HHS Public Access

Author manuscript

Neuron. Author manuscript; available in PMC 2017 October 25.

Published in final edited form as:

Neuron. 2016 June 15; 90(6): 1151–1153. doi:10.1016/j.neuron.2016.05.014.

Human Lesion Studies in the 21st Century

Ralph Adolphs[1,*]

[1]Division of Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125, USA

Abstract

The study of patients with brain lesions has made major historical contributions to cognitive neuroscience. Here I argue for an increased investment in modern lesion mapping, complementing fMRI studies and laying the conceptual and analytic foundations for future techniques that could experimentally manipulate human brain function.

The Challenge and Promise of Lesion Studies

Remarkable alterations in cognition and behavior in lesion patients have been widely reported and popularized, and are often thought to have been the prerequisite for the emergence of cognitive neuroscience (Geschwind, 1997; Damasio and Damasio, 1989). But what exactly do such studies tell us, and what place do they find amidst the vastly more common approach of functional neuroimaging nowadays?

These questions typically generate two very different sets of answers. One set argues for the indispensable value of lesion studies and the dissociations in function they provide. Lesions give us insight into the causally necessary function of brain structures, whereas electrophysiology and fMRI reflect mere correlations with psychological processes. Lesions show us dissociations in cognition we could never have hypothesized, and thus can radically change our model of the architecture of the mind. And careful characterization of the deficits following lesions and their change over time provides clinically valuable information not only about the constellation of impairments produced, but also about their compensation and possible resolution over time. When all these pieces are put together with modern tools, data from lesion patients seem to show us how dynamic brain networks depend on the function of particular components, how cognition can be at least partially decomposed into modules of sorts, and how degeneracy and plasticity come into play. In short, lesion studies continue to offer unique value for cognitive psychology, cognitive neuroscience, clinical neuropsychology, and even cognitive science and philosophy of mind more generally.

A second set of answers is less sanguine. It is instead argued that the famous lesion studies can be counted on the fingers of one hand, do not generalize, and were actually wrong in most of the details. Furthermore, modern lesion studies are no better: every patient is different in idiosyncratic ways, the lesions are far too coarse since they cannot be

[*]Correspondence: radolphs@caltech.edu.

experimentally produced, and fMRI has already made the lesion method obsolete (see Rorden and Karnath, 2004 for a synopsis of these views and counterarguments).

The far more recent history of fMRI research has had its own share of ups and downs in the public image, but nobody concluded from debates about statistical or other technical difficulties that we should abandon fMRI. A huge effort has instead been put into developing more and more refined tools and into accumulating large shared datasets for fMRI analyses. Indeed, much of this development has in fact been driven by the temporary crises around replication and statistical power that the field has weathered. Many of the same problems apply to lesion studies. Happily, so do many of the same solutions. We need better statistical reliability and generalizability, more complex models that consider (for example) white matter and gray matter separately, and multivariate analyses. But these are all eminently surmountable problems. There is no inherent limitation here.

So one answer is that the criticisms outlined above are outdated. Yes, classical studies had single cases and were crude in many ways (so, too, were early neuroimaging studies), but these are problems that should be and can be solved, not inherent barriers to the method. For example, new multivariate algorithms for lesion-symptom mapping are now being applied (Smith et al., 2013; Mah et al., 2014). These algorithms can build in sophisticated information about the statistical anatomical dependencies between voxels in lesion analysis. In the case of a stroke, for example, the likelihood of different voxels being lesioned or not is highly correlated across an anatomically complex shape that results from how blood vessels are distributed in the brain. Multivariate algorithms can incorporate this information into their analysis, can treat gray and white matter separately, and can look for distributed patterns of lesions as the best correspondence to a cognitive process. These approaches address many of the methodological complaints that have been leveled against lesion mapping in the past. Another difficulty is that the lesion interacts with many other variables —so we need models that incorporate these other variables (age, sex, IQ, etc.). Such more complex models are also straightforward to construct in principle. However, these advances all require one key ingredient: a very large sample size. Even hundreds of lesion patients won’t do it; we need more than that. But there is actually no shortage of potential data at all. The challenge is how to collect and coordinate it.

It is estimated that there are about 800,000 inpatients a year in the U.S. with a primary diagnosis of stroke (Mozaffarian et al., 2015). There are a smaller number with brain injury from other causes like tumors or epilepsy surgery, which offer complementary strengths and confounds for comparison. Although the incidence of stroke has been declining somewhat over the years with better preventive measures, and although one would need to add stringent filters to produce quality data, it seems eminently feasible to recruit several thousand stroke patients a year if a comprehensive effort were made, and to aim for a 5-year sample size in the tens of thousands. That number is large enough that we could apply some very powerful machine learning tools to mining structure-function relationships. Of necessity, the sample would be fairly diverse, and indeed representative for the population in which we want to make clinical predictions. But we could not do it at a single hospital. As with fMRI studies, we would need to form consortia that pool data over many sites.

Most stroke patients nowadays will get structural scans of their brains that can be used as is, and with the ready availability of modern MRI scanners in every major hospital, it would be possible to do much better than that. Similarly, most of them will have some cognitive assessment, although here there is more need for improvement and administration of a standard battery of tasks. If healthcare providers and NIH both saw the value in this, a coordinated effort could easily produce research-quality MRI and behavioral data in tens of thousands of patients over a few years. Far from becoming obsolete, lesion mapping would suddenly offer a wealth of data rivaling, and complementing, that of other large efforts such as the human connectome project.

Value to Clinical and Basic Research

Clinical and cognitive neuropsychology part ways in focusing on an assessment of the symptoms and prognosis and in inferring something about the way the mind works, respectively. The history of cognitive neuropsychology has focused on inferring the cognitive architecture of healthy minds from the dissociations that case studies of patients with lesions offer (Caramazza, 1986). In this respect, cognitive neuropsychology is similar to cognitive neuropsychiatry: both make use of dissociations (deficits in circumscribed cognitive domains) to infer causal dependencies between processes. The collection of processes and their causal architecture constitutes our best model of the mind and cognition. It should be noted that this has not been a common goal of functional neuroimaging, which has generally stuck more to describing the neuroscience results than advancing theories of cognition (Coltheart, 2006).

But this is changing rapidly. A slew of recent fMRI studies are exploring the isomorphism between similarity relationships among psychological constructs (for instance, fear and anger may be conceptually more similar to one another than fear and happiness) and similarity relationships among patterns of brain activation (for instance, the pattern evoked when people think about fear may be more similar to the pattern evoked when people think about anger than when people think about happiness). Looking for correspondences between the neural representational space as derived from fMRI data and the psychological similarity space as derived from theories and ratings of stimuli is indeed giving us some initial tests of models of the mind using neuroimaging data (Kriegeskorte and Kievit, 2013). Large-sample lesion studies could offer complementary data: when administered the same battery of tasks, such data could let us see what dissociations between psychological processes are possible and hence serve to separate or combine processes into revised architectures of cognition.

The brain networks we are discovering using resting-state fMRI desperately need some roles in cognition. What do they do, and how do they contribute to the mind? The cognitive architectures informed by lesion studies desperately need detail and constraint. So we should vigorously attempt to merge these two sources of data. Indeed, they can perhaps be merged best with data from the very same subjects, since it is eminently feasible to obtain fMRI data in patients with focal lesions. Such data could show us how large-scale networks will change when one of their nodes is damaged, information critical for understanding the detailed causal relations between the nodes. Even if one obtained only resting-state fMRI, given the large extant databases on resting-state fMRI that are already available in healthy people

(thousands of subjects) and being accumulated also in psychiatric disorders (e.g., the ABIDE network for autism, currently around 500), resting-state fMRI data in lesion patients would be a tremendously valuable addition. Are there lesions whose resting-state networks look similar to those in autism? Are there some lesions in nodes of healthy resting-state networks that severely disrupt those networks, whereas lesions in other nodes can be compensated? Does a patient with a neurodegenerative disease show atypical resting-state networks that resemble those seen with focal lesions of a particular brain structure, implicating that brain structure in the degeneration? The clinical relevance is apparent, but such data would also be highly informative for network models of brain function more generally. Causal discovery modeling, while in its infancy in cognitive neuroscience, would seem ripe for application to such data once sufficiently large datasets have been accumulated (Ramsey et al., 2010; Hyttinen et al., 2013). Essentially every network-level fMRI study, which is most fMRI studies these days, would benefit if the fMRI study were also carried out in patients who have lesions in some of the nodes of the network. One would not even need large samples for this approach to be useful (indeed, we have done it with Ns of 2 and 3; Hampton et al., 2007; Spunt et al., 2015).

These basic-research goals, while of course also informative for clinical questions, would be complemented by direct clinical studies. Multivariate classification of the neuroanatomy of brain lesions could be used to predict clinical outcome and indeed treatment efficacy. Of the 800,000 stroke cases a year, only about 140,000 result in death (and that number is declining notably over the years with better acute treatment). That leaves most for whom to plan for the best care into the future. Redundancy and degeneracy in brain networks, concepts that have been around for a while (Price and Friston, 2002) but that have been tested in only a few studies so far, could be systematically mapped in large and longitudinal samples. Some compensations can be quite dramatic. For example, a rare congenital disorder, the complete absence of the entire corpus callosum, remarkably leaves the functional networks seen with resting-state fMRI entirely unaffected, including their bilateral symmetry (Tyszka et al., 2011). This seems to be telling us something important about the functional principles by which such networks organize.

It is important to realize that the core data on which lesion studies are based are quite objective—we just need good structural scans to map the lesion. Aligning lesions very precisely across brains is less easy, but again we can borrow tools from recent developments in fMRI, which use multimodal data, including functional data, to improve the alignment. There are a host of additional factors that contribute noise and reduce power if not measured, but we can incorporate many of those and reduce the effect of the rest through large sample sizes. Again, the challenge is not with the lesions themselves—those are as objective and close to the brain as you can get. And, more and more, the challenge is also not with the analytic tools, since neuroimaging has already done much of the legwork here. The challenge is mostly in coordinating a sufficiently large and comprehensive consortium to collect the data.

The Future

About a quarter of all stroke patients suffer another stroke, and thus a sufficiently large longitudinal sample would allow us to progressively map the effects of multiple lesions. Add to this the richness of including age among other variables, and it seems clear that an initial 5-year comprehensive lesion mapping effort would nucleate a much larger set of research projects long into the future.

Another important future contribution for lesion studies would be that they lay the groundwork and provide a comparison for experimental manipulations of brain function. There are a lot of attractive ideas spawned by the BRAIN initiative, but their large-scale application in healthy human brains remains unclear. In the meantime, we have naturally occurring lesions, and the tools that we will develop for analyzing such data will be invaluable also for any future techniques that could experimentally manipulate human brain activity. Moreover, many of the same constraints will apply: incorporating background variability, designing efficient tasks, and factoring in compensatory and “off-target” effects that occur even with the most precise circuit manipulations (Otchy et al., 2015).

Even the most optimistic scenario for the future of lesion studies still leaves us with large conceptual challenges that all of cognitive neuroscience faces. What mental functions should we try to map onto the brain in the first place? Clearly, not the ones that phrenology had proposed, but it remains unclear that our ontology of cognitive processes is quite ready for primetime either. There are also deep questions about individual differences. Cognitive neuropsychology generally makes a very strong, but seemingly reasonable, assumption about cognition: that there is a single, typical, human cognitive architecture. But perhaps the mapping from brain function to mind shows more variability, and perhaps the architecture does differ substantially across individuals. It is difficult to know whether this would set limits to what cognitive neuroscience could contribute to our understanding of the mind, but at any rate they would be limits faced by neuroimaging and lesions alike.

Acknowledgments

I thank Antonio Damasio, Hanna Damasio, Daniel Tranel, and Jan Glaescher for helpful comments. Supported in part by funding from NIMH Conte Center P50MH094258.

References

Caramazza A. Brain Cogn. 1986; 5:41–66. [PubMed: 3954906]

Coltheart M. Cortex. 2006; 42:323–331. [PubMed: 16771037]

Damasio, H., Damasio, AR. Lesion Analysis in Neuropsychology. Oxford University Press; 1989.

Geschwind, N. Norman Geschwind: Selected Publications on Language, Epilepsy, and Behavior. Devinsky, O., Schacter, SC., editors. Butterworth-Heinemann; 1997.

Hampton AN, Adolphs R, Tyszka MJ, O’Doherty JP. Neuron. 2007; 55:545–555. [PubMed: 17698008]

Hyttinen A, Eberhardt F, Hoyer PO. Journal of Machine Learning Research. 2013; 14:3041–3071.

Kriegeskorte N, Kievit RA. Trends Cogn Sci. 2013; 17:401–412. [PubMed: 23876494]

Mah YH, Husain M, Rees G, Nachev P. Brain. 2014; 137:2522–2531. [PubMed: 24974384]

Mozaffarian D, Benjamin EJ, Go AS, Arnett DK, Blaha MJ, Cushman M, de Ferranti S, Després JP, Fullerton HJ, Howard VJ, et al. Stroke Statistics Subcommittee. Circulation. 2015; 131:e29–e322. [PubMed: 25520374]

Otchy TM, Wolff SBE, Rhee JY, Pehlevan C, Kawai R, Kempf A, Gobes SMH, Ölveczky BP. Nature. 2015; 528:358–363. [PubMed: 26649821]

Price CJ, Friston KJ. Trends Cogn Sci. 2002; 6:416–421. [PubMed: 12413574]

Ramsey JD, Hanson SJ, Hanson C, Halchenko YO, Poldrack RA, Glymour C. Neuroimage. 2010; 49:1545–1558. [PubMed: 19747552]

Rorden C, Karnath HO. Nat Rev Neurosci. 2004; 5:813–819. [PubMed: 15378041]

Smith DV, Clithero JA, Rorden C, Karnath HO. Proc Natl Acad Sci USA. 2013; 110:1518–1523. [PubMed: 23300283]

Spunt RP, Elison JT, Dufour N, Hurlemann R, Saxe R, Adolphs R. Proc Natl Acad Sci USA. 2015; 112:4827–4832. [PubMed: 25825732]

Tyszka JM, Kennedy DP, Adolphs R, Paul LK. J Neurosci. 2011; 31:15154–15162. [PubMed: 22016549]