UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Efficient Detectives in the Sandbox: Children Demonstrate Adaptive Information-Search Strategies in a Novel Spatial Search Game

Permalink

https://escholarship.org/uc/item/5nh4m0v3

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors

Chai, Kai-Xuan
Stanciu, Oana
Ruggeri, Azzurra

Publication Date

2023

Peer reviewed

Efficient Detectives in the Sandbox: Children Demonstrate Adaptive Information-Search Strategies in a Novel Spatial Search Game

Kai-Xuan Chai [1] **(chai@mpib-berlin.mpg.de)**

Oana Stanciu [2] **(stanciuo@ceu.edu)**

Azzurra Ruggeri [1,2,3] **(ruggeri@mpib-berlin.mpg.de)**
[1] MPRG iSearch, Max Planck Institute for Human Development, Lentzeallee 94, 14195, Berlin, Germany
[2] Department of Cognitive Science, Central European University, Quellenstraße 51, 1100 Vienna, Austria
[3] TUM School of Social Sciences and Technology, Technical University of Munich, Arcisstraße 21, 80333 Munich, Germany

Abstract

Recent studies suggest that children are *ecological active learners* who recognize and exploit the ecology of their learning environment (Ruggeri, 2022). However, when assessed by verbal tasks such as the 20-questions game, systematic search only matures around age 7. The current study examined if even young children can adapt their information-search strategies in a developmentally-appropriate task requiring minimal verbal or conceptual abstraction skills. Three to 7-year-olds ($N = 76$, $M = 5.7$ years) played a search game with a structure analogous to the 20-questions game. We manipulated whether children received predecisional cues about the past location of the solution or not, across two search phases, and further varied whether the cues follow a *Uniform* or *Skewed* distribution. Children adapted their information-search strategies as predicted: They followed a constraint-seeking strategy in the absence of cues, and only switched to hypothesis-scanning when exposed to the Skewed cues.

Keywords: information search; active learning; statistical structure; child-friendly paradigm

Introduction

Children are unrelenting active explorers from very early in life, which undoubtedly contributes to their impressive learning abilities (Alvarez & Booth, 2014). They are not just curious and adventurous explorers, though—their exploration is selective, meaningful and purposeful. Already by 11 months of age, infants prefer to explore surprising events (Stahl & Feigenson, 2015), and a growing body of work has shown that very young children are more likely to explore when presented with confounded (L. E. Schulz & Bonawitz, 2007) or unexpected evidence (Bonawitz, van Schijndel, Friel, & Schulz, 2012). They seek out uncertainty reduction more eagerly than adults (E. Schulz, Wu, Ruggeri, & Meder, 2019; Meder, Wu, Schulz, & Ruggeri, 2021) and are sensitive to (Ruggeri, Swaboda, Sim, & Gopnik, 2019; Ruggeri, Sim, & Xu, 2017) and motivated by the potential information gain of different actions (Jirout & Klahr, 2012).

The development of information search competencies has been formally investigated using the 20-questions game (e.g., Mosher & Hornsby, 1966; Herwig, 1982; Nelson, Divjak, Gudmundsdottir, Martignon, & Meder, 2014; Ruggeri & Feufel, 2015; Ruggeri, Lombrozo, Griffiths, & Xu, 2016; Meder, Nelson, Jones, & Ruggeri, 2019; for a review, see Jones, Swaboda, & Ruggeri, 2020), in which children are tasked to identify the correct answer among a set of given hypotheses by asking as few yes–no questions as possible

(e.g., "why is Toma late to school?"). This game can be solved using either a *hypothesis-scanning* strategy, examining the individual hypothesis one by one (e.g., "Is it because Toma's bike was broken?"), or a *constraint-seeking* strategy, targeting multiple hypotheses at the same time (e.g., "Is it because Toma couldn't find something on the way to school [e.g., books, shoes or jacket]?"). In situations where all hypotheses are equally likely to be correct (i.e., a *uniform* distribution), employing a constraint-seeking strategy is more *efficient*, as it guarantees reaching the solution in a minimal number of steps. However, in situations where one or a few of the considered hypotheses are more likely to be correct (i.e., a *skewed* distribution; e.g., Toma often wakes up late), employing a hypothesis-scanning strategy that focuses on the most likely candidate(s) is more *efficient and effective* (see Ruggeri et al., 2017, 2019; Meder et al., 2019). In this sense, different search strategies are not a-priori better than others. They are like tools in a toolbox: No strategy is suitable for all problems, just as different tasks call for different tools.

Recent developmental work demonstrates that children are *ecological active learners*—despite their limited cognitive and computational resources, they can recognize the statistical structure of the environment and adapt their information-search and learning strategies accordingly (Ruggeri, 2022). For instance, Ruggeri and Lombrozo (2015) showed that 7-year-olds were more likely to ask hypothesis-scanning questions when a few of the given hypotheses were described as very likely to be correct. However, when all hypotheses were presented as equally likely, they predominantly asked constraint-seeking questions. In paradigms in which they do not have to formulate questions from scratch, even 5-year-olds demonstrate the ability to adapt their reliance on different kinds of questions to the statistical structure of the problem presented (Ruggeri et al., 2017).

By implementing a nonverbal version of this search paradigm, Ruggeri et al. (2019) demonstrated that even younger children are ecological active learners. In their task, 3- to 5-year-olds had to find an object hidden in one of four small boxes, contained in two larger boxes. Preschoolers were allowed to open *only one large box* but could shake one or both large boxes first if they wanted to. Crucially, before this test, the children learned either that the target was equally likely to be found in any of the four small boxes (uniform condition) or that it was most likely to be found in one partic-

702

ular small box on the most left/right side (skewed condition). Results showed that children as young as 3 successfully tailored their exploratory actions to the different likelihood distributions: Compared with children in the skewed condition, who had a strong intuition as to where the target would be hidden, children in the uniform condition were more likely to shake a large box first, in order to hear which large box contained the small box with the target without risking opening the wrong large box. Taken together, these results suggest that children as young as 3 years of age already have the computational foundations to support adaptive and efficient information search.

The present study

The current study investigates the emergence and early developmental trajectory of children's information-search strategies using a novel child-friendly and age-appropriate paradigm. On the one hand, compared to previous question-asking studies with children older than 5 years (Ruggeri et al., 2017, 2016; Ruggeri & Lombrozo, 2015), our paradigm minimized vocabulary and conceptual demands (for similar efforts, see also Chai, Xu, Swaboda, & Ruggeri, 2023; Ruggeri et al., 2019; Swaboda, Meder, & Ruggeri, 2022). On the other hand, compared to the box-search study by Ruggeri et al. (2019), this task is more complex, as it involves a much larger hypothesis space (16 potential locations, compared to the 4 from Ruggeri et al., 2019). Finally, this is the first study, to our knowledge, providing a within-participant measure of information-search adaptiveness.

In our game, children were tasked to find a treasure hidden in one of 16 identical buckets filled with sand. To help them locate the treasure, children could use scanners of different sizes, which made a sound if placed on a subset of buckets containing the target bucket. Crucially, the game was designed to manipulate across rounds children's prior beliefs about the likelihood distribution of the treasure across the 16 buckets. The first round (Baseline search phase) served to assess their baseline information-search performance. In a second round (Condition search phase), children were provided with cues suggesting that the treasure could be found in any of the buckets (Uniform distribution) or in one specific bucket (Skewed distribution; see Design).

We hypothesized that children would adapt their information search strategies to the likelihood distribution of the treasure. In particular, we predicted that 1) there will be a between-participants difference in the Condition search phase, as children will be more likely to start by selecting smaller scanners for the Skewed distribution (compared to the Uniform distribution), targeting the "cued" location of the treasure, and will therefore need fewer scanners to reach the solution (since the reward was indeed at this location); and 2) children who were cued with the Uniform distribution in the Condition search phase will consider a uniform prior and will, therefore, employ the same strategy as in the Baseline search phase, which was devoid of any information about the target location. In this sense, we should observe a within-

subject difference between the strategy used in the Baseline versus Condition search phase for children assigned to the Uniform distribution group, but not for children assigned to the Skewed distribution group.

Methods

Participants

Participants were 76 children (38 female; M = 5.7 years, SD = 1.17 years; Range: 3.1–7.9 years). They were recruited by random approach at a local beach resort in Livorno, Italy. An additional 11 children were excluded from the analyses due to experimental error (n = 9) or because they could not complete the task (n = 2). Additionally, data from only one phase was excluded for 6 other children, due to experimental error (Baseline search phase - 2; Condition search phase - 4). The sample size was determined by conducting a-priori power calculations via simulation (based on estimates from a previous search experiment with a similar design) for the planned statistical test concerning the first scanner choice. The most conservative estimate indicates an overall sample of 70 children to detect the estimated effect size (Cohen's h = 0.6) with 85% power using linear mixed-effect models with a 0.05 criterion for statistical significance. A slightly larger sample was collected (∼10%) in order to compensate for potential exclusions.

Written informed consent was obtained from parents prior to participation. Children were also asked for verbal consent and received a small present to thank them for their participation. The study was approved by the Ethics Committee of the Max Planck Institute for Human Development, Berlin, Germany (protocol: Sandbox).

Design

Children were tasked to find a treasure (i.e., glittery stickers in a seashell-shaped box) hidden in one of 16 plastic buckets (diameter = 20 cm, height = 25 cm), filled with sand, arranged in a 4-by-4 grid (see Figure 1a). Children were instructed to use special "scanners" (pieces of cardboard with handles, covered with gray duck tape) to detect the treasure. There were four sizes of scanners corresponding to whether they covered one bucket (Scanner1), two (Scanner2), four (Scanner4) or eight (Scanner8) buckets simultaneously (see Figure 1b). [1] Importantly, in order to encourage the use of efficient search strategies, children were told that they could dig in *one bucket only* so they should be really sure about the treasure location before digging and that each scanner use was costly, i.e, children had to "pay" a sticker to use a scanner (see below). X-shaped cardboard pieces covered with red tape were used as the memory aid, to help children remember when the scanner excluded some buckets as potential targets.

The experimental session consisted of a training phase intended to demonstrate how the scanners work, a *Baseline*

---

[1] The number of buckets covered with a scanner could be equal to or less than the size of the scanner. For example, children could use Scanner4 to test only two buckets but could not use it to test six buckets.
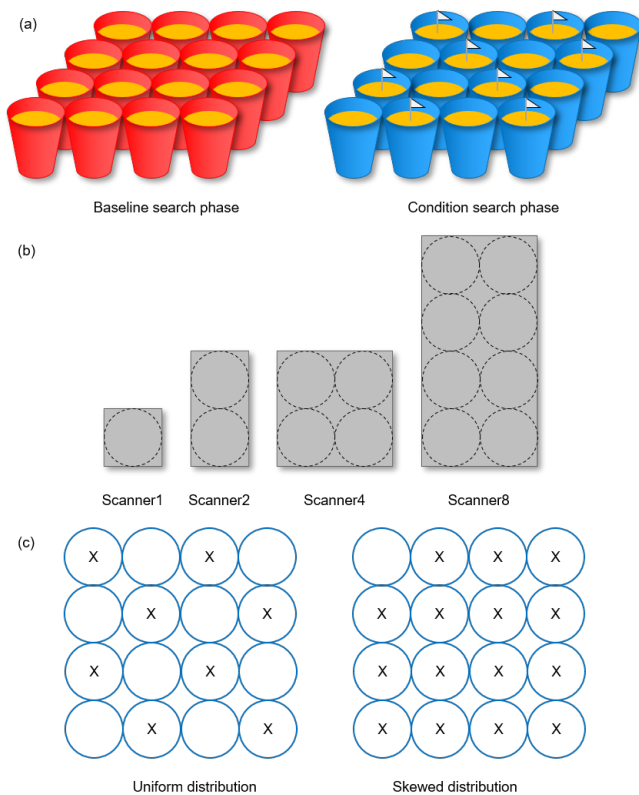
Figure 1: *Schematics of the Experimental Materials and Design.* (a) Sixteen plastic buckets filled with sand were arranged in a 4-by-4 grid, one of which contains a treasure. The red buckets (without flags) are used first in the Baseline search phase; the blue buckets (with flags) are used next in the Condition search phase. (b) Children were instructed to use special scanners (called "radars" in the game) to detect the treasure. Note that the dashed circles here depicted on the scanners for illustration purposes were not actually visible. (c) In the Condition search phase, children were randomly assigned to one of two groups, in which they were cued about the potential location of the treasure. Flags followed either a *Uniform* distribution over buckets (the flags were evenly scattered over half of these 16 buckets) or a *Skewed* distribution (all buckets contain a flag except for one corner bucket).

search phase, and a *Condition* search phase. In the Baseline search phase, children did not receive any information about the location of the treasure. Critically, in the Condition search phase, across two between-subject conditions, children received evidence about the treasure location corresponding to either a sparse uniform probability over buckets (*Uniform* distribution) or a clustered skewed probability that highlights one specific bucket (*Skewed* distribution; see Figure 1c). [2] The session was videotaped for coding purposes.

Procedure

Training phase The experimenter showed the children a golden coin and hid it in one of four (additional) training buckets, arranged in a 2-by-2 grid. The experimenter then introduced the children to two scanners (Scanner1 and Scanner2), saying "Look, here I have two radars which are really useful for finding treasures hidden in the sand! I will show you how they work." The experimenter picked Scanner1/Scanner2 (order counterbalanced) and said, "This radar is able to check only one bucket/two buckets at a time. It works like this: You take the radar, put it on the bucket(s) you want to check, then wait for a few seconds. If the treasure is in (one of) the bucket(s) covered by the scanner, it makes a sound, otherwise it does not make any sound." There was one demonstration with positive feedback and one with negative feedback, saying "Do you remember where I put the coin? Over here, right? If I put the radar here, listen, it makes a sound! But if I put it here (over the bucket[s] without the treasure), nothing happens. Do you want to try?". The experimenter then prompted the children to try a couple of times, and then asked the children to dig into the bucket where *they were sure* the treasure was hidden, saying "Now, when you are sure you know where the treasure is hidden, you can select a bucket and dig to get it. But remember, you can only dig in one bucket, so you have to be really sure! Got it?"

Baseline search phase The experimenter arranged the 16 buckets as shown in Figure 1a and placed the four scanners on the side in an increasing or decreasing size order (counterbalanced). The experimenter explained, "We are looking for a treasure. If you find it, it's yours! The treasure is hidden in one of these buckets. But there are so many of them! And we can only dig in one bucket, just one. So, we really have to be sure where the treasure is hidden. The good thing is that we have radars! Using the radars is expensive. Every time you want to use a radar you have to pay with one of these stickers [the experimenter had previously given children 10 stickers]. You can bring home all the stickers you do not use! Remember, you can only dig in one bucket, and if you do not find the treasure, you lose the game." Then the experimenter asked, "Which radar do you want to use [first]? Go ahead and take it, and give me a sticker to use it. Now place it on the bucket(s) you want to check."

---

[2]We used the term "uniform" here, as the flags were evenly distributed over the buckets grid. In that sense, the location of the treasure could be regarded to follow a (quasi-)uniform distribution.

After the children placed the scanner, the experimenter gave feedback accordingly (note that the target location was pseudo-randomly selected). When the feedback was negative, the experimenter said, "There was no sound, so the treasure is not in this bucket/any of these buckets. To remember that, I am going to place some red crosses on them [marked the ruled-out buckets using the X-shaped cardboard pieces]. The treasure must be in one of the other buckets." When the feedback was positive, the experimenter said, "There was a sound, so the treasure is in this bucket/one of these buckets! To remember that, I am gonna place some red crosses on all the other buckets [marked the ruled-out buckets using the X-shaped cardboard pieces]. The treasure must be in one of these buckets." As long as more than one bucket was left as the potential target, children were prompted to use another radar. Only when one bucket was left as a viable target (i.e., if children received positive feedback after using Scanner1, or if all other buckets had been eliminated as potential targets via negative feedback in previous rounds), the experimenter asked the children to dig out the treasure and then place a small flag on the bucket, "It must be in this bucket! Ready to dig for it? Yeah! Here is the treasure! Let's fill it up again and place a flag to indicate that you found the treasure in this bucket."

Condition search phase Children were then invited to play another round, "The next treasure is even better, and if you find it you can take it home!" The procedure of the *Condition* search phase was identical to the *Baseline* search phase, with one crucial difference. Before the search began, the experimenter drew the children's attention to the flags placed in some of the buckets. Children were randomly assigned to one of two experimental groups: In the *Uniform* distribution group, there were 8 flags evenly distributed across the 16 buckets as shown in Figure 1c (or the complementary configuration; counterbalanced across participants); In the *Skewed* distribution group, there were 15 flags placed in each bucket except one bucket in the corner (counterbalanced across participants). Then the experimenter told the children "Do you see these flags? It means that one child yesterday found the treasure here, another child found it here, [...], and yet another child found it here. [pointing to the flagged locations]" Rewards were always hidden in unflagged corner buckets, in both Condition distribution groups, in order to be consistent with the cover story presented to children, and to control for differences across conditions.

Results

To capture children's overall performance, we examined children's first scanner choices and measured the *effectiveness* of their search by counting how many scanners children required to reach the solution. We also quantified children's search strategies through the average Expected Information Gain (EIG, bits) of their search choices (e.g., Nelson et al., 2014; Ruggeri et al., 2016; Meder et al., 2019) considering a uniform probability over all possible treasure locations. The

EIG measures how much uncertainty, quantified via entropy, each search choice is expected to reduce a priori (Shannon, 1948; Oaksford & Chater, 1994), i.e., to what extent a query narrows down the hypothesis space to converge on the solution. The higher the EIG value of a search choice, the more efficient it is. Specifically, in the Baseline search phase and the Uniform distribution group of the Condition search phase, the most efficient choice is the largest scanner which splits the hypothesis space in half ($EIG = 1$ bit). In the Skewed distribution group, since EIG is calculated with respect to a uniform prior, equivalent EIG values would indicate that children did not adapt their search strategies to the change in the distribution, and lower values would suggest that children's strategies are ill-suited to a Uniform distribution.

Children's First Choice

The majority of children in the Baseline search phase (59.5%) selected the largest scanner to start the search. This proportion dropped to 33.3% in the Condition search phase, as the most frequently chosen scanner was Scanner1 for the Skewed distribution (50%) and Scanner8 for the Uniform distribution (35.3%; see Figure 2).

We fitted a mixed logistic regression using age (in years, not rounded), search phase (Baseline vs. Condition) and distribution (Uniform vs. Skewed) to predict the likelihood that children chose Scanner1 first. The model accounted for a significant amount of variance in the likelihood of first choosing Scanner1 (likelihood ratio [LR] $\chi^2(3) = 10.62$, $p = .014$, $R^2 = .35$). Crucially, controlling for age, which did not predict Scanner1 choice significantly, ($p = .11$), we found a significant main effect of search phase, LR $\chi^2(1) = 6.45$, $p = .011$, showing that the odds of choosing Scanner1 first were 2.91 times larger in the Condition search phase (38.9%) than in the Baseline search phase (21.6%; 95% CI of OR [1.21, 7.02], $p = .017$). This corresponds to the predicted switch from a constraint-seeking strategy to a hypothesis-scanning strategy. Further, consistent with our predictions, in the Condition search phase, when under the Skewed distribution, 57.9% children started the search by testing the unflagged bucket, out of which 45.5% tested it with the smallest, one-bucket scanner.

Effectiveness of information search

Overall, children in the Baseline search phase used on average 5.0 scanners ($n = 74$, $SD = 1.77$). This was significantly more than what is predicted by optimal binary search (i.e., 4 scanners), $t(73) = 5.00$, $p < .001$, Cohen's $d = 0.58$. In the Condition search phase, children used 5.5 scanners for the Uniform distribution ($n = 34$, $SD = 2.71$), which is significantly more than required ($t(33) = 3.29$, $p = .002$, Cohen's $d = 0.56$), and 3.4 scanners for the Skewed distribution ($n = 38$, $SD = 2.42$).

We fitted a Poisson regression to predict the number of scanners (treated as count data) using search phase (Baseline vs. Condition), distribution (Uniform vs. Skewed), their
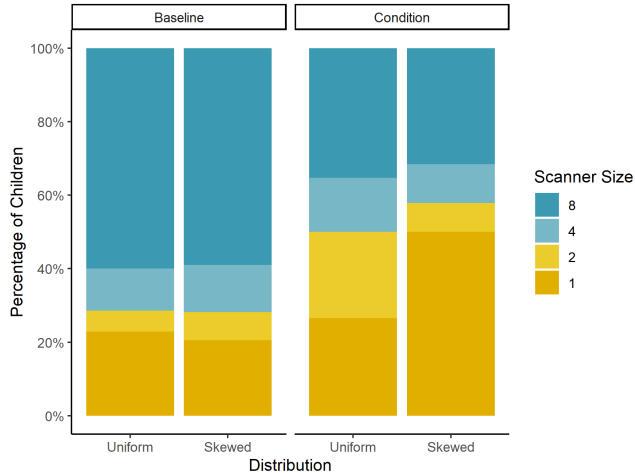
Figure 2: *Participants' First Scanner Choices as a Function of Search Phase and Distribution.*
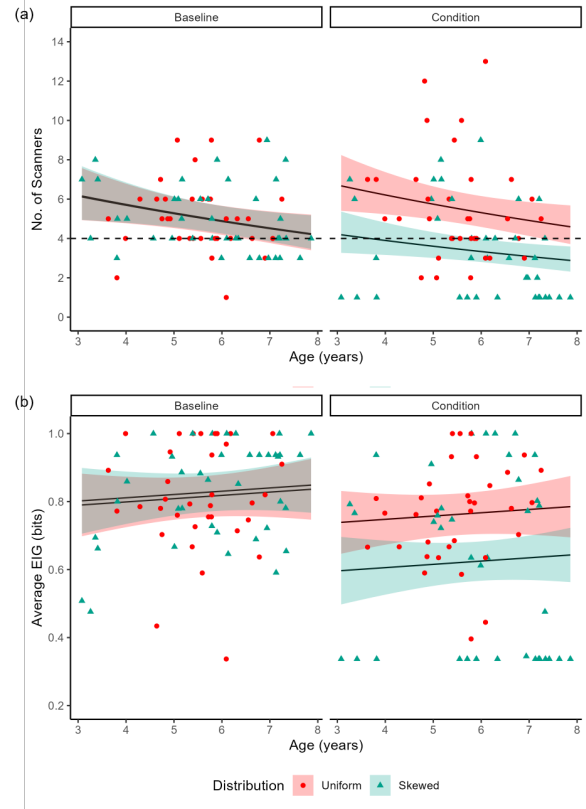

Figure 3: *Children's Performance as a Function of Age, Search Phase and Distribution.* (a) The number of scanners was fitted by a Poisson model, and (b) the average expected information gain was fitted by a linear mixed model. The shaded areas represent 95% confidence bands. The dashed horizontal lines represent the performance of an optimal agent performing a binary search (assuming a uniform prior).

interaction, and age (in years, not rounded). [3] The model accounted for a significant amount of variance in the number of scanners (LR $\chi^2(4) = 27.72$, $p < .001$, Nagelkerke $R^2 = .25$). With an increase of one year of age, the predicted number of scanners children required to solve our task significantly decreased by 8% (95% CI of rate ratio [0.87, 0.99], $p = .019$, see Figure 3a). Crucially, controlling for age, we found a significant interaction of search phase and distribution, LR $\chi^2(1) = 9.44$, $p = .002$, indicating that the within-participant search phase difference was much larger for children in the Skewed distribution. In particular, the decrease in the amounts of scanners that children needed to solve the Skewed distribution task relative to the Baseline condition task was 38% steeper than that for children assigned to the Uniform distribution (95% CI of rate ratio [0.46, 0.84], $p = .002$; see Figure 3a). Indeed, children needed roughly the same number of scanners for the Uniform distribution of the Condition phase and in the Baseline phase (9% difference; 95% CI of rate ratio [0.89, 1.34], $p = .401$). However, significantly fewer scanners were used in the Condition phase relative to the Baseline phase for children assigned to the Skewed distribution (32% fewer; 95% CI of rate ratio [0.54, 0.85], $p = .001$).

A direct comparison of children's scanner usage in the Condition search phase revealed that children solved the task using significantly fewer scanners for the Skewed distribution compared to the Uniform distribution (36% fewer; 95% CI of the rate ratio [0.51, 0.80], $p < .001$).

Information search strategies

Children's average EIG in the Baseline search phase was 0.82 ($n = 74$, $SD = 0.15$). In the Condition search phase, children achieved an average EIG of 0.76 in the Uniform distribution

---

[3] We first fitted a mixed effects Poisson model with random intercepts for each child, but found no significant random intercept variance.

($n = 34$, $SD = 0.15$), and 0.62 in the Skewed distribution ($n = 38$, $SD = 0.26$).

We fitted a linear mixed model using search phase (Baseline vs. Condition), distribution (Uniform vs. Skewed), their interaction and age (in years, not rounded) to predict each child's average EIG, with random intercepts for each child. The model accounted for a significant amount of variance in the average EIG of children's search (LR $\chi^2(4) = 31.53$, $p < .001$, $R^2 = .36$). Performance did not improve with age, $p = .51$. Crucially, controlling for age, we found a significant interaction of search phase and distribution, $\chi^2(1) = 8.07$, $p = .005$. The within-participant search phase difference in average EIG was larger with 0.15 bits for children assigned to the Skewed distribution compared to the Uniform distribution (95% CI [-0.26, -0.05], $p = .004$; see Figure 3b). This was due to the fact that children achieved roughly the same average EIG for the Condition phase with Uniform distribution and in the Baseline phase (0.05 difference; 95% CI [-0.13, 0.03], $p = .187$). At the same time, there was a significant

within-participant effect for children in the Skewed condition (0.21 difference; 95% CI [-0.30, -0.11], $p < .001$).

A direct comparison of children's average EIG between the two distribution groups in the Condition search phase revealed that children solved the task with significantly lower average EIG for the Skewed distribution compared to the Uniform distribution (0.14 lower; 95% CI [-0.24, -0.03], $p < .001$). [4]

Discussion

The present study explored the effectiveness and efficiency of 3- to 7-year-olds' information-search strategies using a child-friendly and age-appropriate search game, which parallels the task structure of the traditional 20-question game, with minimized verbal and conceptual demands. We examined the adaptiveness of children's information-search strategies, that is, whether and how they tailor their strategies to the statistical structure of the given task.

Overall, we found that children were suboptimal at solving the search task, in terms of both the effectiveness and the efficiency of their search strategies. These results are consistent with the pessimistic conclusion of previous work, namely that children do not begin to systematically generate effective questions until age 7 (e.g., Mosher & Hornsby, 1966; Herwig, 1982; Legare, Mills, Souza, Plummer, & Yasskin, 2013; Ruggeri, Walker, Lombrozo, & Gopnik, 2021) and do not demonstrate mature inquiry patterns until around age 10 (e.g., Ruggeri & Feufel, 2015; Ruggeri et al., 2016; Kachergis, Rhodes, & Gureckis, 2017). However, it is crucial to keep in mind that prior work indicates that even adults in similar information-search tasks often fall short of optimality (e.g., Ruggeri et al., 2016; Meder et al., 2019).

Most importantly, in line with our predictions, children's strategies changed significantly as a function of the characteristics of the task they were presented with. Specifically, in the Condition search phase of the task, children adapted their search strategies to the given cues about the likelihood distribution of the target: Compared to children assigned to the *Uniform* distribution, children assigned to the *Skewed* distribution often used smaller sized scanners to target the highly probable bucket, allowing them to solve the task overall more effectively, i.e. with notably less scanner usage. This is analogous to the switch from adopting a constraint-seeking strategy in uniform distribution environments to a hypothesis-scanning strategy in skewed distribution environments observed in 20-questions games (Ruggeri et al., 2017; Ruggeri & Lombrozo, 2015).

We found strong converging evidence for this adaptive strategy switch from the within-participant comparison. Search behavior during the Baseline search phase, where no

information was provided about the target location, differed from that in the Condition phase for the Skewed, but not the Uniform distribution. Together, these findings robustly support the hypothesis that children are ecological active learners, who can recognize the statistical features of the problem space and tailor their learning strategies accordingly (see Ruggeri & Lombrozo, 2015; Ruggeri et al., 2017, 2019; Ruggeri, 2022). Moreover, our pattern of results highlights the need to focus on *adaptability* when examining performance in developmental research, alongside absolute performance benchmarks, to capture the early emergence of children's competencies.

One limitation of this study is that, while we manipulated children's prior beliefs by priming the likelihood distribution of the potential solution, we did not directly measure their actual beliefs about the hypothesis space before searching. First, we assumed that flagging buckets in which other child searchers found treasures would prompt participants to think that flagged buckets are now empty, making unflagged bucket(s) more likely to contain the treasure. However, one can imagine that this cover story could also lead to the opposite intuition, namely that previous searchers tested currently unflagged buckets and found no treasure because these buckets generally do not contain the treasure. If children would have used this second interpretation of the cover story, they should have adapted their behavior by *avoiding* the unflagged bucket, and using a binary search strategy on the remaining buckets. We found children assigned to the Skewed distribution were somewhat more likely to first test the unflagged bucket than any other bucket, which leads us to exclude this alternative interpretation. Moreover, it is still true that our cover story relies on the ability to reason about others' search behaviors, and to make inferences from the absence of information, which may be more difficult for younger children. However, regardless of the nuances in the interpretation of the cover story, we hoped that the visual display would draw children's attention to either one location or multiple, evenly-spread locations, leading to considering differing sets of candidate hypotheses. Future work should test children's initial beliefs about the problem space directly, and try to monitor how beliefs are updated after each search step. For example, experimenters could ask children to mark all the candidate targets before every search decision, or stronger, implicit priors can be established by providing children with extensive training over multiple rounds. In lieu of a direct behavioral test of children's assumptions, simulations of random and rational agents operating under different priors (of different strengths) could be contrasted with the performance of human participants.

In conclusion, we found that a developmentally appropriate (and fun) search task can reveal strong hallmarks of adaptive search strategies in 3- to 7- year-olds. Our results support the ecological active learning framework and suggest a promising avenue for using non-verbal search tasks to draw a clearer and more fair picture of children's early learning.

---

[4]Please note that because EIG was calculated with respect to a uniform prior, the results do not imply that children in the Skewed group were less efficient than children in the Uniform distribution. Instead, results are indicative of the fact that the children's strategy in the Skewed distribution was not well suited to find a uniformly distributed target.

References

Alvarez, A. L., & Booth, A. E. (2014). Motivated by meaning: Testing the effect of knowledge-infused rewards on preschoolers' persistence. *Child Development*, *85*(2), 783–791.

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

Chai, K.-X., Xu, F., Swaboda, N., & Ruggeri, A. (2023). Preschoolers' information search strategies: Inefficient but adaptive. *Frontiers in Psychology*, *13*.

Herwig, J. E. (1982). Effects of age, stimuli, and category recognition factors in children's inquiry behavior. *Journal of Experimental Child Psychology*, *33*(2), 196–206.

Jirout, J., & Klahr, D. (2012). Children's scientific curiosity: In search of an operational definition of an elusive concept. *Developmental Review*, *32*(2), 125–160.

Jones, A., Swaboda, N., & Ruggeri, A. (2020). Developmental changes in question-asking. In L. P. Butler, S. Ronfard, & K. H. Corriveau (Eds.), *The questioning child: Insights from psychology and education* (pp. 118–143). Cambridge University Press New York, NY.

Kachergis, G., Rhodes, M., & Gureckis, T. (2017). Desirable difficulties during the development of active inquiry skills. *Cognition*, *166*, 407–417.

Legare, C. H., Mills, C. M., Souza, A. L., Plummer, L. E., & Yasskin, R. (2013). The use of questions as problem-solving strategies during early childhood. *Journal of Experimental Child Psychology*, *114*(1), 63–76.

Meder, B., Nelson, J. D., Jones, M., & Ruggeri, A. (2019). Stepwise versus globally optimal search in children and adults. *Cognition*, *191*, 103965.

Meder, B., Wu, C. M., Schulz, E., & Ruggeri, A. (2021). Development of directed and random exploration in children. *Developmental science*, *24*(4), e13095.

Mosher, F. A., & Hornsby, J. R. (1966). On asking questions. In J. Bruner, R. Olver, T. Greenfield, J. Hornsby, H. Kenney, & M. Maccoby (Eds.), *Studies in cognitive growth* (pp. 86–102). Wiley.

Nelson, J. D., Divjak, B., Gudmundsdottir, G., Martignon, L. F., & Meder, B. (2014). Children's sequential information search is sensitive to environmental probabilities. *Cognition*, *130*(1), 74–80.

Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, *101*(4), 608.

Ruggeri, A. (2022). An introduction to ecological active learning. *Current Directions in Psychological Science*, *31*(6), 471–479.

Ruggeri, A., & Feufel, M. (2015). How basic-level objects facilitate question-asking in a categorization task. *Frontiers in Psychology*, *6*, 918.

Ruggeri, A., & Lombrozo, T. (2015). Children adapt their questions to achieve efficient search. *Cognition*, *143*, 203–216.

Ruggeri, A., Lombrozo, T., Griffiths, T. L., & Xu, F. (2016). Sources of developmental change in the efficiency of information search. *Developmental Psychology*, *52*(12), 2159.

Ruggeri, A., Sim, Z. L., & Xu, F. (2017). "Why is toma late to school again?" preschoolers identify the most informative questions. *Developmental Psychology*, *53*(9), 1620.

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.

Ruggeri, A., Walker, C. M., Lombrozo, T., & Gopnik, A. (2021). How to help young children ask better questions? *Frontiers in Psychology*, *11*, 2908.

Schulz, E., Wu, C. M., Ruggeri, A., & Meder, B. (2019). Searching for rewards like a child means less generalization and more directed exploration. *Psychological Science*, *30*(11), 1561–1572.

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

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, *27*(3), 379–423.

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

Swaboda, N., Meder, B., & Ruggeri, A. (2022). Finding the (most efficient) way out of a maze is easier than asking (good) questions. *Developmental Psychology*, *58*(9), 1730–1746.