Imagination and the generation of new ideas

Rachel W. Magid [a], Mark Sheskin [b], Laura E. Schulz [c],*

[a] Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA
[b] Department of Psychology, Yale University, New Haven, CT, USA
[c] Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA

ARTICLE INFO

ABSTRACT

A variety of theories have been put forth to explain the function of imagination, most notably that imagination engages and develops children's theory of mind and counterfactual reasoning. Here, we propose that a primary role for imagination is as a cognitive mechanism for efficiently generating new ideas without observing new evidence. Learners must generate hypotheses before they can assess the truth of these hypotheses. Given infinite possibilities, how do learners constrain the process of hypothesis generation? We suggest that learners represent abstract criteria for the solution to a problem and generate solutions that, if true, would solve the problem. As a preliminary test of this idea, we show that, in the absence of any fact of the matter (i.e., when neither prior knowledge nor statistical data distinguishes competing hypotheses), 4–6-year-olds (mean: 63 months) systematically converge on solutions to problems, consistent with an ability to imagine the abstract properties of causal problems and their solutions.

© 2015 Published by Elsevier Inc.

1. Introduction

Imagination pervades human experience. Children begin engaging in pretend play as toddlers (Fein, 1981; Singer & Singer, 1992), and although cultural and parental attitudes affect the amount and content of imaginary play (Gosso, Morais, & Otta, 2007; Haight, Parke, &. Black, 1997), researchers have observed imagination in every culture where they have looked (Farver & Shin, 1997; Farver &

* Corresponding author.
  E-mail address: lschulz@mit.edu (L.E. Schulz).

Wimbarti, 1995; Haight, Wang, Fung, & Williams, 1999). As adults, we are avid consumers and creators of fiction (Harris, 1998; Oatley, 1999), and we respond viscerally and emotionally to imagined scenarios (Carruthers, 2009; Vrana & Lang, 1990). Moreover, we invent fictions even in the pursuit of facts: we confabulate in the face of neurological disorders (Phelps & Gazzaniga, 1992), in defending the bases of our decisions (Nichols & Stich, 2000), and in the construction of autobiographical memory (Kopelman, 1987).

Why do we make things up? Given the uncertainties and complexities of the real world, why do we spend cognitive effort on unreal worlds? Arguably, we do so precisely because the real world is uncertain and complex. Thinking about possible worlds may prepare us for future events in the actual world. Any version of this account, however, must contend with Fodor's farcical endorsement of it:

> . . . what if it turns out that, having just used the ring that I got by kidnapping a dwarf to pay off the giants who built me my new castle, I should discover that it is the very ring that I need in order to continue to be immortal and rule the world? It is important to think out the options betimes, because a thing like that could happen to anyone and you can never have too much insurance. (Fodor, 1998, p. 212)

The sheer fecundity of our imagination poses a problem for functionalist accounts. If the primary role of fantasy is to explore possible realities (the "conundrums we might face someday"; Pinker, 1997, p. 543), shouldn't we have more realistic fantasies?

Still, it does not seem unreasonable to suppose that a universal, early emerging cognitive ability must be good for *something*. One possibility is that imagination provides an attractive package for ordinary cognition (Boyd, 2009). By embedding useful knowledge in extraordinary events with heightened emotional content, learners may be better able to access important cultural skills or facts. Consistent with this, researchers have suggested that imaginative engagement might support a range of cognitive abilities, including creativity, intelligence, problem solving, symbolic reasoning, language development, theory of mind, narrative skills, social skills, causal reasoning, emotional regulation, and executive function. Dismayingly, however, a recent exhaustive review of the literature found little to no evidence that imaginative play supports cognition in any domain for which a benefit has been proposed (Lillard et al., 2013). Of course, the fact that cognitive development is robust to variations in imaginative engagement need not mean imagination is irrelevant to cognition. Cognitive development is also robust to variations in sight and hearing; this does not make perceptual abilities epiphenomenal. Nonetheless, the absence of evidence for any direct causal relationship between pretense and cognitive outcomes makes determining the role of imagination especially challenging.

The challenge is magnified by the polysemy of the central concept. As discussed, imagination may be involved in everything from the play of toddlers (Singer & Singer, 1992) to the confabulatory behavior of neuropsychiatric patients (Kopelman, 1987). We might limit our study of imagination to its manifestation in better-understood aspects of cognition (mental imagery, theory of mind, or counterfactual reasoning). However, although it is clear that our abilities to simulate future states, represent events to which we do not have immediate access, and reason through the consequences of false premises are critical to cognition, it is less clear that describing these abilities as imaginative adds to what we already know about such cognitive processes.

Given this state of affairs, we suggest a new approach. One way to understand the role of imagination in cognition may be to consider it in relationship, not to those aspects of cognition that are relatively well understood, but to other puzzles of cognitive science. Here we focus on the problem of how learners think of new ideas.

At first glance, the topic of how learners generate new ideas might seem like a well-studied problem, and the last place to look for unsolved puzzles. Decades of work in cognitive development have investigated the processes underlying theory change and conceptual change (Carey, 2009; Gopnik & Wellman, 2012; Schulz, 2012). However, in the (understandable) focus on how learners change their beliefs in deep, far-reaching ways, a more commonplace mystery may have been obscured: the mystery of ordinary thought.

To illustrate what we mean, consider two questions, united only in that you probably have not considered them before: (1) What should you name a theater company focusing on new works? (2)

How do manufacturers get the red stripes on peppermints? If you find that you can opine freely on these questions, it is surely not because you were engaged in theory change (Gopnik & Wellman, 2012) or radical conceptual change (Carey, 2009). If asked what you were doing in generating answers, you would probably answer, "Just thinking." However, understanding our ability to "just think" is a hard problem of cognitive science, and one, we suggest, that links the problem of imagination to the problem of learning. Confronted de novo with questions about peppermints or contemporary theater, the learner may have a wealth of relevant prior knowledge, including specific cultural knowledge and an abstract understanding of folk physics and folk psychology. What she does not have is the answers to these questions. She must make them up. Any answers she does make up will presumably be consistent with her existing beliefs. Nonetheless the answers are genuinely new, in that the learner did not have the ideas before she "made them up," and genuinely "made up," in that they emerged from thinking, not from new evidence or testimony.

Before considering how we might solve problems like these, we want to stress that thinking of new ideas is not an optional exercise in creativity; it is fundamental to learning. The peppermint question might be trivial but the ability to answer it is not. Before we can select among competing ideas, we must generate them. Here, we are interested in how, in the absence of new information, a learner can construct a hypothesis space.

The best current answers to this question come from a family of computational models suggesting that learners have a "grammar" of constitutive, causal, and logical relations that obtain over variables in their theories. The grammar (like natural language grammars) lets them generate infinitely many hypotheses (Goodman, Ullman, & Tenenbaum, 2011; Ullman, Goodman, & Tenenbaum, 2012). When existing hypotheses fail to predict observed data, this account suggests that the learner can generate new ideas by proposing random changes to existing hypotheses, constrained by the probability that the theory grammar will generate the new hypothesis. The process of hypothesis generation can be made more efficient by the re-use of existing "templates": grammatical predicates that encode common relationships among variables, like transitivity or equivalence. After the learner generates a new hypothesis through this grammar-based search, she can then test whether the hypothesis applies to the data; if a new hypothesis predicts the data better than the earlier hypothesis, then the new hypothesis is likely to be retained. The authors refer to this approach as "theory-driven stochastic search," and their model involves two search processes: a search through an "outer loop" (in which the model makes random changes to the hypothesis constrained by the theory grammar) and a search through an "inner loop" (in which the model checks the fit between the hypothesis and the data). It thus represents a welcome advance in data driven learning.

Nonetheless, this approach is insufficiently constrained. To understand why, consider a question to which you already know the (putative) answer: why you cannot use your cell phone on an airplane flight. If you do not already know about radio interference (or bureaucratic regulations), you might use theory-driven stochastic search to add or delete simple predicates relating planes and phones. The difficulty is that there are innumerable logical, constitutive, causal, and relational hypotheses consistent with the grammar of your intuitive theories that might connect the two artifacts. You might consider the possibility that planes have numbers on them and so do phones, that planes are bigger than phones, or that planes and phones are both made in China. None of these ideas predicts that you have to turn off your phone, so they will be rejected as explanations. However, the need to generate, only to reject, innumerable irrelevant ideas suggests the difficulties involved in converging on useful new hypotheses by making simple, logical, random, theory-consistent changes to your current beliefs.

Now consider that despite these difficulties, one 4-year-old child heard the pilot's announcement and promptly said, "I know why you have to turn off your cell phone during takeoff. It's because when the plane takes off it's too noisy to hear." Unlike most examples cited by developmental psychologists about children (especially their own), this answer is not particularly cute, or clever, or surprising. Moreover, the answer is wrong. However, this wrong response is vastly better than innumerable possible responses (e.g., that planes are bigger than phones) that are not even wrong. An extraordinary feature of ordinary cognition is that we can generate ideas that are simultaneously "wrong" and "good." We can consider independently the extent to which a hypothesis fulfills the abstract goals of a solution to a problem, and the degree to which a hypothesis fits the data. How might we do this?

4   *R.W. Magid et al. / Cognitive Development xxx (2015) xxx–xxx*

We propose that we can generate new ideas "on the fly" because, well before we know the solution to a problem, we have a lot of abstract knowledge about what a particular solution has to do. A 4-year-old may not know why you need to turn off your phone, but she may recognize that the problem involves not just any relationship between planes and phones, but an unpredicted *incompatibility* between planes and phones. Insofar as this representation of the problem constrains her search for a solution, she need not bother hypothesizing that planes and phones have innumerable commensurable features, nor need she consider other inductively simple, plausible, logical propositions consistent with her prior knowledge. She can generate only hypotheses in which something about phones is in tension with something about planes. In a similar vein, you may not have thought of a great new name for a theater company yet, but we suspect that "McDonald's" and "Asaccharolyticus" were not on your list. You may not know a great new name, but you know some abstract goals for the great new name, such as that it should be novel and pronounceable. Similarly, you may have no idea how peppermints get their stripes, but you know you are looking for something that generates patterns, not randomness; thus you may be more likely to propose some kind of pendulum mechanism than an atomizer spray.

We suggest that the ability to represent the abstract form of a problem and the desiderata of the solution acts as a critical top-down constraint on hypothesis generation. It is beyond the scope of this paper (and beyond the authors' expertise) to provide a formal account of how this process might work. What we will try to do is establish the intuitions underlying this idea and its connection to imagination. We will then provide some empirical evidence that these kinds of abilities are manifest in the reasoning of young children.

The fundamental intuition is that we generate new ideas with goals in mind. These goals, and the criteria for fulfilling them, derive from multiple sources, including the particular kinds of problems we want to solve (e.g., navigation, explanation), our broader epistemic ends (e.g., persuading, illustrating, instructing, deceiving), and ends that are not epistemic at all (e.g., impressing, entertaining). However, although we have a vast range of possible goals, any given goal can only be achieved in a relatively limited number of ways. Insofar as goals specify abstract criteria for their fulfillment, they may make it more likely that we will generate some hypotheses over others (e.g., if our goal is navigation we are likely to generate hypotheses that take the form of 2D maps; if our goal is explanation we are likely to generate hypotheses that take the form of causal chains or tree structures).

Many otherwise puzzling intuitions about human cognition follow from the notion that we generate new ideas informed by an abstract representation of what a good solution would look like. For instance, we often know that we are on the right (or wrong) track in thinking about a problem, even though we cannot account for observed data any better (or worse) than we could before. This may be because we can recognize that the hypotheses we are generating have (or lack) key features of the abstract form of the solution to a problem well before they are in fact solutions to a problem. Similarly, we may think we have struck upon a brilliant new idea well before we have tested or even fully articulated it. Indeed, even when the great new idea turns out to be false, it might only slightly diminish our sense of its brilliance. Again, this makes sense if we evaluate new ideas first with respect to how well they meet the abstract desiderata for the solution to a problem, and only secondarily with respect to their fit to prior beliefs and data. Finally, we often have intuitions about how tractable a problem is, even when tractability is not due to technical or prudential constraints. In such cases, what it might mean for a problem to be tractable is that we have a well-specified representation of the abstract form of the solution, even though we do not yet know what the solution is.

Note that constraining our generation of new hypotheses to those that serve functional goals need not *lessen* our sensitivity to their truth-value. Along the lines of the "inner loop" proposed in theory-guided stochastic search (Ullman et al., 2012), once generated, any hypothesis can be subject to fact-checking by assessing the extent to which it predicts observed data. However, the probability that we adopt a new hypothesis may be a joint function of the degree to which it meets the abstract constraints of our functional goals (by being explanatory, provocative, elegant, or useful) and its actual truth-value.

The comedian Stephen Colbert coined the term "truthiness" – the "gut feeling" that something sounds true without reference to facts – to poke fun at what we might value in ideas that are wrong. However, we suggest that a predilection for "truthiness" may be a feature, not a bug, of human cognition. Along with our inductive inference abilities, having truth-independent criteria for generating and valuing ideas may be critical to our ability to go beyond the data.

In short, we suggest that the ability to constrain hypothesis generation with respect to abstract goals is integral to thinking quite broadly. We further suggest that this ability is fundamental to acts of imagination ranging from pretend play to daydreaming to confabulation. Indeed, a core feature of imagination may be the representation and (in principle) fulfillment of epistemic or affective goals. On this account, imagination is critical not only because it allows us to draw inferences from false premises (Buchsbaum, Bridgers, Weisberg, & Gopnik, 2012; Walker & Gopnik, 2013), but because it provides an efficient way to generate potentially useful premises in the first place.

At this stage, many of our claims for the relationship between imagination and learning are speculative. We can, however, test some of the implications of this account. Here we ask whether children can distinguish good and bad hypotheses by considering relationships between the form of the problem and the form of the solution, even when there are no other differences in the plausibility of the hypotheses and no distinguishing data. Our approach does not require the learner to consider fantastical scenarios or counterfactual worlds. Nonetheless, we suggest that it provides a critical test of the relationship between imagination and learning. If imagination is fundamental to learning, then children should be able to solve problems by imagining abstract relationships that connect the problem and the solution, even when there is no other sense in which the answers are "correct" or "incorrect" with respect to the facts of the world.

2. Experiment 1

In Experiment 1, we showed children two different effects and two candidate causes. We then asked them to decide which of the two causes generates a given effect. The causes were equally plausible given children's prior knowledge of physical mechanisms, and we did not give the children any covariation evidence. However, one cause and effect involved discrete changes of state; the other cause and effect involved continuous changes of state.

If children are unable to imagine an abstract representation of a problem, they should choose at chance (since, indeed, there was no fact of the matter). However, if children represent the problem as one of distinguishing a discrete effect from a continuous effect and infer the causes accordingly, this would suggest that children are able to imagine abstract constraints on hypotheses that support reasoning about new ideas in the absence of new data.

2.1. Method

2.1.1. Participants

Sixteen 4–6-year-olds (mean age 62 months; range 48–82 months) were recruited at a children's museum. All children saw a pair of visual effects and listened to a pair of auditory effects; eight children were asked about the cause of a continuous visual effect and a discrete auditory effect, and eight were asked about the cause of a discrete visual effect and a continuous auditory effect.

2.1.2. Materials

A "machine" was made by disguising a desktop computer ($41 \times 4 \times 21$ cm) with a red cardboard box ($51 \times 51 \times 51$ cm). Only a 5-centimeter-wide diagonal strip across the center of the monitor was visible (see Fig. 1). Two shoeboxes ($27 \times 10 \times 19$ cm) were used as "controllers." A 5×5 cm Lego base was affixed to the box covering the monitor and to each controller. A silver cord with a Lego piece on each end was used to "plug" the controllers into the machine.

Each controller had two parts. On Controller A, the continuous part was a plastic pulley wheel anchored to the shoebox so that the wheel could spin freely; the discrete part was a wooden magnet that could be affixed in one of two distinct locations. On Controller B, the continuous part was a bead that could glide along half of a plastic ring affixed to the outside of the shoebox; the discrete part was a snap-lock bead that could fit inside either one of two other snap-lock beads affixed to the shoebox. The discrete part on Controller A was on the right, while the discrete part on Controller B was on the left. Each controller was associated with either the pair of visual effects or the

Fig. 1. Stimuli used in Experiments 1–4. The continuous part on Controller A was a pulley wheel; its discrete part was a magnet that could be placed in one of two distinct locations. The continuous part on Controller B was a bead that could slide along a plastic ring; its discrete part was a snap-lock bead. Each controller could be "plugged" into the machine (C) with a Lego attachment between the controller's cord and the top of the machine. One controller was used for the visual stimuli and one for the auditory stimuli (order counterbalanced). The visual stimuli were displayed through the opening on the machine; the auditory stimuli emanated from the machine. See text for details.

pair of auditory effects (counterbalanced across children). The effects each lasted 10 s and were as follows:

(1) The continuous visual effect was a spinning rainbow-colored ball (used as an attention-grabber in infant looking-time studies) programmed in MATLAB to move diagonally in a straight trajectory from the bottom left corner of the computer screen to the top right corner of the computer screen, and then back to the bottom left corner.

(2) The discrete visual effect was a static image of the same rainbow-colored ball used in the continuous display. The ball was programmed in MATLAB to appear in the bottom left corner of the computer screen, disappear for 1 s, reappear in the top right corner of the computer screen, disappear for 1 s, and then reappear in the bottom left corner of the computer screen.

(3) The continuous auditory effect was a tone created using the pitch change function in Audacity. The tone started at 225 Hz and increased smoothly in frequency to 900 Hz, and then decreased smoothly back to 225 Hz.

(4) The discrete auditory effect, also created using the pitch change function in Audacity, consisted of a low tone (225 Hz), immediately followed by a high tone (900 Hz), immediately followed by another low tone (225 Hz).

2.1.3. Procedure and coding

Children were tested individually in a quiet room. The experimenter began by introducing the hidden computer as "the big machine" and then showed children either Controller A or Controller B; the order was counterbalanced across participants. The other controller was out of sight. The experimenter then pointed to each part on the controller, showing how each could be manipulated, and said, "Each of these parts can make something different happen on the big machine." The child was encouraged to manipulate the parts. The controller was not "plugged in" during this phase, and no effects were displayed. Children were allowed to manipulate the parts as long as they wished. All children manipulated both parts. After children finished manipulating the parts, the experimenter turned the controller so that the parts no longer faced the child. Then the experimenter plugged the controller into the machine, and told the child, "Now I'm going to use one of the parts on the controller to make something happen on the big machine. Watch!" A felt cloth obscured both parts and the experimenter's hand so the child could not see which part she (supposedly) manipulated.

While children watched the visual effect, the experimenter described the event. For the continuous effect she said, "Do you see the ball? It's moving all the way along. I'm using one of the parts to make the ball move all the way along." For the discrete effect she said, "Do you see the ball? It's moving from the bottom to the top. I'm using one of the parts to make the ball move from the bottom to the top." The effect and the verbal description were repeated three times, and then the experimenter told the child, "I'm going to use the other part on the controller to make something different happen. Watch!" Then the other effect was displayed, accompanied by the appropriate verbal description. Finally, children were reminded that they had seen two different effects (e.g., "Remember, first we saw the ball move all the way along and then we saw the ball move from the bottom to the top."). The experimenter then unplugged the controller from the machine and oriented the controller so that the two parts were facing the child and asked the child to identify the part responsible for the second effect. The dependent measure was the part the child pointed to first.

Next children were introduced to the second controller to be used along with the auditory effect trials. The procedure was identical to that described above. Children were introduced to the other controller's two parts. They were told that when this controller was connected to the machine, each part could make the machine play a different kind of noise. The continuous auditory effect was described as, "Do you hear the noise? The noise is going higher and lower. I'm using one of the parts to make the noise go higher and lower." The discrete auditory effect was described as, "Do you hear the noise? It's going low, high, low. I'm using one of the parts to make the noise go low, high, low."

2.2. Results and discussion

Each child gave two independent responses: one about the cause of the second visual effect and one about the cause of the second auditory effect. Across trials, children chose the corresponding part – the discrete part for the discrete effect and the continuous part for continuous effect – 75% of time (24/32 trials), significantly more often than chance ($p = .007$ by binomial test). Nine of 16 children chose the corresponding part for both trials ($p = .01$ by binomial test). Children were equally good at identifying the corresponding cause of visual (12/16 children) and auditory (12/16) effects as well as continuous (12/16) and discrete (12/16) effects.

These data are consistent with the idea that children can represent abstract criteria for a solution to a problem and use these criteria to identify candidate causes with no additional data. Note that it is not simply the case that children inferred causal relationships when those relationships were in fact imaginary. Researchers have used machines that only appear to activate in response to stimuli for many studies of children's sensitivity to patterns of covariation for causal reasoning (Gopnik & Sobel, 2000; Buchsbaum et al., 2012). In such studies, however, the target inferences are formally rational, given prior knowledge and the pattern of data.

By contrast, in our experiment there is no fact of the matter. The parts on the controllers – pulleys, magnets, rotating beads and snap-lock beads – are equally plausible as physical mechanisms (as attested to by the fact that children chose each equally often across conditions), and no covariation data linked the parts and effects. Thus, there is no strict sense in which we can say that the children's judgments are "correct." What we can say is that children were able to generate an abstract

representation of the relationship between the effects and the causes to converge systematically on a candidate hypothesis.

3. Experiment 2

Although children in Experiment 1 represented abstract properties of effects to converge on candidate causes, the experimenter's descriptive language (e.g., noting that the sound went higher and lower vs. low, high, low) might have highlighted the relevant abstract dimension. To investigate whether children spontaneously generate abstract representations to constrain their evaluation of hypotheses, we omitted descriptive language in Experiment 2.

3.1. Method

3.1.1. Participants

Sixteen 4–6-year-olds (mean age 63 months; range 48–79 months) were recruited at a children's museum. Eight children were asked about the continuous visual effect and discrete auditory effect and eight children were asked about the discrete visual effect and continuous auditory effect. An additional two children were recruited but were not included in the study due to inability to complete the session (1) and experimenter error (1).

3.1.2. Materials and procedure

Materials and procedure were the same as those used in Experiment 1, with the exception of the verbal descriptions of the four effects. Instead of using meaningful words, visual effects were paired with the novel verbs *gazz* and *blick*, and the auditory effects were paired with the novel verbs *flurp* and *dax*. The experimenter said, for example, "Do you see the ball? It's *gazzing*[1]. I'm using one of the parts to make the ball *gazz*." The order of the novel verbs remained constant, while the order of the effects was counterbalanced across participants so that *gazz* described the continuous visual effect for half of the participants while *blick* described the same effect for the other half of participants. At the end of the first trial, the experimenter reminded the child about the two effects, saying, "Remember, first we saw the ball *gazz* and then we saw the ball *blick*." The experimenter oriented the controller so that the two parts were facing the child and asked, "Which part made the ball *blick*?" This process was repeated with the auditory effects and the second controller.

3.2. Results and discussion

Each child gave two independent responses: one to the controller for the visual effects and one to the controller for the auditory effects. As in Experiment 1, children chose the corresponding part significantly above chance: 78% of time (25/32 trials; $p = .002$ by binomial test). Ten children chose the corresponding part on both trials ($p = .002$ by binomial test). Children identified the corresponding parts for continuous (12/16) and discrete (13/16) effects equally well. There was no difference in children's performance between Experiments 1 and 2 (Fisher's Exact Test $p =$ ns). These results suggest that children are able to imagine abstract representations of problems and their solutions without relying on descriptive language that might support the distinction.

4. Experiment 3

If learners evaluate ideas not just based on their truth value, but based on their correspondence to the abstract form of a good solution, then they should be able to distinguish "good" and "bad" ideas, even when they have reason to believe both ideas are wrong. In Experiment 3, we provided children with an alternate explanation for the effects: an invisible magic wand. The children were introduced to two puppets, neither of whom knew about the wand. Each puppet made a different guess about which

---

[1] NB: where the continuous tense (e.g., *gazzing*) was used, it was used to describe both the continuous and discrete effects.

part on the controller produced the effect in question. (We chose an invisible cause as the putative cause to prevent children from choosing a part based on its degree of similarity to the stated cause.) If children believe that both puppets are wrong, but nonetheless converge on the one who guessed the corresponding part, this would suggest that children consider the extent to which a hypothesis fulfills the abstract form of a solution even when they believe the hypothesis is false.

4.1. Method

4.1.1. Participants

Thirty-two 4–6-year-olds (mean age 63 months; range 46–83 months) were recruited at a children's museum. Three additional children were recruited but not included in the study due to inability to complete the session (1) and experimenter error (2).

4.1.2. Materials and procedure

The stimuli were the same as in Experiments 1 and 2, but the design for this experiment was between-subjects; children saw either the pair of visual effects or the pair of auditory effects, with one of the two controllers. In addition, two girl or two boy puppets (matched to the gender of the child) were used. The procedure was identical to Experiment 2 except that, after the presentation of the pair of effects, the experimenter said, "Our friends [the puppets] John and Sam think that one part on the controller made the ball *gazz* (sound *flurp*) and the other part on the controller made the ball *blick* (sound *dax*). I'm putting our friends away for a minute because I have to tell you a secret. Do you know what really makes the machine work? A magic wand! Watch (listen)!" The experimenter ensured that the child could see that she was not touching the controller and then generated the second effect, purportedly by using a magic wand behind the machine. The experimenter reminded children that the puppets did not know about the wand and each believed that one of the parts on the controller produced the effect in question. The puppets returned and were asked what made the ball *blick* or the sound *dax*. One puppet guessed the discrete part and the other guessed the continuous part. The experimenter then said, "Even though we know they are both wrong, who had the better guess about which part made the ball *blick* (sound *dax*)? Was it John [who picked the discrete part] or was it Sam [who picked the continuous part]?" To confirm that children continued to believe the real mechanism was the magic wand, the experimenter concluded by asking, "Can you remind me what really makes the machine go?"

4.2. Results and discussion

The dependent measure in this experiment was which puppet the child chose as having the better guess. Because we were interested in children's ability to evaluate solutions they believed were wrong, we considered data only from 27 of the 32 children who answered that the wand was the real cause of the effects. Twenty of the twenty-seven children (74%) selected the puppet who chose the corresponding part: the discrete cause for the discrete effect and the continuous cause for the continuous effect ($p = .02$ by binomial test)[2]. This result suggests that children can evaluate an idea based on its fit to an abstract form of a solution, even when they believe the idea is wrong.

5. Experiment 4

The abstract relationship between the form of the problem and the form of the solution in Experiments 1–3 is a simple one, requiring only a mapping between the visual/auditory properties of the effect, and the motion/feel of the mechanical properties of the candidate causes. Typically, studies of cross-modal perception look at participants' ability to bind constitutive properties of a single entity—for example, the sight and sound, or sight and feel of a stimulus (Jordan & Brannon, 2006;

---

[2] This result remains significant if all the children are included. Twenty-two of the thirty-two children (69%) picked the puppet who chose the corresponding effect ($p = .05$ by binomial test).

Meltzoff & Borton, 1979). To our knowledge, studies of cross-modal mapping have not looked at causal relations as in the current study. In Experiments 1–3, however, we cannot be sure that the children actually represented the abstract relationship between the problem and the solution (as a continuous cause generating a continuous effect and a discrete cause generating a discrete effect), or whether they did not know how to answer the question and simply guessed, and then used cross-modal mapping, perhaps between the visual or auditory modality and the physical or kinesthetic modality to constrain their guesses.

   To test whether the form of the problem affects children's choice of solution and to ask whether children rely only on cross-modal mapping to identify causes, we modified the test question in Experiment 4. Instead of asking children which cause generated a particular effect, we asked children which part could change the continuous visual effect to a sound (an auditory effect) or which part could change the continuous auditory effect to a movie (a visual effect). If children rely on cross-modal mapping to constrain their guesses, they should be as likely to select the continuous part in response to this question as in response to the question in Experiment 2, given that the continuous display is present. However, although the change in modality can be represented as a continuous transformation (e.g., the visual stimulus fading away and the sound gradually emerging), the change in modality may be more readily represented as a discrete change. If children consider the form of the problem in considering the solution and not just the cross-modal match to the effect, children's pattern of responses should differ between Experiments 2 and 4.

5.1. Method

5.1.1. Participants

   Sixteen 4–6-year-olds (mean age 66 months; range 50–80 months) were recruited at a children's museum. An additional child was recruited but was not included in the study due to inability to complete the session.

5.1.2. Materials and procedure

   The procedure was identical to Experiment 2 with two exceptions. First, rather than asking which part could generate the effect, the experimenter asked children which part could change the modality of the effect. Second, the order of presentation was fixed so that all children were shown the discrete effect followed by the continuous effect and were asked about the continuous effect on both trials. After children saw the continuous visual effect, the experimenter said, "One of the parts on the controller can make the ball that *blicked* become a sound. Which part will make the ball that *blicked* become a sound?" After the continuous auditory effect was played, the experimenter said, "One of the parts on the controller can make the noise that *gazzed* become a movie. Which part will make the noise that *gazzed* become a movie?" Note that we asked children to turn a visual effect into an auditory one or an auditory effect into a visual one, and no mention was made about whether the resulting effect should be continuous or discrete.

5.2. Results and discussion

   Each child gave two independent responses: one to the controller for changing the continuous visual effect to a sound and one to the controller for changing the continuous auditory effect to a movie. We looked at whether the type of problem children were asked to solve influenced children's solutions by comparing the number of trials on which children chose the cross-modal match in Experiment 4 (34%, 11/32 trials) with the number of trials on which they chose the corresponding part in Experiment 2 (78%, 25/32 trials). As predicted, children chose the cross-modal match significantly less in Experiment 4 than in Experiment 2 ($p < .001$ by Fisher's Exact Test). In Experiment 4 only one of the 16 children chose the cross-modal match on both trials. Indeed, within Experiment 4, children were marginally more likely to choose the discrete part than the continuous part (64%, 21/32 trials, $p = .05$ by binomial test, one-tailed). The results of Experiment 4 rule out the possibility that children simply use cross-modal mapping to guess the answer when they have no other grounds for

distinguishing hypotheses. Rather, children seem to search for solutions that might conform to the form of the problem.

6. General discussion

Across four experiments, children were able to use an abstract representation of the relationship between a problem and a solution to converge on candidate causes. In these experiments, children viewed two effects in the visual or auditory modality, one continuous and one discrete. Children performed systematically by choosing the corresponding part for a particular effect both in the presence (Experiment 1) and absence (Experiment 2) of verbal descriptions highlighting the target relationship, and favored the target hypotheses even when they believed the hypotheses under consideration were false (Experiment 3). Moreover, children did not simply look for salient relationships between causes and effects; they took the form of the problem into account. Children distinguished between problems that required them to infer the causes of discrete versus continuous effects from problems requiring them to convert effects from one modality to another (Experiment 4).

In the current studies, the possible representations of both the problem and its solution were tightly constrained, given the contrastive cues and the forced choice between candidate causes. Nonetheless, it is noteworthy that children performed systematically on these tasks despite the absence of any of the information generally posited to help support children's causal inferences: all the candidate causes were equally plausible given prior knowledge about physical mechanisms, and none of the candidate causes was supported by statistical data (Gopnik & Wellman, 2012; Schulz, 2012; Tenenabum, Kemp, Griffiths, & Goodman, 2011). These results suggest that children can generate abstract representations of what might count as a good solution to a problem and use these representations to constrain the hypotheses they consider.

This is a preliminary investigation, and the current data do not support, or even test, all the claims we have advanced here; we will be satisfied if the current work simply provides grounds for believing these ideas are worthy of further research. In particular, our experiments assess only a pre-requisite to goal-constrained hypothesis generation: the ability to distinguish competing hypotheses on the basis of an abstract fit between the form of a problem and a solution. Future work might investigate the extent to which this ability constrains the process by which children generate hypotheses.

Note that in Experiments 1–3, children's choices required them to represent the relational similarity between discrete effects and causes as compared to continuous effects and causes. In this respect, the studies resemble studies of children's analogical reasoning (Christie & Gentner, 2010; Gentner & Markman, 1997). We believe that cases of analogical reasoning are elegant examples of our more general ability to represent abstract criteria for the solution to a problem and use these criteria to constrain the hypotheses we generate. However, here we did not give children base problems and solutions, nor did we provide children with targets to map the base problems onto. Children could not use a process of structural alignment to set up a relational mapping between arguments (Christie & Gentner, 2010). That is, they could not go from a known problem and a known solution to a new problem and a new solution. Rather, children had to infer a representation that might relate the form of a problem to the form of a solution and use this representation to guide their responses. We stress this point not to discount the importance of analogical reasoning, but because we believe a general ability to represent the abstract form of a problem and a solution constrains hypothesis generation even in cases where analogical reasoning is clearly not in play (e.g., in generating new names for theater companies or processes that might put stripes on peppermints).

Although we regret that neither giants nor rings were involved in these studies, we believe this line of investigation may help us better understand the role of imagination in cognition. The establishment of goals and criteria for their fulfillment is arguably common to every act of imagination: the pretend play of preschoolers, the daydreams of adults, the fictions of novelists, and even the confabulations of neuropsychiatric patients. To the degree that the ability to represent the abstract form of a problem and a solution is a constraint on hypothesis generation more broadly, imagination does not just support thinking; it is "just thinking."

Acknowledgements

References

Boyd, B. (2009). *On the origin of stories: Evolution, cognition and fiction*. Cambridge: Harvard University Press.

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, Biological Sciences, 367*, 2202–2212.

Carey, S. (2009). *The origin of concepts*. New York, NY: Oxford University Press.

Carruthers, P. (2009). How we know our own minds: The relationship between mindreading and metacognition. *Behavioral and Brain Sciences, 32*, 121–138.

Christie, S., & Gentner, D. (2010). Where hypotheses come from: Learning new relations by structural alignment. *Journal of Cognition and Development, 11*, 356–373.

Farver, J. A. M., & Shin, Y. L. (1997). Social pretend play in Korean- and Anglo-American preschoolers. *Child Development, 68*, 544–556.

Farver, J. A., & Wimbarti, S. (1995). Indonesian children's play with their mothers and older siblings. *Child Development, 66*, 1493–1503.

Fein, G. G. (1981). Pretend play in childhood: An integrative review. *Child Development, 52*, 1095–1118.

Fodor, J. (1998). *In critical condition: Polemical essays on cognitive science and the philosophy of mind*. Cambridge: MIT Press.

Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist, 52*, 45–56.

Goodman, N. D., Ullman, T. D., & Tenenbaum, J. B. (2011). Learning a theory of causality. *Psychological Review, 118*, 110–119.

Gopnik, A., & Sobel, D. M. (2000). Detecting blickets: How young children use information about novel causal powers in categorization and induction. *Child Development, 71*, 1205–1222.

Gopnik, A., & Wellman, H. M. (2012). Reconstructing constructivism: Causal models Bayesian learning mechanisms, and the theory theory. *Psychological Bulletin, 138*, 1085–1108.

Gosso, Y., Morais, M. L. S., & Otta, E. (2007). Pretend play of Brazilian children: A window into different cultural worlds. *Journal of Cross-Cultural Psychology, 38*, 539–558.

Haight, W. L., Parke, R. D., & Black, J. E. (1997). Mothers' and fathers' beliefs about and spontaneous participation in their toddlers' pretend play. *Merrill-Palmer Quarterly, 43*, 271–290.

Haight, W. L., Wang, X., Fung, H. H., Williams, K., & Mintz, J. (1999). Universal, developmental, and variable aspects of young children's play: A cross-cultural comparison of pretending at home. *Child Development, 70*, 1477–1488.

Harris, P. L. (1998). Fictional absorption: Emotional responses to make believe. In S Braten (Ed.), *Intersubjective communication and emotion in early ontogeny* (pp. 336–353). Paris: Cambridge University Press.

Jordan, K. E., & Brannon, E. M. (2006). The multisensory representation of number in infancy. *Proceedings of the National Academy of the Sciences, 103*, 3486–3489.

Kopelman, M. D. (1987). Two types of confabulation. *Journal of Neurology, Neurosurgery and Psychiatry, 50*, 1482–1487.

Lillard, A. S., Lerner, M. D., Hopkins, E. J., Dore, R. A., Smith, E. D., & Palmquist, C. M. (2013). The impact of pretend play on children's development: A review of the evidence. *Psychological Bulletin, 139*, 1–34.

Meltzoff, A. N., & Borton, R. W. (1979). Intermodal matching by human neonates. *Nature, 282*, 403–404.

Nichols, S., & Stich, S. (2000). A cognitive theory of pretense. *Cognition, 27*, 115–147.

Oatley, K. (1999). Why fiction may be twice as true as fact: Fiction as cognitive and emotional simulation. *Review of General Psychology, 3*, 101–117.

Phelps, E. A., & Gazzaniga, M. S. (1992). Hemispheric differences in mnemonic processing: The effects of left hemisphere interpretation. *Neuropsychologia, 30*, 293–297.

Pinker, S. (1997). *How the mind works*. New York, NY: W.W. Norton & Company.

Schulz, L. (2012). The origins of inquiry: Inductive inference and exploration in early childhood. *Trends in Cognitive Sciences, 16*, 382–389.

Singer, D. G., & Singer, J. L. (1992). *The house of make-believe: Children's play and the developing imagination*. Cambridge: Harvard University Press.

Tenenabum, J. B., Kemp, C., Griffiths, T., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science, 331*, 1279–1285.

Ullman, T. D., Goodman, N. D., & Tenenbaum, J. B. (2012). Theory learning as stochastic search in the language of thought. *Cognitive Development, 27*, 455–480.

Vrana, S. R., & Lang, P. J. (1990). Fear imagery and the startle-probe reflex. *Journal of Abnormal Psychology, 99*, 189–197.

Walker, C., & Gopnik, A. (2013). Causality and imagination. In M. Taylor (Ed.), *The Oxford handbook of the development of the imagination* (pp. 342–358). New York, NY: Oxford University Press.