

Children’s Causal Interventions Combine Discrimination and Confirmation

Yuan Meng

yuan_meng@berkeley.edu
Department of Psychology
University of California, Berkeley

Neil R. Bramley

neil.bramley@nyu.edu
Department of Psychology
New York University

Fei Xu

fei_xu@berkeley.edu
Department of Psychology
University of California, Berkeley

Abstract

Like scientists, children have a sharp sense of when and how to seek evidence, but when it comes to generating causal interventions, their performance often falls short of normative information-theoretic metrics such as the expected information gain (EIG). We looked at whether this deviation resulted from mixing discriminatory strategies such as maximizing EIG with confirmatory strategies such as the positive test strategy (PTS). Thirty-nine 5- to 7-year-olds solved 6 puzzles where they had one opportunity to intervene on a three-node causal system to identify the correct structure from two possibilities. Children’s intervention choices were better fit by a Bayesian model that incorporated EIG and PTS compared to alternative models that only considered a single strategy or selected interventions at random. Our findings suggest that children’s intervention strategy may be a combination of discrimination and confirmation.

Keywords: causal learning; interventions; self-directed learning; Bayesian modeling

Introduction

[I]t is the usual fate of mankind to get things done in some boggling way first, and find out afterward how they could have been done much more easily and perfectly.

— Charles S. Peirce (1882)

Fairy tales often depict swallows as bringers of spring. Given the frequent co-occurrence of the two, no wonder our ancestors suspected a causal link between them. Setting free a flight of swallows, however, is unlikely to end winter—by fixing the value of the variable “whether swallows are present” to “yes” ($do(\text{Swallow} = 1)$), a method known as *interventions*, we soon learn that swallows do not bring forth spring; perhaps it is the warmth of spring that lures them back.

Interventions are a powerful tool for uncovering causal structures, but not all are equally useful for distinguishing among a set of alternatives. Of numerous models that quantify the usefulness of interventions (Nelson, 2005), *information gain* (IG) is currently most widely used (e.g., Bramley, Lagnado, & Speekenbrink, 2014; Oaksford & Chater, 1994; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003). According to the IG model, good interventions make learners feel less *uncertain* (uncertainty is usually measured by “average surprise”, or Shannon entropy, Shannon, 1948) about a causal system: initially, one causal structure may seem just as likely as the next, so guessing which one is correct is like a shot in the dark; after an informative intervention, however, we can use the resulting effect to sift the good from the bad.

The ability to use interventions to learn causal structures emerges early (e.g., Cook, Goodman, & Schulz, 2011; Schulz, Gopnik, & Glymour, 2007) but is still limited in childhood. For instance, when asked to find out whether two gears were independent or connected, only half of the

4- and 5-year-olds generated the most useful evidence on their own (Schulz et al., 2007, Experiment 3). Compared to pairwise relationships between two variables, it is less well-understood how children learn global structures of multiple variables through interventions. In a recent study, McCormack, Bramley, Frosch, Patrick, and Lagnado (2016) looked at whether 5- to 8-year-olds were able to use interventions to learn the structure of a three-node system. Each node A, B, C represented a shape sticking out of a box lid, which could be rotated by hand or by the other shapes via a hidden mechanism inside the box. For a given causal system, children were shown three ways in which the mechanism in the box might work (a common cause: $B \leftarrow A \rightarrow B$; two causal chains: $A \rightarrow B \rightarrow C$ or $A \rightarrow C \rightarrow B$) and were allowed at least 12 opportunities to intervene on this system to figure out which one was the case. Chance levels of IG were established through simulation of randomly selected interventions. Only 7- to 8-year-olds’ intervention quality was above chance for both types of structures. By contrast, 5- to 6-year-olds’ intervention quality was below chance for both; 6- to 7-year-olds’ was above chance for causal chains but not for the common cause.

Consistent with “source preferences” found by Steyvers et al. (2003), McCormack et al. (2016) noted that a considerable proportion of children intervened on the *root node* A, which could not reduce uncertainty since all shapes would rotate regardless of the true structure. Coenen, Rehder, and Gureckis (2015) suggested that this phenomenon may be driven by the *positive test strategy* (PTS): instead of differentiating between alternatives, a PTS user focuses on one hypothesis (e.g., $A \rightarrow B \rightarrow C$) at a time while ignoring everything else. To check whether both links in the working hypothesis exist, the most efficient intervention is to activate A, which examines both links at the same time. More generally, a PTS user favors nodes that can simultaneously test the *largest proportion* of links. This strategy runs parallel to PTS in rule learning where learners favor queries that produce affirmative responses (e.g., “yes”) if the current hypothesis is true (Klayman & Ha, 1989; Wason, 1960). PTS is most useful when evidence can *falsify* the current hypothesis, whereas positive evidence could support just as many hypotheses as before the intervention (although you may have more confidence in the supported ones). PTS is efficient if causal connections are *sparse* (Navarro & Perfors, 2011; Oaksford & Chater, 1994). For instance, to see if pollen X causes allergy Y, you can run tests on pollen X. In theory, you can also check whether people without allergy Y have been exposed to pollen X—however, since allergy Y is rare ($P(Y) < .5$), you need to check frequently; since pollen X is rare ($P(X) < .5$), you may

well find nothing. While causal connections in McCormack et al.’s (2016) study were not sparse, they often are in real life; as a result, children may still use PTS even when it works less well than discriminatory strategies such as IG maximization.

Combining strategies

Coenen et al. (2015) found that when choosing interventions, adults might combine discriminatory (IG) and confirmatory (PTS) strategies. In a series of experiments, they showed participants two possible structures of a three- or four-node causal system and asked them to identify the true structure by using as few interventions as possible. Participants’ (first) intervention choices were better fit by a Bayesian model that incorporated IG and PTS when assigning values to potential interventions compared to models that only considered either IG or PTS, or assigned values indiscriminately.

The new modeling approach above provides a more precise characterization of children’s strategy use, which could potentially explain why they often fail to generate informative interventions. By contrast, other past studies that we know of (e.g., Bramley et al., 2014; McCormack et al., 2016; Steyvers et al., 2003) only quantified the degree to which participants’ performance deviated from the IG model but did not zero in on where such deviation came from. As with Coenen et al. (2015), we addressed this issue by entertaining the possibility that children combine IG and PTS to select interventions. We focused on 5- to 7-year-olds because they did not reliably use IG to select interventions in McCormack et al.’s (2016) study. We adapted Coenen et al.’s (2015) adult task to make it more accessible and engaging for children of this age range. In one experiment, children solved puzzles where they had one opportunity to intervene on a three-node causal system to identify the correct structure among two possibilities. Intervention choices were fit by two single-strategy models (PTS or IG), a random model (all interventions were assigned a value of 1), and a combined model (both PTS and IG).

Modeling strategies

Values of interventions Different strategies assign values to interventions differently. Here we included the expected information gain (EIG), the positive test strategy (PTS), the random strategy, and the combined strategy (EIG and PTS).

1. **Expected Information Gain (EIG)** Learners begin with a set of hypotheses about how a causal system may work. Each hypothesis is formally represented as a *causal Bayesian network*—a directed acyclic graph where nodes represent causal variables and links represent causal relations (Pearl, 2000). Learners’ uncertainty about this causal system $g \in G$ (G is the space of possible graphs and g is each hypothesis) is measured by its Shannon entropy $H(G)$:

$$H(G) = - \sum_{g \in G} P(g) \log_2 P(g). \quad (1)$$

Observing the outcome $o \in O$ of an intervention on a node

$n \in N$ reduces the entropy to $H(G|n, o)$. The reduction of entropy is this intervention’s information gain $IG(n, o)$:

$$IG(n, o) = H(G) - H(G|n, o). \quad (2)$$

The IG model assumes that learners’ goal is to maximize $IG(n, o)$. However, because an intervention’s outcome is unknown beforehand, its *expected information gain (EIG)* (marginalized over all possible outcomes) is used instead:

$$EIG(n) = H(G) - \sum_{o \in O} P(o|n) H(G|n, o). \quad (3)$$

In Eq. (3), the conditional entropy $H(G|n, o)$ is:

$$H(G|n, o) = - \sum_{g \in G} P(g|n, o) \log_2 P(g|n, o). \quad (4)$$

In Eq. (4), for each graph, the prior probability $P(g)$ is the same and the posterior probability $P(g|n, o)$ is given by Bayes’ rule, $\frac{P(o|g, n)P(g)}{\sum P(o|g, n)P(g)}$. $P(o|g, n)$ is an outcome’s likelihood given a hypothesis and an intervention.

2. **Positive Test Strategy (PTS)**

According to the PTS model, learners aim to intervene on the node $n \in N$ that has the largest proportion of links¹ (direct or indirect), with each graph $g \in G$ being considered:

$$PTS(n) = \max_g \left[\frac{\text{DescendantLinks}_{n,g}}{\text{TotalLinks}_g} \right]. \quad (5)$$

3. **Random strategy**

According to the random model, the learner assigns the same value (e.g., 1) to all possible interventions.

4. **Combined strategy**

According to the combined model, the learner assigns a weighted mean of EIG (weight: θ) and PTS (weight: $1 - \theta$) scores to each intervention.

Linking values to interventions Ideal learners always intervene on the node with the highest value $V(n)$; actual learners do so probabilistically—their behavior can be captured by the *softmax choice rule* (Luce, 1959):

$$P(n) = \frac{\exp(V(n)/\tau)}{\sum_{n \in N} \exp(V(n)/\tau)}, \quad (6)$$

where τ is learners’ decision noise: when τ is 0, they behave ideally; when τ approaches $+\infty$, they choose randomly.

The main interests of our study were to 1) examine which model could best capture 5- to 7-year-olds’ intervention selection and to 2) look for possible developmental changes.

¹We tested two alternative PTS models where values are based on the *expectation* or the *total number* of how many nodes can be turned on. Neither model fit our data as well as the one we used.

Adaptations for the current study

We explored scenarios similar to those explored by Coenen et al. (2015) with two key differences.

We set causal connections to be *deterministic*: intervening on a parent node would always affect its children and no background causes existed. This was because children might have difficulties understanding that a parent only had a probability of (for instance) 0.8 to activate its children; incorporating this information into causal reasoning could make it even harder.

We only selected three unique problems² ($A \rightarrow B \rightarrow C$ vs. $B \leftarrow A \rightarrow C$; $A \rightarrow B \rightarrow C$ vs. $B \rightarrow C$; $B \rightarrow A \leftarrow C$ vs. $C \rightarrow A$). To obtain more data, each child faced each problem twice with the roles of A and B counterbalanced ($A \rightarrow C \rightarrow B$ vs. $B \leftarrow A \rightarrow C$; $A \rightarrow C \rightarrow B$ vs. $C \rightarrow B$; $B \rightarrow A \leftarrow C$ vs. $B \rightarrow A$). There was only one informative intervention in each problem, making it easier for us to distinguish children’s performance from the chance level of EIG based on relatively sparse data (there were only a total of 6 problems and children only intervened once to solve each problem). Moreover, EIG and PTS happened to have the least overlap³ in these problems, so our models were better able to capture children’s strategy use.

Experiment

Methods

Participants Thirty-nine 5- to 7-year-olds ($M_{\text{age}} = 79.8$ months, range = 61.8–94.8 months, $SD = 9.7$ months; 18 females) were tested in a laboratory located at University of California, Berkeley or in a local children’s museum.

Equipment and materials An interactive game was developed for this study (see Figure 1). Simple causal systems consisting of a yellow, a green, and a red light bulb were programmed in Scratch (see all causal systems: <https://scratch.mit.edu/users/BayesianBabies/projects/>) and presented on a laptop screen. A response board with three buttons of corresponding colors was connected to the computer via a circuit board to control the light bulbs. In the test phase, possible structures of light bulbs were shown on cards.

Procedure Children first learned how to turn on light bulbs using buttons on the response board. For each light bulb, the experimenter turned it on and off and asked children to turn it back on. In the end, the experimenter summarized the rule, “You just need to click the button that has the same color!”

Next, children practiced describing four basic types of causal structures (common cause: $\text{yellow} \leftarrow \text{green} \rightarrow \text{red}$; common effect: $\text{yellow} \rightarrow \text{red} \leftarrow \text{green}$; causal chain: $\text{green} \rightarrow \text{red} \rightarrow \text{yellow}$; one link: $\text{yellow} \rightarrow \text{red}$) and observed outcomes of all three interventions in

²For three-node causal systems with one or two links, there are 18 unique structures: 3 common-cause structures, 3 common-effect structures, 6 causal-chain structures, and 6 one-link structures, resulting in a total of $\binom{18}{2} = 153$ pairs of graphs. Because the node names are arbitrary, simultaneously swapping nodes in both structures of a pair (e.g., swapping A and C in $A \rightarrow B \rightarrow C$ vs. $A \rightarrow B$) leads to an equivalent pair ($C \rightarrow B \rightarrow A$ vs. $C \rightarrow B$). There are 27 unique pairs left after eliminating the redundant.

³Overlap could not be eliminated in any given pair.




Figure 1: The experiment setup.




Figure 2: 6 puzzles used in the experiment.

each structure. Specific structures used as examples during practice were not included in the test. The trial order was randomized for each child. All three light bulbs were presented simultaneously, along with red arrows denoting causal relationships. On the first trial, the experimenter told children, for instance, “These light bulbs (e.g., $\text{yellow} \rightarrow \text{red}$) have a secret: some give light to others. See this arrow over here? It means that the yellow light bulb gives light to the red light bulb. That is, when yellow turns on, it turns on red, too!” To ensure that children understood causal structures shown in the pictures, they were asked to describe them in their own words on the last three trials. The experimenter asked additional questions if children 1) ignored the red arrows (e.g., Child, “This picture tells us there are three light bulbs.” Experimenter, “Great! Can you tell me what this arrow tells us?”), 2) did not address causal relationships (e.g., Child, “This picture shows us an arrow pointing from yellow to red.” Experimenter, “Good! What does it tell us?”), or 3) made factual errors (e.g., Child, “This picture shows that red turns on yellow.” Experimenter, “Maybe it shows us yellow turns on red? What do you think?”). After describing a picture, children were invited to turn on each light bulb, “Now you can turn on the light bulbs one by one and see what happens!”

In the test, children solved six puzzles (see Figure 2) in a random order. On each test trial, arrows were hidden away and cards with two candidate causal structures were displayed side by side; the relative positions were randomized. Children were told, “These three light bulbs work in a special way. They either work this way (pointing to one card) or this way

(pointing to the other). One picture is correct about how they work and the other one is wrong.” They were asked to find out the true structure based on one intervention, “To solve this puzzle, you can turn on one light bulb and see what happens. From that, you can find out which picture shows us how these three light bulbs really work! Which light bulb do you want to turn on to help you?” After each intervention, children identified the correct structure by putting a smiley face sticker on it. Correct structures were randomly selected and feedback was only provided in the end (to avoid discouragement).

Results

Accuracy We first looked at whether children could identify correct causal structures through interventions. The percentage of correct choices averaged across all children and puzzles⁴ was 54% ($SD = .22$, $MD = 50\%$), which was at chance (50%), $t(38) = 1.02$, $p = .31$, Cohen’s $d = .16$. There was no age difference, $F(1, 37) = .027$, $p = .87$, $\bar{R}^2 = .00$.

Raw scores Before model fitting, we looked at the raw EIG and PTS scores of children’s chosen interventions. The average EIG over all interventions was .39, which was not distinguishable from the chance level⁵ of EIG (.33), $t(38) = 1.23$, $p = .23$, Cohen’s $d = .20$. The average PTS over all interventions was .74, which was above the chance level⁶ of PTS (.55), $t(38) = 5.21$, $p < .001$, Cohen’s $d = .83$. Both the average EIG and the average PTS tended to increase with age, $F(1, 37) = 3.63$, $p = .065$, $\bar{R}^2 = .065$ and $F(1, 37) = 4.13$, $p = .049$, $\bar{R}^2 = .076$, respectively. Both effect sizes were small.

We categorized children by their main strategy and whether they switched strategies⁷ (see Table 1). More children were PTS users ($N = 31$) than EIG users ($N = 6$); two were undecided. Over half (18/31) of the PTS users switched strategies at least once while no EIG users switched from or to PTS.

Comparing models Of central interest to our study was the comparison between four models of children’s strategy use.

We took a hierarchical Bayesian approach to modeling children’s intervention choices. Single-strategy models (Figure 3a) assumed that children noisily maximized EIG or PTS, or assigned a value of 1 to all interventions. A free parameter τ_i captured each child’s decision noise, which was sampled from a population-level gamma distribution with two hyper-parameters α (shape) and β (rate). The combined model (Figure 3b) viewed children’s intervention strategy as a potential mix of of EIG maximization and PTS maximization. An additional free parameter θ_i captured EIG’s weight in each child’s evaluation of interventions, which was sampled from a

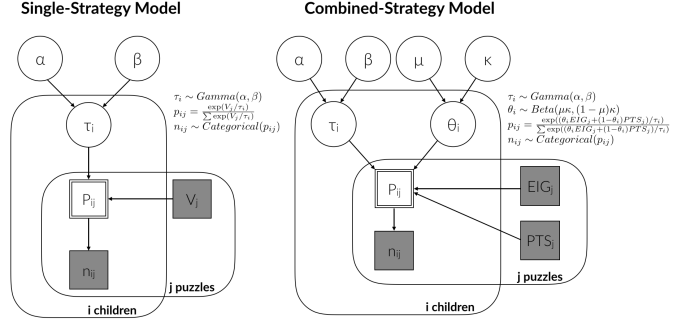


Figure 3: Hierarchical Bayesian models of single (left) and combined (right) strategies. In each puzzle j , each child i chose one node n_{ij} to intervene on. V_j , EIG_j , and PTS_j store the values of three possible interventions in each puzzle. p_{ij} stores probabilities of each child choosing each intervention in each puzzle. τ_i and θ_i capture each child’s decision noise and EIG weight. α and β are population-level hyper-parameters that generate τ_i ; μ and κ generate θ_i .

Table 1: Summary of strategy use and strategy switch.

Strategy	Combined	Single
EIG	0	6
PTS	18	13
Undecided	2	0

Table 2: Comparing four models of intervention strategies.

Model	DIC	τ	θ
Random	523.36	6.02	–
EIG	504.14	6.64	–
PTS	465.61	5.50	–
Combined	431.76	5.11	.24




Figure 4: (A) The best-fitting θ (MLE) for each child. (B) The distribution of the hyper-parameter μ (the mean of θ).

⁴Due to an equipment error, one child’s accuracy on one puzzle was not recorded; this trial was excluded from subsequent analyses.

⁵As mentioned, only one in three interventions was informative.

⁶This was the average PTS over all interventions in all puzzles.

⁷We did so by subtracting each child’s PTS score in each puzzle from their EIG score. If the sum of differences was positive or negative, a child was categorized as an EIG or a PTS user; if the sum was 0, a child was seen as undecided between strategies. If all non-zero differences had the same sign (+ or -), a child was categorized as single-strategy user and otherwise a combined-strategy user.

population-level beta distribution with two hyper-parameters μ (mean) and κ (standard deviation). Uninformative priors were chosen for all hyper-parameters: $\alpha = .001$, $\beta = .001$, $\mu \sim \text{Beta}(.5, .5)$, $\kappa \sim \text{Gamma}(.001, .001)$. Probabilities of possible interventions were determined by values of nodes as well as τ in single-strategy models and θ and τ in the combined model. Finally, interventions were sampled from a categori-

cal distribution of these probabilities.

We estimated parameter values from children’s intervention choices using Markov chain Monte Carlo (MCMC) samples generated by the JAGS program⁸ (Plummer, 2003) and used the deviance information criterion (DIC, Spiegelhalter et al., 2002) for model comparison. Models that can better fit data (smaller posterior mean of the deviance \bar{D}) or are simpler (smaller effective number of parameters p_D) have lower DIC ($= \bar{D} + p_D$). A difference over 10 is considered substantial.

As seen in Table 2, PTS was the best-fitting single-strategy model but the combined model outperformed all single-strategy ones. Across 39 children, the average EIG weight θ was .24 and the average noise τ was 5.11. Each child’s best-fitting θ (Figure 4a) as uncorrelated with τ , $r(37) = -.073$, $p = .66$. Figure 4b shows the distribution of the population-level hyper-parameter μ that captured the mean of θ .

Strategy and accuracy The EIG weight of each child had no effect on their overall accuracy, $F(1, 37) = 1.14$, $p = .29$, $\bar{R}^2 = .0038$. Getting down to the puzzle level, neither the EIG score nor the PTS score of the chosen intervention predicted whether the correct structure was subsequently identified, $\chi^2(1) = .23$ and $\chi^2(1) = .81$, respectively.

Developmental changes The EIG weight did not change with age, $F(1, 37) = .53$, $p = .47$, $\bar{R}^2 = .00$, but the noise decreased with age, $F(1, 37) = 10.59$, $p = .0024$, $\bar{R}^2 = .20$.

Discussion

Our key finding is that 5- to 7-year-olds’ intervention strategy is better described as a combination of discrimination and confirmation than a single strategy or random behavior.

On aggregate, children’s intervention choices lie between nodes that maximize EIG and nodes that maximize PTS. This mixture leans towards PTS: the EIG weight (.24) is much lower than that of PTS (.76). Categorization based on raw scores also suggests that more children were PTS users than EIG users. The EIG weight remains the same from 5 to 7, but the noise decreases with age, suggesting that older children evaluated nodes similarly as younger children but were better at maximizing values. Since only EIG is a reliable guide to informative interventions, one may guess that PTS is random interventions in disguise. However, this is unlikely because the EIG weight is uncorrelated with noise. Rather, PTS may be a genuine strategy and children’s heavy reliance on it may partially explain their difficulties generating informative interventions (McCormack et al., 2016; Schulz et al., 2007).

When it comes to identifying the correct structure, children performed only at chance. Surprisingly, higher reliance on EIG does not predict better structure learning. There may

be a discrepancy between the ability to generate interventions and the ability to learn from them. Similar results were found in several past studies (e.g., McCormack et al., 2015; Schulz et al., 2007). Given the sequential nature of interventions in these studies, by the time of inferring the causal structure, children might have forgotten the evidence. This explanation is unlikely here since children only intervened once and the outcome was present during the inference. However, under the current “one-shot learning” situation, children faced a new challenge: they had much less to draw on compared to at least 12 interventions in McCormack et al. (2016) and 5 minutes’ free play in Schulz et al. (2007). In addition, we noticed that children picked an answer very quickly after each intervention. Perhaps given so little evidence and time, it is even challenging to learn from informative interventions.

General Discussion

Like scientists, children have a sharp sense of when and how to seek evidence (Cook et al., 2011; Schulz & Bonawitz, 2007), but when it comes to actual data generation such as causal interventions, their performance often falls short of normative information-theoretic metrics like the expected information gain (EIG) (McCormack et al., 2016; Schulz et al., 2007). A growing body of work shows that adults’ (Bramley et al., 2014) and children’s (McCormack et al., 2016) intervention selection is neither random nor optimal. However, these studies did not explain why and how that might be the case. Inspired by Coenen et al. (2015), we explored the possibility that children’s suboptimal intervention selection results from combining discriminatory strategies with confirmatory strategies like the positive test strategy (PTS). In our study, 39 5- to -7-year-olds solved 6 puzzles where they intervened on a three-node causal system once in order to identify the correct structure from two alternatives. Like adults, children’s intervention choices were better fit by a Bayesian model incorporating EIG and PTS than models only considering one strategy or choosing interventions randomly. In this combined model, PTS was the main strategy that children relied on.

Granted, little do we know about how exactly children combine different strategies. For instance, do they *integrate* EIG and PTS to solve each puzzle or *switch* between them from one puzzle to another? Our results hint at the latter—PTS users often switched strategies but not EIG users, so perhaps children who first used PTS found their strategy ineffective and tried to switch, whereas those who first used EIG had no such need. If children integrate strategies, do they start with one and then consider another, or simultaneously compute and weight both? Future work is needed to answer these mechanistic questions. However, before digging into mechanisms of intervention selection, we should “guarantee some overall correctness or well-formedness of the computation” (Anderson, 1990). Most past studies only compared human behavior against an optimal benchmark; this combined model may provide a more realistic and nuanced starting point.

⁸We ran MCMC for 100,000 iterations, discarding the first 1000 samples and drawing a sample every 10 iterations. To ensure that samples came from a stationary distribution, we repeated this process 30 times with different initial parameter values and checked if the results from each sequence of samples, or *chain*, converged. Gelman and Rubin’s diagnostic \hat{R} (Gelman & Rubin, 1992) of all parameters was smaller than 1.05, indicating successful convergence.

Relation to past studies

Our findings seem at odds with myriad studies showing that children are good causal learners (see Gopnik & Wellman, 2012, for a review). However, our study differs in several key aspects. First, children choose their own interventions rather than observing those generated by others. Second, children learn the causal structure of multiple variables rather than whether a certain variable has causal power. Lastly, children are asked to identify global structures rather than being tested on pairwise relations. To our knowledge, only one child study (McCormack et al., 2016) has all three features and in that study, 5- to 7-year-olds also did not reliably select informative interventions. Complex causal structures seem to pose a real challenge for young learners, which can be a valuable opportunity to study how adult-like causal learning develops.

Future directions

The current study opens up many future directions to explore.

For developmental psychologists, a crucial question is how this combined strategy develops. Although the EIG weight remains the same from 5 to 7 in our study, people's intervention strategy may change throughout lifetime. Do children initially rely on one strategy and gradually incorporate others? What do the starting and the end points look like? What experiences may contribute to potential changes? To answer these questions, we need to adapt our task and test a much wider age range, such as from toddlers to adolescents.

We discussed one reason why children use seemingly sub-optimal PTS—if you focus on one hypothesis at a time, you should choose the node with the largest proportion of links because it can test your hypothesis most efficiently. However, it is also possible that children simply enjoy turning on as many light bulbs as possible. The current study cannot rule out this motivation. In the future, we can set the default state of light bulbs to “off”: if children prefer root nodes *much less*, it could suggest that it is positive effects that they chase after.

Since causality is central to scientific and everyday thinking, we wish to help children select more informative interventions. Asking them to explain their intervention choices might be one such way. Since explaining *why* can promote comparison and abstraction (see Lombrozo, 2016, for a review), explainers may be better able to *compare* outcomes under different structures and after different interventions and *abstract* away from solving specific puzzles that EIG often works better. It may also be helpful to provide feedback after each test trial and allow children to intervene more than once.

Causal interventions are among many ways to gather information, so are question asking (e.g., Nelson, 2005), exploration (e.g., Cook et al., 2011), hypothesis testing (e.g., Oaksford & Chater, 1994; Wason, 1960), etc.. In these domains, inquiries generated by adults and children are also often better than random but worse than optimal. We can take a similar modeling approach, looking at whether some version of combined strategies may play a role across a wide range of tasks.

Acknowledgment

We thank the research assistants in the Berkeley Early Learning Lab. This research was supported by Berkeley Fellowship for Graduate Study to Y. M. and an NSF grant to F. X. (no. 1640816).

References

- Anderson, J. R. (1990). Introduction. In *The adaptive character of thought* (pp. 1–40). Hillsdale, NJ: Erlbaum.
- Bramley, N. R., Lagnado, D. A., & Speekenbrink, M. (2014). Conservative forgetful scholars: How people learn causal structure through sequences of interventions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*(3), 708–731.
- Coenen, A., Rehder, B., & Gureckis, T. M. (2015). Strategies to intervene on causal systems are adaptively selected. *Cognitive Psychology*, *79*, 102–133.
- Cook, C., Goodman, N. D., & Schulz, L. E. (2011). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*, *120*(3), 341–349.
- Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, *7*(4), 457–472.
- Gopnik, A., & Wellman, H. M. (2012). Reconstructing constructivism: Causal models, bayesian learning mechanisms, and the theory theory. *Psychological Bulletin*, *138*(6), 1085–1108.
- Klayman, J., & Ha, Y.-W. (1989). Hypothesis testing in rule discovery: Strategy, structure, and content. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(4), 596–604.
- Lombrozo, T. (2016). Explanatory preferences shape learning and inference. *Trends in Cognitive Sciences*, *20*(10), 748–759.
- Luce, R. D. (1959). *Individual choice behavior*. New York, NY: John Wiley & Sons, Inc.
- McCormack, T., Bramley, N. R., Frosch, C., Patrick, F., & Lagnado, D. A. (2016). Children's use of interventions to learn causal structure. *Journal of Experimental Child Psychology*, *141*, 1–22.
- McCormack, T., Frosch, C., Patrick, F., & Lagnado, D. (2015). Temporal and statistical information in causal structure learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*(2), 395–416.
- Navarro, D. J., & Perfors, A. F. (2011). Hypothesis generation, sparse categories, and the positive test strategy. *Psychological Review*, *118*(1), 120–134.
- Nelson, J. D. (2005). Finding useful questions: On Bayesian diagnosticity, probability, impact, and information gain. *Psychological Review*, *112*(4), 979–999.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, *101*(4), 608–631.
- Pearl, J. (2000). *Causality*. New York: Oxford University Press.
- Plummer, M. (2003). A program for analysis of bayesian graphical models using gibbs sampling. In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*.
- Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology*, *43*(4), 1045–1050.
- Schulz, L. E., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science*, *10*(3), 322–332.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, *27*, 379–423.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *64*(4), 583–639.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E. J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*(3), 453–489.
- Wason, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology*, *12*(3), 129–140.