Preschoolers search longer when there is more information to be gained

Eric Schulz[a,1], **Madeline Pelz**[b], **Alison Gopnik**[c], **and Azzurra Ruggeri**[d,e]

[a]Department of Psychology, Harvard University
[b]Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology
[c]Department of Psychology, University of California, Berkeley
[d]Max Planck Research Group iSearch, Max Planck Institute for Human Development
[e]School of Education, Technical University, Munich
[*]Corresponding author: ericschulz@fas.harvard.edu

ABSTRACT

What drives human search in situations with sparse rewards? We let preschoolers (age 24-52 months) play a game in which they had to search for an animal and measured their persistence in the absence of explicit rewards. Crucially, we either told children which animal they were searching for (known animal condition), so that they gained no additional information when finding it, or we told them that it could be one out of eight different animals (unknown animal condition), so that they gained additional information when finding it. Although none of the children could actually find an animal, no matter how long they searched, our results indicate that the unknown animal group was more persistent in their search. These results enrich our understanding of children's motivation in scenarios with sparse rewards and inform research on artificially intelligent agents facing similar environments.

Introduction

Many recent advances in Artificial Intelligence have come through methods of reinforcement learning, in which a system tries to maximize rewards. However, these methods have limitations, particularly in scenarios with uncertain outcomes and sparse or no rewards. A number of researchers have proposed that "curiosity-based" systems may do better in such scenarios[1,2]. Though these systems differ, they have in common the idea that an agent is designed so that increases in information are themselves rewarding and motivate actions. These systems often model themselves explicitly on the curiosity-based exploration of young children. Although, intuitively, it does seem like young children are notably curious and that this contributes to their impressive learning abilities[3], there have been no systematic studies yet showing that young children too are motivated by information gain as an intrinsic reward alone.

Already infants prefer to explore surprising events[4], and Kidd et al.[5] demonstrated that this attentional capture can be characterized in terms of information gain, with infants showing the most attention to situations of intermediate information gain. There are also many studies showing that children are more likely to explore when they are presented with confounded[6] or unexpected evidence[7], seek out uncertainty reduction more eagerly than adults[8] and are sensitive to the information gain in different situations[9]. However, although past results do suggest that children are meaningful and effective explorers, they do not address the question of whether, in general, the opportunity to gain information is enough to motivate children's search. To answer this question, and thereby test the above mentioned curiosity-reward based models directly, we would need to contrast children's actions to obtain rewards in conditions that differ only in their information gain. Classic reinforcement learning methods show that agents will be more likely to act and will act more persistently for greater rewards. Will preschoolers be more persistent in situations where the outcome has greater information value?

In the current study, we measure how long 2- and 3-year-olds are willing to keep searching for something that never occurs. In our paradigm, children were told that they had to find an animal hiding behind one of many doors, sequentially presented on a tablet screen. We manipulated whether children knew the animal they were looking for (*known animal* condition; e.g., find the lion) or whether children did not know which of eight possible animals was hiding (*unknown animal* condition). This subtle difference in the instructions impacted the information gain (IG) of the same action (i.e., opening a door) across the two different conditions. Indeed, in the unknown-animal condition, opening a door had the potential to reveal not only the location of the hidden animal, but also its identity. Formally, the IG for finding the animal in the known animal condition was $-\log(1) = 0$, whereas the IG in the unknown animal condition was $-\log(1/8) = 3$.

Our results show that children searched more persistently in the unknown animal condition, where with each action there was more information to be gained. These results extend our understanding of children's motivation, and underline the importance of curiosity-driven algorithms.

Results

As a first indicator of children's persistence, we looked at how many doors they opened before giving up (Fig. 1a). Children in the unknown animal condition searched longer, i.e. opened more doors on average, than children in the known animal condition (25.22 vs. 16.83 doors; $t(47) = 3.43$, $p < .001$, $d = 0.98$, $BF = 50.8$).

As a second indicator, we measured the latency from the start of the experiment to the first time they showed signs of giving up the search (e.g., they left their chair; Fig. 1b). Children in the unknown animal condition searched longer on average than children in the known animal condition (133.22 vs. 91.31 seconds; $t(47) = 3.46$, $p < .001$, $d = 0.99$, $BF = 55.1$).

As a third indicator, we looked at the proportion of participants who kept opening doors until a pre-determined time limit of 4 minutes (Fig. 1c; see Methods). Whereas only 7 out of 24 participants reached the time limit in the known animal condition, 20 out of 25 did so in the unknown animal condition. Thus, significantly more children completed the game in the unknown animal condition ($\chi^2(1, N = 49) = 12.79$, $p < .001$, $BF = 207.9$).

Finally, we regressed participants' age in months and their condition onto the number of opened doors* in a Bayesian linear regression, extracting the posterior estimates of both variables' effects onto the number of doors opened (Fig. 1d). This showed that children opened more doors the older they were ($\hat{\beta} = 3.867$, 95%HPD: $[1.69, 6.10]$), and that the effect of the number of animals persisted even when controlling for age ($\hat{\beta} = 10.02$, 95%HPD: $[5.63, 14.33]$).

Discussion

A challenging problem for intelligent systems is how to behave in scenarios with sparsely occurring or no rewards. Yet already young children successfully cope with this challenge. What motivates children to be persistent information gatherers? One idea is that children perceive information itself as rewarding, and are therefore more motivated in situations where more information can be gained.

We studied children's persistence in a search task without rewards, manipulating across two conditions the expected information gain of the same search actions. Our measures converged in showing that

*Note that we find similar results when regressing onto latency.

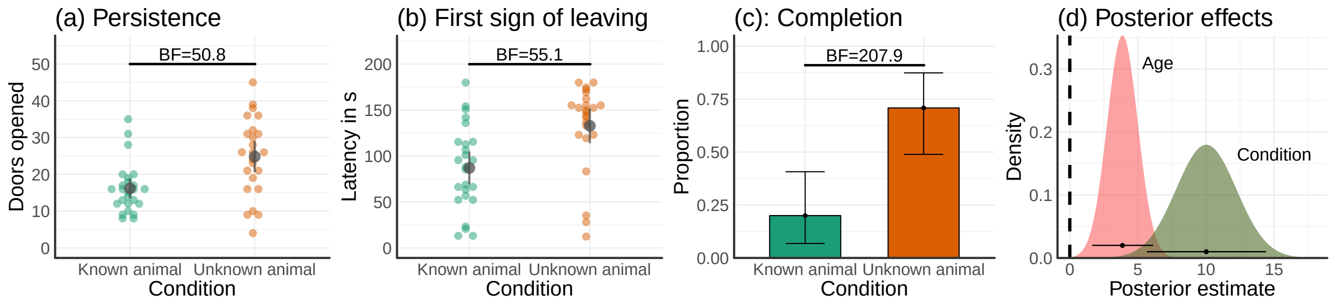

Figure 1. Experimental results. **a)** Total number of doors opened by condition. Dots indicate individual participants. Red dot and error bars indicate the group mean and the 95% confidence intervals. **b)** First sign of leaving, measured as the latency in seconds at which children first seemed to leave the game. Dots indicate individual participants. Black dot and error bars indicate the group mean and the 95% confidence intervals. **c)** Proportion of children completing the experiment without getting distracted and aborting search. Error bars represent the 95% confidence interval of a binomial distribution. **d)** Posterior effects of age and condition onto the number of doors opened. The vertical dashed line indicates an effect of 0, while the black dot indicates the mean effect and confidence intervals show the 95% highest posterior density (HPD).

children were more persistent in their search when the expected information gain of their actions was higher, even if no rewarding outcome ever appeared.

Our results have implications for both developmental psychology and artificial intelligence. Developmentally, our results establish that children can be motivated by the expected information gain of an outcome alone. Thus, one could design educational tools which tell children that the desired outcome could take different shapes, which –according to our findings– would motivate them more strongly. Furthermore, our findings suggest that to build machines that learn like children, one should design algorithms that are motivated by the underlying information gain of their actions. Similar proposals have been put forward already, specifically in research on curiosity and intrinsic motivation[2]. For example, Schmidhuber[1] proposed that the concept of "fun" can be formalized by measuring how much one's model of the world improved from one time step to another. Considering that the expected information gain was higher in the unknown animals condition compared to that in the known animal condition on every step of the game, it could be that children in the eight animal condition persisted longer because they experienced more fun. To further map out all of the details of children's persistent search, we as scientists have to keep searching persistently.

References

1. Schmidhuber, J. Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Auton. Mental Dev.* **2**, 230–247 (2010).

2. Pathak, D., Agrawal, P., Efros, A. A. & Darrell, T. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 16–17 (2017).

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

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

5. Kidd, C., Piantadosi, S. T. & Aslin, R. N. The goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PloS one* **7**, e36399 (2012).

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

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

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

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

Methods

Participants

We recruited 49 young children (range: 24 to 52 months, 26 female, $M_{\text{age}} = 38.06$; SD=7.77) from preschools and museums in the San Francisco Bay Area. Informed consent was obtained from the parents of all children. Ethics approval was obtained by the Max Planck Ethics Review Board.

Design and Materials

The experiment used a between-subjects design, where participants were assigned to either the known animal or the unknown animal condition. The task was presented on a tablet, which presented a (potentially infinite) series of closed doors, displayed one at the time. After clicking on a door, it would open to reveal what was behind it. Afterwards, the next closed door would be displayed. Children were told to find the animal hidden behind one of the doors. In the known animal condition, children were told which animal was hiding, randomly assigned for each child to be either a lion, an elephant, a hippo, a zebra, a crocodile, a bird, a turtle, or a whale. In the unknown animal condition, children were told that the animal hiding could have been one of the eight animals, but were not told which one. In both conditions, the animal would never appear and there was nothing behind each door. The session was ended if the child had not given up searching after 4 minutes. At the end of session, the experimenter encouraged children to open one more door, at which point they would actually find the animal.

Statistical tests

We report all statistics using both frequentist and Bayesian tests. Frequentist tests are presented alongside their effect sizes, i.e. Cohen's d. Bayesian statistics are expressed as Bayes factors (BFs), quantifying the likelihood of the data under the alternative hypothesis compared to the null hypothesis. We apply Bayesian t-test for comparing independent groups, using a Jeffreys-Zellner-Siow prior with its scale set to $\sqrt{2}/2$. We perform Bayesian linear regression with a normal prior on the weights $\hat{\beta} \sim \mathcal{N}(0, 100)$.