Surprise isn't symmetrical: Adults' looking suggests non-perceptual considerations during dishabituation

Qiong Cao[1,2] **(qcao246@mit.edu), Anjie Cao**[3] **(anjiecao@stanford.edu), Gal Raz**[1] **(galraz@gmail.com),**
Joshua Tenenbaum[1] **(jbt@mit.edu), and Shari Liu**[2] **(sliu199@jhu.edu)**
[1]Department of Brain and Cognitive Sciences, MIT,
[2]Department of Psychological & Brain Sciences, Johns Hopkins University
[3]Department of Psychology, Stanford University,

Abstract

People of all ages explore the world through looking. Recently, Raz, Cao et al. (2025) built an image-computable model (RANCH) that predicts adults' and infants' looking behavior to a large stimulus set, including graded responses to changes in pose, animacy, and number. This model succeeded despite having only a perceptual embedding space of stimuli. However, looking may be influenced by non-perceptual considerations. Using the same data, we found that adults' behaviors challenge a key assumption of perceptual-only account: since the perceptual distance between two items is symmetrical, behavior guided only by perceptual space should also be symmetrical. Yet, adults did not treat changes in different directions as mere reciprocal transformations. For instance, adults looked longer at magical appearance than disappearance. We suggest that image-computable models of looking behavior would benefit from representations of objects, in addition to perceptual features of images.

Keywords: dishabituation; looking time; object; number; animacy; cognitive development

Introduction

Imagine sitting in a magic show: A magician places an apple on a wooden table, lowers her velvet top hat over the apple, and raises the hat to reveal a rabbit. For a few moments, the room falls silent while everyone stares. What goes on in our minds during surprising events like this, and what predicts our looking behaviors towards them? This question is relevant not only for open questions surrounding the mechanisms of surprise and curiosity, but also for better understanding the methods that use looking behavior to query the development of perception, attention, and learning, especially in infants.

One major challenge for looking time methods is that looking is a low-dimensional measure that is likely driven by high-dimensional causes—both factors internal to the observer (e.g., their learning goals, mental states, etc.) and those inherent to the stimulus (e.g., how salient it is, how different it is relative to what was there before). How can we tease apart these causes? In the current work, we conducted a secondary analysis of a large dataset of adults to compare shifts in looking time when one stimulus (a rabbit) is replaced by another (an apple), across the direction of the change (rabbit to apple vs. apple to rabbit). We documented systematic asymmetries in adults' looking behaviors that depend on the order of the two stimuli, suggesting that perceptual properties of the stimulus alone are insufficient to explain adults' looking behaviors. Instead, we suggest that representations of objects, beyond the images they project onto our retina, drive adults' looking behaviors.

Drivers of looking in human adults and infants

Given a limited set of mental resources, and an infinite set of potential things to learn, selecting what information to pay attention to is crucial for learners of all ages. Studying looking behavior has offered a unique window into information processing across the lifespan, especially in infants, who have limited motor repertoires. Cognitive scientists have measured habituation (decreasing of looking time as observing repeated stimulus) and dishabituation (regaining of interest in looking after a novel stimulus was presented) in looking behavior to study processes from vision and perceptual learning (Held et al., 1980) all the way to speech processing (Kuhl, 2004) and understanding of agents and objects (Spelke & Kinzler, 2007; Baillargeon et al., 2016; Baillargeon, 1994; Spelke, 2022).

However, a major debate in the field is how to understand and interpret looking behaviors, in light of the many processes it could reveal. Some scholars remain committed to the claim that stimulus-driven and/or non-cognitive processes dominate infant looking (Haith, 1998; Blumberg & Adolph, 2023). On the other hand, others believe that it is possible to infer cognitive processes from looking behavior (Aslin, 2000; Liu et al., 2023; Margoni & Thomsen, 2024). Under this second account, when people look at one stimulus and then another, they are representing more than just the perceptual differences between those two stimuli.

In support of this interpretation, Kunin et al. (2024) found that distinct predictors moderate looking behaviors towards unexpected, but visually familiar stimuli, and visually novel, but expected stimuli, suggesting that exogenous and endogenous factors are empirically separable, and that both drive looking behavior even in young infants. In addition, Lewry et al. (2021) found that adults responded to magic tricks in ways that are predictable from patterns of infants' looking to those same 'tricks' in prior literature: the studies that evoked longer looking in youngest infants were also the events that adults found most interesting and surprising. Lastly, a body of work shows that computational models expressing probabilistic generative models of the physical and social world explain infants' and adults' behavioral responses much better than models that only use stimulus-driven computations (Shu et al., 2021; Smith et al., 2019; Zhi-Xuan et al., 2022; Sto-

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

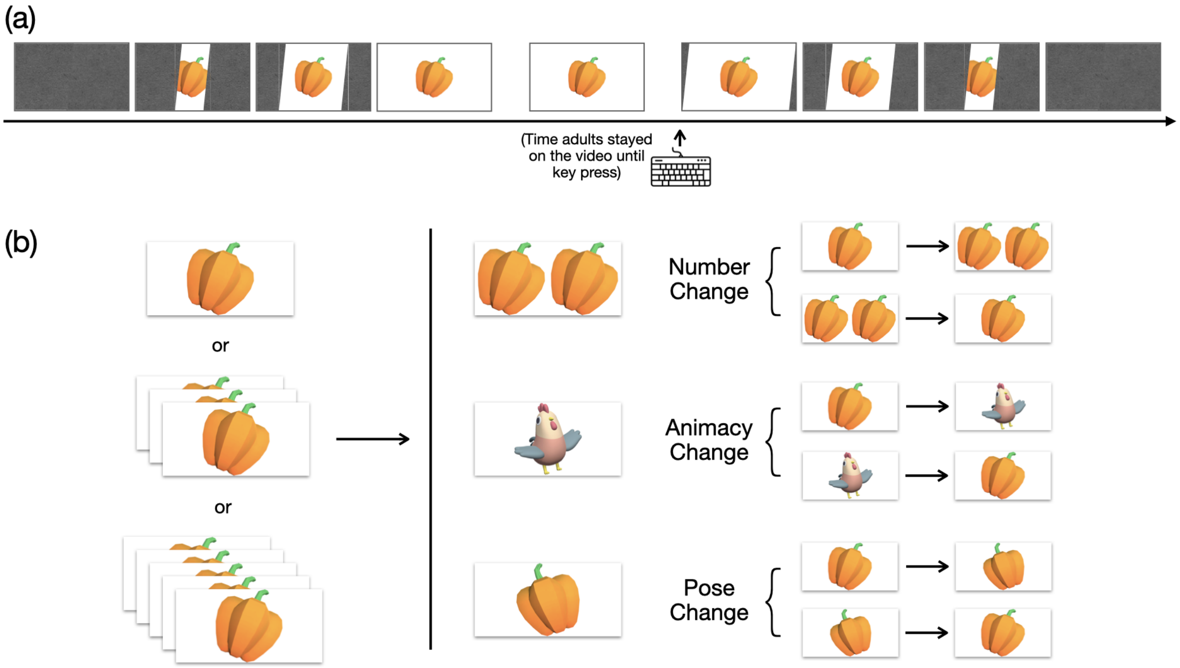

Figure 1: **Examples of stimuli change used in Raz, Cao et al. (2025)** (a) Illustration of a sample trial. In each trial, the curtain opens, revealing an object behind it, and closes again; this repeats on a loop. Adults can watch for as long as they like and proceed to the next trial by pressing a key, and the curtain closes to conclude the trial. The next trial starts immediately after that by reopening the curtain to reveal the next object. (b) Study design in Raz, Cao et al. (2025). On each block, adults first saw 1, 3, or 5 familiarization trials featuring the same stimulus, and then a test trial where a different stimulus was revealed. The object revealed in the test trial differed across conditions: animacy change, number change, and pose change. Two other conditions were used in Raz, Cao et al. (2025) but are not the focus of the current paper.

jnić et al., 2023). This work provides evidence that human learners, from infants to adults, represent the stimuli they see in terms of their relevance for learning, defined over both perceptual and conceptual content (Werchan & Amso, 2017; Raz & Saxe, 2020).

RANCH: A stimulus-computable model of looking behavior

What dimensions of stimuli define their relevance for learning? One challenge for studying this question using computational modeling is that these models should ideally be stimulus computable, so that quantitative comparisons and predictions can be made across a variety of tasks and stimulus sets. Inspired by this challenge, Raz, Cao et al. (2025) developed the Rational Action, Noisy Choice for Habituation (RANCH) model that can predict an observer's looking behavior (specifically, when to keep looking and when to stop given a stream of visual input). In this hierarchical Bayesian model, looking behaviors are modeled as a rational exploration process in which an observer, at each time step, takes a noisy perceptual sample from a visual stimulus, and then estimates the Expected Information Gain (EIG) of the next perceptual sample to decide whether to keep looking or to look away. Raz, Cao et al. (2025) tested the model in an experimental paradigm where a set of curtains repeatedly opened and closed to reveal an item (an animate object, e.g., a chicken; or an inanimate object, e.g., a pumpkin), and then opened to reveal a different item (Figure 1). They defined EIG over the learning goal of learning the visual concept of the item behind the curtain; as observers see the same item over and over, the EIG of continuing to look decreases, and when a new item appears, the EIG of looking increases, signaling that there is something new to learn.

In the empirical studies including the RANCH model, EIG was defined over an perceptual embedding space derived from ResNet-50, a deep convolutional neural network (CNN) that has been widely used in computer vision to account for human's visual representations (He et al., 2016). Using this perceptual embedding space, RANCH has been shown to account for the dynamics of habituation and dishabituation in human adults, preschoolers, and infants (Cao et al., 2023; Raz et al., 2023). In the most recent study, Raz, Cao et al. (2025) demonstrated that RANCH can predict graded habituation and dishabituation looking responses across an array of changes: animacy (e.g., from a chicken to a pumpkin), identity (e.g., from a chicken to a pig), pose (e.g., from facing left to facing right), and number (e.g., from one chicken to two chickens). RANCH could predict how long adults and infants looked to each of these changes (dishabituation responses were greatest for changes in animacy, then identity, then pose, then number), although with some discrepancies.

Present work

Here, we propose that while RANCH expresses the endogenous processes of sampling, learning, and decision-making characteristic of human observers of all ages, its perceptual

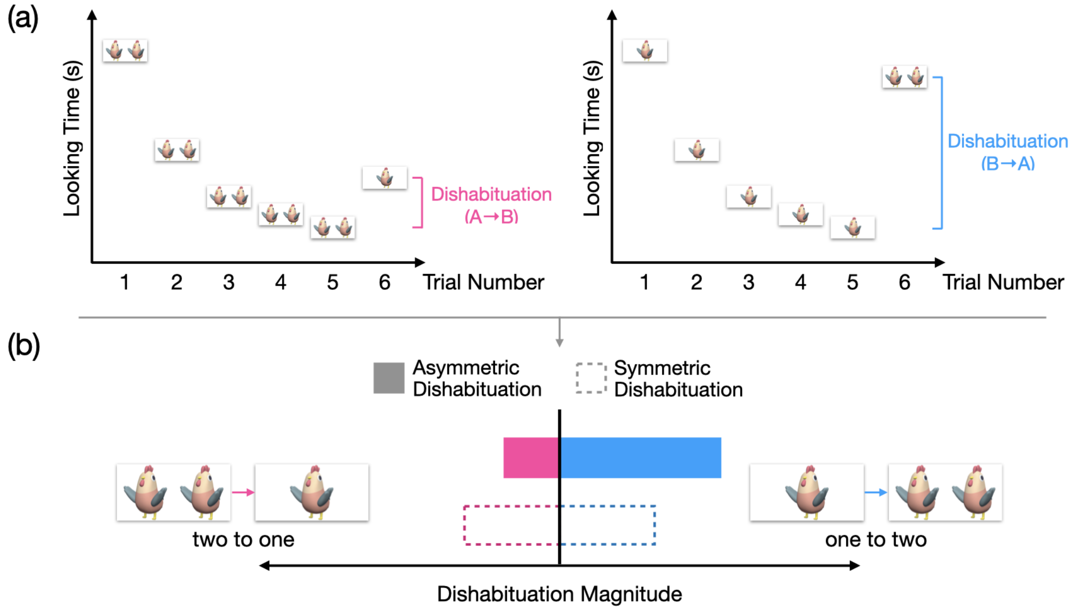

Figure 2: **Illustration of asymmetry of dishabituation (using number change as an example)** (a) Idealized example of asymmetric dishabituation when one object is removed from the scene (left, in pink), compared when one object is added (right, in blue). The x-axis indicates trial number, and y-axis indicates looking time at a trial. Dishabituation is the difference between looking time at the test trial and the last familiarization trial. (b) Another view of data from the same example in (a), showing asymmetric (solid) and symmetric (dotted) patterns of dishabituation. The midpoint of the x axis represents zero dishabituation, with greater distances reflecting larger magnitudes in each corresponding direction.

embedding space cannot fully explain adults' looking behavior. Specifically, we hypothesize that when adults in this experiment see a curtain fall over a pumpkin to reveal a chicken, adults do not merely see one image replaced by another; they also represent each image as a view of a permanent object, track that object under occlusion, and interpret the new revealed object in terms of that representation, in ways that go beyond its visual features.

One distinctive prediction of this hypothesis is that adults, undergoing the experiments carried out by Raz, Cao et al. (2025), should display *asymmetrical patterns of looking for the same pair of objects, depending on which object is occluded and which one is revealed* (Figure 2). In other work, adult observers treat changes in number (appearance and disappearance), and changes in animacy (from objects to agents or vice versa), as more than just reciprocal functions: They assign greater effort to conjuring than to vanishing an object (the operations entailed in a number change (McCoy & Ullman, 2019), and report that objects changing into animates is less plausible than animates changing into objects (Griffiths, 2015). Because these changes lead the observer to entertain different potential explanations, depending on the direction of the change, this may lead adults to look asymmetrically to number changes (singleton to pair vs pair to singleton) and animacy changes (object to agent vs agent to object). In addition, we predict that asymmetries will not appear for changes in pose (facing left, then facing right, vs the inverse), because such changes do not suggest changes to the identity of the object behind the curtain.

We tested these predictions in an exploratory secondary analysis of original data from Raz, Cao et al. (2025). As predicted, adults exhibited asymmetrical looking patterns for the cases of changes in number and animacy, but not in pose. Specifically, a singleton changing to a pair evoked bigger dishabituation than the reverse, and inanimate change to animate evoked bigger dishabituation than the reverse. The results were robust for number and pose when other related factors were accounted for, but asymmetry in animacy change was mitigated by people's baseline interest in the images involved in the change.

Methods

Description of dataset

The dataset we used came from Experiment 2 in Raz, Cao et al. (2025). In this experiment, adults from Prolific (N = 468) participated in a self-paced looking time experiment: On each trial, they viewed short movies of a set of curtains that repeatedly opened to reveal a 3D animated animal or object (Figure 1). Adults pressed a key to move onto the next trial whenever they wanted after a minimum viewing time of 500 ms; this response time was used as a proxy for looking duration, as in the original paper. Items within a block varied in number (a singleton or a pair), animacy (animals for animate stimuli, vegetables or fruits for inanimate stimuli), pose (facing left or right), and identity of the item (e.g., a cow, a penguin, a pumpkin, an orange). Each adult participated in 24 blocks of trials. During each block, adults first saw 1, 3, or 5 familiarization trials of one item, and then in a test trial, they saw either the same stimulus or a novel stimulus that changed on one of the four features. The order of the blocks was semi-randomized. Refer to Raz, Cao et al. (2025) for more details.

The main purpose of the current work was to test for asymmetries in dishabituation during blocks featuring changes in number, animacy, and pose, involving pairs of stimuli with equal perceptual distance. For each condition, we selected blocks where the same two stimuli appeared in the familiarization and test trials, but in opposite orders across blocks. Because stimuli did not repeat across blocks within a participant, no single participant saw the same pair of items swap roles within the study session; therefore, our key analysis relied on comparisons across participants. For each condition, all blocks with data for the last familiarization trial and test trial were included. The final dataset included 1856 pose blocks, 1852 number blocks, and 1856 animacy blocks, with data from all 468 participants featured in each condition. The number of blocks with 1, 3 vs 5 familiarization trials, and featuring each direction of change in each condition were roughly equal (see details in Supplementary Material.

Analysis approach

Dependent measure. We used adults' response time on each trial as a proxy for looking time, as in Raz, Cao et al. (2025). Similar to infants' looking time data (Csibra et al., 2016), adults' raw response time in this dataset also followed a right-skewed distribution rather than a normal distribution. Therefore, to measure how much people shifted their looking to the second item revealed by the curtain, relative to the first item that the curtain occluded, on each block, our main dependent measure was the difference between response times between the test trial, and the last familiarization trial, in log milliseconds.

Predictors. For each block type, we constructed a binary predictor encoding the direction of the change (for number, single to pair vs vice versa; for animacy, inanimate to animate vs vice versa; for pose, right to left vs vice versa). This was our most important predictor; if there is evidence for asymmetry in any of these conditions, then this predictor should explain variance in adults' dishabituation. We also wanted to account for other features of the experiment that likely explain variability in adults' looking behaviors, which were taken directly from the analyses in Raz, Cao et al. (2025): the number of familiarization trials adults saw before test (1, 3 or 5), which could shift the strength of the representation they have for that object prior to test, and the position of the current block (1...24) within the full experiment, which could influence adults' general attentiveness or fatigue.

We also wanted to account for adults' baseline interest in the stimuli (e.g., animals more interesting to look at than objects, pairs taking longer to process than singletons), because such effects in principle could lead to an asymmetrical dishabituation. For instance, a change from a singleton to a pair might appear to elicit longer looking than vice versa because a pair requires more time to process than a singleton. This asymmetry has been observed in oddball paradigms involving stimuli of varying intensity; for instance, deviant stimuli that are longer (vs shorter) in duration, evoke stronger behavioral and neural responses (Kask et al., 2021; Blundon et al., 2017). Thus, we wanted to account for baseline looking in the stimuli in our results. Specifically, we computed the average looking for a given stimulus when it appeared on the first familiarization trial across all blocks in all adults. We then entered this value as a nuisance predictor in subsequent analyses, serving as a proxy for the baseline tendency to look at each specific stimulus.

In all models, a random intercept for participants was included to capture individual variability in dishabituation.

Statistical models measuring asymmetry in dishabituation. We used linear mixed effects models, using `lme4` and `lmerTest` package (Kuznetsova et al., 2017; Bates et al., 2015) in R (Team, 2013), to test for the effect of change direction (i.e., the presence of an asymmetric dishabituation response) in each of the three conditions. To enable direct comparison between the strength of the predictors, all continuous predictors were mean-centered, and categorical predictors were coded using summed contrasts.

We first tested the effect of change direction (`change_direction`) in a simple model where change direction was the only fixed effect.

Next, we entered the predictors that were used in Raz, Cao et al. (2025): number of familiarization trials (`fam_trial_number`) and position of the block (`block_position` in log units). For each of these models, we computed the Akaike Information Criteria (AIC) , and chose the model with the lowest AIC. We also conducted quality assurance on the models (e.g., collinearity between the predictors, normality of the residuals, etc.) using the performance package (Lüdecke et al., 2021).

Lastly, we wanted to address the concern that any effects of asymmetry we observe could be accounted for by adults' intrinsic interest in the stimuli. For example, if adults look longer at animates than inanimates, all things being equal, then that could be sufficient to produce an asymmetry in their dishabituation from inanimates to animates and vice versa. Therefore we added the two predictors capturing baseline looking to the familiarization and test stimuli to the model.

The final model reads: `dishabituation ~ fam_trial_number + block_position_log + change_direction + fam_img_baseline_pref + test_img_baseline_pref + (1|subject)`.

Exploring effects of number, animacy and pose within blocks. Above, we described analyses aiming to estimate asymmetric effects on dishabituation, controlling for other predictors. Because blocks within the number, animacy, and pose conditions also varied along other dimensions (e.g., a number change block could include 1 *object* turning into 2 objects, *or* 1 *agent* turning into 2 agents), it is also possible to study how patterns of dishabituation vary across these changes. We compared the magnitude of the asymmetry qualitatively and reported how asymmetry differs given other features of the image that were not changed within a block, controlling for all of the predictors above.

All data and scripts required to reproduce the results and figures of this paper are available at https://osf.io/hc65m/.

Results

Analyses of asymmetry

Dishabituation in the number, animacy, and pose blocks is shown in Figure 3. When change direction was the only predictor in the model, adults showed asymmetric dishabituation for changes in number ($\beta = 0.08, SE = 0.02, p < .001$) and animacy ($\beta = 0.08, SE = 0.02, p = .001$), but not pose ($\beta = -0.00, SE = 0.02, p = .849$) (Figure 3a). Adults recovered more attention when one item turned into two ($M = -0.04$ log ms, $SD = 0.46$) than the reverse ($M = -0.12$ log ms, $SD = 0.45$), and they recovered more attention when inanimate(s) turned into animate(s) ($M = 0.20$ log ms, $SD = 0.51$) vs the reverse ($M = 0.12$ log ms, $SD = 0.54$. These results held, controlling for the predictors from Raz, Cao et al. (2025) (number: $\beta = 0.09, SE = 0.02, p < .001$; animacy: $\beta = 0.10, SE = 0.03, p = .003$; pose: $\beta = 0.00, SE = 0.02, p = .948$). However, when we then accounted for baseline looking for the two stimuli, this asymmetric effect held for number ($\beta = 0.08, SE = 0.03, p = .001$), but became weaker for animacy ($\beta = 0.10, SE = 0.06, p = .096$). The result for pose was unchanged ($\beta = 0.00, SE = 0.02, p = .953$). In sum, we found evidence that adults treated one singleton turning into a pair as different from a pair turning into a singleton.

Effects of number, animacy, and pose within blocks

We found asymmetry in number change and animacy change, but did other features of the stimuli influence the degree of asymmetry? As shown in Figure 3b, *asymmetry in number change* was stronger for animates (asymmetry $M = 0.12$ log ms, $SD_{pooled} = 0.03$) compared to inanimates (asymmetry $M = 0.04$ log ms, $SD_{pooled} = 0.03$), and the interaction between animacy and change direction was significant ($\beta = 0.10, SE = 0.05, p = .032$; see Supplementary Material for details). As for adults' *asymmetry in animacy change*, descriptively, it was stronger when adults saw a pair (asymmetry $M = 0.09$ log ms, $SD_{pooled} = 0.04$) compared to a singleton (asymmetry $M = 0.06$ log ms, $SD_{pooled} = 0.03$), but we did not find evidence supporting this descriptive effect in model results ($\beta = 0.01, SE = 0.05, p = .785$).

Discussion

In the current study, we conducted a secondary analysis of an adult dataset from Raz, Cao et al. (2025), who developed an image-computable model (RANCH) to predict looking behavior in adults and infants. RANCH model used a perceptual embedding space to define the expected information gain from taking an additional sample by continuing to look at a stimulus, and demonstrated that this space was sufficient to explain differences in adults' responses to changes in identity, number, animacy, and pose. Here, we tested whether there were systematic differences in adults' looking behavior that were not accounted for by the perceptual space used to embed the stimuli in RANCH model. We found evidence that distance between stimuli in perceptual space cannot fully account for adults' behaviors. Adults showed asymmetrical looking responses for changes in number (greater dishabituation when a singleton turned to a pair) and for changes in animacy (greater dishabituation when inanimate turned to animate), although this second result appeared to be partially explained by adults' baseline looking in the stimuli.

These results build on prior work reporting adults' explicit judgments and neural responses to "magical" changes in number (McCoy & Ullman, 2019; Balaban et al., 2023; Smith et al., 2019; Liu et al., 2024) and animacy (Griffiths, 2015). Violations of object solidity, continuity, and support evoke neural activity in brain regions specialized for object processing and endogenous attention, rather than visual processing, suggesting that they do not respond to these events as mere visual changes (Liu et al., 2024). Adults' explicit judgments about magical appearances vs disappearances (McCoy & Ullman, 2019) suggest that these two operations are not merely reciprocal. Here we show that these intuitions also guide adults' free viewing behavior of these events.

But *why* do adults think that making an object appear would require more effort than making it disappear (McCoy & Ullman, 2019), and why do they look longer at magical appearance vs disappearance in the current dataset? One potential mechanism underlying both effects is that adults ascribe object pointers to track objects under occlusion (Balaban et al., 2023; Mitroff et al., 2004; Scholl & Pylyshyn, 1999). When the curtain falls on one object and reveals two copies of the object, the observer has to re-identify the original object and assign a new pointer to its twin. But, when the curtain falls on two objects and reveals just one, the observer merely needs to delete an object pointer. Another account is that when adults see one object become two copies of that object, they are motivated to scan back and forth between the two to check whether they are identical, but do not engage in this process when one object is deleted from that array. These two accounts make distinct empirical predictions. The object pointer account predicts that looking will be sensitive to the cost of re-identifying the second object (e.g., if the object moves locations under occlusion), and the visual comparison account predicts that looking will be sensitive to how confusable the two objects are (e.g., identical, barely distinct, or very distinct in appearance). New experiments could be conducted to test these two accounts directly. For example, in a case where the second object that is revealed is visually distinct from the original, a pointer will still need to be assigned, but checking whether the two objects are identical will no longer be relevant.

Adults report that it is more surprising when a magician turns an object into an agent than vice versa (Griffiths, 2015), but we only found weak evidence for this asymmetry in this study. Why? One possibility is that because both categories

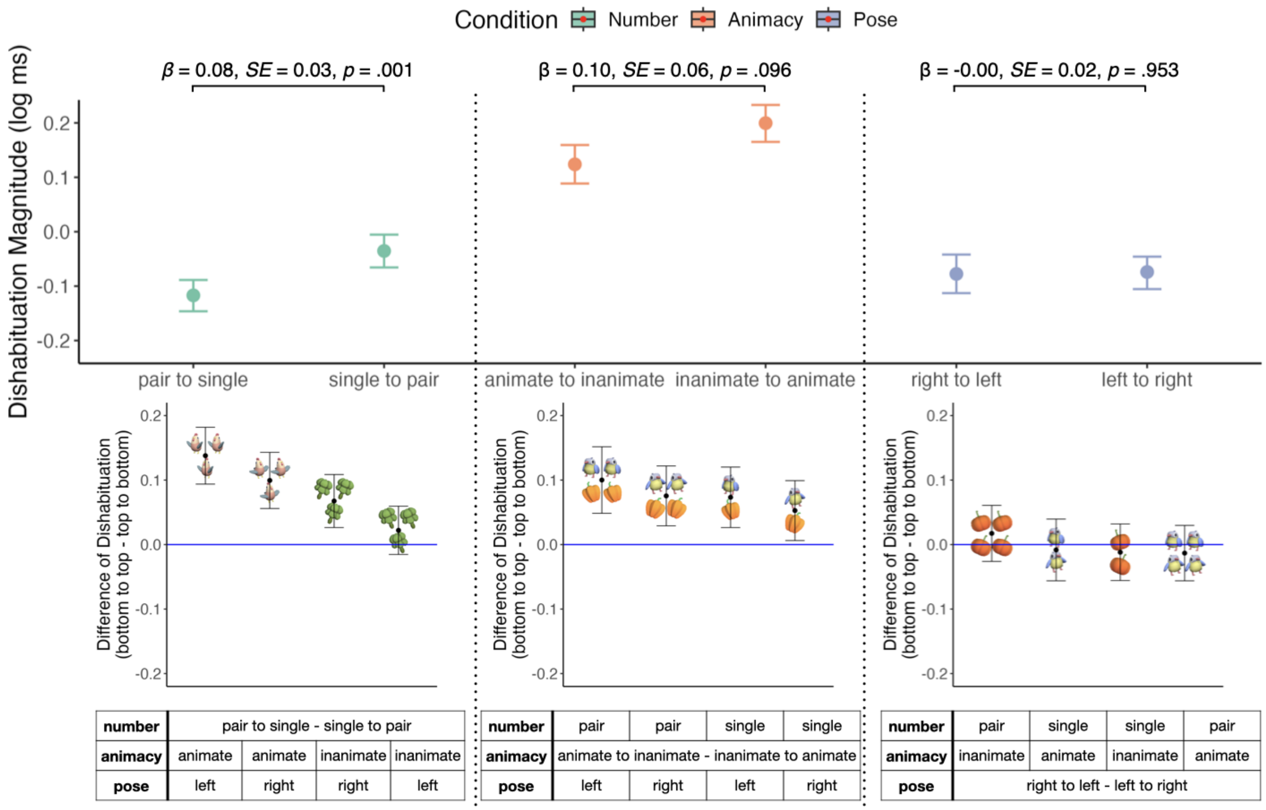

number	pair to single - single to pair			
animacy	animate	animate	inanimate	inanimate
pose	left	right	right	left

number	pair	pair	single	single
animacy	animate to inanimate - inanimate to animate			
pose	left	right	left	right

number	pair	single	single	pair
animacy	inanimate	animate	inanimate	animate
pose	right to left - left to right			

Figure 3: **Asymmetry of dishabituation in number, animacy, and pose blocks.** (a) Dishabituation magnitude given different change directions in each condition based on results of the full models including baseline looking preferences. (b) Asymmetry in number, animacy, and pose blocks, broken down by stimulus features. Y axis represents the magnitude of asymmetry. For example, the first panel shows that increased interest when one item is occluded and two are revealed (vs vice versa) is stronger for animates than objects.

of items moved on their own in Raz, Cao et al. (2025), adults construed changes in animacy as changes in an animate's appearance, rather than ontological status. In other words, adults may have interpreted the "dancing" vegetables and fruits in this study as agents. An experiment that parametrically manipulates the motions of inanimate objects and animate agents could be conducted to test this idea further. Additionally, this highlights a potential limitation of the RANCH model: it does not capture the motion of the item behind the curtain. But even if motion data were incorporated into the perceptual embedding space, that would still not describe adults' potential interpretations of motion (e.g. as indicative of an agent).

The richest interpretation for our results is that adults may be seeking to explain why or how the unexpected outcomes they observed happened. Even when young children see magic tricks similar to those in this paper, they seek explanations for underlying causes (Stahl & Feigenson, 2024; Perez & Feigenson, 2022). To test this idea further, we could introduce a clear explanation for the magical appearances and disappearances using the same methods of the current study (e.g., a box that the second object could have come out of, or disappeared into), and test whether that takes away both the longer looking at the unexpected outcomes, and the asymmetric responses observed in number change and animacy change.

Lastly, it is an entirely open question whether the asymmetry we observed here would also extend to human infants. While Raz, Cao et al. (2025) did test infants on these same changes, they did not test both directions of each change, thus we could not conduct the same analyses of asymmetry in infants. In the future, we could conduct experiments with infants using the same set of stimuli used in Raz, Cao et al. (2025) to test whether, like adults, infants show asymmetrical looking pattern in number change and animacy change. If so, this would support the proposal that infants interpret surprising events as more than merely visually new (Kunin et al., 2024).

Overall, our work highlights the importance of modeling people's conceptual understanding of objects, including a sense of object permanence, to make sense of their dishabituation behavior. Under our account, adults' responses to unexpected changes in number, and potentially changes in animacy, cannot be captured by merely traversing a perceptual embedding space. We propose that in addition to this space (which, together with RANCH, has been shown to capture key aspects of looking), it may be important to incorporate additional representations, like those inspired by probabilistic generative models of physics (Smith et al., 2019) and psychology (Shu et al., 2021). We predict that these additions will allow the model to better reflect how humans see the world, for adult and infant observers alike.

References

Aslin, R. N. (2000). Why take the cog out of infant cognition? *Infancy*, *1*(4), 463–470.

Baillargeon, R. (1994). Physical reasoning in young infants: Seeking explanations for impossible events. *British Journal of Developmental Psychology*, *12*(1), 9–33. doi: 10.1111/j.2044-835X.1994.tb00616.x

Baillargeon, R., Scott, R. M., & Bian, L. (2016). Psychological reasoning in infancy. *Annu. Rev. Psychol.*, *67*(1), 159–186.

Balaban, H., Smith, K., Tenenbaum, J., & Ullman, T. D. (2023). Electrophysiology reveals that intuitive physics guides visual tracking and working memory. *Open Mind*.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using **lme4**. *Journal of Statistical Software*, *67*(1). doi: 10.18637/jss.v067.i01

Blumberg, M. S., & Adolph, K. E. (2023). Protracted development of motor cortex constrains rich interpretations of infant cognition. *Trends Cogn. Sci.*, *27*(3), 233–245.

Blundon, E. G., Rumak, S. P., & Ward, L. M. (2017). Sequential search asymmetry: Behavioral and psychophysiological evidence from a dual oddball task. *PLOS One*, *12*(3), e0173237. doi: 10.1371/journal.pone.0173237

Cao, A., Raz, G., Saxe, R., & Frank, M. C. (2023). Habituation reflects optimal exploration over noisy perceptual samples. *Top. Cogn. Sci.*, *15*(2), 290–302.

Csibra, G., Hernik, M., Mascaro, O., Tatone, D., & Lengyel, M. (2016). Statistical treatment of looking-time data. *Dev. Psychol.*, *52*(4), 521–536.

Griffiths, T. L. (2015). Revealing ontological commitments by magic. *Cognition*, *136*, 43-48.

Haith, M. M. (1998). Who put the cog in infant cognition? is rich interpretation too costly? *Infant Behav. Dev.*, *21*(2), 167–179.

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *In Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.

Held, R., Birch, E., & Gwiazda, J. (1980). Stereoacuity of human infants. *Proc. Natl. Acad. Sci. U. S. A.*, *77*(9), 5572–5574.

Kask, L., Põldver, N., Lippus, P., & Kreegipuu, K. (2021). Perceptual Asymmetries and Auditory Processing of Estonian Quantities. *Frontiers in Human Neuroscience*, *15*. doi: 10.3389/fnhum.2021.612617

Kuhl, P. K. (2004). Early language acquisition: cracking the speech code. *Nat. Rev. Neurosci.*, *5*(11), 831–843.

Kunin, L., Piccolo, S. H., Saxe, R., & Liu, S. (2024). Perceptual and conceptual novelty independently guide infant looking behaviour: a systematic review and meta-analysis. *Nat. Hum. Behav.*, 1–15.

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, *82*(13). doi: 10.18637/jss.v082.i13

Lewry, C., Curtis, K., Vasilyeva, N., Xu, F., & Griffiths, T. L. (2021). Intuitions about magic track the development of intuitive physics. *Cognition*, *214*, 104762.

Liu, S., Lydic, K., Mei, L., & Saxe, R. (2024). Violations of physical and psychological expectations in the human adult brain. *Imaging Neuroscience*, *2*, 1–25.

Liu, S., Raz, G., Kamps, F., Grossmann, T., & Saxe, R. (2023). No evidence for discontinuity between infants and adults. *Trends Cogn. Sci.*, *27*(8), 694–695.

Lüdecke, D., Ben-Shachar, M., Patil, I., Waggoner, P., & Makowski, D. (2021). Performance: An R Package for Assessment, Comparison and Testing of Statistical Models. *Journal of Open Source Software*, *6*(60), 3139. doi: 10.21105/joss.03139

Margoni, F., & Thomsen, L. (2024). How infants predict respect-based power. *Cognitive Psychology*, *152*, 101671.

McCoy, J., & Ullman, T. (2019). Judgments of effort for magical violations of intuitive physics. *PLoS One*, *14*(5), e0217513.

Mitroff, S. R., Scholl, B. J., & Wynn, K. (2004). Divide and conquer: How object files adapt when a persisting object splits into two. *Psychological Science*, *15*(6), 420–425.

Perez, J., & Feigenson, L. (2022). Violations of expectation trigger infants to search for explanations. *Cognition*, *218*, 104942.

Raz, G., Cao, A., Bui, M. K., Frank, M. C., & Saxe, R. (2023). No evidence for familiarity preferences after partial exposure to visual concepts in preschoolers and infants. *PsyArXiv*.

Raz, G., Cao, A., Saxe, R., & Frank, M. C. (2025). A stimulus-computable rational model of habituation in infants and adults. *ELife*.

Raz, G., & Saxe, R. (2020). Learning in infancy is active, endogenously motivated, and depends on the prefrontal cortices. *Annu. Rev. Dev. Psychol.*, *2*(1), 247–268.

Scholl, B. J., & Pylyshyn, Z. W. (1999). Tracking multiple items through occlusion: Clues to visual objecthood. *Cognitive psychology*, *38*(2), 259–290.

Shu, T., Bhandwaldar, A., Gan, C., Smith, K. A., Liu, S., Gutfreund, D., ... Ullman, T. (2021). AGENT: A benchmark for core psychological reasoning. In *Proceedings of the 38th international conference on machine learning* (pp. 9614–9625).

Smith, K., Mei, L., Yao, S., Wu, J., Spelke, E., Tenenbaum, J., & Ullman, T. (2019). Modeling expectation violation in intuitive physics with coarse probabilistic object representations. *Adv. Neural Inf. Process. Syst.*, *32*.

Spelke, E. S. (2022). *What babies know: Core knowledge and composition* (Vol. 1). Oxford University Press.

Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, *10*, 89-96. doi: 10.1111/j.1467-7687.2007.00569.x

Stahl, A. E., & Feigenson, L. (2024). Young children distinguish the impossible from the merely improbable. *Proceedings of the National Academy of Sciences*, *121*(46), e2411297121. doi: 10.1073/pnas.2411297121

Stojnić, G., Gandhi, K., Yasuda, S., Lake, B. M., & Dillon, M. R. (2023). Commonsense psychology in human infants and machines. *Cognition*, *235*, 105406.

Team, R. C. (2013). *R: A language and environment for statistical computing.* R Foundation for Statistical Computing, Vienna, Austria.

Werchan, D. M., & Amso, D. (2017). A novel ecological account of prefrontal cortex functional development. *Psychological Review*, *124*(6), 720.

Zhi-Xuan, T., Gothoskar, N., Pollok, F., Gutfreund, D., Tenenbaum, J. B., & Mansinghka, V. K. (2022). Solving the baby intuitions benchmark with a hierarchically bayesian theory of mind. *arXiv*.