Goal-adaptiveness in children's cue-based information search

Andreas Domberg (domberg@mpib-berlin.mpg.de)
MPRG iSearch, Max Planck Institute for Human Development

Karla Koskuba (k.koskuba.12@ucl.ac.uk)
University College London

Anselm Rothe (rothe@mpib-berlin.mpg.de)
MPRG iSearch, Max Planck Institute for Human Development

Azzurra Ruggeri (ruggeri@mpib-berlin.mpg.de)
MPRG iSearch, Max Planck Institute for Human Development,
Technical University of Munich

Abstract

This paper investigates the emergence and development of children's ability to adapt their information search to different goals. In Study 1, 3- to 7-year-olds had to decide whether to study the arms or legs of two monsters to predict which would succeed at a throwing vs. jumping challenge. Children's ability to adaptively select the relevant piece of information and tailor their search to the given goal increased with age, surpassing chance level around 4;6. Study 2 investigated additional adaptation to distributions of, e.g., long arms in the search domain. Preliminary results confirm the observed developmental trend in search adaptiveness and effectiveness, suggesting an ability to tailor information search to the relevant distributions in the environment. These studies provide first insights into the development of adaptive information search given complex goals, deepening our understanding of this key aspect of learning, judgment and decision-making.
Keywords: information search; ecological learning; decision making; adaptiveness

Introduction

We all probably have some high school memories (or nightmares, depending on how popular and/or sporty you were) of those PE classes where you were supposed to pick or be picked by classmates and make teams for volleyball, soccer and so on. Putting together a good, successful team can be challenging: Who is good at this game? Should you go for particularly strong, fast or smart individuals—or maybe for good team players? What players should be assigned to which positions? What skills are important to be considered depends on the game, and to be effective in finding the best players one must tailor one's information search to a specific goal (e.g. make a volleyball rather than a soccer team) by querying the most relevant skills. In this project we investigate whether 3- to 10-year-old children adapt their information search to a given goal to maximize search effectiveness.

Are children effective when searching for information?

The developmental trajectory of information search has been first investigated using the information board paradigm, in which children are asked to make a choice between several options (e.g., bicycles) based on a set of cues (e.g., state of repair, gears or color) that they can look up. These studies have shown strong developmental improvements in search efficiency between 7 and 14 years, with younger children attending more often to irrelevant cues and searching in a less systematic and more exhaustive manner than older children (Davidson, 1991a, 1991b, 1996; Gregan-Paxton & John, 1997; Gregan-Paxton & John, 1995; Howse, Best, & Stone, 2003).

More recent work also suggests that children younger than 10 years do not adapt their search based on probabilistic or deterministic information about which cues are most reliable. For instance, in a series of studies by Betsch and colleagues, children had to predict where a treasure was hidden, and query informants whose reliabilities were quantified as proportion correct over past predictions (probabilistic reliability, Betsch, Lehmann, Jekel, Lindow, & Glöckner, 2018; Lang & Betsch, 2018; see also Lindow & Betsch, 2018 on deterministic reliability). Only from around age 10 did children start to select reliable cues more often, while still below adult performance. Introducing search costs or providing feedback did not motivate more efficient information search (Lindow & Betsch, 2019).

However, search adaptiveness and effectiveness seem to emerge much earlier in life. Infants already preferentially explore surprising events (Stahl & Feigenson, 2015) and direct attention to events that are neither too chaotic nor too homogeneous, so learning is optimally supported (Kidd, Piantadosi, & Aslin, 2012). A growing body of literature also suggests that preschoolers are more likely to explore when presented with confounded evidence—that is, when they are uncertain about the causal mechanism at work (e.g., Cook, Goodman, & Schulz, 2011; Schulz, Gopnik, & Glymour, 2007)—or when they face evidence that violates their prior beliefs (e.g., Bonawitz, van Schijndel, Friel, & Schulz, 2012; Legare, Gelman, & Wellman, 2010). Toddlers and preschoolers are already able to make informative interventions to disambiguate the causal structure of a system, both in an experimental setting and during spontaneous play (Schulz et al., 2007; Kushnir & Gopnik, 2005; Cook et al., 2011; Sim, Mahal, & Xu, 2017), and the efficiency of these interventions increases with age (McCormack, Bramley, Frosch, Patrick,

& Lagnado, 2016). Indeed, research suggests that children are ecological learners from a very early age, able to adapt their information search and hypothesis testing strategies to different characteristics of the presented task, such as its statistical structure (Ruggeri, Swaboda, Sim, & Gopnik, 2019; Ruggeri, Sim, & Xu, 2017; Ruggeri & Lombrozo, 2015).

The present paper

Building on the contradictory findings reviewed above, we investigate whether children tailor their information search to a given goal and the statistical structure of the environment to maximize efficiency. To this end, we developed a simpler version of the information board paradigm discussed above. In Study 1, 3- to 7-year-old children ($N = 105$) were asked to predict which of two monsters would succeed at a given game (throwing a ball into a bucket or jumping a hurdle). To do that, they had to decide whether to reveal the monsters' arms or legs, knowing that monsters needed long arms to hit the bucket, and long legs to jump the hurdle. In Study 2, for which data collection is still ongoing, we aimed to replicate and extend Study 1, using a similar task to investigate whether 3- to 10-year-old children would also adapt their search to the likelihood distribution of the features in the considered population, that is, how likely a given monster was to have long arms or legs. Children had to put together a team of two monsters with complementary skills (being good at throwing and jumping respectively), and were sequentially presented with monsters drawn from a population in which either both arms and legs had a 50% chance of being long (*uniform* condition), or one of them was only rarely long (*skewed* condition).

Study 1

Participants

Participants were 105 children (59 female, $M_{age} = 5.03$ years; $SD = 13$ months) recruited and tested at local museums in Berlin, Germany. IRB approval was obtained and parents gave informed consent for their children to participate before the study. Four additional children were tested but did not enter the analysis because they were distracted by siblings ($n = 2$), had watched another child being tested before ($n = 1$), or had language difficulties ($n = 1$).

Design and procedure

Children were presented with the pictures of eight monsters, two for each possible combination of short versus long arms and legs (see Figure 1). The experimenter told children that all monsters liked throwing balls and jumping hurdles, but that not all monsters were good at these activities. The experimental session consisted of two blocks, presented in counterbalanced order: The *jump* block and the *throw* block. Each block included a familiarization and a test phase with two trials.

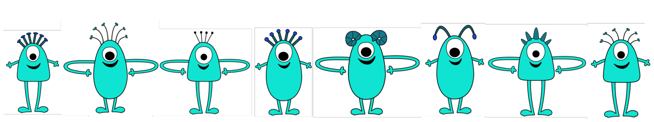

Figure 1: All combinations of long and short arms and legs.

***Jump* block** In the familiarization phase, children observed a short animation of two monsters, one with short and one with long legs, jumping the hurdle one after another (order counterbalanced, see Figure 2).

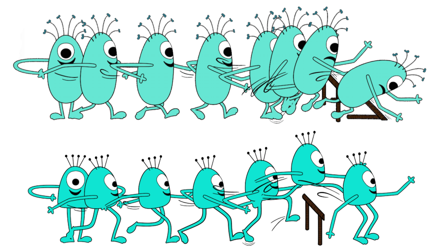

Figure 2: The two animations shown in the *jump* block.

In the animation, the monster with short legs failed, and children were told that this is "because its legs are not long enough. Monsters with short legs are always bad at this game." The monster with long legs succeeded and jumped the hurdle "thanks to its long legs. Monsters with long legs are always good at this game." Children were then presented with two new monsters, one with long and one with short legs, and were asked to select which monster they thought would manage to jump the hurdle. They were told that a correct selection would be rewarded with a sticker. A short animation of one monster failing and the other succeeding at jumping over the hurdle provided children with feedback to their selection.

After this familiarization, children received two test trials. On each, they saw a hurdle and two monsters with their arms and legs covered by grey bars, so their length was unknown (see Figure 3). Children were told that, to win another sticker, they had to find out which of the two monsters could jump over the hurdle. To make this decision, children could decide to reveal either the monsters' legs or arms.

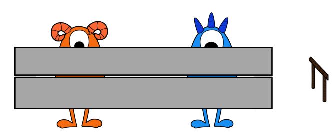

Figure 3: The display presented in the test phase.

***Throw* block** The structure of the *throw* block was identical to that of the *jump* block. However, here monsters had to throw a ball into a bucket and long arms were presented as the critical feature (see Figure 4).

Results

We analyzed children's information queries in the test trials with a Bayesian Generalized Linear Mixed Model with

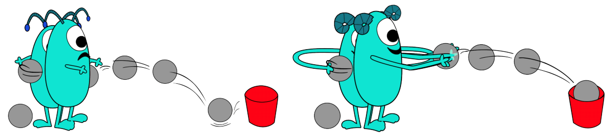

Figure 4: The two animations shown in the *throw* block.

a binomial link function. This model predicts the probability of querying the relevant (score 1) vs. irrelevant information (score 0) in each trial given the participant characteristics age and gender, and the trial characteristics of relevant body part (legs or arms), number of trial within block (first or second, to check for potential learning across trials of the same block), and block (first or second, to check for potential learning across blocks and hence, between the arm and leg domain). We added varying intercepts for each subject, since we collected several observations per child.

We compared the fit and predictive power of this model to alternative models using the Widely Applicable Information Criterion (WAIC; McElreath, 2016), which estimates the expected fit to new, out-of-sample observations. Alternative models included one in which the predictor age is missing, one in which trial number within block is missing, and one in which both are missing. The full model proved to be clearly best for predicting new observations in this experiment (difference to the next best model in WAIC: 13.7; SE of the difference: 7.6).

The model reveals that children's performance improved from chance level at age 3 to close to ceiling at age 7 (see Figure 5 (a)). Children's performance reliably exceeded chance level from around age 4.5 years, as indicated by the lower bound of the 95% Highest Posterior Density Interval (HPDI) of the model prediction of their mean score.

We also observed interesting patterns in children's responses. First, they had a higher average probability of responding correctly to the second trial within each block ($\chi^2(1,210) = 13.77; p < .001$), indicating a tendency either to learn, or to apply a strategy of switching upon failure (see Figure 5 (b)). Second, however, children did not seem to transfer their improvement from the first block (i.e., first body part) to the second. That is, they failed to generalize between task domains, shown by a drop in performance from trial 2 to 3 ($\chi^2(1,105) = 10.86; p < .001$).

More precisely, children's response patterns suggest a qualitative developmental shift from randomness to adaptive competence. To illustrate this transition, we categorized children's patterns as follows: *Rigid* responders did not ever change their responses, querying arms or legs four times; *Slow* adapters responded incorrectly to the very first trial, but were successful in all subsequent trials; and *Adaptive* responders were correct in all four trials. All other response patterns fell into the *Other* category. This category encompasses ten different observed response patterns.

As Figure 6 shows, the *Slow* adapting and *Adaptive* response patterns become more frequent with age, whereas the non-adaptive ones decline. Splitting the sample at the age median, the ratio of the two adaptive patterns versus all other ones compared across the two resulting age groups shows a clear development towards adaptiveness ($\chi^2(1,105) = 11.66, p < .001$).

When simulating new observations from the model and manipulating the covariates gender, body part and block in turn, we found no differences in the mean outcome scores. More precisely, in the middle of the age range, gender accounted for a .03 difference between female and male participants' mean scores (which, recall, can range from 0 to 4); body part (arms vs. legs averaging over block orders) accounted for .14 and block (first vs. second) for .10.

Discussion of Study 1

Results from Study 1 show that children's adaptiveness, i.e., their ability to select the most relevant piece of information and to tailor their search to the given goal, steadily increased with age, becoming reliably better than chance at age 4.5.

Their increasing selectivity and adaptiveness is also reflected in a developmental shift in children's response patterns, which between 3 and 5 years show signs of goal insensitivity. In this age range, 64% of the children either did not adapt their search behavior to the different given goals, or showed altogether unsystematic response patterns. From age 5, however, children already started to reliably adapt their strategies to achieve effectiveness. This result is in alignment with prior work finding that young children can identify and target the information they need (e.g., Schulz et al., 2007; Ruggeri & Lombrozo, 2015). The stark contrast with previous work on cue-based decision making, which found ineffective search even in much older children (Betsch et al., 2018; Lang & Betsch, 2018; Lindow & Betsch, 2018, 2019), could be explained by the much lower memory and other incidental demands made by our elementary version of the information board paradigm.

Also, note that the apparent increase in *Other* response patterns in 6-year-olds (nine children total, Figure 6) is dominated by the pattern "($\checkmark\checkmark | - \checkmark$)". These may plausibly be more cases of *Slow* adapting, although strictly speaking, this pattern is also compatible with two correct guesses and a shift upon error.

In Study 2 we modified the game developed in Study 1 to make it more computationally complex and ecologically valid. In particular, we added the need to adapt not only to a goal, but also to the statistical structure of the population from which participants sampled the monsters.

Study 2

Participants

Data collection for this study is ongoing. The current sample includes 26 participants (10 female, $M_{age} = 6.53$ years; $SD = 7.9$ months), tested at a local museum in Berlin, Germany. For the targeted age span, 4;0 to 10;0 years, we plan to

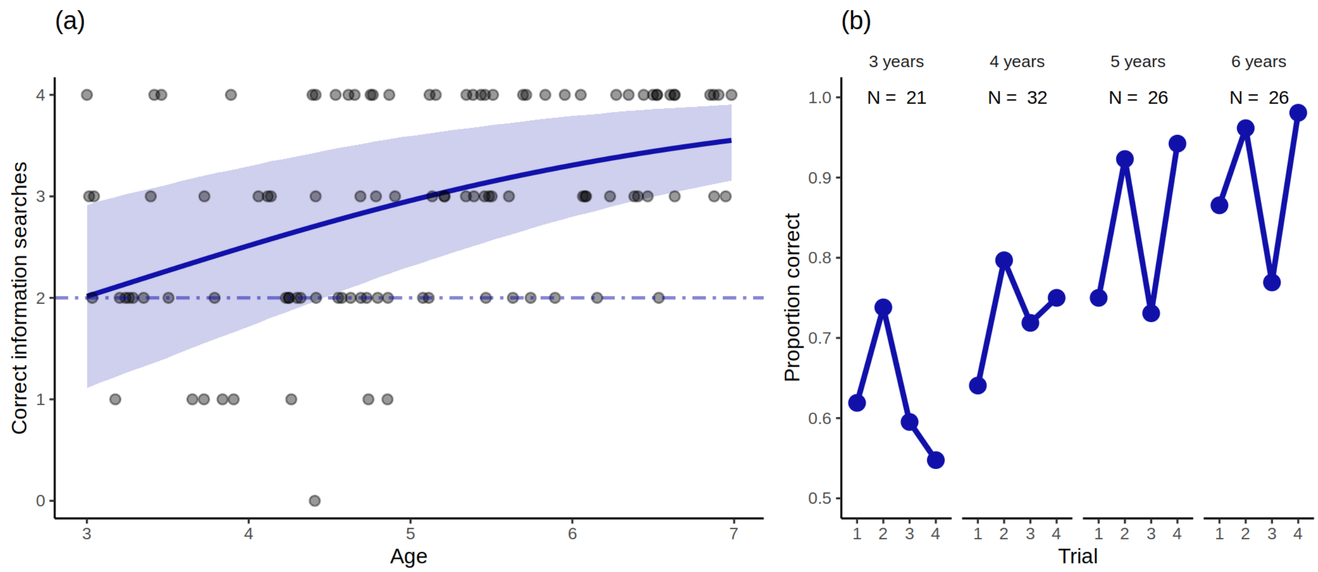

Figure 5: (a) Points: observed correct information searches out of four trials. Line: model predictions with 95% Highest Posterior Density Interval. These predictions are for any new random participant completing four trials under the same conditions as the 105 actual participants, i.e., including variation from age, block (first or second), trial within block (first or second) and body part (arms or legs). These predictions show an increase in identifying relevant information across development. (b) The observed proportion of correct responses in trial 1–4 by age group.

test 30 children for each year, projecting completion of data collection in May 2020.

IRB approval was obtained and parents gave informed consent before the study for their children to participate. One additional participant was excluded from analysis due to experimenter error. We introduced a condition manipulation and of the 26 total participants so far, we tested 15 in the Skewed condition and eleven in the Uniform condition.

Materials

The materials for Study 2 are outlined in Figure 8. Children dealt with a monster population that consisted of 16 (uniform condition) or 15 (skewed condition) physical monster cards. On each card, an individual monster was depicted, whose features (arms and legs) were hidden and could be revealed by pulling up and/or down two flaps (see Figure 8). A laptop with an attached but otherwise inert box ("the monster card reader") served to show the success or failure of monsters in the two activities upon putting a card on the box. A printed picture of a bucket and a hurdle showed two empty slots on each of which a card could be put.

The necessary size and feature distribution of the population in the two conditions was determined by computer simulations to ensure both noticeable distribution differences between conditions and a manageable total number of cards.

Design and procedure

The experimental session consisted of an activity familiarization phase, a population familiarization phase and the test phase.

Activity familiarization phase Children were familiarized with the functional relationship between body features and skills as in Study 1 (i.e., only monsters with longs arms are good at throwing; only monsters with long legs are good at jumping). For instance, focusing on arms, the experimenter showed one monster with short, and one with long arms and put them onto the "card reader", upon which the laptop showed an animation of the same monster being (un-)successful at throwing. In a comprehension check, children had to select a monster that would be successful. All children passed the comprehension check in the first trial. They were then given eight (uniform condition) or ten (skewed condition) small stickers. Children in the skewed condition received more stickers to compensate for the additional queries they typically required to finish the task.

Population familiarization phase Children were presented with the population of monster cards, each featuring short or long arms and short or long legs. Crucially, children were assigned to one of two conditions. In the Uniform condition, any monster picked at random had a 50% chance of exhibiting long or short arms or legs. In the Skewed condition, one of the features was long for only 3 out of the 15 monsters (e.g., only 3 monsters had long arms), while the other feature was uniformly distributed. To help children learn the distribution of features across the population, the experimenter asked them to sort the monsters in two stacks according to the length of their legs and arms (order counterbalanced). At the end of each sorting task, the experimenter pointed out that both stacks were the same size and long and short arms (or

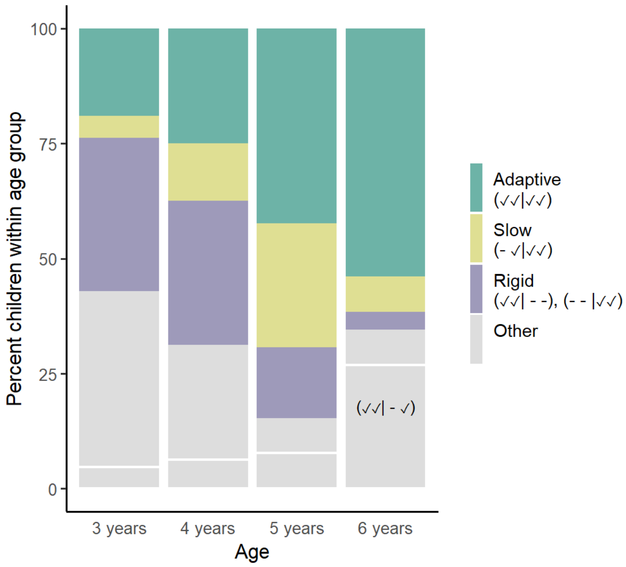

Figure 6: Response patterns binned into age groups and coded into four main categories. "Adaptive": full adherence to what information is necessary at each point. "Slow": adaptive from the second trial. "Rigid": choosing the same body part four times (two possible patterns). "Other": all other patterns. Note: we categorized "(✓✓| − ✓)" as Other because this pattern, while 75% correct, is also compatible with two guesses (as in Rigid) and one simple switch upon error.

legs) were present equally often (uniform condition), or that the piles differed in size, so that one feature was really hard to find (skewed condition).

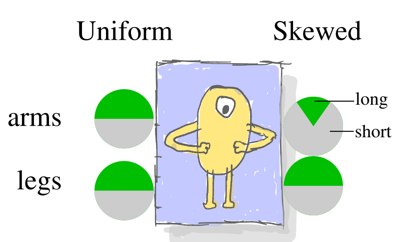

Figure 7: Illustration of the two distributions used in the two conditions. In Skewed, one feature is more frequent in its short form and the long form is hard to find.

Test phase Children were told that they had to pick monsters, one thrower and one jumper, for their team, and that if at the end of the game both their thrower and jumper were successful, they would be rewarded with one large sticker. Also, children learned that they could keep those small stickers that they did not spend in the game. Then the features on the monster cards were hidden by pushing the flaps in, the cards were shuffled, and laid out in a grid.

At every turn, children had a subset of the following four options: they could (1) query the monster's arms if they were hidden; (2) query the monster's legs if they were hidden; (3) assign the current monster card to their team as a thrower or

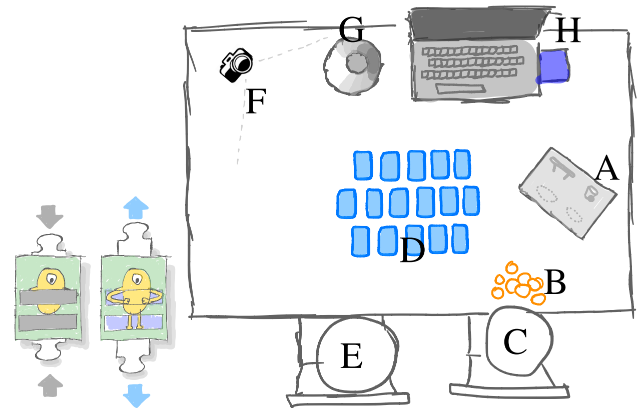

Figure 8: The setup of Study 2 seen from above: (A) the goal picture showing open slots next to a hurdle and bucket; (B) endowed small stickers; (C) child; (D) monster cards; (E) experimenter; (F) camera; (G) paper bin; (H) laptop with "card reader". Lower left: schema of a monster card.

jumper; or (4) discard the current monster card into the paper bin and draw a new one from the population. Options (1) and (2) required children to pay one small sticker each and put it into the paper bin (see Figure 8). Actions could not be undone (e.g., a discarded monster could not be considered again).

We recorded each session on camera and coded children's decisions from the videos.

Results

First, we analyzed children's sensitivity for relevance, which conceptually replicates what we tested in Study 1. In this version of the game, an irrelevant feature is one related to the activity for which a monster has already been chosen, e.g., querying arms when the thrower position has already been filled. Across conditions, children targeted irrelevant cues only in a small proportion of those situations in which it was possible (11.8%).

Second, we analysed whether children adapted their search to the likelihood distributions of the features in the population they were presented with. If children are sensitive and tailor their search to the *base rate* of the given features, they should approach the search task differently across the two conditions. In particular, upon learning in the Skewed condition that, for example, long arms are rare, children's information search should optimally focus on arms first. Indeed, 12 out of 15 children in the Skewed condition first queried the arms of the first considered monster, compared to 3 out of 11 children in the Uniform condition ($\chi^2(1,26) = 7.23, p = .007$). We also calculated for each child the frequency of cards on which they queried the arms first, across all those cards inspected while both positions still had to be filled. This analysis suggests that children's prioritization of the rare cue weakens after the first card, with an average 46% of total first queries per card in the Skewed condition targeted at arms, compared to 36% in the Uniform condition. Children in the Skewed condition ended up assigning their jumper before the thrower nearly consistently (with one exception in 15 children). This way, they prevented situations in which the rare feature would at the same time be irrelevant, thus making analysis of an interaction between cue relevance and likelihood impossible. Of all 26 children, 24 finished the game with correct choices of monsters in their teams, one child assigned a short-armed thrower, one a short-legged jumper.

Discussion of Study 2

In Study 2, we modified the task to replicate and extend the results from Study 1. In particular, we further investigated whether children adapt their information search to the likelihood distribution of the features in the considered population, e.g., how likely a given monster was to have long arms.

The initial data collected suggest a replication of the results of Study 1: Children showed a strong preference for relevant over irrelevant cues. In addition, they responded adaptively to a skewed distribution of the feature of interest in the environment.

General Discussion

The present studies introduce a novel approach to investigate adaptive pre-decisional information search in early childhood. Thanks to our simplified paradigm, we were able to show for the first time that children as young as 5 are already able to dynamically adapt their cue search to the given goal: They are more likely to query the features that are most relevant for the decision to be made, additionally taking into account the likelihood distributions of such features across the population considered. More research is needed to fully trace the emergence and developmental trajectory of goal-adaptive search. First, we cannot exclude that an even simpler paradigm would be able to capture an even earlier emergence of goal adaptiveness, for example operationalized as selective attention or gaze allocation in infants.

Second, so far research has investigated adults' and children's ability to adapt their search to different goals focusing on only a handful of goal-related features. However, many other characteristics of real-world goals can impact the effectiveness of our search, more or less directly. In a new line of research, we are currently exploring how information search is impacted, more or less explicitly or intentionally, by social goals, and how this impact can differ developmentally. On the one hand, are people more effective in their search when the risk of making the wrong move is high—when they need to find the cheater in their team or the mole in their company, or when trying to identify which kind of food is poisonous? On the other hand, what do people do when search effectiveness conflicts with social goals—would they be willing to *give up* on effectiveness, say, not to be impolite?

Finally, future research should try to go beyond identifying developmental shifts in goal-adaptive search, focusing on those factors driving such differences, such as children's ability to reason with numbers and proportions, or their inhibitory control.

References

Betsch, T., Lehmann, A., Jekel, M., Lindow, S., & Glöckner, A. (2018). Children's application of decision strategies in a compensatory environment. *Judgment and Decision Making*, 15.

Bonawitz, E. B., van Schijndel, T. J. P., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology*, *64*(4), 215–234. doi: 10.1016/j.cogpsych.2011.12.002

Cook, C., Goodman, N. D., & Schulz, L. E. (2011). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*, *120*(3), 341–349. doi: 10.1016/j.cognition.2011.03.003

Davidson, D. (1991a). Children's decision-making examined with an information-board procedure. *Cognitive Development*, *6*(1), 77–90.

Davidson, D. (1991b). Developmental differences in children's search of predecisional information. *Journal of Experimental Child Psychology*, *52*(2), 239–255.

Davidson, D. (1996). The effects of decision characteristics on children's selective search of predecisional information. *Acta Psychologica*, *92*(3), 263–281.

Gregan-Paxton, J., & John, D. R. (1995). Are young children adaptive decision makers? a study of age differences in information search behavior. *Journal of Consumer Research*, *21*(4), 567–580. doi: 10.1086/209419

Gregan-Paxton, J., & John, D. (1997). The emergence of adaptive decision making in children. *Journal of Consumer Research*, *24*(1), 43–56. doi: 10.1086/209492

Howse, R. B., Best, D. L., & Stone, E. R. (2003). Children's decision making: the effects of training, reinforcement, and memory aids. *Cognitive Development*, *18*(2), 247–268. doi: 10.1016/S0885-2014(03)00023-6

Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PloS one*, *7*(5).

Kushnir, T., & Gopnik, A. (2005). Young children infer causal strength from probabilities and interventions. *Psychological science*, *16*(9), 678–683.

Lang, A., & Betsch, T. (2018). Children's neglect of probabilities in decision making with and without feedback. *Frontiers in Psychology*, *9*, 191.

Legare, C. H., Gelman, S. A., & Wellman, H. M. (2010). Inconsistency with prior knowledge triggers children's causal explanatory reasoning. *Child Development*, *81*(3), 929–944. doi: 10.1111/j.1467-8624.2010.01443.x

Lindow, S., & Betsch, T. (2018). Child decision-making: On the burden of predecisional information search. *Journal of Cognition and Development*, *19*(2), 137–164.

Lindow, S., & Betsch, T. (2019). Children's adaptive decision making and the costs of information search. *Journal of Applied Developmental Psychology*, *60*, 24–34.

McCormack, T., Bramley, N., Frosch, C., Patrick, F., & Lagnado, D. (2016). Children's use of interventions to learn causal structure. *Journal of Experimental Child Psychology*, *141*, 1–22. doi: 10.1016/j.jecp.2015.06.017

McElreath, R. (2016). *Statistical rethinking: A bayesian course with examples in r and stan* (1st ed.). Chapman & Hall/Crc Texts in Statistical Science.

Ruggeri, A., & Lombrozo, T. (2015). Children adapt their questions to achieve efficient search. *Cognition*, *143*, 203–216. doi: 10.1016/j.cognition.2015.07.004

Ruggeri, A., Sim, Z. L., & Xu, F. (2017). "why is toma late to school again?" preschoolers identify the most informative questions. *Developmental Psychology*, *53*, 1620–1632. doi: 10.1037/dev0000340

Ruggeri, A., Swaboda, N., Sim, Z. L., & Gopnik, A. (2019). Shake it baby, but only when needed: Preschoolers adapt their exploratory strategies to the information structure of the task. *Cognition*, *193*, 104013. doi: 10.1016/j.cognition.2019.104013

Schulz, L. E., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science*, *10*(3), 322–332. doi: 10.1111/j.1467-7687.2007.00587.x

Sim, Z. L., Mahal, K. K., & Xu, F. (2017). Is play better than direct instruction? learning about causal systems through play. In *Cogsci*.

Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, *348*(6230), 91–94. doi: 10.1126/science.aaa3799