



BOX 3.4

Then and now: Measuring brain activity through blood flow

In the late nineteenth century, the Italian physiologist Angelo Mosso observed a brain-injured patient and, based on his observations, made a connection between mental activity and blood flow in the brain. He later devised a “human balancing device” on which he tested his sense of this connection by conducting non-invasive studies of healthy individuals. The subject lay on a horizontal platform with the head on one side of a pivot and feet on the other, with the two sides perfectly balanced (see **Figure 3.8A**). Mosso assigned the subject tasks that called for various degrees of mental effort, in order to see whether this mental effort would cause the head to tip lower than the feet—a presumed consequence of increased blood flow to the brain. As Sandrone et al. (2013) describe:

Mosso nicknamed his device “the machine to weigh the soul.” He reported that the balance tipped towards the head when subjects were given more complex tasks; for instance, more head-tipping occurred while reading a page from a mathematics or philosophy text than when reading a novel. He also claimed to see effects of emotionally charged stimuli. For instance, he reported that the balance tipped toward the head immediately when one of his subjects read a letter from his spouse, and another read a note from an upset creditor. Media hype was just as present in the day of Mosso’s balance as with today’s fMRI studies, with a French newspaper reporting in 1908 that the device would “soon fully explain the physiology of the human brain” and lead to new treatments for neurological and mental illnesses.

Mosso’s method was primitive, but it’s worth remembering that it shares the same starting assumptions as our current, highly sophisticated brain-imaging techniques. Based on the assumption that active brain regions will display higher levels of blood flow and blood

oxygen than inactive regions, modern fMRI machines use magnetic field differences to detect and record brain activity (see **Figure 3.8B**).

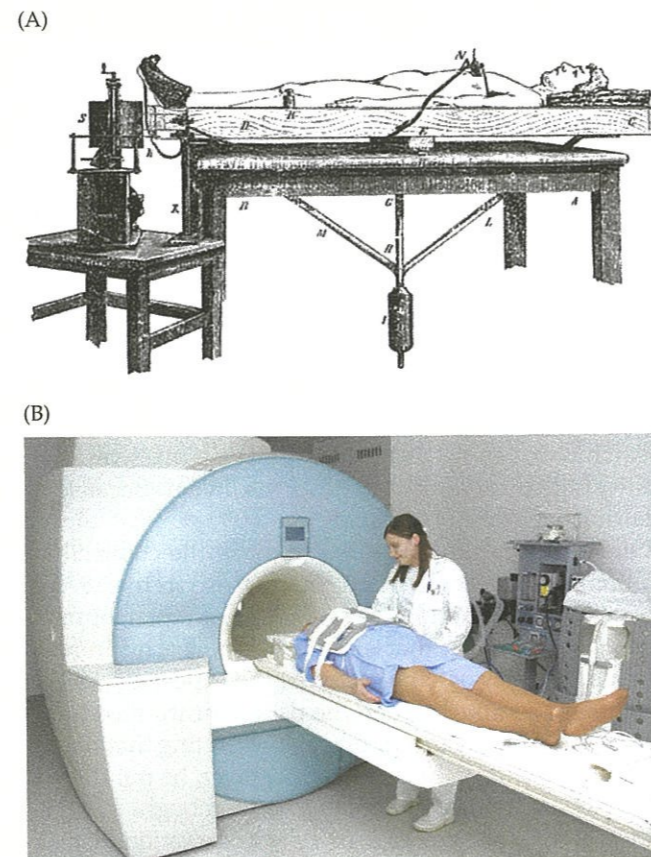


Figure 3.8 (A) Mosso’s balance for measuring blood flow. (B) A successor to Mosso’s balance, a modern fMRI brain scanner. (A reprinted from Sandrone et al., 2013; B © Shutterstock.)

So, the first assumption that neuroscientists make is that there’s a principled connection between hemodynamic measurements and brain activity. The second important assumption is that if changes in blood flow are consistently seen in certain areas of the brain shortly after the presentation of a certain stimulus, this is because the brain is recruiting those areas to process that type of stimulus. Relying on these two assumptions, how would we go about detecting the “language areas” of the brain in an fMRI experiment?

It’s not quite enough just to show someone in a scanner an image of a word or sentence, or have her hear a snippet of speech, and then see which brain regions show a change in blood flow. First of all, hemodynamic changes hap-

pen even in a brain that’s at rest (whatever *that* might mean), so these changes need to be factored out somehow (see **Method 3.2**). A more subtle and difficult point is this: How do we know that the active areas of the brain are engaged in processing the *linguistic* aspects of the stimulus? In reading a word, for example, there will be areas of the brain that are involved in very basic aspects of visual processing that have nothing to do with language—processes that would be just as active in, say, looking at an abstract painting, or recognizing a couch. Or, the word may trigger non-linguistic memories, associations, or thoughts,

METHOD 3.2

Comparing apples and oranges in fMRI


The pictures of activated brain regions that you see in published fMRI studies don’t represent a snapshot of the activity of any one brain for the task in question. They’re more sensibly read as graphs rather than photos, and they typically represent the *difference* between the experimental condition of interest and some chosen comparison condition, as averaged over many subjects. The dark areas in the picture don’t mean that those areas of the brain weren’t active while the task was being accomplished. They simply mean that those areas weren’t *more* active—to a statistically meaningful degree—than they were during the comparison condition. This means that it’s always worth thinking about what the comparison condition is, because the conclusions can only be stated in terms of this difference. A larger or smaller number of brain areas can show up as statistically different depending on the choice of the comparison condition. Let’s consider some of the issues that might come up with a language task and various comparison conditions we might opt for.

A common comparison condition is to instruct subjects to close their eyes and think about nothing in particular. Suppose we wanted to use this condition as a baseline for a task in which people listened to sensible conversations. What would people be likely to do in the “think about nothing in particular” baseline condition? If a good portion of the subjects actually lay there replaying the morning’s conversation with a girlfriend, or running a recent lecture through their minds in preparation for midterms, there would be a good chance that important language areas of the brain would be involved. The activity in these areas would then become subtracted from the actual language condition, which might give the impression that certain key regions are not activated for language, simply because they were actually activated in *both* the critical language condition and the baseline comparison condition.

Instead of a “resting” baseline condition, researchers sometimes use a control condition that focuses the subject’s attention on a specific task that is presumed to involve different computations than the condition of interest. For example, we might compare listening to words (linguistic input) with listening to single tones (non-linguistic input). The hope would be that the differences in activation (see **Figure 3.9**) would reflect the processing of spoken linguistic input as opposed to the processing of non-linguistic auditory input. But other unexpected differences might emerge. For example, it might be

Continued on next page

Noise



Speech sounds




Figure 3.9 These fMRI scans are composites from several subjects that, when combined, indicate areas of peak activation. Pure tones or “noise” (top scans) activate a relatively small region of auditory cortex. When speech sounds are heard (lower two scans), strong activity appears in many areas of the dorsal and ventral auditory pathways. Both the left (L) and right (R) cerebral hemispheres are shown. (From Binder et al., 1994.)

METHOD 3.2 (continued)

that the words are more complex examples of auditory stimuli and that they activate regions that are associated with processing complex sequences of sounds, whether linguistic or otherwise. Or, it may be that the words are more interesting and have an effect on brain areas that are linked to heightened attention. In this case, the results might show a difference in activation for a very large number of brain regions, even though only some of them are involved specifically in language.

To take one more example, let's say we compared listening to sentences with listening to musical melodies, based on the logic that both involve processing complex strings of auditory units but only one of them is linguistic. Let's also suppose that, unknown to us, processing melodies actually requires some of the same computations as unraveling the structure of spoken sentences, and involves some of the same areas of the brain. These regions would not show up in the results. The remaining areas that would be identified as being active in the language condition might well provide some answers to the question of which brain regions are devoted to language

and not music. But what if we operated under the false assumption that language and music are fundamentally distinct in their computations and use of brain resources (other than what's required for basic auditory processing)? We might wrongly conclude that our results answered the question of which brain regions are recruited for the purpose of linguistic processing. By assuming too great a distinction between linguistic and musical processing, we might have missed out on identifying some important brain areas that are common to both types of stimuli.

With any luck, over a large number of studies and using a variety of comparison conditions, we'd start to get a clearer picture of how to isolate language-relevant brain regions. But in reading the results of any single study, it's important to realize that it's cutting corners to say, "This study revealed activation in region X for task Y." Statements like this should really be understood as an abbreviation for, "This study revealed greater activation in region X for task Y as compared with task Z." And this understanding should lead us to spend at least a little time thinking about the relationship between tasks Y and Z.

activating the same areas of the brain that would be engaged in non-linguistic tasks like silently reminiscing or looking at a photograph. The task itself may incite boredom or arousal, mental states that have certain brain activation patterns. A reasonable strategy for isolating the language areas is to come up with a comparison condition that's as similar as possible to the target stimulus except that it doesn't require language. The brain regions that show activity over and above the control task can then more plausibly be attributed to the linguistic aspect of the stimulus.

Now that neurolinguists are equipped with an anatomical map in one hand and imaging techniques for brain function in the other, what have we learned about language in the brain? Keeping in mind that there are literally thousands of studies out there, the next sections provide very broad outlines of two key conclusions.

Language function is distributed throughout the brain in complex networks

Here's one way to think about the connection between brain regions and their function: we might conceive of important regions as dedicated processing centers, responsible for specific kinds of activities—for instance, visual processing, or language comprehension. A useful analogy might be to think of the regions as self-contained factories that take in raw material as input and produce certain products as output. Each factory has its own structural organization and sets of procedures that are independent from those in other factories, though some commonalities might crop up just because different factory operations settle on similar efficient solutions. This is an easy and intuitive way to think about brain localization, and it's probably made even more intuitive by the type

of language that's often used in media reports of neuroimaging studies, with references to notions like "the pleasure center" or headlines like "Scientists Locate Sarcasm in the Brain."

But even some of the earliest proponents of brain localization argued that this picture of the brain as a collection of independent processing centers was overly simplistic. For instance, Brodmann himself doubted that any of the brain regions he identified would turn out to be encapsulated dedicated processors. In his 1909 seminal work, he warned:

Mental faculties are notions used to designate extraordinarily involved complexes of elementary functions. . . . One cannot think of their taking place in any other way than through an infinitely complex and involved interaction and cooperation of numerous elementary activities. . . . Thus, we are dealing with a physiological process extending widely over the whole cortical surface and not a localized function within a specific region. We must therefore reject as a quite impossible psychological concept the idea that an intellectual faculty or a mental event or a spatial or temporal quality or any other complex, higher psychic function should be represented in a single circumscribed cortical zone, whether one calls this an "association centre" or "thought organ" or anything else.

In fact, if we turn to someone like Carl Wernicke, working early in the history of neuroscience, we see a similarly subtle view. Far from viewing Wernicke's area as something equivalent to the "language comprehension organ," Wernicke conceived of it as a critical piece in a larger network that linked information from different sensory modalities to information about the acoustic quality of words (see **Figure 3.10**).

Instead of thinking of the brain as an assortment of dedicated processing centers or independent factories, here's another possible scenario, one that is more in keeping with the speculations of Brodmann and Wernicke. Imagine the brain as a highly coordinated complex of commercial activity in which the makers of different products have arranged to share resources and their workers' expertise whenever possible. (For instance, the same factory space would handle the production of both fish sticks and chicken fingers, given that they rely on similar procedures. The packaging of many different kinds of goods might take place in another area, bringing together all kinds of frozen foods that go into boxes, including fish sticks, chicken fingers, miniature quiches, and hamburger patties.) In this industrial complex, the production of a specific

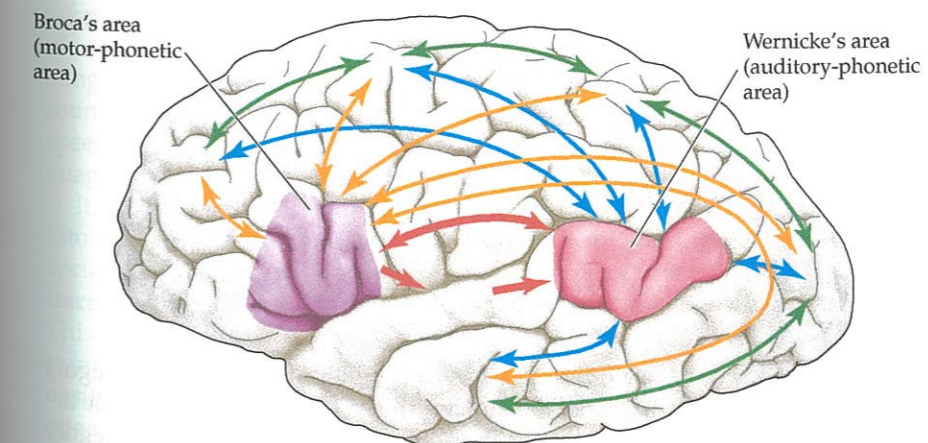


Figure 3.10 Wernicke's view of language involved a broadly distributed network. Thick red arrows connect the "motor-phonetic" or Broca's area and the "auditory-phonetic" or Wernicke's area. The blue arrows show connections between Wernicke's area and areas that store non-verbal information in "acoustic," "visual," "tactile," and "motor imagery" areas. The orange arrows represent connections between Broca's area and these various sensory areas. The green arrows show connections among the various sensory areas outside of the language network. (After Ross, 2010.)

product wouldn't take place within an isolated factory—instead, its trajectory from start to finish could be described as a path through the complex, making use of whichever resources were suitable for the production process. Some areas within the complex might be highly specialized, with a very small number of specific products (or perhaps even just one type of product) moving through them, while others would perform general tasks that apply to a great number of different products. One consequence of this kind of arrangement would be that products might have to travel large distances from one area of the complex to another, depending on what specific operations they needed to undergo.

From the very earliest work applying brain imaging to the study of language, results have lined up better with this second view of distributed brain function than with the first view of brain regions as dedicated processing centers. In the rest of this section, I'll touch on just a small subset of relevant examples.

In 1978, Bo Larsen and colleagues used a technique that was a precursor to PET and fMRI to identify the regions of the brain that were active while subjects listened to speech, as opposed to the regions that were active while "resting." Surprisingly, in the language-listening task, they found activity not just in Broca's and Wernicke's areas, but also throughout much of both the left and right hemispheres. They concluded that conversation was "likely to involve not only the cortical areas of importance for speech, but practically the whole brain, the left as well as the right side."

The fact that language-related functions are scattered throughout the brain is a testament to the great variety of separate tasks that need to be accomplished in the course of regular, daily language use. Many of the right-hemisphere functions seem to be quite different and complementary to those in the left, perhaps focusing on taking into account how something was said rather than decoding what was said. For example, the processing of information about intonation appears to be mainly housed in the right hemisphere (e.g., Ross & Monnot, 2008). The right hemisphere may also play an important role in how individual sentences are linked together into a coherent story or discourse (e.g., St. George et al., 1999).

The spatial distribution of language in the brain, though, isn't just due to the fact that a great variety of separate tasks are involved. Some of the diffusion also comes from the fact that language is entangled with non-linguistic knowledge. One of the most striking demonstrations of this is the pattern of brain activity that researchers see when they study the recognition of words. It doesn't seem unreasonable, as a first guess, to propose that word recognition might be associated with a certain pattern of brain activity—perhaps there's a location that corresponds to a "mental dictionary," or a general connection path between a "sounds" region of the brain and a "meaning" area. But in fact, you can get quite different patterns of activation for the following three categories of words:

(A)	(B)	(C)
kick	type	lick
step	throw	speak
walk	write	bite
tiptoe	grasp	smile
jump	poke	chew

Did you figure out what each category has in common? The words in category A refer to actions that involve the feet or legs; the words in category B name actions that require the use of fingers, hands, or arms; and the words in category

C describe actions accomplished via movements of the mouth and face. As demonstrated by Olaf Hauk and colleagues (2004), simply *reading* words from these categories activates some of the same brain regions involved in actually carrying out the movements, and reading words from different categories activates different brain regions (reading *kick* activates some of the brain regions involved in moving the feet, etc.; see **Figure 3.11**). Some of the more typical "language-y" areas are engaged as well, but, as Wernicke so astutely predicted at the dawn of modern neuroscience, fMRI data provide visible evidence that the language representations are connected with information in various other regions of the brain that are responsible for storing information about movement and the senses.

The functional neuroanatomy of language

Thinking about language function in terms of many distinct (but often overlapping) networks can help explain some otherwise mystifying data. For example, some patients with brain lesions do poorly on speech perception tests that require them to discriminate between two different syllables. You might predict that this would lead to great difficulty in recognizing words as well—but, while that's true for many patients, it's not necessarily the case. Some patients with poor speech perception skills are easily able to recognize the meanings of words, though they often have a great deal of trouble with language *production*. Conversely, there are other patients who have trouble recognizing words, but pass tests of basic speech perception with flying colors. It seems that it's possible to find cases of double dissociation between the processing of sequences of speech sounds and the recognition of words. What could possibly be going on, since (presumably) you can't easily figure out what a word is without having processed its individual sounds?

Greg Hickok and David Poeppel (2007) have argued that these puzzling findings start to make more sense if you think of the two tasks as belonging to different language-related networks. According to Hickok and Poeppel, word recognition recruits a network that maps speech input onto representations of meaning. Performing tasks like identifying individual syllables, on the other hand, leans more heavily on a different network that maps the acoustic information about sounds onto the articulatory gestures that produce them (this would be the kind of mapping that babies are learning during the babbling stage, when they spend countless hours uttering strings of meaningless sounds, as described in Chapter 2.) This would explain why trouble with simple speech perception tasks can be more directly connected to impairments in language *production* than to difficulties in understanding the meanings of words.

It might seem weird that knowledge of speech sounds would split apart into two separate networks like this. But other modalities show similar dissociations. It's now well known that visual recognition of physical objects fractures into knowledge of *what* objects are and of *how* they are to be used. This can lead to bizarre cases in which, for example, a brain-damaged patient is unable to visually recognize what a comb is or describe its purpose, but can easily demonstrate how to use it. It's more intuitive to think of our knowledge of objects (or sounds) as falling into one bin, but in fact, there's strong

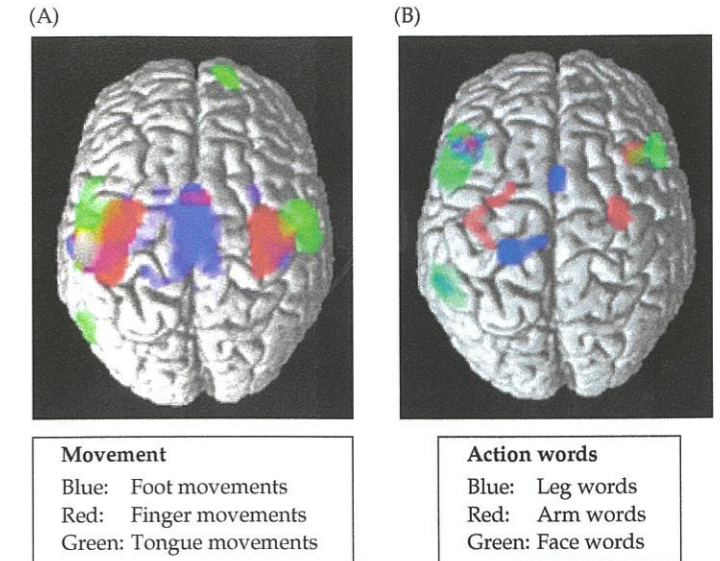


Figure 3.11 Results from a study of action words. (A) Activation of brain areas following instructions to move particular parts of the body. (B) Activation of brain areas during silent reading of action words involving three different parts of the body. In a comparison (baseline) condition, subjects saw meaningless rows of hatch marks, averaging the same length as the action words. (From Hauk et al., 2004.)




Figure 3.12 Analogous ventral and dorsal streams for (A) vision and (B) language in the left hemisphere of the brain. (B adapted from Gierhan 2013.)

evidence that separate knowledge streams exist for processing “what” and “how” information.

The separation of distinct “what” and “how” networks in the brain seems to be a basic way of organizing knowledge across a number of different domains, governing not just vision, but auditory perception and memory as well. Memory researchers, for instance, have long distinguished between declarative and procedural memory. **Declarative memory** refers to memory for facts and events (whether real or fictional) and includes bits of information such as the date on which World War I began, the names of Snow White’s seven dwarves, and the object of your first crush. **Procedural memory**, on the other hand, refers to memory for actions, such as how to thread a sewing machine or play your favorite guitar riff. If you’ve ever forgotten a familiar phone number, only to be able to dial it correctly when given a keypad, then you’ve directly experienced the disconnect that can happen between the two kinds of memory.

There’s now considerable evidence that language, too, is organized in two streams, and that these streams have clearly distinct locations in the brain. As with vision, processing the first type of information (the “what” knowledge) is organized into a network known as the **ventral stream**; the second type of information (the “how” knowledge) takes place in the **dorsal stream** (see **Figure 3.12** and **Box 3.5**). A good deal of research is being conducted with the aim of identifying exactly what kind of information is shuttled along each highway (a 2013 review by Sarah Gierhan provides an overview). The dorsal pathways seem to be involved in information that’s relevant for the detailed processing of sounds, for the planning of articulation, and for the repetition of words. The ventral pathways specialize in information about word meanings; damage to these connections, for example, can lead to trouble in understanding the meanings of words, or in retrieving words from memory. Both networks appear to be involved in the processing of syntactic information, though some researchers have suggested that each system is responsible for different kinds of syntactic information, with the processing of very complex structures taking place along the dorsal network.

Much of the emerging evidence supporting the existence of dorsal and ventral pathways is the result of new approaches and techniques that allow researchers to take the next step beyond simply identifying which regions of the brain are active during language tasks. They can now also investigate the ways in which the various language-related regions of the brain are connected to each other by

declarative memory Memory for facts and events (whether real or fictional) that can be spoken of (“declared”).

procedural memory Memory for physical actions and sequences of actions.

ventral stream Theoretical “knowledge stream” of ventral neural connections (i.e., located in the lower portion of the brain) that process knowledge about “what.”

dorsal stream Theoretical “knowledge stream” of dorsal neural connections (i.e., located in the upper portion of the brain) that process knowledge about “how.”

long bundles of neural fibers (*axons*; see Section 3.4) collectively called **white matter**. White matter tracts act as the brain’s road networks, allowing products from one processing area to be shuttled to another area for further processing or packaging. (Fun fact: The average 20-year-old human possesses between 150,000 and 175,000 kilometers of white matter fibers, as estimated by Lisbeth Marnar and her colleagues in 2003. That’s a lot of road.) White matter fiber tracts can be visualized in the living brain by using **diffusion magnetic resonance imaging (dMRI)**, which tracks how water molecules diffuse through the brain. Since water dif-

white matter Bundles of neural tissue (axons) that act as the brain’s information network, allowing products (signaling molecules) from one processing area to be shuttled to another area for further processing.

diffusion magnetic resonance imaging (dMRI) Neuroimaging technique that tracks how water molecules are diffused in the brain, providing a view of the brain’s “white matter highway.”



BOX 3.5
The functional neuroanatomy of language

The language areas of the cerebral cortex (the outer layer of neural tissue that covers the cerebral hemispheres) are diagrammed in **Figure 3.13**.

The **STG** (superior temporal gyrus) and the posterior portion of the **STS** (superior temporal sulcus) are involved in the phonological stages of spoken-word recognition—for example, in distinguishing between the important sounds in *bear* versus *pear*. This function seems to be bilaterally organized. That is, damage to only the left hemisphere does not result in great difficulties in processing the details of sound, but damage to both hemispheres (bilateral damage) results in “word deafness,” in which hearing is preserved but understanding of speech is badly impaired.

The anterior temporal lobe region labeled **ATL** is involved in accessing and integrating semantic knowledge across modalities, and within a syntactic structure. Damage to this area leads to difficulties in understanding complex or ambiguous sentences. Also in the anterior temporal

lobe, the **MTG** (middle temporal gyrus), **ITG** (inferior temporal gyrus), and anterior portions of the **STS** play a role in mapping sound to meaning and are also involved in accessing the meaning of written words. The representation of the meanings of words is widely distributed throughout the cerebral cortex (see Figure 3.11), but some researchers have argued that there is a more organized “hub” for word meanings in the anterior temporal region.

The left dorsal **STG** and **SMG** (supramarginal gyrus), along with the primary auditory cortex (**Aud**) and areas of the primary **motor cortex**, play a role in speech production, which involves integrating auditory information with a set of motor sequences for speech. Unlike speech perception, speech production seems to be heavily lateralized in the left hemisphere.

The **Spt** (Sylvian parietal temporal) region may play a role in sensory-motor integration for the vocal tract,

Continued on next page




Figure 3.13 This contemporary view of areas of the brain that contribute to language function, as organized into dorsal (green arrows) and ventral networks (red arrows; see Figure

3.12). Note that the networks extend into the right as well as the left cerebral hemisphere, although the left-hemisphere structures predominate. (Adapted from Hickok, 2009.)

BOX 3.5 (continued)

including “auditory imagery” of speech and non-speech vocal sounds (for example, humming music), whether the task involves producing sounds out loud or simply imagining them. This region shows heightened activity if auditory feedback from speech is disrupted (for instance, by delays). It is also likely involved in short-term verbal memory, which keeps sound-based information about words active in memory (for example, mentally “rehearsing” a phone number so you don’t forget it before you get a chance to dial it). This region also supports the learning of new, unfamiliar words.

Broca’s area (Brodmann areas 44 and 45) supports the production and understanding of syntactic structure.

In addition to the language areas of the cerebral cortex shown in Figure 3.13, language may also involve subcortical (internal) areas of the brain. For example, the **basal ganglia**, a collection of structures deep inside the brain (see Figure 3.14), have a key role in regulating bodily movement but also appear to be connected to the dorsal auditory stream. Some researchers argue that the basal

ganglia play an important role in the sequencing of sounds and syntactic units.




Figure 3.14 The basal ganglia, located deep within the forebrain, consist of several brain nuclei (clusters of anatomically discrete neurons, seen here in different shades of blue and lavender) and appear to have functions in the language pathway as well as their better-known functions in the motor pathway.

fuses in a direction that runs parallel to the white matter fiber bundles, dMRI provides a view of the brain’s “white matter road” (see Figure 3.15) and, crucially, some insight into how information moves between various regions of the brain—including the dorsal and ventral information processing “highways.”

Brain organization for language is both specialized and flexible

Broca’s area, which is implicated in language production, is conveniently located next to the part of the motor cortex that controls movement of the mouth and lips, while Wernicke’s area, which is important for comprehension, sits next door to the auditory cortex (see Figure 3.16). This makes sense, as there would likely be many connections between these adjacent areas. But not all language is spoken. Sign languages involve making movements with the hands rather than with the tongue and lips (though much of the face can be heavily involved); no hearing is necessary, with comprehension relying instead on visual-spatial processes. So here’s a trick question: Where would you find the




Figure 3.15 A view of the brain using dMRI, which tracks the movement of water molecules through the brain. Water diffuses in a manner that parallels the white matter tracts that carry neural signals. This imaging technique can provide insights into how information moves between various regions of the brain. (Courtesy of Patric Hagmann.)

language-related networks for people who grew up with a sign language as their native language? Would Broca’s and Wernicke’s areas be involved? Or would language establish its base camps in other regions? A logical place to look for this hypothetical “sign language network” might be near the part of the motor cortex that controls hand movements, or over in the right hemisphere, which takes on a good portion of visual-spatial processing.

In order to think about this question, let’s revisit our metaphor of the brain as a complex commercial network that makes many different kinds of products. Having an area like Wernicke’s next to the auditory cortex is a lot like setting up a fish stick factory near a fishing port—sensible, as the main ingredients don’t need to travel far in order to get to the processing plant. But what if, instead of making fish sticks, we decided to make chicken fingers? The ingredients are different, but it turns out that the machinery needed is very similar, as are the various steps in the production process. While it might make sense to build our chicken finger factory near a chicken farm, what if there’s already a facility in place near a fishing port that’s ideally set up for making products like fish sticks and chicken fingers? Even though it might require shipping the raw ingredients over a greater distance, it might still make more sense to use that facility than to build a whole new facility. So, one way to think about the question of localization of brain function is like this: does the brain’s organization reflect mostly the raw ingredients that it uses (spoken sounds versus hand movements), or does it specialize for the various processes (that is, the specific computations) that the raw ingredients have to undergo?

The answer is that, at least much of the time, the brain specializes for processing rather than for the ingredients. This can be seen from a number of studies of sign language users. For example, Greg Hickok and colleagues (2001) worked with a number of patients with aphasia who were American Sign Language (ASL) users and found that, just like hearing folks, there were deaf aphasic patients who had trouble producing signs but could comprehend them reasonably well, while others could produce signs but had trouble understanding them. The deaf patients had brain damage in exactly the areas usually found for aphasic hearing patients—in the areas known as Broca’s and Wernicke’s, respectively.

Evidence from imaging confirms that the brain organization of ASL signers looks a lot like that of speakers of sound-based languages despite the fact that a completely different modality is being used (for a review, see MacSweeney et al., 2008). This is interesting because in the last chapter, we saw that when gesture is used *linguistically* by homesigners and inventors of new sign languages, it has deeply different properties from pantomime gesture—a fact that had been lost on hearing observers for many years. The distinction between linguistic and non-linguistic gesture also shows up in brain-imaging studies, as found by Karen Emmorey and her colleagues (2011) when they compared brain activation patterns for ASL signs with those for pantomime gestures. To people who don’t know ASL, signs can sometimes *look* like pantomime because a number of signs have their origins in a pantomimed gesture that became conventionalized. For example, the ASL signs used to communicate the concepts of hammering or of pouring syrup are a lot like what you’d do if you were asked to

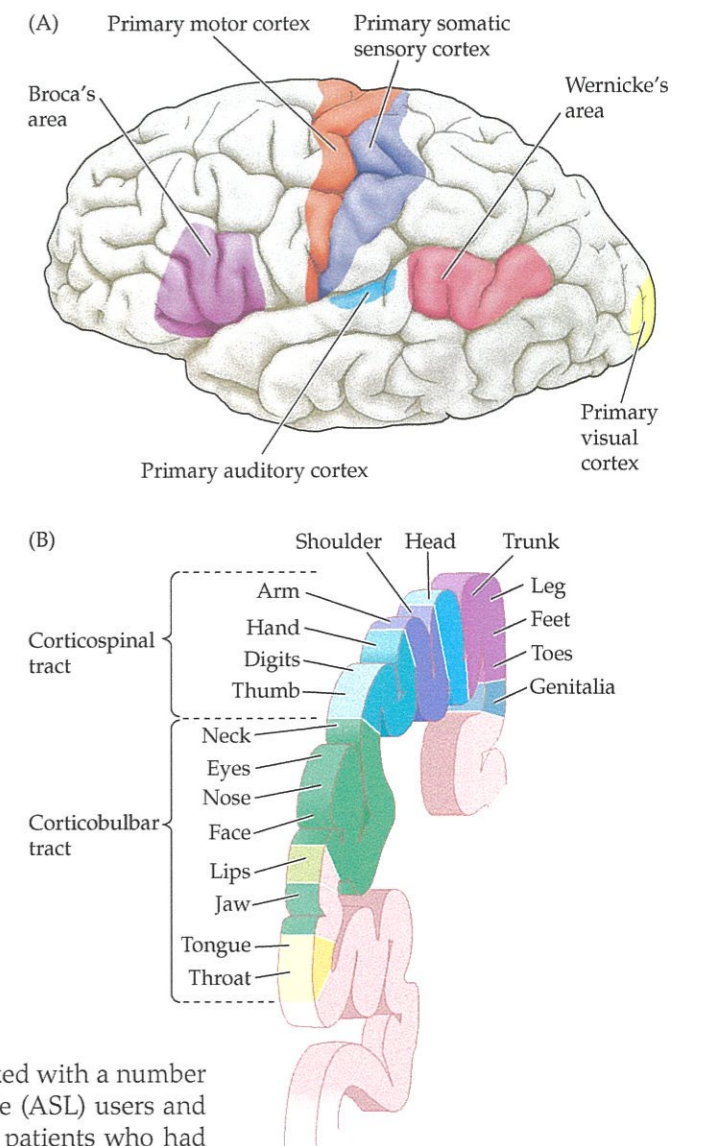


Figure 3.16 (A) This drawing illustrates the proximity of the motor cortex to Broca’s area, and of the auditory cortex to Wernicke’s area. (B) A schematic illustration of the organization of the primary motor cortex. The areas that control movements of the mouth and lips are located near Broca’s area, while the areas controlling movements of the hands, arms, and fingers are more distant.

LANGUAGE AT LARGE 3.2 (continued)

noted, for example, that when subjects who'd rated Hillary Clinton unfavorably saw photographs of her, they "exhibited significant activity in the anterior cingulate cortex, an emotional center of the brain that is aroused when a person feels compelled to act in two different ways but must choose one. It looked as if they were battling unacknowledged impulses to like Mrs. Clinton." Elicitation of disgust was attributed to viewing the candidate John Edwards, based on high levels of activation in the insula, and Mitt Romney was claimed to have elicited anxiety, based on an active amygdala among viewers.

But many neuroscientists responded to this study with a strong warning that the conclusions were premature. Sure, activity in the amygdala *could* represent anxiety upon seeing the face of Mitt Romney. But since the amygdala is known to become active under a number of different emotional states, not just anxiety, it could also mean that viewers were responding with happiness, anger, or sexual excitement—or perhaps a mental state that hasn't yet been discovered to be associated with the amygdala.

A hefty portion of potentially misleading or mistaken conclusions (including the above example) come from what's known as a "reverse inference," which has the following pattern of logic:

Previous studies have shown that during process X, brain area Y was active.

In the current study, we see stimulus A activates brain area Y.

Hence, processing stimulus A must involve process X.

Such inferences may seem intuitively appealing in many cases, but they are not valid conclusions. Other alternative explanations have not yet been properly tested and ruled out—and given that we now know that any one brain area can easily be implicated in a number of different processes, this is not a step that researchers can afford to skip. Certainly, to dispense "expert" (and expensive) advice about a candidate's political strategy or a company's marketing practices, based on results like these, risks becoming a form of pickpocketing through the scalp.

On the other hand, it's worth remembering that phrenology had some good, even revolutionary ideas. Had these ideas been subjected to proper scientific scrutiny, the accurate ones might eventually have been sifted from the less accurate ones, and we might remember the discipline as laying the groundwork for crucial breakthroughs in scientific thinking about how brain function is localized. In the end, a good part of distinguishing between science and pseudoscience amounts to being able to tell the difference between end points and starting points: does a particular result lend itself to a confident conclusion—or does it hint at an intriguing hypothesis to be tested further?

languages can differ in how they use tone, or changes in pitch. In languages like English, tone has a **paralinguistic use**—that is, it doesn't directly contribute to the composition of words or sentences, but it can be used for emphasis, to clarify the speaker's intended meaning, or for emotional color. But some languages (Mandarin and Vietnamese among them) use pitch as a third type of acoustic building block for words, along with consonants and vowels. In English, producing one consonant instead of another usually results in saying a completely different word—try, for instance: *pan*, *ban*, *can*, *man*, *tan*, and *Dan*. But varying the pitch on a word like *pan* doesn't turn it into a different word. In Mandarin, though, the same string of consonants and vowels—for instance, the simple sequence *ma*—can mean different things depending on the pitch contour you lay over it. Say it with a high tone, and you've uttered "mother"; with a rising tone, and you've said "hemp"; with a low falling tone, you've meant "horse"; and with a high falling tone, you've made a scolding sound.

For many English speakers, it's preposterously difficult to learn to attend to different distinctions in pitch as part of a word's core identity. Yet, there's nothing *inherently* difficult about making distinctions in pitch—you can easily tell if someone is saying your name in order to scold, query, exclaim with pleasure, or warn of impending danger, and yet all of this requires you to discriminate among pitch contours. So why is it so hard to use exactly the same information for a different purpose? Brain research provides a clue. When

paralinguistic use The use or manipulation of sounds for emphasis, clarification of meaning, or emotional color but not as an element in the composition of words or sentences.

tone is used linguistically to mark the identity of a word, it appears to recruit left-hemisphere regions in the brain. But non-linguistic pitch information is normally processed in the right hemisphere, and sure enough, speakers of languages like English, which uses tone paralinguistically, process tone primarily in the right hemisphere (e.g., Wang et al., 2001). Brain localization reflects the fact that tone has a different job description in the two language types. The process of learning a language like Mandarin isn't as simple as learning to perceptually distinguish tone; it involves some fairly deep reorganization of how information about tone is handled, and through which brain networks it's routed.

This statement about Mandarin likely generalizes to other languages that use sound in seemingly exotic ways. For example, a number of African languages use a set of sounds known as "clicks" that are made with a sucking action of the tongue, much like the sound you might make to spur on a horse. These languages use a variety of different clicks in the same way English uses consonants—that is, as basic building blocks of words. English speakers, it turns out, also use clicks, but as with tone, clicks are put to work in paralinguistic tasks; for English speakers, clicks seem to function a lot like paragraph breaks, to signal a shift in conversational topic. For instance, in studying the paralinguistic function of clicks in English, researcher Melissa Wright (2011) noted the following snippet of telephone dialogue:

Nor: You leave Wincanton about three o'clock and get back about two in the morning...

Les: Oh.

Nor: ...and work full-time on top of that.

Les: Oh dear.

Nor: But it's a lot easier now huh.

Les: Yes I'm sure. Hm. [click] Okay well I'll tell Gordon and uhm, I'm sure he was going to give you a ring anyway.

Despite being regular clickers themselves, many English speakers are fascinated by the linguistic use of clicks in other languages. To English ears, inserting these sounds inside words sounds a bit like highly skilled beatboxing (more on this in Chapter 4). I've had friends ask, "How do they talk and click *at the same time*?" To a speaker of a click language, this is a very odd question, much like asking, "How do they talk and make the consonant *b* sound *at the same time*?"

The take-home message from this line of work is that specific patches of neurons seem to be predisposed to handle certain kinds of tasks, but there is a fair bit of flexibility in how these core tasks can be adapted to handle a variety of different situations. This flexibility, together with the distributed nature of the brain's organization, no doubt makes research life a bit more challenging than it would be if the brain were set up as a cluster of independent "factories" or processing centers. Back in 1793, before scientists had any real evidence that different functions lived in different clumps of neural tissue, a medical doctor (Leny, 1793) wrote about the brain:

There is no part of the human body concerning which we possess so limited a knowledge. Its intricacy is great, and to that, our ignorance seems to bear proportion. We know neither the manner in which it



WEB ACTIVITY 3.4

Linguistic or paralinguistic? In this exercise, you'll see examples of how the same types of sounds have been recruited for different purposes by different languages.

performs its functions, nor the share which each of its parts have in their performances. The former is perhaps for wise purposes, placed above human understanding; and the latter, though it appears more within our reach, has hitherto eluded the research of inquirers.

We now know that it's not going to be possible to identify the "share which each of its parts have" in various functions of the brain without a much deeper understanding of "the manner in which it performs its functions." You can't just point to a brain region and say, "Vision occurs here" or, "Speech sounds are processed there." You have to have some ideas about how speech might be produced, in order to tease apart the various components that go into it; without this knowledge, you won't be able to figure out where in the brain the separate operations might be taking place, and how these operations might or might not generalize to non-speech stimuli. As you'll see throughout the rest of this book, brain localization research gets a lot more precise and interesting when it's tightly interwoven with ideas about the specific kinds of computations that take place there. In the best-case scenarios, evidence from brain imaging can help resolve some long-standing debates about the nature of cognitive processes. Though there's still an enormous amount we don't know, happily, researchers today need not feel that it's "above human understanding" to eventually arrive at a deep knowledge of how our own brains work.

3.4 The Brain in Real-Time Action

Brain-imaging techniques help scientists answer *where* in the brain language computations take place, and this question goes hand in glove with the question of what these computations look like. But another way to understand how language works is to ask *when* in the brain linguistic action takes place. For example, are some words retrieved sooner than others? Does the timing of language processes happen on a fixed schedule, or does it depend on the context? When we understand language, do we first sort out how the sentence is structured before we tackle its meaning? Answering specific questions like these would allow us to build a much more detailed model of what it means to use language.

Neuroimaging techniques like fMRI are not the ideal tools to use if you want to get a sense of the timing of brain processes, rather than their location. As you learned in the previous section, fMRI doesn't measure brain activity itself—instead, it taps into the changes in blood flow or blood oxygen levels that are the *result* of increased brain activity. These hemodynamic changes lag behind the brain activity that triggers them by several seconds, and as you'll soon see, all the interesting language-related activity takes place on a smaller time scale than this. There have been technical improvements in fMRI methods that allow researchers to mathematically factor out the time lag between brain activity and the fMRI signal, but for researchers who are mostly interested in very precise recordings of the timing of brain activity, it's possible to use methods that measure brain activity more directly.

Measuring electrical brain activity in response to language

Neurons communicate with each other through electrical signaling. In neurons, signaling occurs when electrically charged particles called **ions** move across a neuron's membrane. Each neuron is equipped with **dendrites**, which

ions Electrically charged particles; the charge can be positive or negative. Ions that are especially important in neural signaling include sodium (Na^+), potassium (K^+), calcium (Ca^{2+}), and chloride (Cl^-).

dendrites Neuronal extensions that receive informational "input" from other neurons.

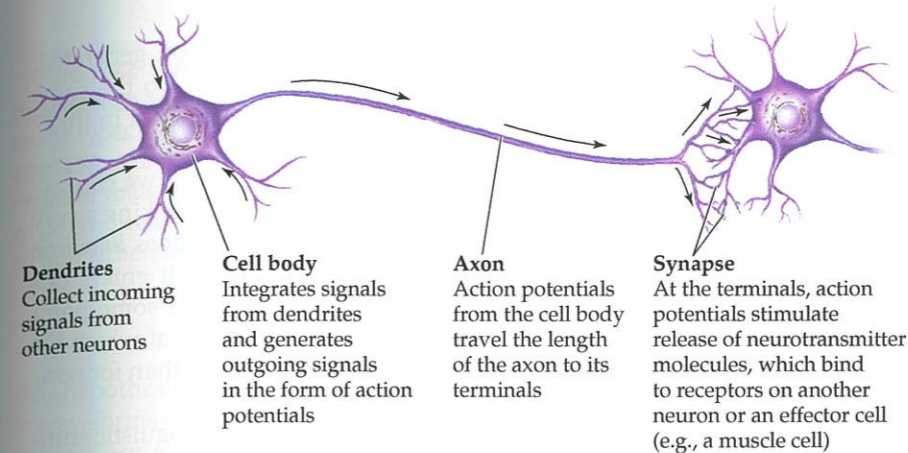


Figure 3.20 Electrical activity in a neuron. Dendrites collect electrical signals from other neurons. These signals are integrated in the cell body, and if the resulting voltage exceeds a threshold, an outgoing signal—an action potential—is sent along the axon, releasing neurotransmitters that have the capacity to alter the voltage of connected neurons.

are "input" sites that process the information from these signaling molecules. At the "output" end is the **axon**, which extends from the neuron's nucleus and ends in a number of **synapses**, where the axon connects with and passes information to the dendrites of other neurons (see **Figure 3.20**). At rest, neurons have a negative electrical voltage, which changes if they are stimulated. If a neuron's voltage rises above a certain threshold, it fires an electrical pulse—an **action potential**—that travels down the axon to the synapses, resulting in the release of chemical signaling molecules called **neurotransmitters**. These neurotransmitters in turn can allow ions to pass through the membranes of connected neurons, altering their electrical voltage.

The action potentials of individual cells can be measured by placing probes near the target cells. But this technique is too invasive to be used with human subjects, so scientists rely on **electroencephalography (EEG)**, using electrodes placed on the scalp to measure the changes in the electrical voltage over large numbers of neurons (see **Figure 3.21**). Electrodes used in this way are highly sensitive to the timing of voltage changes. But because they're picking up the brain's electrical activity through the skull, information about the precise locations of the voltage changes is blurred, providing only very approximate data about where in the brain this activity is taking place. A related technique, known as **magnetoencephalography**, or **MEG**, detects changes in magnetic fields that are caused by the brain's electrical activity. MEG provides better information about where this activity is taking place, but since the technique is much more expensive than EEG, there are many more research studies using EEG than MEG.

Using ERPs to learn the timing of brain processes

For studying language processes, researchers are interested in seeing how the brain's activity changes in response to a particular linguistic stimulus, so they usually look at EEG waveforms that are lined up to the onset of that stimulus. This way of looking at brain activity is known as an **event-related potential (ERP)**—the "event" in question being the presentation of the relevant stimu-

axon Extension of a nerve cell (neuron) along which informational "output" travels to another neuron.

synapse Site of connection between the axon terminal of a neuron and the receptors of another neuron or a muscle cell.

action potential An electrical pulse that travels down the axon of a neuron to a synapse, resulting in the release of neurotransmitters.

neurotransmitter Molecules produced by a neuron and released across a synapse in response to an action potential. Neurotransmitters bind to receptors on a receiving cell (another neuron or a muscle cell), producing a response in the second cell.

electroencephalography (EEG) The use of electrodes placed on the scalp to measure changes in electrical voltage over large numbers of neurons in the brain, thus obtaining information about the timing of responses in the brain.

magnetoencephalography (MEG) A technique related to electroencephalography that detects changes in magnetic fields caused by the brain's electrical activity.

event-related potential (ERP) The change in electrical voltage (the potential) over large numbers of brain neurons, measured with EEG and lined up with the presentation of a relevant stimulus (the event).




Figure 3.21 A research participant with EEG electrodes placed over the scalp.