

The massive redeployment hypothesis and the functional topography of the brain

Michael L. Anderson
Institute for Advanced Computer Studies
Program in Neuroscience and Cognitive Science
University of Maryland
College Park, MD 20742

Abstract

This essay introduces the massive redeployment hypothesis, an account of the functional organization of the brain that centrally features the fact that brain areas are typically employed to support numerous functions. The central contribution of the essay is to outline a middle course between strict localization on the one hand, and holism on the other, in such a way as to account for the supporting data on both sides of the argument. The massive redeployment hypothesis is supported by case studies of redeployment, and compared and contrasted with other theories of the localization of function.

Keywords

localization, functional localization, holism, brain imaging, modularity, fMRI

1. Introduction

The brain, it needs hardly be said, has many functions. Even focusing on the so-called cortical functions—and thus leaving out the largely unconscious, automatic, “lower” functions like homeostatic regulation, balance, and the like—leaves quite a large number of things for the brain to do, from moving one’s fingers (say, to turn the pages of this essay), to reading and understanding these sentences (because of which the pages will soon need turning). Naturally, there is a great deal of interest in how the brain does these things, and that has generally involved interest in *where* the brain does them. This interest in the functional topography of the brain is at least as old as Gall’s phrenology (Gall, 1798), and has undergone a recent resurgence, due in large measure to advances in various imaging technologies. The main purpose of the current essay is to outline an idea

regarding what we should generally expect the functional topography of the brain to be, an idea I call the massive redeployment hypothesis.

In the next section, I will discuss what I take to be the current, orthodox position on the functional topography of the brain—something I call strict localization—as well as the main alternative to this view, functional holism (Uttal 2001, Lloyd 2000). These will serve as opposing contrasts to the massive redeployment hypothesis (henceforth MRH), also introduced in section two. As will become clear, MRH offers a middle ground between strict localization on the one hand, and holism on the other, and in such a way that offers an explanation of—or at least does some justice to—the evidence on both sides of the debate. Section three will provide some evidence for and illustrations of MRH, in the form of three case studies of the re-deployment of brain areas in different functions. Section four outlines some more broad-based evidence for MRH, by way of an empirical review of 135 brain-imaging experiments, and section five will provide a detailed account of MRH. Finally, section six will briefly discuss some of the implications of, and possible objections to, MRH.

2. Cognitive functions and brain areas.

Roughly speaking, a cognitive function is a process of cognitive or psychological interest (e.g. that supports or causes some cognitive effect like intentionally moving one's fingers, recognizing a face, or remembering a list of words) that can be specified in terms of inputs, outputs, and the (often mysterious) transformation of the former into the later. In this I am following Fodor (2000) in the pragmatic definition of a (cognitive) function

as whatever appears in one of the boxes in a psychologist's diagram of cognitive processing.¹ The project of functional topography, simply put, is to map these boxes onto brain areas. In imaging studies, this generally involves getting subjects to engage in some cognitive task, and seeing what "lights up", that is, finding brain areas, the activation of which is reliably and uniquely correlated to the cognitive activity in question, once noise and background processing are subtracted out.

So far, so good. What the brain topographer wants to know is which brain areas participate in which functions. Let us say a brain area *participates* in a function if activity or processing in that area supports the transformation of inputs to outputs that define the function. A given function may (and typically does) have more than one participant; among these, there may be those the failure of which would cause the failure of the function; let us call these the *necessary* participants.² In addition, there may be areas that participate in only one function; let us call these the *exclusive* participants. Finally, let us call the collection of participants in a certain function a *functional complex*.

Now, let us assume, for the sake of the discussion, that the brain topographer is interested only in the necessary participants in cognitive function. And let us assume further that many functions have exclusive participants (something I think is possibly true, but need not be). It does not follow from these assumptions that most or all of the necessary participants in a function are exclusive participants, nor that most or all functions *have* exclusive participants. Yet these are clearly the predominant working assumptions of brain topography; let's call it the *strict localization hypothesis*.

2.1 Strict localization and scientific practice

Do working scientists actually accept the strict localization hypothesis? One might be prepared to doubt it, if only in light of the caution most scientists employ when stating their conclusions. Consider the following, multiply hedged passage from Grasby *et. al.* (1993: 1): “The brain areas identified in these comparisons define a number of the neuroanatomical components of a distributed system for signal processing and storage relevant to auditory-verbal memory function.” And yet the impression that they are in fact committed to strict localization is encouraged by the many available brain maps—such as might be found in any text of cognitive neuroscience—showing the functional breakdown of the brain in terms of neat, contiguous, non-overlapping brain areas. And, of course, for every cautious statement can be found a corresponding less cautious one.

The following are cited by Lloyd (2000):

The implications of these results are discussed, and it is argued that they are consistent with localization of a lexicon for spoken word recognition in the middle part of the left superior and middle temporal gyri, and a lexicon for written word recognition in the posterior part of the left middle temporal gyrus. (Howard *et al.* 1992: 1769).

“These data localize the vigilance aspects of normal human attention to sensory stimuli . . .”(Pardo *et al.* 1991: 61).

Indeed, even when there are controversies regarding the function of a given area—as for instance in the ongoing controversy about how to explain the fact that the so-called “fusiform face area” (Allison, 1994; Haxby, *et. al.* 1994; Kanwisher, *et. al.* 1997; McCarthy, *et. al.* 1997) seems to be involved in processing stimuli other than faces (Gauthier, *et. al.* 2000; Tarr & Gauthier, 2000)—the arguments are typically directed not toward identifying the multiplicity of functions in which the area(s) of interest

participate, but rather toward properly defining *the* unique function in which a given area participates.

Moreover—and perhaps more tellingly—even if few scientists would embrace the strict localization hypothesis when it is put starkly before them, their methodology is at least somewhat predisposed to produce data consistent with it. What I mean is this: a typical design for an imaging experiment involves subjects who perform a number of trials of an assigned task, along with control trials that generally involve the subject engaging in a different but closely related task (but may also consist of resting periods, e.g., a low-level reference). The experimenters image the brain during all these periods, and the data analysis involves comparing—within subjects, between subjects, or both—on-task activations with control activations, and *subtracting out* what is common; the images that result are *difference* images. I intend no broad attack on this method,³ nor will I suggest that the results garnered from its use are in any way invalid. That such experiments tell us something important is not in question. There is a serious question, however, regarding how to *interpret* fMRI data, and what kinds of conclusions they in fact support.⁴ In my view, the brain areas that are shown to be active in the on-task trials are indeed participants, and in most cases *necessary* participants, in the task in question. However, it also appears that necessary but *non-exclusive* participants will often be, or are at least at risk of being, subtracted out (and the more closely related the experimental and control task, the higher the risk is likely to be).

Why this is so is easy to see. Imagine that the following table represents the average⁵ level of activity in five brain areas, in the task and control conditions.

Brain area	A	B	C	D	E
Activation					
Task	5	6	2	1	0
Control	2	6	5	1	0
Difference	3	0	-3	0	0

Table 1: Imaginary brain area activations for task and control conditions

A brain image created using the subtraction method (based on the difference line in Table 1) would strongly suggest that area A was the brain area responsible for the main experimental task (and perhaps that area C was involved in the control task, or inhibited by the experimental task). But note the high activation in area B for both task and control. There are a number of possible explanations for this: for instance, that area B has generally high activation having nothing to do with either the task or the control; that area B has high activation because it participates in the processing of the *inputs* for both the task and control; or that area B is a necessary (but non-exclusive) participant in both the experimental *and* the control task. In the first two cases, it would be both legitimate and desirable to subtract out the activation of area B; but in the third case the effect of subtraction is the loss of valuable information, and potentially misleading support for strict localization.

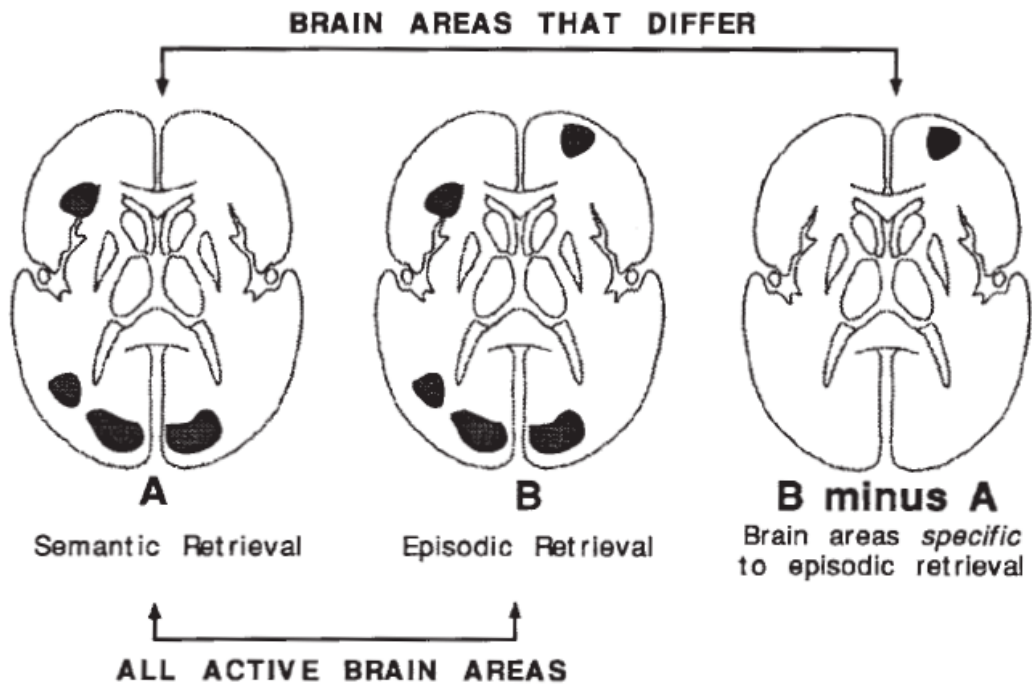


Figure 1: Representative activations are displayed for two different memory tasks (from Buckner, Petersen, et al., 1995), one relying on semantic retrieval (left section, labeled A) and one relying on episodic retrieval (middle section, labeled B). (Reprinted with permission)

This worry is not merely abstract, but reflects an interpretive choice faced by working scientists, one that also imposes an interpretive constraint. In a review of work on the role of prefrontal cortex in long-term memory, in which both semantic and episodic memory was being investigated, Randy Buckner writes:

[T]he data can be discussed and interpreted in one of two ways, both of which are correct (Figure [1]). One way to examine brain areas active during semantic and episodic retrieval tasks is to compare the tasks with a low-level control task that does not make any memory demands. Such comparisons reveal all the brain areas active during the tasks—those which overlap as well as those which are unique to the two tasks. Alternatively, if one wanted to isolate brain areas specific to episodic retrieval, one might compare it with a similar task that required only semantic retrieval. This approach would isolate areas selective for episodic retrieval but would not necessarily reveal all the areas activated by the episodic retrieval task. Areas shared in common by both the semantic and the episodic retrieval tasks would be missed. There are tradeoffs between the two kinds of comparisons. Well-controlled comparisons serve to better isolate cognitive processes and, presumably, isolate brain areas differentially involved in those processes. Such comparisons, however, potentially miss important brain areas being activated by the task, simply because the reference control task is also activating those brain areas. Comparisons involving low-level reference tasks

identify more completely the brain pathways activated but are difficult to interpret because they are underconstrained. (Buckner, 1996: 154)

So difference images, especially when they involve subtractions of closely related tasks (as in Figure 1), are likely to remove the necessary, non-exclusive participants in the task. Thus, we need to be mindful of this possibility, and not conclude that a given difference image necessarily reveals the *only* participants in a given task or function. But can we nevertheless conclude, from subtraction images produced by narrowly focused, well-designed studies, that the areas shown to be active are the *exclusive* participants in the task under consideration? Buckner appears to endorse this reading, noting that the B minus A image allows one to “isolate brain areas specific to episodic retrieval” (Buckner, 1996: 154).⁶

By way of returning to the question with which this section began—whether working scientists accept the strict localization hypothesis—it is worth noticing that, even while sounding a cautionary note regarding the interpretation of imaging data, Buckner is nevertheless willing to describe the B minus A image as revealing the brain areas specific to, or specialized for, episodic retrieval. And as can be seen by again glancing at the localization quotes above, Buckner is far from alone in a tendency to interpret such data as revealing the exclusive participants in the tasks under consideration. The answer to the question would seem to be “yes”. Still, whether any given scientist believes in strict localization is less important than whether strict localization is in fact justified by the data. And as we will see in the next section, even highly focused, well-designed studies in no way establish strict localization, and, taking a somewhat broader view, one can easily find independent reasons to doubt it.

2.2 Redeployment in the brain

So, why can't we conclude from the studies summarized above that a particular area of right prefrontal cortex is specific to (an exclusive participant in) episodic retrieval? The answer is simple: because there are in fact many *other* functions in which the right prefrontal cortex participates, including the management of anxiety (Wang, et al., 2005) and the generation of hypotheses in unconstrained situations (Vartanian and Goel, 2005). Indeed, the areas identified as participating in these three quite different functions have significant overlap.⁷ There appear to be areas of right prefrontal cortex, then, that participate in at least three significantly different functions: episodic retrieval, anxiety management, and hypothesis generation. Thus, while difference images can show areas that participate in one task and not another, they cannot show that the area is *limited to* that task—and, in point of fact, the area under discussion is *not* limited to the tasks highlighted by any one these studies. As I would like to put it, in the course of its normal operation, this area of right prefrontal cortex is *redeployed* to support (perhaps many) different functions.

As we will see in much more detail below, right prefrontal cortex is in no way unique in this regard. But even from what little we have seen so far, the outlines of an alternative to strict localization should be clear. It is widely recognized that cognitive functions typically have several necessary participants. But the data we have reviewed above also suggests that individual brain areas can be participants in several cognitive functions. That is, (i) a typical cognitive function requires the participation of more than one brain

area, and (ii) each brain area may be a participant—may be redeployed—in support of other cognitive functions. Let us call any position on brain function that accepts (i) and (ii) above a *redployment hypothesis*. In contrast, we'll call any position that denies (ii) a *localization hypothesis*.

The main purpose of this essay is to defend redeployment, and question localization, as the fundamental organizational principle of the brain. However, before getting to the *evidence* for redeployment, we need to get a bit clearer about the principles defining redeployment generally, because the second point—that individual brain areas can be participants in more than one cognitive function—admits of at least two interpretations. On one interpretation, the redeployed brain area does the same thing (at some level of description) in each instance of redeployment, and differences in function are the result of differences in the structure and dynamics of the functional complex as a whole. On the other interpretation, by contrast, the reason it is possible for brain areas to be redeployed to support different cognitive functions is because the area *does something different* in each case of redeployment; in becoming part of a different functional network, the area becomes a non-dissociable part of the overall implementation of the network's function. This latter interpretation is the defining claim of functional holism, and so as the last step in situating the position being offered here, we should try to say specifically how MRH (or any RH) differs from holism.

2.3 Holism

Anyone familiar with the case for functional holism will have noticed that many of the arguments I have used above are quite similar to those deployed by the holists. Given that we both object to strict localization, this should not be too surprising. Thus, for instance, Dan Lloyd (2000) questions the plausibility of localization by criticizing the subtractive method along the same lines as I have above, and by performing a review of 35 PET experiments, each investigating a different cognitive function, from movement to memory to language understanding. He shows that each of the thirty-five different tasks activated, on average, 3.3 different Brodmann areas; likewise, each of the thirty-one Brodmann areas participated, on average, in 3.4 different tasks. He writes:

[Brain] areas are often involved across stimulus and response types, subsuming a variety of perceptual and cognitive tasks. This multifunctionality of the average Brodmann area undermines modularity. Each area has more than one job to perform. [Thus], any hypothesis of functional localization must be severely hedged. For each conclusion of the form, “Subnet S computes function *f*,” we must substitute “Subnet S computes function *f*, *among others*.” This is not a trivial emendation. (Lloyd, 2000: 98)

One weakness of this study is that each Brodmann area is relatively large,⁸ and it could well be that the areas of activation in the same Brodmann area do not in fact overlap (especially given the small number of experiments reviewed). Likewise, the study does not appear to control for laterality (e.g. whether the Brodmann area is in the right or left hemisphere). Nevertheless, the review is striking, and certainly defies the spirit of strict localization—and its essential findings are confirmed by the larger and more carefully designed empirical review discussed in section four.

Similarly, discussing a delayed-response experiment by Jacobsen (1935), Uttal writes:

My contention is that the brain localization aspects of this experiment should be interpreted differently from the way they usually are. The prefrontal region should be considered, not as the “locus” of the mechanisms responsible for the delayed-response behavior, but rather as one possible region among many whose manipulations may affect, control, or influence that particular behavior. (Uttal 2001: 17-18)

But while we all reject strict localization, Uttal and Lloyd appear to endorse holism as the proper alternative. As mentioned already above, I define holism with respect to brain function as the claim that the same region of the brain can do different things at different times. One can imagine a radical holism, whereby everything the brain does is done with all the brain—along the lines Lashley’s mass action hypothesis (Lashley 1950)—but probably no one would defend this position.⁹ The real question is whether a more limited form of holism—some parts of the brain do different things some of the time, but no part does all things at all times—is the right alternative to localization. Both Uttal and Lloyd appear to think so. Uttal is a bit difficult to pin down in this regard, as he spends most of his book arguing that we have very little evidence for localization—because of numerous technical, epistemic, and methodological limitations that (he argues) are too often played down or ignored when interpreting neuroscientific data—and much less time offering a specific alternative. It is clear that he does not believe that the brain is an equipotent mass, but while he thinks there is *some* functional organization to the brain (no part of the brain does everything, nor is everything done with all the brain), he also seems to think that the brain does a lot of things with a lot of its parts. He writes, for instance:

An alternative, and perhaps more realistic, point of view to the notion of isolatable cognitive-neural modules postulates a complex mind-brain system instantiated as a unified entity in which the various parts interact too strongly to be isolated. That they cannot be isolated from each other has to do with their fundamental nonlinear nature and heavy interconnectedness, and not with inadequate research tools or incomplete data. (Uttal, 2001:204)

Likewise, Dan Lloyd offers a theory of sparsely distributed networks, whereby “...anatomically defined brain regions are *multifunctional*. A region may be recruited to join a subnetwork to compute one function, and later recruited to a different subnet to compute a different function.” (Lloyd, 2000: 95) If his hypothesis is that brain regions do different things (compute different functions) at different times, then his position, too, is a form of holism. Nevertheless, a note of caution is in order here: it *could* be that Lloyd’s claim is not that the *region* computes a different function in each of the subnetworks of which it is a part, but rather that the *subnetwork* computes a different function, utilizing the resources of the region (which does the same thing in each sub-network).¹⁰ If *this* is his position, then Lloyd may be offering a redeployment hypothesis, and not a form of (limited) holism. However, insofar as Lloyd is committed to a connectionist architecture, it is not at all clear how it would be possible for him to make this latter claim. In a given connectionist network, is not generally possible to specify *what* a given part is doing, never mind claim that it is doing the *same* thing when the part is included in different networks. A network generally processes as a whole. Still, I must acknowledge the possibility that our views are more similar than I am allowing here.

One reason for the uncertainty of interpretation here is that the localization-holism debate has generally been presented in terms of a choice between whether cognitive functions are typically instantiated by a few and closely grouped neural participants, or by many and widely distributed ones, and although both Uttal and Lloyd make an effort to frame the debate in somewhat different terms, they nevertheless largely follow this convention. By this measure, Lloyd, Uttal, and I are *all* holists—and yet this is pretty clearly not the

right distinguishing factor between localization and holism, for as Mundale (2002) persuasively argues, the belief that cognitive functions typically have many and widely distributed participants is perfectly compatible with localization.¹¹ Thus, to help further clarify these issues, I am explicitly identifying two other questions that can be asked: (1) are brain areas largely dedicated to (exclusive participants in) the cognitive function(s) in which they participate?; and (2), when a brain area participates in more than one cognitive function, is it doing the same thing in each case? The believer in localization answers “yes” to both questions (although question 2 does not really arise), whereas the holist answers “no”.

In contrast both to both localization and holism—limited or otherwise—a redeployment hypothesis splits the difference, answering “no” to question 1, and “yes” to question 2. That is, a redeployment hypothesis claims that parts of the brain *are* specialized, in that they do the same thing each time they are activated. However, the thing that they do—the function they compute or transformation they effect—does not line up with any specific cognitive function. Rather, brain areas must work in concert with other areas to do anything interesting, and are generally deployed in many different functional complexes, which do many different (interesting) things. The main motivating reason for rejecting holism is that it seems that one can offer an evolutionary reason for redeployment as an architectural feature of the brain *only* if brain areas do roughly the *same* thing for each of the functional complexes in which they participate. As new brain functions develop, one might well expect opportunistic reuse of existing functional components,¹² but it seems that this would only be effective insofar as the existing components already did

something that could easily become a useful part of a functional complex supporting the new function. Too little initial compatibility would make the incorporation of existing components into a new functional complex quite puzzling, and too much alteration in the functional structure of the existing component could cause problems with the *other* functions it supports.

For imagine that component c computes function f , and that it does do because it (and its participants) compose a circuit of a particular description. If that's the right explanation—that a component and an area do what they do in virtue of their physical layout, the nature of each part and the details of their relations—then we have an easy sort of story to tell about how component d , which shares some of c 's participants, can compute a different function, g , just so long as each participant, although doing the same thing, is put into such relations with other participants so as to produce a different outcome. This technique of design re-use is employed every day in building new generations of electronic hardware and software by layering new functions over old. But a cardinal rule of re-use is never to break prior functionality (backwards compatibility), and one obeys this rule by not changing prior implementations. For suppose that one can get function g to work by changing the role of participant a , shared by component c . By hypothesis, this means changing its physical/functional properties, thereby altering the functional characteristics of component c , which introduces the possibility that it will no longer compute function f (or that it will fail in circumstances where it did not fail before). There can be little doubt that such fitness regressions sometimes occur (they certainly occur in technology), and may even remain if the benefits are greater than the

loss, but this can hardly constitute the evolutionary norm. Some of the examples I will detail below, such as the use of a phonological loop to support working memory, offer a nice illustration of the redeployment of existing components to support a novel task, in such a way that little or no modification of that original function is required (it might be, for instance, that covert or silent rehearsal was a later adaptation of existing function, developed to support memory).

This brings us back to the main task of this essay—to introduce and support MRH. MRH, obviously, is a redeployment hypothesis, but what makes it *massive* (as opposed, perhaps to mild, moderate, meek, or modest)? There are three things that distinguish a *massive* redeployment hypothesis from its alternatives. First are its expectations regarding the degree to which redeployment is used in the brain: MRH predicts that non-exclusive participation will turn out to be the *norm* when it comes to the functional topography of the brain (a more moderate hypothesis might predict *occasional* instances of redeployment). Another factor is whether redeployment respects the traditional boundaries between cognitive domains (e.g. perception, motor control, language, memory, etc.). MRH predicts significant redeployment both within and between cognitive domains. The final distinguishing feature of MRH is a proposal for specific, 3-tier architecture for the functional topography of the brain. As this proposal would make little sense without first introducing a great deal of empirical evidence, it is to this evidence that we now turn. We will return to the details of MRH in section five.

3. Case studies for massive redeployment

In this section, I will discuss three different instances involving the apparent redeployment of brain areas to support multiple functions. The case studies both provide some evidentiary support for MRH and, perhaps more importantly, illustrate how such redeployment works, and why it might have evolved.

3.1 The organization of M1

The first case study I would like to discuss involves the organization of the primary motor cortex, or M1. One of the brain maps with which nearly everyone is familiar is the motor homunculus (see Figure 2).




Figure 2: “Motor homunculus” from Penfield and Rasmussen (1950).

The somatotopic organization of M1 has long been part of the standard account of the functional topography of that region. In its classic form, Penfield’s homunculus specified

distinct, non-overlapping regions for motor control down to the level of individual fingers and joints. It is a clear product of the strict localization hypothesis.¹³ However, over the past few decades, evidence has been mounting that the areas of M1 controlling the various body parts are in fact distributed and overlapping. Recently, Marc Schieber (2001) has reviewed this evidence, and found six factors constraining the somatotopic organization in M1:

- 1) *Convergent* output from a large M1 territory controls any particular body part, joint, or muscle.
- 2) *Divergent* output of many single M1 neurons reaches multiple spinal motoneuron pools.
- 3) *Horizontal connections* interlink the cortex throughout a major body part region.
- 4) *Widely distributed activity* appears in a major body part region whenever any smaller body part is moved.
- 5) *Partial inactivation* of a major region affects multiple smaller body parts simultaneously.
- 6) *Plasticity* limits the degree to which control of a specific body part can be assigned to a particular piece of cortex. (p. 2125, emphasis in original)

For the purposes of this essay, I will be focusing on findings 1, 2 and 4. Findings 3, 5 and 6, while compatible with MRH and interesting in their own right, nevertheless have implications somewhat orthogonal to the main elements of MRH I am trying to support.

The clear implication of convergence is that there are multiple, not necessarily spatially contiguous areas that share in the motor control of a given muscle or body part. Using intracortical microstimulation (ICMS), a technique that limits the possibility that the stimulus will accidentally spread to larger areas of cortex, Asanuma and Rosen (1972) found multiple small areas controlling the same movement of, or contracting the same muscle in, a monkey's forelimb. Moreover, they found that these areas were intermixed with areas controlling other movements or muscles such that, although a gross somatotopic organizational *trend* could be observed (arm movements controlled by this

general area, leg movements in that, face movements over here), boundaries between large areas were not necessarily clear, and the somatotopic organization did not extend to the fine-grained structure of the cortex. More recent studies have confirmed this finding, and shown further that, as the stimulation of small cortical areas is increased in intensity or duration, responses are evoked in increasing numbers of muscles or joints (an effect that cannot be adequately explained in terms of stimulation spreading or leaking to neighboring cortical areas). This brings us to the issue of *divergence*.

Anatomic evidence indicates that a single neuron from M1 can terminate in different spinal segments, connecting to different motoneuron pools (Shinoda *et al.*, 1981), and functional studies demonstrate that these connections can affect muscle groups across different body parts, as widely separated as finger and shoulder (McKiernan *et al.*, 1998). In an especially striking demonstration of the possible utility of such connections Graziano *et al.* (2002) showed that the stimulation of individual cortical areas could evoke complex coordinated movements of a monkey's forelimb, such as reaching, grasping, or adopting a defensive posture. They found further that these areas were not *somatotopically* organized, but rather showed a *spatial* and *postural* organization, roughly corresponding to the locations at which the movements were directed (the endpoint of the motion in ego-centric space), and the limb posture resulting from the action.

Such evidence for convergence and divergence alone does not necessarily suggest MRH. Convergence, it might be argued, merely shows *redundancy* of function, while the

evidence for divergence is compatible with the strict localization claim that cortical areas are functionally specialized and dedicated, so long as the functions in question are characterized in a complex way: reaching to a given spot, rather than contracting a single muscle.¹⁴ However, the defining claim of MRH is that the same cortical area can play a role in supporting *multiple* functions, however complex their characterizations. Perhaps the clearest evidence for both distribution and redeployment comes from single neuron recordings of monkeys performing individuated finger and wrist movements (Schieber & Hibbard, 1993). Schieber and Hibbard found that the general territories of M1 involved in finger control were virtually coextensive. Moreover, while each neuron was consistently related to at least one movement, there were multiple, spatially distributed neurons involved in each movement, most of which were related to multiple different finger and/or wrist movements. Imaging studies in humans confirm extensive overlap in the areas of activation in M1 corresponding not just to finger movements (Sanes *et al.*, 1995), but also to thumb, index finger, wrist, elbow and shoulder movements (Kleinschmidt *et al.*, 1997).

What is attractive about the evidence from M1 is that it is such an extensively studied area. Insofar as the emerging picture of its functional organization indicates reliance on multiple, distributed and non-exclusive participants in motor control functions, then given the extent of the evidence, the hypothesis needs to be taken quite seriously. On the other side of the coin, given that the evidence is restricted to M1 and motor control, extensive redeployment might not seem all that surprising. The support for MRH coming from the study of M1, while strong, is also somewhat narrow. Thus, the next two case studies

showcase some rather more radical and surprising instances of apparent redeployment. The evidence for these examples is somewhat less strong, but the implications are far broader.

3.2 Sensorimotor coding in working memory

One instance of redeployment on which there has been a fair amount of work is in the apparent use of sensorimotor resources to support working memory. As the evidence has been reviewed in detail by Margaret Wilson (2001), I'll only provide a brief summary. The experiments in question typically involve the presentation of multiple items (words or letters) either visually or auditorily, with the task being to remember these items, in order. The question of interest is what kind of processing supports this ability, and there is a great deal of evidence supporting some version of the Baddeley and Hitch model of working memory, which posits that working memory has both verbal and visuospatial components, among others (Baddeley & Hitch, 1974; 1994, Baddeley, 1986; 1995). Basically, the Baddeley and Hitch model says that one strategy for remembering such lists involves (silently) saying them to one's self (producing a "phonological loop"), which engages brain areas typically used both in speech production and in audition. Another strategy for remembering words is the visual representation of their form or meaning (especially for abstract nouns). Wilson notes that this latter strategy is not particularly effective for maintaining an *ordered* list, and that therefore a strategy involving some version of the phonological loop is more typically employed.

A pattern of findings supports the existence of a phonological loop, a strategy that engages both inner “speaking” and inner “hearing” to support working memory. First, there is poor recall of similar sounding terms; second, there is poor recall of longer words; third, there is poor recall if the subject is made to speak during the maintenance period; and fourth, there is poor recall when the subject is exposed to irrelevant speech during the maintenance period. Moreover, imaging studies have found that such memory tasks cause activation in areas typically involved in speech production (Broca’s area, left premotor cortex, left supplementary motor cortex, and right cerebellum) and in phonological storage (left posterior parietal cortex) (Awh *et al.*, 1996). Imaging data also tends to support the use of sensorimotor strategies in visuospatial working memory, showing activation of right hemisphere, including areas of visual and prefrontal cortex (Smith, 2000).

Although these findings will not be at all surprising to anyone who has ever tried to remember multiple things, only to be foiled by having to say, or listen to, something unrelated, the broad implications are nevertheless significant. As Wilson writes, in this case it appears that:

...sensorimotor processes are run covertly to assist with the representation and manipulation of information, in the temporary absence of task-relevant input or output. Such an arrangement would make sense, given our evolutionary heritage from creatures whose neural resources were devoted largely to perceptual and motor processes. Indeed, given that we have such resources, it would be odd if we did not exploit them whenever possible to assist in off-line cognitive processing. (pp. 44-5)

3.3 The use of motor simulations in language understanding

Finally, the last case I would like to consider is an even more striking example of the redeployment of resources in apparently disparate functions: the action-sentence compatibility effect (Glenberg & Kaschak, 2002), which suggests the involvement of the motor system in language understanding. To demonstrate this interesting interaction between comprehension and motor control, Glenberg and Kaschak asked subjects to indicate whether a given sentence made sense or not by making a response that required a movement either toward or away from their bodies (e.g. reaching for a button). They found that response times were longer in cases where the required movement ran counter to a movement suggested by the sentence itself (e.g. where the response required a movement toward the body, and the sentence, e.g., “Close the drawer” indicated a movement away from the body, or vice-versa), and that this was true even when the “movement” indicated by the sentence was abstract, as in the transfer of information from one party to another (e.g. “You told Ann about the party.”). A general explanation of this effect would be that the comprehension of the sentences involved a motor simulation of the action they describe, thus “priming” the system to move in one way, rather than another. More particularly, Glenberg and Kaschak posit that understanding language involves combining the affordances of the sentence elements, and judging the “doability” of the action corresponding to the meshed set of affordances. A doable action indicates a comprehensible sentence.

These results are intriguing and highly suggestive, yet, as Glenberg and Kaschak readily admit, there is much more work to be done.

In summary, our results demonstrate that the understanding of imperative, double-object and dative constructions is grounded in action. Given that language almost certainly arose to facilitate coordination of action, it is not surprising that there is an observable remnant of that history. The results also raise the intriguing possibility that much, if not all, language comprehension is similarly grounded. Although substantial work needs to be done to secure that possibility, that work may well be rewarded by an account of language and meaning firmly anchored in human experience. (p. 564)

One kind of evidence that is currently missing for this effect is neural imaging data. To help address this lacuna, I hope in the near future to run an MEG experiment featuring the Glenberg-Kaschak task. MEG evidence, especially given its temporal resolution, might help rule out the most obvious alternate explanation of the data, that it is a *post-understanding* simulation of the action that is interfering with the response, rather than a simulation implicated in the understanding itself. Although it is true that it is difficult to use this alternative to explain the effect in the case of *abstract* transfers (for there is little reason to believe that a post-understanding simulation of abstract transfers would implicate movements toward or away from the subject, even if it involved simulating the actions used in the transfer, such as speaking), MEG data might help settle the matter.

There are nevertheless other kinds of evidence available that appear to support the general finding that motor control and language understanding are intertwined with one another. For instance, patient KJ-1360, who has a lesion in left premotor cortex, shows an impairment in verb retrieval, but has otherwise normal linguistic abilities (Damasio and Tranel, 1993). Martin et al. (1995, 1996) confirm this basic finding that areas associated

with motor control are involved in verb retrieval, and also show that naming colors and animals involved visual processing areas, suggesting that language use and comprehension involves the re-use of many other areas of the brain besides motor areas, and, moreover, that this redeployment is content specific, with verbs reusing motor control areas, and certain nouns like animal and color names reusing visual processing resources. That there is a large amount of redeployment of *sensory* processing areas in linguistic and conceptual tasks is another striking case of redeployment worth pursuing in its own right (Barsalou, 1999), but we will focus here on the relation between language use and motor areas.

One particularly interesting part of the brain in this regard is Broca's area (left Brodmann areas 44 and 45). Broca's area has long been associated with language processing, but what has recently begun to emerge is its functional complexity (Hagoort, 2005). For instance, it has been shown that Broca's area is involved in many different action-related tasks, including movement preparation (Thoenissen, Zilles, & Toni, 2002), action sequencing (Nishitani, et al., 2005), action recognition (Decety et al., 1997; Hamzei et al., 2003, Nishitani, et al., 2005), imagery of human motion (Binkofski, et al., 2000), and action imitation (Nishitani, et al., 2005). In other words, language processing involves (much) more than one region of the brain, *and* the regions of the brain associated with language processing are involved in many other tasks, of which we have listed just a few. Note, however, that it does not appear to be the case that brain areas are redeployed haphazardly; rather, the contributions they make are useful in more than one situation. In the case of Broca's area, for instance, it is not surprising that an area of the brain that

plays a role in action sequencing would be useful in language processing and production, since this, too, requires action sequencing. Likewise, that verb retrieval/comprehension would involve motor simulation is unsurprising, so long as we suppose that our ability to understand verbs is closely connected to our experience of acting in the world.

Returning, then, to our central theme, a main distinguishing feature of MRH is the claim that the functional complexes of the brain make heavy use of nonexclusive participants, not just within, but across classically specified domains. The three case studies above offer some evidence for this claim—and, just as importantly, help to illustrate what redeployment *does* for the brain, and why it makes sense as an organizational principle. But it must of course be admitted that this evidence in no way *proves* MRH, and certainly does not establish redeployment as the *norm*. Thus, in the next section we turn to a different kind of evidence that can help do just that.

4. Further evidence for MRH

The evidence for MRH is in no way restricted to the few brain areas or cognitive functions listed above. In fact, a recent empirical review by Cabeza and Nyberg (2000) strongly suggests that rather rampant redeployment to be the norm. Cabeza and Nyberg survey 275 fMRI and PET experiments, arranging them by task category (attention, perception, imagery, language, working memory, episodic memory encoding, episodic memory retrieval, etc.). For each task, they catalog the participants in that task, from a list of 31 different brain areas (28 Brodmann, and three subcortical areas), each divided into four different parts: left lateral, right lateral, left medial and right medial. Although

Cabeza and Nyberg do not do any statistical analysis of this data (their primary interest is in examining/establishing the consistency of findings across different experiments on similar tasks), the results of even a simple analysis are striking.

For simplicity and brevity, I focus here on only four of the ten categories of tasks surveyed: attention, perception, imagery, and language. The data on the other task categories is consistent with what I report here. Cabeza and Nyberg looked at 39 attention-related tasks, 42 perception-related tasks, 18 imagery-related tasks, and 36 language-related tasks, for a total of 135 tasks in these four categories. The attention tasks included things like tone detection and Stroop tasks (naming colored words); perception tasks included such things as object identification and facial recognition; the various imagery tasks include mental rotation and landmark visualization; and the language tasks included reading out loud and silently, lexical decision tasks (discriminating words from non-words), and the like.

As mentioned already above, Cabeza and Nyberg divided each brain area into 4 parts; however, their coding scheme forces a decision between lateral and medial activation, such that it is not possible to show a left medial and a left lateral activation in a given area for a given task. Instead, the possible activations for each brain area are left lateral (LL), right lateral (RL), bilateral lateral (BL); left medial (LM), right medial (RM), bilateral medial (BM). Thus, for instance, they list the following activations for a task involving hearing words vs. a resting condition (Muller et al., 1997): an LL activation in Brodmann area 47, and BL activations in areas 21 and 22.

For the purposes of counting participants in a task, I treated bilateral activations of an area as two participants, one left and one right (medial or lateral). Thus, the language task above would have five participants, three LL participants (areas 47, 21 and 22) and two RL participants (areas 21 and 22). For the purposes of counting redeployments (areas activated by more than one task), I matched LL activations in an area to other LL activations of that area, as well as to BL activations, and I matched RL activations in an area to other RL activations of that area, as well as to BL activations. I followed the same procedure for medial activations. I did not match bilateral activations to each other.

Task Category	Average number of participants per task
Attention	5.26, SD 4.23
Perception	4.88, SD 3.55
Imagery	6.39, SD 3.29
Language	7.81, SD 6.56
Total	5.97, SD 4.80

Table 2: Average number of participants per task, by task category

The data show that, on average, each of the 135 tasks has 5.97 participants (SD 4.80), with somewhat more than that for the language tasks and slightly less for attention and perception (see Table 2). More importantly, they show that each area was typically a participant in more than one task, although interestingly, there is a significant difference

between medial and lateral activations in this regard. Thus, each LM area that was a participant in at least one task was, on average, a participant in 3.87 different tasks (SD 3.34); likewise, each RM area that was a participant in at least one task was, on average, a participant in 3.29 tasks (SD 1.77). In contrast (and providing incredibly striking evidence against strict localization), each LL area that was a participant in at least one task was a participant in an average of **14.29** different tasks (SD 9.20), and each RL area was a participant in **10.41** (SD 7.96). Put differently, an average LL area participated in nearly one in nine (10.6%) of the tasks studied, and the average RL area participated in one in thirteen (7.7%).¹⁵

The participation of areas in multiple tasks was not restricted to only closely related tasks. In fact, of the 28 LL areas that participated in at least one task in one of the four task categories, 26 (93%) were *also* participants in at least one other task in a different task category. Moreover, 23 of those areas (82%) participated in tasks in at least three categories, and 15 (54%) participated in tasks in all four categories. The numbers are similar for RL activations, and while the numbers for medial activations were somewhat lower, the data overall undermine strict localization, and strongly suggest widespread redeployment throughout the cortex (see Tables 3 and 4).

Activation Type	Number of areas with activations in <i>at least</i> :			
	1 task category	2 task categories	3 task categories	4 task categories
Right Lateral	29	26	22	11
Left Lateral	28	26	23	15

Right Medial	14	10	6	0
Left Medial	15	9	6	2

Table 3: The number of brain regions activated (out of 31), with activations in at least the number of task categories listed, out of the four categories surveyed.

Activation Type	Number of areas with activations in <i>exactly</i> :			
	1 task category	2 task categories	3 task categories	4 task categories
Right Lateral	3	4	11	11
Left Lateral	2	3	8	15
Right Medial	4	4	6	0
Left Medial	6	3	4	2

Table 4: The number of brain regions activated (out of 31), with activations in exactly the number of task categories listed, out of the four categories surveyed.

It appears to be the norm not just that a given task has many participants, but that each participant contributes to many distinct and very different cognitive tasks. To underline this last point, and to give some more detailed sense of the *distribution* of activations across task categories, Tables 5-8 (given in the appendix) show the number of activations, arranged by area and task category and normalized with respect to the number of tasks in each category. Although the data is consistent with brain areas having

some categorical *preferences* or *tendencies*, there is nevertheless in nearly all cases a significant distribution of activation across multiple task categories. Thus, for instance, an average of 46% of LL activations and 40% of RL activations occur in task categories *other* than the category with the highest number of activations for a given area (the numbers are 28% and 29% for LM and RM, respectively), and in 15 of the 28 LL areas with at least one activation, the task categories with the highest number of activations accounted for less than 51% of the total activations for that area.

Finally, there *were* nine tasks of the 135 examined that activated a total of ten areas not activated by any other task. Since other tasks not examined might also activate these areas, we cannot conclude on this basis alone that these nine tasks have exclusive participants—and, in fact, eight of these ten areas *are* known to be involved in tasks in categories not surveyed here. Even if the remaining two areas turn out to be exclusive to their two tasks, this still would mean that fewer than 1.5% of the tasks examined had exclusive participants. Only one of these two tasks had a single participant, not activated by any other task; the other had 8 other participants, none of which were exclusive.

Overall, the picture could hardly be clearer: nearly every brain area participates in multiple cognitive functions, and each cognitive function utilizes many participants, *very* few of which are exclusive. The data on brain function is decidedly *not* consistent with the strict localization hypothesis.¹⁶

5. Massive redeployment and the meaning of function

Now it is certainly possible, when faced with such data, to try to maintain strict localization—most or all participants are exclusive to a task, and most or all tasks have exclusive participants—by *redefining* the function of a given area to encompass and account for the various tasks that it seems to support. As was mentioned already, this has been one strategy employed in trying to account for the participation of the fusiform face area in recognition tasks not involving faces,¹⁷ and there may well be future proposals for the “true” function of right prefrontal cortex, and each of the thirty-one other areas surveyed, that account for its participation in the various tasks listed above. Insofar as this move involves a recognition that cognitive functions, as defined by psychology, are unlikely to map one-to-one with particular *brain* functions, defined by the specific, identifiable mechanisms (integration, filtering, and other such transforms) implemented in our neural architecture, then I think it can be counted as a step forward (Bechtel 2002, 2005a, 2005b). On the other side of the coin, unless this move is accompanied by a call to redefine *cognitive* functions in terms of area-specific *brain* functions (that is, to question the validity of the ontology of cognitive psychology, a move for which I currently see no compelling grounds), we would still be left with a many-to-many mapping of cognitive functions to brain areas, and this is a significant step away from strict localization.

Since we have just introduced a new term—brain function—let us step back for a bit to examine it. As stated briefly, above, I am using the term to refer to a specific, identifiable mechanism implemented in the brain for transforming signals and/or information from one form to another. Signal processing functions might include

filtering, smoothing, integrating, enhancing, shifting waveform or periodicity, interpolating, and the like. Information processing functions might be described in mathematical terms (adding, subtracting) or in terms of data processing (store, search, sort). There may also be another, more appropriate vocabulary for brain functions, not yet discovered, developed or widely accepted. In point of fact, we can have little confidence that the field of cognitive science has defined a reliable list of the various functions (elementary operations) implemented in the brain, and my point here is not to recommend any such list, but only to give the reader some notion of the *sort* of thing I mean by the term *brain function*. The real issue, for the purposes of the current article, is to come to some understanding of how these functions are likely to relate to cognitive functions, on the one hand, and to brain areas on the other.

So, let me suggest first of all that, given what has been shown and argued so far, there can be little doubt that it will be a very rare to have a single brain function implement a single cognitive function. Likewise the reverse, it will be rare for a single cognitive function to be the only user of a given brain function.¹⁸ Thus the arrangement here is likely to be many-many. Now, as I noted above, we cannot be confident that we know what the various brain functions are—we have not yet discovered the basic functional vocabulary of the brain. So, one approach to developing this vocabulary is to *define* a brain function as whatever it is that some brain area is discovered to do. In my view, defining into existence a one-to-one correspondence between function and area is too dogmatic, especially when the evidence seems to indicate that brain areas *don't* always do things by themselves, but generally act in concert with other areas.¹⁹ Moreover, we

need to be aware of the fact that it is the purpose of many brain areas, not to actually participate in any of the processing leading to a cognitive result, but rather to *inhibit* or *suppress* other areas; just because an area becomes active during a task, as shown by some imaging method, it cannot be concluded that this area is a direct contributor to the processing required by the task.²⁰ Its role, instead, might be inhibitory—a very necessary role, to be sure, since often an area will be inhibiting another area that itself normally inhibits the very processing necessary in the current task—but of course this sense of “function” is not the same as the that under discussion.²¹ In any case, it further shows that brain function, in that sense, cannot be strictly tied to brain area, but that a brain function is in fact often the product of cooperation between pieces of brain matter (whether they are spatially contiguous or not) playing distinctly different roles.

To try to avoid confusion, and to follow the lead of (Bechtel, 2002), let us call a brain function, in the sense defined above (a specific, identifiable mechanism implemented in the brain for transforming signals and/or information from one form to another) a *component function*, and use the term *role*, or *area role*, for whatever it is determined that a particular, contiguous bit of brain matter actually does. To introduce equivalent vocabulary to refer to anatomy instead of function, let’s call the entirety of participants in a cognitive function a *functional complex*, and the sum of the participants that implement a component function a *component*; an area will still be called an area. Relating all these levels of anatomy and function, we can say that so long as it seems possible (a) that some number of brain areas could cooperate to implement some component function, (b) that the component function in question need not be identical to any cognitive function, and

(c) that the individual area roles need not be identical to the component function (as in the case of components that require more than one brain area for processing, and/or involve inhibitory areas), then a more flexible approach to the functional topography of the brain is called for, one that neither insists that cognitive functions map to single component functions, nor requires that component functions map to single brain areas, but allows for many-to-many relations between all three levels (see figure 3).

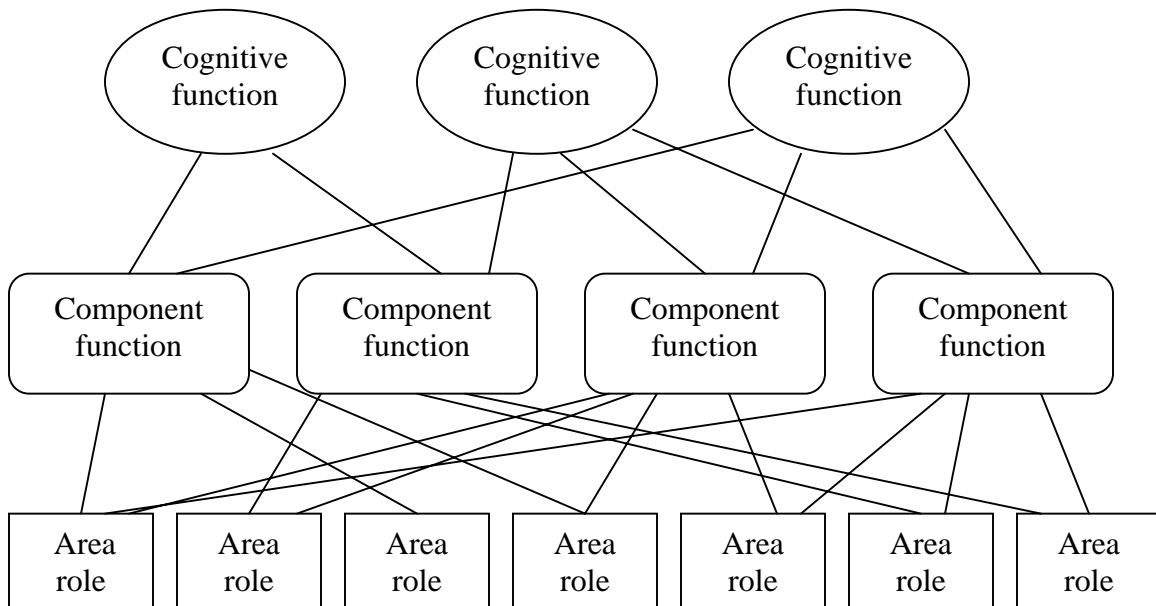


Figure 3: A 3-tier architecture showing a many-to-many relationship between each level.

Saying the same thing with our anatomical vocabulary, we should expect each functional complex to have more than one component, each of which in turn will involve more than one area; likewise, we should expect areas to be members of more than one component,

and components to be members of more than one functional complex, and we should *not* expect that such cross-participation will respect traditional functional-anatomical boundaries (see figure 4).²² This, stated in the most detailed form that will be offered in this paper, is the massive redeployment hypothesis.

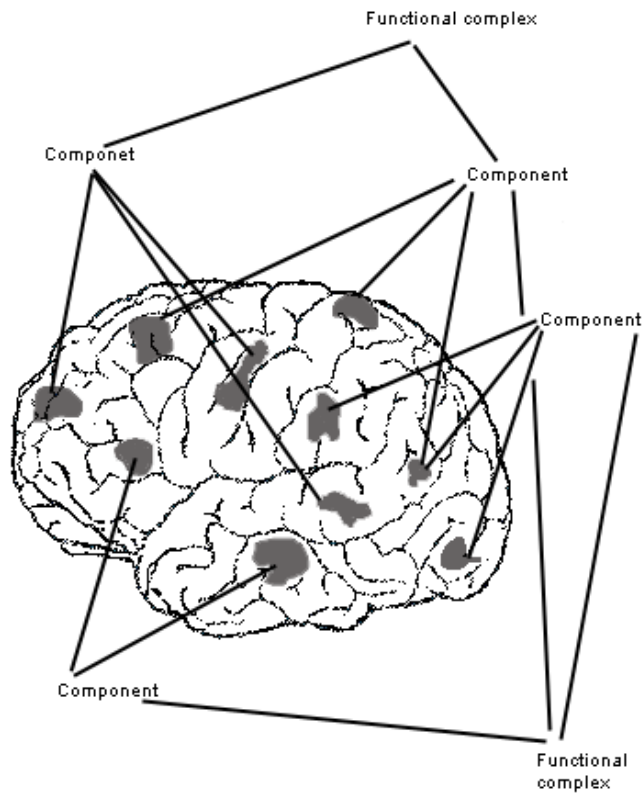


Figure 4: Anatomical illustration of 3-tier architecture, with many-to-many relationships between levels

6. Implications, objections, and replies

In this section I will address some of the implications of MRH, thereby further clarifying the position, by answering some of the questions most likely to occur to the reader.

6.1 *Since you claim that brain areas have unique roles, why isn't this just strict localization by another name?*

First, and most simply, because area roles are not likely to be very interesting by themselves; they only start to have much functional or cognitive significance in combination with other brain areas to form components and functional complexes, the elements of which might be widely scattered, and participants in multiple components and complexes. Second, because localization has typically meant being able to assign domain-specific roles to brain areas, and the prospect for this looks very dim. And finally, what may be among the most important practical implications of MRH, because if MRH is right, then one *cannot* determine what any brain area does by looking at its participation in (activation by) any individual task or task category, as been the normal practice. Rather, one must begin to look at the participation of brain areas across multiple tasks in multiple domains, and try to discern what role it could be playing to account for its functional promiscuity. It must be admitted that this deeply violates the spirit of localization as it has been traditionally understood. Nevertheless, it is true that MRH is compatible with an anemic version of localization that claims simply that individual brain areas do some thing, and the same thing, however low-level, simple, or cognitively uninteresting, whenever they are activated.

6.2 *Are you denying modularity?*

No, although I must add the usual caveat that it depends what one *means* by modularity. MRH is compatible with a limited form of functional localization (see 6.1), insofar as localization does not require all localized participants in a brain function to be exclusive participants in that function. Since I am obviously denying the *strict* localization

hypothesis, I also deny any version of modularity that requires it. What I am proposing is an hypothesis regarding the functional organization of the brain that posits overlapping functional complexes; that is, I expect the entities that implement different cognitive and brain functions to share functional elements (“participants”). Any version of modularity on which modules could be components or functional complexes (or even organized groups of functional complexes) could thereby be compatible with MRH.²³

One aspect of the modularity hypothesis that, in its strong form, sits uneasily with MRH is the claim that modules are *domain specific*. I have reviewed evidence that motor areas are participants in language-related and memory-related tasks, and that language-related areas are participants in motor-related tasks. One result of this, as Prinz (2005) notes, is that focal brain lesions can produce deficits across multiple domains, and genetic language disorders often manifest non-linguistic problems. Moreover, it appears that most brain areas participate in functions across several task categories. Thus, although functional complexes might be domain specific as organized, they are not thereby composed of domain-specific parts. This also casts into doubt the oft-made claim for modularity that many modules can be run simultaneously, helping avoid bottlenecks and computational constraints. The more overlap between modules, the less likely they can be active simultaneously, a fact that accords well with what we know about interference between different cognitive tasks. Although strong interpretations of localization, domain specificity, and simultaneous operation are not strictly speaking *necessary* to the modularity hypothesis, Prinz argues that insofar as we have to weaken our interpretation of its various tenets, we should discard the hypothesis, and focus on functional

decomposition more broadly construed. This may well be the right position, but I'll not pursue the matter further here.

6.3 Can't there be parts of the brain that conform to the strict localization hypothesis?

Of course there can be some brain areas that are best understood in terms of strict localization of function. But the evidence is that re-deployment is the norm. Two specific (and to my mind novel and exciting) predictions of MRH would be (1) that evolutionarily older cognitive functions would tend to use fewer and more tightly grouped participants, whereas newer functions would use more, and more widely scattered ones, and (2) that older cortical *areas* will be participants in more, and more diverse, cognitive functions, and newer areas the opposite. These are empirical questions, and will eventually be settled by the data (I am beginning to perform the required statistical analyses).

6.4 Why didn't you discuss topic Y?

There are a large number of topics, the discussion of which could conceivably have enhanced this essay. Here is a partial list: the various conceptions and critiques of functional localization;²⁴ the modularity hypothesis (massive or otherwise);²⁵ the organization of other well-studied brain regions;²⁶ evolutionary accounts of cognition;²⁷ embodied cognition.²⁸ However, it was my hope to keep this essay focused on the relatively narrow task of introducing MRH, and offering some evidence in support of it. In maintaining this focus I have had to sacrifice some breadth.

7. Conclusion

This essay introduced the massive redeployment hypothesis, an account of the functional organization of the brain that gives pride of place to the fact that brain areas are typically employed to support numerous functions, with little respect for traditional categorical boundaries. Although I think that the three case studies in which there appears to be redeployment of brain areas to support very different functions, together with the empirical review that suggests such redeployment is the norm, strongly support MRH, this is not likely to be, nor is it intended as, the last word on brain organization. Still, an hypothesis can be prove useful even (or perhaps especially) in the course of being disproved and discarded, insofar as it offers a way to help (re-)organize old data and interpret new information, and may suggest novel experimental inquiries. I hope for no more than this from MRH.

Acknowledgements

Thanks are due to the participants of the workshop for early career researchers, sponsored by the McDonnell Project in Philosophy and the Neurosciences, and especially to John Bickle, Valerie Gray Hardcastle, Ben Hardy, David Kaplan, Anthony Landreth, and Bill Seeley. Without the discussions we had, I probably would not have thought to write this essay. Thanks are also due to William Bechtel, Peter Carruthers, Tony Chemero, and Dan Lloyd for helpful advice in revising and greatly improving the presentation.

Notes

¹ Many theorists of embodied cognition might object immediately to this definition, since, if cognition is embodied, then cognitive functions are certainly not definable simply in terms of neural inputs and outputs. (Note that it follows immediately that cognitive functions are not localizable, and the game is over before it begins.) But even on the EC model, there is something that the brain is doing in the course of contributing

to a cognitive function, and even if we believe that isolating this bit of processing leaves us with something quite incomplete from the standpoint of defining a cognitive function, it is still a legitimate question to ask whether and to what degree the brain's contribution is itself localizable. I will be arguing for a particular understanding of functional localization that denies strict localization, but does not embrace holism (as many EC theorists seem wont to do). However, the evidence with respect to localization of function is such that should give comfort to theorists of embodied cognition, since it indicates that higher-order cognitive functions are very often supported by brain areas typically involved in tasks like motor control and sensory processing. See Anderson (in press) for further discussion.

² There is a complication here that is worth noting, but does not require solving: redundancy in some brain areas, and cooperation among others, may mean that actually specifying the necessary participants would be logically complex; for instance, it could be that what is necessary is *either* area A *or* areas (B & C) together, or some such. But for our purposes, we will not need to spell out any such relations; it is enough to agree that there is indeed some set of necessary participants for each brain function, however difficult to specify their form and identity. Note further that the necessary participants do not exhaust the necessary *conditions*—nor even the necessary brain processing—for the success of the function. In the case of understanding a sentence, for instance, there is a great deal of processing necessary to prepare the inputs to the comprehension function (if such there be), the failure of which would prevent the success of the comprehension function. Because we are focusing on the processes supporting the post-input *transformation*, these sorts of issues are not material.

³ Although it is worth noting that there are other sorts of analytical methods available, for instance multi-variate methods including multiple regression, discriminant analysis, and principle component analysis, all of which are designed to reveal the multiple contributors to a given effect or outcome. It is likely that imaging experiments would be designed somewhat differently from the usual “blocked” designs, were the data to be analyzed using multi-variate methods.

⁴ And it should also be stressed that this issue is not *specific* to imaging studies, but is faced whenever one is engaged in comparative studies, i.e. in producing analyses of *differences*.

⁵ Activation results are typically averaged over many trials, thereby filtering out noise.

⁶ Later he uses a more cautious formulation: “allows one to observe brain areas specialized for the processes that differ between the two tasks.” (Buckner, 1996: 155).

⁷ To be precise, the activations for anxiety management and hypothesis generation do not overlap with the activations shown in Figure 1—they overlap in the X and Y coordinates, but, being near $Y = -20 \pm 4$, are considerably below the horizontal section shown in the figure, which is somewhere in the neighborhood of $Y = 12$ in the Talairach and Tournoux (1988) coordinate system. However, they *do* overlap with some of the activations reported in the papers from which Buckner gathered the data for his survey of memory retrieval (see, e.g., Andreasen, et al., 1995). Thus the statement that these three functions have overlapping brain activations is true. There is also another complication here that should be noted, but does not affect the conclusions. The papers cited in the Buckner survey report coordinates based on the Talairach and Tournoux (1988) atlas; the more recent work I am citing uses MNI coordinates (Evans, et al. 1993). I have compared these two systems using the transforms described by Matthew Brett at <http://www.mrc-cbu.cam.ac.uk/Imaging/Common/mnispace.shtml>

⁸ Although, as Lloyd points out, the spatial accuracy of PET is not so great that it is an inappropriate choice.

⁹ Although some of the quantum theories of mind may come close to espousing a radical holism, at a different level of organization.

¹⁰ The text is somewhat equivocal with respect to these possibilities. Here is a fuller quotation: “Third, a brain might be constructed of *sparsely distributed networks*. Here anatomically defined brain regions are *multifunctional*. A region may be recruited to join a subnetwork to compute one function, and later recruited to a different subnet to compute a different function. Thus, subnetworks would overlap in their anatomy. This form of distribution is sparse, however, insofar as particular brain regions are not omnifunctional. That is, each function is computed by a subset of regions, rather than the whole brain. The engaged subnetworks overlap, but the adaptability of each region is limited to a fixed list of functions.” (Lloyd, 2000: 95).

¹¹ My arguments for MRH are largely compatible with those in (Mundale, 2002) for a limited form of localization. There, Mundale argues that the localization hypothesis allows for the possibility that brain functions utilize multiple participants in complex combinations—and, furthermore, she implies that this is

what localization has meant all along. However, while allowing for the possibility of multiple necessary participants in a given brain function, her argument does *not* address the difference between necessary and exclusive participation, and thus she does not discuss the possibility of a form of localization in which the (localized) participants in one function *also* participate in other functions. Someone else theorizing in the same ballpark as MRH is Jesse Prinz, who centrally features evidence that brain regions are used in multiple functions in his critique of modularity (Prinz, 2005). I thus expect he would be sympathetic to the account offered here, but (Prinz, 2005) does not develop such an account.

¹² This is one reason there is—or we should expect there will turn out to be—so much redeployment of sensorimotor resources, as these are some of the oldest. For some discussion of this issue, and its relevance to embodied cognition, see (Anderson, in press).

¹³ This is a potentially interesting case study of the effect of prior assumptions in science, because both (Penfield and Rasmussen, 1950) and (Woolsey *et al.* 1952) admit that the pictures they offered did not reflect the complexity of the data, and the degree to which their evidence indicated a great deal of overlap of the areas controlling different body parts. It took a certain amount of interpretive effort to make their findings consistent with strict localization—for instance, focusing only on the *primary* movement evoked by the electro-stimulation, and discounting some movements as the effect of spreading activation from the primary target site. This is *not* to say that these interpretive moves were unreasonable or illegitimate. See (Schieber, 2001) for some of the interesting details.

¹⁴ It might be thought that I am understating the case in saying that evidence like that presented in (Graziano *et al.*, 2002) is *compatible* with strict localization; rather it seems to positively cry out for such interpretation. After all, stimulation of each of the cortical sites studied did exactly one thing, which was to move the monkey's arm to a spatially well-defined posture. I have two responses to this claim. The first involves a clarification of the scope of MRH: the hypothesis does *not* amount to the claim that *no* part of the brain is organized according to strict localization; the part of the cortex studied by Graziano *et al.* may turn out to be one of those areas. However, MRH predicts (and the evidence presented here strongly suggests) that strict localization will not turn out to be the norm of functional organization. This brings us to the second response, which is to point out that the electrical stimulation of brain areas is a *highly* unnatural event. It does not follow from the fact that the stimulation of an area provokes a given movement that this movement is evoked in the course of an organism's natural behavior by the selective activation of that same brain area. To demonstrate this latter claim would require observing brain activation during behavior elicited under the control of the organism, and the evidence gathered *this* way suggests a more complex picture, involving both distributed activation in the control of individual movements, and redeployment of single areas for the control of different movements.

¹⁵ Cabeza and Nyberg, whose work was the basis of this analysis, *do* notice that there is apparently a great deal of redeployment in the brain, but they decline to offer any hypotheses about the significance of this finding. Interestingly, they observe that researchers tend to interpret activations in terms of the domain within which they are working: "Area 7 activations, for instance, were usually attributed to attentional processes in attention studies, to perceptual processes in perception studies, to working memory processes in working memory studies, and so on." (Cabeza and Nyberg, 2000: 31) This tendency would serve to mask the prevalence of redeployment from those disinclined to look for it.

¹⁶ Note that the data presented do *not* rule out holism as an alternative. I have already indicated why I favor redeployment over holism (see section 2.3), but it must be admitted that, until we have much better awareness of what individual brain areas contribute to cognitive functions (something that imaging data alone will not provide), it will be difficult to definitively rule out holism as an alternative explanation of the functioning of the brain.

¹⁷ It has been proposed that FFA is in fact involved in automatic visual processing tasks correlated with expertise, and this is why it is involved in face recognition (Tarr & Gauthier, 2000).

¹⁸ According to MRH, this will be true for both tokens and types of brain functions.

¹⁹ It goes without saying that different brain areas might implement the same, or a very similar, brain function.

²⁰ This is yet another reason to be cautious about concluding that the active brain areas in an imaging experiment are those responsible, or specialized for, the task being examined.

²¹ Inhibition is an incredibly important mechanism in the brain. Common uses include inhibition of inhibition (thereby releasing some function normally inhibited), lateral inhibition to emphasize contrast

borders, and inhibition that can toggle the function of a component between two possibilities (saccade left vs. saccade right).

²² It appears that this MRH architecture is a somewhat more flexible version of what was proposed by Bechtel (2002) and Petersen and Fiez (1993), with the difference being that they seem to imply that component functions will be strictly localized (whereas I am claiming that only area roles will be). Still, the current proposal is clearly compatible with this work in spirit. For instance, Petersen and Fiez write:

“The areas involved in performing a particular task are distributed in different locations in the brain, but the processing involved in task performance is not diffusely distributed among them. Each area makes a specific contribution to the performance of the task, and the contribution is determined by where the area resides within its richly connected, parallel, distributed hierarchy.” (pp. 513-4)

Mesulam (1990, 1998) also suggests a position similar to MRH: “[M]any cortical nodes are likely to participate in the function of more than one network. Conceivably, top-down connections from transmodal areas could differentially recruit such a cortical node into the service of one network or another.” (1998, p.1040). However, Mesulam’s development and defense of this common basic idea is significantly different from that offered here. Detailing those differences, and the respective strengths and weaknesses of each approach, while a potentially worthwhile task, is well beyond the scope of the present paper.

²³ Another possibility is a version of modularity on which modules were identified with the individual participants in functional complexes. But this would make for some very low-level, uninteresting modules, and so is not likely to be an attractive option for the modularity theorist.

²⁴ On this see, e.g. Mundale (2002), Bechtel and Richardson (1993), and Zola-Morgan (1995).

²⁵ There are many, many works on modularity. Here are a few: Carruthers (2003; 2005; 2006), Fodor (1983; 2000), Pinker (2005), Prinz (2005), Sperber (2002).

²⁶ See, for instance, Bechtel (2001), Dagher et al. (1999).

²⁷ See Barkow, Cosmides & Tooby (1992), Sperber (1996).

²⁸ See Clark (1997; 1998), Wilson (2002), Anderson (2003). For a brief account of the relation between embodied cognition and the redeployment hypothesis, see Anderson (in press).

References

Allison, T., Ginter, H., McCarthy, G., Nobre, A. C., Puce, A., Luby, M. and Spencer, D.D. (1994). Face recognition in human extrastriate cortex. *J. Neurophysiol.* 71, 821-825.

Anderson, M. L. (2003). Embodied cognition: A field guide. *Artificial Intelligence*, 149 (1): 91-103.

Anderson, M.L. (in press). How to study the mind: An introduction to embodied cognition. In: F.Santoianni and C. Sabatano, eds. *Brain Development in Learning Environments: Embodied and Perceptual Advancements*.

Andreasen, N. C., O'Leary, D. S., Arndt, S., Cizadlo, T., Hurtig, R., Rezai, K., Watkins, G. L., Boles Ponto, L. L., & Hichwa, R. D. (1995). Short-term and long-term verbal memory: A positron emission tomography study. *Proceedings of the National Academy of Science*, 92, 5111-5115.

Asanuma, H. and Rosen, I. (1972). Topographical organization of cortical efferent zones projecting to distal forelimb muscles in the monkey. *Experimental Brain Research*, 14: 143-56.

Awh, E., Jonides, J., Smith, E.E., Schumacher, E.H., Koeppel, R.A., and Katz, S. (1996). Dissociation of storage and rehearsal in verbal working memory: Evidence from positron emission tomography. *Psychological Science*, 7: 25-31.

Baddeley, A.D. (1986). *Working Memory*. Oxford: Oxford University Press.

Baddeley, A.D. (1995). Working memory. In: M.S. Gazzaniga, ed. *The cognitive neurosciences*. Cambridge, MA: The MIT Press. 755-64.

Baddeley, A.D. and Hitch, G. (1974). Working memory. In; G.H. Bower, ed. *The Psychology of Learning and Motivation*. Hillsdale, NJ: Earlbaum. 647-67.

Baddeley, A.D. and Hitch, G. (1994). Developments in the concept of working memory. *Neuropsychology*, 8: 485-93.

Barkow, J., Cosmides, L., and Tooby, J. (1992). *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. New York: Oxford University Press.

Barsalou, L. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22: 577-660.

Bechtel, W. (2001a). The compatibility of complex systems and reduction: A case analysis of memory research. *Minds and Machines*, 11: 483-502.

- Bechtel, W. (2001b). Decomposing and localizing vision: An exemplar for cognitive neuroscience. In: W. Bechtel, P. Mandik, J. Mundale, and R.S. Stoffelbeam, eds. *Philosophy and the Neurosciences: A Reader*. Basil Blackwell, Oxford, 225-49.
- Bechtel, W. (2002). Decomposing the mind-brain: A long-term pursuit. *Brain and Mind*, 3: 229-242.
- Bechtel, W. (2005a). The challenge of characterizing operations in the mechanisms underlying behavior. *Journal of the Experimental Analysis of Behavior*.
- Bechtel, W. (2005b). Mental mechanisms: What are the operations? Proceedings of the 27th Annual Meeting of the Cognitive Science Society.
- Bechtel, W. and Abrahamsen, A. (2005). Explanation: A Mechanistic Alternative. *Studies in History and Philosophy of the Biological and Biomedical Sciences*, 36, 421-441.
- Bechtel, W. and Richardson, R.C. (1993). *Discovering Complexity: Decomposition and Localization as Scientific Research Strategies*. Princeton University Press, Princeton, N.J.
- Binkofski, F., Amunts, K., Stephan, K. M., Posse, S., Schormann, T., Freund, H.-J., Zilles, K., and Seitz, R.J. (2000). Broca's region subserves imagery of motion: A combined cytoarchitectonic and fMRI study. *Human Brain Mapping* 11:273–285.
- Buckner, R.L. (1996) Beyond HERA: Contributions of specific prefrontal brain areas to long-term memory retrieval. *Psychonomic Bulletin and Review* 3(2): 149-158.
- Buckner, R. L., Petersen, S. E., Ojemann, J. G., Miezin, F. M., Squire, L. R., & Raichle, M. E. (1995). Functional anatomical studies of explicit and implicit memory retrieval tasks. *Journal of Neuroscience*, 15: 12-29.
- Cabeza, R. and Nyberg, L. (2000). Imaging Cognition II: An empirical review of 275 PET and fMRI studies. *Journal of Cognitive Neuroscience*, 12: 1-47.
- Carruthers, P. (2003). Moderately massive modularity. In A. O'Hear (ed.), *Mind and Persons*. Cambridge.
- Carruthers, P. (2005). The case for massively modular models of mind. In R. Stainton (ed.), *Contemporary Debates in Cognitive Science*. Blackwell.
- Carruthers, P. (2006). Simple heuristics meet massive modularity. In P. Carruthers, S. Laurence and S. Stich (eds.), *The Innate Mind: Culture and Cognition*. Oxford University Press.
- Clark, A. (1997). *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA: MIT Press.

Clark, A. (1998). Embodied, situated, and distributed cognition. In W. Bechtel and G. Graham (eds.), *A Companion to Cognitive Science*. Malden, MA: Blackwell Publishers.

Dagher, A., Owen, A., Boecker, H., and Brooks, D. (1999). Mapping the network for planning. *Brain*, 122: 1973-1987.

Damasio, A. R. and D. Tranel (1993) Nouns and verbs are retrieved with differently distributed neural systems. *Proceedings of the National Academy of Sciences of the USA*, 90: 4957-4960.

Decety, J., Grézes, J., Costes, N., Perani, D., Jeannerod, M., Procyk, E., Grassi, F., & Fazio, F. (1997). Brain activity during observation of actions. Influence of action content and subject's strategy. *Brain*, 120, 1763-1777.

Evans, A. C., Collins, D. L., Mills, S. R., Brown, E. D., Kelly, R. L. and Peters, T. M. (1993). 3D statistical neuroanatomical models from 305 MRI volumes, *Proc. IEEE-Nuclear Science Symposium and Medical Imaging Conference*, 1813-1817.

Fodor, J. (1983). *The Modularity of Mind*. Bradford Books, Cambridge, MA.

Fodor, J. (2000). *The Mind Doesn't Work That Way*. The MIT Press, Cambridge, MA.

Gall, F. J. (1798). Letter from Dr. F. J. Gall, to Joseph Fr[eiherr] von Retzer, upon the Functions of the Brain, in Man and Animals. *Der neue Teutsche Merkur*, 3: 311-332.

Gauthier, I., Skudlarski P., Gore, J.C. and Anderson, A.W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience*, 3: 191-197.

Glenberg, A. and Kaschak, M. (2002). Grounding language in action. *Psychonomic Bulletin and Review* 9: 558-565.

Grasby P.M., Frith C.D., Friston K.J., Bench C., Frackowiak R.S. and Dolan R.J. (1993). Functional mapping of brain areas implicated in auditory—verbal memory function. *Brain*, 116: 1, 1-20.

Graziano, M.S.A., Taylor, C.S.R., Moore, T., and Cooke, D.F. (2002). The cortical control of movement revisited. *Neuron*, 36: 1-20.

Hagoort, P. (2005). On Broca, brain and binding. In: Y. Grodinsky and K. Amuts (eds.) *Broca's Region*. Oxford: Oxford University Press.

Hamzei, F., Rijntjes, M., Dettmers, C., Glauche, V., Weiller, C., & Büchel (2003). The human action recognition system and its relationship to Broca's area: an fMRI study. *Neuroimage*, 19, 637-644.

- Haxby, J.V., Horwitz, B., Ungerleider, L.G., Maisog, J.M., Pietrini, P., and Grady, C.L. (1994). The functional organization of human extrastriate cortex: A PET-RCBF study of selective attention to faces and locations. *J. Neurosci.*, 14: 6336-6353.
- Howard, D., Patterson, K., Wise, R., Brown, W., Friston, K., Weiller, C., & Frackowiak, R. (1992). The cortical localization of the lexicons. *Brain*, 115: 1769-1782.
- Jacobsen, C. F. (1935). Functions of the frontal association areas in primates. *Archives of Neurology and Psychiatry, Chicago*, 33: 558-69.
- Kanwisher, N., McDermott, J. & Chun, M.M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *J. Neurosci.* 17, 4302-4311.
- Kleinschmidt, A., Nitschke, M.F., and Frahm, J. (1997). Somatotopy in the human motor cortex hand area: A high-resolution functional MRI study. *European Journal of Neuroscience*, 9: 2178-86.
- Lashley, K.S. (1950). In search of the engram. *Society of Exp. Biology Symp.*, No. 4
- Lloyd, D. (2000). Terra cognita: From functional neuroimaging to the map of the mind. *Brain and Mind*, 1: 93-116.
- Martin, A., J. V. Haxby, F. M. Lalonde, C. L. Wiggs and L. G. Ungerleider (1995) Discrete cortical regions associated with knowledge of color and knowledge of action. *Science* **270**, 102-105.
- Martin, A., C. L. Wiggs, L. G. Ungerleider and J. V. Haxby (1996) Neural correlates of category-specific knowledge. *Nature*, 379: 649-652.
- McCarthy, G., Puce, A., Gore, J. C. & Allison, T. (1997). Face-specific processing in the human fusiform gyrus. *J. Cogn. Neurosci.*, 9: 605-610.
- McIntosh, A.R. (1999) Mapping cognition to the brain through neural interactions. *Memory*, 7 (5/6) 523-548.
- McKiernan, B.J., Marcario, J.K., Karrier, J.H., and Chenery, P.D. (1998). Corticomotoneuronal postspike effects in shoulder, elbow, wrist, digit, and intrinsic hand muscles during a reach and prehension task. *Journal of Neurophysiology*, 21: 1132-39.
- Mesulam, M.-M. (1990). Large-scale neurocognitive networks and distributed processing for attention, language, and memory. *Annals of Neurology*, 28, 597-613.
- Mesulam, M.-M. (1998). From sensation to cognition. *Brain*, 121, 1013-1052.

- Muller R.A., Rothermel R.D., Behen M.E., Muzik O., Mangner T.J., and Chugani H.T. (1997). Receptive and expressive language activations for sentences: a PET study. *Neuroreport*, 8(17):3767-70.
- Mundale, J. (2002), Concepts of localization: Balkanization in the brain, *Brain and Mind*, Vol. 3, No. 3, pp. 313-330.
- Nishitani, N., Schürmann, M., Amunts K., and Hari, R. (2005). Broca's Region: From Action to Language. *Physiology*, 20: 60-69.
- Pardo, J., Raichle, M., & Fox, P. (1991). Localization of a human system for sustained attention by positron emission tomography. *Nature*, 349: 61-63.
- Penfield, W. and Rasmussen, T. (1950). *The Cerebral Cortex of Man*. New York: MacMillan.
- Petersen, S.E., and Fiez, J.A., 1993: The processing of single words studied with positron emission tomography, *Annual Review of Neuroscience*, 16: 509–530.
- Pinker, S. (2005). So how *does* the mind work? *Mind and Language*, 20 (1): 1-24.
- Prinz, J. (2005). Is the mind really modular? In: R. Stainton, ed. *Contemporary Debates in Cognitive Science*. Blackwell, New York.
- Sanes, J.N., Donoghue, J.P., Thangaraj, V., Edelman, R.R. and Warach, S. (1995). Shared neural substrates controlling hand movements in human motor cortex. *Science*, 268: 1775-77.
- Schieber, M.H. (2001). Constraints on somatotopic organization in the primary motor cortex. *Journal of Neurophysiology*, 86: 2125-43.
- Schieber, M.H. and Hibbard, L.S. (1993). How somatotopic is the motor cortex hand area? *Science*, 261: 489-92.
- Shinoda, Y., Yokota, J., and Futami, T. (1981). Divergent projection of individual corticospinal axons to motoneurons of multiple muscles in the monkey. *Neuroscience Letters*, 96: 7-12.
- Smith, E.E. (2000). Neural bases of working memory. *Current Directions in Psychological Science*, 9: 45-9.
- Sperber, D. (1996). *Explaining Culture: A Naturalistic Approach*. London: Blackwell.
- Sperber, D. (2002). In defense of massive modularity. In E. Dupoux, ed. *Language, Brain and Cognitive Development*. The MIT Press: Cambridge, MA.

- Tarr, M.J. and Gauthier, I. (2000). FFA: a flexible fusiform area for subordinate-level visual processing automatized by expertise. *Nature Neuroscience* 3: 764-769.
- Thoenissen, D., Zilles, K., & Toni, I. (2002). Differential involvement of parietal and precentral regions in movement preparation and motor intention. *Journal of Neuroscience*, 22, 9024-9034.
- Uttal, W. R. (2001) *The New Phrenology: The Limits of Localizing Cognitive Processes in the Brain*. Cambridge: MIT Press.
- Vartanian, O. and Goel, V. (2005). Task constraints modulate activation in right ventral lateral prefrontal cortex. *NeuroImage*, 27: 927-33.
- Wang J, Rao H, Wetmore GS, Furlan PM, Korczykowski M, Dinges DF, Detre JA (2005). Perfusion functional MRI reveals cerebral blood flow pattern under psychological stress. *Proc Natl Acad Sci U S A*. 102(49):17804-9.
- Wilson, M. (2001). The case for sensorimotor coding in working memory. *Psychonomic Bulletin and Review* 8: 44-57.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin and Review* 9 (4): 625-36.
- Woolsey, C.N., Settlage, P.H., Meyer, D.R., Sencer, W., Hamuty, T.P., and Travis, A.M. (1952). Patterns of localization in precentral and "supplementary" motor areas and their relation to the concept of a premotor area. *Res Pub Assoc Res Nerv Ment Dis*, 30: 238-264.
- Wright, C. and Bechtel, W. (in press). Mechanisms and psychological explanation. In P. Thagard (ed.), *Philosophy of Psychology and Cognitive Science* (Volume 4 of the *Handbook of the Philosophy of Science*). New York: Elsevier.
- Zola-Morgan, S. (1995). Localization of brain function: The legacy of Franz Joseph Gall (1758-1828). *Annual Review of Neuroscience*, 18: 359-83.

Appendix 1: Data tables 5-8

Task category	Brain areas																														
	10	9	46	11	47	45	44	6	8	4	32	24	23	31	7	40	39	38	ins	42	22	21	20	mt	37	19	18	17	bg	th	cb
Attention	0.00	1.08	3.23	1.08	2.15	2.15	3.23	7.54	3.23	2.15	1.08	0.00	0.00	0.00	10.77	8.62	3.23	0.00	2.15	1.08	3.23	3.23	0.00	1.08	4.31	8.62	7.54	0.00	2.15	2.15	4.31
Perception	2.00	0.00	3.00	2.00	7.00	3.00	2.00	1.00	2.00	0.00	0.00	0.00	0.00	0.00	6.00	3.00	1.00	0.00	6.00	0.00	0.00	6.00	1.00	5.00	11.00	17.00	7.00	3.00	1.00	0.00	6.00
Imagery	7.00	7.00	4.67	0.00	2.33	4.67	4.67	14.00	4.67	7.00	0.00	0.00	0.00	0.00	14.00	16.33	4.67	0.00	7.00	0.00	4.67	9.33	4.67	4.67	9.33	11.67	7.00	0.00	0.00	0.00	0.00
Language	1.17	4.67	1.17	1.17	11.67	11.67	7.00	5.83	3.50	3.50	0.00	0.00	0.00	0.00	0.00	2.33	9.33	4.67	7.00	11.67	26.83	19.83	5.83	3.50	8.17	9.33	9.33	1.17	4.67	4.67	4.67
Average	2.54	3.19	3.02	1.06	5.79	5.37	4.22	7.09	3.35	3.16	0.27	0.00	0.00	0.00	7.69	7.57	4.56	1.17	5.54	3.19	8.68	9.60	2.88	3.56	8.20	11.65	7.72	1.04	1.96	1.71	3.74
SD	3.08	3.23	1.44	0.82	4.52	4.32	2.15	5.37	1.09	2.94	0.54	0.00	0.00	0.00	6.09	6.49	3.52	2.33	2.31	5.68	12.26	7.26	2.81	1.78	2.84	3.79	1.11	1.42	2.01	2.22	2.60

Table 5: Left lateral activations by task category and brain area, normalized with respect to the number of tasks in each category

Task category	Brain areas																														
	10	9	46	11	47	45	44	6	8	4	32	24	23	31	7	40	39	38	ins	42	22	21	20	mt	37	19	18	17	bg	th	cb
Attention	2.15	6.46	5.38	1.08	3.23	1.08	5.38	7.54	2.15	0.00	1.08	0.00	0.00	0.00	10.77	9.69	3.23	0.00	1.08	2.15	4.31	0.00	1.08	1.08	1.08	5.38	9.69	0.00	1.08	0.00	3.23
Perception	3.00	1.00	2.00	5.00	3.00	4.00	1.00	3.00	0.00	0.00	0.00	0.00	0.00	0.00	6.00	4.00	1.00	0.00	1.00	1.00	0.00	2.00	1.00	5.00	18.00	20.00	10.00	1.00	1.00	1.00	3.00
Imagery	0.00	4.67	2.33	0.00	0.00	7.00	2.33	2.33	4.67	2.33	0.00	0.00	0.00	0.00	4.67	4.67	2.33	0.00	4.67	0.00	2.33	2.33	2.33	7.00	7.00	9.33	7.00	4.67	0.00	0.00	0.00
Language	0.00	1.17	0.00	1.17	4.67	2.33	0.00	4.67	0.00	3.50	0.00	0.00	0.00	1.17	1.17	1.17	1.17	3.50	3.50	5.83	16.33	11.67	0.00	1.17	4.67	4.67	5.83	3.50	4.67	4.67	8.17
Average	1.29	3.32	2.43	1.81	2.72	3.60	2.18	4.38	1.71	1.46	0.27	0.00	0.00	0.29	5.65	4.88	1.93	0.88	2.56	2.25	5.74	4.00	1.10	3.56	7.69	9.85	8.13	2.29	1.69	1.42	3.60
SD	1.53	2.69	2.22	2.19	1.96	2.56	2.34	2.32	2.22	1.75	0.54	0.00	0.00	0.58	3.97	3.55	1.05	1.75	1.82	2.55	7.28	5.21	0.96	2.93	7.29	7.07	2.04	2.16	2.05	2.22	3.38

Table 6: Right lateral activations by task category and brain area, normalized with respect to the number of tasks in each category

Task category	Brain areas																																
	10	9	46	11	47	45	44	6	8	4	32	24	23	31	7	40	39	38	ins	42	22	21	20	mt	37	19	18	17	bg	th	cb		
Attention	1.08	2.15	0.00	0.00	0.00	0.00	0.00	1.08	0.00	0.00	4.31	3.23	1.08	1.08	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.08	1.08	0.00	6.46	1.08	
Perception	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	1.00	1.00	0.00	0.00	1.00
Imagery	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2.33	0.00	0.00	0.00	2.33	0.00	2.33	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2.33
Language	2.33	0.00	0.00	0.00	0.00	0.00	0.00	3.50	1.17	0.00	8.17	4.67	0.00	3.50	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.17	1.17	1.17	2.33	
Average	0.85	0.54	0.00	0.25	0.00	0.00	0.00	1.73	0.29	0.00	3.37	2.81	0.27	1.73	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.25	0.00	0.00	0.52	0.81	0.29	1.91	1.69	
SD	1.11	1.08	0.00	0.50	0.00	0.00	0.00	1.52	0.58	0.00	3.69	1.54	0.54	1.52	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.50	0.00	0.00	0.60	0.54	0.58	3.09	0.75	

Table 7: Left medial activations by task category and brain area, normalized with respect to the number of tasks in each category

Task category	Brain areas																																
	10	9	46	11	47	45	44	6	8	4	32	24	23	31	7	40	39	38	ins	42	22	21	20	mt	37	19	18	17	bg	th	cb		
Attention	0.00	1.08	0.00	0.00	0.00	0.00	0.00	1.08	1.08	0.00	5.38	3.23	1.08	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.08	0.00	3.23	1.08
Perception	0.00	1.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	1.00	0.00	
Imagery	0.00	0.00	0.00	0.00	0.00	0.00	0.00	4.67	0.00	0.00	0.00	0.00	2.33	2.33	2.33	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	4.67	7.00	0.00	2.33	
Language	0.00	0.00	0.00	0.00	0.00	0.00	0.00	3.50	1.17	0.00	2.33	0.00	0.00	1.17	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2.33	1.17	0.00	1.17	1.17	
Average	0.00	0.52	0.00	0.25	0.00	0.00	0.00	2.31	0.56	0.00	1.93	0.81	1.10	1.13	0.58	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.58	1.73	2.00	1.35	1.14	
SD	0.00	0.60	0.00	0.50	0.00	0.00	0.00	2.15	0.65	0.00	2.55	1.62	0.96	0.96	1.17	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.17	2.03	3.37	1.36	0.95	

Table 8: Right medial activations by task category and brain area, normalized with respect to the number of tasks in each category