11 Language and the brain

The frontiers of present-day brain science

Chapter 4 presented the functionalist approach to language in the brain: studying the abstract organization of the patterns of language, without worrying too much about how the brain actually encodes these patterns. Most linguists work within the functionalist approach, and all our discussion up to this point has been from the functionalist point of view.

But this still leaves a major mystery. How does the brain do it? What do the patterns of mental grammar look like in the neurons? How does an English speaker's brain differ from a French speaker's— or an ASL speaker's? When (from the functionalist point of view) speakers put together collections of mental patterns in Lego fashion to form long and intricately structured sentences, what are their neurons doing?

These sorts of questions pertain to the Argument for Mental Grammar. Similar questions arise for Argument for Innate Knowledge and its implications for language learning. What does Universal Grammar, the knowledge that a child brings to the language learning process, look like in terms of brain structure? When a child learns a language, how does the brain change? How do patterns of growth in brain structure over time account for the observed stages of language acquisition, and how does the brain spontaneously change by the end of the critical period so as to make language learning more difficult?

The past twenty or thirty years have seen remarkable advances in brain science, and from the media one might get the impression that a full understanding of how the brain works is not too far off. At the risk of seeming unduly pessimistic, though, I have the sense that the main lines of research in neuroscience, exciting as they are, are still not ready to approach the kinds of questions about language I've just posed. Let me try to summarize how far these approaches have gotten.

First the good news. It has become abundantly clear that the brain is not a big general-purpose device—its many functions are highly localized. Even vision, which on the face of it is a homogeneous undifferentiated process, proves to be broken up into numerous subfunctions (or *modules* in the sense of Chapter 4), each of which has its own area in the brain. There are separate brain areas devoted to the detection of the location, the shape, and the color of objects in the environment, all fed by the neurons of the optic nerve. There is even an area whose main function seems to be the recognition of familiar faces, a function whose purpose we'll discuss in Chapter 15.

It is also clear that these areas of the brain are not located where they are by chance. Though people vary to some extent in the precise location of particular brain areas, we find, for example, that the area for shape recognition is always in the temporal lobe, and that for location recognition is always in the parietal lobe—never the other way around. And brain areas with analogous functions are often found in analogous places in lower primates. So it looks as though this micro-differentiation of brain structure and function is genetically driven.

How are these areas identified? The most highly publicized current techniques are CAT (computerized axial tomography) scans, MRI (magnetic resonance imagery) scans, and PET (positron emission tomography) scans. CAT and MRI scans can provide exquisite images of details of brain anatomy; PET scans, while showing less detail, can detect what parts of the brain are most active while one is carrying out different tasks such as looking, listening, speaking, or doing mental arithmetic. PET scans reveal considerable differentiation even among apparently similar tasks: for example, counting aloud shows activity in different areas from counting to oneself. PET scans can also be used to show abnormalities in patterns of brain activity, for instance in schizophrenia.

Older techniques for detecting brain localization include EEG (electro-encephalogram) and, more gruesomely, electrical stimulation of different parts of a conscious patient's brain while it is open for an operation. (Since the brain itself doesn't have pain receptors, this doesn't hurt.) Animal experiments provide the further option of implanting electrodes in the brain, from which the activity of single neurons can be recorded while the animal is roaming around doing things.

But by far the most common approach to brain localization is

to work with people who have sustained brain damage through injury, a stroke, a tumor, or a brain operation (to remove the focus of epileptic seizures, for instance). The area of damage can now be identified by means of a CAT scan or an MRI scan; in the old days it had to await an autopsy, which made research much more difficult. Numerous experimental techniques have been devised to discover what brain-damaged people can and cannot do.

Some of the deficits discovered by these techniques are fairly amazing. For instance, at the very back of the brain is the striate cortex, one of the first visual areas to receive input from the eyes. Damage to this region on one side of the brain produces blindness in the opposite half of the visual field. If people with such damage are shown things in such a way that they have to detect them with the blind part of the visual field, they will fervently deny that they are seeing anything. But suppose the experimenter says "Well, just for fun, imagine there *is* something there. Where do you think it might be? I know this is stupid, but would you point to it?" Under such prodding, they do surprisingly well. The recent discovery of this phenomenon of "blindsight" has stimulated a great deal of excitement among not only neuroscientists but philosophers as well, for it challenges our common-sense notions of what it means to see and to be conscious.

Much more widely known are the unfortunate people who, in order to control seizures, have undergone an operation that severs the corpus callosum, the large bundle of fibers that connects the left and right hemispheres of the cortex. If such people are shown a salacious picture in the left-hand visual field, they will blush or giggle appropriately. But if asked what they saw, they will honestly say "Nothing," for the language areas of the brain, in the left hemisphere, can't get information from the areas in the right hemisphere responsible for the left visual field. How are we to understand the behavior of such people? Do they have two independent minds, or two independent seats of consciousness? Or is the right hemisphere unconscious? It's hard to know what to say, and reputable people come down on all different sides of the issues.

I'd like to add that many popularizations of brain science focus almost exclusively on the asymmetry of the left and right hemispheres, speaking of the left hemisphere as "analytic" and the right as "holistic" and even "emotional." While there is some truth to this, at the same time each hemisphere is heavily broken down into further specializations, of which some are asymmetrical and some are not. For instance, language is usually concentrated in the left hemisphere,

face recognition in the right. But low-level visual processes, tactile sensation, and motor control are divided pretty much symmetrically between the two hemispheres. And in addition to the left–right asymmetry, neuroscientists talk about differences of "style" between the top and bottom of the brain, and between the front and the back. So the highly publicized left–right distinction is only one of many.

Such a picture of brain localization is, of course, altogether congenial to the overall view we have reached in studying language from the functionalist approach. If the brain in general is made up of a lot of specialized modules, and if these specialties develop under genetic control, then it is hardly outrageous to use the Argument for Innate Knowledge as evidence for a genetically determined specialization for language.

On the other hand, establishing that language (or any other function) is specialized in the brain only gets us down to a certain fairly crude level of detail. It doesn't tell us how those functions *work*. Imagine someone explaining how a TV set works: "There are a number of specialized devices in the set. There is an area that tunes in the signal from the antenna and responds to your changing the channel. It sends a signal to another area, which splits the signal into sound and picture. The sound signal goes to another area, which. . ." We couldn't be blamed for complaining, "This is helpful up to a point, but *how* does the first area tune in a signal, *how* does the next area split the signal? How do they work?"

This is about the way I feel about brain localization studies. How does the specialized linguistic part of the brain combine speech sounds into syllables? How does it combine adjectives and nouns into noun phrases? How do the neurons store the word "banana"—its phonological structure and the fact that it is a noun? What do nouns have in common in the way they are stored in the brain that makes them different from verbs? And so forth. The fact that brain functions are localized may tell us that nouns and verbs are stored somewhere different in the brain from, say, smells and tastes—or even that nouns and verbs are stored in different places from each other. But that doesn't help us much on the question of how nouns and verbs work neurally within their specialized areas, or of how the neurons bring nouns and verbs together to produce and understand sentences.

Another active area of brain research concerns the chemical substances called neurotransmitters, which affect different types of neurons, leading to overall changes in mood, attention, and body control. A well-known case is the interaction of the neurotransmitter

dopamine with the basal ganglia that control voluntary movement; a disruption or depletion of dopamine leads to the movement disorders of Parkinson's disease. This sort of research, while it is extremely exciting and of great medical importance, is also at too coarse a scale for our purposes. We need to know more than whether the speech centers are generally stimulated or inhibited. The operation of mental grammar has to involve the fine structure within the speech areas of the brain.

What about research on the fine structure of the nervous system? A lot is known about how an individual neuron works: what goes on when a neuron fires, how neurons communicate with each other through their synaptic connections, and how a neuron can change its pattern of behavior in reaction to new patterns of input (that is, the neural basis of learning). But again, while this research is fascinating, it does not immediately help us with the question of how the neurons accomplish language behavior, other than to rule out certain oversimplistic hypotheses. Any facet of language has to involve more than a single neuron and its synapses. It's not as though there's a particular cell for the word "icecream" that fires whenever you hear the word or think of it, or that there's a single cell for "noun" that fires whenever you use a noun. The problem is how larger assemblies of cells are wired up to perform these functions. So this kind of research is *too* fine-scale for our purposes.

To develop a neural account of mental grammar, we'll have to understand the combinatorial behavior of assemblies of neurons— how each one of a group of neurons (how many? hundreds? thousands? millions?) is reacting to the others and to inputs presented to the system as a whole. Such work is in its infancy. To understand a system of a couple of dozen interacting neurons in detail is, for now, a real *tour de force.* Larger systems of so-called "neural networks" are being extensively modeled on computers (this line of research is also called "connectionism" or "parallel distributed processing"). But on one hand these models leave out many properties of real neurons, and on the other hand they don't approach the richness of language that has been described under the functionalist approach.

This is not to denigrate the many important advances made by neuroscience over the past decades. It is just that there still seems to be a long way to go before we will be able to answer the kinds of questions about the brain that are raised by the study of the structure and learning of language. That should not discourage us from asking such questions, or from keeping our eyes open for hints from neuroscience for where answers might come from.

Language deficits due to brain damage

With all these caveats in place, let me turn to a brief survey of what *is* known about language and the brain, and what sorts of conclusions might be drawn.

The study of language deficits due to brain damage has a long history. In fact, the localization of language in the brain was one of the earliest examples of localization to be discovered, and it served as a major impetus to other brain research.

In 1864, the French surgeon Paul Broca showed that patients with a particular set of linguistic deficits proved upon autopsy to have damage to a particular area in the left frontal lobe. He also pointed out that corresponding right-hemisphere damage showed little effect on speech. The area in question has come to be known as *Broca's area*, and the set of symptoms as *Broca's aphasia*. Broca's aphasics seem to understand what is said to them, but their speech is slow, effortful, and poorly articulated. They have evident difficulty in finding words. Here are two representative quotes from Broca's aphasics (which in the transcripts I have available do not record the problems in pronunciation):

(1) *a* Me . . . build–ing . . . chairs, no, no cab–in–nets. One, saw . . . then, cutting wood . . . working . . .
    b Cookie jar . . . fall over . . . chair . . . water . . . empty . . . ov . . . ov . . . [Examiner: "overflow"] Yeah.

Notice that these quotes are not just slowed-down sentences. They are missing a great deal of the grammatical tissue that holds normal speech together—things like articles, auxiliary verbs, and tenses. For this reason Broca's aphasia is also called *agrammatism*. (I should add that sometimes Broca's aphasics can still swear fluently, and they can often sing.)

A quite different set of symptoms was identified by Karl Wernicke in 1874. In people with *Wernicke's aphasia*, Broca's area is intact, but there is damage in the left temporal lobe, in an area now called *Wernicke's area*. The speech of Wernicke's aphasics is altogether fluent—if anything, it tends to come out in a big rush. Taken a few words at a time it often makes a little sense, but the larger parts don't fit at all, and there are often insertions of nonsense words:

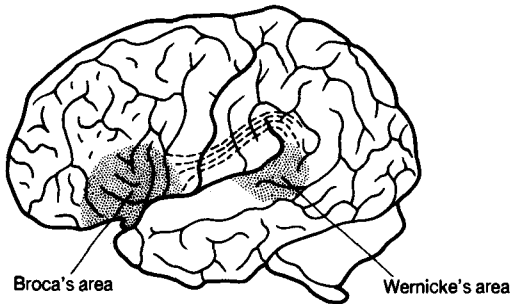(2) *a* [Examiner: "What kind of work have you done?"] We,

Figure 11.1 *The areas of the left hemisphere relevant to language*

    the kids, all of us, and I, we were working for a long
    time in the . . . you know . . . it's the kind of space, I
    mean place rear to the spedwan . . . [Examiner:
    "Excuse me, but I wanted to know what work you
    have been doing."] If you had said that, we had said
    that, poomer, near the fortunate, forpunate, tamppoo,
    all around the fourth of martz. Oh, I get all confused.
b Well, this is . . . mother is away here working out
    o'here to get her better, but when she's working, the
    two boys looking in the other part. One their small
    tile into her time here. She's working another time
    because she's getting, too.

Not only don't Wernicke's aphasics make much sense, they don't
seem to understand much of what you say to them. For instance, for
the most part they don't follow instructions well. By contrast with
Broca's aphasics, who are painfully aware of their deficits, Wernicke's
aphasics often seem unaware that they aren't making sense, and they
become irritated at people who can't understand them.

    There are many other sorts of language deficits due to brain
damage. *Anomic aphasics* have word-finding deficits, more extensive
versions of the experiences we all have occasionally in not being able
to think of a word or a name. *Conduction aphasics* have relatively
fluent meaningful speech, but with many errors of pronunciation;
and, oddly, they have great difficulty repeating sentences spoken to
them. A whole class of deficits involves reading: there are patients

who can speak but not read, read but not write, write but not read (even what they've just written!), and other combinations. One of the strangest afflictions (for my taste) is *deep dyslexia*: if a deep dyslexic is asked to read the word "apple," he may say "fruit" or "flower"—something related but not the same. The word seems to get in and influence what comes out, but the person doesn't (and possibly can't) check whether the response is correct.

In practice, these symptoms are rarely entirely "clean": a random stroke seldom affects exactly one area and spares everything around it. In addition, it isn't clear that the boundaries of any of the areas in the brain are altogether "clean" either. But, along with most of the literature, we'll accept the conventional idealizations for now.

What has gone wrong in these aphasias? One widespread approach seeks to view aphasia as a disorder of general conceptual thinking. There are indeed afflictions of *dementia* that result in generalized diminutions of cognitive functions, usually as a result of widespread brain deterioration due for instance to Alzheimer's disease. These may simultaneously affect naming, ability to panto-mime, purposeful movement, memory, and reasoning. But aphasias are not like this: they often leave the nonlinguistic abilities intact. And, conversely, the ability to produce grammatical sentences is often spared in the face of other substantial cognitive loss due to brain damage. This points to language ability being quite distinct from general-purpose cognitive functioning.*

Wernicke himself proposed a different account of aphasia. He observed that Wernicke's area is near the area of the brain involved in hearing, and that Broca's area is adjacent to the area that controls motor movements of the vocal tract. He suggested, therefore, that Wernicke's area stores the auditory memories of words, and Broca's area stores the memories for how to pronounce them. This nicely explains the fact that Wernicke's aphasics can articulate language but can't understand it (even their own), and that the reverse is true of Broca's aphasics.

However, while Wernicke's account acknowledges the special-ization of language ability, it is still inadequate. Language doesn't just consist of knowledge of the sounds of words and how to pronounce them. Rather, auditory and motor abilities are only the most superficial part of language. Most of the interest is in the abstract

---

* Notice that this evidence for the specialization of the language capacity parallels the evidence drawn from the different varieties of genetic brain impairment discussed in Chapter 9.

phonological and syntactic patterns that organize both perception and production. Wernicke's theory has nothing useful to say about how these abilities are manifested in the brain, and how they could be disrupted.*

A serious difficulty with Wernicke's approach was discovered in the 1970s. Up to that time, it had been thought that Broca's aphasics understand what is said to them, and just have trouble speaking. So Broca's aphasia was viewed purely as a problem with production. It turns out, though, that most Broca's aphasics have problems understanding language as well. Their difficulties, in fact, turn on just the sorts of things they most typically leave out in their speech: the parts of language that signal grammatical structure.

Here's one kind of experiment that reveals this comprehension deficit. Suppose we show Broca's aphasics pictures of a boy hitting a girl and of a girl hitting a boy. Then we ask them: "Which of these pictures goes with the sentence I'm now going to say to you?" If we say (3a), they get the right answer, but if we say (3b), they choose randomly.

(3) *a* The boy hit the girl.
    b The boy was hit by the girl.

Apparently, the "was" and "by" of the passive sentence (3b), both functional words, are causing them difficulty.

Similarly, Broca's aphasics have trouble telling the difference between sentences (4a) and (4b). The reason, evidently, is that they cannot take account of the functional word "the," so both sentences sound to them like (4c).

(4) *a* He showed her baby the pictures. (Who saw the
                                        pictures? The baby did)
    b He showed her the baby pictures. (She did)
    c He showed her baby pictures. (ambiguous)

(Incidentally, notice that Broca's aphasics are not simply *ignoring* functional words; if they did, they would always guess that (3b) means the same as (3a). So there is some more subtle story here.)

Since Broca's aphasia usually affects both comprehension and

---

* You can't blame Wernicke; hardly anything was known about the psychology of language in 1874. In the light of what is now known about language, though, it is hard today to condone explaining aphasia by means of Wernicke's hypothesis, as was done in a recent public television series on the brain.

production, it looks as though it is a deficit in a part of the mental grammar that is shared between the two—in the processing of phonological or syntactic structure. By contrast, Wernicke's aphasia appears to be a disruption of the linkage between language and thought—also in both comprehension and production. Returning to our functional diagram of the crude organization of language, we can place these aphasias approximately in the regions designated in Figure 11.2.
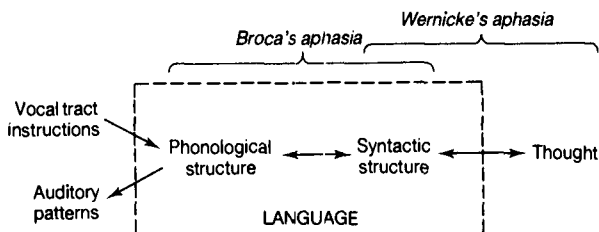

Figure 11.2 *The areas of linguistic information affected by Broca's and Wernicke's aphasias*

There have been lots of fascinating experiments that attempt to pin this down further—to develop theories about *exactly* what part of mental grammar is disrupted in Broca's aphasia. None of these theories has yet proven entirely satisfactory, though we have learned a lot about Broca's aphasia from them. It would take us too far afield here to explain the various theories and how they differ. However, given that Broca's aphasics have both phonological and syntactic problems—and that more specific deficits in either phonology or syntax are not so frequently attested—it may not make sense to seek a single grammatical problem in Broca's aphasia. Rather, some researchers are coming to the conclusion that Broca's aphasia is a general deficit in the ability to process fine details of linguistic structure. In addition, it may be that phonological and syntactic abilities are closely interspersed in Broca's area (in adjacent layers of cortical tissue, say), so that a lesion to one almost inevitably affects the other as well.

In any event, even for the intensively studied Broca's aphasia, it's still hard to draw a clear relation between the deficit we observe, the affected brain area, and the functional organization of grammar.

So much the worse for the other aphasias, which have received considerably less attention from neurolinguists.

Aphasia in ASL

In one of the most spectacular findings of the last ten years, it has turned out that precise analogues of Broca's and Wernicke's aphasias occur in brain-damaged speakers of ASL—and parallel brain areas are affected! An ASL Broca's aphasic signs slowly and leaves out all the grammatical inflections of location and style of movement; an ASL Wernicke's aphasic signs fluently but confusingly, and shows obvious comprehension problems. Not only that: sign language aphasics can produce and comprehend pantomime despite their language deficits, and they're fine at using their hands for purposes other than language.

And not only *that*. There exist *right*-hemisphere deficits that impair one's understanding of space, typically producing a so-called *left neglect*. People with these deficits fail to see things in the left half of their visual field, fail to draw the left-hand side of pictures, and fail to dress the left-hand side of their bodies. But if they happen to be ASL speakers, they still use the left side of the space in front of them *just for the purpose of signing ASL syntax*. Another kind of right-hemisphere damage can lead to loss of the ability to produce facial expressions. But despite such damage, ASL speakers can produce the aspects of facial expression that are relevant to ASL grammar, using the very same muscles.

In each of these cases, the differential character of the deficit shows that it is not a case of motor paralysis: the muscles can still be activated. What is damaged is the part of the brain that organizes the use of the muscles into coherent actions. And it turns out that the part of the brain organizing *ASL* action is not the same as the part that organizes ordinary action.

You can see why these results are so exciting. They confirm just about all the main points we made about sign language on the basis of the functionalist approach, but now adding evidence in terms of brain damage.

1. ASL is a language, not a collection of pantomimes and facial expressions. And it is localized in the language areas of the brain, in a different place from pantomimes and facial expressions.