Representations of Abstract Relations in Early Childhood

Nicole H. Coates, Max H. Siegel, Joshua B. Tenenbaum, Laura E. Schulz
{nhcoates, maxs, jbt, lschulz}@mit.edu
Department of Brain and Cognitive Sciences, MIT

Abstract

In science, we use common graphical representations to indicate changes in events over time, independent of domain. Are children also sensitive to abstract patterns in the ways events change over time? In a series of four experiments, we show that young children (range: 48-84 months) distinguish different function families (Exp 1). Children can also distinguish specific function types within function families (e.g., between linear and sigmoid monotonic functions; Exp 2); nonetheless, they will group different function types within a family together, rather than with functions in a different family (Exp 3). Finally, we show that children's sensitivity to functions is abstract, allowing them to match observable causes to verbal descriptions of their effects (Exp 4). These results suggest that although some aspects of function understanding, like learning how to interpret graphs, requires formal education, the ability to identify abstract functional relationships is intuitive and early emerging.

Keywords: children; data representations; abstract concepts; functions

Introduction

That the world has structure is apparent in the ways that patterns change over time. This is true in the natural world; whether we look at strata in rocks to tell us where they are from and when they were formed, or tree rings that tell us how old a tree is, there are regularities that underlie the ways that things change. Thinking about these regularities is necessarily abstract – for if not, how could we observe similarities between different phenomena? – but must allow specialization, for otherwise we would be unable to tell when a regularity describes a particular instance. In science, we use common representations to indicate certain kinds of changes, independent of domain. We can talk about monotonic increases in bird populations, hospital billing charges, the growth rate of proteins, or the progression of diseases. We can think about U-shaped curves, such as the impact of cortisol on memory, the impact of sleep on health, the percentage of top talent on a basketball team and team scores, or age and measures of life satisfaction. And we can understand periodic changes, whether we are talking about cell expression, cardiac rhythms, the global economy, or the appearance of sunspots. The ability to reason about change across domains is undoubtedly essential to science and much of everyday life. But its origins are less clear. Is this an ability shared by all humans, or is it a product of education or culture? Does it build upon other cognitive abilities, or is it itself a foundation for later thought? Are even young children able to distinguish monotonic, U-shaped, and periodic functions?

Many of the cognitive traits supporting scientific inquiry emerge at a very young age (Gopnik & Meltzoff, 1997; Schulz, 2012; Weisberg & Sobel, 2022). Without any formal training, young children attend to statistical patterns in data (Saffran, Aslin, & Newport, 1996) and to whether evidence is sampled randomly or selectively (Xu & Denison, 2009); they selectively explore when evidence is surprising (Stahl & Feigenson, 2015; Legare, 2012) or confounded (Schulz & Bonawitz, 2007), and they attend to patterns of covariation in data (Gweon, Tenenbaum, & Schulz, 2010) and distinguish merely associated variables from genuine causes (Gopnik, Sobel, Schulz, & Glymour, 2001). Thus there is reason to think that sensitivity to abstract patterns in the ways events change over time might also be part of early, intuitive reasoning.

On the other hand, even much older learners – middle schoolers, high school students and college students – sometimes have difficulty understanding functional relationships in data. For instance, students have difficulty abstracting information from graphs and connecting abstract functions to actual changes in the world (McDermott, Rosenquist, & Van Zee, 1987; Rodrigues, 1994; Ciccione et al., 2022). However, in educational contexts, students are asked to reason about complex, unfamiliar content and are often asked to learn both the content material and the graphical representations of this content at the same time. Here we test children's ability to draw abstract inferences, not by looking at their ability to reason about a novel cultural technology, like a graph, but by asking whether they can use commonalities in functional relationships to match perceptually distinct stimuli to each other, or to verbal descriptions.

Relatively little work has looked at preschoolers' ability to learn abstract functions, although considerable work has looked at children's ability to make analogical mappings based on abstract relationships. Studies show for instance that four and five-year-olds can match small-big-small patterns in the size of circles to squares (Kotovsky & Gentner, 1996). Interestingly, four and five-year-olds fail these relational match-to-sample tasks when they are asked to match across different dimensions or polarities. More critically, these tasks involve static images that vary according to size, texture, layout or other observable features; they test children's ability to make

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society.*

analogical mappings between patterns but they do not test children's ability to represent functions: dynamic changes over time.

To our knowledge, only two studies have looked at whether children are sensitive to dynamic changes over time. One study introduced preschoolers to causal scenarios where a wand magically transformed object properties. For instance, a small apple could become a big apple, or a single object could become five objects. Children were able to predict the kind of transformation that would occur in a novel object given prior examples (Goddu, Lombrozo, & Gopnik, 2020). However, the tasks involved a discrete transformation from one state to another, rather than continuous functions, and what was at stake was whether the causal context facilitated children's ability to make analogical mappings ("the wand turned this object from one object to five thus it will have the same effect on this object") rather than to distinguish the functions themselves. Other work looked at children's understanding of discrete versus continuous functions in a causal context, and showed that children distinguished these two types of functions, matching discrete causes to discrete effects and continuous causes to continuous effects (Magid, Sheskin, & Schulz, 2015).

Building on this prior literature, here we study children's ability to represent and reason with a broader range of abstract functional types. We focus on three classes of continuous functions which we refer to as "function families": monotonic, U-shaped, and periodic functions. But of course, within each function family, there are many different specific types of functions. For instance, monotonic functions can be linear or sigmoid; U-shaped functions can be quadratic or cube root; and periodic functions can be sinusoidal or exponentiated triangular wave functions. Our experiments use these functions in particular (see Figure 1), but they reflect just two possible instances of three representative abstract function classes that, we hypothesize, children may be able to represent and reason about in causal settings.

In Experiment 1 we ask the very simplest version of this question: Do children distinguish these different function families and match identical functions to each other, across different concrete cause and effect variables? In Experiment 2, we ask whether children distinguish function types within a function family. In Experiment 3, we ask whether they nonetheless can match distinguishable function types within the same family to each other, rather than to a different function type from a different family. Finally, in Experiment 4, we ask whether children's representations of functions are sufficiently abstract that they can match observed causal functions to verbal descriptions of their effects.

Previous research suggests that children are better at inferring abstract relationships in causal rather than arbitrary contexts (i.e., "Which makes this one go?" rather than "Which goes with this one?"; see e.g., Goddu et al. (2020)). Thus although our core interest is in children's understanding of abstract functional relationships, in our experiments we always asked children to match causes and effects. In each experiment, children were introduced to an alien greenhouse and were told, "In this greenhouse, different sets of lights make the flowers bloom in different ways". Children were shown two sets of lights and two sets of flowers on each trial; one set of lights dimmed and brightened following a particular function; then a second set of lights dimmed and brightened following a different function. Children then saw one set of flowers open and close following one of the two functions; and then saw the second set of flowers open and close following the other function. No other perceptual cues linked the candidate causes and effects. Children were asked to decide which set of lights made each set of flowers bloom. Critically, participants never saw the lights and flowers changing together; each set of flowers and each set of lights was presented individually, so children must track and remember the functional dynamics of both candidate causes and both effects to identify the matching pairs.

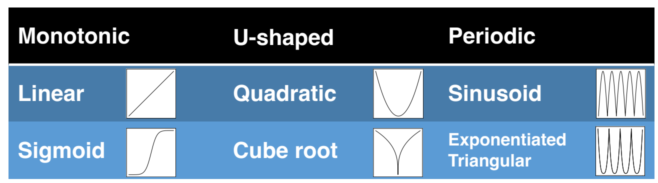

Figure 1: Each function used in our experiments, organized by function family. In Experiment 1, the first function in each family type (linear, quadratic and sinusoidal) were tested. In Experiment 2, each function within a family were tested against each other. In Experiment 3, functions within function families were tested against each other. In Experiment 4, the first in each function family were again tested.

Experiment 1

In Experiment 1, we studied the simplest and starkest contrast between functional forms that we expected children might be sensitive to. We contrasted a single monotonic function (linear), a single U-shaped function (quadratic), and a single periodic function (sinusoidal), in pairwise contrasts across six trials.

Participants Twenty-nine children of 32 pre-registered[1] (mean age: 65.6 months, range: 50 months - 83 months; `https://osf.io/gs6nk?view_only=None` were recruited through the asynchronous online testing platform Lookit (Scott & Schulz, 2017) and through a direct email to families in a database for studies in a social cognition lab and from a list of parents recruited on Prolific. Fifteen additional children were excluded for failing the inclusion question.

[1]Recruitment was suspended early due to a previous conference deadline but will resume for a complete final sample, with an explicit caveat that the data were observed at n = 29. The results were sufficiently robust that the three remaining children are unlikely to change them.

Stimuli & Procedure We generated animations of flashing lights and blooming flowers using the P5.js animation library (`https://www.p5js.com`); the flowers were based on the open source flower generation code made public at `https://github.com/anokhee/botanicals`. We set the luminosity of the lights and extension of the flower petals over time according to linear, quadratic, or sinusoidal functions. See Figure 2d for a visualization of these functions. Children were tested online in their homes; stimuli were presented in the form of a keynote movie with a recorded narration. No experimenter was present. Parents were instructed that they could click to report the response if the child had trouble using the trackpad or mouse and pointed or responded verbally. They were reminded that we were interested in how children of different ages perceived the stimuli and that they were not to prompt their child or interfere with their responses. Children were first introduced to an alien greenhouse and told that special lights make alien flowers bloom. Children were exposed to the set up of each trial: lights appear on the left and right sides of the screen, whereas flowers appear on the top and bottom, and at test are only shown one set of flowers with two sets of lights. They were then prompted to answer an inclusion trial, which showed a light flash once and another light flash ten times; it was intended to be obvious enough to get correct, but in the same format as test trials to ensure children understood the task. Following the inclusion trial, the 6 test trials were presented (3 concepts, 2 trials each) in random order. For each trial, children were first shown lights animated on the left, then lights animated on the right. Children were asked "Can you see the ways these sets of lights are different?" Then, as the lights stayed on the screen, one set of flowers were presented on top, and one set at the bottom so that all four stimuli were on the screen. Children were asked "Can you see the way these two sets of flowers are different?" At test, children were shown one set of flowers, and two sets of lights. Children were asked "Can you point to the set of lights that make these flowers bloom?" (See Figure 2a). There were two buttons below the stimuli for children to press to give their response. The flowers and lights for each trial type varied in color for each concept. Each of the two kinds of flowers was the target on one trial. After every trial, children saw a picture of an alien that said "Good job!" but they did not receive feedback on their performance.

Results Children successfully used abstract correlations between functions to match the candidate cause and effects for all six functions according to a mixed effects regression with random intercepts of participant and trial type to account for the within-subjects design ($\beta = 1.41$, $z = 5.20$, $p < 0.001$). Adding a fixed effect of age did not significantly improve model fit ($\chi^2(1) = 2.31$, $p = 0.13$). See Figure 3a.

Experiment 2

In Experiment 2, we ask whether children distinguish specific types of functions within function families. For instance, can children distinguish a linear function from a sigmoid function even though they are both monotonic? To test this, children were again introduced to the alien greenhouse. This time, children saw three sets of flowers blooming and closing back up or three sets of lights dimming and brightening; each set was introduced one at a time. Two sets bloomed following identical functions and one set was the odd one out. Across six trials, we asked if children could distinguish two types of monotonic functions (linear, sigmoid), two types of U-shaped functions (quadratic, cube root), and two types of periodic functions (sinusoid, exponentiated triangular wave).

Participants A pre-registered sample of thirty-two children (mean age: 66 months, range: 48 months - 83 months) were recruited through Lookit. Ten children were excluded and replaced due to failing the inclusion trials, and one child was excluded and replaced due to oversampling. Analysis and procedures were pre-registered on OSF (`https://osf.io/uxtrn/?view_only=8ce6a8f5cee54c9f921764527fe0d99a`) and were within participants.

Stimuli & Procedure The setup was similar to Experiment 1. Children were told that they would see flowers blooming and closing back up and lights going brightening and dimming. They were told that in each set of three, two of the sets changed in the exact same way, and one set was just a little bit different. Three trials used lights and three used flowers but for simplicity, we'll refer to flowers throughout. See Figure 2b for the setup. Children were first given two inclusion trials where the odd one out was very obvious (e.g., two sets just bloomed and closed once; another bloomed and closed eight times) to ensure children understood the task. Following the inclusion trials, the 6 test trials were presented (3 function family contrasts, 2 trials each; for instance, one trial might have two sinusoidal functions and one exponentiated triangular wave function) in random order. For each trial, children were shown the sets of flowers one at a time, and then the narrated voice said "Let's see that again!". Children were shown all three movies again, and at test were asked, "Can you point to the set of flowers that change a little bit differently than the others?" Three buttons appeared under the stimuli for children (or their parents if children could not click) their answer. After every trial, children saw a picture of an alien that said, "Good job!" but they did not receive feedback on their performance. Children saw six trials, two of each contrast (linear/sigmoidal; quadratic/cube root; sinusoidal/exponentiated triangular wave); and within each contrast, each function was the target once.

Results and Discussion We fit a logistic mixed effects model predicting children's choice of the target response, with random effects for child and trial type. The model estimated the likelihood of giving a correct response on any one trial as 52% of the time (95% CI $[0.43, 0.51]$), which was significantly different from the chance response of 33% ($z = 4.04$, two-tailed, $p < .001$). Adding a fixed effect of age did significantly improve model fit ($\chi^2(1) = 4.87$, $p =$

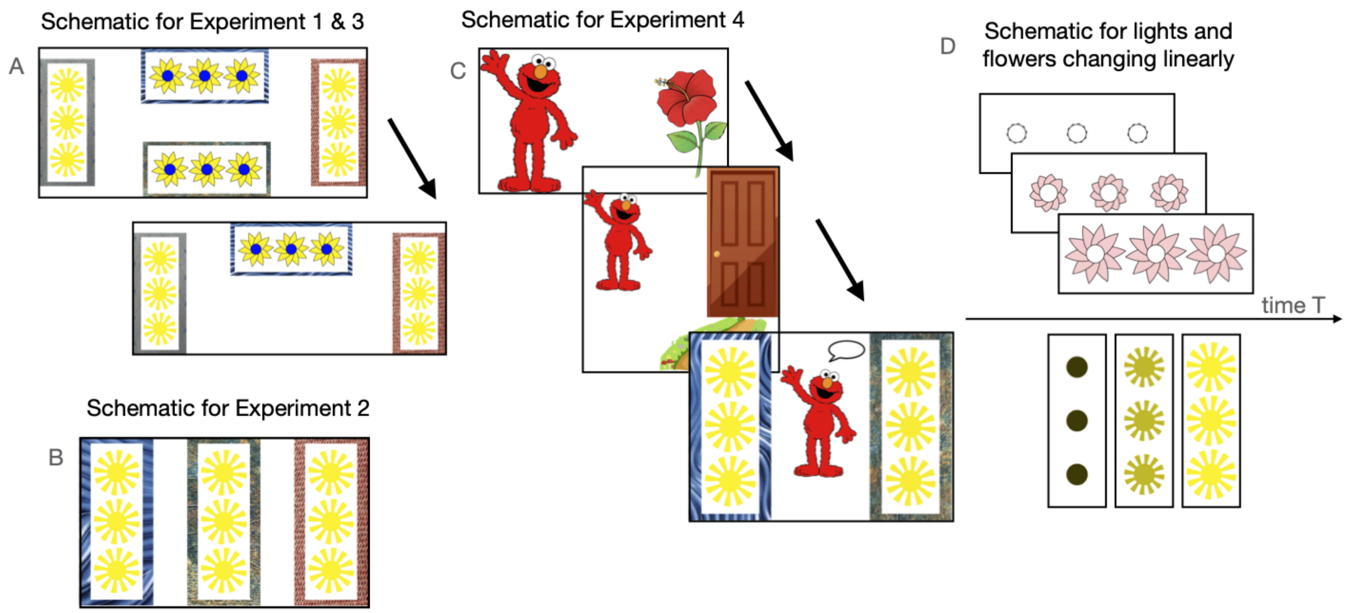

Figure 2: (A) Schematic of the design in Experiments 1 & 3: Children see one set of lights on the left and one set on the right and one set of flowers on the top and another set on the bottom of the screen. Each set is played individually as a movie, exhibiting the functional changes. (See text for details.) At test, both sets of lights are presented again as static images and only one set of flowers (top or bottom) is presented. Children were asked which set of the lights made the target set of flowers bloom. (B) Schematic of Experiment 2. Three sets of flowers or lights were presented and the movies illustrating the functional changes were played, one at a time from left to right. At test, children were asked to pick the set that was the odd one out. (C) Schematic of Experiment 4. Children were shown a character and a picture of flowers. A door came up, obscuring the child's view of the flowers (but not the character's). Children then saw a set of lights on the left and a set of lights on the right. Each set of lights was played as movie, exhibiting the functional changes. Then the character described the changes in their flowers. At test, children were asked to pick which set of lights made that character's flowers bloom. (D) Schematic of one functional change (a linear change): The lights change in luminescence while the flowers change in petal size.

0.01).See Figure 3b. Experiment 2 thus suggests that children can distinguish within functions of the same family. Given this, will they nonetheless generalize across different types of functions within function families?

Experiment 3

In Experiment 1, we showed that children represent abstract functions insofar as they match identical functions to each other across different perceptual stimuli. In Experiment 3 (See Figure 2a), we ask whether children can make an even more abstract inference: generalizing across different function types within function families while distinguishing function types across families (e.g., matching a linear and sigmoid function with each other but not with a sinusoidal function). As in Experiment 1, we asked children to match causes and effects. Again, critically, children never see the lights and flowers changing together; each set of flowers and each set of lights is presented on their own. Children must track and remember the functional dynamics of each candidate cause and effect to link them.

Participants A pre-registered sample of 32 children (mean age: 53 months, range: 46 months - 84 months) were recruited through Lookit. Twelve children were excluded due to failing the inclusion trials, and one child was excluded due to oversampling. Analysis and procedures were pre-registered on OSF (`https://osf.io/ct39q/?view_only=8c24ff4fcd154563b787b8a2d62f04bf`) and were within participants.

Stimuli & Procedure The setup for Experiment 3 was the same as Experiment 1 except here, no two functions were exactly alike. Children had to match causes and effects within function families. Children saw six trials, two of each contrast (monotonic/U-shaped; monotonic/periodic; U-shaped/periodic); a different specific function within each function family was the target once for each contrast.

Results and Discussion Children successfully generalized across function families for all three function families according to a mixed effects regression with random intercepts of participant and trial type to account for the within-subjects design ($\beta = 0.734$, $z = 3.31$, $p < 0.001$). Adding a fixed effect of age did significantly improve model fit ($\chi^2(1) = 6.55$, $p = 0.01$). See Figure 3c. Experiment 3 thus found that children are sensitive to higher-order functions; they were able

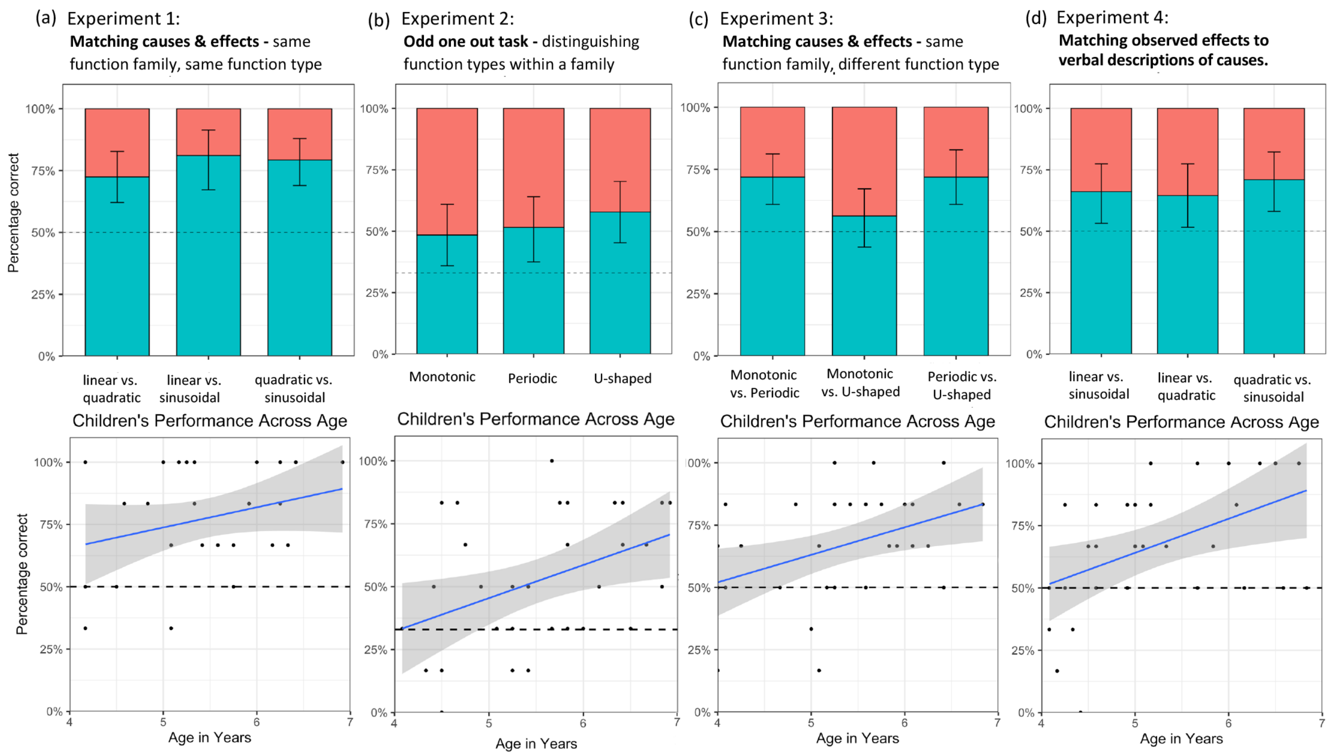

Figure 3: Results for all experiments. Graphs show percentage correct by trial-type (e.g., "linear vs. quadratic" indicates trials where children saw linear and quadratic lights). Note that the dashed lines on all plots indicate chance performance level: In Experiments 1, 3 & 4, chance is 50%. In Experiment 2, chance is 33%.

to generalize across function types within families and distinguish function types across function families. Children's ability to match cause and effect according to functions even in a context in which no two functions were identical suggests that they are capable of abstract reasoning about function kinds. However, thus far, all our experiments have used visual stimuli. Are children's abstract representations limited to observable changes, or can children match an observable function to a verbal description of that function?

Experiment 4

In Experiment 4, we ask whether children distinguish functions within function families, given descriptions of the functions. This experiment probes both the flexibility of children's representations and the degree to which children have explicit, linguistic access to these representations.

To test children's sensitivity to descriptions of functions, children were again introduced to an alien greenhouse, and were told, "In this greenhouse, there are special lights. The lights make the flowers bloom. Each set of lights controls a different set of flowers. They were then introduced to three characters: Elmo, Grover, and Cookie Monster who can each see a different set of flowers. However, an animated door then appears and obscures the child's view of the flowers, while preserving the character's view On each trial, children see two sets lights dim and brighten; a single character then describes the set of flowers they are looking at. We ask if children can use that character's verbal description of the flowers they see blooming to decide which set of lights made those flowers bloom. Children did not hear the verbal description and watch the videos of lights at the same time; children had to remember the verbal description they heard and map that to one of two sets of lights they had observed.

Participants A pre-registered sample of thirty-two children (mean age: 52 months, range: 48 months - 81 months) were recruited through Lookit. Nine children were excluded due to failing the inclusion trials. Analysis and procedures were pre-registered on OSF [`https://osf.io/beupq/?view_only=89feb524452f4e0baebfc88f6a7fa261`] and were within participants.

Stimuli & Procedure The set up was similar to the previous experiments: children were again introduced to an alien greenhouse, and told that special lights make flowers bloom; each set of lights controls a different set of flowers. Children were then shown three sets of lights, playing one at a time, to familiarize them with the lights. They were then introduced to three characters: Elmo, Grover and Cookie Monster, individually, each with a different colored flower. An animated door appeared and obscured the flowers from the child's view but (apparently) preserved the line of sight for the character. The narrated voice then said, “Each of our friends is looking at a different set of flowers. Remember the lights control the way the flowers bloom. We can’t see the flowers but our friends can and they are each going to tell us about the set of flowers they are looking at”. Children then heard each character describe their set of flowers. Elmo always described a monotonic function (e.g. “Elmo says the flowers he is looking at start low and go up and up over time”); Grover always described a U-shaped function (”Grover says the flowers he is looking at start high and go down and then go up again.”) and Cookie Monster always described a periodic function (”Cookie Monster says the flowers he is looking at go back and forth; they go up, then go down, and then go up and go down again a few times.”). Next, two sets of lights appeared (See Figure 2c), and children are told that for the rest of the game, one of the characters will describe the flowers he sees, and that their job is to decide which of the two sets of lights is making that character’s flowers bloom the way they do. After all six test trials were completed, children received an inclusion trial. The inclusion trial was an “Odd One Out” task in which children saw three sets of flowers, and were prompted to choose the set of flowers that “changes a little bit differently” than the other two to ensure that children were attentive.

Results and Discussion Children successfully matched functions with verbal description for all six functions (mixed effects regression with random intercepts of participant and trial type to account for the within-subjects design; $\beta = 0.87$, $z = 3.40$, $p < 0.001$). See Figure 4. Adding a fixed effect of age did significantly improve model fit ($\chi^2(1) = 7.01$, $p = 0.008$). See Figure 3d. This suggests that children are able to represent highly abstract aspects of functions, at the level captured by natural language descriptions. That is, children are not only able to perceptually match functions within a type or a family defined by visual features, they are also able to match the perceptual change of a light turning on and off with the way characters merely talked about a flower changing. Thus children were able to match simple linguistic descriptions of functions to observed changes in variables.

General Discussion

Across four experiments, we showed that children as young as four represent abstract functions. Children can distinguish linear, quadratic and sinusoidal functions (Exp. 1) and distinguish different types of functions even when they belong to the same family (e.g. sigmoid vs. linear monotonic functions) (Exp. 2). Children can also generalize across these function types within a family, distinguishing them from function types from a different family (Exp. 3). Finally, children’s ability to form these abstract representations is not restricted to data which is directly observable; children can match functions to appropriate verbal descriptions as well (Exp. 4). These results suggest that the ability to think about how variables change over time – despite surface differences in the specific features that vary – emerges early in development. These abilities may lie at the foundation of cultural advances in our ability to think about functional relationships in more complex contexts, including those involved in scientific inquiry.

Of course, the task we gave children here was relatively simple: children were presented with forced choice tasks and given contrastive cues. Future work might look at whether children can represent functions in more open-ended contexts. And while we tested a relatively broad range of functions here, it is still a tiny subset of possible functions; it would be interesting to know the limits, as well as the capabilities, of children’s sensitivity to functional forms.

As discussed, we used a causal context here because previous work has shown that children are more sensitive to abstract correlations in these contexts (Goddu et al., 2020). The causal context was of course artificial (e.g., in the real world, lights shining at a constant rate can cause a monotonic increase in the growth of flowers). It would be interesting to know to what extent children’s prior knowledge of actual causal mechanisms might enable or inhibit their representation of these functions, and whether children can use commonalities among functional changes in variables not just to identify but to generate causal hypotheses.

Why did children as young as four and five succeed in this task although they struggle in other relational match to sample tasks when the perceptual dimensions of the stimuli don’t readily align (Kotovsky & Gentner, 1996)? As noted, the causal context may have helped. We also note however, that our task may have actually provided more structural alignment than typical match to sample tasks. In the standard tasks, children are given a target and asked to choose between the match and the distractor. The consequence of course, is that the target and the match are well aligned but the distractor is genuinely a distractor. By contrast, in our task, each of the two causes matched with one of the two effects. Thus, we believe our results are consistent with the idea that children benefit from structural alignment in recognizing abstract relationships.

Perhaps most importantly, it would be interesting to know how precisely children are representing these functions. Mathematically, it is a relatively simple process to discover that say, linear functions correlate better with each other than with, say, sinusoidal functions, but consider what it demands of children psychologically: children would have to remember the changes they observed in each attribute over time and then – in memory – and then compare them with each other. It is not impossible that children do this. However, an alternative possibility is that children use the data they observe to infer the kind of function that describes the data. Consistent with this, children were able to match the observed functions simply to the abstract verbal descriptions. Of course, children don’t have access to scientific terms for these functions (linear, U-shaped, or periodic) but they themselves might be able to spontaneously describe observed changes as, for instance, "getting bigger"; "getting small and then getting big again" or "going back and forth again lots of times". Such intuitive descriptors may allow children to represent two functions as being members of the same kind. Although the computations that underlie such inferences may be sophisticated, such inferences may be easier than remembering and trying to correlate changes in many different kinds of events directly. These inferences can be understood as rational ones, under the hierarchical Bayesian view (Gopnik & Wellman, 2012; Tenenbaum, Kemp, Griffiths, & Goodman, 2011) which has proven useful for understanding the structure and dynamics of children's causal theory learning in other contexts. A sensitivity to abstract functions would be another instance of the "blessing of abstraction" (Goodman, Ullman, & Tenenbaum, 2011). Indeed, children may be more able to identify causal structure at an abstract level (with respect to common functional changes) than at more concrete levels. Here for instance, children were able to represent functional relationships between changes in the lights and flowers even though they presumably knew very little about the actual biophysics connecting light and flowers. Insofar as abstract functional relationships between causes and effects obtain in the real world, the ability to represent these relationships may help constrain and support more specific causal learning. This may contribute to humans' ability to infer causal relations accurately and reliably and from sparse and apparently impoverished patterns of data.

References

Ciccione, L., Sable-Meyer, M., Boissin, E., Josserand, M., Potier-Watkins, C., Caparos, S., & Dehaene, S. (2022). Graphicacy across age, education, and culture: a new tool to assess intuitive graphics skills. *bioRxiv*, 2022–10.

Goddu, M. K., Lombrozo, T., & Gopnik, A. (2020). Transformations and transfer: Preschool children understand abstract relations and reason analogically in a causal task. *Child development*, *91*(6), 1898–1915.

Goodman, N. D., Ullman, T. D., & Tenenbaum, J. B. (2011). Learning a theory of causality. *Psychological review*, *118*(1), 110.

Gopnik, A., & Meltzoff, A. N. (1997). *Words, thoughts, and theories*. Mit Press.

Gopnik, A., Sobel, D. M., Schulz, L. E., & Glymour, C. (2001). Causal learning mechanisms in very young children: two-, three-, and four-year-olds infer causal relations from patterns of variation and covariation. *Developmental psychology*, *37*(5), 620.

Gopnik, A., & Wellman, H. M. (2012). Reconstructing constructivism: Causal models, bayesian learning mechanisms, and the theory theory. *Psychological bulletin*, *138*(6), 1085.

Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences*, *107*(20), 9066–9071.

Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development*, *67*(6), 2797–2822.

Legare, C. H. (2012). Exploring explanation: Explaining inconsistent evidence informs exploratory, hypothesis-testing behavior in young children. *Child development*, *83*(1), 173–185.

Magid, R. W., Sheskin, M., & Schulz, L. E. (2015). Imagination and the generation of new ideas. *Cognitive Development*, *34*, 99–110.

McDermott, L. C., Rosenquist, M. L., & Van Zee, E. H. (1987). Student difficulties in connecting graphs and physics: Examples from kinematics. *American Journal of Physics*, *55*(6), 503–513.

Rodrigues, S. (1994). Data handling in the primary classroom: Children's perception of the purpose of graphs. *Research in Science Education*, *24*, 280–286.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*(5294), 1926–1928.

Schulz, L. E. (2012). The origins of inquiry: Inductive inference and exploration in early childhood. *Trends in cognitive sciences*, *16*(7), 382–389.

Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Developmental psychology*, *43*(4), 1045.

Scott, K., & Schulz, L. E. (2017). Lookit (part 1): A new online platform for developmental research. *Open Mind*, *1*(1), 4–14.

Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, *348*(6230), 91–94.

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *science*, *331*(6022), 1279–1285.

Weisberg, D. S., & Sobel, D. M. (2022). *Constructing science: Connecting causal reasoning to scientific thinking in young children*. MIT Press.

Xu, F., & Denison, S. (2009). Statistical inference and sensitivity to sampling in 11-month-old infants. *Cognition*, *112*(1), 97–104.