

Development of directed and random exploration in children

Björn Meder¹  | Charley M. Wu²  | Eric Schulz³  | Azzurra Ruggeri⁴ 

¹Health and Medical University Potsdam and Max Planck Institute for Human Development, Berlin, Germany

²University of Tübingen and Max Planck Institute for Human Development, Berlin, Germany

³Max Planck Institute for Biological Cybernetics, Tübingen, Germany

⁴Max Planck Institute for Human Development and Technical University Munich, Berlin, Germany

Correspondence

Björn Meder, Health and Medical University, Potsdam, Olympischer Weg 1, 14471 Potsdam, Germany.

Emails: bjoern.meder@health-and-medical-university.de; meder@mpib-berlin.mpg.de.

Funding information

Deutsche Forschungsgemeinschaft, Grant/Award Number: EXC 2064/1 – 390727645; Bundesministerium für Bildung und Forschung, Grant/Award Number: FKZ: 01IS18039A

Abstract

Are young children just random explorers who learn serendipitously? Or are even young children guided by uncertainty-directed sampling, seeking to explore in a systematic fashion? We study how children between the ages of 4 and 9 search in an explore–exploit task with spatially correlated rewards, where exhaustive exploration is infeasible and not all options can be experienced. By combining behavioral data with a computational model that decomposes search into similarity-based generalization, uncertainty-directed exploration, and random exploration, we map out developmental trajectories of generalization and exploration. The behavioral data show strong developmental differences in children's capability to exploit environmental structure, with performance and adaptiveness of sampling decisions increasing with age. Through model-based analyses, we disentangle different forms of exploration, finding signature of both uncertainty-directed and random exploration. The amount of random exploration strongly decreases as children get older, supporting the notion of a developmental “cooling off” process that modulates the randomness in sampling. However, even at the youngest age range, children do not solely rely on random exploration. Even as random exploration begins to taper off, children are actively seeking out options with high uncertainty in a goal-directed fashion, and using inductive inferences to generalize their experience to novel options. Our findings provide critical insights into the behavioral and computational principles underlying the developmental trajectory of learning and exploration.

KEYWORDS

directed exploration, exploration-exploitation dilemma, generalization, multi-armed bandit task, random exploration, search

1 | DEVELOPMENT OF DIRECTED AND RANDOM EXPLORATION IN CHILDREN

Children are natural born explorers. While exploration and active learning are quintessential features of development and maturation, they also pose fundamental challenges to children and adults alike. In particular, efficiently searching for information and rewards requires

balancing the dual goals of exploring unknown options to learn something new, and exploiting familiar options to obtain known rewards. At a restaurant, should you go with your usual favorite or should you try the chef's latest creation? As a child, should you play your favorite game again or try out something new? Exploring novel options can potentially reveal new and even better rewards, but could also lead to disappointment. Known as the *explore–exploit dilemma*, this

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2021 The Authors. *Developmental Science* published by John Wiley & Sons Ltd.

fundamental problem contrasts the goals of gaining knowledge to reduce uncertainty with immediately acquiring rewards.

Optimal solutions to explore–exploit dilemmas are unattainable in all but limiting cases (Bellman, 1952; Gittins & Jones, 1979), making heuristic strategies an active area of research in many fields, including cognitive and developmental psychology. Whereas many studies have investigated how adults balance exploration and exploitation (for reviews, see Cohen et al., 2007; Hills et al., 2015; Mehlhorn et al., 2015), less is known about the developmental processes that shape learning and exploration during childhood. Studying how children, who have fewer cognitive resources and less experience, approach such problems can provide critical insights into the computational and behavioral principles that drive learning and development more generally. Here, we investigate developmental trajectories in learning and exploration between the ages of 4 and 9, an age range where substantial changes in children's exploration behavior have been observed across different tasks (Betsch et al., 2016; Ronfard et al., 2018; Ruggeri, Markant, et al., 2019; Ruggeri, Xu, et al., 2019). To map out developmental trajectories, we combine behavioral data from a spatial search task with predictions from a computational model that disentangles different forms of exploration. Consistent with previous theories (Gopnik et al., 2017), our results show that the exploration patterns of young children are characterized by high levels of random sampling, which decreases with age. However, even at the youngest age range, children do not rely solely on random exploration, but they actively seek out options with high uncertainty (directed exploration) and use inductive inferences to predict unobserved rewards (generalization).

1.1 | How to explore: Random exploration, directed exploration, and generalization

Research on explore–exploit problems typically contrasts two distinct classes of exploration strategies (Gershman, 2018; Wilson et al., 2014). *Random exploration* models exploration by adding noise to the decision process (Luce, 1959; Thompson, 1933). Instead of only making reward-maximizing decisions, this added randomness can lead to the incidental exploration of new options and (better or worse) rewards. This exploration strategy is often also referred to as *undirected exploration*, because it is not goal oriented but merely relies on adding more randomness to the search process. Related to this strategy, it has been recently suggested that children's exploration behavior is characterized by “higher temperature” (i.e., noisier) sampling, which “cools off” with age (Gopnik et al., 2017). The idea behind the *temperature* analogy evokes methods such as *simulated annealing* (Kirkpatrick et al., 1983), which is an optimization algorithm that uses a time-dependent reduction of randomness to avoid getting stuck in a local optimum. Higher temperatures produce more randomness during the search process. Over time, the algorithm cools off, implementing a gradual decrease in the amount of random exploration of possible solutions. On this view, young children exhibit high amounts of random sampling, which results in exploration

Research Highlights

- We investigate developmental trajectories in random and uncertainty-directed exploration in children between 4 and 9 years, using a complex explore–exploit dilemma with spatially correlated rewards.
- Children adapt their search to the structure of the environment but also exhibit a tendency to explore more than beneficial for the goal of maximizing rewards.
- We find a reliable decrease of random exploration between age 4 and 9, as well as substantial levels of uncertainty-directed exploration even in the youngest age range.
- As random exploration begins to taper off, children are already engaging in more sophisticated forms of exploration and generalize their experiences to novel options.

of a larger set of possibilities compared to adults (Cauffman et al., 2010; Mata et al., 2013). As children grow older, temperature decreases, yielding a stronger focus on reward maximization, leading to less diverse sampling behavior (Bonawitz et al., 2014).

Directed exploration (Schulz & Gershman, 2019; Wilson et al., 2014) is an alternative strategy, which relies on representing one's uncertainty about the world and then assigning an intrinsic value toward actively reducing this uncertainty (Gottlieb & Oudeyer, 2018). Instead of adding more variability through random (noisy) sampling, directed exploration actively seeks out uncertainty. According to this view, obtaining information is rewarding in and of itself, and the value of an option is inflated through an *uncertainty bonus* (Auer, 2002). By valuing uncertainty positively, directed exploration encourages sampling options with promising but uncertain rewards, rather than focusing merely on exploiting known high-reward options. Computationally, directed exploration is more demanding, since it requires a richer representational structure that encodes both expected rewards and the underlying uncertainty. However, already infants have been shown to value the exploration of uncertain options positively (Schulz, 2015), 6- and 7-year-olds can integrate prior beliefs and obtained evidence in simple learning and exploration tasks (Bonawitz et al., 2012), and children aged 7–11 have been shown to rely more on directed exploration than adults when searching for rewards (Schulz et al., 2019).

In addition to random and directed exploration, the ability to *generalize* (Shepard, 1987) is another important cognitive capacity for navigating the exploration–exploitation dilemma. In particular, generalization provides traction for exploring large problem spaces by making predictions about novel options. For instance, when Italian immigrants came to the United States around 1900, they brought with them knowledge and love of the classic Neapolitan pizza. In their search for creating similarly rewarding dishes, they explored a variety of novel, but similar options—giving the world Chicago-, New York-, and California-style pizza, as well as several other new

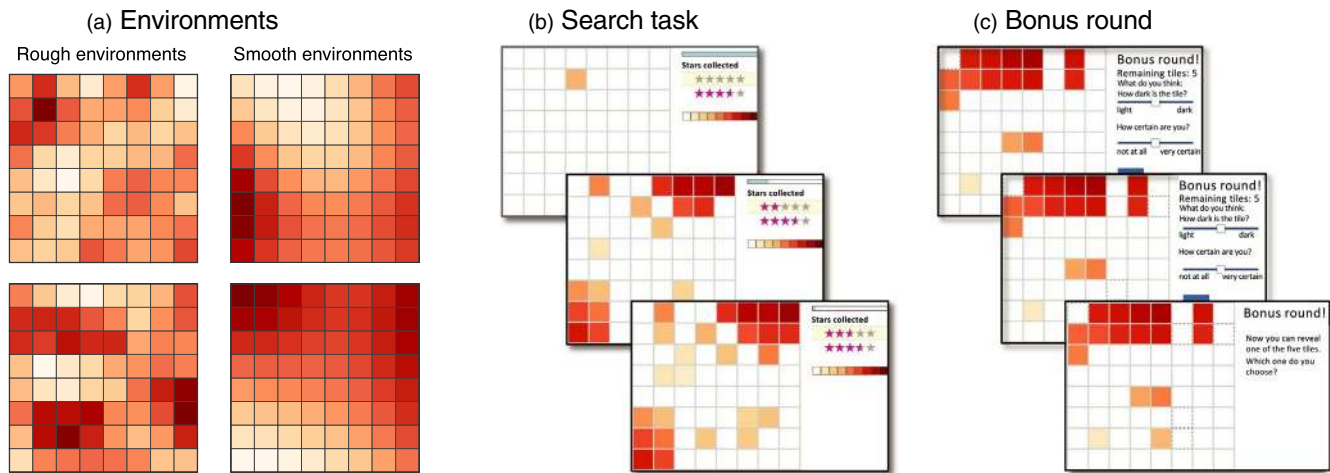


FIGURE 1 Example environments and screenshots from experiment. (a) Two rough environments with low spatial correlation and two smooth environments with high spatial correlation. Darker shades of red correspond to higher rewards. (b) Exploration task, in which children had 25 clicks in each round to obtain as many stars as possible by finding darker (i.e., more rewarding) tiles. (c) Bonus round judgments, in which children predicted the rewards for five previously unobserved tiles (tile with dashed border) and made a confidence judgment about their prediction

variations. A child encountering a new toy can predict whether or not it will be fun by comparing it to other toys it has encountered. If it appears similar to other fun toys, there is a good chance this new toy is also fun. Thus, generalization provides critical guidance for *which* options to explore—namely those which are similar to known high-reward options. On this view, developmental differences in exploration are tightly connected to the ability to make inductive inferences about unexplored options based on prior experience. As cognitive functions and memory develop, they enable more complex cognitive processes and representations (Blanco et al., 2016), thereby supporting more effective generalization for guiding exploration. For instance, changes in search behavior over the life span may be due to the accumulation of knowledge, with adults having stronger inductive biases than children, who seem to weigh new evidence more strongly (Gopnik et al., 2015).

2 | GOALS AND SCOPE

While random and directed exploration are conceptually different, they are not mutually exclusive. Research shows that both types of exploration strategies contribute to search and decision-making in adolescent and adult participants (Gershman, 2018; Somerville et al., 2017; Wilson et al., 2014), with dissociable neural signatures underlying the two forms of exploration (Zajkowski et al., 2017). In addition, both children and adults rely on generalization to learn about the environment and make inferences from experienced to not-yet-explored options (Schulz et al., 2018, 2019; Wu et al., 2018).

The goal of the present paper is to investigate how young children, aged 4–9 years, balance random and directed exploration, using a spatial search task with correlated rewards. In particular, we trace age-related differences in learning and exploration using a computational model that combines similarity-based generalization

with both directed and random exploration (Wu et al., 2018). Our data enable a direct test of the “cooling off” hypothesis and offers empirical evidence for the trajectory with which random sampling decreases over the course of childhood development.

Previous studies have shown reliable signatures of generalization and directed exploration in adults, with relatively little random exploration (Wu et al., 2018; Wu, Schulz, Gershman, 2020). In a comparison of children aged 7–11 and adults, Schulz et al. (2019) found no age-related differences in random exploration. Rather, children differed from adults by having higher levels of directed exploration and narrower generalization. While the lack of differences in random exploration does not support the idea of a “cooling off” process over the lifespan, it could also be the case that children aged 7–11 had already transitioned to a lower temperature and had already developed the capacity for directed exploration. Therefore, our goal is to investigate a younger age range to search for the developmental stage where random exploration diminishes and directed exploration emerges.

3 | EXPERIMENT

We used a simplified version of the spatially correlated multi-armed bandit paradigm (Wu et al., 2018) to investigate how children learn and search for rewards on a grid world by clicking on different tiles (Figure 1). Each tile had a different reward distribution, where the goal was to accumulate as many rewards as possible within a limited search horizon (i.e., a fixed number of clicks). Rather than displaying rewards numerically, as in previous experiments (Schulz et al., 2019), here the value of rewards was indicated using different shades of red to be interpretable by children as young as 4 (Figure 1). In this task, rewards were spatially correlated, such that nearby options had a similar mean reward. Thus, participants could use generalization

from a sparse number of observations to guide their exploration toward promising regions of the search space. Importantly, the number of available clicks (25) was much smaller than the number of available options (64), requiring searchers to balance clicking novel tiles to discover new rewarding options (exploration) with re-clicking tiles already known to provide high rewards (exploitation).

3.1 | Methods

3.1.1 | Participants

We recruited 102 children between age 4 and 9 years. There were 54 children whose age was below or equal to the median of 82 months, and 48 children who were older than the median age. We refer to the group of younger children henceforth as 6-year-olds ($M = 72.6$ months, $SD = 7.6$, range 51 – 82 months, 24 female), and to the group of older children as 8-year-olds ($M = 93.1$ months, $SD = 6.5$, range 84 – 108 months, 23 female) from public museums in Berlin, Germany. In addition to comparing these age groups, we also conducted analyses that treat age as a continuous variable. Fourteen additional children were excluded from analysis because they failed the instruction check ($n = 9$), did not want to play anymore ($n = 1$), were not native speakers ($n = 2$), or because their parents intervened during the experiment ($n = 2$). The study was approved by the ethical review board of the Max Planck Institute for Human Development in Berlin. Informed consent was obtained from children's legal guardians prior to participation; average duration was about 12 min.

3.1.2 | Materials, design, and procedure

Children played six rounds of a spatial search game on a tablet, in which they were presented with an 8×8 grid world with spatially correlated rewards (Figure 1). The expected reward across all environments was identical (i.e., average reward over all tiles of a grid); what differed between environments was the spatial correlation among rewards. The strength of the spatial correlations was manipulated between subjects, with *smooth environments* having stronger spatial correlations than *rough environments*. For each class of environments, we generated 40 different environments using a radial basis function kernel (see Equation 1) with either $\lambda_{smooth} = 4$ or $\lambda_{rough} = 1$. Each environment defined a bivariate reward function on the grid, such that each tile location was mapped to a reward value. Intuitively, smooth environments had smoother reward functions that varied gradually over the grid, whereas rough environments had rougher reward functions that varied more suddenly (Figure 1). On each round, a new environment was sampled without replacement from the set of 40 environments for the respective class.

At the beginning of each round, one random tile was revealed and children could sequentially sample 25 tiles. On each trial, they could either click a new tile or re-click a tile they had already selected before (clicking was done by touching the desired tile on the

tablet). Clicking a tile for the first time revealed its color, with darker colors indicating higher rewards along a continuous, linearly scaled color range (Figure 1). The color (i.e., underlying reward) of the revealed tiles remained visible for the entire duration of the round. Re-clicked tiles could show small variations in the observed color due to normally distributed noise, $\epsilon \sim \mathcal{N}(0, 1)$, with the revealed color indicating the most recent observation (Figure 1b).

To avoid having the global maximum immediately recognizable when revealed, we randomly sampled a different maximum value in each round from a uniform distribution $\sim \mathcal{U}(0.7, 0.9)$. Color values were re-scaled in each round such that the lowest value corresponded to 10% of the darkest value and the highest value corresponded to the randomly sampled maximum (between 70% and 90% of the darkest value). Note that because of the noise applied to observations, sampled rewards could be below 10% or above 90% darkness, hence the additional range in our color scale. Reward values reported throughout the paper are arbitrarily scaled to the range [0,50] to be consistent with previous work (Schulz et al., 2019).

Children were awarded up to five stars at the end of each round (e.g., 4.6 out of 5; see Figure 1b), based on the ratio of their average reward to the global maximum of the given grid. At the beginning of a round, the stars were empty, then they continuously filled up in accordance with each obtained reward. The instructed goal was to collect as many stars as possible in each round; at the end of the game, children received a number of stickers proportional to the average number of stars earned in each round.

In total, children played six rounds of the spatial search game. The first round was a tutorial round, in which children were familiarized with the goal of the game, the spatial correlation of rewards, the maximum number of clicks allowed per round, and the possibility of re-clicking tiles. Specifically, children were told that before each click they would have to decide whether to reveal a novel tile or re-click an already revealed tile. Both actions were explicitly demonstrated by the experimenter. After the tutorial, children were required to answer three comprehension questions. These questions pertained to the instructed task, that stars could be collected both by revealing new tiles and re-clicking previously revealed tiles, and the distribution of tiles in the grid (Appendix D, Figure D1 bottom right). If they failed to answer any of the questions correctly, the relevant part of the instructions was repeated and the questions were asked again. If a child failed again, they continued with the experiment, but were later excluded from the analyses. Children were not explicitly told that the expected reward of individual tiles was constant in each round, or that the expected reward across all options was the same in each environment. However, we also never suggested otherwise (e.g., that rewards might change or reverse over time).

Rounds two to five comprised the actual exploration task, where in each round children had 25 clicks to find rewards on the grid. The sixth and last round was a bonus round, in which children sampled for 15 trials and then made reward predictions for five randomly chosen and previously unobserved tiles (Figure 1c). This was explained to them before the bonus round started. Judgments were made using a continuous slider, asking children to indicate the

darkness of the target tile, with the end points labeled as “light” and “dark.” When moving the slider, the target tile changed its color accordingly. The underlying reward scale was continuous, ranging from 0 to 50. To assess the level of confidence associated with the reward predictions, children were asked how certain they were about the predicted darkness, using a slider from 0 to 10 in steps of 1, with the endpoints labeled as “not certain at all” to “very certain.” After judging five tiles, children were asked to select one of them. They received the corresponding reward and then continued the round until the search horizon was exhausted.

4 | BEHAVIORAL RESULTS

We first analyze the behavioral data in terms of performance and exploration behavior. These analyses exclude the tutorial and bonus rounds, leaving a total of 100 search decisions (4 rounds \times 25 trials) for each of the 102 participants. We then report the results of the bonus round, where we analyze children's reward predictions and confidence judgments. The behavioral data are complemented by model-based analyses, where we disentangle generalization, directed exploration, and random exploration. We report both frequentist statistics and Bayes factors (*BF*) to quantify the relative evidence of the data in favor of the alternative hypothesis (H_A) over the null hypothesis (H_0) (see Appendix A for details).

4.1 | Exploration task: Performance

Whereas both smooth and rough environments had the same expected rewards, the stronger spatial correlations in the smooth environment facilitated better performance for both age groups (6-year-olds: $M_{smooth} = 29.9$ vs. $M_{rough} = 26$, $t(52) = 3.3$, $p = 0.002$, $d = 0.9$, $BF = 22$; 8-year-olds: $M_{smooth} = 34.3$ vs. $M_{rough} = 28$, $t(46) = 6.4$, $p < 0.001$, $d = 1.8$, $BF > 100$; Figure 2a). Thus, regardless of age, children were able to leverage the spatial correlation of rewards

in the environment, and performed better in more correlated environments. Performance was more variable in smooth compared to rough environments (6-year-olds: $F(29, 23) = 3.8$, $p = 0.002$; 8-year-olds: $F(21, 25) = 2.7$, $p = 0.002$), indicating individual differences in the ability to learn about and harness the environmental structure when searching for rewards.

Eight-year-old children obtained higher rewards than 6-year-olds in both rough ($M = 28$ vs. $M = 26$, $t(48) = 2.6$, $p = 0.012$, $d = 0.7$, $BF = 4.1$) and smooth environments ($M = 34.3$ vs. $M = 28.9$, $t(50) = 3.3$, $p = 0.002$, $d = 0.9$, $BF = 19$). Age-related performance differences were also found when treating age as continuous variable (Figure 2b), with performance increasing with age in both rough (Pearson's $r = .36$, 95% CI = [.09, .58], $p = 0.011$, $BF = 6.0$) and smooth environments ($r = .39$, 95% CI = [.14, .60], $p = 0.004$, $BF = 14$).

Figure 2c shows the learning curves (average reward over trials; first aggregated within and then across participants). Consistent with the overall performance, learning curves increased more strongly in smooth compared to rough environments. In rough environments, 8-year-olds performed slightly better than 6-year-olds, but generally there was only little improvement over trials. In smooth environments, older children learned more quickly than younger children and consistently outperformed them. A notable finding is that in smooth environments, toward the end of the search, the average obtained rewards tended to decrease again, in both age groups, suggesting a tendency to continue exploration even at the cost of foregone rewards.

4.2 | Exploration task: Search trajectories

Rather than only comparing performance, we also looked for behavioral patterns in how children searched for rewards, by analyzing the distance between consecutive choices and how this was affected by the magnitude of rewards and the subsequent search decisions. Figure 3a shows the distribution of Manhattan distances between consecutive

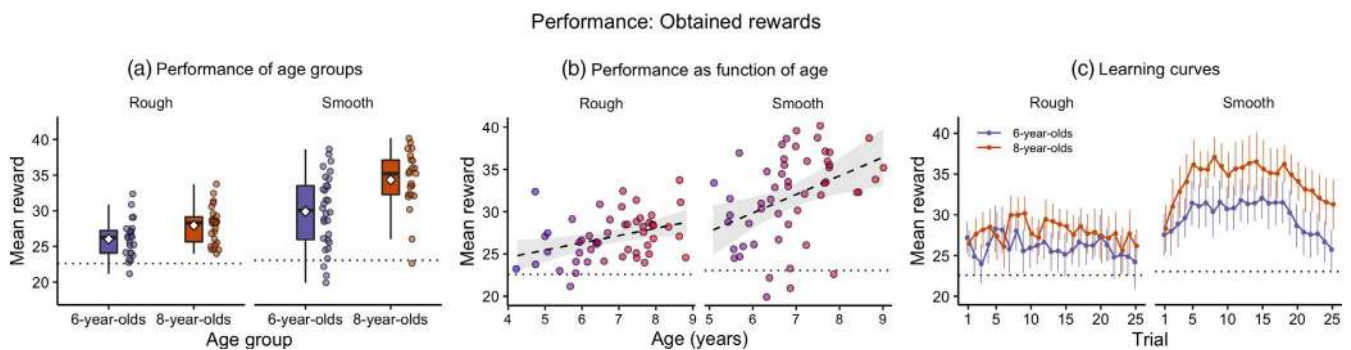


FIGURE 2 Obtained rewards measured in arbitrary units in the range [0,50]. (a) Tukey box plots of the distribution of obtained mean rewards, separately for each age group and environment. Each dot is a participant-wise mean, the horizontal line in the box shows the group median and the diamonds indicate group means. Dotted line is random performance. (b) Average obtained rewards as a function of age in smooth and rough environments. Each dot represents one participant, the dashed line shows a linear regression ($\pm 95\%$ CI); dotted line is random performance. (c) Learning curves showing the average rewards over trials, first averaged within participants and then aggregated across participants; error bars are 95% CIs

choices. For 8-year-olds, the mean distance was smaller in smooth than in rough environments ($M_{smooth} = 2.04$ vs. $M_{rough} = 2.69$, $t(46) = -3.1$, $p = 0.003$, $d = 0.9$, $BF = 13$), indicating they searched more locally in the presence of strong spatial correlations. For 6-year-olds, there was no difference between environments ($M_{smooth} = 2.11$ vs. $M_{rough} = 1.93$, $t(52) = 1.0$, $p = 0.31$, $d = 0.3$, $BF = .42$), suggesting a more limited capability to adapt to environmental structure.

We also analyzed search decisions (Figure 3b) by computing the proportions of *repeat choices*, corresponding to re-clicking the previously revealed tile, *near choices*, corresponding to searching a neighboring tile (i.e., distance of 1), and *far choices*, corresponding to clicking tiles with a distance larger than 1. Older children tended to search more locally in smooth compared to rough environment, while conversely making more far choices in rough compared to smooth environments.

This pattern was not observed for 6-year-olds, indicating that younger children did not adapt their search patterns to the correlation structure of rewards in the environment. Notably, the number of repeat clicks is overall rather low, regardless of age group and environment (see Section 6). This may also explain the learning curves (Figure 2c), which tended to decrease toward the end of each round in smooth environments. This demonstrates that children generally show higher levels of exploration when searching for rewards, and thus less exploitation of high-value options that have already been observed.

Finally, we analyzed the relation between the value of a reward obtained at time t and the search distance on the subsequent trial $t + 1$. If a large reward was obtained, searchers should search more locally, while conversely, if a low reward was obtained, searchers should be more likely to search farther away. Using hierarchical Bayesian regression analyses, we predicted search distance using the reward obtained on the previous step, age group, and their interactions as population-level ("fixed") effects, while treating participants as random intercepts. Figure 3c shows how the reward obtained from the previous choice related to subsequent search distance (see Table B1 in Appendix B for detailed results). Both 6- and 8-year-olds tended to search more locally when high rewards were obtained and searched further away when low rewards were obtained. The two age groups were differentially influenced by the obtained rewards, such that 8-year-olds more markedly increased the distance following low rewards compared to 6-year-olds, in both smooth and rough environments. Taken together, these findings indicate that the magnitude of rewards influenced search distance, but 8-year-olds were more responsive in adapting their search behavior than 6-year-olds.

4.3 | Bonus round judgments

The last round was a bonus round in which children made 15 search decisions and then predicted the expected rewards for five random, unrevealed tiles. Additionally, they were also asked how confident they were about the predicted reward (i.e., darkness of tile).

Figure 4a shows the mean absolute error between children's estimates and the true underlying expected reward. Overall, 8-year-olds had lower prediction error than 6-year-olds ($M = 11.5$ vs. $M = 16.5$, $t(100) = 3.9$, $p < 0.001$, $d = 0.8$, $BF > 100$). The difference between age groups was found in both environments, albeit less pronounced in rough ($M = 11.5$ vs. $M = 15.5$; $t(48) = 2.4$, $p = 0.019$, $d = 0.7$, $BF = 2.9$) compared to smooth environments ($M = 17.2$ vs. $M = 11.5$; $t(50) = 3.0$, $p = 0.004$, $d = 0.8$, $BF = 9.1$). Aggregating both age groups, we found no effect of environment on prediction error ($M_{rough} = 13.4$ vs. $M_{smooth} = 14.8$; $t(100) = -1.0$, $p = 0.32$, $d = 0.2$, $BF = .32$). We constructed a random baseline by sampling 10,000 random values from the reward interval $[0, 50]$ and 10,000 samples (with replacement) from the true reward values in the bonus round environments that children experienced. We then computed the absolute error between each random guess and the bootstrapped true values, and finally computed the mean absolute error across all samples. Compared to this random baseline, 6-year-olds performed worse than chance level ($t(53) = 2.7$, $p = 0.009$, $d = 0.4$, $BF = 4.2$), whereas 8-year-olds were better than chance ($t(47) = -3.1$, $p = 0.003$, $d = 0.4$, $BF = 9.6$). Younger children's performance below chance level can be traced to a tendency to frequently make extreme judgments, a tendency that has also been observed in other studies (Chambers, 2002; Meder et al., 2020). Out of the 270 judgments, 83 (31%) times 6-year-olds predicted a reward of 0 or 50, whereas this was much less frequent in 8-year-olds (22 out of 240, 9%). Since the true rewards in the experienced bonus environments were normally distributed (with a mean around 22), this bias substantially increased prediction error in younger children, resulting in below chance level performance.

Looking at prediction error as a function of age in months (Figure 4), we found that in both rough and smooth environments children's prediction error declined with age (rough: $r = -.40$, $p = 0.004$, $BF = 14$, smooth: $r = -.46$, $p < 0.001$, $BF = 57$). Across all judgments and children, we found no systematic relation between confidence and prediction error (Kendall's rank correlation: $r_t = .07$, $p = 0.04$, $BF = .67$). A Bayesian regression with confidence, age group, and their interaction as predictors and subject-wise random intercept also showed no reliable relationship (see Table B2 in Appendix B).

We also analyzed whether the distance to previously revealed tiles was related to participants' reward predictions and confidence. For each participant, we computed the average (Manhattan) distance of each of the five target tiles to the 15 previously revealed tiles. We then computed subject-wise correlations between distance and either prediction error or confidence, respectively. Seventeen children gave the same confidence judgment to all five predictions, such that the correlation was undefined and were omitted from these analyses. Generally, more proximal target tiles tended to produce lower prediction error (mean correlation: $M_r = .12$) and higher confidence ($M_r = -.07$). However, there were substantial variation between age groups and environments. The prediction error of 8-year-olds decreased more strongly with spatial proximity than 6-year-olds in both environments, although age-related differences were only reliable in rough environments (Rough: mean correlation $M_{8-year-olds} = .24$

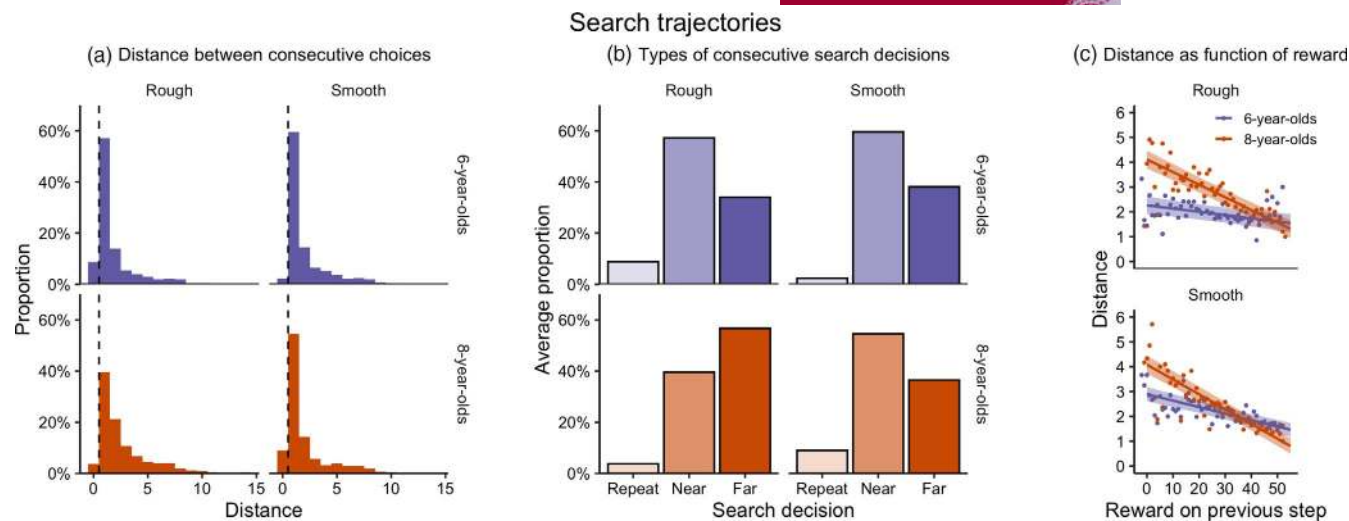


FIGURE 3 Search trajectories. (a) Histogram of distances between consecutive search choices. A distance of zero indicates a repeat click; a distance of 1 corresponds to clicks on neighboring tiles; distances >1 correspond to other clicks on the grid. The vertical dashed line marks the difference between a repeat click and selecting any other tile. (b) Average proportion of search decisions by age group and environment. Repeat clicks correspond to re-clicking a previously revealed tile, near clicks correspond to directly neighboring tiles, and far clicks are sampling decisions with a distance >1 . (c) Search distance as function of reward obtained on the previous trial. The lines visualize the relation between search distance and previous reward for each age group and environment, obtained from a Bayesian regression ($\pm 95\%$ CI). The dots show the observed mean distances given previous rewards, aggregated across all decisions and children. One outlier has been removed from the lower plot, but is included in all statistical analyses.

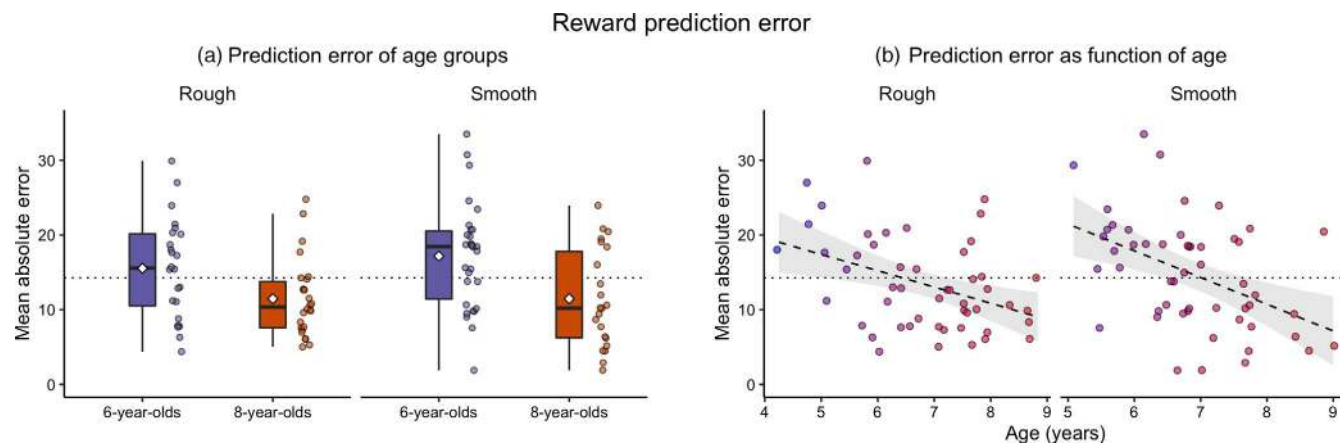


FIGURE 4 Bonus round judgments. (a) Mean absolute prediction error for 6- and 8-year-olds. (b) Mean absolute prediction error as function of age. Each dot is one participant, the dashed line shows a linear regression ($\pm 95\%$ CI). Dotted line is random performance

vs $M_{6\text{-year-olds}} = -.07$, two-sample t test $t(48) = -2.2$, $p = 0.034$, $d = 0.6$, $BF = 1.9$; Smooth: $M_{8\text{-year-olds}} = .17$ vs $M_{6\text{-year-olds}} = .12$, $t(50) = -0.3$, $p = 0.754$, $d = 0.1$, $BF = .29$). Similarly, the confidence ratings of 8-year-olds were higher for more proximate targets than for 6-year-olds, but the age-related differences were not reliable (Rough: $M_{8\text{-year-olds}} = -.21$ vs $M_{6\text{-year-olds}} = .05$, $t(38) = 1.8$, $p = 0.079$, $d = 0.6$, $BF = 1.1$; Smooth: $M_{8\text{-year-olds}} = -.12$ vs $M_{6\text{-year-olds}} = .01$, $t(43) = 0.8$, $p = 0.420$, $d = 0.2$, $BF = .39$). These findings indicate that older children's reward predictions and confidence tended to be more strongly influenced by the spatial distance to known options than 6-year-olds' judgments, but the age-related differences were not consistent.

To analyze selected and nonselected options, we first averaged the predicted reward and confidence of the not-chosen tiles within subjects, and then compared chosen and not chosen options. Selected tiles tended to have higher predicted rewards ($M_{\text{chosen}} = 32$ vs $M_{\text{nonchosen}} = 28.9$, $t(101) = 2.4$, $p = 0.018$, $d = 0.3$, $BF = 1.7$), and there was also a tendency to select options where participants were more confident in their reward predictions ($M_{\text{chosen}} = 7.59$ vs. $M_{\text{nonchosen}} = 7.04$, $t(101) = 2.2$, $p = 0.028$, $d = 0.2$, $BF = 1.2$). Selected tiles also tended to have a higher true reward than nonselected tiles, but the difference was not reliable ($M_{\text{chosen}} = 23.75$ vs. $M_{\text{notchosen}} = 21.95$, $t(101) = 2.0$, $p = 0.048$, $d = 0.3$, $BF = .74$). Thus, children tended to choose options they

expected to have high rewards and for which they were confident in their predictions.

In summary, 8-year-olds obtained higher rewards than 6-year-olds, with both groups performing better in smooth compared to rough environments, facilitated by stronger spatial correlations. Participants adapted their search patterns in response to the magnitude of obtained rewards, searching locally upon finding rich rewards, and searching farther away upon finding poor rewards. The responsiveness of this adaptive search pattern was mediated by age, where 8-year-olds exhibited a stronger relationship between reward value and search distance than 6-year-olds. Lastly, prediction accuracy increased reliably with age, but there was no relation between children's subjective confidence in their reward judgments and their prediction error.

5 | A COMPUTATIONAL ANALYSIS OF DIRECTED AND RANDOM EXPLORATION IN CHILDREN

The behavioral data presented above show strong and systematic differences between the exploration behavior of 6- and 8-year-old children. We next present a computational model that captures key aspects of generalization and sampling strategies in order to map the developmental trajectory of learning and exploration. In particular, the model provides a clear computational framework for estimating to what extent children generalize about the spatial correlation of rewards, and how their sampling behavior can be decomposed into directed and random exploration.

5.1 | The Gaussian process upper confidence bound (GP-UCB) model

Our model consists of three building blocks: a *learning model* that makes predictions about the distribution of rewards in the environment, a *sampling strategy*, which maps these predictions onto valuation of options, and a *choice rule*, which converts value into choice probabilities. We now briefly describe these components, with further details provided in Supplement S1.

5.1.1 | Learning model

To model learning about rewards in the environment we use *Gaussian Process* (GP) regression as a form of Bayesian function learning (Rasmussen & Williams, 2006). The GP uses the principles of Bayesian inference to adaptively learn a value function, mapping the location of each option onto rewards. Generalization about novel options is thus accomplished through interpolation or extrapolation from previous observations (rewards and their locations). This approach has been shown to account for how adults explicitly learn functions (Lucas et al., 2015), and has been successfully applied to

model the behavior of children and adults in a wide range of learning and search tasks (Schulz et al., 2017, 2019; Wu, Schulz, Garvert, et al., 2020; Wu, Schulz, & Gershman, 2020; Wu et al., 2018).

Formally, a GP defines a distribution over functions $f \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$, where each function can be interpreted as a candidate hypothesis about the relationship between spatial location and expected rewards. The GP prior is determined by a mean function $m(\mathbf{x})$ and a kernel function $k(\mathbf{x}, \mathbf{x}')$. We follow the convention of setting the mean function to zero, while using the kernel function to encode the covariance structure. Put simply, the kernel provides an inductive bias about how points in the input space are related to each other as a function of distance (i.e., spatial similarity). A common choice for the kernel is the *radial basis function* (RBF):

$$k(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\lambda^2}\right), \quad (1)$$

where \mathbf{x} and \mathbf{x}' denote two inputs (e.g., coordinates of tiles on the grid) and λ is the *length-scale* parameter governing the extent of generalization. Put simply, the RBF kernel models generalization as an exponentially decaying function of the distance between inputs \mathbf{x} and \mathbf{x}' . This kernel is closely related to Shepard's (1987) universal law of generalization, which models generalization as an exponentially decaying function of similarity, where similarity is the inverse of distance. In the present study, the λ parameter specifically pertains to generalization about the extent of spatial correlation of rewards in the environment, where higher λ values correspond to stronger spatial correlations. For instance, $\lambda = 1$ indicates that the rewards of two neighboring tiles are assumed to be correlated by $r = .61$; if options are further than three tiles away, the correlation decays to effectively zero. Smaller values of λ indicate that the assumed correlation decays more rapidly as a function of distance, while larger values of λ indicate stronger spatial correlations. Thus, this parameter represents how strongly participants generalize across options (tiles) based on their spatial proximity.

In the present task, GP regression generates normally distributed beliefs about the rewards for any tile \mathbf{x} , summarized as expectation $\mu(\mathbf{x})$ and uncertainty $\sigma(\mathbf{x})$. These predictions are modulated by the length-scale parameter λ , which defines the extent to which rewards are assumed to be correlated as a function of distance. For instance, $\lambda = 1$ corresponds to the assumption that the rewards of two neighboring tiles are correlated by $r = 0.6$, and that due to the exponential decay this correlation effectively decreases to zero for options further than three tiles apart. We treat λ as a free parameter, which we estimate for each individual participant. This enables us to assess each child's tendency to generalize.

5.1.2 | Sampling strategies

Given a learner's belief about expected reward $\mu(\mathbf{x})$ and estimated uncertainty $\sigma(\mathbf{x})$, we use a sampling strategy to map these

beliefs onto a valuation for each option. Specifically, we use *Upper Confidence Bound* (UCB) sampling (Auer, 2002) to model directed exploration as a simple weighted sum:

$$\text{UCB}(\mathbf{x}) = \mu(\mathbf{x}) + \beta\sigma(\mathbf{x}) \quad (2)$$

where μ is the mean expected reward and β represents the extent to which uncertainty σ (measured in terms of the standard deviation of \mathbf{x}) is valued positively. The parameter β is an “uncertainty bonus,” since it optimistically inflates expected rewards by their degree of uncertainty. UCB provides an effective sampling strategy for balancing the exploration-exploitation dilemma, by mediating between exploring novel options to reduce uncertainty while also prioritizing the exploitation of high-value options.

To illustrate this sampling strategy, consider two options (tiles) \mathbf{x}_1 and \mathbf{x}_2 . Option \mathbf{x}_1 has expected reward of $\mu(\mathbf{x}_1) = 50$ and uncertainty $\sigma(\mathbf{x}_1) = 5$. Option \mathbf{x}_2 has expected reward of $\mu(\mathbf{x}_2) = 45$ and uncertainty $\sigma(\mathbf{x}_2) = 15$. Thus, option \mathbf{x}_1 has higher expected reward than \mathbf{x}_2 , but \mathbf{x}_2 is more uncertain. UCB sampling takes into account both reward and uncertainty to balance the explore-exploit trade-off. For instance, if $\beta = 1$, $\text{UCB}(\mathbf{x}_1 | \beta = 1) = 50 + 5 = 55$ and $\text{UCB}(\mathbf{x}_2 | \beta = 1) = 45 + 15 = 60$, meaning that option \mathbf{x}_2 is more attractive than option \mathbf{x}_1 . By contrast, if $\beta = 0.2$, then $\text{UCB}(\mathbf{x}_1 | \beta = 0.2) = 50 + 1 = 51$ and $\text{UCB}(\mathbf{x}_2 | \beta = 0.2) = 45 + 3 = 48$. In this case, option \mathbf{x}_1 is valued higher than \mathbf{x}_2 , making it more likely to click this tile. Thus, the higher β , the stronger a searcher values uncertainty positively, nudging them toward sampling uncertain options. Conversely, when $\beta \rightarrow 0$ the value of an option is dominated by its expected reward, regardless of the attached uncertainty. In our model, we estimate β for each learner based on their individual search behavior, to assess their level of uncertainty-directed exploration.

5.1.3 | Choice rule

The final component of the model is the choice rule, which translates UCB values into choice probabilities with a softmax function:

$$p(\mathbf{x}) = \frac{\exp(\text{UCB}(\mathbf{x})/\tau)}{\sum_{j=1}^N \exp(\text{UCB}(\mathbf{x}_j)/\tau)}. \quad (3)$$

Importantly, the softmax choice contains a temperature parameter τ that governs the amount of randomness in the choice probabilities. This enables us to quantify the amount of random exploration for each learner. Higher temperature sampling corresponds to noisier predictions, where as $\tau \rightarrow \infty$, all options have an equal probability of being chosen. Conversely, lower temperatures produce choice probabilities that are more concentrated on high-value options, where as $\tau \rightarrow 0$, it becomes an argmax choice rule (i.e., always choosing the option with the highest value). In our model, τ is estimated from the data, to assess the amount of random exploration for each child.

5.1.4 | Model summary

In sum, the GP-UCB model combines (i) a learning component that generalizes from limited observations to unobserved options, (ii) a UCB sampling strategy that inflates expectations of reward by the associated uncertainties to perform directed exploration, and (iii) a softmax choice rule that converts UCB values into choice probabilities and adds decision noise as a form of random exploration. Each model component has a single free parameter that we estimate through cross-validation from children's search decisions: the length-scale parameter λ indicates the extent of generalization, the uncertainty bonus β defines the level of directed exploration, and the temperature parameter τ captures the amount of random exploration. Careful analyses of these parameters provides a window into the computational principles of learning and exploration, enabling us to identify age-related changes.

5.2 | Model comparison

We contrast the predictive accuracy of the GP-UCB model with a Bayesian reinforcement learning model (*Mean Tracker*; MT). Both models share the same uncertainty bonus β and temperature parameter τ , but in place of the GP λ parameter, the MT uses an error variance parameter θ_ϵ^2 , which can be interpreted as inverse learning sensitivity. Thus, both models have three free parameters, where the MT model uses the same UCB and softmax components, but does not generalize. Instead, it learns independent reward distributions about each option using the principles of associative learning (see Supplement S1 and S2 for details and extended model results including additional sampling strategies).

We used cross validation to assess how well the models predict each searcher's sampling decisions, where—as before—we omit the tutorial round and bonus round. Specifically, we iteratively split each child's data into a training set consisting of three of the four rounds, and holding out the remaining round as a test set. We computed the maximum-likelihood estimates for each model's parameters (range $[\exp(-5), \exp(4)]$) using differential evolution (Mullen et al., 2011) and then evaluated each model's predictive accuracy on the held-out test set. This procedure was repeated for each participant for all rounds.

We can describe the objective performance of our models using *predictive accuracy* as a pseudo- R^2 , comparing the summed out-of-sample log loss for each model k against a random model (i.e., choosing all options with equal probability):

$$R^2 = 1 - \frac{\log \mathcal{L}(M_k)}{\log \mathcal{L}(M_{rand})}, \quad (4)$$

where $\log \mathcal{L}$ represents log loss. Intuitively, $R^2 = 0$ indicates chance-level predictions and $R^2 = 1$ indicates theoretically perfect predictions.

Figure 5a shows the predictive accuracy of the two models for both age groups. The GP-UCB model had higher predictive accuracy

than the MT-UCB model overall ($t(101) = 6.6, p < 0.001, d = 0.7, BF > 100$), and also for each age group (6-year-olds: $t(53) = 3.4, p = 0.001, d = 0.5, BF = 22$; 8-year-olds: $t(47) = 6.1, p < 0.001, d = 1.0, BF > 100$). In total, 73 out of 102 participants were best described by the GP-UCB model: 34 out of 54 six-year-olds (63%) and 39 out of 48 eight-year-olds (81%). These results demonstrate the importance of generalization, since this component was not present in the MT learning model.

5.3 | Developmental differences in parameter estimates

To map the developmental trajectories of learning and search, we analyzed the parameter estimates of the GP-UCB model (Figure 5b). There was no difference in the level of generalization (λ parameter) between 6- and 8-year-olds (Mann-Whitney U -test: $U = 1093, p = 0.18, r_\tau = -.11, BF = .42$). We also analyzed whether the estimate for the generalization parameter differed between smooth and rough environments. The mean λ estimates were higher in smooth than in rough environments ($M_{smooth} = 2.8$ vs. $M_{rough} = 0.56$), in line with the difference in ground truth ($\lambda_{smooth} = 4$ vs. $\lambda_{rough} = 1$). This difference was observed for both age groups (6-year-olds: $M = 3.3$ vs. $M = 0.4$ and 8-year-olds: $M = 2.1$ vs. $M = 0.53$). However, no reliable difference between environments was found when performing a comparison of median parameter values (Mann-Whitney U -test:

$Md_{smooth} = 0.42$ vs. $Md_{rough} = 0.41, U = 1425, p = 0.405, r_\tau = .07, BF = .33$). Generally, there was a tendency to undergeneralize; a finding that echoes related research with adults (Wu et al., 2018) and older children (Schulz et al., 2019). While this may indicate a potential limitation in the ability to harness the amount of spatial correlation in the environment, simulations show that undergeneralization tends to produce better performance than overgeneralization, and can in fact sometimes lead to better performance than precisely matching the true amount of spatial correlation in the environment (Wu et al., 2018).

While there was little difference between age groups regarding their extent of generalization, we found systematic developmental differences in directed and random exploration. Younger children had higher estimates than older children for both the exploration bonus β ($U = 1602, p = 0.041, r_\tau = .17, BF = 1.6$) and temperature τ ($U = 1688, p = .009, r_\tau = .21, BF = 2.2$), with a stronger age-related decrease for the latter. These results indicate that 6-year-olds exhibited a stronger tendency toward both directed and random exploration than 8-year-olds.

Figure 5c-f provide a more detailed analysis of these findings by treating age as a continuous variable. First, Figure 5a shows that the predictive accuracy of the GP-UCB model increased with age (Kendall's $r_\tau = .27, p < 0.001, BF > 100$). Second, consistent with the group-based analyses, there were little changes in the generalization parameter λ as a function of age ($r_\tau = .10, p = 0.14, BF = .39$). In contrast, both the uncertainty bonus parameter β and in particular the

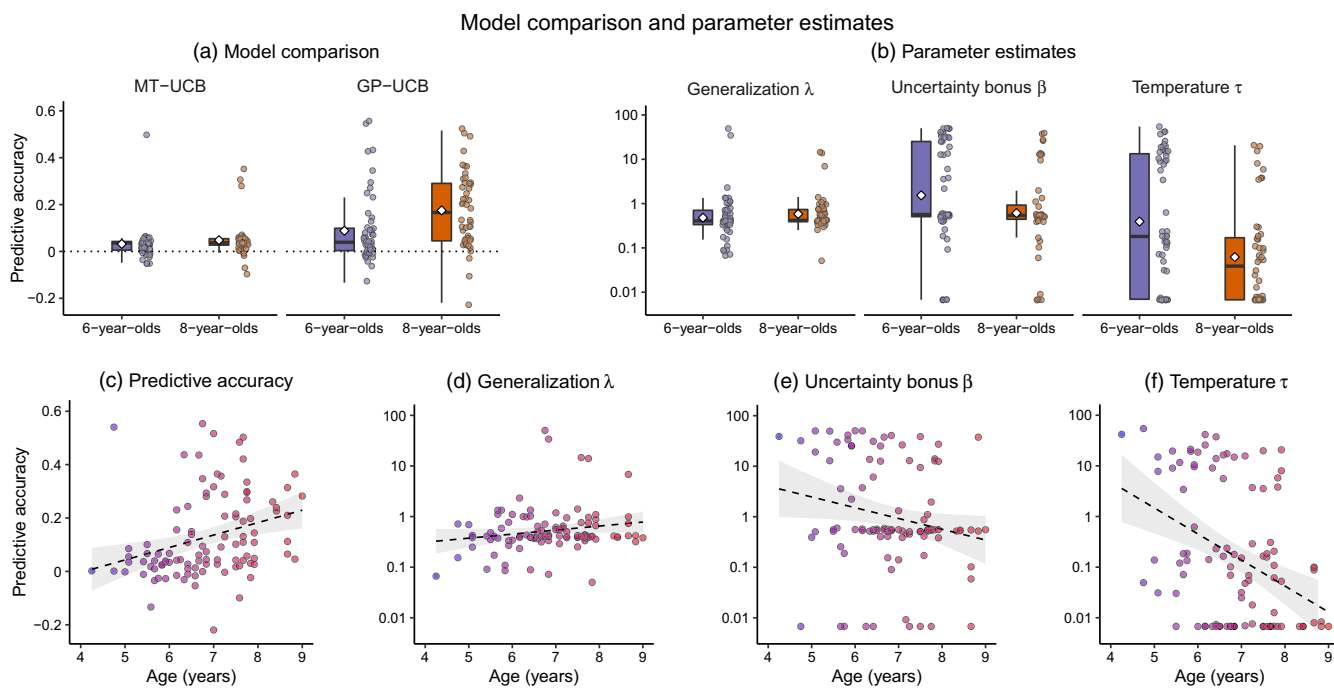


FIGURE 5 Model comparison and parameter estimates of the GP-UCB model. (a) Predictive accuracy (pseudo- R^2) of mean tracker (MT) and Gaussian process (GP) learning model combined with upper-confidence bound (UCB) sampling. Each dot represents one participant with the mean out-of-sample accuracy across rounds (excluding practice and bonus round). Box shows IQR, the line is the median and the diamond is the mean. (b) Individual parameter estimates of the GP-UCB model by age group. (c) Predictive accuracy of the GP-UCB model as function of age. (d-f) Parameter estimates of the GP-UCB model as function of age. Each dot represents one child with their cross-validated median parameter estimates. Dashed line indicates a linear regression ($\pm 95\%$ CI)

temperature parameter τ of the softmax function decreased with age. Younger children tended to have higher values of β ($r_\tau = -.14$, $p = 0.043$, $BF = 1.0$), indicating a somewhat larger value placed on reducing uncertainty, and thus more directed exploration. Whereas the age-related change in directed exploration were rather weak, there was a marked decrease in the temperature parameter τ ($r_\tau = -.23$, $p < 0.001$, $BF = 46$). Thus, the amount of random sampling decreased with age. These same changes in parameters as a function of age also hold when controlling for the predictive accuracy of the GP-UCB model (see Figure B2 and Table B2 in Appendix B), although these analyses find a slightly stronger increase in λ as a function of age, indicating broader generalizations as children grow older. We additionally analyzed parameter estimates for β and τ separately for children best accounted for by the GP-UCB and MT-UCB model, respectively (Appendix C). The same qualitative trends for β and τ were obtained within both subgroups as in the overall analysis, with a strong decrease for the random exploration parameter τ and weaker age-related differences for the directed exploration parameter β . This was the case regardless of whether children's behavior was overall better described by the GP-UCB model or the MT-UCB model. Thus, the overall trends do not result from aggregating across subgroups with qualitatively different exploration strategies.

Taken together, these analyses provide a window into the developmental trajectories of exploration behavior, showing how both directed and, in particular, random exploration decrease as children get older.

5.4 | Parameter estimates and performance

The extent of generalization λ was positively correlated with performance in both age groups (6-year-olds: $r_\tau = .19$, $p = 0.041$, $BF = 1.4$; 8-year-olds: $r_\tau = .25$, $p = 0.011$, $BF = 4.4$). The stronger correlation for 8-year-olds suggests that, compared to 6-year-olds, they were better able to use generalization about the spatial correlation of rewards to achieve higher performance.

Both the uncertainty bonus β (6-year-olds: $r_\tau = -.26$, $p = 0.005$, $BF = 8.7$; 8-year-olds: $r_\tau = -.29$, $p = 0.003$, $BF = 13$) and the random exploration parameter τ (6-year-olds: $r_\tau = -.38$, $p < 0.001$, $BF > 100$; 8-year-olds: $r_\tau = -.28$, $p = 0.005$, $BF = 8.8$) were negatively correlated with performance, showing how too much exploration can hurt performance within the demands of the experiment. For directed exploration, high values of β can lead to excessive exploration at the cost of forgoing options with high expected rewards, and is a direct outcome of the explore-exploit trade-off defined by UCB sampling (Eq. 2). For random exploration, the higher the temperature τ , the more behavior tends toward random choice and random performance, regardless of the learning mechanisms (GP vs. MT) or the UCB trade-off between exploitation and exploration that enter the softmax choice rule (Equation 3).

Thus, one key mechanism underlying the age-related performance differences is that younger children were characterized by higher levels of both directed and random exploration compared to

8-year-olds, who therefore were better able to harness the spatial correlation of rewards in the environment. Yet the optimal level of exploration ultimately depends on the demands of the environment, particularly the available time horizon. Over long horizons, high initial exploration can pay dividends when there are ample opportunities for exploration down the road. Our participants may have been better calibrated to the long-horizon of their lifespan, than the short-horizon of our task.

Differences in exploration also allow us to explain some age-related differences in performance. The observed performance differences in 6-year-olds between smooth and rough environments can be at least partially attributed to differences in the amount of random exploration in the two types of environment, since too much random exploration typically hurts performance in structured environments (Schulz et al., 2019; Wu, Schulz, Garvert, et al., 2020). Although they did not adapt their search trajectory to the same extent as 8-year-olds did (Figure 3b), 6-year-olds still achieved better performance in smooth compared to rough environments (Figure 2a). Indeed, 6-year-olds showed a higher amount of random exploration (i.e., higher temperature parameter τ in rough compared to smooth environments ($Md_\tau = 2.02$ vs $Md_\tau = 0.11$), although the difference was not statistically reliable ($U = 268$, $p = 0.112$, $r_\tau = .18$, $BF = .67$). Eight-year-olds also showed slightly higher levels of random exploration in rough environments, but the difference was less pronounced than for 6-year-olds ($Md_\tau = 0.05$ vs $Md_\tau = 0.02$; $U = 229$, $p = 0.245$, $r_\tau = -.14$, $BF = .59$). However, the analysis of the search trajectories shows that 8-year-olds were generally better able to adapt their search trajectories to the structure of the environment (e.g., by searching more locally in smooth environments), helping them to better exploit the correlation between rewards.

6 | GENERAL DISCUSSION

We investigated how 6- and 8-year-old children search for rewards in a spatial version of the explore-exploit dilemma, focusing on disentangling how generalization, random exploration, and directed exploration contribute to age-related changes. Although general performance increased with age, we found that even younger children could successfully generalize the observed spatial correlations and use this knowledge to guide their search for rewards. Children adapted their exploration behavior depending on the rewards they obtained, with 8-year-olds showing a stronger relationship between obtained rewards and search distance. Finally, while prediction accuracy in the bonus round increased with age, there was no relation between children's confidence and their prediction error.

The model-based analyses showed that the GP-UCB model provided a better account of children's behavior than the MT-UCB model, highlighting the importance of similarity-based generalization. A key finding is a strong age-related decrease of random exploration, represented by the τ parameter of the softmax choice rule, consistent with the hypothesis that children's temperature "cools off" as they get older (Gopnik et al., 2017).

However, children's exploration behavior was not solely driven by random exploration, but also by a high amount of uncertainty-directed sampling, as indicated by high levels of the uncertainty-bonus parameter β . The valuation of uncertainty also tended to decrease with age, but this trend was much weaker compared to the tapering off of random exploration.

Our findings extend the developmental investigation of children's exploration behavior, complementing previous research with older children (Schulz et al., 2019), as well as adolescent and adult participants, who also show signatures of both types of exploration strategies (Wilson et al., 2014; Wu et al., 2018). Table 1 provides an overview of children and adults' model parameters across different studies using similar versions of the multi-armed spatially correlated bandit paradigm. The comparison shows that children up to around age 11 show higher levels of directed exploration than adult subjects, whereas adults tend to generalize more strongly. High levels of random exploration were only observed in 6-year-olds, indicating that this form of exploration diminishes earlier in development than uncertainty-guided exploration. Future studies should systematically investigate an even broader age range (e.g., from childhood through adolescence to adulthood, ideally in a longitudinal design) to identify changes in exploration and generalization over the lifespan.

Children are keen explorers—but are they good exploiters? One peculiar finding we obtained was the low number of exploitation decisions (i.e., repeat clicks; Figure 3b). Across all children and rounds (excluding tutorial and bonus round), the proportion of repeat clicks was about 7% (6-year-olds: 6.8%, 8-year-olds: 7.5%). While this proportion was comparable to participants in a similar age range as reported in other studies (e.g., Schulz et al., 2019, reported 5.6% repeat clicks for 7- to 8-year-olds and 6.4% for 9- to 11-year-olds), this contrasts with the behavior of adults, who typically show a higher proportions of repeat clicks; 12% in Wu et al. (2018, averaged across three experiments) and 32.1% in the study by Schulz et al. (2019). Lower exploitation rates for children have also been observed in simpler bandit tasks with fewer options and independent reward distributions (Blanco & Sloutsky, 2019).

The tendency to over-explore might be responsible for the decrease of children's average rewards toward the end of the search horizon (Figure 2c). Indeed, given a fixed search horizon, it is typically better at some point to start exploiting the found high-reward options, rather than keeping on searching for even better options. It is likely that this behavior was driven by the high amount of both random and directed exploration, as captured by a high temperature parameter τ , leading to increased random sampling, and a high uncertainty bonus β , leading children to optimistically inflate expected rewards of unobserved tiles. While this tendency to over-explore impaired performance in our task, it may nevertheless be adaptive in some settings (Sumner et al., 2019), by allowing children to discover changes that are not obvious and are overlooked by adults (Gopnik et al., 2015; Lucas et al., 2014). It could be especially adaptive in dynamic environments where reward structures change over time (Behrens et al., 2007; Speekenbrink & Konstantinidis, 2015). In such nonstationary environments, previously rewarding options may no longer be valuable at a later point in time, thereby benefiting continuous exploration.

Another factor contributing to the drop-off in performance toward the end of rounds might be genuine curiosity about discovering the rewards associated with unrevealed tiles. Indeed, anecdotal evidence suggests that some children were prone to an additional "novelty bonus" based on expressions of excitement when revealing new tiles. One interpretation of this is that the objectively same reward from a known tile might be valued less compared to the very same reward obtained from revealing a novel tile. This is in line with various theories of curiosity (Berlyne, 1950; Dubey & Griffiths, 2019; Gottlieb & Oudeyer, 2018; Gottlieb et al., 2013; Kidd & Hayden, 2015) that posit intrinsic rewards from novel stimuli, which have been linked to the lifespan goal of self-development (Lopes & Oudeyer, 2012). In the present study, curiosity and the intrinsic reward signals associated with revealing new tiles might have contributed to the declining performance toward the end of the search round. Despite the instructed aim of the experiment,

TABLE 1 Comparison of predictive accuracy and GP-UCB parameter estimates across different studies with children and adults, using the spatially correlated multi-armed bandit paradigm

Age group	Accuracy R^2	Generalization λ	Uncertainty bonus β	Randomness τ
Current study				
6-year-olds (N = 54)	0.09	0.41	0.57	0.18
6-year-olds (N = 54)	0.18	0.42	0.54	0.04
Schulz et al. (2019)				
6-year-olds (N = 54)	0.17	0.44	0.51	0.01
6-year-olds (N = 54)	0.26	0.53	0.50	0.02
6-year-olds (N = 54)	0.39	0.83	0.24	0.03
Wu et al. (2018)				
6-year-olds (N = 54)	0.26	0.74	0.40	0.03

Note: R^2 is the mean predictive accuracy of the GP-UCB model. Model parameters λ , β , and τ are the median values of the cross-validated estimates. We report the mean across three experiments from Wu et al. (2018), which used both 1D (Exp. 1) and 2D spatially correlated bandits (Exps. 2 and 3), with similar smooth and rough environments (Exp. 1 and 2) or natural environments defined by agricultural data (Exp. 3).

children may find it more rewarding to try out novel options, even to the detriment of accumulating higher rewards through exploiting known options. Therefore, an important avenue for future research is to integrate theories and models of curiosity with generalization, directed exploration, and random exploration (Brändle et al., 2020).

Another critical question for future research concerns the representation of uncertainty in learning and exploration. In our task, the spatial correlation of rewards favors a more complex representation of uncertainty structured around generalization, but in other tasks simpler representations of uncertainty may provide a better account. For instance, count-based exploration strategies operate on simpler representations of uncertainty solely based on the number of experiences with a certain stimulus (e.g., the number of times a tile has been visited; Bellemare et al., 2016; Cogliati Dezza et al., 2019). This representation of uncertainty can be used to implement a variant of the GP-UCB model, where the posterior uncertainty $\sigma(\mathbf{x})$ is replaced with a count-based representation of uncertainty (Supplement S1). Exploratory analyses with a GP count-based model with our data suggest promising results (Supplement S2), yet also present a crucial limitation. Specifically, the uncertainty estimates of the count-based model are decoupled from the generalization component, producing identical uncertainty estimates for all unobserved options. This holds for both near and distant options, disregarding the level of spatial proximity to previous observations. This is also the case for time-based representations, where uncertainty is assumed to increase the longer an option has not been chosen (Blanco & Sloutsky, 2019). In this sense, the count-based account is similar to the MT model, where both the estimates of reward and uncertainty are updated only when a tile is observed. When using a count-based representation of uncertainty, reward estimates are influenced by generalization, but not the uncertainty of rewards which is solely a function of previous visits. By contrast, the GP-UCB model generalizes both reward expectations and attached uncertainty by exploiting the correlation structure of rewards in the environment. In fact, research with adults has shown that confidence judgments are systematically related to the uncertainty estimates predicted by the GP (Wu, Schulz, Garvert, et al., 2020; Wu, Schulz, & Gershman, 2020), as opposed to being uniform across all unobserved options. (We observed a similar relation for 8-year-olds in our study, but the data were rather noisy, so a cautious interpretation is warranted; see Appendix B). Future research should contrast different representations of uncertainty in their ability to predict children's and adults' confidence judgments about expected rewards of novel options, to gain a better understanding of possible developmental trends in the representation of uncertainty across the lifespan.

7 | CONCLUSIONS

Our study provides important new insights into the developmental origins and trajectory of learning and exploration, revealing some of its underlying computational principles. Being able to disentangle the role of generalization, and directed versus random exploration

enriches our understanding of how children learn about the world they live in (Buchsbaum et al., 2011; Gopnik et al., 2001) and the people they interact with (Bridgers et al., 2019; Jara-Ettinger et al., 2016). It is also important to extend this computational approach to investigate the exploration behavior of even younger preschoolers, toddlers, and infants, to identify a more comprehensive developmental trajectory and potentially account for individual differences. Finally, connecting this line of work with the growing body of research and theories on curiosity promises to bring us one step closer to identifying the key to children's impressively successful early learning.

ACKNOWLEDGMENTS

We thank all families who participated in this research, Calvin Paulus and Jeanette Blümel for collecting the data, and Federico Meini for help with programming the experiment. CMW is supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A and funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2064/1 – 390727645.

CONFLICT OF INTEREST

The authors declare that they had no conflicts of interest with respect to their authorship or the publication of this article.

DATA AVAILABILITY STATEMENT

All data and code for reproducing the analyses is available at <https://osf.io/eq2bk/>.

ORCID

Björn Meder  <https://orcid.org/0000-0002-9326-400X>

Charley M. Wu  <https://orcid.org/0000-0002-2215-572X>

Eric Schulz  <https://orcid.org/0000-0003-3088-0371>

Azzurra Ruggeri  <https://orcid.org/0000-0002-0839-1929>

REFERENCES

- Auer, P. (2002). Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3, 397–422.
- Behrens, T. E., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, 10, 1214–1221.
- Bellemare, M., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., & Munos, R. (2016). Unifying count-based exploration and intrinsic motivation. *Advances in Neural Information Processing Systems*, 1471–1479.
- Bellman, R. (1952). On the theory of dynamic programming. *Proceedings of the National Academy of Sciences of the United States of America*, 38, 716–719.
- Berlyne, D. E. (1950). Novelty and curiosity as determinants of exploratory behaviour. *British Journal of Psychology*, 41, 68–80.
- Betsch, T., Lehmann, A., Lindow, S., Lang, A., & Schoemann, M. (2016). Lost in search: (Mal-)adaptation to probabilistic decision environments in children and adults. *Developmental Psychology*, 52, 311–325.
- Blanco, N. J., Love, B. C., Ramscar, M., Otto, A. R., Smayda, K., & Maddox, W. T. (2016). Exploratory decision-making as a function of lifelong experience, not cognitive decline. *Journal of Experimental Psychology: General*, 145, 284–297.

- Blanco, N. J., & Sloutsky, V. M. (2019). Systematic exploration and uncertainty dominate young children's choices. *PsyArXiv*. <https://doi.org/10.31234/osf.io/72sfx>
- Bonawitz, E., Denison, S., Griffiths, T. L., & Gopnik, A. (2014). Probabilistic models, learning algorithms, and response variability: Sampling in cognitive development. *Trends in Cognitive Sciences*, 18, 497–500.
- Bonawitz, E., van Schijndel, T. J., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology*, 64, 215–234.
- Brändle, F., Wu, C. M., & Schulz, E. (2020). What are we curious about? *Trends in Cognitive Sciences*, 24, 685–687.
- Bridgers, S., Jara-Ettinger, J., & Gweon, H. (2019). Young children consider the expected utility of others' learning to decide what to teach. *Nature Human Behaviour*, 4, 144–152.
- Buchsbaum, D., Gopnik, A., Griffiths, T. L., & Shafto, P. (2011). Children's imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition*, 120, 331–340.
- Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, 80, 1–28.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76, 1–32.
- Cauffman, E., Shulman, E. P., Steinberg, L., Claus, E., Banich, M. T., Graham, S., & Woolard, J. (2010). Age differences in affective decision making as indexed by performance on the Iowa gambling task. *Developmental Psychology*, 46, 193–207.
- Chambers, C. T. (2002). Developmental differences in children's use of rating scales. *Journal of Pediatric Psychology*, 27, 27–36.
- Cogliati Dezza, I., Cleeremans, A., & Alexander, W. (2019). Should we control? The interplay between cognitive control and information integration in the resolution of the exploration-exploitation dilemma. *Journal of Experimental Psychology: General*, 148, 977–993.
- Cohen, J. D., McClure, S. M., & Angela, J. Y. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 362, 933–942.
- Dubey, R., & Griffiths, T. L. (2019). Reconciling novelty and complexity through a rational analysis of curiosity. *Psychological Review*, 127, 455–476.
- Gershman, S. J. (2018). Deconstructing the human algorithms for exploration. *Cognition*, 173, 34–42.
- Gittins, J. C., & Jones, D. M. (1979). A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika*, 66, 561–565.
- Gopnik, A., Griffiths, T. L., & Lucas, C. G. (2015). When younger learners can be better (or at least more open-minded) than older ones. *Current Directions in Psychological Science*, 24, 87–92.
- Gopnik, A., O'Grady, S., Lucas, C. G., Griffiths, T. L., Wente, A., Bridgers, S., Aboody, R., Fung, H., & Dahl, R. E. (2017). Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood. *Proceedings of the National Academy of Sciences of the United States of America*, 114, 7892–7899.
- 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, 620–629.
- Gottlieb, J., & Oudeyer, P.-Y. (2018). Towards a neuroscience of active sampling and curiosity. *Nature Reviews Neuroscience*, 19, 758–770.
- Gottlieb, J., Oudeyer, P.-Y., Lopes, M., & Baranes, A. (2013). Information-seeking, curiosity, and attention: Computational and neural mechanisms. *Trends in Cognitive Sciences*, 17, 585–593.
- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., Couzin, I. D., & the Cognitive Search Research Group. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, 19, 46–54.
- Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The naïve utility calculus: Computational principles underlying commonsense psychology. *Trends in Cognitive Sciences*, 20, 589–604.
- Kidd, C., & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. *Neuron*, 88, 449–460.
- Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220, 671–680.
- Lopes, M., & Oudeyer, P.-Y. (2012). The strategic student approach for life-long exploration and learning. *2012 IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL)*, 1–8.
- Lucas, C. G., Bridgers, S., Griffiths, T. L., & Gopnik, A. (2014). When children are better (or at least more open-minded) learners than adults: Developmental differences in learning the forms of causal relationships. *Cognition*, 131, 284–299.
- Lucas, C. G., Griffiths, T. L., Williams, J. J., & Kalish, M. L. (2015). A rational model of function learning. *Psychonomic Bulletin & Review*, 22, 1193–1215.
- Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. Wiley.
- Ly, A., Verhagen, J., & Wagenmakers, E.-J. (2016). Harold Jeffreys's default Bayes factor hypothesis tests: Explanation, extension, and application in psychology. *Journal of Mathematical Psychology*, 72, 19–32.
- Mata, R., Wilke, A., & Czienskowski, U. (2013). Foraging across the life span: is there a reduction in exploration with aging? *Frontiers in Neuroscience*, 7, 53.
- Meder, B., Mayrhofer, R., & Ruggeri, A. (2020, Aug). Developmental trajectories in the understanding of everyday uncertainty terms. *PsyArXiv*. <https://doi.org/10.31234/osf.io/sfgzu>
- Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., Hausmann, D., Fiedler, K., & Gonzalez, C. (2015). Unpacking the exploration-exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, 2, 191–215.
- Mullen, K., Ardia, D., Gil, D. L., Windover, D., & Cline, J. (2011). DEoptim: An R package for global optimization by differential evolution. *Journal of Statistical Software*, 40, 1–26.
- Rasmussen, C. E., & Williams, C. (2006). *Gaussian processes for machine learning*. MIT Press.
- Ronfard, S., Zambrana, I. M., Hermansen, T. K., & Kelemen, D. (2018). Question-asking in childhood: A review of the literature and a framework for understanding its development. *Developmental Review*, 49, 101–120.
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16, 225–237.
- Ruggeri, A., Markant, D. B., Gureckis, T. M., Bretzke, M., & Xu, F. (2019). Memory enhancements from active control of learning emerge across development. *Cognition*, 186, 82–94.
- Ruggeri, A., Xu, F., & Lombrozo, T. (2019). Effects of explanation on children's question asking. *Cognition*, 191, 103966.
- Schulz, E., & Gershman, S. J. (2019). The algorithmic architecture of exploration in the human brain. *Current Opinion in Neurobiology*, 55, 7–14.
- Schulz, E., Konstantinidis, E., & Speekenbrink, M. (2017). Putting bandits into context: How function learning supports decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44, 927–943.
- Schulz, E., Wu, C. M., Huys, Q. J., Krause, A., & Speekenbrink, M. (2018). Generalization and search in risky environments. *Cognitive Science*, 42, 2592–2620.
- 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, 1561–1572.
- Schulz, L. E. (2015). Infants explore the unexpected. *Science*, 348, 42–43.

- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, 237, 1317–1323.
- Somerville, L. H., Sasse, S. F., Garrad, M. C., Drysdale, A. T., Abi Akar, N., Insel, C., & Wilson, R. C. (2017). Charting the expansion of strategic exploratory behavior during adolescence. *Journal of Experimental Psychology: General*, 146, 155–164.
- Speekenbrink, M., & Konstantinidis, E. (2015). Uncertainty and exploration in a restless bandit problem. *Topics in Cognitive Science*, 7, 351–367.
- Sumner, E., Li, A. X., Perfors, A., Hayes, B., Navarro, D., & Sarnecka, B. W. (2019). The exploration advantage: Children's instinct to explore allows them to find information that adults miss. *PsyArXiv*. <https://doi.org/10.31234/osf.io/h437v>
- Thompson, W. R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25, 285–294.
- van Doorn, J., Ly, A., Marsman, M., & Wagenmakers, E.-J. (2018). Bayesian inference for Kendall's rank correlation coefficient. *The American Statistician*, 72, 303–308.
- van Doorn, J., Ly, A., Marsman, M., & Wagenmakers, E.-J. (2020). Bayesian rank-based hypothesis testing for the rank sum test, the signed rank test, and Spearman's ρ . *Journal of Applied Statistics*, 47, 2984–3006.
- Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, 143, 2074–2081.
- Wu, C. M., Schulz, E., Garvert, M. M., Meder, B., & Schuck, N. W. (2020). Similarities and differences in spatial and non-spatial cognitive maps. *PLOS Computational Biology*, 16, e1008149.
- Wu, C. M., Schulz, E., & Gershman, S. J. (2020). Inference and search on graph-structured spaces. *Computational Brain & Behavior*. <https://doi.org/10.1007/s42113-020-00091-x>
- Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature Human Behaviour*, 2, 915–924.
- Zajkowski, W. K., Kossut, M., & Wilson, R. C. (2017). A causal role for right frontopolar cortex in directed, but not random, exploration. *Elife*, 6, e27430.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section.

How to cite this article: Meder B, Wu CM, Schulz E, Ruggeri A. Development of directed and random exploration in children. *Dev Sci*. 2021;24:e13095. <https://doi.org/10.1111/desc.13095>

APPENDIX A

STATISTICAL ANALYSES

We report both frequentist statistics and Bayes factors (BF) to quantify the relative evidence of the data in favor of the alternative

hypothesis (H_A) over the null hypothesis (H_0). All model specifications and R-code are available online at <https://osf.io/eq2bk/>

Group comparisons

Frequentist tests are reported as t tests for parametric comparisons, and Mann–Whitney U -test or Wilcoxon signed-rank test for nonparametric comparisons. Bayes factors are based on the default two-sided Bayesian t test for either independent or dependent samples, using a Jeffreys–Zellner–Siow prior with its scale set to $\sqrt{2}/2$ (Rouder et al., 2009). All statistical tests are nondirectional as defined by a symmetric prior. Bayes factors for the Mann–Whitney U -test are based on performing posterior inference over the test statistic (Kendall's r_τ), assigning a prior using parametric yoking (van Doorn et al., 2020). Bayes factors for nonparametric comparisons are based on performing posterior inference over the test statistics (Kendall's r_τ for the Mann–Whitney U -test and standardized effect size $r = \frac{Z}{\sqrt{N}}$ for the Wilcoxon signed-rank test), assigning a prior using parametric yoking (van Doorn et al., 2020). The posterior distribution for Kendall's r_τ or the standardized effect size r yields a Bayes factor via the Savage–Dickey density ratio test, where the null hypothesis posits that parameters do not differ between groups and the alternative hypothesis posits an effect and assigns an effect size using a Cauchy distribution with the scale parameter set to $1/\sqrt{2}$.

Correlations

Linear correlations are tested with Pearson's r , the corresponding Bayesian test is based on Jeffrey's test for linear correlation assuming a shifted, scaled beta prior distribution $B\left(\frac{1}{k}, \frac{1}{k}\right)$ for r , where the scale parameter is set to $k = \frac{1}{3}$ (Ly et al., 2016). For testing rank correlations with Kendall's tau, the Bayesian test is based on parametric yoking to define a prior over the test statistic (van Doorn et al., 2018). Bayesian inference is performed to compute a posterior distribution for r_τ , and the Savage–Dickey density ratio test is used to produce an interpretable Bayes Factor.

Bayesian multilevel regressions

Regression analyses were performed in a Bayesian framework with Stan (Carpenter et al., 2017), accessed via R-package brms (Bürkner, 2017). In all models, participants were treated as a random intercept, the remaining predictors were implemented as population-level (“fixed”) effects. For population-level effects, we used a normal prior with a mean of 0 and standard deviation of 10; for group-level (“random”) effects, we used a half student- t prior with 3 degrees of freedom, a mean of 0, and a scale parameter of 10; for the intercept a student- t prior with 3 degrees of freedom, a mean of 1, and a scale parameter of 10. All models were estimated over four chains of 4000 iterations, with a burn-in period of 1000 samples.

APPENDIX B

BAYESIAN REGRESSION ANALYSES

Search distance as function of reward on previous step

We ran separate regression analyses for each environment to assess the influence of reward obtained at trial t on search distance at $t + 1$, with population-level (“fixed”) effects for previous reward, age group, and their interaction, and by-participant random intercepts. Figure 3c illustrates the population-level effects; Table B1 provides a summary

TABLE B1 Bayesian regression results: search distance as function of reward on previous step

Predictor	Rough environment		Smooth environment	
	Estimate	95% HDI	Estimate	[95% HDI]
Intercept	2.26	[1.90–2.63]	2.89	[2.6–3.19]
Previous reward	–0.01	[–0.02––0.01]	–0.03	[–0.03––0.02]
Age group	1.85	[1.31–2.34]	1.19	[0.73–1.64]
Previous reward × age group	–0.04	[–0.05––0.03]	–0.03	[–0.04––0.02]
Random effects				
σ^2	0.48		0.29	
τ_{00}	4.84		4.14	
N	50		52	
Observations	5000		5200	
Bayesian R^2	0.16		0.13	

Note: Both models were implemented in *brms* (Bürkner, 2017). We report the posterior mean estimates for the coefficients, followed by a 95% uncertainty interval in brackets (“highest density interval”, *HDI*). σ^2 indicates the individual-level variance and τ_{00} indicates the variation between individual intercepts and the average intercept. For categorical variable age group, 6-year-olds are the reference level.

TABLE B2 Bayesian regression results: prediction error and confidence

Predictor	Estimate	95% HDI
Intercept	13.72	[10.04–17.51]
Confidence	0.35	[–0.09–0.77]
Age group	–2.12	[–7.57–3.29]
Confidence × age group	–0.38	[–1.07–0.30]
Random effects		
σ^2	25.09	
τ_{00}	81.36	
N	102	
Observations	510	
Bayesian R^2	0.3	

Note: The model was implemented in *brms* (Bürkner, 2017). We report the posterior mean estimates for the coefficients, followed by a 95% uncertainty interval in brackets (“highest density interval”, *HDI*). σ^2 indicates the individual-level variance and τ_{00} indicates the variation between individual intercepts and the average intercept. For variable age group, 6-year-olds are the reference level.

of the results. For both environments, these analyses showed an effect of previously obtained reward on search distance (i.e., lower rewards lead to higher subsequent search distances), an effect of age group (i.e., 8-year-olds showed higher search distances overall), and an interaction (i.e., the search distance of 8-year-olds was stronger influenced by obtained rewards than that of 6-year-olds).

Judgments

In the bonus round, children made reward predictions for five previously unseen tiles and rated their confidence in their predictions. To assess the relation between prediction error (mean absolute deviation between judged and true reward value) and confidence we ran a Bayesian linear regression with prediction error as dependent variable, and confidence, age group and their interaction as population-level (“fixed”) effects, and a random intercept for participants. Children's confidence judgments were elicited using an 11-point (0–10) slider with the endpoints labeled as “not at all” and “very sure.”

Table B2 provides a summary of the results; Figure B1 show the population-level (fixed) effects of the model, excluding the group-level effects (random intercepts over participants). These data show no systematic relation between children's subjective confidence in their predictions.

Regression analyses for age-related trends in parameter estimates

To control for the effect of predictive accuracy R^2 on the age-related changes in the GP-UCB parameter estimates, we ran

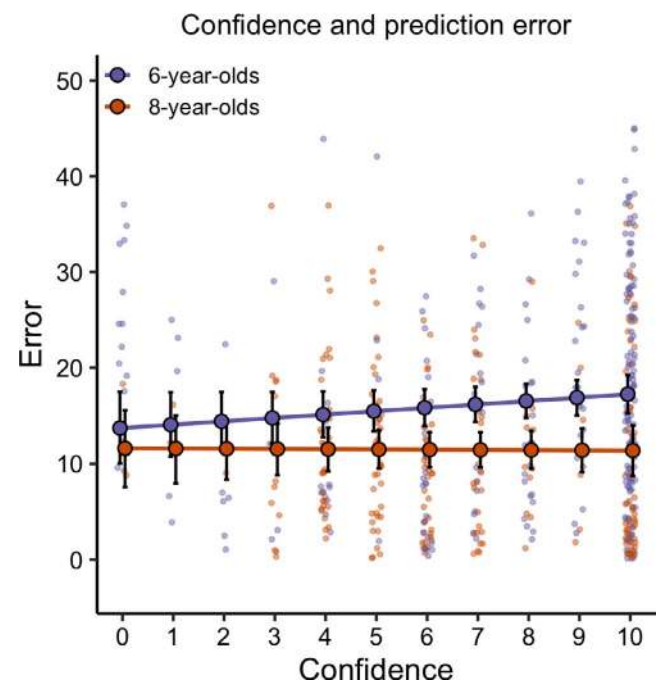


FIGURE B1 Confidence and prediction error in the bonus round. The lines visualize the expected values of the posterior predictive distribution of a Bayesian regression ($\pm 95\%$ CI); the dots show the raw data

TABLE B3 Bayesian regression results: parameter estimates with age and R^2 as predictors

Predictor	Generalization λ (log)		Uncertainty bonus β (log)		Temperature τ (log)	
	Estimate	95% HDI	Estimate	95% HDI	Estimate	95% HDI
Intercept	-2.83	[-4.55--1.12]	2.98	[-0.77-6.70]	3.76	[0.21-7.56]
Age (in months)	0.03	[0.01-0.05]	-0.03	[-0.08-0.02]	-0.05	[-0.10--0.01]
R^2	5.49	[-3.05-13.95]	-7.75	[-22.98-7.55]	-5.94	[-20.91-8.96]
$R^2 \times$ Age (in months)	-0.08	[-0.18-0.03]	0.04	[-0.14-0.22]	-0.06	[-0.24-0.13]
Observations	102		102		102	
Bayesian R^2	0.08		0.13		0.69	

Note: All models were implemented in brms (Bürkner, 2017). We report the posterior mean estimates for the coefficients, followed by a 95% uncertainty interval in brackets ("highest density interval", HDI).

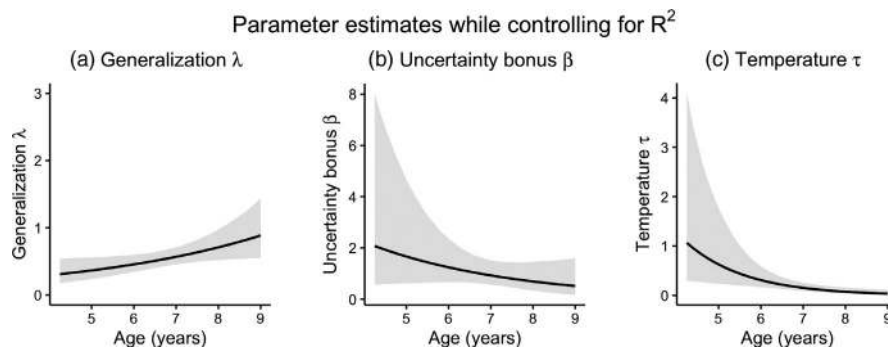
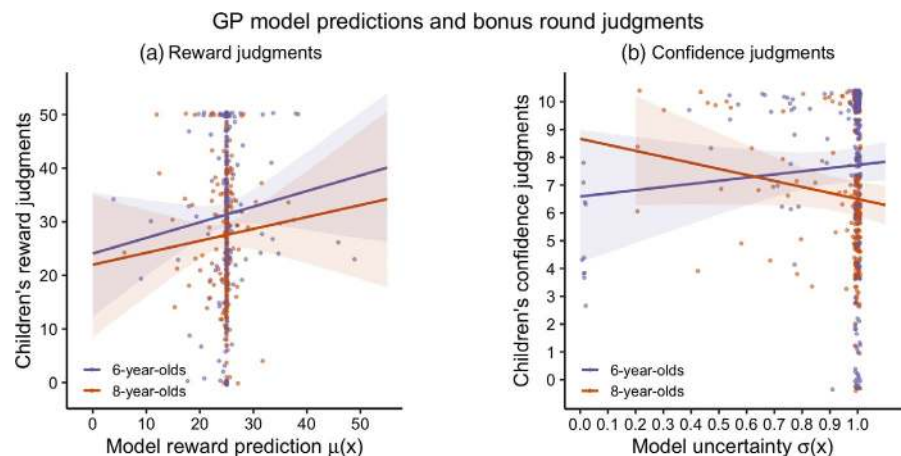


FIGURE B2 Effect of age on GP-UCB parameters, derived from a Bayesian regression with age (in months), individual model R^2 , and their interactions, as predictor for the (log-transformed) median parameter estimates. For plotting we converted the regression models' predictions back to the original scale by exponentiating the parameter estimates, such that all parameters are non-negative

FIGURE B3 GP model predictions for judgments. The lines visualize the means of the posterior predictive distribution of the Bayesian regression ($\pm 95\%CI$); the dots show the raw data points. (a) Relation between GP model predictions of reward and children's reward judgments. (b) Relation between GP model uncertainty about expected rewards and children's confidence about their reward judgments



regression analyses for each parameter with age (in months), individual R^2 , and their interaction as predictors for the individual median parameter estimates. Since λ , β , and τ are defined as non-negative, we log-transformed them for the regressions; for plotting the influence of age on parameters we converted the regression models' predictions back to the original scale by exponentiating them, such that all parameters are non-negative. Table B3 shows the results of the regression analyses; Figure B2 visualizes the effects of age on the GP-UCB parameter estimates while taking into account R^2 .

GP model predictions and judgments of reward and confidence

We assessed the relation between GP model predictions and participant judgments about expected reward and confidence in the bonus round. In the bonus round, participants selected 15 tiles and then made reward predictions for five unseen tiles and judged their confidence in their predictions. The MT model, which learns independent reward distributions, makes identical predictions for all unseen tiles, as it does not generalize. By contrast, the GP model makes specific predictions for novel options, taking into account the data obtained so far and the spatial correlation of the search ecology.

TABLE B4 Bayesian regression results: GP model predictions and bonus round judgments

Predictor	Reward judgments		Confidence judgments	
	Estimate	95% HDI	Estimate	95% HDI
Intercept	24.08	[12.43–35.46]	6.59	[4.23–8.99]
GP predictions	0.29	[–0.15–0.75]	1.14	[–1.32–3.55]
Age group	–2.09	[–16.13–12.33]	2.08	[–1.19–5.42]
GP predictions × age group	–0.07	[–0.65–0.51]	–3.3	[–6.73–0.03]
Random effects				
σ^2	32.69		3.49	
τ_{00}	168.25		4.53	
N	102		102	
Observations	510		510	
Bayesian R^2	.19		.49	

Note: Both models were implemented in brms (Bürkner, 2017). We report the posterior mean estimates for the coefficients, followed by an 95% uncertainty interval in brackets ("highest density interval", HDI). σ^2 indicates the individual-level variance and τ_{00} indicates the variation between individual intercepts and the average intercept. For categorical variable age group, 6-year-olds are the reference level.

For each participant, we used parameters estimated from rounds 2 to 5 in order to generate individual GP model predictions (estimated mean reward and variance) for the five randomly selected tiles in the bonus round. These predictions were conditioned on the 15 individual choices and observations made by each child and were generated using each individuals' median λ estimates. This represents a type of out-of-task prediction, where we used parameters estimated from search decisions to prediction out-of-sample judgments. We use the mean reward predictions of the GP model (posterior $\mu(\mathbf{x})$ of tile) as a prediction for each child's judgment about expected reward and the GP's uncertainty estimates (posterior σ) as a prediction of each child's confidence judgments, where we treat uncertainty as the inverse of confidence.

GP predictions were somewhat correlated with participant predictions ($r_{\tau} = .08$, $p = 0.013$, $BF = 1.5$), although this disappeared when separating participants into age groups (6-year-olds: $r_{\tau} = .06$, $p = 0.182$, $BF = .22$; 8-year-olds: $r_{\tau} = .08$, $p = 0.054$, $BF = .57$). GP uncertainty estimates were negatively correlated with confidence for 8-year-olds ($r = -.18$, $p = 0.005$, $BF = 7.5$), but not for 6-year-olds ($r = .06$, $p = 0.330$, $BF = .23$). This suggests that the confidence judgments of 8-year-olds were somewhat accounted for by the GP model, but not those of 6-year-olds.

To analyze these findings in more detail, we conducted Bayesian regression analyses to predict children's reward and confidence judgments based on the outputs of the GP model. Specifically, we used GP model predictions, age group, and their interaction as population-level ("fixed") effects, and by-participant random intercept (Table B4). In the first model (*Reward judgments*), participant reward judgments in the range [0,50] for novel options \mathbf{x} (tiles) were predicted from the GP posterior means of rewards, $\mu(\mathbf{x})$. The second model (*Confidence judgments*) used the GP posterior uncertainty, $\sigma(\mathbf{x})$ to predict children's confidence judgments in the range [0,10]. All GP predictions were computed based on individual participant λ -values and the 15

search decisions they made prior to providing their judgments for five random novel options.

Table B2 provides a summary of the results; Figure B3 visualizes the population-level (fixed) effects of the model, excluding the group-level effects (random intercepts over participants). The results show a positive but rather weak relation between the GP model's reward predictions and children's reward judgments about unobserved tiles (Figure B3a). The trends for the relation between model uncertainty and children's confidence judgments mirror the overall correlations. For 6-year-olds, there's a weak relation in the wrong direction (i.e., they tend to be more confident when the GP model is more uncertain). By contrast, for 8-year-olds there is a fairly strong trend in that children's confidence declined with increasing model uncertainty. However, the raw data are very noisy and unevenly distributed, so a cautious interpretation of these results is warranted.

APPENDIX C

SUBGROUP ANALYSES

The majority of participants (73 out of 102 children), were best predicted by the GP-UCB model. The proportion of children best described by the GP-UCB model was somewhat lower for 6-year-olds (63%) than for 8-year-olds (81%), raising the possibility that the observed developmental trends in the parameters representing directed and random exploration are due to aggregating across subgroups exhibiting qualitatively different patterns of parameter estimates.

Figure C1 shows the β parameter representing uncertainty-directed exploration and the τ parameter representing random exploration separately for children best predicted by the GP-UCB and MT-UCB model, respectively. These are the two components shared by the two models, therefore one should expect similar trends as for the overall analyses.

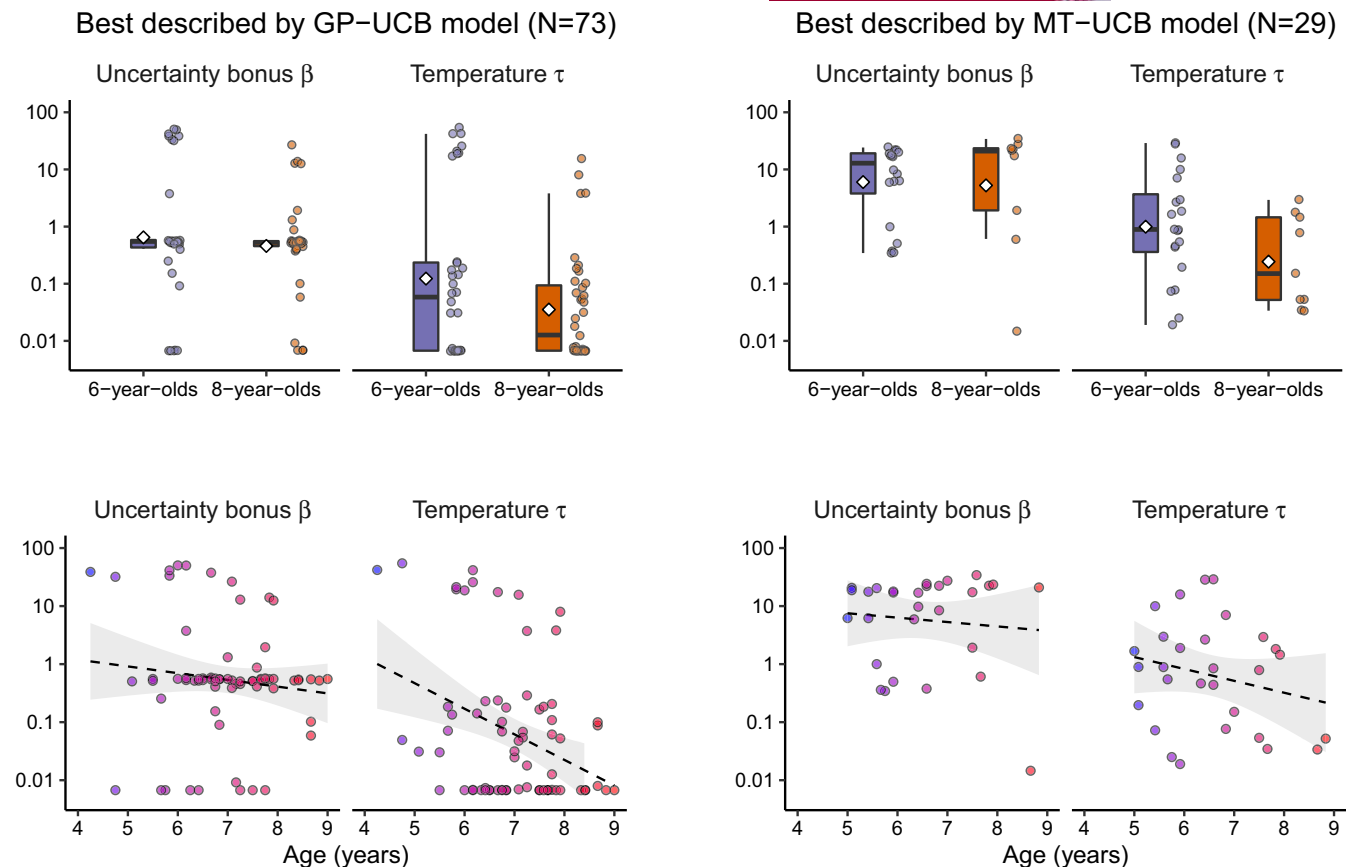


FIGURE C1 Model parameters separately for children best accounted for by the GP-UCB model and MT-UCB-model

In fact, the qualitative pattern for the two parameters mirror the overall analyses, with weak age-related differences for β and stronger difference for τ . This pattern hold regardless of whether children's behavior was overall better described by the GP-UCB model or the MT-UCB model. Importantly, for children best predicted by GP-UCB model, the amount of random exploration decreased as a function of age, $r_{\tau} = -.18$, $p = 0.03$, $BF = 1.6$. However, while the other trends show the same age-related pattern as the overall analyses, the other comparisons within these subgroups were not statistically reliable. Nevertheless, the persistence of the observed developmental trajectories on the subgroup level refutes the possibility that the overall trends are due to aggregating across subpopulations with qualitatively different parameter estimates.

APPENDIX D

INSTRUCTIONS

The experiment was implemented on a tablet, where children could touch the screen to select tiles. Below are screenshots from the tutorial (translated from German); example screenshots from the task

are shown in Figure 1b,c. To account for individual and developmental differences in reading ability, the experimenter always read out loud the instructions displayed on screen. Additional information was given verbally during the tutorial. For instance, after explaining the game and the goal (i.e., selecting tiles to collect stars; top left) and before practicing the search for rewards (top right), children were told that before each choice they would have to decide whether they would like to reveal a novel tile or re-click a previously revealed tile. The experimenter demonstrated both actions before the child completed the tutorial round. After the tutorial round and prior to the instruction test (bottom right) participants could also ask questions to clarify the given task and instructions (Figure D1).



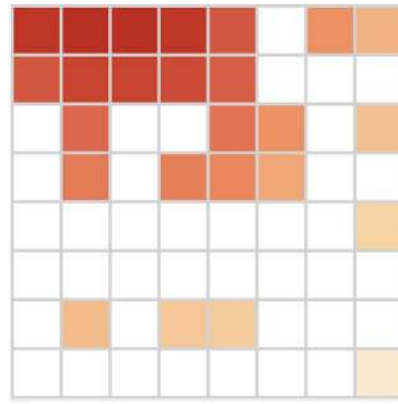
In this game I will show you 6 such grids. At the beginning, there is always one tile that has already been revealed.

The goal of the game is to collect as many stars as possible, by finding as many dark tiles as possible. The darker a tile is, the more stars you get. The more stars you get, the more stickers you will receive at the end of the game.

If you want to reveal another tile, you can click on it and the color of that tile will be shown. Below you see the colors you can observe.

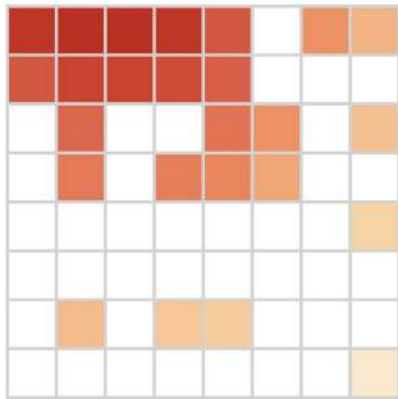
In each round, you have 25 clicks, to find as many dark tiles as possible!

Continue



Before we start, here's a hint: On the grids, dark tiles are frequently located close to other dark tiles, and light tiles tend to be close to other light tiles. Thus, the color of a tile depends on where it is located on the grid.

Continue



Great!
You did a great job!
In the next rounds you can collect up to 5 stars on each grid. The more dark tiles you find, the more stars you will get. The more stars you get, the more stickers you will receive at the end of the game?

Do you have any questions?

Continue



Please answer a few questions before we start with the game.

What is your task?

- Learn colors
- Find as many dark tiles as possible
- Find the darkest tile
- No idea

How can you collect stars?

- Only by reaving new, white tiles
- Only by re-clicking revealed, colored tiles
- By clicking new, white tiles or by re-clicking revealed, colored tiles

How are the dark tiles distributed?

- Randomly
- Dark tiles are **never** close to each other
- Dark tiles are **always** in close to each other
- Dark tiles are **frequently** close to each other

Only when you correctly answer all questions we can start with the game.

Check answers

FIGURE D1 Screenshots and instruction test from tutorial