# Language in Mind

## An Introduction to Psycholinguistics

**JULIE SEDIVY** *University of Calgary*

**SINAUER ASSOCIATES, INC.** • Publishers • Sunderland, Massachusetts • USA

**Figure 3.1** Rembrandt's *The Anatomy Lesson of Dr. Nicolaes Tulp*, painted in 1632, depicts an anatomy demonstration of that time. Such lessons were open to the public for a fee.

100 billion of them). Some clues can be gleaned from the more obvious physical structure of the brain. For example, the left and right hemispheres are largely physically separate from each other, so it seems reasonable to ask whether the two sides *do* different things. More subtle clues can be discerned by looking at brain tissue under a microscope; if two different regions of the brain have a different cellular makeup, this suggests that they may take on somewhat different tasks. But even today, the connection between the structure of brain tissue and the functions those tissues serve is far from clear. The brain poses significant challenges simply because it's a physical object whose function is not easily understood from its anatomical form. This fact is a big reason why an understanding of the brain has historically lagged far behind our understanding of the other organs in the human body.

A second and even greater difficulty comes from the sheer number, variety, and complexity of the brain's functions. In order to understand how a physical object "works," you need to have a clear idea of what it does. Sometimes this is trivially easy; for example, the function of a (non-digital) clock is to move the clock's hands around in a way that is consistently linked to units of time. When you look at the wheels and gears inside of a clock, it's with the aim of understanding how it accomplishes this specific function. A car is a bit more complicated. Sure, its ultimate purpose is to "drive," but peering under the hood is going to be a lot more informative if you've first been able to break down that larger purpose into component tasks. It helps to start with the idea that a number of different sub-tasks are involved, with the expectation that these map onto different mechanical "systems." For example, in order to "drive," your car's engine has to be able to start, the wheels have to turn in specific directions, the vehicle has to move forward *and* backward while transforming fuel into energy, the speed has to be modulated, and the car needs to be able to be brought to a stop—to name just a few sub-tasks. The different systems that accomplish

**neurolinguists** Scientists who study how the physical brain relates to language behavior.

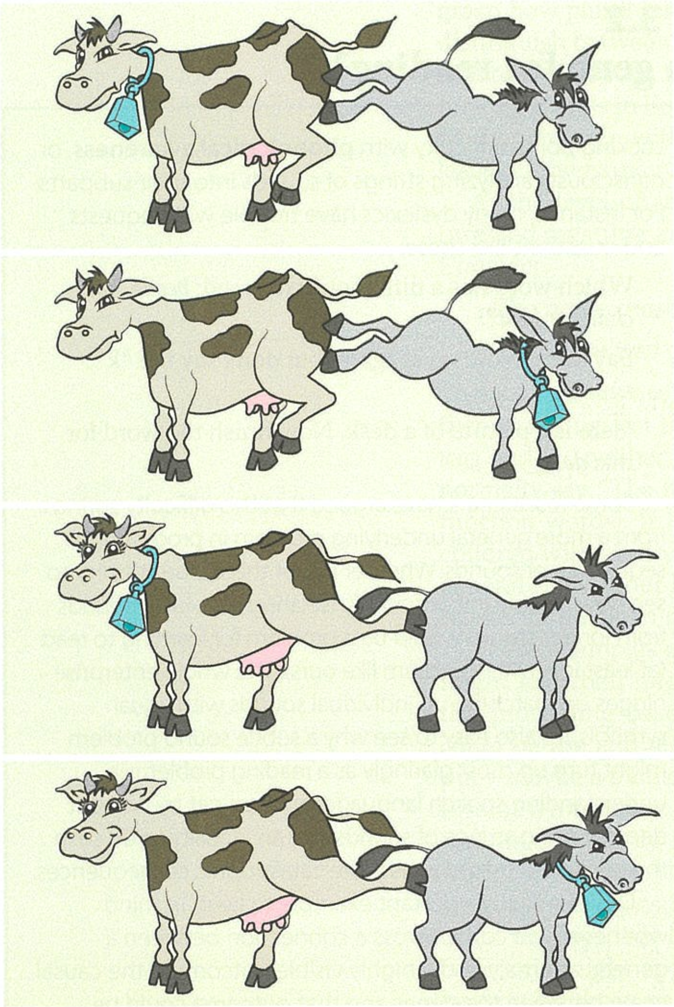

**Figure 3.2** Sample stimuli from a sentence picture verification task. Children are asked to point to the picture that matches the sentence "The donkey that kicked the cow has a bell."

### Testing the right thing: Method is important

There's an important methodological issue to take into consideration when trying to figure out how linguistic skills relate to non-linguistic ones: in order to measure language function, we have to rely on some appropriate test. But the test itself may depend on cognitive skills over and above the specific linguistic skills that are being targeted. For example, a common way to test how well children understand syntactic structure is to present them with a series of test trials involving complex sentences that differ in subtle ways, such as "The donkey that kicked the cow has a bell" versus "The donkey kicked the cow that has a bell." Children are shown several pictures and are asked to choose which picture best goes with the sentence they just heard (see **Figure 3.2**). In order to perform reliably on this test, children need to have intact syntactic skills. But they also need to have several other things: the perceptual skills to make fine distinctions among similar images; the ability to relate visual images to representations of similar events; the memory capacity to keep track of which pictures differ how; the memory capacity to remember exactly what sentence the experimenter uttered; the motivation to repeatedly pay attention to a series of test trials; and so on. This test—intended to probe for syntactic understanding—is hardly purely linguistic. So, let's suppose we find that children who have especially short memory spans do worse on this test than those with roomier memory spans. Does this mean that the children's difficulty with syntax can be explained as originating in problems with working memory? Not necessarily—it may just be that *this particular test* relies heavily on working memory, creating a false connection between memory and syntactic performance. Ideally, we'd want to check to see if the relationship holds across a number of different tests probing for syntactic understanding and memory, using tests that vary in the ways in which they tax non-linguistic cognitive functions.

### We need more knowledge about how language works

The title of this section is "What can genetic disorders tell us about brain systems?" Perhaps it's time to take a stab at an answer, based on the research survey so far. The fact that there's a variety of different genetic disorders, with strikingly different effects on both language and general cognition, shows that there is some degree of specialization in the brain, and that genes can affect how these specialized skills develop. At the same time, evidence from language disorders doesn't offer us an easy picture, with a clear division between language and the rest of the brain. Instead, it looks as if we'll need to look

WEB ACTIVITY 3.1

**Cognitive demands in language tests** In this activity, you'll explore several tests that have been used to test language functioning in SLI. You'll consider what other cognitive skills might be necessary to succeed at the task, in addition to the targeted linguistic skill.

25-year-old railroad worker. In 1848 Gage was the unfortunate victim of an accidental explosion that drove an iron rod into his left cheek and out the top of his head, landing about 25 meters away. Incredibly, Gage not only survived, but moments after the accident, sat up and chatted, and a short while later was able to relate the details of the accident to a doctor. He survived for more than 12 years with most of his capacities seemingly intact—his language and motor function, for example, appeared to be fine. The doctor who cared for him noted that Gage's survival was surely due in part to the fact that "the portion of the brain traversed, was, for several reasons, the best fitted of any to sustain the injury." But he also noted that the accident had caused some deep changes; evidently Gage's personality took a turn for the worse, and he was never able to function as well as he had before the accident (see **Box 3.3**).

## BOX 3.3

## Phineas Gage and his brain

Dr. John Martyn Harlow was practicing in Cavendish, Vermont, near where Gage's accident occurred in 1848. He treated Gage at the time and followed his patient's progress until Gage's death in 1860. Harlow then prepared a detailed summary of the case (he even obtained and studied Gage's skull), which was published in 1868 and describes Gage's altered personality:

> *The equilibrium or balance, so to speak, between his intellectual faculties and animal propensities, seems to have been destroyed. He is fitful, irreverent, indulging at times in the grossest profanity (which was not previously his custom), manifesting but little deference for his fellows, impatient of restraint or advice when it conflicts with his desires, at times pertinaciously obstinate, yet capricious and vacillating, devising many plans of future operations, which are no sooner arranged than they are abandoned in turn for others appearing more feasible. A child in his intellectual capacity and manifestations, he has the animal passions of a strong man. Previous to his injury, although untrained in the schools, he possessed a well-balanced mind, and was looked upon by those who knew him as a shrewd, smart businessman, very energetic and persistent in executing all his plans of operation. In this regard his mind was radically changed, so decidedly that his friends and acquaintances said he was "no longer Gage."*

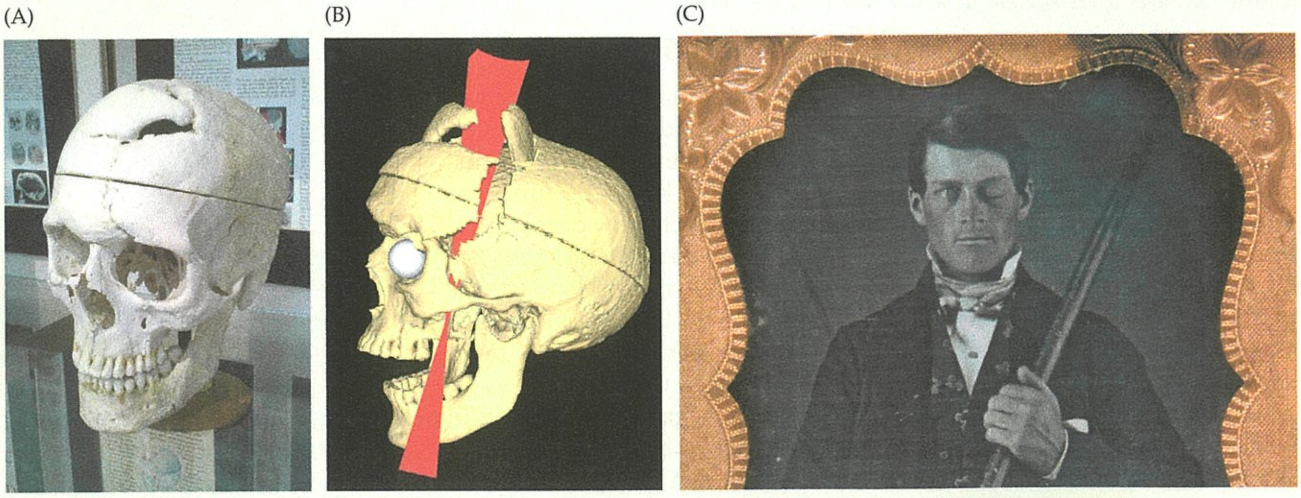

**Figure 3.3** (A) Phineas Gage's skull is on display at the Warren Anatomical Museum at Harvard Medical School. (B) Reconstruction of the pathway of the iron rod through Gage's skull. (C) A recently discovered photograph of Gage (holding the iron rod), taken some time after his accident. (A,B from Van Horn et al., 2012; C from The Jack and Beverly Wilgus Collection.)

**cerebral cortex** The outer covering of the brain's cerebral hemispheres.

**aphasia** Any language disruption caused by brain damage.

**Broca's aphasia** Aphasia characterized by halting speech and tremendous difficulty in choosing words, but fairly good speech comprehension. Also called motor aphasia or expressive aphasia.

**Wernicke's aphasia** Aphasia associated with fluent speech that is well articulated but often nonsensical, and enormous difficulty in understanding language. Also called sensory or receptive aphasia.

a particular swear word, the syllable *tan* was the only set of speech sounds he'd managed to eke out for 21 years. The patient died a few days after their meeting, and as Broca was aware that scientists were beginning to explore claims about the localization of language, he decided to autopsy Leborgne's brain. He considered language to be a good test case for the more general hypothesis that the various functions of the brain were compartmentalized into different physical regions. He discovered extensive damage to the frontal lobe on the left side of Leborgne's brain, providing some of the earliest hard evidence of localization in the brain (Broca, 1861).

Based on his observations, Broca argued that the faculty of language was further split apart into subfunctions, an idea that was consistent with many earlier reports of language loss due to brain damage. He noticed that Leborgne seemed to understand language much better than you'd expect from his utter lack of ability to speak—for example, when asked how long he'd been hospitalized, he flashed four sets of five fingers and then a single finger, to indicate 21. To Broca, this suggested that he'd lost the ability to produce spoken language (despite maintaining reasonable dexterity of his tongue and mouth) but that other aspects of language functioning were better preserved. Following this famous case, Broca autopsied the brains of a number of patients whose language was impaired after stroke or other brain damage, and he found that a significant portion of them had damage to the same part of the **cerebral cortex** (the brain's outer layer of neurons), specifically on the left side of the frontal lobe.

Shortly after Broca's discovery, neurologist Carl Wernicke studied a patient who had suffered a stroke and, though able to speak fluently, didn't seem to understand anything that was said to him. A later autopsy revealed a lesion, or evidence of brain damage, on the left side of the cerebral cortex—but the lesion was farther back than the region Broca had described, in the temporal lobe rather than the frontal lobe (see **Figure 3.4**).

In 1874, Wernicke published an influential text in which he explored his ideas about **aphasia**, the clinical term for language disruption caused by brain damage. Even though scientists and clinicians had long suspected that language loss came in at least two distinct varieties, the pioneering work of Broca and Wernicke established that the distinct forms of aphasia were related to different areas of the brain. **Broca's aphasia** (also called motor or expressive aphasia) is characterized by halting speech, if any at all, and tremendous difficulty in choosing words, but fairly good comprehension. **Wernicke's aphasia** (also called sensory or receptive aphasia) is associated with fluent speech that is well articulated but often nonsensical, and enormous difficulty in understanding language. (See **Table 3.2** for examples of speech by patients with Broca's and Wernicke's aphasias.)

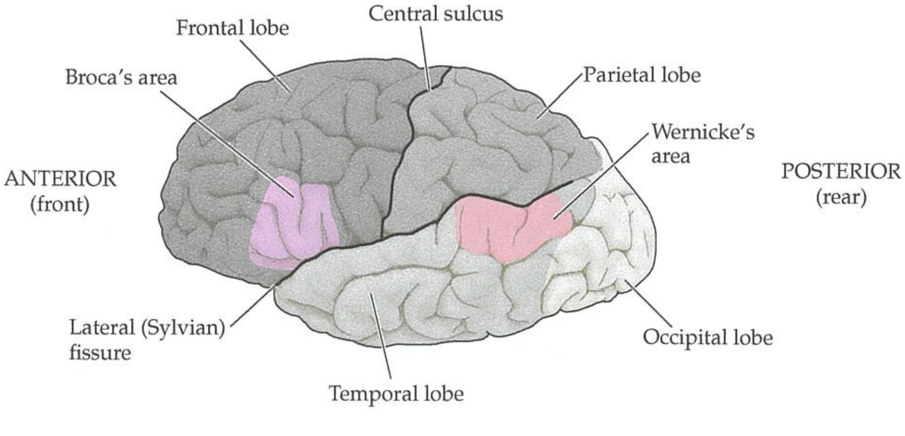

**Figure 3.4** A side view of the surface of the brain's left hemisphere. The four lobes of the cerebral cortex are indicated in shades of gray, with Broca's area and Wernicke's area shown in color.

the brain). In fact, even the dramatic language impairment of Broca's famous patient named Leborgne may have resulted from more extensive damage than Broca originally thought: since the good doctor had the incredible foresight to preserve Leborgne's brain for future scientists, researchers were recently able to image the brain using modern magnetic resonance imaging (MRI) techniques. They found evidence of deep damage to the brain not just in the frontal lobe on the left side, but also in subcortical areas and throughout the superior longitudinal fasciculus, a bundle of neurons that connects the front and back areas of the cerebral cortex (see **Figure 3.5**).

It's apparent that the divide between comprehension and production is not a tidy one. On closer inspection, most patients with Broca's aphasia have trouble with some aspects of comprehension as well as devastating difficulties with language production. Especially irksome for these patients are sentences that rely on subtle or complex syntactic structure without any helpful clues about meaning. (For example, a Broca's patient might readily figure out the meaning of *The mouse was chased by the cat* but not *The boy was chased by the girl*. For the first example, but not the second, the meaning of the sentence can be plausibly assembled if all you can figure out are the word meanings.) Symptoms like these have prompted researchers to offer various proposals about additional duties of Broca's area. Some have argued that certain kinds of syntactic structures are computed in this region; others have suggested that it's an important site for working memory processes, or for mechanisms that resolve the tension between conflicting linguistic cues. This rethinking of the nature of aphasia is driven in part by more detailed techniques for studying the brain. But it also comes from much more detailed theories about all of the mental operations that are involved in producing and understanding language. And as these theories become richer and more complex, so do ideas about how language function maps onto areas of the brain.

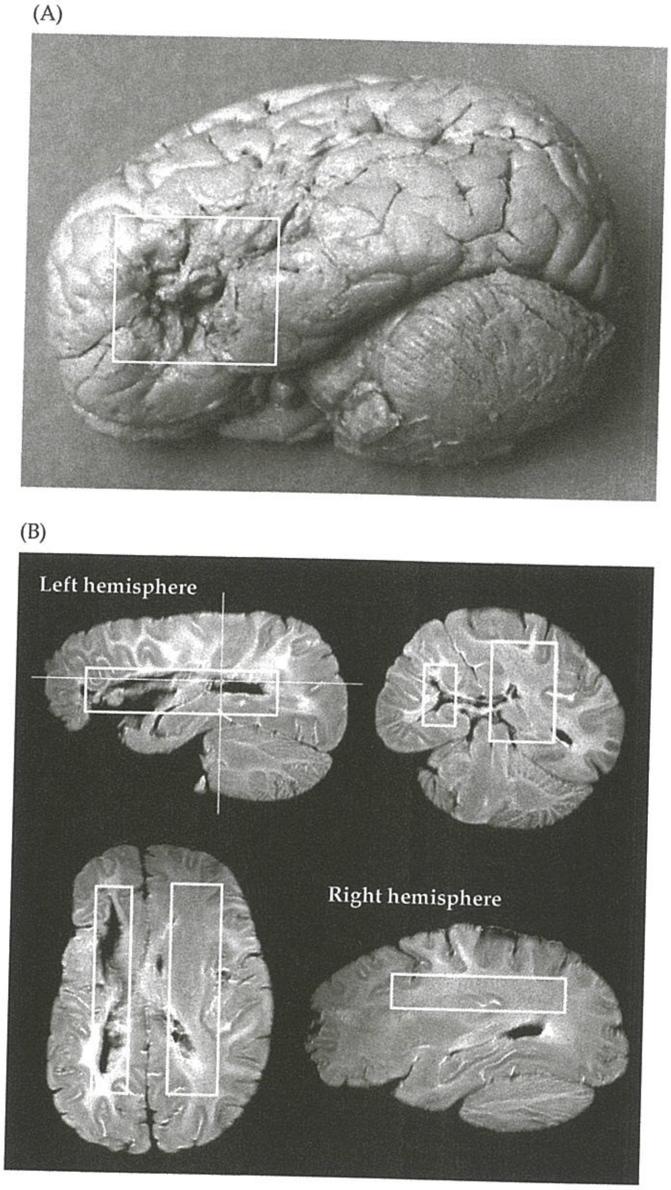

**Figure 3.5** (A) Photograph of Leborgne's brain, with damage clearly visible in the inferior frontal lobe of the left hemisphere. (B) MRI images showing extensive damage throughout the left hemisphere. Boxes are drawn around comparable areas in the left and right hemispheres. (From Dronkers et al., 2007.)

## *Brain lateralization*

The fact that Broca's and Wernicke's areas were both traced to the left side of the brain led to the first inkling that the brain might be organized differently in its two hemispheres—a possibility that initially came as a great surprise to Paul Broca. But since Broca's time, additional evidence of **brain lateralization** (that is, evidence that the right and left cerebral hemispheres are specialized for different functions) has come from many corners, and has involved somewhat exotic brain conditions as well as clever studies of people with uninjured brains.

The best-known studies of brain lateralization were done by Roger Sperry and Michael Gazzaniga in the 1960s, about a hundred years after the pioneering work of Broca and Wernicke. The studies involved a number of "split-brain" patients who had undergone a radical, last-resort treatment to prevent the spread of epileptic seizures from one side of the brain to the oth-

**brain lateralization** The specialization of the brain's right and left cerebral hemispheres for different functions.

**corpus callosum** A bundle of neural fibers that connects and transfers information between the two hemispheres of the brain.

er. These patients submitted to a surgery that severed the **corpus callosum**, the bundle of neural fibers that connects the two hemispheres of the cerebral cortex in a high-speed "superhighway." The surgery was approved as a treatment after studies by Roger Sperry showed that the procedure in monkeys resulted in very little change in the monkeys' behavior—and indeed, human split-brain patients were able to function surprisingly well even though their two hemispheres had lost the ability to share information with each other.

But using clever experimental tests, the researchers were able to demonstrate some bizarre consequences of the disconnection. The experiments required finding some way to present information to only one side of the brain. For example, to present information to the left hemisphere, sensory input needs to come from the right side of the body because the brain is wired in such a way that it receives input from, and sends motor commands to, the opposite side of the body. "Split-brain" patients used their right hands to handle objects that were hidden behind a barrier, so that only the left hemisphere had access to information gleaned from touching the objects (see **Figure 3.6**). In other versions of the experiments, patients sat in front of a screen and were told to look

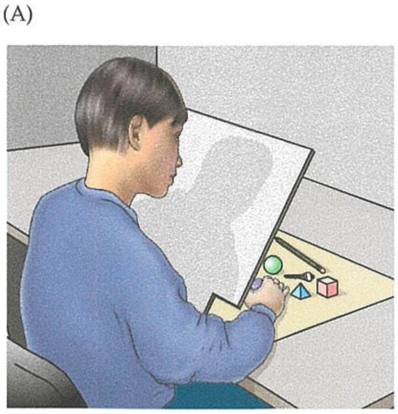

| Left hemisphere functions | Right hemisphere functions |
|---|---|
| Analysis of right visual field | Analysis of left visual field |
| Stereognosis (right hand) | Stereognosis (left hand) |
| Lexical and syntactic language | Emotional coloring of language |
| Writing | Spatial abilities |
| Speech | Rudimentary speech |

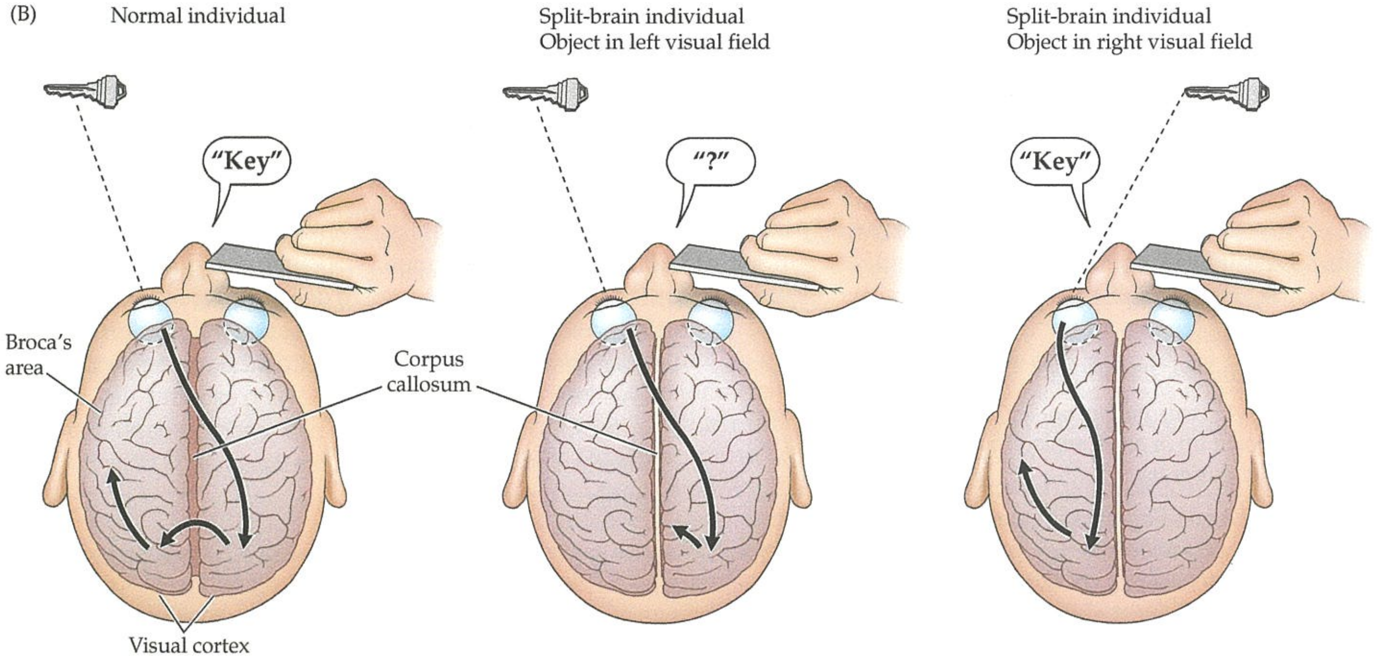

**Figure 3.6** (A) A split-brain patient handles an object behind the screen with his right hand. (B) Presenting visual information in just the left or right visual field has different effects on individuals with normal versus split brains. When the corpus callosum is intact, information presented in the left visual field is processed in the right hemisphere but can be relayed to crucial language areas in the left hemisphere. In a split-brain individual, only information presented in the right visual field is able to reach the language areas in the left hemisphere.

reorganize itself, and that even within a few weeks of a stroke, there's evidence that brain function has been rerouted in significant ways. If a function that was previously accomplished by a now-damaged area becomes taken over by a healthy part of the brain, it makes it hard to know what the original organization of brain function was like. There are other more practical challenges that come with relying on individuals with brain damage as the primary research participants. There's a relatively small number of them, which limits how much researchers can generalize to the broader population. It also constrains the amount of research that can be carried out; many individuals with brain damage are extraordinarily generous with their time in helping researchers make progress in the field, but there's a limit to how many hours any one person can spend in a lab performing tests—those who are recovering from a stroke, in particular, may tire easily, or they may show inconsistent performance partly because of their brain injury. Being able to test hypotheses within the general population was necessary in order for the field to make rapid progress and gain greater confidence in its findings.

### *Localizing language: Brain mapping techniques*

Although the possibility of large-scale testing of brain function in healthy humans had to wait until the advent of modern imaging techniques, some groundbreaking contributions to the science of brain localization were made more than a century ago. Among the most influential was the brain-mapping work of German neurologist Korbinian Brodmann, published in 1909. Brodmann believed that the study of brain function had to be grounded in a solid understanding of how the brain was built, so he set about meticulously analyzing the cellular composition of countless slices of brain tissue from animals and human cadavers. Based on his work, he created a "map" of areas in the human cerebral cortex that were anatomically distinct from each other (see **Figure 3.7**). His reasoning was that areas that differed in their physical structure were likely to be responsible for different functions. These **Brodmann areas** have guided much of the exploration of brain function, and are still commonly referred to in current cognitive neuroscience.

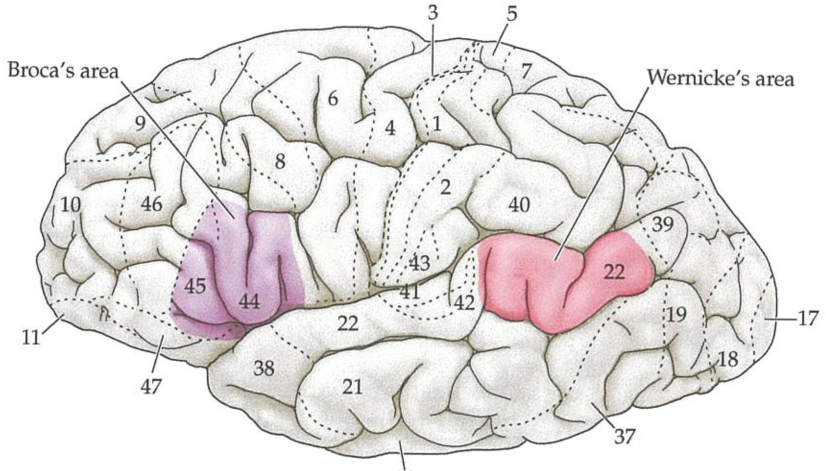

**Figure 3.7** The Brodmann areas of the brain mark distinctions in cell composition in the various layers of tissue in these regions. Broca's area corresponds approximately to Brodmann areas (BA) 44 and 45, while Wernicke's area corresponds to BA 22.

**Brodmann areas** Areas of the human cerebral cortex that are distinct from each other anatomically and in cellular composition, as determined by Korbinian Brodmann.

## 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**).

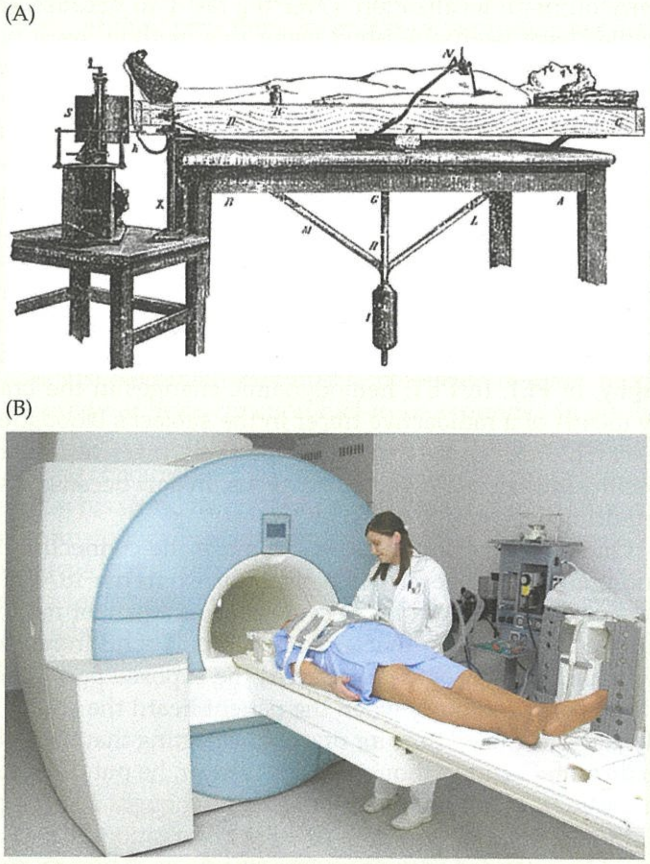

**Figure 3.8** (A) Mosso's balance for measuring blood flow. (B) A successor to Mosso's balance, a modern fMRI brain scanner. (A reprinted from Sandrone et al., 2013; B © Shutterstock.)

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 happen 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

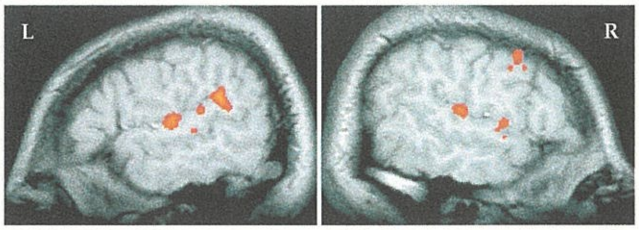

Speech sounds

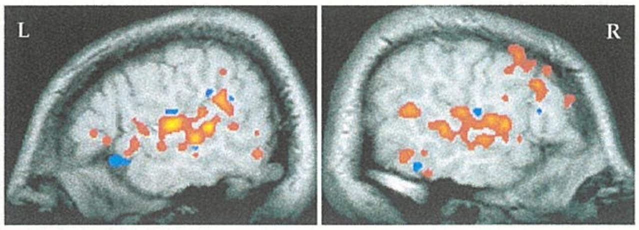

**Figure 3.9** These fMRI scans are composites from several subjects that, when combined, indicate areas of peak activation. Pure tones or "noise" (top scans) activate a relatively small region of auditory cortex. When speech sounds are heard (lower two scans), strong activity appears in many areas of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

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

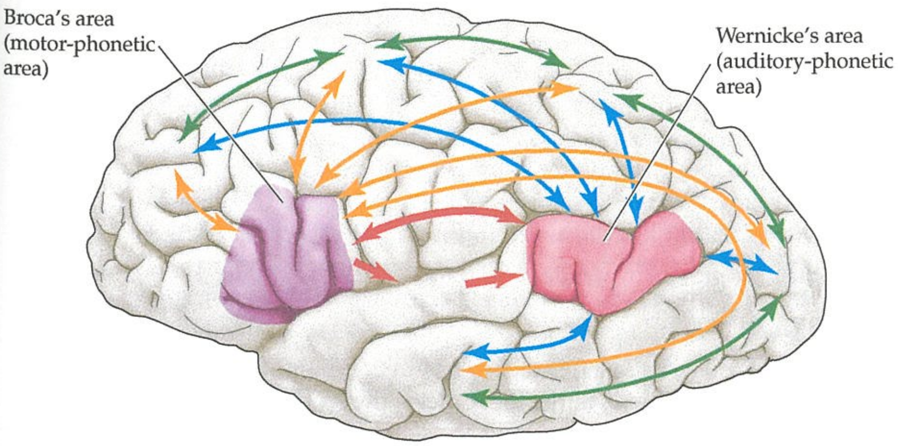

**Figure 3.10** Wernicke's view of language involved a broadly distributed network. Thick red arrows connect the "motor-phonetic" or Broca's area and the "auditory-phonetic" or Wernicke's area. The blue arrows show connections between Wernicke's area and areas that store non-verbal information in "acoustic," "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.)

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.

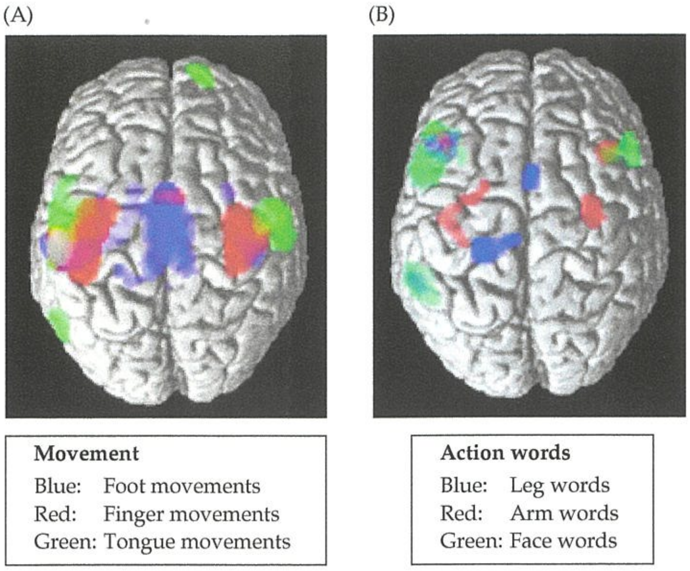

**Figure 3.11** Results from a study of action words. (A) Activation of brain areas following instructions to move particular parts of the body. (B) Activation of brain areas during silent reading of action words involving three different parts of the body. In a comparison (baseline) condition, subjects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

### *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

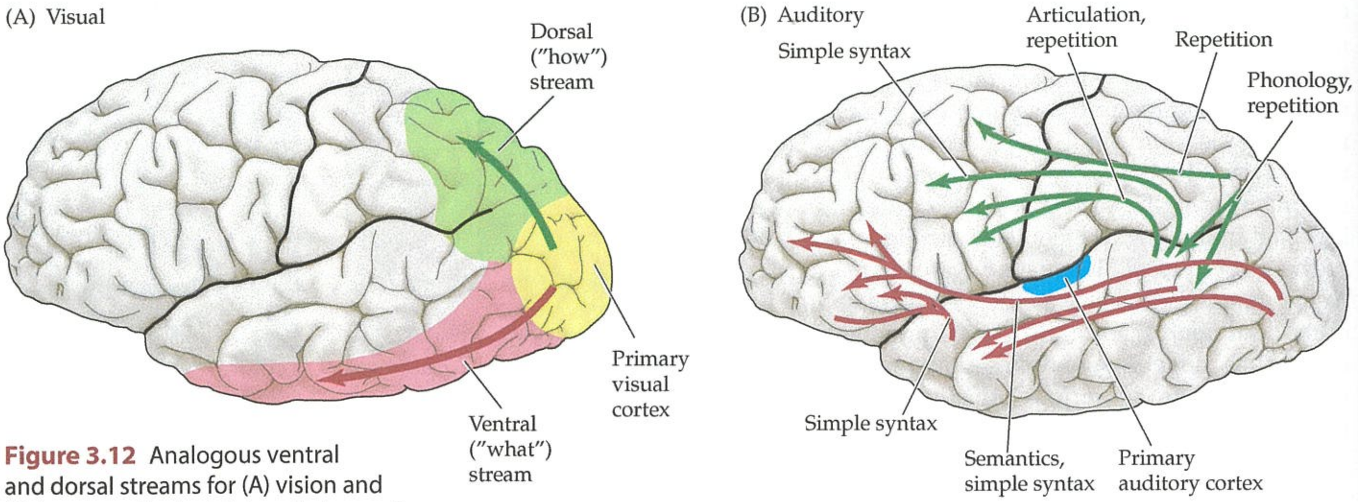

**Figure 3.12** Analogous ventral and dorsal streams for (A) vision and (B) language in the left hemisphere of the brain. (B adapted from Gierhan 2013.)

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 Marner 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*

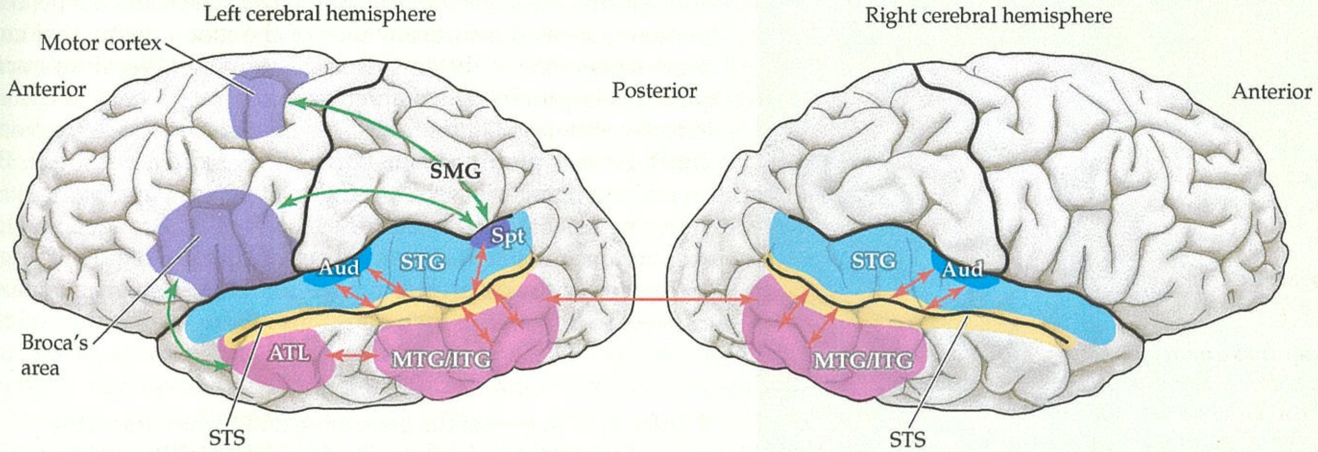

**Figure 3.13** This contemporary view of areas of the brain that contribute to language function, as organized into dorsal (green arrows) and ventral networks (red arrows; see Figure 3.12). Note that the networks extend into the right as well as the left cerebral hemisphere, although the left-hemisphere 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.

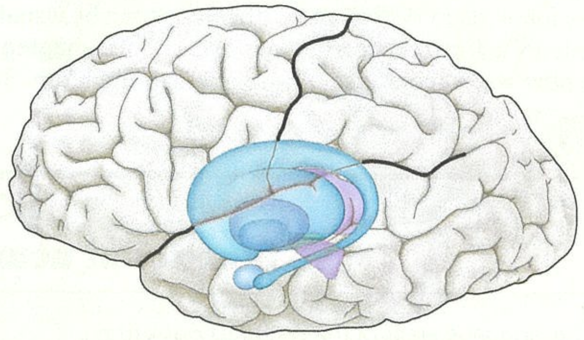

Figure 3.14 The basal ganglia, located deep within the forebrain, consist of several brain nuclei (clusters of anatomically discrete neurons, seen here in different shades of blue and lavender) and appear to have functions in the language pathway as well as their better-known functions in the motor pathway.

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."

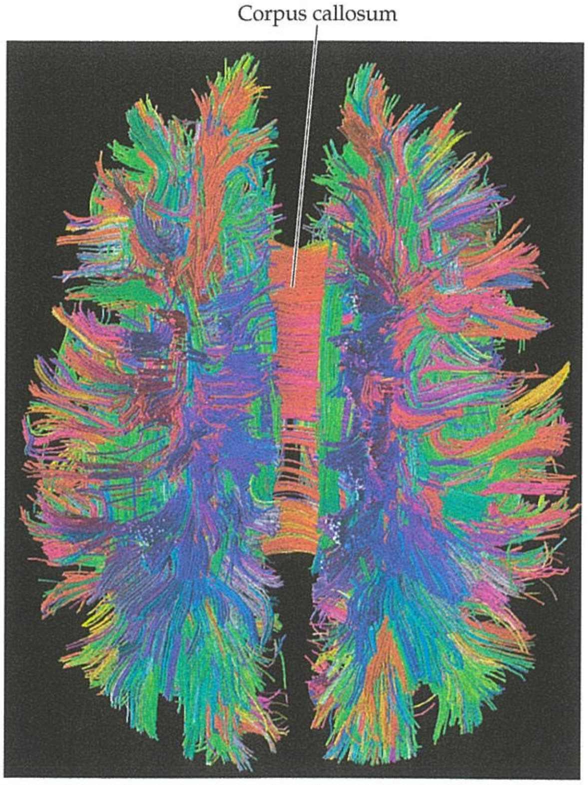

Figure 3.15 A view of the brain using dMRI, which tracks the movement of water molecules through the brain. Water diffuses in a manner that parallels the white matter tracts that carry neural signals. This imaging technique can provide insights into how information moves between various regions of the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

## *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 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.

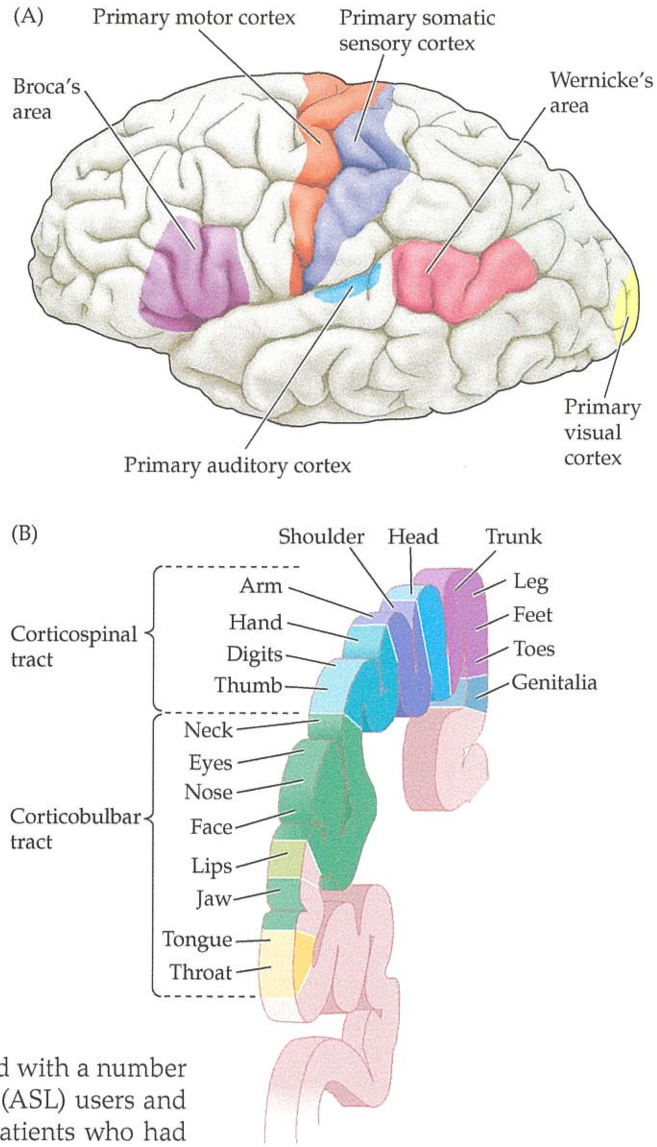

**Figure 3.16** (A) This drawing illustrates the proximity of the motor cortex to Broca's area, and of the auditory cortex to Wernicke's area. (B) A schematic illustration of the organization of the primary motor cortex. The areas that control movements of the mouth and lips are located near Broca's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

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 3.17** Examples of ASL verbs produced in response to the pictured objects.

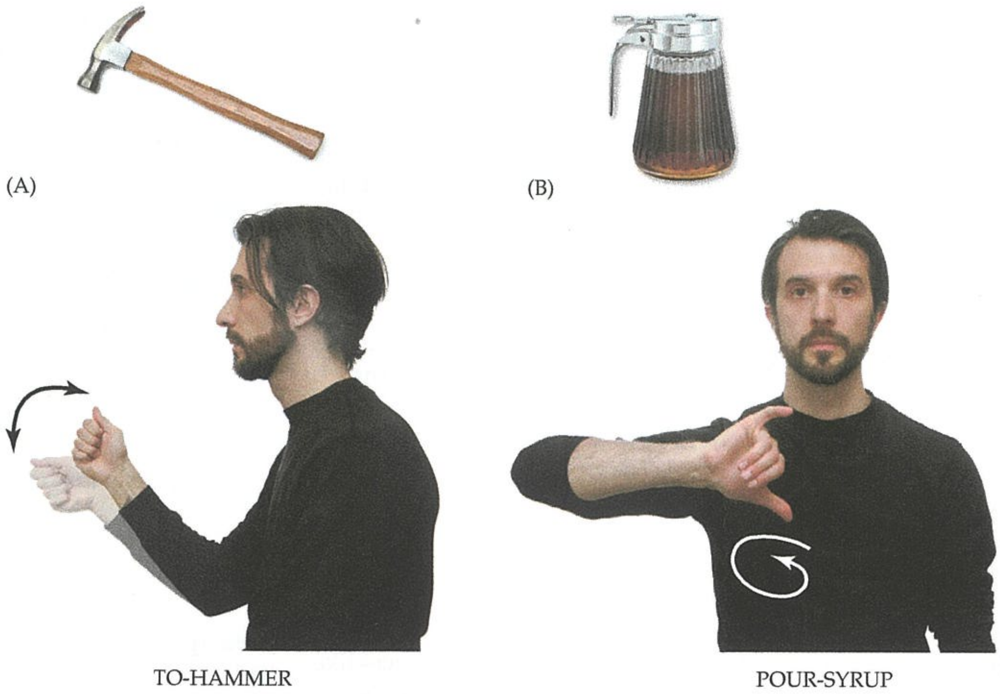

pantomime the actions rather than convey them linguistically (see **Figure 3.17**). Emmorey and her colleagues decided to look specifically at iconic signs like these, in order to see whether producing them would activate different brain regions than would pantomiming gestures, even though the hand motions for the two are actually very similar.

To elicit a linguistic sign, the researchers showed native ASL signers a picture of an object, such as a hammer or a bottle of syrup, and asked the signers to generate a verb related to that object. If pantomime gestures were being elicited, subjects were asked to gesture to show how they would use that object. **Figure 3.18** shows data from brain scans for ASL signers producing verbs and from hearing subjects who were gesturing rather than using language. As you can see, the patterns of activation are quite different; the ASL verbs resulted in more activity in the frontal lobe, home of Broca's area, while pantomime gestures triggered more activity in the parietal lobe.

Sign language studies show that when it comes to brain localization, it's not just the raw ingredients of your language that matter; it's also what you do with them. Language networks in the brain readily adapt to a slew of different materials that could be used for linguistic purposes. This is apparent in spoken languages too. For example, lan-

(A) Deaf (handling verbs)

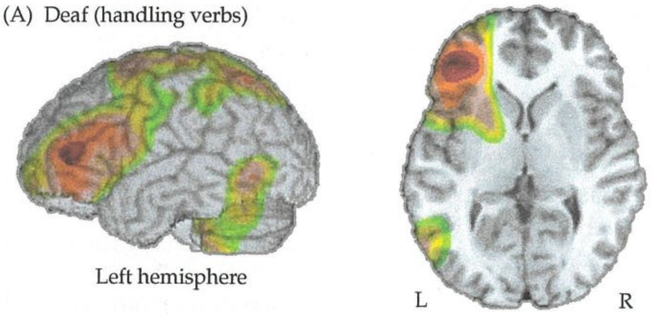

(B) Hearing ("handling verbs" pantomime)

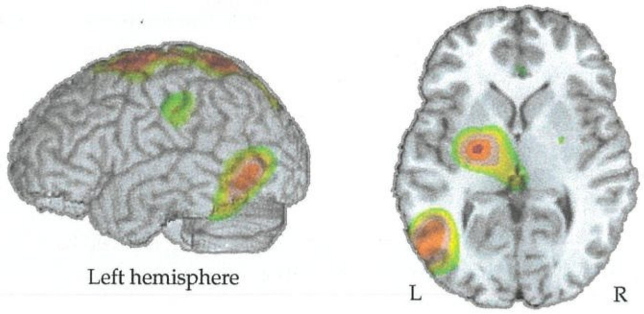

**Figure 3.18** Averaged data for (A) brains scans of deaf subjects producing ASL signs and (B) hearing subjects producing pantomime gestures in response to the same stimuli. These scans plot comparisons with a baseline condition in which subjects saw pictures of objects and gave a thumbs-up to indicate that the objects could be manually handled or waved their hands to indicate that they couldn't. (From Emmorey et al., 2011.)

## LANGUAGE AT LARGE 3.2

## Brain bunk: Separating science from pseudoscience

Phrenology, the practice of analyzing a person's character by feeling or measuring the shape of his skull, is now known as a disgraced discipline, a pseudoscience. But it was based on a set of ideas that were perfectly reasonable at the time (the early 1800s), some of which even turned out to be correct.

Phrenology's main theoretical claim held that the brain was the home of the mind, and since the mind appeared to be made up of a number of distinct faculties (which were conceived of as traits such as time and space perception, language ability, hopefulness, benevolence, acquisitiveness, etc.), these faculties must have corresponding organs in the brain (see **Figure 3.19**). It seemed logical to think that the size of any one of these organs would determine the strength of the corresponding trait for an individual, and that people might vary in which faculties were stronger than others (and hence, which of their brain organs would be bigger than others). The final piece of reasoning was that the skull formed to accommodate the shape of the underlying mental organs and that it was possible to discern a person's mental traits from the shape of the skull.

Phrenology's problem was not with the content of these ideas, all of which were interesting, testable hypotheses; it was with how people went about testing them. Instead of scientifically testing each of the major premises in a systematic way, phrenologists tended to fit the data to match their preconceived theories. The initial charts connecting features of the skull to specific traits were developed by examining people whose traits were already known, and these charts were "confirmed" by additional examinations that were biased by the pre-existing ideas. The great American humorist Mark Twain poked fun at such shoddy practices when he anonymously visited a phrenologist, only to be told that a "cavity" in his skull revealed that his "humor organ" was entirely lacking. He returned a few months later under his own name, and the very same phrenologist, not remembering their earlier encounter but now knowing him to be the famous humorist Mark Twain, examined the author and found "the loftiest bump of humor he had ever encountered in his life-time!" (Lopez, 2002).

Phrenology was eventually discredited, but not before it became wildly popular, with people paying substantial sums of money to phrenologists who would "read" their character and give them advice about which careers or marriage partners they were best suited for. In step with Mark Twain, humorist Ambrose Bierce defined phrenology as "the science of picking the pocket through the scalp" (Bierce, 1911).

Many parallels have been drawn between the pseudoscience of phrenology and the use of fMRI techniques by researchers or consultants who claim to be able to detect, on the basis of the activation of certain brain regions, whether someone will buy a particular product, or vote for a certain candidate. In one highly publicized study (Iacoboni et al., 2007), researchers tucked prospective voters into fMRI scanners and collected brain images in response to images of various candidates, or to words referring to political parties. Based on the results, they drew a number of concrete inferences. They

*Continued on next page*

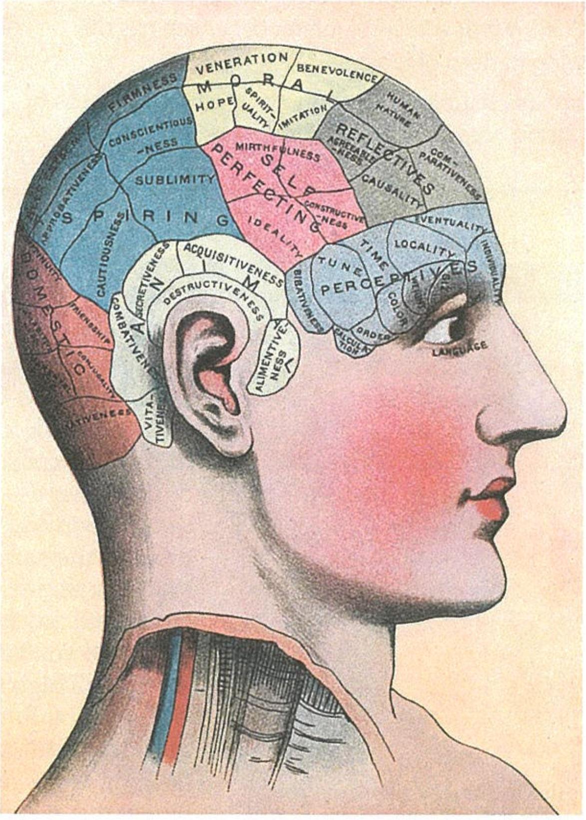

**Figure 3.19** A phrenologist's "map" of faculties believed to be associated with certain brain regions.

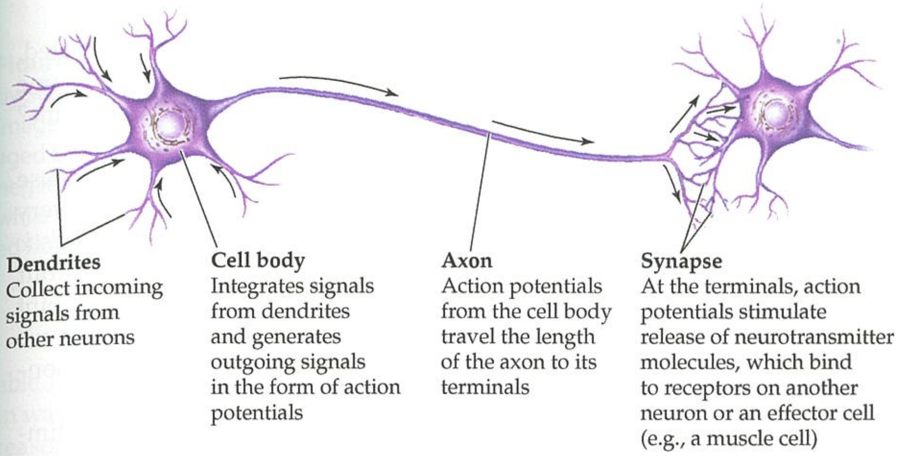

**Figure 3.20** Electrical activity in a neuron. Dendrites collect electrical signals from other neurons. These signals are integrated in the cell body, and if the resulting voltage exceeds a threshold, an outgoing signal—an action potential—is sent along the axon, releasing neurotransmitters that have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

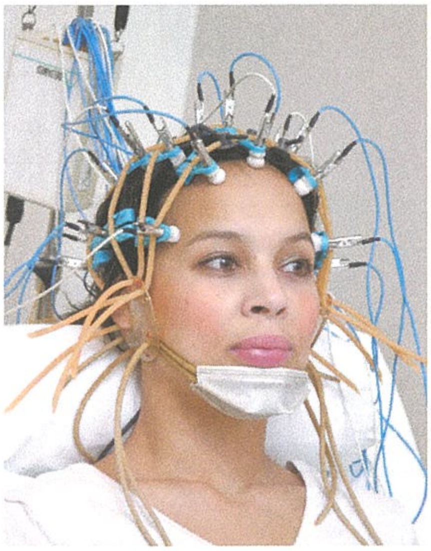

**Figure 3.21** A research participant with EEG electrodes placed over the scalp.

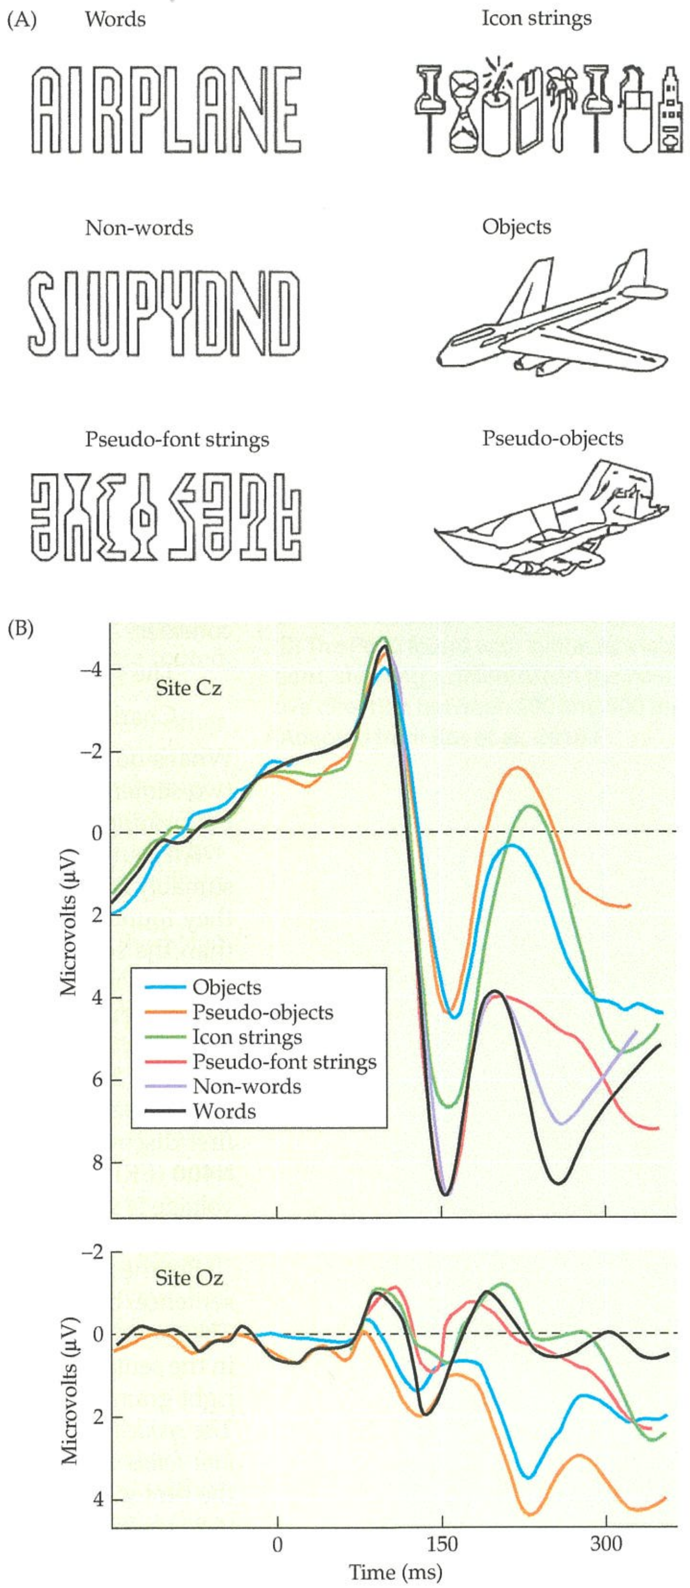

**Figure 3.22** Sample stimuli and ERP data from experiments by Schendan and colleagues. (A) Examples of the six different image types that subjects saw in random sequence. (B) Averaged ERP data from two recording sites (Cz and Oz). Note that negative voltage is plotted on the upper portion of the *y*-axis, while positive voltage is plotted on the lower portion. (Adapted from Schendan et al. 1998.)

able hypothesis, but others exist as well. The difference in waveforms could instead arise from other, more subtle reasons. For example, the real words contained clusters of letters that are more commonly found together, while the non-words contained letter clusters that are less commonly seen together. So, the difference between a word like *spark* and a non-word like *ctuik* could simply be that the letter sequences *spa* and *ark* are ones that people have seen very often before (for example, *spat, spare, span, spam, spackle, spartan; bark, lark, shark, mark, dark, embark,* and so on). On the other hand, how often have you met the sequences *ctu* or *uik*? Recognition of familiar sequences of letters doesn't necessarily mean that the word itself has been retrieved and recognized, and it could be that recognizing familiar letter strings is all that the brain is doing at 200 ms after the word's presentation. To tease apart these two alternative explanations, we need to set up yet another experiment with just the right contrasting conditions so we can test the hypotheses more precisely.

Sure enough, later ERP work by a team of French researchers (Bentin et al., 1999) did just this and compared the brain's activity in response to real French words, pronounceable pseudo-words (for example, *lartuble*), and unpronounceable processions of consonants (for example, *rtgdfs*). The pronounceable pseudo-words (*lartuble*) contained letter sequences that were common in French, while still not being real words; the unpronounceable consonant strings (*rtgdfs*), on the other hand, contained highly improbable letter sequences. The researchers found that the separation of these types of stimuli occurred in two stages. First, the waveforms showed a difference between the improbable consonant strings, on the one hand, and the real words and pseudo-words, on the other hand, showing that the brain is in fact sensitive to the combinations of individual letters. Only later did the waveforms show a distinction between the pseudo-words and the real words, with the brain activity peaking at about 350 ms after the stimuli were first seen. This is the earliest point at which we can confidently say that real words are in fact being recognized.

This set of meticulous comparisons serves as an important reminder to both researchers and smart consumers of ERP research: what might look like the most obvious difference between two types of stimuli isn't necessarily what

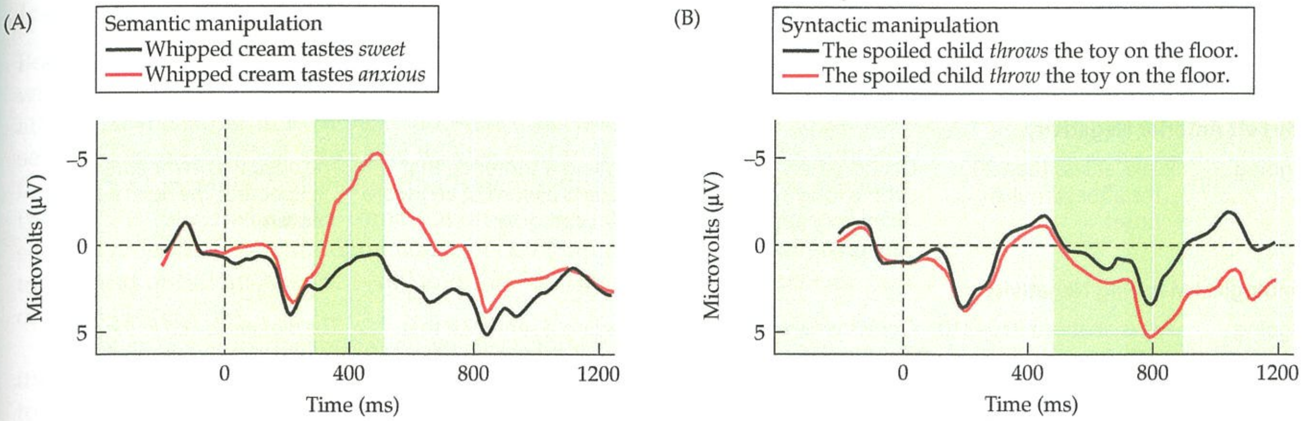

**Figure 3.23** N400 and P600 effects (shaded areas) recorded at site Pz. (A) The N400 effect found with incongruous words in a sentence frame, showing a difference in the negative direction between 300 and 500 ms. (B) The P600 found with syntactic violations, showing a difference in the positive direction between 500 and 900 ms. (Adapted from Kos et al., 2010.)

These findings led researchers to suggest that the N400 reflected the processing of meaning, while the P600 was a marker of processing syntactic structure. This proposal stirred up some excitement among psycholinguists because it hinted at the possibility that ERPs could be used to isolate different aspects of language processing and study their relative timing in the brain. The dramatically different waveforms seemed to show that meaning and syntax are routed through different processing streams in the brain.

The N400 and P600 also appeared to be different from a handful of other ERP components that had been found for processing non-linguistic information (see **Table 3.3**). This opened up the possibility that ERP research might be able to identify language-specific operations in the brain. Needless to say, this would add some important evidence to the heated debate about whether language consists of mental operations that are fundamentally different from other cognitive operations.

Remember, though: It may obvious that the difference between the N400 effect and the P600 effect is that one deals with weird meanings and the other deals with syntactic glitches. But this distinction is just the starting hypothesis. Despite the stirring implications of the discovery of these two components, the research community had to hunker down and do much more probing before it could determine whether the N400 does in fact correspond to the processing of meaning in language, or whether the P600 is the brain's signature of syntactic computation.

In the decades since the N400 was first discovered, we've learned a lot about the conditions under which it turns up in an EEG study. A word doesn't have to be nonsensical in order to trigger the N400; it just needs to be somewhat improbable or unpredictable, as measured in any one of a number of ways. For example, the N400 can be found for sentences that describe *unlikely* events, even if they're not outright nonsensical:

> He planted string beans in his car.

The N400 can also be found for uncommon words, even when they're perfectly sensible within their sentence frames. Moreover, repeating a word within the experiment leads to a smaller N400 effect for the second occurrence than for the first one. These findings suggest that the N400 isn't a direct marker of the incongruity of meaning; maybe instead it reflects the brain's efforts at retrieving a word and its meaning from memory. This retrieval process gets harder if the word is incongruous in the context of the sentence. But the accessibility of a word can also be affected by factors like how common or rare it is, or how recently it's been seen or heard. In fact, words don't even need to

## BOX 3.6

## A musical P600 effect

You don't have to be a musician to have developed very sharp cognitive expectations about music. You only need to have normal music perception and have been exposed to structured music throughout your life. In Western music, much of our musical experience centers around the structures of major and minor scales. In their 1998 ERP study, researchers led by Aniruddh Patel had subjects listen to musical sequences set in a particular key. The researchers varied whether they produced a target chord in the same key, a nearby key, or a distant key (**Figure 3.24A**). In terms of perception, in-key sounds are the most predictable, while chords from a distant key are the most jarring.

The ERP data over a number of recording sites show that compared with in-key chords, the less expected sounds elicited positive-going activity beginning at 300 ms and continuing for several hundred milliseconds (**Figure 3.24B**). Chords from distant keys showed the largest positive amplitude, while in-key sounds showed the least. When the waveforms were compared with those elicited by hearing unexpected syntactic structures, they were found to be statistically indistinguishable (Patel et al., 1998).

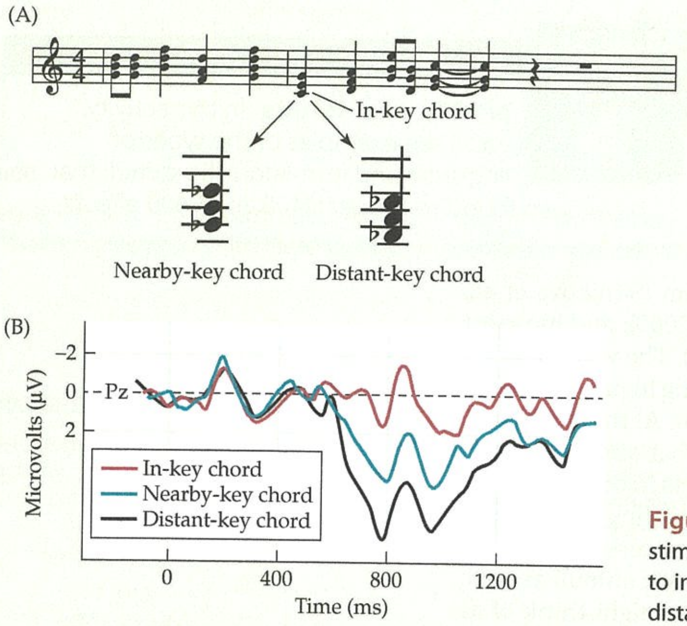

**Figure 3.24** (A) Musical notation showing Patel's sample stimuli. (B) Averaged ERP data showing responses at site Pz to in-key sounds, chords from nearby keys, and chords from distant keys. (Adapted from Patel et al., 1998.)

simply be that processing language and music both require access to shared cognitive processes at some point, even if many of their computations are carried out separately. But finding evidence of similar ERP patterns has provided a provocative launching point for further exploration. In Digging Deeper, we'll spend a bit more time looking at evidence for the neural overlap between music and language.

In short, we still don't know precisely what's going on in the brain when effects like the N400 or P600 turn up. This might strike you as vaguely depressing, given that 30-plus years and more than a thousand studies have accumulated since the N400 was first discovered. But as you'll see in some of the later chapters, ERPs have turned out to be highly valuable research tools in testing some very specific theories about the order in which various linguistic processes take place as people understand language. Even though it doesn't provide instant or magical insights about how electrical activity in the brain translates into thought, EEG research contributes some unique pieces to the overall puzzle of how language works in the brain.

GO TO
sites.sinauer.com/languageinmind
for web activities, further readings, research updates, new essays, and other features