

## **UC Merced**

### **Proceedings of the Annual Meeting of the Cognitive Science Society**

#### **Title**

Children's Acquisition of the Concept of Antonym Across Different Lexical Classes

#### **Permalink**

<https://escholarship.org/uc/item/9730v1k5>

#### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

#### **Authors**

Ionescu, Amalia

Lu, Hongjing

Holyoak, Keith

et al.

#### **Publication Date**

2022

Peer reviewed

# Children’s Acquisition of the Concept of Antonym Across Different Lexical Classes

Amalia Ionescu ([amaliaionescu@g.ucla.edu](mailto:amaliaionescu@g.ucla.edu))

Hongjing Lu ([hongjing@g.ucla.edu](mailto:hongjing@g.ucla.edu))

Keith J. Holyoak ([holyoak@lifesci.ucla.edu](mailto:holyoak@lifesci.ucla.edu))

Catherine M. Sandhofer ([sandhofer@psych.ucla.edu](mailto:sandhofer@psych.ucla.edu))

Department of Psychology, 502 Portola Plaza  
Los Angeles, CA 90095 USA

## Abstract

Understanding abstract relations, and reasoning about various instantiations of the same relation, is an important marker in human cognition. Here we focus on development of understanding for the concept of antonymy. We examined whether four- and five-year-olds ( $N=67$ ) are able to complete an analogy task involving antonyms, whether language cues facilitate children’s ability to reason about the antonym relation, and how their performance compares with that of two vector-based computational models. We found that explicit relation labels in the form of a relation phrase (“opposites”) improved performance on the task for five-year-olds but not four-year-olds. Five-year-old (but not four-year-old) children were more accurate for adjective and verb antonyms than for noun antonyms. Two computational models showed substantial variability in performance across different lexical classes, and in some cases fell short of children’s accuracy levels. These results suggest that young children acquire a solid understanding of the abstract relation of opposites, and can generalize it to various instantiations across different lexical classes. These developmental results challenge relation models based on vector semantics, and highlight the importance of examining performance across different parts of speech.

**Keywords:** analogy; relational learning; cognitive development; computational modeling.

## Introduction

Understanding abstract semantic relations between concepts expressed as words (e.g., synonym, antonym, category membership), and using them to reason by analogy, is a fundamental component of typical lexical and conceptual development. Antonyms are a unique semantic relation: they involve pairs of closely associated words yet differ maximally, typically on a single bipolar dimension (e.g., *hot: cold, long : short, rich : poor, love : hate*). Children are typically taught the antonym relation as a formal concept in elementary school, although some evidence suggests that they begin to form an understanding of antonyms much earlier (Phillips & Pexman, 2015). Because the acquisition of the antonym relation seems to reflect an important milestone in semantic development, investigating its origins and development trajectory can help elucidate how humans learn and represent abstract semantic relations.

Empirical research assessing children’s understanding of the concept of *opposite* reflects two general methodological approaches: discourse studies and metalinguistic studies. Discourse studies largely center on spontaneous usage of antonyms by children as young as age two (Tribushinina et al., 2013). There is evidence that young children’s production

of explicit contrasts, indicative of opposites (e.g., “this car is big and that one is small”), is strongly associated with parents’ production of explicit contrasts. Reasoning about contrasts may facilitate attention to the various dimensions on which antonyms can be evaluated.

Metalinguistic studies have evaluated children’s ability to understand and use metalinguistic vocabulary related to the concept of opposition. Paradigms primarily involve verbal games in which children respond to questions such as, “What is the opposite of X?” Other studies of this type have used free association tasks. Such studies have found that prior to five years of age, children tend to respond with a word that is closely associated with the stimulus word (e.g., *dark-night*), whereas older children tend to respond with a word that is semantically opposite to the stimulus (e.g., *dark-light*) (Entwistle, Forsyth, & Muuss, 1964).

The verbal emphasis in metalinguistic studies might explain why the findings suggest that the antonym relation becomes salient to children only around five years of age. More recent studies that have reduced the verbal component in antonym relation tasks have found that children have an understanding of the *opposite* relation somewhat earlier, around four years of age. For example, using a non-verbal task, Phillips and Pexman (2015) found that labeling the objects shown, as well as providing a label for the opposite relation, helped four- and five-year-old children identify the antonym match of various adjectives.

## Language as a facilitator of relational reasoning

Past research suggests that providing relational labels (e.g., “in front of”) helps children notice and manipulate relational similarities, comparable to how labels help children learn categories (Rattermann & Gentner, 1998). For example, in a mapping task in which children were shown the hidden location of an object in one situation and then had to find the hidden location of a second object in another nearly identical situation, only children who were provided with a label (language condition) were successful in finding the object (Loewenstein & Gentner, 1998). It seems likely that language could also facilitate children’s ability to reason about the antonym relation, perhaps by making the relation more salient.

It therefore seems plausible that providing a label and using relational language might support children’s understanding of semantic relations. For example, although four-to-five-year-olds might not have learned the words “opposite” or

“antonym,” providing a label to represent the relation may make it easier for children to identify it.

### Variability across parts of speech

When considering the acquisition of antonym understanding, it is important to examine possible variability across different lexical classes. Most studies on antonym relation learning have focused on adjective pairs (e.g., *big* : *small*); however, nouns dominate children’s early lexicons compared to verbs and adjectives (Gentner, 1978, Nelson, 1973, Sandhofer & Smith, 2007; Phillips & Pexman, 2015). These findings raise the possibility of similar variability in how children are able to reason about antonyms based on different parts of speech. For example, perhaps children may show earlier success with noun pairs instantiating antonym relation (e.g., *king* : *queen*). However, though nouns are learned earlier than adjectives, nouns are semantically richer because they can hold multiple meanings and share more than one relation with other words, which could make it more difficult for young children to evaluate nouns as compared to adjectives. For example, to generate the opposite of “short” (“tall”), one evaluates the concepts on a single dimension of length (height); however, to generate the opposite of “king,” one could produce “queen” if evaluating based on gender, or “peasant” if evaluating based on economic status.

It is therefore possible that reasoning about antonyms across various parts of speech may not follow the same developmental pattern as acquisition of individual words. Studies assessing performance of computational models of verbal analogy (e.g., Mikolov et al., 2013; Lu, Wu, & Holyoak, 2019) have compared different semantic relations, but not performance across different lexical classes within a single semantic relation of interest. Accordingly, one of the goals of the current study is to examine differences in both human and model performance across three parts of speech: adjectives, nouns, and verbs.

Although there is evidence that young children are able to identify pairs of words that fit the antonym relation, we do not yet know at what age they are able to solve analogy problems using pairs of antonyms. Solving analogies involves evaluating *pairs* of antonymous words based on different dimensions; therefore, examining whether young children can solve such problems can help assess their ability to reason about different instantiations of the same abstract relation.

The current work focuses on addressing the gaps in previous research on antonym learning in children, while also using two computational methods to further investigate how reasoning with antonyms varies across parts of speech. The first part of the current study focuses on children aged four-five years. This age range is particularly important, as previous research has demonstrated that children as young as age three can begin to solve analogies using complex relations (Shao & Gentner, 2018), as well as successfully transfer what they learn in one situation to analogous problems (Brown, Kane, & Long, 1989; Holyoak, Junn, & Billman, 1984).

Because substantial development in language occurs during the age range we examine, the first part of the current study is intended to increase understanding of children’s knowledge of the words being used, their meanings, and the semantic relations that they share. In order to reason about abstract relations, one must first learn the meaning of the words being related and the nature of that relationship. Comprehending the antonym relation involves having an understanding of the concept of “opposite,” which makes it possible to identify an indefinite number of instantiations of the same abstract relation. For instance, we can understand that “love” and “hate” are related to each other in the same way that “rich” and “poor” are, even though we evaluate these pairs of words on different dimensions (in this case, emotion vs. economic status). Such variations in the dimensions relevant to antonyms may also be a source of difficulty for computational models of verbal analogy. For example, both Word2vec (Mikolov et al., 2013) and *Bayesian Analogy with Relational Transformations* (BART; Lu, Wu, & Holyoak, 2019) seem to perform less well on analogy problems based on antonyms than on problems involving other abstract semantic relations, such as category membership.

The second part of the present study focuses on how two vector-based computational models, Word2vec and BART, perform on the same dataset as that used with children. These models are intended to model adult-level performance on analogy tasks involving semantic relations; however, deviations from children’s performance (particularly if children surpass the accuracy levels of the models) may potentially reveal limitations of the models. Examining possible differences across parts of speech might elucidate whether these models exhibit the same patterns of difficulty as young children who are just beginning to learn this abstract relation.

### Children’s Performance on an Antonym Analogy Task

In order to eliminate the constraints of a verbal task, we used a pictorial analogy task intended to measure children’s ability to solve analogy problems between pairs of antonyms.

### Methods

**Participants** 30 four-year-old ( $M = 4.28$ ,  $SD = .82$ ) and 37 five-year-old ( $M = 5.51$ ,  $SD = .26$ ) children were recruited through the Language and Cognitive Development Lab at the University of California, Los Angeles (UCLA) either online or through the university child database. Only children whose parents granted formal consent participated, in accordance with the UCLA Institutional Review Board. Data collection was completed entirely online using Zoom.

**Measures** Parents completed a language survey in which they were instructed to identify the words that their children produce. This survey included words that were used in the analogy task (e.g., “opposite”) in order to determine whether the children had prior knowledge of the words used in the

study, and whether word knowledge would be related to their performance on the analogy task.




Figure 1: Examples of three trials on the pictorial antonym analogy task, illustrating the three lexical classes used in the task. A: An adjective source pair exemplifying a contrastive relation (*big* : *small*), with a distractor pair (*surprised* : *sad*) on the left and correct option (*happy* : *sad*) on right. B: A noun source pair (*boy* : *girl*), with correct option (*friends* : *enemies*) on the left and distractor pair (*friends* : *mother*) on right. C: A verb source pair (*cry* : *laugh*), with correct option (*smile* : *frown*) on the left and a distractor pair (*frown* : *hate*) on right.

**Materials** The pictorial analogy task followed the format of a Relational Match-to-Sample (RMTS) task (see Figure 1). Children were allowed to simultaneously compare a source pair exemplifying a contrast relation to a target pair also exemplifying a contrast relation, but on a different dimension than the source (e.g., size versus cleanliness), and a distractor pair that did not exemplify a contrast. Thus, the relational match between the source and target was at the abstract level of antonym rather than at the level of a more specific relational contrast.

The pairs corresponding to antonyms, as well as the distractor pairs, were pictures of people, familiar animals, and/or objects. As illustrated in Figure 1, the objects used in the target and distractor were from the same category, which differed from the category used for the source objects. The contrastive relations used in the task could be expressed as either adjectives (e.g., *happy* : *sad* :: *dry* : *wet* or *tired* : *dry*), nouns (e.g., *friends* : *enemies* :: *teacher* : *mother* or *teacher* : *student*), or verbs (e.g., *open* : *close* :: *build* : *destroy* or *build* : *stop*).

**Stimulus validation** The antonym word pairs were sourced from educational websites and subsequently verified on WordBank, a database of children’s vocabulary development. Word pairs were chosen by selecting only those

known to over 80% of 30-month-olds. In order to determine which pairs of antonyms were appropriate to use, we conducted a Google Form survey with adults to validate which words are considered “opposites.” Two forms were created, each of which included one of the words in each pair. A set of twenty-five UCLA undergraduates were asked to generate the antonym for each word on a list, and only pairs with reliability over 95% were chosen for the final list.

**Procedure** Children received three training trials (one per part of speech) and thirty test trials (ten per part of speech), all within-subjects. Children were assigned to one of two conditions (between subjects): the Label and No-Label conditions. On each trial in the Label condition, children were told that the animals/objects/humans depicted in the source pair were “opposites” of each other (e.g., “This is dirty, this is clean. Dirty and clean are *opposites*”) in both practice and test trials. The words used to describe objects were either adjectives, nouns, or verbs. In the No-Label condition, children were not given a label for the abstract relation in any of the trials. Instead, children were only provided with verbal descriptions of the individual objects (e.g., “This is dirty and this is clean”).

For each condition, we created five versions of the task to semi-randomize which source pairs were matched with which target/distractor pairs. Because some of the picture pairs were repeated across trials, combinations were semi-randomized so that a target pair never appeared earlier as a source pair. For example, if a pair based on *big/small* was used as the source pair in the first test trial, that pair was never used afterward as a target pair. The part of speech was kept consistent among the source, target, and distractor pairs for every trial. In addition, the display position of the target/distractor pair was randomized between trials such that the correct pair appeared on the left side of the screen for half of the trials and on the right side of the screen for the other half. Children were randomly assigned to one of the five versions within each condition.

**Practice trials** To begin, children were shown a source picture showing two animals, objects, or humans depicting a pair of antonyms (e.g., a big balloon and a small balloon; see Figure 1). The experimenter labeled the pictures, emphasizing the words that depicted the contrastive relation (e.g., “This is *big*, this is *small*. Big and small”). Afterward, simultaneously, the experimenter provided two more images that respectively depicted either a target pair of antonyms (e.g., a *clean* pig and a *dirty* pig) or a distractor pair of semantically-unrelated words, one of which was kept consistent with the antonym pair (e.g., a *clean* pig and a *sad* pig). The experimenter always described each of the pictures, emphasizing the key words (e.g., *clean* and *dirty*). The participants were asked, “Which one is like this one (pointing to the source picture)?” Children were given feedback: either told that they were correct or told the correct answer if the child provided an incorrect one.

**Test trials** The format of test trials was identical to that of training trials, except that the children were not given any feedback regarding their answers on each trial. The

animals/objects/humans shown in the target/distractor pictures were always kept consistent in color and category (animals, objects, or humans) but were different in these respects from the source picture. These constraints ensured that children could not simply choose the picture that was most similar to the source picture based on features of individual objects.

### Behavioral Task Results

To analyze children’s accuracy in selecting the correct target, we implemented a Bayesian logistic regression model using the R package brms (Burkner, 2018). We tested hypotheses by fitting a logistic regression model predicting responses on the analogy task based on the interaction between condition (Reference = No Label) and age (Reference = four-year-olds). This model included group-level effects of subject and item and allowed for heterogeneity in the intercepts of the effects of condition and age. The model also included a grouping of the item types into three parts of speech in order to analyze differences between problems based on nouns, verbs, and adjectives. For the prior distributions in our model, we used a uniform (i.e., uninformative) distribution for the main effects and interaction coefficient, and used a  $t(3,0,2.5)$  for the random intercepts and their standard deviation. Specified in brms syntax, the model is:

$$\text{Response} \sim \text{Condition} * \text{Age} + \text{PartofSpeech} + (1 | \text{Subject}) + (1 | \text{Item})$$

These analyses revealed that being in the older age group and being given the relation label of “opposite” predicted higher accuracy on the analogy task ( $b = 0.82$ , 95% CI [-0.05, 1.66]) (see Table 1 and Figure 2). Moreover, the pattern of results suggested that labeling the antonym relation was particularly effective for five-year-old children, but did not make a difference for four-year-old children. Although previous research indicates that four-year-olds do have some understanding of the antonym relation when provided with a label (Phillips & Pexman, 2015), the current findings indicate that this abstract analogy task is too difficult for four-year-olds to solve with above-chance accuracy, even with a verbal label for the relation. Similarly, reliable performance on other versions of RMTS problems is not observed prior to age five (Hochmann et al., 2017).

Table 1: Estimates of Posteriors for Bayesian Logistic Regression Model

Population-Level Effects	Estimate	Est. Error	Lower 95% CI	Upper 95% CI
Intercept	-0.98	1.30	-3.56	1.59
Condition	-3.16	1.98	-7.06	0.79
Age	0.37	0.28	-0.18	0.92
Part of Speech – Noun	-0.30	0.24	-0.77	0.15
Part of Speech – Verb	-0.18	0.24	-0.66	0.28
Condition*Age Interaction	0.82	0.44	-0.05	1.66

An item analysis revealed that there were no reliable differences among individual items within a lexical class, indicating that no lexical class was systematically more challenging than others for children. Moreover, there were no

differences across parts of speech for four-year-olds, regardless of condition.




Figure 2: Percentage of correct responses across all parts of speech tested in the pictorial analogy task as a function of age (four- to five-year-olds), separately for the Label and No-Label conditions.

In contrast, for five-year-olds, differences were found across parts of speech. In particular, five-year-old children performed more accurately on trials involving adjectives ( $M = .751$ ,  $SD = .179$ ) than nouns ( $M = .627$ ,  $SD = .223$ ) ( $t(36) = 4.029$ ,  $p < .001$ ), and more accurately on trials involving verbs ( $M = .724$ ,  $SD = .032$ ) than nouns ( $t(36) = 2.97$ ,  $p = .005$ ) (see Figure 3). There were no differences between how five-year-olds performed on adjective and verb trials ( $t(36) = 1.137$ ,  $p = .26$ ). In addition, regardless of condition, there were no differences between four- and five-year-old children’s performance on noun trials. These results suggest that providing five-year-olds with a label aided them on analogy problems involving adjectives and verbs, but not those involving nouns, for which they perform as poorly as do younger children.




Figure 3: Mean correct responses for Label and No-Label conditions across three lexical classes, separately for four- and five-year olds.

### Performance of Computational Models on Antonym Analogy Task

We implemented two computational models of verbal analogy, Word2Vec (Mikolov et al., 2013) and BART (Lu et al., 2019), to compare model predictions with children’s

performance. Both models operate on vector representations (*embeddings*) of individual word meanings. However, as illustrated in Figure 4, the two models operate in different representation spaces. Word2Vec is based on a semantic space for individual words, in which words with similar meanings are clustered together. In contrast, BART represents word meanings in a relation space, in which each dimension codes a specific relation. Accordingly, word pairs instantiating similar relations are located close together in the BART relation space. Based on their representations of the two words in each pair, each model computes the dissimilarity of a source word pair with a target word pair, and selects the option with the smaller dissimilarity value as the predicted correct response.




Figure 4: Illustration of Word2vec semantic space for individual words, and BART relation space for word pairs.

Word embeddings produced by Word2vec (Mikolov et al., 2013) were used to represent the meanings of each of the words included in the test trials of the pictorial analogy task (90 word pairs, with 180 total word embeddings). Word2vec-diff is a measure defined as the difference between the vectors of each word in a pair: i.e.,  $f_A - f_B$  for the word pair  $A:B$ . The dissimilarity between two pairs is then defined by the cosine distance between the difference vectors for the two pairs:

$$D_{W2V-diff} = \cos(f_A - f_B, f_C - f_D).$$

The second model, BART, is trained on a set of specific relations, including 79 abstract relations from the SemEval-2012 Task-2 dataset (Jurgens et al., 2012) and additional 56 relations in (Popov et al., 2017). For each of those relations, BART was trained with less than 100 examples, including a small number (10 or 20) of positive examples instantiating this relation, and some negative examples (~70) that instantiate other relations.

After learning explicit representations of each semantic relation, BART encodes the specific relation between any pair of words ( $A, B$ ) using distributed representations expressed as a relation vector  $R_{AB}$ , in which each element indicates the probability that this pair of input words instantiates each of the learned relations. The relation vector is in the size of 270 dimensions (including the 135 relations in the training datasets and their corresponding converses). To solve an analogy problem, the model computes dissimilarity as the cosine distance between corresponding relation vectors based on the two word pairs, and selects the answer with smaller dissimilarity:

$$D_{BART} = \cos(R_{AB}, R_{CD}).$$

## Simulation Results

For both models, the dissimilarity between the word pairs was computed using cosine distance between the vectors representing each pair. If the cosine distance between the source pair and the target pair was less than that between the source pair and the distractor pair, we considered that the models had correctly answered the analogy problem. Note that neither model (W2vec-diff and BART) is sensitive to the presence of relation labels. Accordingly, we focus on comparing model predictions and children’s performance in the No-Label conditions.

BART performed most accurately on adjective antonym pairs (.80 correct), followed by noun and verb pairs (.60 and .40, respectively (see Figure 5). Word2vec also performed most accurately on adjective pairs (.70 correct), followed by verbs and nouns (.60 and .50, respectively).




Figure 5: Percentage of correct responses for each part of speech for four- and five-year-old children on the analogy task in the No-Label condition, with the performance of two computational models, BART and Word2vec (W2V).

Overall, both models showed variability across the different parts of speech. Both models yielded levels of accuracy approximating (or higher than) that of five-year-olds in the No-Label condition for antonyms based on adjectives and nouns; but for verb antonyms, the models (particularly BART) fell well short of the level achieved by five-year-olds.

## General Discussion

The present study applied both developmental and computational methods to examine the solution of analogy problems based on antonyms. Using a verbal analogy task (with picture illustrations), we demonstrated that by age five years—before the antonym relation is formally taught in school—children are able to reliably solve analogies based on antonyms, especially when the antonym relation is given a verbal label (“opposites”).

Although we found no differences in performance across lexical classes for four-year-olds (who performed at chance on all three types), five-year-olds were more accurate on analogy problems based on adjectives and verbs as compared to nouns. These findings suggest that developmental trends

in reasoning about the antonym relation do not coincide with children's lexical development, given that children tend to acquire nouns earlier than adjectives and verbs (Nelson, 1973; Sandhofer & Smith, 2007). A possible explanation is that nouns can be compared on a wider range of dimensions than either adjectives or verbs, making it more difficult to determine the basis for an antonymy relation for nouns. A future step to address this issue would be to examine variability in how both children and adults generate antonym pairs across different lexical classes. Perhaps responses will be especially variable for antonyms based on nouns.

We also compared children's performance to that of two vector-based models of verbal analogy, Word2vec (Mikolov et al., 2013) and BART (Lu et al., 2019), using the same set of problems. These models are based on embeddings derived from training on corpora of adult language; however, both models fell short of children's level of accuracy, particularly for solutions to problems based on verb antonyms by five-year-olds. In addition, neither model is sensitive to the provision of verbal labels for the antonym relations. Additional computational work will be required to address these apparent shortcomings of current models. Finally, the present study sets the stage to study the development trajectory of other semantic relations, such as cause-effect and category-membership, in analogical reasoning.

### Acknowledgments

We thank the children and families who participated in this research, the research assistants from the LCD lab for help with data collection, and members of the LCD and Reasoning Lab for their insightful comments. Preparation of this paper was supported by NSF Grant BCS-1827374.

### References

- Brown, A. L., Kane, M. J., & Long, C. (1989). Analogical transfer in young children: Analogies as tools for communication and exposition. *Applied Cognitive Psychology, 3*(4), 275-293.
- Bürkner, P. C. (2018). Advanced Bayesian multilevel modeling with the R package brms. *The R Journal, 10*(1), 395-411. doi: 10.32614. RJ-2018-017.
- Entwisle, D. R., Forsyth, D. F., & Muuss, R. (1964). The syntactic-paradigmatic shift in children's word associations. *Journal of Verbal Learning and Verbal Behavior, 3*, 19-29.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science, 7*(2), 155-170.
- Gentner, D. (2010). Bootstrapping the mind: Analogical processes and symbol systems. *Cognitive Science, 34*(5), 752-775.
- Gentner, D., & Markman, A. B. (1994). Structural alignment in comparison: No Difference without similarity. *Psychological Science, 5*(3), 152-158.
- Gentner, D., & Medina, J. (1998). Similarity and the development of rules. *Cognition, 65*(2-3), 263-297.
- Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology, 15*, 1-38.
- Hochmann, J.-R., Tuerk, A. S., Sanborn, S., Zhu, R., Long, R., Dempster, M., & Carey, S. (2017). Children's representation of abstract relations in relational/array match-to-sample tasks. *Cognitive Psychology, 99*, 17-43.
- Holyoak, K. J., Junn, E. N., & Billman, D. O. (1984). Development of analogical problem-solving skill. *Child Development, 55*, 2042-2055.
- Jurgens, D. A., Mohammad, S. M., Turney, P. D., & Holyoak, K. J. (2012). SemEval-2012 Task 2: Measuring degrees of relational similarity. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics (\*SEM)* (pp. 356-364). Montreal, Canada: Association for Computational Linguistics.
- Loewenstein, J., & Gentner, D. (1998). Relational language facilitates analogy in children. In *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 615-620).
- Lu, H., Wu, Y. N., & Holyoak, K. J. (2019). Emergence of analogy from relation learning. *Proceedings of the National Academy of Sciences, USA, 116*, 4176-4181.
- Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development, 67*(6), 2797-2822.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In C. J. C. Burges, L. Bottou, & M. Welling, *Advances in Neural Information Processing Systems, 26*, 3111-3119.
- Nelson, K. (1973). Structure and strategy in learning to talk. *Monographs of the Society for Research in Child Development, 38*. Chicago, IL: University of Chicago Press.
- Phillips, C. I., & Pexman, P. M. (2015). When do children understand "opposite"? *Journal of Speech, Language, and Hearing Research, 58*(4), 1233-1244.
- Popov, V., Hristova, P., & Anders, R. (2017). The relational luring effect: Retrieval of relational information during associative recognition. *Journal of Experimental Psychology: General, 146*, 722-745.
- Rattermann, M.J., & Gentner, D. (1998). More evidence for a relational shift in the development of analogy: children's performance on a causal-mapping task. *Cognitive Development, 13*, 453-478.
- Sandhofer, C., & Smith, L. B. (2007). Learning adjectives in the real world: How learning nouns impedes learning adjectives. *Language Learning and Development, 3*(3), 233-267.
- Shao, R., & Gentner, D. (2018). Same/different relation detection and word production in 4-year-olds. In *Proceedings of the 40th Annual Conference of the Cognitive Science Society*.
- Tribushinina, E., van den Bergh, H., Kilani-Schoch, M., Adsu-Koc, A., Dabasinskiene, I., Hrzica, G., . . . Dressler, W. U. (2013). The role of explicit contrast in adjective acquisition: A cross-linguistic longitudinal study of adjective production in spontaneous child speech and parental input. *First Language, 33*, 594-616.