

## **UC Merced**

### **Proceedings of the Annual Meeting of the Cognitive Science Society**

#### **Title**

Predicting children's and adults' preferences in physical interactions via physics simulation

#### **Permalink**

<https://escholarship.org/uc/item/4j58p1r2>

#### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

#### **Authors**

Kachergis, George  
Radwan, Samaher F  
Long, Bria  
[et al.](#)

#### **Publication Date**

2021

#### **Copyright Information**

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# Predicting children’s and adults’ preferences in physical interactions via physics simulation

George Kachergis<sup>1,\*</sup>, Samaher Radwan<sup>1</sup>, Bria Long<sup>1</sup>, Judith Fan<sup>2</sup>,  
Michael Lingelbach<sup>1,3</sup>, Daniel M. Bear<sup>1,3</sup>, Daniel L. K. Yamins<sup>1,3</sup>, and Michael C. Frank<sup>1</sup>

<sup>1</sup>Department of Psychology, Stanford University

<sup>2</sup>Department of Psychology, University of California at San Diego

<sup>3</sup>Department of Computer Science, Stanford University

\*kachergis@stanford.edu

## Abstract

Curiosity is a fundamental driver of human behavior, and yet because of its open-ended nature and the wide variety of behaviors it inspires in different contexts, it is remarkably difficult to study in a laboratory context. A promising approach to developing and testing theories of curiosity is to instantiate them in artificial agents that are able to act and explore in a simulated environment, and then compare the behavior of these agents to humans exploring the same stimuli. Here we propose a new experimental paradigm for examining children’s – and AI agents’ – curiosity about objects’ physical interactions. We let them choose which object to drop another object onto in order to create the most interesting effect. We compared adults’ (N=155) and children’s choices (N=66; 3-7 year-olds) and found that both children and adults show a strong preference for choosing target objects that could potentially contain the dropped object. Adults alone also make choices consistent with achieving support relations. We contextualize our results using heuristic computational models based on 3D physical simulations of the same scenarios judged by participants.

**Keywords:** curiosity; novel objects; object interactions; intuitive physics

## Introduction

Curiosity is a hallmark aspect of human intelligence. From infants exploring the objects in their environment to scientists exploring the frontiers of our solar system, humans are highly motivated to seek out new knowledge and experiences. Although such exploratory behavior has long been recognized as a critical component of human learning (James, 1983) and cognitive development (Gopnik, Meltzoff, & Kuhl, 2009; Piaget, 1952), formal theories that explain human curiosity and how it drives exploratory behavior have remained elusive (Kidd & Hayden, 2015). Moreover, extant theories have rarely provided precise enough predictions to be directly compared to empirical measurements of curiosity-driven behavior in humans, in part because curiosity tends to result in complex, open-ended behaviors that are difficult to quantify.

The goal of the current paper is to help close this gap by proposing a novel paradigm to advance our theoretical understanding of curiosity, specifically within the domain of physical object interactions (Bramley, Gerstenberg, Tenenbaum, & Gureckis, 2018; Kubricht, Holyoak, & Lu, 2017). First, we present an empirical investigation of the pattern of choices taken by children and adults in a novel task. We then present a set of heuristic models of this behavior based on a variety of metrics of simulated physical interactions (e.g., likelihood of

a dropped object coming to rest on a target object), and test the degree to which these features predict children and adults’ behaviors in the same task.

We take inspiration from a large body of prior work in developmental psychology investigating the development of knowledge about physical objects, their properties, and how they interact (Baillargeon, 2007; Hespos & VanMarle, 2012; Smith, Jayaraman, Clerkin, & Yu, 2018). Children both spend a remarkable amount of time across exploring how different objects and surfaces interact during naturalistic play (Fenson, Kagan, Kearsley, & Zelazo, 1976) and spend longer time exploring objects that appear to violate physical laws (e.g., an object that appears to pass through a wall; Stahl & Feigenson, 2015; Baillargeon, 2007). In other words, children seem to actively learn about physical interactions by intervening on the world and observing the consequences of their actions (Allen, Smith, & Tenenbaum, 2020; Gopnik et al., 2009; Gureckis & Markant, 2012; Needham, 2000). However, relatively little work has either examined what types of physical interactions children are most interested in testing, or linked children’s actions to formal theories of exploration or of physics learning. Instead, most work that has investigated exploratory behavior in children has done so by examining how they play with relatively complex objects – for example, documenting the number of functions discovered while playing with a novel toy (e.g., Cook, Goodman, & Schulz, 2011; Bonawitz, Schijndel, Friel, & Schulz, 2012; Gweon, Pelton, Konopka, & Schulz, 2014; Hoicka et al., 2016).

We thus developed a novel physical exploration task in which children, adults, or AI agents can choose which series of physical experiments to perform, and then demonstrate how this paradigm can be used to test theories of curiosity about physical interactions. First, we measured people’s choices in a novel task, in which participants were told that they would drop a given object (e.g., a sphere, a torus; see Figure 1) onto one of two target objects (e.g., a dumbbell, a pentagonal prism; see Figure 2), with the goal of the creating the most interesting physical interaction. We recorded adults’ and children’s (3–7 years of age) choices for a set of trials in which some of the drop objects could plausibly end up *supported* by one of the target objects, and for a set of trials in which some of the drop objects could end up *contained* by one of the targets – although the other potential

targets likely also offered affordances of interest. After examining children’s and adult’s preferences, we tested the predictive power of a variety of heuristic features derived from the simulated physics of this task in a 3D environment. To preview our results, we found a surprising amount of consistency in how both children and adults chose to explore the relationships between different objects – but that this consistency was not easily explained by any of the simple heuristic features that we tested. We propose that this type of physical exploration task is a promising test for adjudicating between different embodied models of curiosity.

### Experiment 1: Adults

To investigate the systematicity of people’s preferences for physical interactions between objects, we began by studying adults, whom we might expect to be less idiosyncratic and thus more consistent in their choices than children. Our design is motivated by the results of a pilot study conducted in-person in January, 2020, in which we asked adults to select which of a pair of 3D-printed toy blocks (see Figure 1) they would like to drop on a given target object from the set, or vice-versa: on which of a given pair of target objects would they like to drop a given object. In the pilot study (N=15), the pairs of target or drop objects were chosen essentially at random, but were the same for all participants. We were surprised to find consistency in adults’ preferences for many of the trials: especially when given a drop object and asked to choose one of two target objects, adults were quite often (75-90%) targeting the object that would either *contain* or *support* the dropped object (i.e., the pipe could contain the cone; the pentagonal prism could support the octahedron). This consistency in their choices is remarkable given that there are many other possible goals that people might choose in order to make something “interesting” happen: they might attempt to make the dropped object roll or bounce far from the target, or rebound in an unexpected direction, but in fact most people settled on attempting support or containment relations. Thus, we set out to examine the development of these preferences, first in a large adult sample in order to have high reliability for later model comparisons, and then in young children.

### Method

**Participants** 200 adults were recruited via Amazon Mechanical Turk and were paid \$1 for participating.

**Materials** Stimuli were images of 3D objects produced using Blender 3D-modeling software. The nine objects, depicted in Figure 1, were bowl, cone, dumbbell, octahedron, pentagon (pentagonal prism), pipe, pyramid, torus, triangular prism, and ball (sphere; not pictured).

**Design** The experiment consisted of 20 drop trials, completed by each participant in one of four pseudorandom orders. Participants were randomly assigned to each order. Each drop trial displayed an object X that is to be dropped on one of two other objects (Y or Z), as described in the following procedure. In the 20 trials, each of the above 3D ob-




Figure 1: (Top) Nine of the 10 3D objects used for dropping and as targets (not pictured: ball). (Bottom) Example support and containment relations.

jects was used twice as the drop object (X), and appeared four times as potential targets. Instead of randomly sampling the target objects (Y or Z), each trial was designed to allow for one of two goal relations (or affordances): containment, or support. That is, the target objects were chosen so that one of them would be capable of either supporting or containing the dropped object (see examples: Figure 1, bottom). For example, if asked to drop the cone on either the torus or the octahedron, a containment relation could be achieved between the cone and the torus, but not the octahedron (and support is unlikely due to the small surface area of the octahedron). An example support relation is choosing to drop the pentagonal prism on the (equal diameter) pipe rather than the pyramid. 10 of the 20 trials were designed to afford containment relations, and the other 10 afforded support relations. We refer to the objects affording the designed containment and support relations as *targets*, and the alternative objects as *distractors*. Distractors for each trial were the same for all participants, and were selected to 1) not afford containment/support of the dropped object, and 2) roughly equate the frequency of appearance of all 10 objects.

**Procedure** Participants were instructed that they would be helping choose which toys to include in a new set of children’s blocks in order to make them the most interesting. Participants were then given a practice trial (see Figure 2),




Figure 2: Example of a trial given to adults in Experiment 1.

in which they were told that they should imagine dropping a given object (the torus) on each of two toys (the dumbbell or the pentagonal prism) in a bin. They were then asked to choose which toy to drop the torus on, and prompted to choose the most interesting or surprising combination. After making their choice on the practice trial, participants were shown a 10-second video recording of 3D-printed plastic toys carrying out their chosen interaction (e.g., if they chose the dumbbell, they would see the torus dropped on the dumbbell). This was done to ensure that they understood the consequences of the choice they made, although it was pointed out to them that the video they saw showed only one possible outcome. After the example video, participants were given a sequence of 20 more drop trials asking them to choose which toy (X or Y) they would like to drop object Z on. Four catch trials were interspersed among the 20 drop trials, which asked them to indicate which object they had just dropped on the previous trial (3-alternative forced choice). The catch trials were meant to encourage attention to the objects, and participants who were incorrect on two or more of the catch trials were excluded.

## Results

We analyzed data from the 155 adults who completed the experiment and answered at least 3 of the 4 catch trials correctly (45 participants were excluded for not meeting this criterion). Averaging each subject's binary responses (1=target relation, 0=alternative) for each trial type revealed a stronger preference for containment relations ( $M = 0.87$ ) than support relations ( $M = 0.63$ ; paired  $t(154) = 12.69$ ,  $p < .001$ ,  $d = 1.02$ ). Table 1 shows the proportion of participants that chose the designed target relation for each of the 20 trials. Based on the binomial distribution, any of the trials on which more than 90

of the 155 participants agree (i.e.,  $>0.59$  or  $<0.41$ ) significantly differ from chance. As can be seen in Table 1, adult participants significantly preferred the target relation for all 10 of the containment trials, and significantly preferred five targets on the 10 support trials; the other five support trials did not significantly differ from chance, suggesting that many participants found the prospect of dropping on the distractor object at least equally enticing.

Further, these trial effects were quite reliable across participants. We examined the split-half reliability of adults' choices by repeatedly splitting the data and testing the correlation between the halves. With 100 random samples, reliability for adults was  $r = 0.95$  ( $sd=0.02$ ). Reliability was  $r = 0.85$  ( $sd=0.07$ ) for containment trials and  $r = 0.9$  ( $sd=0.05$ ) for support trials.

relation	drop	target	child	adult
contain	bowl	pyramid	0.72	0.96
contain	cone	torus	0.73	0.94
contain	dumbbell	pipe	0.67	0.92
contain	octahedron	pipe	0.58	0.90
contain	pentagon	bowl	0.69	0.88
contain	pipe	cone	0.58	0.88
contain	pyramid	torus	0.71	0.87
contain	sphere	bowl	0.80	0.86
contain	torus	cone	0.67	0.81
contain	trig prism	bowl	0.65	0.71
support	bowl	torus	0.71	0.85
support	cone	trig prism	0.45	0.82
support	dumbbell	pentagon	0.38	0.74
support	octahedron	pentagon	0.45	0.64
support	pentagon	pipe	0.57	0.62
support	pipe	torus	0.60	0.56
support	pyramid	trig prism	0.52	0.55
support	sphere	pipe	0.58	0.54
support	torus	bowl	0.70	0.53
support	trig prism	pentagon	0.62	0.46

Table 1: Children's and adults' target preferences per trial.

## Experiment 2: Children

Experiment 1 demonstrated that adults show consistent preferences for dropping objects on target objects that afford containment relations, and to a lesser extent support relations. Experiment 2 investigates whether these strong preferences are present even in young children, recognizing that children may find different types of physical interactions (e.g., rolling, bouncing, unpredictability) of greater interest than adults, or may simply show more idiosyncratic choice patterns.

## Method

We adapted the same materials and procedure used in Experiment 1 for use in an online experiment with children, in order to directly compare the results for children and adults.

**Participants** Participants were 73 children recruited online via outreach through a local nursery school and targeted Facebook ads over the course of 6 months. Participant exclusions were made based on cases where i) the participant did not complete more than half of the trials or ii) the parent did not consent for video recording of the session. After exclusions, results from 66 children were analyzed, including 17 3-year-olds, 15 4-year-olds, 16 5-year-olds, 16 6-year-olds, and 2 7-year-olds.

**Materials** The materials were the same as those used in Experiment 1, except that the trials were adapted for presentation via Zoom in a slide presentation as shown in Figure 3.




Figure 3: Zoom screen configuration for children in Experiment 2. Parents were asked to place the experimenter’s video feed in the bottom center so the experimenter could point up to the drop object and left and right at each target choice.

**Procedure** After the parent provided informed consent, children were assigned to one of two pseudorandom trial orders in counterbalanced order. To accompany the practice trial, children were shown 3D-printed instances of the three physical objects held by the researcher, to ensure that children understood their physicality. Children were asked to verbally select target object Y or Z, and then asked to confirm whether they wanted to select the object on the red or green side of the screen, in order to prevent left/right confusion. The sides on which target objects and colors appeared were counterbalanced across trials.

## Results

Similar to adults, averaging each child’s binary responses for each trial type revealed a stronger preference for containment relations ( $M = 0.68$ ) than support relations ( $M = 0.56$ ; paired  $t(154) = 12.69, p < .001, d = 1.02$ ). Unlike adults, however, children did not show a significant preference for choosing the support relation. Figure 4 shows the proportion of participants’ choice of the designed target relation broken down by relation type and age group. Preferences increased with age and were overall stronger for containment.

We next examined trial-level effects in children. Table 1 shows the proportion of children (“child” column) choosing the designed target relation for each of the 20 trials, alongside the adult choice proportions from Experiment 1. Based

on the binomial distribution, any of the trials on which more than 41 of the 66 participants agree (i.e.,  $>0.64$  or  $<0.36$ ) significantly differ ( $p < .05$ ) from chance. Children significantly preferred the designed relation on eight of the 10 containment trials, while only significantly preferring two of the 10 support relation choices. On the remaining two containment trials and the other eight support trials, children’s preferences also did not significantly differ from chance against the designed relation. These item effects were relatively reliable, though noisier than that of adults: the split-half reliability for children’s choices was  $r = 0.61$  ( $sd=0.09$ ). The split-half reliability of children’s choices on containment trials was  $r = 0.32$  ( $sd=0.22$ ), and on support trials was  $r = 0.64$  ( $sd=0.14$ ).




Figure 4: Children’s vs. adults’ choices for each type of trial, with bootstrapped 95% confidence intervals.

**Comparison of Children’s vs. Adults’ Preferences** Figure 5 shows adults’ (Experiment 1) vs. children’s (Experiment 2) choice proportions for the target relations, colored by relation type (visualizing the same data presented in Table 1, but with 95% CIs). If adults’ and children’s judgments corresponded perfectly, they would fall along the dotted  $y = x$  line, but children’s preferences mostly fall short of that line, lying closer to the chance line ( $y = 0.5$ ). At a glance, these results seem to support the hypothesis that “children are noisy adults” in this task as well as suggesting some consistency in the trial-level effects across development. This consistency also confers additional motivation for attempting to understand what drives human interest in particular physical interactions in this task through computational models.

## Comparison to Physics Simulations

We next attempted to predict adults’ and children’s judgments using a 3D model of the physical interactions instantiated in our paradigm. We constructed a simulation of each drop in-




Figure 5: Comparison of children’s vs. adults’ preferences on each trial (dot), with bootstrapped 95% confidence intervals.

teraction in each trial using a physics engine (Gan et al., 2020) and 3D models of the objects (the same models that we had printed and used in the videos and photos in the experimental paradigm). We then assessed how well a variety of heuristic features based on these simulated physical interactions could predict choice data.

### Physical interaction model and features

For each trial’s two possible drop object choices, we ran 250 simulated drops, and for each drop we calculated a variety of features measuring the state of the model after the drop was completed. We selected features that we thought may provide good metrics for what people could find interesting, for instance, the mean amount of time before both objects come to rest ( $M(\text{Move Time})$ ), as well as the standard deviation of that time ( $SD(\text{Move Time})$ ). Some features were calculated separately for drop and target objects, such as how variable each object’s final positions tend to be (e.g., the inverse SD of each object’s final position ( $1/SD(\text{Drop Obj Pos})$  and  $1/SD(\text{Target Pos})$ ), the mean and maximum final distance of each object from the drop location (e.g.,  $M(\text{Target Dist})$  and  $Max(\text{Target Dist})$ ), and how variable the velocity of each object tends to be after the first collision (e.g.,  $SD(\text{Target Vel})$ ). We also tested a variety of features based on the likelihood of the dropped object coming to rest atop the target object— $Pr(\text{support})$ —and the robustness of this likelihood to small perturbations in drop position ( $Supp. Sharpness$ ). Note that the model’s definition of support does not explicitly distinguish support from containment.

We averaged each of these features across simulations and then used the resulting values to generate relative preferences

for the drop choices on each trial. A model’s preference on any given trial was assumed to be proportional to the relative magnitude of the feature values for the possible physical interactions on that trial, scaled by a softmax parameter  $\beta$ . For each model (feature), we also optimized  $\beta$  separately to find best-fitting values for children’s and adults’ choice proportions, with the objective of minimizing mean squared error (MSE) between model and human choice proportions across all 20 trials.

### Results

Overall, a few of these features were modestly correlated with people’s choices. Table 2 shows the correlation between model preferences (with softmax  $\beta = 1$ ) vs. adults’ and children’s choice proportions, both to all trials and separately for containment and support trials. When all trials were considered together, the feature with the strongest correlation to adults’ choices was the velocity of the target object ( $SD(\text{Target Vel})$ ,  $r = .33$ ), while the feature with the strongest correlation to children’s choices was *Support Sharpness* (see Methods,  $r = .18$ ). For support trials, both adults’ ( $r = .39$ ) and children’s ( $r = .41$ ) choices were best correlated with a feature indexing the variability in the target object’s final position ( $1/SD(\text{Target Pos})$ ). For containment trials, adults’ choices were most correlated with the velocity of the dropped object ( $SD(\text{Drop Obj Vel})$ ,  $r = .54$ ), which also showed a modest correlation with children’s choices ( $r = .41$ ); however, children’s choices were most correlated with the maximum distance of the target object ( $Max(\text{Target Dist})$ ,  $r = .43$ ). Many of the other features showed weak or even moderate negative correlations with people’s choices.

Examination of the models with  $\beta$  optimized, in contrast, revealed that none achieved a good fit to human preferences. While the best-fitting feature,  $Pr(\text{support})$ , captured both adults’ and children’s overall mean preference for the designed target relations, it explained essentially no trial-level variation and was in fact negatively correlated with choices. This was not due to noisy behavioral data: while the noise ceiling on choice preferences for all trials was  $r = 0.99$  for adults and  $r = 0.87$  for children, the best-fitting feature accounted for very little of that.

### General Discussion

This study aimed to 1) measure the consistency of children’s and adults’ preferences for particular object interactions, and 2) determine whether people’s interest was predicted by particular features of the imagined physical interactions between objects. In an online study that gave adults the opportunity to make drop choices that could result in either a likely support or likely containment relation on each trial, we found that adults consistently chose targets likely to contain the dropped object, and to a lesser degree chose target objects that were likely to support the dropped object. Experiment 2 found the same pattern in children, but with attenuated preferences, especially in the youngest children. One interpretation of these results is that infants’ early-emerging interest in containment

Table 2: Correlations ( $r$ ) of model and human responses.

Feature	Adult: All	Child: All	Adult: Support	Child: Support	Adult: Contain	Child: Contain
SD(Target Vel)	<b>0.33</b>	-0.11	0.1	-0.51	0.34	-0.09
1/SD(Target Pos)	0.14	-0.08	<b>0.39</b>	<b>0.41</b>	0.16	-0.52
M(Target Dist)	0	0.13	0.11	0.22	0.42	0.42
Max(Target Dist)	-0.01	0.17	0.11	0.33	0.41	<b>0.43</b>
Supp. Sharpness	-0.02	<b>0.18</b>	-0.26	0.14	0	0.18
M(Move Time)	-0.07	0.04	-0.08	-0.08	-0.02	0.29
SD(Drop Obj Vel)	-0.1	0.17	-0.27	0.3	<b>0.54</b>	0.41
SD(Move Time)	-0.18	-0.19	-0.01	-0.29	-0.37	0.13
Pr(support)	-0.36	-0.48	-0.48	-0.54	-0.39	-0.51
Max(Drop Obj Dist)	-0.45	-0.28	-0.39	-0.29	-0.09	0.18
M(Drop Obj Dist)	-0.47	-0.31	-0.45	-0.35	-0.04	0.2
1/SD(Drop Obj Pos)	-0.48	-0.32	-0.49	-0.33	-0.1	0.12

(Casasola, Cohen, & Chiarello, 2003) may be related to older children’s and adults’ interest in stochastic interactions that have some possibility of producing containment relations.

Given the high reliability of choices on this task, particularly for adults ( $r = .95$ , children  $r = .61$ ), we tested how well heuristic features of the physical interactions could predict these choices. In light of findings that people often single out one of several possible causes as “the” cause in judgements of physical causation (e.g., Gerstenberg & Icard, 2020), we might expect to find single features that are predictive of people’s choices. We explored a wide range of features that we thought might be proportional to people’s interest in particular object interactions (e.g., how far a dropped object might end up from the drop location), but not one of these features explained much of the trial-level variance in people’s choices. Instead, we found a few features that were partially correlated with people’s choices: for example, the best predictor of both children’s and adult’s choices on support trials was the inverse standard deviation of the target’s final position, but the strength of this correlation was only moderate ( $r \approx 0.4$ ). For containment trials, the normed velocity of the drop object after first collision was most correlated with adults’ choices and moderately with children’s, but children’s choices were also correlated with the maximum and mean distance traveled by the target. On balance, although much systematic variation in this dataset is unaccounted for, it is somewhat encouraging that people’s interest in physical interactions can be in part accounted for by simple heuristic features based on simulations of the same physical task we asked participants to merely imagine.

At the same time, the fact that these results are not fully explainable by simple heuristic models suggest that this paradigm may fruitfully serve as a test bed for evaluating more complex computational theories of curiosity. For example, in one proposal learners prefer to explore stimuli that are “moderately discrepant” in relation to their current knowledge state, thereby providing an opportunity to learn (Kinney

& Kagan, 1976), and curiosity changes as the gap between the learner’s knowledge develops and the state of the world closes (Loewenstein, 1994). These theories are beginning to be implemented in deep neural networks, which are now capable of learning forward and inverse physical dynamics (i.e., “intuitive physics”) from images when given the ability to “poke” the objects in the scene (Agrawal, Nair, Abbeel, Malik, & Levine, 2016), and have more recently been used to test which “curiosity” policies for generating actions result in robust and effective learning of intuitive physics in deep reinforcement learning agents (Haber, Mrowca, Fei-Fei, & Yamins, 2018).

Such embodied deep reinforcement learning agents are often composed of two models: a *world model* which attempts to predict the consequences of the agent’s actions (i.e., the forward dynamics—physics—of the world), while the *self model* attempts to predict the errors of the world model. The predictions of the self model can be used to implement different curiosity policies for choosing actions (e.g., picking actions that are expected to challenge its world model; Haber et al., 2018). In future work, we aim to examine the development of curiosity in such curious artificial agents, and compare whether these agents behave similarly to people when asked to “drop it like it’s hot” (Dogg, 2004). Will these agents show choices consistent with potential containment/support relations, or do these human-like preferences require representation and selection of possible goals? How might the agents’ choices evolve over time, as the world model matures? Might the agents’ curiosity policies need to change over development in order to match people’s choices, or do these goals emerge through the dynamics of learning?

The multifaceted nature of curiosity has fascinated researchers for decades, but only recently have we developed the means to factorize and test these theories. We propose that the development of parallel tasks for children, humans, and embodied agents is a promising way forward towards modeling curiosity in the real world.

## Acknowledgements

This work was funded by a Stanford University Human-centered Artificial Intelligence grant.

## References

- Agrawal, P., Nair, A. V., Abbeel, P., Malik, J., & Levine, S. (2016). Learning to poke by poking: Experiential learning of intuitive physics. In *Advances in neural information processing systems* (pp. 5074–5082).
- Allen, K. R., Smith, K. A., & Tenenbaum, J. B. (2020). Rapid trial-and-error learning with simulation supports flexible tool use and physical reasoning. *Proceedings of the National Academy of Sciences*, *117*(47), 29302–29310.
- Baillargeon, R. (2007). The acquisition of physical knowledge in infancy: A summary in eight lessons. In *Blackwell handbook of childhood cognitive development* (pp. 47–83). Blackwell Publishers Ltd.
- Bonawitz, E. B., Schijndel, T. van, Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology*, *64*, 215–234.
- Bramley, N. R., Gerstenberg, T., Tenenbaum, J. B., & Gureckis, T. M. (2018). Intuitive experimentation in the physical world. *Cognitive Psychology*, *105*, 9–38.
- Casasola, M., Cohen, L. B., & Chiarello, E. (2003). Six-month-old infants' categorization of containment spatial relations. *Child Development*, *74*(3), 679–693.
- Cook, C., Goodman, N. D., & Schulz, L. E. (2011). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*, *120*, 341–349.
- Dogg, S. (2004). *Drop it like it's hot. R&G (Rhythm & Gangsta): The Masterpiece*.
- Fenson, L., Kagan, J., Kearsley, R. B., & Zelazo, P. R. (1976). The developmental progression of manipulative play in the first two years. *Child Development*, *232*–236.
- Gan, C., Schwartz, J., Alter, S., Schrimpf, M., Traer, J., De Freitas, J., ... others. (2020). Threedworld: A platform for interactive multi-modal physical simulation. *arXiv Preprint arXiv:2007.04954*.
- Gerstenberg, T., & Icard, T. (2020). Expectations affect physical causation judgments. *Journal of Experimental Psychology: General*, *149*(3), 599.
- Gopnik, A., Meltzoff, A. N., & Kuhl, P. K. (2009). *The scientist in the crib: Minds, brains, and how children learn*. HarperCollins.
- Gureckis, T. M., & Markant, D. B. (2012). Self-directed learning: A cognitive and computational perspective. *Perspectives on Psychological Science*, *7*(5), 464–481.
- Gweon, H., Pelton, H., Konopka, J. A., & Schulz, L. E. (2014). Sins of omission: Children selectively explore when teachers are under-informative. *Cognition*, *132*, 335–341.
- Haber, N., Mrowca, D., Fei-Fei, L., & Yamins, D. L. (2018). Learning to play with intrinsically-motivated self-aware agents. In *Advances in neural info. Processing systems*.
- Hespos, S. J., & VanMarle, K. (2012). Physics for infants: Characterizing the origins of knowledge about objects, substances, and number. *Wiley Interdisciplinary Reviews: Cognitive Science*, *3*(1), 19–27.
- Hoicka, E., Mowat, R., Kirkwood, J., Kerr, T., Carberry, M., & Bijvoet-van den Berg, S. (2016). One-year-olds think creatively, just like their parents. *Child Development*, *87*(4), 1099–1105.
- James, W. (1983). *Talks to teachers on psychology and to students on some of life's ideals*. Harvard University Press.
- Kidd, C., & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. *Neuron*, *88*(3), 449–460.
- Kinney, D. K., & Kagan, J. (1976). Infant attention to auditory discrepancy. *Child Development*, *47*, 155–164.
- Kubricht, J. R., Holyoak, K. J., & Lu, H. (2017). Intuitive physics: Current research and controversies. *Trends in Cognitive Sciences*, *21*(10), 749–759.
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psych. Bulletin*, *116*, 75–98.
- Needham, A. (2000). Improvements in object exploration skills may facilitate the development of object segregation in early infancy. *Journal of Cognition and Development*, *1*(2), 131–156.
- Piaget, J. (1952). The origins of intelligence in children, 8.
- Smith, L. B., Jayaraman, S., Clerkin, E., & Yu, C. (2018). The developing infant creates a curriculum for statistical learning. *Trends in Cognitive Sciences*, *22*(4), 325–336.
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, *348*(6230), 91–94.