


**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

## BOX 14C NEUROMARKETING



Most applications of cognitive neuroscience have been within the medical domain. Research on the neuroscience of decision making, however, has become increasingly relevant for applications to real-world consumers, leading to the emergence of **neuromarketing**. Applied research in neuromarketing seeks to develop algorithms that can use neuroscience data from a small group of individuals—in a sense, a focus group consisting of human brains—to predict consumer choices for real-world products. This may seem like an unnecessary expense, in that marketers do not need to measure brain activity to determine whether a particular advertisement causes more people to buy their product; they can simply measure purchasing behavior, as has been done for decades. Yet there has been an explosion of interest in neuromarketing research, and many companies have used it for product development and marketing. If neuroscience simply duplicates basic measurements of behavior—at greater apparent expense—then why has neuromarketing grown so rapidly?

To understand why neuromarketing has attracted such attention, consider the key goals of commercial marketing, which extend well beyond measuring

Activation in the ventromedial prefrontal cortex tracks subjective preferences. Functional MRI activation in the ventromedial prefrontal cortex (top) was different in subjects anticipating a squirt of Coke or Pepsi (bottom, y-axis), depending on whether the subject tended to prefer Coke or Pepsi in a prior taste test (bottom, x-axis). (After McClure et al. 2004.)

purchasing behavior. Modern marketing research attempts to improve all aspects of the consumer product cycle: developing new products and advertisements, assessing the factors that shape consumer preferences (e.g., understanding how a given brand influences choices), and getting the right messages to the right consumers through market segmentation. Neuroscience methods have been applied in each of these areas.

Think about an advertisement for a consumer good (e.g., clothing) on television. A typical 15-second commercial contains about 10 separate video segments, each with multiple images, all targeted toward a single message. Traditional methods for assessing and optimizing advertisements involve showing the commercial to potential consumers and then soliciting evaluative data



studying decision making; it will be a useful tool for many applications, but it will not replace traditional methods for economic or psychological analysis. Thus, a key challenge for cognitive neuroscientists who study decision making lies in building productive links between neuroscience and other disciplines.

Moreover, like other complex cognitive functions, decision making does not exist in isolation. Real-world choices rely on the processes of attention, memory, and emotion described in other chapters. People rarely, for example, make decisions of the sort studied in typical laboratory paradigms, such as choices between monetary gambles with known probability and fixed outcomes. Most decisions involve more tangible and meaningful outcomes, for which we often have substantial prior knowledge. As decision-making paradigms become increasingly realistic, researchers will need to incorporate these other cognitive functions in their models.

Even with these caveats, however, research on the cognitive neuroscience of decision making can serve as a bridge between basic decision science and some real-world applications (Box 14C). For instance, economists are increasingly interested in populations whose decision making is systematically maladaptive, such as addicts, habitual criminals, sociopaths, and others with behavioral pathologies. Because different neural deficits characterize each group, a policy of intervention that might be appropriate for one group would not necessarily be appropriate for others. Maladaptive decision making is especially evident in drug addiction. Addicts make repeated poor decisions that contribute to