CHAPTER 13

BRAIN IMAGING AND RELATED METHODS

David H. Zald and Clayton Curtis

Modern neuroimaging techniques enable researchers to noninvasively assess brain structure and function in humans. The knowledge gained from these techniques has led to a revolution in our understanding of brain–behavior relationships and has dramatically altered the psychological sciences. Several brain imaging techniques are currently in wide use, including computerized tomography (CT), magnetic resonance imaging (MRI), positron emission tomography (PET), single photon emission computed tomography (SPECT), magnetoencephalography (MEG), and near infrared optical imaging. In this chapter we focus on functional MRI (fMRI) and PET techniques because of their enormous impact on the psychological sciences. Although they are most often used in isolation, both PET and fMRI are adaptable to a multitrait–multimethod (MTMM) approach toward assessment. Indeed, it may be argued that these techniques require integration within a broad multimethod framework if they are to reach their full scientific potential. This chapter provides a brief primer on fMRI and PET imaging, followed by a discussion of the benefits of placing neuroimaging data within a MTMM approach.

Depending on the specific technique used, PET and MRI scanners can assess a number of different statewise and traitwise characteristics of the brain. These include measurement of brain structure (MRI), neurotransmitter functioning (PET and magnetic resonance spectroscopy [MRS]), glucose metabolism (PET), blood oxygenation (PET and fMRI), and blood flow (PET and fMRI). Because changes in neural activity are accompanied by changes in metabolism, blood oxygenation, and blood flow (Raichle, 1988), PET and fMRI measurements of these physiological variables allow researchers to index changes in brain functioning in relationship to specific perceptual, cognitive, and behavioral tasks. However, PET and fMRI take very different approaches to these measurements. It therefore is useful to first discuss how these measurements are made in each technique.

fMRI PHYSICS AND PHYSIOLOGY

When biological tissue is placed within a strong externally applied magnetic field, denoted B_0, the axis of individual nuclei, like hydrogen, tend to align with the field. Nuclei line up with the field because this results in the lowest energy state of the system. Outside the magnetic field, the alignment of all nuclei tends to be randomly oriented and produce no net magnetic field. However, when placed in a strong magnetic field, the nuclei align in the same direction as the field. This alignment produces a net magnetization, referred to as M, which represents the sum of all of the magnetic moments of the individual hydrogen nuclei (see Haake, Brown, Thompson, & Venkatesan, 1999, for a full review of MRI physics).

Hydrogen nuclei consist of a single positively charged particle, the proton, which spins around its axis. An individual proton not only spins around its axis, but also precesses (revolves) about the external magnetic field, much like a top both

http://dx.doi.org/10.1037/11383-013
Handbook of Multimethod Measurement in Psychology, edited by M. Eid and E. Diener

spins around its axis and precesses about the direction of gravity's magnetic field. Importantly, each type of atomic nuclei precesses at a characteristic frequency, the resonance or *Larmor frequency*, which is directly proportional to the strength of the applied magnetic field. This proportional dependence of the resonance frequency on the applied magnetic field forms the basis for MRI. Specifically, by spatially manipulating the field strength and measuring resonance frequencies, it becomes possible to resolve the source and location of signals from the brain.

When all the nuclei in a sample are at the resting equilibrium state, the net magnetization of the nuclei are aligned with the field, and no MR signal can be detected because each of the nuclei precesses at the same rate, but out of phase with one another. Magnetic resonance occurs when a radiofrequency (RF) pulse is transmitted to the sample at the Larmor frequency of a specific type of nuclei. For instance, hydrogen (H) precesses at a frequency of 64 MHz in a 1.5 Tesla (T) magnetic field (standard clinical scanners possess a 1.5 T field strength, whereas research dedicated scanners often use higher field strengths, such as 3 T, 4 T, or even 7 T). When an RF pulse is applied at the Larmor frequency of H, energy is selectively absorbed by H nuclei, exciting their spins from their lower resting state to an unstable higher energy state. The RF pulse also deflects the net magnetization of the nuclei away from the direction of the external magnetic field and causes each precessing nuclei to precess in phase with one another (i.e., they become phase coherent). At the point in time when the RF field is extinguished, the nuclei are in an excited, high-energy state because the axes of their small magnetic fields are not oriented with that of the strong external field. This unstable state decays quickly as the nuclei begin to realign with the external field. The precessing nuclei radiate the energy that they absorbed from the RF pulse as the phase coherence exponentially decays and the net magnetization of the nuclei realign with the external magnetic field. The energy that is emitted during this brief process induces a detectable current (known as the free induction decay or FID) and is detectable by an RF coil placed around the stimulated sample (i.e., the subject's head). This is the MR signal and forms the basis of all MRI techniques.

When the application of the RF energy is terminated, the system reapproaches equilibrium, a process known as *relaxation*. Different types of tissue have different rates of relaxation, which is why we can obtain MR images that can distinguish between gray and white matter, bone, cerebrospinal fluid, and vasculature. For most functional MRI studies, the critical source of contrast derives from changes in the oxygen content of cerebral vasculature, typically referred to as *Blood Oxygen Level Dependent* (BOLD) signal (Bandettini, Wong, Hinks, Tikofsky, & Hyde, 1992; Kwong et al., 1992; Ogawa, Lee, Kay, & Tank, 1990; Ogawa et al., 1992).

The BOLD Signal

The fMRI signal is a function of the metabolic demands of local neural activity. However, the coupling between the measured BOLD signal and the underlying neural activity is neither direct nor straightforward (Heeger & Ress, 2002; Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001). As neural activity increases, there are changes in both the amount of blood flow to the region and a change in the concentration of oxygenated and deoxygenated forms of hemoglobin. The oxy- and deoxyhemoglobin have different magnetic properties (diamagnetic vs. paramagnetic) and because of this behave differently within a magnetic field. The paramagnetic properties of deoxyhemoglobin lead it to have a greater interaction with the magnetic field than oxyhemoglobin such that shifts in the concentration of oxy- and deoxyhemoglobin cause changes in the MR signal. Specifically, as the concentration of oxyhemoglobin increases in response to neural metabolic demands, the BOLD signal increases. Importantly, the BOLD signal does not convey an *absolute* value—it is only a *relative* measure. Therefore, one rarely sees attempts to compare the BOLD signal between individuals. Instead, research focuses on the location and magnitude of relative changes in BOLD during different task conditions.

There are three characteristic phases of the hemodynamic response to a neural event (Figure 13.1a).

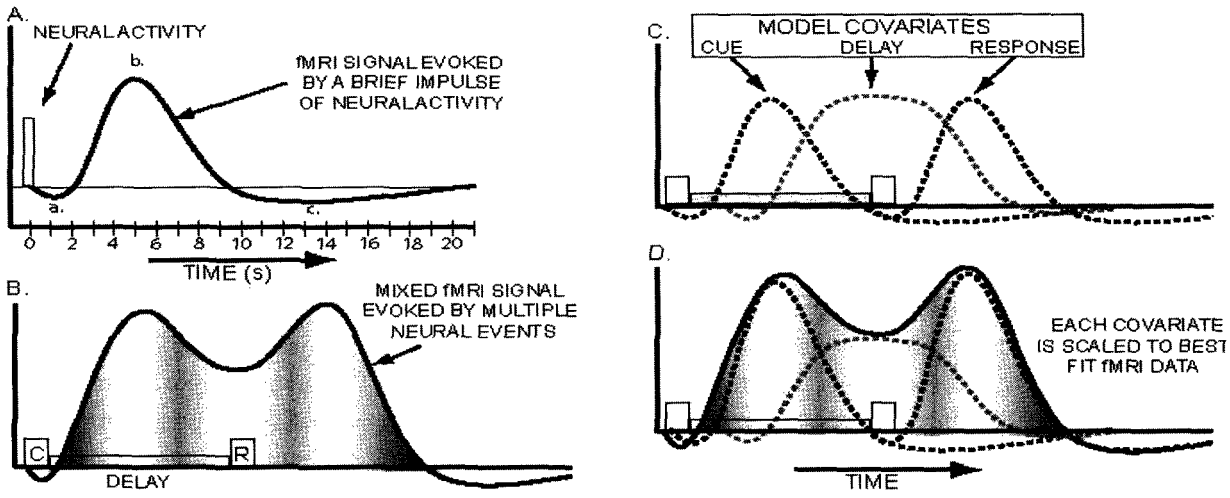

FIGURE 13.1. Modeling fMRI signals. A. In response to a single brief impulse of neural activity, the fMRI BOLD response lags the neural activity by about 5 seconds and is characterized by 3 epochs: a) the initial dip, b) the positive hemodynamic response, and c) the postresponse undershoot. B. Hypothetical neural activity during a delayed-response task, where C is a cue to be remembered, and R is the response occurring after an imposed delay. The evoked fMRI BOLD response involves a mixture of signals emanating from more than one time and more than one trial component. The gradient under the curve schematically represents the mixing or temporal overlap of the various signal components. Whiter regions reflect purer (less colinear) BOLD signal, and darker regions reflect highly colinear signal. For example, the white region at the peak of the first hump is almost exclusively evoked from neural processing during the cue phase of the task. However, just a few seconds later, in the darker portion just to the right, the signal is a mixture of processing at the cue phase and the beginning of the delay period. C. To resolve the individual components of the mixed fMRI signal, ideal hemodynamic response functions (which take into account the lag and spread of the BOLD response) are used to model within-trial components. In this case, a separate covariate is used to model the cue, delay, and response phase of the trial. D. The covariates are entered into the modified GLM of the fMRI time-series data, and a least-squares procedure is used to derive parameter estimates (i.e., beta values) that scale with the degree to which a given covariate accounts for the variance in the observed data. For example, the height of the delay covariate can be used as an index of the amount of delay-period activity.

First, in response to transient increases in neuronal oxygen consumption, the BOLD signal decreases because the ratio of oxy-/deoxyhemoglobin in blood decreases. This transient decrease has been termed the *initial dip* and is currently under increased scrutiny because it may provide greater spatial localization than subsequent responses (Ugurbil et al., 1999; Yacoub & Hu, 2001). Second, a large increase in signal above baseline is observed beginning approximately 2 seconds and peaking 4 to 6 seconds after the onset of a brief impulse of neural activity (the precise latency and time course of the response can vary depending on the individual, the brain region, or the length of the neural activation; Aguirre, Zarahn, & D'Esposito, 1998; Buckner, 1998). This increase is caused by a local increase in blow flow that actually overcompensates for the amount of oxygen consumed. Thus, the ratio of oxy-/deoxyhemoglobin increases in the vasculature near the site of neural activity. Most fMRI studies primarily focus on this positive phase. Finally, there is a decrease in signal that most often falls below baseline and can require tens of seconds to return to baseline.

Echo Planar Imaging

In a completely uniform magnetic field, the spatial location of the measured MR signal cannot be determined. However, by temporarily imposing a separate magnetic field that varies linearly across a

volume, known as a *gradient*, it becomes possible to temporarily alter the strength of the external magnetic field in a spatially specific manner. This *spatial encoding* is achieved by the application of gradient fields in three dimensions at critical times in relation to the RF pulse. Gradients affect which portions of the brain receive the RF energy (*slice selection gradient, z-direction*), the phase in which the excited nuclei are precessing (*phase encoding gradient, y-direction*), and the frequency in which the excited nuclei are precessing at the time that the emitted RF energy is reradiated (*frequency or "read-out" gradient, x-direction*). The application of these gradients allows for the transformation of an acquired free induction decay signal into an image in Cartesian space, where the matrix size of the volume is a function of the number of steps in the three gradients. For example, if an image was created using 20 *z*, 128 *y*, and 128 *x* encoding steps, it would result in a volume of 20 slices containing 128 × 128 pixels or voxels on each slice (a pixel refers to a distinguishable square of information within a two-dimensional image, whereas a voxel corresponds to a box of information in a three-dimensional image).

The type of information that is obtained by an MRI scan depends on how and when magnetic gradients and RF pulses are applied (commonly referred to as the pulse sequence). Fast echo planar imaging (EPI) is by far the most dominant pulse sequence technique used for fMRI (Buxton, 2002; Jezzard, Matthews, & Smith, 2001; Moonen & Bandettini, 1999). EPI differs from other standard imaging methods in that it acquires multislice volumes of MR images very rapidly. It does this by applying a rapid cycling phase encoding gradient where all the phase encoding steps are done in a single *repetition time* (TR) after a single RF pulse. In contrast, more traditional structural MRI techniques apply a single phase encoding gradient step per TR. Depending on the slice thickness and matrix size, fMRI studies using EPI may obtain whole brain coverage every 1 to 3 seconds, (TR = 1–3 s). A number of alternative pulse sequences, such as spin echo or spiral sequences, can be used to detect changes in BOLD signal (Haacke et al., 1999; Noll, Cohen, Meyer, & Schneider, 1995). These techniques vary in terms of aspects of the RF pulse or the ordering of gradient steps and have both advantages and disadvantages relative to EPI (Kennan, 1999; Noll, Stenger, Vazquez, & Peltier, 1999).

FMRI EXPERIMENTAL DESIGN AND ANALYSIS

Because images can be collected at the level of seconds (TR for a single slice can be < 1 s, whole brain coverage < 3 s), it becomes possible to collect hundreds of images consecutively, with the primary limiting factors being subject to fatigue or movement over time, or hardware processing constraints. This allows a wide range of study designs. The prototypical fMRI experimental design involves a "boxcar" in which two behavioral tasks alternate over the course of a scanning session, and the fMRI signal between the two tasks or between a task and a resting condition is compared. In the most typical application of this *block* design, subjects will perform multiple trials of the stimulation (i.e., experimental) task (say for 20 seconds) and then multiple trials of the control task (say for the next 20 seconds), and these conditions will repeatedly alternate over time. The primary analysis essentially involves a subtraction in which one condition is subtracted from the other.

Event-related designs provide the primary alternative to the block design (Buckner et al., 1996; D'Esposito, Zarahn, & Aguirre, 1999). In these studies, individual trials are treated as discrete events, rather than being grouped together as a block of trial. The trials can either be performed in a temporally discrete manner, such that the hemodynamic response is allowed to return to baseline between each trial, or trials can be performed in a manner in which the hemodynamic responses temporally overlap, but are separated enough that the responses can be modeled in relation to a reference function. If responses have significant temporal overlap, as is the case with rapid event-related designs, successful estimation of the evoked hemodynamic responses rely on random presentation of stimuli (i.e., trial Type A is followed by Type B as often as B is followed by A) and highly jittered intertrial-interval durations (Buckner et al., 1996).

Block designs have an advantage over event-related designs in that they provide strong signal detection characteristics over relatively brief times (a single functional scan on the level of 4–7 minutes is often sufficient to detect a substantial BOLD change) (Liu, Frank, Wong, & Buxton, 2001). However, the interpretational power of this design is limited because it cannot disambiguate differential contributions of events occurring within a block or trial (see Figure 13.1b). As described following, event-related designs provide a far more powerful tool in separating the different components of a task.

Consider a spatial delayed response task. The task has three main epochs; a cue period where stimuli to be remembered are presented (say the location of a briefly appearing dot), an unfilled retention period where the location of the dot must be retained in memory, and finally a response period where a memory-guided response is required (say a saccade to the remembered location). In a typical block design, a control condition (not requiring maintenance but attempting to control for other sensory and motor features) is subtracted from the delayed response condition. Because the requirements of the experimental and control tasks have similar visual and motor attributes, but differ in the attribute of interest (i.e., maintenance of the location), subtracting these two blocks is reasoned to yield areas active during memory maintenance. The inferential framework of *cognitive subtraction* attributes differences in neural activity between the two tasks to the specific cognitive process (i.e., maintenance; Friston et al., 1996; Posner, Petersen, Fox, & Raichle, 1988). However, the assumptions required for this method may not always hold (Zarahn, Aguirre, & D'Esposito, 1999) and could produce erroneous interpretation of functional neuroimaging data. Cognitive subtraction relies on the assumption of *pure insertion*—that a cognitive process can be added to a preexisting set of cognitive processes without altering the other processes. If pure insertion fails as an assumption, then a difference in the BOLD signal between the two tasks might be observed, not because a specific cognitive process was engaged in one block and not the other, but because the added cognitive process and the preexisting cognitive processes interact. Continuing with our delayed-response example, the insertion of a maintenance requirement may directly impact the other encoding and retrieval/response processes (e.g., visual encoding; why encode the cue if it will not be used to guide the response made after the delay?). The result is a failure to meet the assumption of cognitive subtraction. Thus, inferences drawn from the results of such blocked experiments may fail to specifically isolate maintenance-related activity.

Event-related designs allow researchers to statistically disambiguate the hemodynamic signals specifically related to encoding the cue stimulus and generating memory-guided responses from the maintenance-related activity present in the retention interval (Aguirre & D'Esposito, 1999). Event-related designs model each component of the trial independently (e.g., cue, delay, and response; see Figures 13.1c and 13.1d). Task designs are often complicated due to the sluggish hemodynamic response, but are feasible as long as different components of the task are temporally varied in relation to each other so that separate aspects of the task can be modeled. Such designs allow separate identification of brain regions involved in encoding spatial locations, maintaining that information across the retention interval, and making the memory-guided response. The ability to model maintenance separately from other task components thus makes it possible to avoid assumptions of pure insertion.

Image Preprocessing

Analysis of fMRI data (or PET data) is almost never performed without first preprocessing the raw data. These preprocessing steps variably include *temporal filtering* to remove signal jitter across adjacent scans, removal of linear trends in signal intensity, *spatial filtering* (also called spatial smoothing), filtering to remove sources of periodic signal fluctuation related to vascular pulsation or breathing, *intrasubject spatial alignment* to remove movement across scans, and *coregistration* of BOLD data to the subject's structural MRI to allow visualization of the images. Finally, *intersubject alignment* and *warping* (resizing) to a common stereotactic space are frequently performed. This final stage allows group statistical analyses on a voxelwise basis, but comes

at the cost of spatial resolution and an understanding of individual variability (Brett, Johnsrude, & Owen, 2002).

Statistical Analysis

A wide range of techniques has been applied to look at changes in brain activity (Lange, 1999), with most using some variant of the *general linear model* (GLM; Friston et al., 1994). The change in BOLD signal intensity over time represents the dependent variable in fMRI studies. Typically, a reference time series is created that denotes what type of event and when it happened during a scanning session. This time series is then convolved with a standard or empirically derived hemodynamic response function (which incorporates the sluggish nature of the hemodynamic response) yielding a suitable estimate or model of the predicted BOLD signal. Each of the independent variables (e.g., different types of task events) is represented by a set of covariates that are shaped like hemodynamic responses and are shifted in time to account for the lag. These can be either categorical variables (such as presence or absence of a task demand) or quantitative variables (such as number of stimuli presented at a time). The covariates are entered into the modified GLM with the fMRI time-series data, and a least-squares procedure is used to derive parameter estimates (i.e., beta values) that scale with the degree to which a given covariate accounts for the variance in the observed data. Similar to traditional statistical methods, these parameter estimates, when normalized by estimates of noise, are used to compute inferential statistics such as *t* values or F-statistics. These inferential statistics are calculated on a voxel-by-voxel (voxelwise) basis to create statistical parametric brain maps.

Some researchers alternatively perform analyses based on structurally defined regions of interests. This can have advantages when investigators have a specific hypothesis about a specific brain region. However, most investigators prefer a voxelwise approach because it is not constrained by preconceived ideas regarding the volume or location of expected activations. The primary drawback with the voxelwise approach involves the large number of voxels in the brain, causing a high risk of Type I statistical error. Thus one needs to perform an adjustment for the number of independent comparisons in each analysis. Because neighboring voxels are correlated and there exists temporal autocorrelation over time, it is overconservative to apply a simple Bonferroni correction to these data sets. Instead, investigators typically apply corrections based on an estimate of the number of independent resolution elements (RESELs) or adjust the degrees of freedom to account for the nonuniformity in the noise. For instance, an estimate may be made for the number of independent spatial resolution elements by correcting the total size of the volume of interest by the Full-Width at Half-Maximum estimate of spatial resolution (Worsley et al., 1996).

In addition to looking at individual activations, increased attention is being paid to the functional relationships between different brain regions (Mesulam, 1990). Because most psychological phenomena are not mediated by single brain regions, but instead involve networks of brain regions, it becomes essential to understand how these brain regions interact, when their activity is functionally coupled or uncoupled, and the extent to which these changes in functional connectivity are related to experimental variables of interest. Toward this end, researchers have used a number of strategies, ranging from correlation analysis to principle components analysis and structural equation modeling (McIntosh, 1999).

PET PHYSICS AND PHYSIOLOGY

Functional neuroimaging with PET predates the development of fMRI. PET imaging takes advantage of the fact that unstable elements (such as $^{15}0$, ^{11}C, or ^{18}F, which possess too few neutrons relative to protons) go through a rapid process of decay involving the release of a positron (positively charged electron) from the nucleus. Once released, the positron collides with an electron, which causes the annihilation of both the electron and the positron and the production of two high-energy (511 keV) photons that travel at 180° from each other (see Figure 13.2). PET cameras consist of rings of crystals that produce light scintillation

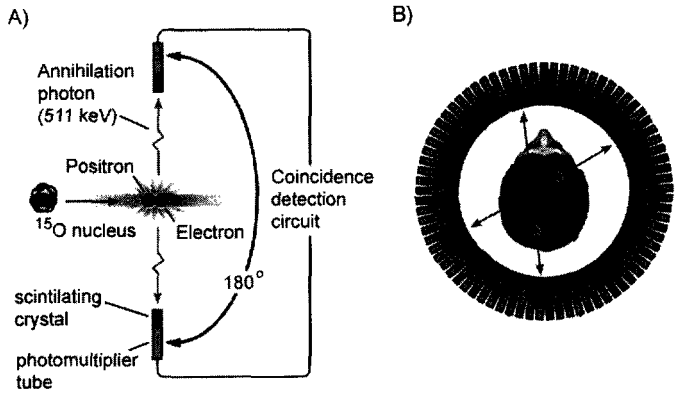

FIGURE 13.2. Measurement of positron emissions. A) An unstable $^{15}0$ nucleus emits a positron that collides with an electron, releasing a pair of high-energy 511 keV photons at a 180° angle. The photons are detected by an annihilation photon detector, which is comprised of crystals that scintillate when struck by a photon and photomultiplier tubes that transform the light emitted by the crystals into an electrical impulse. When two detectors 180° apart are activated, the coincidence is registered and sent on for signal processing and image reconstruction. B) Rings of annihilation photon detectors are arrayed around a subject's head. PET scanners have multiple rings arrayed in parallel, allowing multislice data collection. Detection of coincident scintillations at 180° angles within a ring (in 2-D imaging) or across rings (in 3-D imaging) allows identification of the approximate location from which the positron emitting radiotracer is located. Figure from pages 62 and 63 of *Images of Mind* by Michael I. Posner and Marcus E. Raichle. Copyright 1994, 1997 by Scientific American Library. Reprinted by permission of Henry Holt and Company, LLC.

when penetrated by photons (Raichle, 1983). This scintillation is then converted to electrical impulses that can be amplified and analyzed. Within the crystals' range of sensitivity, there exists a direct relationship between the concentration of radiotracer present in a brain region and the level of photon detections arising from that region.

PET allows assessment of multiple aspects of brain functioning depending on the radiotracer used. Importantly, PET can be used to measure regional cerebral blood flow (rCBF). Recall that when neurons in a brain region become active, they increase their oxygen consumption, which is compensated for by a substantial increase in rCBF (Figure 13.1a). This increase in blood flow exceeds the oxygen consumption demanded by the neurons, making rCBF a particularly robust index of regional neural activity (Fox & Raichle, 1986; Fox, Raichle, Mintun, & Dence, 1988). When unstable $^{15}0$ is attached to H_2, it can be injected directly into the bloodstream. Once in the bloodstream, $H_2{}^{15}0$ will travel wherever the blood travels, such that areas with the highest levels of rCBF will emit the most positrons (Herscovitch, Markham, & Raichle, 1983). Thus, by measuring $^{15}0$ positron emissions, we can index neural activity. Indeed, the measurement of rCBF with ^{15}O PET represents a far more simple and direct index of neural activity than the BOLD response, which is influenced by several different features of the hemodynamic response (i.e., blood volume, flow rate, and oxyhemoglobin–deoxyhemoglobin ratios). The directness of the

relationship also makes PET less sensitive to some of the artifacts associated with fMRI discussed at greater length later.

The temporal window measured in PET studies is directly linked to the speed at which the radiotracer decays. To get adequate signal-to-noise ratios, the detected positron annihilations are aggregated over time. With ^{15}O, which decays rapidly, one typically scans for 30 seconds to 90 seconds to achieve adequate signal-to-noise ratios. The data from these scans thus represents the aggregate of activity during this window, with the largest weighting occurring earlier in the scan when positron emissions are highest (Silbersweig et al., 1993). Therefore, the minimum temporal resolution of ^{15}O PET is on the level of about 30 seconds. In cognitive studies, this dictates that tasks need to Engage Brain regions for a substantial portion of a 30-second to 90-second scan window if they are to produce robust changes in rCBF measurements.

In addition to measuring rCBF, PET can also be used to measure glucose metabolism in the brain, which provides an even more direct index of neural activity. Glucose metabolism is assessed by labeling a deoxygenated form of glucose with ^{18}F. ^{18}F-deoxyglucose (FDG) is injected into the bloodstream, and the FDG is taken up by brain regions in direct proportion to their metabolic demands (Raichle, 1988). In contrast to ^{15}O, the slower decay of ^{18}F requires imaging over substantially longer temporal windows, requiring tasks to be carried on for 20 minutes or longer. Because of this, ^{15}O provides the primary tool for studying brain activations, whereas FDG is more frequently used to make baseline (resting) comparisons between different subject populations.

Radiotracers can also be created by tagging ligands or precursors for various neurotransmitter systems with ^{18}F or ^{11}C (Fowler, Ding, & Volkow, 2003). These radiotracers allow for the assessment of many of the major neurotransmitters systems, providing the ability to detect individual or group differences in neuroreceptor density, transporter density, and even neurotransmitter synthesis (see Table 13.1). This has proved highly useful both for research and, in some cases, clinical diagnosis such as assessment of the dopaminergic system in patients with Parkinson's disease (Kaasinen & Rinne, 2002).

TABLE 13.1

PET Tracers

Tracer	System measured
^{15}O (half life = 2.1 min.)	
$H_2{}^{15}O$	Blood flow/oxygen extraction fraction
^{15}O–CO_2	Blood flow
^{18}F (half life = 109 min.)	
deoxyglucose	Glucose metabolism
Dopa	Dopamine synthesis
Fallypride	Extrastriatal D2 receptors
FLB 457	Extrastriatal D2 receptors
CFT	Dopamine transporter
spiperone	5HT2a receptors
altanserin	5HT2a receptors
setoperone	5HT2a receptors
^{11}C (half life = 20.4 min.)	
raclopride	Striatal D2 receptors
N-methylspiperone	Striatal D2 receptors
cocaine	Dopamine transporter
altropane	Dopamine transporter
SCH23390	Striatal D1 receptors
carfentanil	Mu Opioid receptors
diprenorphine	Opioid receptors
flumazenil	Benzodiazipine receptors
WAY100635	5HT1a receptors
spiperone	5HT2a receptors
MDL100907	5HT2a receptors
McN 5652	Serotonin transporter
ketamine	NMDA receptors

The binding potential (level at which ligands bind) is affected by factors such as synaptic competition, receptor internalization, and changes in affinity states of receptors following stimulation (Laruelle & Huang, 2001). This allows PET imaging to assess the effects of medications on the functioning of different neurotransmitter systems and has provided insights into the occupancy levels necessary to achieve therapeutic effects (Fowler et al., 1999). Moreover, in some cases it has become possible to examine the degree to which behavioral tasks cause endogenous release of neurotransmitters (Koepp et al., 1998; Zald et al., 2004; Zubieta et al., 2001).

PET EXPERIMENTAL DESIGN AND DATA ANALYSIS

The range of experimental design strategies for PET is limited by the physical constraints associated with delivering and measuring radioisotopes. At a minimum, ^{15}O studies involve an aggregation of activity occurring over 30 seconds or more, and for metabolism and neurotransmitter studies the aggregation covers 20 to 60 minutes. This temporal resolution precludes event-related types of designs and makes it difficult to dissociate the different processes involved in a task. The total number of radiotracer injections (and hence scans) is limited by radiation exposure and the need to allow previously administered radiotracers to decay substantially before starting the next scan (this takes hours for ^{11}C and ^{18}F and about 8 to 10 minutes for $^{15}0$). With $^{15}0$, one is typically limited to about 12 scans in a 2-hour scanning session, and with ^{18}F and ^{11}C, one is typically limited to 2 to 4 scans (usually scheduled on different days because the subjects would need to spend hours waiting for the isotope to decay between scans). However, because the data in each scan is an aggregate of activity over time, a single contrast between two PET scans can be informative, whereas a contrast between two individual BOLD images (one phase each) has little value.

The most common PET analysis involves a simple subtraction paradigm. As with block designs in fMRI, these analyses are efficient and straightforward, but often depend on the problematic assumption of pure insertion. Parametric designs, where a variable is quantitatively manipulated across different conditions, and factorial designs are also frequently implemented. All of the preceding designs can be analyzed within the framework of the general linear model (Friston et al., 1995). Studies examining the covariance of activity across regions and application of techniques to assess functional connectivity are also possible, although the statistical power to apply such techniques is often restricted by the limited number of scans (Friston, Frith, Liddle, & Frackowiak, 1993; Zald, Donndelinger, & Pardo, 1998).

CONVERGENCE ANALYSIS TO DEAL WITH METHOD VARIANCE WITHIN NEUROIMAGING STUDIES

Although rarely attributed to Campbell and Fiske's pioneering descriptions of the MTMM approach, the core ideas inherent in the MTMM concept can be seen in the convergence approaches that are used in many neuroimaging studies. Convergence approaches have become increasingly popular in functional neuroimaging in response to one of the core problems in the field. Specifically, a multitude of stimulation tasks or procedures exist that can be used to engage a particular psychological construct or brain region, with each variation possessing slightly different properties. In other words, each stimulation paradigm comes with its own method variance, and not surprisingly, substantial inconsistencies emerge in the literature. To deal with this, many neuroimaging researchers have begun to use procedures to look at the convergence of responses across procedures. In its simplest form, this is accomplished with a simple logistic analysis in which each effect is transformed voxel by voxel into a binary representation of whether the voxel was activated above a certain threshold. These binary representations are then summed or multiplied across contrasts to produce a spatial map of areas activated in more than one condition.

A convergence approach also helps deal with the problem of pure insertion. As noted earlier, in a simple subtraction design it is impossible to determine if a change in brain activity relates to the inserted cognitive component or to changes in other components that arise as a consequence of the inserted component. However, by using multiple stimulation–control contrasts it becomes possible to more clearly parse the component in question from its effects on other task components. Imagine, for instance, a judgment task in a given sensory modality that is contrasted with a passive task in which the stimulus is presented but no judgment is made. It is difficult to know if changes in brain activity are related to the judgment itself or if the act of making the judgment caused modality-

specific changes in sensory processing because of increased attention to the stimulus rather than the act of making the judgment. Now, if we run similar experiments in other sensory modalities, we can analyze them to determine common vs. modality-specific activations. The areas that are active in all tasks can be considered modality independent processes and cannot be attributed to factors such as increased attention to a specific stimulus category. Thus, even if the assumption of pure insertion fails in a given task, it becomes possible to separate activations related to the component of interest (the judgment) from changes in other processes (modality-specific attention) that arise as a consequence of the task insertion. Of course, a delineation of the common activations may fail to detect sensory-specific processes that are directly related to the component in question. Nevertheless, the remaining modality insensitive, common regions of activation will be more clearly attributable to the component of interest.

Price and Friston (1997) referred to this approach of examining the commonalities between activations arising in different contrasts as "cognitive conjunction analysis." It is worth noting that when applied in neuroimaging, particularly among researchers using the popular SPM program (Department of Cognitive Neurology, London, UK), the conjunction refers to the presence of a main effect in the absence of differences in simple effects at a given voxel. The analysis is performed by taking the sum of all activations [(stimulation$_1$ – control$_1$) + (stimulation$_2$ – control$_2$) . . .] and eliminating voxels where there exist significant differences among the individual contrasts [(stimulation$_1$ – control$_1$) – (stimulation$_2$ – control$_2$) . . .].

Convergence and Divergence Across Brain Regions

Because neuroimaging experiments provide data on multiple brain regions simultaneously, neuroimaging data need not be limited to a single entry in an MTMM matrix. Rather, different brain regions can be sampled to examine the extent to which activity converges or diverges across brain regions. For instance, in considering a measure related to attention, it may be useful to know that task performance correlates with activity within the frontal eye field, parietal cortex, and anterior cingulate (all areas involved in attention), but not with activity in the temporal lobe or Broca's area (which is not a component of the system). In this situation, the different brain regions can be equated with different traits in the MTMM matrix. The question becomes, Do anatomically connected or functionally related brain regions (i.e., related traits) show convergence, whereas functionally or anatomically unconnected (i.e., unrelated traits) show divergence? Applying a network approach, one can treat functional couplings (covariance) between regions as separate traits. We can then ask whether different tasks produce convergent or divergent effects on the functional connectivity between regions.

Considered in this framework, neuroimaging is highly compatible with the MTMM approach, allowing the assessment of convergence and divergence across stimulation methods and the brain regions activated by those methods. However, this approach is rarely formally applied in the neuroimaging field. This in part reflects the difficulty in ascribing brain activations in a given region to a specific function. For instance, although the dorsolateral prefrontal cortex frequently activates during working memory tasks, it also activates during tasks that are not specifically related to working memory (D'Esposito, Ballard, Aguirre, & Zarahn, 1998). Indeed, the multitude of functions proposed for the prefrontal cortex makes it unlikely that a single discrete process can explain all the varied tasks that lead to increased activity in the region (Duncan & Owen, 2000). Thus, it would be unwise to assume that activation of the dorsolateral prefrontal cortex (or other brain regions involved in working memory tasks) necessarily indicate the involvement of working memory in a given task. On the other hand, if we have three areas, each of which are engaged by multiple cognitive tasks, but that only show simultaneous activation during working memory, then the multiregion approach could prove highly useful.

Integrating Neuroimaging Data With Other Data in a Multitrait–Multimethod Framework

Assessment of psychological constructs has traditionally focused on behaviors that are either directly

observable by a researcher or can be reported by the examinee. Neuroimaging can supplement these methods of assessment by providing information at a neural level. Although, one might be tempted to view this at a causal level (i.e., the brain activity causes the behavior, or the behavior causes the brain activity), it need not be viewed as such. Rather, neuroimaging data can be viewed as just another indicator or correlate of a psychological process or trait. However, neuroimaging data are qualitatively different from most other types of measures in psychological research in that the brain's response can be measured without requiring the subject to make a behavioral response or use introspection. Thus response may be measured uncontaminated by requirements to self-monitor or control a motor act (both of which may add method variance in psychological studies).

Imagine, for instance, the assessment of a personality trait. A number of investigators have found neural correlates of personality either in terms of resting data or the degree of activation during stimulation (see Canli, Sivers, Whitfield, Gotlib, & Gabrieli, 2002; Gusnard et al., 2003; Zald et al., 2004). By combining neuroimaging data with other self-report, observer rating, or experimental performance measures, we may increase accuracy in assessment. In such a paradigm, levels of regional brain activity would be predicted to converge with self-report and objective ratings of the trait of interest, but not other traits.

The MTMM approach can similarly be applied to the assessment of a psychological process. Imagine you are testing subliminal processing of visual stimuli using a tachistoscopic method. The presence of subliminal processing is traditionally tested by having subjects "guess" about stimulus features in the absence of an explicit awareness of having seen the stimulus. Performance significantly above chance provides evidence for subliminal processing. Now, if we simultaneously scan subjects with fMRI and see BOLD responses that are temporally linked to the presentation of the stimuli, we could use the fMRI data as a second source of evidence that subliminal processing occurred. Because neither measure is likely to be 100% sensitive or selective, the combination of the two types of data may dramatically increase predictive power.

A critical problem must be resolved before including functional imaging data in a MTMM matrix. Specifically, the precise relationship between activations and behavioral performance cannot always be predicted in advance. In some cases, higher activations may reflect greater performance or ability level. However, in some cases, subjects with lower ability may have to activate a region more to perform a task at an equivalent level to a more skilled person. This issue has been particularly salient in the psychiatric imaging literature, where researchers attempt to draw conclusions about the relationship between functional activations and the neural substrates of psychiatric conditions. This is essentially an empirical question. Once we understand the nature of performance-activation relationships, it becomes reasonable to consider the neuroimaging data in a MTMM matrix.

Unfortunately, because of the expense of collecting neuroimaging data, it seems unlikely that neuroimaging data will be routinely used as a component in MTMM matrices. However, its utility may be appraised in terms of a cost–benefit analysis. In situations where the neuroimaging data has significantly greater sensitivity or selectivity than other forms of data, then the benefit of its inclusion may outweigh the costs. Plus, with the advent of data sharing through the fMRI data center (http://www.fmridc.org), which is a public repository of peer-reviewed published fMRI data, researchers can potentially pool data from dozens of studies that fit key cells in MTMM matrices.

Simultaneous Measurement of Other Variables to Enhance Understanding of Neuroimaging Data

Although the preceding discussions have focused on the ability of neuroimaging data to provide information on psychological constructs, a multimethod approach can also prove extremely useful in directing the interpretation of neuroimaging data. Group statistical analyses often proceed on the assumption that all subjects performed a cognitive task in a similar manner or responded similarly to procedures aimed at inducing a specific psychological state.

Unfortunately, verification of this assumption is often difficult. For instance, if we wish to study fear, it is important that we verify that we indeed induced fear and not disgust or other negative emotions. If we lack certainty that the intended state was provoked, then we cannot confidently assume that the brain responses occurred in relationship to the cognitive process or psychological state in question. The solution to this problem is to triangulate on the desired response using multiple methods, including, for example, measurement of task performance, self-report, and psychophysiological recording. As convergent evidence verifies the induction of the intended process or state (and not an unintended state), confidence in interpreting brain responses increases. This triangulation strategy is an example of the multilevel analytic approach described in the preceding chapter by Berntson and Cacioppo, in which information from different levels is used to mutually tune and calibrate data or concepts across different levels of analysis.

Unfortunately, a problem arises in trying to integrate fMRI data with simultaneous collection of other types of data. Specifically, fMRI is both sensitive to artifacts caused by psychophysiological recording devices and causes interference in those same devices. Nevertheless, it is possible to implement psychophysiological recordings such as galvanic skin response, heart rate, blood pressure, and eye tracking within the fMRI environment (Savoy, Ravicz, & Gollub, 1999). These measures are all easily implemented in the PET environment as well. Similarly, measures of hormonal responses such as cortisol can be collected in the scanner environment. The large differences in time scales of these various measures can cause interpretational issues when moving across levels. Nevertheless, the benefit of collecting such measures should be increasingly apparent.

METHOD VARIANCE AND NEUROIMAGING

Neuroimaging researchers have often highlighted sources of method variance associated with specific technical steps in neuroimaging, such as the effects of using different techniques for spatial normalization, movement correction, or modeling the hemodynamic response. In contrast, because neuroimaging data has usually been collected in relative isolation, much less attention has been paid to overall sources of method variance when attempting to include neuroimaging data as part of a larger multimethod approach. In such a context, attention to the temporal, spatial, and other methodological limitations of PET and fMRI become paramount. Because these limitations substantially influence both the level of noise in the data and the ability to detect relevant activations, they will directly influence the utility of including neuroimaging in a MTMM matrix. The following section describes six important sources of method variance in neuroimaging studies: (a) temporal resolution, (b) the nature and source of the signal change, (c) spatial resolution, (d) anatomical variability, (e) imaging artifacts, and (f) influences on functional activations unrelated to brain processes.

Temporal Resolution

Because of the nature of both radioisotope decay and the slow time course of the hemodynamic response, the temporal resolution of neuroimaging is inherently limited. The sluggish nature of the hemodynamic response prohibits the detection of numerous events that occur on a millisecond time scale and may be conceptualized as a low-pass filter that prevents detection of higher frequency information. The temporal resolution is particularly poor for PET, which is largely insensitive to transient responses unless they are sustained or of large magnitude. The different temporal limitations of PET and fMRI almost certainly lead to situations in which the results of the two techniques disagree with each other, leading to nontrivial differences in conclusions (Zald, 2003). Similarly, both techniques may fail to converge with data from techniques such as single-cell recordings, event-related potentials, and near infrared optical imaging, which are sensitive to changes at the millisecond level.

The Nature and Source of the Signal Change

Neuroimaging studies measure changes in signal magnitude. However, many processes in the brain

may be characterized by changes in firing patterns or synchronization among neurons, rather than changes in overall firing rates (Lestienne, 2001; Neuenschwander, Castelo-Branco, Baron, & Singer, 2002). Neuroimaging studies will often be insensitive to such changes.

Equally important is a consideration of the source of rCBF/BOLD changes, which in addition to the neuronal output signals are also significantly associated with input to a region and local processing within the activated region (Logothetis, 2002). Thus, when an area shows increased rCBF or BOLD signal, the finding may not directly inform us about the region's output. This differs from many electrophysiological techniques that solely examine a brain region's output.

Spatial Resolution

The inherent resolution of high-quality, commercially available PET scanners is around 4 to 7 mm (full-width half-maximum; DeGrado et al., 1994; Spinks et al., 2000). This is high enough to measure activity in most cortical and subcortical regions, but limits the ability to look at subnuclei and often leads to difficulties in determining the exact origin of foci that occur near the boundaries between regions. By comparison, fMRI is capable of higher spatial resolution. However, many fMRI studies are performed with parameters that provide no higher spatial resolution than that produced by high-quality PET cameras. Moreover, draining vein effects often lead to mislocalization of the source of fMRI signals (Lai et al., 1993), thus lowering the effective resolution for localizing responses. Both fMRI and PET images are usually filtered to a lower spatial resolution after the data is collected. This filtering serves several purposes. First, it reduces noise and hence improves signal to noise characteristics. Second, it lowers the number of resolution elements and hence reduces correction factors for multiple comparisons. Third, it improves the detectability of large-volume activations (Poline & Mazoyer, 1994). However, this comes at the cost of restricting the ability to detect more discrete focal activations and therefore biases the methods toward detections of large-volume activations. This bias is at its most extreme in older PET studies, but remains a bias in the fMRI literature as well.

Anatomic Variability

Imaging data is also smoothed to remove minor differences in anatomical variability across subjects. Anatomical variability is a constant issue faced in neuroimaging. Different warping algorithms and landmark systems have been proposed to optimize coregistration in different regions of the brain. It is clear, though, that humans show variability in both the structural and functional topography of the brain that cannot be overcome by coregistration. Most researchers naturally focus on group analyses to report common areas of activation. However, this approach fails to capture more idiosyncratic activations. Moreover, in regions of high structural variability and in tasks that localize to variable locations, the group analysis may fail to detect relevant activations. For instance, portions of the fusiform gyrus are responsive to faces, but the precise location varies across subjects (Kanwisher, McDermott, & Chun, 1997). In such a situation, group analyses could easily fail to detect the presence of this region.

Imaging Artifacts

All neuroimaging techniques are sensitive to artifacts associated with data collection that can appear as either false positives (signal change unrelated to brain activity) or false negatives (failure to detect real changes in brain activity). With PET, the most significant artifacts are associated with subject movement and variability in the timing or amount of radiotracer delivery. Many of these can be measured and adjusted for, but they may nevertheless impact the quality of PET results.

Functional MRI is more prone than PET to artifacts, with even small movements producing large changes in signal within individual voxels. Signal changes caused by periodic motion from breathing or cardiac pulsation can similarly hinder detection of changes in brain activation (Hu, Le, Parrish, & Erhard, 1995). Indeed, in many cases artifactual changes in MRI signals are substantially larger than the changes in BOLD signal associated with neural activation. Moreover, certain areas of the brain are

extremely difficult to measure with fMRI because of signal dropout caused by boundaries between brain tissue and air (Farzaneh, Riederer, & Pelc, 1990). Because of this signal dropout, PET can detect changes in certain brain regions where many fMRI studies will produce false-negative results (particularly in ventromedial frontal and anteriormedial temporal regions). Many techniques exist to address these problems, but the quality of the data must be considered on a case-specific basis, especially when considering negative findings.

Influences on Functional Activations Unrelated to Brain Processes

When we see individual or group differences in the magnitude of activations in brain-imaging studies, we frequently assume that these differences arise from differences in the level of brain activity in a given region. However, this assumption can be problematic. For instance, in fMRI, the magnitude of the BOLD response to visual stimulation is associated with levels of hematocrit in the blood (Levin et al., 2001). Because individuals differ in hematocrit levels, and men have higher overall hematocrit levels than women, these differences can easily confound interpretation of differences in BOLD magnitude. Attention to such variables becomes especially important if functional activations are going to be used as an assessment measure.

PSYCHOMETRIC PROPERTIES OF NEUROIMAGING DATA

The selectivity, sensitivity, criterion validity, and test–retest reliability can be calculated for both PET and MRI studies. Establishing the test–retest reliability of most baseline PET measures is relatively straightforward (Ball, Fox, Herscovitch, & Raichle, 1988; Nyberg, Farde, & Halldin, 1996; Schmidt et al., 1996), although this literature remains surprisingly small considering the increasing use of these measures in clinical diagnosis. Establishing the test–retest reliability of activations caused by stimulation paradigms is a trickier issue. Reliability in these paradigms will always be task and region specific, making it impossible to make generalizable statements about reliability. Nevertheless, there are increasing attempts to define the test–retest reliability of the activations associated with specific cognitive and motor tasks (Fernandez et al., 2003; Kiehl & Liddle, 2003; Maitra, Roys, & Gullapalli, 2002; Specht, Willmes, Shah, & Jancke, 2003). This issue has proved particularly important when fMRI is used as part of presurgical planning for intractable epilepsy. Obviously, a neurosurgeon needs to know which measures (neuropsychological data, WADA procedure, etc.) provide the most valid and reliable information about functional localization of cognitive tasks (particularly language tasks) before choosing to remove part of a patient's cortex. However, determination of the psychometric properties of neuroimaging data is complicated by the fact that the data sets include information on magnitude of change (or the degree of temporal correlation) and location of activation. For instance, imagine performing a receptive language study on a patient on two occasions. In both cases the subject demonstrates activation in the left superior temporal gyrus, but the emerging foci, although within 5 mm of each other, do not overlap. Depending on one's criteria, this could be viewed as a replication or a failure to replicate. When viewed loosely (for instance, in terms of hemispheric asymmetries within the temporal or frontal lobe), such tasks have typically shown good reliability (Fernandez et al., 2003; Rutten, Ramsey, van Rijen, & van Veelen, 2002). In contrast, when viewed on a voxelwise basis, the overlap between activations across sessions tends to be much lower (Fernandez et al., 2003).

Attempts to use functional neuroimaging for diagnoses have also provided information regarding the sensitivity and selectivity of this information. This has received particularly strong attention in the diagnosis of early Alzheimer's disease (Petrella, Coleman, & Doraiswamy, 2003). Research along similar lines will clearly need to be performed if functional neuroimaging is to reach its full potential as an assessment tool, regardless of whether it is used in isolation or as part of a MTMM matrix.

IMPLICATIONS BEYOND THE FORMAL MTMM APPROACH

The examples given in earlier sections of this chapter have described how neuroimaging data can be

used within a formal application of the MTMM approach. However, the general approach toward looking for convergence and discrepancies across methods and traits can be applied as an evaluative strategy, even in situations where it is not possible to use the same methods in the same subjects. In such a situation one cannot produce a covariance matrix across methods, but one can nevertheless use an emphasis on convergence and divergence for evaluating hypotheses.

Sarter, Cacioppo, Berntson, and colleagues (Cacioppo et al., 2003; Sarter, Berntson, & Cacioppo, 1996) have articulated the importance of understanding the type of information that functional neuroimaging studies provide relative to other types of neuroscientific data. Specifically, most neuroimaging studies provide information on the probability that a given brain area activates as a function of a cognitive process (i.e., the experimenter performs a task aimed at inducing a specific cognitive process and determines whether the task leads to activity in a specific brain region). In contrast, such studies do not typically provide information on the probability that a given cognitive process arises as a function of activation of a specific brain region (although researchers frequently make the erroneous interpretation that the results provide this information). Such a conclusion would only be true if there is a one-to-one correspondence between the brain region's activity and the cognitive process, and we rarely possess evidence for such a one-to-one correspondence. Sarter et al. argue that to fully understand the bidirectional relationship between brain activity and cognitive processes, one needs to integrate other types of paradigms (such as lesion or electrical stimulation data) that allow direct manipulation of brain regions and thus provide information on the probability of a cognitive process given activity (or lack of activity) within a specific brain region.

The preceding analysis parallels a classic distinction in the neurobehavioral field between brain areas that are activated in a task and brain areas that are necessary for performance of the task. Taken alone, neuroimaging typically only addresses the question of what is activated and fails to address whether that activation is necessary. In contrast, neuropsychological studies of patients address what is necessary, but not what is activated. Thus, to answer the question of what is both engaged and necessary in a task, one needs to use both methods. The greatest clarity arises when both methods converge to show that an area is both necessary for and engaged by a task involving a given psychological process, but is not necessary or engaged by tasks that do not require that psychological process.

Considered in this light, it also becomes necessary to expand the MTMM approach to include data from other species. Specifically, most techniques that allow us to look at the causative effects of manipulating brain regions can only ethically be carried out in nonhuman populations. These animal studies typically proceed on the assumption that (a) there are "homologous" brain regions across species, (b) these regions perform the same tasks, and (c) the regions perform the tasks in the same way. However, despite many features that are conserved across species, even a cursory study of neuroanatomy reveals substantial interspecies differences. Given these potential cross-species differences, we need evidence of convergence and divergence across methods used in different species. It thus may prove useful to take a multitrait–multimethod–multispecies approach to evaluating brain–behavior relationships. In summary, the core logic articulated by Campbell and Fiske provides an extremely useful overall strategy for placing neuroimaging research within the larger field of psychology and neuroscience, even in situations where formal MTMM analyses are not feasible.