THE ABSTRACT NATURE OF SYNTACTIC REPRESENTATIONS:

CONSEQUENCES FOR A THEORY OF LEARNING

Jeffrey Lidz[1]

University of Maryland

Introduction

The modern study of syntax begins with the observation that people can produce and understand sentences that they have never heard before (Chomsky 1957). From this observation, we conclude that linguistic knowledge must be in the form of a generative symbol system. The term *generative*, in this context, refers to the property that a finite number of symbols describes a potentially infinite set of sentences. When someone knows a language he possesses a system of mental representations and computations that allows for the generation of an unbounded number of novel expressions. Of course, not any novel arrangement of words counts as a sentence and so the system must be constrained to distinguish the possible from the impossible. In short, syntactic knowledge consists of a system of rules and representations that allows for the generation of all and only the sentences that are possible in a given language.

[1] This paper was supported in part by grants from the National Science Foundation (BCS-0418309) and the National Institutes of Health (R03-DC006829). Thanks also to Marilyn Shatz, Sandy Waxman, Colin Phillips, Tonia Bleam, Kristen Syrett, Ann Bunger, Misha Becker, Justin Halberda, Rebecca Baier, and the students in Child Language at Northwestern University in Winter 2005 for helpful input.

The field of generative syntax aims to identify the nature of this symbol system. Lurking behind this agenda, however, lies the more fundamental problem of determining how such a symbol system is acquired. How does a learner exposed to sentences acquire a symbol system? In order to gain traction on this problem, we require an understanding of the kinds of representations and computations that undergird adult languages. But even before turning to syntactic representations, I would like to discuss the nature of symbol systems in general, so as to provide a framework from which to view syntactic systems in particular.

A symbol system consists of three components: the symbols, rules for manipulating these symbols, and a way to interpret these symbols. To take a simple example, in the domain of 4^{th} grade mathematics the numbers and the operators (+, -, x, ÷, =) are the symbols, the rules for manipulating these symbols are the rules of arithmetic, and the interpretation of the symbols is the relation between the symbols and what they represent. We take the computation in (1) to be a valid computation because (a) it follows the rules for manipulating the symbols and (b) the conclusion drawn from manipulating the symbols is in accord with the conclusion drawn from performing the addition with real quantities.

(1) 3 + 4 = 7

Conversely, we do not take the computations in (2) to be valid because, there, the rules for manipulating the symbols are not followed.

(2) a. + 3 = 4 7

b. 3 + 4 = 9

We believe that we understand addition because the results of manipulating the symbols are identical to the results of adding actual quantities.

Turning now to the symbol systems that describe human syntactic competence, the symbols are the linguistic representations that are proposed to explain syntactic structure. The rules for manipulating these symbols are the principles that determine how sentences are built and related to sound and meaning. The interpretation of the symbols is the relation between these symbols and the sentences of a language.

Consider, for example, the grammar fragment in (3). This fragment generates only the sentences in (4). The sentences in (5) are predicted to be impossible.

(3) a) S →NP VP

b) VP → V NP

c) NP → Det N

d) N = cat, dog

e) V = chased, caught

f) Det = the

(4) The cat chased the dog; The cat chased the cat; The dog chased the dog; The dog chased the cat; The cat caught the dog; The cat caught the cat; The dog caught the dog; The dog caught the cat

(5) The dog chased; Chased the cat the dog; Caught; Cat chased dog, …

This isomorphism between the consequences of the operations of the symbol system and the sentences of the language is what is meant by the psychological reality of linguistic theory. We take the symbol system to be a valid representation of syntactic knowledge because manipulations of the symbols yield results that are identical to the results of trying to make the relevant sentences. If it is true that the sentences in (4) are possible and that the sentences in (5) are not, then we take the grammar fragment in (3) to be an accurate representation of (a portion of) the knowledge underlying the ability to speak English.

That said, two points are obvious. First, we understand the symbol system of simple arithmetic considerably better than we understand the symbol systems that underlie human languages. Second, the symbol systems underlying human languages are enormously more complex than the grammar fragment in (3). Consequently, questions about how syntactic systems are acquired lead to the dual challenge of, first, ensuring that the representations that the researcher assumes must be acquired are the right ones, and, second, determining how these representations are actually acquired. Given the general uncertainty and widespread disagreement about precisely what representations and computations characterize syntactic systems, one might conclude that it is pointless to ask about language acquisition. After all, if we don't have a full understanding of *what* is to be acquired, how could we possibly build a theory of *how* it is acquired?

The difficulty of building a theory of how syntactic symbol systems change across development is directly reflected in research on children's syntax. This research has largely focused on the character of children's representations and only rarely has addressed questions of development (see Guasti 2001). Speaking broadly, this research generally finds that children's representations do not differ in kind from those of adults and that in cases where children behave differently from adults, it is rarely because they have the wrong representations. Instead, differences between children and adults are often attributed to task demands (Crain and Thornton 1998), computational limitations

(Bloom 1990; Grodzinsky and Reinhart 1993), and the problems of pragmatic integration (Thornton and Wexler 1999) but only rarely to representational differences between children and adults (Radford 1995). Even in the cases where differences between children's and adults' syntactic representations have been found, very little work has examined the transition to maturity (Wexler 1990).

One clear example of these trends in the literature can be seen in the study of children's subjectless sentences. Bloom 1970 observed that 2- and 3-year-old children learning English produce sentences lacking overt subjects, as shown in (6), despite the fact that the adult language requires a subject in all sentences.

(6) a. play it (Bloom 1970, p. 108)
b. eating cereal
c. shake hands
d. see window

This difference between children and adults has been attributed to a wide range of factors, with debates centering around two questions. First, do children's errors derive from their having nonadult representations or from a difficulty in expressing these representations through their behavior? Second, if the problem is representational in nature, what are the erroneous representations like? Similarly, if the problem is based in linguistic performance, which performance factors stand in the way of adult-like production?

Among those who argue for a representational account of children's errors, we find several distinct hypotheses. One kind of approach takes the typology of natural language syntax as a framework in which to view the errors that children make in acquiring a particular language (Hyams 1986, Crain and Thornton 1998). For example, Hyams (1986) proposes that children drop subjects because their grammars contain features that are appropriate for languages like Italian, but not for English. Others propose similar explanations, but make the child grammar more like Chinese than Italian or English (Hyams 1991, Valian 1991, Rizzi 1994). This style of proposal attributes to children what would be the correct grammatical representations for some language, but not for the language they happen to be learning.

Alternative representational approaches attribute to children incomplete phrase structure representations (Radford 1990; Vainikka 1993). Radford (1990) proposes that subjectless sentences derive from children having built only partial syntactic representations. They have learned how to construct some phrases of their language, but are lacking the feature of syntactic representation that makes subjects obligatory.

Performance accounts, on the other hand, propose that children have the same representations as adults, but that other factors, for example, prosodic properties of pronouns (Gerken 1991) or processing difficulties at the beginnings of long sentences (Bloom 1990), are responsible for the lack of subjects in children's productions.

Because the jury is still out on how these various factors might contribute, singly or in combination, to the production of subjectless sentences, very little research has been able to ask how the transition to maturity takes place. Explanations of the course of development will have a different character depending on what the starting place is and so it is vital to understand the grammatical and extragrammatical contributors to children's errors as a first step in building an explicit model of development.

With these points in mind, it should now be clear that a theory of syntactic development begins with an understanding of children's syntactic representations and the degree to which they are like those of adults. In this paper, I discuss three features of syntactic symbol systems that are shared across all languages. I will show not only that adult grammars exhibit these properties but also that children's grammars do. First, syntactic representations are hierarchically structured. Second, rules of grammar make reference to abstract relations defined over these hierarchically structured representations. Third, certain abstract properties of syntactic representations contribute to the behavior of wide ranges of syntactic phenomena.

Accepting these observations as valid, I argue, places limits on models of the learning systems that drive the acquisition of syntax. Any learning system must, at the very least, yield as its output these abstract syntactic representations. Theories of syntactic development must also be structured in a way that allows for children both to come to the wrong representations and to explain how these errors can be overcome. While children may have syntactic representations that differ from those of the adults around them, in many cases these representations are constrained to conform to the abstract nature described herein.

The chapter proceeds as follows. In section 2, I review evidence showing that sentences are represented as hierarchical structures and that 18-month-old infants are sensitive to this kind of structure. In section 3, we move to a higher level of abstraction showing that the rules of natural language syntax are stated in terms of relations defined over hierarchical representations. Children also show evidence of a sensitivity to these abstract relations. Moreover, children show a sensitivity to abstract hierarchical relations even in constructions for which these relations are not used by adults, suggesting that children are biased to encode syntactic structure in these terms. In section 4, we move to an even higher level of abstraction, showing that certain abstract features of grammatical representation contribute to wide ranges of syntactic phenomena that do not appear to be alike on the surface. In this domain, we find significant correlations in the acquisition of superficially distinct syntactic constructions, suggesting that abstract features drive the acquisition of syntax.

Taken together these three kinds of results indicate that children's knowledge of syntax must be stated in abstract terms from the beginning of development. Abstract symbolic representations are not just the output of syntactic learning but rather form the core out of which knowledge of a particular syntactic symbol system is built. Finally, in section 5, I briefly discuss how these conclusions about children's syntactic representations can guide future research into syntactic development. To preview this discussion, a theory of syntactic development, which traces the evolution of syntactic knowledge throughout infancy and childhood, must begin with the question of how children map their linguistic experience onto the kinds of abstract representations that appear to drive the acquisition of language.

Syntactic representations are hierarchical

Perhaps the most fundamental property of linguistic representations is their hierarchical nature. In this section, I review evidence showing first that this is the correct characterization of adult syntactic systems and second that at the earliest stages of syntactic development, this is also the correct characterization of children's syntactic representations.

A sentence like (7) is not just a string of words, but rather can be represented as a hierarchical tree structure in which words combine to form phrases and these phrases combine to form larger phrases, as in (8).

(7) The lion chased the zebra

(8)

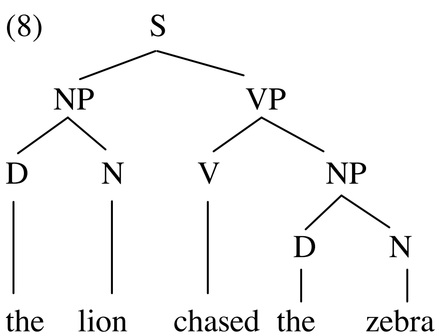

Two kinds of evidence support the claim that sentences have internal hierarchical structure. First, a wide range of phenomena treat the subparts of sentences as units. Either of the phrases labeled NP above can be pronominalized, moved, or coordinated with another phrase of the same type:

(9) a. He chased the zebra

b. The lion chased him

(10) a. The lion is what chased the zebra

b. The zebra is what the lion chased

(11) a. The lion and the tiger chased the zebra

b. The lion chased the zebra and the tiger

Second, the hierarchical nature of syntactic structures can be detected by examining ambiguous expressions like (12):

(12) John saw the man with binoculars

This sentence has two possible interpretations. On the instrumental interpretation, John used binoculars in order to see the man. On the modificational interpretation, what John saw was a man who had binoculars. Given that the sentence has only one set of words, it must be that the arrangement of the words into subgroups (i.e., the sentence's hierarchical structure) explains the ambiguity. In this case, each of the interpretations arises from a unique syntactic representation. The instrumental interpretation of the prepositional phrase in (12) derives from the structure in (13a) whereas the modificational interpretation of the prepositional phrase derives from the structure in (13b).

(13)

a.

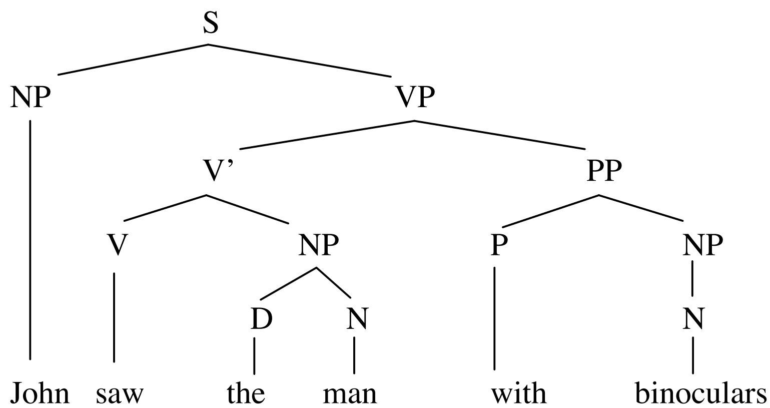

b.

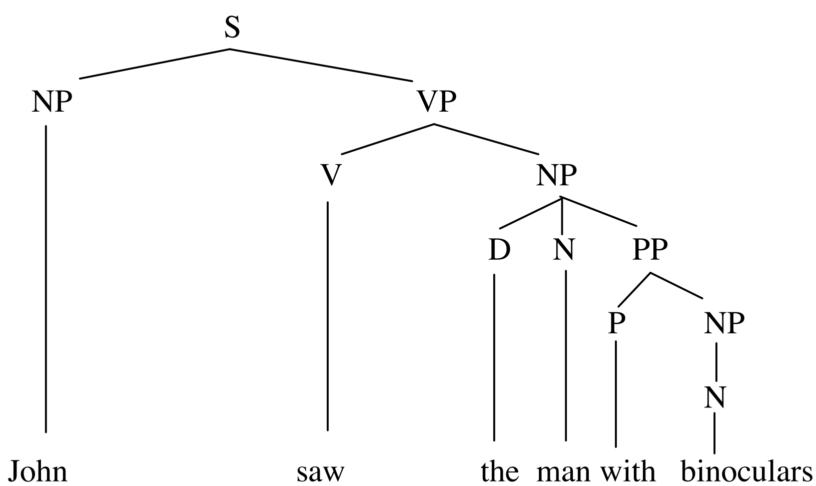

These analyses are supported by the diagnostics we observed above.

(14) a. John saw him with binoculars

b. the man is what John saw with binoculars

c. the man with binoculars is what John saw

Only in (13a) is the string *the man* a constituent and so only the corresponding instrumental interpretation is available when we pronominalize that string as in (14a) or front it, as in (14b). Similarly, when we front *the man with binoculars*, as in (14c), only the modificational interpretation is available since only that interpretation treats that string as a constituent.

Children's knowledge of hierarchical structure

Hierarchical syntactic structure has been observed experimentally in children as young as 18-months of age (Lidz, Waxman and Freedman 2003). Consider two hypotheses for the structure of NP, given in (15). Both would, in principle, be consistent with the input that children receive.

(15) a. Flat structure hypothesis b. Hierarchical structure hypothesis

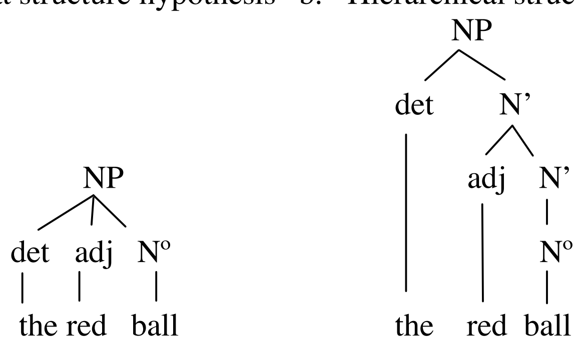

We know, on the basis of anaphoric substitution, that for adults (15b) is the correct representation. In (16), the element *one* refers anaphorically to the constituent [*red ball*].

(16) I'll play with this red ball and you can play with that one.

Since anaphoric elements substitute only for constituents and since it is only under the nested structure hypothesis that the string *red ball* is represented as a constituent (i.e., a single node containing only that string), it follows that (15b) is the correct structure.

Lidz, Waxman and Freedman (2003) tested 18-month-old infants in a preferential looking study (Hirsch-Pasek and Golinkoff 1996) in order to determine whether children represent strings like *the red ball* as containing hierarchical structure. Each infant participated in 4 trials, each consisting of two phases. During the familiarization phase, an image of a single object (e.g., a yellow bottle) was presented three times, appearing in alternating fashion on either the left or right side of the television monitor. Each presentation was accompanied by a recorded voice that named the object with a phrase consisting of a determiner, adjective and noun (e.g., "Look! A yellow bottle."). During the test phase, two new objects appeared simultaneously on opposite sides of the television monitor (e.g., a yellow bottle and a blue bottle). Both objects were from the same category as the familiarization object, but only one was the same color. Infants were randomly assigned to one of two conditions which differed only in the linguistic stimulus. In the control condition, subjects heard a neutral phrase ("Now look. What do you see now?"). In the anaphoric condition, subjects heard a phrase containing the anaphoric expression *one* ("Now look. Do you see another one?").

The assumption guiding the preferential looking method is that infants prefer to look at an image that matches the linguistic stimulus, if one is available. Given this methodological assumption, the predictions were as follows. In the control condition, where the linguistic stimulus does not favor one image over the other, infants were expected to prefer the novel image (the blue bottle), as compared to the now-familiar image (the yellow bottle). In the anaphoric condition, infants' performance should reveal their representation of the NP. Here, there were two possible outcomes. If infants represent the NP with a flat structure, and therefore interpret *one* as anaphoric to the category N^o, then both images would be potential referents of the noun (*bottle*). In this case, the linguistic stimulus is uninformative with regard to the test images, and so infants should reveal the same pattern of performance as in the Control condition. However, if infants represent the NP with a nested structure, and interpret *one* as anaphoric to N', then they should reveal a preference for the (only) image that is picked out by N' (the yellow bottle).

Subjects in the control condition revealed the predicted preference for the novel image, devoting more attention to it than to the familiar image. This preference was reversed in the anaphoric condition, where infants devoted more attention to the familiar than to the novel image. This constitutes clear evidence for the hypothesis that by 18 months, infants interpret *one* as anaphoric to the category N' and thus that they represent the NP with a nested, hierarchical, structure. More generally, we learn from this study, among others (e.g., Valian 1986, Hamburger and Crain 1984, Hirsch-Pasek et al 1986), that at the earliest stages of syntactic development, children's syntactic representations are hierarchically structured, just like those of adults.

Syntactic Relations are defined over hierarchical representations

At this point, we have seen that linguistic representations in adults and infants are hierarchically structured. In this section, I show that the rules and constraints that apply to linguistic expressions also make reference to this structure. Moreover, I show that both adults and children make use of abstract relations defined over hierarchical structures. In addition, we will see that children sometimes apply these relations more liberally than adults, suggesting that their use of these abstract relations is not a consequence of acquisition but rather acts as a guiding force in syntactic development.

As we have seen, various rules of grammar target constituents, indicating that the structural representation of a sentence contains internal hierarchical structure. More abstractly, however, many rules of grammar are defined in terms of relations that are defined over these hierarchical structures. Consider the following:

(17) a. *Hillary* decided that *she* will run for president

b. * *She* decided that *Hillary* will run for president

c. After *she* moved to New York, *Hillary* decided to run for president

The contrast between (17a) and (17b) illustrates that there are constraints determining when a pronoun, like *she*, and a referential expression, like *Hillary*, can refer to the same entity. This constraint is not based on precedence, as can be seen from examining (17c) in which the pronoun precedes the referential expression. The appropriate generalization becomes clear when we examine the structure underlying (17a) and (17b).

(18)

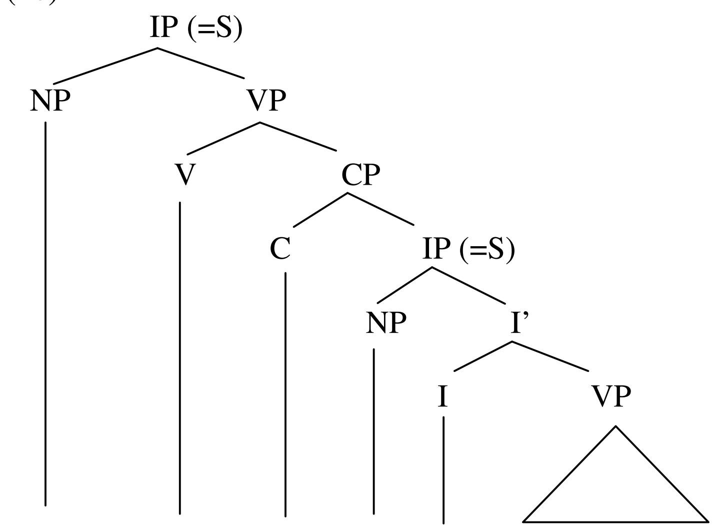

Hillary decided that she will run for president (=17a)

* She decided that Hillary will run for president (=17b)

Informally, what rules out (17b) is a constraint barring a pronoun from being interpreted as identical in reference to an expression that is contained in the smallest (nontrivial) constituent containing the pronoun itself. Here, the smallest constituent containing the subject of the sentence is the entire sentence and so a pronoun in that position cannot be interpreted as coreferential with anything else in the sentence. In the structure underlying (17a), the smallest constituent containing the pronoun is the embedded sentence (IP). Since *Hillary* is not contained within this constituent, coreference with the pronoun is allowed.

This generalization is supported by the observation that if we put the pronoun inside a branching constituent in subject position, coreference is allowed:

(19) *Her* husband thinks *Hillary* will run for president

(20)

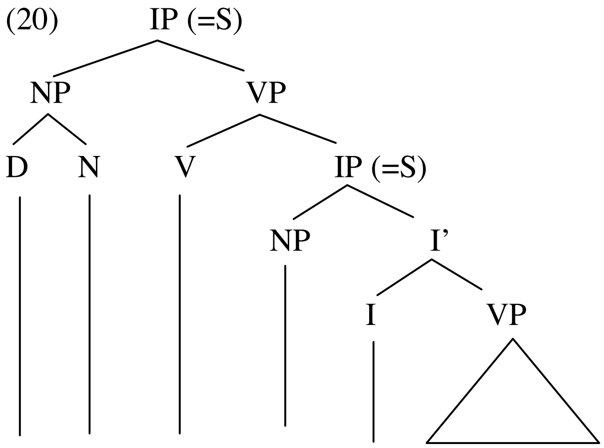

Her husband thinks Hillary will run for president

Here, the smallest constituent containing the pronoun is the subject NP. Since there is no coreferential expression within this constituent, the pronoun can corefer with anything else in the sentence.

The relation expressed in the previous discussion has been formalized under the notions of c-command and binding (Reinhart 1976; Chomsky 1981).

(21) x c-commands y iff

- a) the first branching node dominating x also dominates y
- b) x does not dominate y
- c) $x \neq y$

(22) x binds y iff

- a) x c-commands y
- b) x and y are coreferential

The constraint on coreference, called Principle C by Chomsky 1981, can now be stated as (23):

(23) Principle C: A referring expression cannot be bound

Because the expression *Hillary* is c-commanded by and coreferential with the pronoun in (17b), coreference is blocked by (23). The crucial observation for our purposes is that the explanatory predicate "bind" is based on a relation defined over hierarchical representations, namely c-command. It is only in defining relations over hierarchical representations that we find an explanation for the observed phenomenon, and so we have evidence that these representations exist as part of adults' implicit knowledge of syntax. Importantly, the c-command relation has been found to bear a significant explanatory burden in a wide range of seemingly unrelated phenomena including the interpretation of pronouns (Reinhart 1976; Chomsky 1981), the placement of question words (Fiengo 1977), the computation of syntactic locality (Rizzi 1990), the scope of syntactic operators (May 1985) and the assignment of thematic roles (Chomsky 1981), among others. The prevalence of the c-command relation in syntactic explanation suggests that this relation plays a primitive role in defining syntactic competence (Frank and Vijayshanker 2000; Kayne 1995)

Children's knowledge of c-command

We can now ask whether children's grammars include relations that are defined over abstract hierarchical representations. A wide range of evidence reveals children's knowledge of the c-command relation across different constructions and in different languages (Crain and McKee 1985, McKee 1992, Chien and Wexler 1990, McDaniel, Cairns and Hsu 1991). In this section, I review two kinds of evidence showing that preschool aged children represent relations defined over hierarchical structures as part of their syntactic knowledge.

Children correctly apply the c-command relation: Backwards anaphora and Principle C. Crain and McKee 1985 examined English learning preschoolers' knowledge of Principle C (23), asking whether children know that a pronoun can precede its antecedent but cannot c-command it. We refer to cases in which a pronoun precedes its antecedent as "backwards anaphora." In a truth-value judgment experiment, children were presented with sentences like (24a) and (24b):

(24) a. While he was dancing, the Ninja Turtle ate pizza
b. He ate pizza while the Ninja Turtle was dancing

In this task, participants observe a story acted out by the experimenter with toys and props. At the end of the story a puppet makes a statement about the story. The participant's task is to tell the puppet whether he was right or wrong. Crain and McKee (1985) presented children with these sentences following stories with two crucial features. First, the Ninja Turtle ate pizza while dancing. This makes the interpretation in which the pronoun (*he*) and the referring expression (*the Ninja Turtle*) are coreferential true. Second, there was an additional salient character who did not eat pizza while the Ninja Turtle danced. This aspect of the story makes the interpretation in which the pronoun refers to a character not named in the test sentence false. Thus, if children allow coreference in these sentences, they should accept them as true, but if children disallow coreference, then they should reject them as false. The reasoning behind this manipulation is as follows. If children reject the coreference interpretation, then they must search for an additional extrasentential antecedent for the pronoun. Doing so, however, makes the sentence false. The theoretical question is whether children know that backwards anaphora is possible in sentences like (24a) but not (24b).

Crain and McKee found that, in these contexts, children as young as 3-years-old accepted sentences like (24a), but overwhelmingly rejected sentences like (24b). The fact that they treated the two sentence types differently, rejecting coreference only in those sentences that violate Principle C, indicates that by 3 years of age, English learning children respect Principle C, and therefore that they use the c-command relation in their syntactic computations.

The observation that children use the c-command relation in determining whether two noun phrases are coreferential illustrates that this relation plays a role in defining their syntactic representations. This conclusion raises the question of the origin of the c-command relation and the rules and constraints that make reference to it. The fact that children as young as 3-years of age behave at adult-like levels in rejecting sentences that violate Principle C is often taken as strong evidence not just for the role of c-command in children's representations, but also for the innateness of Principle C itself (Crain 1991). The reasoning behind the argument is that Principle C is a constraint on what is possible in language. It says that a given pairing between certain sentences and certain meanings is *im*possible. But, given that children do not have access to explicit evidence regarding what is *not* a possible form-meaning pairing in their language (cf. Marcus 1993 for a review), their acquisition of Principle C must be driven by internally generated constraints and not by experience alone (cf. Gelman & Williams 1998).

In recent years, however, the possibility that children can learn in an indirect fashion on the basis of Bayesian learning algorithms has gained some prominence (Tenenbaum & Griffiths 2001; Regier & Gahl 2003, inter alia). On this view, the absence of a given form-meaning pairing might be informative about the structure of the grammar. In the particular case under discussion, the learner might observe that she has never seen a situation in which a pronoun c-commands its antecedent and thus remove from the hypothesis space any grammar in which pronouns are allowed to c-command their antecedents. Now, it is important to observe that this kind of approach would have to assume that the learner is representing sentences in hierarchical terms and that she is keeping track of c-command relations among possibly coreferential expressions. Without this assumption, the indirect learner could not even begin. Nonetheless, it does bring up the possibility that the existence of a constraint against certain form-meaning pairings is not by itself evidence of the innateness of that constraint.

Kazanina and Phillips (2001) addressed this issue by looking at the acquisition of backwards anaphora in Russian. Like every language, Russian obeys Principle C. Importantly, however, Russian exhibits a further constraint against backwards anaphora when the pronoun is contained in certain adverbial clauses but not others. These facts are illustrated in (25):

(25) a. Pux$_i$ s"el jabloko, poka on$_i$ čital knigu.
Pooh ate.PERF apple while he read.IMP book
Pooh ate an apple while he was reading a book.

b. Poka Pux$_i$ čital knigu, on$_i$ s"el jabloko.
while Pooh read.IMP book he ate.PERF apple
While Pooh was reading a book, he ate an apple.

c. * On$_i$ s"el jabloko, poka Pux$_i$ čital knigu.
he ate.PERF apple while Pooh read.IMP book
He ate an apple while Pooh was reading a book.

d. * Poka on_i čital knigu, Pux_i s"el jabloko.

while he read.IMP book Pooh ate.PERF apple

While he was reading a book, Pooh ate an apple.

In (25) we see that forwards anaphora is completely free, as in English, but that backwards anaphora is more restricted than in English. In (25c), the pronoun both precedes and c-commands its antecedent and so the sentence is ruled out by Principle C. But, in (25d), the pronoun does not c-command its antecedent, but still the sentence is ungrammatical, unlike its English counterpart. The restriction on backwards anaphora appears to be tied to certain adverbial clauses, as illustrated above with the temporal adverbial *poka*, 'while.' With different temporal adverbials, backwards anaphora is possible, as shown in (26):

(26) Do togo kak ona pereehala v Rossiyu, Masha zhila vo Fancii

before she moved.PERF to Russia, Masha was.living.IMP in France

'Before she moved to Russia, Masha lived in France.'

However the restriction on backwards anaphora in *poka*-clauses is to be formulated, it is clear that it is a Russian-specific constraint, since it does not hold in English (see Kazanina 2005 for details). The existence of two kinds of constraint against backwards anaphora allows us to ask about the origins of Principle C. In particular, the existence of language particular constraints undermines the argument that Principle C must be innate simply because it is a constraint. The existence of constraints like the Russian *poka*-constraint, therefore, makes a Bayesian approach to constraint learning more plausible. Nonetheless, Kazanina and Phillips asked whether children learning Russian demonstrate the same knowledge of Principle C as their English learning counterparts and whether they also demonstrate knowledge of the *poka*-constraint.

These researchers found a developmental dissociation between Principle C and the *poka*-constraint in Russian. While 3-year-olds demonstrated adult-like knowledge for Principle C violating sentences, children at this age appeared not to know the *poka*-constraint. By 5-years-of age, however, the Russian children had acquired the *poka*-constraint (Figure 1).

Because Principle C is a universal constraint but the constraint against backwards anaphora in Russian *poka*-clauses is specific to that language, Kazanina and Phillips suggest that their dissociation in acquisition derives from how they are learned. Principle C is a universal, innate, constraint on possible grammars and so does not need to be learned. Consequently, the effects of this constraint are visible in children at the earliest possible experimental observations. The *poka*-constraint, on the other hand, is specific to Russian and so must be learned from experience, perhaps on the basis of indirect negative evidence, as discussed above.

If the learning of both constraints were based on indirect negative evidence, then we would expect *prima facie* that both would be acquired concurrently. That is, children's experience with sentences that fail to obey either of these constraints is equal: they have encountered no such sentences. Consequently, a learner using indirect negative evidence should acquire both constructions at the same rate, contrary to fact. However, a proponent of the indirect learning approach might argue that the base rate of relevant observations is higher for Principle C configurations than for *poka*-constraint configurations and so acquisition of Principle C precedes acquisition of the *poka*-

constraint. Testing this variant of the indirect learning hypothesis would require a measure of the relative frequency of the two kinds of configurations, a project that has yet to be carried out.

However these constraints are learned, the data on the acquisition of backwards anaphora minimally leads us to conclude that children keep track of abstract relations defined over hierarchical structures, one of the hallmarks of adult syntactic knowledge.

Children over-apply the c-command relation: Quantifier scope. While the data from young children's adherence to principle C cross-linguistically provide compelling evidence that their grammars make use of abstract relations defined over hierarchical structures, we have yet to determine whether this sensitivity to the c-command relation is due to a responsiveness to the data they have been exposed to or whether it derives from a more basic representational constraint on syntactic relations. In this light, Lidz and Musolino examined children's errors in scope assignment preferences in order to ask whether their errors are driven by an attention to c-command where none is called for in the adult language. This study illustrates that children use the c-command relation as a constraint on syntax even when adults do not, suggesting that the c-command relation plays a role in defining how children interpret their linguistic experience. While children's adherence to Principle C shows that they can acquire a constraint that is based on the c-command relation, Lidz and Musolino show that children sometimes build erroneous constraints based on this relation. This observation leads to the conclusion that children's sensitivity to c-command does not derive from their experience, but rather guides them through it. Before we can describe the study motivating this conclusion, however, we must first review the background context motivating it.

Musolino et al. (2000) examined children's interpretations of ambiguous sentences containing quantifiers and negation. Such sentences permit readings which do not directly follow from an isomorphic mapping of surface form to semantic interpretation (Horn 1989; Jackendoff 1972; Büring 1997 among others). Consider the following.

(27) Every horse didn't jump over the fence

a. $\forall x$ [horse (x) $\rightarrow \neg$ jump over the fence (x)]

every horse is such that it didn't jump over the fence (i.e., none jumped)

b. $\neg \forall x$ [horse (x) $\rightarrow$ jump over the fence (x)]

not every horse jumped over the fence (i.e., some jumped and some didn't)

Sentences like (27) are scopally ambiguous. On one interpretation, (27) means that none of the horses jumped over the fence. Here, the universally quantified subject takes scope over negation (abbreviated every>not), as illustrated by the logical representation given in (27a). We call this an isomorphic interpretation because in this case the scope relation between the universally quantified NP and negation coincides with their surface positions. Another possible interpretation of (27) is that not all of the horses jumped over the fence. In this case, negation takes scope over the quantified subject (abbreviated not>every), as shown in the logical representation given in (27b). We call this a non-isomorphic interpretation because here, negation takes scope over the whole sentence, i.e., in a position different from the one it occupies in surface syntax.

Musolino et al. (2000) tested children's comprehension of quantificationally ambiguous sentences. They found that while adults can easily access the non-isomorphic interpretations of such sentences, 4-year-olds systematically assign such sentences an isomorphic interpretation only. This was true also for sentences like (28) in which the isomorphic reading is the opposite from that of (27). While 4-year-olds do not assign sentences like (27) a not>every interpretation, they do assign that interpretation to (28).

(28) The Smurf didn't buy every orange

The finding that children systematically assign examples like (27-28) isomorphic interpretations to conclude that young children, unlike adults, systematically interpret negation and quantified NPs on the basis of their position in overt syntax.

Musolino et al.'s findings, however, do not tell us the nature of the constraint underlying children's resistance to nonisomorphic interpretations. One possibility is that children's overly isomorphic interpretations reflect the linear arrangement between the quantified NPs and negation. Alternatively, children's interpretations may be constrained by the surface c-command relations holding between these elements. These alternatives arise because c-command and linear order are systematically confounded in the materials used by Musolino et al. As can be seen in the tree diagrams below, the subject position always precedes and falls outside the c-command domain of negation while the object position always follows and falls within the c-command domain of negation.

(29)

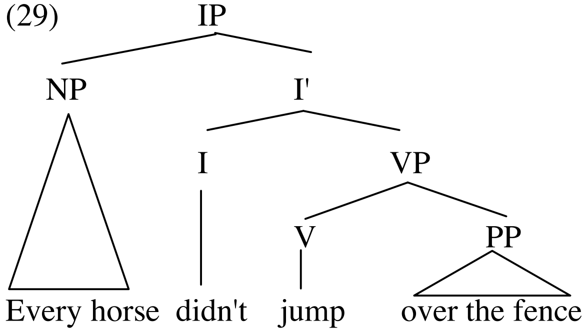

(30)

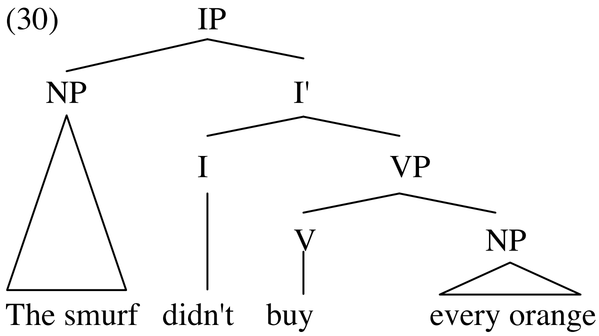

Thus, there is no way to know from these data whether children's behavior is driven by linear order or hierarchical relations.

In order to tease these possibilities apart, Lidz and Musolino (2002) examined English and Kannada-learning children's understanding of sentences like (31).

(31) The smurf didn't catch two guys

Kannada is a Dravidian language spoken by approximately 40 million people in the state of Karnataka in south-western India. The canonical word order in Kannada is Subject-Object-Verb (SOV) and Kannada displays the same kind of scope ambiguities as English with respect to negation and quantified NPs (Lidz in press). These properties are illustrated in (32), which can be interpreted as meaning that it is not the case that I read two books, a narrow scope interpretation of the numeral, or that there are two books that I didn't read, a wide scope interpretation of the numeral.

(32) naanu eraDu pustaka ood-al-illa

I-NOM two book read-INF-NEG

'I didn't read two books'

The crucial difference between Kannada and English is that in Kannada, linear order and c-command are not confounded. Consider the representations in (33). In English, negation both precedes and c-commands the object position. In Kannada, however, negation c-commands the object but *does not* precede it.

(33)

Kannada English

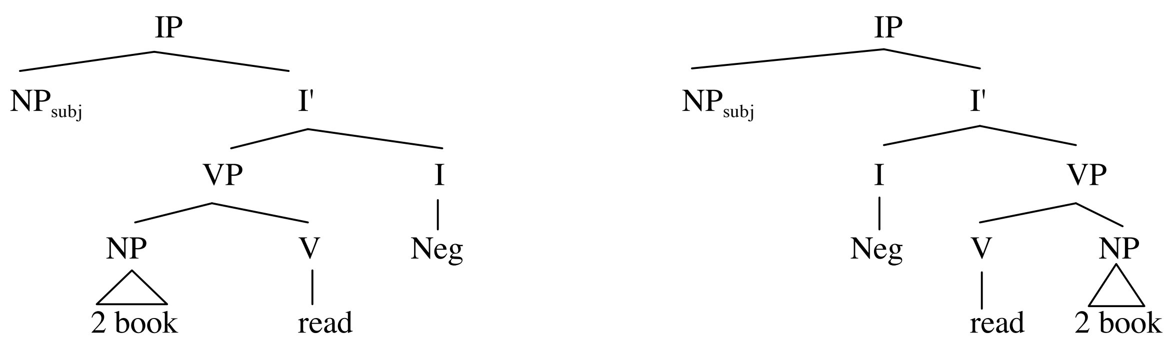

Kannada therefore provides the ideal language to tease apart the contribution of linear and hierarchical relations in children's assignment of scope. To the extent that Kannada-speaking children are restricted to one of the two possible interpretations of sentences like (32) in the way that English-speaking children are, linear order and hierarchical relations make opposite predictions. If children's interpretations of scope relations are constrained by linear order, then Kannada-learning children will display a preference for the wide scope reading of the quantified object with respect to negation. On the other hand, if children's interpretations of scope relations are constrained by c-command relations between negation and the quantified object, then Kannada-learning children will display a preference for the narrow scope reading of the object with respect to negation.

Lidz and Musolino found that children assigned a narrow scope reading to the numeral independent of the language being acquired (Figure 2). This finding illustrates that children’s scope assignment preferences are determined by the hierarchical relation of c-command and not by linear order. Crucially, it is children’s nonadult behavior that allows us to see that their representations are of the same character as adults'. The fact that children's scope interpretations differ from adults' enabled us to determine that their limitations derive from the same kinds of representations as we find in mature linguistic systems. Children use the c-command relation defined over hierarchical representations in determining scope, even though adult grammars are more flexible in this regard.

More broadly, we have determined that relations defined over hierarchical structures play an explanatory role not only in the characterization of adult knowledge but also in the characterization of children's knowledge. The fact that children apply the c-command relation in structures to which they shouldn't suggests that children prioritize this relation in their syntactic representations, using it to guide their acquisition.

Links between cross-linguistic variability and syntactic acquisition

Although there is great diversity in human languages, languages do not vary at random (Greenberg 1963). Certain surface features of a language are predictive of other features. For example, if verbs precede their objects in some language, then auxiliary verbs precede main verbs in that language. Likewise, if main verbs follow their objects, then auxiliary verbs follow main verbs (Baker 2003). These kinds of typological generalizations point towards abstract principles of grammar that unify constructions with apparently unrelated surface properties. These principles are stated in terms of abstract properties of syntactic representations, so that one point of variability in the representations can lead to a wide range of surface differences (Chomsky 1986, Baker 2003).

If this perspective on cross-linguistic variation is correct, then it makes a strong prediction about language acquisition: if two superficially distinct constructions share a certain piece of representational structure, then the acquisition of those two constructions should be significantly correlated (Snyder 1995, 2001; Snyder and Stromswold 1997; Baker 2005). Importantly, if this prediction is borne out, then it supports a view of children's syntactic knowledge as being stated over abstract properties of representations.

The prediction has been tested in a range of studies conducted by William Snyder and his colleagues. Snyder and Stromswold (1997) examined the acquisition of a wide range of "complex predicate" constructions, given in (34):

(34) a. John painted the house red. (resultative)

b. Mary lifted up the box/lifted the box up. (verb-particle)

c. Fred made Barney leave. (*make*-causative)

d. Fred saw Barney leave. (perceptual report)

e. Bob put the book on the table. (*put*-locative)

f. Alice sent the letter to Sue. (*to*-dative)

g. Alice sent Sue the letter. (double-object dative)

These constructions are all alike in that they involve multiple predications but only a single tense specification. For example, (34a) contains the predicates *paint* and *red*, (34b) contains the predicates *lift* and *up*, etc. Languages differ with respect to whether they exhibit this set of "complex predicate" constructions. English and other Germanic languages typically exhibit all of these constructions, whereas the Romance languages lack direct counterparts of them.

The possibility that these constructions form a family related by an abstract syntactic property, and that this property plays an explanatory role in the acquisition of these constructions, was examined by Snyder and Stromswold (1997). These researchers examined longitudinal transcripts of the spontaneous speech of 12 children from the CHILDES database (MacWhinney and Snow 1985). For each of the constructions in (34b-g), Snyder and Stromswold determined the first clear use of that construction and used that as a measure of acquisition. By this measure, every child in the sample acquired all of these constructions as a group. In addition, Snyder and Stromswold showed that the tight correlations between age of acquisition of these constructions were not a consequence of general syntactic ability, MLU, or specific lexical properties. These results suggest that the unit of acquisition in syntax is not the construction, but rather a more abstract component of representation that ties together constructions with different surface properties.

Snyder (2001) argues that these constructions share the property of involving syntactic compounding of the two predicates at an abstract level of representation. On this view, (34a) is syntactically represented as in (35), although the tight relation between the two predicates is ultimately obscured in the word order. (Larson 1991)

(35) [john [[painted red] [the house]]

As support for this view, Snyder notes that only languages with productive noun-noun compounding exhibit complex predicate constructions. English has productive compounding, as in (36), whereas Spanish does not, as in (37). In languages like Spanish, noun-noun compounds are formed only through conscious coinage.

(36) English
a. banana box (= a box for bananas, a banana shaped box, etc.)

b. box banana (= a banana shaped like a box, a banana that came from a box, etc.)

(37) Spanish
a. * platano caja

banana box

b. * caja platano

box banana

Similarly, English but not Spanish has complex predicate constructions like those in (38).

(38) a. John hammered the iron flat

b. Juan golpeò el hierro (*plano)

c. John lifted the box up

d. Juan levantò la caja (*arriba)

In a typological survey drawing from twelve language families, Snyder (2001) found that only languages that allow productive noun-noun compounding also allow resultatives and separable particle constructions. This pattern indicates that noun-noun compounding is a necessary condition for complex predication. The abstract feature of syntactic representations that licenses noun-noun compounding is therefore implicitly tied to the possibility of complex predicate constructions, explaining why these constructions cooccur cross-linguistically.

Snyder (2001) also showed that the acquisition of complex predicate constructions, like resultatives and separable particle constructions, is directly tied to the acquisition of noun-noun compounding. Looking again at first clear use, the age of acquisition of compounding and complex predicate constructions was shown to be nearly identical in almost all children, ranging from 1.8-2.6 years of age. Putting this together with the typological data, we have strong evidence that there is a representational link between compounding and complex predicate constructions.

An additional test of the relation between compounding and complex predicates was performed by Sugisaki and Isobe (2000). These authors observe that in Japanese, novel compounding is acquired later than in English, between 3- and 4-years of age. This relative delay in acquisition allowed them to test experimentally for a relationship between resultatives and compounding in three- and four-year old children. As a test of the compounding parameter, these authors asked whether only Japanese children who can produce novel noun-noun compounds can also understand resultatives in an adult-like fashion.

Each child received an elicited production test of novel compounds (as in 39), and a comprehension test of transitive resultatives (as in 40b) as compared to transitive sentences with attributive adjectives (40a).

(39) *kame pan* 'turtle bread' (i.e., bread in the shape of a turtle)

(40) a. Pikachu-wa aka-i isu-o nutte-imasu. (attributive)

'Pikachu is painting the red chair.'

b. Pikachu-wa aka-ku isu-o nutteiru. (resultative)

'Pikachu is painting the chair red.'

As illustrated in Table 1, there was a significant contingency between passing the test on resultatives and passing the test on compounding. Performance on one test predicted performance on the other.

These results lend further support to the hypothesis that compounding and complex predicate constructions share a certain piece of representational structure. Learning that abstract piece of syntactic structure leads to the concurrent acquisition of superficially distinct constructions.

More generally, this kind of data points towards a theory of children's syntactic knowledge that is stated in terms of abstract syntactic representations. Such a theory entails that the object of acquisition for a syntactic learner is this kind of abstract representation and not something more closely tied to surface forms. This kind of data also invites more thorough integration of work in comparative syntax with work in syntactic acquisition. Identifying the range and limits of cross-linguistic variation can provide useful hypotheses about children's representations and can thereby restrict the class of learning theories that are consistent with a given phenomenon (Lidz, in press).

Consequences for learning

In the preceding sections, we have identified three properties of syntactic representations that are continuous across children and adults. First, children's syntactic representations are hierarchical from the earliest stages of acquisition. Second, children compute and keep track of abstract relations defined over hierarchical structures, like the c-command relation, and use such relations in building their grammatical knowledge. Third, children's syntactic representations contain abstract features that lead to concurrent acquisition of superficially distinct constructions.

Having identified these properties puts us in position to ask about syntactic learning. At least three possible explanations for this continuity can be developed. First, the strong nativist position would hold that the reason we find this kind of representational continuity is that this is simply the vocabulary over which the learning algorithm is defined. On this view, when children hear sentences they automatically encode them in terms of abstract hierarchical representations with consequences for wide ranges of surface phenomena (Chomsky 1975, Wexler 1996, Baker 2003). A second, more weakly nativist, position would hold that this kind of representational continuity results from constraints on computations. The representations themselves are not innate, but rather they are the only possible output of a certain kind of learning algorithm (Newport 1990, Newport and Aslin 2000). A third option would be even less nativist in its orientation, taking the unity of linguistic representations across children and adults to be a consequence of the kinds of information that can be extracted from linguistic experience by a general purpose learning mechanism across time (Elman 1993; Thomas & Karmiloff-Smith 2005). On this view, certain features of a grammar are unable to be noticed by the learner until other features are already in place. Moreover, the nature of what can be learned at a given stage of development constrains the kinds of representations that are built at later stages of development. On this view, the uniformity of syntactic representations across languages and across populations results from a cascading learning procedure that becomes more constrained as it develops.

Various combinations of these proposals can and have been put together. For example, some of the mathematically inspired language learnability work (Wexler and Culicover 1980; Lightfoot 1991) proposes that limited computational resources guide the learner through an innate space of possible representations. Similarly, proposals concerning parameter setting (Baker 2005, Dresher and Kaye 1991) and experimental work on vocabulary acquisition (Gleitman et al 2005, Waxman and Lidz, in press) argue that learning is guided through an innate hypothesis space on the basis of the degree to which prior representational commitments must be made before some feature of the language can be learned. On this view, the order of acquisition of certain features is driven by changes in the learner's representational space. Once a given feature is acquired, it makes available new resources for the acquisition of subsequent features.

Whatever the appropriate learning theory for syntax is, understanding the kinds of representations that make up adult languages and the kinds of representations children have at a given stage of development places limits on hypotheses concerning this theory. At a minimum, a researcher asking about the acquisition of some syntactic phenomenon must ask three questions. First, what is the representation that is acquired for that phenomenon? Second, how is this phenomenon expressed in speech to children? Third, how could the representation be acquired on the basis of experience?

Notice that even the most staunch nativist must provide an answer to the last question. Even if we view learning as triggering (i.e., the representations are fully specified at birth but need only to be activated by experience), the learner faces the problem of identifying which forms in the input correspond to her innate syntactic representations. Simply having a representational vocabulary is insufficient for acquisition. Rather, the learner must be able to link data from experience with those representations (Pinker 1989, Tomasello 2000).

A clear example of how experience could play a role in identifying innate representations comes from recent work by Misha Becker (in press), which examines the acquisition of raising and control structures like those in (41).

(41) a. Chris seems to be happy. (raising)

b. Chris wants to be happy. (control)

The problem presented by these kinds of sentences is that although their structural representations are distinct, their surface properties are identical. In both sentences the main clause has an overt subject and the embedded infinitival clause does not. Despite these similarities, a number of syntactic tests illustrate that these kinds of sentences have distinct structural descriptions. For example, only the subject of raising verbs can be a pleonastic:

(42) a. It seems to rain every time I go to Paris

b. * It wants to rain every time I go to Paris

c. There seems to be a riot every time I go to Paris

d. * There wants to be a riot every time I go to Paris

Similarly, phrasal idioms can be separated by raising verbs but not by control verbs:

(43) a. The shit seems to hit the fan every time I go to Paris

b. * The shit wants to hit the fan every time I go to Paris

These differences (among others) derive from the fact that the subject of a control verb is a participant in the event denoted by that verb (in addition to the event denoted by the embedded verb) whereas the subject of a raising verb is a participant only in the event denoted by the embedded verb. If learners knew the meanings of the relevant verbs upon hearing a sentence like (41a) or (41b), then they would know what structure to assign. However, it is unlikely that learners could identify the relevant meanings independent of the syntactic structures they encounter the verb in (cf. Gleitman 1990). Consequently, upon hearing these kinds of sentences with previously unknown verbs, the learner is provided with no information about whether to assign a raising or a control structure, since the surface realizations of these structures are alike.

With these surface similarities in mind, Becker (in press) proposes that there are certain probabilistic cues to these structures that derive from, but are not entailed by, the structural differences between raising and control. For example, because control verbs exert selectional restrictions (typically having to do with animacy or volitionality) on their subjects but raising verbs do not, Becker proposes that learners should treat the animacy of the subject as a probabilistic cue to the control analysis of a sentence with an infinitival complement. She goes on to show that adults and 5-year-old children do treat the animacy of the subject as a cue to the control structure, despite the fact that it is only probabilistically present across the whole language. Interestingly, 3- and 4-year-old children were less sensitive to animacy as a cue to control than were 5-year-olds, suggesting that the initial analysis of a novel infinitival complement taking verb is the raising analysis and that this analysis is overcome over time on the basis of experience with the relevant probabilistic cues.

This kind of analysis is important because it reveals that learners who are equipped with a set of possible syntactic structures might be able to predict what kinds of sentences are likely to instantiate those structures. Crucially, these predictions need not be entailments of the structure. In Becker's example, animacy is not a trigger for control, since animate subjects are possible with both raising and control verbs, but it is a cue for control since control verbs are more likely to insist that their subjects be animate. This cue-based theory of learning illustrates how learners with rich innate structures might go about identifying the strings that realize those structures (see Fodor 1998, Dresher and Kaye 1991, Lightfoot 1999, Yang 2004 for related proposals).

Speaking more generally, the solution to the problem of how a grammatical system becomes increasingly specified on the basis of experience will require an understanding of how the sentences falling outside of the current representational space are encoded and how the existing representations can be updated on the basis of that experience (Valian and Casey 2003). Experience, not surprisingly, plays a critical role in understanding syntax acquisition. But, the learner's reliance on the input may be most successful when learners are constrained in the kinds of representations they consider. In sum, an emphasis on syntactic representations as both the target of acquisition and as a contributor to syntactic development will allow us to more precisely formulate the significant role of input in syntactic acquisition.

The general framework outlined here, with its emphasis on children's representational systems for syntax across time, makes it possible to bridge developmental and theoretical research in language acquisition. As noted above, very little work in syntactic acquisition, if any, has attempted to ask what drives syntactic change in actual learners. This dearth of research is probably explained in part by the difficulty of formulating hypotheses about representational change. As noted by Fodor 1998, there is an inherent paradox in syntactic development. The problem is that learners can only parse strings for which they have representations. For those sentences there is nothing to be learned. On the other hand, if the parser can't assign a structure to a sentence, then how could the learner possibly learn from it? A cue-based learning approach like that described above gives one possible solution to this paradox since a sensitivity to the surface cues that point to a structure could aid parsing even for those sentences that fall outside of the learner's current grammar.

A second contributor to the scarcity of developmental hypotheses in language acquisition is the problem of identifying the child's representations at a given stage, as noted in the introduction to this chapter. Only when we have some understanding of the successive representational systems built by children over time can we begin to ask about the mechanisms that drive the elaboration of syntactic knowledge. I believe we have now reached a stage where we have a great deal of information about the character of children's syntactic representations. As we have seen in this chapter, these representations are highly abstract, with rules and relations defined over hierarchical structures playing a critical role in a wide range of superficially distinct constructions. Future work in the field must take advantage of these conclusions in order to ask both how these representations are learned and how they contribute to learning.

References

Baker, M. (2003). *The Atoms of Language*. Basic Books, NY.

Baker, M. (2005). Mapping the Terrain of Language Acquisition. *Language Learning and Development,* 1, 93-129.

Baker, M. (2005). *Lexical Categories: Verbs, Nouns and Adjectives*. Cambridge, Mass.: Cambridge University Press.

Becker, M. (in press). There began to be a learnability problem. *Linguistic Inquiry*.

Bloom, L. (1970). *Language development: Form and function in emerging grammars*. Cambridge, Mass.: MIT Press.

Bloom, P. (1990) Subjectless sentences in child language. *Linguistic Inquiry* 21, 491-504.

Büring, D. (1997) The great scope inversion conspiracy. Linguistics and Philosophy 20: 175-194.

Chien, Y.C. and Wexler, Kenneth (1990) . Children's knowledge of locality conditions in binding as evidence for the modularity of syntax and pragmatics. *Language Acquisition* 1, 225–295.

Chomsky, N. (1957) *Syntactic Structures*. Mouton.

Chomsky, N. (1981). Lectures on government and binding. Dordrecht, the Netherlands Foris.

Chomsky, N. (1975). *Reflections on language*. New York: Pantheon.

Chomsky, N. (1986) *Knowledge of Language*. Praeger.

Crain, S. 1991. Language Acquisition in the absence of experience. *Behavioral and*

Brain Science 14, 597-650.

Crain, S. and McKee, Cecile (1985). The acquisition of structural restrictions on anaphora. Proceedings of NELS 15, 94 –110.

Crain, S. and R. Thornton (1998) Investigations in Universal Grammar: A Guide to Research on the Acquisition of Syntax and Semantics, The MIT Press, Cambridge, MA.

Dresher B.E., and J.D. Kaye. (1990). A computational learning model for metrical phonology. *Cognition,* 34, 2,137-195.

Elman, J.L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, 48, 71-99.

Fiengo, R. (1977). On trace theory. *Linguistic Inquiry*, 8, 35–61.

Fodor, J. D. (1998). Unambiguous triggers. *Linguistic Inquiry,* 29, 1-36.

Frank, R. and K. Vijay-Shanker. (2001). Primitive C-Command. *Syntax,* 4, 164-204.

Gelman, R. and Williams, E. (1998). Enabling constraints for cognitive development and learning: Domain specificity and epigenesis. In D. Kuhn and R. Siegler, (Eds.). Cognition, perception and language. Vol. 2. Handbook of Child Psychology (Fifth Ed). (pp. 575-630). New York: John Wiley and Sons.

Gerken, L. A. (1991). The metrical basis for children's subjectless sentences. Journal of Memory and Language, 30, 431-451.

Gleitman, L. (1990). Structural sources of verb learning. *Language acquisition*, 1, 1 – 63.

Greenberg, J. 1963. Some universals of grammar with special reference to the order of words. in Greenberg, J. (ed.) Universals of Language. MIT Press: Cambridge.

Grodzinsky, Y., & Reinhart, T. (1993). The innateness of binding and coreference.

Linguistic Inquiry, 24 (1), 69–101.

Hamburger, H., & Crain, S. (1984). Acquisition of cognitive compiling. Cognition, 17, 85 – 136.

Hirsch-Pasek, K., and R.M. Golinkoff. (1996). *The origins of grammar*. Cambridge, Mass.: MIT Press.

Horn, L. R. (1989) A natural history of negation. University of Chicago Press, Chicago.

Hyams, N. (1986). *Language acquisition and the theory of parameters*. Dordrecht:Reidel.

Hyams, N. 1992 A reanalysis of null subjects in child language. In J. Weissenborn, H. Goodluck and T. Roeper, eds. Theoretical Issues in Language Acquisition, Hillsdale, N.J.: Lawrence Erlbaum.

Jackendoff, R. (1972) Semantic Interpretation in Generative Grammar. Cambridge, MA: MIT Press.

Kayne, R. 1995. The antisymmetry of syntax. MIT Press: Cambridge

Kazanina, N. 2005 The acquisition and processing of backwards anaphora. Ph.D. dissertation, University of Maryland.

Kazanina, N. and C. Phillips 2001. Coreference in Child Russian: Distinguishing Syntactic and Discourse Constraints. In Anna H.-J. Do, Laura Domínguez, and Aimee Johansen (eds.) *Proceedings of the 25th annual Boston University Conference for Language Development*. Somerville, MA: Cascadilla Press. pp. 413-424.

Larson, R. 1991. "Some issues in verb serialization," in C. Lefebvre (ed), *Serial Verbs*. Benjamins: Philadelphia.

Lidz, J. in press. The grammar of accusative case in Kannada. *Language*.

Lidz, J. et al., 2004. Quantifier Raising in 4-year-olds. Proceedings of the Boston University Conference on Language Development. Somerville, MA., Cascadilla Press.

Lidz, J. in press. Verb learning as a probe for children's grammars. in K. Hirsh-Pasek and R. Golinkoff, eds. *When Action Meets Words*. Oxford University Press.

Lidz, J. and Musolino, J. (2002) 'Children's command of quantification' Cognition 84, 113-154.

Lidz, J., S. Waxman and J. Freedman 2003. What infants know about syntax but couldn't have learned. Cognition, 89, B65-B73.

Lightfoot, D. (1991). How to set parameters: Arguments from language change. Cambridge, MA: MIT Press.

Lightfoot, D. 1999. The development of language: acquisition, change and evolution. Blackwell: Cambridge.

MacWhinney, B., and C. Snow. (1985). The Child Language Data Exchange System. *Journal of Child Language, 12, 271-296.*

Marcus, G. F. (1993). Negative evidence in language acquisition. *Cognition,* 46, 53-85.

May, R. (1985) *Logical form: Its structure and derivation.* Cambridge, Mass.: MIT Press.

McDaniel, D., H.S. Cairns, and J.R Hsu. (1991). Control Principles in the grammars of young children. *Language Acquisition*, 1, 297-335.

McKee, C. (1992). A comparison of pronouns and anaphors in Italian and English acquisition. *Language Acquisition*, 1, 21-55.

Musolino, J. and J. Lidz, in press. Why children are not universally successful with quantification. *Linguistics.*

Musolino, J., S. Crain and R. Thornton. (2000). Navigating Negative Quantificational Space. *Linguistics*, 38, 1-32.

Newport, E.L. (1990). Maturational constraints on language learning. *Cognitive Science, 14, 11-28.*

Newport, E. and R. Aslin 2000. Innately constrained learning: Blending old and new approaches to language acquisition. In S.C. Howell, S.A. Fish, and T. Keith-Lucas (Eds.), *Proceedings of the 24th Annual Boston University Conference on Language Development*. Somerville, MA: Cascadilla Press.

Pinker, S. (1989). *Learnability and Cognition: the acquisition of argument structure*. Cambridge, Mass.: MIT Press.

Radford, A. (1990). Syntactic theory and the acquisition of English syntax: The nature of early child grammars in English. Oxford, UK: Blackwell.

Radford, A. (1995) Phrase Structure and Functional Categories. in P. Fletcher & B. McWhinney (eds) *The Handbook of Child Language*, Blackwell, Oxford.

Regier, T. & S. Gahl 2003. Learning the unlearnable: the role of missing evidence. Cognition, 93, 147-155.

Reinhart, T. (1976). The syntactic domain of anaphora. Doctoral dissertation, MIT, Cambridge, MA.

Rizzi, L. (1990). *Relativized Minimality*. Cambridge, Mass.: MIT Press.

Rizzi, L. (1993/1994). Some notes on linguistic theory and language development: The case of root infinitives. *Language acquisition*, 3, 371-393.

Snyder, W. 1995. Language acquisition and language variation. Ph.D. Dissertation, MIT.

Snyder, W. 2001. On the nature of syntactic variation: evidence from complex predicates and complex word formation. Language, 77, 324-342.

Snyder, W., and K. Stromswold. (1997). The structure and acquisition of English dative constructions. *Linguistic Inquiry*, 28, 281-317.

Sugisaki, K. and M. Isobe (2000). Resultatives result from the compounding parameter. Proceedings of WCCFL.

Tenenbaum, J. & T. Griffiths 2001. Generalization, Similarity and Bayesian Inference. *Behavioral and Brain Sciences*, 24, 629-640.

Thomas, M. and A. Karmiloff-Smith (2005) Can developmental disorders reveal the component parts of the human language faculty? *Language Learning and Development*, 1, 65-93.

Thornton, R. and K. Wexler. (1999). *Principle B, VP Ellipsis, and Interpretation in Child Grammar*. Cambridge, Mass.: MIT Press.

Tomasello, M. (2000). Do children have adult syntactic competence? *Cognition*, 74, 209-253.

Vainikka, A. (1994). Case in the development of English syntax. *Language Acquisition,* 3, 257-325.

Valian, V. 1986. Syntactic categories in the speech of young children. Developmental Psychology, 22, 562-579.

Valian, V. 1991. Syntactic subjects in the early speech of American and Italian children. *Cognition*, 40, 21-81.

Valian, V. and L. Casey 2003. Young children's acquisition of wh-questions: the role of structured input. Journal of Child Language, 30, 117-143.

Viau, J. J. Lidz and J. Musolino 2005. LF priming in 4-year-olds. Paper presented at CUNY conference on psycholinguistics.

Waxman, S. and J. Lidz, in press. Word learning. in D. Kuhn and R. Siegler, eds. Handbook of Child Psychology.

Wexler, K. 1990. Innateness and maturation in linguistic development. *Developmental Psychobiology*, 23, 645-660.

Wexler, K. 1996. Very early parameter setting and the unique checking constraint. Lingua, 106.

Wexler, K., & Culicover, P. W. (1980). Formal principles of language acquisition. Cambridge, MA: MIT Press.

Yang, C. (2002) Knowledge and Learning in Natural Language. Oxford University Press. Oxford, UK.

		Resultatives	
		Pass	*Fail*
Compounding	*Pass*	10	2
	Fail	2	6

Figure captions

Figure 1. Backwards anaphora in Russian (from Kazanina and Phillips, 2001)

Figure 2. Scope interpretations of Kannada- and English-learning 4-year-olds (from Lidz and Musolino, 2002)

Table caption

Table 1. Contingency between understanding of resultatives and producing compounds (adapted from Sugisaki and Isobe, 2000)

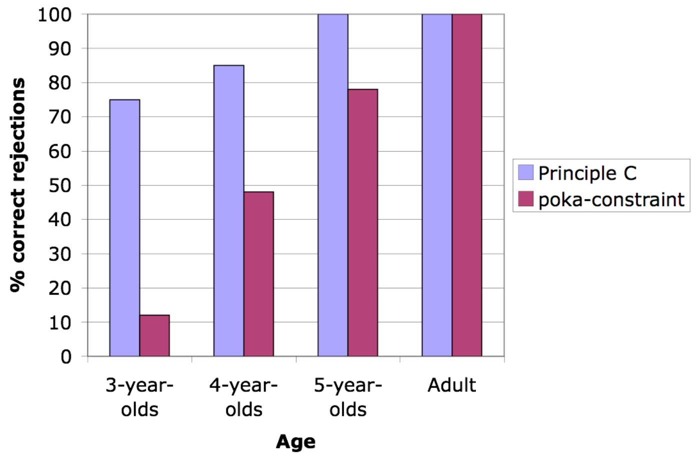
100
90
80
70
60
50
40
30
20
10
0
% correct rejections
3-year-olds
4-year-olds
5-year-olds
Adult
Age
Principle C
poka-constraint

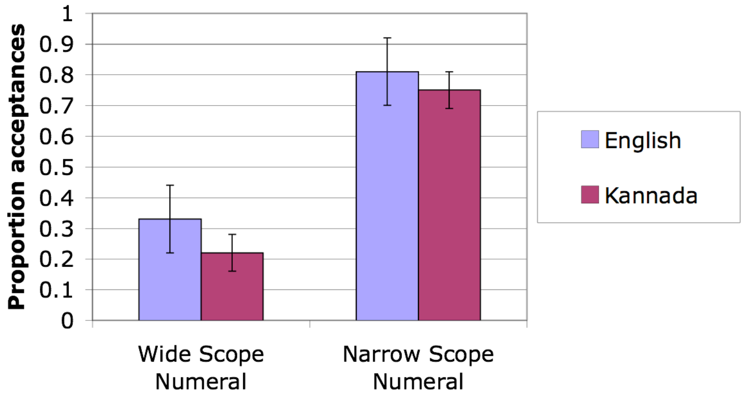
Proportion acceptances
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Wide Scope Numeral
Narrow Scope Numeral
English
Kannada