

Aligning Multiple Research Techniques in Cognitive Neuroscience: Why Is It Important?

William Bechtel[†]

University of California, San Diego

The need to align multiple experimental procedures and produce converging results so as to demonstrate that the phenomenon under investigation is real and not an artifact is a commonplace both in scientific practice and discussions of scientific methodology (Campbell and Stanley 1963; Wimsatt 1981). Although sometimes this is the purpose of aligning techniques, often there is a different purpose—multiple techniques are sought to supply different perspectives on the phenomena under investigation that need to be integrated to answer the questions scientists are asking. After introducing this function, I will illustrate it by considering three of the major techniques in cognitive neuroscience for linking cognitive function with neural structure.

1. Independent Epistemic Support or Complementation? The conception that demonstrating convergence between results procured with different techniques serves to provide epistemic support for each of them stems in part from recognition that research techniques are fallible. Typically, techniques work by introducing major alterations in the phenomena under study and rest on assumptions that are untested. It is often not difficult for critics to mount a challenge that a result generated by just one technique is an artifact. Converging results are seen as supporting each of the converging techniques since it is unlikely that by chance two techniques that involve different interventions and make different assumptions would produce convergent results. But there are reasons to think this is not the main goal of aligning results from different techniques.

In a previous examination of how scientists vindicate new techniques (Bechtel 2000), I found that their judgments more frequently turned on

[†]William Bechtel, Department of Philosophy, 0119, University of California, San Diego, CA 92093; e-mail: bill@mind-brain.net.

whether new techniques produced data that supported plausible theories (which may themselves draw support from different types of data) than on convergence between the techniques themselves. More troubling for the view that epistemic support is the prime purpose of seeking convergence is that making comparisons between techniques, especially when they involve radically different interventions in the phenomenon, itself often requires a host of untested and contested assumptions. Even more serious is the fact that the alternative techniques are frequently employed to calibrate the interpretation of a given technique. This is due to the fact that the data resulting from a given technique are often not directly linked to the phenomena and researchers need to determine how to relate the data back to the phenomenon. Here the ability to make the data from two different techniques converge provides a way to use the first technique to link the data from the second to the phenomenon. An example in the cognitive neurosciences is the use of comparisons between results of cell recording on primates and neuroimaging on humans to identify areas comparable to primate brain regions (e.g., MT) in humans. In this research, the area active in a neuroimaging study for a given task is identified as the homologue to the area that elicits activity in individual cells in other primates. When techniques are used in such a way to calibrate each other, they are no longer independent in the manner required for them to provide independent evidence.

A different purpose for aligning techniques, and the one I will highlight in this paper, is to obtain complementary information about the phenomenon under investigation. This is extremely important in circumstances in which any given technique can only provide a very selective and distorted perspective on the phenomenon. In such cases it is only by relating results to those acquired with other techniques that researchers are able to address the questions being posed. In the cognitive neurosciences, where a major goal is to relate neural structures to cognitive operations, no one technique can reveal what cognitive operation is performed by a given brain area, but integrating the results from multiple techniques can provide a much better understanding. To illustrate this, I focus on three techniques commonly invoked in cognitive neuroscience—lesion, cell recording, and neuroimaging.

Underlying all of the research techniques I examine is the assumption that different brain areas perform different information processing operations (often referred to as computations) and that an explanation of a cognitive performance involves both decomposing an overall task into component information processing activities and determining what brain area performs each activity (Bechtel and Richardson 1993). The fact that it is component operations, not whole mental tasks, that are localized distinguishes this endeavor from phrenology (Zawidzki and Bechtel 2002).

The contrast between localizing a *task* in a brain area and localizing an information processing *operation* is critical and is well explicated by Peterson and Fiez:

[E]lementary operations, defined on the basis of information processing analyses of task performance, are localized in different regions of the brain. Because many such elementary operations are involved in any cognitive task, a set of distributed functional areas must be orchestrated in the performance of even simple cognitive tasks. . . . A functional area of the brain is not a task area: there is no “tennis forehand area” to be discovered. Likewise, no area of the brain is devoted to a very complex function; “attention” or “language” is not localized in a particular Brodmann area or lobe. Any task or “function” utilizes a complex and distributed set of brain areas. (Peterson and Fiez 1993, 513)

The idea that different brain areas perform different information processing operations is partly grounded in the fact that different areas have different neuroarchitectures: different types of neurons, different thicknesses of cortical layers, different patterns of projection, etc. (Mundale 1998). Brodmann ([1909] 1994), whose cartography of the brain is still a reference point in attempts to localize brain function, employed only the first two criteria but more recent researchers have taken advantage of newer techniques to trace connectivity (Carmichael and Price 1994; Felleman and van Essen 1991). (Mundale notes that the different criteria for delineating brain areas are not always equally applicable and do not always yield comparable results, requiring researchers to arbitrate between them in deciding what really are brain areas. Nonetheless, the differing criteria are all considered, in those cases where they can be applied, because each provides a different perspective on brain organization. Thus, the same point I am making in focusing on the mapping between structure and function applies to the techniques for identifying structural units.)

Sometimes the specialized neuroanatomy of a brain area may be suggestive as to its information processing operation. For example, in the hippocampus the sparse firing of granule cells in the dentate gyrus and sparse connectivity from the dentate gyrus to the CA3 fields in the hippocampus are strongly suggestive of a component involved in separating patterns, whereas the high level of interconnections between CA3 neurons suggests that it serves to create attractors to further separate the input patterns (O’Reilly and McClelland 1994; Rolls and Treves 1998). But this case is exceptional, and even in this case, as noted below, there is a fierce controversy over what information is being processed.

The challenge researchers confront is that within a reasonably well adapted complex system of many interacting parts, one cannot simply

look to find out what operation each part is performing. The different activities are so coordinated with each other that their individual identity is not readily detected. Rather, one must proceed indirectly, and until recently the two most common approaches were either to inhibit or to excite one of the components and then infer its contributions to the normal process from the affects of the intervention on the overall behavior of the system (Bechtel and Richardson 1993). More recently, tools have been developed that permit one to correlate a measure of the activity of a component with the types of tasks in which it might figure. But each of these approaches is seriously limited in its ability to fix the information processing activity performed by an area, giving rise to the need to invoke multiple approaches. In what follows I will begin with the first technique that provided suggestions of the information processing function of brain areas—lesion studies—and show both what they can and cannot demonstrate about the functions of brain areas. I will then show how two other techniques—cell recording and neuroimaging—complement lesion studies and permit a more detailed account of the information processing operations performed by brain areas.

2. Lesion Studies: What They Can and Cannot Establish. The exemplar of the use of lesion studies to assign a cognitive function to a brain area was Broca's (1861) assignment of the deficit in articulate speech in his patient Leborgne (Tan) to damage in the third frontal lobe in the left hemisphere. Although Broca's accomplishment is justly celebrated, one should note that what he did was assign a task to a brain area, not an elementary operation. Moreover, the claim that the location in the third frontal lobe alone performed the processes required in articulate speech involved an inference that went far beyond the actual data, which showed only that articulate speech depended on Broca's area, not what operations required in articulate speech were performed there.

Subsequently lesions, either naturally occurring or due to surgery to relieve symptoms of illness in humans or ones experimentally produced in other animals, and examination of the resulting deficits, have been one of the most widely used tools for relating structure and function. Among a very large set of examples in late twentieth century research that drew primarily on lesion research is Ungerleider and Mishkin's (1982) differentiation of what and where pathways in the visual system, Scoville and Milner's (1957) attribution of a role in encoding long term declarative memories to the hippocampus and O'Keefe and Nadel's (1978) attribution of a role in tracking spatial location in allocentric maps to the same structure, and Goldman-Rakic's (1987) attribution of working memory processes to frontal areas.

Most researchers who utilize lesion induced deficits are quite cognizant

of their limitations. Some limitations are practical. Naturally occurring lesions in humans do not respect brain cartography, removing a given functional area and leaving others untouched. Moreover, given the adaptation of other brain areas, a lesioned brain is not just a whole brain minus the lesioned site, but a brain in which processes performed in other areas may be altered. Finally, and most critically for my purposes, the loss of a cognitive ability following a lesion does not directly reveal what information processing operation that area performed. If the neural system is a large, complexly integrated system, the loss of a single component, which only performs a part of the activity, may make the system unable to perform the activity at all. For example, destruction of the lateral geniculate nucleus will eliminate the ability to see but does not reveal what information processing function cells in the LGN are performing in the process of seeing. (A clearly mistaken inference is to assume that if a cognitive ability is lost with a lesion, then the area lesioned itself performed the whole function. Researchers who make such inferences nonetheless often contribute importantly by motivating further research that delimits more fully what operation the area does perform (see Bechtel and Richardson 1993).

Sometimes detailed analysis of the abilities that are spared as well as those lost after a lesion can suggest the information processing activity the lesioned component performed. For example, HM, following Scoville's resection of his hippocampus, exhibited both anterograde and graded retrograde amnesia (Scoville and Milner 1957). The fact that HM could still remember events from many years before his surgery shows that the hippocampus was not the site of engrams for long-term memories. HM's inability to acquire new memories suggested that it played a role in encoding long-term memories. But the fact that HM also lost memory of events in the years immediately prior to surgery suggests that the hippocampus required a prolonged period to encode memories into cortex. This idea that has been spelled out more fully in O'Reilly and McClelland's (1994) hypothesis that the hippocampus temporarily stores information and then trains other areas of cortex by repeatedly reinstantiating the information to be learned.

Even this example, however, reveals the difficulties of attributing a particular information processing operation on the basis of lesion results alone. The putative role of the hippocampus in reinstantiating patterns to be learned is supported by neuroanatomical analysis of the peculiar hippocampal architecture and mathematical and computational analyses of how it could perform this task. Even so, though, these claims about encoding long-term memories are disputed by other researchers, many of whom argue that the primary activity of the hippocampus is in encoding allocentric spatial information (Nadel 1994; see Bechtel and Abrahamsen 2002).

Thus, while lesion studies can establish that a brain area plays a necessary role in a cognitive task, and even provide clues as to the information processing function it performs, they cannot establish in detail what information processing it performs. As well, they fail to show what other areas are involved in performing the activity, information that can be very useful for figuring out what different information processing operations go into the overall activity. As we shall see, cell recording studies can provide additional clues as to the information processing an area performs and neuroimaging is proving very useful in revealing networks of areas engaged in various tasks.

3. Complementing Lesion Studies with Cell Recording Studies. Although neuroimaging has attracted much attention in recent years, and in many respects was the development that catalyzed the establishment of *cognitive neuroscience* as its own field of research, the capacity to record from individual cells and to correlate their firing rate with features of stimuli (or tasks) to which they were responsive has provided the most direct clues to the information processing operations performed by specific neural tissue.

Kuffler (1953) pioneered this line of research when he demonstrated that some cells in LGN respond most strongly (produce the greatest number of action potentials) to a bright circular stimulus situated in a dark surround (an *on-center* cell) while others respond most strongly to a dark stimulus in a bright surround (an *off-center* cell). Hubel and Wiesel were working in Kuffler's laboratory when they began extending this approach to cortical cells in area 17. They met with little success when they initially tried to elicit responses with circular stimuli such as Kuffler had used. By accident, though, they discovered that bars of light would elicit rapid firing in these cells. Hubel reports that, as they were inserting a glass slide into their projecting ophthalmoscope, "over the audiometer the cell went off like a machine gun" (Hubel 1982, 438). Further exploration revealed that "as the glass slide was inserted its edge was casting onto the retina a faint but sharp shadow, a straight dark line on a light background" (439). In subsequent research they discovered cells in this area that responded to lines only at particular locations and orientations within their receptive fields (they named these *simple cells*) and others that responded to lines anywhere in their receptive fields or to lines moving in a particular direction through their receptive fields (they named these *complex cells*) (Hubel and Wiesel 1962, 1968).

The research strategy initiated by Kuffler, Hubel, and Wiesel was continued by numerous researchers in subsequent decades. Focusing on extrastriate areas (areas adjacent to and in front of area 17), Zeki (1973, 1974) determined that cells in V4 responded to light of specific wave lengths whereas those in V5 (MT) responded to motion. Gross, Rocha-Miranda, and Bender (1972) found cells in inferotemporal cortex that re-

spond to stimuli of different shapes while Goldberg and Robinson (1980) found cells in posterior parietal cortex that are specially sensitive to location. Although the strategy is most naturally applied to cells engaged directly in sensory or motor processing, it can also be applied to memory tasks by showing which cells are particularly active during intervals when animals must retain information in order to complete a task.

The fact that a given cell is particularly active in response to stimuli with specific features suggests that the firing of the cell is carrying information about those features. However, it requires a further inference to conclude that the cell is *processing* that information rather than just conducting it. However, if areas that project to a given area are not responsive to a particular type of information (e.g., if cells in area 17 are not sensitive to illusory contours), that provides a potent clue to what is processed in an area, for it suggests that a computation has been performed to extract the particular information about the stimulus feature from the information earlier cells were carrying. This led Hubel and Wiesel to propose a computational model of how simple cells in area 17 computed information about bars of light (edges) from information about the circular stimuli detected by LGN cells.

Recording from individual cells provides powerful clues but cannot alone settle what information processing occurs in a given brain area. One problem is that the technique only identifies stimuli researchers think to test. As I noted, Hubel and Wiesel found stimuli to excite area 17 cells only serendipitously. Luck also figures in Gross's identification of areas in inferotemporal cortex—a major breakthrough occurred when, in frustration, researchers waved their hands in front of the monkey from which they were recording. Secondly, cell recording rests upon the assumption that a cell carries information about the stimulus that causes it to fire maximally. The cell can thus be construed as detecting that feature. Recently this assumption has been challenged by investigators who note that most cells are responsive to a range of stimuli at varying rates, and who suggest that the cell is performing information processing even when it is firing at lower rates. Van Essen and Gallant (1994) therefore propose that cells should be viewed as filters, not detectors. Finally, although it is reasonable to assume that an area where cells first respond to a feature of the stimulus is responsible for processing that information, it could be processed yet further downstream and the activity in the selected area may only reflect feedback from processing in these later areas.

Many of these difficulties can be partly overcome by relating single-cell recording results to lesion results. Deficits revealed when an area is lesioned can provide clues to the information that is being processed in the area, which can then guide the choice of stimuli in cell-recording studies. Lesion studies can also shed light on the feature detector/filter controversy.

If cells are filters then lesioning them should have more widespread effects than if they were simply feature detectors. (Detecting such diffuse deficiencies through lesion studies, though, will require attending to subtle patterns of deficits.)

Lesion and cell-recording clearly complement each other. But they also share a common limitation in that researchers can only focus on one or a few brain areas at a time. (Techniques for simultaneously recording from many cells at once are in use, but they are complex and still only give information about areas where one chooses to implant an electrode.) If, though, brain areas are performing elementary information processing operations, multiple areas must be engaged together in the performance of a task. One of the major accomplishments of neuroimaging techniques has been to help identify networks of brain areas engaged in specific tasks.

4. Complementing Lesion and Cell Recording Studies with Neuroimaging.

Early neuroimaging studies tended to emphasize a single area involved in a cognitive operation. This was probably largely a consequence of reliance on the subtractive method that was adapted from reaction time studies to interpret imaging results. Using this method, researchers imaged subjects while performing a pair of tasks that were thought to differ in just a single information processing operation. They then subtracted the blood flow measured while performing the simpler task from that recorded during the more complex task (Posner and Raichle 1994). However, the most probative early studies drew upon the ability of PET to identify multiple areas of increased blood flow (something the increased resolution of fMRI further facilitates). In their classic study using the verb-generate task to identify areas involved in semantic processing, Petersen et al. (1988) found increased activation not only in the left prefrontal cortex, but also in the cerebellum. Similarly, in their study of episodic memory processes, Squire et al. (1992) found increased activation not only near the hippocampus but also in right prefrontal cortex. Perhaps due to the fact that the subtraction technique leads researchers to look for just one area corresponding to the one additional information processing operation thought to be required, there was a tendency to discount these other activations (although both of the additional areas identified have subsequently turned out to be very important).

There is another way the subtraction technique may conceal the involvement of multiple brain areas. In studying episodic memory encoding and retrieval with PET, Nyberg, Cabeza, and Tulving (1996) subtracted areas active in semantic memory retrieval (which are the same as those involved in episodic encoding) from those active in episodic retrieval to argue for localizing encoding in left prefrontal cortex and retrieval in right prefrontal cortex (the Hemispheric Encoding/Retrieval Asymmetry hy-

pothesis or HERA). Subsequently, Buckner (1996) showed that subtraction concealed the fact that in episodic retrieval there is activation both in the left prefrontal areas associated with semantic retrieval (episodic encoding) and in the right prefrontal areas in which Tulving localized episodic retrieval. This suggests that the processing involved in episodic retrieval relies on a network including those involved in encoding.

Roediger, Buckner, and McDermott (1999) provide yet another illustrative example of how neuroimaging may reveal networks of brain areas engaged in any cognitive performance. They compare three tasks involving word stem completion (e.g., completing COU with *couple*). In the first, subjects are simply asked to complete the stems with the first words that come to mind. Imaging results reveal a network of active areas including visual cortex, left frontal-opercular cortex, anterior cingulate, supplemental motor areas and premotor and motor cortex. In the second task, subjects were first presented a list of words and then asked to complete stems with the first words that came to mind (they could use the previously presented words, but were not instructed to do so). The imaging results were very similar except for less activation in visual cortex, which Roediger et al. attribute to the fact that the priming effect of previous exposure reduced the need for visual processing. The third task was like the second except that subjects were required to complete the stems with words from the previously heard list. This resulted in activation in all the previously active areas plus two additional ones—medial parietal cortex and right anterior prefrontal cortex. Such a focus on networks of processing areas activated in a task has recently become common in neuroimaging. In Cabeza and Nyberg's (2000) review of 275 recent PET and fMRI studies they found that imaging studies involving attention, language, and various forms of memory all found increased activation in a broad range of brain areas.

Identifying a network of brain areas, though, raises the challenge of figuring out which information processing activity each area performs. The comparison of different tasks can provide suggestive clues—by analyzing the details of the tasks that elicited different activation patterns, one may be able to determine what types of information processing were required in one condition but not in another. The additional active areas are good candidates for performing these additional tasks. But this approach relies critically on a task analysis that is generally arrived at a priori by trying to imagine a set of operations that would perform the task. Sometimes more probative evidence about what information processing operations these areas perform can be procured by aligning these results with results of lesion and, when possible, cell-recording studies.

5. Conclusion. None of the techniques available to researchers for linking

cognitive processes with brain areas are themselves capable of identifying the information processing performed in the brain area. The three techniques I have considered—lesion, cell-recording, and neuroimaging—can each provide suggestive clues. But each needs to be complemented by the results of the other techniques to resolve difficulties that it cannot address on its own. Even together they do not provide a definitive account of the processing in a particular area. But as additional information is provided, plausible interpretations are constrained. It is this potential, rather than the epistemic objective of evaluating the reliability of techniques and avoiding artifacts, that is often what drives scientists to try to align the results of different research techniques.

REFERENCES

- Bechtel, William (2000), "From Imaging to Believing: Epistemic Issues in Generating Biological Data", in R. Creath and J. Maienschein (eds.), *Biology and Epistemology*. Cambridge: Cambridge University Press, 138–163.
- Bechtel, William and Adele Abrahamsen (2002), *Connectionism and the Mind: Parallel Processing, Dynamics, and Evolution in Networks*. Second Edition. Oxford: Blackwell.
- Bechtel, William and Robert C. Richardson (1993), *Discovering Complexity: Decomposition and Localization as Scientific Research Strategies*. Princeton: Princeton University Press.
- Broca, Paul (1861), "Remarque sur le siège de la faculté du langage articulé; suivies d' une observation d' aphémie", *Bulletins de la Société d'Anthropologie* 6: 343–357.
- Brodman, Korbinian ([1909]1994), *Vergleichende Lokalisationslehre der Grosshirnrinde*. Translated by L. J. Garvey. Leipzig: J. A. Barth.
- Buckner, Randy L. (1996), "Beyond HERA: Contributions of Specific Prefrontal Brain Areas to Long-term Memory Retrieval", *Psychonomic Bulletin and Review* 3: 149–158.
- Cabeza, Roberto and Lars Nyberg (2000), "Imaging Cognition II: An Empirical Review of 275 PET and fMRI studies", *Journal of Cognitive Neuroscience* 12: 1–47.
- Campbell, Donald T. and Julian C. Stanley (1963), *Experimental and Quasi-experimental Design for Research*. Boston: Houghton Mifflin.
- Carmichael, S. T. and Joel L. Price (1994), "Architectonic Subdivision of the Orbital and Medial Prefrontal Cortex in the Macaque Monkey", *Journal of Comparative Neurology* 346: 366–402.
- Felleman, Daniel J. and David C. van Essen (1991), "Distributed Hierarchical Processing in the Primate Cerebral Cortex", *Cerebral Cortex* 1: 1–47.
- Goldberg, M. E. and D. L. Robinson (1980), "The Significance of Enhanced Visual Responses in Posterior Parietal Cortex", *Behavioral and Brain Sciences* 3: 503–505.
- Goldman-Rakic, Patricia S. (1987), "Circuitry of Primate Prefrontal Cortex and Regulation of Behavior by Representational Memory", in J. M. Brookhart, V. B. Mountcastle, and S. R. Geiger (eds.), *Handbook of Physiology: The Nervous System*, vol. 5. Bethesda, Maryland: American Physiological Society, 373–417.
- Gross, Charles G., C. E. Rocha-Miranda, and D. B. Bender (1972), "Visual Properties of Neurons in Inferotemporal Cortex of the Macaque", *Journal of Neurophysiology* 35: 96–111.
- Hubel, David H. (1982), "Evolution of Ideas on the Primary Visual Cortex, 1955–1978: A Biased Historical Account", *Bioscience Reports* 2: 435–469.
- Hubel, David H. and Torsten N. Wiesel (1962), "Receptive Fields, Binocular Interaction and Functional Architecture in the Cat's Visual Cortex", *Journal of Physiology (London)* 160: 106–154.
- (1968), "Receptive Fields and Functional Architecture of Monkey Striate Cortex", *Journal of Physiology (London)* 195: 215–243.

- Kuffler, Stephen W. (1953), "Discharge Patterns and Functional Organization of Mammalian Retina", *Journal of Neurophysiology* 16: 37–68.
- Mundale, Jennifer (1998), "Brain Mapping", in William Bechtel and George Graham (eds.), *A Companion to Cognitive Science*. Oxford: Basil Blackwell.
- Nadel, Lynn (1994), "Multiple Memory Systems: What and Why, an Update", in Daniel L. Schacter and Endel Tulving (eds.), *Memory Systems 1994*. Cambridge, MA: MIT Press, 39–63.
- Nyberg, Lars, Roberto Cabeza, and Endel Tulving (1996), "PET Studies of Encoding and Retrieval: The HERA Model", *Psychonomic Bulletin and Review* 3: 135–148.
- O'Keefe, John and Lynn Nadel (1978), *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press.
- O'Reilly, Randall C. and James L. McClelland (1994), "Hippocampal Conjunctive Encoding, Storage and Recall: Avoiding a Tradeoff", *Hippocampus* 4: 661–682.
- Petersen, Stephen E. and Julie A. Fiez (1993), "The Processing of Single Words Studied with Positron Emission Tomography", *Annual Review of Neuroscience* 16: 509–530.
- Peter T. Fox, Michael I. Posner, Mark Mintun, and Marcus E. Raichle (1988), "Positron Emission Tomographic Studies of the Cortical Anatomy of Single-word Processing", *Nature* 331: 585–588.
- Posner, Michael I., and Marcus E. Raichle (1994), *Images of Mind*. San Francisco: Freeman.
- Roediger III, Henry L., Randy L. Buckner, and Kathleen B. McDermott (1999), "Components of Processing", in J. K. Foster and M. Jelicic (eds.), *Memory: Systems, Process, or Function*. Oxford: Oxford University Press, 32–65.
- Rolls, Edmund T. and Alessandro Treves (1998), *Neural Networks and Brain Function*. Oxford: Oxford University Press.
- Scoville, William B. and Brenda Milner (1957), "Loss of Recent Memory after Bilateral Hippocampal Lesions", *Journal of Neurology, Neurosurgery, and Psychiatry* 20: 11–21.
- Squire, Larry, Jeffrey G. Ojemann, Francis M. Miezin, Stephen E. Petersen, Tom O. Videen, and Marcus E. Raichle (1992), "Activation of the Hippocampus in Normal Humans: A Functional Anatomical Study of Memory", *Proceedings of the National Academy of Sciences* 89: 1837–1841.
- Tulving, Endel, S. Kapur, Fergus I. M. Craik, Morris Moscovitch, and S. Houle (1994), "Hemispheric Encoding/Retrieval Asymmetry in Episodic Memory: Positron Emission Tomography Findings", *Proceedings of the National Academy of Sciences (USA)* 91: 2016–2020.
- Ungerleider, Leslie G. and Morten Mishkin (1982), "Two Cortical Visual Systems", in D. J. Ingle, M. A. Goodale, and R. J. W. Mansfield (eds.), *Analysis of Visual Behavior*. Cambridge, MA: MIT Press, 549–586.
- van Essen, David C., and Jack L. Gallant (1994), "Neural Mechanisms of Form and Motion Processing in the Primate Visual System", *Neuron* 13: 1–10.
- Wimsatt, William C. (1981), "Robustness, Reliability, and Overdetermination in Science", in M. Brewer and B. Collins (eds.), *Scientific Inquiry and the Social Sciences*. San Francisco: Jossey-Bass, 124–163.
- Zawidzki, Thadeus and William Bechtel (2002). "Gall's Legacy Revisited: Decomposition and Localization in Cognitive Neuroscience", in Christina E. Erneling and David M. Johnson (eds.), *Mind as a Scientific Object: Between Brain and Culture*. Oxford: Oxford University Press, (in press).
- Zeki, Semir M. (1973), "Colour Coding of the Rhesus Monkey Prestriate Cortex", *Brain Research* 53: 422–427.
- (1974), "Functional Organization of a Visual Area in the Posterior Bank of the Superior Temporal Sulcus of the Rhesus Monkey", *Journal of Physiology* 236: 549–573.