When correlation equals causation: A behavioral and computational account of second-order correlation learning in children

Conflicts of interests: none

Abstract

We examined 2- and 3-year-old children's ability to use second-order correlation learning—in which a learned correlation between two pairs of features (e.g., A and B, A and C) is generalized to the noncontiguous features (i.e., B and C)—to make causal inferences. Previous findings showed that 20- and 26-month-olds can use second-order correlation learning to learn about static and dynamic features in category and non-category contexts. The present behavioral study and computational model extend these findings to show that 2- and 3-year-old children can detect the second-order correlation between an object's surface feature and its capacity to activate a novel machine but only if the children had encoded the first-order correlations on which the second-order correlation was based. These results have implications for children's developing information-processing capacities on their ability to use second-order correlations to infer causal relations in the world.

Keywords: causal learning, computational modeling, second-order correlation learning, causal reasoning, causal representations, neural networks

Causal reasoning is a fundamental part of child development. Children must learn to encode causal structure to make predictions, generate explanations, and reason counterfactually. For example, inferring that flipping a light switch to the *on* position will cause a light bulb to illuminate requires recognizing that light switches cause light bulbs to turn on and not the other way around and that light bulbs generally do not illuminate when their switches are *off*.

There is now considerable evidence that shows that infants possess statistical learning mechanisms that support learning across a variety of domains (e.g., Haith, 1993; Kirkham et al., 2002; Saffran et al., 1996). For example, infants between 4½ and 10 months of age can use statistical information and different patterns of spatial and temporal cues to perceive causality (Cohen & Amsel, 1998; Leslie & Keeble, 1987; Oakes & Cohen, 1990; Rakison & Krogh, 2012). By 2 years of age, children make causal inferences that extend beyond the statistical regularities that they observe (e.g., Luchkina, Sommerville, & Sobel, 2018; Meltzoff, Waismeyer, & Gopnik, 2012; Sobel & Kirkham, 2006; Waismeyer & Meltzoff, 2017; Waismeyer, Meltzoff, & Gopnik, 2015; Walker & Gopnik, 2014; Walker et al., 2017). These findings suggest that infants might initially perceive and encode statistical regularity among events, but how these capacities support causal reasoning in children is unclear.

One promising mechanism that may assist in making causal inferences from statistical regularity is *second-order correlational learning*. This refers to the ability to make inferences about features that are correlated but are separated both in space and time. For example, it is commonplace for infants to register that hands (A) convey information about goal-directed action (B) (e.g., Woodward, 1998). But objects with hands (A) often also have legs (C). Can infants infer that objects with legs are goal-directed? That is, given a correlation between A and B and between A and C, do infants also register that B and C should co-occur? Second-order

correlation is a powerful domain-general learning process because learners need not be exposed to the full space of correlations for inferences to be made. For example, a learner who is presented with the two sets of aforementioned relations (i.e., A and B and A and C) may infer—through second-order correlation learning—that a novel object with legs is goal-directed

Several studies have demonstrated that young children can detect second-order correlations among the static, non-causal features of objects (Cuevas, Rovee-Collier, & Learmonth, 2006; Yermolayeva & Rakison, 2016). Rakison and Benton (2019) showed that toddlers also make such inferences about the dynamic features of objects. They habituated 20- and 26-month-olds to separate correlations between features of objects. First, they presented toddlers with stationary objects of different shapes that had different surface features (e.g., blue squares had yellow hearts inside; red circles had white crosses inside). The second involved a dynamic relation between an object and its motion trajectory (featureless blue squares and blue pentagons moved rectilinearly; featureless red circles and triangles moved curvilinearly). The question they asked was whether toddlers extracted the second-order correlation between yellow hearts and rectilinear motion and white crosses and curvilinear motion.

Toddlers were then shown a test event that was consistent with these second-order correlations (a novel object with a yellow heart inside moved rectilinearly) and a test event that was inconsistent with these correlations (a novel object with a yellow heart inside moved curvilinearly). Both age groups detected the second-order correlation: the 20-month-olds looked longer at the inconsistent test event, whereas the 26-month-olds looked longer at the consistent test event. Given that the features were presented equally during habituation, greater interest to either test event could only occur if infants learned the second-order correlation between features and motion paths. Rakison and Benton (2019) also showed that a computational (connectionist)

model provided a good description of children's inferences, including the developmental difference in looking times. These results indicated that children (and the computational model) could detect the correlation between features that are not presented simultaneously together.

The goal of the present investigation is to demonstrate that toddlers make similar inferences about the causal properties of objects. We examined whether 2- to 3-year-olds could detect a second-order correlation between objects' surface features and their causal efficacy. Critically, examining this age group allows us to expand on the Rakison and Benton (2019) findings and to address an important facet of children's causal reasoning. Numerous studies on toddlers' causal reasoning suggest that various information-processing demands influence children's ability to make certain inferences (e.g., Luchkina et al., 2018; Walker et al., 2014). To detect a second-order correlation, toddlers must be able to detect the first-order ones that form the basis of the inference. We examine whether children's inferences are influenced by this capacity in a behavioral study. We then present a computational model similar to that of Rakison and Benton's (2019). The computational model enabled us to determine to what extent second-order correlation learning for causal stimuli emerges in an associative-learning system that embodies certain information-processing constraints such as constraints on the speed of learning or constraints on the amount of information that is retained across the phases of the experiment.

Behavioral Study

Two- and 3-year-olds were shown two objects, each with a distinct correlated feature (e.g., Object₁ with feature A and Object₂ with feature B). Children then observed that Object₁ without feature A activated a novel machine, whereas Object₂ did not. Children were then asked to choose which of two new objects—one with feature A and one with feature B—would activate the machine. We predicted that children would choose the test object with feature A if

they detected the indirect relation between that feature and the machine's activation but that this choice would depend critically on their encoding of the first- and second-order associations.

Method

Participants. Participants were 32 two-year-olds (17 boys, 15 girls; $M = 29.21$ months, range = 24-35 months) and 32 three-year-olds (17 boys, 15 girls; $M = 39.83$ months, range = 36-47 months). Sample size for all analyses was determined based on a power analysis for χ^2 tests ($\alpha = .05$, Power = 0.80, and effect size = 0.50). Three additional children were tested but not included in the final analysis: two because of experimenter error and one because the child refused to participate in the experiment. All participants were tested in a quiet room in a local children's museum. No explicit information about ethnicity or SES was collected, but previous studies at the museum suggest the following breakdown: 14% Hispanic, 3% African American, 5% Asian American, 1% Native American, 58% Caucasian, 7% mixed race, and 12% no response.

Materials. This experiment used a version of the *blicket detector* (Gopnik & Sobel, 2000). The machine was a 12.7 cm x 17.8 cm x 7.6 cm black plastic box with a translucent white top. Pressing a button on a wireless remote control activated the machine, and the experimenter activated the machine as soon as an object contacted it. This created the illusion that objects activated the machine. Six 3.8 cm x 2.54 cm wooden blocks and two unique stickers were also used (Figure 1).

Procedure. Children were tested in a quiet room in a local museum and were seated across the table facing a male experimenter. The experiment consisted of two training trials and a test phase.

Training trials. In the static training trial, children were first shown a red cube and green cylinder, each positioned next to each other on the table. One object had a yellow circle sticker on it (e.g., the green cylinder); the other had a purple diamond sticker on it (e.g., the red cube). The correlation between the blocks and the stickers was counterbalanced across participants. Children were shown each block one at a time and were told, “See this toy? This toy has a yellow (or red) thing on it,” This was done twice. The order of presentation was random.

In the causal training trial, two new objects were placed on the table. These were sticker-less objects that were identical to the first pair of objects. The experimenter then introduced the machine. Children were told that the machine “lights up and plays music when some toys are put on it but not when other toys are put on it.” The two blocks were then placed on the machine one at a time. One block made the machine go (e.g., the red cube); the other did not (e.g., the green cylinder; counterbalanced across participants).

Test trial. Two new, identical blue triangular blocks were brought out. One had the same sticker that was paired with the red cube in the static training trial; the other had the same sticker that was paired with the green cylinder on the static training trial. The experimenter then said, “Here are two new toys. Can you tell me which toy makes the machine go?” If children inferred an association between the red cube and the yellow sticker during the static phase, and then an association between the red cube (without the sticker) and the machine’s activation during the causal phase, then at test they should use the blue block with the yellow sticker to make the machine go and not the other block. This is because the yellow sticker was associated—via a second-order correlation—with the machine’s activation (see Figure 1).

All children were then given a memory check in which children were shown the static-phase blocks and stickers and asked to place the stickers on the blocks to which they were initially appended.

Results

All p -values were supplemented with Bayes Factors (BFs) to quantify the relative support for the alternative hypothesis compared to the null hypothesis. BFs close to 1 indicate equal support for both hypotheses, whereas $BFs \geq 3$ or < 10 indicate, respectively, moderate and strong evidence for the alternative hypothesis (Lee & Wagenmakers, 2013). Table 1 shows the frequency with which children passed the memory control and chose the object consistent with detecting the second-order correlation between both age groups. Preliminary analyses showed that neither children's sex nor any of our counterbalancing measures were significantly related to their choice of the test object, all $\chi^2(1, N = 64)$ -values < 1.91 , all p -values $> .16$.

We next examined the relation between children's age and their choices on the test question and memory control. All 64 participants were used in these analyses. Although there was no relation between children's age (in months) and their choice of object at test, $r_s(62) = .02$, $p = .91$, there was a significant relation between children's age and their performance on the memory control, $r_s(62) = .40$, $p = .001$. A logistic regression corroborated these results. This analysis revealed that age did not significantly predict choices on the test question, Odds Ratio = 1.01, $p = .98$, 95% CI [-1.41, 3.44], $BF = 0.13$ but did significantly predict performance on the memory control, Odds Ratio = 7.13, $p < .01$, 95% CI [0.11, 14.14], $BF = 25.00$. Children who passed the memory control were more likely to choose the correct test object than children who did not pass it, $\chi^2(1, N=64) = 6.15$, $p = .01$, $BF = 5.37$. Follow-up comparisons indicated that children who passed the memory control chose the consistent test object 68% of the time (32 out

of 47 children), which is greater than expected by chance, Binomial test, $p = .02$, $BF = 5.37$. In contrast, children who did not pass the memory control chose the consistent test object 59% of the time (10 out of 17), which is not significantly different from chance, Binomial test, $p = .63$, $BF = 0.62$. This analysis suggests that 2- and 3-year-olds do not differ in their ability to use second-order correlational information but potentially differ in their ability to encode the first-order correlations on which the second-order correlation is built.

Discussion

Two- and 3-year-olds registered second-order correlations to make causal inferences but only if they remembered the feature-block pairings from the static phase. The ability to remember the initial correlation in this procedure develops between the 2 and 3 years of age. This result is consistent with other findings that suggest that even younger children may be capable of encoding these correlations (Cuevas et al., 2006; Yermolayeva & Rakison, 2016). In those studies, the initial association between the stimuli were presented for much longer, which may have allowed for better encoding of the information.

We have suggested that second-order correlational learning is a domain-general associative-learning mechanism that requires sufficient information-processing capacities. To examine this hypothesis, we constructed a similar computational model (i.e., an autoencoder connectionist model) to that of Rakison and Benton (2019) to assess whether it can explain children's second-order correlational learning and performance on the memory control. The success of the model in simulating children's second-order correlation learning performance as well as their performance on the memory control would be important because it would suggest

that an associative-learning mechanism with sufficient information-processing capacities mechanism can form second-order correlations for causal stimuli.

Computational Model

Method

Network architecture. We used a three-layer autoencoder neural network that was trained using backpropagation and momentum. The learning rate, momentum, weight decay, and number of hidden units, respectively, were set to 0.08, 0.9, 0.001, and 15 to correspond to children who passed the memory check and to 0.001, 0.9, 0.005, and 13 to correspond to children who failed the memory check. This implemented a simple model of development, capable of simulating the difference in choices on the test question between children who passed and failed the memory check. Thirty-two networks—each initialized with small (distribution range = +/- 0.8), random weights—were run for each simulation. This number paralleled the number of participants in the behavioral study. Finally, both network types received 50 training epochs (or trials) because children who passed and failed the memory check in the behavioral study received an equal number of training trials. Fifty training epochs was chosen to mirror Rakison and Benton (2019), who used the number of training epochs as a way of simulating the “age” of the model. Fifty training epochs simulated the models for the 26-month-olds (which would be akin to the youngest participants tested here). Nonetheless, the simulation results reported below held regardless of whether models were trained for as few as 15-20 epochs or as many as 200-250 epochs.

The input to the network consisted of patterns of activity across three different groups of input units: a shape, color, and feature input group (see Figure 2). The shape group consisted of

distributed activity across twelve units, whereas activity was encoded locally across three units for the color group and across two units for the feature group. Critically, the similarity between any two input patterns within each of the three groups was orthogonal to ensure unbiased output responses. Each input group projected to a group of hidden units, which in turn were connected to output groups that were copies of each input group and an additional single-unit output group that took on a value of 1 whenever a causal object was presented and a value of 0 otherwise.

We represented the task in terms of three input groups for three reasons. First, these dimensions would be available to low-level perceptual processes that segment tasks and stimuli into component parts. Second, representing the task and stimuli in terms of three feature groups would be the absolute minimum number of input groups needed to simulate successful second-order correlation learning. Third, the simulations were designed to account for children's second-order correlation learning and memory-control performance rather than to capture the precise stimuli that children attended to during training and test. We do not suggest that these simulations are full cognitive models of second-order correlation learning in children. Rather, they serve as proof of concept that these models can register second-order correlations from only the statistical regularity in the data.

Training. Consistent with the sample size of the behavioral study, networks were trained. Thirty-two simulations were run to simulate the successful networks and 32 for the unsuccessful networks. Individual networks were simulated by initializing a fresh set of small, random weights. In the static training phase, networks were trained to predict an internal circle shape on the output when given as input a stationary red object and an internal diamond shape on the output when given a stationary green object. The precise feature-block pairings presented to networks during the static training phase were counterbalanced. In the causal training phase,

networks were trained to predict that either a featureless red cube or green cylinder made the machine go (counterbalanced).

Testing. Networks were then presented with two identically shaped objects. One of the test objects possessed the feature that was indirectly associated with the machine's activation and the other object possessed the feature that was indirectly associated with the machine's inactivation. Correct test performance was indicated by activity in the single activity output group that was above 0.5 for the object that was designated as the correct test object, whereas chance or incorrect performance was indicated by equal activity in the activity output unit for both test objects.

Results

All of the networks with parameters that were more consistent with passing the memory control placed greater weight on the consistent test object (32 out of 32 networks), Binomial test, $p < .001$, $BF = 9105.00$. Networks with parameters that were more consistent with not passing the memory control placed greater weight on the consistent test object 50% of the time (16 out of 32), which was not significantly different from chance, Binomial test, $p = 1.00$, $BF = 0.40$. Critically, the number of consistent choices on the test question for networks that passed the memory control ($N = 32$) did not differ from that for children who passed the memory control ($N = 47$), $\chi^2(1, N=79) = 0$, $p = 1.00$, $BF = 1.01$. Likewise, networks that did not pass the memory control ($N = 32$) did not differ from children who did not pass it in their choice of the consistent object at test ($N = 17$), $\chi^2(1, N=49) = 1.38$, $p = .24$, $BF = 0.76$.

To ensure that these results were not idiosyncratic to the precise parameter values used in the reported simulation, we ran a separate set of 48 other simulations in which the original parameter values were modified. For networks whose parameters were more consistent with

passing the memory control, the number of hidden units was changed from 15 to 13, weight decay was changed from .001 to .002, and momentum was changed from .9 to .875. For networks whose parameters were more consistent with not passing the memory control, the number of hidden units was changed from 13 to 11, weight decay from .005 to .006, and momentum from .9 to .875. These parameter-value changes were consistent with those made to Rakison and Benton's (2019) computational model. The results of these changes are shown in Table 2. Critically, these additional simulations show that the same pattern of results as that in the original models and children regardless of how differences in memory are represented.

Discussion

The simulations showed that networks that possessed more robust memory capacities weighted the consistent test object more heavily than networks that possessed less robust memory capacities. Critically, this result mirrored that in the behavioral study. The improved performance of networks that corresponded to children who passed the memory control relative to those who failed it resulted from their increased processing speed (e.g., higher learning rate), improved memory retention (e.g., lower weight decay and momentum), and larger memory storage (e.g., more hidden units). The greater learning rate for networks that reflected the children who passed the memory control enabled them to encode the first-order correlations within the allotted training trials. In contrast, the lower learning rate for networks reflected the children who failed the memory control was insufficient to allow these networks to register the first-order correlations during training.

Combined with a higher learning rate, the lower weight decay and momentum and a larger set of hidden units allowed the networks with more robust information-processing capacities to form an internal (hidden) representation—through the set of hidden units—that

captured the first-order correlations (e.g., yellow sticker and red cube; red cube and machine activation) as well as the second-order correlation between them (e.g., yellow stickers and machine activation). Crucially, this ability to form a hidden representation that captured both the first- and second-order correlations served to facilitate second-order correlation learning only in those children who passed the memory control.

General Discussion

The present investigation demonstrated that young children can encode second-order correlations to make causal inferences but only if they recalled the feature-block pairings from the static phase. The computational model showed that an associative-learning mechanism combined with sufficient information-processing capacities is critical for registering second-order correlations among causal stimuli. Together, this study extends previous ones on this topic by showing that children with sufficient information-processing capacities can use second-order correlation learning to infer causal relations.

Children's developing information-processing abilities may also explain why second-order correlation learning emerges at different ages in studies on this topic (e.g., Rakison & Benton, 2018; Yermolayeva & Rakison, 2016). For example, in Cuevas et al.'s (2006) study, 6-month-olds showed second-order correlation learning but only after extended exposure to the stimuli. In contrast, children in the current study were given two, brief exposures to the first-order correlations that constituted the second-order correlation. This difference in familiarization time might have affected the current children's ability to recall the initial association. Critically, age was not a critical factor for encoding the second-order correlations in the current study. Instead, the results indicated a relation between children's age and their developing memory

demands. The computational model simulated how this difference in memory capacity might contribute to differences in one's ability to detect second-order correlations. This suggests that second-order correlational is an important learning mechanism that depends children's developing memory and information-processing capacities. Importantly, the "age" of the model (represented by the number of epochs used to train it) was not a relevant factor in the simulation. The results were similar regardless of whether the model was trained for 15 epochs or 200 epochs.

More generally, these results suggest that children might be using second-order correlation learning to interpret correlational data in causal ways. To illustrate, consider a study by Saxe, Tzelnic, and Carey (2007). Nine-month-olds were habituated to alternating events in which a bean bag was tossed either from behind an opaque screen located on the left-hand side of a stage or an opaque screen on the right-hand side of the stage. The screens were then lowered to reveal a hand or puppet with human-like eyes on one side of the stage and a toy truck on the other side of the stage. Saxe et al. (2007) found that 9-month-olds looked longer at the event in which the beanbag ostensibly was tossed from the side on which the toy truck was located. They interpreted these data as suggesting that infants understand that people, but not inanimate objects, are agents that can cause ballistic motion. An alternative view is that infants' conception of people as agents may well have derived from encoding two separate correlations. The first is that people have hands and/or eyes. The second is that people move on their own. Based on these separate associations, infants might have understood that hands/eyes and ballistic motion are correlated through a second-order association. Future studies should look carefully at the explanatory nature of second-order correlations as a mechanism for children's causal inference, particularly in interpreting investigations with infants.

To conclude, the present experiment demonstrates that second-order correlation learning for causal stimuli is present by 2 years of age and that encoding the first-order relations—on which the second-order correlations are built—is critical for second-order correlation learning. The present results also suggest that second-order correlation learning may be grounded in associative processes and represents a powerful mechanism by which infants and children can form broad generalizations that span domains and content areas.

Acknowledgments. We thank the children and their caregivers for their participation.

References

- Cuevas, K., Rovee-Collier, C., & Learmonth, A. E. (2006). Infants form associations between memory representations of stimuli that are absent. *Psychological Science*, 17(6), 543-549.
- Cohen, L. B., & Amsel, G. (1998). Precursors to infants' perception of the causality of a simple event. *Infant behavior and development*, 21(4), 713-731.
- Gopnik, A., & Sobel, D. M. (2000). Detecting blickets: How young children use information about novel causal powers in categorization and induction. *Child development*, 71(5), 1205-1222.
- Haith, M. M. (1993). Future-oriented processes in infancy: The case of visual expectations. *Visual perception and cognition in infancy*, 235-264.
- Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition*, 83(2), B35-B42.
- Lee, M. D., & Wagenmakers, E. J. (2013). Bayesian data analysis for cognitive science: A practical course.
- Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality?. *Cognition*, 25(3), 265-288.
- Luchkina, E., Sommerville, J. A., & Sobel, D. M. (2018). More than just making it go: Toddlers effectively integrate causal efficacy and intentionality in selecting an appropriate causal intervention. *Cognitive Development*, 45, 48-56.

- Meltzoff, A. N., Waismeyer, A., & Gopnik, A. (2012). Learning about causes from people: Observational causal learning in 24-month-old infants. *Developmental psychology*, 48(5), 1215.
- Oakes, L. M., & Cohen, L. B. (1990). Infant perception of a causal event. *Cognitive Development*, 5(2), 193-207.
- Rakison, D. H., & Benton, D. T. (2018). Second-Order Correlation Learning of Dynamic Stimuli: Evidence from Infants and Computational Modeling. *Infancy*, 24(1), 57-78.
- Rakison, D. H., & Krogh, L. (2012). Does causal action facilitate causal perception in infants younger than 6 months of age?. *Developmental science*, 15(1), 43-53.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926-1928.
- Saxe, R., Tzelnic, T., & Carey, S. (2007). Knowing who dunnit: Infants identify the causal agent in an unseen causal interaction. *Developmental psychology*, 43(1), 149.
- Sobel, D. M., & Kirkham, N. Z. (2006). Blickets and babies: the development of causal reasoning in toddlers and infants. *Developmental psychology*, 42(6), 1103.
- Waismeyer, A., & Meltzoff, A. N. (2017). Learning to make things happen: Infants' observational learning of social and physical causal events. *Journal of experimental child psychology*, 162, 58-71.
- Waismeyer, A., Meltzoff, A. N., & Gopnik, A. (2015). Causal learning from probabilistic events in 24-month-olds: an action measure. *Developmental science*, 18(1), 175-182.

- Walker, C. M., & Gopnik, A. (2014). Toddlers infer higher-order relational principles in causal learning. *Psychological science*, 25(1), 161-169.
- Walker, C. M., Lombrozo, T., Williams, J. J., Rafferty, A. N., & Gopnik, A. (2017). Explaining constrains causal learning in childhood. *Child development*, 88(1), 229-246.
- Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, 69(1), 1-34.
- Yermolayeva, Y., & Rakison, D. H. (2016). Seeing the unseen: Second-order correlation learning in 7-to 11-month-olds. *Cognition*, 152, 87-100.

Table 1. Frequency of children choosing the test object consistent with registering the second-order correlation

	Two-year-olds	Three-year-olds
Responded Correctly on Memory Question (n = 47)	10 out of 17	22 out of 30
Responded Incorrectly on Memory Question (n = 17)	9 out of 15	1 out of 2
Total	19 out of 32	23 out of 32

Table 2. The effects of changing parameters on the computational model

Networks corresponding to children who passed the memory control		
	Binomial Test Results	<i>BF</i>
Learning rate: .08 to .06	(8 out of 8), $p < .01$	9
Hidden units: 15 to 13	(8 out of 8), $p < .01$	9
Weight decay: .001 to .002	(8 out of 8), $p < .01$	9
Momentum: .9 to .875	(8 out of 8), $p < .01$	9
Training Trials: 15	(9 out of 10), $p < .05$	5.49
Networks corresponding to children who did not pass the memory control		
	Binomial Test Results	<i>BF</i>
Learning rate: .001 to .0008	(4 out of 8), $p = 1$	0.65
Hidden units: 13 to 11	(4 out of 8), $p = 1$	0.65
Weight decay: .005 to .006	(5 out of 8), $p = .73$	0.75
Momentum: .9 to .875	(5 out of 8), $p = .73$	0.75

The numbers inside the parentheses correspond to the proportion of networks for each parameter-value change that responded with the consistent object at test to the total number of networks that responded with both types of objects.

Figure 1.

Figure 2.

Figure Captions

Figure 1. Examples of the stimuli used in the study. Children were shown the static training trial first, followed by the causal training trial. Two novel identical objects were presented at test. The inconsistent test object had the surface feature that was indirectly associated with machine's nonactivation. The consistent test object had the surface feature that was indirectly associated with the machine's activation. Children were asked which object made the machine go.

Figure 2. A schematic of the model that was used to simulate the results from the behavioral study. The three input groups were connected to one hidden layer, which in turn was connected to three output groups that mirrored the three input groups. The task was to recreate the input patterns along the three input groups on each of the corresponding output groups as well as to respond with a 0 or 1 in the Activity group if a causally inefficacious or efficacious object was presented, respectively.