families who made this research possible. This research was funded by a Hellman's Fellowship, awarded to C. Walker.

References

- Ahl, R. E., & Keil, F. C. (2017). Diverse effects, complex causes: children use information about Machines' functional diversity to infer internal complexity. *Child development*, 88(3), 828-845.

- Allcott, H., & Mullainathan, S. (2010). Behavior and energy policy. *Science*, 327(5970), 1204-1205.
- Allen, S. (2004). Designs for learning: Studying science museum exhibits that do more than entertain. *Science Education*, 88(S1), S17-S33.
- Bonawitz, E., Denison, S., Griffiths, T. L., & Gopnik, A. (2014). Probabilistic models, learning algorithms, and response variability: Sampling in cognitive development. *Trends in cognitive sciences*, 18(10), 497-500.
- Bonawitz, E., & Griffiths, T. (2010). Deconfounding hypothesis generation and evaluation in Bayesian models. In S. Ohlsson & R. Catrambone (Eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 2260–2265). Austin, TX: Cognitive Science Society.
- Buchsbaum, D., Gopnik, A., Griffiths, T. L., & Shafto, P. (2011). Children's imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition*, 120(3), 331-340.
- Butler, L. P., & Markman, E. M. (2012). Preschoolers use intentional and pedagogical cues to guide inductive inferences and exploration. *Child development*, 83, 1416-1428.
- Carstensen, A. & Walker, C.M. (2017). The paradox of relational development is not universal: Abstract reasoning develops differently across cultures. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. Davelaar (Eds.), *Proceedings of the 39th Annual Conference Cognitive Science Society*, pp. 1721-1726. London, UK: Cognitive Science Society.
- Christie, S., & Gentner, D. (2010). Where hypotheses come from: Learning new relations by structural alignment. *Journal of Cognition and Development*, 11(3), 356-373.
- Christie, S. & Gentner, D. (2014). Language helps children succeed on a classic analogy task. *Cognitive Science*, 38(2), 383-397.
- DiYanni, C., & Kelemen, D. (2008). Using a bad tool with good intention: Young children's imitation of adults' questionable choices. *Journal of experimental child psychology*, 101(4), 241-261.
- Dougherty, M. R., & Hunter, J. E. (2003). Hypothesis generation, probability judgment, and individual differences in working memory capacity. *Acta psychologica*, 113(3), 263-282
- Flin, R., Slaven, G., & Stewart, K. (1996). Emergency decision making in the offshore oil and gas industry. *Human Factors*, 38(2), 262-277.
- Gergely, G., Bekkering, H., & Kiraly, I. (2002). Rational imitation in preverbal infants. *Nature*, 415, 755.
- Goodman, N. (1955). *Fact, fiction, and forecast*. Cambridge, MA: Harvard University Press.
- Holland, R. W., Hendriks, M., & Aarts, H. (2005). Smells like clean spirit: Nonconscious effects of scent on cognition and behavior. *Psychological Science*, 16(9), 689-693.
- Keil, F. C. (1992). *Concepts, kinds, and cognitive development*. MIT Press.
- Kelemen, D. (1999). Function, goals and intention: Children's teleological reasoning about objects. *Trends in Cognitive Sciences*, 3(12), 461-468.
- Kelemen, D., Seston, R., & Saint Georges, L. (2012). The designing mind: Children's reasoning about intended function and artifact structure. *Journal of Cognition and Development*, 13(4), 439-453.
- Kemp, C., Perfors, A., & Tenenbaum, J. B. (2007). Learning overhypotheses with hierarchical Bayesian models. *Developmental science*, 10(3), 307-321.
- Klein, G.A. (1993). A Recognition-Primed Decision (RPD) Model of Rapid Decision Making. In G. A. Klein, J. Orasanu, R. Calderwood, and C. Zsombok (Eds.), *Decision Making in Action: Models and Methods*. Norwood, NJ: Ablex Publishing Corp., 138–147.
- Koehler, D. J. (1994). Hypothesis generation and confidence in judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(2), 461.
- Magid, R. W., Sheskin, M., & Schulz, L. E. (2015). Imagination and the generation of new ideas. *Cognitive Development*, 34, 99-110.
- Norman, D. (1988). *The design of everyday things*. New York: Basic Books.
- Schunn, C.D., & Klahr, D. (1993) Self vs. Other-Generated Hypotheses in Scientific Discovery. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*.
- Shin, H., Park, E. J., & Kim, C. J. (2014). Learning affordances: Understanding visitors' learning in science museum environment. In *Topics and Trends in Current Science Education* (pp. 307-320). Springer, Dordrecht.
- Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in cognitive sciences*, 10(7), 309-318.
- Thaler, Richard H., Sunstein, Cass R. (2008) *Nudge: improving decisions about health, wealth, and happiness* New Haven: Yale University Press.
- Thorndike, A. N., Sonnenberg, L., Riis, J., Barraclough, S., & Levy, D. E. (2012). A 2-phase labeling and choice architecture intervention to improve healthy food and beverage choices. *American Journal of Public Health*, 102(3), 527-533.
- Ullman, T. D., Goodman, N. D., & Tenenbaum, J. B. (2012). Theory learning as stochastic search in the language of thought. *Cognitive Development*, 27(4), 455-480.
- van Nieuw-Amerongen, M. E., Kremers, S. P. J., De Vries, N. K., & Kok, G. (2011). The use of prompts, increased accessibility, visibility, and aesthetics of the stairwell to promote stair use in a university building. *Environment and Behavior*, 43(1), 131-139.
- Walker, C. M., Bridgers, S., & Gopnik, A. (2016). The early emergence and puzzling decline of relational reasoning: Effects of knowledge and search on inferring abstract concepts. *Cognition*, 156, 30-40.

- Walker, C.M. & Gopnik, A. (2014). Toddlers infer higher-order relational principles in causal learning. *Psychological Science*, 25(1): 161-169.
- Walker, C.M., & Gopnik, A. (2017). Discriminating relational and perceptual judgments: Evidence from human toddlers. *Cognition*, 166, 23-37.
- Walker, C. M., Lombrozo, T., Legare, C. H., & Gopnik, A. (2014). Explaining prompts children to privilege inductively rich properties. *Cognition*, 133(2), 343-357.
- Walker, C.M., Rett, A., & Bonawitz, E. (in prep). Design drives discovery in causal learning. *Manuscript in preparation*.
- Walker, C.M., Walker, J.C., & Gopnik, A. (under review). Toddlers generalize abstract representations of *same* and *different*.
- Weber, E.U., Böckenholt, U., Hilton, D.J., Wallace, B. (1993). Determinants of diagnostic hypothesis generation: effects of information, base rates, and experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19, 1151–1164
- Weisberg, D. S., Hirsh-Pasek, K., Golinkoff, R. M., & McCandliss, B. D. (2014). Mise en place: Setting the stage for thought and action. *Trends in Cognitive Sciences*, 18(6), 276-278.
- Williams, J. J., & Lombrozo, T. (2010). The role of explanation in discovery and generalization: Evidence from category learning. *Cognitive Science*, 34, 776–806. doi:10.1111/j.1551-6709.2010.01113.x
- Wineman, J. D., & Peponis, J. (2010). Constructing spatial meaning: Spatial affordances in museum design. *Environment and Behavior*, 42(1), 86–109.