13 Concepts and Principles of Clinical Functional Magnetic Resonance Imaging

AVRAM J. HOLMES AND ANGUS W. MACDONALD III

The human brain is comprised of a dense web of interdigitated functional networks. Understanding how the brain's complex functions give rise to human cognitive abilities in both health and disease depends on unraveling the carefully coordinated interactions between networked brain regions and their responses to environmental change (Holmes & Patrick, 2018). Historically, substantial progress was made delineating this intricate architecture through postmortem dissections in humans and tract tracing and lesion studies in animals. Yet there remained many gaps in our understanding of how the brain influences behavior, particularly psychiatric illnesses. The limitations of these labor-intensive approaches have receded over the past 40 years with the advent of *in vivo* imaging approaches such as positron emission tomography (PET), electroencephalography (EEG), electrocorticography (ECoG), and magnetoencephalography (MEG; see Raichle, 2009, for a historical overview). The introduction of functional magnetic resonance imaging (fMRI), in particular, has sparked spectacular growth in psychiatry research.

Fueled by rapid methodological and analytic advances, fMRI has come to dominate the clinical literature, allowing us to study brain function in a rapid and noninvasive manner across ever larger samples. These technological developments have made it easy to become exceedingly optimistic about the future of clinical neuroscience. Although fMRI methods have evolved rapidly since the first brain scans in the 1980s, there remain core approaches and theoretical principles that can be used to understand the current state of the field and anticipate future innovations. Here we take a critical look at how fMRI measures can inform our understanding of brain functions in psychopathology. To help researchers select appropriate methods, we will cover fMRI study design, analysis, and interpretation and discuss some of the advantages and disadvantages of each design and analytic choice.

EXPERIMENTAL APPROACHES

A central goal of clinical cognitive neuroscience is to understand how common cognitive and neural systems may differ in people with, or at risk for, psychopathology. Among cognitive mechanisms we include perceptual processes, such as stimulus detection or facial recognition; salience-related processes, such as reward or threat detection; executive processes, such as cognitive control, emotion regulation or decision making; and motor processes, such as response initiation. Of course, there are many more, with names and descriptions that often overlap (Poldrack, 2010). The brain implements these cognitive processes in a manner that is evident at different levels of analysis, from molecular mechanisms (e.g., neurotransmission) to large-scale networks (e.g., spike-timing-dependent plasticity through which neurons assemble). Clinical cognitive neuroscience uses an array of methods to disentangle the complexities of psychopathology into finer, potentially discrete, deficits in specific aspects of brain biology. This approach is often most effective when it builds on established behavior-brain research that informs an understanding of individual differences. For example, a good deal of work on working memory impairments in psychosis patients extends upon foundational working memory studies in humans (Park, Holzman, & Goldman-Rakic, 1995) and nonhuman primates (e.g. Funahashi, Bruce, & Goldman-Rakic, 1989).

Measuring the Brain When Performing Tasks

It is intuitive to ask how the brains of people with mental illnesses differ while thinking. Pioneers such as Ingvar and Franzen (1974) used a forerunner of PET imaging to study resting cerebral blood flow (rCBF) in patients with schizophrenia during a cognitively demanding task. Although overall rCBF levels were similar to controls, in postcentral sulcus it was relatively higher in patients with schizophrenia, whereas in prefrontal cortex it was relatively lower. Thus began an era of function-based neuroimaging efforts that have increased in power and sophistication in the subsequent decades.

Table 13.1 delineates the four general approaches to task-evoked fMRI that are commonly used, and Figure 13.1

https://doi.org/10.1017/9781316995808.016

Table 13.1 Experimental approaches using functional magnetic resonance imaging

Method	Application	Strengths	Limitations	Key references
Block design (task evoked)	Contrasts conditions within an ongoing task or between an ongoing task and rest	Efficient data collection; maximizes potential activation differences	Difficult to attribute activation to a specific cognitive mechanism when contrasted tasks differ in several demands; difficult to exclude errors from analysis; requires task development	(Amaro & Barker, 2006)
Slow event-related design (task evoked)	Contrasts different trial types or cognitive demands sparsely timed	Different events may be intermixed in an unpredictable manner; relatively few assumptions about the nature of the hemodynamic response; few constraints on the interactions (or dependencies) between different cognitive demands; behavioral results often generalize to faster-paced versions; trials that fulfill a criterion (e.g., error trials) can be examined	Because hemodynamic response must return to baseline between trials, fewer trials can be collected; fewer trials for the evaluation of behavioral performance; boredom due to slow pace; requires task development	(Amaro & Barker, 2006)
Fast event-related design (task evoked)	Contrasts different trial types or cognitive demands more densely timed	As with slow event-related designs, different events may be intermixed; behavioral results are more robust because more trials are available; trials that fulfill a criterion (e.g., error trials) can be examined if sufficiently independent	Hemodynamic response function modeling required; events must be sufficiently independent; jittering time between events or using partial (catch) trials to make events independent may affect performance; requires task development	(Amaro & Barker, 2006; Ollinger, Shulman, & Corbetta, 2001)
Hybrid block/ event-related designs (task evoked)	Nests an event-related design within a larger block design to allow multiple analyses	Allows for robust analyses present in block designs; additionally, trials that fulfill a criterion (e.g., error trials) can be examined separately	Hemodynamic response function modeling required; events only occur within the context the block; requires task development	(Braver, Reynolds, & Donaldson, 2003)
Intrinsic function ("resting state")	Examines on-going activity in the absence of specifically timed tasks or cognitive demands	Shorter development cycle; shorter training and fewer task demands facilitating data collection in special-needs populations; applicable to patients asleep or under sedation; easier to harmonize across sites and easier to combine data sets post hoc facilitating larger sample sizes	Brain activity cannot be related to specific cognitive events (but see Smith et al., 2009); group/individual differences findings may reflect differences in habitual thought patterns rather than ability to activate a region; connectivity metrics affected by subtle head motion	(Biswal et al., 1995; Smith et al., 2013)

illustrates how these approaches differ in stimulus presentation and associated models of the blood-oxygen-level-dependent (BOLD) response, which reflects the concentration of deoxygenated, relative to oxygenated hemoglobin. A ratio that is altered by local neural activity (Huettel, Song, & McCarthy, 2004; Logothetis et al., 2001). *Block designs* are far and away the most robust, and are therefore the most efficient, strategy for obtaining maps of where the BOLD response is occurring in the brain. A block design study consists of discrete "on" and "off" periods, each lasting from tens of seconds to minutes in duration. During the "on" times a stimulus is presented or a behavior is elicited. These blocks are contrasted with "off" periods that consist of rest, baseline, or alternate task

https://doi.org/10.1017/9781316995808.016

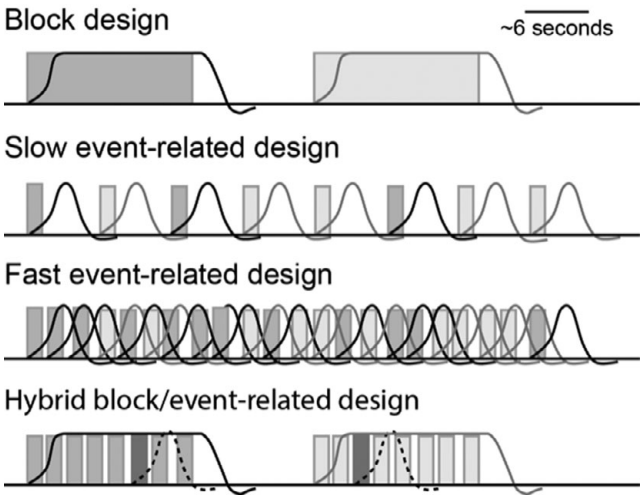

Figure 13.1 Different task conditions (gray boxes) and the corresponding convolution with a prototypical hemodynamic response function predict the rise and fall of the BOLD time course in an activated voxel across various experimental designs. Time proceeds from left to right. Note, the subtraction method of analysis (and derived brain maps) often compares the difference between the extent to which a voxel's time course resembles the black predicted time course relative to the lighter grey.

states. For example, the Human Connectome Project (HCP) collected seven tasks twice in block designs lasting 2–5 minutes (Barch et al., 2013). During the HCP working memory task, participants switched between performing a 2-back task (which required a running memory load of two items) and a 0-back attentional control task for 25 seconds each. All brain areas with metabolic demands when performing the task showed a rising BOLD response that reached a steady state within 5–6 seconds of the beginning of the block and then declined within 5–6 seconds when the cognitive load was removed. Because the metabolic demands associated with observing and responding to the stimuli were similar in the 2-back and 0-back conditions, a comparison between them would likely show little differential visual or motor activity. Instead, the biggest changes appeared in places where there were greater demands when maintaining a running load of two items. This strategy assumes that all items in a given condition are similarly difficult and that no aspect of the task (e.g., the occurrence of a repeated item) is of particular interest. The robustness of block designs can also be used to examine changes in activation across populations (e.g., case vs. comparison samples) or treatment conditions. Haut, Lim, and MacDonald (2010) compared 2-back to 0-back activity, this time in people with schizophrenia, to examine how cognitive training tasks changed activity more than controls in several regions of prefrontal cortex. Despite these advantages, block designs come with a number of potential interpretive problems for clinical research. For example, brain regions activated more in the 2-back than in the 0-back may be involved in many processes besides a higher working memory load, such as updating the stimuli after each trial, suppressing interference, monitoring conflict, expecting and preparing for another trial of the same kind, and experiencing frustration or even futility. For reasons we'll discuss further below in "Interpretation," performance differences between groups on the different blocks can also present a challenge, as the analysis of error trials is mixed together with the analysis of accurate trials.

In order to examine discrete trials and address constraints of block designs, many experimenters have employed event-related (later called "slow event-related") designs (Huettel, 2012) that leverage the hemodynamic time-course associated with local regional neural activity. This approach is characterized by large gaps in time between stimuli, allowing the BOLD response to rise and fall before presenting the next trial (which may have different tasks demands). Because this allows one to disentangle more components of task performance, many investigators gravitated toward this technique. For example, in a study of context processing-related deficits associated with the genetic liability to schizophrenia, MacDonald, Becker, and Carter (2006) differentiated between the task demand of maintaining a task representation, which involved dorsolateral prefrontal cortical (DLPFC) and was impaired in patients and first-degree relatives, and that of overcoming conflict, which evoked the anterior cingulate and was impaired in patients but not their relatives. What's more, these differences were observed on correct trials, thereby partially controlling for individuals' fluctuating task engagement. However, due to the slow nature of the hemodynamic response, fewer trials can be included using this approach, which provides reduced statistical power relative to a block design.

One way to introduce more trials is to group them closer together without allowing the hemodynamic response to resolve fully, which is commonly called the *fast event-related design*. This strategy owes its existence to several key contributions. As illustrated in Figure 13.1, convolution models for fMRI analysis (Friston et al., 1994) combined the time a stimulus occurred with an expectation about how the BOLD signal would respond. Subsequently, it was found that the BOLD response could be summated across successive trials even with short intertrial intervals (Dale & Buckner, 1997). This property of the BOLD signal allows task-relevant activations to be predicted based on the expected response to closely spaced stimuli and events, if the trials – or events within the trial – are sufficiently independent. This provides the opportunity to analyze closely spaced trials, or events within a task, drastically shortening the study collection times and the associated burden on research participants. For example, Poppe and colleagues (2016) took advantage of this design using a paradigm that required cue maintenance to control the response to a subsequent probe. In this case, if the cue was an A, then one would respond left to the probe, but if the cue was a B, then a right response to the probe was correct. To disentangle (to the degree possible) the relationships between the cues and probes, jitter (i.e., pseudorandom variation) was

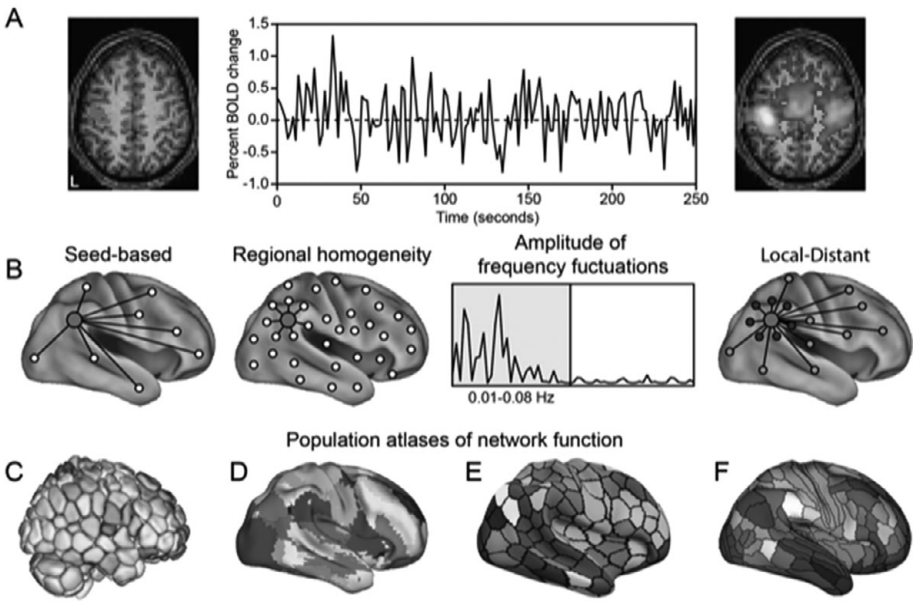

Figure 13.2 A summary of intrinsic connectivity analyses methods and cortical parcellations based on in-vivo brain imaging. (A) Intrinsic fluctuations in the fMRI BOLD signal exhibit patterns of covariation within functionally connected brain networks in the absence of overt task performance. Map of motor network from the seminal work by Biswal and colleagues (1995) as adapted by Vincent et al. (2006). (B) Selection of common analyses methods for intrinsic connectivity analyses (see Table 13.3). (C–E) Intrinsic fluctuations can be used to derive *in-vivo* brain parcellation. (C) Shen et al. (2013), (D) Power et al. (2011), and (E) Schaefer et al. (2018). (F) Multimodal parcellation using intrinsic connectivity, relative myelin mapping, cortical thickness and task-based fMRI (Glasser et al., 2016).

introduced into the interstimulus and intertrial intervals to facilitate modeling these independently.

Hybrid designs are used when the robustness of a block design is desirable, but where some questions can also be addressed using the more specific information that comes from examining individual trials. In such instances, one can blend block and event-related designs. This technique could be used to model individual trials within a block to identify and remove error trials, examine how different trials within a category interact with each other, or track different time courses across regions. The overarching point is that even as each design has various strengths and limitations, they need not be mutually exclusive. With planning, a given task might be conceptualized and analyzed from several perspectives, taking advantage of the associated strengths of each.

Measuring the Brain without Tasks: The "Resting State"

In addition to responding to stimuli, the brain also shows reliable patterns of activity in the absence of explicit task states or in so-called *resting state designs* (e.g., Biswal et al., 1995). As the brain is not particularly good at resting, we prefer the term *intrinsic function* or *intrinsic functional connectivity* for what is being measured when the mind is not directed to a particular task. The earliest clinical cognitive neuroscience studies examining patients' intrinsic brain functioning generally used PET to measure glucose metabolism. For example, individuals with schizophrenia display lower levels of rCBF in prefrontal and temporal regions (Farkas et al., 1984). This raised the question as to whether these differences reflected an inability to use those brain regions (a direct result of illness) or a disposition to use those brain regions less, perhaps because due to distraction, fatigue or some other factor (a downstream result of illness). This ambiguity led this approach to fall out of favor for a period of time, yet a number of advantages as well as promising empirical observations have reestablished resting-state, or intrinsic functioning, as a mainstay of neuroimaging.

Work in this domain by Biswal and colleagues (1995; Figure 13.2A) suggested an intrinsic organization to the brain that mirrored task-related functions. Their foundational study revealed that even when the brain was not engaged in a motor task, signal fluctuations in the motor cortex were highly correlated with neighboring voxels as well as spatially distinct regions associated with motor functioning. Functional connectivity between brain regions, such as reported by Biswal and colleagues, is usually analyzed in terms of correlation, signal coherence, or other temporal similarities in BOLD fluctuations. Providing converging evidence for intrinsic approaches to the study of brain functions, Koch, Norris, and Hud-Georgiadis (2002) combined diffusion-based and functional methods to reveal that intrinsic correlations between brain regions may depend on anatomic projections. This principle was expanded by Smith and

colleagues (2009), who used a statistical algorithm called independent components analysis to identify a number of distributed functional networks (about 20 different such "components") and showed that the structure of these networks mapped quite closely to activation patterns from a large-scale meta-analysis of broad task categories. For example, meta-analysis indicated that regions of left prefrontal cortex and posterior parietal cortex frequently coactivated, along with a region of anterior cingulate cortex. The tasks most likely to coactivate these regions were working memory, explicit memory and language tasks. The observed locations closely resembled a network of voxels that coactivated at rest among a much smaller cohort of volunteers (a set of findings that has been replicated, e.g., Wisner et al., 2013), suggesting that intrinsic fluctuations may reflect coactivation among the regions with shared profiles of task-evoked function (Deco, Jirsa, & McIntosh, 2011).

Intrinsic approaches provide a new complement to task-based study designs, and the ease of collection and flexibility of intrinsic analyses has led to a rapid rise in their popularity. This is particularly true for clinical researchers, as the derived markers of intrinsic network function are more widely applicable than traditional measures of task-based fMRI. Since intrinsic network function can be assessed during sleep and under anesthesia, this functional mapping approach may be widely implemented in diverse populations including children, non-English-speaking participants, developmentally delayed patients, and patients who are under sedation. The promise of resting-state scanning has undergirded the accumulation of large-scale data sets that would be difficult if not impossible to obtain through more traditional task-based approaches. Open-access samples in the thousands are now widely available to the broader scientific community, such as the 1000 Functional Connectomes Project (Biswal et al., 2010), the Brain Genomics Superstruct Project (Holmes et al., 2015), the Human Connectome Project (Van Essen et al., 2013), and the UK Biobank (Ollier, Sprosen, & Peakman, 2005).

Despite the putative simplicity of intrinsic function studies, there remain a number of outstanding questions when acquiring and interpreting these data: what are the implications of collecting data with eyes open or closed (Van Dijk et al., 2010), with or without eye tracking, or if performing a low-level periodic response task (Krienen, Yeo, & Buckner, 2014)? How much data need to be collected to measure properties such as connectivity strength or network coherence reliably (Zuo et al., 2014)? This is a dynamic area of work with constantly emerging findings. Whatever the answers may be, intrinsic functional connectivity does have heritable characteristics (Ge et al., 2018) that may be informative about the nature of psychopathology (Baker et al., 2014).

Measurement and Neurometrics

Whether observing the brain during task performance or not, it is important to consider several general points as we move from asking "basic" questions regarding *how the brain works* to individual differences questions like *how do brains function differently*. Individual differences analyses evoke the 4Rs of measurement: robustness, repeatability, reliability, and replicability. Robustness is the likelihood a given analytic approach will provide a consistent answer. Repeatability is the likelihood that the same pattern of findings will occur if the same group is measured again. Reliability is the extent to which the participants are at the same point in a distribution when measured again, showing a consistent pattern of individual differences. Finally, replicability is the likelihood that the same pattern of findings will occur in a new sample.

The latter three, repeatability, reliability, and replicability, may be affected by various factors, such as caffeine (Laurienti et al., 2002), nicotine (Thiel & Fink, 2007), and ethanol (Seifritz et al., 2000) intake, but also by more subtle variables such as glucose levels (Anderson et al., 2006) and cardiac variability (Shmueli et al., 2007). Before becoming overwhelmed, the investigator should consider the extent to which these factors will affect *changes* in the pattern of evoked BOLD response or intrinsic connectivity for which they will be searching. Similarly, are these factors going to be a source of noise (reducing power) or biases (introducing confounds)?

ANALYSIS

Analytic approaches for fMRI data have rapidly increased in complexity since the initial discovery of the BOLD contrast and the early efforts to use it to map human mental operations (for review see Raichle, 2009). Nevertheless, there are several principles that can be used to understand the current state of the field, as well as guide our anticipation of what advances may wait on the horizon. In this section, we will introduce the methods most frequently used in event-related and intrinsic fMRI analyses, briefly discussing key advances that have shaped the field. For ease of interpretation, the analytic techniques, as well as their associated strengths and limitations, are presented in Tables 13.2 and 13.3. Critically, while these approaches can be applied individually, often two or more will be utilized within the same set of analyses.

Subtraction, Correlation, and Contrast Analyses

The evolution of block, fast-event-related, and hybrid fMRI designs was closely followed by the development of associated analytic methods. Initially, fMRI researchers leveraged approaches adapted from positron emission tomography where signal quality was greatly enhanced if participants were placed in a standard stereotaxic, or common, physical space (Fox et al., 1988). Once the individual participant data is registered to a common reference space the most straightforward and broadly applied method for obtaining results is to perform a simple

https://doi.org/10.1017/9781316995808.016

Table 13.2 Analytic techniques to examine task-evoked data sets

Method	Application	Strengths	Limitations	Key references
General linear model	Estimates the contribution of known predictors to BOLD signal fluctuations	Mathematically simple, relatively easy to interpret, available in standard analysis packages; can include multiple independent variables (e.g., scanner drifts, participant motion, etc.)	Relies on assumptions including a consistent hemodynamic response throughout the brain and the temporal stability of noise terms	(Friston et al., 1994)
Psychophysiological interaction	Examines the interaction between a task contrast of interest and the functional coupling between brain areas	Can reveal a task-specific change in correlation between areas that may not be evident through a shared effect of task	Can only examine a single source area; causal relations cannot be inferred	(Friston et al., 1997)
Structural equation model	Assesses the degree to which experimental manipulations influence the functional connectivity of brain regions	Can be used for both exploratory and confirmatory testing; based on prior knowledge of brain structure/function; can estimate causal relations and be used across multiple regions simultaneously	Can require a priori assumptions about causality, potentially obscuring other relations; lacks temporal information; assumes linearity	(McIntosh & Gonzalez-Lima, 1991)
Dynamic causal model	Uses biologically plausible neuronal models of the BOLD response to estimate the influence of experimental context on the functional coupling among brain regions	Uses hidden interactions at the neuronal level to study observable shifts in BOLD response; models bidirectional and modulatory interactions (for a comparison of SEM and DCM approaches see Penny et al., 2004)	Relies on pre-specified models and the inferences provided are only as valid as the priors used in the estimation procedure	(Friston, Harrison, & Penny, 2003)
Granger causality model	Assesses the degree to which one time series can predict another	Does not rely on a priori assumptions (e.g., regions of interest and associated connection)	Assumes (local) stationarity, incorrect inferences can result from measurement noise and/or hemodynamic response latencies across brain	(Kamiński et al., 2001)
Meta/mega-analysis	Assesses relations across multiple imaging data sets. Meta-analysis refers to the pooled analysis of published results; mega-analysis refers to the pooled analysis of raw data	Can increase power due to the large number of studies/participants available for analysis; other approaches (e.g., estimates of effective connectivity) can utilize meta/mega-analysis defined regions of interest	Experimental designs may not be uniform and/or adequately sample the full spectrum of behavior and function; meta-analyses often consider the distribution of activation peaks, rather than each study's/contrast's distributed pattern of activity; relies on traditional contrast analyses, this can serve as a confound if the process of interest is not successfully isolated	(Fox, Parsons, & Lancaster, 1998; Laird et al., 2005)

Table 13.3 Analytic techniques of intrinsic brain function

Method	Application	Strengths	Limitations	Key references
Seed-based correlations	Estimates the correlation between the BOLD signal in a predefined regions of interest with other regions, or rest of the brain	Mathematically simple and easy to interpret	Requires the a priori selection of regions; may provide illusory specificity	(Biswal et al., 1995)
Regional homogeneity	Uses Kendall's coefficient concordance to assess the similarity of the time series of a given voxel to those of its nearest neighbors	Mathematically simple and easy to interpret	Requires the a priori selection of regions; sensitive to spatial smoothing and the size of the region of interest	(Zang et al., 2004)
Local-distant	Takes into account local regional connections as well as remote or distant connections outside of a defined area	Allows for the analyses of relative weighting of local or distant connectivity in a region	Can conflate real cortical/anatomical distance with Euclidean distance	(Sepulcre et al., 2010)
Principal component analysis	Creates uncorrelated variables from best-fitting linear combinations of the variables in the raw data; reduces the dimensionality of complex data types	Can reveal hidden, simplified, features in high-dimensional data; does not require a priori task models or estimates of BOLD response	Based on a strong assumption of linearity and orthogonality in the resulting components; sensitive to noise and assumes a high signal-to-noise ratio in the data	(Friston et al., 1993; Viviani, Grön, & Spitzer, 2005)
Independent component analysis	An extension of principal component analysis that separates data into spatially or temporally independent patterns of activity	Few a priori assumptions; not restricted to deriving orthogonal components	Components are assumed to be statistically independent	(McKeown & Sejnowski, 1998; Calhoun et al., 2001)

subtraction across conditions of interest. Subtraction techniques, or more generally correlation analyses, are based on the expectation that the voxels/brain regions participating in a psychological or cognitive process should show dissociable functional responses during the completion of associated tasks. Rather than revealing absolute levels of cerebral blood flow or metabolism linked to a cognitive process, contrast analyses reveal relative changes in BOLD response across conditions. By averaging the time points acquired during an experimental condition and subtracting the average of all the time points associated with a control condition, differing in only one property, the brain regions associated with a cognitive process of interest can be identified.

While a growing proportion of fMRI studies go beyond subtraction logic to include parametric effects where the independent variable has a number of levels (e.g., task difficulty, stimulus intensity, monetary rewards), simple subtraction techniques are a powerful analytic approach. With an appropriate task design, they can be applied to preprocessed fMRI time courses using standard statistical techniques. Historically, subtraction analyses have provided foundational discoveries, characterizing the aspects of brain function that support key facets of cognition and behavior across health and disease. For example, consistent with a role in the modulation of affective functions, differential amygdala responses have been observed during the visual processing of emotional and neutral facial expressions in healthy populations (Breiter et al., 1996). Dysregulated amygdala response to emotional stimuli is hypothesized to underlie the onset and maintenance of affective illness (Mayberg, 1997). In line with these theories, in patient populations subtraction techniques have revealed abnormal amygdala responses in disorders marked by affective impairments (Price & Drevets, 2012) and in populations at increased genetic risk for onset (Smoller et al., 2014).

General Linear Models

The introduction of single-trial or event-related fMRI designs provided researchers the opportunity to separate mental operations into discrete moments in time, allowing for the differentiation of their associated fMRI signals (Huettel, 2012). The associated shift from representing

https://doi.org/10.1017/9781316995808.016

BOLD responses as static across blocks of time to considering moment-to-moment fluctuations allowed researchers to leverage dynamic analysis methods. In this area, general linear models (GLMs; introduced for fMRI analyses by Friston et al., 1994) are the primary analysis approach utilized in task-based research. The events in event-related designs often occur so rapidly that their associated BOLD responses overlap. GLM analyses assume that the observed BOLD signal is comprised of a linear combination of experimental factors (thereby allowing overlapping responses) and an uncorrelated noise term. A GLM analysis identifies voxels where the signal changes in response to experimental conditions, or events, calculating the significance/extent of effects based on how well the observed data fits the predicted model. These GLM-based approaches form the theoretical scaffolding that underlies most forms of fMRI data analysis, for instance the regression, prediction, and data exploration approaches detailed below. Importantly, GLM analyses are typically conducted in a mass univariate manner across each voxel, and there are several assumptions to keep in mind when utilizing GLM-based approaches that can constrain our interpretation of the results (see Monti, 2011). These include the use of a single model (design matrix) throughout the brain, that noise varies consistently across all time points (e.g., baseline relative to a contrast of interest), and the independence of associated statistical tests.

Multivariate Modeling and Predictive Approaches

An important limitation of traditional GLM-based analytic techniques is that they treat each voxel as independent, assessing if the signal within these discrete data points fluctuates in response to a task condition of interest. They do not account for the possible contribution of complex multivariate relations linking multiple voxels. As the field has developed beyond this mass univariate approach, an increased emphasis has been placed on computationally sophisticated approaches for identifying spatially distributed patterns of brain activity (e.g., multivoxel pattern analysis). For a more thorough treatment of the techniques from the field of machine learning and analytic approaches where specific mental states or task contexts are decoded from distributed activity patterns readers are referred to Chapter 34. In brief, these multivariate approaches are typically implemented in a two-step process. First, a classifier is trained to distinguish the occurrence of events for different conditions within a subset of the available data. Second, the trained model is then applied to an independent or held out sample where the classifier attempts to predict the events of interest. These approaches hold promise as a potential diagnostic tool for psychiatric illness, and their flexibility allows for their integration with other complementary processing and analysis techniques (Rosenberg, Casey, & Holmes, 2018). For instance, machine-learning approaches have been used to discriminate male from female participants accurately (Chekroud et al., 2016), identify dissociable cognitive trajectories in Alzheimer's disease (Zhang et al., 2016) with gross morphometric estimates of brain anatomy, and predict individual participant attentional capacity, disease status (e.g., ADHD; Rosenberg et al., 2016), or symptomology through analyses of large-scale network function (e.g., presence of psychosis; Reinen et al., 2018).

Network Modeling

Recently, researchers have begun to shift their emphasis from the study of the specialization or segregation of brain functions in isolated regions toward an analytic framework that targets functional integration, working to characterize how signals covary across spatially distinct regions (for review see Sporns, 2014). These distributed processing models of brain function provide a powerful method to explain complex cognitive functions, individual variation, and the behavioral expression of psychiatric illness. Network models allow researchers to represent brain systems as distributed sets of neural elements and their associated interconnections. The generation of these network models requires partitioning or parcellation aspects of the brain into regions, or nodes, which share a consistent set of features. Broadly, brain networks reflect two different categories. Structural networks that describe the anatomical wiring properties of the brain, and functional networks reflect interactions among time series (e.g., correlations) across anatomical parcels or regions of interest. Unsurprisingly, network approaches encompass much of the current research on brain functions ranging from the biophysical modeling of task data through the estimation of the integrity of resting-state networks. Figure 13.2 displays a collection of population-intrinsic network parcellation schemes. Readers should note that the distinction between event-related (task-evoked) and resting-state (intrinsic) analytic techniques is in many ways arbitrary. These methods each probe specific features of brain function; with an appropriate study design they can be applied across data types.

Task-Evoked Functional Connectivity

These analyses can be broadly separated into two classes. The first examines functional connectivity, or the temporal correlation of observed BOLD responses between remote neural areas, similar to the intrinsic techniques detailed below. The second consists of model-based approaches that assess effective connectivity, or the putative influence one brain system or region may exert on another. Prototypical effective connectivity analyses include psychophysiological interaction, structural equation (McIntosh & Gonzalez-Lima, 1991), dynamic causal, and Granger models (Friston et al., 2003).

Psychophysiological interaction, for example, assesses whether connectivity varies between spatially distant brain regions in different psychological/task contexts (Friston et al., 1997). The presence of a psychophysiological interaction suggests that regional responses in the source area to an experimental or psychological factor are modulated by signals from a distal brain region. These approaches have revealed a host of key discoveries, for example, the aberrant development of amygdala-prefrontal connectivity following maternal deprivation, potentially reflecting an ontogenetic adaptation in response to early adversity (Gee et al., 2013).

Intrinsic Functional Connectivity

The convergence of new imaging technologies and increased computational resources has provided tools to map both local and distant connections in the brain (Holmes & Yeo, 2015). Recent work in this domain has established a strong correspondence between the structure of intrinsic (resting state) and extrinsic (coactivation) brain networks, suggesting that the brain's functional architecture at rest is closely linked to cognitive function (Smith et al., 2009; Tavor et al., 2016). Aberrant patterns of connectivity within these networks are evident across many major mental disorders, indicating that their breakdown can lead to diverse forms of psychological dysfunction (Buckholtz & Meyer-Lindenberg, 2012). For instance, impaired connectivity within the frontoparietal control network, which encompasses portions of the dorsolateral prefrontal, dorsomedial prefrontal, lateral parietal, and posterior temporal cortices, as well as corresponding aspects of the striatum and cerebellum (Yeo et al., 2011), is believed to underlie executive functioning deficits in psychotic illness (Baker et al., 2014; Reinen et al., 2018). A growing literature implicates frontoparietal network impairments as transdiagnostic markers of psychopathology (Cole, Repov, & Anticevic, 2014). A set of relationships may emerge through the generation of symptoms that are domain-specific (e.g., impaired executive function), but cut across many pathologies (Buckholtz & Meyer-Lindenberg, 2012).

There are myriad ways that network functions can be probed with intrinsic approaches (Table 13.3; Figure 13.2). From flexibility in the definition of networks of interest, including the use hypothesis derived "seed" regions defined through meta-analyses of task data (Yarkoni et al., 2011) and population atlases of network function (Schaefer et al., 2018; Yeo et al., 2011), through the use of complex dynamic (Hutchison et al., 2013; Reinen et al., 2018) and graph theoretical techniques (Sporns, 2014). Approaches that allow researchers to map functional network topography down within a single person, for example, are critical for clinical intervention and the study of individual differences (Kong et al., 2019; Wang et al., 2015). Research in this domain has led to the development of cortical parcellation methods to accurately map the brain's intrinsic functional organization at the individual level. Functional networks mapped by these techniques are highly reproducible among participants and effectively capture intersubject variability (Wang et al., 2015). Providing converging evidence for the use of intrinsic connectivity when defining participant-specific network topographies, these approaches have been validated by invasive cortical stimulation mapping in surgical patients, suggesting potential for use in clinical applications.

One key factor to consider across all fMRI analyses, but particularly those that examine functional connectivity, is the impact of participant motion. This is a concern for clinical researchers who frequently have to contend with study populations that differ markedly in terms of both disease status and data quality. In the area of intrinsic analyses, for instance, motion generates nonlinear effects on functional connectivity that can either artificially induce or obscure hypothesized results (Van Dijk, Sabuncu, & Buckner, 2012). While these effects cannot simply be regressed out, there are processing approaches that can limit the impact of motion on substantive findings (e.g., motion scrubbing; Ciric et al., 2018; Power et al., 2014). Additionally, given the availability of large-scale fMRI databases that measures of brain structure and function as well as multiple domains of cognition, behavior, and genetics (e.g., Holmes et al., 2015), some research groups have elected to carefully match patient and healthy comparison samples on the basis of data quality (e.g., Baker et al., 2014). The influence of data quality on connectivity analyses is a key point of consideration when interpreting case-control analyses, as patient populations often move more than healthy comparison samples. In the next section, we'll turn to several additional problems faced by clinical cognitive neuroscience.

INTERPRETATION

If brain functions are involved in mental illness, it would seem that methods akin to taking pictures of the living brain and then developing them would provide an objective, biological perspective on how that occurs. Of course, there are any number of reasons this simplistic optimism may not hold, but six criticisms of clinical cognitive neuroscience studies using neuroimaging stand out as particularly important to avoid. We hope the reader will note the challenge of satisfying all these constraints within a single study.

Mechanistic Specificity

The challenges of behavioral experimental psychopathology transfer quite directly to clinical cognitive neuroscience and neuroimaging in particular. A prominent challenge is the difficulty of demonstrating that a deficit in performance on a task is mechanistically relevant to the

https://doi.org/10.1017/9781316995808.016

disorder and not an epiphenomenal, or secondary, effect associated with the presence of illness. Deficits observed in isolation are uninterpretable. For example, when patients perform worse on a facial affect recognition task and have reduced fusiform gyrus activation we are tempted to conclude that these two features are linked to the pathology. Yet the link may be tenuous. Rather than having a role in the symptoms of the disorder, the association between performance and brain activity may result from an earlier perceptual impairment, attention lapses, reduced effort on the task, or any of a number of other failures. More compelling would be to show that patients are worse on facial affect recognition *relative to another task demand measured with equivalent discriminating power*, or ability to distinguish between the groups being measured (see Salem, Kring, & Kerr, 1996, for this particular comparison). Such deficits have been called differential deficits, mechanism-specific or specific cognitive deficits (for review, see Macdonald, 2015). Experiments using one condition, without a second condition that has similar levels of discriminating power, are obviously not up to this standard of evidence. More subtly in experiments with multiple conditions, if the condition of interest is measured with more discriminatory power, then patients may perform worse on it without actually tapping a mechanism related to the disorder. That is, the difference between patients and controls may derive from a nonspecific raft of difficulties patients face when performing behavioral tasks.

Causality Confound

This confound refers to the concern that group differences in brain activity that occur when one group shows differences in performance may not be interpretable (Gur & Gur, 1995). In this case we wish to conclude that the difference in brain activation causes the observed performance differences, however we must also rule out the possibility that both the activation and performance differences reflect another, perhaps unmeasured, impairment. For example, lower motivation, compliance or visual acuity, misunderstanding the task, higher distress, or any number of other failures could also impair performance and reduce task-related brain engagement. This challenge has caused a great deal of aggravation in clinical research because it runs counter to the goal of demonstrating mechanistic specificity (discussed above), which alone is quite a daunting task.

Several approaches to this conundrum have been suggested, none of which fully addresses all of the potential concerns. The least satisfying approach has been to use tasks on which patients are unimpaired but which tap into a known deficiency (such as using a very easy working memory task). This is generally accomplished by taking advantage of a ceiling effect, rather than making the unmeasured impairment irrelevant. Three other approaches match performance in other ways. One way to match performance is to select patients and controls from their broader population distributions based on who performs at a comparable level. This solution falters because of the problem of generalization to the populations of interest. Another way to match performance is to train participants differently so that those who struggle more with the task receive more practice than those who naturally perform it better. This solution can be critiqued in so far as tasks that have become more automated often use different brain areas compared to more novel tasks. The third way to match performance is to titrate the difficulty of the task so all participants, and therefore groups, perform equally. This may be an ideal solution in many cases, however it means that group differences in activation reflect, in part, differences in the tasks they are performing. A final approach we have used is to examine only accurate trials using an event-related analysis, suggesting that the participant was engaged in the task during a given trial. One criticism of this approach is that to the extent that participants can respond accurately simply by chance then some proportion of those trials may still reflect an unmeasured spurious impairment. A second criticism is that it is overly conservative, insofar as part of the deficit of interest is the inability to respond accurately, and in this case there will be no group differences in brain regions that may be generally more difficult for patients to engage. These strategies and critiques are all an extension of concerns that come from using a quasi-experimental methodology, with both within- and between-subject effects. Whereas the quest for mechanistic specificity leads us to test within-subject effects, we are still hampered from strong causal claims by the challenges of differences in performance.

Diagnostic and Symptom Specificity

Diagnostic specificity refers to the extent to which an impairment is disorder-specific versus common across a number of disorders. This concern arises especially from testing patients sampled from a single categorical diagnosis. For example, early findings that patients with schizophrenia showed impairments in dorsolateral prefrontal cortical functioning were greeted with excitement (Ingvar & Franzen, 1974). Subsequent findings of DLPFC dysfunction in many other disorders, from depression (Goodwin, 1997) to substance abuse (Goldstein et al., 2004), may suggest that DLPFC dysfunction is less a cause of psychotic symptoms and perhaps more of a general psychopathology liability factor. Strikingly, a recent meta-analysis comparing task-related brain activation in people with a psychiatric illness and healthy controls reported few differences between the diagnostic constructs in terms of the distribution of case-control effects across the brain (537 studies, total n = 21,427; Sprooten et al., 2018). The challenge of diagnostic specificity is not

https://doi.org/10.1017/9781316995808.016

limited to categorical discrimination, however. In the era of metastructural approaches to diagnosis (Holmes & Patrick, 2018), specificity refers to showing that an impairment relates more closely to a particular branch of psychopathology (e.g., thought disorder) than to other branches (e.g., externalizing or internalizing), or to general psychopathology (Lahey et al., 2012). This interpretive challenge can be addressed rigorously in a quasi-experimental design. Using a between-subjects design one can show that patients with an equal level of dysfunction with a different diagnosis show either performance or brain activation differences. Using a within-subject design, one can show that performance or brain activation differences correlate significantly more with one symptom factor relative to another using a Meng's Z or other appropriate test for correlated correlation coefficients. These complementary approaches allow researchers to demonstrate the presence of case/control differences in an aspect of brain biology and provide evidence that associated patterns of individual variability link with shifts in associated behaviors.

Forward and Reverse Inference

The vast majority of neuroimaging research uses an approach termed "forward inference" when probing the underlying biological architecture that supports cognitive functions (Henson, 2006). For example, when researchers manipulate stimuli to determine how the brain responds, forward inference proposes that a given experimental condition causes changes in local brain activity. Thus, dissociable BOLD responses can be used to distinguish between competing cognitive functions or theories. Critically, because forward inference is a correlational approach (see "Subtraction, Correlation, and Contrast Analyses" above), researchers cannot infer that the observed patterns of brain activity are either necessary or sufficient to support the associated cognitive process. However, as noted below, these shortcomings can be addressed through the integration of complementary methodology across levels, for instance the optogenetic modulation of neural activity within freely moving animals.

"Reverse inference" is a different inferential strategy utilized by much of the field, at least informally, and fraught with controversy. Here, researchers make a claim about the engagement of a specific cognitive process based on the activation of a given brain region (Poldrack, 2006). As an example, a researcher might observe that patients with schizophrenia exhibit heightened amygdala responses to images of scenes (e.g., mountains, plains, forests), leading them to erroneously conclude that scene viewing is associated with the experience of fear in psychotic illness. This sort of inference is common within the clinical literature where the core cognitive processes underlying psychiatric illnesses remain unknown. Reverse inference provides a useful deductive tool for expanding our understanding of the underlying brain mechanisms supporting behavior. However, this is a particularly weak standard of evidence, insofar as brain regions and networks generally activate in response to many different demands (Poldrack, 2006).

The issues pertaining to reverse inference are a widespread concern. Clearly, researchers should be cautious when making claims regarding their results, particularly when the functional properties of a given region have yet to be fully established, or in the absence of converging evidence from other methods. Even factors outside of a researcher's control can influence the accuracy of inferences, such as the number of voxels in a region of interest (ROI) or the selectivity of response in a given region of interest. Despite these limitations, reverse inferences can be exceptionally useful when applied judiciously, allowing researchers to relate cognitive processes across distinct theories and experimental contexts (Henson, 2006). Reverse inference can also be used to generate hypotheses, particularly when based on real data. Critically, both reverse and forward inferences can be formalized within a probabilistic framework. They can then be used for meta-/mega-analysis where they provide the opportunity for researchers to map links across diverse neural, cognitive, and disease states. These models provide the field with a powerful tool when coupled with the meta-analytic databases resulting from the recent development of text-mining and machine-learning techniques (Yarkoni et al., 2011).

Regional Differences in Sensitivity

Whenever we write that patients are impaired in one brain region, we imply that they are *not* impaired in the other brain regions examined. However, this implication is only true of other regions that we have measured at least as accurately, or sensitively, as the region where we found the group difference. The extent to which these other brain areas actually are measured as sensitively is largely ignored in the clinical imaging literature. Signal loss and susceptibility artifacts arise as a result of magnetic field inhomogeneities. In BOLD images, the decay in recoverable signal is exacerbated in regions where the brain is adjacent to air (e.g., sinus cavities). Clear spatial variation in voxel-level temporal SNR (the mean of the signal at each voxel over the BOLD run divided by the variance) is evident across the cortical mantle (Holmes et al., 2015). The associated problem is simple to illustrate: brain regions A and B are both impaired in patients, however brain region A (say the anterior cingulate) is measured with very good signal to noise and region B (say the orbitofrontal cortex, subject to susceptibility artifact) is measured with low signal to noise. In reporting our findings without acknowledging these differences in SNR, we end up implying that region B is unimpaired. We see the growing interest in neurometrics, the study of imaging

https://doi.org/10.1017/9781316995808.016

measurement akin to psychometrics, as an important development in clinical cognitive neuroscience (Poppe et al., 2013). A ready-to-hand check on this assumption is to examine signal-to-noise maps across the brain to see that the areas implicated in group differences are not simply those with the highest signal-to-noise.

Cross-Modality Integration

fMRI provides a remarkably powerful technique for researchers to measure and map the functional networks in the human brain in both health and disease, albeit with the limitations inherent to all non-invasive approaches. For instance, recent fMRI work has demonstrated correspondence across the topographic structure of intrinsic and task-evoked functional networks of the human brain, suggesting that the features of the resting brain are closely linked to cognition (Crossley et al., 2013). Yet an integrated understanding of the complex neurobiological architecture of the human brain, from molecules through cells, circuits, and functional networks, will not be possible with a single method or approach. Rather, progress in clinical neuroscience will be made through the combined efforts of researchers working across levels of analyses and species (Holmes & Patrick, 2018). In this regard, work that can join the heterogeneous information provided through distinct analytic approaches, including genetics, brain metabolism, anatomy, electrophysiology, and behavior, has the potential to provide deep insights into the pathophysiology of psychiatric illness. The incorporation of methods that directly manipulate brain function, for instance lesion and optogenetic approaches in animal models or transcranial magnetic stimulation in humans, can allow researchers to test the causal relations between brain and behavior observed in fMRI (e.g., Deng, Yuan, & Dai, 2018). Coupling molecular and genetic approaches with fMRI, as another example, can nominate gene profiles that preferentially associate with functional brain networks (Anderson et al., 2018; Richiardi et al., 2015), revealing the molecular machinery of network communication.

SUMMARY AND FUTURE DIRECTIONS

This chapter sought to bridge between the basic experimental world of cognitive neuroscience and that of clinical research. We hope that basic researchers will find in it links to questions they want to resolve when entering the correlational science of individual differences and clinical problems. Clinical researchers, in turn, should find here the tools to inform a clinical cognitive neuroscience approach to their populations, or the ideas needed to be informed consumers of such research. But whether the reader is more at home with a basic or a clinical perspective, clinical cognitive neuroscience remains an uncanny domain in which the most important achievements seem to be just over the horizon.

On the one hand, advances in methodology and our understanding of brain functions seem to be advancing at an unprecedented speed. Within the last several years, there have been developments in spatial and temporal resolution of MRI equipment, larger samples allowing us to observe subtle effects, a growing number of algorithms to identify meaningful signals, and studies of the effects of genes on brain functioning, promising to remake the landscape of clinical cognitive neuroscience. At the same time, much of this excitement is familiar from previous episodes in which the field was enthusiastic about the potential of widespread noninvasive imaging (in the 1990s) and ever-increasing magnetic field strength (in the 2000s). While technological advances will continue to allow us to ask new questions, we should be sober about how these changes will affect our understanding of, and ultimately our ability to help, people with mental illness. The field is uncanny because that "just over the horizon" feeling drives us forward, but at the same time we need to gird ourselves for the likelihood that new insights may only fill in a few more pieces of a very large puzzle.

Ultimately, our understanding of psychopathology will not come from MRI, a high-throughput genetic chip, or a sophisticated data-mining algorithm. While these will continue to provide new and suggestive leads – and may even ultimately provide crucial elements for diagnosis and prognosis – such technologies cannot bridge the final gap between the biological measurement and the fundamental experience of distress, threat, or craving that make up the core of psychopathology. Researchers who are well-studied in these experiences, and those with firsthand knowledge, will need to work on both sides of the ledger – with these new sources of data, but also with the broad array of people's thoughts, feelings, and experiences, to assemble the final pieces of the puzzle of mental illness.

REFERENCE

Amaro, E., Jr., & Barker, G. J. (2006). Study Design in fMRI: Basic Principles. *Brain and Cognition, 60*(3), 220–232.

Anderson, A. W., Heptulla, R. A., Driesen, N., Flanagan, D., Goldberg, P. A., Jones, T. W., … Gore, J. C. (2006). Effects of Hypoglycemia on Human Brain Activation Measured with fMRI. *Magnetic Resonance Imaging, 24*, 693–697.

Anderson, K. M., Krienen, F. M., Choi, E. Y., Reinen, J. M., Yeo, B. T. T., & Holmes, A. J. (2018). Gene Expression Links Functional Networks across Cortex and Striatum. *Nature Communications,* 9, 1428.

Baker, J. T., Holmes, A. J., Masters, G. A., Yeo, B. T. T., Krienen, F. M., Buckner, R. L., & Öngür, D. (2014). Disruption of Cortical Association Networks in Schizophrenia and Psychotic Bipolar Disorder. *JAMA Psychiatry, 71*(2), 109–118.

Barch, D. M., Burgess, G. C., Harms, M. P., Petersen, S. E., Schlaggar, B. L., Corbetta, M., … Van Essen, D. C. (2013).

Function in the Human Connectome: Task-fMRI and Individual Differences in Behavior. *NeuroImage, 80*, 169–189.

Biswal, B. B., Mennes, M., Zuo, X.-N., Gohel, S., Kelly, C., Smith, S. M., … Milham, M. P. (2010). Toward Discovery Science of Human Brain Function. *Proceedings of the National Academy of Sciences, 107*(10), 4734–4739.

Biswal, B. B., Yetkin, F. Z., Haughton, V. M., & Hyde, J. S. (1995). Functional Connectivity in the Motor Cortex of Resting Human Brain Using Echo-Planar MRI. *Magnetic Resonance in Medicine, 34*(4), 537–541.

Braver, T. S., Reynolds, J. R., & Donaldson, D. I. (2003). Neural Mechanisms of Transient and Sustained Cognitive Control during Task Switching. *Neuron, 39*(4), 713–726.

Breiter, H. C., Etcoff, N. L., Whalen, P. J., Kennedy, W. A., Rauch, S. L., Buckner, R. L., … Rosen, B. R. (1996). Response and Habituation of the Human Amygdala during Visual Processing of Facial Expression. *Neuron, 17*(5), 875–887.

Buckholtz, J. W., & Meyer-Lindenberg, A. (2012). Psychopathology and the Human Connectome: Toward a Transdiagnostic Model of Risk for Mental Illness. *Neuron, 74*(6), 990–1004.

Calhoun, V. D., Adali, T., Pearlson, G. D., & Pekar, J. J. (2001). A Method for Making Group Inferences from Functional MRI Data Using Independent Component Analysis. *Human Brain Mapping, 14*, 140–151.

Chekroud, A. M., Ward, E. J., Rosenberg, M. D., & Holmes, A. J. (2016). Patterns in the Human Brain Mosaic Discriminate Males from Females. *Proceedings of the National Academy of Sciences, 113*(14), E1968.

Ciric, R., Wolf, D. H., Power, J. D., Roalf, D. R., Baum, G. L., Ruparel, K., … Satterthwaite, T. D. (2018). Benchmarking of Participant-Level Confound Regression Strategies for the Control of Motion Artifact in Studies of Functional Connectivity. *NeuroImage, 154*, 174–187.

Cole, M. W., Repov, G., & Anticevic, A. (2014). The Frontoparietal Control System: A Central Role in Mental Health. *The Neuroscientist, 20*(6), 652–664.

Crossley, N. A., Mechelli, A., Vértes, P. E., Winton-Brown, T. T., Patel, A. X., Ginestet, C. E., … Bullmore, E. T. (2013). Cognitive Relevance of the Community Structure of the Human Brain Functional Coactivation Network. *Proceedings of the National Academy of Sciences of the United States of America, 110*(28), 11583–11588.

Dale, A. M., & Buckner, R. L. (1997). Selective Averaging of Rapidly Presented Individual Trials Using fMRI. *Human Brain Mapping, 5*(5), 329–340.

Deco, G., Jirsa, V. K., & McIntosh, A. R. (2011). Emerging Concepts for the Dynamical Organization of Resting-State Activity in the Brain. *Nature Reviews Neuroscience, 12*(1), 43–56.

Deng, C., Yuan, H., & Dai, J. (2018). Behavioral Manipulation by Optogenetics in the Nonhuman Primate. *The Neuroscientist;24*(5), 526–539.

Farkas, T., Wolf, A. P., Jaeger, J., Brodie, J. D., Christman, D. R., & Fowler, J. S. (1984). Regional Brain Glucose Metabolism in Chronic Schizophrenia: A Positron Emission Transaxial Tomographic Study. *Archives of General Psychiatry, 41*(3), 293–300.

Fox, P. T., Mintun, M. A., Reiman, E. M., & Raichle, M. E. (1988). Enhanced Detection of Focal Brain Responses Using Intersubject Averaging and Change-Distribution Analysis of Subtracted PET Images. *Journal of Cerebral Blood Flow and Metabolism, 8*(5), 642–653.

Fox, P. T., Parsons, L. M., & Lancaster, J. L. (1998). Beyond the Single Study: Function/Location Metanalysis in Cognitive Neuroimaging. *Current Opinion in Neurobiology, 8*(2), 178–187.

Friston, K.J., Frith, C. D., Liddle, P. F., & Frackowiak, R. S. J. (1993). Functional Connectivity: The Principal-Component Analysis of Large (PET) Data Sets. *Journal of Cerebral Blood Flow and Metabolism, 13*(1), 5–14.

Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J.-P., Frith, C. D., & Frackowiak R. S. J. (1994). Statistical Parametric Maps in Functional Imaging: A General Linear Approach. *Human Brain Mapping, 2*(4), 189–210.

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.

Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic Coding of Visual Space in the Monkey's Dorsolateral Prefrontal Cortex. *Journal of Neurophysiology, 61*(2), 331–349.

Ge, T., Holmes, A. J., Buckner, R. L., Smoller, J. W., & Sabuncu, M. R. (2018). Heritability Analysis with Repeat Measurements and Its Application to Resting-State Functional Connectivity. *Proceedings of the National Academy of Sciences, 114*(21), 5521–5526.

Gee, D. G., Gabard-Durnam, L. J., Flannery, J., Goff, B., Humphreys, K. L., Telzer, E. H., … Tottenham, N. (2013). Early Developmental Emergence of Human Amygdala-Prefrontal Connectivity after Maternal Deprivation. *Proceedings of the National Academy of Sciences of the United States of America, 110*(39), 15638–15643.

Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., … Van Essen, D. C. (2016). A Multi-Modal Parcellation of Human Cerebral Cortex. *Nature, 536*(7615), 171–178.

Goldstein, R. Z., Leskovjan, A. C., Hoff, A. L., Hitzemann, R., Bashan, F., Khalsa, S. S., … Volkow, N. D. (2004). Severity of Neuropsychological Impairment in Cocaine and Alcohol Addiction: Association with Metabolism in the Prefrontal Cortex. *Neuropsychologia, 42*(11), 1447–1458.

Goodwin, G. M. (1997). Neuropsychological and Neuroimaging Evidence for the Involvement of the Frontal Lobes in Depression. *Journal of Psychopharmacology, 11*(2), 115–122.

Gur, R. C., & Gur, R. E. (1995). Hypofrontality in Schizophrenia: RIP. *Lancet, 345*(8962), 1338–1340.

Haut, K. M., Lim, K. O., & MacDonald, A. (2010). Prefrontal Cortical Changes following Cognitive Training in Patients with Chronic Schizophrenia: Effects of Practice, Generalization, and Specificity. *Neuropsychopharmacology, 35*(9), 1850–1859.

Henson, R. (2006). Forward Inference Using Functional Neuroimaging: Dissociations versus Associations. *Trends in Cognitive Sciences, 10*(2), 64–69.

Holmes, A. J., & Patrick, L. M. (2018). The Myth of Optimality in Clinical Neuroscience. *Trends in Cognitive Sciences, 22*(3), 241–257.

Holmes, A. J., & Yeo, B. T. T. (2015). From Phenotypic Chaos to Neurobiological Order. *Nature Neuroscience, 18*(11), 1532–1534.

Holmes, A. J., Hollinshead, M. O., O'Keefe, T. M., Petrov, V. I., Fariello, G. R., Wald, L. L., … Buckner, R. L. (2015). Brain Genomics Superstruct Project Initial Data Release with Structural, Functional, and Behavioral Measures. *Scientific Data, 2*, 150031.

Huettel, S. A. (2012). Event-Related fMRI in Cognition. *Neuroimage, 62*(2), 1152–1156.

Huettel, S. A., Song, A. W., & McCarthy, G. (2004). *Functional Magnetic Resonance Imaging* (Vol. 1). Sunderland, MA: Sinauer Associates Sunderland.

Hutchison, R. M., Womelsdorf, T., Allen, E. A., Bandettini, P. A., Calhoun, V. D., Corbetta, M., … Chang, C. (2013). Dynamic Functional Connectivity: Promise, Issues, and Interpretations. *NeuroImage, 80*, 360–378.

Ingvar, D. H., & Franzen, G. (1974). Distribution of Cerebral Activity in Chronic Schizophrenia. *Lancet, 304*(7895), 1484–1486.

Kamiński, M., Ding, M., Truccolo, W. A., & Bressler, S. L. (2001). Evaluating Causal Relations in Neural Systems: Granger Causality, Directed Transfer Function and Statistical Assessment of Significance. *Biological Cybernetics, 85*(2), 145–157.

Koch, M. A., Norris, D. G., & Hund-Georgiadis, M. (2002). An Investigation of Functional and Anatomical Connectivity Using Magnetic Resonance Imaging. *Neuroimage, 16*(1), 241–250.

Kong, R., Li, J., Sun, N., Sabuncu, M. R., Schaefer, A., Scholz, M., … Yeo, B. T. T. (2019). Spatial Topography of Individual-Specific Cortical Networks Predicts Human Cognition, Personality and Emotion. *Cerebral Cortex, 29*(6), 2533–2551.

Krienen, F. M., Yeo, B. T. T., & Buckner, R. L. (2014). Reconfigurable Task-Dependent Functional Coupling Modes Cluster around a Core Functional Architecture. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences, 369*(1653), 20130526.

Lahey, B. B., Applegate, B., Hakes, J. K., Zald, D. H., Hariri, A. R., & Rathouz, P. J. (2012). Is There a General Factor of Prevalent Psychopathology during Adulthood? *Journal of Abnormal Psychology, 121*(4), 971–977.

Laird, A. R., Mickle Fox, P., Price, C. J., Glahn, D. C., Uecker, A. M., Lancaster, J. L., … Fox, P. T. (2005). ALE Meta-Analysis: Controlling the False Discovery Rate and Performing Statistical Contrasts. *Human Brain Mapping, 25*(1), 155–164.

Laurienti, P. J., Field, A. S., Burdette, J. H., Maldjian, J. A., Yen, Y.-F., & Moody, D. M. (2002). Dietary Caffeine Consumption Modulates fMRI Measures. *Neuroimage, 17*, 751–757.

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.

MacDonald, A. W., Becker, T. M., & Carter, C. S. (2006). Functional Magnetic Resonance Imaging Study of Cognitive Control in the Healthy Relatives of Schizophrenia Patients. *Biological Psychiatry, 60*(11), 1241–1249.

MacDonald, A. W. (2015). Differential Deficit. In R. Cautin & S. Lilienfeld (Eds.), *The Encyclopedia of Clinical Psychology* (1st edn.). Hoboken, NJ: John Wiley.

Mayberg, H. S. (1997). Limbic-Cortical Dysregulation: A Proposed Model of Depression. *Journal of Neuropsychiatry and Clinical Neurosciences, 9*(3), 471–481.

McIntosh, A. R., & Gonzalez-Lima, F. (1991). Structural Modeling of Functional Neural Pathways Mapped with 2-Deoxyglucose: Effects of Acoustic Startle Habituation on the Auditory System. *Brain Research, 547*(2), 295–302.

McKeown, M. J., & Sejnowski, T. J. (1998). Independent Component Analysis of fMRI Data: Examining the Assumptions. *Human Brain Mapping, 6*, 368–372.

Monti, M. M. (2011). Statistical Analysis of fMRI Time-Series: A Critical Review of the GLM Approach. *Frontiers in Human Neuroscience, 5*(28), 1–13.

Ollier, W., Sprosen, T., & Peakman, T. (2005). UK Biobank: From Concept to Reality. *Pharmacogenomics, 6*(6), 639–646.

Ollinger, J. M., Shulman, G. L., & Corbetta, M. (2001). Separating Processes Within a Trial in Event-Related Functional MRI I: The Method. *NeuroImage, 13*(1), 210–217.

Park, S., Holzman, P. S., & Goldman-Rakic, P. S. (1995). Spatial Working Memory Deficits in the Relatives of Schizophrenic Patients. *Archives of General Psychiatry, 52*(10), 821–828.

Penny, W. D., Stephan, K. E., Mechelli, A., & Friston, K. J. (2004). Modelling Functional Integration: A Comparison of Structural Equation and Dynamic Causal Models. *Neuroimage, 23*(Suppl. 1), S264–274.

Poldrack, R. A. (2006). Can Cognitive Processes Be Inferred from Neuroimaging Data? *Trends in Cognitive Sciences, 10*(2), 59–63.

Poldrack, R. A. (2010). Mapping Mental Function to Brain Structure: How Can Cognitive Neuroimaging Succeed? *Perspectives on Psychological Science, 5*, 753–761.

Poppe, A. B., Wisner, K., Atluri, G., Lim, K. O., Kumar, V., & MacDonald, A. W. (2013). Toward a Neurometric Foundation for Probabilistic Independent Component Analysis of fMRI Data. *Cognitive, Affective, & Behavioral Neuroscience, 13*(3), 641–659.

Poppe, A. B., Barch, D. M., Carter, C. S., Gold, J. M., Ragland, J. D., Silverstein, S. M., & MacDonald, A. W. (2016). Reduced Frontoparietal Activity in Schizophrenia Is Linked to a Specific Deficit in Goal Maintenance: A Multisite Functional Imaging Study. *Schizophrenia Bulletin, 42*(5), 1149–1157.

Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Barnes, K. A., Church, J. A., … Petersen, S E. (2011). Functional Network Organization of the Human Brain. *Neuron, 72*(4), 665–678.

Power, J. D., Mitra, A., Laumann, T. O., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2014). Methods to Detect, Characterize, and Remove Motion Artifact in Resting State fMRI. *NeuroImage, 84*, 320–341.

Price, J. L., & Drevets, W. C. (2012). Neural Circuits Underlying the Pathophysiology of Mood Disorders. *Trends in Cognitive Sciences, 16*(1), 61–71.

Raichle. M. E. (2009). A Brief History of Human Brain Mapping. *Trends in Neurosciences, 32*(2), 118–126.

Reinen, J. M., Chen, O. Y., Hutchison, R. M., Yeo, B. T. T., Anderson, K. M., Sabuncu, M. R., … Holmes, A. J. (2018). The Human Cortex Possesses a Reconfigurable Dynamic Network Architecture That Is Disrupted in Psychosis. *Nature Communications, 9*, 1157.

Richiardi, J., Altmann, A., Milazzo, A.-C., Chang, C., Chakravarty, M. M., Banaschewski, T., … Greicius, M. D. (2015). Correlated Gene Expression Supports Synchronous Activity in Brain Networks. *Science, 348*(6240), 1241–1244.

Rosenberg, M. D., Finn, E. S., Scheinost, D., Papademetris, X., Shen, X., Constable, R. T., & Chun, M. M. (2016). A Neuromarker of Sustained Attention from Whole-Brain Functional Connectivity. *Nature Neuroscience, 19*(1), 165–171.

Rosenberg, M. D., Casey, B. J., & Holmes, A. J. (2018). Prediction Complements Explanation in Understanding the Developing Brain. *Nature Communications, 9*, 589.

Salem, J. E., Kring, A. M., & Kerr, S. L. (1996). More Evidence for Generalized Poor Performance in Facial Emotion Expression in Schizophrenia. *Journal of Abnormal Psychology, 105*(3), 480–483.

Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X.-N., Holmes, A., … Yeo, B. T. (2018). Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. *Cerebral Cortex*, 28, 3095–3114.

Seifritz, E., Bilecen, D., Hänggi, D., Haselhorst, R., Radü, E. W., Wetzel, S., … Scheffler, K. (2000). Effect of Ethanol on BOLD Response to Acoustic Stimulation: Implications for Neuropharmacological fMRI. *Psychiatry Research Neuroimaging, 99*(1), 1–13.

Sepulcre, J., Liu, H., Talukdar, T., Martincorena, I., Yeo, B. T. T., & Buckner, R. L. (2010). The Organization of Local and Distant Functional Connectivity in the Human Brain. *PLoS Computational Biology, 6*(6), e1000808.

Shen, X., Tokoglu, F., Papademetris, X., & Constable, R. T. (2013). Groupwise Whole-Brain Parcellation from Resting-State fMRI Data for Network Node Identification. *Neuroimage, 82*, 403–415.

Shmueli, K., van Gelderen, P., de Zwart, J. A., Horovitz, S. G., Fukunaga, M., Jansma, J. M., & Duyn, J. H. (2007). Low-Frequency Fluctuations in the Cardiac Rate as a Source of Variance in the Resting-State fMRI BOLD Signal. *Neuroimage, 38*(2), 306–320.

Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Mickle Fox, P., Mackay, C. E., … Beckmann, C. F. (2009). Correspondence of the Brain's Functional Architecture during Activation and Rest. *Proceedings of the National Academy of Sciences, 106*(31), 13040–13045.

Smith, S. M., Vidaurre, D., Beckmann, C. F., Glasser, M. F., Jenkinson, M., Miller, K. L., … Van Essen, D. C. (2013). Functional Connectomics from Resting-State fMRI. *Trends in Cognitive Sciences, 17*(12), 666–682.

Smoller, J. W., Gallagher, P. J., Duncan, L. E., McGrath, L. M., Haddad, S. A., Holmes, A. J., … Cohen, B. M. (2014). The Human Ortholog of Acid-Sensing Ion Channel Gene ASIC1a Is Associated with Panic Disorder and Amygdala Structure and Function. *Biological Psychiatry, 76*(11), 902–910.

Sporns, O. (2014). Contributions and Challenges for Network Models in Cognitive Neuroscience. *Nature Neuroscience, 17*(5), 652–660.

Sprooten, E., Rasgon, A., Goodman, M., Carlin, A., Leibu, E., Lee, W. H., & Frangou, S. (2018). Addressing Reverse Inference in Psychiatric Neuroimaging: Meta-Analyses of Task-Related Brain Activation in Common Mental Disorders. *Human Brain Mapping, 38*(4), 1846–1864.

Tavor, I., Parker Jones, O., Mars, R. B., & Smith, S. M. (2016). Task-Free MRI Predicts Individual Differences in Brain Activity during Task Performance. *Science, 352*(6282), 216–220.

Thiel, C. M., & Fink, G. R. (2007). Visual and Auditory Alertness: Modality-Specific and Supramodal Neural Mechanisms and Their Modulation by Nicotine. *Journal of Neurophysiology, 97*(4), 2758–2768.

Van Dijk, K. R. A., Hedden, T., Venkataraman, A., Evans, K. C., Lazar, S. W., & Buckner, R. L. (2010). Intrinsic Functional Connectivity as a Tool for Human Connectomics: Theory, Properties, and Optimization. *Journal of Neurophysiology, 103*(1), 297–321.

Van Dijk, K. R. A., Sabuncu, M. R., & Buckner, R. L. (2012). The Influence of Head Motion on Intrinsic Functional Connectivity MRI. *Neuroimage, 59*(1), 431–438.

Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., & Ugurbil, K. (2013). The WU-Minn Human Connectome Project: An Overview. *Neuroimage, 80*, 62–79.

Vincent, J. L., Snyder, A. Z., Fox, M. D., Shannon, B. J., Andrews, J. R., Raichle, M. E., & Buckner, R. L. (2006). Coherent Spontaneous Activity Identifies a Hippocampal-Parietal Memory Network. *Journal of Neurophysiology, 96*(6), 3517–3531.

Viviani, R., Grön, G., & Spitzer, M., (2005). Functional Principal Component Analysis of fMRI Data. *Human Brain Mapping, 24*, 109–129.

Wang, D., Buckner, R. L., Fox, M. D., Holt, D. J., Holmes, A. J., Stoecklein, S., … Liu, H. (2015). Parcellating Cortical Functional Networks in Individuals. *Nature Neuroscience, 18*(12), 1853–1860.

Wisner, K. M., Atluri, G., Lim, K. O., & MacDonald, A. W. (2013). Neurometrics of Intrinsic Connectivity Networks at Rest Using fMRI: Retest Reliability and Cross-Validation Using a Meta-Level Method. *Neuroimage, 76*, 236–251.

Yarkoni, T., Poldrack, R. A., Nichols, T. E., Van Essen, D. C., & Wager, T. D. (2011). Large-Scale Automated Synthesis of Human Functional Neuroimaging Data. *Nature Methods, 8*(8), 665–670.

Yeo, B. T. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., … Buckner, R. L. (2011). The Organization of the Human Cerebral Cortex Estimated by Intrinsic Functional Connectivity. *Journal of Neurophysiology, 106*(3), 1125–1165.

Zang, Y., Jiang, T., Lu, Y., He, T., & Tian, L. (2004). Regional Homogeneity Approach to fMRI Data Analysis. *NeuroImage, 22*, 394–400.

Zhang, X., Mormino, E. C., Sun, N., Sperling, R. A., Sabuncu, M. R., & Yeo, B. T. T. (2016). Bayesian Model Reveals Latent Atrophy Factors with Dissociable Cognitive Trajectories in Alzheimer's Disease. *Proceedings of the National Academy of Sciences of the United States of America, 113*(42), E6535–E6544.

Zuo, X.-N., Anderson, J. S., Bellec, P., Birn, R. M., Biswal, B. B., Blautzik, J., … Milham, M. P. (2014). An Open Science Resource for Establishing Reliability and Reproducibility in Functional Connectomics. *Scientific Data, 1*, 140049.