Building Interconnected Networks of Word Knowledge Over Time

Layla Unger (layla.unger@york.ac.uk)
Department of Psychology, University of York &
Department of Psychology, Ohio State University

Olivera Savic (o.savic@bcbl.eu)
Basque Center on Cognition, Brain and Language

Vladimir M Sloutsky (sloutsky@psy.ohio-state.edu)
Department of Psychology, Ohio State University

Abstract

To become fluent language users, children must learn not only individual words but also connections between them. For example, connections are vital for understanding that "apples" are "yummy", something you can "eat" and similar to "oranges". To date, there is evidence that children develop increasingly sophisticated abilities to form these connections from regularities in the way that words co-occur with other words that are ubiquitous in everyday language. Yet despite the fact that children repeatedly experience such regularities day-to-day, existing evidence focuses just on what children learn from a single experience. We used a multi-session approach to examine the connections children build from repeated exposures over time. We found that from age four to six, children not only improve in their formation of connections between words from regularities in language, but also in building increasingly richly interconnected knowledge from one experience to the next.

Keywords: word learning; distributional semantics; language models

Introduction

From everyday language experiences, children not only learn thousands of words, but also form meaningful connections between words. For example, a child may learn that "oranges" are "yummy", something you can "eat", and similar to "apples". These connections come to play a vital role in language fluency (Borovsky, Ellis, Evans, & Elman, 2016; Neuman, Newman, & Dwyer, 2011; Nation & Snowling, 1997). For example, they can allow a listener to understand that "I'd love some of those oranges, they're so yummy" expresses a desire to "eat" oranges, rather than play or otherwise interact with them. Likewise, they can support generalizations between words similar in meaning, such as allowing a listener who has heard an "orange" described as "ripe" to understand and produce utterances about "ripe apples", even if they have never heard such utterances before. Although children might learn something about words and connections between them from a single language experience (Carey & Bartlett, 1978), such learning typically unfolds *over time* as children accumulate day-to-day language experiences (Bolger, Balass, Landen, & Perfetti, 2008; Hills, Maouene, Riordan, & Smith, 2010; Roy, Frank, DeCamp, Miller, & Roy, 2015; McMurray, Horst, & Samuelson, 2012; Unger & Fisher, 2019). Indeed, knowledge about words and their meanings becomes increasingly richly interconnected with development (Unger & Fisher, 2016; Crowe & Prescott, 2003; Storm, 1980). How do children transform everyday language experiences into this interconnected word knowledge?

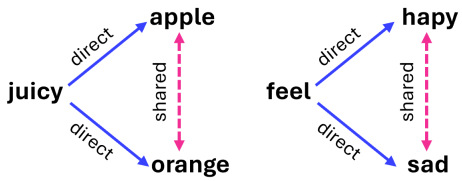

Figure 1: Examples of direct and shared co-occurrence regularities that can support connections between words.

Children's everyday language experiences provide a rich source of information for learning connections between words because they are full of regularities in the way words occur with other words (Figure 1). For example, both "orange" and "apple" tend to co-occur with "yummy". In principle, it is therefore possible to connect "orange" and "apple" to "yummy" based on their tendency to directly co-occur with each other in sentences and utterances. Moreover, it is possible to integrate across these direct co-occurrence regularities to connect "orange" to "apple" based on their shared co-occurrence with "yummy". The *potential* power of these regularities has been amply demonstrated. Increasingly sophisticated language models have successfully captured much of human interconnected word knowledge just from learning the way words occur with other words (Landauer & Dumais, 1997; Jones & Mewhort, 2007; Mikolov, Chen, Corrado, & Dean, 2013; Pennington, Socher, & Manning, 2014), even when trained on relatively small samples of language input to children (Huebner & Willits, 2018; Fourtassi, 2020).

If word co-occurrence regularities are so useful for modeling interconnected word knowledge, perhaps children also acquire increasingly richly interconnected word knowledge from repeated exposure to word co-occurrences in language. This possibility rests on two assumptions: that children are indeed sensitive to regularities of word use, and that children build up interconnected word knowledge from the regularities of word use they experience from one day to the next. Yet our understanding of this possibility comes from studies of the knowledge children pick up from a *single* bout of exposure to input containing regularities. This evidence suggests that some sensitivity to co-occurrence regularities is present early in childhood and undergoes substantial developmental changes. Specifically, an ability to form connections based on direct co-occurrence appears to emerge early

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

and improve across childhood. This ability is only gradually supplemented by an ability to form connections based on shared co-occurrence (Savic, Unger, & Sloutsky, 2022; Bauer & San Souci, 2010; Bauer, Cronin-Golomb, Porter, Jaganjac, & Miller, 2020; Schlichting, Guarino, Schapiro, Turk-Browne, & Preston, 2017; Schlichting, Guarino, Roome, & Preston, 2021). This pattern suggests that the increasingly richly interconnected word knowledge that children acquire over the course of development may come not just from accumulating exposure to regularities, but also from developmental changes in children's abilities to transform these regularities into interconnected knowledge.

We thus have two potentially powerful contributors to children's interconnected word knowledge: their accumulated exposure to regularities of word use over time, and developmental changes in their abilities to transform this exposure into interconnected word knowledge. Yet our only insight into children's transformation of regularities in language into interconnected word knowledge comes from studies of what children learn from a single bout of exposure. We do not know whether children do indeed capitalize on repeated exposures to regularities to build increasingly richly interconnected word knowledge, nor do we know whether children's use of repeated exposures changes with development.

In what follows, we first evaluate what is known about the development of children's abilities to extract interconnected knowledge from exposure to regularities in language and other domains. We then present a multi-day study in which we systematically investigated the development of children's formation of connections between words from exposure to regularities in language over time.

Forming Connections from a Single Experience

Sensitivity to regularities during development has been studied most extensively in the field of *statistical learning*. This field has focused on sensitivity to direct co-occurrence regularities, such as regularities in which sets of two or three speech sounds or shapes always occur together in the same order. In a typical experiment, participants receive a single bout of exposure to these regularities (often for just a few minutes), and are then tested on their ability to distinguish between sets of stimuli that reliably co-occurred versus those that did not. There is extensive evidence that some sensitivity to these regularities emerges in infancy (e.g., Saffran, Aslin, & Newport, 1996; Romberg & Saffran, 2010). Moreover, a handful of studies suggest that at least for some types of stimuli (e.g., non-speech sounds and shapes), this sensitivity improves across years of childhood (Arciuli & Simpson, 2011; Raviv & Arnon, 2018; Shufaniya & Arnon, 2018).

Research in a few domains has also examined the *relative* developmental trajectories for sensitivity to direct versus shared co-occurrence. The typical approach has been to present participants with triads of items, in which one item (A) separately co-occurs with two other items (B and C). Subsequently, participants are tested for connections between items that directly co-occur (A-B, A-C) and share co-occurrence (B-C). This approach has been used to study the formation of connections between arbitrary pictures such as faces, scenes and objects (Schlichting et al., 2017, 2021), facts (Bauer & San Souci, 2010; Bauer et al., 2020) and in one recent series of studies, words (Savic et al., 2022).

This work has revealed that in early childhood, there is some ability to pick up on direct co-occurrence. With age, this ability: (1) improves, and (2) is supplemented by an ability to pick up on shared co-occurrence that likewise improves with age. Importantly, this developmental pattern highlights that merely getting better at forming connections based on direct co-occurrence is necessary but insufficient for forming connections based on shared co-occurrence. Instead, this evidence points to an independent developing ability to extrapolate or integrate across instances of direct co-occurrence to form connections based on shared co-occurrence.

These developmental changes may allow children to build increasingly richly interconnected word knowledge from their everyday language experiences. Yet, our understanding is limited by its focus on what children pick up from a single experience, as we discuss below.

Forming Connections from Repeated Experiences

Exposure to language is an everyday experience. Children thus have repeated opportunities to both pick up on regularities in language in the first place, and build upon what they have already learned from previous exposures to regularities. Yet the contribution of repeated exposure to forming connections between words remains little understood. To the best of our knowledge, only one study in adults and one study in children shed light on how repeated exposure to regularities in language foster connections between words.

In a study with adults, McNeill (1963) used the A-B-C triad approach described above to embed triads of novel pseudowords into sentences. Within a single session, participants received two, four or six repetitions of each triad before their formation of connections was assessed. Given two or four repetitions, participants primarily connected pseudowords that directly co-occurred; only with six repetitions did participants robustly connect words that shared co-occurrence. In adults, repeated exposures may thus build increasingly richly interconnected word knowledge.

These findings are in stark contrast with the one study of repeated exposures in children. In Savic et al. (2022)'s Experiment 2, four-year-old children repeatedly heard short stories each embedded with 10 repetitions of A-B-C triads across three days. For adults in a preceding experiment (Experiment 1), the repetitions within a single session had been enough for them to form connections based on both direct and shared co-occurrence. For children, further repetition of the stories across days only supported children's formation of increasingly strong connections based on direct co-occurrence.

Together, these studies hint at an interplay between experience and development. Repeated exposure to regularities in language may provide opportunities to build increasingly richly interconnected word knowledge. Yet abilities to cap-

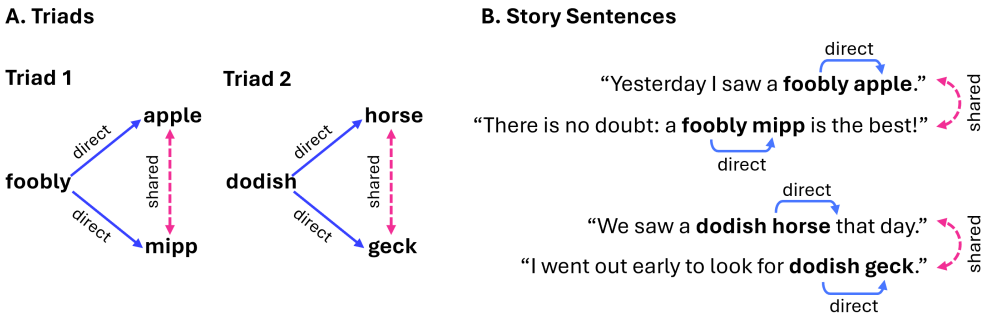

Figure 2: Triads (A) and examples of words from embedded in story sentences.

italize on these repeated opportunities may themselves undergo substantial development across childhood. This hinted interplay suggests that the emergence of increasingly interconnected word knowledge (Unger & Fisher, 2016; Crowe & Prescott, 2003; Storm, 1980) comes from accumulating language experience and developing abilities to form connections from experience.

Present Study

The goal of the present study was to illuminate how interconnected word knowledge emerges from the interplay between repeated exposures to regularities in language and developmental changes in abilities to pick up on these regularities. We used a version of the A-B-C approach that prior work used to demonstrate developmental differences in sensitivity to word co-occurrence.

In this study, children aged four, five and six heard short stories embedded with triads of words that could support the formation of connections based on both direct and shared co-occurrence (Figure 2). In these triads, a novel pseudoadjective (e.g., "foobly") directly co-occurred with a highly familiar noun (e.g., "apple") in some sentences, and another novel pseudonoun (e.g., "mipp") in others. The inclusion of pseudowords in these triads was designed to ensure that the connections children could form were entirely new. From these regularities, children could form connections based on direct co-occurrence ("foobly"-"apple", "foobly"-"mipp"). In addition, they could form connections between the familiar noun and pseudonoun ("apple"-"mipp") based on their shared co-occurrence with the same pseudoadjective ("foobly").

Children repeatedly heard these stories across three sessions that took place on separate days over the course of a month. During each session, children's formation of connections between words was assessed. This approach allowed us to tackle key unknowns about the interplay of experience and development in building interconnected word knowledge.

From previous studies of what children learn from a single bout of exposure, we expected that the very first session will reveal that the formation of connections based on direct and shared co-occurrence improves asynchronously with age, with connections based on direct co-occurrence emerging earlier than connections based on shared co-occurrence. Critically, repeated exposure provides additional opportunities to strengthen these connections. Do children capitalize similarly on these opportunities, or do older children more successfully harness these opportunities to form increasingly strong connections?

Methods

Participants

Participants were 4, 5, or 6-year olds recruited via social media posts in the US. Following exclusions based on performance on "control" trials used to assess attentiveness and task comprehension (see below), sample sizes consisted of 27 4-year-olds, 29 5-year-olds, and 32 6-year-olds.

Materials

Triads. The primary materials were two triads of words (Figure 2A). Each triad contained a novel pseudoadjective (e.g., "foobly") that was paired with both a familiar noun (e.g., "apple") and a novel pseudonoun (e.g., "mipp"). Thus, each triad contained two pairs of words that were embedded in a short story as described below, for a total of four word pairs.

Story. Pairs of words from the triads were embedded within sentences that were combined in a short story (Figure 2B) in which an alien travels to different planets. Sentences were designed to be uninformative about the meanings of novel pseudowords, such as "Yesterday I saw a foobly apple", "On the next day, I looked for a foobly mipp". Each of the four Triad word pairs was embedded within 10 sentences, for a total of 40 sentences within the story.

Pilot work suggested that children in the youngest age group might struggle to learn even the simpler connections based on direct co-occurrence when all four word pairs were intermixed the story, probably because this introduces four novel words at the same time. We therefore presented the story in two parts, in which Part 1 presented 20 sentences containing the pairs in which the pseudoadjective directly co-occurred with the familiar noun, and Part 2 presented 20 sentences containing pairs in which the pseudoadjective directly co-occurred with the pseudonoun. The full story was recorded by a female native speaker of English.

Muffled Test Sentences. We developed additional sentences for use in a "Muffled Test" designed to test connections

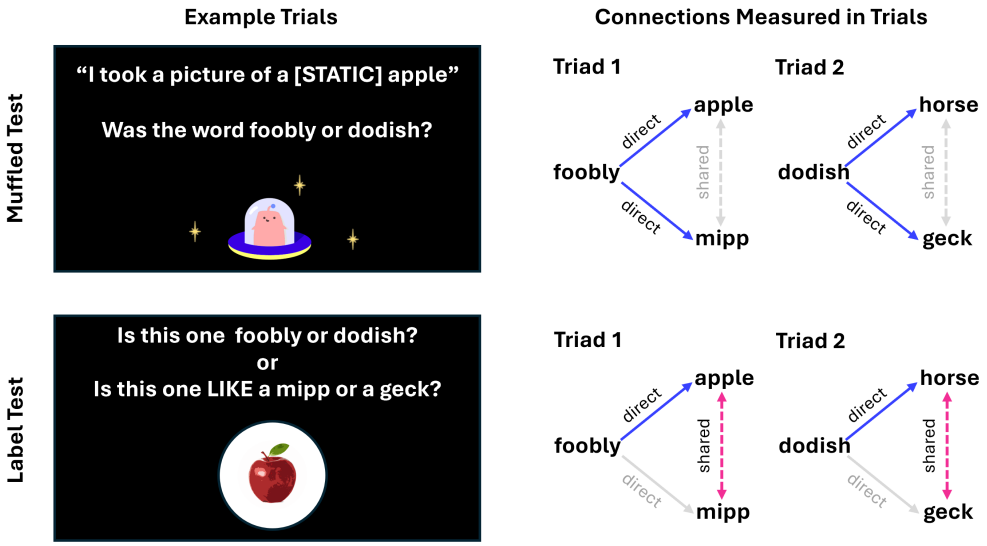

Figure 3: The top panel shows an example of a Muffled Test sentence used with a prompt in a trial (left) and the connections that Muffled Test trials assessed. The bottom panel shows an example of a Label Test picture that could be presented with one of two prompts in a Label Test trial (left) and the connections that Label Test trials assessed.

based on direct co-occurrence (see Procedure and Figure 3). These sentences were similar to story sentences, but in each sentence, one of the words in a pair was replaced by an audio clip of radio static, as in "I'd love to see a [STATIC] mipp". We developed four sentences to test the formation of connections between each pair, for a total of 16 sentences.

Label Test Pictures. We used pictures of the familiar nouns in triads for a "Label Test" designed to test whether participants connected the familiar noun to both the pseudoadjective in the same triad based on direct co-occurrence, and the pseudonoun in the same triad based on shared co-occurrence. As described below and shown in Figure 3, each picture was presented with one prompt to assess connections based on direct co-occurrence, and another prompt to assess connections based on shared co-occurrence. In total, we used four pictures of each familiar noun.

Control Items. We developed a set of "control" items for use in both the Muffled and Labeled tests to check that participants were attentive and understood the tests. The tests assessed whether children had formed connections between a pseudoadjective and a noun in the same triad, or a pseudonoun and a noun in the same triad. Accordingly, control items consisted of a familiar adjective and two familiar nouns that are already likely to be connected in young children's vocabularies, such as "furry"-"dog" and "cat"-"dog". Items were frequent words that reliably co-occur with each other in children's everyday language input (MacWhinney, 2000).

Procedure

Participants took part in three sessions, each on a separate day within one month (mean gap between sessions was 6 days). Sessions were hosted online on the Gorilla platform, took place on Zoom and were guided by an experimenter. Each session followed the same procedure to provide children with exposure to regularities in language, and assess their formation of connections based on these regularities.

Introduction and Warmup Muffled Test. Participants were first introduced to an alien called Izzy who likes to tell travel stories. Participants then completed a warmup version of the Muffled Test that was subsequently used to assess children's formation of connections based on direct co-occurrence (see below). To introduce the Muffled Test, children were told that Izzy was going to say some things over the radio in her spaceship, but that the radio sometimes makes a funny sound. The experimenter would give the child two word choices, and the child would pick the one that Izzy probably said when the radio made the funny sound. The experimenter then played an example sentence, in which Izzy said "My mom saw a [STATIC] butterfly", and the experimenter explained that if the word choices were "pretty" and "sharp", the best answer would be "pretty", because "pretty butterfly" makes sense. Children then completed the warmup Muffled Test containing four control item trials, such as hearing "I'd really love a [STATIC] cookie" and being asked whether the muffled word was "yummy" or "wet".

Story and Muffled Test. Participants were told that Izzy would tell a story about a visit to an alien planet using some fun alien words. Participants then completed the Story and Muffled Tests in two parts. At the beginning of each part, the experimenter played audio clips of the pseudowords that would be presented in the subsequent story, and prompted the child to repeat each one. The participant then heard Part 1 or Part 2 of the story. To make listening more interactive, children were intermittently prompted to click the screen to trigger an animation in which Izzy travelled part-way across a map of an alien landscape and then said four sentences.

Participants then completed Part 1 or 2 of the Muffled Test consisting of 8 trials. The Muffled Test was introduced by explaining that the child's job was to choose a word to fill in the funny radio sound that Izzy would use. Each trial presented one of the Muffled Test sentences in which an audible word from a triad was accompanied by an audio clip of static. Children chose between two options to fill in the muffled word: one from the same triad, and one from the opposite triad. For example, when the familiar noun was audible, the choices were the pseudoadjective from the same versus different triad as the audible familiar noun. Accordingly, the Muffled Test assessed children's formation of connections based on the direct co-occurrence of words from the same triad.

Label Test. The experimenter first explained that the child's job was to help Izzy describe pictures she had taken on her travels. On each trial, participants saw a picture and were asked to label it with one of two options: the two pseudoadjectives, or the two pseudonouns. Trials in which the two options were pseudoadjectives measured connections based on direct co-occurrence from the accuracy with which participants labelled pictures of a familiar noun (e.g., apple) using the pseudoadjective with which it directly co-occurred (e.g., "foobly"). The same logic applies for measuring connections based on shared co-occurrence from accuracy on pseudonoun trials (e.g., using "mipp" to label apples). To ensure that children were asked to label the pictures in a way that is grammatical in English, for pseudoadjectives, children were asked, "Is this one [pseudoadjective 1] or [pseudoadjective 2]?", and for pseudonouns, children were asked, "Is this one like a [pseudonoun 1] or a [pseudonoun 2]?". Control trials used the same format to assess attentiveness and task comprehension by prompting children to label pictures with familiar adjectives and nouns, such as labeling a dog as "furry" rather than "yummy", and like a "cat" rather than "ice cream".

Results

We first assessed attentiveness and task comprehension based on accuracy on control trials in the Muffled and Label Tests, and excluded children whose performance was poor on these trials across multiple sessions. Each session contained four Muffled Test control trials and six Label Test control trials. Our threshold for including data from a session was accurate performance on the majority of control trials (three Muffled Test control trials and four Label Test control trials). We excluded participants who met these criteria for fewer than two out of three sessions. The numbers of participants excluded were 16 four-year-olds, 8 five-year-olds and 2 six-year-olds.

Our analyses were designed to measure the contributions of age and repeated exposure across sessions to the formation of connections based on direct and shared co-occurrence. All analyses were performed as generalized linear mixed effects models on raw accuracy data using the lme4 package in R [REF]. Because the maximum random effects structure that achieved convergence across all analyses only included a random intercept for participant, we used this random effects structure for consistency across analyses.

Muffled Test Analysis: Formation of Connections from Direct Co-Occurrence

The Muffled Test allowed us to measure the formation of connections from the direct co-occurrences of a pseudoadjective with the familiar noun and pseudonoun in the same triad. We considered contributions and interactions between the following factors: Age (4-, 5-, or 6-year-old), Session (1, 2, or 3) and Pair Type (pseudoadjective-familiar noun or pseudoadjective-pseudonoun). We included this final factor to account for the possibility that children might more easily form connections that involve at least one familiar word.

Muffled Test performance is shown in Figure 4, top panel. An initial omnibus analysis revealed significant main effects for all factors (all ps $< .0001$) and significant interactions between Age and both Session ($\chi^2(1) = 17.30$, $p < .0001$) and Pair Type ($\chi^2(1) = 18.80$, $p < .0001$). We therefore conducted follow-up analyses testing the effects of Session, Pair Type and their interaction within each Age group.

In 4-year-olds, the analysis revealed only a main effect of Pair Type, in which children formed stronger connections between directly co-occurring pairs that included a familiar word versus those that did not ($\chi^2(1) = 6.02$, $p = .01$). There was no effect of Session, nor an interaction between Session and Pair Type (ps $> .09$). 5- and 6-year-olds showed a similar effect of Pair Type (ps ¡ .0001), but additionally showed significant effects of Session (ps $< .01$), in which connections strengthened across sessions. Thus, all children more readily formed connections between words when at least one of them was familiar, but only for older children did these connections strengthen with repeated exposure.

Label Test Analysis: Formation of Connections from Direct and Shared Co-Occurrence

The Label Test allowed us to measure connections based on the familiar noun's direct co-occurrence with the pseudoadjective and shared co-occurrence with the pseudonoun from the same triad. We considered contributions and interactions between: Age (4-, 5-, or 6-year-old), Session (1, 2, or 3) and Regularity Type (direct versus shared).

The patterns captured in these analyses are depicted in Figure 4, bottom panel. An initial omnibus analysis revealed significant main effects for all factors (all ps $< .0001$) and significant interactions between Age and both Session ($\chi^2(1) = 11.93$, $p = .0006$) and Regularity Type ($\chi^2(1) = 3.87$, $p = .049$). We therefore conducted follow-up analyses testing the effects of Session, Regularity Type and their interaction within each Age group.

In 4-year-olds, the analysis revealed only a main effect of Regularity Type, in which children formed stronger connections based on direct versus shared co-occurrence ($\chi^2(1) = 5.15$, $p = .02$). In contrast, there was no effect of Session, nor an interaction between Session and Regularity Type (ps ¿ .5). 5-year-olds showed a significant effect of Session ($\chi^2(1) = 8.88$, $p = .003$) that did not interact with Regularity Type

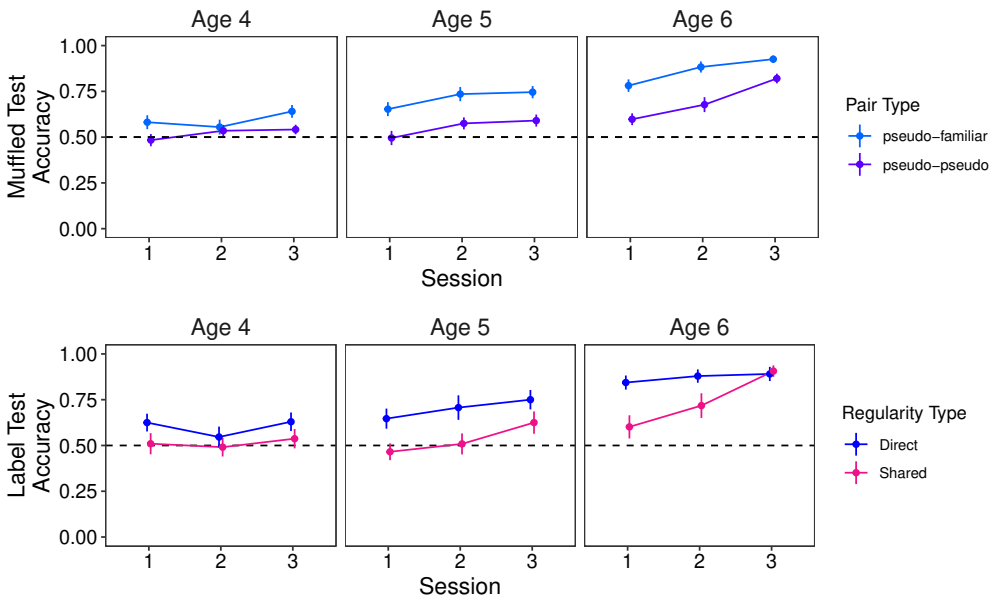

Figure 4: The formation of connections across age and sessions measured from the Muffled (top) and Label (bottom) Tests. The Muffled Test measured connections from the direct co-occurrence of a pseudoadjective with a familiar noun and pseudonoun. The Label Test measured connections based on a familiar noun's direct co-occurrence with the pseudoadjective and shared co-occurrence with the pseudonoun.

($\chi^2(1) = 0.15$, $p = .70$), indicating that connections based on both direct and shared co-occurrence improved across sessions. Finally, 6-year-olds showed a significant effect of session ($\chi^2(1) = 25.95$, $p < .0001$) that interacted with Regularity Type ($\chi^2(1) = 8.60$, $p = .003$), such that only connections based on shared co-occurrence improved across sessions. This is likely because connections based on direct co-occurrence were already very strong starting from the first session (see Figure 4 bottom). Thus, children of all ages more readily connected words that directly co-occurred than those that shared co-occurrence, but only for older children did these connections strengthen with repeated exposure.

Building Shared Co-Occurrence Connections

Prior single-session studies in multiple domains suggest that in addition to direct co-occurrence connections, forming shared co-occurrence connections also depends on some maturing ability to integrate across direct co-occurrence. To further test this possibility, we tested whether the strength of shared co-occurrence connections was pedicted by: direct co-occurrence connections between the pseudoadjective and familiar noun, direct co-occurrence conections between the pseudoadjective and pseudonoun, Session and Age. The strength of shared co-occurrence conections was predicted by all factors ($ps < .05$) except direct co-occurrence connections between the pseudoadjective and familiar noun ($p = .72$). Thus, shared co-occurrence connections may emerge from the conjunction of direct co-occurrence connections, repeated exposure, and developmental improvements in abilities to integrate across direct co-occurrence.

General Discussion

Children's language experiences are full of regularities in the way words co-occur with other words. In principle, accumulating exposures to these regularities could support the formation of the increasingly rich connections between words that emerge across childhood (Unger & Fisher, 2016; Crowe & Prescott, 2003; Storm, 1980). Alternatively, accumulating exposure might not be enough: instead, prior evidence points to developmental changes in children's abilities to transform co-occurrence regularities into interconnected knowledge (Savic et al., 2022). Yet this evidence comes from studies of learning from a single exposure. Any such developmental improvement poses a key question: are younger children really incapable of doing what older children can do, or do they just need more learning opportunities? We thus examined the development of abilities to build up connections from such repeated exposures.

From the first session, we replicated prior evidence that older children more readily form connections from exposure to regularities. Critically, we further found that older children also more successfully build increasingly rich connections from one exposure to the next. Indeed, although children in the youngest age group did form some connections starting from their first exposure, they showed no evidence of building upon these connections with repeated exposures. These findings suggest that the increasingly interconnected word knowledge that emerges in development comes in part from the interplay between repeated exposures to regularities of word use and developmental improvements in abilities to capitalize on these exposures.

Acknowledgments

This work was supported by National Institutes of Health Grants R01HD111458 awarded to LU and VS and R01HD078545 awarded to VS. The authors would like to thank the participating families and Madison Meares for facilitating the research.

References

Arciuli, J., & Simpson, I. C. (2011). Statistical learning in typically developing children: The role of age and speed of stimulus presentation. *Developmental Science*, *14*(3), 464-473.

Bauer, P. J., Cronin-Golomb, L. M., Porter, B. M., Jaganjac, A., & Miller, H. E. (2020). Integration of memory content in adults and children: Developmental differences in task conditions and functional consequences. *Journal of Experimental Psychology: General*, *150*, 1259–1278.

Bauer, P. J., & San Souci, P. (2010). Going beyond the facts: Young children extend knowledge by integrating episodes. *Journal of Experimental Child Psychology*, *107*(4), 452-465.

Bolger, D. J., Balass, M., Landen, E., & Perfetti, C. A. (2008). Context variation and definitions in learning the meanings of words: An instance-based learning approach. *Discourse Processes*, *45*(2), 122–159.

Borovsky, A., Ellis, E. M., Evans, J. L., & Elman, J. L. (2016). Lexical leverage: Category knowledge boosts real-time novel word recognition in 2-year-olds. *Developmental Science*, *19*(6), 918–932.

Carey, S., & Bartlett, E. (1978). Acquiring a single new word. *Proceedings of the Stanford Child Language Conference*, *15*, 17–29.

Crowe, S. J., & Prescott, T. J. (2003). Continuity and change in the development of category structure: Insights from the semantic fluency task. *International Journal of Behavioral Development*, *27*(5), 467-479.

Fourtassi, A. (2020). Word co-occurrence in child-directed speech predicts children's free word associations. In *Proceedings of the workshop on cognitive modeling and computational linguistics* (p. 49-53).

Hills, T. T., Maouene, J., Riordan, B., & Smith, L. B. (2010). The associative structure of language: Contextual diversity in early word learning. *Journal of Memory and Language*, *63*(3), 259-273.

Huebner, P. A., & Willits, J. A. (2018). Structured semantic knowledge can emerge automatically from predicting word sequences in child-directed speech. *Frontiers in Psychology*, *9*.

Jones, M. N., & Mewhort, D. J. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, *114*(1), 1-37.

Landauer, T. K., & Dumais, S. T. (1997). A solution to plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*(2), 211-240.

MacWhinney, B. (2000). *The CHILDES project: The database* (Vol. 2). Psychology Press.

McMurray, B., Horst, J. S., & Samuelson, L. K. (2012). Word learning emerges from the interaction of online referent selection and slow associative learning. *Psychological Review*, *119*(4), 831.

McNeill, D. (1963). The origin of associations within the same grammatical class. *Journal of Verbal Learning and Verbal Behavior*, *2*(3), 250-262.

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

Nation, K., & Snowling, M. (1997). Assessing reading difficulties: The validity and utility of current measures of reading skill. *British Journal of Educational Psychology*, *67*(3), 359-370.

Neuman, S. B., Newman, E. H., & Dwyer, J. (2011). Educational effects of a vocabulary intervention on preschoolers' word knowledge and conceptual development: A cluster-randomized trial. *Reading Research Quarterly*, *46*(3), 249-272.

Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (emnlp)* (p. 1532-1543).

Raviv, L., & Arnon, I. (2018). The developmental trajectory of children's auditory and visual statistical learning abilities: Modality-based differences in the effect of age. *Developmental Science*, *21*(4), e12593.

Romberg, A. R., & Saffran, J. R. (2010). Statistical learning and language acquisition. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*(6), 906-914.

Roy, B. C., Frank, M. C., DeCamp, P., Miller, M., & Roy, D. (2015). Predicting the birth of a spoken word. *Proceedings of the National Academy of Sciences*, *112*(41), 12663–12668.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*(5294), 1926-1928.

Savic, O., Unger, L., & Sloutsky, V. M. (2022). Experience and maturation: The contribution of co-occurrence regularities in language to the development of semantic organization. *Child Development*, Available from: https://doi.org/10.31234/osf.io/5tpf2.

Schlichting, M. L., Guarino, K. F., Roome, H. E., & Preston, A. R. (2021). Developmental differences in memory reactivation relate to encoding and inference in the human brain. *Nature Human Behaviour*, 1–14.

Schlichting, M. L., Guarino, K. F., Schapiro, A. C., Turk-Browne, N. B., & Preston, A. R. (2017). Hippocampal structure predicts statistical learning and associative inference abilities during development. *Journal of Cognitive Neuroscience*, *29*(1), 37-51.

Shufaniya, A., & Arnon, I. (2018). Statistical learning is not age-invariant during childhood: Performance improves

with age across modality. *Cognitive Science*, *42*(8), 3100-3115.

Storm, C. (1980). The semantic structure of animal terms: A developmental study. *International Journal of Behavioral Development*, *3*(4), 381-407.

Unger, L., & Fisher, A. V. (2016). Semantic contamination of visual similarity judgments [Book Section]. In *Proceedings of the 38th annual conference of the cognitive science society.* Philadelphia, PA.

Unger, L., & Fisher, A. V. (2019). Rapid, experience-related changes in the organization of children's semantic knowledge. *Journal of Experimental Child Psychology*, *179*, 1-22.