2 The argument for mental grammar

The communicative situation

Let's start with a fairly crude picture of the communicative situation—what goes on when one person says something to another.

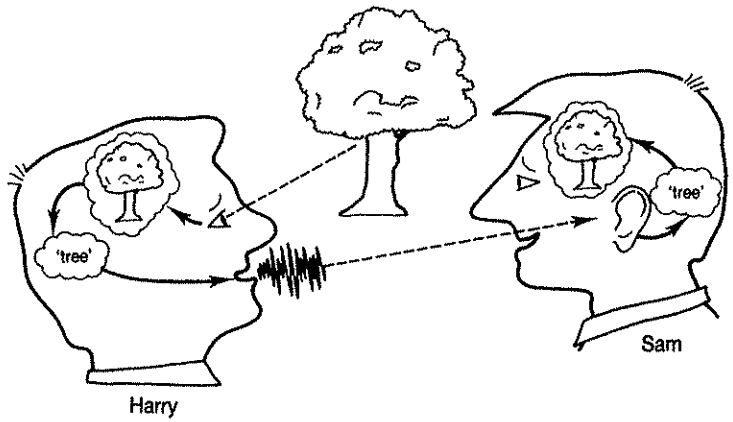

Figure 2.1 *The communicative situation: Harry tells Sam about a tree*

In this picture (Figure 2.1), a pattern of light reflected off of a tree strikes the eyes of the person on the left (let's call him Harry). As a result of activity in Harry's nervous system, he comes to see the tree out there in the world. This is indicated in the picture by a tree inside a little cloud in Harry's head. Of course we know there are no clouds or trees in people's heads, and eventually (Chapter 13) we'll ask what's really there, but let this stand for the moment.

Once Harry has perceived a tree, it may occur to him that the

word "tree" describes what he has seen—that is, the word "tree" is evoked from his memory. (If Harry spoke French rather than English, of course, the word "arbre" would be evoked instead.) This is indicated in the picture by the word "tree" in another little cloud in Harry's head. Again, we know there are really no little clouds, but this will have to do for now (we'll come back to it in Chapter 4).

Perhaps Harry decides to say something about the tree to the person on the right (let's call him Sam). Then Harry's nervous system causes his lungs to expel air, his vocal cords to tighten, and his tongue and jaw and lips to go through some gyrations. As a result, he produces some sound waves which travel through the air, striking Sam's ears, Sam's eyes, the furniture, and everything else.

But unlike Sam's eyes and the furniture, Sam's ears react to these sound waves by activating Sam's nervous system, so that he comes to perceive Harry uttering the word "tree." Assuming Sam also speaks English, his nervous system very likely goes on to produce a visual image of a tree—Sam is able to imagine what Harry sees, though probably not in many of its particulars.

Even this little dissection of the obvious has revealed quite a lot of complexity. There are a lot of parts to this simple communicative act, and each one of them involves tough puzzles. (For example: What is really in Harry's and Sam's brains instead of the little clouds? Exactly what gyrations of Harry's tongue, jaw, and lips take place? What happens in Sam's ears?) But we still haven't seen the full difficulty of the problem.

Suppose Harry wants to say something a little more interesting about what he sees than the single word "tree." Here are some things he might say (I'll put numbers and letters in front of example sentences so we can refer back to them later):

(1) *a* There's a bird in the tree.
b A bird was in the tree yesterday.
c Are there any birds in that tree?
d A bird might be in the tree.
e Birds like that tree.
f That tree looks like a bird.

This time it isn't so easy to draw pictures in little clouds that depict what Harry has in mind. What difference can we make in the pictures in order to distinguish sentences (1a), (1b), (1c), and (1d)? (If we start putting question marks and writing in a picture, that's cheating—it's not just a picture anymore!) For sentence (1e), how do we show that the birds *like* the tree rather than, say, merely swarm around it?

In all the cases so far, the picture at least has both a bird and a tree in it, whatever its other failings. But what about sentence (1f)? What seems to come to mind is something like a bird-shaped tree. But such a picture has only one object corresponding to both words—yet another complication.

These examples illustrate some of the *expressive variety* of language—the number of different things we can say by combining words in different ways. Moreover, this expressive variety in many respects can't be conveyed by pictures, whether on a piece of paper or in the head. That is, significant parts of the messages that language conveys are abstract, or nonsensory, in nature.

Here we see a significant difference between human language and any of the forms of animal communication. To be sure, many kinds of animals convey information to each other. But in none of the known systems—birds, bees, whales, nonhuman primates, or whatever—is there an inventory of elements like words that can be combined and recombined in limitless new ways to express new messages. There are no elements that indicate points in time ("yesterday"), a desire for information ("are there ... "), or possibility ("might"). Animals may have a way to indicate their own desires or feelings, but they can't convey someone else's, as in "Birds like that tree." Nor can animal communication systems explicitly draw resemblances among different objects, as in "That tree looks like a bird."

So, although people often speak loosely of animal communication as a kind of language, in fact the way animals communicate is orders of magnitude different from the way humans do. To make this distinction clear, I will adopt the policy of using the word "language" to mean only "human language (Spanish, Chinese, Navajo, etc.)," and I will use the more general word "communication" for any means by which information is conveyed, including both language and animal systems. (I will mention some attempts to teach human languages to apes in Chapter 10.)

The argument for mental grammar: The expressive variety of language use implies that a language user's brain contains unconscious grammatical principles

The expressive variety of language is the springboard for the first of the Fundamental Arguments. Any normal human being can understand and create an indefinitely large number of sentences in his or her native language. Aside from stereotyped utterances like "Hi, how are you?" and "Please pass the salt," most of the sentences we speak in the course of a day are sentences we have never heard or spoken in their entirety before. The same is true of most of the sentences we hear. For example, I doubt that you have ever heard or spoken any of the sentences on this page before. Yet you have no difficulty understanding them.

Let's think about what must be going on in your head that makes this possible. In the previous section we assumed that Harry and Sam could simply pull the word "tree" out of their memories when needed. Could this be true of whole sentences as well?

No. The number of sentences we are capable of using is just too large to store them individually. Let me run up the number in some rather stupid ways, just as a sample. Consider this series of sentences, all of which are perfectly comprehensible.

(2) Amy ate two peanuts.
Amy ate three peanuts.
Amy ate four peanuts.
...
Amy ate forty-three million, five hundred nine peanuts.
...
...

There are as many sentences in this series as there are nameable integers. The biggest number name listed in my Webster's Collegiate is a vigintillion (10^{63} in US/French usage; 10^{120} in British/German usage). With all the numbers up to this at our disposal, we can create more sentences in this series than there are elementary particles in the universe.

Here's another way to make lots of sentences. There are at least some tens of thousands of nouns in English. Let's be conservative and say we know ten thousand (10^4). Now let's construct all the sentences we can by putting in different nouns for X and Y in "An X is not a Y." Here are some of them.

(3) A numeral is not a numbskull.
A numeral is not a nun.
A numeral is not a nunnery.
...
A numbskull is not a numeral.
A numbskull is not a nun.
A numbskull is not a nunnery.
...

A nun is not a nursery.

. . .

An oboe is not an octopus.

. . .

These are all completely absurd, but they *are* sentences of English nevertheless. There will be something like $10^4 \times 10^4$ of them = 10^8. Now let's put pairs of these sentences together with "since," like this:

(4) Since a numeral is not a numbskull, a numbskull is not a nun.
Since a numeral is not a numbskull, a numbskull is not a nunnery.
Since a numeral is not a numbskull, a numbskull is not a nuptial.
. . .
Since a numeral is not a nursery, a numbskull is not a nun.
. . .
Since an oboe is not an octopus, a numeral is not a numbskull.
. . .

And so on it goes, giving us $10^8 \times 10^8 = 10^{16}$ absolutely ridiculous sentences. Given that there are on the order of ten billion (10^{10}) neurons in the entire human brain, this divides out to 10^6, or one million sentences per neuron. Thus it would be impossible for us to store them all in our brains, in the unlikely event that we should ever want to use or understand any of them. But still, you did just understand a sampling of them. And these lists are only a minute proportion of the sentences you can understand. What lists include the sentences of this paragraph, for instance?

In short, we can't possibly keep in memory all the sentences we are likely to encounter or want to use—not to mention all the unlikely ones such as the sentences in (2)–(4). On the other hand, we are apparently ready to encounter them—we seem to know what the possibilities are.

The way the brain seems to achieve expressive variety is to store not whole sentences, but rather words and their meanings, plus *patterns* into which words can be placed. For example, it is only by using patterns that we can reasonably store the sets of sentences of which (2), (3), and (4) form a tiny sample: the pattern for the sentences in (2) is "Amy ate N peanuts"; that for the sentences in (3) is "An X is not a Y"; and that for the sentences in (4) is "Since an X is not a Y, a Z is not a W." With such patterns, plus a list of words to insert into them, we can specify a large number of possibilities at minimal cost in storage. Moreover, such a system is prepared for *novelty*: it can recognize or create examples of the pattern on the spur of the moment, whether or not they have been encountered before.

But even using these kinds of fixed patterns isn't quite good enough. Consider the list of sentences in (5).

(5) *a* Bill thinks that Beth is a genius.
b Sue suspects that Bill thinks that Beth is a genius.
c Charlie said that Sue suspects that Bill thinks that Beth is a genius.
d Jean knows that Charlie said that Sue suspects that Bill thinks that Beth is a genius.
. . .

This sequence can be extended as long as we have the patience—that is, it is effectively infinite. (To be more precise, there is no longest sentence in this sequence, because we can always add one more.) As a result, we can't specify a single pattern for this list the way we could for the lists sampled in (2)–(4). Rather, each sentence has to come from a different pattern, and the patterns get longer and longer. (6) shows the first three of these patterns; the term "Verbs" stands for one of the words "thinks," "suspects," "knows," and so forth.

(6) X Verbs that Y is a Z.
W Verbs that X Verbs that Y is a Z.
T Verbs that W Verbs that X Verbs that Y is a Z.
. . .

Can we store all these patterns in our heads? Again, no, because no matter how many we store, there is always a longer one. On the other hand, there is clearly a more basic pattern involved: given any declarative sentence, we can make another declarative sentence by placing "X Verbs that . . . " in front of it. For instance, we can apply this pattern to any of the sentences in (2)–(4) above to get whole new classes of sentences. Here are some of them (marking in italics the sentence we started with): "Bill knows that *Amy ate two peanuts*," "Wolfgang realizes that *an oboe is not an octopus*," "Ludwig suspects that *since a numbskull is not a nunnery, a nun is not a nuptial*," and so on. This pattern can be summarized as the formula given in (7).

(7) X Verbs that S. (where *S* is any declarative sentence)

Going back to the sequence of sentences in (5), we can apply

formula (7) to the sentence "Beth is a genius" to get "Bill thinks that *Beth is a genius*," sentence (5a). And then comes the fun: we can use our new sentence as the sentence *S* in formula (7), giving us "Sue suspects that *Bill thinks that Beth is a genius*," sentence (5b); then we can use *this* sentence as *S* in (7), giving us "Charlie said that *Sue suspects that Bill thinks Beth is a genius*," and so on as long as we want. That is, we get longer and longer sentences by applying formula (7) over and over to its own output, or *recursively*. What makes (7) different from the earlier patterns is that it contains another pattern within it: instead of just putting words into the slots in the pattern, we insert another pattern—in this case a whole declarative sentence.

This is a typical case of what we find in the course of investigating the expressive variety of language. The sequences in (8) and (9) show two more patterns with patterns inside them; as in (5), we can go on applying them recursively till our patience runs out.

(8) *a* Ben's father is a linguist.
b Ben's father's older brother is a linguist.
c Ben's father's older brother's best friend is a linguist.
d Ben's father's older brother's best friend's former lover is a linguist.
. . .

(9) *a* This is the house that Jack built.
b This is the refrigerator that sits in the house that Jack built.
c This is the cheese that fell out of the refrigerator that sits in the house that Jack built.
d This is the mold that grew on the cheese that fell out of the refrigerator that sits in the house that Jack built.
. . .

In short, in order for us to be able to speak and understand novel sentences, we have to store in our heads not just the words of our language but also the patterns of sentences possible in our language. These patterns, in turn, describe not just patterns of *words* but also patterns of *patterns*. Linguists refer to these patterns as the *rules* of language stored in memory; they refer to the complete collection of rules as the *mental grammar* of the language, or *grammar* for short.

This demonstration of the expressive variety of English, complete with recursive patterns, can be reproduced in any of the human languages of the world. The particular patterns of mental grammar may not be the same from one language to the next, but

patterns of comparable complexity can always be found. In this respect, there is no difference between the languages of contemporary Western societies, those of present-day "primitive" cultures, and those of the distant past that can be recovered from written records. (An important exception arises in "pidgin" languages, to be discussed in Chapter 10.)

Clarifying the notion of mental grammar

The notion of a mental grammar stored in the brain of a language user is *the* central theoretical construct of modern linguistics. So it's important to make it as clear as possible before going on. Let me engage in a dialogue with an imaginary skeptic who raises some of the most common questions and objections.

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.

Why don't they make sense?

Well, these are sentences I've never heard before.

But look: You never heard the sentences in (2)–(5) before either, and even so, they "make sense" (albeit of a stupid sort).

What's the difference?

The difference is that the sentences in (2)–(5) are examples of patterns of English that we know, and the strings of words in (10) are not. 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.

This is not to say that conformity to the patterns of English is the only factor involved in "making sense." Lots of sentences conform to the grammatical patterns of English but still don't "make sense."

(11) Colorless green ideas sleep furiously.
Bill elapsed three times this month.

I'm memorizing the score of the sonata I hope to compose someday.
The harvest was clever to agree with you.

These examples (drawn from early writings of Chomsky's) are certainly nonsense. But they do conform to the grammatical patterns of English, as we can see by substituting one or two more "sensible" words in each one:

(12) Large green lizards sleep soundly.
Bill sneezed three times this month.
I'm memorizing the score of the sonata I hope to perform someday.
The lawyer was clever to agree with you.

On the other hand, notice that if we exchange adjacent words in the sentences in (11), so that they violate the grammatical patterns of English, they sound far worse: "Colorless green sleep ideas furiously," "Bill three elapsed times this month," etc. In this case, it's not that they have strange meanings; rather, it's hard to say what they mean at all. So the mental grammar seems to be involved even in sentences like (11) that don't make sense.

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 . . .

These lines are clearly an example of the same pattern as the following, which contains all real words:

'Twas evening, and the slimy toads
Did squirm and wiggle in the cage . . .

This shows that the patterns themselves have a degree of life independent of the words that make them up. Indeed, if you start exchanging words at random in *Jabberwocky*, again the patterns fall apart altogether.

Why do you want to call it a **grammar** *that I store in my head? Why couldn't I just have a bunch of* **habits** *that I follow in speaking and understanding English?*

My return question is: What is a habit anyway? It's something stored in memory that guides behavior on appropriate occasions. If

the "habitual" behavior varies from occasion to occasion, as it does in the case of language, what is stored in memory has got to be a pattern. Why? Because the brain can't store all the individual examples—and even if it could, there would be no reason to call this random collection of behaviors a unified "habit."

Once we realize that habits must themselves be stored patterns, we shouldn't have a problem acknowledging that the "habits" of speaking English involve storing the patterns of English. That is, claiming that our knowledge of English is a kind of habit doesn't eliminate the need for us to have grammars in our heads.

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.

The concept of "mental grammar" provides a different perspective on this issue. The mental grammar in our heads is what enables us to put words together into sentences. So it has to specify not just which patterns are socially acceptable and which are not, but *all* the patterns of the language. This includes some patterns that are much more basic than they ever had to teach us in school, for instance that the subject precedes the verb in English, or that adjectives precede the nouns they apply to ("ripe banana," not "banana ripe"). But it also includes some patterns that are much more complex than those taught in school, as we will see in the next few chapters.

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." In other

words, there are principles that govern the use of "incorrect" English too, even if it violates the canons of school grammar.

This means that such speakers don't *lack* a mental grammar; they just have a mental grammar that is slightly different from that of speakers of "correct" English. Setting aside the issue of social approbation, the situation is exactly parallel to the difference between "proper" British and American English. Speakers of these two dialects have slightly different mental grammars, so the patterns they produce don't match up exactly. Consequently, each sounds somewhat exotic (or sloppy) to the other.

In short, although my imaginary critic may wish to deplore certain people's language from the point of view of school grammar, it is hard to deny that they have a mental grammar in their heads that governs their patterns of speech.

> *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?*

The answer to this question is potentially the most troubling. Here's the situation. We have just seen that an explanation of language ability demands that the patterns of language be stored in our memory somehow. We're now faced with the apparently conflicting fact that our memory reveals no such patterns to us. So something has got to give.

Can we give up the idea of a mental grammar? No: I've tried to convince you that just about any other way of thinking about the expressive variety of language amounts to the same thing. So let's grasp the other horn of the dilemma, and explore the hypothesis that the rules of language are not conscious, and are not available to introspection.

What could such a hypothesis mean? In this post-Freudian age, we are certainly accustomed to speaking of unconsciously (or subconsciously) guided behavior: "Willy has low self-esteem because he unconsciously identifies with his father." The premise of Freudian analysis, as well as most subsequent forms of psychotherapy, is that unconscious beliefs of this sort can be made conscious through suitable therapeutic procedures, and that in becoming conscious they cease to exert the same pernicious influence on one's experience and behavior.

Freud's notion that parts of the mind are not accessible to consciousness challenges the standard Cartesian identification of the

mind with *consciousness*: there is a lot more going on in our minds than we are ever aware of. This is upsetting not only because it goes against intuition ("I *know* what I think!") but also because it tells us we are not altogether in conscious control of our behavior. What's more, the Freudian unconscious is full of dark and uncomfortable motives. (Freud stressed the sexual underpinnings of those motives, perhaps because of his social milieu; modern psychodynamic theory recognizes many other themes as well.)

In a way, the unconsciousness of mental grammar is still more radical than Freud's notion of the unconscious: mental grammar isn't available to consciousness under *any* conditions, therapeutic or otherwise. On the other hand, an unconscious mental grammar that guides our behavior is a good deal less personally threatening than an Oedipus Complex or a Death Instinct. Unlike these Freudian constructs, mental grammar doesn't have pernicious effects. On the contrary, we couldn't speak without it, except in terms of stereotyped fixed expressions. It is mental grammar that makes possible the expressive variety of our language.

> *You're telling me that a mental grammar is present in my mind but that I'll never find it by looking there? Aren't you trying to pull a fast one?*

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.

If mental grammar can't be studied by introspection, then we have to find some other, less direct way to study it. I will take up this problem in Part II, showing how the investigation of mental grammar is an experimental science, and describing some of the organization that has been revealed by linguistic research. For now, the point is that if at least some other processes in the mind are not open to consciousness, it shouldn't be too distasteful to say that parts of language ability are unconscious too.

This, then, is our first inference about human nature on the basis of the nature of language. In order to account for the human

ability to speak and understand novel sentences, we must ascribe to the speaker's mind a mental grammar that specifies possible sentence patterns. But in order to account for the fact that we have no direct access to this mental grammar, we must admit the possibility that some essential and highly structured parts of our abilities are completely unconscious.

3 The argument for innate knowledge

The character of language acquisition

We now turn to the preliminaries to the second Fundamental Argument. Suppose, following the discussion of the previous chapter, that we have mental grammars in our heads. The next question is: How did they get there?

Observation: All normal human children end up being able to speak whatever language is spoken in the community where they grow up. (If more than one language is spoken regularly, they usually end up speaking them all—but let's stick to the monolingual case for now.) And the language they speak has nothing to do with where their parents came from: a child of American parents growing up in Israel as part of a Hebrew-speaking community will become a native speaker of Hebrew; a Vietnamese baby adopted in Holland will become a native speaker of Dutch. So it's pretty obvious that children learn their language from the other speakers around them.

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. Nor is the children's knowledge necessarily a result of teaching in school—and of course in nonliterate