How Nature Meets Nurture: Universal Grammar and Statistical Learning

Jeffrey Lidz[1] and Annie Gagliardi[2]

[1]Department of Linguistics, University of Maryland, College Park, Maryland 20742;
email: jlidz@umd.edu

[2]Department of Linguistics, Harvard University, Cambridge, Massachusetts 02138

Abstract

Evidence of children's sensitivity to statistical features of their input in language acquisition is often used to argue against learning mechanisms driven by innate knowledge. At the same time, evidence of children acquiring knowledge that is richer than the input supports arguments in favor of such mechanisms. This tension can be resolved by separating the inferential and deductive components of the language learning mechanism. Universal Grammar provides representations that support deductions about sentences that fall outside of experience. In addition, these representations define the evidence that learners use to infer a particular grammar. The input is compared with the expected evidence to drive statistical inference. In support of this model, we review evidence of (*a*) children's sensitivity to the environment, (*b*) mismatches between input and intake, (*c*) the need for learning mechanisms beyond innate representations, and (*d*) the deductive consequences of children's acquired syntactic representations.

1. INTRODUCTION

The theoretical frameworks used to approach the problem of language acquisition have traditionally taken two forms. In the input-driven tradition, learning is essentially a form of memory. Learners track patterns of co-occurrence in the environment and store them in a summarized format. The stored summary representation of experience allows for the production and comprehension of sentences beyond those that have been experienced (Elman et al. 1996, Tomasello 2000). Although researchers vary in whether they see the key to generalization as essentially statistical (e.g., Elman et al. 1996, Thelen & Smith 2006, Aslin & Newport 2012) or based in extralinguistic conceptual cognition (Goldberg 2006, Lieven & Tomasello 2008), these approaches share the perspective that generalizing beyond the input is driven by properties of the learner that are not fundamentally linguistic in nature. Instead, generalization relies on cognitive capacities that are applicable across domains and form the basis for understanding learning and social interactions more generally.

The alternative, knowledge-driven, tradition views learning as inference (Chomsky 1965, Pinker 1979, Gleitman 1990, Crain 1991, Lightfoot 1991). The learner, on this view, is endowed with a system of knowledge, typically called Universal Grammar (UG), that specifies a space of possible grammars. Experience provides information that allows for the selection of features from that space in the acquisition of a particular grammar. Domain-specific representations provide the foundation for generalization beyond experience. Experience provides the information that is relevant to the identification of those representations.

These two perspectives differ in two fundamental respects. First, they differ in how they see linguistic representations developing. On the input-driven view, abstract linguistic representations are arrived at by a process of generalization across specific cases. Abstract representations are therefore the output of a learning process that goes from specific to general (Tomasello 2000, Goldberg 2006). On the knowledge-driven view, abstract linguistic representations are part of the inherent linguistic capacity of the learners. Learners do not need to arrive at abstract representations via a process of generalization, but rather use their input to identify the realization of those abstract representations in the surface form of the exposure language (Chomsky 1965, Pinker 1989, Viau & Lidz 2011).

Second, the two views differ with respect to the role that input plays in learning. On the input-driven view, what is learned is a recapitulation of the inputs to the learner (Elman et al. 1996). The acquired representations are a compressed memory representation of the regularities found in the input. Thus, new inputs evoke past experiences via some metric of similarity. New sentences are possible to the extent that they are similar to past experiences. On the knowledge-driven view, the learner searches the input for cues to help choose an abstract representation. The abstract representations themselves are innate and define the cues that learners use to identify them (Fodor 1998, Lightfoot 1999). Once those representations are identified, what the learner carries forward is not a recapitulation of the input. Rather, it is a representation that may bear no obvious relation to the input that triggered it. Thus, this representation allows for specific generalizations that may not be at all similar to the experience of the learner (Chomsky 1981, Snyder 2007).

Where the two approaches do <u>not</u> differ is in acknowledging that input plays a significant role in language acquisition. In both theories, acquisition of a particular grammar is guided to a large extent by experience. Where we expect differences is in the relation between the features of experience and the acquired linguistic knowledge.

In this review, we present a model that incorporates key insights from both traditions, highlighting the link between input and outcomes while maintaining the essential insights of the knowledge-driven perspective on language acquisition. This model separates language

acquisition into three components: (*a*) The intake component identifies the features of the information to which learners are sensitive; (*b*) the UG component identifies the class of representations that shape the nature of human grammatical systems; and (*c*) the inference component identifies how learners combine the intake with UG in selecting a particular grammar.

After briefly describing this model, we review data that show the very powerful role that input plays in shaping language outcomes. Although these data are often used to underwrite input-driven theories (e.g., Goldberg 2006, Huttenlocher et al. 2010), once we understand the components of a UG-based knowledge-driven learning mechanism, we can see why they are fully compatible with knowledge-driven approaches. Next, we show that the statistical sensitivities of the learner are sometimes distinct from ideal-observer measures of informativity, and that these differences may be revealing about the role learners play in selecting relevant input to drive learning. Then, we show that the presence of UG does not remove the need for a learning theory that explains the relation between input and the acquired grammar. Finally, we show that the dissimilarity between the input to the learner and the acquired representations supports a UG-based knowledge-driven approach to language acquisition. In the end, we believe that deconstructing a language acquisition device into separate components for input sensitivity, a rich space of representations and rational statistical inference can explain many key results from both traditions of research in language acquisition.

2. INSIDE THE LANGUAGE ACQUISITION DEVICE

2.1. The Poverty of the Stimulus and the Motivations for Universal Grammar

Language learners acquire knowledge that falls outside of what they experience. This knowledge is evident in the capacity to produce and understand novel sentences, distinguish possible from impossible sentences, and distinguish possible from impossible interpretations of novel sentences (Chomsky 1959, 1965). Consider, for example, the interpretation of the reflexive pronouns in examples 1*a*–*d*.

(1a) Sally knew that Jane put a picture of herself on the wall.
(1b) Sally knew that Jane was very proud of herself.
(1c) Sally knew which picture of herself Jane put on the wall.
(1d) Sally knew how proud of herself Jane was.

In example 1*a*, *herself* is contained in an argument noun phrase (NP) and must be interpreted as referring to Jane and not to Sally. Similarly, in example 1*b*, where *herself* is contained in a predicate adjective, only Jane is a possible antecedent. When *herself* is contained in a fronted argument wh- phrase, as in example 1*c*, its interpretation is less restricted, allowing either Sally or Jane as antecedent. Finally, in example 1*d*, *herself* is contained in a fronted predicate wh-phrase and is again restricted to taking only Jane as antecedent (Huang 1993, Heycock 1995, Barss 2001). In sum, we see an interaction between the phrase containing the reflexive (argument versus predicate) and the position of that phrase (fronted versus in situ) whereby the interpretation is less restricted only in the fronted argument environment (see Leddon & Lidz 2006 for experimental evidence confirming these facts in adults). Where does the knowledge of this interaction come from?

Leddon & Lidz (2006) examined 10,000 wh- phrases in child-directed speech in the CHILDES database (MacWhinney 2000) and found none containing a reflexive pronoun. Consequently, they concluded that whatever English speakers come to know about the possible interpretation of reflexive pronouns in embedded questions such as examples 1*c* and *d*, it must come from

a projection beyond their experience. Moreover, one can imagine many ways to project beyond experience in these cases. It might have turned out that both examples 1*c* and *d* were ambiguous or that neither was. Leddon & Lidz (2006) go on to show experimentally that by age 4 years, knowledge of this asymmetry is in place. The uniform acquisition of this pattern, despite the lack of directly relevant evidence, calls for explanation.

The poverty of the stimulus resides in the mismatch between experience and the acquired grammar. The argument is an invitation for future research that aims to explain that mismatch. The typical response to this invitation is to define the innate representations from which the observed interaction could follow. For example, Huang (1993) relates the predicate–argument asymmetry in examples 1*c* and *d* to the nature of the subject–predicate relation, claiming that predicates carry information about their subjects with them in the syntax. Other responses to this invitation might focus less on the innateness of the representations but rather on general learning mechanisms that could give rise to the relevant interactions, although we are aware of no proposals from outside the knowledge-driven tradition.

UG represents one possible response to the invitation from the poverty of the stimulus, linking diverse sets of facts to a small set of highly abstract representations. These abstract representations drive the language learner's capacity to project beyond experience in highly specific ways. The emphasis of most work in generative linguistics has been the specification of these representations, with less focus on how learners use their experience. But this emphasis should not be confused with a claim that experience is irrelevant to language acquisition.

2.2. Beyond a Specification of Possible Grammars

UG is thus taken to be one component of a knowledge-driven learning mechanism. It defines the character of the acquired representations, which in turn allows the learner to have knowledge of the structure and interpretation of sentences that fall outside of their experience. Importantly, UG must also be embedded in an architecture that allows learners to extract information from the input. This information is used to identify which of the resources defined by UG is applicable to any given sentence or sentence type. Such a model is illustrated in **Figure 1**.
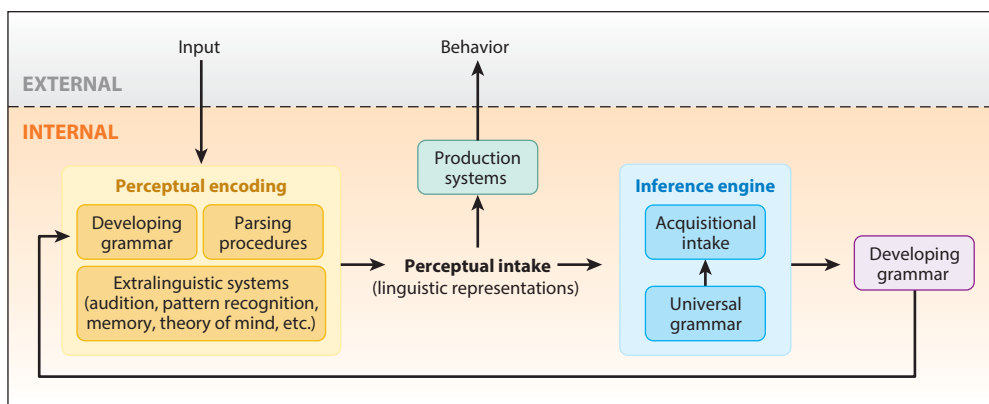

Figure 1

Inside the language acquisition device. The dashed line represents the division between what happens in the external world and what happens within the child's learning mechanism. The arrow from Universal Grammar to acquisitional intake represents the predictive process about what the learner should expect to find in the environment.

This model of the language acquisition device has several components, which we explore in more detail below. The input feeds into a perceptual encoding mechanism, which builds an intake representation. This perceptual intake is informed by the child's current grammar, along with the linguistic and extralinguistic information-processing mechanisms through which a representation from that grammar can be constructed (Omaki & Lidz 2014). To the extent that learners are sensitive to statistical–distributional features of the input, as discussed below, that sensitivity will be reflected in the perceptual intake representations.

The developing grammar contains exactly those features for which a mapping has been established between abstract representations and surface form. A learner's perceptual intake representation does not contain all of the information that an adult represents for the same sentence. If it did, then there would be nothing to learn (Valian 1990, Fodor 1998). What it does contain, however, is information that is relevant to making inferences about the features of the grammar that produced that sentence. This perceptual intake feeds forward into two subsequent mechanisms: the inference engine and the production systems that drive behavior.

The inference engine is based on the idea that every possible grammar defined by UG makes predictions about what the learner should expect to find in the environment (Gleitman 1990, Lightfoot 1991, Fodor 1998, Tenenbaum & Griffiths 2001, Yang 2002, Regier & Gahl 2003, Pearl & Lidz 2009). The acquisitional intake compares those predicted features against the perceptual intake representation, driving an inference about the grammatical features responsible for the sentence under consideration. This inference updates the grammatical knowledge, which is added into the developing grammar. For the purposes of learning, the updated grammar feeds the perceptual processes through which subsequent sentences are analyzed.

The perceptual intake also feeds forward into the systems that support comprehension and behavioral responses. Thus, children's behavior in response to any given sentence is a function of the developing grammar and the mechanisms that yield a perceptual intake representation for that sentence.

In the next sections, we discuss evidence for each of these components of the learning mechanism. We demonstrate first that language learners are sensitive to statistical features of their environment, but that these statistical sensitivities do not always match objective measures of informativity. These observations highlight the role of perceptual intake mechanisms in language acquisition. Second, we show that UG does not by itself define a learning mechanism, but rather must be paired with an inference mechanism that links the representations provided by UG with the data of experience. Finally, we show that some of the statistical sensitivities of the learner are defined by UG, allowing for the inferences from which learners build a grammar of the target language, licensing knowledge about phenomena that fall outside of experience. This last component demonstrates how UG not only contributes a space of possible grammatical representations, but also defines the data that allow the inference engine to connect the perceptual intake to a specific grammar.

3. THE SENSITIVITIES OF THE LEARNER

3.1. Individual Differences in Input and Outcomes

A growing body of research has identified substantial individual differences in the quality and quantity of caregiver speech and corresponding individual differences in children's speech (Huttenlocher et al. 1991, Hart & Risley 1995, Hoff 2003, Rowe 2012). Hart & Risley (1995) examined in-home conversations in families of different socioeconomic status (SES) with children under 2 years of age. They found that in an average week, children of high-SES (professional)

parents hear roughly 215,000 words, whereas children of middle-SES (working-class) parents hear 125,000 words and children of low-SES (public-assistance) parents hear only 62,000 words. The high-SES children heard not only more total words than the low-SES children but also a greater variety of words. Correspondingly, SES-related differences in vocabulary size were evident at the earliest measurements and increased over time. High-SES 3-year-olds had an average cumulative vocabulary of approximately 1,000 words, whereas low-SES 3-year-olds had an average cumulative vocabulary of approximately 500 words. Moreover, Hart & Risley (1995) found that although SES predicted both the parental input and children's vocabulary growth, actual measures of input were stronger predictors of child outcome than was SES.

Similarly, Hoff (2003) examined vocabulary growth as a function of maternal speech characteristics in high- and middle-SES children, finding that SES-related differences in children's vocabulary growth were fully explained by variability in parents' speech. The diversity of word types and the mean length of utterance of the parents accounted for the variability in the rate of vocabulary growth in the children (also see Huttenlocher et al. 1991, Naigles & Hoff-Ginsberg 1998, Hoff & Naigles 2002).

Although these studies clearly demonstrate a role for input quantity in vocabulary development, quantity is not the whole story. Rowe (2012) showed that (controlling for differences in SES) the quantity of speech is the strongest predictor of vocabulary growth in the second year of life but that the diversity of vocabulary used by parents predicts vocabulary growth in the third year of life, and the complexities associated with narratives and decontextualized speech predict vocabulary growth in the fourth year of life. These findings, in particular, suggest that input has different effects on language development as a function of what the child already knows. Consequently, models of language development must be careful to recognize that children extract different information from the input at different points of development, highlighting the variable nature of perceptual intake across developmental time (Gagliardi & Lidz 2014).

Relatedly, Weisleder & Fernald (2013) show that variability in child-directed speech, but not overheard speech, is a significant predictor of vocabulary growth in low-SES Spanish-speaking children in the United States (also see Shneidman & Goldin-Meadow 2012, Shneidman et al. 2013). However, they also show that the effects of input are mediated by differences in language processing. Differences in input affect children's processing speed, which in turn contributes to the rate of vocabulary growth. This result suggests that the information-processing mechanisms through which children understand their input play a critical role in distinguishing the input from the usable input (Omaki & Lidz 2014).

In the domain of syntax, we find similar effects of input on outcomes. Children from higher-SES backgrounds produce more complex utterances in spontaneous speech as toddlers (Arriaga et al. 1998) and at age 5 (Snow 1999), and they perform significantly better on measures of productive and receptive syntax at age 6 (Huttenlocher et al. 2002). Moreover, variation in the syntactic complexity of maternal speech substantially explains SES-related differences in the syntactic complexity of 5-year-old children's speech (Huttenlocher et al. 2002). Huttenlocher et al. (2010) found that higher diversity in caregiver speech predicted several properties of their children's speech, including lexical diversity, diversity within constituents, and diversity of clause types.

This research establishes a link between specific characteristics of the speech that children hear and the speech that they come to produce. At the same time, however, effects of input seem to express themselves in the frequency with which children use complex structures rather than in the initial acquisition of those structures (Huttenlocher et al. 2002), suggesting that perhaps the effect of input has more to do with the mechanisms through which knowledge of grammar is deployed, rather than with that knowledge per se (Weisleder & Fernald 2013).

It is important to recognize, however, that all of these effects are equally compatible with input-driven and knowledge-driven accounts of language acquisition. For the input-driven theorist, what gets acquired is essentially a compressed memory representation of the exposure language. Consequently, properties of children's speech recapitulate properties of the input. For the knowledge-driven theorist, the learner's task is to use the input to identify the surface expression of certain abstract features supplied by UG. Consequently, recognition of those abstract features and the ability to process them efficiently are predicted by their presence in the input. Indeed, although the relevant individual-difference work has not been done within the knowledge-based tradition, the idea that the input should predict the timeline of acquisition is well represented in that literature (Newport et al. 1977, Yang 2002, Valian & Casey 2003, Legate & Yang 2007).

Legate & Yang (2007) examined the acquisition of tense in the productions of child learners of English, Spanish, and French. Children learning these (and many other) languages go through a stage during which they produce "root infinitives," that is, main clauses lacking in tense, unlike in the target language (Poeppel & Wexler 1993, Phillips 1995). However, children learning languages with richer verbal morphology tend to have shorter root infinitive stages than do children learning languages with poorer verbal morphology (Phillips 1995, Guasti 2002).

Further, Legate & Yang (2007) propose to link this variability with UG and the idea that different languages express the same syntactic function differently. In particular, some languages have a morphosyntactic expression of tense, whereas others (e.g., Chinese, Japanese) do not. The learner's task from this perspective is to determine which kind of grammar the language they are exposed to exhibits. Within the languages that express tense in the morphosyntax, there is variability in the degree to which tensed clauses can be distinguished from untensed clauses morphologically. Thus, it is predicted that the age of acquisition of a grammar with tense is a function of the degree to which the input unambiguously supports that kind of grammar over one without tense.

Legate & Yang (2007) show that roughly 60% of child-directed utterances in Spanish unambiguously support a grammar with tense, whereas 39% of child-directed utterances in French do, and only 6% of child-directed utterances in English do. Correspondingly, the root infinitive stage appears to end in Spanish children at around 2 years of age, in French children at around 2 years 8 months, and in English children after 3 years 5 months. Root infinitives are, on this view, an expression of children maintaining the possibility that the target grammar is tenseless. The more the input is consistent with this possibility, the longer these sentences persist in children's speech.

In sum, the existence of correspondences between properties of the input and properties of the acquired grammar is perfectly consistent with both input-driven and knowledge-driven accounts of language acquisition. Both theories predict tight correlations between the exposure language and the timeline of acquisition, although for different reasons. Consequently, we should be careful not to take such data as prima facie evidence for one approach over another. Rather, where the approaches differ concerns the way that learners project beyond their experience, as we discuss below. More generally, one can view children's sensitivity to frequency and statistical–distributional properties of the input as an expression of the intake mechanisms that feed forward for inferences about grammatical structure.

3.2. Statistical Sensitivity in Artificial Language Paradigms

A large body of work explores infants' abilities to extract generalizations about the grammars of artificial languages on the basis of statistical patterns in the data. In particular, learners track statistical information in uncovering dependencies among segments, syllables, words, and phrases.

Saffran et al. (1996) examined the use of adjacent dependencies in word segmentation by exposing infants to sequences of words from an artificial language made up of three-syllable

"words." Although there were no pauses between words in the speech stream, words could be segmented on the basis of transitional probabilities between syllables, that is, the probability of one syllable given the previous one. The language was designed such that transitional probabilities were high within words, but low across word boundaries. Eight-month-old infants exposed to these stimuli learned to segment words, as evidenced by their differentiation of words from nonwords during a test phase.

Although this ability is impressive, it appears to be somewhat limited to the extremely simplified "toy" quality of the artificial language used. Johnson & Tyler (2010) carried out a similar experiment but with words that varied in syllable number. They found that infants were not able to perform word segmentation in this more realistic situation.

Not all added complexity thwarts infants' ability to extract statistical patterns, however. Gómez & Gerken (1999) showed infants' sensitivity to transitional probabilities between words at the sentence level. In their experiment, infants exposed to an artificial language were able to distinguish grammatical strings from ungrammatical ones after only 2 minutes of exposure.

Infants also appear able to go beyond tracking the adjacent dependencies recoverable from transitional probabilities, extracting generalizations about nonadjacent dependencies as well. Gómez (2002) showed that infants can learn dependencies between nonadjacent words, but only when dependencies between adjacent words are not detectable. Further probing this ability, Gómez & Maye (2005) found that including higher variability in the space between the elements in a nonadjacent dependency promotes better learning. Similarly, Lany & Gómez (2008) showed that first acquiring adjacent dependencies makes acquiring nonadjacent dependencies possible.

This body of work highlights children's sensitivity to statistical patterns in their linguistic environment. However, because this research fails to address how the relevant dependencies are represented and how learners might project beyond their input, it is silent with respect to distinguishing input-driven and knowledge-driven models.

Consider, for example, constituent structure in syntax and the data that learners use to acquire it. Constituent structure representations provide explanations for (at least) three kinds of facts. First, constituents provide the units of interpretation. Second, the fact that each constituent comes from a category of similar constituents (e.g., NP, VP) makes it such that a single constituent type may be used multiple times within a sentence.

(2) [$_{IP}$ [$_{NP}$ the cat] [$_{VP}$ ate [$_{NP}$ the mouse]]]

Third, constituents provide the targets for grammatical operations such as movement and deletion.

(3a) I miss [the mouse]$_i$ that the cat ate ___$_i$
(3b) The cat ate the mouse before the dog did [$_{VP}$ ~~eat the mouse~~]

Thompson & Newport (2007) make a powerful observation about phrase structure and its acquisition: Because the rules of grammar that delete and rearrange constituents make reference to structure, these rules leave a statistical signature of the structure in the surface form of the language. The continued co-occurrence of certain categories and their consistent appearance and disappearance together ensure that the co-occurrence likelihood of elements from within a constituent is higher than the co-occurrence likelihood of elements from across constituent boundaries.

Thompson & Newport (2007) show that adult learners can use this statistical footprint in assigning constituent structure to an artificial language with a flat constituent structure. This result was extended by Takahashi & Lidz (2008) and Takahashi (2009), who showed that this statistical

footprint could be used by both adults and 18-month-old infants to acquire a grammar with hierarchically nested constituent structure (also see Morgan & Newport 1981, Saffran 2001). One probe of this knowledge came from participants' ability to distinguish moved constituents from moved nonconstituents. Because only constituents move in natural language, evidence that learners learned which substrings of a sentence could enter into movement relations provided evidence of their knowledge of which substrings were constituents.

But still these observations do not illuminate the role that the statistical footprint plays in acquisition. Does statistical sensitivity feed a memory representation that supports the use of new sentences that are appropriately similar to those in the input, as under an input-driven theory of acquisition? Or does it feed an inferential process whereby a learner with expectations about how phrase structure is organized compares those expectations with the statistics to select a representation, as under a knowledge-driven theory?

Critically, the exposure provided to the learners in Takahashi's (2009) experiments included sentences containing movement. Although the particular sentences tested were novel, they exhibited structures that had been evident during the initial exposure to the language. But where did the knowledge that only constituents can move come from? Did it come from the training, or did it come from expectations about the link between constituency and movement?

To ask this question, Takahashi (2009) created a corpus of sentences from an artificial language. This corpus exhibited a statistical signature of the constituent structure of the language, but provided no evidence of the possibility of movement rules. After exposure to this corpus, both adults and 18-month-old infants successfully distinguished sentences with moved constituents from those with moved nonconstituents, thus displaying knowledge of the constraint that only constituents can move. Given that the exposure did not include any evidence of movement, some of what participants acquired on the basis of statistical information was not itself reflected in the input statistics. This finding is consistent with a knowledge-based approach to learning because the constraint that only constituents can move was not reflected in the exposure language.

Although there are many possible explanations for this effect, among the contenders is the idea that learners know antecedently that human languages exhibit movement relations and that these relations are restricted to applying to constituents. Thus, identifying a string as a constituent makes it a possible target of movement rules. And, for the purposes of this section, the availability of such an explanation illustrates the independence of statistical sensitivity and UG.

Learning mechanisms based on statistical association as well as those based on inferential hypothesis testing predict learners to be sensitive to distributional features in their environment (Frank & Tenenbaum 2011). Thus, the existence of statistical sensitivity plays no role in distinguishing these models (contra Saffran et al. 1996, Huttenlocher et al. 2010). Rather, to distinguish these approaches, we must be careful to differentiate statistical sensitivity from the consequences of that sensitivity. It is only through the consequences for sentences outside of the training data that we can draw conclusions about the role that the input plays in language acquisition (Frank et al. 2013). When we see learners generalizing to structures that are sufficiently dissimilar from their input experience, we find clear candidates for the contribution of UG while still recognizing the statistical sensitivities of language learners.

4. DISTINGUISHING INPUT FROM INTAKE: NOUN CLASSES IN TSEZ

The model in **Figure 1** distinguishes the input from two levels of intake. The input is passed through the perceptual mechanisms of the learner to build a perceptual intake representation. This representation feeds into the inference engine, which selects specified features of that representation (the acquisitional intake) to derive conclusions about the grammatical representations underlying

the target language. In this section, we present evidence that the information used to drive learning is only a subset of what the learner can perceive in the ambient language. Thus, we find a mismatch between objective measures of informativity in the input and the weighting that children assign to different cues during the process of language acquisition (also see Lidz et al. 2003, Hudson-Kam & Newport 2009).

The relevant phenomenon comes from the acquisition of noun classes in Tsez. Noun classes (grammatical gender) can be characterized in two ways: by distributional information external to the noun and by distributional information internal to the noun. On one hand, noun-external distributional information is expressed mainly through other words (e.g., verbs, adjectives) that agree with an NP. The form of the agreement varies depending on the class of the head of the agreeing NP. Such information is highly regular, as it is this behavior that forms the basis for assigning nouns to classes and, consequently, is highly informative. If a noun triggers a particular agreement pattern, this is diagnostic of the class of that noun.

Noun-internal distributional information, on the other hand, consists of properties of the nouns themselves: semantic and phonological features that can (but do not necessarily) correlate with class. This information works probabilistically, as some features provide very good predictions of class while others may make weaker predictions or no predictions at all. Gagliardi & Lidz (2014) measured noun-internal distributional information in Tsez by tagging all of the nouns in a corpus of Tsez child-directed speech for potentially relevant semantic and morphophonological cues. They used decision-tree modeling to determine which features were most predictive of each of the four noun classes (Plaster et al. 2013). Although the features vary in terms of how well they predict a class, at least one feature can predict each class. In general, biological semantic features, which perfectly predict the class of a noun, make the strongest predictions about class. Other semantic features, such as being clothing or made of paper, can also predict the class, but less reliably. Some phonological features, in particular the first and last segments, are also predictive of class, but to a lesser degree than the biological semantic features.

Because the noun-internal distributional information was measured from a corpus of Tsez child-directed speech, it serves as a good characterization of the information Tsez learners are exposed to. By comparing the reliability of an information source with the learners' sensitivity to this information in noun class acquisition, we can determine how learners encode their experience and hence how the inference engine uses this perceptual encoding to drive learning. To the extent that learners are biased to attend to certain information sources, we observe the contribution of intake mechanisms. Furthermore, because the information perceived at different stages of development changes, we can observe the residue of this perceptual intake as a bias toward sensitivity to certain information sources over others. Likewise, among the information that the learner can perceive, a bias in favor of one source over another may reflect the contribution of knowledge-driven acquisitional intake mechanisms.

Gagliardi & Lidz (2014) conducted a classification experiment to determine the degree to which Tsez speakers (and learners) are sensitive to noun-internal distributional information. The words used for classification were either real words that had the predictive features discovered in the corpus (or certain combinations of the features) or nonce words invented to have these features (as well as a set of nonce words without any predictive features). Speakers were tested on nouns with each of the single features, as well as on certain feature combinations making different predictions about class (e.g., a biological semantic cue for class 2 but a phonological cue for class 4). Speakers (adults, older children 8–12 years old, and younger children 4–7 years old) were introduced to test nouns in the form of labeled pictures of objects, and sentences containing agreeing verbs were elicited. As noun class is evident in verbal agreement, speakers' classification of the test nouns was visible through such agreement.

For real words, all groups of speakers tended to correctly classify each word type. However, there were two exceptions to this pattern. First, words with weak (nonbiological) semantic features were less likely to be correctly classified by both groups of children. Second, when the real words' class conflicted with the class predicted by a phonological noun-internal feature, children were more likely to assign the word to the incorrect class, relying inappropriately on the phonological cue.

In classifying nonce words, participants generally displayed appropriate sensitivity to the relevant features. However, this was not the case for all features. First, neither group of children used the weak semantic features at all. Adults were only mildly sensitive to these features, but more so than children. Second, children, especially in the younger group, were more likely to use phonological features instead of semantic features when the two made conflicting predictions, despite the greater informativity of the biological semantic features.

Overall, although speakers were sensitive to both semantic and phonological features that predict class, they used them out of proportion to their statistical reliability, even ignoring certain highly predictive semantic features. The intuition that children's classification behavior does not match the predictions based on the statistical data available in the input is confirmed by the predictions of an optimal naïve Bayesian classifier (Gagliardi et al. 2012). This model predicts classification in line with the semantic features when both semantic and phonological features are present, and provides a representative estimate of adult behavior.

The preference for using phonological information over semantic information likely reflects perceptual intake in the initial stages of noun class learning. Phonological information can be tracked before the learner has assigned a meaning to a given word form, but a learner can be sure of the presence of semantic information only after acquiring the meaning. These differences in how reliably a learner can perceive the different types of features may be reflected in the perceptual intake, causing the learner to more reliably represent phonological information than semantic when building a lexicon. Alternatively, the acquisitional intake may put greater weight on phonological information in forming noun classes, perhaps reflecting an innate bias to treat formal categories as formally (and not semantically) conditioned. This bias in the acquisitional intake would account for the learner's tendency to draw inferences from each feature type unequally.

The apparent blindness to nonbiological semantic features, despite their high reliability, may also stem from the acquisitional intake. Given that noun class systems nearly universally implicate semantic features such as animacy, humanness, or natural gender, it is likely that the learner has expectations about the role of such features in structuring the noun lexicon. The other semantic features implicated in noun class systems vary greatly across languages and thus may be more difficult for learners to perceive as relevant.

In summary, Tsez-acquiring children are sensitive to certain predictive noun-internal features out of proportion to their predictive reliability in the input. This finding highlights the need to distinguish both input from intake and perceptual intake from acquisitional intake in our model of language acquisition.

5. WHY UNIVERSAL GRAMMAR IS NOT SUFFICIENT TO DRIVE LEARNING

It has been recognized at least since Chomsky 1965 that UG is only one piece of a language acquisition mechanism (Pinker 1979, 1989; Wexler & Culicover 1980; Gibson & Wexler 1991; Fodor 1998; Yang 2002; Sakas & Fodor 2012). In some sense this should be obvious, because the elements of UG are abstract but the form of a language is concrete. Thus, learners must discover a mapping between the abstract principles of grammar and their realization in the surface form of the particular language they are acquiring.

Suppose, for example, that UG defines an innate category of reflexive pronoun, paired with constraints on interpretation. These constraints would require that a reflexive pronoun take an antecedent, that the antecedent c-command it, and that there be some domain of locality for finding that antecedent, with room for cross-linguistic variation in defining that domain (Koster & Reuland 1991, Cole et al. 2001). Even with this innate structure, the learner would have to identify which forms in the language fall into that category, to be able to identify structures over which c-command relations obtain, and to define alternative locality domains (Sutton et al. 2012, Pearl & Sprouse 2013). This identification may require a complex set of statistical inferences based on partial information about intended interpretations and the syntactic environment of the pronoun (Orita et al. 2013).

Similarly, suppose that UG defines a narrow class of possible dependency relations, including, for example, the relations (*a*) between a wh- phrase and the variable it binds, (*b*) between a quantifier and a pronoun it binds, (*c*) between an agreeing head and the trigger for agreement, or (*d*) between a head and its complement—each with different properties. Although certain properties of these relations will be a consequence of the UG principles that define them, their expression in the surface form of the language will not be so defined. Consequently, the learner must still use surface-based evidence from the input to identify possible dependency relations and to determine which type of dependency must be invoked for representing any particular case (A. Gagliardi, T. Mease & J. Lidz, manuscript submitted; A. Omaki, N. Orita & J. Lidz, manuscript submitted).

An extreme example of the need for a learning mechanism that includes more than UG comes from the interpretation of quantificational sentences in Korean. Han et al. (2007) examined the relative scope of a universal quantifier in object position in a clause containing negation, as in examples 4*a* and *b*, which illustrate two different morphosyntactic realizations of negative sentences.

(4a) Khwuki monste-ka motun khwuki-lul mek-ci ani ha-yess-ta
    cookie monster-NOM every cookie-ACC eat-CI NEG do PST-DECL
    'Cookie Monster didn't eat every cookie.' (long negation)

(4b) Khwuki monste-ka motun khwuki-lul an mek-ess-ta
    cookie monster-NOM every cookie-ACC NEG eat-PST-DECL
    'Cookie Monster didn't eat every cookie.' (short negation)

In principle, such sentences could be compatible with two interpretations. If the universal takes scope over negation, then this would mean that Cookie Monster ate none of the cookies. If negation takes scope over the universal, then it would mean that Cookie Monster ate fewer than all of the cookies. The latter interpretation is compatible with Cookie Monster eating some cookies, whereas the former is not.

Han et al. (2007) related the interpretations of these sentences to theories of verb placement in the clause, but what is important here is that interpretive judgments in the literature about Korean grammar vary considerably. To probe the source of this variability, Han and colleagues tested adult and 4-year-old speakers of Korean to determine which interpretations of such sentences were licensed. Consistent with the variability in the literature, they found that people differed systematically in their judgments. Roughly half of both populations allowed only the interpretation where the universal takes scope over negation, and roughly half allowed the interpretation where negation takes scope over the universal. This was true for both long and short negation.

Moreover, this variability was not simply due to the difficulty of computing interpretations; when the universal was in subject position, all speakers allowed only the interpretation where the

universal takes wide scope. In addition, in other languages, judgments about these kinds of sentences are stable in both adults and children (Musolino et al. 2000; Lidz & Musolino 2002, 2006). Finally Han et al. (C.H. Han, J. Lidz & J. Musolino, manuscript submitted) found that speakers were consistent in their judgments across multiple testing sessions and that they were consistent in their scope assignments for both kinds of negation in Korean. This pattern supports the view that each speaker controls only one grammar, with the variability in scope judgments following from which grammar that is.

Han et al. argued that this variability results from a sparseness of relevant evidence. Sentences like examples 4a and b are extremely rare (Gennari & MacDonald 2006), and even if they did occur frequently, their intended interpretation in context would be unlikely to be transparent without knowledge of the grammar. Thus, even if UG defines a narrow space of possible grammars, the input to the learner may still be insufficient for deciding which grammar is responsible for the target language. In this case, the absence of evidence leads learners of Korean to pick a grammar at random. In the vast majority of sentences that learners encounter, this choice has no consequences. But in the case of the relative scope of an object quantifier and negation, the effects of this choice can be seen.

To explore this possibility further, Han et al. (C.H. Han, J. Lidz & J. Musolino, manuscript submitted) tested a group of children with their parents. They found again that roughly half of both populations allowed the universal to scope over negation. However, they found no correlation between parents and their children, consistent with the view that grammar selection in this case is not driven by the environment.

With respect to the model in **Figure 1**, the Korean scope data highlight the need for a mechanism distinct from UG that allows learners to select a grammar. In most cases, the selection of an alternative will be driven by the input, and acquisition will be driven by the degree to which learners can identify evidence in favor of one grammar over another (Gibson & Wexler 1991, Frank & Kapur 1996, Yang 2002, Pearl & Weinberg 2007, Pearl 2011, Sakas & Fodor 2012). In the extreme case that there is no relevant evidence, this mechanism chooses a grammar at random from the set of structures licensed by UG. This result illustrates (a) the potential independence of the acquired grammar and the language of the environment and (b) the partial independence of UG from the grammar selection mechanism, given that in this case UG specifies the options but fails to determine the acquired grammar.

6. ON THE DISSIMILARITY OF INTAKE STATISTICS AND ACQUIRED KNOWLEDGE

A critical feature of the knowledge-driven learning model illustrated in **Figure 1** is that it separates (a) the ability to track distributional features in the input and use them to select a grammatical representation from (b) the deductive consequences of those representations. One part of the learning mechanism is inferential, using distributional evidence to license conclusions about the abstract representations underlying the language. The second part is deductive, yielding a wide range of empirical consequences from the acquired abstract representations. To the extent that the empirical consequences of the representations do not resemble properties of the input that gave rise to those representations, we have evidence in favor of this kind of model. In this section, we examine such a case, capitalizing on structural features of ditransitive constructions cross-linguistically.

6.1. A Poverty of the Stimulus Problem in Kannada

Kannada ditransitives exhibit a flexible word order, allowing either order of the dative and accusative arguments. In addition, Kannada also optionally displays a benefactive affix (shown in

boldface below) on the verb in ditransitives. Putting these two features together allows for four possible ditransitives, illustrated in examples *5a–d*.

(5a) hari   rashmi-ge   pustaka-vannu   kalis-id-a
     hari   rashmi-DAT   book-ACC   send-PST-3SM
     'Hari sent the book to Rashmi.'

(5b) hari   rashmi-ge   pustaka-vannu   kalis-i-**koTT**-a
     hari   rashmi-DAT   book-ACC   send-PP-**BEN.PST**-3SM
     'Hari sent the book to Rashmi.'

(5c) hari   pustaka-vannu   rashmi-ge   kalis-id-a
     hari   book-ACC   rashmi-DAT   send-PST-3SM
     'Hari sent the book to Rashmi.'

(5d) hari   pustaka-vannu   rashmi-ge   kalis-i-**koTT**-a
     hari   book-ACC   rashmi-DAT   send-PP-**BEN.PST**-3SM
     'Hari sent the book to Rashmi.'

In addition, we find asymmetries with respect to binding across these constructions (Lidz & Williams 2006, Viau & Lidz 2011). In the benefactive construction, the dative argument can bind into the accusative argument independent of word order (examples *6a* and *b*). However, in the nonbenefactive construction, the dative can bind into the accusative only when the dative comes first (examples *6c* and *d*).

(6a) Q-DAT$_x$ ACC$_x$ V+BEN
     Rashmi pratiyobba hudugan-ige   avan-a   kudure-yannu tan-du-**koTT**-aLu
     Rashmi every   boy-DAT   3SM-GEN horse-ACC   return-PPL-**BEN.PST**-3SF
     'Rashmi returned every boy his horse.'

(6b) ACC$_x$ Q-DAT$_x$ V+BEN
     Rashmi avan-a   kudure-yannu pratiyobba hudugan-ige tan-du-**koTT**-aLu
     Rashmi 3SM-GEN horse-ACC   every   boy-DAT   return-PPL-**BEN.PST**-3SF
     'Rashmi returned his horse to every boy.'

(6c) Q-DAT$_x$ ACC$_x$ V
     Rashmi pratiyobba hudugan-ige   avan-a   kudure-yannu tan-d-aLu
     Rashmi every   boy-DAT   3SM-GEN horse-ACC   return-PST-3SF
     'Rashmi returned every boy his horse.'

(6d) *ACC$_x$ Q-DAT$_x$ V
     *Rashmi avan-a   kudure-yannu pratiyobba hudugan-ige tan-d-aLu
      Rashmi 3sm-GEN horse-ACC   every   boy-DAT   return-PST-3SF
     'Rashmi returned his horse to every boy.'

With respect to the accusative argument binding into the dative, the pattern is reversed. Here, in the benefactive construction the accusative can bind into the dative only when the accusative comes first (examples *7a* and *b*). But in the nonbenefactive construction the accusative can bind into the dative independent of word order (examples *7c* and *d*).

(7a) *DAT$_x$ Q-ACC$_x$ BEN

    *sampaadaka adar-a   lekhan-ige pratiyondu lekhana-vannu kaLis-i-**koTT**-a

    editor   it-GEN   author-DAT every   article-ACC send-PP-**BEN.PST**-3SM

    'The editor sent its author every article.'

(7b) Q-ACC$_x$ DAT$_x$ BEN

    sampaadaka   pratiyondu lekhana-vannu adar-a   lekhan-ige   kaLis-i-**koTT**-a

    editor   every   article-ACC   it-GEN   author-DAT   send-PP-**BEN.PST**-3SM

    'The editor sent every article to its author.'

(7c) DAT$_x$ Q-ACC$_x$ unaffixed

    sampaadaka adar-a   lekhan-ige pratiyondu lekhana-vannu kaLis-id-a

    editor   it-GEN   author-DAT every   article-ACC send-PST-3SM

    'The editor sent its author every article.'

(7d) Q-ACC$_x$ DAT$_x$ unaffixed

    sampaadaka   pratiyondu lekhana-vannu adar-a   lekhan-ige   kaLis-id-a

    editor   every   article-ACC   it-GEN   author-DAT   send-PST-3SM

    'The editor sent every article to its author.'

In sum, we see the interaction of three factors: word order, morphology, and the grammatical function of the quantifier. When the benefactive morpheme is present on the verb, the dative argument behaves as if it is syntactically prominent for binding, hence indifferent to word order. But when the benefactive morpheme is absent, it is the accusative argument that behaves as if it is syntactically prominent for binding, hence indifferent to word order.

With respect to learning, these patterns are not reflected in the input. Viau & Lidz (2011) conducted two large-scale corpus analyses and observed that ditransitive sentences in which one internal argument is a quantifier and the other contains a pronoun that matches that quantifier in phi-features almost never occur. The few cases that do occur would not provide enough variability to license conclusions about which binding configurations are licensed across constructions. Consequently, the data from which children come to acquire these patterns must involve projections from other facts that are more readily available.

6.2. Comparative Syntax as a Window into the Contribution of Universal Grammar

Harley 2002 (building on Freeze 1992, Kayne 1993) makes three important observations about ditransitive constructions cross-linguistically. First, ditransitives differ cross-linguistically in whether the theme or the goal behaves as though it is syntactically prominent for the purposes of binding. Many languages, like Kannada and English, exhibit both goal-prominent and theme-prominent ditransitives. But some, like Irish and Diné, exhibit only theme-prominent ditransitives. Second, goal prominence is typically paired with a more restricted interpretation on the goal, such that it must be a possible possessor of the theme. For example, in both English double-object constructions and Kannada benefactive ditransitives, the goal argument must be interpreted as a possible possessor of the theme argument, whereas such restrictions do not apply to the theme-prominent ditransitives in these languages (or in languages with only theme-prominent ditransitives).

(8a) *The editor sent New York the book.

(8b) The editor sent the book to New York.

(9a) *hari   bangaluur-ige   pustaka-vannu   kalisi-**koTT**-a
      hari   banaglore-DAT   book-ACC   send-**BEN.PST**-3SM
      'Hari sent Bangalore the book.'

(9b) hari   bangaluur-ige   pustaka-vannu   kalis-id-a
      hari   bangalore-DAT   book-ACC   send-PST-3SM
      'Hari sent the book to Bangalore.'

Finally, goal prominence occurs only in languages that have possession constructions in which the possessor is syntactically higher than the possessed. For example, Irish possessives do not allow the possessor to bind a pronoun inside the possessed. Likewise, in Irish ditransitives the goal cannot bind into the theme, illustrating the presence of only theme-prominent structures. In contrast, Kannada possession constructions allow the possessor to bind into the possessee, and that language also displays goal-prominent ditransitives independent of word order, as shown above.

6.3. Two Contributions of Universal Grammar to Learning

In light of these cross-linguistic patterns in ditransitives, Viau & Lidz (2011) identified two potential contributions for UG in the acquisition of the binding facts illustrated in examples 6 and 7. The first contribution, as discovered by Harley (2002), is in defining the space of possible languages, linking the syntax of possession to the syntax of ditransitives. The second contribution is in allowing that syntax to define the acquisitional intake from which statistical inference can proceed.

Regarding the first contribution, UG makes a complex set of facts follow from a single representational parameter concerning the syntax of possession relations. If a language exhibits possession constructions in which the possessed is higher than the possessor, it can recruit that structure in certain ditransitives, treating the goal argument as a possessor and making it syntactically prominent. This syntactic prominence explains the binding asymmetries. Importantly, the level at which the cross-linguistic generalization applies is highly abstract and thus is not detectable directly in the surface form of the language.

Regarding the second contribution, UG helps to define the kind of information that children should use in determining whether a given ditransitive utilizes the goal-prominent or theme-prominent syntax. The surface realization of ditransitives varies considerably cross-linguistically. In English, the two kinds of ditransitives are distinguished in word order (Oehrle 1976). In Kannada, they are distinguished by an affix on the verb (Lidz & Williams 2006) but not by word order. In Spanish, they are distinguished by clitic doubling of the dative argument but not by word order (Uriagereka 1988, Bleam 2001). Given this divergent surface realization, the mapping between the structure and surface form is opaque.

Viau & Lidz (2011) argued that matching the strings with their underlying structures can be achieved by tracking the kinds of NPs that occur as the dative argument in each surface form. Because the dative in a goal-prominent argument is restricted to being a possible possessor of the theme, the kinds of NPs that fill that role are expected to be more restricted. In particular, possessors tend to be animate, so learners should expect a relatively high proportion of animates as the dative argument in a goal-prominent ditransitive. The perceptual intake must therefore consist of a representation of morphological variability, word-order variability, and the grammatical functions of each argument. The acquisitional intake is a representation of the relative proportion of animate to inanimate datives for each of the morphological and word-order variants. If the learner sees a construction that is statistically biased toward animate goals, that skew in the distribution would support the inference that that construction involves the goal-prominent syntax.

Summarizing this section, we have identified two roles of UG in language acquisition: (*a*) explaining the specific ways that children project beyond their input and (*b*) defining the acquisitional intake of the learner, making statistical–distributional evidence informative about grammatical structure. In the case of Kannada ditransitives, the former explains children's knowledge of binding patterns across novel sentence types, whereas the latter explains how observations of the distribution of animate datives can be linked to those binding patterns.

7. CONCLUSIONS

In this review, we outline a model of language acquisition with several important properties. This model integrates the statistical sensitivity of learners to a UG-based, knowledge-driven approach to language learning. In addition, by recognizing a distinction between input and intake, we can understand why learners' sensitivities do not match up perfectly with observer-neutral measures of informativity. In some cases it is the learner's expectations about how languages are structured that define the information in the environment that, in turn, drives learning (Lidz et al. 2003, Viau & Lidz 2011).

Above, we try to emphasize that a theory of UG is not equivalent to a theory of acquisition. UG defines a space of possible representations that, when mapped to appropriate strings, license a rich set of conclusions about sentences that are highly dissimilar from those of experience. But this space of representations only sets the initial conditions for learning. Beyond that, we must also have mechanisms for mapping sentences onto those representations and for defining the environmental inputs that guide that mapping process.

We hope that the model described here will help dissolve the traditional nature–nurture dialectic that has polarized the cognitive science of language, and that it will allow the field to bridge the contributions of the environment with those of the learner. Input undoubtedly plays a critical role in shaping language development. By the same token, learners are able to project far beyond their experience in acquiring a language. The perspective adopted here makes it possible to unify these observations in a single framework.

SUMMARY POINTS

1. Effects on language acquisition of statistical–distributional patterns in the input are expected under any theory of language acquisition, even those with a large contribution of innate knowledge.
2. The hypothesized existence of UG does not remove the need for a theory of learning that explains how experience contributes to language acquisition.
3. We deconstruct the language acquisition device into three parts: intake mechanisms, UG, and inference mechanisms.
4. UG explains the dissimilarity between experience and acquired knowledge. It also allows learners to draw inferences from statistical–distributional evidence to the grammatical representations responsible for producing that evidence.
5. Intake mechanisms explain how learners filter their input to identify critical information to learn from. Input may be filtered by information processing mechanisms, prior knowledge, or expectations associated with particular grammatical hypotheses.
6. Inference mechanisms connect UG with the intake to determine the appropriate mapping from abstract representations to surface form.

FUTURE ISSUES

1. What is the role of online information processing mechanisms in shaping perceptual intake?
2. What are the features of UG that allow it to shape the acquisitional intake?
3. Can hypotheses about the contents of UG be evaluated by examining the predictions they make about the surface features that allow particular grammatical representations to be selected?
4. Does the child's developing, incomplete grammar allow for the identification of unambiguous information to support grammar selection in a way that overcomes problems of overlapping grammatical hypotheses?
5. Can we build explicit computational models of inference that link hypotheses about UG with real input data to better understand the contributions of intake, UG, and inference?
6. How should the models of inference that drive the identification of grammatical features in the target language be evaluated with respect to the algorithms that implement them?

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

LITERATURE CITED

Arriaga RK, Fenson L, Cronan T, Pethick SJ. 1998. Scores on the MacArthur Communicative Development Inventory of children from low- and middle-income families. *Appl. Psycholinguist.* 19:209–25

Aslin RN, Newport EL. 2012. Statistical learning: from learning items to generalizing rules. *Curr. Dir. Psychol. Sci.* 21:170–76

Barss A. 2001. Syntactic reconstruction effects. In *Handbook of Contemporary Syntactic Theory*, ed. M Baltin, C Collins, pp. 670–96. Cambridge, MA: Blackwell

Bleam T. 2001. Properties of the double object constructions in Spanish. In *A Romance Perspective on Language Knowledge and Use*, ed. R Nuñez-Cedeño, L López, R Cameron, pp. 235–52. Amsterdam: Benjamins

Chomsky N. 1959. Review of B.F. Skinner's *Verbal Behavior*. *Language* 35:26–58

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

Chomsky N. 1981. *Lectures on Government and Binding*. Dordrecht, Neth.: Foris

Cole P, Huang CTJ, Hermon G. 2001. *Long-Distance Reflexives*. New York: Academic

Crain S. 1991. Language acquisition in the absence of experience. *Behav. Brain Sci.* 14:597–612

Elman JL, Bates EA, Johnson M, Karmiloff-Smith A, Parisi D, Plunkett K. 1996. *Rethinking Innateness: A Connectionist Perspective on Development*. Cambridge, MA: MIT Press

Fodor JD. 1998. Parsing to learn. *J. Psycholinguist. Res.* 27:339–74

Frank R, Kapur S. 1996. On the use of triggers in parameter setting. *Linguist. Inq.* 27:623–60

Frank R, Mathis D, Badecker W. 2013. The acquisition of anaphora by simple recurrent networks. *Lang. Acquis.* 20:181–227

Frank MC, Tenenbaum JB. 2011. Three ideal observer models for rule learning in simple languages. *Cognition* 120:360–71

Freeze R. 1992. Existentials and other locatives. *Language* 68:553–95

Gagliardi A, Feldman NH, Lidz J. 2012. *When suboptimal behavior is optimal and why: modeling the acquisition of noun classes in Tsez*. Presented at Annu. Conf. Cogn. Sci. Soc., 34th, Sapporo, Jpn.

Gagliardi A, Lidz J. 2014. Statistical insensitivity in the acquisition of Tsez noun classes. *Language* 90:1–32

Gennari SP, MacDonald MC. 2006. Acquisition of negation and quantification: insights from adult production and comprehension. *Lang. Acquis.* 13:125–68

Gibson E, Wexler K. 1991. Triggers. *Linguist. Inq.* 25:403–54

Gleitman LR. 1990. The structural sources of verb meanings. *Lang. Acquis.* 1:3–55

Goldberg A. 2006. *Constructions at Work*. Oxford, UK: Oxford Univ. Press

Gómez RL. 2002. Variability and detection of invariant structure. *Psychol. Sci.* 13:431–36

Gómez RL, Gerken LA. 1999. Artificial grammar learning by one-year-olds leads to specific and abstract knowledge. *Cognition* 70:109–35

Gómez RL, Maye J. 2005. The developmental trajectory of nonadjacent dependency learning. *Infancy* 7:183–206

Guasti MT. 2002. *Language Acquisition: The Growth of Grammar*. Cambridge, MA: MIT Press

Han CH, Lidz J, Musolino J. 2007. Verb-movement and grammar competition in Korean: evidence from quantification and negation. *Linguist. Inq.* 38:1–47

Harley H. 2002. Possession and the double object construction. *Linguist. Var. Yearb.* 2:31–70

Hart B, Risley T. 1995. *Meaningful Differences in the Everyday Experience of Young American Children*. Baltimore: Brookes

Heycock C. 1995. Asymmetries in reconstruction. *Linguist. Inq.* 26:547–70

Hoff E. 2003. The specificity of environmental influence: Socioeconomic status affects early vocabulary development via maternal speech. *Child Dev.* 74:1368–78

Hoff E, Naigles L. 2002. How children use input to acquire a lexicon. *Child Dev.* 73:418–33

Huang J. 1993. Reconstruction and the structure of VP: some theoretical consequences. *Linguist. Inq.* 24:103–38

Hudson-Kam CL, Newport EL. 2009. Getting it right by getting it wrong: when learners change languages. *Cogn. Psychol.* 59:30–66

Huttenlocher J, Haight W, Bryk A, Seltzer M, Lyons T. 1991. Early vocabulary growth: relation to language input and gender. *Dev. Psychol.* 27:236–48

Huttenlocher J, Vasilyeva M, Cymerman E, Levine S. 2002. Language input at home and at school: relation to child syntax. *Cogn. Psychol.* 45:337–74

Huttenlocher J, Waterfall H, Vasilyeva M, Vevea J, Hedges LV. 2010. Sources of variability in children's language growth. *Cogn. Psychol.* 61:343–65

Johnson EK, Tyler M. 2010. Testing the limits of statistical learning for word segmentation. *Dev. Sci.* 13:339–45

Kayne R. 1993. Toward a modular theory of auxiliary selection. *Stud. Linguist.* 47:3–31

Koster J, Reuland E. 1991. *Long-Distance Anaphora*. Cambridge, UK: Cambridge Univ. Press

Lany JA, Gómez RL. 2008. Twelve-month-olds benefit from prior experience in statistical learning. *Psychol. Sci.* 19:1247–52

Leddon E, Lidz J. 2006. Reconstruction effects in child language. In *Proceedings of the 30th Annual Boston University Conference on Language Development*, ed. D Bamman, T Magnitkaia, C Zaller, pp. 328–39. Cambridge, MA: Cascadilla

Legate J, Yang C. 2007. Morphosyntactic learning and the development of tense: a new approach to root infinitives. *Lang. Acquis.* 14:315–44

Lidz J, Gleitman H, Gleitman LR. 2003. Understanding how input matters: the footprint of universal grammar on verb learning. *Cognition* 87:151–78

Lidz J, Musolino J. 2002. Children's command of quantification. *Cognition* 84:113–54

Lidz J, Musolino J. 2006. On the quantificational status of indefinites: the view from child language. *Lang. Acquis.* 13:73–102

Lidz J, Williams A. 2006. On the lack of reflexive benefactives in Kannada. In *Yearbook of South Asian Languages and Linguistics*, ed. T Battacharya, pp. 237–54. Berlin: Mouton de Gruyter

Lieven EV, Tomasello M. 2008. Children's first language acquisition from a usage-based perspective. In *Handbook of Cognitive Linguistics and Second Language Acquisition*, ed. P Robinson, N Ellis, pp. 168–96. New York: Routledge

Lightfoot D. 1991. *How to Set Parameters*. Cambridge, MA: MIT Press

Lightfoot D. 1999. *The Development of Language: Acquisition, Change and Evolution*. Cambridge, MA: Wiley-Blackwell

Morgan JL, Newport EL. 1981. The role of constituent structure in the induction of an artificial language. *J. Verbal Learn. Verbal Behav.* 20:67–85

Musolino J, Crain S, Thornton R. 2000. Navigating negative quantificational space. *Linguistics* 38:1–32

MacWhinney B. 2000. *The CHILDES Project: Tools for Analyzing Talk*. Mahwah, NJ: Erlbaum

Naigles LR, Hoff-Ginsberg E. 1998. Why are some verbs learned before other verbs? Effects of input frequency and structure on children's early verb use. *J. Child Lang.* 25:95–120

Newport EL, Gleitman H, Gleitman LR. 1977. Mother, I'd rather do it myself: some effects and noneffects of maternal speech style. In *Talking to Children: Language Input and Acquisition*, ed. CE Snow, CA Ferguson, pp. 109–50. Cambridge, UK: Cambridge Univ. Press

Oehrle R. 1976. *The grammatical status of the English dative alternation*. PhD thesis, MIT, Cambridge, MA. 289 pp.

Omaki A, Lidz J. 2014. Linking parser development to the acquisition of syntactic knowledge. *Lang. Acquis*. In press

Orita N, McKeown R, Feldman N, Lidz J, Boyd-Graber J. 2013. *Discovering pronoun categories using discourse information*. Presented at Annu. Conf. Cogn. Sci. Soc., 35th, Berlin

Pearl LS. 2011. When unbiased probabilistic learning is not enough: acquiring a parametric system of metrical phonology. *Lang. Acquis.* 18:87–120

Pearl LS, Lidz J. 2009. When domain general learning fails and when it succeeds: identifying the contribution of domain specificity. *Lang. Learn. Dev.* 5:235–65

Pearl LS, Sprouse J. 2013. Syntactic islands and learning biases: combining experimental syntax and computational modeling to investigate the language acquisition problem. *Lang. Acquis.* 20:23–68

Pearl LS, Weinberg A. 2007. Input filtering in syntactic acquisition: answers from language change modeling. *Lang. Learn. Dev.* 3:43–72

Phillips C. 1995. Syntax at age two: cross-linguistic differences. *MIT Work. Pap. Linguist.* 26:325–82

Pinker S. 1979. Formal models of language learning. *Cognition* 7:217–83

Pinker S. 1989. *Learnability and Cognition*. Cambridge, MA: MIT Press

Plaster K, Polinsky M, Harizanov B. 2013. Noun classes grow on trees: noun classification in the North-East Caucasus. In *Language Typology and Historical Contingency: In Honor of Johanna Nichols*, ed. B Bickel, LA Grenoble, DA Peterson, A Timberlake, pp. 153–70. Amsterdam: Benjamins

Poeppel D, Wexler K. 1993. The full competence hypothesis of clause structure in early German. *Language* 69:1–33

Regier T, Gahl S. 2003. Learning the unlearnable: the role of missing evidence. *Cognition* 93:147–55

Rowe M. 2012. A longitudinal investigation of the role of quantity and quality of child-directed speech in vocabulary development. *Child Dev.* 83:1762–74

Saffran J. 2001. The use of predictive dependencies in language learning. *J. Mem. Lang.* 44:493–515

Saffran J, Aslin RN, Newport EL. 1996. Statistical learning by 8-month-old infants. *Science* 274:1926–28

Sakas WG, Fodor JD. 2012. Disambiguating syntactic triggers. *Lang. Acquis.* 19:83–123

Shneidman L, Goldin-Meadow S. 2012. Language input and acquisition in a Mayan village: How important is directed speech? *Dev. Sci.* 15:659–73

Shneidman L, Arroyo M, Levine S, Goldin-Meadow S. 2013. What counts as effective input for word learning? *J. Child Lang.* 40:672–86

Snow CE. 1999. Social perspectives on the emergence of language. In *The Emergence of Language*, ed. B MacWhinney, pp. 257–76. Mahwah, NJ: Erlbaum

Snyder W. 2007. *Child Language*. Oxford, UK: Oxford Univ. Press

Sutton M, Fetters M, Lidz J. 2012. Parsing for Principle C at 30 months. In *Proceedings of the 36th Boston University Conference on Language Development*, ed. AK Biller, EY Chung, AE Kimball, pp. 577–89. Cambridge, MA: Cascadilla

Takahashi E. 2009. *Beyond statistical learning in the acquisition of syntax*. PhD thesis. Univ. Md., College Park. 240 pp.

Takahashi E, Lidz J. 2008. Beyond statistical learning in syntax. In *Proceedings of GALA 2007: Language Acquisition and Development*, A Gavarró, MJ Freitas, pp. 444–54. Cambridge, UK: Cambridge Sch.

Tenenbaum JB, Griffiths TL. 2001. Generalization, similarity, and Bayesian inference. *Behav. Brain Sci.* 24:629–40

Thelen E, Smith LB. 2006. Dynamic systems theories. In *Handbook of Child Psychology*, Vol. 1: *Theoretical Models of Human Development*, ed. W Damon, pp. 258–312. New York: Wiley. 6th ed.

Thompson SP, Newport EL. 2007. Statistical learning of syntax: the role of transitional probability. *Lang. Learn. Dev.* 3:1–42

Tomasello M. 2000. Do young children have adult syntactic competence. *Cognition* 74:209–53

Uriagereka J. 1988. *On government*. PhD thesis, Univ. Conn., Storrs. 555 pp.

Valian V. 1990. Logical and psychological constraints on the acquisition of syntax. In *Language Processing and Language Acquisition*, ed. L Frazier, J Villiers, pp. 119–45. Dordrecht, Neth.: Kluwer

Valian V, Casey L. 2003. Young children's acquisition of wh- questions: the role of structured input. *J. Child Lang.* 30:117–43

Viau J, Lidz J. 2011. Selective learning in the acquisition of Kannada ditransitives. *Language* 87:679–714

Weisleder A, Fernald A. 2013. Talking to children matters: Early language experience strengthens processing and builds vocabulary. *Psychol. Sci.* 24:2143–52

Wexler K, Culicover P. 1980. *Formal Principles of Language Acquisition*. Cambridge, MA: MIT Press

Yang C. 2002. *Knowledge and Learning in Natural Language*. Oxford, UK: Oxford Univ. Press

Contents

Errata

An online log of corrections to *Annual Review of Linguistics* articles may be found
at http://www.annualreviews.org/errata/linguistics