CHAPTER 34

GROUP STUDIES IN EXPERIMENTAL NEUROPSYCHOLOGY

Lesley K. Fellows

WHY NEUROPSYCHOLOGY?

A fundamental assumption of neuropsychology, and of cognitive neuroscience more generally, is that behavior has a biological basis—that it results from processes that are executed in the nervous system. Following from this assumption, emotions, thoughts, percepts, and actions can be understood in neurobiological terms. This premise was advanced by the philosophers of ancient Greece, supported, in part, by observations of patients with brain injury (Gross, 1995). The fact that damage to the brain could lead to paralysis, disorders of sensation, or even disruptions of consciousness suggested that this organ was *the seat* of such abilities, although the broad claim that brain function underlies behavior was not without controversy over the centuries that followed (Crivellato & Ribatti, 2007).

The 19th century saw, on the one hand, major developments in understanding the anatomy and physiology of the brain and, on the other, more systematic descriptions of behavioral changes resulting from neurological diseases. These advances laid the groundwork for current thinking about the brain. Here again, clinical observations provided an important impetus, as did analyses of individual differences in normal behavior: Neurologists such as Paul Broca and Carl Wernicke reported that focal brain injury to specific areas within the left hemisphere disrupted particular aspects of language (Feinberg & Farah, 2006). Their thinking was influenced, in part, by Franz-Joseph Gall and others who developed the concept of phrenology in about the same period. Phrenology was based on observations of specific individual differences in skull shape (explicitly thought to be a proxy for underlying brain structure) in relation to individual differences in behavior. Complex traits like *benevolence* and *wit* were thus related to particular parts of the brain. Although the methods are clearly flawed to the eye of the modern reader, the underlying concept of localization, that brain structure and function are related, had a major impact on the development of clinical neurology and of experimental neuropsychology.

The work of Broca, Wernicke, and other 19th- and early 20th-century neurologists illustrated how observation in clinical populations can (a) provide insights into how a complex behavior (like language) can be segmented into simpler components (e.g., production and comprehension) and (b) how such components can be related to specific regions of the brain. Both defining the components of behavior and relating these components to the brain can be done on the basis of a single, carefully studied case (see Chapter 33 of this volume). However, the limitations of clinical observations in humans were also apparent in these early days. Clinicians were (and are) acutely aware of wide variability in the clinical presentation of a particular pathological condition, determined both by differences in premorbid individual characteristics (e.g., age, education, or health status) and differences in the specific details of the pathological process. Case series and group studies provide an important means of determining the generalizability of inferences that can be drawn from individual observations.

DOI: 10.1037/13620-034
APA Handbook of Research Methods in Psychology: Vol. 2. Research Designs, H. Cooper (Editor-in-Chief)

Gordon Holmes, a British neurologist whose work on the effects of penetrating brain injury in World War I soldiers helped to establish the retinotopic organization of primary visual cortex, poetically captured the limitations of clinical observation in a lecture delivered in 1944:

> My own work on the visual cortex has been limited to observation in man. . . . This has required the collection of a large number of observations, for while the physiologist can rely on experiments when he can select and control, . . . the clinician must depend on the analysis of observations which are rarely so simple or clear cut.. . . The physiologist may be compared with the builder in . . . hewn stones which can easily be fitted together, the physician resembles the mason who has to use irregular rubble and therefore requires more time and labour to attain his end. But in some branches of neurology, the "rubble" collected and put together by the clinician is essential. (Holmes, 1944/1979, pp. 440–441)

Holmes underlined two key points: (a) that the limitations inherent to studying the effects of brain injury in humans can be minimized by gathering data from many subjects and by interpreting these data in the context of converging evidence from other methods and (b) that the limitations are offset by the fact that these observations provide crucial insights that may not be acquired in any other way. As this chapter will describe, there have been many logistical, technical, and analytic advances in human lesion studies over the past century. However, Holmes's comments on the core advantages and limitations remain as pertinent as ever.

INFERENTIAL STRENGTHS OF LESION STUDIES

Research on effects of brain injury on behavior addresses two main issues: First, it can establish that a particular region of the brain is necessary for the expression of a particular behavior, in turn, supporting the inference that it is critical for a particular cognitive process (Fellows et al., 2005; Rorden & Karnath, 2004). In principle, this is a powerful form of evidence because it addresses causality. Although cognitive neuroscience now has many other methods available to investigate brain–behavior relations, most provide correlational data. Standard functional neuroimaging methods, for example, reveal brain regions in which blood-oxygen-level dependent (commonly referred to as BOLD) signal (itself a correlate of neural activity) is correlated with a behavioral process of interest. These findings can be informative, but alone they are insufficient to establish that the brain regions so identified are in fact necessary for the behavior in question (Fellows et al., 2005; Rorden & Karnath, 2004).

These inferential considerations are particularly relevant in the study of complex behaviors, and in new areas of enquiry. Consider risky decision making as an example. Imagine yourself at the blackjack table, deciding how much to stake on the next card. Several correlated processes are likely under way in your brain. You may be calculating the odds of winning, integrating your recent history of wins and losses, and weighing these factors to reach a decision. You may be imagining how you would spend your winnings, or how you would explain a loss to your spouse. It is likely that you are experiencing substantial changes in arousal and autonomic tone: A pounding heart and sweaty palms often accompany a risky choice. Whether all of these putative processes are distinct, important to the decision, or simply correlated epiphenomena are empirical questions. Interpreting functional magnetic resonance imaging (fMRI) activations in this situation is not easy—for example, is a given area more active because it is critically involved in risky decision making or is it important in central autonomic control, mediating the changes in sympathetic nervous system outflow that result in the pounding heart? Although careful design can help to minimize these uncertainties of interpretation, the nature of correlational evidence means that they can never be eliminated entirely. Converging evidence from loss-of-function methods, such as lesion studies, can help test necessity claims. If we take a hypothetical "risky decision" brain area as an example, a study of patients with damage to that area could directly test

whether it was critical for the decision, for the autonomic changes that accompany that decision, or for both. Such an experiment would shed light both on the critical components of decision making (do autonomic changes influence choice?) and on the brain substrates of the critical processes (e.g., Critchley et al., 2003).

More profoundly, the study of patients can provide biological constraints to psychological theory. The usual form this has taken is that of dissociation of cognitive processes. Two putative psychological constructs may be considered distinct if brain injury disrupts one and not the other—establishing what is termed a *function dissociation*. As will be described, experiments of this kind have been influential, but how these are best designed and interpreted is not without controversy.

WHAT DO GROUP STUDIES ADD TO THE ANALYSIS OF SINGLE CASES?

Group studies address two potential problems of interpretation that plague single cases: One is that observed deficits in a single patient may be due to premorbid differences in function—normal individual differences may thus be misattributed to the lesion. This may be implausible for some deficits: Common sense dictates that major hemiparesis or visual field defects are outside the range of normal variation and can generally be safely linked to the brain injury. But other aspects of behavior, such as executive functions, emotional, or social processes, may differ substantially across healthy individuals, making it more likely that such a difference will be found by chance in a brain-injured patient. Idiosyncrasies in brain organization, in structure–function mapping, or in recovery from brain injury can also contribute to exceptional performance in a single case.

Even if we can safely assume that the patient's brain, function, and brain–function relations were representative before the injury, a second source of variability would make group studies important. For obvious reasons, brain lesions in human subjects are not under experimental control. As a result, there is substantial inherent variability in the extent and causes of brain injury. Group studies help to exclude potential lesion-related confounds that might explain the observations in a single case; for example, they can establish that it is the site of damage, rather than its etiology, that underlies the behavior change.

Another common and related problem that can be addressed by group studies is that lesions are often more extensive, or less precisely located, than is ideal for testing a given structure–function hypothesis. If the function is disrupted but the lesion is large, the conclusions cannot be specific. If a group of patients with lesions varying in extent but overlapping in some smaller area are found to have a common impairment in function, one can infer that the function likely relies on the region of overlap that is common across patients. Recent methods have built on this logic to allow statistical tests of structure–function relations at the voxel level and will be discussed in more detail later in this chapter. Thus, lesion extent can limit structure–function mapping in single case studies but can at least be addressed, and maybe turned to advantage, in group studies.

Group studies in neuropsychology are strictly observational rather than experimental. Like case-control studies in epidemiology, they are vulnerable to confounds and biases, but they can nevertheless offer important insights. These biases are predictable and generally can be avoided with careful design or addressed with appropriate analyses. The observational nature of this approach to lesion–function mapping means that there is no imposed directionality: Studies may begin from either lesion or function. These two perspectives are discussed in turn.

DESIGNS DRIVEN BY BEHAVIOR

A major challenge in both psychology and cognitive neuroscience is to define the architecture of behavior. One way or another, the complexity of behavior needs to be parsed into analyzable constituents, whether these are conceptualized as modules, processes, or interacting networks (Dunn & Kirsner, 2003). The challenge is to identify the appropriate constituent parts and then to understand how they interact from both a psychological and a neural point of view.

Arguably, this enterprise has been most successful when it has been closely linked to neurobiology. For example, we now have a detailed understanding of visual processing that begins from response properties of single neurons in the retina, moves to how these are combined in the initial stages of cortical visual processing, and continues from there to the computations that support object or face recognition (Farah, 2004; Van Essen, Felleman, DeYoe, Olavarria, & Knierim, 1990). This enterprise obviously requires data gathered with a variety of methods. Studies in patients with brain injury can provide important insights into the biologically relevant lines of cleavage for a given (complex) behavior, by helping to identify associations and dissociations between putative component processes.

When behavior is treated as the independent variable, patients are selected on the basis of the presence of some behavioral manifestation—either a clinical syndrome or performance on a particular task. Additional behavioral measures aiming to isolate putative component processes are then administered to determine whether these processes are, in fact, distinct (i.e., dissociable). A single dissociation refers to a situation in which subjects are impaired on a task that presumably assesses a particular ability but are unimpaired on another task that assesses a separate ability. Single dissociations are evidence in favor of a hypothesis that the tasks measure distinct component processes (Damasio & Damasio, 1989; Shallice, 1988). However, there are practical issues that make alternative explanations for such patterns quite likely: As one example, dissociations assume that the tasks being used are approximately equally difficult. An easy task and a hard task tapping the same component process would show apparent dissociation because at least some patients would fail the hard task but pass the easy task (Shallice, 1988).

This potential explanation is less likely if a double dissociation can be demonstrated: Here, one set of patients fails Task A but does well on Task B, whereas another set shows the opposite pattern. The explanatory power and experimental elegance of double dissociation has been a touchstone since the early days of experimental neuropsychology (Teuber, 1955).

How Do You Know a Dissociation When You See One?

The logic of dissociation is clear in principle but can be challenging to operationalize in practice. (See Dunn & Kirsner, 2003, for a more detailed analysis of these challenges.) How intact must a group be in Task A? How impaired in Task B? What is the likelihood of such dissociations occurring by chance, in a given population, and for any given pair of tasks? One common approach is to test for a crossover interaction in the performance of two tasks, across two groups, but other patterns may be as or more important, depending on the relations between the tasks and on the relations between a given cognitive process and performance on the task that is meant to measure it (Bates, Appelbaum, Salcedo, Saygin, & Pizzamiglio, 2003; Dunn & Kirsner, 2003; Shallice, 1988).

Conceptual Precision

An important first step in any experiment of this sort is to start from a position of conceptual clarity: A model of the component processes of interest that is well-justified will dictate the appropriate analyses. A priori hypotheses might come from existing experimental work in humans using lesion or other methods, from animal studies, from computational models, or from combinations of these sources.

Measurement Reliability

Once processes of interest are identified, tasks are needed to measure the relevant behavior as specifically as possible. Ideally, such measures will have good psychometric properties: no ceiling or floor performance, good test–retest reliability, and performance that is minimally influenced by demographic or education factors (Laws, 2005). Brain-injured patients are typically older and less educated, on average, than the convenience samples of healthy undergraduates often used in the development of new measures. Because the time and energy of these patients are limited, it is wise to pilot new tasks in healthy subjects who are otherwise demographically similar to the target patient population. That said, the appropriate reference population for the actual experiment may not be healthy subjects. Depending on the hypothesis, patients with brain injury may

provide more relevant comparison data, and such comparisons may be less affected by the ceiling effects that can be a problem in healthy reference groups.

Measurement variation, that is, the extent to which task performance will vary if the same subject is tested repeatedly, is a source of noise that in principle is under the experimenter's control. It may have important influences on the analysis and should be minimized to the extent possible (Bates, Appelbaum, et al., 2003). In addition to piloting in demographically relevant healthy populations, attention needs to be given to particular challenges that may arise in patient populations. Depending on the patient population of interest, relevant issues might include (a) difficulty understanding instructions, (b) difficulty with motor or perceptual aspects of the tasks that are related to the lesion but not of interest (e.g., because of weakness interfering with responding, or disruption of primary sensory processing), and (c) nonspecific changes in arousal or attention related to the injury or to psychoactive medications (e.g., anticonvulsants) that may be more commonly taken by those in the target group than in the reference group. Some of these issues can be addressed in the experiment. For example, patients may need simplified instructions, additional practice, or modifications of how stimuli are presented or responses collected. These problems also apply to single case studies, but their solutions may be different in group studies. In single cases, there may be more flexibility in optimizing the details of the task to accommodate patient-specific factors. In groups, there is a trade-off between using a consistent measure across all participants (and so allowing the results to be easily pooled) and adapting the task to individual restrictions.

Interindividual Variability

A second source of variability relates not to the measurement tools but rather to the individuals being measured. Individual differences in group lesion studies can be conceptualized as arising from three potential sources, and these differences can be of no interest or of major interest. The first source is individual differences of the same sort that one finds in healthy populations. People differ in their cognitive capacities. This may be particularly true for certain cognitive capacities. This variation is generally only a nuisance in lesion studies, increasing the variance across both experimental and reference groups. It can become a confound, however, if such individual differences are not randomly distributed across groups. This can be due to sampling error, systematic sampling bias, or nonindependence of lesion-related variables. The simplest form of sampling error is that, by chance, more subjects from one end of the normal range are present in one group than another. This risk is minimized by increasing the sample size (indeed, avoiding this risk is an important motivation of group versus single case studies), but patient studies have practical limits on sample size that make this a challenge.

Sampling bias may occur for other reasons. For example, some normal individual difference may also make it more likely that an individual would suffer a particular neurological injury. This problem is illustrated, for example, in studying the links between impulsivity and the frontal lobes: If patients who have suffered frontal lobe damage from traumatic brain injury are found to be more impulsive, does this establish that the frontal lobes are important in impulse control? Or do impulsive people get into situations in which they suffer such injuries more often than the less impulsive, so that this normal individual difference ends up overrepresented in the patient group?

A second source of individual variability relates to inevitable variation in the nature and extent of brain injury within a group. Furthermore, variation in lesion location is not the only lesion-related determinant of this kind of variability: Factors such as comorbidity and medication use may differ systematically with lesion etiology or location and so be another source of bias.

Do Lesion Data Matter in Behavior-Driven Designs?

One can test the hypothesis that two processes are functionally dissociable in a patient population without ever considering the details of their lesions. In principle, dissociation, particularly double dissociation, addresses the issue. In practice, however,

there are many nuances in determining what the thresholds might be for establishing dissociations, including the need to consider departures from correlations across tasks as well as (or instead of) absolute performance in each of two tasks. Even if the experimental goal is purely to understand the architecture of cognitive processes, rather than their relation to the brain, lesion analysis can provide external validation of claims of dissociation. Consistent lesion location–function mapping bolsters the argument that what impaired (or unimpaired) patients have in common is disruption of a specific system, rather than some demographic or task-related confound (Robertson, Knight, Rafal, & Shimamura, 1993).

LESION-DRIVEN DESIGNS

It is equally possible to study structure–function relations in the human brain with the brain injury treated as the independent variable. Rather than aiming to discern how behavior can be dissected, the starting point is to determine the cognitive processes for which a given brain region is necessary. Of course, these two aims converge on the same central questions.

Characterizing Lesions

In the early days of neuropsychology, lesion characterization was based on neurosurgical sketches, plain X-rays of the skull, or the results of autopsies. Computerized tomography (CT) and, more recently, structural magnetic resonance imaging (MRI), have dramatically improved the quality of anatomical data.

The first step in characterizing lesions is thus to acquire either MRI or CT images of each patient's brain. MRI is preferred because it offers better resolution, and in many cases better sensitivity, than CT, and it avoids exposing the participant to ionizing radiation. However, MRI may be contraindicated in patients with pacemakers or surgical clips, for example, or not tolerated because of claustrophobia. Ideally, high-resolution imaging should be acquired in the whole patient sample using standard parameters and equipment, as close to the time of behavioral testing as possible. That said, it may be much more practical to use the most recently available clinical imaging. This is less resource intensive, minimizes patient inconvenience, and often provides lesion data that are of more than adequate resolution for testing a given hypothesis. Regardless of approach, the quality (and so anatomical precision) of the lesion characterization needs to be considered in the analysis.

The simplest way of presenting lesion data is to reproduce the imaging as-is for each patient. This works well for single cases but becomes awkward in group studies. Indeed, it is only appropriate to reproduce individual scans if the behavioral data are also presented for each individual, that is, in case series format. If behavioral data are presented as group means, imaging data also need to be presented in a form that allows insights into what is common in the group. This can be achieved simply—for example, by tabulating the number of subjects with damage to particular regions. However, modern imaging data are acquired in digital form, permitting group lesion data to be presented as brain images that are much more accessible, and also are more easily related to the wider literature, particularly fMRI studies.

Individual lesions first need to be represented in a common space. This can be achieved in two main ways: either by manually tracing the lesion onto some common template (Damasio & Damasio, 1989; Kimberg, Coslett, & Schwartz, 2007) or by manually or automatically defining the lesion on the individual patient's anatomical scan and then warping the brain (and the lesion) onto a standard template. The first method is labor intensive and requires substantial expertise. The second method relies on the same algorithms used to warp individual scans into common space for fMRI analysis in healthy subjects and can be more automatized. However, the anatomical distortions caused by the presence of a lesion lead to particular technical issues that need to be addressed thoughtfully if this approach is taken (Nachev, Coulthard, Jager, Kennard, & Husain, 2008; Rorden & Brett, 2000). Regardless of approach, defining the boundaries of lesions always involves some judgment and so is a potential source of error.

Registering individual lesions to a common template allows these data to be shown in aggregate—most commonly as overlap images (generated by

representing the arithmetic sum of damage in each voxel, across the group), which show the degree to which damage affects common brain structures for a given group of patients (Frank, Damasio, & Grabowski, 1997; Makale et al., 2002; Rorden & Brett, 2000). Such images can also demonstrate the absence of common damage in two groups that are meant to be anatomically distinct. Digitized lesion data that are represented in a common space can be used in more complex computations, including statistical tests of structure–function relations (see the section Finer Grained Lesion–Symptom Mapping), and are more readily linked to other sources of data that are also expressed in common brain coordinates, notably fMRI studies.

REGION OF INTEREST DESIGNS

When there is an a priori hypothesis about the functional role of a particular brain region, region of interest (ROI) designs are appropriate. Here, participants are identified on the basis of the presence of damage affecting (or sometimes restricted to) some specified region of the brain. Behaviors of interest are measured with one or several tasks, and performance is compared with appropriate reference groups. If impairment is identified, it is evidence that the brain region plays a necessary role in task performance and, by extension, in the cognitive process of interest. The major advantage of this approach is its hypothesis-driven design and the statistical power that accompanies such designs. This power means that relatively small sample sizes may be adequate, particularly because effect sizes in lesion studies are often quite large. Such designs may have directional hypotheses, making one-tailed statistical tests appropriate.

There are several important design issues to consider in these focused studies. The first is the appropriate reference group. One common approach is to compare participants with damage to a particular region to a healthy group, matched on demographic characteristics. This provides some control over potential demographic confounds (although perhaps not as much as one might think, depending on the sample size and the variance in these demographic characteristics). However, one cannot unequivocally conclude that the effects are due to damage in a particular region. They may relate to some effect of brain damage more generally, including effects of confounding factors that may be more common in those with brain injury than in those without. Furthermore, it may be difficult to avoid ceiling effects in a healthy control group. To address these problems, many studies include a comparison group with brain injury that spares the region of interest. If the aim is to exclude generic effects of brain damage (or confounds more likely to be present in ill than in healthy participants) in the interpretation of the findings, then any site of damage that spares the region in question is fine. This is something of a missed opportunity, however. If the lesioned comparison group is selected so that the lesions affect a second, specific brain region, then that group can serve double duty: both controlling for nonspecific effects of injury, and establishing that the other region is not necessary for the process in question. Such a targeted approach also assesses the possibility that the lesioned reference is impaired on some other task, providing insurance against the claim that the reference group is somehow less impaired for whatever reason. In the end, one is left with a focused double dissociation.

This elegant design is not easy to achieve and, when achieved, is still potentially susceptible to the problems described for functional dissociations. The main practical difficulty is in recruiting an appropriate lesioned comparison group, whether it involves patients with anatomically common or disparate lesions, matched to the experimental group on both clinical and demographic variables. Systematic recruitment methods, such as patient registries, can make this more feasible (Fellows, Stark, Berg, & Chatterjee, 2008).

If it is impossible to recruit a lesioned comparison group, then the next best approach is to thoroughly characterize the relation between demographic variables and task performance in a healthy reference group, in which adequate sample sizes are much more feasible, and then use that information to inform analyses of the patient data. For example, demographic characteristics that differ between patient groups can be shown not to substantially influence task performance in a large

healthy reference group, or the influence can be characterized sufficiently to allow these contributions to be covaried out in the primary analysis. A common approach along these lines is to express performance of each lesioned participant as a percentile or *z* score on the basis of performance in a larger healthy reference sample (e.g., Gläscher et al., 2009; Tsuchida & Fellows, 2009). Although desirable, this may not always be feasible, depending on the reference data available for a given task.

It is important to consider what an ROI design does not do: It does not necessarily impose a true anatomic boundary. Also, it is obvious that nothing will be learned about the potential contributions of brain regions outside that boundary. Perhaps less obviously, there can be a risk of not detecting effects that are, in fact, related to damage within the boundary. This can happen if the region is much larger than the actually critical brain area; effects caused by damage in the smaller area are diluted by normal performance in those with damage affecting the larger, but not critical, area. The group as a whole will have variable performance, and the statistical analyses may fail to detect effects. Finally, effects that are detected with a given ROI may nevertheless have been better captured by a different anatomical boundary.

STATISTICS FOR ROI DESIGNS

When the data are considered as group means, the same statistical approaches used for comparing groups in any study are appropriate. ROI designs commonly are limited to small samples and may involve skewed behavioral data (either because of ceiling effects in the control group, or substantial variability in the patient group, or both). These issues obviously need to be taken into account when planning the analysis, if they cannot be avoided in the design. Sometimes group studies are better analyzed as a series of single cases. This approach may be suitable when the group of patients varies widely on relevant demographic or other variables, or indeed, in task performance. Sometimes this approach is taken post hoc, in which case, the results should be considered with particular caution, given the ease with which confounds other than lesion location may explain observed effects.

PITFALLS IN ROI DESIGNS

Recruitment Bias

The observational nature of human lesion studies, in general, requires particular care to minimize potential bias. When testing a structure–function hypothesis with an anatomical ROI design, efforts should be made to include all subjects with damage to the region in question. It is common to undertake what might be called a *hybrid* study—for example, including patients with both left hemisphere damage *and* aphasia and then asking a more specific question about language processes. This runs the risk of distorting the results. At the least, it may magnify apparent structure–function relations, by picking desired patients with both lesion and dysfunction. It may give spurious findings as well because subjects without impairment provide important constraints in lesion–function mapping (Rorden, Fridriksson, & Karnath, 2009). At the other end of the spectrum, patients may be excluded because they are too impaired to perform the experimental tasks. This is unavoidable but important to keep in mind. For example, it might be difficult to study the neural processes related to the control of behavior. If damage to some key structure resulted in severe agitation, such patients are unlikely to be approached (and if approached, to agree) to participate in cognitive neuroscience research. Similarly, severe aphasia often precludes informed consent, so such patients may be systematically excluded.

Control Groups

As in epidemiologic case-control studies, the reference group is important in lesion studies. In principle, those in the control group should differ from the patient group only in that they have not suffered a brain injury. Subjects drawn from the same population are optimal. Practically, this is often challenging to achieve, but it can be accomplished in several ways. If healthy subjects are needed, then individuals with similar demographic profiles should be recruited, perhaps even friends or family members of the patients. Patients with damage to other brain regions (but due to the same causes as damage in the group of interest) are often better choices because they are more likely to be matched on

potentially confounding variables, such as medication usage, or nonspecific psychological effects of serious illness. However, depending on the size of the groups and the anatomical specificity of the hypotheses, it can be harder to match such subjects on demographic variables.

All Lesions Are Not Created Equal

Lesions do not occur at random. There are systematic biases in who suffers a brain injury, in the extent to which an injury that affects one part of the brain will be accompanied by damage to other parts of the brain, in the destructiveness of a given injury, and in the time course and mechanisms of recovery from that injury. For example, ischemic stroke damages parts of the brain that are supplied by particular blood vessels. These vascular territories mean that damage to one area, for example, inferior frontal lobe, will be more commonly associated with damage to another area in the same territory (e.g., insula, inferior parietal lobe). Conversely, such damage will almost never be associated with damage to the areas that are supplied by other blood vessels, such as the other hemisphere, the frontal pole, or the occipital lobe. Furthermore, some vascular territories are more commonly affected than others. Injury to the areas supplied by the middle cerebral artery, for example, will be overrepresented in a given series of unselected stroke patients. These regularities have implications for interpreting the results of lesion studies (Rorden & Karnath, 2004) and constrain the brain regions that can be readily studied by lesion methods.

Lesion etiology can also affect the accuracy of the lesion mapping, and the observable structure–function relations: Slow-growing tumors push normal brain tissue aside without necessarily disrupting function, which can lead to lesions that appear quite large but have much milder functional effects. Cortical resections in epilepsy have precise margins and spare the white matter—two advantages—but the cortex that is resected is often not normal and may have been abnormal for a long time.

Relatedly, the degree to which compensation can occur depends on the time course over which the brain disorder develops, its extent, and the time since injury. There is no doubt that the functional effects of stroke evolve over time (Fruhmann Berger, Johannsen, & Karnath, 2008). The effects of brain damage can be studied at any time point, but this is a highly relevant variable and must be considered in both study design and interpretation. The development of MRI sequences that can delineate ischemic brain tissue very shortly after the onset of acute stroke has provided the opportunity to do hyperacute lesion–function mapping (Marsh & Hillis, 2008; Newhart, Ken, Kleinman, Heidler-Gary, & Hillis, 2007). Such work can identify regions that are normally necessary for a given function. In contrast, studies examining chronic impairments after brain damage are perhaps better thought of as identifying regions that are necessary for the recovery of a given function: Deficits still present months or years after an injury are, by definition, resistant to compensatory mechanisms.

Finer Grained Lesion–Symptom Mapping

Regions of interest can, in principle, be any size. In practice, there is a lower limit of resolution imposed by the volume of brain tissue that is injured in individual subjects, the extent to which those volumes overlap in a given sample, or the resolution of the imaging methods that are used to characterize the injury. Lesion volume, rather than imaging resolution, is typically the limiting factor in the MRI era. The upper limit of resolution is determined by conceptual issues; determining that some function is related to the integrity of the whole brain, for example, is likely to be of limited interest. That said, many core concepts in neuropsychology began with regions of interest encompassing entire cerebral hemispheres, and defining such broad structure–function relations may still be important as cognitive neuroscience tackles new areas of study, such as in social or affective domains.

Converging evidence argues that structure–function relations are considerably more discrete than is captured by examining hemispheric, or even lobar, effects. There are practical limits to the regional specificity that can be attained in group studies with ROI designs. If the study is restricted to patients with damage to some specific and small brain area, an adequate sample is unlikely to be recruited in a reasonable time. An alternative is to

enroll patients with variable damage to a relatively broad region—even one hemisphere or the whole brain—and then undertake analyses to establish which subregion contributes to the observed deficits in function.

There are three main approaches to analyzing data from patients who have variable damage in a large brain region. The one with the longest history involves a secondary analysis in a standard ROI study. Having established that some anatomically defined group is impaired and observing the usual variability in that impairment, one may ask whether there is an anatomical basis to that variability. That is, whether damage to a specific subregion is a main determinant of task performance. This can be addressed qualitatively by examining the pattern of lesions in the impaired and unimpaired subgroups, in essence, carrying out a behavior-driven analysis nested in the original ROI study. Lesion overlap methods are often used to this end: The lesion overlap image for impaired and unimpaired subgroups can be examined visually, or lesion extent can be subtracted across these groups to identify the potentially critical subregion (e.g., Milne & Grafman, 2001). Alternative but analogous methods include tabulating the presence or absence of injury to Brodmann areas and comparing the outcome in patients with and without behavioral impairment.

There are drawbacks to this approach. First, it is important to realize that it is usually undertaken post hoc. Any finding needs confirmation in a new experiment that is designed to test the specific ROI a priori. Selection bias and confounding factors can easily influence the results. Such subregion analyses usually involve very small sample sizes, and it can be impossible to properly account for other (e.g., demographic) contributors to observed effects. This approach is also prone to problems because of the nonindependence of damage (see the section All Lesions Are Not Created Equal). Results from such analyses should be treated with particular caution when the a priori, ROI-based analysis does not establish significant differences between groups. A multiple ROI approach can be applied a priori. Several studies have taken this approach (e.g., Picton et al., 2007; Stuss, Murphy, Binns, & Alexander, 2003). The main difficulty, beyond the perennial limitations of sample size, is in determining how to appropriately correct for multiple comparisons.

Recently, statistical methods that were developed for fMRI have been adapted for examining structure–dysfunction relations at a voxel-by-voxel level. This is a natural extension of multi-ROI designs, with the (potential) advantage of principled control of multiple comparisons. Once lesions volumes are registered to a common template, univariate statistics can be applied to test whether the performance of patients with damage to a given voxel differs from performance of patients with damage that spares the voxel. This results in a statistical map showing the strength of association between damage and dysfunction in anatomical space.

This approach, commonly referred to as voxel-based lesion–symptom mapping (VLSM), does not require imposing potentially arbitrary (or somehow "wrong") ROI boundaries and allows task performance to be considered either as a dichotomous (intact–impaired) or continuous variable. The latter avoids having to impose a second, potentially arbitrary, boundary on the data. VLSM also has the potential to map networks, that is, to identify several regions that may contribute to task performance within a single experiment. Several variations of this method, using different statistical approaches, have been developed (see Bates, Wilson, et al., 2003; Chen, Hillis, Pawlak, & Herskovits, 2008; Kinkingnéhun et al., 2007; Rorden et al., 2009; Rorden & Karnath, 2004; Rorden, Karnath, & Bonilha, 2007; Solomon, Raymont, Braun, Butman, & Grafman, 2007).

These advantages come with trade-offs. As with fMRI analysis, this massively univariate approach requires conservative correction for multiple comparisons, which in turn demands a substantial sample size. The number of subjects is not the only consideration; lesion overlap and distribution are also important determinants of a study's power. Methods exist to estimate the anatomical extent of adequate power in a given sample, and this is an important adjunct in interpreting VLSM analyses (Kimberg et al., 2007; Rudrauf et al., 2008). Systematic approaches to patient recruitment are also critical in acquiring a suitable sample size, and in ensuring that the sample has been appropriately characterized (Fellows et al., 2008).

White Matter Damage and Disconnection Effects

With the exception of certain neurosurgical resections, lesions are rarely confined to a single structure and often disrupt the white matter leading into or away from a given gray matter region or fibers of passage (i.e., adjacent tracts that may have nothing to do with the damaged gray matter beyond physical proximity). This can pose challenges in interpreting lesion studies. Observed behavioral effects might be due to the white matter damage, which would be particularly misleading if it involves fibers of passage. Modern neuroimaging can assess the extent of white matter injury, either in standard structural scans, or by using tract-specific imaging such as diffusion tensor imaging. Furthermore, white matter atlases are becoming increasingly sophisticated. Thus, methods exist to address possible white matter contributions and are beginning to be applied to structure–function mapping (Catani, Jones, & ffytche, 2005; Karnath, Rorden, & Ticini, 2009; Philippi, Mehta, Grabowski, Adolphs, & Rudrauf, 2009; Rudrauf, Mehta, & Grabowski, 2008; Thiebaut de Schotten et al., 2008; Urbanski et al., 2008).

Developments in image analysis to study network properties of the brain, whether captured by structural or functional measures, may prove useful as adjuncts to the lesion approaches discussed so far (Dosenbach, Fair, Cohen, Schlaggar, & Petersen, 2008; He, Dagher, et al., 2009; He, Wang, et al., 2009). At the least, these techniques draw attention to network-oriented conceptual frameworks.

Clinical Conditions With Diffuse Damage

Brain–behavior relations can also be studied in clinical conditions that have multifocal or diffuse damage. Traumatic brain injury, multiple sclerosis, and degenerative dementias are examples. Imaging methods can quantify regional cortical and white matter changes, even when these are subtle or diffuse, and such changes can be correlated with behavior. Most of the pitfalls that have been discussed also apply to such studies. There are additional challenges in interpreting anatomical data when multiple areas are dysfunctional in a more or less correlated (and more or less detectable) way, and in interpreting behavioral data when multiple cognitive functions that may be necessary for a given task are also degraded in more or less correlated ways.

CONCLUSION

Studies of disrupted function can provide important insights into the architecture of cognitive processes and can identify the brain substrates critical for these processes. Lesion studies in humans have particular inferential strengths, explaining their long and fruitful history in neuroscience and psychology, and recent advances in anatomical imaging and statistical analysis contribute to the continued relevance of such work (Catani & ffytche, 2010; Chatterjee, 2005). The ability to learn about brain function from the experience of people with brain injury has intrinsic worth beyond its inferential logic: Patients and families can provide rich descriptions of how brain damage has affected their lives, which can lead to unexpected insights beyond the laboratory context. Such anecdotal evidence can provide interesting starting points for hypothesis-driven experiments, the results of which may in turn be directly relevant to patient care. The observational nature of these studies requires thoughtful experimental design, and this chapter has aimed at providing an overview of the main factors to be considered in such designs.

References

Bates, E., Appelbaum, M., Salcedo, J., Saygin, A. P., & Pizzamiglio, L. (2003). Quantifying dissociations in neuropsychological research. *Journal of Clinical and Experimental Neuropsychology*, *25*, 1128–1153. doi:10.1076/jcen.25.8.1128.16724

Bates, E., Wilson, S. M., Saygin, A. P., Dick, F., Sereno, M. I., Knight, R. T., & Dronkers, N. F. (2003). Voxel-based lesion-symptom mapping. *Nature Neuroscience*, *6*, 448–450.

Catani, M., & ffytche, D. H. (2010). On "the study of the nervous system and behaviour." *Cortex*, *46*, 106–109. doi:10.1016/j.cortex.2009.03.012

Catani, M., Jones, D. K., & ffytche, D. H. (2005). Perisylvian language networks of the human brain. *Annals of Neurology*, *57*, 8–16. doi:10.1002/ana.20319

Chatterjee, A. (2005). A madness to the methods in cognitive neuroscience? *Journal of Cognitive*

Neuroscience, *17*, 847–849. doi:10.1162/089892 9054021085

Chen, R., Hillis, A. E., Pawlak, M., & Herskovits, E. H. (2008). Voxelwise Bayesian lesion-deficit analysis. *NeuroImage*, *40*, 1633–1642. doi:10.1016/ j.neuroimage.2008.01.014

Critchley, H. D., Mathias, C. J., Josephs, O., O'Doherty, J., Zanini, S., Dewar, B. K., . . . Dolan, R. J. (2003). Human cingulate cortex and autonomic control: Converging neuroimaging and clinical evidence. *Brain: A Journal of Neurology*, *126*, 2139–2152. doi:10.1093/brain/awg216

Crivellato, E., & Ribatti, D. (2007). Soul, mind, brain: Greek philosophy and the birth of neuroscience. *Brain Research Bulletin*, *71*, 327–336. doi:10.1016/ j.brainresbull.2006.09.020

Damasio, H., & Damasio, A. R. (1989). *Lesion analysis in neuropsychology*. New York, NY: Oxford University Press.

Dosenbach, N. U., Fair, D. A., Cohen, A. L., Schlaggar, B. L., & Petersen, S. E. (2008). A dual-networks architecture of top-down control. *Trends in Cognitive Sciences*, *12*, 99–105. doi:10.1016/j.tics.2008.01.001

Dunn, J. C., & Kirsner, K. (2003). What can we infer from double dissociations? *Cortex*, *39*, 1–7. doi:10.1016/S0010-9452(08)70070-4

Farah, M. J. (2004). *Visual agnosia* (2nd ed.). Cambridge, MA: MIT Press.

Feinberg, T. E., & Farah, M. (2006). A historical perspective on cognitive neuroscience. In T. E. Feinberg & M. Farah (Eds.), *Patient-based approaches to cognitive neuroscience* (pp. 3–20). Cambridge, MA: MIT Press.

Fellows, L. K., Heberlein, A. S., Morales, D. A., Shivde, G., Waller, S., & Wu, D. H. (2005). Method matters: An empirical study of impact in cognitive neuroscience. *Journal of Cognitive Neuroscience*, *17*, 850–858. doi:10.1162/0898929054021139

Fellows, L. K., Stark, M., Berg, A., & Chatterjee, A. (2008). Establishing patient registries for cognitive neuroscience research: Advantages, challenges, and practical advice based on the experience at two centers. *Journal of Cognitive Neuroscience*, *20*, 1107–1113. doi:10.1162/jocn.2008.20065

Frank, R. J., Damasio, H., & Grabowski, T. J. (1997). Brainvox: An interactive, multimodal visualization and analysis system for neuroanatomical imaging. *NeuroImage*, *5*, 13–30. doi:10.1006/nimg.1996.0250

Fruhmann Berger, M., Johannsen, L., & Karnath, H. O. (2008). Time course of eye and head deviation in spatial neglect. *Neuropsychology*, *22*, 697–702. doi:10.1037/a0013351

Gläscher, J., Tranel, D., Paul, L. K., Rudrauf, D., Rorden, C., Hornaday, A., . . . Adolphs, R. (2009). Lesion mapping of cognitive abilities linked to intelligence. *Neuron*, *61*, 681–691. doi:10.1016/j.neuron.2009. 01.026

Gross, C. G. (1995). Aristotle on the brain. *Neuroscientist*, *1*, 245–250. doi:10.1177/107385849500100408

He, Y., Dagher, A., Chen, Z., Charil, A., Zijdenbos, A., Worsley, K., & Evans, A. (2009). Impaired small-world efficiency in structural cortical networks in multiple sclerosis associated with white matter lesion load. *Brain: A Journal of Neurology*, *132*, 3366–3379. doi:10.1093/brain/awp089

He, Y., Wang, J., Wang, L., Chen, Z. J., Yan, C., Yang, H., . . . Evans, A. (2009). Uncovering intrinsic modular organization of spontaneous brain activity in humans. *PLoS ONE*, *4*(4), e5226. doi:10.1371/ journal.pone.0005226

Holmes, G. (1979). The organization of the visual cortex in man. In C. G. Phillips (Ed.), *Selected papers of Gordon Holmes* (pp. 438–451). Oxford, England: Oxford University Press. (Original work published 1944)

Karnath, H. O., Rorden, C., & Ticini, L. F. (2009). Damage to white matter fiber tracts in acute spatial neglect. *Cerebral Cortex (New York, NY)*, *19*, 2331–2337. doi:10.1093/cercor/bhn250

Kimberg, D. Y., Coslett, H. B., & Schwartz, M. F. (2007). Power in voxel-based lesion-symptom mapping. *Journal of Cognitive Neuroscience*, *19*, 1067–1080. doi:10.1162/jocn.2007.19.7.1067

Kinkingnéhun, S., Volle, E., Pelegrini-Issac, M., Golmard, J. L., Lehericy, S., du Boisgueheneuc, F., . . . Dubois, B. (2007). A novel approach to clinical-radiological correlations: Anatomo-Clinical Overlapping Maps (AnaCOM): Method and validation. *NeuroImage*, *37*, 1237–1249. doi:10.1016/j.neuroimage.2007.06.027

Laws, K. R. (2005). "Illusions of normality": A methodological critique of category-specific naming. *Cortex, 41*, 842–851.

Makale, M., Solomon, J., Patronas, N. J., Danek, A., Butman, J. A., & Grafman, J. (2002). Quantification of brain lesions using interactive automated software. *Behavior Research Methods, Instruments, and Computers*, *34*, 6–18. doi:10.3758/BF03195419

Marsh, E. B., & Hillis, A. E. (2008). Dissociation between egocentric and allocentric visuospatial and tactile neglect in acute stroke. *Cortex*, *44*, 1215–1220. doi:10.1016/j.cortex.2006.02.002

Milne, E., & Grafman, J. (2001). Ventromedial prefrontal cortex lesions in humans eliminate implicit gender stereotyping. *Journal of Neuroscience*, *21*(12), RC150.

Nachev, P., Coulthard, E., Jager, H. R., Kennard, C., & Husain, M. (2008). Enantiomorphic normalization of focally lesioned brains. *NeuroImage*, *39*, 1215–1226. doi:10.1016/j.neuroimage.2007.10.002

Newhart, M., Ken, L., Kleinman, J. T., Heidler-Gary, J., & Hillis, A. E. (2007). Neural networks essential

for naming and word comprehension. *Cognitive and Behavioral Neurology*, *20*, 25–30. doi:10.1097/WNN.0b013e31802dc4a7

Philippi, C. L., Mehta, S., Grabowski, T., Adolphs, R., & Rudrauf, D. (2009). Damage to association fiber tracts impairs recognition of the facial expression of emotion. *Journal of Neuroscience*, *29*, 15089–15099. doi:10.1523/JNEUROSCI.0796-09.2009

Picton, T. W., Stuss, D. T., Alexander, M. P., Shallice, T., Binns, M. A., & Gillingham, S. (2007). Effects of focal frontal lesions on response inhibition. *Cerebral Cortex (New York, NY)*, *17*, 826–838. doi:10.1093/cercor/bhk031

Robertson, L. C., Knight, R. T., Rafal, R., & Shimamura, A. P. (1993). Cognitive neuropsychology is more than single-case studies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 710–717. doi:10.1037/0278-7393.19.3.710

Rorden, C., & Brett, M. (2000). Stereotaxic display of brain lesions. *Behavioural Neurology*, *12*, 191–200.

Rorden, C., Fridriksson, J., & Karnath, H. O. (2009). An evaluation of traditional and novel tools for lesion behavior mapping. *NeuroImage*, *44*, 1355–1362. doi:10.1016/j.neuroimage.2008.09.031

Rorden, C., & Karnath, H. O. (2004). Using human brain lesions to infer function: A relic from a past era in the fMRI age? *Nature Reviews Neuroscience*, *5*, 812–819. doi:10.1038/nrn1521

Rorden, C., Karnath, H. O., & Bonilha, L. (2007). Improving lesion-symptom mapping. *Journal of Cognitive Neuroscience*, *19*, 1081–1088. doi:10.1162/jocn.2007.19.7.1081

Rudrauf, D., Mehta, S., Bruss, J., Tranel, D., Damasio, H., & Grabowski, T. J. (2008). Thresholding lesion overlap difference maps: Application to category-related naming and recognition deficits. *NeuroImage*, *41*, 970–984. doi:10.1016/j.neuroimage.2007.12.033

Rudrauf, D., Mehta, S., & Grabowski, T. J. (2008). Disconnection's renaissance takes shape: Formal incorporation in group-level lesion studies. *Cortex*, *44*, 1084–1096. doi:10.1016/j.cortex.2008.05.005

Shallice, T. (1988). *From neuropsychology to mental structure*. New York, NY: Cambridge University Press. doi:10.1017/CBO9780511526817

Solomon, J., Raymont, V., Braun, A., Butman, J. A., & Grafman, J. (2007). User-friendly software for the analysis of brain lesions (ABLe). *Computer Methods and Programs in Biomedicine*, *86*, 245–254. doi:10.1016/j.cmpb.2007.02.006

Stuss, D. T., Murphy, K. J., Binns, M. A., & Alexander, M. P. (2003). Staying on the job: The frontal lobes control individual performance variability. *Brain: A Journal of Neurology, 126*, 2363–2380. doi:10.1093/brain/awg237

Teuber, H. L. (1955). Physiological psychology. *Annual Review of Psychology*, *6*, 267–296. doi:10.1146/annurev.ps.06.020155.001411

Thiebaut de Schotten, M., Kinkingnéhun, S., Delmaire, C., Lehericy, S., Duffau, H., Thivard, L., . . . Bartolomeo, P. (2008). Visualization of disconnection syndromes in humans. *Cortex*, *44*, 1097–1103. doi:10.1016/j.cortex.2008.02.003

Tsuchida, A., & Fellows, L. K. (2009). Lesion evidence that two distinct regions within prefrontal cortex are critical for n-back performance in humans. *Journal of Cognitive Neuroscience*, *21*, 2263–2275. doi:10.1162/jocn.2008.21172

Urbanski, M., Thiebaut de Schotten, M., Rodrigo, S., Catani, M., Oppenheim, C., Touze, E., . . . Bartolomeo, P. (2008). Brain networks of spatial awareness: Evidence from diffusion tensor imaging tractography. *Journal of Neurology, Neurosurgery, and Psychiatry*, *79*, 598–601. doi:10.1136/jnnp.2007.126276

Van Essen, D. C., Felleman, D. J., DeYoe, E. A., Olavarria, J., & Knierim, J. (1990). Modular and hierarchical organization of extrastriate visual cortex in the macaque monkey. *Cold Spring Harbor Symposia on Quantitative Biology*, *55*, 679–696.