Looking without seeing: Children do not distinguish efficient from inefficient means to achieve a goal

Ori Ossmy¹, Danyang Han¹, Brianna Kaplan¹, Melody Xu¹, Catherine Bianco¹, Roy Mukamel²,
Karen E Adolph¹

¹ Department of Psychology and Center for Neural Science, New York University, New York, NY, United States.

² School of Psychological Sciences and Sagol School of Neuroscience, Tel-Aviv University, Tel Aviv, Israel.

Abstract

Observing actions provides important information about other people's goals and the means they use to achieve their goals. Preschoolers ($N=22$) and adults ($N=22$) watched video-recorded actors use efficient and inefficient means of grasping a hammer to pound a peg. Eye tracking showed that participants at both ages looked equally long at the goal target (the peg) but adults looked longer than children at the means—how actors grasped the hammer. Deep learning analysis of participants' eye gaze distinguished observation of efficient from inefficient grasps for adults, but not for children. Moreover, only adults showed differential physiological responses while observing efficient versus inefficient grasps in action-related neural activity (EEG) and pupil dilation. Thus, children can actively direct their gaze to look at goal-directed actions without seeing whether the means are efficient or not. Moreover, findings suggest that the development of action perception is built from children's own motor experiences.

Keywords: Action perception; development; EEG; eye tracking; deep learning

Introduction

Perceiving Goals and Means

Observing others is a powerful way to learn about goal-directed actions. But observational learning entails more than noticing the movement trajectory. It provides information about the actor's intended goal and the means to achieve it. For example, watching another person hammer a peg informs the observer about the actor's goal (to flatten the peg), and the means to achieve it (how to grasp and wield the hammer).

Even infants can perceive action goals (for review, see Elsner & Adam, 2020). For example, infants looked longer at actors reaching to a new object after seeing them reach repeatedly to an old object, suggesting infants perceived the actor's intention (Cannon, Woodward, Gradeback, von Hofsten, & Turek, 2012; Woodward, 1998). Similarly, infants perceive the end-goal of two-step actions. They looked longer at an actor pulling a cloth to retrieve a toy after seeing the same action with a different toy (Sommerville, Hildebrand, & Crane, 2008; Sommerville & Woodward, 2005; Woodward, 1998). Other measures support the conclusion that infants perceive action goals. Infants showed increased pupil dilation—an index of enhanced visual attention—when viewing an unexpected change in a goal (moving food-laden spoon to the hand rather than mouth; Gredebäck & Melinder, 2011). Infants also showed increased EEG activity in the gamma frequency band over primary motor cortex (Reid, Hoehl, & Striano, 2006) for incomplete actions (end goal not achieved) compared to completed actions (end goal achieved).

In addition to perceiving the goal, observing others' actions can yield information about the means to achieve the goal and whether the action is efficient or inefficient. However, previous work focused primarily on infants and yields discrepant findings. In some studies, infants showed evidence of perceiving the means to achieve a goal. For example, 9-month-olds looked longer when grip aperture of a reaching hand (the means) and the size of a cup (goal

target) were unmatched versus matched (Daum, Vuori, Prinz, & Aschersleben, 2009). Likewise, eye-tracking showed that 6- to 20-month-olds anticipate the goal based solely on how actors wield a tool (Hunnius & Bekkering, 2010; Paulus, Hunnius, & Bekkering, 2011). Moreover, 12- and 14-month-olds imitated unusual means to achieve a goal (turning on a light with their head rather than their hand), but only when the action was more efficient because the actor's hands were occupied (Gergely, Bekkering, & Király, 2002; Zmyj, Daum, & Aschersleben, 2009).

However, several studies found that infants and children do not attend to the efficiency of means to achieve a goal. Ten- to 12-month-olds showed no difference in looking time when actors changed the means to achieve a goal by pulling a different cloth or lifting a different box (Sommerville & Woodward, 2005; Woodward & Sommerville, 2000). Similarly, 12- and 18-month-olds imitated the end goal of putting a mouse in a house (Carpenter, Call, & Tomasello, 2005), but not the means to get there (hopping motions and distinctive sounds). Three- to 5-year-olds (Bekkering, Wohlschläger, & Gattis, 2000) imitated only the observed action goal (touching hand to ear) and not the means to achieve it (which hand touched target ear). Given the importance of observational learning for acquiring knowledge about goal-directed action, the question is still open as to whether young children notice efficient means to achieve a goal.

Theories on the Development of Action Perception

What contributes to children's perception of means for achieving a goal? "Motor-resonance" theory proposes that observers' own motor experiences contribute to perception of others' actions. In adults, action observation activates the neural motor program used in action production (Mukamel, Ekstrom, Kaplan, Iacoboni, & Fried, 2010; Ossmy & Mukamel, 2018). Thus, motor experience teaches children to attend to the relevant features of others' actions during action perception, just as it teaches children to attend to relevant sources of information when producing actions. Although previous work did not focus on perception of means, grasping experience facilitated 3-month-olds' perception of goals when watching actors grasping different toys (Sommerville, Woodward, & Needham, 2005). Infants do not make predictive eye

movements toward an action goal before they can perform the action themselves (Falck-Ytter, Gredeback, & von Hofsten, 2006). Twelve-month-olds showed anticipatory eye movements when watching actors move objects into a basket; but 6-month-olds (who have not yet mastered grasping and transporting actions) did not.

Alternatively, the “rational-action” theory proposes that infants and children come innately equipped with teleological reasoning about others’ actions (Gergely & Csibra, 2003). Using the principle of rational action, infants understand others’ means to achieve a goal by expecting the most efficient action given the situational constraints (Csibra & Gergely, 1998). For example, after infants repeatedly viewed an agent use high-amplitude “jumping” movements to clear a barrier, they were tested with the same high-amplitude movement without the barrier and a new direct-path movement also without the barrier (Biro & Leslie, 2007; Csibra, 2008; Sodian, Schoeppner, & Metz, 2004). Infants looked longer at the familiar, high-amplitude movement, presumably because it was inefficient in the barrier-free situation, compared to the more efficient, direct-path action.

But what happens when children’s own motor experience involves using irrational, inefficient means to achieve their goals? That is, do children perceive the efficiency of another person’s means when motor resonance violates the principle of rational action? Specifically, children habitually use inefficient, poorly planned (“irrational,” if you will) means for multi-step actions, where efficient action planning requires the initial movement to take the future end goal into account (Rosenbaum et al., 1990). For example, to pound a peg when the handle of a hammer points away from their dominant hand, as shown in Fig. 1A, adults use an atypical, underhand grip to grasp the hammer to enable a smooth transition to efficient hammering (Comalli et al., 2016). In contrast, children repeatedly use their habitual overhand grip, so they end up holding the handle in an awkward grasp to hammer the peg, forcing them to subsequently adjust their grip (Fig. 1B).

Current Study

Thus, we exploited children's inefficient performance in the hammering task to contrast motor-resonance and rational-action theories. We compared children's and adults' ability to obtain and register visual information about the means for producing multi-step actions. We tested children at an age (3-5 years) when they typically use inefficient means themselves. Children and adults observed video displays of actors using efficient and [inefficient means](https://nyu.databrary.org/volume/321/slot/44641/-?asset=231554) (Fig. 1A) to grasp a hammer [to grasp a hammer](https://nyu.databrary.org/volume/321/slot/44643/-?asset=231573) (Fig. 1B) to pound a peg. The logic of the experiment was simple: If both children and adults direct their gaze to the actor's initial grip and their physiological responses differentiate efficient from inefficient grips, then we can infer that participants applied the principle of rational action. Alternatively, the evidence would point to motor resonance if children fail to look at the actors' initial grip and/or fail to register the information obtained by looking.

We used eye tracking to assess visual attention—where participants looked and when. In particular, we asked whether both children and adults looked at the relevant part of the scene at the appropriate time to observe the means for pounding the peg (i.e., hammer and hand during actors' reach and grasp). We used deep artificial neural networks and probability maps to test whether participants' looking patterns distinguished efficient from inefficient grasps. Moreover, we used pupil dilation and electroencephalography (EEG) to test whether participants' physiological responses provided evidence of registering differences between efficient and inefficient means (see Krüger, Bartels, & Krist, 2019), that is whether participants “see” as well as look at the means.

Method

Videos of the displays, de-identified looking videos of the eye-tracking recordings, and EEG data are publicly available in the Databrary web-based library <https://nyu.databrary.org/volume/321>. With participants' permission, third-person videos of

their behaviors, individualized eye-tracking and EEG, and demographic data are also shared in Databrary with authorized researchers. The analysis codes are shared on Github (<https://anonymous.4open.science/r/321/>).

Participants

We tested 22 children from 3.09 to 5.49 years of age ($M = 4.06$ years; 11 girls) and 22 adults from 19.37 to 26.40 years of age ($M = 22.02$ years; 14 women). Children were recruited from a pool of families in the NYC area who had expressed interest in participating in research. Adults were recruited through word of mouth. Most participants were white and middle class. Children received a robot toy, photograph magnet, and tote bag as souvenirs of participation, and adults received a photo magnet. Data from 10 additional children were not analyzed because they did not attend to the displays (6 children) refused to wear the EEG cap (4 children). Data from 7 adults and 10 children were excluded due to experimenter error in syncing the EEG with eye tracking and displays (see participants labeled “excluded” in the shared Databrary volume).

All participants were healthy with normal vision. We verified hand dominance so as to match the actor’s dominant hand in the video displays to each participant’s dominant hand. All adults reported one dominant hand (right = 20, left = 2). All children also had one dominant hand (right = 19, left = 3) based on parents’ report of which hand children used to brush their teeth, cut with scissors, and eat with a spoon, and on our observations of children cutting with scissors and drawing a line.

Videos of Multi-Step Actions

Participants watched a set of “test videos” to determine their perception of multi-step actions, interspersed with “localizer” videos to determine their baseline neural responses to motor actions (see display videos at <https://nyu.databrary.org/volume/321>); all videos recorded at 30 Hz. Test videos showed three adult actors grasping a hammer and pounding a peg. In half

of the videos, actors used an efficient, underhand initial grip that led to an efficient, overhand grip to pound the peg (Fig. 1A). In the other half of the test videos, actors used an inefficient, initial grip that led to an awkward ulnar grip to pound the peg (Fig. 1B). To ensure that other elements of the display could not explain differences between actions, we varied the actor (several actors performed each type of action) and movement trajectories (differed across actors) across displays while keeping constant the visual saliency (identical environments across actions), target (all actors pounded the peg), and outcome (actors were always successful).

All videos began with the actors' hands placed flat on the table and the hammer handle pointing toward the actors' non-dominant hand. Actors began moving 2500 ms after the display onset to obtain reliable baseline measures of pupil size and the EEG signal. Each video was 5.4 seconds long. Videos were edited so that grasping the hammer and pounding the peg occurred at the same time in each clip, regardless of the actor. Moreover, we maintained similar luminance across displays to avoid artifacts in the pupil dilation analysis (Winn, Whitaker, Elliott, & Phillips, 1994). Right-handed participants viewed actors pounding the peg with their right hand and left-handed participants viewed actors pounding the peg with their left hand.

Localizer videos portrayed two additional actions (with unusual grips) to identify features in the neural signal (electrodes, frequencies, and times) that show sensitivity to action observation (see localizer videos at <https://nyu.databrary.org/volume/321/slot/44645/-?asset=231603>). In one localizer video, the action started with an overhand/radial grip and ended with an overhand/radial grip by way of a rotatory motion in which the wrist turned out away from the body and the fingers pointed downward toward the body (Fig. S1-A). In the other localizer video, the action started with an overhand/ulnar grip and ended with an underhand/ulnar grip (Fig. S1-B).

Procedure

As shown in Fig. 1C, participants sat in a child or adult-sized chair in front of a 60-cm widescreen LCD monitor (1920 × 1200 resolution) with the height and orientation of the monitor adjusted to participants' eye level. We dimmed the room light to improve the quality of the eye-tracking recording. A desk-mounted SMI eye tracker (SensoMotoric Instruments, RED, 120 Hz) recorded gaze location and pupil dilation. We used a 4-point routine to calibrate the tracker, then validated the calibration with a second 4-point routine, and finally used a third 4-point routine at the end of the session to ensure the calibration did not deviate (Franchak, Heeger, Hasson, & Adolph, 2016).

Next, we fitted participants with a child- or adult-sized, wireless, 32-channel EEG cap (Neuroelectronics, ENOBIO 32, 500 Hz). Data were recorded from 32 scalp electrodes at locations of the extended 10-20 system and from two electrodes on the right mastoid. The single-ended voltage was recorded between each electrode site and CMS/DRL electrodes.

Each participant watched 96 videos in total (48 test videos and 48 localizer videos), presented in 8 blocks of 12 videos interspersed with animated cartoons to maintain participants' motivation and attention. Each block contained the same grasp-to-pound action performed by all three actors. Actors were pseudorandomized within blocks (no more than 3 displays by the same actor consecutively) and the order of action blocks was randomized within age groups.

Eye-Tracking: Data Processing and Analyses

Areas of interest. We defined 5 rectangular areas of interest (AOIs) for each video frame using SMI BeGaze software (Fig. 2A): the hammer (average size of area within and across videos = 12,397 pixels); the pegboard (30,223 pixels); the actor's dominant and non-dominant hand (13,618 pixels for right hand; 12,961 pixels for left hand); and the actor's face (25,112 pixels). The AOIs for the hammer, peg, and dominant hand are essential to understand the planned action (yellow, red, and orange areas in Fig. 2A), whereas the actor's nondominant hand and face are areas known to attract observers' gaze (Foulsham, Cheng, Tracy, Henrich, &

Kingstone, 2010; Klin, Jones, Schultz, Volkmar, & Cohen, 2002; Shepherd, Steckenfinger, Hasson, & Ghazanfar, 2010), but are not essential to capture the action (blue areas in Fig. 2A).

To determine the percentage of time participants' gaze tracked the action in the videos, we conducted dynamic area-of-interest analyses using the SMI BeGaze software. For each video frame, participants could look at one AOI, two or more AOIs (e.g., when the actor grasped the hammer or pounded the peg), or irrelevant areas (table, wall, etc.). Conflicts between overlapping AOIs were resolved by attributing gaze to the front-most AOI based on depth order. Percent "dwell" time was the total time gaze was on an AOI divided by the total tracking time. Percent "fixation" time was the total time gaze was on an AOI for at least 100 ms divided by the total fixation time. Frequency of "revisits" was the number of times participants shifted their gaze to an AOI from outside the AOI.

Scanning patterns. To test participants' scanning patterns, we calculated moment-to-moment gaze shifts from one AOI to another. Gaze shifts occurred when participants fixated an AOI and then fixated a different AOI within 100 ms. We counted the number of gaze shifts between each pair of AOIs and averaged them across participants in each age group to obtain a 5x5 matrix per age group, where each cell represents the average number of gaze shifts between each pair of AOIs.

Temporal probability maps. Looking patterns from movement onset to the initial grasp reflect participants' attention to the means for achieving the goal. Thus, for each participant, we determined the probability of looking at the actor's reach trajectory and grasp of the hammer handle by dividing the number of video frames each pixel was fixated (within a 75-pixel radius of participants' gaze location to account for detection resolution) by the total number of video frames.

Averaging the probabilities across participants in each age group yielded two "temporal probability maps." We then compared the maps using an unequal variance *t*-test for each pixel

and corrected for multiple comparisons by controlling the False Discovery Rate (FDR; Benjamini & Hochberg, 1995) with a threshold of $q(\text{FDR}) < 0.05$.

Deep-learning analysis. We used a deep-learning video classifier (Yue-Hei Ng et al., 2015) to determine whether participants gathered sufficient visual information to distinguish between the efficient and inefficient action plans. For each participant and trial, we created a “looking video” to represent the visual information gathered in each trial. That is, for each video frame in the looking video, the area within a 75-pixel radius of participants’ gaze was identical to the original display and the rest of the frame was black (Fig. S2; looking videos at <https://nyu.databrary.org/volume/321>). The resulting looking videos were used as input to the classifier. If participants paid sufficient visual attention to action means, the classifier should be able to correctly distinguish efficient from inefficient actions based on the looking videos.

We applied the classifier separately for each participant. To classify the test videos based on the derived looking videos, we created a deep-learning neural network that combines a pre-trained image classification model, convolutional neural network (CNN), and long-short term memory (LSTM) network. We used the Matlab Deep Learning Toolbox to (1) load a pre-trained GoogleNet model; (2) convert each video to a feature sequence using CNN; (3) create an LSTM network that classifies feature sequences; and (4) assemble a unified classification network using layers from both CNN and LSTM.

For each participant, we randomly chose one trial from a test video with an efficient action plan and one trial from a test video with an inefficient action plan as a “test set.” The remaining 46 trials (23 efficient and 23 inefficient trials) were used as a “training set” for the classifier. We assessed the accuracy of the classifier by comparing its assigned labels to the trials which it was not trained on and the actual labels of the original test videos. Classification accuracy was assessed as the average performance level across 500 randomly chosen trials for the test-set (“leave-one-out” procedure). Significance was assessed by comparing the

accuracy levels to a null distribution ($p < 0.05$) obtained by shuffling the labels of the trials (1000 shuffles) and performing the same classification procedure as on the original data.

Pupil Dilation: Data Processing and Analyses

We analyzed pupil dilation as in van Rij, Hendriks, van Rijn, Baayen, and Wood (2019). We used SMI BeGaze software to remove artifacts (blinks and saccades) from the pupil dilation signal. Visual inspection of the data corroborated automatic artifact rejection.

To analyze pupil response in relation to the actor's movements in the videos, we down-sampled the pupil data to 30 Hz and aligned it to movement onset. "Baseline" pupil size was the average pupil size during the first 1250-ms time window (in the period before movement began, as is customary). For each video frame of each trial, we defined pupil response as a percent of that trial's baseline pupil size.

EEG: Data Processing and Analyses

Data were analyzed offline using the EEGLAB tool for MATLAB (Delorme & Makeig, 2004). Raw EEG data were band-pass filtered offline between 1 and 45 Hz (Butterworth filter, 24 db), and re-referenced offline to the digital average of the two mastoids. The continuous data were segmented into epochs from -1.5 s to +2.9 s relative to movement onset in each video.

Time-frequency analysis. Event-related spectral perturbation (ERSP) was computed using a continuous Morlet wavelet transform. For each participant, we computed the logarithm of the power (from 0 s to 2.9 s post-movement onset) relative to power during baseline (from -1.5 s to 0 s before movement onset). We defined suppression indices in the various frequencies as the average log-ratio values. A negative log-ratio indicates a suppression in the power of EEG oscillations relative to baseline, whereas positive log ratios indicate enhanced power.

Action observation localizer. Participants' brains differ as do participants' neural activity while observing actions. Thus, we used an individualized, data-driven approach to determine specific differences in neural activity due to viewing the efficient versus inefficient test video. The EEG signal to the localizer videos provides the frequencies, timing, and channels of

neural activity related to action observation for each participant. We focused on significant suppression in the signal during the time between movement onset and hammering in the 6-20 Hz frequency range (frequencies previously reported as associated with action observation in children and adults; (Aridan, Ossmy, Buaron, Reznik, & Mukamel, 2018; Liao, Acar, Makeig, & Deak, 2015). To this end, a nonparametric cluster analysis was performed on the power in bilateral sensorimotor sites (FC1-5, C1-6, CP1-6) and bilateral occipital sites (O1 and O2) during action observation. Statistical significance was assessed with a well-established procedure that accounts for multiple comparisons (Maris & Oostenveld, 2007). Electrodes with significant suppression were used for classification of data recorded during observation of the test-videos.

Classification analysis. We examined whether the efficient and inefficient planning could be distinguished based on the neural activity related to action observation for each participant. Thus, for each participant, we constructed a support vector machine (SVM) classifier. The classifier was provided with the label of each trial (efficient or inefficient planning) and the corresponding power suppression in the time and frequency windows showing significant suppression as detected from the localizer data. We used a Matlab implementation of binary SVM (Chang & Lin, 2011) and applied a least-squares cost function ($C=1$).

To determine the classification accuracy for each participant, we used a standard leave-one-out classification procedure in which we randomly chose one trial of an efficient and inefficient test video to serve as the “test set” and the classifier was trained on the remaining dataset. Following training, classification performance was assessed by comparing the classifier’s assigned labels to the test set (the trials which it was not trained on) and their real labels. The average performance level across 500 permutations was assigned as the classification accuracy. Similar to the deep learning analysis, significance was assessed by comparing the accuracy levels to a null distribution ($p < .05$) obtained by shuffling the labels of

the data (1000 shuffles) and performing the same classification procedure as on the original data.

Results

Overall, children and adults looked equally long at the displays during the reach and initial grasp and during the pounding of the peg ($ps > .10$ for all looking measures). Likewise, children and adults looked equally long to displays of the efficient and inefficient actions (all $ps > .10$). Below, we report group differences in looking patterns—where participants looked and when—and physiological responses to the two types of displays.

Looking at Areas of Interest

For most children, looking at task-relevant areas of interest (warm colors in Fig. 2A) was lower than adult levels (red lines in Fig. 2B). However, every measure of looking—fixation time, dwell time, and gaze revisits—increased with age; $rs(20) > .43$, $ps < .05$; scatters in Fig. 2B.

As shown in Fig. 2C, children and adults looked equally long at the final goal target (peg), but children looked less at the task-relevant areas for the means (hammer and dominant hand). ANOVAs on fixation time, percent dwell time, and number of revisits, using 2 (ages) $\times$ 2 (actions: efficient and inefficient) $\times$ 5 (AOIs) as factors, showed main effects only for age and AOI, and interactions between age and AOI (Table 1). Sidak-corrected post hoc comparisons confirmed higher percentages in all looking measures in adults than children at the hammer and dominant hand, and no age differences in looking at the peg and non-relevant AOIs (face and non-dominant hand).

Adults also showed more gaze shifts between task-relevant AOIs than children. Each vertex in Fig. 2D represents an AOI, and the thickness of each black line represents the average number of transitions between the pair of AOIs in both directions (e.g., a gaze shift from the dominant hand to the hammer and a gaze shift from the hammer to the dominant hand count as 2 transitions between these AOIs). A 2 (ages) $\times$ 10 (pairs of AOIs) ANOVA confirmed

main effects for age, pair, and an interaction between age and pair (Table 2). Sidak-corrected post hoc tests showed that the hammer-to-peg and hammer-to-dominant-hand transitions were higher in adults compared to children.

Looking at the Right Thing at the Right Time

Accumulated measures of looking do not necessarily reflect the temporal dynamics of looking in the key moments of the action. Observers must look at the task-relevant AOIs at the right time in the action sequence. As shown in the temporal probability maps in Fig. 3A, adults did so more than children. Specifically, adults were more likely than children to fixate the handle of the hammer when the actor reached and grasped it.

Fig. 3A visually illustrates the average temporal probability map across all video frames (movement onset to initial grip) for each age group by varying the opacity of each pixel. If the average probability of fixating a specific pixel is 0 (i.e., no participant fixated that pixel), the pixel is completely opaque (as in the outer regions of the child and adult maps). The pixel becomes less opaque as the probability of fixation increases (as in the hammer region of the adult map). When the probability of fixating a specific pixel is 1 (i.e., every participant fixated that pixel), the pixel is completely transparent, and the display is clearly visible in the figure. Group comparisons for each pixel showed significant differences inside the circular region in the difference map (map at far right of Fig. 3A); $t_{(42)} > 3.58$, $p < .05$, false discovery rate corrected for multiple comparisons). In other words, adults looked more than children at the pixels inside the red circle when the actor reached out and grasped the handle.

Finally, deep-learning analysis successfully distinguished efficient from inefficient planning based on the "looking-videos" of every adult. In contrast, the analysis successfully distinguished test videos only for a small subset of children (6 of 22). As shown in Fig. 3B, the mean accuracy across participants of the CNN-LSTM classifier for the type of action (efficient vs. inefficient) was $M = 76.2\%$, $SD = 11.59$, when it was trained on adults' looking patterns, but only $M = 64.9\%$, $SD = 8.32$, when it was trained on children's looking patterns. Classification

accuracy in the children was unrelated to their age $r(20) = .21, p = .34$. Taken together, our findings show that children and adults gathered different visual information as expressed by differences in the temporal dynamics of their looking patterns.

Pupil Dilation

Pupil dilation reflects physiological responses to the displays. Adults showed greater pupil dilation when observing inefficient planning than when observing efficient planning at $M = 0.7$ - 1.4 seconds after the moment of initial grip (Fig. 4, gray area); $t_s(42) > 2.02, p_s < .05$. Increased dilation was absent in children (Fig. 4; $t_s(42) < 0.83, p_s > .40$). Given that the known delay in pupil dilation is 1 - 1.5 seconds after the visual event (Krüger et al., 2019; van Rij et al., 2019), this finding indicates a difference in adults' attention to the different grips. The lack of difference in children's pupillary response corroborates their lack of differentiation in patterns of looking.

Neural Activity

As with pupil dilation, EEG and classification analyses showed differential physiological responses in adults, but not children. We analyzed the neural data in two stages—localization and classification. In the localization stage, we used the data obtained during observation of two additional videos to identify frequencies, times, and electrodes in which neural activity is sensitive to action observation. We used a data-driven approach using nonparametric cluster analysis to examine significant event-related spectral perturbation (ERSP) locked to movement onset in all participants and electrodes. Table 3 shows the frequencies of the individual EEG signal in which oscillation power in sensorimotor sites was significantly suppressed relative to baseline. The average multi-participant ERSP maps for each group (Fig. 5A) is consistent with previous studies (Fox et al., 2016) and confirm an activity suppression in the mu band (children: 6 - 9 Hz, $M = -.94$ $SD = .86$, $t(21) = 5.13, p < .00$; adults: 8 - 13 Hz, $M = -.42$ $SD = .6$, $t(21) = 3.23, p < .00$; unequal variance t -test compared to zero). These findings indicate increased event-related desynchronization of the EEG signal in both groups, following observation of others' actions.

In the classification stage (Fig. 5B), we examined whether neural signals evoked by the test videos in the areas defined by the individual subject localizer, distinguish between action types. Significant, above-chance classification performance was obtained in adults ($M = 79.2\%$; Fig. 5C). However, children did not show significant neural differentiation between the two actions ($M = 57.3\%$).

Discussion

We used a novel combination of methods (eye tracking, pupillary responses, EEG, and machine learning) to test whether children differentiate the efficiency of means to achieve a goal in a task where they do not use efficient means to perform the action themselves—grasping a hammer to pound a peg. Preschoolers and adults watched actors use efficient and inefficient means of grasping a hammer to pound a peg. On every measure, adults showed clear evidence of differentiating efficiency of means, but preschoolers did not.

Motor Resonance or the Principle of Rational Action?

Our displays were designed to disentangle the motor-resonance and rational-action theories of action perception. In prior work on action production, adults performed the efficient grasping action nearly uniformly: They used overhand grips when the hammer handle pointed toward their dominant hand and underhand grips when the handle pointed toward their non-dominant hand. In contrast, preschoolers rarely performed the efficient, underhand grip when the hammer handle pointed toward their non-dominant hand (Comalli et al., 2016). Instead, they used the same overhand grip on most trials, so their motor experience includes only the inefficient action. If, as rational-action theory proposes, children use the “principle of rational action” to assess action efficiency (Gergely & Csibra, 2003), they should differentiate efficient and inefficient actions, regardless of their own motor experiences. But they did not. Instead, only adults showed differential looking patterns and neural activity. Thus, our findings are consistent with the motor resonance account—that observed actions activate observers’ own motor

programs built from their motor experiences, and so motor resonance generated differential responses only in adults.

Looking Without Seeing

Use of eye tracking, pupillary responses, and EEG allowed us to obtain evidence about both where participants directed their gaze and what they “saw” or differentiated. Preschoolers looked at the end goal (peg) just as much as adults did. However, adults gathered *more* relevant visual information about the means for achieving the end goal evidenced by longer looking to the performing hand and hammer, more gaze shifts between them, and more fixations on key areas that distinguished efficient from inefficient means. Moreover, adults exhibited differences in pupillary response and action-related neural activity when they observed efficient versus inefficient grasping actions, but preschoolers did not.

These findings suggest two related explanations for children’s failure to distinguish the means, both consistent with motor resonance theory. First, children do not look at the right thing at the right time. Second, even if children had gathered the relevant visual information at the moments surrounding the grasp (and a handful of children did), their brains failed to process it adequately to notice differences in means. In other words, looking is necessary but not sufficient to ensure that the relevant information is processed. Indeed, our findings show that children’s action-related neural activity while observing actions does not guarantee that they registered the way the action was performed. In many prior studies, children showed mu suppression when watching others, suggesting that the observed actions are simulated in their own motor system (e.g., Fox et al., 2016). We also found significant mu suppression in both groups. However, unlike adults, children’s patterns of neural activity did not differ *between* efficient and inefficient means. This finding points to an important, unanswered question regarding action-related neural activity: What features of the observed action are encoded in children’s mu suppression and how these features change with development.

The Cost of Anticipatory Looking

We offer an additional suggestion. Children's failure to notice the means might result from their prediction of the end goal. Similar to previous eye-tracking studies (Cannon et al., 2012; Falck-Ytter et al., 2006; Gredebäck & Melinder, 2011), we found that children shift their gaze toward the actor's goal target (here, the peg) *before* the action is completed. Possibly, early gaze shifts come with a cost. If young children look at the goal target too soon—without frequent gaze shifts between peg and hammer as do adults, children may fail to notice the means. Thus, anticipatory eye movements reveal children's ability to understand others' motivations and intentions but may also indicate a failure in perceiving how others performed the action.

Conclusion

We provide evidence that children fail to register critical elements in others' multi-step actions. They do not distinguish efficient from inefficient actions. Children also show deficits in planning and executing such actions themselves. Our findings suggest that motor resonance (i.e., children's own motor experience), and not teleological reasoning, underlies the development of perception of action means.

Acknowledgments

This work was supported by NSF/SBE-BSF grant #1627993 and DARPA grant N66001-19-2-4035 to Karen Adolph and BSF grant #20168

References

- Aridan, N., Ossmy, O., Buaron, B., Reznik, D., & Mukamel, R. (2018). Suppression of EEG mu rhythm during action observation corresponds with subsequent changes in behavior. *Brain research*, 1691, 55-63.
- Bekkering, H., Wohlschlager, A., & Gattis, M. (2000). Imitation of gestures in children is goal-directed. *The Quarterly Journal of Experimental Psychology: Section A*, 53(1), 153-164.
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57, 289-300.
- Biro, S., & Leslie, A. M. (2007). Infants' perception of goal-directed actions: development through cue-based bootstrapping. *Developmental Science*, 10(3), 379-398.
- Cannon, E. N., Woodward, A. L., Gradeback, G., von Hofsten, C., & Turek, C. (2012). Action production influences 12-month-old infants' attention to others' actions. *Developmental Science*, 15, 35-42.
- Carpenter, M., Call, J., & Tomasello, M. (2005). Twelve-and 18-month-olds copy actions in terms of goals. *Developmental Science*, 8(1), F13-F20.
- Chang, C. C., & Lin, C. J. (2011). LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2, 27.
- Comalli, D. M., Keen, R., Abraham, E., Foo, V. J., Lee, M. H., & Adolph, K. E. (2016). The development of tool use: Planning for end-state comfort. *Developmental Psychology*, 52, 1878-1892.
- Csibra, G. (2008). Goal attribution to inanimate agents by 6.5-month-old infants. *Cognition*, 107(2), 705-717.
- Csibra, G., & Gergely, G. (1998). The teleological origins of mentalistic action explanations: A developmental hypothesis. *Developmental Science*, 1(2), 255-259.

- Daum, M. M., Vuori, M. T., Prinz, W., & Aschersleben, G. (2009). Inferring the size of a goal object from an actor's grasping movement in 6- and 9-month-old infants. *Developmental Science*.
- Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of neuroscience methods*, 134(1), 9-21.
- Elsner, B., & Adam, M. (2020). Infants' Goal Prediction for Simple Action Events: The Role of Experience and Agency Cues. *Topics in Cognitive Science*.
- Falck-Ytter, T., Gredeback, G., & von Hofsten, C. (2006). Infants predict other people's action goals. *Nature neuroscience*, 9(7), 878-879.
- Foulsham, T., Cheng, J. T., Tracy, J. L., Henrich, J., & Kingstone, A. (2010). Gaze allocation in a dynamic situation: Effects of social status and speaking. *Cognition*, 117, 319-331.
- Fox, N. A., Bakermans-Kranenburg, M. J., Yoo, K. H., Bowman, L. C., Cannon, E. N., Vanderwert, R. E., . . . Van IJzendoorn, M. H. (2016). Assessing human mirror activity with EEG mu rhythm: A meta-analysis. *Psychological Bulletin*, 142(3), 291-313.
- Franchak, J. M., Heeger, D. J., Hasson, U., & Adolph, K. E. (2016). Free-viewing gaze behavior in infants and adults. *Infancy*, 21, 262-287. doi:10.1111/infa.12119
- Gergely, G., Bekkering, H., & Király, I. (2002). Rational imitation in preverbal infants. *Nature*, 415(6873), 755-755.
- Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naive theory of rational action. *Trends in Cognitive Sciences*, 7(7), 287-292.
- Gredebäck, G., & Melinder, A. (2011). Teleological reasoning in 4-month-old infants: pupil dilations and contextual constraints. *PLoS ONE*, 6(10), e26487.
- Hunnius, S., & Bekkering, H. (2010). The early development of object knowledge: A study of infants' visual anticipations during action observation. *Developmental Psychology*, 46(2), 446.

- Klin, A., Jones, W., Schultz, R., Volkmar, F., & Cohen, D. (2002). Visual fixation patterns during viewing of naturalistic social situations as predictors of social competence in individuals with autism. *Archives of general psychiatry*, 59, 809-816.
- Krüger, M., Bartels, W., & Krist, H. (2019). Illuminating the Dark Ages: Pupil Dilation as a Measure of Expectancy Violation Across the Life Span. *Child Development*, 00, 1-16.
- Liao, Y., Acar, Z. A., Makeig, S., & Deak, G. (2015). EEG imaging of toddlers during dyadic turn-taking: Mu-rhythm modulation while producing or observing social actions. *Neuroimage*, 112, 52-60.
- Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG-and MEG-data. *Journal of neuroscience methods*, 164(1), 177-190.
- Mukamel, R., Ekstrom, A. D., Kaplan, J., Iacoboni, M., & Fried, I. (2010). Single-neuron responses in humans during execution and observation of actions. *Current Biology*, 20(8), 750-756.
- Ossmy, O., & Mukamel, R. (2018). Perception as a route for motor skill learning: perspectives from neuroscience. *Neuroscience*, 382, 144-153.
- Paulus, M., Hunnius, S., & Bekkering, H. (2011). Can 14-to 20-month-old children learn that a tool serves multiple purposes? A developmental study on children's action goal prediction. *Vision Research*, 51(8), 955-960.
- Reid, V. M., Hoehl, S., & Striano, T. (2006). The perception of biological motion by infants: An event-related potential study. *Neuroscience Letters*, 395(3), 211-214.
doi:10.1016/j.neulet.2005.10.080
- Rosenbaum, D. A., Marchak, F., Barnes, H. J., Vaughan, J., Slotta, J., & Jorgensen, M. (1990). Constraints for action selection: Overhand versus underhand grips. In M. Jeannerod (Ed.), *Attention and Performance XIII: Motor representation and control* (pp. 321-342). Hillsdale, NJ: Lawrence Erlbaum Associates.

- Shepherd, S. V., Steckenfinger, S. A., Hasson, U., & Ghazanfar, A. A. (2010). Human-monkey gaze correlations reveal convergent and divergent patterns of movie viewing. *Current Biology*, 20, 649-656.
- Sodian, B., Schoeppner, B., & Metz, U. (2004). Do infants apply the principle of rational action to human agents? *Infant Behavior and Development*, 27(1), 31-41.
- Sommerville, J. A., Hildebrand, E. A., & Crane, C. C. (2008). Experience matters: The impact of doing versus watching on infants' subsequent perception of tool-use events. *Developmental Psychology*, 44, 1249-1256.
- Sommerville, J. A., & Woodward, A. L. (2005). Pulling out the intentional structure of action: The relation between action processing and action production in infancy. *Cognition*, 95, 1-30.
- Sommerville, J. A., Woodward, A. L., & Needham, A. W. (2005). Action experience alters 3-month-old infants' perception of others' actions. *Cognition*, 96, B1-B11.
- van Rij, J., Hendriks, P., van Rijn, H., Baayen, R. H., & Wood, S. N. (2019). Analyzing the time course of pupillometric data. *Trends in Hearing*, 23, 1-22.
- Winn, B., Whitaker, D., Elliott, D. B., & Phillips, N. J. (1994). Factors affecting light-adapted pupil size in normal human subjects. *Investigative ophthalmology & visual science*, 35(3), 1132-1137.
- Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, 69(1), 1-34.
- Woodward, A. L., & Sommerville, J. A. (2000). Twelve-month-old infants interpret action in context. *Psychological Science*, 11(1), 73-77.
- Yue-Hei Ng, J., Hausknecht, M., Vijayanarasimhan, S., Vinyals, O., Monga, R., & Toderici, G. (2015). *Beyond short snippets: Deep networks for video classification*. Paper presented at the Proceedings of the IEEE conference on computer vision and pattern recognition.
- Zmyj, N., Daum, M., & Aschersleben, G. (2009). The development of rational imitation in 9-and 12-month-old infants. *Infancy*, 14(1), 131-141.

Figures

Fig. 1. Displays and experimental set up. Participants observed an adult actor pound a peg, using a hammer and pegboard. Video displays were designed to distinguish perception of efficient versus inefficient means of using a hammer to pound a peg. In all displays, the action started with the hammer handle pointing away from the dominant hand. **(A)** In the “efficient” displays, the action began with an underhand, radial grip that led to a comfortable overhand, radial grip to pound the peg. This is the grip used by adults in previous work. **(B)** In the “inefficient” displays, the action began with a habitual, overhand, ulnar grip that led to inefficient tool use. This is the grip used by children in previous work. **(C)** Experimental set up for children with remote eye tracker and wireless EEG. Adults sat on an adult-sized chair.

Fig. 2. Looking at areas of interest (AOIs). (A) Areas of interest for eye-tracking analysis. We calculated amount of looking to 5 AOIs during each video: 3 task-relevant AOIs—dominant hand (orange), hammer (yellow), and peg (red), and 2 task-irrelevant AOIs—the face and the non-dominant hand (blue). Right-handed participants observed displays of actors using their right hand to pound the peg, and vice versa for left-handed participants. (B) Correlations between children's age and measures of looking to the AOIs: percent fixation time, percent dwell time, and number of revisits ($p < .05$). Red dotted lines indicate adult averages. (C) Measures of looking to AOIs for each group averaged across videos and trials. In terms of fixation time, dwell time and number of revisits, adults looked significantly more at the hammer and the dominant hand (all $p < .01$) compared to children. (D) Transitions between each pair of AOIs. Line thickness denotes the average number of transitions between AOIs. Compared to children, adults shifted their gaze more between the dominant hand and the hammer ($p < .01$), and between the hammer and the peg ($p < .05$).

Fig. 3. Looking patterns. **(A)** Probability maps representing the likelihood that children (left) and adults (middle) fixated each pixel during different intervals of the action. If the pixel is white, then the group's probability of fixating that pixel is 0. As the probability increased, the pixel is clearer. When the probability is 1, the pixel is shown in its entirety. The right map represents the difference between the two probability maps. If the pixel is white, there is no significant difference between the two groups. If the video frame is shown in the pixel, then adults fixated more on this pixel than children did. The red circle denotes significant differences after correction for multiple comparisons. **(B)** Deep-learning classification of efficient and inefficient "looking videos" of individual participants. In the adult group, classification was significant in all participants whereas in the child group, classification was significant only in 6 participants. Participants are ordered by accuracy. Black lines indicate range of shuffle accuracy for each participant. Each dot represents accuracy based on the real labels.

Fig. 4. Pupillary response. Children's (left panel) and adults' (right panel) pupillary responses to each display after movement onset. The change in children's pupils was not significantly different during observation of videos of efficient versus inefficient actions. Adults' pupil size was larger when observing inefficient planning compared to efficient planning about a second after the initial grip. Gray area denotes significant differences (FDR corrected for multiple comparisons).

Fig. 5. EEG results. (A) Event-Related Spectral Perturbation map representing average changes in oscillation power across sensorimotor electrodes of children (left) and adults (right), collapsed across localizer videos. (B) Classification of action-related neural responses. For each participant, we used a SVM classifier to differentiate neural responses evoked by the two types of displays. We used the channels, frequencies and times of suppression that were identified in the localizer stage. (C) Classification analysis between efficient and inefficient action videos revealed that in the adult group, every participant showed significantly greater than chance differences in neural activity (right panel). In contrast, classification levels in all children were not significantly different from chance (left panel). Participants are ordered by accuracy. Black lines indicate range of shuffle accuracy for each participant. Each dot represents accuracy based on the real labels.

Supplemental Figure Captions

Fig. S1. Examples of action displays used for EEG localization. **(A)** One action started with an underhand/radial grip and ended with an overhand/radial grip with a rotatory motion in which the wrist turns out away from the body and the fingers point downward and in toward the body. **(B)** Another action started with an initial overhand/ulnar grip and ended with an underhand/ulnar grip.

Fig. S2. Looking videos were created from the original displays based on participants' looking patterns in each frame. These videos were used as input in the deep learning analysis.

Tables

	Fixation time		Dwell time		Revisits	
Age	$F(1,440) = 26.31$	$p < .00^*$	$F(1,440) = 21.63$	$p < .00^*$	$F(1,440) = 74.57$	$p < .00^*$
Action	$F(1,440) = .61$	$p = .43$	$F(1,440) = .20$	$p = .65$	$F(1,440) = .72$	$p = .39$
AOI	$F(3,440) = 103.81$	$p < .00^*$	$F(3,440) = 86.92$	$p < .00^*$	$F(3,440) = 136.95$	$p < .00^*$
Age $\times$ action	$F(1,440) = .00$	$p = .98$	$F(1,440) = .00$	$p = .97$	$F(1,440) = .00$	$p = .99$
Age $\times$ AOI	$F(3,440) = 9.70$	$p < .00^*$	$F(3,440) = 8.24$	$p < .00^*$	$F(3,440) = 30.77$	$p < .00^*$
AOI $\times$ action	$F(3,440) = .12$	$p = .94$	$F(3,440) = .10$	$p = .98$	$F(3,440) = .44$	$p = .72$
Age $\times$ action $\times$ AOI	$F(3,440) = .52$	$p = .66$	$F(3,440) = .39$	$p = .81$	$F(3,440) = .13$	$p = .94$

Table 1. AOI analysis. 2 (ages) $\times$ 2 (actions) $\times$ 5 (AOIs) ANOVA results for looking measures (fixation time, percent dwell time, and number of revisits). Asterisks denote significant differences.

Age	$F(1,440) = 19.68$	$p < .00^*$
AOI pairs	$F(9,440) = 133.37$	$p < .00^*$
Age $\times$ AOI pairs	$F(9,440) = 5.62$	$p < .00^*$

Table 2. Gaze-shifts. 2 (ages) $\times$ 10 (pairs of AOIs) ANOVA results for gaze shifts between AOIs.

Children			Adults		
<i># channels</i>	<i># clusters</i>	<i>Frequency range</i>	<i># channels</i>	<i># clusters</i>	<i>Frequency range</i>
1	2	6.40 - 8.19	3	4.66	5.85 – 25.00
1	3	6.03 - 13.75	4	1.75	5.85 - 12.78
4	2	5.85 - 13.7	1	2	8.8 - 17.23
1	7	6.68 - 19.51	4	2.5	7.03 - 22.24
1	3	6.85 - 19.65	5	5.4	6.03 - 23.07
1	1	5.85 - 7.03	10	3.9	6.03 – 25.00
2	8.5	5.85 - 18.95	4	2.75	6.07 – 25.00
7	7.71	5.85 - 25.00	2	7.5	6.53 - 24.7
2	7	5.85 - 24.81	1	1	11.45 - 13.16
1	1	6.44 - 7.61	1	2	12.32 – 25.00
1	1	6.97 - 7.78	1	3	6.85 - 18.13
1	1	7.78 - 9.34	2	5	6.85 - 21.6
2	7.5	5.85 - 25.00	1	3	5.85 – 25.00
1	18	5.85 - 25.00	2	4	5.98 - 15.9
1	5	7.50 - 21.29	5	3.2	5.85 - 24.6
1	4	6.03 - 11.79	1	11	6.25 - 25.00
3	2.33	5.85 - 22.08	8	3.12	6.85 - 12.97
2	2.5	6.03 - 13.06	7	4.85	5.85 - 20.53
1	4	7.85 - 20.68	1	7	13.16 - 23.07
1	6	5.85 - 21.92	4	8.5	6.85 - 25.00
2	4	6.63 - 18.54	1	5	6.85 - 25.00
1	4	6.44 - 21.14	1	4	12.41 - 23.58

Table 3. Localizer results. Cluster details for each participant in each group. *# channels* denotes the number of motor channels with at least one significant cluster. *# of clusters* denotes the amount of clusters per channel. *Frequency range* is the minimal and maximal frequencies for each participant's clusters.