

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Simple Mechanisms, Rich Structure: Statistical Co-Occurrence Regularities in Language Shape the Development of Semantic Knowledge

Permalink

<https://escholarship.org/uc/item/8s86m6xw>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 42(0)

Authors

Unger, Layla
Savic, Olivera
Sloutsky, Vladimir

Publication Date

2020

Copyright Information

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

Peer reviewed

Simple Mechanisms, Rich Structure: Statistical Co-Occurrence Regularities in Language Shape the Development of Semantic Knowledge

Layla Unger (unger.114@osu.edu)

Olivera Savic (savic.7@osu.edu)

Vladimir M. Sloutsky (sloutsky@psy.ohio-state.edu)

Ohio State University, Department of Psychology, 1835 Neil Avenue
Columbus, OH 43210 USA

Abstract

Many hallmarks of human intelligence including language, reasoning, and planning require us to draw upon knowledge about the world in which concepts, denoted by words, are organized by meaningful, *semantic* links between them (e.g., juicy-apple-pear). The goal of the present research was to investigate how these organized semantic networks may emerge in development from simple but powerful mechanisms sensitive to statistical co-occurrence regularities of word use in language. Specifically, we tested whether a mechanistic account of how co-occurrence regularities shape semantic development accurately predicts how semantic organization changes with development. Using a sensitive, gaze-based measure of the semantic links organizing knowledge in children and adults, we observed that developmental changes in semantic organization were consistent with a key role for statistical co-occurrence regularities.

Keywords: semantic organization; semantic development; statistical learning; taxonomic; association

Introduction

We rely on our knowledge about the world to achieve many vital, every-day cognitive tasks. For example, our knowledge of apples can allow us to use language to express and comprehend ideas about eating apples, retrieve knowledge from memory that apples are healthy to achieve a goal to eat a healthy snack, plan for lunch by packing an apple, and generate new ideas such as “pears are healthy” by generalizing what we know about apples to other fruits. These feats rely on knowledge that is not a jumble of facts, but instead an organized semantic network of linked concepts, such as apples, pears, eating, and healthy. How does such vital semantic organization emerge and change in the course of development?



Figure 1: Direct and shared co-occurrence regularities that can form associative and taxonomic links.

The many prior accounts of knowledge organization development have focused on how we form taxonomic links between members of the same, stable category, such as links between pigeon and duck that belong to the category of birds (Gelman & Markman, 1986; Inhelder & Piaget, 1964; Lucariello, Kyratzis, & Nelson, 1992; Sloutsky, 2010). Research into taxonomic link development may provide valuable insights into the development of semantic organization. However, semantic organization is much richer than just these links, encompassing a variety of taxonomic and non-taxonomic links between a multitude of concepts. The goal of the present research is to evaluate how this richer semantic structure may be driven in part by simple but powerful mechanisms that form semantic links based on statistical co-occurrence regularities of word use in language.

Co-Occurrence in Semantic Development

The potentially fundamental roles for co-occurrence that we will investigate are outlined in the recently proposed Co-Occurrence Account (Sloutsky, Yim, Yao, & Dennis, 2017). According to this account, sensitivity to co-occurrence initially fosters *associative* links between concepts whose labels reliably occur close together in language (adjacent or separated by intervening words), such as juicy-apple. Henceforth, these regularities are referred to as *direct co-occurrence*. These associative links form a key non-taxonomic facet of semantic networks that supports knowledge-dependent cognition from early development onward. For example, upon hearing a new word such as *dax* accompanied by words associated with animal such as *furry*, both young children and adults infer that *dax* means *animal* (Sloutsky et al., 2017).

Critically, sensitivity to co-occurrence can also foster taxonomic links, because words for members of taxonomic categories (e.g., “apple” and “pear”) reliably *share* overlapping patterns of direct co-occurrence with other words (e.g., “juicy”, Jones, Willits, & Dennis, 2015). However, unlike direct co-occurrence, shared co-occurrence cannot be immediately gleaned from language input. For example, to form a shared co-occurrence-based link between apple and pear as shown in Figure 1, the learner must

experience direct co-occurrences between both “juicy” and “apple”, and “juicy” and “pear”, then form a link between “apple” and “pear” based on their overlapping direct co-occurrence with “juicy”¹. Therefore, taxonomic links may develop more gradually (Bauer & Larkina, 2017; Schlichting, Guarino, Schapiro, Turk-Browne, & Preston, 2017).

Together, the learning processes proposed in the Co-Occurrence account outline how simple co-occurrence regularities may build both associative and taxonomic links between any concepts denoted by words. Moreover, this account makes specific predictions about how these semantic links emerge in the course of development. In what follows, we describe evidence from prior research supporting this account, then present an experiment designed to evaluate its predictions about the development of semantic organization.

Support for the Co-Occurrence Account. The Co-occurrence account is motivated by extensive evidence that linguistic input is rich in regularities from which semantic links can be formed. First, much of the variance in the strength of semantic links in adult semantic networks can be predicted by regularities with which words directly co-occur or share co-occurrence in language (Hofmann, Biemann, Westbury et al., 2018; Spence & Owens, 1990). Moreover, computational models that form word representations based on co-occurrence statistics in language simulate semantic networks that predict complex semantic phenomena, from semantic priming effects to the typical vocabulary growth rate of school children (Jones et al., 2015; Landauer & Dumais, 1997; Sahlgren, 2008). Together, these findings ground the Co-occurrence account’s proposal that co-occurrence regularities are important drivers of semantic organization development.

Importantly, a key proposal of the Co-Occurrence account is that *developing humans* form semantic links from co-occurrence regularities in language. Extensive evidence supports the possibility that humans form links based on *direct co-occurrence* starting early in development. Much of this evidence comes from statistical learning research, which has shown early-developing abilities to form direct co-occurrence-based links between stimuli such as speech sounds (Saffran, Aslin, & Newport, 1996) and images (Fiser & Aslin, 2002). Recently, the formation of direct co-occurrence-based links between words has also been observed in toddlers (Wojcik & Saffran, 2015) and young children (Matlen, Fisher, & Godwin, 2015).

Evidence that humans form shared co-occurrence-based links comes from a smaller body of research. For example, adults in Preston, Zeithamova and colleagues’ studies (Zeithamova, Dominick, & Preston, 2012) who explicitly memorize pairs of images also link images that were never paired, but instead share each other’s pairing with the same image (see also Hall, Mitchell, Graham, & Lavis, 2003;

Schapiro, Rogers, Cordova, Turk-Browne, & Botvinick, 2013). Moreover, this ability may develop only gradually. For example, Schlichting et al. (2017) observed that the ability to link images based on their shared pairing with another image improved substantially from age six to adulthood. Similar evidence for gradual development comes from studies conducted by Bauer and colleagues, in which participants were given two *stem facts* that both link information to a shared concept, such as “dolphins talk by clicking and squeaking” and “dolphins live in groups called pods”. The ability to integrate across stem facts to derive a new fact such as “pods talk by clicking and squeaking” is poor at age four, and substantially increases over childhood.

This prior evidence motivates and supports Co-occurrence account’s proposals about how sensitivities to direct and shared co-occurrence regularities may build semantic organization. However, evaluating the Co-occurrence account critically involves testing the specific predictions it makes about the development of semantic organization. Specifically, the Co-Occurrence account predicts that associative links between concepts that can form from direct co-occurrence regularities should emerge early in development. With development, associative links should become gradually supplemented by taxonomic links that can be formed from shared co-occurrence regularities. The goal of the present experiment was to evaluate these predictions.

Present Experiment

The present experiment tested the Co-Occurrence account’s prediction that associative links that can be learned from direct co-occurrence develop early, and are gradually supplemented by taxonomic links that can be learned from shared co-occurrence. We therefore measured the development of associative and taxonomic links between familiar concepts from early childhood (4-year-old children) to adulthood. To target associative and taxonomic links, we measured the strength of semantic links between “associated” concepts whose labels regularly co-occur in child language input (MacWhinney, 2000), and “taxonomically related” concepts similar in meaning (“About wordnet,” 2010).

We designed our measurement of semantic links to fulfill two important criteria. First, instead of merely measuring whether a certain type of semantic link is present or absent in a given age group, we acquired *fine grained, sensitive measures of the strength of semantic links*. Such sensitive measures are important for tracking the *gradual*, developmental emergence of semantic links. Second, we designed our measurement to primarily capture developmental changes in semantic links, rather than in other cognitive processes such as reasoning.

To fulfill these criteria, we used a Visual World paradigm. This paradigm capitalizes on the fact that people tend to look

¹ The *mechanism(s)* that form links between inputs that share patterns of co-occurrence remain unknown. Candidates have been proposed and investigated in multiple fields, including conditioning (Honey & Hall, 1989), hippocampal memory formation (Schapiro, Turk-Browne, Botvinick, & Norman, 2017), and semantic

organization (McNeill, 1963). For the present research, the key point is that direct co-occurrence can be directly experienced from language input, whereas shared co-occurrence-based links can only be derived by integrating across separate episodes of direct co-occurrence.

at images that they perceive as related to language that they hear. Therefore, we can measure the strength of semantic links from the degree to which hearing a label for one concept prompts looking at an image of another concept *over time*. This measure is both fine-grained, and because it is based on spontaneous looking behavior, should be relatively uncontaminated by other cognitive processes.

In our Visual World paradigm, participants saw a pair of unrelated Target pictures (e.g., bed and fish), and heard either: (1) An Associate Prime for one of the Targets (e.g., pillow or water), (2) A Taxonomic Prime for one of the Targets (e.g., chair or bird), or (3) An Unrelated prime that was neither associated with nor taxonomically related to either Target (e.g., stick). We measured the strength of associative and taxonomic links based on the degree to which participants looked more at Targets *over time* following Associate or Taxonomic versus Unrelated Primes. Importantly, our use of the same pairs of Target pictures in all Prime conditions meant that looking differences across Prime conditions can be attributed to semantic links between Prime and Target concepts, rather than the salience, visual properties, or subjective appeal Targets themselves.

Method

Participants

Informed consent was obtained from parents/guardians of child participants and from adult participants prior to participation. The sample included 41 4-year-olds and 37 adults. Children were recruited from families, daycares, and preschools and adults were recruited from the undergraduate population at a public university in the same city.

Stimuli

The stimuli for this experiment were Sets of words consisting of a Target, an Associate Prime, and a Taxonomic Prime generated according to the following criteria.

Associate criteria. Associate Primes were selected as words that reliably co-occur with Targets in corpora of child speech input (CHILDES database; MacWhinney, 2000). Using scripts developed in-lab, we measured the degree to which word pairs co-occurred more frequently within a 7-word window across 25 CHILDES corpora (O) than the frequency with which they would be expected to co-occur by chance, based on their respective frequencies (E). The larger the difference between observed versus expected frequency, the more reliably words in a pair co-occur. This ratio is captured by the following “t.score” formula:

$$t.score = \frac{O - E}{\sqrt{O}}$$

Candidate Target-Associate pairs were pairs of nouns with t-scores of > 2.5 (Baayen, Davidson, & Bates, 2008). In addition, Associates *could not* meet the taxonomic criteria described below.

Taxonomic criteria. Candidate Taxonomic Primes for Targets were identified based on their membership in the

same taxonomic category (e.g., clothing, foods) and similarity in meaning in WordNet (a database of word definitions composed by lexicographers; 2010). In WordNet, words are hierarchically organized such that more specific words (e.g., dog) are subsumed within less specific words (e.g., animal). We used Resnik similarity as a measure of meaning similarity in WordNet, which is based on identifying the most specific subsumer of a pair of words: The more specific the subsumer, the higher the similarity. For example, dog and cat are subsumed within carnivore, whereas dog and rat are subsumed within mammal; because carnivore is more specific than mammal, Resnik similarity is higher between dog and cat versus dog and rat. Candidate Taxonomic Primes had Resnik similarities to Targets of > 5, and did not meet the Associate criterion.

Composition of Set Pairs. We used the Associate and Taxonomic criteria to compose 22 Sets each consisting of a Target, Associate, and a Taxonomic Prime. Importantly, Targets were neither taxonomically related to Associate Primes, nor associated with Taxonomic Primes. All words also met a familiarity criterion of being produced by at least 55% of 36-month-old children (approximately one year younger than children in our youngest sample). Words within Sets were balanced for this criterion.

We organized these 22 Sets into 11 “Set Pairs”. Within Set Pairs: (1) Targets were unrelated and equivalently familiar, and (2) Primes for one Target were unrelated to the other Target. To each Set Pair we added an Unrelated Prime that: (1) Met the familiarity criterion, and (2) Met *neither* Associate *nor* the Taxonomic criteria for both Targets.

Materials. All words in Set Pairs were recorded by a female speaker using child-friendly speech. Targets were presented as pictures subtending ~5.3° of visual angle.

Apparatus

This experiment used an EyeLink Portable Duo eye tracking system with a sampling rate of 500Hz, and a button box that participants used in a cover task (see Procedure).

Procedure

Adults were tested in a quiet lab room, and children were tested either in a quiet lab room, or at their preschool or daycare. The procedure was similar for adults and children, with the exception that children completed one block of trials, and adults completed two blocks (i.e., repeated the same block twice with randomized trial orders).




Figure 2: Sequence of events in experimental trials.

Following eye tracker calibration, the experiment consisted of trials in which participants were shown the two Target pictures from a Set Pair (e.g. apple and bottle), one presented on the left side of the screen, on a yellow background, and one on the right side of the screen, on a blue background. As shown in Figure 2, the two Targets appeared alone for 500ms, and then participants heard a word.

Participants first completed familiarization trials, then the main experiment, which included a mix of cover task and experimental trials. In familiarization and cover task trials, participants heard “yellow” or “blue”, and clicked a button of the same color on a button box to complete the trial. These trials were designed to keep participants engaged in a task on non-experimental trials.

Experimental trials were similar to familiarization/cover task trials. However, instead of “yellow”/“blue”, participants heard Primes from the Set Pairs. Participants were instructed not to respond to these trials, which instead ended automatically 2000ms following Prime onset.

Across trials, each pair of Target pictures (e.g., apple and bottle) was presented with the five Primes from their Set Pair: (1) The Associate Prime for one of the two Targets (e.g., tree or baby), (2) The Taxonomic Prime for one of the two Targets (e.g., grapes or bowl), or (3) The Unrelated Prime (e.g., door). Thus, there was a total of 55 experimental trials within a block: 22 Associate, 22 Taxonomic and 11 Unrelated. These trials were mixed with 22 cover task trials (one “yellow” and one “blue” trial for each Set Pair). The assignment of Target pictures in each Set Pair to appear on the yellow background on the left or blue background on the right was counterbalanced across experimental and cover task trials. This design ensured that looking on experimental trials was not contaminated by response-related behavior.

Results

To test the contributions of association and taxonomic relatedness, the data from this experiment were used to compare the time course of looking at Targets accompanied by Associate or Taxonomic Primes versus Unrelated Primes in each age group. To conduct this comparison, we first generated outcome variables of interest.

Outcome Variables

Data from practice and filler trials were excluded from analyses. The raw eye tracking data consisted of the position of gaze on the screen sampled every 2ms within experimental trials, which was identified as falling within an AOI for the image on the left, an AOI for the image on the right, or neither AOI. We removed data from the 500ms prior to Prime onset, then divided the remaining two seconds into 100ms time bins. We used these data to generate two outcome variables.

Target Dwell Time. We first calculated a Target Dwell Time variable that captures the amount of time spent looking at each Target in each time bin when accompanied by its Associate, Taxonomic, or Unrelated Prime. We used this variable to test whether looking dynamics for Targets

differed when accompanied by their Associate or Taxonomic versus Unrelated Primes.

Difference from Unrelated. This variable captured the degree to which looking in the Associate and Taxonomic Prime conditions each deviated from the Unrelated Prime condition. We calculated this value by subtracting the Unrelated Target Dwell Time for a Target/time bin from both the corresponding Target Dwell Time in the Associate condition, and the in the Taxonomic condition. We used this variable to test for differences between the effects of Associate versus Taxonomic (relative to Unrelated) Primes.

Analysis of Looking Behavior

We followed the Growth Curve Analysis (GCA) approach developed by Mirman and colleagues (Mirman, Dixon, & Magnuson, 2008). GCA involves generating hierarchical mixed effects models, starting with a “base” model with temporal terms that captures how looking changes over time, without considering variation across conditions, individuals, or items. In the base model, the intercept captures the average value of the looking outcome variable, a linear term captures monotonic changes in the value of the outcome variable over time, and a quadratic term captures the sharpness of the peak in looking. Finally, cubic and quartic terms capture changes in asymptotic tails of looking over time that are not typically informative about the effects of experimental conditions.

To analyze the effects of experimental conditions, the base model is supplemented with: Fixed effects of experimental conditions and their interaction with the temporal terms, random intercepts for participants and/or items, and random slopes for effects of experimental conditions within participants and/or items. Effects of experimental conditions are interpreted from their interactions with temporal terms. For example, an interaction between a fixed effect of condition and the linear term reveals that condition influences the monotonic increase or decrease in looking over time.

Target Dwell Time Analysis. We first tested whether the temporal dynamics of looking at Targets differed when accompanied by Associate or Taxonomic Primes in comparison to when accompanied by Unrelated Primes. Specifically, we generated separate models of Dwell Times for Targets in each time bin for each age group that both supplemented the base model with a fixed effect of Prime condition (with Unrelated as the reference level to which Associate and Taxonomic were compared). These models

Table 1: Target Dwell Time GCA results. Estimates are relative to the Unrelated condition. Non-significant estimates are in *italics*.

Term	Age	Associate	Taxonomic
		Est. (SE)	Est. (SE)
Intercept	Child	9.599 (1.852)	6.389 (1.852)
Linear	Child	31.351 (5.603)	13.229 (5.603)
Quadratic	Child	<i>-4.702 (5.057)</i>	<i>-9.130 (5.057)</i>
Intercept	Adult	8.986 (2.605)	6.768 (2.605)
Linear	Adult	21.462 (7.149)	14.384 (7.149)
Quadratic	Adult	<i>-23.151 (5.097)</i>	<i>-18.836 (5.097)</i>

additionally included random intercepts for participant and item, and random slopes for the effect of Prime condition within participants and within items.

Parameter estimates and their significance are reported in Table 1. Both children and adults looked more overall at Targets upon hearing an Associate or a Taxonomic versus an Unrelated Prime (as shown by significant effects on the Intercept). Associate and Taxonomic Primes also affected changes in looking at a given Target over time, including the rate at which looking at the Target increased (Linear term) and/or the sharpness of the peak in Target looking time (Quadratic term). Taken together, these results show that concepts depicted by Targets were activated by both Co-Occur and Taxonomic Primes in both adults and children.




Figure 3: Difference from Unrelated values in the Associate and Taxonomic conditions in Children and Adults, plotted with lines depicting the fitted values from the models. Error bars show standard errors of the mean.

Difference from Unrelated. This analysis assessed differences in degree to which Associate and Taxonomic Primes activated Targets, relative to Unrelated Primes. Specifically, we generated separate models of Difference from Unrelated values for children and adults that supplemented the base model with a fixed effect of Relatedness condition (Associate and Taxonomic), random intercepts for participant and item, and random slopes for the effect of Relatedness condition within participants and within items. Figure 3 depicts the Difference from Unrelated data and the corresponding fitted data from the models.

The parameter estimates and their significance levels are reported in Table 2. In children, Associate Primes produced greater rates of increased looking at Targets (relative to Unrelated Primes) than Taxonomic Primes. In contrast, in adults, no such differences were observed: Associate and

Taxonomic Primes affected looking at Targets relative to Unrelated Primes to equivalent extents.

Discussion

The present experiment revealed that, in young children, associative links between concepts whose labels reliably co-occur are initially stronger than taxonomic links between concepts whose labels often share patterns of co-occurrence. By adulthood however, associative and taxonomic links were similar in strength. This trajectory is consistent with the Co-Occurrence account prediction that associative links emerge early, and are supplemented by taxonomic links with development. These results therefore highlight how rich semantic structure may emerge from simple but powerful sensitivities to co-occurrence statistics.

However, the trajectory of semantic organization development cannot be inferred from the present experiment alone. To contextualize these findings, we next evaluate the degree to which this developmental trajectory is consistent with evidence from prior research on semantic development. In this evaluation, we highlight how the present findings are both consistent with, and expand upon much of the large body of prior semantic development research.

Contribution of Co-Occurrence

Although a role for co-occurrence throughout semantic organization development has been overlooked (or posited to be transient) in the majority existing semantic development accounts, the present evidence supporting this role is consistent with many prior findings. Specifically, numerous studies with children (e.g., Blaye, Bernard-Peyron, Paour, & Bonthoux, 2006; Lucariello et al., 1992) and a handful of studies with adults (e.g., Lin & Murphy, 2001) have observed the presence of links in semantic organization that may be learned from co-occurrence, such as *schematic* and *thematic* relatedness. Moreover, in contrast with schematic and thematic relatedness, which are constructs subjectively defined by researchers, the present findings highlight co-occurrence regularities as a measurable source of input in the environment that may shape these semantic links.

Contribution of Taxonomic Relations

This experiment revealed an influence of taxonomic relatedness that was initially weaker than the influence of co-occurrence, but reached similar strength by adulthood. Contextualizing this finding within prior research is complicated by the fact that it has yielded conflicting findings. One body of findings suggest that taxonomic relations only gradually emerge starting from mid-to-late childhood (e.g., age 6-7, Blaye et al., 2006; Lucariello et al., 1992). In contrast, a similarly large body of findings suggests that taxonomic relations are strong and robust starting early in development (Gelman & Markman, 1986; Waxman & Namy, 1997). In spite of these apparently contradictory bodies of evidence, we suggest that the present findings can be reconciled and shed new light on both.

Table 2: Difference from Unrelated GCA results. Estimates are for the Associate versus the Taxonomic condition. Non-significant estimates are in italics.

Associate vs Taxonomic		
Term	Age	Est. (SE)
Intercept	Child	3.210 (1.970)
Linear	Child	18.122 (5.938)
Quadratic	Child	4.428 (5.094)
Intercept	Adult	2.217 (2.493)
Linear	Adult	7.078 (7.541)
Quadratic	Adult	-4.315 (5.169)

Gradual Taxonomic Development. Numerous studies using a variety of behavioral paradigms have found that taxonomic relations only begin to gradually contribute to semantic organization with starting in mid-to-late childhood (e.g., Blaye et al., 2006; Lucariello et al., 1992). The present findings are consistent with evidence for the gradual development of taxonomic relations, and suggest that sensitive measures (such as those used in the present experiment) can capture this gradual development starting earlier in childhood.

Early Taxonomic Onset. Another large body of findings suggests early, robust taxonomic organization. Many such studies have assessed semantic organization using match-to-sample paradigms, in which participants choose to match a sample item (e.g., dog) with one of two other items (e.g., elephant and bone) that are related to the sample in different ways. In some studies using variants of this paradigm (e.g., Bauer & Mandler, 1989; Gelman & Markman, 1986; Waxman & Namy, 1997), young children chose taxonomic matches throughout the study or under specific conditions.

Our findings also suggest the presence of taxonomic relations in young children. Only the notion that these prior findings indicate *robust* taxonomic knowledge starting in early childhood conflicts with the present evidence that taxonomic relations are initially weak. However, this contradiction can be resolved by considering how additional information that could support taxonomic choices was available in prior studies showing “robust” taxonomic relations in young children. For example, in some prior studies, many target items are likely to have been visually similar to (e.g., car and jeep, pot and skillet) and/or co-occurring with (e.g., chair and table) their taxonomic matches. Moreover, targets and taxonomic matches were sometimes given either identical labels, which may act as perceptual features that contribute to similarity in young children (Sloutsky & Fisher, 2004), or co-occurring labels (e.g., puppy and dog), such that taxonomic choices could be based on co-occurrence. The availability of co-occurrence and/or perceptual similarity in addition to taxonomic relatedness also characterizes stimuli used in many studies of semantic knowledge in infants (e.g., Willits, Wojcik, Seidenberg, & Saffran, 2013).

Asynchronous Development of Associative and Taxonomic Relations

The Co-Occurrence account predicts that associative links emerge early because they can be formed from direct co-occurrence regularities that can be directly gleaned from language input. For example, hearing “I’d like a juicy apple” can immediately contribute to a semantic link between “juicy” and “apple”. By the same token, the Co-occurrence account predicts that taxonomic links emerge later because they may rely on shared co-occurrence regularities that can only be *derived* by integrating across separate episodes of direct co-occurrence. However, the Co-Occurrence account does not currently specify a precise reason for why associative links based on direct co-occurrence may form

earlier in development than taxonomic links based on shared co-occurrence. Instead, the present evidence for the Co-Occurrence account highlights explanations for this developmental asynchrony to be explored in future research.

One possibility is that taxonomic links develop more slowly simply because they require more language input: Hearing “juicy” and “apple” directly co-occurring can immediately contribute to an associative link, whereas the learner must separately hear “juicy” with both “apple” and “pear” to form a shared co-occurrence-based taxonomic link. Alternatively, abilities to form links between inputs based on direct and shared co-occurrence statistics may themselves develop asynchronously. This possibility is supported by the contrast between extensive statistical learning evidence that even infants can form links between inputs that directly co-occur (Fiser & Aslin, 2002; Saffran et al., 1996; Saffran, Johnson, Aslin, & Newport, 1999), and a handful of evidence for the more gradual development of abilities to form links based on shared co-occurrence (Bauer & San Souci, 2010; Schlichting et al., 2017). Disentangling these possibilities can shed light on the underlying processes that drive developmental changes in semantic organization.

Conclusions

Organized semantic knowledge plays a fundamental role in many facets of human intelligence. The present experiment provides evidence supporting the possibility that this organization emerges in part from the operation of simple but powerful learning mechanisms that form semantic links from statistical regularities in language.

References

- Princeton University. (2010). About wordnet. <http://wordnet.princeton.edu>
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59, 390-412.
- Bauer, P. J., & Larkina, M. (2017). Realizing relevance: The influence of domain-specific information on generation of new knowledge through integration in 4-to 8-year-old children. *Child Development*, 88, 247-262.
- Bauer, P. J., & Mandler, J. M. (1989). Taxonomies and triads: Conceptual organization in one-to two-year-olds. *Cognitive Psychology*, 21, 156-184.
- Bauer, P. J., & San Souci, P. (2010). Going beyond the facts: Young children extend knowledge by integrating episodes. *Journal of Experimental Child Psychology*, 107, 452-465.
- Blaye, A., Bernard-Peyron, V., Paour, J.-L., & Bonthoux, F. (2006). Categorical flexibility in children: Distinguishing response flexibility from conceptual flexibility. *European Journal of Developmental Psychology*, 3, 163-188.
- Fiser, J., & Aslin, R. N. (2002). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences*, 99, 15822-15826.

- Gelman, S. A., & Markman, E. M. (1986). Categories and induction in young children. *Cognition*, 23, 183-209.
- Hall, G., Mitchell, C., Graham, S., & Lavis, Y. (2003). Acquired equivalence and distinctiveness in human discrimination learning: Evidence for associative mediation. *Journal of Experimental Psychology: General*, 132, 266.
- Hofmann, M. J., Biemann, C., Westbury, C., Murusidze, M., Conrad, M., & Jacobs, A. M. (2018). Simple co-occurrence statistics reproducibly predict association ratings. *Cognitive Science*, 42, 2287-2312.
- Honey, R. C., & Hall, G. (1989). Acquired equivalence and distinctiveness of cues. *Journal of Experimental Psychology: Animal Behavior Processes*, 15, 338.
- Inhelder, B., & Piaget, J. (1964). *The early growth of logic in the child*. New York: Norton.
- Jones, M. N., Willits, J., & Dennis, S. (2015). Models of semantic memory. In J. Busemeyer & J. Townsend (Eds.), *Oxford handbook of mathematical and computational psychology* (pp. 232-254). New York, NY: Oxford University Press.
- 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, 211.
- Lin, E. L., & Murphy, G. L. (2001). Thematic relations in adults' concepts. *Journal of Experimental Psychology: General*, 130, 3-28.
- Lucariello, J., Kyratzis, A., & Nelson, K. (1992). Taxonomic knowledge: What kind and when? *Child Development*, 63, 978-998.
- MacWhinney, B. (2000). *The chldes project: The database* (Vol. 2): Psychology Press.
- Matlen, B. J., Fisher, A. V., & Godwin, K. E. (2015). The influence of label co-occurrence and semantic similarity on children's inductive generalization. *Frontiers in Psychology*, 6, 1146.
- McNeill, D. (1963). The origin of associations within the same grammatical class. *Journal of Verbal Learning and Verbal Behavior*, 2, 250-262.
- Mirman, D., Dixon, J. A., & Magnuson, J. S. (2008). Statistical and computational models of the visual world paradigm: Growth curves and individual differences. *Journal of Memory and Language*, 59, 475-494.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926-1928.
- Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, 70, 27-52.
- Sahlgren, M. (2008). The distributional hypothesis. *Italian Journal of Disability Studies*, 20, 33-53.
- Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., & Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure. *Nature Neuroscience*, 16, 486-492.
- Schapiro, A. C., Turk-Browne, N. B., Botvinick, M. M., & Norman, K. A. (2017). Complementary learning systems within the hippocampus: A neural network modelling approach to reconciling episodic memory with statistical learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372, 20160049.
- 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, 37-51.
- Sloutsky, V. M. (2010). From perceptual categories to concepts: What develops? *Cognitive Science*, 34, 1244-1286.
- Sloutsky, V. M., & Fisher, A. V. (2004). Induction and categorization in young children: A similarity-based model. *Journal of Experimental Psychology: General*, 133, 166-187.
- Sloutsky, V. M., Yim, H., Yao, X., & Dennis, S. (2017). An associative account of the development of word learning. *Cognitive Psychology*, 97, 1-30.
- Spence, D. P., & Owens, K. C. (1990). Lexical co-occurrence and association strength. *Journal of Psycholinguistic Research*, 19, 317-330.
- Waxman, S. R., & Namy, L. L. (1997). Challenging the notion of a thematic preference in young children. *Developmental Psychology*, 33, 555-567.
- Willits, J. A., Wojcik, E. H., Seidenberg, M. S., & Saffran, J. R. (2013). Toddlers activate lexical semantic knowledge in the absence of visual referents: Evidence from auditory priming. *Infancy*, 18, 1053-1075.
- Wojcik, E. H., & Saffran, J. R. (2015). Toddlers encode similarities among novel words from meaningful sentences. *Cognition*, 138, 10-20.
- Zeithamova, D., Dominick, A. L., & Preston, A. R. (2012). Hippocampal and ventral medial prefrontal activation during retrieval-mediated learning supports novel inference. *Neuron*, 75, 168-179.