Précis of *How Grammars Grow: Argument Structure and the Acquisition of Non-Basic Syntax*

Laurel Perkins
PhD 2019

Overview

What we can learn depends on what we already know. Just as a child who can't count cannot learn arithmetic, a child who can't segment words cannot learn properties of verbs in her language. Learning in any domain depends on how the data for learning are represented: as input representations change over the course of development, this impacts how learners form categories and draw generalizations. This perspective on how learners use data from their environment is orthogonal to the question of whether learners have innate knowledge that guides the learning process. Any complete theory of learning must account for how prior knowledge makes contact with the data of experience.

In the domain of language, theories since Chomsky (1965) have typically abstracted away from development by idealizing language acquisition as an instantaneous process of mapping a child's linguistic input onto a grammar (e.g. Wexler & Culicover, 1980); see Fig. 1. Many modern approaches make the same abstraction by asking whether the data as a whole support grammar selection (Clark & Lappin, 2013; Perfors, Tenenbaum, & Regier, 2006; Yang, 2002).

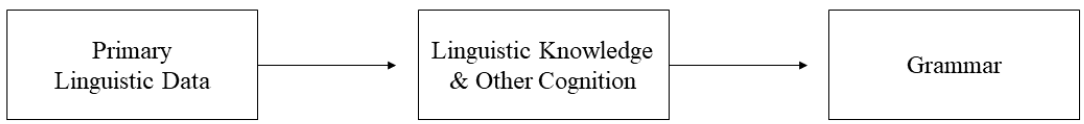

Fig. 1 Traditional model of language acquisition

While illuminating the problem of language learnability from a global perspective, these approaches leave out the problem of how children experience their input over development. In my dissertation, I explore how we can move to a developmental model of language acquisition (Fig. 2). In this model, children represent their input using the knowledge of their language that they currently have available— their developing grammar—together with other developing cognitive abilities. These input representations are immature at early stages of development, but they then become the data that children use to gain new knowledge and draw new inferences about their grammar (Lidz & Gagliardi, 2015). And this process iterates: children's updated grammars allow them to represent their input more richly, enabling further inferences about their grammar, until they incrementally arrive at the target state.

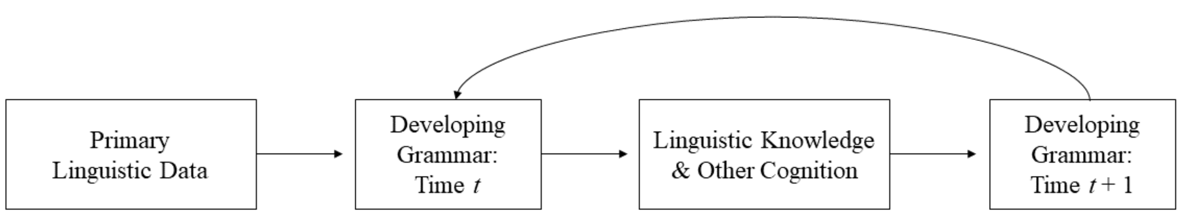

Fig. 2 Developmental model of language acquisition

Understanding how children incrementally converge on the correct theory of their language requires us to specify two types of resources that learners bring to the acquisition task. First, children need resources for representing their input before they have acquired the target grammar. And second, they need resources for drawing generalizations from input representations that are noisy and incomplete.

My dissertation investigates how children at the earliest stages of syntax acquisition represent their input and how they learn from it. Because children's earliest grammatical knowledge is acquired before they are able to express it in their own speech (e.g. Hirsh-Pasek & Golinkoff, 1996), we must look earlier in development—in infancy— to see how that knowledge arises. To this end, I adapt methods from developmental psychology that are tailored to the behaviors that a 1-year-old can control, such as looking time towards a stimulus. I supplement these behavioral studies with computational models of the mechanisms that infants might use to learn from their input. This approach is novel in three ways. First, it develops novel strategies for identifying syntactic representations in infancy. Second, it models the acquisition of syntax developmentally, taking into account what infants know and represent at different points in time. Third, it develops precise theories of learning that allow the acquisition of new knowledge from incomplete representations of the input.

As a case study, I focus on the domain of argument structure: how subjects, objects, and other clause arguments co-occur with verbs. This is one of the most basic syntactic properties to be acquired in infancy. In the behavioral experiments in **Chapter 2**, I show that infants very early in syntactic development differentiate clause arguments by their grammatical relations, using relations like 'subject' and 'object' to draw inferences about verb meanings. In **Chapter 3**, I show that infants recognize local argument relations as early as 15 months and that they recognize non-local ones, like those in *wh*-questions, at 18 months. I argue that the ability to identify the structure of non-basic clauses like *wh*-questions depends on prior acquisition of verb argument structure, and I investigate this process computationally in **Chapters 4-5**. The computational models tell us about the kinds of mechanisms that could allow learners to acquire both of these phenomena from developing representations of their input. This case study provides a window into how learners leverage their current linguistic knowledge and extra-linguistic abilities to identify the structure of sentences they hear and of their language in general.

Chapter 2: Bootstrapping from Transitivity

Clause transitivity—whether a clause has both a subject and an object—is robustly correlated with clause meaning cross-linguistically, making it a useful cue for learners trying to bootstrap into the grammatical system of their language (Fisher, Gertner, Scott, & Yuan, 2010; Gleitman, 1990; Lidz & Gleitman, 2004; Pinker, 1984). Young learners are sensitive to these correlations: 19-month-olds reliably infer a causative meaning for a verb in a transitive frame, and younger infants do so as well under the right experimental settings (see Fisher et al., 2010 for a review). A key question is how learners represent subjects and objects in their input in order to enable these inferences, and what inferences they draw.

Past experimental results do not tell us how infants represent transitivity: whether as an asymmetrical relation between a subject and an object, or as merely a set of two noun phrases. One influential hypothesis proposes that infants use a heuristic based only on the number of noun phrases in a sentence, taking those to be the arguments and assuming that they match one-to-one the number of participants in an event the clause describes, as we naturally perceive it (Fisher et al., 2010; Lidz & Gleitman, 2004; Naigles, 1990). This "Participant-to-Argument Matching" strategy (PAM) is attractive because it provides strong constraints on verb learning without requiring infants to know which argument is the subject and which is the object. However, because PAM does not characterize the relation between clause structure and meaning in adult grammars, it could only be an initial heuristic to be abandoned over the course of development.

Prior results are also consistent with a more sophisticated bootstrapping strategy. If infants do differentiate the subject and object of a clause, their behavior may be driven by finer-grained inferences about the participant relations of those arguments, rather than argument number. This "Thematic Linking" strategy would link particular grammatical relations to particular thematic relations, e.g. transitive subject to agent and transitive object to patient (e.g. Pinker, 1984; Williams, 2015). Although this strategy requires more sophisticated linguistic knowledge than the number-based alternative, it implies an easier pathway to the adult grammar. Because many linking principles between grammatical and thematic relations hold robustly cross-linguistically, children would not need to abandon them later in development: their learning would be informed by generalizations that are consistent with the grammar they are acquiring.

In this chapter, I empirically distinguish these two hypotheses. When PAM and Thematic Linking would lead to different inferences about verb meanings, which strategy do infants use?

Answering this question requires identifying how many participants infants readily perceive when viewing a particular stimulus event, in the absence of language. Prior work developed a new habituation-based task to norm infants' representations of a scene in which a girl takes a toy truck from a boy (Fig. 3), and confirmed that they view this scene as a 3-participant TAKING (Knowlton, Perkins, Williams, & Lidz, 2018; Perkins, Knowlton, Williams, & Lidz, 2018). This result allows us to ask whether infants think that an event viewed with 3 participants can be described by a novel verb with fewer than 3 arguments:

(1) a. The girl *pimmed* the truck.
b. The truck *pimmed*.

Fig. 3 Girl takes truck from boy

PAM predicts that neither sentence can describe a 3-participant TAKING, because both have the wrong number of arguments. (1a) must describe a 2-participant event involving only the girl and the truck, such as the girl's moving or grabbing of the truck. Likewise, (1b) must describe a 1-participant event involving only the truck, such as the truck's motion. By contrast, Thematic Linking predicts that the transitive clause in (1a) can be a description of a 3-participant TAKING, provided the girl is represented as an agent and the truck as a patient; it does not matter that a third participant (the boy) is not mentioned. However, the intransitive clause in (1b) describes a different perspective on this scene, in which the truck is a patient undergoing a change. An infant who is sensitive to this difference would infer that *pimming* in this context is unlikely to describe TAKING, and more likely to describe the truck's motion.

Results are consistent with the predictions of Thematic Linking. After familiarizing 20-month-olds with the TAKING video described by either (1a) or (1b), we asked them to find another instance of *pimming* in the context of two side-by-side candidates: another TAKING scene, and a MOVING scene in which the girl moves the truck without the boy present. We found that infants who heard transitive *pimming* showed significantly above-baseline preferences for the TAKING video, but infants who heard intransitive *pimming* did not (Fig. 4). Thus, infants appear to allow a 2-argument but not a 1-argument description of a 3-participant TAKING. This tells us that learners at the age previously tested do not rely on PAM as their primary bootstrapping strategy. Instead, they appear to use a more flexible strategy, linking argument and participant relations in a way that is sensitive to the different meanings that transitive and intransitive clauses express.

Fig. 4 Looks to TAKING vs. MOVING video

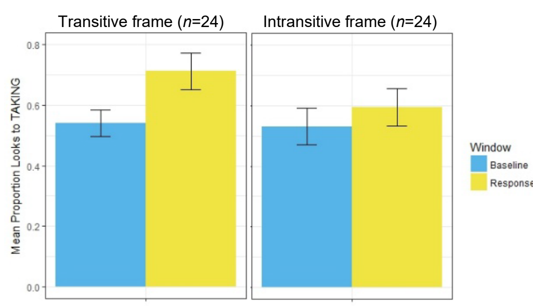

This work has implications for the resources that children have for representing arguments in a clause at early stages in syntactic development. 20-month-olds appear to differentiate subjects and objects in a format that allows for strikingly adult-like inferences about the meanings that those arguments can express. Although these data cannot tell us precisely how those clause arguments are represented, they nonetheless provide further evidence for the independent contribution of both conceptual and linguistic structure in bootstrapping. Infants attempt to relate a sentence to the conceptual representations under which they perceive the world around them, but simultaneously take the linguistic form of the sentence as evidence for which conceptual representation is being tokened (Gleitman 1990). This case study shows how a bootstrapping theory can be enriched by taking into account infants' representational resources in these two domains.

Chapter 3: Representing Transitivity in Non-Basic Clauses

Learners need to recognize subjects and objects not only in their canonical argument positions, but also when they have been displaced, as in (2a). In this sentence, the *wh*-phrase is satisfying the verb's requirement for an object through a non-local *wh*-dependency. These dependencies take various forms cross-linguistically, so learners must discover their form in the specific language they are exposed to. Failure to do so might lead to faulty inferences about where arguments can and cannot occur, and about the specific argument-taking requirements of particular verbs. A learner who does not know that a *wh*-dependency is present in (2a) might mistakenly infer that this clause is intransitive and that the verb *bump* does not require a direct object.

(2) a. Which dog should the cat bump?
    b. *Which dog should the cat bump him?

(3) a. *A dog! The cat should bump.
    b. A dog! The cat should bump him.

Previous literature has noted that so-called non-basic clauses, i.e. those with noncanonical argument realization patterns, pose a chicken and egg problem for early grammar learning (Gleitman, 1990; Pinker, 1984). If learners knew that arguments had been moved in these sentences, they would be able to draw the right inferences about the verb's argument structure. Conversely, if they knew the argument structure of particular verbs, they would be able to infer that these sentences involve some kind of displacement. Which comes first? Using behavioral methods, I show that verb argument structure knowledge precedes the acquisition of non-basic clauses. And, in Chapters 4 and 5, using computational methods, I explore how learners circumvent this chicken-and-egg problem.

Wh-dependencies like those in (2a) comprise about 15% of children's total input (Newport, Gleitman, & Gleitman, 1977). But prior literature argues that infants only begin to represent these sentences in an adult-like way around the age of 20 months (Gagliardi, Mease, & Lidz, 2016). On this hypothesis, apparent comprehension of these dependencies at earlier ages (e.g. 15 months) is not due to adult-like representations, but rather to a "gap-driven" heuristic based on developing verb knowledge. A child who knows that *bump* requires a direct object might notice when this argument is missing after the verb in (2a), and infer that this missing object is the answer to the question. Thus, a 15-month-old may succeed in a preferential looking task simply by detecting an argument gap, without representing *which dog* as a fronted argument. In prior work, we supported this account by showing that 15-month-olds' ability to identify the correct response depends on their vocabulary, a likely index of their verb knowledge (Perkins & Lidz, 2020).

The argument for this gap-driven heuristic is somewhat indirect. In this chapter, I provide more direct evidence for this approach. The gap-driven hypothesis predicts that 15-month-olds should react similarly whenever a transitive verb occurs with a direct object following it, as in (2b) and (3b), because in both cases the verb's requirement for a direct object is locally satisfied. By the same token, they should react similarly to a locally missing direct object, as in (2a) and (3a), despite the fact that the *wh*-phrase satisfies that relation non-locally in (2a). In the adult-grammar, having a direct object gap leads to a different pattern of acceptability in the *wh*-questions in (2) as compared to the declaratives in (3). But learners won't respond according to that pattern until they can represent the *wh*-phrase as an argument of the verb. We asked infants younger than 20 months how acceptable they find these sentences.

To do so, we measured infants' listening time to blocks of auditorily presented sentences, either *wh*-questions (2) or declaratives (3). Infants heard alternating trials of sentences with and without direct objects after the verb while watching an abstract unrelated video, and we recorded looking time towards the screen as a measure of infants' interest in these sentences. 15-month-olds listened longer to sentences with object gaps, regardless of whether they were declaratives or *wh*-questions (Fig. 5). However, 18-month-olds showed different preferences for *wh*-questions and declaratives, listening longer to grammatical sentences of each type. This tells us that 15-month-olds were aware that the verbs in our study required objects and could detect when they were

Fig. 5 Listening time preference for sentences with direct object gaps vs. overt objects

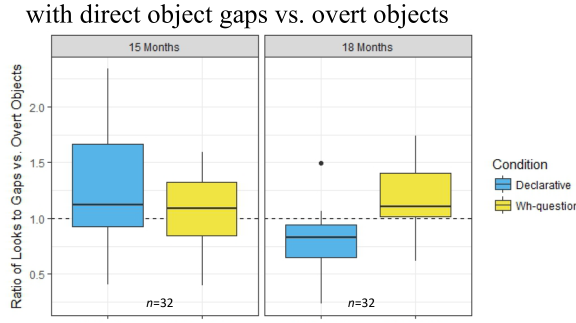

locally missing, but were unaware that the *wh*-phrase satisfied that requirement non-locally. Only the 18-month-olds appeared to represent the *wh*-phrase as an object of the verb.

There are two key findings in this chapter. First, it provides the youngest evidence to date of when infants can represent non-local argument dependencies. Second, it shows that these dependencies are acquired after local argument structure dependencies. By de-confounding syntactic structure and interpretation, our method allows us to find the hallmarks of syntactic knowledge at younger ages than previously tested.

Chapter 4: Filtering Input for Transitivity Acquisition

The results of Chapter 3 provide an answer to the chicken-and-egg question above. English learners appear to acquire verb transitivity knowledge before they recognize moved arguments in non-basic clauses. However, they do not tell us how infants managed to learn verb transitivity in the face of 15% *wh*-questions in their input. Why don't those questions lead learners to think that every transitive verb has an intransitive alternate?

Prior literature proposed that learners need to filter non-basic clauses from their input to avoid faulty inferences about clause structure and argument structure (Gleitman, 1990; Pinker, 1984). These proposals assumed that learners had some way of identifying non-basic clauses in order to filter them out, but our empirical data raises questions about this solution. I offer a new computational solution to this problem. I instantiate a Bayesian learner that filters its input for verb learning, without knowing in advance which sentences to filter out. This learner assumes that it occasionally parses sentences erroneously, and it learns how much of its parses it should treat as signal vs. noise for the purpose of learning verb transitivity. The model only learns from verbs' distributions with and without direct objects, and does not track any additional cues that might signify non-basicness—to this learner, a *wh*-object question is indistinguishable from an intransitive clause. In simulations on child-directed speech, the model learned how to filter erroneous parses in order to categorize the majority of transitive, intransitive, and alternating verbs (Fig. 6).

Fig. 6 Proportions of verbs categorized correctly

	Model	*No-Filter Baseline*	*Oracle*
Transitive	0.67	0.00	0.77
Intransitive	0.83	0.00	0.83
Alternating	0.63	1.00	0.54

By introducing a mechanism for a learner to filter erroneous parses of its input, this model helps answer how learners avoid drawing faulty inferences about grammar and meaning, at stages of development when they do not perceive their linguistic input veridically. Previous approaches either assume that filtering is possible without identifying how (Gleitman, 1990; Pinker, 1984), or presuppose that linguistic perception is veridical (Alishahi & Stevenson, 2008; Barak, Fazly, & Stevenson, 2014; Parisien & Stevenson, 2010; Perfors, Tenenbaum, & Wonnacott, 2010). Crucially, this filtering mechanism can flexibly adapt over the course of a learner's development. As learners more knowledge of their grammar, their syntactic percepts will change: they will be learning from more complete and accurate sentence representations. This means that amount of signal vs. noise in their parses will reduce over time, and they will not need to filter as much of their data for learning. This filtering mechanism may also generalize to other cases in language acquisition where learners must ignore misleading data in order to draw correct inferences about their language (Adriaans & Swingley, 2012; Hudson Kam & Newport, 2009; Pearl & Lidz, 2009; Singleton & Newport, 2004).

Chapter 5: Learning the Surface Forms of Argument Movement

The results of Chapter 4 show us that it is in principle possible for children to identify verb argument structure without parsing non-basic clauses. In Chapter 5, I investigate the mechanisms by which children then identify the structure of those non-basic clauses.

The observation that argument structure is acquired before argument movement may reflect a causal relationship. I pursue the hypothesis that verb argument structure bootstraps the acquisition of non-basic syntax via a three-step mechanism of "Gap-Driven Learning." First, learners use their knowledge of argument structure to detect when predicted arguments are not in their expected positions. Second, they identify what surface properties are correlated with these missing arguments. Third, they infer what kinds of syntactic dependencies are responsible for those correlations.

The learner in this chapter performs the first two steps of this hypothesis. The model tracks not only the distributions of direct objects but also the other morphosyntactic properties of sentences in its input (sentence-initial function words, subject-auxiliary inversion, etc.). It jointly categorizes sentences based on their shared feature distributions and uses its knowledge of verb transitivity to infer which "categories" of sentences contain argument gaps. When it sees a sentence that violates its expectations about verb transitivity, it infers that that sentence contains a gap, and that all sentences in that category do so as well. When the learner was tested on child-directed speech, it out-performed several baselines in identifying sentences with movement, particularly those with object movement (Fig. 7). Thus, it may be possible for a learner to use prior verb knowledge to guide distributional learning in order to categorize non-basic clauses in the target language.

Fig. 7 Proportion of sentences with object movement correctly identified

Verb Type	***Model***	***No-Category Baseline***
Transitive	0.81	0.76
Intransitive*	0.93	0.36
Alternating	0.86	0.50
Total	**0.85**	**0.55**

*Rare or ungrammatical uses of intransitive verbs in transitive frames

This model stops short of the final step of learning under the Gap-Driven Learning hypothesis: inferring which underlying syntactic dependencies are responsible for the observed correlations between argument gaps and surface forms in these sentences. Many spurious correlations exist in the learner's data, and the distributional learning mechanism cannot by itself determine which of these correlations are due to movement vs. other non-movement dependencies. This shows where the limits of distributional learning lie. In order to perform the final step of inference to identify particular movement dependencies in the language, learners may need to make use of not only the formal distributions they observe, but also additional information— potentially from prosody and pragmatics— about the likely dependencies in a given sentence and the ways those dependencies might be realized (Christophe, Millotte, Bernal, & Lidz, 2008; Gleitman, Gleitman, Landau, & Wanner, 1988; Hacquard & Lidz, 2018; Morgan, 1986). This invites further investigation into how the learning mechanisms for grammar acquisition might be fed by informed statistical learning working in concert with knowledge of possible syntactic dependencies, and knowledge of how those dependencies relate to speakers' goals in discourse.

Conclusion

The work reported in my dissertation explores argument structure acquisition as a window into the incremental nature of language learning. I posit that the acquisition of verb argument structure and non-local dependencies like *wh*-movement are tightly connected, and jointly inform our understanding of the representational and learning resources that children bring to the acquisition task. Children may be able to infer aspects of verb meanings on the basis of their syntactic distributions, but these inferences rely on the ability to recognize the core arguments of a clause, and may involve a sophisticated understanding of the relations between conceptual and linguistic structure. Languages vary in how particular syntactic dependencies are realized, but detecting those properties depends on first identifying the language's argument structure profile.

More broadly, the fact that children use their input in the learning process is orthogonal to the question of whether language acquisition is primarily knowledge-driven or data-driven. Both approaches must still explain how knowledge, whether domain-general or domain-specific, guides learners in using their data to draw inferences about their grammar. Domain-general statistical learning may help learners extract signal from noise in their own input representations, and identify formal regularities that are present. But statistical learning may need to be guided by prior grammatical knowledge in order for learners to draw the right generalizations to explain those regularities. In addition, understanding a speaker's communicative intent in using a particular sentence may help constrain the structure and interpretation that a learner assigns to that sentence. But this pragmatic information is not by itself constraining enough to provide a complete parse; a learner must also have available a partial syntactic representation for which this top-down information could be useful. Thus, studying grammatical development provides a window into how linguistic knowledge interacts with the rest of cognition in order to enable learning from input.

Finally, this work has novel implications not only for theories of language acquisition, but also for learning in general. I have proposed novel mechanisms by which a language learner might draw the right generalizations from partial and immature input representations, offering a new perspective on the use of data in learning. Under this approach, learners infer the regularities underlying a particular phenomenon in their input by jointly inferring what data to use in order to best identify those regularities. Learning succeeds through a combination of specific hypotheses that guide the learner to the relevant evidence in her data, and the assumption of a noisy relationship between her data and the hypotheses she is evaluating. The flexibility in this approach invites further investigation into how broadly these mechanisms might generalize beyond language learning. Understanding when learners choose to learn from their input, and when they choose not to learn, may help illuminate how learning succeeds in many other domains in which learners must generalize from incomplete or unreliable representations of data.

Selected References

Adriaans, F., & Swingley, D. (2012). Distributional learning of vowel categories is supported by prosody in infant-directed speech. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *34*.

Alishahi, A., & Stevenson, S. (2008). A computational model of early argument structure acquisition. *Cognitive Science*, *32*(5), 789–834.

Barak, L., Fazly, A., & Stevenson, S. (2014). Learning verb classes in an incremental model. *Proceedings of the Fifth Workshop on Cognitive Modeling and Computational Linguistics*, 37–45.

Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.

Christophe, A., Millotte, S., Bernal, S., & Lidz, J. (2008). Bootstrapping Lexical and Syntactic Acquisition. *Language and Speech*, *51*(1–2), 61–75.

Clark, A., & Lappin, S. (2013). Complexity in language acquisition. *Topics in Cognitive Science*, *5*(1), 89–110.

Fisher, C., Gertner, Y., Scott, R. M., & Yuan, S. (2010). Syntactic bootstrapping. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*(2), 143–149.

Gagliardi, A., Mease, T. M., & Lidz, J. (2016). Discontinuous development in the acquisition of filler-gap dependencies: Evidence from 15-and 20-month-olds. *Language Acquisition*, *23*(3), 1–27.

Gleitman, L. R. (1990). The structural sources of verb meanings. *Language Acquisition*, *1*(1), 3–55.

Gleitman, L. R., Gleitman, H., Landau, B., & Wanner, E. (1988). Where learning begins: Initial representations for language learning. *Linguistics: The Cambridge Survey*, *3*, 150–193.

Hacquard, V., & Lidz, J. (2018). Children's attitude problems: Bootstrapping verb meaning from syntax and pragmatics. *Mind & Language*, *34*(1), 73–96.

Hirsh-Pasek, K., & Golinkoff, R. M. (1996). The intermodal preferential looking paradigm: A window onto emerging language comprehension. In D. McDaniel, C. McKee, & H. S. Cairns (Eds.), *Methods for assessing children's syntax* (pp. 105–124). Cambridge, MA: The MIT Press.

Hudson Kam, C. L. H., & Newport, E. L. (2009). Getting it right by getting it wrong: When learners change languages. *Cognitive Psychology*, *59*(1), 30–66.

Knowlton, T., Perkins, L., Williams, A., & Lidz, J. (2018). *Getting a grip on infants' event representations: participant number in TAKE and PICK-UP*. Poster presented at the The XXI Biennial International Congress of Infant Studies (ICIS), Philadelphia.

Lidz, J., & Gagliardi, A. (2015). How nature meets nurture: Universal grammar and statistical learning. *Annual Review of Linguistics*, *1*(1), 333–353.

Lidz, J., & Gleitman, L. R. (2004). Argument structure and the child's contribution to language learning. *Trends in Cognitive Sciences*, *8*(4), 157–161.

Morgan, J. L. (1986). *From simple input to complex grammar.* The MIT Press.

Naigles, L. R. (1990). Children use syntax to learn verb meanings. *Journal of Child Language*, *17*(2), 357–374.

Newport, E., Gleitman, H., & Gleitman, L. (1977). Mother, I'd rather do it myself: Some effects and non-effects of maternal speech style. In C. E. Snow & C. A. Ferguson (Eds.), *Talking to children: language input and acquisition*. Cambridge: Cambridge University Press.

Parisien, C., & Stevenson, S. (2010). Learning verb alternations in a usage-based Bayesian model. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *32*.

Pearl, L., & Lidz, J. (2009). When Domain-General Learning Fails and When It Succeeds: Identifying the Contribution of Domain Specificity. *Language Learning and Development*, *5*(4), 235–265.
Perfors, A., Tenenbaum, J. B., & Regier, T. (2006). Poverty of the stimulus? A rational approach. *Proceedings of the Cognitive Science Society*, *28*.
Perfors, A., Tenenbaum, J. B., & Wonnacott, E. (2010). Variability, negative evidence, and the acquisition of verb argument constructions. *Journal of Child Language*, *37*(3), 607–642.
Perkins, L., Knowlton, T., Williams, A., & Lidz, J. (2018). *Matching number vs. linking roles: Using 3-participant scene percepts to understand infants' bootstrapping*. Poster presented at the 43rd Boston University Conference on Language Development (BUCLD 43), Boston, MA.
Perkins, L., & Lidz, J. (2020). Filler-gap dependency comprehension at 15 months: The role of vocabulary. *Language Acquisition*, *27*(1), 98–115.
Pinker, S. (1984). *Language Learnability and Language Development*. Cambridge, MA: Harvard University Press.
Singleton, J. L., & Newport, E. L. (2004). When learners surpass their models: The acquisition of American Sign Language from inconsistent input. *Cognitive Psychology*, *49*(4), 370–407.
Williams, A. (2015). *Arguments in syntax and semantics*. Cambridge: Cambridge University Press.
Yang, C. (2002). *Knowledge and learning in natural language*. Oxford: Oxford University Press.