

do not register well on the magnetometer. In contrast, the electrodes that pick up EEG and ERP signals detect voltage fluctuations produced by the return-volume currents, which are not affected by this limitation. Accordingly, EEG picks up voltage fluctuations from sources in both cortical gyri and sulci, although it tends to be more sensitive to the former.

Although MEG's insensitivity to gyral sources renders it less complete than EEG methods, what MEG *does* pick up (i.e., mainly sulcal activity) is easier to localize—for two main reasons: First, EEG currents suffer some distortion as they are conducted through the variable resistances of the skull and other tissues to reach the recording electrodes on the scalp. The magnetic fields in MEG recordings are much less affected by this problem. In addition, because MEG picks up mainly sulcal sources (and best from those nearer the magnetometer), the signal distributions it records tend to be less complex than those from ERP signals, thereby simplifying source estimation. Given their complementary sensitivities, recording by MEG and EEG methods together can offer a more complete and effective way to estimate sources and study mechanisms than does either method alone.

Positron emission tomography (PET) imaging

Probably the most popular techniques today for assessing brain activity related to cognitive functions are those that rely on measuring changes in metabolism and blood flow to visualize active areas of the brain. The attractiveness of these functional brain imaging methods arises from several factors, but derives mainly from their ability to produce images that localize brain activity with high spatial resolution.

The brain utilizes a remarkably large fraction (20 percent) of the body's energy resources, despite representing only about 2 percent of body mass, and at any given moment the most active nerve cells use more of these and other metabolites than do relatively quiescent neurons. To meet the increased metabolic demands of active neurons, the local flow of blood to the relevant brain area increases. Detecting and mapping these local changes in cerebral metabolism and blood flow is the hemodynamic basis of two functional imaging techniques that have been widely used: positron emission tomography (PET) and functional magnetic resonance imaging (fMRI). Because these techniques reveal patterns of activity in the whole intact brain with good spatial resolution, they have greatly advanced our ability to study cognitive functions and their neural underpinnings.

In **positron emission tomography**, or **PET**, unstable positron-emitting isotopes are synthesized in a cyclotron by bombarding elements such as oxygen, carbon, or fluorine with protons. Because of its short half-life, the isotope most commonly used for PET functional mapping has been ^{15}O (oxygen-15, with a half-life of 2 minutes), which is incorporated into water molecules for this activity-mapping use. When the radiolabeled compounds are injected into the bloodstream, they become distributed in the brain according to current physiological needs, accumulating preferentially in more metabolically active areas.

As the unstable isotope decays, the extra proton breaks down into a neutron and an emitted positron. The emitted positron travels several millimeters, on average, until it collides with an electron. The collision of a positron with an electron destroys both particles, emitting two gamma rays that travel in opposite directions from the site of the collision. Gamma-ray detectors placed around the subject's head are arranged to register a "hit" only when two detectors 180 degrees apart react simultaneously (**Figure 2.11A,B**). By reconstructing the density of these collision lines using computer algorithms, the location of the active regions can be imaged. The PET images of activity can then be superimposed onto magnetic resonance structural images from the same subject(s)




Figure 2.11 Positron emission tomography (PET) (A) In PET, molecules labeled with radioactive probe atoms such as ^{15}O are injected into the bloodstream. These agents travel preferentially to areas of increased neural activity because of increased metabolism and thus blood flow to those areas. Radioactive decay ultimately emits two “annihilation photons” (gamma rays) in opposite directions. (B) The annihilation photons are detected by sensors surrounding the head of the subject lying in the PET scanner. (C) A PET activation image overlaid on a structural magnetic resonance image. (C courtesy of David Madden.)

to provide spatial information about specific brain areas involved in cognitive and other functions (Figure 2.11C).

The mean distance that the positrons travel in brain tissue before the gamma-producing collision with an electron limits the theoretical resolution of PET scanning to several millimeters, although true resolution is typically considerably lower because of spatial smoothing and other steps in the analysis. An even greater limitation of PET, however, is the time required to accumulate an adequate signal, which precludes its use in many cognitive paradigms. In particular, PET activation imaging requires accumulating signal over the period of about a minute or so (i.e., the signal to obtain a usable image must be compiled over this time frame), and thus it has essentially no temporal resolution. As a result, functional imaging experiments using PET require a **blocked design**. In this design, the brain activity measured reflects neural activity integrated over an extended period (a block) while the subject performs a task or is presented with a particular stimulus condition. The resulting activation pattern can then be compared to the activity recorded during a different block of time in which the subject is not doing the task or is under a different condition. The need to integrate activity across time in this way thus substantially limits the specificity with which brain activation can be associated with particular cognitive processes.

The various limitations of PET (the short half-life of the reagents, the need for a nearby cyclotron to create the reagents, the use of radioactivity, and the very poor temporal resolution) have led to its replacement in most cognitive studies today by functional magnetic resonance imaging.

Functional magnetic resonance imaging (functional MRI or fMRI)

Functional magnetic resonance imaging (fMRI; Figure 2.12A) is based on the fact that oxyhemoglobin (hemoglobin carrying a bound oxygen molecule) and deoxyhemoglobin (the oxygen-depleted form of hemoglobin) have different magnetic resonance signals. As noted already, active brain areas use more oxygen than relatively inactive areas, and thus require more local blood flow. Within a second or two after an area has been activated by a specific cognitive event or task, the related microvasculature responds to the resulting local oxygen depletion by increasing the flow of new, oxyhemoglobin-rich, arterial blood to the active area (Figure 2.12B). The relative concentration of deoxyhemoglobin decreases over the next few seconds, leading to localized changes in the magnetic resonance (MR) signal from that part of the brain (Figure 2.12C). These

(A)






Figure 2.12 Functional magnetic resonance imaging (fMRI) (A) A subject is moved into an MRI scanner for functional MR imaging. The machine itself is not different from the scanner shown and described in Box 2A. The subject is wearing goggles in order to view visual stimuli for the experiment to be undertaken. (B) The blood oxygenation level-dependent (BOLD) signal. The vascular system supplies blood containing oxyhemoglobin to active regions of the brain. The influx of oxygenated blood to regions that become active reduces the local concentration of deoxyhemoglobin, which increases the BOLD signal; the difference in the signal provides a measure of local neuronal activity. (C) There are several ways to display functional MRI data on corresponding structural MRI scans, including overlaying onto an image slice; overlaying onto a three-dimensional image of the cortical surface extracted computationally from the MR structural image data; or overlaying onto the three-dimensional image of an extracted cortical surface that has been “inflated” by a computer algorithm. (B after Huettel et al. 2009.)

(B)




(C)


fMRI on MR slice



fMRI on folded brain



fMRI on “inflated” brain



MR changes, known as the **blood oxygenation level-dependent (BOLD)** signal, are the basis of most current forms of fMRI.

Functional MRI offers better spatial localization than PET (several millimeters compared to approximately a centimeter in PET), as well as much better temporal resolution (a few seconds compared to a minute or more in PET). In addition, fMRI uses endogenous signals intrinsic to normal brain function rather than signals originating from exogenous, radioactively labeled probes. But probably the greatest advantage of fMRI over PET for the study of cognitive processes is its much higher temporal resolution.

Figure 2.13 Experimental designs in PET and fMRI studies (A) In PET functional imaging studies, which necessarily employ a block design, the subject performs a cognitive task continuously over approximately 1 minute, with different tasks being performed in each block (in this example either task A or B, or, in the resting block, neither). If ^{15}O -labeled water is used as the tracer, a 10-minute break between blocks is necessary to allow the radioactively labeled tracers to decay sufficiently before the next block. Accumulated activity for each block is then measured, and functional activity is inferred by comparing the different blocks. (B) The block design for fMRI is similar, but with multiple samples of brain activity in each block, much shorter blocks, and no need for a 10-minute break between blocks. (C) Because of the much faster activity sampling possible with fMRI (typically every 1.5 to 2 seconds), the fMRI responses to individual events can be captured in an event-related way. Shown is a representative fMRI response to a single event type (event type A). (D) This paradigm allows the time-locked responses to different event types within the same run to be selectively extracted (i.e., event types A and B).

Most initial fMRI studies in the early 1990s followed the same blocked design as PET studies (Figure 2.13A), although the blocks were only around 20 to 30 seconds, as compared to a minute, and there was no need for a long waiting period between blocks (Figure 2.13B). However, investigators soon realized that fMRI had enough temporal resolution to allow an **event-related design**. Today, samples of fMRI activity throughout the whole brain can be acquired in a second or two, allowing the average fMRI hemodynamic response to a single stimulus or cognitive event type to be captured. The fMRI response to a stimulus or cognitive event begins at about 1 to 2 seconds, peaks at about 5 to 6 seconds, and returns to baseline at about 12 to 15 seconds (Figure 2.13C). Although this sequence is very slow relative to the millisecond resolution of electrophysiological methods such as EEG, MEG, or single-unit recording, it is fast enough to allow fMRI to be applied in an event-related manner, at least when stimuli are presented at a moderately slow rate. As shown in Figure 2.13D, this capability allows selective averaging of the responses to different event types *within* a run, much as is done with electrophysiological methods. Moreover, using temporal randomization of the stimulus or task events and modern signal processing algorithms that effectively disentangle the overlapping responses from successive stimuli, it is now possible to present stimuli at faster rates and still extract an event-related response (Figure 2.14).

The development of event-related methods for fMRI was a highly significant advance, as it enabled many cognitive paradigms that had not been possible in the early days of fMRI. Moreover, improved analytic techniques in event-related paradigms made it possible to link the *behavioral* measures from trial to trial to the





Figure 2.14 Event-related fMRI responses Responses are from a visual stimulation paradigm obtained at fast stimulus rates (two events per second). Unilateral visual stimuli were presented in random order to the left and right visual fields while subjects fixated on a central cross. At this high stimulation rate, there is severe overlap of the fMRI responses to successive stimuli in the sequence. The overlap was eliminated in this example by a subtraction technique, yielding clean average hemodynamic fMRI responses (right panels) in the visual sensory areas in contralateral occipital cortex (left panels). (After Burock et al. 1998.)



corresponding fMRI *neural* responses in different brain regions. In particular, variations in a behavioral measure (e.g., the reaction time) could be examined with respect to how they co-varied with the amplitudes of the corresponding fMRI responses from trial to trial. The development of event-related paradigms using fMRI has also allowed researchers to compare the response images with the activation patterns of high-temporal-resolution ERP and ERF responses to the same types of events. In addition, modern fMRI techniques allow the use of “hybrid,” or “mixed,” designs that have both block- and event-related experimental manipulations, the effects of which can be disentangled by appropriate analyses.

As a result of its many advantages, fMRI has emerged today as the dominant technology for imaging brain activation in studies of both normal and abnormal cognitive function.

Using fMRI to analyze activation patterns within a brain area

The bulk of fMRI research has focused on activations of various brain areas as a function of a specific manipulated variable, such as the presentation of stimuli of a particular type, under particular cognitive conditions. Yet, as later chapters make plain, the cortex possesses considerable local organization within particular areas, and these organizational characteristics are often functionally important. Although such local characteristics may be reflected in intervoxel differences in fMRI, standard fMRI analysis does not typically take them into consideration. Indeed, researchers usually apply spatial smoothing to their fMRI data to increase sensitivity, and they require activations to be of a specified minimum number of voxels to be meaningful. Both of these strategies tend to obscure differences between nearby voxels.


A number of approaches aimed at addressing this problem are based on what are known as **pattern classification** algorithms. In particular, within fMRI, researchers use an analysis scheme called **multivoxel pattern analysis (MVPA)** for such purposes, which looks at patterns of activation across voxels that consis-

tently correspond to a particular stimulus type or event type, rather than the overall increase or decrease in activation of a large brain region. For example, in an area of the lateral occipital cortex that is generally responsive to objects, images of cars may activate the area about as much as houses and tools do, at least collapsed across all the voxels of the “object analysis” area. But one type of visual object may induce a particular pattern of activation among those voxels that is consistently different from the pattern that objects of a different class would induce. MVPA specifically analyzes whether such trial-related activation patterns are indeed consistently observed. In addition, these sorts of analyses can be used to predict the patterns of activation for new examples of the different stimulus types. Such approaches have thereby enabled, for example, inferring from brain activation data the particular type of stimulus a participant is likely looking at. Thus, these sorts of pattern analysis approaches are becoming increasingly popular in cognitive neuroscience fMRI research.

Repetition suppression is another way researchers have endeavored to assess functional characteristics of a brain area that respond similarly to different stimulus classes in overall activation magnitude (Figure 2.15). This technique rests on the tendency of areas of the brain to respond less to the repeated occurrence of an identical (or closely similar) stimulus than to the first presentation of the stimulus. In fMRI, the version of this approach most commonly used is known as **fMRI adaptation**. The use of fMRI adaptation to enhance the understanding of a brain region is based on the assumption that if that region shows repetition suppression for the second stimulus in a pair (e.g., the target in a particular prime-target combination), then the two stimuli must share a process supported in some way by that region. Thus, by manipulating the relationship between primes and targets, it is possible to constrain hypotheses regarding the function of a particular brain region. For example, the lateral occipital region mentioned in the previous example as being activated during the processing of visual objects showed repetition suppression for repeated objects (e.g., an image of a table) even when the two stimuli were presented in different formats (e.g., as a drawing versus as a photo) or with different angular views. This finding suggests that the lateral occipital complex is concerned with processing an object’s abstract properties rather than the details of its sensory features.

In some instances, fMRI adaptation can provide additional information about the processing role of a brain region that standard fMRI methods cannot. For example, in conventional fMRI, comparing two stimulus conditions A and B may fail to reveal differences between them because the populations of neurons sensitive to these conditions (or two stimulus types) are mixed within the same brain region being sampled (see Figure 2.15A). In contrast, in the fMRI adaptation paradigm, the subpopulations of neurons sensitive to stimulus A may show adaptation when this stimulus is repeated, giving rise to a lower fMRI signal than that elicited for stimulus B (see Figure 2.15B).

Figure 2.15 Repetition suppression in fMRI studies (A) A conventional fMRI experiment may fail to distinguish between responses to stimulus A and stimulus B because the two stimuli activate different populations of neurons (orange versus blue in the figure) within the same brain region sampled; hence the total fMRI signal may be similar. (B) Using a form of repetition suppression (fMRI adaptation), it may be possible to detect a difference in fMRI signal within this brain region. When the subpopulation of neurons sensitive to stimulus A adapts, the signal is reduced for a repetition of stimulus A but strong for stimulus B. (After Kourtzi and Grill-Spector 2005.)



Using fMRI to examine activity relationships between brain areas

Most neuroimaging studies in cognitive neuroscience build on the concept of localization of function—namely, that different brain regions support different types of information processing. However, brain regions do not participate in cognitive brain function in isolation. Information flows between regions via the action potential signals transmitted by axons, which are bundled together into large **fiber tracts**. Although an enormous amount of knowledge has been gained about such anatomical connections, knowledge of the structural connections between brain regions can provide only a limited picture of information flow in the brain. Understanding of the **functional connectivity**, or how the activity of one brain region varies with the activity in another brain region, is also needed.

Functional MRI can help elucidate the functional relationships between brain regions because of its ability to acquire functional activation data both from across the whole brain and from across time. The simplest type of relationship between regions is **coactivation**, in which two or more brain regions change similarly in response to an experimental condition. A straightforward and rather complementary way to show that two (or more) regions share functional connectivity comes from studies of **resting-state connectivity**, which identify fMRI signal fluctuations that co-vary across brain regions while subjects lie in the MRI scanner but do not perform an experimental task. These co-variations are thought to reflect spontaneously occurring operations of the brain rather than the processing of specific external stimuli.

One way to extract the distributed patterns of common activation that occur during rest is to use data exploration techniques such as *principal component analysis* (PCA) or *independent component analysis* (ICA), which use algorithms for extracting similarly behaving portions of a complex data set. Alternatively, one can choose a *seed voxel* or group of voxels in a specific region of interest and determine which other regions exhibit similar fluctuations in activation over time. Using these approaches, researchers have found co-varying activation in groups of regions that have been implicated in other studies as being involved in visual, auditory, memory, and attentional functions. In general, regions that are coactivated during cognitive tasks also tend to show resting-state connectivity, suggesting that brain regions with common functionality tend to express similar patterns of spontaneous activity. Combining resting-state and task-related connectivity within the same study can strengthen conclusions about the relationships between brain regions, by building on the findings of one to explore the other. Moreover, explicit information on the structural connections between brain areas in individual subjects, such as can be obtained with diffusion tensor imaging (DTI; see Box 2B) can also help with these analyses.

Although finding that two brain areas are both activated during a particular task indicates that the two regions are related, it does not tell us much about the form of that relationship. For example, consider a task entailing attending to a visual scene to search for faces. Such a task (relative to a baseline task) will elicit enhanced activity in the dorsolateral prefrontal cortex, the inferior parietal sulcus, and face-processing areas in ventral occipital cortex. Such a coactivation might suggest that these various areas are involved in the task, but we would not know how they work together to accomplish that task. For example, top-down influences from the prefrontal cortex might be guiding the activity in the other regions. Another possibility is that visual processing in the ventral occipital cortex is triggering activity in the other areas in a bottom-up fashion. On the other hand, the activity in all of these regions may be controlled by signals from another region of the brain.

A variety of methods have been developed to gain further insight into the relationships between brain areas and their activation patterns. In one approach

for combining fMRI and behavioral data, **psychophysiological interaction (PPI) analysis** is used. Like some other connectivity approaches, PPI analysis uses the time course of activity in one seed region to predict changes in activity in another region. However, the technique analyzes how these interactions differ as a function of the cognitive task being performed. For example, it analyzes whether the correlation in activity between the two areas, rather than activity itself, differs in one task versus another.

A significant shortcoming of fMRI (and indeed most other neural activity monitoring methods) is its difficulty in establishing *causality* in the interactions between brain areas—that is, whether activity in one brain area caused activity in another. Suppose that two regions, A and B, are coactivated in an fMRI study in a particular task. Is it possible to infer from the fMRI data whether activation in Region A caused changes in Region B, activation in Region B caused changes in Region A, or activation in both of these regions was induced by signals from another brain area?

Some modeling approaches, such as **structural equation modeling**, have endeavored to infer causality from fMRI activation patterns by incorporating known anatomical connections, comparing models with different hypothesized sets of directional influences between areas, and identifying which model best accounts for the observed data. For example, an early application of this approach to studying attention found that the best model for the activation patterns suggested that attention-related signals from the prefrontal cortex increased the connectivity between the extrastriate visual areas and the posterior parietal cortex. Related approaches include **dynamic causal modeling**, which creates models of the functional connections between brain regions, and then evaluates how that connectivity changes as a result of experimental manipulations.

In summary, delineating the functional connectivity relationships between brain regions shown to be activated in fMRI tasks is critical for understanding how these regions work together to accomplish cognitive functions. The ongoing development and application of these various approaches for analyzing functional connectivity relationships will continue to improve the inferences that can be made from fMRI data.

Optical brain imaging

Another way to measure neural activity in both animals and humans is based on optical imaging techniques. These approaches derive from the fact that active brain tissue transmits and/or reflects light differently than does inactive brain tissue, and these differences can be picked up and imaged by optical recording devices.

One such approach is, like fMRI, based on hemodynamic changes in response to neural activity. When populations of neurons become active, thereby inducing local hemodynamic changes, the amount of light reflected from the cortical surface changes as well. This optical signal, as with the BOLD signal in fMRI, is thought to arise from changes in the relative concentrations of oxyhemoglobin and deoxyhemoglobin that occur as neuronal activity varies in a given region in response to a stimulus. To detect these changes requires that the skull be opened and the cortical surface illuminated with a red light in the 580- to 700-nanometer range. The reflected light is picked up by a sensitive video camera, and the changes in responses to stimulation are averaged over a series of trials. The resulting images can provide a highly detailed map of the cortical patterns of activity.

These hemodynamically based optical imaging approaches have been most successfully applied to studying sensory map representations in the sensory cortices in animals, particularly in visual sensory cortex. For example, this type of imaging has produced highly detailed maps of the patterns of relative

sensitivity across primary visual cortex to specific line orientations (see Chapter 3). These methods have also been useful in showing the relationships between the mapping configurations for different features, such as the relationship between line orientation preference maps and ocular dominance maps. Because the technique is invasive, it has been used mostly in experimental animals, and the results have thus far tended to be more informative in the context of neurobiology than cognitive neuroscience. Nonetheless, it has been used in a few circumstances in human patients undergoing neurosurgery (e.g., to map language areas). Because of its simplicity and relatively low cost, such optical imaging would seem to have considerable promise for future studies of cognitive functions in both experimental animals and neurosurgical patients.

Another technique, known as **event-related optical signals (EROS)**, uses optical methods based on a very different activity-dependent mechanism and can be applied noninvasively. When brain tissue is illuminated, even through the skull, the amount of light that is transmitted versus scattered varies as a function of whether the neuronal tissue is electrically active. This differential light diffusion is thought to arise from changes that occur in neuronal membranes and their immediate surroundings during electrical activity. By shining light into the brain from a number of different sources around the head and measuring the intensity of the reflected light with multiple external sensors, these activity-dependent changes can be detected and imaged using three-dimensional, computer-assisted reconstruction algorithms.

Like ERPs, EROS has high temporal resolution but relatively low spatial resolution. However, determination of the sources of the measured activity is not subject to the inverse problem inherent in EEG and MEG source analyses. Although the relatively low signal-to-noise characteristics have limited the use of this method, EROS holds significant promise as a useful method for noninvasively imaging cognition-related cortical activity.

Assembling Evidence and Delineating Mechanisms

A major goal in the quest to understand the neural bases of cognition is to establish links between localized brain structures and neural activity on the one hand, and cognitive functions or processes on the other. To establish these links and delineate underlying mechanisms, cognitive neuroscientists often assemble evidence from multiple studies and data gleaned through the implementation of multiple methodologies.

Associations and dissociations

One approach to linking neuroscience and cognition is apparent in the methods already discussed: to experimentally associate specific cognitive functions with the neural structures and/or local neural activity that underlie them. A complementary approach is to determine cognitive functions and neural processes that do *not* seem to be related—that is, that are dissociated. For example, finding that damage to a particular brain area disrupts the performance of Task A but not Task B establishes not only an association between the damaged area and Task A, but also a dissociation between the lesion and Task B (**Figure 2.16A**).

An even firmer link between a specific cognitive function and a particular neural structure or process can be made by establishing a **double dissociation**. For example, suppose Task A is designed to engage mainly System A (say, the neural system for identifying a visually presented face), whereas Task B is designed to engage mainly System B (the neural system for identifying the emotional expression of a visually presented face). If Systems A and B are relatively distinct, one would expect to be able to find neurological patients




Figure 2.16 Associations, dissociations, and double dissociations (A) The relative performance on two tasks (A and B) for neurological patients with a common brain lesion in Region 1. Impairment in performing Task A but not Task B indicates an association between this region and performing Task A, and a dissociation between this region and performing Task B. (B) Two groups of patients, one with lesions as Region 1, as in panel (A), plus another group of patients with lesions in a different brain region (Region 2) who show the opposite pattern of impairment in the two tasks. The results from the two patient groups taken together are a double dissociation providing strong evidence for the specific involvement of the two regions in the two different tasks. (C) Analogous double dissociation determined from brain activation data (e.g., from fMRI). In this case, subjects engaged in Task A show activation of Region 1 but not Region 2; during Task B this pattern is reversed.

with a particular brain lesion who are impaired in performing Task A (recognizing faces), but essentially normal when performing Task B (recognizing emotions), and to find patients with a lesion in a different brain region who showed the opposite pattern (normal performance on Task A but impaired performance on Task B). This pattern of effects constitutes a classic double dissociation (Figure 2.16B).

Double dissociations provide more definitive evidence for relatively separate mechanisms or functions than do single dissociations. This is because a single dissociation could arise from a general factor (such as the relative difficulty of the tasks) rather than from the engagement of separable functions. However, when another type of lesion produces the opposite, or at least quite different, characteristics, the case for the engagement of functionally distinct neural systems becomes much stronger. It often turns out, however, that the neural systems being examined are only partially separable; that is, performance on Task A for a certain patient group is seriously impaired, but performance on Task B is also disturbed to some degree (or vice versa). This situation would suggest that the functional systems in question interact to some degree—which of course is useful information as well.

Although these examples focus on establishing dissociations using brain lesions, double dissociations are also important in methods that measure brain activity during cognitive tasks (Figure 2.16C). For example, a neuroimaging experiment might indicate that a specific brain area (say, Region 1) is *activated* during Task A but not during Task B. As with brain lesion studies, however, this neuroimaging pattern could be due to differing difficulty of the two tasks or to the degree to which a neural system is being engaged by the task. Thus, establishing a double dissociation in the context of neuroimaging (i.e., finding that another brain area, Region 2, is activated during Task B but not during Task A) is also important for more firmly linking specific brain activity measures to specific cognitive functions.

TABLE 2.1 Summary of the Major Imaging Techniques Used in Cognitive Neuroscience

TECHNIQUE	MAIN ADVANTAGES	MAIN DISADVANTAGES
Naturally occurring lesions	Can strongly implicate a region as being essential for a task Occur naturally	Still generally need double dissociation to strongly confirm selectivity of area Not specific to functional areas; variable in distribution and extent Do not identify a network No temporal resolution Relatively few available subjects; subjects often from heterogeneous groups Effects of recovery unknown or complex
Directed lesions	Can strongly implicate a region as being essential for a task Can be much more selective than naturally occurring lesions Can be timed (e.g., before or after training)	Can generally be done only in animals; ethical concerns Very limited temporal resolution
Intracranial stimulation	Can provide rather specific neural perturbation	Mostly limited to animals and rare clinical circumstances In humans, clinical concerns, limited locations that can be stimulated
Transcranial magnetic stimulation (TMS)	Advantages of lesions, but transient and noninvasive (or at least nonsurgical) With single shot, can get some temporal resolution	Mostly can do only superficial cortex Not very focused; stimulates other areas nearby and above the target area Even for some superficial brain regions, is too uncomfortable Some safety issues, particularly for repetitive TMS (rTMS)
Single-unit recordings	High spatial and temporal resolution Very specific (single neurons)	Picks up only some neurons (typically larger ones) Very invasive; almost completely limited to animals Typically from only one brain area, thus does not identify a network or network interactions
Electroencephalography (EEG)	Good temporal resolution Good for state effects (e.g., arousal, sleep stages) Noninvasive Inexpensive, fast, and easy recording procedures	Coarse or problematic spatial resolution Not very specific for information processing or cognitive function

Multimethodological approaches

All of the individual methods described here provide a way of linking cognitive processes to underlying brain processes. Each has advantages and disadvantages. The spatial and temporal ranges of these approaches were depicted in Figure 1.6. The most important advantages and disadvantages of each method are described in [Table 2.1](#).