Cognitive Science Online, Vol.1, pp.22-33, 2003 http://cogsci-online.ucsd.edu

A discussion and review of Uttal (2001) *The New Phrenology* *

Edward M. Hubbard

Department of Psychology and
Interdisciplinary Program in Cognitive Science
University of California, San Diego
edhubbard@psy.ucsd.edu

Abstract

In his recent book, William R. Uttal suggests that the modern attempts to understand the localization of cognitive functions are misguided. According to Uttal, attempts to localize cognitive functions will fail because the to-be-localized cognitive processes simply do not exist. In this review, I will defend the position that the mind can be meaningfully divided into cognitive parts and, given that the brain is non-homogenous, we may be able to localize these cognitive parts to brain parts. However, this view does not posit the simple localization that Uttal argues against. Instead, I argue that mental functions, broadly construed, are distributed, but are composed of multiple sub-components that can be localized to specific brain regions.

Introduction

I am neither a universalizer nor a localizer…In consequence I have been attacked as a universalizer and also as a localizer. But I do not remember that the view I really hold as to localization has ever been referred to. If it is, it will very likely be supposed to be a fusion of, or a compromise betwixt recent doctrines.

John Hughlings Jackson, (1882/1932)

One of the fundamental problems in philosophy, psychology, neuroscience and cognitive science is the extent to which mental functions can be decomposed, and once decomposed, how closely they can be tied to specific brain locations. Descartes famously argued that the mind was a unified entity and that it could not be divided. However, with increasing knowledge of the mind, Descartes' position has largely been abandoned in favor of the view that the mind can be decomposed into a number of functional parts and that these mental parts have their basis in different parts of the brain. The first (and most notorious) of these ideas was phrenology, first proposed by Franz Joseph Gall in 1796. This view has rightly been criticized on many grounds, but the essential insight, that the mind is composed of distinct functional parts, and that

* Uttal, W.R. (2001). *The New Phrenology: The Limits of Localizing Cognitive Processes in the Brain.* Cambridge, MA : MIT Press, 255 p., ISBN 0262210177.

these parts can be localized to specific parts of the brain, has remained an essential component of modern thinking about the relation between the mind and the brain. Since the time of Gall, there has been tremendous progress in our knowledge of both the mind and the brain, but whether, and how precisely, cognitive functions can be localized in the brain is still open to debate.

In *The New Phrenology*, William R. Uttal (2001) explores the fundamental assumptions underlying the attempt to localize cognitive functions in specific brain areas (what Uttal refers to as the "localizationist approach"). In his book, Uttal concludes, contrary to prevailing opinion, that the localizationist enterprise is beset by such deep fundamental problems that any attempt to localize cognitive functions in the brain is bound to failure. Uttal's book is organized around three main questions:

1. Can the mind be subdivided into components, modules or parts?

2. Does the brain operate as an equipotential mass or is it also divisible into interacting but separable functional units?

3. Can the components, modules, or parts of the mind, if they exist in some valid psychological sense, be assigned to localized portions of the brain?

Uttal's answer to question 2 is a limited yes, the answer to question 1 is a decided no, and therefore the answer to question 3 must also be no. That is, there can be no meaningful relation between the parts of the mind and the parts of the brain because there are no meaningful parts of the mind to speak of.

To make his case that the answer to question 1 is "no" Uttal draws on both the history of neuroscience and psychology and current findings in cognitive science. Here, Uttal divides his critique into three main parts: In the first part, he reviews the history of localization attempts; In the second part, he critiques the idea that the mind can be divided into functional parts, especially for the higher cognitive processes, and in the third part, he lists a number of problems with the technologies that have been used to infer localization of mental function from localization of brain processes.

While many of Uttal's criticisms are valid and should be borne in mind when evaluating extreme claims of localization, I find that Uttal's skeptical conclusion is unwarranted. Contrary to Uttal's claims, there has been progress in our understanding of both cognitive functions and their location in the brain, and there is every reason to believe that the problems involved with localization can be overcome, although not by the simplistic model of localization Uttal critiques.

Organization of the Book

Uttal's book is organized around three broad themes, illustrated with numerous examples. In the first chapter, Uttal outlines the general problem, as he sees it, beginning with a discussion of historical attempts to localize brain function. To Uttal, the apparent unity of the mind poses a *prima facie* problem for theories of the mind that assume that it can be divided into functional parts that can then be localized to specific brain regions.

In chapter two, Uttal briefly digresses from his main theme to explain the technologies that have been used in cognitive neuroscience. Uttal describes methods from the level of single unit recording and surgical techniques to ERPs (which he refers to as “evoked brain potentials, EVBPs”) computerized tomography (CT) scans, positron emission tomography (PET) and blood oxygen level dependent functional magnetic resonance imaging (BOLD fMRI). This section is relatively solid overall, and it is probably a good reference for introducing people to the techniques of modern cognitive neuroscience, but a number of minor flaws suggest that Uttal is either not careful in his reading, or that he is not completely prepared to discuss these issues. For example, the use of EVBP instead of ERPs suggests that Uttal is not reading current literature. Additionally, in his discussion of techniques to inactivate cortex, Uttal does not include recent developments (such as the use of the reversible GABA-agonist muscimol), which do not suffer many of the problems Uttal describes (see below). However, the most obvious error is that Uttal refers to “the central or Sylvian sulcus” (p. 30). However, the central sulcus is also known as the Rolandic fissure, while the Sylvian fissure is also known as the lateral sulcus. These are the two largest and best-known landmarks in the brain, and so Uttal should have taken extra care to get these details right. Overall, the chapter is adequate, but I wouldn’t use it to study for an exam.

In chapter three, Uttal gets into the heart of his argument, laying out his reasons for being skeptical that a taxonomy of cognitive processes is even possible. Since a great deal of the burden of Uttal’s argument lies here, I will dedicate the greatest part of this review to this portion of Uttal’s book. To briefly summarize the argument here, Uttal claims that there has been, and can be, no progress on the problem of developing a taxonomy of cognitive processes, and therefore there can be no hope of localizing cognitive processes in the brain.

In the fourth chapter, Uttal focuses on technical problems that limit the inferences about localization that can be drawn on the basis of current methods in neuroscience. Many of the criticisms in this chapter are well-known to practicing neuroscientists, but the assemblage presented here is impressive. I will argue, however, that Uttal underestimates the power of converging research methods and the self-correcting nature of science to address problems in individual studies.

In the fifth and final chapter, Uttal argues that, instead of fractionating cognitive processes into parts, and then attempting to localize these parts to specific parts of the brain, we should return to a “molar level” analysis of behavior, focusing on input-output relations, instead of postulating unobservable cognitive processes. I will briefly discuss each of these sections, addressing some concerns that Uttal raises along the way, finally arguing that Uttal’s concerns, while cause for caution, do not cause the entire enterprise of localization in the brain to unravel, and that, even more so, they do not necessitate Uttal’s proposed return to behaviorism.

A Brief History of the Parts of the Mind

Uttal’s main concern is that there has been, and can be, no progress on the problem of developing a taxonomy of cognitive processes, and therefore there can be no hope of localizing cognitive processes in the brain. “The preeminent problem in achieving a general solution to the localization issue lies in defining the psychological processes and mechanisms for which loci are being sought” (p.16).

To demonstrate this, Uttal dedicates over 40 pages to reviewing two thousand years of theories on the various different cognitive processes. While considerations of space do not allow me to reiterate Uttal's review here, anyone interested in psychology, cognitive science, or systems level neuroscience should consider reading this section.

In the fundamental taxonomy of cognitive psychology, Neisser proposed a number of categories of the mind, such as iconic storage, pattern recognition, focal attention, visual memory, speech perception and memory and thought (Neisser, 1967; Matlin, 1994). For Uttal, the fact that decades later, the main components of Neisser's taxonomy remained unchanged, simply serve to indicate that a pedagogical tool had become reified into the taxonomy of what could reasonably be sought in the brain. One alternative interpretation, not explored by Uttal, is the possibility that this taxonomy remains relatively unchanged not because it has been uncritically reified, but rather that it has provided a coherent account of mental activity that leads to testable predictions, and which has survived experimental disconfirmation. Compare this with other such theories, such as the ideas of earlier faculty psychology (the form of psychology that the early Phrenologists used), which of course are completely out of favor. Patricia and Paul Churchland have been especially vociferous advocates of the idea that our mental taxonomy is a corrigible theory, and that many ideas about the mind we currently hold will eventually turn out to be like the now discredited ideas of phlogiston or élan vital (see, e.g., P.S. Churchland, 1986; P.M. Churchland, 1989). Much future work remains to be done, but there is, again, no in principle reason to believe that we cannot develop an adequate taxonomy of the mind suitable for the task of localizing to parts of the brain.

It should also be noted that, while the fundamental taxonomy of the mind has remained relatively unchanged, the exact manner in which we have explained and modeled these mental abilities has undergone dramatic revision. In the early days of cognitive science, the mind was conceived of as a collection of black boxes, each autonomously performing computations on representations that were in some way a "language of thought" guided by sentence-like rules. With the advent of connectionism and alternative approaches to cognition, we now realize that mental computations are more likely carried out by interconnected networks of neurons, in which information is stored as patterns of synaptic strength, and computation can be thought of as transformations of high-dimensional neural vectors (see, e.g., Churchland & Sejnowski, 1992). These drastic changes of perspective at the microlevel have been mirrored by changes at the macrolevel by the realization that a great deal of cognitive activity takes place in situated groups of agents, interacting with their environment (Hutchins, 1995). One of Uttal's concerns (p. 143) is that the rigid, serial models of cognitive psychology are implausible. However, much work in cognitive science, building from neural networks, also denies this rigid, serial model, and yet attempts to localize cognitive functions to parallel distributed networks in the brain.

Another main concern for Uttal is the ever-greater number of cognitive parts that have been posited by localizationists. However, a careful analysis of the literature demonstrates that the path to progress in cognitive science lies in subdividing the components of the mind, which we can then attempt to localize to parts of the brain. For example, consider the case of memory. Early psychological investigations of memory treated it as a unitary construct. However, more recent analyses of memory, combining both psychological and neuroscientific approaches have suggested that "memory" really should be thought of as a variety of separable components with different proposed neural substrates (although the exact taxonomies of memory are subject to debate). One such classical example comes from Larry Squire's work on patients with hippocampal damage, in which the ability to form new "declarative" memories is impaired, while the ability to learn new "non-declarative" memory for skills and abilities is spared (for a review, see Squire & Kandel, 1999). While a great deal of work, both empirical and theoretical, remains to be done, the basic lesson that progress in cognitive science will come through subdivision of cognitive functions is one that seems to have been completely lost on Uttal.

Modularity and Localization

Uttal, following Fodor (1983), argues that certain input and output functions are modular, but central cognitive processes may not be. On Fodor's account, there are several important distinctions between the input and output processes and the cognitive processes. First, information in the input and output processes can be meaningfully encapsulated. Second, each of the input processes is clearly linked to a specific anchor of a specific class of physical stimulus. Third, the dimensions of the physical stimulus can be clearly identified, and therefore the dimensionality of the mental experience can also be clearly identified.

Based on these results, Uttal claims that, even if it were to turn out that certain input and output functions could be localized, it is not clear that cognitive functions could be localized because the taxonomy of mental processes seems to be open ended and there is no way to achieve closure in such a list of possible faculties. However, the discussion of memory, above suggests that it may be possible to localize certain cognitive capacities (e.g., memory) in the same way that we have localized certain perceptual capacities. Additionally, much recent evidence, such as that coming from embodied cognitive linguistics, has suggested that the strict distinction between perceptual-motor and cognitive functions may be illusory (see e.g., Lakoff & Johnson, 1999; Johnson, 1987) thus making it less clear exactly where Uttal would draw the dividing line between those processes that can be localized and those that cannot.

Uttal further argues, even if perceptual and motor functions can be identified with specific brain regions, "there is considerable evidence that…any but the simplest (sensory or motor) cognitive function involves large and distributed regions of the brain" (p. 155). However, the problem here for cognitive science is only apparent. One of the fundamental strategies in cognitive science is the technique of recursive decomposition (Palmer & Kimchi, 1986), in which a single component (e.g., the mind) is divided into smaller and smaller functional units (e.g., perception, attention, memory or action). These components can then be further recursively decomposed into smaller and smaller parts (e.g., processing of color, motion or orientation), which can then be localized to specific parts of the brain (e.g., cortical visual areas V4, MT and V1).

However, it is not only perceptual functions like vision that can undergo recursive decomposition. Consider the case of mathematics, a paradigmatic case of a cognitive function. To define the problem more carefully, where in the brain do arithmetic calculations occur? There is no one place in the brain where arithmetic calculations occur. Instead, arithmetic calculations activate a widespread network of brain regions, consistent with Uttal's arguments. However, it is also possible to decompose the process of performing arithmetic problems into several key components. First, a visually presented number (a grapheme) must be recognized. Second, the numerical magnitude of the grapheme must be recognized, and finally, computations must be invoked to determine the correct solution for the problem at hand. Additionally, this process may draw on stored math facts such as the overlearned knowledge that "2 + 2 = 4." This model, dubbed the "triple-code" model, (Dehaene, 1992; 1997) was derived from a variety of different sources of evidence, including reaction time studies, neuropsychological deficits, and early imaging (PET) studies.

Not only does this model propose a specific taxonomy of the cognitive processes involved, but it also suggests that specific regions of the brain would be essential for the performance of these functions. Based on previous findings and the new, refined framework in which to study arithmetic calculation, the recognition of the visual grapheme could be expected to depend on structures in the inferior temporal lobe (fusiform and lingual gyri) while the ability to perform the numerical calculation would likely depend on the angular gyrus, and the ability to retrieve stored mathematical facts would depend on frontal lobe structures. This model has since received further verification from studies making use of a variety of techniques. For example, using cortical recording electrodes Allison *et al*. (1994) find that numbers and letters are represented in fusiform regions near other regions associated with the recognition of visual objects, while Whalen *et al*., (1997) find that stimulating frontal cortex causes a temporary disruption of stored arithmetic facts. Göbel *et al*., (2001) report that transcranial magnetic stimulation (TMS) over angular gyrus disrupts the ability to perform novel computations but does not impair the ability to report the answer to overlearned problems like 2 + 2. Finally, several recent fMRI studies have confirmed and extended this overall model (Pesenti *et al*. 2000; Rickard *et al*. 2000), although as in the case of memory, there remain unresolved questions.

One of the main points of this extended example is that it matters what you are attempting to localize. That is, there is an interaction between how cognitive processes are defined and how they are localized. Attempts to localize mathematical cognition would fail because performing arithmetic problems depends on a distributed network of brain areas. However, if we can appropriately define the subcomponents of mathematical cognition, it may be possible to identify different cognitive parts and localize them to specific brain parts. Returning to the epigraph at the beginning of this review, brain functions are neither localized, nor distributed throughout the whole brain, and this is the modal view in cognitive neuroscience today. By considering the interaction between the cognitive parts that we postulate and our ability to localize them, we see that claims that cognitive functions are distributed is an argument against something of a localizationist straw man, a view that very few (if any) serious researchers in this area actually believe.

Technical Problems in Lesion Analysis

In order to further support his claim that cognitive functions will never be reduced to specific neural substrates, Uttal describes a number of technical problems that limit the inferences about localization that should be drawn on the basis of the lesion method (both natural lesions in humans and surgical ablation techniques in animals).

For example, Uttal points out that the various ablation techniques (ways of removing brain tissue to see what happens when that tissue is removed) all suffer different forms of limitations, such as effects due to inadvertently severed fibers of passage or due to recovery of function. However, neuroscientists are aware of these difficulties and are always searching for newer and better techniques to overcome these difficulties. Recently, there have been a family of antagonists developed that are able to reversibly, selectively inactivate certain populations of neurons (for example, the GABA agonist muscimol described above), overcoming some of the difficulties Uttal has discussed with this class of experiments.

Another problem for the localizationist is the heterogeneity of lesions and the individual variability in brain organization. Although it has been at the core of the localizationist enterprise since the earliest days, the lesion method still suffers from numerous problems of interpretation. After the downfall of phrenology, the idea of localization of function was revived by Paul Broca's studies of patients with brain lesions in the 1860s (Greenblatt, 1995; Star, 1989). Broca studied numerous patients with deficits in their ability to produce language (aphasia), independent of any motor problems (dysarthria) and concluded that the left inferior frontal cortex was critically involved in language production. Similarly, Carl Wernicke proposed that an area in the left posterior superior temporal region was critical for comprehension of language.

Most people are familiar with this basic taxonomy, which is still in use today: Broca's area is critical for language production and grammar while Wernicke's area is crucial for comprehension. However, more recent evidence has demonstrated that lesions to insula lead to nonfluent aphasias more so than Broca's area (Bates *et al*., submitted; Dronkers, 1996) and that many different lesions can cause agrammatism (Dick *et al*., 2001). Additionally, recent research suggests that Broca's area may house the human mirror neuron system (Iacoboni *et al*., 1999), although this may not necessarily be a serious contradiction, as the mirror neuron system may be critical for the evolution of language (Arbib & Rizzolatti, 1996). Similarly, Wernicke's area is believed to be critical for speech comprehension and phonological encoding, but others have found (Saygin *et al*., in press) that Wernicke's area lesions can also cause deficits in the comprehension of nonverbal environmental sounds (even significantly more than speech sounds) and thus processing in this region is probably not specific language, but instead may be involved in any complex auditory analysis. Again, these problems demonstrate that how you divide up and refine the processes within a cognitive task/domain is important, and we should always be aware of these problems when pursuing the localizationist goals. However, a look at the study of language localization also demonstrates that conceptual revision and progress is still occurring, sometimes top-down, as we come to better understand language, and sometimes bottom-up, as we explore its neural basis (see also, Bates & Dick, 2000).

Technical Problems with Functional Imaging

In addition to the problems with the lesion method, Uttal lists a number of concerns about one of the primary methods used in cognitive science, functional imaging. This concern about the interpretation of imaging data is, of course, not new (for a review of the problems with fMRI and the inferences that one can and cannot make on the basis of the BOLD signal, the interested reader is referred to Cabeza and Kingstone, 2001 or to any of the numerous online texts discussing potential MRI artifacts see, e.g., Hornak, 1999).

One of the concerns that most clearly relates to our current concerns about localization is the problem of the arbitrary threshold in fMRI. As Uttal notes, "Evidence of sharply defined and highly localized artificial boundaries arising from a poor choice of a threshold could easily lead to an erroneous conclusion about the cerebral localization or nonlocalization of a psychological process" (p. 167-168). This is especially true in the case of certain classes of statistical analyses.

However, these types of threshold analyses are not the only methods for analyzing fMRI data available, and as new and more powerful techniques are developed, new analyses (such as independent components analysis, principal components analysis or network connectivity analysis) will be performed that will provide an alternative to the arbitrary threshold setting methods. For example, in much current vision research, one technique is to use a "reference" scan against which to have a baseline or standard and then to compare activation relative to that standard. This does not entirely avoid the problem of thresholds, but can make them a little less arbitrary. As these types of techniques find their way into higher cognitive processes, the problem of the baseline will become less arbitrary.

An additional change in imaging methods is a return to region of interest (ROI) analyses. In the early days of functional imaging, ROI analyses were common, as researchers were interested in simply confirming the validity of the new techniques. However, as researchers became more confident in their techniques, imaging became more and more commonly used as an exploratory tool. When used as an exploratory tool, researchers often simply look for the "hottest" spot in the brain and leave their analysis at that. More recently, however, there has been a return to the ROI technique to explore activity in specific brain regions, on the basis of a prior hypothesis, and to quantify the manner in which the brain responds to various different types of stimulation and in different tasks. This hypothesis driven usage of fMRI can be used to avoid many of the problems associated with the use of arbitrary thresholds.

The Scientist in a Vacuum

There is, however, a deeper problem with Uttal's analysis, that I will call the "scientist in a vacuum fallacy." Throughout the book, Uttal raises valid concerns about the extent to which we can infer localization from any one experiment or any one methodology. For example, from the fact that damage to one region of the brain causes a particular deficit, we cannot infer that the part of the brain damaged is necessarily the seat of that function (if you remove a transistor from a radio and it starts to squawk, it doesn't follow that the transistor is a squawk suppressor). Or, finding a particular peak of activity in the brain doesn't imply that the locus of that particular peak is "the" region of the brain that performs this particular function.

No scientist works in a vacuum. Rather, every scientist has colleagues and competitors who reinforce or challenge his or her results (see, for example, Latour 1987). These colleagues and competitors may replicate (or fail to replicate) any of the results that a given researcher may arrive at. So, for example, if a given researcher were to find a unique spot of activity by chance alone that actually had no bearing on the cognitive functions being investigated, then other future researchers would fail to find a similar spot of activity, and the results would therefore be regarded as an error or an anomaly. That is, although Uttal has raised the specter that noise may contaminate the results of a single experiment, or even a family or experiments, he has not demonstrated that there is a clear bias in the methods of cognitive neuroscience, writ large, that would lead to *systematic* errors regarding whether cognitive functions are localized or not.

Conversely, if the same localization was to be supported by similar results from deficits after brain lesions, deficits induced by TMS, and increased activity as indicated by the fMRI BOLD signal, we would then feel justified in concluding that there was in fact a real neural basis for the proposed localization. This is because different techniques have different strengths and weaknesses and different potential artifacts. The possibility that each method would independently hit upon the same location simply due to these technical issues is vanishingly small. For example, in the case of mathematical cognition above, many researchers, making use of many different techniques have reached similar conclusions regarding the localization of the cognitive processes involved in arithmetic cognition. The collaborative, self-correcting nature of scientific research and the use of multiple converging research methodologies make it likely that future localizationist ventures will be more and more accurate.

Conclusions

The history of localizationist research has long been fraught with bitter battles as to whether the brain, like the other organs in the body, can be divided into subparts, which perform different functions. Uttal would have us believe that the answer to this question is no. Instead, he argues that the mind is a unitary phenomenon, and that attempts to subdivide the mind suffer from severe logical difficulties, especially since the mental faculties that we are attempting to localize can only be observed indirectly. However, this problem is no greater than the problem faced by other domains of science. For example, particle physicists cannot directly observe sub-atomic particles. Instead, they must infer their existence from the behavioral effects on other physical measuring devices, such as a cloud chamber. Therefore, there is nothing uniquely privileged about mental states.

Since the mind can be meaningfully decomposed into subprocesses (e.g., memory, language comprehension), and can be inferred from the behavior of the system, there is no logical, in principle, barrier to developing an appropriate taxonomy of the mind even though there is still a great deal of debate about the details of this taxonomy. Uttal's historical arguments, contrary to his conclusions, do demonstrate that significant progress has been made in our understanding of the taxonomy of the mental, and further suggest that a growing consensus may be at hand. With this high level description of the mind in hand, the attempt to localize these processes in specific parts of the brain may move forward rather rapidly.

Uttal's specific concerns about the techniques used to explore the neural basis of these cognitive parts are hardly unique. These concerns are shared (and indeed have often been raised) by the very people that are attempting to localize cognitive processes to specific parts of the brain. I would argue that Uttal's concerns are, in part due to an outdated view of what the localizationist enterprise is attempting to accomplish. To localize macrolevel cognitive processes like memory or mathematics to specific neural substrates is clearly an untenable goal (as Hughlings-Jackson noted over 100 years ago) and very few researchers in cognitive neuroscience are still pursuing this simplistic goal. Additionally, Uttal's arguments suffer from a very unrepresentative view of a scientist in a vacuum, working without benefit of colleagues and competitors and restricted to only one methodology. If this were the true situation of the scientist, it would be a dire situation indeed. However, given that scientists do not work in a vacuum, these concerns should not limit the possibility of the localizationist enterprise in toto. Instead, it should merely limit our enthusiasm for specific experimental results.

It should also be noted that Uttal's approach assumes a purely top-down answer to the question of localization. That is, once we have an accurate cognitive taxonomy in hand, then, and only then, will it be possible to even begin to attempt to localize mental parts in brain parts. However, the study of brain parts can also provide a bottom-up constraint on our theories of the mental (cf. P.S. Churchland, 1986). Just as placing a couch in a living room begins to constrain the possible locations of the rest of the furniture in the room, so too does an understanding of the brain begin to constrain the theories about cognitive parts that we entertain.

In the final analysis, the conclusions of this book are truly disappointing. Uttal argues that, instead of fractionating cognitive processes into parts, and then attempting to localize these parts to specific parts of the brain, we should return to a "molar level" analysis of behavior, focusing on input-output relations, instead of postulating unobservable cognitive processes (in essence, a thinly veiled return to behaviorism). Uttal's concerns, while cause for caution, should not cause the entire enterprise of localization in the brain to unravel, and that, even more so, they do not necessitate a return to behaviorism.

In contrast to Uttal's claims, I believe that it is quite probable that we will, piece by piece, come to an understanding not only of the parts of the mind, but also of their localization to the parts of the brain. In contrast to Uttal's conclusion that we need to return to a molar level description of behavior focusing on input-output relations, I believe that progress in localizing cognitive functions to brain areas will occur when we divide the mind more and more finely, not less. Currently, we have but the dimmest glimmer of what such a theory might look like, but its outlines are beginning to take shape. As the form becomes clearer and we part the mists that obscure our view of the relation between the mind and the brain, we will one day be able to answer that ultimate question of who we are, and what makes us think.

References

Allison, T., McCarthy, G., Nobre, A., Puce, A. (1994). Human extrastriate visual cortex and the perception of faces, words, numbers, and colors. *Cerebral Cortex, 4* (5):544-554.

Arbib, M. A. & Rizzolatti, G. (1996). Neural expectations: A possible evolutionary path from manual skills to language. *Communication & Cognition*, *29*(3-4):393-424.

Bates, E., & Dick, F. (2000) Beyond phrenology: Brain and language in the next millennium. *Brain and Language, 71*, 18-21.

Bates, E., Wilson, S., Saygin, A.P., Dick, F. Sereno, M., Knight, R. T. & Dronkers, N. F. (submitted). Assessing brain-behavior relationships using voxel-based lesion-symptom mapping.

Binder, J. R. & Rao, S. M. (1994). Human brain mapping with functional magnetic resonance imaging. In A. Kertesz (Ed.) *Localization and neuroimaging in neuropsychology*. (pp. 185-212) San Diego, CA: Academic Press.

Cabeza, R. & Kingstone, A. (Eds, 2001) *Handbook of functional neuroimaging of cognition*. Cambridge, MA: MIT Press.

Churchland, P. M. (1989). *A neurocomputational perspective : the nature of mind and the structure of science.* Cambridge, MA: MIT Press.

Churchland, P. S. (1986). *Neurophilosophy: toward a unified science of the mind-brain.* Cambridge, MA: MIT Press.

Churchland, P. S. & Sejnowski, T. J. (1992). *The computational brain*. Cambridge, MA: MIT Press.

Dehaene, S. (1992). Varieties of numerical abilities. *Cognition, 44* (1-2):1-42.

Dehaene, S. (1997). *The number sense: How the mind creates mathematics*. New York: Oxford University Press.

Dick, F., Bates, E., Wulfeck, B., Utman, J., Dronkers, N. & Gernsbacher, M.A. (2001). Language deficits, localization and grammar: evidence for a distributive model of language breakdown in aphasics and normals. *Psychological Review, 108*(4), 759–788.

Dronkers, N. F. (1996). A new brain region for coordinating speech articulation. *Nature*, *384*(6605): 159-161.

Fodor, J. A. (1983). *The modularity of mind: An essay on faculty psychology*. Cambridge, MA: MIT Press.

Gazzaniga, M. S., Ivry, R. B. & Mangun, G. R. (1998). *Cognitive neuroscience: The biology of the mind*. New York: W.W. Norton.

Göbel, S., Walsh, V. & Rushworth, M. F. S. (2001). The mental number line and the human angular gyrus. *NeuroImage, 14*(6): 1278-1289.

Greenblatt, S.H. (1995). Phrenology in the Science and Culture of the 19th Century. *Neurosurgery*, *37*:790-805.

Hornak, J.P. (1999) *The basics of MRI*. http://www.cis.rit.edu/htbooks/mri/mri-main.htm.

Iacoboni, M., Woods, R. P., Brass, M., Bekkering, H. (1999). Cortical mechanisms of human imitation. *Science, 286*(5449):2526-2528.

Jackson, J.H. (1882/1932). On some implications of dissolution of the nervous system. *Medical Press and Circular 2*; Reprinted in J. Taylor, (Ed. 1932) *Selected Writings of John Hughlings Jackson*. (pp. 29-44) New York: Basic Books.

Johnson, M. (1987). *The body in the mind : The bodily basis of meaning, imagination, and reason*. Chicago: University of Chicago Press.

Hutchins, E. (1995). *Cognition in the wild.* Cambridge, MA: MIT Press.

Lakoff, G., & Johnson, M. (1999). *Philosophy in the flesh : The embodied mind and its challenge to Western thought.* New York : Basic Books.

Latour, B. (1987). *Science in action: How to follow scientists and engineers through society*. Cambridge, MA: Harvard University Press.

Matlin, M. W. (1994). *Cognition* (3rd ed.). New York: Harcourt Brace.

Neisser, U. (1967). *Cognitive psychology.* New York: Appleton-Century-Crofts

Palmer, S. E. & Kimchi, R. (1986). The information processing approach to cognition. In T.J. Knapp & Lynn C. Robertson, (Eds.) *Approaches to cognition: Contrasts and controversies*. (pp. 37-77) Hillsdale, NJ: Lawrence Erlbaum Associates.

Pesenti, M., Thioux, M., Seron, X., De Volder, A. (2000). Neuroanatomical substrates of Arabic number processing, numerical comparison, and simple addition: A PET study. *Journal of Cognitive Neuroscience*, *12*(3):461-479.

Rickard, T. C., Romero, S. G., Basso, G., Wharton, C., & Grafman, J. (2000). The calculating brain: An fMRI study. *Neuropsychologia*, *38*(3):325-335.

Saygin, A. P., Dick, F., Wilson, S., Dronkers, N. & Bates, E. (in press). Neural resources for processing language and environmental sounds: Evidence from aphasia. *Brain*, *126*.

Star, S. L. (1989). *Regions of the mind: brain research and the quest for scientific certainty*. Stanford, CA: Stanford University Press.

Squire, L. R. & Kandel, E. R. (1999). *Memory: From mind and molecules*. New York, NY: Scientific American Books.

Uttal, W. R. (2001). *The new phrenology: the limits of localizing cognitive processes in the brain*. Cambridge, MA: MIT Press.

Whalen, J., McCloskey, M., Lesser, R.P. & Gordon, B. (1997). Localizing arithmetic processes in the brain: Evidence from a transient deficit during cortical stimulation. *Journal of Cognitive Neuroscience*, *9*(3):409-417.