

A Grounded Approach to Modeling Generic Knowledge Acquisition

Deniz Beser (beser@isi.edu)

Joe Cecil (cecil@isi.edu)

Marjorie Freedman (mrf@isi.edu)

Jacob Lichtefeld (jacobl@isi.edu)

Information Sciences Institute, University of Southern California
Marina del Rey, CA 90292, USA

Mitch Marcus (mitch@cis.upenn.edu)

Department of Computer and Information Science, University of Pennsylvania
Philadelphia, PA 19104, USA

Sarah Payne (paynesa@sas.upenn.edu)

Charles Yang (charles.yang@ling.upenn.edu)

Departments of Linguistics and Computer and Information Science, University of Pennsylvania
Philadelphia, PA 19104, USA

Abstract

We introduce and implement a cognitively plausible model for learning from generic language, statements that express generalizations about members of a category and are an important aspect of concept development in language acquisition (Carlson & Pelletier, 1995; Gelman, 2009). We extend a computational framework designed to model grounded language acquisition by introducing the *concept network*. This new layer of abstraction enables the system to encode knowledge learned from generic statements and represent the associations between concepts learned by the system. Through three tasks that utilize the concept network, we demonstrate that our extensions to ADAM can acquire generic information and provide an example of how ADAM can be used to model language acquisition.¹

Keywords: Language Acquisition, Generics, Cognitive Modeling

Introduction

Generics, statements in the form of “*bananas are yellow*” or “*birds fly*”, express generalizations about members of a category, and are frequent in everyday language (Carlson, 1977; Carlson & Pelletier, 1995). These expressions refer to a category as a whole (e.g. *birds*), as opposed to a single instance of a category (e.g. *a bird*), and while referring to a conceptual category as a whole, generic statements may assert information that while typical, does not cover all instances, and hence are not necessarily invalidated by counterexamples (McCawley, 1981; Gelman, 2004). For example, the statement “*birds fly*” is not invalidated by the existence of penguins, a bird that cannot fly. Moreover, experience with only a single instance of a conceptual category can be sufficient to acquire generic knowledge (Carlson & Pelletier, 1995; Prasada, 2000). Besides the prevalence and the expressive power of generics in language, generics also play an important role in child language acquisition. During conceptual development, generic statements complement observational learning and help construct conceptual knowledge, al-

lowing children to learn hierarchical information that cannot be learned from the world with experience alone (Cimpian & Markman, 2009; Gelman, 2009).

Prasada (2000) emphasized the acquisition problem posed by the imprecise nature of generic language and our ability to learn generic knowledge from minimal experience, outlining that the requirements of a formal system for acquiring generic knowledge are to complement other learning mechanisms and to establish relationships between categories and properties. By age 2, children manifest the capabilities of such a system, as they can recognize category labels (Graham, Kilbreath, & Welder, 2004) and make use of names to classify objects, even if an instance is atypical for a given category (Nazzi & Gopnik, 2001; Jaswal & Markman, 2007). Considering how generic utterances such as “*penguins are birds*” can override interpretation of perceptual features and allow learning hierarchical knowledge even from a young age (Gelman, 2009; Cimpian, Brandone, & Gelman, 2010), it is clear that language can inform perceptual learning, making generics a critical component of language acquisition worth modeling.

Inspired by the plethora of research on generics and child language development, we introduce and implement a developmentally plausible model for learning concepts from generic language that assumes no prior conceptual knowledge, learns meanings of words and concepts in a manner similar to infants (from the ground up), and demonstrates the capacity to make inference with the learned concepts. The importance of generic language has been recognized in the artificial intelligence and natural language processing community for tasks that involve knowledge acquisition, ontology development, and semantic inference (e.g., Reiter & Frank, 2010; Friedrich & Pinkal, 2015; Sedghi & Sabharwal, 2018). These approaches generally make use of large-scale resources and employ methods such as supervised learning that are not suitable for modeling child language and conceptual development. We base our system on ADAM, a

¹Our code is available at <https://github.com/isi-vista/adam>

software platform for modeling early language acquisition (Gabbard et al., 2021). Through pairs of perceptual representations of situations and linguistic utterances describing these situations, ADAM learns interpretable representations for *concepts* such as objects, attributes, relations, and actions (see Model and Representation section for further detail). ADAM’s cognitively plausible design choices regarding perceptual representations (Marr, 1982; Biederman, 1987) and simple word/pattern learning mechanisms (Webster & Marcus, 1989; Yu & Smith, 2012; Stevens, Gleitman, Trueswell, & Yang, 2017) lay the groundwork for modeling higher level semantics in child language acquisition. Given the interplay between generic statements, concepts, and observational learning in natural language (Carlson, 1977; Carlson & Pelletier, 1995; Gelman, 2009) and ADAM’s power to learn meanings of concepts and identify them across situations, we consider ADAM to be a promising system for modeling acquisition of generic knowledge and the resultant semantic category inference.

We expand ADAM’s modeling capabilities to capture semantic associations and generics by introducing a generic learner module and combining ADAM’s representations with an additional layer of abstraction, a network data structure called *the concept network*. The concept network organizes the associations between concepts and learns hierarchies and properties about the concepts through observation as well as generic utterances. Through three tasks that use generic language across different learning curricula, we demonstrate that ADAM, coupled with the generic learner module and the concept network, can acquire generic knowledge, establish semantic certainty with generics language, and make category inference. Our demonstrations provide an example of how ADAM can be used to model language acquisition.

Model and Representation

Meaning Representation in ADAM

ADAM is a software platform for experiments in child language learning. The system can process a range of expressions covering a very young learner’s vocabulary and grammar, such as objects (“a ball”), adjectives (“a red ball”), prepositions (“a ball on a table”) and actions (“Mom rolls a ball”). ADAM uses *perception graphs* to represent situations (e.g. a cup sitting on a table) that it perceives. Perception graphs consist of perceptually plausible components, such as geons (Biederman, 1987), body parts, colors, and regions. For example, when the model observes a cup on the table, the model perceives a structured graph that consists of the cup’s color, hollow shape, its handle, and its position relative to the table. The model learns patterns over observed perception graphs. The patterns represent hypotheses about the meaning of individual *concepts* and are consolidated throughout the learning process. For example, the learner could perceive perception graphs and utterances from multiple situations that involve a bear, such as hearing “a brown bear” while observing a bear by itself, hearing “a bear sits” while observ-

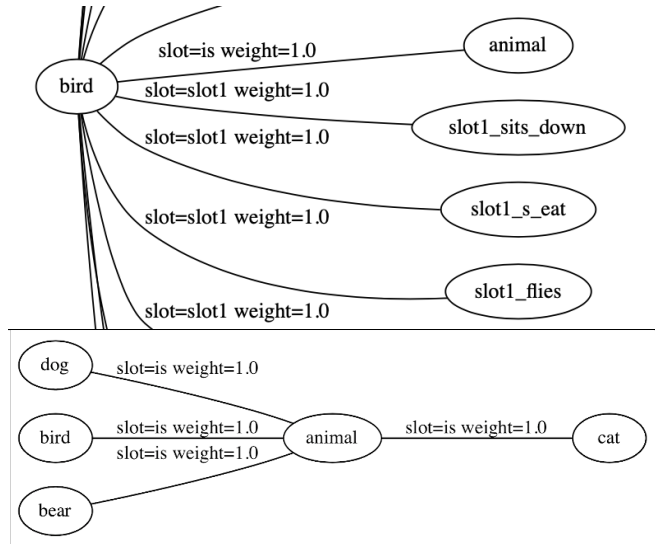


Figure 1: Examples of a concept network. Nodes represent concepts (e.g. *bird*, *animal*, *sitting down*). The edges are labeled by the semantic relation between the concepts (slot), and the association strength (weight). The top shows a portion of the concept graph for *bird*. The bottom shows *animal* and its neighbors.

ing a bear sitting, and hearing “two bears” while observing two bears, and eventually learn a representation of “bear”’s meaning. ADAM uses the learned mapping between linguistic structures and meaning representations to describe new scenes. Before learning, we define a configurable curriculum that pairs observations and descriptions. During testing, the perception input is generated without the descriptions.

The Concept Network

We extend ADAM’s representations of concepts and patterns with the concept network, a graph-based data structure that represents learned concepts (e.g. objects, attributes, and actions) as nodes and the semantic associations between them as edges. The concept network enables one to see whether any two concepts are related, what their relation is, and how strong this relation is, by looking at the edge connecting the two concepts. Contrary to ADAM’s perception graphs and patterns that are used to describe the components of a scene or perceptual properties of a perceived object, the concept network represents the overall semantic and conceptual knowledge of the learner learned over time. We represent each concept as a single node in the network; for example, *bird* is represented by a single concept node in the concept network, and other semantically related concepts are its neighbors. Figure 1 (top) visualizes a section of the learned concept network structure for the *bird* concept. Using the concept network, we can infer that a concept such as *bird* is an *animal*, and that it is associated with the concept *fly*. Since these associations are formed through the learner’s observations of occurrences

of concepts throughout learning, the association between the *bird* and *fly* concepts would suggest that a bird was observed while flying. Similarly, the network can be used to observe the categories formed by the learner; Figure 1 (bottom) shows the *animal* node and its neighbors. To construct the nodes and edges in this network, while perceiving concepts and semantic relations between them as the learner observes situations as described in Gabbard et al. (2021), the learner also simultaneously creates nodes for concepts (e.g. *bird*, *fly*) and edges to associate concepts in the network. The concept network consists of about 250 concepts when generated with ADAM’s standard curriculum (see Gabbard et al. (2021) for a complete list of content descriptors).

Edges and Semantic Associations The edges in the concept network represent the semantic relation between the two concepts and the association strength between them. For action concepts, the relation often denotes an argument relation with the neighboring concept; the argument relation is labeled with *slots* that describe a slot an argument can take in a phrase. For instance, given “*Mom drinks juice*”, *Mom* has argument slot position *slot-1* with *drink*, while *juice* has *slot-2*. The strengths in the network start from 0 and have a maximum of 1.0. Throughout learning, the association strength between two concepts is updated with each co-occurrence of the two concepts, using the plateauing update function $a = a + 0.2 * (1 - a)$. The edges that represent semantic associations initialized with generic statements have maximum association strength, since we tend to believe what we are told is true. These edges are also marked with a slot label that matches the argument relation of the statement, or the *is* label if the statement is a predicate.

Matrix Representation We represent the network as a weighted adjacency matrix to facilitate rich operations such as vector similarity between concepts. In order to preserve argument relation information in the edges in the matrix, each action concept in the concept network is first translated to multiple concepts that include the argument relation (e.g. *drink slot 1* to represent drink in *Mom drinks* and *drink - slot 2* to represent drink in “*drinking juice*”).

Concept Similarity We use concept similarity for evaluation. Each concept is represented as an adjacency vector extracted from the weighted adjacency matrix. Vectors for *category concepts* (e.g. *animal*) are represented as the average of the vectors of all members of that category. Similarity between pairs of concepts is measured using cosine similarity. Measuring concept similarity provides insight into the hierarchical structures that naturally emerge throughout learning, as visualized in Figure 2 which shows a clustered heatmap of the vector representations of concepts that are learned through a curriculum of objects and actions, without generics. We see the formation of categories that correspond to body parts, liquids, and animate objects.




Figure 2: A qualitative demonstration of how hierarchical categories can be measured with concept similarity. We see the formation of categories in the concept space, e.g. animate objects (*cat*, *dog*, *Mom*, *baby*, *Dad*), liquids (*water*, *juice*, *milk*), and body parts (*head*, *hand*). The vector representations of learned concepts were clustered with hierarchical clustering method *clustermap* (Müllner, 2011; Waskom & the Seaborn Development Team, 2020).

Generic Utterances

To teach generic language and explicit categories to the learner, we present simple scenes paired with generic utterances such as “*bears sit*” and “*bears are brown*”. Since generic language encodes generic knowledge (e.g. “*birds fly*” stipulates that birds generally can fly), inputs to the learner that are in generic form maximize the corresponding association strength when observed by the learner. For instance, while the non-generic utterance “*a bear sits*” increases the association strength between *bear* and *sit*, observing “*bears sit*” maximizes it.

Generic Learner Module The ADAM system is built of learning modules. Each module targets a specific type of learning, such as objects, actions, and relations. We build a generic learner module to enable learning from generic utterances. The generic learner first verifies that the utterance is a generic by checking whether all the recognized nouns in the utterance are bare plurals, an indicative property of generics in English (Lyons, 1977; Gelman, 2004). Once confirmed, the learner recognizes the concepts mentioned in the utterance, and forms a semantic connection between these concepts in the concept network. The association strength of this connection is maximal, which implies a semantic certainty that is learned from generic input.

In the special case where the generic statement is a predicate containing a previously unknown category, such as *animal* in “*dogs are animals*,” we create a new category concept node in the concept network and associate the object concept (*dog*) with the category concept (*animal*). If a novel object

concept appears in a generic statement as a member of a category concept, such as *wug* in “*wugs are animals*,” we create a new object concept and associate it with the category concept as well as the features of the members of the category. Overall, the module can interpret statements in the form of “*birds fly*,” “*birds are animals*” and “*wugs are animals*”.

Model Evaluation

We present three tasks to illustrate the behavior of the system across different learning conditions.

Task 1: Generic Color Predicates

The generic color predicates task shows how the system can learn generics by establishing stronger associations between learned concepts. In other words, we measure how the generic input can change the learner’s understanding of the world by influencing semantic connections. In this task, we first teach objects and arbitrary colors to the system with non-generic utterances (e.g. *a red truck*) based on the standard object and color curricula provided in the ADAM system. We then evaluate the associations between the object and color concepts that the system has learned. Then, we pick colors that are typical for the object concepts and teach them with generic statements (“*watermelons are green*,” “*papers are white*,” “*cookies are light brown*”), and evaluate the associations between the object and color concepts again. In the non-generic training phase, each object appears with a variety of arbitrary colors (e.g. cookies with blue, green, and light-brown colors), which we expect will lead the model to learn associations between an object and many colors. We expect the presentation of generic input to yield stronger associations between the object and the color used in a generic utterance.

Results and Discussion We measure the association strengths between objects and all associated colors before and after observing generic color predicates. The results are plotted in Figure 3. Prior to hearing generic input, objects have associations to many arbitrary colors, reflecting the initial non-generic curriculum. Once the generic color predicates are observed, we see that each object has a significantly stronger association with the appropriate color as stated in the generic utterances. Upon observing generics, only the association strengths for the correct colors are updated. These results demonstrate that our model interprets and uses generics as designed, to establish strong associations between an object and a prototypical version of some feature.

Task 2: Category Inference

The category inference task demonstrates how the model forms categories, e.g. *animals* and *food*. As visualized in Figure 2, concept categories form naturally in the representational space throughout learning, as similar objects have similar semantic roles in actions; for instance, only animate objects *eat*, and only food objects are *eaten*. In this task, we label these implied categories using generic statements

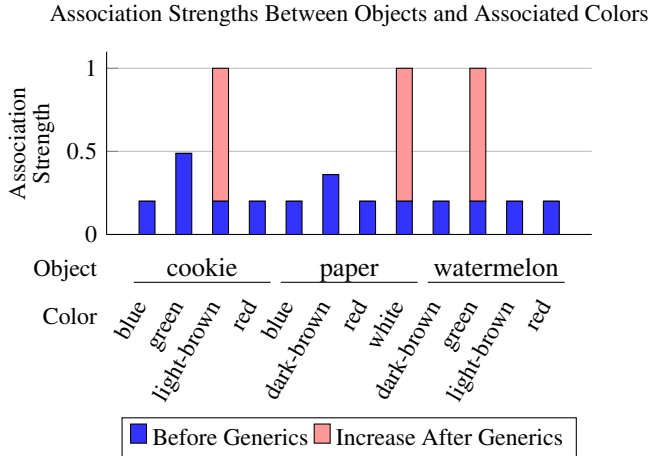


Figure 3: Results of the generic color predicates task, plotting the associations strengths between objects and colors that are associated with them, before and after observing generic input. At first, each object is associated with an arbitrary set of colors that reflect the curriculum. The only associations that increase are correctly the ones observed through generics.

(e.g. “*cats are animals*,” “*dogs are animals*”). Then, we compare the novel concepts placed in these categories with each possible category. We first teach the learner a particular curriculum, then a set of previously unknown objects in known categories (e.g. “*wugs are animals*”), and finally evaluate the similarity of the new item (e.g. *wug*) to each possible category. We repeat this process for different curricula with increasing semantic complexity, by including more complicated content such as *plurals* and *actions*. Table 1 shows different learning curricula (top), example utterances (center), and examples with unknown categories (bottom). We expected that, while the model should be able to form categories regardless of the curricula, the formed categories would be most distinct in the less complex conditions.

Table 1: Inputs to the learner in the category inference task: (top) learning curricula; (center) examples of situations and utterances; (bottom) inputs presented to the learner. Animal and people categories are treated as separate categories.

Learning Curriculum	
1	Objects, categories
2	Objects, categories, generics
3	Objects, colors, actions, plurals, categories, generics

Results and Discussion The results of the category inference task are plotted in Figure 4. In the easiest curriculum setting containing only objects and categories, every object is most similar to the category in which it was learned, i.e. *wugs* are most similar to *animals*, *vonks* to *foods*, and *snarps* to *people*. Expectedly, we see some similarity between animals and foods, because *chicken* is a member of both categories. The

Similarity of Novel Objects to Each Category across Curricula

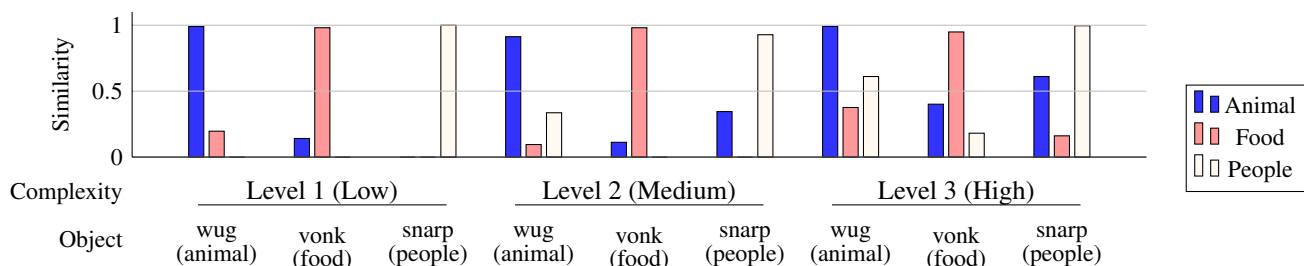


Figure 4: Results of the category inference task, plotting the similarity of novel objects (e.g. *wug*) to animal, food, and people categories across curricula of varying complexity. The correct categories of objects, i.e categories in which they were presented in the generic input, are labeled in parentheses. Regardless of the curriculum complexity and concept category, the learner associates new objects most strongly with the correct category. Incorrect associations increase with curriculum complexity.

Curriculum	Examples
Objects	a house; a dog
Colors	a red truck; papers are white
Actions	a baby drinks milk from a cup
Plurals	two balls; many cookies
Category Generics	bears are animals
Action Generics	cats walk; Moms eat

New Object	Category	Utterance
wug	animal	wugs are animals
vonk	food	vonks are foods
snarp	people	snarps are people

second curriculum condition, which includes objects, categories, and generics, shows that while the correct trend holds, there is some association between animals and people due to generic statements that apply to both categories, such as *sitting*. Finally, in the most complex curriculum condition that includes objects, some actions, plurals, colors, categories, and generics, we see that the associations across categories are increased, but the correct category still has the highest association strength. There is some similarity between animals and foods, and people and animals due to shared concepts such as *chicken* and *eating* respectively. Overall, while associations between object concepts and the incorrect categories increase with curriculum complexity, regardless of the curriculum complexity and category of the novel object, objects are associated most strongly with the correct category in which they were learned through generic statements.

Task 3: Joint Category

The goal of the joint category task is to demonstrate how the ADAM system learns categories across curricula with different contents. Specifically, we create learning curricula for each of the four conditions shown in Table 2 and evaluate the model behavior. We use combinations of *chicken*, *beef*, and *cow* objects; *chicken* is used as an example of a lexical

item that is shared by two very similar yet distinct concepts (*chicken* as an animal, and *chicken* as food) and hence is a member of both the *food* and *animal* categories. While *beef* and *cow* could refer to the same thing at a certain semantic level, they are regarded as disjoint examples of a food and an animal category respectively. We hypothesized that while observing *chicken* would cause some semantic association between food and animal categories, the model should not show any semantic association between food and animal categories when it observes just *beef* and *cow* and no *chicken*.

To execute the task, we run four different versions of the ADAM system, each one trained with one of the four curriculum conditions shown in Table 2. Then, similar to the category inference task, each system is presented a set of previously unknown objects in known categories, e.g. “*wugs are animals*”. Finally, we evaluate the similarity of the new object concept to animal and food categories.

Table 2: Curricula and contents for the joint category task

	Test Objects Included in the Learning Curriculum
1	None (no chicken, beef, or cow)
4	Beef (<i>food</i>) and cow (<i>animal</i>)
3	Chicken (<i>food</i> and <i>animal</i>)
4	Chicken, beef, and cow

Results and Discussion Figure 5 plots the results of the joint category task, showing the similarity of the *wug* concept to animal and food categories across curricula with different contents. In the first condition, with a curriculum that does not include *chicken*, *beef*, or *cow*, the similarity to animals is strong, but the similarity level for food category indicates an absence of similarity between animals and foods. Likewise, in the second curriculum condition, when the curriculum includes *beef* and *cow*, but not *chicken*, the learner does not learn any association between food and animal categories. However, when we introduce *chicken* into the cur-




Figure 5: Results of the joint categories task. Through generic input, a novel object, *wug* is learned as an *animal*. The plot shows the similarity between the *wug* concept and the animal and food categories. *Wug* is more similar to foods when the curriculum includes chicken. The similarity is lower when beef and cow are added and is zero without chicken.

riculum, the association between animals and foods grows. Finally, including *chicken*, *beef*, and *cow* in the curriculum leads to some similarity between foods and animals, but less so than when only *chicken* was included. That these changes match our expectations suggests that ADAM can successfully capture differences in the curricula and the contents in them. While doing so, the system maintains its robustness as indicated by how the similarity to the correct category outcores the similarity to the incorrect category in every condition.

Conclusion & General Discussion

We have illustrated that the grounded language acquisition system of ADAM, coupled with the learner module and concept network, can be successfully used for modeling semantics of generic learning. We demonstrated that the model learns from generic language, makes desired semantic associations between learned concepts, and forms semantic categories that reflect the contents of the training curricula.

The concept network makes it possible to perform operations on learned concepts and on the semantics of generics, enabling generic utterances to establish associations between concepts and consequently form semantic categories. In the future, we plan to further examine the role of association weights in category formation, and how the system performs on different languages. We also hope to explore how we can use the matrix representation of the concept network to integrate ADAM’s structured and interpretable meaning representations with distributional models that operate on continuous space, such as neural networks. Moreover, while we did not include the meaning patterns of the learned concepts in the concept network, the flexible nature of the network makes this possible, creating the potential for further semantic analyses that utilize perceptual properties of learned concepts.

Overall, our system’s ability to acquire word meanings and robustness in capturing desired semantic properties across diverse learning curricula shows its promising capacity to

model different aspects of language acquisition and to investigate the question of how children discover the expression of generics in their own languages.

Acknowledgements

Approved for public release; distribution is unlimited. This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Agreement No. HR00111990060. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government.

References

- Biederman, I. (1987). Recognition-by-components: a theory of human image understanding. *Psychological review*, 94(2), 115.
- Carlson, G. N. (1977). Reference to kinds in english.
- Carlson, G. N., & Pelletier, F. J. (1995). *The generic book*. University of Chicago Press.
- Cimpian, A., Brandone, A. C., & Gelman, S. A. (2010). Generic statements require little evidence for acceptance but have powerful implications. *Cognitive science*, 34(8), 1452–1482.
- Cimpian, A., & Markman, E. M. (2009). Information learned from generic language becomes central to children’s biological concepts: Evidence from their open-ended explanations. *Cognition*, 113(1), 14–25.
- Friedrich, A., & Pinkal, M. (2015). Discourse-sensitive automatic identification of generic expressions. In *Proceedings of the 53rd annual meeting of the association for computational linguistics and the 7th international joint conference on natural language processing (volume 1: Long papers)* (pp. 1272–1281).
- Gabbard, R., Beser, D., Lichtefeld, J., Cecil, J., Marcus, M., Payne, S., ... Freedman, M. (2021). Adam: A sandbox for implementing language learning. *arXiv preprint arXiv:2105.02263*.
- Gelman, S. A. (2004). Learning words for kinds: Generic noun phrases in acquisition. *Weaving a lexicon*, 445–484.
- Gelman, S. A. (2009). Learning from others: Children’s construction of concepts. *Annual review of psychology*, 60, 115–140.
- Graham, S. A., Kilbreath, C. S., & Welder, A. N. (2004). Thirteen-month-olds rely on shared labels and shape similarity for inductive inferences. *Child development*, 75(2), 409–427.
- Jaswal, V. K., & Markman, E. M. (2007). Looks aren’t everything: 24-month-olds’ willingness to accept unexpected labels. *Journal of Cognition and Development*, 8(1), 93–111.
- Lyons, J. (1977). *Semantics* (2 volumes). Cambridge etc.: Cambridge UP.

- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*, Henry Holt and Co. Inc., New York, NY, 2(4.2).
- McCawley, J. D. (1981). *Everything that linguists have always wanted to know about logic... but were ashamed to ask*. University of Chicago Press.
- Müllner, D. (2011). Modern hierarchical, agglomerative clustering algorithms. *arXiv preprint arXiv:1109.2378*.
- Nazzi, T., & Gopnik, A. (2001). Linguistic and cognitive abilities in infancy: When does language become a tool for categorization? *Cognition*, 80(3), B11–B20.
- Prasada, S. (2000). Acquiring generic knowledge. *Trends in cognitive sciences*, 4(2), 66–72.
- Reiter, N., & Frank, A. (2010). Identifying generic noun phrases. In *Proceedings of the 48th annual meeting of the association for computational linguistics* (pp. 40–49).
- Sedghi, H., & Sabharwal, A. (2018). Knowledge completion for generics using guided tensor factorization. *Transactions of the Association for Computational Linguistics*, 6, 197–210.
- Stevens, J. S., Gleitman, L. R., Trueswell, J. C., & Yang, C. (2017). The pursuit of word meanings. *Cognitive science*, 41, 638–676.
- Waskom, M., & the Seaborn Development Team. (2020, September). *mwaskom/seaborn*. Zenodo. Retrieved from <https://doi.org/10.5281/zenodo.592845> doi: 10.5281/zenodo.592845
- Webster, M., & Marcus, M. (1989, June). Automatic acquisition of the lexical semantics of verbs from sentence frames. In *27th annual meeting of the association for computational linguistics* (pp. 177–184). Vancouver, British Columbia, Canada: Association for Computational Linguistics. Retrieved from <https://www.aclweb.org/anthology/P89-1022> doi: 10.3115/981623.981645
- Yu, C., & Smith, L. B. (2012). Modeling cross-situational word–referent learning: Prior questions. *Psychological review*, 119(1), 21.