

Trinity University

Digital Commons @ Trinity

Psychology Faculty Research

Psychology Department

2016

Early Verb Learning: How Do Children Learn How to Compare Events?

Jane B. Childers

Trinity University, jchilder@trinity.edu

Rebecca Parrish

Trinity University, rparrish@trinity.edu

Christina V. Olson

Trinity University

Clare Burch

Trinity University, cburch@trinity.edu

Gavin Fung

Trinity University, gfung@trinity.edu

See next page for additional authors

Follow this and additional works at: https://digitalcommons.trinity.edu/psych_faculty



Part of the [Psychology Commons](#)

Publication Details

Journal of Cognition and Development

Repository Citation

Childers, J. B., Parrish, R., Olson, C. V., Burch, C., Fung, G., & McIntyre, K. P. (2016). Early verb learning: How do children learn how to compare events? *Journal of Cognition and Development*, 17(1), 41-66. doi: 10.1080/15248372.2015.1042580

This Post-Print is brought to you for free and open access by the Psychology Department at Digital Commons @ Trinity. It has been accepted for inclusion in Psychology Faculty Research by an authorized administrator of Digital Commons @ Trinity. For more information, please contact jcostanz@trinity.edu.

Authors

Jane B. Childers, Rebecca Parrish, Christina V. Olson, Clare Burch, Gavin Fung, and Kevin McIntyre

Early Verb Learning: How Do Children Learn How to Compare Events?

Jane B. Childers¹, Rebecca Parrish¹, Christina V. Olson¹, Clare Burch¹, Gavin Fung¹,
Kevin McIntyre¹

¹Trinity University

Correspondence concerning this article may be addressed to Jane B. Childers, Trinity University, Department of Psychology, One Trinity Place, San Antonio, TX 78212
Email: jchilder@trinity.edu

Jane Childers, Rebecca Parrish, Christina Velasquez Olson, Clare Burch, Gavin Fung and Kevin McIntyre, Department of Psychology, Trinity University. Funding for this research was provided by a grant from the National Institute of Health (2R15 HD044447), Howard Hughes Medical Institute summer research program, McNair Program and the support of Trinity University. The project was supported by the Eunice Kennedy Shriver National Institute of Child Health and Human Development, and the content is solely the responsibility of the authors and does not necessarily represent the views of NICHD or NIH.

We wish to thank Tyler Howard, Angel Bottera, Alexandra Nancarrow, Rebecca Zelaya, Kristin Benavides Malone, Thania Galvan, Elizabeth Stephens, Melissa Flores, Sheena Yazdandoost, Katy Poettcker, Clarissa Terrell, Emily Roberson & Olivia Hunt for their assistance in data collection and coding. We thank the parents and children who participated in the studies. We also thank the teachers, parents and children at Laurel Heights United Methodist Church Child Development Center, First Presbyterian Church Child Development Center and the University Presbyterian Church Children's Center in San Antonio, Texas, and we thank the Witte Museum for allowing us to conduct some of our study sessions in the Science Treehouse.

Abstract

An important problem verb learners must solve is how to extend verbs. Children could use cross-situational information to guide their extensions, however comparing events is difficult. Two studies test whether children benefit from initially seeing a pair of similar events ('progressive alignment') while learning new verbs, and whether this influence changes with age. In Study 1, 2 ½- and 3 ½-year-old children participated in an interactive task. Children who saw a pair of similar events and then varied events were

able to extend verbs at test, differing from a control group; children who saw two pairs of varied events did not differ from the control group. In Study 2, events were presented on a monitor. Following the initial pair of events that varied by condition, a Tobii x120 eye tracker recorded 2 ½-, 3 ½- and 4 ½-year-olds' fixations to specific elements of events (AOIs) during the second pair of events, which were the same across conditions. After seeing the pair of events that were highly similar, 2 ½-year-olds showed significantly longer fixation durations to agents and to affected objects as compared to the all varied condition. At test, 3 ½-year-olds were able to extend the verb, but only in the progressive alignment condition. These results are important because they show children's visual attention to relevant elements in dynamic events is influenced by their prior comparison experience, and they show that young children benefit from seeing similar events as they learn to compare events to each other.

KEYWORDS: Language acquisition; word learning; eye-movements; verbs

An important problem for verb learning researchers is to explain how speakers of a language go beyond an initial learning context and extend verbs in new ways. Children appear to solve this problem by the time they are 4 or 5 years old, but researchers are only beginning to understand how this is accomplished. Recent evidence has shown children can use the range of sentences (e.g., Scott & Fisher, 2009), or range of events (Childers, 2011; Childers & Paik, 2009), linked to a new verb to extend verbs to new events. However, how children learning verbs use the information available across

situations is still poorly understood. The focus of the present studies is to investigate how children learn how to extract information across a range of events.

These studies are framed by a specific theory of cross-situational learning, structural alignment (e.g., e.g., Gentner, 1983; 1988; Gentner & Markman, 1997), which predicts that children benefit from ‘progressive alignment’ or access to examples that are highly similar early in a learning phase. Previous studies have shown that 3- to 5-year-old children learning a new part term (Gentner, Loewenstein & Hung, 2007), and 4-year-olds learning verbs (Haryu, Imai & Okada, 2011), show enhanced learning when test items are high in similarity to a target item. The present studies extend these findings by asking whether high similarity helps during the learning phase, by including multiple methodologies, and by testing younger ages. Study 1 investigates this question using an interactive task and, in Study 2, a video pointing task is used to ask whether children’s looking to events differs after they experience a similar pair of events as compared to a more varied pair. To preview our results, both studies show a benefit for children who saw a similar pair of events initially while learning a new verb. In addition, the looking data begins to address *why* seeing a similar pair of events may help children by showing increased looking to relevant elements in events following this experience.

The extension of new words is particularly difficult in verb learning because there is no single strategy that will be effective across verbs. Early noun uses could be governed by one or more simple strategies including (perhaps) a ‘shape bias’ (e.g., Landau, Smith & Jones, 1988; Landau, Smith & Jones, 1998). For example, a child (or adult) learning the

English noun ‘cat’ could use this type of bias to extend the word accurately to other similarly shaped cats, and could also accurately predict that the word ‘cat’ should not be extended to differently shaped horses. No similar single strategy will be as useful in verb learning because verbs differ from each other in the ways in which they refer to different aspects of events, meaning that children learning verbs have to deduce which package of elements in an event fit a particular verb’s meaning or solve a “packaging problem” (e.g., Gleitman & Gleitman, 1992). For example, a child learning the verb ‘open’ (in English) in the context of opening a door can extend the verb to opening a toy box or opening a present, but should not extend it to opening a watermelon or opening a light (switch). (Adults face a similar problem when learning new verbs like ‘facebook me’ or ‘Did she tweet it?’). Additionally, the verb category itself differs across languages such that languages vary in the types of patterns they exhibit (e.g., in the proportion of verbs that refer to path or manner aspects of events, Talmy, 1975). Thus, predictions speakers may make in English when extending verbs may not be useful for extending verbs in other languages. For these reasons, the problem of how to extend new verbs is a complex one.

Prior Evidence Of Children’s Verb Extensions

One way children could deal with this problem (at least for a time) is to avoid it. Indeed, early in development (and perhaps for newly learned verbs), children do seem to restrict their uses of verbs (Tomasello, 1992; 2000; Huttenlocher, Smiley & Charney, 1983; Roberts, 1983). However, at least two studies show that extending a newly learned verb to include a different agent is the earliest type of verb extension that children make. In one study 34 month-old children seeing colored drawings of Sesame Street characters

performing familiar (clapping) and novel actions (doing a split) were able to extend verbs to different characters enacting the action (Golinkoff, Jacquet, Hirsh-Pasek & Nandakumar, 1996). In a preferential looking task, 20- and 26-month-old toddlers extended new verbs to changes in agent, but did not extend verbs when other changes are tested (Forbes & Poulin-Dubois, 1997). Although these studies show 2-year-old children can extend verbs to include new agents, individual verbs can refer to a range of agents (e.g., “eat” can include both people and pets), instruments (e.g., one can eat with a spoon, chop sticks, one’s fingers), and affected objects (e.g., soup, dog food, pizza) for example. Thus, there is much left to explain.

When researchers have tested children’s ability to extend verbs to scenes with differing elements other than agents, evidence suggests extensions are difficult. For example, Forbes and Farrar (1993) showed that three-year-olds are more conservative in their extension of newly learned verbs than are older children and adults, particularly in terms of extending verbs to events with different instruments (but see Behrend, 1990 for evidence that they most readily extended verbs to new instruments). And, a study of Japanese-speaking children (Imai, Haryu & Okada, 2005) found children did not extend new verbs until 5 years. Kersten & Smith (2002) also showed English-speaking 3-year-olds have difficulty extending new verbs when the objects in the new events are novel (e.g., bug-like creatures), but do not have a similar difficulty extending nouns.

In sum, there is mounting evidence to suggest that young 2-year-olds have difficulty extending new verbs (e.g., Huttenlocher et al. , 1983), and that the ability to extend verbs

to other objects may develop gradually between 2 and 4 years (Forbes & Farrar, 1993; Imai et al., 2005; Kersten & Smith, 2002). Thus, there must be one or more mechanisms that underlie children’s verb extensions, perhaps emerging over development. Given the nature of the problem, it seems likely that mechanisms for verb learning will support active learning—that is, children appear to need to construct individual verb meanings, and intuitions about how to use those words in new situations and sentences, from the environment.

How Do Children Extend New Verbs?: Theories & Evidence Of Cross-Situational Learning

One strategy children could use is to compare information across multiple contexts when learning verbs (e.g., see Fisher, Hall, Rakowitz & Gleitman, 1994; Pinker, 1989), and recently verb researchers have begun to consider the central role of cross-situational information in verb learning (e.g., Scott & Fisher, 2012). Cross-situational information could be especially helpful because the range of instances linked to a verb in the past provides important clues to how the verb should be used in the future. This usefulness holds whether the “instances” include sentences in which the verb is heard, events with which the verb has been linked, or both.

At present, there are three major theories that address how learners could glean information across a range of examples: learners may use associationist processes (e.g., Smith & Yu, 2008; Yu & Smith, 2007), they may form a hypothesis that they later test as they see multiple examples (e.g., Trueswell, Medina, Hafri & Gleitman, 2013), or they

may structurally align examples from one to the other and draw conclusions from alignments (e.g., Gentner, 1983; 1988). Therefore, not only are verb researchers largely in agreement that cross-situational learning is central in verb acquisition, but they are beginning to build and test multiple models to explain how it is accomplished. Yet, to date, there are very few studies that have tested any of these theories as they apply specifically to verb learning. In addition, even if children are using these mechanisms to deal with the cross-situational information they must process, even fewer studies have asked how they *learn how to compare* multiple events (or sentences). Yet, instead of arguing about which mechanism truly underlies children's verb learning, in the following paragraphs we would like to propose a way that these mechanisms may work together.

In the associationist account, learners who hear two words while seeing two objects will not be sure which word/object pairing holds until they hear the same new word a second time. Once they have seen a second pair of objects while hearing the same new word, they can track the probability of the word/object pairing across examples and learn the word (similar to statistical learning in speech segmentation, Saffran, Aslin & Newport, 1996, and artificial grammar learning, Gomez & Gerken, 1999). Smith and Yu (2008) showed that 12- and 14-month-old infants could learn six object-word pairings over thirty training trials. A question that arises then is can sheer associations of events with particular verbs explain verb learning? In a study of verb learning (Scott & Fisher, 2012), children saw two actions and heard two novel verbs. In that study, 2 ½-year-old children saw three pairs of events in six presentations, and during each pair they heard two different novel verbs. Their looking behavior to each event was coded frame-by-

frame as the sentence with the two verbs unfolded. Starting in the second trial, or once children had heard a specific verb twice, children looked longer at a specific event that co-occurred with that verb as opposed to a distractor event. This is a complex task for these young children, and thus their performance is impressive. However, a second study shows that this level of accuracy is present only for events that are simple body movements; if events are shown with an agent and an object, only high vocabulary children at this age succeed. In addition, these events were whole events, prepackaged for the child. In every day contexts, children see dynamic sets of events that must be parsed, and specific verbs are used to refer to different sets of elements within these events.

If children are associating an entire unconstrained memory of a dynamic event, *how* do they compare it to the new event before them that also co-occurs with the same verb? How do children or adults converge on relevant and irrelevant elements of actions across scenes? If children form a collection of scene to verb links, the other two mechanisms that have been proposed may explain how they get beyond whole scenes and focus on relevant objects and relations.

A second main mechanism proposed for cross-situational learning is the hypothesis testing approach (Medina, Snedeker, Trueswell & Gleitman, 2011). In this view, learners propose a hypothesis the first time they hear a new word and see an object (or event). When they hear additional uses of that word, they test that hypothesis (verify) to determine whether it holds in the new context and, if not, they revise it, abandoning

hypotheses that fail. In a recent paper, two studies with adults show that, for adults learning new nouns, the order of the examples influences their learning (Medina et al., 2011). This is used as evidence for the hypothesis testing view since order should not be as influential if learners are computing associations. What type of hypotheses are formed, and how do these differ from nouns and verbs? In this paper, adults and children who were shown a video with audio that had omitted words seemed to thought to be guessing which familiar word applied in which context. In a subsequent paper, three studies show adults learning new nouns performed in ways consistent with the hypothesis testing account, using both explicit and looking measures (Trueswell, Medina, Hafri & Gleitman, 2013). This paper also does not contain an explicit example focusing on verb learning. However, this mechanism could interact with associations if children formed an initial link between an event and a new verb, and then created a hypothesis about which specific elements or actions in the event could be especially important.

In the third structural alignment view, the observer compares two events by aligning specific elements of one event to elements of another based on the relational structure of each event (e.g., Gentner, 1983; 1988; Gentner & Markman, 1997). For example, children seeing a scene in which a soccer player kicks a ball (Event 1) and then a scene in which a football player punts a ball (Event 2) could recognize that there is a kicker in Event 1 and align it with the kicker in Event 2, and recognize there is something kicked in Event 1 and align it with the kicked object in Event 2. These alignments are initially guided by the perceptual similarity of objects across the two events, and there can be different numbers of elements in each event (nonalignable differences, see

Gentner & Markman, 1997). A benefit of this view is that it specifically describes the mental processes observers would use to compare elements and parts of events across different examples. Again, this structural alignment mechanism could be initiated after an initial associations of a set of dynamic scenes with a particular verb, and it would then be a way children could compare these multiple instances to each other.

There is prior evidence that children as young as 2 ½ years can compare events when learning new verbs, perhaps by aligning the events. For example, in one study, two ½-year-olds appeared to extract the common element across a set of three events and use that information to direct their enactments, preserving either the action or the result when given a chance to extend the verb (Childers, 2011). In a related study (Childers & Paik, 2009), both English-speaking and Korean-speaking children who saw three varied events used varied objects to perform a new action at test whereas children seeing three events with the same objects did not include the varied objects as often at test. A third set of studies showed 2 ½-year-olds learning verbs perform as well at test after seeing a set of events that can be compared as they do after receiving direct instruction about a verb from an experimenter (Childers, Heard, Ring, Pai & Sallquist, 2012). In addition, 3 ½-year-olds who heard contrastive sentences with a new verb (e.g., “Look! I’m meeking it. Look! I’m not meeking it.”) are unable to extend the verb unless these statements are applied to more than one set of events and compared (Childers, Hirshkowitz & Benavides, 2014). These studies suggest children benefit from the opportunity to compare multiple events when learning a verb and can use information gleaned across a set of events to guide verb extensions.

The present study uses a specific prediction from structural alignment theory: children develop the ability to align objects and relations from experiences with high similarity examples, or ‘progressive alignment’. Children asked to initially compare two similar instances improve in their ability to solve an analogical reasoning task (Kotovsky & Gentner, 1996) and a model room task (Loewenstein & Gentner, 2001). There is also evidence that children shown objects at test that are more similar to a target object are better able to learn a part term than are children shown varied objects at test (Gentner et al., 2007). In addition, one study shows that children learning verbs are better at extending a new verb to a relevant event with a similar object than to an event with a dissimilar object (Haryu, Imai and Okada, 2011). In in this paper, 3- and 4-year-olds were shown a single novel event in which an agent performed a repeated action with a novel object. After seeing this event once, they were shown a pair of events, in one the agent used a new object to perform the same action (Action Same) and in the other, the agent performed a different action using the same object (Object Same). To extend the verb, children should chose the Action Same event at test, and 4-year-olds did so (but not 3-year-olds), but only when this event included a similar object. These results suggest that the objects in events have important influences on verb extensions, however comparisons were only between initial and test events.

In a second study, 4 ½-year-olds were shown a target event, then a pair of test events across multiple trials. What was varied across children was whether the first four blocks of trials depicted similar objects in the AS test event, or differently shaped objects. The

question asked then in this study is whether repeated target-test experience with high object similarity helps children in later trials with lower object similarity. Results showed 4 ½-year-olds did benefit from high object similarity, and were more successful than the other group when they were given lower object similarity events in the second set of trials. Thus, practice extending verbs to events with high similarity can benefit later verb learning. The present study extends this work by including younger children, by including multiple examples during a learning phase that can be compared, and by including multiple measures, to see whether we can also show that children learn to compare through experience with high similarity examples.

Our Studies

In these studies, we tested whether 2 ½- and 3 ½-year-old children benefit from seeing two highly similar event pairs before two varied pairs as compared to children who saw only varied pairs or to children who only saw a single event before test. In Study 1, we presented 2 ½- and 3 ½-year-old children with ‘live’ events in which an experimenter enacted events using objects, and then children enacted the events using a new set of objects. We predicted children seeing multiple events would produce more extensions than would children seeing a single event, based on prior research showing that multiple event experience is useful in verb learning (e.g., Childers, 2011; Childers & Paik, 2009). Study 2 uses a video procedure and a Tobii x120 eye tracker to better understand observers’ visual attention to elements in events they see during the learning phase of the study, and to seek converging results across different methodologies.

STUDY 1

Method

Participants

Thirty-three 2½ year-old children ($M = 2;7$, range: 2;3– 2;11) and thirty-six 3 ½ -year-old children ($M = 3;7$; range: 3;4– 3;10) participated in this study, 36 girls and 33 boys. Most children were from middle income or upper middle income homes. Of the families who provided ethnicity information, 30 reported their ethnicity as Caucasian, 24 self-identified as Hispanic, and 11 were members of two or more ethnic groups. Children were included only if their parents reported exposure to English at least 80% of the time, and were excluded if teachers reported a speech delay. Parents who brought their children to the laboratory completed the verb vocabulary section (103 verbs) from the MacArthur-Bates Communicative Development Inventory for Words and Sentences (Fenson et al., 1993). The younger children's verb vocabulary was $M = 79$ verbs (range: 36-103, $n = 19$ reporting) and in the older group, children's verb vocabulary was $M = 98$ verbs (range: 75-103 words, $n = 20$ reporting). Additional children participated but were excluded from the final sample because there was an experimenter error (2), the child was extremely distracted (2), or the child failed to complete the study (3).

Design

There were two between subjects factors in this study: Age group (2 ½ or 3 ½ years) and Condition (3: PA, All Far and Control). Within each age group, there were 11-13 participants who were randomly assigned to each condition. In the Progressive Alignment condition, children saw two highly similar events then two varied ones before

test, while in the All Far condition, children saw four varied events before test. Similar events were considered to be similar to each other because they included the same number of objects, with same shapes, engaged in the action in the same way. Events were categorized as varied because the two events in a pair included differently shaped objects, different numbers of objects, and different movements that were used to accomplish the same result.

In the construction of the trials, every child saw the same pair of varied events immediately before the test trials. Thus, the only difference between the two experimental conditions was in the initial event pairs that were shown. In a control group, children saw a single repeated event before test. To ensure that the results from the control group did not rely on a single event presented before test, half of the control group in each age group was shown one event, and the other half saw a different repeated event before test.

To confirm that the initial pair of events in the PA condition was more similar than the pair of events shown in the All Far condition, we asked a sample of adults ($N = 21$; aged 18- 22 years) to rate pairs of events in both conditions on a 7 point Likert scale, with 1 = very dissimilar and 7 = very similar. Participants randomly chose whether to rate the Similar First or All Far sets first. A repeated measures ANOVA with video (PA, AF) and learning pair (1st pair, 2nd pair) as within-subjects factors showed a main effect of video, $F(1, 20) = 11.18, p < .003$, a main effect of learning pair, $F(1, 20) = 147.79, p < .001$, and a Video x Learning Pair interaction, $F(1, 20) = 51.60, p < .001$. The events did differ in

rated similarity in the first learning phase, with pairs of events in the PA condition rated as more similar to each other ($M = 5.88$, $SE = .24$) than in the All Far condition ($M = 4.24$, $SE = .29$), $t(20) = -6.47$, $p < .001$; as designed, the events in the second learning phase were not rated differently across conditions (PA, $M = 3.38$, $SE = .23$; AF, $M = 3.60$, $SE = .25$).

Materials

Two novel verbs were chosen for the study and each verb was paired with a set of comparison events and two test sets. It may be more difficult to learn verbs linked to these types of events than it is for them to learn verbs linked to events that do not cause a change of state (Huttenlocher et al., 1983; Scott & Fisher, 2012). However, they conform to Slobin's (1981; 1985) prototypical (causative) event in which an animate agent causes a salient change in a patient, and they are similar to events used in other studies of 2-year-olds' verb learning (e.g., Tomasello & Barton, 1984). One verb, 'tam' corresponded to a 'squishing' event, while the other, 'gorp' corresponded to a hiding event (see Appendix A).

Two test trials were created for each novel verb. To explore the range of extensions children could produce, one test set included objects that were similar in shape to objects the experimenter acted on, and one set included objects that were more varied.

For example, in 'tam', in the PA condition, children saw the experimenter squish a blue sponge using a vertically held roller, and then squish a pink and orange ball using a

vertical green roller. These two structurally similar events were followed by an event in which the experimenter squished a blue and yellow sponge against a suspended yellow board using a pink box, and used his/her fingers to squish a white sponge against a green heart. In the All Far condition in this event set, the two initial events were squishing a blue ball using a white roller and squishing an orange sponge into a transparent measuring cup using a black potato masher. The second two events were the same as in the PA condition. In the Control condition, half of the children saw the experimenter squish a blue and yellow sponge against a suspended yellow board using a pink box, which was repeated before test, and the other half saw the experimenter use his/her fingers to squish a white sponge against a green heart (see Appendix A).

Test objects for ‘tam’ were 1) a sponge, red/yellow roller, and black spool and 2) a plush turtle, plastic blue baby bottle, and a hard plastic carrot. As the second test set was designed to “push” children to extend the verb to include even more varied objects, the similarities between learning and test in object shape were more distant.

Procedure

Upon arrival at an on campus laboratory (or after leaving their classroom at the preschool and going to a separate, quiet room), the child sat on the floor with an experimenter who played with unrelated toys to build rapport. A second experimenter introduced the consent form to the parent in the lab (or, in the preschools, the consent forms had been returned to the teacher). Once the child had become accustomed to the situation, the child was asked to sit in a small chair at a child-sized table across from an experimenter

(in the preschools, experimenters and children usually sat on the floor). The second experimenter recorded the children's responses using a video camera mounted on a tripod in the corner of the room, and also coded children's responses during the session using a score sheet. The experimenter used a script to be sure the presentation of the stimulus sentences was uniform across subjects. In addition, each set of objects for a given event set was stored in its own opaque box. The child randomly chose the order of the two verb sets by choosing a box; the box contained a notecard with information about the stimuli for each set for the experimenter, again to insure consistency across participants.

The experimental procedure began with an initial play phase to give participants a chance to explore the test objects prior to the learning or test phase. Children were given the first set of test objects and asked to act on them ("Look at these things. What can you do with these?"). This process was repeated for this second set of test objects. Children's responses were observed to ensure that they did not spontaneously enact the extension actions.

The learning phase followed this initial play phase. The experimenter pulled out objects for two events at the same time and placed the two sets of objects that she would act on side-by-side on the table (or floor). She then enacted the first event once while producing the stimulus sentences ("Look! I'm going to it." (before the action), "I'm ing it." (during the action), "I ed it." (after the action ended)). After putting those objects back on the table/floor, she turned to the second set of objects and repeated the same sentences while performing that action. The experimenter then put those two sets away

in the box, and took out the second two sets of comparison objects, repeating the same enactment process. After completing the two pairs of events, and putting all of the objects back in the box, the experimenter said “Can you see why they’re all ing? Can you say ?”. Prior research suggests that the first sentence of this pair is an important invitation to compare before test (Gentner, 2002). Children in the control condition heard the same sentences as in the experimental conditions, but only saw a single event repeated four times before the test trials.

The experimenter next took out the first set of test objects and placed them on the table/floor in a random order in front of the child while saying (“It’s your turn to play. Can you it? Can you play the game?”). This set of prompts was repeated once if necessary. Once the child acted, the experimenter asked the child to produce the verb (“What are you doing?”). This process was repeated for a second set of test objects while the experimenter said “You get one more turn. Can you it again? How else can you it?”). The learning and test phase formed a single block of trials. The entire process was repeated in a second block of trials presented for a second novel verb.

Coding

For each test trial, we coded whether a participant reproduced the action that had been seen in the learning phase using new objects. If they did, this was coded as an Extension. If instead children performed an off task response (e.g., feeding the carrot to the bunny), that was coded as Other. We allowed children to produce both an extension response and an off task response on each test trial, but did not code the number of times children

repeated a particular type of response. Even so, some children only produced one type of response (only Extension or only Other), and some produced both Extension and Other responses. Thus, we created proportions with the number of extensions/ (extensions + other); these were computed across both verbs.

Experimental sessions were initially coded by a live observer and recorded using a video camera. All sessions were later coded by a second observer from video files. This second coding was the coding used in the final data set unless the recording of the session could not be coded ($n=2$), because the second coder could pause the videotape and review enactments if needed. Interrater reliability computed across these two coders was found to be 88% with a Cohen's kappa = .76 (substantial agreement, Landis & Koch, 1977).

RESULTS

Children learned two verbs and their performance across these two verbs was averaged to ensure that the results are not restricted to a single event or verb. As part of the design of the study, there were two dependent variables: children's response on the first test trial which contained objects similar in shape to the learning phase (close extensions), and children's response in the second test trial, which contained more varied objects (far extensions).

To examine the number of close extensions, we conducted a 2 (Age: 2 ½, 3 ½ years) x 3 (Condition: PA, All Far, Control) factorial ANOVA.¹ The results revealed a

near significant main effect of age, $F(1, 63) = 3.91, p = .053, \eta^2 = .06$, with 2 ½ year olds producing fewer close extension responses than 3 ½ year olds (2 ½: M_{prop} extensions = .55 ($SE = .06$), 3 ½: M_{prop} extensions = .70 ($SE = .05$)), but no main effect of Condition or Age x Condition interaction.

Next, to examine the number of far extensions, we conducted a 2 (Age: 2 ½, 3 ½ years) x 3 (Condition: PA, All Far, Control) factorial ANOVA, which revealed a significant main effect of Age group, $F(1, 63) = 5.88, p = .018, \eta^2 = .09$, with 2 ½-year-olds producing fewer far extensions ($M_{prop} = .33, SE = .06$) than 3 ½-year-olds ($M_{prop} = .52, SE = .06$). As predicted, there was also a significant main effect of Condition, $F(2, 63) = 3.46, p = .037$. To interpret this effect, we conducted Tukey post-hoc analyses, which revealed that children in the PA condition ($M_{prop} = .54, SE = .07$) performed significantly more verb extensions than did children in the Control condition ($M_{prop} = .28, SE = .07$), $p = .037$, while the All Far ($M_{prop} = .46, SE = .07$) and Control condition did not differ, $p = .15$ (see Figure 1). Finally, this analysis revealed no significant interaction between Age group and Condition.

DISCUSSION

Two main findings emerged in this study. One was that there was evidence for developmental change in children's ability to extend new verbs between 2 ½- and 3 ½-years. This difference was suggested in the easier verb extension test trial, but was most clearly seen when children were asked to extend a new verb to events that included objects that varied in shape and category from those in the learning set. This evidence for

developmental change suggests that children are progressively overcoming any initial conservatism that is present early in verb learning, either as they develop in their cognitive skills with age, or with additional experience learning new verbs. It also fits well with other research that has shown that young children can have difficulty extending newly learned verbs (Forbes & Farrar, 1993; Imai et al., 2005; Kersten & Smith, 2002).

The second main finding, which fits predictions made by structural alignment theory and supports other empirical findings (e.g., Loewenstein & Gentner, 2001), was that children who initially saw a pair of similar events produced more verb extensions than did children seeing only a single repeated event. Interestingly, children seeing pairs of varied events performed almost as well as the PA group, but did not differ from the control group. In addition, the apparent advantage conferred by the similar first experience only helped boost verb extensions when the most varied extensions were considered. Put another way, access to more than one event in the learning phase produced the largest benefit, and there was a smaller difference in performance when comparing the two types of variation we presented.

Perhaps the benefit of the PA condition would have been clearer if children in the control condition had performed more poorly. When we designed the study, our view was that some variation during a learning phase provides information to the child about the range of events to which a new verb might extend, and thus the absence of variation should be problematic. However, there is evidence that comparing even across two presentations of the same stimulus can be useful. For example in the model room study

(Loewenstein & Gentner, 2001), children who saw two identical model rooms were better able to search in a large test room than were children without this experience. Thus, it may be helpful to see a repeated example, even when it does not add additional information. Secondly, in the Haryu et al. (2011) study, children who saw a test trial with more similar objects were better able to extend the verb, even though all children only saw a single event. Because the first test trial in the present study always included more similar objects, it may have optimized children's ability to extend in all conditions (a prediction that could be tested in future studies). Given these concerns, it is important that a significant difference still emerged between the PA and Control conditions.

This study and the results reported in Haryu et al. (2011) are the only results thus far (to our knowledge) that show that children benefit from similarity when extending new verbs. These converging results are important because they build a case for a specific types of experience as important as children learn how to learn from cross-situational examples. However, evidence is more convincing if similar results can be found across studies using different methodologies. Thus, in Study 2, we extended Study 1 using a video event procedure instead of the 'live' interaction procedure, and added a Tobii x120 eye tracker to track children's eye movements during the comparison phase of the study. Although many studies of verb learning have relied on looking time measures (e.g., Naigles, 1990), to our knowledge, only two previous published studies have examined specific eye movements produced by children in a verb learning task (though see Knoeferle, Crocker, Scheepers & Pickering, 2005, and Papafragou, Hulbert & Trueswell, 2008 for related eye tracking results from adults in verb tasks). In Waxman,

Lidz, Braun and Lavin (2009), 24-month-olds learned verbs while seeing comparison trials and a contrast trial. An eye tracker was used to show 24-month-old children could direct their looking to a target event with a new object vs. a distractor event as the stimulus sentence was heard. In addition, as described, Scott and Fisher (2012) shows children can direct their visual attention to a repeated verb-event pair while learning new verbs, particularly if the events depict body movements.

Note that in these two prior studies, looking to an entire event was measured. In the present study, we sought to use an eye tracker to examine looking to specific elements in events. In addition, the eye tracker we used (Tobii x120) includes software that allowed us to designate dynamic areas of interest (AOIs), and thus we were able to track children's looking to specific parts of an event as those elements were shown in motion. This provides a powerful new tool for verb researchers interested in how children attend to objects in events (and perhaps the relations between objects). Given the paucity of this type of data, and its importance for understanding mental processes that may underlie verb *learning*, the present study provides important new data to language researchers.

EXPERIMENT 2

Method

Participants

Twenty-four 2 ½-year-old children ($M= 2;7$; range: 2;2-2;11; 12 girls, 12 boys), twenty-four 3 ½-year-old children ($M= 3;6$; range: 3;0-3;11; 15 girls, 9 boys), and twenty-three 4 ½-year-old children ($M= 4;6$; range: 4;2-4;10; 9 girls, 14 boys) participated. Of families

who responded to ethnicity questions, 34 self-identified as Caucasian, 18 as Hispanic, 1 as Asian, and 8 as having two or more ethnicities. Parents who came to the on campus laboratory filled out the verb vocabulary section from the MacArthur-Bates CDI. A few of the participants participated at a local science museum. In the 2 ½- year-old age group, children's verb vocabulary was $M= 87$ verbs (range: 19-103, $n= 13$ reporting), and in the 3 ½-year-old age group it was $M= 96$ verbs (range: 59-103 words, $n = 14$ reporting) (the 4 ½-year-olds were at ceiling). Additional children participated but were excluded because they were exposed to English less than 80% of the time (2), there was an experimenter error (5), the eye tracker failed to capture the observer's eye movements (9), children refused to watch the videos or only watched and did not point during any test trials (21), or they failed to point during one or more test trials (3)².

Design

There were two between subjects factors in this study: Age group (3: 2 ½, 3 ½, 4 1/2) and Condition (2: PA, All Far). In each age group, there were 11-12 participants in each condition.³ Each participant produced two types of dependent variables (within-subjects): total fixation duration to specific Areas of Interest (AOIs) during the learning phase and pointing data at test. We describe each of these separately in a later section. The particular verb set shown first was counterbalanced within each condition and age group so that half of the participants learned one verb as their first verb and the other half learned that verb as their second verb.

Materials

Two sets of video events were created for this study. One of the events used in Study 1 was included in Study 2 ('tam'/an agent squishes an affected object) and the other event in Study 1 (an agent puts small objects into a container) was redesigned because the elements of the action often overlapped visually on the screen. Thus, this event was replaced with a new event (an agent made an imprint on an affected object). A second difference between Studies 1 and 2 was that in Study 2, events were seen on a video display, and thus events could be presented simultaneously on the left and right halves of the screen.

As in Study 1, in the PA condition, participants initially saw a pair of events with agents, instruments and affected objects that were the same size and shape whereas in the All Far condition, the first pair of events included events that varied in the number of elements, the shape of objects and the way in which the result was accomplished (see Appendix B). In the second learning trial shown before test, *in both conditions* the pair of events shown during the learning phase *was the same* and for both conditions, the pair included events that varied in multiple ways. All of the video events depicting the squishing event ("tam") were 8 seconds long, and all of the video events depicting the imprinting ("zim") were 12 seconds long. Preliminary analyses did not reveal differences between the events.

As in Study 1, two test trials were created, the first including objects that were similar to those seen in the learning phase, and the second test trial including more varied

objects. Two objects were chosen from the test set in Study 1 and events were filmed showing an adult using the two objects to enact the target action (e.g., in the squishing event, the agent used a black spool to squish a sponge) and an adult using these same objects to enact a distractor action (e.g., the agent used the spool to push a sponge forward). The second test trial contained more varied elements (e.g., the actor squished a plush turtle using a baby bottle vs. the actor fed the turtle using the bottle). All of the test events lasted 12 seconds.

Areas Of Interest

Areas of interest were drawn by hand using the tools available in the Tobii Studio software. The software allows users to designate ‘dynamic AOIs’ which uses keyframes corresponding to a particular point in the timeline. We defined multiple keyframes for each AOI frame by frame, and the software then allowed that AOI to move smoothly across multiple keyframes. AOIs were drawn only for scenes shown in the learning phase and frames in which AOIs began to overlap (typically towards the ends of events) were excluded. Three regions were identified: the agent (which included the hand, arm and torso of the female actor), the instrument or tool, and the affected object. AOIs traced the shape of each element with some allowance for the immediately surrounding region adjacent to each element (see Appendix B).

Experimental Set-Up

Participants were tested individually in a quiet room in an on campus laboratory, in a quiet office in a local science museum (the Witte Museum), or in a quiet room in their

child care center. Participants sat in front of a 21 inch flat screen video monitor. In order to minimize head movement and distractions, children sat in a car seat with stabilizing wings alongside the position of the head, which was attached to an adjustable office chair. A Tobii X120 eye tracker device was placed below the video screen, which was connected to a laptop. The distance between the table holding the monitor and tracker and the chair was approximately 16 inches, with some variation to maximize and individual participant's calibration. Eye movements were measured by the eye tracker using a corneal reflectance tracking technique. A near infrared light source was directed at the participant, undetectable to the naked eye, and the reflection of the light on the cornea was recorded as the participant watched the video stimulus on a monitor.⁴

As in Study 1, two experimenters were present. One interacted with the laptop that controlled the eye tracker and used a script to produce the stimulus sentences. The second experimenter recorded the children's responses using a webcam mounted on the top of the video monitor, and coded children's pointing responses using a score sheet. Parents sat behind the child and were asked to refrain from talking or assisting their child.

Procedure

As in Study 1, children who came to the lab (or who were met at the Witte museum) were greeted by two research assistants. One focused on developing rapport with the child using unrelated toys and the other explained the consent form and vocabulary checklist to the parent. In children's child care center, parents had returned a signed form to their child's teacher. An experimenter played with the child in his/her classroom before taking

them to the quiet room in the center. When the child seemed comfortable, he or she was asked to sit in the chair in front of the video monitor and the experimental session began.

The experimental session began with a calibration process. In preparation for the calibration, the experimenter adjusted the eye tracker and the height of the seat to make sure that the reflections of both eyes were centered in the eye-tracking camera's field of view. We used a procedure that measured 5 calibration points; the Tobii Studio infant calibration stimuli (a cartoon cat) was presented in each corner of the screen and in the center. The software displays a graphic of looking to each calibration area and we recalibrated one or more points if the results of the tracking did not cluster around a calibration point. We did not include a child's eye tracking data if his/her looking could not be calibrated. The calibration procedure took approximately 1 minute before practice trials were initiated.

The experimental session began with a warm-up phase to allow participants to practice pointing to the screen. In the first warm-up pair, observers saw a person using a wooden spoon to stir something in an opaque bucket and a person moving a stuffed bunny up and down while hearing "Look at these things. Can you point to the jumping bunny?". They then saw a video of a person rolling a fire truck and a person moving a stuffed tiger side-to-side while they heard, "Look at these things. Can you point to the moving truck?". In these trials, one correct answer was on the left side of the split screen image and one was on the right.

After the warm-up phase, the experimental phase began. Observers first heard “Now watch. We are going to play a game!” The video started with a still photo of both of the key events, and then one event was seen in motion while the experimenter produced a set of 3 stimulus sentences (“Look! She’s going to it.” (before the action), “She’s ing it.” (during the action), “She ed it.” (after the action ended)). Once that event ended, the second event on the other side of the screen was shown in motion and the same set of sentences was produced for this event. Next children had the opportunity to see both events in motion at the same time on the split screen, while they heard a sentence prompting them to compare the events to each other (i.e., “Can you see why these are both ing? Can you say ?”). This process was repeated for a second pair of events before the test trials.

As in Study 1, each participant responded to two test trials for each novel verb presented. First, during a black screen, the experimenter said, “Now it’s your turn to find ing.” Then a pair of actions came on the screen and the experimenter said “Point to ing. Can you point to the one who is ing? Good job!” This split screen video showed the target event enacted with new objects or a distractor event with the same new objects. This process was repeated for a second pair of events while the experimenter said “You get one more turn. Can you point to ing? Which one is ing?”). The learning and test phase formed a single block of trials. The entire process was repeated in a second block of trials presented for a second novel verb.

Coding

Research assistants coded whether a participant pointed to the target or distractor event. Responses were initially coded by a live observer; later, 24% of the children who participated were coded by a second observer from video files and interrater reliability across the coders was computed. Coders agreed on 94% of the trials with a Cohen's kappa = .88.

RESULTS

One key question addressed in this study is whether results at test, using a different method of measurement (pointing vs. enactment), would support the results found in Study 1. Many verb learning studies have used pointing to one of two videos as a measure, and thus this methodology is well tested (e.g., Maguire, Hirsh-Pasek, Golinkoff & Brandone, 2008). Thus, we start with report of our test trials.

Analyses Of Pointing Results At Test

These analyses investigate whether differences by condition emerged at test, and whether they were in the predicted direction, with children in the PA condition benefitting from that experience at test. Given the results from the “live” experiment in Study 1, which showed differences in the second test trial with more varied objects, we included Test Trial (close extension, far extension) as a within-subjects factor in the analysis. The dependent variable was the total number of correct points produced across the two verbs.

To examine the number of correct extensions, we conducted a 3 (Age: 2 ½, 3 ½, 4½ years) x 2 (Condition: PA, All Far) x 2 (Test Trial: close extension, far extension) mixed model ANOVA, with age and condition as between-subjects factors and test trial as a within-subjects factor. The results revealed a main effect of test trial, $F(1, 65) = 5.60, p = .021, \eta^2 = .08$, such that children were more successful on the far extension test trial ($M = 1.49, SE = .07$) than on the close extension test trial ($M = 1.24, SE = .09$). The results also revealed a significant effect of age group, $F(2, 65) = 5.81, p = .005, \eta^2 = .15$, which was qualified by a significant Age group x Condition interaction, $F(2, 65) = 4.01, p = .023, \eta^2 = .11$. Because test trial did not interact with age and condition, we collapsed the results across the two test trials to interpret this significant interaction and follow up simple main effects tests revealed that there was no difference between the PA ($M = 2.00, SE = .28$) and All Far ($M = 2.58, SE = .28$) conditions for 2 ½-year-olds, $F(1, 65) = 2.17, p = .14$, who responded at chance levels. Similarly, there was no difference between the PA ($M = 3.50, SE = .28$) and All Far ($M = 3.00, SE = .29$) conditions for 4 ½-year-olds, $F(1, 65) = 1.53, p = .22$, because they universally succeeded on the pointing task. However, 3 ½-year-olds performed significantly more correct verb extensions in the PA condition ($M = 3.08, SE = .28$) than the All Far condition ($M = 2.25, SE = .28$), $F(1, 65) = 4.43, p = .039$, and only responses in the PA condition differed from chance, $t(11) = 4.73, p = .001$ (see Figure 2).

Eye Tracking Results

Analyses By Areas Of Interest (Aois)

The scenes shown in the first pair of events differed across conditions, and thus looking could differ by condition for merely perceptual reasons. Thus, the focus of the analysis of observers' looking behavior was to investigate whether the total fixation duration to each area of interest (AOI) in the second pair of events differed by condition. If looking varies by condition in this trial, which depicted the *same pair of events*, this would suggest that their looking was influenced by their prior experience of seeing a similar event pair or a varied pair. Total fixation duration to an AOI (e.g., agent) was averaged across the two scenes in the second learning trial and across the two verb sets, yielding a total fixation duration value that was *not specific* to a particular scene or verb.

We predicted that what experience seeing a similar event pair could do is influence children's visual attention (and perhaps processing of) one or more key elements in events. This was predicted because within structural alignment theory, observers first align objects across examples and then their relations. Thus, the following analyses examine whether looking to a particular AOI (Agent, Tool, Affected object) is influenced by condition. We computed a univariate ANOVA for each AOI type with Age group (3: 2 ½, 3 ½, 4 ½) and Condition (2: PA, All Far) as between subjects factors and total fixation duration as the dependent variable.

The analysis examining children's total fixation to the Agent AOI revealed a main effect of Age group, $F(2, 65) = 4.84, p = .011, \eta^2 = .13$, with Tukey post hocs revealing

that 2 1/2-year-olds looked longer at the agent ($M = .34$, $SE = .05$) than 4 1/2-year-olds ($M = .18$, $SE = .03$), whereas 3-1/2 year-olds ($M = .25$, $SE = .04$) did not differ from either of the other age groups ($ps > .23$). This effect was qualified by a significant Age group x Condition interaction, $F(2, 65) = 3.55$, $p = .034$, $\eta^2 = .10$. Simple main effects analyses showed that, only in the 2 1/2-year-old group, total fixation duration was greater to the Agent in the PA ($M = .42$, $SE = .05$) than the All Far condition ($M = .26$, $SE = .05$), $F(1, 65) = 4.04$, $p < .05$ (see Figure 3).

The analysis examining the Affected Object AOI showed a significant main effect of Condition, $F(1, 65) = 4.06$, $p = .048$, $\eta^2 = .06$, with participants looking longer at the affected object in the PA condition ($M = .74$, $SE = .11$) than in the All Far condition ($M = .46$, $SE = .08$). This effect was qualified by a significant Age group x Condition interaction, $F(2, 65) = 3.46$, $p = .037$, $\eta^2 = .10$. Simple main effects analyses showed that, again only for 2 1/2-year-olds, looking to the Affected object was greater in the PA ($M = .86$, $SE = .16$) than the All Far condition ($M = .24$, $SE = .16$), $F(1, 65) = 7.02$, $p = .01$ (see Figure 4).

The analysis including the Tool AOI revealed no significant effects across age or condition.

Gaze Plot Analysis

A second type of data available using the Tobii Studio software is a gaze plot. This static gaze plot shows the order in which fixations to AOIs occur within a particular time frame

(see Figure 6). In structural alignment theory, observers should mentally align objects between two examples based on their common relational structure. Thus, we asked whether observers, when looking from one event to another, looked at a matching element in the second scene. We only coded gaze plots during the second pair of events, which was the same across conditions, and we only coded eye movements that switched from one event to the other.

Research assistants viewed a randomly selected subset ($n=19$) of participants' static gaze plots in Tobii Studio. They then coded by hand whether, when gaze switched from one event to the other, children fixated on the AOI in the second event that corresponded to the AOI they had been fixating on immediately prior to the switch in the first event (e.g., if fixating on the agent in Event 1, they then fixated on the agent in the event in Event 2). We counted the number of times children looked to a matching AOI when looking from one event to the other, and the number of times they looked to a non-matching AOI (in addition to instances in which children looked at different AOIs when looking from one event to the other, non-matches included instances in which children looked off screen or to areas outside of AOIs in one event, which did not correspond to an AOI fixation in the second event). We then computed a mean proportion score of looking to the match with number of matches/(matches + nonmatches). This revealed some evidence that children looked at matching elements when they were looking across scenes (AF $Mp = .40$, $SE = .09$; PA $Mp = .27$, $SE = .10$), but a univariate ANOVA with Condition (AF, PA) as an independent variable and proportion looking to the match as the dependent variable showed no difference across conditions ($F(1, 18) = .94$, ns).

DISCUSSION

In Study 2, important results emerged across different types of measures, and these results support the findings in Study 1. Specifically, 3 ½-year-old children in the PA condition were able to point to a correct depiction of a verb extension more often than were children in the All Far condition, and they exceeded the rate expected by chance. (If children simply preferred one test event over the other, an effect of condition would not have emerged.) This finding supports the finding in Study 1 in which children in the PA condition performed significantly more enactments than did children in the Control condition. A difference between the two experimental conditions emerged more strongly in this study than was seen in Study 1, although in this study, the benefit was seen only for the 3 ½ year-old children whereas in Study 1, the finding applied across age. Yet, it is important that converging evidence emerged across studies as it was highly possible that results in the enactment and pointing task could differ.

Also, as in Study 1, in this study children were more successful extending new verbs in the second trial than they were in the first test trial. In both studies, the second test trial included objects that differed more in shape, material or other properties from the objects seen in the learning trials. Thus, we predicted that this trial would be more difficult for children than the first trial. As both studies showed better performance in this difficult trial, it seems likely that experience with the first test trial, even though children often failed to extend the verb successfully in that trial, may have helped them succeed in the later trial. Future studies counterbalancing the test trial order will be needed to test whether this conjecture is true.

Looking at an element of an event is the first step in mentally processing information about that element. Two-½-year-olds in the PA condition (and perhaps 3 ½-year-olds, though their result only approached significance) increased their looking to agents and to affected objects only in the condition in which they had initially seen highly similar events. Recall that this result was from the second pair of events shown in the learning set in which observers in both conditions saw the same pair of events. This type of effect is just the sort that would be predicted by structural alignment theory: a progressive alignment experience should help naïve learners learn how to compare events. Prior research has shown that 2-year-old children learning new verbs expect verbs to refer to intentional and not accidental actions (Tomasello & Barton, 1984), and thus attending to the agent could be linked to generating hypotheses about her intentions. Behrend (1990) showed 3-year-olds exhibit a ‘result verb’ bias, or a bias to expect actions with different results to be linked to different verbs. Thus, attention to the affected object may be linked to attention to the result of the event. Because an increase in looking to the agent and affected object only occurred in the PA condition, one interpretation is that the PA experience helped children focus on these relevant objects; thus, is the first empirical evidence that this type of experience influences children’s precise eye movements. At the same time, we had hoped that the gaze plots would show that children were looking back and forth between the two events and seeking out matches between elements in one event with elements in another event. We failed to find that this occurred more in one condition than another, though we did find some visual aligning occurred.

At the same time, it was the 3 ½-year-olds who showed the clearest benefit from a PA experience at test. This suggests that children at the youngest age had begun to exhibit looking behaviors that could be of benefit to them in comparing events and learning new verbs. However, the test trials were sufficiently difficult that this benefit was not revealed at test in this youngest age, but emerged in children who were 1 year older. Further studies are needed to test this possible account fully.

GENERAL DISCUSSION

This paper describes two studies with multiple ages and different methodologies which both show that children benefit from seeing similar events that can be compared when learning new verbs. This paper describes data from the youngest age group ever shown to benefit from this experience in verb learning, and it provides new evidence of children's visual attention as they compare events, which is also important. Across studies, our data show that the ability to compare events and extend new verbs increases over development, with key developmental shifts occurring between 2 1/2- and 4 ½ years. For example, in Study 1, 2 ½-year-olds performed significantly fewer verb extensions than did 3 ½-year-olds but across age, children were more successful at extending new verbs in the PA condition than in the Control condition. Thus, there were developmental differences between 2 ½ and 3 ½-year-olds, but there were also commonalities across these ages. In Study 2, although there were developmental differences with 2 ½-year-olds in the PA condition showing increased looking to the agent and affected object, the 3½ year-olds showed the greatest benefit from the PA

condition at test. Taken together, this set of results suggests a gradual increase in the ability to compare events and extend verbs across age (or experience).

Additionally, a strength of these studies is that the same ages were included, and similar events and test trials were used across studies, and thus the influence of methodology (enactment vs. pointing to video) can be explored. Study 2 shows a benefit of progressive alignment at test, but only in 3 ½-year-olds, while Study 1 shows a benefit of progressive alignment at test, collapsed across age, but only in the more difficult second test trial. In some ways, these results suggest to us that the enactment procedure may be more sensitive to 2 ½-year-olds' verb knowledge than was our pointing procedure; 3 ½ year-olds (in the appropriate PA condition) appeared to fare well in both procedures. This conclusion, if true, is also interesting to consider because it is much easier for children to succeed purely by chance in a two choice pointing task than in an enactment task. Yet overall, even though conclusions about the usefulness of seeing similar pairs of events first was influenced somewhat by methodology, a benefit was

These results are consistent with a particular mechanism that has been shown to influence children's and adults' ability to solve analogies (e.g., Gentner, 1988)— the structural alignment of elements across two instances—which could also describe the mental processes observers use when they are comparing multiple dynamic events. A key prediction of structural alignment theory is that surface similarity across examples matters. That is, children (or naïve observers in a new task) use perceptual similarity to

guide what should be compared and how two instances could be compared. One byproduct of this attention to surface features initially, is that observers are more likely to notice opportunities to compare when they encounter two instances that are highly similar than when the instances differ from each other, and also may be more successful in aligning the elements in two examples effectively in this case. Previous research has shown that a progressive alignment experience helps children learn new part terms (Gentner et al., 2007) and solve a spatial mapping task (Loewenstein & Gentner, 2001). The present study, and Haryu et al. (2011), are the only studies that show this type of experience may also help children faced with comparing events while they try to learn new verbs. The present results are especially important because they reveal a benefit for progressive alignment experience in age groups younger than have been previously revealed (at 2½ and 3 ½ years) and also show when most children have moved beyond a need for progressive alignment experience (4 ½ years).

As mentioned in the introduction, structural alignment is not incompatible with the other two mechanisms that have been proposed for early verb learning. Children may initially associate whole events with particular verbs, and then use structural alignment to compare those events to each other. This process of aligning may help them create good hypotheses of a new verb's meaning (or this step may not be needed if the alignment leads the child to glean enough information from the comparisons). Future studies will be needed to test how these multiple mechanisms may interact during verb learning. However the present studies add to this body of knowledge by showing how children could learn how to align and compare multiple events, which is an important question.

Comparing multiple events to each other is only part of the verb learning problem as well. Further studies also will be needed to explore whether similar mechanisms could underlie cross-sentence comparisons, and how they may extend differently to different types of verbs (e.g., verbs of perception, mental verbs) or different types of events. In addition, comparisons in everyday life often will involve comparing a memory of a previous related event with a present event, and studies are needed to explore how memory processes interact with comparison processes. Nevertheless, these studies provide important new evidence of the benefit children may accrue from exposure to sets of similar events, which could be used by researchers interested in developing new interventions for children experiencing language delay (e.g., Schwarz, 2013). The problem of productively extending new verbs will be a difficult problem for researchers to solve, but is obviously a problem children solve by the age 4 or 5 years. That children have a remarkable ability to be productive and creative users of language is a hallmark of human language, and should be celebrated as well as adequately explained.

REFERENCES

- Behrend, D. A. (1990). The development of verb concepts: Children's use of verbs to label familiar and novel events. *Child Development*, *61*, 681-696. doi: 10.1111/1467-8624.ep5858725
- Bowerman, M. (1985). Beyond communicative adequacy: From piecemeal knowledge to an integrated system in the child's acquisition of language. In K. E. Nelson (Ed.), *Children's Language, Vol. 5* (pp. 369-398). Hillsdale, NJ England: Lawrence Erlbaum Associates.

- Childers, J. B. (2011). Attention to multiple events helps two-and-a-half-year-olds extend new verbs. *First Language, 31*(1), 3-22. doi:10.1177/0142723710361825
- Childers, J. B., Heard, M. E., Ring, K., Pai, A., & Sallquist, J. (2012). Children use different cues to guide noun and verb extensions. *Language Learning and Development, 8*(3), 233-254. doi: 10.1080/15475441.2011.585285
- Childers, J. B. & Paik, J. H. (2009). Korean- and English-speaking children use cross-situational information to learn novel predicate terms. *Journal of Child Language, 36*, 201-24. doi:10.1017/S0305000908008891
- Corbetta, D., Guan, Yu & Williams, J. L. (2011). Infant eye-tracking in the context of goal-directed actions. *Infancy, 17*, 102-125. doi: 10.1111/j.1532-7078.2011.00093.x
- DeLoache, J. S. (1987). Rapid change in the symbolic functioning of very young children. *Science, 238*, 1556-1557. doi:10.1126/science.2446392
- Fenson, L., Dale, P. S., Reznick, J. S., Bates, E., Thal, D., Hartung, J., et al. (1993). *Guide and technical manual for the MacArthur Communicative Development Inventories*. San Diego, CA: Singular Press.
- Fisher, C., Hall, D. G., Rakowitz, S., & Gleitman, L. (1994). When it is better to receive than to give: Syntactic and conceptual constraints on vocabulary growth. *Lingua, 92*, 333-375.
- Forbes, J. N. & Farrar, M. J. (1993). Children's initial assumptions about the meaning of novel motion verbs: Biased and conservative? *Cognitive Development, 8*, 273-290.
- Forbes, J. N. & Poulin-Dubois, D. (1997). Representational change in young children's understanding of familiar verb meaning. *Journal of Child Language, 24*, 389-406.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170. doi:10.1037/0003-066X.52.1.45

Gentner, D. (1988). Metaphor as structure mapping: The relational shift. *Child Development*, 59, 47-59. doi:10.2307/1130388

Gentner, D. (2002). Why we're so smart: Language and thought in early development. Paper presented at the 32nd Annual Carnegie Symposium on Cognition, Carnegie Mellon University, Pittsburgh, PA, June 7-9, 2002.

Gentner, D., Loewenstein, J. & Hung, B. (2007). Comparison facilitates children's learning of names for parts. *Journal of Cognition and Development*, 8, 285-307. doi: 10.1080/15248370701446434

Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist*, 52, 45-56.

Gleitman, H., & Gleitman, L. (1991). Language use and language judgment. In W. Kessen, A. Ortony, F.M. Craik (Eds.), *Memories, thoughts, and emotions: Essay in honor of George Mandler* (pp. 99-122). Hillsdale, NJ England: Lawrence Erlbaum Associates, Inc.

Golinkoff, R. M., Jacquet, R. C., Hirsh- Pasek, K., & Nandakumar, R. (1996). Lexical principles may underlie the learning of verbs. *Child Development*, 67, 3101-3119. doi:10.1111/1467-8624.ep9706244850

Gomez, R. L., & Gerken, L. (1999). Artificial grammar learning by 1-year-olds leads to specific and abstract knowledge. *Cognition*, 70, 109-135.

- Haryu, E., Imai, M., & Okada, H. (2011). Object similarity bootstraps young children to action-based verb extension. *Child Development, 82*, 674-86. DOI: 10.1111/j.1467-8624.2010.01567.x
- Huttenlocher, J., Smiley, P. & Charney, R. (1983). Emergence of action categories in the child: Evidence from verb meanings. *Psychological Review, 90*, 72-93. doi: 10.1037/0033-295X.90.1.72
- Imai, M., Haryu, E., & Okada, H. (2005). Mapping novel nouns and verbs onto dynamic action events: Are verb meanings easier to learn than noun meanings for Japanese children?. *Child Development, 76*, 340-355. doi: 10.1111/j.1467-8624.2005.00849.x.
- Kersten, A. W., & Smith, L. B. (2002). Attention to novel objects during verb learning. *Child Development, 73*(1), 93-109.
- Knoeferle, P, Crocker, M. W., Scheepers, C. & Pickering, M. J. (2005). The influence of the immediate visual context on incremental thematic role assignment: evidence from eye-movements in depicted events. *Cognition, 95*, 95-127. doi: 10.1016/j.cognition.2004.03.002
- Kotovsky, L. & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development, 67*, 2797-2822. doi: 10.2307/1131753
- Landau, B., Smith, L. B., & Jones, S. S. (1988). The importance of shape in early lexical learning. *Cognitive Development, 3*, 299-321. doi: 10.1016/0885-2014(88)90014-7
- Landau, B., Smith, L. B., & Jones, S. S. (1998). Object shape, object function, and object name. *Journal of Memory and Language, 38*, 1-27. doi: 10.1006/jmla.1997.2533
- Landis, J. R.; & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics, 33*, 159–174. doi: 10.2307/2529310

- Loewenstein, J., & Gentner, D. (2001). Spatial mapping in preschoolers: Close comparisons facilitate far mappings. *Journal of Cognition and Development, 2*, 189-219. doi: 10.1207/S15327647JCD0202_4
- Loewenstein, J., Thompson, L. & Gentner, D. (2003). *Analogical learning in negotiation teams: Comparing cases promotes learning and transfer*. Academy of Management Learning and Education, 2, 119-127. doi: 10.5465/AMLE.2003.9901663.
- Maguire, M. J., Hirsh- Pasek, K., Golinkoff, R. M., & Brandone, A. C. (2008). Focusing on the relation: Fewer exemplars facilitate children's initial verb learning and extension. *Developmental Science, 11*, 628-634. doi: 10.1111/j.1467-7687.2008.00707.x
- Markman, E. M. (1989). *Categorization and naming in children: Problems of induction*. Cambridge, MA US: The MIT Press.
- Medina, T. N., Snedeker, J., Trueswell, J. C., & Gleitman, L. R. (2011). How words can and cannot be learned by observation. *Proceedings of the National Academy of Sciences, 108*, 9014-9019.
- Netter, J., Wasserman, W., & Kutner, M. H. (1985). *Applied linear statistical methods: Regression, analysis of variance, and experimental designs*. Homewood, IL: Richard D. Irwin.
- Papafragou, A., Hulbert, J. & Trueswell, J. (2008). Does language guide event perception? Evidence from eye movements. *Cognition, 108*, 155-184. doi: 10.1016/j.cognition.2008.02.007
- Pinker, S. (1989). *Learnability and Cognition*. MIT Press: Cambridge, MA.
- Roberts, K. (1983). Comprehension and production of word order in Stage I. *Child Development, 54*, 443-449.

- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926-1928. DOI: 10.1126/science.274.5294.1926
- Schwarz, A. L. (2013). *Do acting out verbs and receiving structurally aligned examples increase verb comprehension in 30-month-olds during familiar social routines?* Unpublished dissertation.
- Scott, R.M., & Fisher, C. (2009). Two-year-olds use distributional cues to interpret transitivity-alternating verbs. *Language and Cognitive Processes*, 24, 777-803. doi:10.10180/01690960802573236
- Scott, R. M., & Fisher, C. (2012). 2.5-Year-olds use cross-situational consistency to learn verbs under referential uncertainty. *Cognition*, 122(2), 163-180. doi: 10.1016/j.cognition.2011.10.010
- Slobin, D. I. (1981). The origins of grammatical encoding of events. In W. Deutsch (Ed.) *The child's construction of language*. New York: Academic Press.
- Slobin, D. I. (1985). Crosslinguistic evidence for the language-making capacity. In D. I. Slobin (Ed.), *The crosslinguistic study of language acquisition, Volume 2: Theoretical Issues*, Hillsdale, NJ: Lawrence Erlbaum Associates.
- Smith, L., & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition*, 106(3), 1558-1568. doi: 10.1016/j.cognition.2007.06.010
- Talmy, L. (1975). Semantics and syntax of motion. In J. P. Kimball (ed), *Syntax and Semantics: Volume 4*. New York: Academic Press.
- Tomasello, M. (1992). *First verbs: A case study of early grammatical development*. New York, NY: Cambridge University Press. doi:10.1017/CBO9780511527678

Tomasello, M. (2000). Do young children have adult syntactic competence? *Cognition*, 74, pp. 209-253. doi: 10.1016/S0010-0277(99)00069-4

Tomasello, M., & Barton, M. E. (1994). Learning words in nonostensive contexts. *Developmental Psychology*, 30, 639-650. doi: 10.1037/0012-1649.30.5.639

Trueswell, J. C., Medina, T. N., Hafri, A., & Gleitman, L. R. (2013). Propose but verify: Fast mapping meets cross-situational word learning. *Cognitive Psychology*, 66, 126-156. doi: 10.1016/j.cogpsych.2012.10.001

Waxman, S. R., Lidz, J. L., Braun, I. E., & Lavin, T. (2009). Twenty four-month-old infants' interpretations of novel verbs and nouns in dynamic scenes. *Cognitive Psychology*, 59, 67-95. doi: 10.1016/j.cogpsych.2009.02.001

Yu, C. & Smith, L. B. (2012). Modeling cross-situational word-referent learning: Prior questions. *Psychological Review*, 119, 21-39. doi: 10.1037/a0026182

Yu, C., & Smith, L. B. (2007). Rapid word learning under uncertainty via cross-situational statistics. *Psychological Science*, 18, 414-420. doi: 10.1111/j.1467-9280.2007.01915.x

NOTES

¹Because proportional data may be subject to instability of error term variances, arc sine transformations were applied to all proportional data (Netter, Wasserman, & Kutner, 1985). These analyses with transformed data showed the same patterns as reported here.

²One reason children may have failed to point on one or more trials was the test trials as presented could not be paused using the Tobii Studio software. Most (but not all) of the children who were excluded because they refused point at test were in the youngest 2 ½-year-old age group, with approximately equal representation from both conditions.

³Study 2 did not include a control group because the task of recognizing an action at test (even with new objects) seemed so easy to perform that we did not predict that the advantage of seeing multiple events vs. a single event would emerge. Given this study design, it will not be possible to conclude from the results that the comparison of multiple events is more useful than is seeing a single event, but it will be possible to show whether the types of events that are compared influence performance.

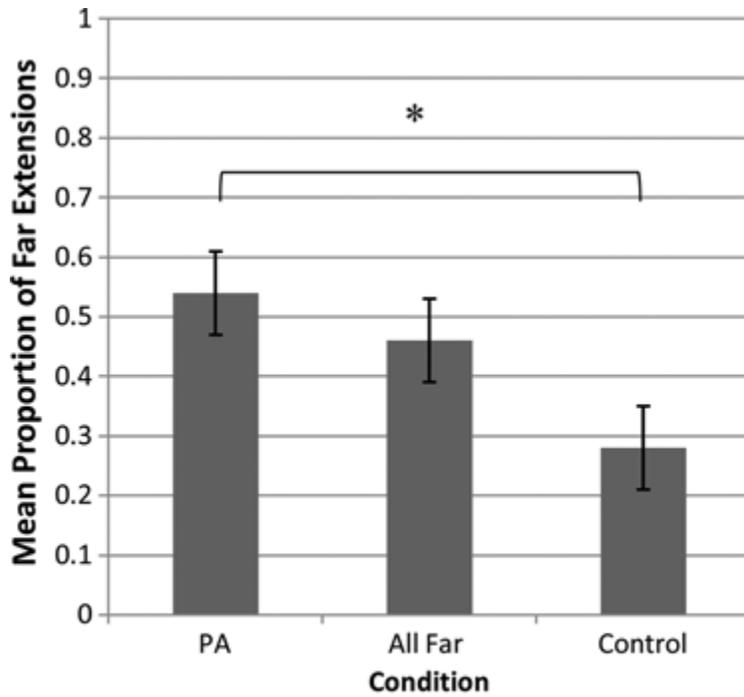
⁴Additional eye tracking details:

Size of stimuli. The scenes fit into a 24 cm x 14 cm area (visual angle: 32.9 x 19.6 degrees). The agent was approximately 18 x 7 cm (visual angle: 25 x 9.85 degrees), and the objects ranged in size, including objects that were 3 x 2 cm (4.33 x 2.8), 9 x 1 cm (12.6 x 1.4 degrees), 4 x 3 cm (5.6 x 4.2 degrees) and 5 x 2.5 cm (7 x 3.5 degrees).

Fixation filter: We used the standard I-VT fixation filter in the Tobii Studio software, which has been set to yield accurate fixation data for the most common eye tracking uses. The noise reduction setting was disabled. The minimum fixation duration was set at 60 ms, which is a conservative setting to allow for complex visual behavior.

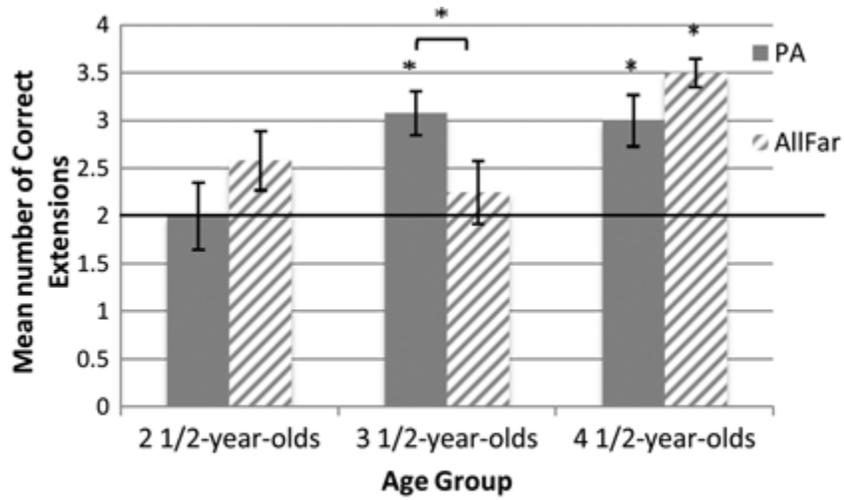
Processing of eye-movement data: As reported, we averaged individuals' fixation duration for a specific AOI across scenes to be sure a particular scene or object did not have a major effect on the results. We also excluded participants as reported in the Participant section. However, beyond these considerations, and our use of the standard I-VT fixation filter, we did not further process the eye tracking data.

Figure 1. Study 1: Results. Note: Graph shows mean proportion of far extensions by condition, error bars = SEM; * $p < .05$.



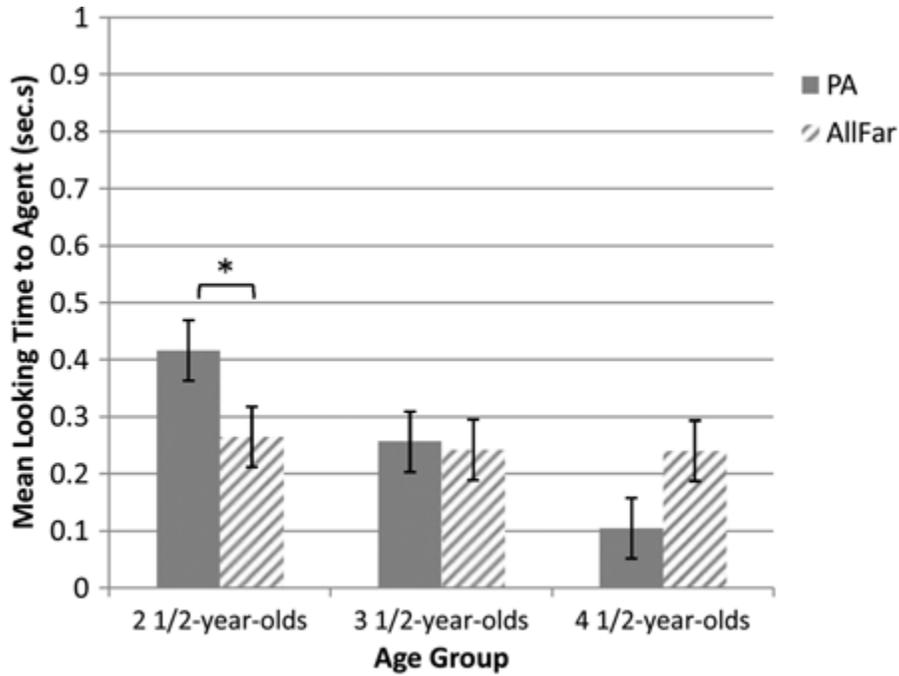
Note: Graph shows mean proportion of far extensions by condition, error bars = SEM; * $p < .05$.

Figure 2. Study 2: Pointing results. Note: Graph shows mean number of correct verb extensions by Age Group and Condition, error bars = SEM. Line represents mean expected by chance; * $p < .05$.



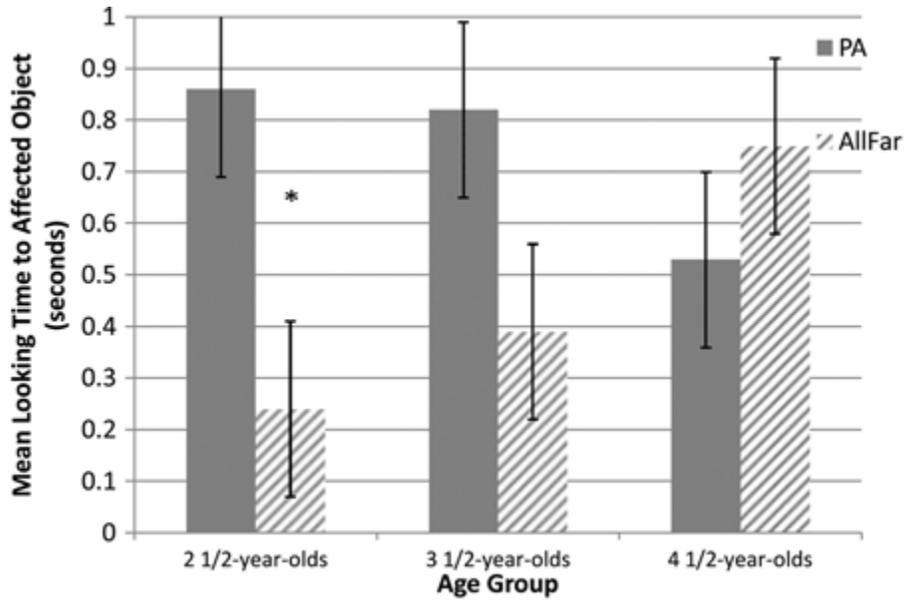
Note: Graph shows mean number of correct verb extensions by Age Group and Condition, error bars = SEM. Line represents mean expected by chance; * $p < .05$.

Figure 3. Study 2: Total fixation duration in the second learning trial: Agent AOI. Note: Graph shows mean fixation duration by age group and condition, error bars = SEM; $p < .05$.



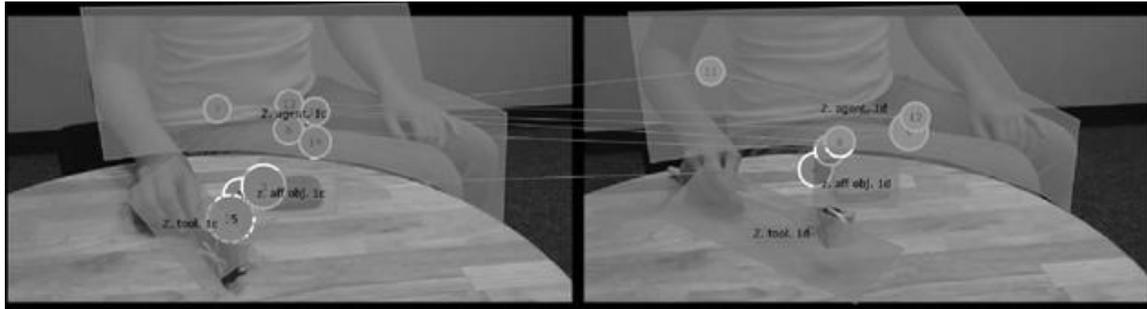
Note: Graph shows mean fixation duration by age group and condition, error bars = SEM; $p < .05$.

Figure 4. Study 2: Total fixation duration in the second learning trial: Affected Object AOI. Note: Graph shows mean fixation duration by age group and condition, error bars = SEM; * $p < .05$.

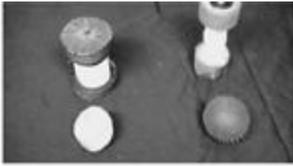


Note: Graph shows mean fixation duration by age group and condition, error bars = SEM; * $p < .05$.

Figure 5. Example of AOIs with a gaze plot.



Appendix A: Study 1, stimulus set: 'tam'.



PA Condition, Learning trials

All Far Condition, Learning trials

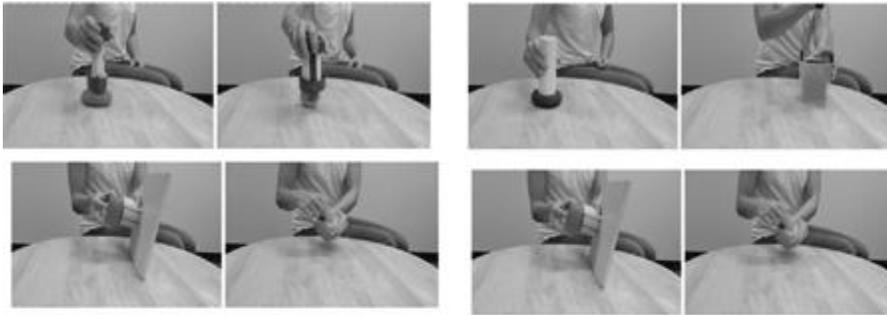


Test Trial 1



Test Trial 2

Appendix B: Study 2, stimulus set: 'tam'.



PA Condition, Learning trials

All Far Condition, Learning trials



Test Trial 1



Test Trial 2