Children Spontaneously Design Curricula to Tackle Challenging Tasks

Annya Dahmani* (adahmani@berkeley.edu)
Department of Psychology, 2121 Berkeley Way
Berkeley, CA 94704 USA

Eunice Yiu* (ey242@berkeley.edu)
Department of Psychology, 2121 Berkeley Way
Berkeley, CA 94704 USA

Alison Gopnik (gopnik@berkeley.edu)
Department of Psychology, 2121 Berkeley Way
Berkeley, CA 94704 USA

Abstract

We study how children develop a causal curriculum to achieve a challenging goal that is not solvable at first. Adopting the Procgen environments that include various challenging game tasks, we found that 5- to 7-year-old children actively used their current level competence to determine their next step in the curriculum and made improvements to their performance during this process as a result. This suggests that children treat their level competence as an intrinsic reward, and are motivated to master easier levels in order to do better at a more difficult one, even without explicit reward. However, our findings also suggest that children's self-designed curricula may not always be the most effective design. Rather, repeatedly practicing on the difficult target task may be sufficient. Notably, when constrained to stay on the target task instead of crafting their own curricula, more children actually succeeded and made greater progress in the game, suggesting that children perceive a curriculum as beneficial, even when focusing on a singular difficulty might prove more effective.

Keywords: curriculum learning; exploration; cognitive development

Introduction

Humans are remarkable learners, especially when they are faced with challenging tasks. We are able to craft personalized curricula that shape our experiences in ways that allow us to aquire new knowledge and skills (Alfieri et al., 2011; Baranes et al., 2014; Gopnik et al., 1999; Khan et al., 2011; Loyens et al., 2008; Poli et al., 2020; Schmidt et al., 2007; Schulz, 2012; Ten et al., 2021).

The question arises: How do humans learn to reason about their own learned capabilities and use this information to bootstrap the future knowledge they need to learn to overcome their current limitations? Causal learning may be crucial for enabling us to efficiently explore various levels of task difficulty and complexity within an environment (Jiang et al., 2023; Ke et al., 2021; J. X. Wang et al., 2021). In particular, first mastering the causal relations that are involved in a simpler task can allow humans and artificial agents to solve more complex tasks that involve similar relations (Florensa et al., 2018; Jiang et al., 2023; Ke et al., 2021; Schölkopf et al., 2021; J. X. Wang et al., 2021). We build causal models to help guide our exploration, rather than randomly varying policies and observing the results. Causal models are well-designed precisely to afford a wide range of novel actions and interventions in the world (Glymour et al., 1991; Pearl, 2000; Peters et al., 2017; Spirtes et al., 2000). The ability to collect data from causal interventions can allow an agent to construct a new causal model, leverage that model to make decisions, and repeat this process for improvement (Jiang et al., 2023; Ke et al., 2021; Schölkopf et al., 2021; J. X. Wang et al., 2021).

Related work in cognitive science suggests that child learning follows the "Goldilocks" principle: children opt for information that is neither too easy nor too complex, but moderately predictable (Cubit et al., 2021; Kidd et al., 2012, 2014). Furthermore, children are able to monitor the "zone of proximity" between their current capabilities and the goal at hand. This enables them to progress from what they cannot do at all to what they can learn to do with the assistance of an adult or a teacher (Vgotsky, n.d.). Infants allocate their visual attention based on surprise, predictability and learning progress (Poli et al., 2020). 4- to 6-year-old children use their improvement over time to decide whether to persist on a challenging goal on their own (Leonard et al., 2023). And by age 7, children ask questions that yield higher information gain when problems are more difficult (Ruggeri & Lombrozo, 2015). However, there are no systematic studies showing that young children can indeed spontaneously construct an appropriate curriculum in order to master complex goal-directed tasks. In particular, we show that children use intrinsic rewards based on level competence to determine a curriculum.

Research has shown that young children's learning decisions are often scaffolded by external guidance, typically from teachers or adults. For instance, six- to eight-year-old children, rely on teachers' beliefs to calibrate their explore-exploit decisions (Bass et al., 2023). Similarly, four-year-olds exhibit persistent exploratory play and engage in more difficult interventions following success in a task that is framed as "difficult" rather than "easy" by the experimenter (Doan et al., 2020). Furthermore, adults' feedback about children's performance trajectories significantly impacts their motivation: four- to six-year-old children who were told that their performance was improving were more likely to persist on a challenging task than those whose performance was described as constant, even when their final performance outcomes were actually the same (Leonard et al., 2023). These findings highlight the critical role of external pedagogical inputs in shaping children's

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

learning and exploration.

However, it remains unclear whether children can autonomously design effective learning curricula for themselves, without any feedback from others. Some evidence suggests that this is possible. Young children are sensitive to task difficulty. For example, four- and five-year-olds tend to skip tasks that are both difficult and associated with low reward probabilities (J. Wang & Bonawitz, 2023). Additionally, four- to eight-year-olds selectively practice on difficulty levels that they expect to encounter in subsequent testing, and opt to train on the more difficult level when the test difficulty is unknown (Serko et al., 2024). However, these decisions typically involve a simple comparison between two task levels rather than complex curricula with multiple difficulty levels. Can children continuously choose task environments that align with their current level of competence and make further progress? On the one hand, children under seven years of age overestimate their performance instead of anticipating improvement across trials (Zhang et al., 2023); on the other hand, five- to ten-year-olds would design easier goals for themselves when they wree motivated to win (Rule et al., 2023). It is therefore an open question whether children can and will independently craft an effective curriculum to help solve challenging tasks, particularly when the task requires sustained effort and strategic planning across multiple trials.

Term	Definition
Level Competence	An indicator of success in an episode of game play on a specific level, reaching 100% when the specific level is successfully solved in that episode
Global Progress	Difference in competence between the initial attempt and the final attempt at the target level

Table 1: Definitions of terms in the paper

Study 1: Can children design their own curriculum?

To systematically analyze curriculum learning for human players, we selected a subset of Procgen (Cobbe et al., 2020) environments and tailored them by adjusting game difficulty along a single parameter or variable. Procgen is a procedurally generated environment that develops a wide range of reinforcement learning games with varying levels of difficulty. Our experiments focused on three distinct Procgen games: *Leaper*, *Climber*, and *Heist* (Figure 1) We ask how children scaffold their own learning by engaging in varying difficulty levels and analyze their curriculum decisions.

Methods

We recruited 22 five- to seven-year-old children (10 females, 12 males), with a mean age of 5.55 years ($\sigma = 0.74$ years) at public museums in the Bay Area, California, USA. Participants were told that they would be rewarded with a sticker if they won a target game level (the target level was always Level 3 of the game). The target level was intentionally designed to be too difficult to succeed in a single attempt, requiring repeated efforts and incremental learning to achieve success. Children were then given the opportunity to autonomously select different difficulty levels of the game (Levels 1-4, with Level 3 being the target game level). Children could play either until they won the target level, or for up to a total of 10 rounds, whichever came first.

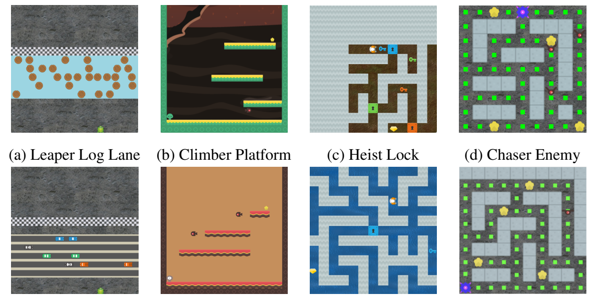

(a) Leaper Log Lane (b) Climber Platform (c) Heist Lock (d) Chaser Enemy
(e) Leaper Car Lane (f) Climber Enemy (g) Heist Maze Size (h) Chaser Maze Size

Figure 1: In Study 1, we selected 3 Procgen environments with 2 different axes each and varied the level of difficulty across the axes within a game. In Study 2, we focused on *Leaper* for children and included a new game *Chaser* for adults. (a)-(h) show the goal levels for the games.

Game	Axis	Level 1	Level 2	Level 3 (target)	Level 4 (extra difficult)
Leaper	Log Lane	1 log lane	3 log lanes	5 log lanes	7 log lanes
Leaper	Car Lane	1 car lane	3 car lanes	5 car lanes	7 car lanes
Climber	Enemy	0 enemies	1 enemy	2 enemies	3 enemies
Climber	Platform	1 platform	2 platforms	3 platforms	4 platforms
Heist	Maze Size	small	medium	large	extra large
Heist	Lock	1 lock	2 locks	3 locks	4 locks
Chaser	Maze Size	small	medium	large	extra large
Chaser	Enemy	1 enemy	2 enemies	3 enemies	4 enemies

Table 2: Levels of *Leaper*, *Climber*, *Heist* for children, as well as *Chaser* for adults in Study 2. These games were chosen to vary certain aspects of each game based on a particular axis. The levels increase in difficulty from Level 1 through Level 4. The goal level to complete is Level 3.

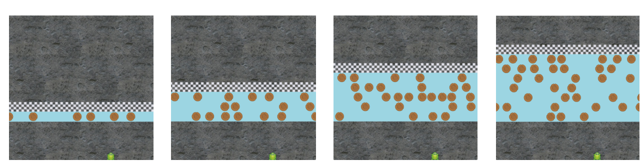

(a) Level 1 (b) Level 2 (c) Level 3 (d) Level 4

Figure 2: (a)-(d) Example levels of difficulty in *Leaper*, represented by the number of log lanes as the measure of difficulty.

Experimental Procedure

The study was performed on a computer. Participants were randomly assigned to play one of the three Procgen games.

Participants first underwent a familiarization trial where they practiced exploring an empty environment of the game with a video-game controller (e.g., they explored *Leaper* without any lanes or obstacles between the starting point and the finishing line), so that they could get used to the controller and gaming interface.

Next, participants were told to play the goal level of the game and were told they would be rewarded a sticker if they won. The rules of the game were not revealed to the participants. Since the experiment aimed to measure curriculum learning in a case where the goal was too challenging to be attained outright, participants had to fail the goal level to continue with the experiment. If a participant passed the goal level, experimenters reassigned them to play a different game. After the participant failed the goal level, the experimenter restated to them that the goal was to win that particular level, and once they did, they would get a sticker. The experimenter asked if the participant could tell them how the game worked and what the participant would have to do to get the sticker.

Then, participants were asked which level of the game they wanted to play in the next round and why. They were shown images of varying difficulty levels that quantitatively varied along a single game axis on a tablet, similarly to the ones in Figure 2, but were not explicitly told the relative difficulty of each level. Participants were presented with a total of three levels of difficulty: the goal level, two levels that were incrementally easier than the goal level,. A subset of 7 participants were further presented with an additional level 4 that was more difficult than the goal level.

At the end of each round, participants could visually observe their position on the computer screen and were also informed by the experimenter whether they passed their selected level. However, they were not given explicit verbal feedback on their performance or proximity to the goal level. On every other trial, they were reminded they would only win a sticker if they succeeded on the goal level.

This procedure continued until the participant passed the goal sticker level or up to a total of 10 trials, whichever was earlier. Participants who did not pass the goal level by the tenth trial were invited to play the goal level again and then the experiment concluded.

Results

Overall, children made an average global progress toward the goal of $\mu = 58.4\%$ with a standard error $SE = 11.8\%$ towards the goal (see Table 1 for definition), which was significantly different from 0%, $t(15) = 4.96$, $p < .001$. After failing the initial goal level, 72.7% participants started their curricula by selecting an easier level: 9 chose Level 1 and 7 chose Level 2. Only 1 (of the 7 presented with Level 4) chose Level 4 - the even more difficult level. If children were selecting their curricula at random, all four levels should be equally likely to be selected, resulting in a 25% chance of selecting Level 4. Our finding suggests that children recognized that spending extra effort on Level 4 would not increas their chance of a reward.

Game	Level Competence Calculation
Leaper	the number of lanes crossed normalized by the total number of lanes in the episode
Climber	the number of platforms reached normalized by the total number of platforms in the episode
Heist	the number of objectives achieved (gathering a key, unlocking a lock, or gathering the jewel) normalized by the total number of objectives to achieve
Chaser	the fraction of total green orbs collected

Table 3: Calculation of level competence (a percentage of completion) in each game.

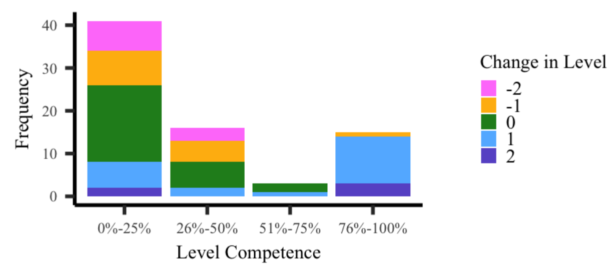

Figure 3: Level adjustments based on children's level competence on the current level. The x-axis represents current level competence (%) and the y-axis shows the frequency of level adjustments. A level change of +1 indicates selecting a harder level, while -1 represents choosing an easier one. Children tended to stay on the same level or move to easier levels when their competence was low, and moved to harder levels when their competence was high.

Next, we computed level competence as a measure of how far participants progressed within a single episode of gameplay on a specific level (see Table 1 for definition and Table 3 for a specific breakdown of how level competence is calculated for each game). We examined the change in levels selected by children across their automated curricula, given their competence on their current level (see Figure 3). A change in level of 0 indicates that the participant chose to remain on the same level, a change in level of -1 indicates that they chose the next easiest level, and so on. We found that children were more likely to remain on the same level or go to easier levels if their progress was low, and were more likely to move to more challenging levels if their competence was high. More specifically, level change (z-scored) was positively predicted by level competence (z-scored) in a linear mixed effects model with participant as a random effect, $\beta = .37$, $t(75) = 3.38$, $p < .01$. Thus, children used competence information to avoid overly easy tasks and advance to more difficult levels that were closer to the target level in their curriculum.

Study 2: Are children's and adults' curricula better than just focusing on the target task?

In Study 1, we found that children systematically designed their own curricula to tackle a challenging game task. However, it remains unclear whether their curricula are more effective than simply practicing the target level directly, as well as how their learning behavior compares to that of adults.

In Study 2, we introduced an additional target-only condition for children, where participants practiced exclusively on the same target level. We also examined adult participants' gameplay across both the target-only and self-initiated curriculum conditions. Furthermore, to expand our performance measures beyond level competence, we surveyed participants after each attempt to assess their causal understanding of the game—specifically, their knowledge of the gaming mechanics and what they had to do in order to win the game.

Methods

We recruited 33 five- to seven-year-old children (19 females, 14 males) with a mean age of 5.97 years ($\sigma = 0.77$ years) from public museums in the Bay Area, California, USA. Additionally, we recruited 44 adults aged 18 to 40 years, with a mean age of 21.68 years ($\sigma = 3.96$ years), from a college campus. Adult participants received class credit upon completing the experiment, whereas child participants continued to receive stickers upon succeeding on the target game level.

The methods in this study were largely the same as in Study 1, but with some modifications. For children, they were only assigned to play the game *Leaper*. Based on the survey results from Study 1, this game was found to be the most enjoyable and engaging for children. In contrast, adults were assigned a different game, *Chaser* (see Figure 1), which involves collecting green orbs and avoiding specific types of enemies in a maze. This game was chosen because it is more structurally complex and involves additional rules, ensuring it was challenging enough that the adults could not solve it too easily.

Participants were randomly assigned to one of two experimental conditions. In the self-initiated curriculum condition, participants were allowed to design their own learning pathways as in Study 1. In the target-level condition, participants practiced exclusively on the target level until they either successfully solved it or reached the maximum of 10 attempts. These modifications allowed us to explore whether the effectiveness of self-initiated curricula differed from focused practice on the target level across age groups.

In both conditions, participants were asked after every other round to explain how the game worked and what they needed to do in order to win the game. All open-ended responses were hand-coded by two double-blind researchers, using a point system to evaluate the completeness of their explanations. Scores were assigned based on the number of key observations participants identified, with a perfect score reflecting a full causal understanding of the game.

Participants were also asked a multiple choice question at the end as a causal rule understanding metric. Children

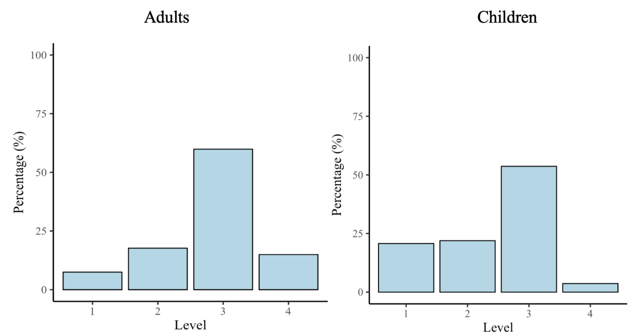

Figure 4: Distribution of levels selected by adults and children. Both groups selected the target Level 3 most frequently in their curricula; children selected the extra difficult Level 4 less than adults did.

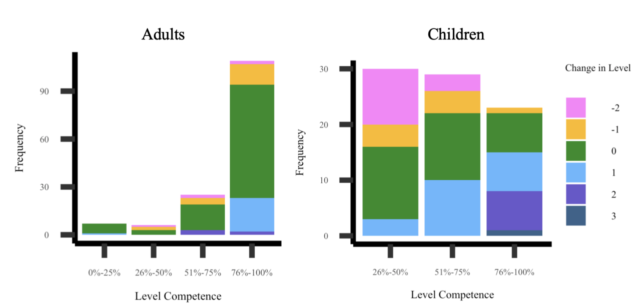

Figure 5: Level adjustments determined by adults and children based on their current level competence. Children's choices largely replicated Study 1, with change in level positively correlated with level competence. By contrast, adults did not determine their subsequent change in level based on their current level competence and mostly decided to remain on the same level independent of level competence.

(N=32) were asked "What do you think you had to do to win the game?", and given the choices of a) Avoid the logs, b) Swim in the water, and c) Hop from one log to another (correct answer). Adults were asked "Which statement is true about the maze?", and given the choices of a)You should collect all the stars and the ghosts, the green orbs don't matter, b) You should collect all the green orbs and the ghosts, the stars don't matter, and c) You cannot collect or destroy the static green eggs that appear and disappear in the maze (correct answer). They were also asked an additional multiple choice question "Which statement is true about the enemies?" and given the choices of a) You should avoid both the bat-looking enemies and the ghost-looking enemies, b) If you collide with a ghost-looking enemy, it will die and never come back again, and c) When you collide with a ghost-looking enemy, a static green egg spawns somewhere else on the map which then later hatches into a new enemy (correct answer).

Results

Children and adults in the self-initiated learning condition Overall, both adults and children invested over half of their attempts on the target Level 3 (59.86% and 53.65% respectively, see Figure 4). Children selected Level 4 only 3.66% time, much less frequently than adults (14.97%), again indicating that they were less likely to spend extra effort that would not cause them to win the extrinsic reward.

The data of children in the self-initiated curriculum condition in Study 2 largely replicates that in Study 1. After failing the initial goal level, 71.4% participants started their curriculum by selecting an easier level: 8 chose Level 1, 2 chose Level 2 and 4 chose Level 3. Furthermore, there is again a significantly positive correlation between level competence (z-scored) and level change (z-scored) in a linear mixed effects model with participant as a random effect, $\beta = .55$, $t(82) = 5.74$, $p < .0001$ (Figure 5, right), suggesting that children were selecting their next level contingent on how well they were performing on the current level. By contrast, adults largely chose to remain on the same level independent of their level competence. As observed in Figure 5 (left), adults' level competence was largely between 76% and 100% already, implying that their game was relatively easier to them than *Leaper* was for children. As a result, they might have thought that building a curriculum with incremental difficulty levels was not necessary for reaching their target level.

Self-initiated curriculum vs. Target-only condition Overall, children made significantly less average global progress in the self-initiated curriculum condition ($\mu = 2.97\%$, $SE = 1.71\%$, 78.6% succeeded on the target level) compared to the target-only condition ($\mu = 35.60\%$, $SE = 8.58\%$, 94.4% succeeded on the target level), $t(18) = 3.05$, $p = 0.00678$. Adults did not show a difference in global progress between the self-initiated curriculum condition ($\mu = 23.65\%$, $SE = 8.45\%$, 70% succeeded on the target level) and the target-only condition ($\mu = 23.21\%$, $SE = 4.90\%$, 100% succeeded on the target level). Both children and adults in the self-initiated curriculum condition spent more rounds to successfully solve the target level ($\mu = 6.8$ rounds vs. $\mu = 3.5$ rounds, $t(21) = 3.09$, $p = 0.00549$ in adults, $\mu = 6.82$ rounds vs. $\mu = 5.06$ rounds, $t(17) = 1.89$, $p = 0.075$ in children).

Both groups of participants generally developed a good understanding of the rules and mechanics of their respective games. Adults actually demonstrated greater causal understanding in their verbal descriptions of the game in the self-initiated curriculum condition than in the target-only condition, $t(37) = 2.32$, $p = 0.0259$. More specifically within the self-initiated curriculum condition, the more rounds that were played, the greater was the increase in causal understanding, $t(18) = 1.95$, $p = 0.0664$. By contrast, there was not a significant difference between the two conditions for children, $t(17) = 2.59$, $p = 0.0190$.

General Discussion

Children can systematically design their own curricula to achieve challenging goals based on their current level of competence, which they are able to assess independently without external feedback from the experimenter. In fact, they will systematically shift to easier levels, even though this may not actually be the most effective strategy for reaching the target level, as demonstrated in Study 2. Simply continuing to attempt the most difficult level, in spite of failing at first, may actually be more effective Beyond the extrinsic reward (a sticker) for solving the target level, children may also find intrinsic satisfaction in fully succeeding at a game level and gaining information about that level. The target level of *Leaper*, which children were assigned to play, is also relatively simple in complexity and relies heavily on motor control for success. *Chaser*, on the other hand, involved more complicated rules and adults in the self-initiated curriculum condition gained a deeper understanding of the game when they had a curriculum. A curriculum approach might prove more advantageous in games with greater causal complexity.

An interesting direction, is examining how learning generalizes across different game environments, specifically comparing a self-generated curriculum learning approach to a target-based approach. Perhaps learning through a curriculum can lead to better generalization. It's possible that learning under noisy, adaptive conditions fosters more flexible and transferable skills, which could enable participants to adapt more effectively to novel challenges. In contrast, a target-based condition might encourage optimization within a specific context but limit broader transfer. If the curriculum condition leads to stronger performance in a new game, this would suggest that exposure to dynamic, uncertain learning environments enhances adaptability and generalization. This finding could have implications for designing more effective training paradigms in AI and human learning, where structured yet unpredictable experiences lead to more robust skill acquisition. We leave this to future work.

Limitations

The games we provided may have been too easy for participants. Further work can use more difficult games that have more causal structure to be discovered.

Future directions

Building on our finding that children employed level competence as a metric to assess advancements in curriculum-based learning, we hope to further explore whether self-generated curriculum learning necessarily outperforms curriculum learning in a prescribed sequence based on , in a random sequence, or in a sequence designed by another individual (yoked control) (Atkinson, 1972; Kornell and Metcalfe, 2006) for both children and adults.

We are also interested in formalizing this work using computational methods. How exactly should an agent decide what to learn next? That is, should an agent decide to challenge itself by choosing something difficult and risky or stick to what it knows by choosing something easy and safe? Or perhaps an agent should aim for something in between – something that is "just right"? There remains a gap in understanding on how the Goldilocks Principle influences learning and exploration in humans and its extension to intelligent systems. A question remains of how to define "just right" both computationally and empirically.

While this task focuses on prediction – where participants decide which level to play next – similar tasks could be used to make inferences about other agents. That is, participants could observe another player's level choices without knowing their competence and estimate how far the player progressed in a given level based on their subsequent level selection. This approach would provide insight into how people infer progress and competence from decision-making patterns.

Finally, we are also currently running vision and large language model experiments to investigate whether multi-modal models could spontaneously generate a curriculum to solve difficult tasks.

Conclusion

Children can create a systematic curriculum to reach a goal that is too challenging to attain in a single attempt. We found that 5- to 7-year-old children used level competence to decide what level to play in order to eventually solve the most challenging level. Participants who were less competent at their current level were more likely to choose to play the same level again or an easier level – participants who were more competent at their current level were more likely to choose a more difficult level. Interestingly, unlike adults, children created a systematic curriculum based on their current competence, even though their performance might have actually been better by simply training on the challenging level itself.

References

Alfieri, L., Brooks, P. J., Aldrich, N. J., & Tenenbaum, H. R. (2011). Does discovery-based instruction enhance learning? *Journal of educational psychology*, *103*(1), 1.

Atkinson, R. C. (1972). Optimizing the learning of a second-language vocabulary. *Journal of experimental psychology*, *96*(1), 124.

Baranes, A. F., Oudeyer, P.-Y., & Gottlieb, J. (2014). The effects of task difficulty, novelty and the size of the search space on intrinsically motivated exploration. *Frontiers in neuroscience*, *8*, 317.

Bass, I., Mahaffey, E., & Bonawitz, E. (2023). Children use teachers' beliefs about their abilities to calibrate explore–exploit decisions. *Topics in Cognitive Science*.

Cobbe, K., Hesse, C., Hilton, J., & Schulman, J. (2020). Leveraging procedural generation to benchmark reinforcement learning. *International conference on machine learning*, 2048–2056.

Cubit, L. S., Canale, R., Handsman, R., Kidd, C., & Bennetto, L. (2021). Visual attention preference for intermediate predictability in young children. *Child development*, *92*(2), 691–703.

Doan, T., Castro, A., Bonawitz, E., & Denison, S. (2020). "wow, i did it!": Unexpected success increases preschoolers' exploratory play on a later task. *Cognitive development*, *55*, 100925.

Florensa, C., Held, D., Geng, X., & Abbeel, P. (2018). Automatic goal generation for reinforcement learning agents. *International conference on machine learning*, 1515–1528.

Glymour, C., Spirtes, P., & Scheines, R. (1991). Causal inference. *Erkenntnis*, *35*(1-3), 151–189.

Gopnik, A., Meltzoff, A. N., & Kuhl, P. K. (1999). *The scientist in the crib: Minds, brains, and how children learn.* William Morrow & Co.

Jiang, C., Ke, N. R., & van Hasselt, H. (2023). Learning how to infer partial mdps for in-context adaptation and exploration. *arXiv preprint arXiv:2302.04250*.

Ke, N. R., Didolkar, A., Mittal, S., Goyal, A., Lajoie, G., Bauer, S., Rezende, D., Bengio, Y., Mozer, M., & Pal, C. (2021). Systematic evaluation of causal discovery in visual model based reinforcement learning. *arXiv preprint arXiv:2107.00848*.

Khan, F., Mutlu, B., & Zhu, J. (2011). How do humans teach: On curriculum learning and teaching dimension. *Advances in neural information processing systems*, *24*.

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

Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2014). The goldilocks effect in infant auditory attention. *Child development*, *85*(5), 1795–1804.

Kornell, N., & Metcalfe, J. (2006). Study efficacy and the region of proximal learning framework. *Journal of experimental psychology: Learning, Memory, and Cognition*, *32*(3), 609.

Leonard, J. A., Cordrey, S. R., Liu, H. Z., & Mackey, A. P. (2023). Young children calibrate effort based on the trajectory of their performance. *Developmental Psychology*, *59*(3), 609.

Loyens, S. M., Magda, J., & Rikers, R. M. (2008). Self-directed learning in problem-based learning and its relationships with self-regulated learning. *Educational psychology review*, *20*, 411–427.

Pearl, J. (2000). Causal inference without counterfactuals: Comment. *Journal of the American Statistical Association*, *95*(450), 428–431.

Peters, J., Janzing, D., & Schölkopf, B. (2017). *Elements of causal inference: Foundations and learning algorithms*. The MIT Press.

Poli, F., Serino, G., Mars, R., & Hunnius, S. (2020). Infants tailor their attention to maximize learning. *Science advances*, *6*(39), eabb5053.

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

Rule, J., Goddu, M., Chu, J., Pinter, V., Reagan, E. R., Bonawitz, E., Ullman, T., et al. (2023). Fun isn't easy: Children choose more difficult options when "playing for fun" vs."trying to win".

Schmidt, H. G., Loyens, S. M., Van Gog, T., & Paas, F. (2007). Problem-based learning is compatible with human cognitive architecture: Commentary on kirschner, sweller, and. *Educational psychologist*, *42*(2), 91–97.

Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., & Bengio, Y. (2021). Toward causal representation learning. *Proceedings of the IEEE*, *109*(5), 612–634.

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

Serko, D., Leonard, J., & Ruggeri, A. (2024). Children strategically decide what to practice.

Spirtes, P., Glymour, C. N., & Scheines, R. (2000). *Causation, prediction, and search*. MIT press.

Ten, A., Kaushik, P., Oudeyer, P.-Y., & Gottlieb, J. (2021). Humans monitor learning progress in curiosity-driven exploration. *Nature communications*, *12*(1), 5972.

Vgotsky, L. S. (n.d.). Mind in society: The development of higher psychological processes. *London: Harvard University Press*.

Wang, J. X., King, M., Porcel, N., Kurth-Nelson, Z., Zhu, T., Deck, C., Choy, P., Cassin, M., Reynolds, M., Song, F., et al. (2021). Alchemy: A benchmark and analysis toolkit for meta-reinforcement learning agents. *arXiv preprint arXiv:2102.02926*.

Wang, J., & Bonawitz, E. (2023). Children's sensitivity to difficulty and reward probability when deciding to take on a task. *Journal of Cognition and Development*, *24*(3), 341–353.

Zhang, X., Carrillo, B. A., & Leonard, J. A. (2023). Children's developing understanding of learning as improvement over time. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *45*(45).