Papassotiropoulos, A., Stephan, D. A., Huentelman, M. J., Hoerndli, F. J., Craig, D. W., Pearson, J. V., et al. (2006). Common Kibra alleles are associated with human memory performance. *Science, 314*, 475–478.

Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences USA, 98*(2), 676–682.

Rao, H., Gillihan, S. J., Wang, J., Korczykowski, M., Sankoorikal, G. M., Kaercher, K. A., et al. (2007). Genetic variation in serotonin transporter alters resting brain function in healthy individuals. *Biological Psychiatry, 62*(6), 600–606.

Rusch, B. D., Abercrombie, H. C., Oakes, T. R., Schaefer, S. M., & Davidson, R. J. (2001). Hippocampal morphometry in depressed patients and control subjects: Relations to anxiety symptoms. *Biological Psychiatry, 50*(12), 960–964.

Sheline, Y. I., Gado, M. H., & Kraemer, H. C. (2003). Untreated depression and hippocampal volume loss. *American Journal of Psychiatry, 160*(8), 1516–1518.

Surtees, P. G., Wainwright, N. W., Willis-Owen, S. A., Luben, R., Day, N. E., & Flint, J. (2005). Social adversity, the serotonin transporter (5-HTTLPR) polymorphism and major depressive disorder. *Biological Psychiatry, 59*, 224–229.

Weaver, I. C., Cervoni, N., Champagne, F. A., D'Alessio, A. C., Sharma, S., Seckl, J. R., et al. (2004). Epigenetic programming by maternal behavior. *Nature Neuroscience, 7*(8), 847–854.

14

Functional Magnetic Resonance Imaging in the Affective and Social Neurosciences

Tom Johnstone
M. Justin Kim
Paul J. Whalen

Over the last 10 years, there has been an unprecedented expansion in brain imaging research. With the availability of magnetic resonance imaging (MRI) scanners steadily growing, and with the refinement of functional MRI (fMRI) techniques, the pool of brain imaging researchers has expanded beyond neuroscience and medical imaging departments into psychology departments throughout the world. Early adopters of fMRI in psychology departments were visual psychophysicists interested in mapping the visual processing areas of the brain. Researchers of perception, language, and cognition soon followed, often building upon the theoretical framework provided by animal researchers and neuropsychologists. In addition, affective and social neuroimagers have begun to offer data to make their case that functional brain imaging can add valu-

able insights to the existing experimental descriptions of human behavior already put forward by affective and social psychologists. Ideally, these insights will enable us to ask new questions to enrich these fields further. The purpose of this chapter, then, is to identify ways in which functional brain imaging may be usefully employed in the social and affective sciences, and to highlight both the advantages and the potential pitfalls of functional brain imaging as a technique for studying the human brain within a social and affective context.

What Does fMRI Measure?

It is generally understood that fMRI is not a measure of neural activity. It is a measure of the ratio of deoxygenated to oxygenated blood, commonly referred to as "blood-oxygenation-level-dependent" (BOLD) contrast. Through some good fortune, oxygenated blood and deoxygenated blood have slightly different magnetic properties; this difference makes it possible to detect the ratio by combining a large electromagnet and radio waves (for details of the physics of BOLD measurements, see Moonen, Bandettini, & Aguirre, 1999). What is clear is that fMRI is based on the experimentally verified assumption that changes in oxygenation in a given region are useful proxies for changes in neuronal activation in that region.

Still not entirely understood is what exactly this oxygenated-to-deoxygenated ratio reflects in terms of neural activity. Is it the activity of neurons within this very region communicating with other brain regions? Is it the activity of other regions communicating with this observed region? Or could it be the activity of neurons within this region communicating with other neurons within this same region? One line of research indicates that the BOLD signal scales roughly linearly with the local summed activity of dendrites, known as the "local field potential" (Goense & Logothetis, 2008; Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001; Logothetis & Wandell, 2004). These data suggest that BOLD contrast may represent an indirect and approximate measure of dentritic activity averaged over a certain neural region. This would mean that BOLD contrast is not a measure of neuron spiking, but instead that it reflects signals coming into a particular neural region, as well as signals being passed between neurons within the region (though not outgoing signals to other parts of the brain). As research concerning the nature of the BOLD signal continues, the best we can do as we interpret observed BOLD activations is to understand that these different possibilities exist.

Some Advantages of fMRI

The major advantage of fMRI over other techniques of measuring brain activity is its combined spatial and temporal resolution. fMRI can measure BOLD contrast with a spatial resolution of a few millimeters and with a temporal resolution of a few seconds. This compares favorably with electroencephalography (EEG), which can only measure the effects of brain activity at the scalp surface with a resolution of a few centimeters, and with positron emission tomography (PET), which has a temporal resolution of minutes. Thus fMRI is ideally suited to measurement of various brain processes. For example, fairly extensive areas of primary and secondary auditory cortex show sustained activity in response to common sounds, including speech. Large visual processing areas in the occipital cortex and adjacent parts of the parietal and temporal lobes show sustained activity to different types of visual stimuli. Where the temporal and spatial resolution of fMRI comes in handy is in more fine-grained distinctions between subregions that specialize in one particular type of visual or auditory processing. For example, Epstein and Kanwisher (1998) and O'Craven and Kanwisher (2000) have found that distinct parts of the inferior occipital and temporal lobes are involved in processing information about faces and people, as opposed to places. Similar regions in auditory cortex have been identified that respond selectively to human speech sounds (Belin, Zatorre, & Ahad, 2002; Belin, Zatorre, Lafaille, Ahad, & Pike, 2000).

Sustained neural activity in delineated areas of prefrontal cortex and parietal cortex has also been measured during working memory tasks involving maintenance of items in memory. These sorts of processes should thus be amenable to examination with fMRI, as has in fact been demonstrated in many neuroimaging studies (Buckner & Koutstaal, 1998; Cohen et al., 1997; D'Esposito et al., 1995; Feredoes & Postle, 2007; McCarthy et al., 1994). Furthermore, the temporal resolution of fMRI has permitted the separate examination of encoding, maintenance, and retrieval stages in a working memory task. In such a way, fMRI potentially provides researchers in social psychology with a means to examine the neural processes that support the various stages of social behavior.

Some Limitations of fMRI

Many neural processes fall outside of this spatial and temporal resolution, however. For example, many neural processes involved in perception

and cognition last just a few milliseconds, and so will not be captured by fMRI unless the experiment is designed in such a way as to allow the accumulation of several such processes. Batty and Taylor (2003) used EEG to measure event-related potentials (ERPs) to emotional facial expressions. They demonstrated effects of emotion (vs. nonemotion) as early as 90 msec after stimulus onset, and differences between emotion expressions after just 140 msec. Importantly, later ERPs showed a different pattern across emotion expressions, presumably due to the engagement of later-stage processing of a different type. The same experiment performed with fMRI would not be able to identify these different stages of processing; rather, the measured response to a given expression would be the cumulative result of all processing stages. For the measurement of fast neural responses, EEG/ERPs or magnetoencephalography would thus be a more appropriate choice—although these techniques do not offer the three-dimensional spatial resolution of fMRI, since they are based on measurements from the scalp surface, and thus reflect the distal aggregate electromagnetic effects of underlying neural activity.

A further consideration with fMRI is that BOLD contrast is a relative and arbitrarily scaled measurement, as explained above. Although BOLD varies with dentritic activity, we have no way of quantifying exactly how much activity is occurring, or even how much more activity (as a difference or a ratio) occurs in one experimental condition than in another. This has a number of implications for emotion and social psychology research, particularly for studies involving individual differences or special populations. Perhaps the most important is that although all measured changes in fMRI signal obviously occur with respect to some underlying baseline neural activity, fMRI cannot tell us anything about the level of baseline activity (e.g., Stark & Squire, 2001). Furthermore, fMRI researchers must make the assumption that the amount of BOLD signal change in response to an experimental manipulation is unrelated, or at most linearly related, to the (unmeasurable) level of baseline neural activity. This assumption is questionable on basic neurophysiological grounds (Hyder, Rothman, & Shulman, 2002; Vazquez et al., 2006), as well as on psychological grounds. Consider, for example, a researcher who wishes to study the neural mechanisms underlying anxiety disorders by comparing brain activation in a clinical group with that in a group of low-anxiety individuals. It is reasonable to expect that upon entering the scanner, anxious individuals will have elevated activity in the brain regions involved in anxiety. To what extent, then, might the researcher expect their fMRI BOLD responses to increase further with the experimental manipulations? Might there be a ceiling effect in the anxious participants' fMRI BOLD responses, due to their elevated baseline activity? Is it valid to compare fMRI BOLD signal changes between

a group with low resting baseline and one with higher resting baseline? Without the means of measuring the baseline level of neural activity, the researcher has no way of answering these questions (see Whalen et al., 2008).

fMRI measurements also drift slowly over time, in such a way that fMRI cannot typically be used to measure any change to neural activity that occurs over a time period of more than a few minutes (though some newer fMRI techniques, such as arterial spin labeling [ASL], promise to address this limitation; Wang et al., 2003). So if a researcher's interest is in long-term changes in neural activity, or characterizing resting neural activity as a function of social context or individual differences in mood, then fMRI is probably not a good measurement choice. A further consequence of the relative nature of fMRI BOLD contrast is that comparisons of brain activation between different brain regions are difficult. This is because the magnitude of BOLD response in different neural regions to a given experimental manipulation will depend not only on changes in neural activity, but also on characteristics of the underlying brain anatomy and physiology. Even if we were able to measure those aspects of anatomy and physiology, their relationship with the BOLD response is not yet well enough understood to enable accurate comparisons of BOLD between different brain regions. As a consequence, any claim based on fMRI that, for example, auditory cortex was more or less active than visual cortex in response to a social stimulus would be methodologically questionable. In certain cases, then, PET would be a more appropriate brain imaging choice, since it can be used for absolute baseline and change measurements of cerebral metabolism—though its more invasive nature (involving radioactive tracers) precludes its use in many studies. Alternatively, ASL and other more absolutely quantifiable fMRI techniques are under development and are likely to be applicable to such research designs in the near future (Detre & Wang, 2002; see also this chapter's "Future Directions . . . " section).

Experimental Design Considerations

In order to get the most from an fMRI experiment, great consideration must be given to the ordering, duration, and timing of experimental conditions and stimuli. All of the same basic rules for psychology experiments also apply to brain imaging experiments. Conditions need to be properly counterbalanced or randomized to control for habituation, fatigue, or other order confounds. A suitable number of repetitions of each experimental condition is needed to produce reliable data. The exact number will depend not only on the inherent reliability of fMRI,

but also on the reliability of the behavior under observation. In addition to more general rules for experimental design, fMRI poses some further constraints, which we now discuss.

Block Designs versus and Event-Related Designs

A brief burst of neural activity corresponding to presentation of a short discrete stimulus or event will produce a more gradual BOLD response lasting about 15 sec. Due to noisiness of the BOLD signal, multiple repetitions of each condition are required in order to achieve sufficient reliability and statistical power. Block designs and event-related designs achieve this through slightly different means.

In a block design, multiple stimulus repetitions from a given experimental condition are strung together in a condition block, which alternates with one or more other condition blocks or control blocks. The advantage of such an approach is that the BOLD signal from multiple repetitions of closely spaced stimuli scales additively with the number of repetitions (Boynton, Engel, Glover, & Heeger, 1996); the more repetitions, the bigger and more reliable the signal. Because successive trials within a condition block are all the same condition, there is no fMRI-imposed lower limit to the intertrial interval (ITI), meaning that more repetitions can be fitted within a block of a given duration than would be the case with stimuli of different conditions. Block designs thus remain the most statistically powerful designs for fMRI experiments, all other considerations being equal (Bandettini & Cox, 2000). When a researcher is adopting a block design, three decisions need to be made:

- How long should each block be?
- How many blocks of each condition are needed?
- Should a resting baseline be included?

Based on knowledge of the noise characteristics of fMRI BOLD contrast, the answer to the first question is in the range of 20–40 sec (Skudlarski, Constable, & Gore, 1999; Smith, Jenkinson, Beckmann, Miller, & Woolrich, 2007). With blocks shorter than 20 sec, the statistical power advantage of a block design starts to disappear, particularly if psychological/behavioral constraints require a reasonably long ITI. With blocks longer than 40 sec, the random, slow fluctuations in fMRI BOLD signal start to compromise the reliability of the mean activation measured over the duration of a block. Of course, the psychological phenomena under question will add their own constraints, so this ideal range of block durations should be seen as a guideline only.

Because each block of a block design typically includes multiple repetitions of the same stimulus condition, it is not necessary to include many blocks. Many block design experiments are successful with between two and four blocks per condition. Once again, more blocks may be required for behavioral reasons (i.e., to achieve a reliable behavioral effect).

Finally, we strongly recommend including a resting baseline condition. Such a condition usually consists of having subjects look at a fixation cross on an otherwise blank screen. The amount of time spent in this condition is best when it matches the amount of time subjects are engaged in other experimental conditions of interest. The importance of including a resting baseline has been highlighted in some recent research showing that contrasts between an experimental condition and a resting baseline can be more reliable than contrasts between two experimental conditions. The idea is that a well-controlled resting baseline will be less variable than responses to many experimental stimuli. For example, Johnstone and colleagues (2005) reported that BOLD contrast in the amygdala showed greater test–retest reliability when calculated between responses to viewed fear faces and a passive fixation baseline than between responses to fear faces and responses to neutral faces. Obviously such condition–baseline contrasts do not have the task specificity of condition–condition contrasts, so inclusion of at least two active conditions is usually a good idea. In this way the condition–baseline contrast can be used for maximum sensitivity and reliability, whereas the condition–condition contrast can be used to answer questions about specific stimulus types or conditions. Inclusion of a resting baseline also makes basic fMRI quality control easier, since it is easy to check in primary sensory cortices for basic stimulus–baseline activation. For example, in a paradigm that uses visual stimuli, lack of stimulus–baseline activation in visual cortex would alert the experimenter to basic signal problems with the scanner, as well as to participants who were not paying attention, had their eyes closed, or fell asleep (not an uncommon problem in fMRI experiments, particularly with sleep-deprived undergraduate participants).

Although block designs are more statistically efficient (offering approximately 50% better signal-to-noise ratio than event-related designs; Bandettini & Cox, 2000), event-related designs are more suitable and often necessary in many experimental situations. In an event-related design, presentations of trials from different experimental conditions are interspersed in a random order, rather than being blocked together by condition (Buckner, 1998). This type of design avoids potential problems of habituation to specific stimulus types that might occur in a block design. In the affective and social sciences, event-related designs also avoid the problem of cumulative effects of emotional or social stimuli on a participant's mood. Moreover, event-related designs are necessary

when it is important that the participants not know in advance what type of stimulus will be presented. Finally, event-related designs allow subsequent analysis on a trial-by-trial basis, using behavioral measures such as judgment times, subjective reports, or physiological responses to correlate with measured BOLD responses.

The minimum number of events, or trials per condition, needed in an event-related design must above all else be determined by considering how many such trials would be required to achieve a reliable behavioral effect in a nonimaging study. In addition, the low signal and high noise of fMRI mean that experimenters would do well to include as many trials as possible within a scan session. For example, Huettel and McCarthy (2001) have shown that even in a fairly simple cognitive task, 50 repetitions of a given condition only resulted in 50% of activated voxels being detected—though with a greater number of participants this number can fall to as low as 20 (Thirion et al., 2007), or perhaps even lower for paradigms or stimuli that are particularly potent, such as pain (e.g., Salomons, Johnstone, Backonja, & Davidson, 2004). The statistics underlying the analysis of slowly varying hemodynamic signals impose further constraints on the timing of trials in an event-related fMRI design. Because the BOLD response is relatively slow, if trials are spaced too close together, then the BOLD responses from successive trials will overlap. This is a strength of block designs, because overlapping BOLD responses add together, leading to a bigger aggregate response. In an event-related design, however, it can prevent the BOLD responses from different trials, and different trial types, from being separately quantifiable. The easiest and most obvious solution to this is to space successive trials at least as far apart as the expected duration of the BOLD response to a single stimulus (e.g., 12 sec apart for brief stimuli, but longer for stimuli lasting more than about 1 sec; Bandettini & Cox, 2000).

Some simple math shows that this will not always lead to an ideal experimental design. Take, for example, an experiment in which we want to examine the differential brain responses to 1-sec presentations of happy, fearful, or neutral facial expressions, each of which shows eyes directed straight ahead or averted. Thus the experiment consists of six conditions. Let's say that from previous behavioral pilot testing, we know that we need 30 repetitions of each condition to achieve suitable statistical power to detect an expression × eye gaze interaction in a 3 (expression) × 2 (eye gaze) within-subjects analysis of variance. Allowing 12 sec per trial, this would amount to $12 \times 6 \times 30$ sec, or 36 min for the experiment. Although this experiment length might be reasonable for some participant samples, it may be too long for adolescent or older participants, for participants with social anxiety, or for other specific participant groups. In order to get the most "bang for the buck," many

researchers will also want to run two or more tasks in the one scanning session. Thus finding a way to reduce the time for a single task is invaluable.

The accepted way to do this for event-related designs is to use variable-length ITIs (i.e., "jitter"). By doing so, we can use the fact that overlapping BOLD responses add together in an approximately linear fashion, and so can be linearly separated if we use variable ITIs with known timing. Although the math underlying this technique is beyond the scope of this chapter (Birn, Cox, & Bandettini, 2002; Dale, 1999; Liu, Frank, Wong, & Buxton, 2001; Smith et al., 2007), the upshot is that trials can be spaced apart by an average ITI much shorter than would otherwise be necessary. In fact, to detect the difference in activation between two conditions, the shorter the mean ITI the better, down to a limit imposed by nonlinear saturation effects when successive trials are very close to one another. Mean ITIs of as low as 2–3 sec can lead to maximum statistical power with some designs (Friston, Zarahn, Josephs, Henson, & Dale, 1999); however, it is important to note than such short ITIs allow for the measurement of the difference in activation between conditions, but limit the ability to detect brain activation common to all conditions. To detect both differential and common activation, a somewhat longer ITI is preferable. Using our example of the face expressions experiment, we could use ITIs randomly varying between 2 and 10 sec, with an average ITI of 6 sec, which would result in an experiment length of $6 \times 6 \times 30$ sec, or 18 min—less than half the duration required for a slow event-related design—while retaining similar statistical power.

Three things are worth noting with respect to experimental design and timing. The first is that even if overall experimental time is not a major consideration, using variable ITIs makes it possible to increase statistical power. In our example, if time were not an issue, we could use variable ITIs to double the number of trials per condition (relative to a well-spaced design) without lengthening the experimental duration, and thus increase the reliability of our data. More sophisticated methods for choosing optimal experimental timing can yield even greater statistical power (Birn et al., 2002; Buracas & Boynton, 2002; Murphy, Bodurka, & Bandettini, 2007; Murphy & Garavan, 2005; Hagberg, Zito, Patria, & Sanes, 2001; Wager & Nichols, 2003). The second point is one that we repeatedly return to: The timing of an event-related design must make sense in a behavioral context. If it doesn't make sense to space trials closely or use variable ITIs in a behavioral experiment, then it almost certainly doesn't in an fMRI experiment either. The final point is that the ordering and timing of trials are critical if the results are to be statistically analyzable. Table 14.1 presents a summary of factors to consider in designing different types of fMRI experiments.

TABLE 14.1. Considerations for Different Types of fMRI Designs

Design type	Comments
Block designs	Have the strongest statistical power, compared to other types.
	Length of each block: 20–40 sec. Number of blocks per condition: Minimum 2–4 blocks. Resting baseline condition is highly recommended.
	Collection of behavioral and other measures is possible for correlation/regression analysis across participants.
Event-related designs	Number of events per condition depends on behavioral effect, but typically minimum of 20.
	Minimum ITI for evenly spaced events: 12 sec. Minimum ITI for jittered events: As low as 2–3 sec.
	Using jittered ITIs and randomized event order can increase statistical power, improve ecological validity, and reduce habituation/anticipation.
	Collection of behavioral and other measures is possible for correlation/regression analysis within and across participants.
Mixed designs	Useful when separating sustained from transient brain activity.
	Each event is presented using jittered ITIs within blocks; thus considerations for both block and event-related designs should be taken into account.
	Collection of behavioral and other measures is possible for correlation/regression analysis within and across participants.

Mixed Block and Event-Related Designs

More recently, researchers have recognized the need to take into account two distinct types of neural processes that occur in the brain during fMRI tasks. One type of brain activity is sustained throughout the trials of a task, regardless of individual trials; the other type of brain activity is specifically evoked by each trial of a task. Imagine that someone is taking a test such as the Graduate Record Examination. Once the clock starts ticking and the person begins reading the questions, certain neural circuits may enter a mode or a state that is optimized for written exams. However, in addition to this, other neural regions are reacting to each individual question on the test sheet. Naturally, brain regions turned on for the task itself rather than for each individual trial will show prolonged, sustained activations, while brain regions reacting to individual trials will exhibit short, transient activations. Reflecting the difference in the time course of brain activity, the former is often called "sustained activity," while the latter is dubbed "transient activity" (Visscher et al., 2003).

Since a simple block design or event-related design cannot separate these two types of brain activity, researchers have developed ways to combine block and event-related designs in fMRI. These "mixed designs" have been used for the purpose of dissociating sustained and transient trial-related activity in the brain (Donaldson, 2004; Donaldson, Petersen, Ollinger, & Buckner, 2001; Dosenbach et al., 2006; Visscher et al., 2003). In order to separate these two types of processes in fMRI data, trials are presented within task blocks interspersed with control blocks just as in a block design; within the task blocks, however, the onset of each trial is jittered, just as it would be in an event-related design. Having control blocks allows estimation of sustained brain activity separately from transient effects, and jittered trials enable dissociation of transient brain activity from sustained effects. Importantly, the interaction of sustained and transient manipulations can also be examined by measuring the difference in event-related responses between different blocks of sustained activity.

Mixed block and event-related designs are likely to be particularly suitable to social and affective neuroscience studies, in which the sustained social or emotional context has a direct impact on more transient socioaffective behaviors or responses. For example, Smith, Stephan, Rugg, and Dolan (2006) used a mixed design to examine the sustained effect of task requirements on retrieval of emotional information. Retrieval trials were presented as variable-ITI events embedded within blocks of one of two tasks—one requiring emotional information in order to make an accurate response, and one for which the emotional information was irrelevant. Smith and colleagues were able to demonstrate that connectivity among the hippocampus, amygdala, and orbitofrontal cortex during retrieval of emotional cues depend on whether the emotional cues were relevant or irrelevant to the task demands.

In another example, Bishop, Jenkins, and Lawrence (2007) used a mixed block and event-related design to examine how perceptual load (high vs. low, manipulated in blocks) would modulate activation of the amygdala to anxiety-relevant information (manipulated in an event-related manner with fearful vs. neutral facial expressions). They found that trait anxiety interacted with perceptual load in the engagement of prefrontal circuits thought to regulate amygdala output.

Studies of social interaction may also benefit from such mixed designs, enabling researchers to model sustained activation corresponding to attention to certain types of social stimuli (e.g., facial expressions of emotion, emotional speech) separately from transient activation

corresponding to processing specific stimulus features (e.g., presence or absence of a given facial or vocal feature).

Combining fMRI with Other Online Measures: Behavior and Physiology

fMRI reveals activation at a spatial and temporal scale too coarse to enable us to say much about the fundamental neural processes involved. fMRI is also fundamentally correlational, telling us which brain regions coactivate with a particular behavior or response, but not which brain regions were responsible for producing the response. But fMRI data do not exist in isolation, and when combined with other complementary techniques can lead to valuable insights into the social–affective brain.

Take, for example, the question of the neural mechanisms that underlie social interaction problems in autism. Individuals with autism have a spectrum of symptoms, including severe disruption of social contact and interaction (Kanner, 1943). It has been observed that individuals with autism perform poorly on tasks involving the perception of emotional facial expressions (Celani, Battacchi, & Arcidiacono, 1999). Early brain imaging studies comparing individuals with autism to controls in facial expression tasks found a lack of activation in the fusiform gyrus, a brain region thought to be important for processing faces (Pierce, Muller, Ambrose, Allen, & Courchesne, 2001). The obvious conclusion was that autism involves loss of function of this brain region, with consequent poor performance in facial expression tasks, which would be expected to have a negative impact on social interaction skills more generally.

In a follow-up to these studies, however, Dalton and colleagues (2005) questioned whether the lack of fusiform activation was the cause of performance deficits, or whether it was itself caused by dysfunction in some other brain region, perhaps reflecting an attentional or affective deficit. To test this hypothesis, Dalton and colleagues measured the pattern of eye gaze fixations while individuals viewed pictures of different emotional facial expressions. In individuals with autism, not only fusiform gyrus but also amygdala activation was found to correlate with the amount of time individuals fixated on the eye region of the faces—one of the principal regions that conveys emotional information in the face (see Plate 14.1). The implication is that faces of strangers, which are particularly aversive stimuli to individuals with autism, provoke strong amygdala activation; this may underlie the tendency of those with autism to avoid focusing on the eye region, and thus may explain the prior autism findings of hypoactivation in the fusiform gyrus.

Individual Differences in Social and Affective Style, Personality, and Temperament

Many fMRI studies have two or more experimental conditions, and look for main effects of these conditions. This is especially true in the field of traditional visual and cognitive neuroscience, where individual differences are usually negligible, or considered as a confounding factor. However, in social and personality psychology studies, individual differences are often the main focus. For example, researchers conducting such a study may use a scale that measures empathy to collect a range of empathy scores across all subjects. These empathy scores can then be used to categorize subjects into subgroups (e.g., low empathizers and high empathizers), or entered as a covariate in a general linear model analysis. The same approach can be applied to fMRI data. Social neuroscientists interested in the neural correlates of empathy can collect empathy scores just as they would do in a behavioral experiment, and analyze them in conjunction with fMRI data collected during a task that evokes empathy. As with behavioral experiments, brain activity can be contrasted between groups (low vs. high empathizers), or empathy scores can be used as a regressor to identify brain regions within which activation correlates with the scores. Obviously, the word "empathy" in this example could be exchanged with any characteristic of social or affective style, personality, or temperament of interest. And since we would have a single value per subject, this analysis could be performed on fMRI data regardless of whether a block or an event-related design is used to collect them.

Another way of using fMRI to study individual differences is to parametrically model each response to a stimulus within each subject. To put it another way, when analyzing fMRI data, we could create a general linear model with subjective ratings for each and every stimulus entered as a regressor, convolved with an ideal hemodynamic response function. This would allow us to identify brain regions that show activation covarying with the magnitude of the subjective rating we have used. For example, let's assume that we have presented pictures of human faces during an fMRI experiment and collected attractiveness ratings for each face from all of the subjects. We could analyze the fMRI data by creating a general linear model as we would in a conventional subtraction analysis, but in this case we would also enter the (zero-meaned) attractiveness ratings as a regressor. As a result, we would have a statistical map for each individual showing brain regions in which activation linearly tracks the magnitude of attractiveness ratings for each subject. We could then use a further (second-level) analysis to identify brain regions that

show such attractiveness activation across subjects. Here lies the major difference between this type of analysis and the former one: The former regression is performed during the second level (i.e., "group"-level) analysis only.

Functional and Effective Connectivity

Complex social behaviors involve the coordinated activity of multiple brain regions, so testing for fMRI activation in single brain regions is likely to be of limited use in social–affective neuroscience. Instead, we may be interested in investigating distributed networks of neural circuitry and determining how it is engaged or disengaged in different type of tasks. For example, we may want to know whether the activation we see in the amygdala during emotional face perception is functionally coupled to other regions of the brain, such as the orbitofrontal cortex, which is involved in learning and/or storage of hedonic information (Ghashghaei, Hilgetag, & Barbas, 2007; Iidaka et al., 2001; LoPresti et al., 2008).

Two different approaches to studying brain connectivity with fMRI are functional connectivity analysis and effective connectivity analysis. "Functional connectivity" is the temporal correlation between spatially remote neurophysiological events, whereas "effective connectivity" is defined as the influence one neuronal system exerts over another (Friston, 1994). By definition, output from functional connectivity analysis is purely correlational; even if two different brain regions show similar activation profiles across time, it does not necessarily mean that they are connected and directly influencing each other. In contrast, effective connectivity analysis attempts to infer directional influences of one region on another, at least to a certain degree.

Concretely, one simple way of applying functional connectivity analysis to our emotional face expression example would involve extracting time course data from the amygdala and performing a voxelwise cross-correlation analysis across the whole brain. More sophisticated techniques involve testing models of connectivity between multiple brain regions, using structural equation modeling (SEM; McIntosh & Gonzalez-Lim, 1994), dynamic causal modeling (Friston, Harrison, & Penny, 2003), or other related analytic techniques. Researchers have successfully confirmed that functional connectivity maps generated from voxels that were active during sensory or motor tasks corresponded to actual statistical maps of the auditory (Biswal, Yetkin, Haughton, & Hyde, 1996), visual (Lowe, Mock, & Sorenson, 1998), and motor

(Cordes et al., 2001; Lowe et al., 1998; Xiong, Parsons, Gao, & Fox, 1999) cortices, validating this approach. As with standard tests of brain activation, functional connectivity analyses can proceed voxelwise for the whole brain, though if we have predictions or hypotheses concerning a set of specific brain regions, it would be desirable to reduce the search volume to these a priori regions of interest (ROIs). Using an ROI-based functional connectivity analysis, Kim and colleagues (2004) found that the amygdala and the rostral anterior cingulate cortex were positively coupled during the processing of surprised facial expressions when combined with relevant contextual information. Thus this analysis proved useful in a task where the authors hypothesized that the amygdala would require assistance from the prefrontal cortex to use contextual information to understand the meaning of a given facial expression.

Effective connectivity analysis involves more advanced methods, such as psychophysiological interaction (PPI; Friston et al., 1997) analysis and SEM (McIntosh & Gonzalez-Loma, 1994). PPI analysis shows significant changes in the contribution of one brain region to another, as a function of experimental (psychological) context (Friston et al., 1997). Like functional connectivity analysis, PPI analysis can be performed across the whole brain without a priori regions of interest; however, once again, it is desirable to have a strong hypothesis and use PPI as a confirmatory rather than an exploratory analysis. This issue becomes especially important in SEM techniques. SEM takes into account the covariance matrix among multiple a priori regions of interest, and estimates the overall validity of the model as well as the strength and direction of the paths between the brain regions. Without a valid hypothesis for a brain network, SEM can lead to identifying a strongly significant but anatomically implausible model.

Increasing numbers of social and affective neuroscience researchers are taking advantage of these functional and effective connectivity analyses. As long as the researchers are aware of their caveats, these analyses have the potential to complement conventional statistical analysis by providing a more in-depth characterization of the neural mechanisms underlying a given social or affective process.

Considering Model Habituation or Familiarity Effects

As briefly mentioned earlier in this chapter, some brain regions tend to show a lesser degree of activation after repeated exposure to the same or similar stimuli. One prime example is the amygdala in response to fearful facial expressions. Amygdala activation to fearful facial expressions has

been observed to show rapid habituation, regardless of explicit knowledge of perceiving the faces (e.g., use of masked or nonmasked facial stimuli) or study design (block or event-related; Breiter et al., 1996; Hare et al., 2008; Whalen et al., 1998). Moreover, the degree of amygdala habituation to fearful facial expressions has been found to be correlated with trait anxiety levels (Hare et al., 2008). Thus the rate of habituation itself could be a particularly interesting aspect of fMRI data.

The question arises of how we can quantify this habituation effect. Recall that a typical fMRI experiment consists of multiple repetitions of each event or block type, often distributed over multiple scan runs. Depending on how we define a habituation effect, there are two ways to calculate this: (1) early trials versus late trials within each run, or (2) early runs versus late runs. The former method is useful when we predict that the degree of brain activity will decline as a function of repetition or time within each run, whereas the latter method is more suitable if we think that habituation will occur throughout the entire experiment. We can certainly take advantage of both methods as well. For example, Somerville, Kim, Johnstone, Alexander, and Whalen (2004) calculated habituation effects of the amygdala in response to happy and neutral facial expressions within runs (e.g., block 1 > block 2 > block 3) and across runs (run 1 > run 2). Another possibility is to explicitly include a habituation factor in the general linear model by modifying each condition regressor by some function of time (e.g., an exponential decay; Tabert et al., 2007).

A remaining question, however, pertains to the underlying mechanism of this apparent habituation effect in the brain. The declination of brain activity over time may be solely due to habituation, but an alternative explanation would involve inputs from other brain areas that modulate brain activity in this region. In this case, we could achieve more complete understanding of our data by using the functional and effective connectivity methods described earlier to analyze the temporal profiles of these brain activations.

Future Directions in Social and Affective Brain Imaging

MRI of the human brain is a fairly recent development, and its application to social and affective neuroscience is even more recent. As such, the constant stream of new MRI techniques and analysis strategies announced at conferences and in journals can be quite daunting to the social neuroscientist. In this section, we highlight a few developments that may find their way into the social neuroscience toolbox in the near future.

Perfusion Imaging

Perfusion imaging, or ASL imaging, is a technique by which blood flow in arteries (which supply the brain with blood) is imaged—as opposed to blood flow in veins (which drain blood), as is the case with traditional BOLD imaging. ASL involves magnetically "tagging" the blood flowing into the brain by subjecting it to a brief radiofrequency burst. As the tagged blood perfuses into the small arteries in particular neural regions, it gives rise to an altered MRI signature. One advantage of ASL is that arterial blood flow changes that are associated with brain activation are more local than changes to venous draining blood flow changes, which may reflect downstream effects of large areas of neural tissue. ASL thus promises to provide better spatial resolution than BOLD imaging (Detre & Wang, 2002). In addition, ASL MRI can (at least in principle) be absolutely quantified, and thus ASL changes from different brain regions, different individuals, and different experimental sessions should be directly comparable. Even though absolute quantification of ASL images is not yet practical in most experimental contexts, ASL does have better temporal stability and test–retest reliability than BOLD (Leontiev & Buxton, 2007; Wang et al., 2003), thus making it attractive for the imaging of longer-term brain activation changes or studies involving repeated imaging sessions. The disadvantage of ASL is that it is less sensitive than BOLD, showing smaller signal change for a given experimental manipulation.

Diffusion Tensor Imaging

Diffusion tensor imaging (DTI) exploits the fact that water molecules diffuse at different rates through different types of brain tissue. In particular, water diffusion along the myelinated axons in white matter tracts (bundles of nerve fibers) is greater than across the tracts. DTI can thus be used to map the white matter tracts of the brain—the paths of nerve fibers that connect different neural regions (Le Bihan, 2003). DTI is a structural rather than a functional technique. In social and affective neuroscience, DTI is likely to be useful in interpreting data from functional and effective connectivity analyses—for example, by posing constraints on which neural regions are interconnected. Without DTI, knowledge of the neural connections between different brain regions must be based on studies with animals and on histological studies on postmortem human brains, neither of which is ideal for the study of healthy human brains. It may also be possible to use DTI to examine long-term changes to the neural connectivity between different brain

regions associated with specific social–affective disorders (e.g., autism; Alexander et al., 2007; Barnea-Goraly et al., 2004), or in the study of the development of the human social brain.

Magnetic Resonance Spectroscopy

Most functional brain imaging studies address changes to the activity of different brain regions resulting from some experimental manipulation. Magnetic resonance spectroscopy (MRS) is a technique that can be used to measure concentrations of specific metabolites within the brain, including *N*-acetyl aspartate, lactate, creatine, and choline. MRS works by identifying magnetic resonance peaks at frequencies that correspond to known molecular resonances. Because the signal-to-noise ratio of MRI is low, MRS can currently only be applied to measure molecular concentrations of a small number of molecules in an a priori selected region of the brain. For the sensitive imaging of the concentrations of a different range of molecules across the whole brain, PET is a more suitable technology. Nonetheless, MRS may find uses in studies where individual differences in the concentration of a specific neurotransmitter in a given neural region are thought to underlie some affective behavior or disorder (see Pitman, Shin, & Rauch, 2001; Steingard et al., 2000; van Elst et al., 2001).

Summing Up: How Can fMRI Contribute to Social Neuroscience?

As we have seen in this chapter, fMRI can be applied in numerous ways to study the neural basis of human behavior. Standard BOLD imaging can measure increases in brain activation resulting from an experimental manipulation or corresponding to a specific behavior to within a few millimeters. Connectivity analyses can inform us about how different parts of the brain interact in social perception, expression and behavior. Many of the weaknesses of fMRI, such as its relative lack of temporal resolution, can be overcome with clever experimental design or through combination with other techniques (such as EEG).

As with all relatively new scientific approaches, a great deal of caution is required in designing, performing, and analyzing fMRI experiments. In particular, researchers need to be wary of using fMRI for fMRI's sake—that is, performing fMRI experiments merely because the scanner and funding to do so are available. Much of the research reported in the neuroimaging literature might be criticized on this point. An experiment that tells us which isolated brain regions are activated in a given task may give us no information at all about what processes are involved in that task, or how they are instantiated in the brain.

The problem of "inverse inference," or inferring process from brain activation, is also a potential pitfall. For inverse inference to be valid, we need to know enough about the detailed functioning of a given neural region to start with, yet this is rarely the case. For example, just because a particular brain region such as the anterior cingulate is involved in processing pain, this does not imply that activation of the same region during some other task must involve pain processing. Far more likely is that the anterior cingulate serves some other, more fundamental subprocess of pain processing that might equally be involved in processing any number of other types of stimuli.

The limitations of fMRI are all the more important to keep in mind for two reasons. One is the cost of fMRI studies. Researchers (and grant reviewers) need to ask whether the cost involved in brain imaging studies is justified by the added scientific value. Is an fMRI replication of a well-known social psychology experiment justified if the only conclusion that can be drawn is that "process *X* resides in the brain" (surely only an ardent dualist would need brain imaging proof of that), or that "process *X* activates regions *A*, *B*, and *C* of the brain" (telling us little or nothing about what regions *A*, *B*, and *C* do, except the trivial answer "process *X*")?

The second consideration is the weight that brain imaging findings seem to carry with the general public, and also parts of the research community. A recent study (McCabe & Castel, 2008) showed that experimental findings, when accompanied by brain imaging pictures, were given more credence than the same study without the fancy graphics. A related study (Weisberg, Keil, Goodstein, Rawson, & Gray, 2008) showed that non-neuroscience experts were more likely to accept a flawed explanation of human behavior when it was dressed up in pseudoscientific neuroscientific jargon than when the identical theory was presented in non-neuroscientific terms. Not only researchers, but also grant reviewers and journal editors and reviewers, need to keep this very much in mind, and ask themselves this question: "What does brain imaging really add to the study in question?"

Despite these limitations, there is much to be optimistic about. fMRI, when combined with well-thought-out, well-controlled, and theoretically motivated experimental designs and appropriate complementary measures, offers researchers an unprecedented tool for understanding the brain and all the behaviors it supports. The emerging fields of social and affective neuroscience promise to bring about a profound change in our understanding of our social and emotional lives.

References

Alexander, A. L., Lee, J. E., Lazar, M., Boudos, R., DuBray, M. B., Oakes, T. R., et al. (2007). Diffusion tensor imaging of the corpus callosum in autism. *NeuroImage, 34*(1), 61–73.

Bandettini, P. A., & Cox, R. W. (2000). Event-related fMRI contrast when using constant interstimulus interval: Theory and experiment. *Magnetic Resonance in Medicine, 43*(4), 540–548.

Barnea-Goraly, N., Kwon, H., Menon, V., Eliez, S., Lotspeich, L., & Reiss, A. L. (2004). White matter structure in autism: Preliminary evidence from diffusion tensor imaging. *Biological Psychiatry, 55*(3), 323–326.

Batty, M., & Taylor, M. J. (2003). Early processing of the six basic facial emotional expressions. *Brain Research: Cognitive Brain Research, 17*(3), 613–620.

Belin, P., Zatorre, R. J., & Ahad, P. (2002). Human temporal-lobe response to vocal sounds. *Brain Research: Cognitive Brain Research, 13*(1), 17–26.

Belin, P., Zatorre, R. J., Lafaille, P., Ahad, P., & Pike, B. (2000). Voice-selective areas in human auditory cortex. *Nature, 403*, 309–312.

Birn, R. M., Cox, R. W., & Bandettini, P. A. (2002). Detection versus estimation in event-related fMRI: Choosing the optimal stimulus timing. *NeuroImage, 15*(1), 252–264.

Bishop, S. J., Jenkins, R., & Lawrence, A. D. (2007). Neural processing of fearful faces: effects of anxiety are gated by perceptual capacity limitations. *Cerebral Cortex, 17*(7), 1595–1603.

Biswal, B. B., Yetkin, F. Z., Haughton, V. M., & Hyde, J. S. (1996). Functional connectivity in the auditory cortex studied with fMRI. *NeuroImage, 3*(Suppl.), 305.

Boynton, G. M., Engel, S. A., Glover, G. H., & Heeger, D. J. (1996). Linear systems analysis of functional magnetic resonance imaging in human V1. *Journal of Neuroscience, 16*(13), 4207–4221.

Breiter, H. C., Etcoff, N. L., Whalen, P. J., Kennedy, W. A., Rauch, S. L., Buckner, R. L., et al. (1996). Response and habituation of the human amygdala during visual processing of facial expression. *Neuron, 17*(5), 875–887.

Buckner, R. L. (1998). Event-related fMRI and the hemodynamic response. *Human Brain Mapping, 6*(5–6), 373–377.

Buckner, R. L., & Koutstaal, W. (1998). Functional neuroimaging studies of encoding, priming, and explicit memory retrieval. *Proceedings of the National Academy of Sciences USA, 95*(3), 891–898.

Buracas, G. T., & Boynton, G. M. (2002). Efficient design of event-related fMRI experiments using M-sequences. *NeuroImage, 16*(3, Pt. 1), 801–813.

Celani, G., Battacchi, M. W., & Arcidiacono, L. (1999). The understanding of the emotional meaning of facial expressions in people with autism. *Journal of Autism and Developmental Disorders, 29*(1), 57–66.

Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., et al. (1997). Temporal dynamics of brain activation during a working memory task. *Nature, 386*, 604–608.

Cordes, D., Haughton, V. M., Arfanakis, K., Carew, J. D., Turski, P. A., Moritz, C. H., et al. (2001). Frequencies contributing to functional connectivity in the cerebral cortex in "resting-state" data. *American Journal of Neuroradiology, 22*(7), 1326–1333.

Dale, A. M. (1999). Optimal experimental design for event-related fMRI. *Human Brain Mapping, 8*(2–3), 109–114.

Dalton, K. M., Nacewicz, B. M., Johnstone, T., Schaefer, H. S., Gernsbacher, M. A., Goldsmith, H. H., et al. (2005). Gaze fixation and the neural circuitry of face processing in autism. *Nature Neuroscience, 8*(4), 519–526.

D'Esposito, M., Detre, J. A., Alsop, D. C., Shin, R. K., Atlas, S., & Grossman, M. (1995). The neural basis of the central executive system of working memory. *Nature, 378*, 279–281.

Detre, J. A., & Wang, J. (2002). Technical aspects and utility of fMRI using BOLD and ASL. *Clinical Neurophysiology, 113*(5), 621–634.

Donaldson, D. I. (2004). Parsing brain activity with fMRI and mixed designs: What kind of a state is neuroimaging in? *Trends in Neurosciences, 27*(8), 442–444.

Donaldson, D. I., Petersen, S. E., Ollinger, J. M., & Buckner, R. L. (2001). Dissociating state and item components of recognition memory using fMRI. *NeuroImage, 13*(1), 129–142.

Dosenbach, N. U. F., Visscher, K. M., Palmer, E. D., Miezin, F. M., Wenger, K. K., Kang, H. C., et al. (2006). A core system for the implementation of task sets. *Neuron, 50*(5), 799–812.

Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature, 392*, 598–601.

Feredoes, E., & Postle, B. R. (2007). Localization of load sensitivity of working memory storage: Quantitatively and qualitatively discrepant results yielded by single-subject and group-averaged approaches to fMRI group analysis. *NeuroImage, 35*(2), 881–903.

Friston, K. J. (1994). Functional and effective connectivity in neuroimaging: A synthesis. *Human Brain Mapping, 2*, 56–78.

Friston, K. J., Buechel, C., Fink, G. R., Morris, J., Rolls, E., & Dolan, R. J. (1997). Psychophysiological and modulatory interactions in neuroimaging. *NeuroImage, 6*(3), 218–229.

Friston, K. J., Harrison, L., & Penny, W. (2003). Dynamic causal modelling. *NeuroImage, 19*(4), 1273–1302.

Friston, K. J., Zarahn, E., Josephs, O., Henson, R. N., & Dale, A. M. (1999). Stochastic designs in event-related fMRI. *NeuroImage, 10*(5), 607–619.

Ghashghaei, H. T., Hilgetag, C. C., & Barbas, H. (2007). Sequence of information processing for emotions based on the anatomic dialogue between prefrontal cortex and amygdala. *NeuroImage, 34*(3), 905–923.

Goense, J. B., & Logothetis, N. K. (2008). Neurophysiology of the BOLD fMRI signal in awake monkeys. *Current Biology, 18*, 631–640.

Hagberg, G. E., Zito, G., Patria, F., & Sanes, J. N. (2001). Improved detection of event-related functional MRI signals using probability functions. *NeuroImage, 14*(5), 1193–1205.

Hare, T. A., Tottenham, N., Galvan, A., Voss, H. U., Glover, G. H., & Casey, B. J. (2008). Biological substrates of emotional reactivity and regulation in

adolescence during an emotional go/no-go task. *Biology Psychiatry, 63*, 927–934.

Huettel, S. A., McCarthy, G. (2001). The effects of single-trial averaging upon the spatial extent of fMRI activation. *NeuroReport, 12*, 2411–2416.

Hyder, F., Rothman, D. L., & Shulman, R. G. (2002). Total neuroenergetics support localized brain activity: Implications for the interpretation of fMRI. *Proceedings of the National Academy of Sciences USA, 99*, 10771–10776.

Iidaka, T., Omori, M., Murata, T., Kosaka, H., Yonekura, Y., Okada, T., et al. (2001). Neural interaction of the amygdala with the prefrontal and temporal cortices in the processing of facial expressions as revealed by fMRI. *Journal of Cognitive Neuroscience, 13*(8), 1035–1047.

Johnstone, T., Somerville, L. H., Alexander, A. L., Oakes, T. R., Davidson, R. J., Kalin, N. H., et al. (2005). Stability of amygdala BOLD response to fearful faces over multiple scan sessions. *NeuroImage, 25*(4), 1112–1123.

Kanner, L. (1943). Autistic disturbances of affective contact. *Nervous Child, 2*, 217–250.

Kim, H., Somerville, L. H., Johnstone, T., Polis, S., Alexander, A. L., Shin, L. M., et al. (2004). Contextual modulation of amygdala responsivity to surprised faces. *Journal of Cognitive Neuroscience, 16*, 1730–1745.

Le Bihan, D. (2003). Looking into the functional architecture of the brain with diffusion MRI. *Nature Reviews Neuroscience, 4*(6), 469–480.

Leontiev, O., & Buxton, R. B. (2007). Reproducibility of BOLD, perfusion, and CMRO2 measurements with calibrated-BOLD fMRI. *NeuroImage, 35*(1), 175–184.

Liu, T. T., Frank, L. R., Wong, E. C., & Buxton, R. B. (2001). Detection power, estimation efficiency, and predictability in event-related fMRI. *NeuroImage, 13*(4), 759–773.

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature, 412*, 150–157.

Logothetis, N. K., & Wandell, B. A. (2004). Interpreting the BOLD Signal. *Annual Review of Physiology, 66*(1), 735–769.

LoPresti, M. L., Schon, K., Tricarico, M. D., Swisher, J. D., Celone, K. A., & Stern, C. E. (2008). Working memory for social cues recruits orbitofrontal cortex and amygdala: A functional magnetic resonance imaging study of delayed matching to sample for emotional expressions. *Journal of Neuroscience, 28*(14), 3718–3728.

Lowe, M. J., Mock, B. J., & Sorenson, J. A. (1998). Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. *NeuroImage, 7*(2), 119–132.

McCabe, D. P., & Castel, A. D. (2008). Seeing is believing: The effect of brain images on judgments of scientific reasoning. *Cognition, 107*(1), 343–352.

McCarthy, G., Blamire, A. M., Puce, A., Nobre, A. C., Bloch, G., Hyder, F., et al. (1994). Functional magnetic resonance imaging of human prefrontal cortex activation during a spatial working memory task. *Proceedings of the National Academy of Sciences USA, 91*(18), 8690–8694.

McIntosh, A. R., & Gonzalez-Lim, F. (1994). Structural equation modeling and its application to network analysis in functional brain imaging. *Human Brain Mapping, 2*, 2–22.

Moonen, C. T. W., Bandettini, P. A., & Aguirre, G. K. (1999). *Functional MRI*. New York: Springer.

Murphy, K., Bodurka, J., & Bandettini, P. A. (2007). How long to scan?: The relationship between fMRI temporal signal to noise and necessary scan duration. *NeuroImage, 34*(2), 565–574.

Murphy, K., & Garavan, H. (2005). Deriving the optimal number of events for an event-related fMRI study based on the spatial extent of activation. *NeuroImage, 27*(4), 771–777.

O'Craven, K. M., & Kanwisher, N. (2000). Mental imagery of faces and places activates corresponding stimulus-specific brain regions. *Journal of Cognitive Neuroscience, 12*(6), 1013–1023.

Pierce, K., Muller, R. A., Ambrose, J., Allen, G., & Courchesne, E. (2001). Face processing occurs outside the fusiform "face area" in autism: Evidence from functional MRI. *Brain, 124*(Pt. 10), 2059–2073.

Pitman, R. K., Shin, L. M., & Rauch, S. L. (2001). Investigating the pathogenesis of posttraumatic stress disorder with neuroimaging. *Journal of Clinical Psychiatry, 62*(Suppl. 17), 47–54.

Salomons, T. V., Johnstone, T., Backonja, M. M., & Davidson, R. J. (2004). Perceived controllability modulates the neural response to pain. *Journal of Neuroscience, 24*, 7199–7203.

Skudlarski, P., Constable, R. T., & Gore, J. C. (1999). ROC analysis of statistical methods used in functional MRI: Individual subjects. *NeuroImage, 9*(3), 311–329.

Smith, A. P., Stephan, K. E., Rugg, M. D., & Dolan, R. J. (2006). Task and content modulate amygdala–hippocampal connectivity in emotional retrieval. *Neuron, 49*(4), 631–638.

Smith, S., Jenkinson, M., Beckmann, C., Miller, K., & Woolrich, M. (2007). Meaningful design and contrast estimability in fMRI. *NeuroImage, 34*(1), 127–136.

Somerville, L. H., Kim, H., Johnstone, T., Alexander, A. L., & Whalen, P. J. (2004). Human amygdala responses during presentation of happy and neutral faces: Correlations with state anxiety. *Biology Psychiatry, 55*(9), 897–903.

Stark, C. E. L., & Squire, L. R. (2001). When zero is not zero: The problem of ambiguous baseline conditions in fMRI. *Proceedings of the National Academy of Sciences USA, 98*, 12760–12766.

Steingard, R. J., Yurgelun-Todd, D. A., Hennen, J., Moore, J. E. C., Moore, C. M., Vakili, K., et al. (2000). Increased orbitofrontal cortex levels of choline in depressed adolescents as detected by *in vivo* proton magnetic resonance spectroscopy. *Biological Psychiatry, 48*(11), 1053–1061.

Tabert, M. H., Steffener, J., Albers, M. W., Kern, D. W., Michael, M., Tang, H., et al. (2007). Validation and optimization of statistical approaches for modeling odorant-induced fMRI signal changes in olfactory-related brain areas. *NeuroImage, 34*(4), 1375–1390.

Thirion, B., Pinel, P., Mériaux, S., Roche, A., Dehaene, S., & Poline, J.-P. (2007). Analysis of a large fMRI cohort: Statistical and methodological issues for group analyses. *NeuroImage, 35*(1), 105–120.

van Elst, L. T., Thiel, T., Hesslinger, B., Lieb, K., Bohus, M., Hennig, J., et al. (2001). Subtle prefrontal neuropathology in a pilot magnetic resonance spectroscopy study in patients with borderline personality disorder. *Journal of Neuropsychiatry and Clinical Neuroscience, 13*(4), 511–514.

Vazquez, A. L., Cohen, E. R., Gulani, V., Hernandez-Garcia, L., Zheng, Y., Lee, G. R., et al. (2006). Vascular dynamics and BOLD fMRI: CBF level effects and analysis considerations. *NeuroImage, 32*(4), 1642–1655.

Visscher, K. M., Miezin, F. M., Kelly, J. E., Buckner, R. L., Donaldson, D. I., McAvoy, M. P., et al. (2003). Mixed blocked/event-related designs separate transient and sustained activity in fMRI. *NeuroImage, 19*(4), 1694–1708.

Wager, T. D., & Nichols, T. E. (2003). Optimization of experimental design in fMRI: A general framework using a genetic algorithm. *NeuroImage, 18*(2), 293–309.

Wang, J., Aguirre, G. K., Kimberg, D. Y., Roc, A. C., Li, L., & Detre, J. A. (2003). Arterial spin labeling perfusion fMRI with very low task frequency. *Magnetic Resonance in Medicine, 49*(5), 796–802.

Weisberg, D. S., Keil, F. C., Goodstein, J., Rawson, E., & Gray, J. R. (2008). The seductive allure of neuroscience explanations. *Journal of Cognitive Neuroscience, 20*(3), 470–477.

Whalen, P. J., Johnstone, T., Somerville, L. H., Nitschke, J. B., Polis, S., Alexander, A. L., et al. (2008). A functional magnetic resonance imaging predictor of treatment response to venlafaxine in generalized anxiety disorder. *Biological Psychiatry, 63*(9), 858–863.

Whalen, P. J., Rauch, S. L., Etcoff, N. L., McInerney, S. C., Lee, M. B., & Jenike, M. A. (1998). Masked presentations of emotional facial expressions modulate amygdala activity without explicit knowledge. *Journal of Neuroscience, 18*(1), 411–418.

Xiong, J., Parsons, L. M., Gao, J. H., & Fox, P. T. (1999). Interregional connectivity to primary motor cortex revealed using MRI resting state images. *Human Brain Mapping, 8*(2–3), 151–156.

Index

Page numbers followed by *f* indicate figure, *t* indicate table