

Indirect Evidence and the Poverty of the Stimulus: The Case of Anaphoric *One*

Stephani Foraker,^a Terry Regier,^a Naveen Khetarpal,^a
Amy Perfors,^b Joshua Tenenbaum^b

^a*Department of Psychology, The University of Chicago*

^b*Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology*

Received 13 November 2007; received in revised form 17 July 2008; accepted 15 August 2008

Abstract

It is widely held that children's linguistic input underdetermines the correct grammar, and that language learning must therefore be guided by innate linguistic constraints. Here, we show that a Bayesian model can learn a standard poverty-of-stimulus example, anaphoric *one*, from realistic input by relying on indirect evidence, without a linguistic constraint assumed to be necessary. Our demonstration does, however, assume other linguistic knowledge; thus, we reduce the problem of learning anaphoric *one* to that of learning this other knowledge. We discuss whether this other knowledge may itself be acquired without linguistic constraints.

Keywords: Language acquisition; Poverty of the stimulus; Indirect evidence; Bayesian learning; Syntax; Anaphora

1. Introduction

Language-learning children are somehow able to make grammatical generalizations that are apparently unsupported by the overt evidence in their input. Just how they do this remains an open question. One influential proposal is that children succeed in the face of impoverished input because they bring innate linguistic constraints to the task. This argument from *poverty of the stimulus* (e.g., Chomsky, 1965:58; 1980:34) is a long-standing basis for claims of innate linguistic knowledge. An alternative solution is that the learner instead relies on indirect evidence and domain-general learning (e.g., Landauer & Dumais,

Correspondence should be sent to Dr. Stephani Foraker, Department of Psychology, Buffalo State College, 1300 Elmwood Ave., Buffalo, NY 14222. E-mail: forakesm@buffalostate.edu

1997; Lewis & Elman, 2001; Reali & Christiansen, 2005; Regier & Gahl, 2004) and does not require specifically linguistic innate constraints.

We pursue this idea here, focusing on what can be learned by noting that certain forms are systematically absent from the input. Hypotheses about the grammar should gradually lose support as the evidence they predict consistently fails to appear. Some discussions have explicitly considered this possibility (Chomsky, 1981:9), but many presentations of the argument from poverty of the stimulus do not (Baker, 1978:416; Hornstein & Lightfoot, 1981:18–20; Lidz, Waxman, & Freedman, 2003). We argue that this domain-general principle can support language learning, despite the apparent poverty of the stimulus. In particular, we show that a standard poverty-of-stimulus example can be learned in this manner without a linguistic constraint that has been held to be necessary, and necessarily innate.

1.1. Noun phrase structure and anaphoric one

An established example of the argument from poverty of the stimulus concerns the anaphoric use of the word *one* (Baker, 1978; Hornstein & Lightfoot, 1981; Lidz, Waxman, & Freedman, 2003). This is the use of *one* as a referential pronoun, not as a count label (e.g., *one, two, three*) or without a specific antecedent (e.g., *one must drink to survive.*). The learning challenge is to determine the antecedent of *one* within a hierarchically structured noun phrase. Concretely, the learner may encounter a sentence such as (1).

1. Here’s a yellow bottle. Do you see another *one*?

As illustrated in Fig. 1, the antecedent of the anaphor *one* is ambiguous: *one* could take as its antecedent three different levels of noun phrase structure: (a) the upper N', referring to some *yellow bottle*, (b) the lower N', referring to a *bottle* of some unspecified color, or (c) N⁰, also referring to a *bottle* of some unspecified color.

What is the correct answer, then? What level of syntactic structure does *one* substitute for? It is commonly accepted that in general *one* can take any N' constituent as its antecedent but cannot substitute for N⁰ (Lidz & Waxman, 2004; Radford, 1988). The reason for this

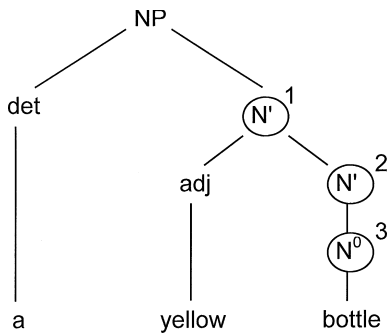


Fig. 1. Structure of the noun phrase *a yellow bottle*. Numbered circles show the three possible antecedents of anaphoric *one*.

can be seen by comparing the linguistic behavior of two different types of prepositional phrase that may follow a noun, namely complements and modifiers. The core distinction between the two is that a complement is necessarily conceptually evoked by its head noun. For instance, *member* necessarily evokes the organization of which one is a member; so, in *member of congress*, the phrase *of congress* is a complement. In contrast, a modifier is not necessarily evoked by its head. The word *man* does not necessarily evoke conceptually where the man is from, so in *man from Rio*, the phrase *from Rio* is a modifier, not a complement (Baker, 1989; Bowen, 2005; Keizer, 2004; Taylor, 1996). While there are more subtle intermediate cases, this is the conceptual core of the distinction.

As shown in (2), it is ungrammatical for *one* to be anaphoric with a complement-taking noun (*piece*) without its complement (*of cheese*).

2. *I'll have a piece of cheese and you can have one of apple.

Such unacceptable complement structures contrast with modifiers. In (3), which has a noun with a postnominal modifier, it is grammatical for *one* to be anaphoric with the noun, *ball*, without its modifier, *with stripes*.

3. I want the ball with stripes and you can have the one with dots.

Fig. 2 shows standard syntactic representations for the noun phrases in these two sentences. Note that the sister of a complement is N^0 , while the sister of a modifier is N' . These representations account for the different behavior of complements and modifiers in connection with anaphoric *one*, provided one postulates that *one* may only take N' as antecedent. The correct hypothesis then, on this standard analysis, is that *one* substitutes for any N' constituent, but not for N^0 .

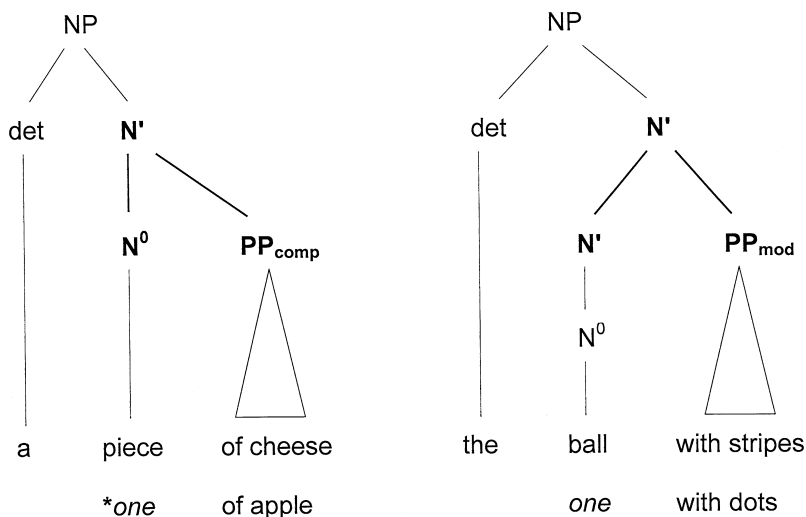


Fig. 2. Structure of a complement phrase (left panel) and of a modifier phrase (right panel).

1.2. Where does this knowledge come from?

How does a child come to know this constraint on the antecedents of *one*? Lidz, Waxman, and Freedman (2003) showed that 18-month-old infants, given sentences like (1), looked longer at a yellow bottle than at a bottle of a different color, which they interpreted to mean that the infants knew that *one* took the upper N' as its antecedent, rather than the lower N' or N⁰. Lidz et al. (2003) also searched through a child language corpus for input that, in a single occurrence, could unambiguously rule out incorrect hypotheses. They found that such input was effectively absent from the corpus. They argued that this poverty of the stimulus implicated innate syntactic knowledge: children know something about language that they could not have learned from the input, so at least part of the knowledge must be innate. In particular, they argued that the N⁰ hypothesis is innately excluded from consideration.

This account has been challenged (Akhtar, Callanan, Pullum, & Scholz, 2004; Regier & Gahl, 2004; Tomasello, 2004), and the question remains unresolved. Of particular relevance to the present study, Regier and Gahl (2004) showed that a simple Bayesian model could learn the upper N' solution for such sentences, given linguistic input of the form shown in (1) and information about the color of the referent, without prior exclusion of the N⁰ hypothesis. Thus, this model can account for Lidz et al.'s (2003) empirical findings without any innate exclusion of the N⁰ hypothesis. In response, Lidz and Waxman (2004) argued that this model is inadequate since it learns to support only the upper N' hypothesis, while as we have seen, more generally *one* may take any N' as its antecedent. It is true that Regier and Gahl's (2004) model is limited in this way—but importantly, the same criticism also applies to the Lidz et al. (2003) study, which itself focused on the upper N' and did not demonstrate that children know the more general any-N' hypothesis. It is in this respect that the question remains open.

Critically, there is a fundamental limitation that affects both the Lidz et al. (2003) experiments and the Regier and Gahl (2004) model. Both studies relied on *referential* information such as the color of the referent. And referentially, nothing distinguishes a situation in which the antecedent of *one* is the lower N' from a situation in which it is N⁰, since the referenced object in either case can be of any color. This is a problem because the lower N' situation is consistent with the correct any-N' hypothesis while the N⁰ situation is not.

1.3. The child as linguist

Since referential evidence will not suffice to learn—or even test—the correct any-N' hypothesis, what sort of evidence might? We propose that the child may learn about the antecedent of *one* in much the same way a linguist does: by noticing the different behavior of complements and modifiers. Concretely, if the child came to realize that *one* never substitutes for a complement-taking noun without its complement, while *one* does sometimes substitute for a modified noun without its modifier, that observation could lead the child to the knowledge that *one* can substitute for N' but not for N⁰.

This idea inverts a standard linguistic test (e.g., Radford, 1988:175), in which the acceptability or unacceptability of substituting *one* in an NP is used to determine whether a given postnominal phrase within the NP is a complement or a modifier, as in examples (2) and (3). There is an apparent circularity in this: we can use the acceptability of substituting anaphoric *one* to determine whether a phrase is a complement or a modifier—but we need the distinction between complements and modifiers to learn the correct use of anaphoric *one* in the first place. This circularity is only apparent, however, since complements may be distinguished from modifiers on conceptual grounds. As discussed earlier, a complement is necessarily conceptually evoked by its head noun, while a modifier is not. Crucially, if the language-learning child has grasped this conceptual distinction between complements and modifiers, that distinction could serve as a basis for learning about anaphoric *one*.

In this paper, we show that a Bayesian model given realistic input can learn the correct any- N' hypothesis based on these principles, without innately excluding the false N^0 hypothesis. Thus, a piece of linguistic knowledge that has been held to be necessarily innate need not be, if attention is given to what is absent from, as well as what is present in, the input. We intend this demonstration to show that learning is in principle possible given the data available to a child. The question of whether, or how, this idealized account might be implemented in the mind of a child is a separate question that we do not pursue here.

2. Model

We assume a learner that assesses support for hypotheses on the basis of evidence using Bayes' rule:

$$p(H|e) \propto p(e|H)p(H)$$

Here H is a hypothesis in a hypothesis space, and e is the observed evidence. The likelihood $p(e|H)$ is the probability of observing evidence e given that hypothesis H is true, and the prior probability $p(H)$ is the a priori probability of that hypothesis being true. To flesh this general framework out into a model, we next specify the sort of evidence that will be encountered, our general assumptions, the hypothesis space, the prior, and the likelihood.

2.1. Evidence

The model observed a series of noun phrases, drawn from child-directed speech. Each noun phrase was represented without hierarchical structure, as a sequence of part-of-speech tags (e.g., *the big ball* would be coded as “determiner adjective noun”), followed by a code for a postnominal modifier or complement, if applicable (e.g., *side of the road* would be coded as “noun complement”). This source of evidence was chosen because children receive a steady supply of such input, and following our discussion above, it is data of this sort that can discriminate among the relevant hypotheses.

2.2. Assumptions

We made four major assumptions about the knowledge available to the language-learning child. First, following Lidz et al. (2003), we assumed that the child is able to recognize anaphoric uses of *one*. Second, we assumed that the child knows how noun phrases are hierarchically structured in English, such that the central learning problem is determining which constituent within this structure is the antecedent of *one*. Third, we assumed that the child knows a fact about pronouns generally, including *one*: that a pronoun *substitutes* for its antecedent and must therefore be of the same syntactic type as the antecedent. As a corollary, we also assumed that the antecedent of *one* could be of only one syntactic type, for example, only N' , or only N^0 , and not both. Thus, if the pronoun occupies an N' position within its noun phrase, the antecedent must similarly occupy an N' position in its noun phrase, for otherwise the pronoun would not be able to substitute for the antecedent. By the same token, if the pronoun occupies an N^0 position, the antecedent should, too. Critically, given this assumption, the problem of determining whether the antecedent of *one* is N' or N^0 reduces to the problem of determining whether *one* itself, within its own NP, takes the role of N' or N^0 . We felt justified in making this assumption since knowledge of substitutability concerns pronouns generally, and it could be learned by observing the behavior of pronouns other than *one*.

Finally, we assumed that the child is able to recognize and distinguish between complements and modifiers when they appear in the child's linguistic input. To our knowledge, there are no studies that have tested whether young children are indeed sensitive to this distinction, but we felt justified in making this assumption since the core distinction between complements and modifiers can be captured in semantic or conceptual terms, and thus, could in principle be learned without innate specifically syntactic constraints. We return in the discussion to the question of just how much of our argument hangs on these assumptions.

2.3. Hypothesis space

We assumed a hypothesis space containing two hypotheses that addressed the question, "Which of the constituents of the NP does anaphoric *one* take as its antecedent?" The two hypotheses are [any N'] and [N^0]. Thus, we chose the simplest possible hypothesis space that includes both the correct hypothesis [any N'] and the incorrect [N^0] hypothesis that Lidz et al. (2003) argued must be innately excluded if learning is to succeed. If our learner can learn the correct answer given this hypothesis space and realistic input, that outcome will indicate that the posited innate exclusion of [N^0] is unnecessary.

Each hypothesis takes the form of a grammar that generates a string of part-of-speech tags corresponding to a noun phrase. The two grammars both capture standard NP structure and are identical except for one rule. Each grammar contains the following productions, with options separated by ' | '.

NP → Pro | Nbar | Det Nbar | Poss Nbar
 Poss → NP ApostropheS | PossPronoun

Nbar \rightarrow Poss Nbar | Adj Nbar | Nbar Mod
 Nbar \rightarrow Nzero | Nzero Comp
 Det \rightarrow determiner
 Adj \rightarrow adjective
 PossPronoun \rightarrow possessive-pronoun
 ApostropheS \rightarrow apostrophe-s
 Mod \rightarrow modifier
 Comp \rightarrow complement
 Nzero \rightarrow noun
 Pro \rightarrow pronoun

In addition, the grammar for the [any N'] hypothesis contains the production:

Nbar \rightarrow anaphoric-*one*,

while the [N⁰] hypothesis instead contains the production

Nzero \rightarrow anaphoric-*one*.

Thus, the two grammars embody, in their last production, the link between *one* and either N' or N⁰. The grammars were designed to be able to parse noun phrases in a child-language corpus. Each production in each grammar has a production probability associated with it, and these probabilities may be adjusted to fit the observed corpus.

2.4. Prior

The two grammars are equally complex: they differ only in one production, which is itself equally complex in the two cases. Since grammar complexity gave no grounds for assigning either hypothesis greater prior probability, we assigned the two hypotheses equal prior probability $p(H)$: $p(N^0) = p(\text{any } N') = .5$. Thus, all discrimination between hypotheses was done by the likelihood.

2.5. Likelihood

Given a hypothesis H in the form of a grammar, and evidence e in the form of a corpus of noun phrases, we used the inside-outside algorithm¹ to obtain a maximum likelihood fit of the grammar to the corpus. This algorithm iteratively reestimates production probabilities in the grammar so as to maximize the probability $p(e|H)$ that the corpus e would be generated by the grammar H . Given the prior and likelihood, we then obtained the posterior probability of each grammar given the corpus, $p(H|e)$, using Bayes' rule.

Both grammars were designed to be consistent with all noun phrases in our corpus. What differs between the grammars is the *expected* observations given that a hypothesis is true.

To see why this is the case, consider the interaction of two rules from the N^0 grammar: $[Nbar \rightarrow Nzero \text{ Comp}]$ and $[Nzero \rightarrow \text{anaphoric-one}]$. Together, these two rules produce strings of the form “*one + complement*,” as in (2) above. Thus, the N^0 hypothesis predicts that such strings will be encountered in the input. But since such strings are ungrammatical, that expectation will not be fulfilled. In contrast, the N' hypothesis does not give rise to this false expectation, since it lacks the second rule. This difference between the two grammars is captured in their likelihoods. If no instances of “*one + complement*” appear in the input, the N^0 grammar will progressively lose support, and the learner will select the N' grammar as the correct hypothesis.

3. Data

We selected as our input data source the Nina corpus (Suppes, 1974) in the CHILDES database (MacWhinney, 2000), which is one of the corpora that Lidz et al. (2003) consulted. The corpus was collected while the child was 23- to 39-months-old and consists of just over 34,300 child-directed mother utterances, containing approximately 60,000 noun phrases. We identified the noun phrases preliminarily using a parser,² and then from those noun phrases we selected a random 5%, 10%, and 15% cumulative sample for further coding by one of the authors (SF). These percentages yielded samples large enough for our purposes, yet small enough to code by hand. We found no ungrammatical noun phrases or ungrammatical uses of anaphoric *one* in the 15% cumulative sample.

We coded each noun phrase as a sequence of part-of-speech tags, without hierarchical structure, for example, “determiner adjective noun,” “determiner noun,” “pronoun,” etc., of the sort generated by the above grammars. We used these corpus data in two ways. First, we designed both grammars to accommodate these data. Second, we provided the data as input to the model. While both grammars were designed to be *consistent* with the data, we were interested in finding which grammar fit the data most closely.

We also coded whether a complement or modifier (or neither) was present, following the noun. For example, the noun phrase *a piece of cheese* would be coded “determiner noun complement,” while *crackers with cheese* would be coded “noun modifier.” We limited ourselves to post-head complements and modifiers in the form of prepositional phrases or clauses (Bowen, 2005; Keizer, 2004; Radford, 1988). To identify a complement or a modifier we used the conceptual intuition described earlier: as noted by several sources, the head noun that takes a complement presupposes some other entity which must be expressed (Huddleston & Pullum, 2002:221; Taylor, 1996:39; see also Fillmore’s, 1965 “inalienably possessed nouns”) or inferable from context (Bowen, 2005:18, Keizer, 2004). To classify post-head strings that followed anaphoric *one*, such as *the one on the table*, we identified the head noun of the antecedent NP from the transcript and applied the same test as for nouns. Note that the head noun is the same regardless of whether N' or N^0 is the correct hypothesis. Another author (NK) also classified material following a noun as a complement or modifier. Inter-rater reliability was 91%, with disagreement on only a small number of nouns (e.g., *name*, *story*, *color*). In cases of disagreement, the classifications of the first rater were retained.

We did not use substitution of anaphoric *one* as a test for classifying the post-head forms, to avoid the circularity alluded to earlier. We restricted ourselves to the conceptual distinction that could lead a child to the complement-modifier distinction without requiring prior syntactic knowledge of anaphoric *one*. However, we did find post hoc that the anaphoric *one* test for count nouns³ yielded results consistent with the conceptual criteria we adopted.

Table 1 shows a summary of the types of noun phrase forms available in the input, focusing on the post-head structure (none, complement, modifier). Following a head noun we found complements (*piece of a puzzle*, *your side of the road*) and modifiers (*food for the llamas*, *the ocean beach with big waves*), while following anaphoric *one* we found only modifiers (*the one with the girl*, *the other ones you like*), consistent with an adult-state grammar. The complements were all prepositional phrases, while the modifiers were prepositional phrases or subordinate clauses.

We also created two reduced variants of the 15% sample, in a form of “targeted impoverishment of the stimulus,” to explore what data are critical to learning anaphoric *one*. The first variant (“no-ones” corpus) was stripped of all NPs containing anaphoric *one*, while the other variant (“no-complements” corpus) was stripped of all NPs containing a complement. We refer to the 15% sample as “the full corpus.” The no-ones, no-complements, and full corpora contained 8,763, 8,750, and 8,816 NPs, respectively. Thus, the manipulations removing instances of *one* or complements each eliminated only a very small proportion (.6% and .8%, respectively) of the full corpus.

4. Method

We first calculated the probability of each hypothesis ($[N^0]$, $[any\ N']$), given the 5%, 10%, and 15% samples, corresponding to 4.2, 8.4, and 12.6 hr of mother input, respectively. We then recalculated the same probabilities on the no-ones and no-complements corpora. In each case, we computed these probabilities using Bayes’ rule, with uniform prior, and likelihood computed by the inside-outside algorithm, as specified above. The production probabilities returned by the inside-outside algorithm were intuitively reasonable in that they closely reflected the relative frequencies of the corresponding forms in the corpus.

Table 1
Frequency counts of postnominal structures in the input for the 5%, 10%, and 15% cumulative samples

Noun Phrase Forms	5%	10%	15%
Noun	1,253	2,478	3,752
Pronoun	1,605	3,213	4,784
Anaphoric <i>one</i>	13	32	49
Noun + complement	29	47	65
Noun + modifier	66	113	161
Noun + complement + modifier	0	1	1
Anaphoric <i>one</i> + modifier	3	3	4




Fig. 3. Posterior probability of hypotheses given varying amounts of mother input.

5. Results

Fig. 3 shows that the probability of the correct [any N'] hypothesis is equal to that of the incorrect [N^0] hypothesis prior to observing any data (0 hr of input). However, as more data are seen, the probability of the correct hypothesis given those data grows steadily higher, at the expense of the incorrect hypothesis. This indicates, contra the poverty-of-stimulus argument, that a learner can discover that the antecedent of *one* is N' , even though N^0 is not innately excluded from consideration during learning.

Why does this happen? The N^0 hypothesis falsely predicts that the input will include strings containing “*one* + complement,” while the N' hypothesis does not. Thus, the likelihood of the observed data is higher for the N' hypothesis. As we have seen, the false N^0 prediction arises from the interaction of two rules in the N^0 grammar: [$N\text{bar} \rightarrow N\text{zero Comp}$], and [$N\text{zero} \rightarrow \text{anaphoric-one}$]. The first production is shared with the N' grammar, but the second one is not. The probabilities of both productions must be substantial, since the corpus contains complements, which require the first rule, and instances of *one*, which require the second. However, if the corpus lacked either *one* (i.e., the no-ones corpus), or complements (i.e., the no-complements corpus), the corresponding rule would receive 0 probability in a maximum likelihood fit to the data. In such cases the N^0 hypothesis would *not* make the false “*one* + complement” prediction, and there would be nothing to distinguish N^0 from N' . These expectations were confirmed, as shown in Fig. 4.

Thus, the successful learning we see on the full corpus is dependent on the interaction of anaphoric *one* and complements. When both appear in the corpus, even in very modest quantities as was true here, the N^0 hypothesis falsely predicts the unattested “*one* + complement” pattern and is penalized for its absence.⁴ This interaction supports learning without the innate exclusion of N^0 .




Fig. 4. Targeted impoverishment of the stimulus: posterior probability of hypotheses given the full corpus, and the same corpus stripped of strings with *one*, or strings with complements.

6. Discussion

We have shown that a simple model can learn the behavior of anaphoric *one* without a linguistic constraint that has been held to be necessary, and necessarily innate. Our demonstration relies on learning from the absence of predicted input patterns, a form of indirect evidence that is broadly consistent with other recent work emphasizing the power of indirect evidence in countering standard poverty-of-stimulus arguments (Landauer & Dumais, 1997; Lewis & Elman, 2001; Reali & Christiansen, 2005; Regier & Gahl, 2004).

We anticipate several objections to our demonstration. First, we have adopted relatively simple representations of the relevant aspects of English grammar, and it may prove harder to learn anaphoric *one* using more realistic grammars.⁵ We have made these simplifying assumptions so as to engage earlier work that argues for nonlearnability of anaphoric *one* using similarly simplified representations. We have also not concerned ourselves here with the question of noisy (ungrammatical) input, as our corpus did not require it. We hope that our broader conceptual points, about the learnability of grammar and the poverty of the stimulus, will generalize to learning with more complex and realistic grammatical formalisms, and occasional ungrammatical input. We grant that this is an open question, however.

Another possible objection is that our hypothesis space is very restricted, containing just two hypotheses. One might concede our point that N^0 need not be excluded from consideration, contra the standard argument, but then counter that we ourselves use a very constrained space. Thus, perhaps the fundamental “constrained space” idea is correct, even if the excluded- N^0 proposal was not. We consider it self-evident that the space must be

constrained; the critical question is whether the constraints are specifically linguistic. Will our demonstration scale up in a hypothesis space that is constrained only by nonlinguistic general cognitive considerations? We consider that an open and interesting question, which is preliminarily addressed by some of our discussion below.

A related possible objection is that we have assumed a good deal of linguistic knowledge: for instance, knowledge of the hierarchical structure of the noun phrase, knowledge that pronouns substitute for their antecedents, and awareness of the complement-modifier distinction. This is true. Our primary contribution has been to *reduce* the problem of learning anaphoric *one* to the problem of learning this other knowledge. The critical subsequent question is whether this assumed knowledge can itself be learned without innate linguistic constraints. For example, Perfors, Tenenbaum, and Regier (2006) have shown that the hierarchical nature of language may be learned without specifically linguistic prior bias, by relying on a domain-general preference for parsimonious description. If the other linguistic knowledge we have assumed can likewise be learned with only domain-general constraints, a standard poverty-of-stimulus example will have been shown to be learnable without specifically linguistic constraints. If not, the example of anaphoric *one* will retain its status as an argument for innate linguistic knowledge—but we will have shown that the critical linguistic constraints lie elsewhere than traditionally imagined.

At the same time, some of our assumptions highlight ways in which our account can be empirically tested. We assume that children learn the correct antecedent of *one* by building on the complement/modifier distinction. Thus, on our account, the complement/modifier distinction must be available to children before they know the correct antecedent of *one* (namely, any N'). To our knowledge, this question is empirically untested.

Perhaps the broadest potential objection is that it may seem wrong-headed, or paradoxical, to argue against the nativist poverty-of-stimulus claim while using structured linguistic representations of exactly the sort commonly proposed by nativists. We see no problem here. We consider ourselves to be working “from the inside out.” We start with linguistic representations that a nativist should recognize and show that domain-general principles support learning of the nominally correct grammar, contra specific unlearnability claims in the literature. This allows us to engage the poverty-of-stimulus argument in its own representational terms while working “outwards” to domain generality. In contrast, connectionist studies that also question the poverty of the stimulus (e.g., Lewis & Elman, 2001; Reali & Christiansen, 2005) work “from the outside in.” They start with domain-general representations and learn linguistic behavior similar to that of a grammar. The two approaches complement each other: the starting point for connectionist studies is undeniably domain-general, while in our case that which is learned is undeniably a grammar.

Notes

1. We used the code made publicly available by Mark Johnson at <http://www.cog.brown.edu/~mj/Software.htm>.

2. We used the Stanford Parser, version 1.5.1, available at <http://nlp.stanford.edu/downloads/lex-parser.shtml>.
3. Anaphoric *one* does not substitute for mass nouns, like *sand*.
4. This idea is related to the “subset principle” (Berwick, 1986; Pinker, 1995:172–175). The subset principle is a domain-general learning principle by which a learner assumes the narrowest (most specific) hypothesis that is consistent with the data seen so far.
5. For example, so-called “picture NPs,” in which the head noun denotes a picture, are a known exception to the *one* = [any N’] generalization: unlike other NPs, they allow the form *one* + complement, e.g., ‘Here’s a picture of Paul and here’s one of Sally.’ Picture NPs are thus a challenge for theories of anaphor resolution generally, as they are sensitive to constraints other than those of structural binding (Kuno, 1987; Pollard & Sag, 1994; Reinhart & Reuland, 1993; Runner, Sussman & Tanenhaus, 2006). As it happens, our data did not contain any instances of *one* substituting for a picture noun.

Acknowledgments

We thank Lisa Pearl and two anonymous reviewers for their comments and suggestions. This research appeared in the Proceedings of the 29th Annual Conference of the Cognitive Science Society. It has also benefited from feedback received by four anonymous reviewers for that conference and attendees of the conference. We also thank Susanne Gahl for helpful comments. This research was partially supported by NIH fellowship HD049247 to Stephani Foraker.

References

- Akhtar, N., Callanan, M., Pullum, G. K., & Scholz, B. C. (2004). Learning antecedents for anaphoric *one*. *Cognition*, 93, 141–145.
- Baker, C. L. (1978). *Introduction to generative transformational syntax*. Englewood Cliffs, NJ: Prentice Hall.
- Baker, C. L. (1989). *English syntax*. Cambridge, MA: MIT Press.
- Berwick, R. C. (1986). Learning from positive-only examples: The subset principle and three case studies. In R. S. Michalski, J. C. Carbonell & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 2, pp. 625–645). Los Altos, CA: Morgan Kaufmann.
- Bowen, R. (2005). Noun complementation in English: A corpus-based study of structural types and patterns. In G. Florby (Ed.), *Gothenburg Studies in English 91* (pp. 1–270). Göteborg, Sweden: Acta Universitatis Gothoburgensis.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Chomsky, N. (1980). *Rules and representations*. New York: Columbia University Press.
- Chomsky, N. (1981). *Lectures on government and binding: The Pisa lectures*. Berlin: Mouton de Gruyter.
- Fillmore, C. J. (1965). *Indirect object constructions in English and the ordering of transformations*. The Hague: Mouton.
- Hornstein, N., & Lightfoot, D. (1981). Introduction. In N. Hornstein & D. Lightfoot (Eds.), *Explanation in linguistics* (pp. 9–31). London: Longman.

- Huddleston, R., & Pullum, G. K. (2002). *Cambridge grammar of the English language*. Cambridge: Cambridge University Press.
- Keizer, E. (2004). Postnominal PP complements and modifiers: A cognitive distinction. *English Language and Linguistics*, 8, 323–350.
- Kuno, S. (1987). *Functional syntax: Anaphora, discourse and empathy*. Chicago: University of Chicago Press.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The Latent Semantic Analysis theory of acquisition, induction and representation of knowledge. *Psychological Review*, 104 (2), 211–240.
- Lewis, J. D., & Elman, J. L. (2001). Learnability and the statistical structure of language: Poverty of stimulus arguments revisited. In B. Skarabela, S. Fish & A. H. -J. Do (Eds.), *Proceedings of the 26th annual Boston University conference on language development* (pp. 359–370). Somerville, MA: Cascadilla Press.
- Lidz, J., Waxman, S., & Freedman, J. (2003). What infants know about syntax but couldn't have learned: Evidence for syntactic structure at 18-months. *Cognition*, 89, B65–B73.
- Lidz, J., & Waxman, S. (2004). Reaffirming the poverty of the stimulus argument: A reply to the replies. *Cognition*, 93, 157–165.
- MacWhinney, B. (2000). *The CHILDES project: Tools for analyzing talk* (3rd ed.). Mahwah, NJ: Erlbaum.
- Perfors, A., Tenenbaum, J., & Regier, T. (2006). Poverty of the stimulus? A rational approach. In R. Sun (Ed.), *Proceedings of the 28th annual conference of the cognitive science society* (pp. 663–668). Mahwah, NJ: Erlbaum.
- Pinker, S. (1995). Language acquisition. In L. Gleitman & M. Liberman (Eds.), *Language: An invitation to cognitive science* (2nd ed., Vol. 1, pp. 135–182). Cambridge, MA: MIT Press.
- Pollard, C., & Sag, I. (1994). *Head-driven sentence structure grammar*. Chicago: CSLI/University of Chicago Press.
- Radford, A. (1988). *Transformational grammar: A first course*. New York: Cambridge University Press.
- Reali, F., & Christiansen, M. (2005). Uncovering the richness of the stimulus: Structure dependence and indirect statistical evidence. *Cognitive Science*, 29, 1007–1028.
- Regier, T., & Gahl, S. (2004). Learning the unlearnable: The role of missing evidence. *Cognition*, 93, 147–155.
- Reinhart, T., & Reuland, E. (1993). Reflexivity. *Linguistic Inquiry*, 24, 657–720.
- Runner, J. T., Sussman, R. S., & Tanenhaus, M. K. (2006). Processing reflexives and pronouns in picture noun phrases. *Cognitive Science*, 30, 193–241.
- Suppes, P. (1974). The semantics of children's language. *American Psychologist*, 29, 103–114.
- Taylor, J. R. (1996). *Possessives in English: An exploration in cognitive grammar*. Oxford, UK: Oxford University Press.
- Tomasello, M. (2004). Syntax or semantics? Response to Lidz et al. *Cognition*, 93, 139–140.