Psych 229: Language Acquisition

Lecture 2

Introduction to Language Acquisition Continued

Possible objections to mental grammar

Why should I believe that I store a grammar in my head? I just understand sentences because they make sense.

In reply I ask you: Why do some combinations of words "make sense" and others not? For instance, if we interchange adjacent words in the sentences in (2)–(5), to form chains of words like (10), we find that the sentences don't "make sense" anymore.

(10) Amy two ate peanuts.
A is numeral not a numbskull.
Bill that thinks Beth is a genius.
etc.

That is, "making sense" involves, among other things, conformity to known patterns. In other words, the mental grammar plays some sort of role after all.

In fact, we can recognize patterns of English even if not all the words are real English words. This is the basis of Lewis Carroll's famous poem *Jabberwocky*:

'Twas brillig, and the slithy toves
Did gyre and gimble in the wabe . . .

Possible objections to mental grammar

What about people who speak ungrammatically, who say things like "We ain't got no bananas"? They don't have grammars in their heads.

This question points up an important difference between the ordinary use of the term "grammar" and the linguists' theoretical construct "mental grammar." In ordinary usage, "grammar" refers to a set of rules taught in school that tell us how we should speak in order to conform to the norms of polite (roughly, educated middle-class) society. "Proper grammar" frowns on the use of "ain't," the use of "got" for "have," and the use of double negatives; the "proper" way to say this sentence is "We don't have any bananas" or "We have no bananas." In the sense of "school grammar," then, speaking ungrammatically is a violation of a social norm, sort of like spitting in public.

What about the people who don't speak "correct English"? A moment's reflection suggests that their speech does in fact fall into consistent patterns. Someone who says "We ain't got no bananas" still doesn't produce monstrosities like "ain't no we got bananas" or "no got ain't bananas we"; the words come in a well-defined order. More subtly, such a speaker won't substitute the so-called correct term "have" for "got," saying "We ain't have no bananas."

Possible objections to mental grammar

When I talk, the talk just comes out—I'm not consulting any "grammar in my head." If I look into my mind, I may find some scraps of school grammar, but you're trying to tell me that's not what mental grammar is supposed to be. So what is it supposed to be?

Well, consider: there are lots of other things going on in our brains of which we aren't conscious either. Think about getting from an intention such as "I think I'll wiggle my fingers now" into commands to be sent to the muscles, so that our fingers wiggle. Just how do we do it? From the point of view of introspection, the experience is entirely immediate: we decide to wiggle the finger, and the finger wiggles, unless there is some obstruction or paralysis. How the mind actually accomplishes this is entirely opaque to awareness. In fact, without studying anatomy, we can't even tell which muscles we've activated. So it is, I want to suggest, with the use of mental grammar.

The argument for innate knowledge

Suppose we have mental grammars in our heads - how did they get there?

How do children do it? Many people immediately assume that the parents taught it. To be sure, parents often engage in teaching *words* to their kids: "What's this, Amy? It's a *BIRDIE*! Say 'birdie,' Amy!" But language learning can't be entirely the result of teaching words. For one thing, there are lots of words that it is hard to imagine parents teaching, notably those one can't point to: "Say 'from,' Amy!" "This is *ANY*, Amy!"

Think also about children of immigrants, say the Americans who move to Israel. The adults often never feel comfortable with the language of the adopted country. They speak with an accent, they express themselves with hesitation, they admit to not quite following the news on television, and so forth. Yet their children become fully fluent native speakers of the new language. Evidently the children have learned something their parents don't know. So the parents couldn't have taught them.

The argument for innate knowledge

Although children often learn *words* as a result of parental instruction, it is less clear that they learn *grammatical patterns* this way. Anyone who has attempted to correct a two-year-old's grammar will know that it can't be done. The following dialogue, recorded by the linguist David McNeill, is a famous illustration.

CHILD: Nobody don't like me.
MOTHER: No, say "nobody likes me."
CHILD: Nobody don't like me.
. . .
(eight repetitions of this dialogue)
. . .
MOTHER: No, now listen carefully; say "*nobody likes me.*"
CHILD: Oh! Nobody don't likes me.

It is true that certain grammatical patterns *are* taught as part of school grammar, for example the rule that a preposition is something you must never end a sentence with. However, English speakers violate this rule all the time, and have for hundreds of years. I just did, two sentences ago. The idea that a preposition shouldn't occur at the end of a sentence seems to have arisen during the eighteenth century, when for the first time "authorities on English usage" sought to determine the "correct" way to speak, on the basis of the models of the classical languages Latin and Greek.

The argument for innate knowledge

Such proscriptive teaching of grammar, which evidently doesn't work very well, contrasts strikingly with aspects of English sentence patterns that probably nobody has ever thought to teach. Here's an example. Look at the four sentences in (1).

(1) *a* Joan appeared to Moira to like herself.
b Joan appeared to Moira to like her.
c Joan appealed to Moira to like herself.
d Joan appealed to Moira to like her.

Without thinking about it consciously, you have automatically inferred that each of these sentences has a different combination of who is to like whom. In (1a), Joan likes Joan; in (1b), Joan likes Moira or some unspecified third party; in (1c), Moira is to like Moira; in (1d), Moira is to like Joan or a third party.

How do we come to understand these sentences this way? It obviously depends somehow on the difference between ordinary pronouns such as "her" and reflexive pronouns such as "herself," and also on the difference between the verbs "appear" and "appeal." But how? Whatever reasons there may be, I'm sure no one is ever taught about contrasts like this by their parents or teachers or anyone else. Yet this aspect of English grammatical patterns is deeply ingrained, much more so than the taught prohibition against ending a sentence with a preposition.

The argument for innate knowledge

I can't resist another example, because it's so striking. There is an alteration called "expletive infixation" that many speakers perform on words of English under conditions of extreme exasperation, as in (2).

(2) How many times do I have to tell you? I'm not talking about the *Allegheny* River! Can't you get it into your stupid head that I'm talking about the *Susque-goddam-hanna*?

Even if you're too refined ever to use an expression like this, I'm sure you recognize it. Now the interesting thing is that we have pretty clear intuitions about how to use this infix. It sounds natural in the examples in (3), but decidedly odd in those in (4).

(3) uni-goddam-versity
manu-fuckin-facturer
(4) Jacken-bloody-dotff
ele-goddam-phant

In addition, for those words that allow us to use the infix, we are very particular about where it has to go. If we try moving the infix to different places in the words in (3) ("un-goddam-iversity," "manufac-fuckin-turer," etc.) we can see that only the versions given in (3) sound at all acceptable.

I'm fairly certain none of us was ever *taught* the principle (or pattern) that says where it is possible to insert an expletive infix into English words. Yet we readily use this principle to make intuitive judgments about new cases. At the same time, the principle is not so obvious to conscious introspection.

The argument for innate knowledge

Since adults aren't consciously aware of the principles of mental grammar (and the examples just presented provide further illustration), they certainly can't explain these principles to children—if children could understand the explanations in any event!

In fact, the most an adult can do is supply the child with *examples* of the patterns, in the form of grammatical sentences, or corrections to the child's sentences. For instance, notice that in the dialogue I quoted above, the mother isn't saying " 'Nobody' and 'not' are both negative words, and you shouldn't use two negatives in a sentence." She is just supplying the child with a correct *form*. This means that the child has to *figure out* the patterns of the language—that is, *the child has to construct his or her own mental grammar*. How?

Even simpler phenomena show the disparity between children's command of language and their *conscious* command of it. For instance, by the age of three or four, children can be taught to count syllables in a word, but they are certainly making use of syllables long before that.

A question

Discussion question: What does it mean to have a fixed grammar, and how does this relate to children's mistakes…since they obviously waver around a lot? (Jennifer)

The argument for innate knowledge

As a result, we can draw another conclusion about human nature: *We can acquire unconscious patterns unconsciously, with little or no deliberate training.* Perhaps we shouldn't even call such a process "learning," but for lack of a better word, let's leave the terminology alone.

Here is what makes the child's acquisition of language even more remarkable. Thousands of linguists throughout the world have been trying for decades to figure out the principles behind the grammatical patterns of various languages, the very same grammatical principles that children acquire unconsciously. But any linguist will tell you that we are nowhere near a complete account of the mental grammar for any language. In other words, an entire community of highly trained professionals, bringing to bear years of conscious attention and sharing of information, has been unable to duplicate the feat that every normal child accomplishes by the age of ten or so, unconsciously and unaided. This contrast is so striking and so fundamental that it deserves a name. I like to call it the Paradox of Language Acquisition.

The argument for innate knowledge

There are three steps involved in escaping the Paradox. The first two have already been touched on. First, as shown in Chapter 2, what the child ends up with is a mental grammar that is completely inaccessible to consciousness. Hence adult linguists can't figure out the principles of mental grammar just by looking into their minds. Second, as shown in the last section, a substantial part of the language-learning process is also unconscious, so linguists can neither directly observe it nor ask children about it.

But to escape the Paradox, a third step is needed. Remember: children can't just "absorb" mental grammar from the surroundings. All they can hear in the surroundings are sentences; they must (unconsciously) discover for themselves the patterns that permit them both to understand these sentences and to construct new sentences for other people to respond to. Whether this process of discovery goes on unconsciously in the child or consciously in the linguist, the very same problems have to be solved. That is, doing it unconsciously still gives the child no advantage over the linguist.

About the only way anyone has devised to overcome this difficulty is to suppose that *children have a head start on linguistics*: children's unconscious strategies for language learning include some substantial hints about how a mental grammar ought to be constructed. These hints make it relatively easy for them to figure out principles that fit the examples of language they are hearing around them. (Only relatively easy: it still takes them eight or ten years!)

The argument for innate knowledge

Somewhat more technically, the claim is that all of us as children come to the task of language learning equipped with a body of innate knowledge pertaining to language. Using this knowledge, children can find patterns in the stream of language being beamed at them from the environment, and can use these patterns as a mental grammar. Because this innate knowledge must be sufficient to construct a mental grammar for any of the languages of the world, linguists call it *Universal Grammar* or *UG*.

1. What do children know (unconsciously) about language in advance of language learning? That is, what is Universal Grammar?
2. How do they use Universal Grammar to construct a mental grammar?
3. How do they acquire Universal Grammar?

Some objections

The child's acquisition of language clearly depends on exposure to language in the environment. So why should we believe that it is genetically determined?

The answer is that one's language ability is a complex combination of nature and nurture. A biological comparison may be helpful here. Our bone structure is obviously genetically determined, but it can't develop properly without nourishment and exercise. In this case it's clear that environmental interaction complements genetic endowment: both are necessary. Why shouldn't the same be true of the brain structure that supports language, where "nourishment" includes a sufficient quantity and variety of incoming information, and "exercise" includes the opportunity to converse with people?

Some objections

I don't mind the idea of a genetic component to learning, so long as none of it is specifically linguistic, that is, if it consists only of general-purpose learning strategies such as stimulus–response learning or principles of association or analogy. But why do you insist that there is a genetic component of learning that has to do specifically with language?

The problem is that general-purpose learning strategies alone can't solve the Paradox of Language Acquisition. Adults, including adult linguists, have access to plenty of general purpose strategies, but they can't figure out the organization of mental grammar. We have to suppose that children know something more, something specifically about *language*. Still, this is not to say that language acquisition doesn't make use of more general learning strategies. It is just that this cannot be all there is.

Additional Discussion:

Domain-general vs. domain-specific? What's the deal with Universal Grammar? (Emily)

Universal Grammar

Mental Grammar = Innate Part (Universal Grammar) + Learned Part

1. If the language in question is different from other languages in some respect, the child must be able to acquire this difference, so it must fall into the learned part.
2. If certain aspects of all languages we have examined are alike, these aspects are *likely* to fall into the innate part. Of course, there is always the possibility that they are alike purely by accident. In practice, this can be checked out by examining more languages, preferably unrelated ones.
3. Suppose there is some aspect of language that children couldn't possibly figure out from the evidence in the speech they hear around them. Then this aspect can't be learned; it has to fall under the innate part of the language.

The last of these criteria has been called the "poverty of the stimulus" argument. Its use requires a certain amount of care, and in fact there is a running debate on what sorts of evidence children are capable of using.

Poverty of the Stimulus: We Know More than We See

Lightfoot (1999): examples of PoS

Consider some subtleties that people are not consciously aware of. The verb *is* may be used in its full form or its reduced form: people say *Kim is happy* or *Kim's happy*. However, certain instances of *is* never reduce: for example, the underlined items in *Kim is happier than Tim is* or *I wonder where the concert is on Wednesday*. Most people are not aware of this, but we all know subconsciously not to use the reduced form here. How did we come to know this? As children, we were not instructed to avoid the reduced form in certain places. Yet, all children typically attain the ability to use the forms in the adult fashion, and the ability is quite independent of intelligence level or educational background. Children attain this ability early in their linguistic development. More significantly, children do not try out the nonoccurring forms as if testing a hypothesis, in the way that they "experiment" by using forms like *goed* and *taked*. The ability emerges perfectly, as if by magic.

Poverty of the Stimulus: We Know More than We See

1 (a) Jay hurt his nose.
(b) Jay's brother hurt him.
(c) Jay said he hurt Ray.
(d) Jay hurt him.

As adults we generalize that a pronoun may refer to a preceding noun *except* under very precise conditions (1d). But then, how did we all acquire the right generalization, particularly knowledge of the exception?

Problem: because we were not informed about what cannot occur, our childhood experience provided no evidence for the "except" clause, that pronouns sometimes do not co-refer. That is, we had evidence for generalizations like "*is* may be pronounced *z*" and "pronouns may refer to a preceding noun," but no evidence for where these generalizations break down.

Possible solution: reliable imitation?

One is to say that children do not overgeneralize, because they are reliable imitators. That is, children do not produce the reduced *is* in the wrong place or use a pronoun in (1d) wrongly to refer to Jay, because they never hear language being used in this way. In other words, children acquire their native language simply by imitating the speech of their elders. We know that this approach is not tenable, because everybody constantly says things they have never heard. We express thoughts with no conscious or subconscious consideration of whether we are imitating somebody else's use of language. This is true of the most trivial speech: in saying *I wanna catch the 3:25 p.m. bus, which leaves from outside Border's bookstore*, I am using a sentence that I have almost certainly not heard.

[illegible] emerges through an interaction between our genetic inheritance and the linguistic environment to which we happen to be exposed. English-speaking children learn from their environment that the verb *is* may be pronounced iz or z, and native principles prevent the reduced form from being used in the wrong places. Likewise, children learn from their environment that *he, his*, etc. are pronouns, and native principles dictate where pronouns may not refer to a preceding noun. The interaction of the environmental information and the native principles accounts for how the relevant properties emerges in an English-speaking child.

Back to grammars

If asked to say quite generally what is now known about the linguistic genotype, I would say that it permits finite grammars, because they are represented in the finite space of the brain, but that they range over infinity. Finite grammars consist of a limited, precise set of operations which allow for infinite variation in the expressions that are generated.

To see the kind of compositionality involved, consider how words combine. Words are members of categories like noun (N), verb (V), preposition (P), adjective/adverb (A), etc. If two words combine, then the grammatical properties of the resulting phrase are determined by one of the two words, which we call the "head." So, if we combine the verb *visit* with the noun *Chicago*, the resulting phrase *visit Chicago* has verbal, not nominal, properties. It occurs where verbs occur and not where nouns occur: *I want to visit Chicago*, but not **the visit Chicago*, or **We discussed visit Chicago*. So the expression *visit Chicago* is a verb phrase (VP), where the verb *visit* is the head projecting to VP.

(a) $_{VP}$[$_{V}$[visit] $_{N}$[Chicago]]

(b) 

Grammars & Compositionality

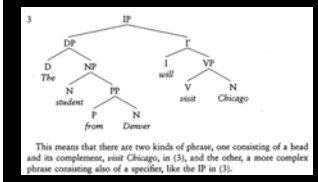

This means that there are two kinds of phrase, one consisting of a head and its complement, *visit Chicago*, in (3), and the other, a more complex phrase consisting also of a specifier, like the IP in (3).

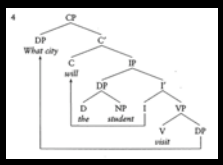

Comparison with Vision (Jackendoff 1994)

Next I'd like to show that vision, too, is governed by a mental grammar. What could this mean? Don't we just see things out there? I'll try to make clear what the problem is, drawing on a long and rich tradition of research on visual perception.

Remember in Chapter 2, when we had Harry looking at a tree, and we represented what he saw by drawing a little tree inside a cloud in his head? Here it is again.

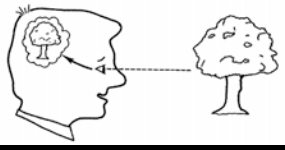

Comparison with Vision (Jackendoff 1994)

We might be tempted to answer that the brain *interprets* the images in the head. However, the brain can't interpret the images by *seeing* them—it doesn't have visual organs.

We have to think about the process somewhat differently. When light strikes the eyes, the lenses focus it to produce images on the retinas. But nobody looks at those images either. Rather, the retinas convert the light into patterns of neural impulses, and it's nothing but neural impulses from there on out.

Some areas of the brain reproduce the spatial arrangement of the neural impulses coming from the retinas; these areas are often said to have "retinotopic maps." But again, these maps aren't images as such; they're not at all like pictures in the head. They're merely more patterns of neural impulses, and they have to be processed further to arrive at what I'll call *visual understanding*.

Comparison with Vision (Jackendoff 1994)

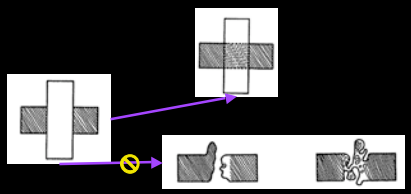

In order for us to use this pattern to understand Figure 13.4, it has to be there somewhere in our brains. And it has to be different from a visual image (or a retinotopic map), because it is more general than any image can be—it applies to vast numbers of different situations. It is a tool with which we *comprehend* visual situations.

Comparison with Vision (Jackendoff 1994)

Impossible figure: Snakes

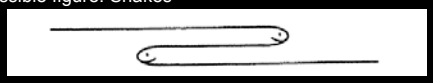

In this case the principle is, very roughly, that a contour (or boundary) in the visual field has to separate the inside of a region from the outside background in a consistent way. (For convenience, I'll call this the Contour Principle.)

Intuitively, the oddness here has to do with insides and outsides: somehow the inside of the top snake merges indistinguishably into the outside of the bottom snake, and vice versa, violating the Contour Principle.

Comparison with Vision (Jackendoff 1994)

It should be clear that I'm going through these examples in order to construct a visual analogue of the Argument for Mental Grammar. Our being able to comprehend an unlimited number of visual situations depends on our having in our brains a set of unconscious patterns and principles that can analyze a visual image and create an interpretation. We can call this set of patterns and principles a visual grammar—though it will probably bear little substantive resemblance to a linguistic grammar. (I don't care much whether we actually use the term "grammar," which some readers may find strained. The point is that the principles in the brain abstract away from particular visual images and help to organize what we see.)

The "invisible contours" in Figures 13.4 and 13.7, the visual ambiguity in Figure 13.9, and the strangeness of Figure 13.10 are certainly part of our visual experience. But there is no sense in which they are actually in the physical world—any more than words are physically present in an acoustic signal. Yet we can't help experiencing these figures the way we do: unconscious processes in the visual system make use of stored visual patterns and principles to construct an optimal understanding of the optical signal presented to the eyes. Every time we see something behind something else, and every time we see something standing out against its background, we are making use of principles of visual understanding illustrated here.

Comparison with Vision (Jackendoff 1994)

Innate knowledge

The evidence that has accumulated suggests that an overwhelming proportion of the basic principles of visual perception (for instance, the Contour Principle) are innate.

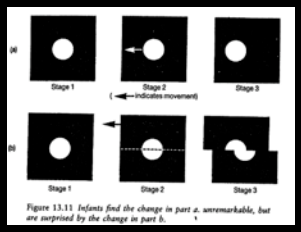

Figure 13.11 *Infants find the change in part a unremarkable, but are surprised by the change in part b.*

Comparison with Vision (Jackendoff 1994)

The Argument for Mental Grammar: In each of these domains, our ability to make sense of novel stimuli is supported by a set of abstract patterns that are specialized for that domain.

The Argument for Innate Knowledge: In each of these domains, we learn the patterns we do in part because our brains are genetically programmed with substantial aspects of these patterns in advance. Learning is not "soaking up" of patterns, but rather active tuning and elaboration of the innate specialized mental "proto-patterns." That is, a great deal of nature lies behind learning through nurture.