

# Massive redeployment, exaptation, and the functional integration of cognitive operations

Michael L. Anderson

Published online: 7 September 2007  
© Springer Science+Business Media B.V. 2007

**Abstract** The massive redeployment hypothesis (MRH) is a theory about the functional topography of the human brain, offering a middle course between strict localization on the one hand, and holism on the other. Central to MRH is the claim that cognitive evolution proceeded in a way analogous to component reuse in software engineering, whereby existing components—originally developed to serve some specific purpose—were used for new purposes and combined to support new capacities, without disrupting their participation in existing programs. If the evolution of cognition was indeed driven by such exaptation, then we should be able to make some specific empirical predictions regarding the resulting functional topography of the brain. This essay discusses three such predictions, and some of the evidence supporting them. Then, using this account as a background, the essay considers the implications of these findings for an account of the functional integration of cognitive operations. For instance, MRH suggests that in order to determine the functional role of a given brain area it is necessary to consider its participation across multiple task categories, and not just focus on one, as has been the typical practice in cognitive neuroscience. This change of methodology will motivate (even perhaps necessitate) the development of a new, domain-neutral vocabulary for characterizing the contribution of individual brain areas to larger functional complexes, and direct particular attention to the question of how these various area roles are integrated and coordinated to result in the observed

---

M. L. Anderson (✉)  
Department of Psychology, Franklin & Marshall College, Lancaster, PA, USA  
e-mail: michael.anderson@fandm.edu

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

cognitive effect. Finally, the details of the mix of cognitive functions a given area supports should tell us something interesting not just about the likely computational role of that area, but about the nature of and relations between the cognitive functions themselves. For instance, growing evidence of the role of “motor” areas like M1, SMA and PMC in language processing, and of “language” areas like Broca’s area in motor control, offers the possibility for significantly reconceptualizing the nature both of language and of motor control.

**Keywords** Brain · Cognition · Computation · Cortex · Evolution · Function

## 1 Introduction

The massive redeployment hypothesis (Anderson 2006, 2007a,c) is both a theory about the functional topography of the human brain, and also an account of how and why it got that way. As a theory of the functional topography of the brain, it is perhaps most easily understood as a middle ground between strict localization, on the one hand, and holism on the other. 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 circuits, or by many and widely distributed ones. Yet as Mundale (2002) persuasively argues, the belief that cognitive functions typically utilize many and widely distributed neural circuits is perfectly compatible with localization. Just so long as one can (more or less clearly) distinguish the neural tissues that support a given cognitive function from those that do not, one might claim to have thereby “localized” the function.

But few holists would deny the possibility of doing *that* for cognitive functions, for to do so would be to claim that the all of the brain is involved in everything it does. Although this is logically possible (albeit not currently scientifically plausible), it does not seem to accurately represent the holistic view. What, then, offers a more illuminating contrast between localization and holism? I would like to suggest that the appropriate distinction can be found in the answers to the following two questions: (1) are the brain areas that support a given function typically dedicated to—that is, are they not just *necessary* to, but also *exclusive* to—the cognitive function in question?; and (2) if a brain area supports more than one cognitive function, is it doing the same thing in each case? The believer in strict localization answers “yes” to both questions, whereas the holist answers “no”.

In contrast to both localization and holism, 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 therefore not generally deployed in support of only a single

function, but are instead *redeployed*<sup>1</sup> in many different functional complexes, which do many different (interesting) things.

The reader will immediately wonder what would count as doing “the same thing” in or for different functional complexes. I would like to defer this question for now, for two reasons. First, because one of the implications of the hypothesis is that we need to develop a new vocabulary for characterizing cognitive functions and their parts, and it is (only) within the framework of this as-yet-undeveloped vocabulary that “same” and “different” could even be adjudicated. Thus, although I will say something about this issue below, in point of fact defining what it would mean for a brain area to do the same thing in different contexts (and even establishing *that they in fact do* the same thing) must be considered a *goal* of this research, and not its starting point. Second, because this particular research-guiding commitment of the redeployment hypothesis is one for which there are reasons (although no direct empirical evidence), and those reasons, which must be given later, will help frame and clarify the provisional answer I will provide.

Putting this open issue to one side for now, if answering “no” to question 1, and “yes” to question 2 defines a *redeployment* hypothesis, what then is a *massive* redeployment hypothesis? I use the term *massive* (as opposed perhaps to mild, moderate, meek, or modest) for two reasons. First, the massive redeployment hypothesis (MRH) holds that redeployment is the *norm* when it comes to the functional topography of the brain (a more moderate hypothesis might predict occasional instances of redeployment). Second, MRH expects significant redeployment both within and between traditional cognitive domains (e.g. perception, motor control, language, memory, etc.), that is, it suggests that most neural circuits are not domain specific.

To understand these two commitments returns us to the second aspect of MRH mentioned above, its account of how and why the brain came to be organized the way it is. MRH proposes that cognitive evolution proceeded in a way analogous to component reuse in software engineering (Heineman and Council 2001), whereby existing components—originally developed to serve some specific purpose—are used for new purposes and combined to support new capacities, without disrupting their participation in existing programs. This would parallel the evolution of other capabilities wherein existing structures, evolved for other purposes, are re-used and built upon in the course of continuing evolutionary development (“exaptation”: Gould and Vrba 1982). There is psychological support for exaptation in cognition (Cosmides 1989; Cruse 2003; Glenberg and Kaschak 2002; Gould 1991; Lakoff and Nuñez 2000; Riegler 2001; Wilson 2001), theoretical reason to expect it (Anderson 2003, 2007b, in press), and neuroanatomic evidence that the brain evolved by preserving, extending, and combining existing network components, rather than by generating complex structures *de novo* (Sporns and Kötter 2004). MRH integrates these various perspectives, bringing exaptive accounts of the evolution of the brain at both the functional and neuroanatomic levels together in the realm of cognitive neuroscience.

<sup>1</sup> At this stage in the explanation, it might seem better to call this a *multiple* deployment hypothesis, rather than a *re*-deployment hypothesis. But the current functional topography of the human brain has an evolutionary history, and that history is characterized by the initial development, and subsequent *re-use*, of specialized neural circuitry (see below).

If the evolution of cognition was indeed driven by a kind of exaptation in which older functionality was retained as new functions developed, then we should be able to make some specific empirical predictions regarding the resulting functional topography of the brain, and it is these predictions that give rise to the *massive* element of MRH. Most generally, we should expect a typical brain region to support numerous cognitive functions in diverse task categories. Evidence to the contrary would tend to suggest, instead, that the brain evolved by developing dedicated circuits for each new functional capacity. More interestingly, there should be a correlation between the phylogenetic age of a brain area and the frequency with which it is redeployed in various cognitive functions; older areas, having been available for exaptation for longer, are *ceteris paribus* more likely to have been integrated into later-developing functions. Finally, there should be a correlation between the phylogenetic age of a cognitive function and the degree of localization of its neural components. That is, more recent functions should generally use more, and more widely scattered brain areas than evolutionarily older functions, since the later a function is developed, the more likely that there will already be useful neural circuits that can be incorporated into the developing functional complex, and there is little reason to suppose that the useful elements will reside in neighboring brain regions. An account of the evolution of the brain that instead expected the continual development of new, largely dedicated neural circuits would predict that the resulting functional complexes would remain tightly grouped, as this would minimize the metabolic cost of wiring the components together and communicating between them.

## 2 Evidence in support of MRH

Although the main purpose of this essay is to discuss the implications of MRH for understanding the functional integration of cognitive operations, how seriously one considers those implications will of course depend on the plausibility of MRH itself. Thus, in this section I will highlight some of the more striking empirical findings in support of the hypothesis, and the three predictions presented above. For more complete accounts of the evidence and arguments for (and against) MRH, the reader is directed to (Anderson 2007a,c).

For easy reference, and to guide our evaluation, here is a list of the commitments and predictions of MRH mentioned so far:

- (1) each brain area is typically redeployed in support of other cognitive functions, and such redeployment will not respect traditional domain boundaries,
- (2) nevertheless, redeployed areas play the same “role” in each of the functional complexes they support,
- (3) more recent cognitive functions will utilize more, and more widely scattered brain areas, and
- (4) evolutionarily older brain areas will be deployed in more cognitive functions.

As mentioned already above, clause 2 will be supported by reasons rather than empirical evidence, so initially we will focus on 1, 3 and 4. What sort of data might illuminate these predictions, and thus the plausibility of MRH? Suppose we were given

a large set of cognitive functions in different domains (perception, language, etc.) for which we had already determined, via functional imaging, the network of brain areas supporting each function. In order to evaluate the empirical predictions suggested by MRH, we would need to extract the following information from that set: For prediction (1) we'd want to know (a) the number of individual functions a typical brain region supports (i.e., how often it is activated), and (b) the distribution of those activations across the cognitive domains; for prediction (3) we'd need to know (c) the number of brain regions activated for a typical cognitive function, along with any significant differences in this number between typical functions in different cognitive domains, and (d) the degree of scatter in the network of brain regions supporting a typical cognitive function, along with any significant differences in scatter between typical functions in different cognitive domains, and for predictions (3) and (4) we'd need to know (e) how this data—and especially data relating to differences in these measures between different brain regions and functional domains—relates to information about the phylogenetic age of the region and cognitive domain.

As it happens, Cabeza and Nyberg (2000) compiled just such a set of cognitive functions, providing detailed information about the brain regions activated by each. Their interest was to evaluate the degree to which different functional imaging labs were producing consistent results, but it is quite possible to adapt their data to our purposes. Cabeza and Nyberg provide data for 275 fMRI and PET experiments, in ten task domains. To evaluate the plausibility of MRH, I focused on the 135 tasks in the four most dissimilar domains: attention, perception, imagery, and language. As the other six domains all involve various types of memory, evidence for redeployment across these domains would presumably have limited persuasive force. There were 39 attention tasks, involving things like tone detection and Stroop tasks (naming colored words); 42 perception tasks, involving such things as object identification and facial recognition; 18 imagery tasks, including mental rotation and landmark visualization; and 36 language tasks, including reading out loud and silently, lexical decision tasks (discriminating words from non-words), and the like.

Cabeza and Nyberg report activations using a list of brain areas including 26 numbered Brodmann areas, plus the insula and MT, and three subcortical areas—basal ganglia, thalamus and cerebellum—for each hemisphere. Each area was divided into a lateral and medial segment, for a total of 124 brain regions. Note that the activations reported by Cabeza and Nyberg do not represent the full network of brain areas activated by a given cognitive task, but only those remaining after the relevant control/comparison tasks have been subtracted out. That is, the areas identified in the studies are understood to be those specifically responsible for the cognitive function under investigation. It is also perhaps worth noting that their coding scheme forces a choice between recoding a lateral and a medial activation; a given task cannot show both a lateral activation and a medial activation in the same brain area. Even if there were two separate regions of activation in the same Brodmann area, or one large region of activation covering both lateral and medial segments, this would be reported as one single activation (medial or lateral, depending on the judgment of the coder).

So, how can we use this list of tasks and activation to get the information we need? To get data on (a) and (c)—the number of tasks activating each brain region, and the number of brain regions activated per task—is simply a matter of counting. As it

**Table 1** Illustrations of category diversity for selected Brodmann areas

Area	Normalized proportion of activations by category				Category diversity
	Attention	Imagery	Language	Perception	
BA46R	0.55	0.24	0.00	0.21	0.80
BA18L	0.26	0.21	0.28	0.24	0.97
BA38L	0.00	0.00	1.00	0.00	0.57

happens, only 86 of the 124 brain regions were activated by at least one of the 135 tasks; on average, these 86 regions supported 9.36 different tasks (SD 8.62). Ignoring the division into medial and lateral regions gives an average of 13.00 tasks per area (SD 8.44), nearly one in ten of the tasks surveyed. Looking at regions by task, we find that each of the 135 tasks activated an average of 5.97 regions (SD 4.80). Perceptual tasks activated 4.88 ( $n = 42$ , SD 3.55), attention 5.26 ( $n = 39$ , SD 4.23), imagery 6.39 ( $n = 18$ , SD 3.29) and language 7.81 ( $n = 36$ , SD 6.56).

To evaluate (b), the distribution of activations across the different cognitive domains, we can borrow a measure known as category diversity, used to calculate such things as the ethnic diversity of neighborhoods. When used with four categories, the values for category diversity range from 0.57 for completely imperfect diversity (all instances in a single category) to 1.00 for perfect diversity (25% of instances in each category). In our sample of activations by task, the 86 brain regions have a mean category diversity of 0.76 (SD 0.11); ignoring the medial/lateral division gives 0.81 (SD 0.09).<sup>2</sup> As shown in Table 1, an average category diversity of 0.81 indicates that in a typical brain area there is a fairly even distribution of activations across cognitive domains (supporting prediction 1).

What about (d), the scatter of brain regions involved in a given cognitive task? To calculate these values, one can construct an adjacency graph of the cortex, where each node represents a given Brodmann area, and a line linking two nodes means that the areas share a border in the brain. Nodes that are directly linked have a “minimum graph distance” of one, and other nodes have a minimum graph distance equal to the smallest number of lines one must travel to get from one node to the other in the graph. For each task, then, one can map the Brodmann areas activated by the task onto such a graph, determine the minimum graph distance between each of the nodes, and use this to calculate the *average* minimum graph distance (AMGD) between each of the nodes in the network activated by the task. A larger AMGD indicates that the nodes are more distant from one another—more scattered in the brain—whereas a smaller AMGD indicates the network is limited to a smaller area.

The average minimum graph distance between the Brodmann areas activated by each of the 135 tasks is 3.89 (SD 2.00). That means that for a typical task, each of the six areas activated by that task are separated from the others by about four lines

<sup>2</sup> Note that for the purpose of calculating category diversity, the activation counts in each category were normalized to  $n = 42$ .

**Table 2** Comparison of network size and AMGD across cognitive domains

Categories being compared	Difference in average number of regions activated per task	Difference in average minimum graph distance of activated regions
Language vs. Perception	2.93, $p = 0.0165^*$	1.11, $p = 0.0121^*$
Language vs. Attention	2.55, $p = 0.0475^*$	1.69, $p = 0.0003^*$
Language vs. Imagery	1.42, $p = 0.3922$	0.85, $p = 0.0998$
Perception vs. Attention	0.38, $p = 0.6618$	0.58, $p = 0.2002$
Perception vs. Imagery	1.51, $p = 0.1285$	0.26, $p = 0.6317$
Attention vs. Imagery	1.13, $p = 0.3214$	0.84, $p = 0.1402$

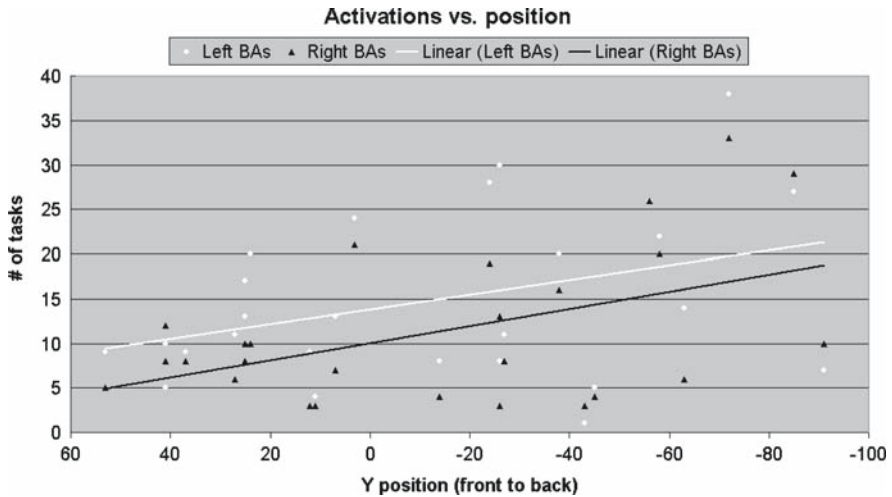
in the graph—that is, by three (unactivated) brain regions. Broken down by cognitive domain, we get attention 3.13 (SD 2.06), perception 3.71 (SD 1.98), imagery 3.97 (SD 1.75), and language 4.82 (SD 1.76).

With these basic data in front of us, we are now in a position to evaluate (e), relating these values, and especially any significant differences between them, to the phylogenetic age of our cognitive domains, and brain regions, respectively. The first comparison is relatively straightforward. Assuming that language is the most recently evolved cognitive domain, and perception and attention are the oldest domains, MRH would predict (3) that language tasks would use more, and more widely scattered brain areas than either. And indeed, this is just what we find: for the mean number of areas activated, language is greater than perception by 2.93 ( $p = 0.0165$ ) and greater than attention by 2.55 ( $p = 0.0475$ ). For average minimum graph distance, language is greater than perception by 1.11 ( $p = 0.0121$ ) and greater than attention by 1.69 ( $p = 0.0003$ ). Differences between other categories are not significant (Table 2).

The last piece of evidence we need—to support prediction (4), a direct relationship between the phylogenetic age of a brain area and its frequency of redeployment—is the most difficult to accurately establish, for the simple reason that there is little consensus on the relative evolutionary age of various cortical structures. Thus, for this measure, we employ an admittedly crude simplifying assumption: all things being equal, structures in the front of the brain will be evolutionarily more recent than structures in the back of the brain. Given this, we can plot the number of activations for a given brain area versus its Y-position in the brain (Fig. 1). The result reveals a significant linear correlation  $R = -0.4121$ ,  $p \leq 0.00244$  ( $t = -3.198$ ,  $DF = 50$ ). Although based on an admittedly questionable indirect measure, it is nevertheless quite striking to have found such a correlation.

Together, these data suggest a picture of the evolution of cognition where redeployment has indeed played a significant role. As predicted by MRH, we see correlations between the phylogenetic age of brain areas and the frequency of their activation by cognitive functions, and between the age of cognitive functions and their degree of localization. We also saw that the typical brain area is a diverse instrument, supporting functions distributed across multiple cognitive domains. The massive redeployment hypothesis thus appears to be both empirically supported, and consistent with the evidence for evolution by exaptation in both psychology and neuroanatomy.





**Fig. 1** Plot of number of activations for each Brodmann area versus the Y-position of the area

Still, we have yet to defend one specific commitment of MRH—number (2), above, that redeployed areas “do the same thing” (in my preferred vocabulary, “play the same role”) in each of the functional complexes they support. The reasoning behind this commitment is relatively simple, but involves several considerations, which I will try to lay out as clearly as I can. Let’s assume first that the functional properties of a neural circuit are determined by its configural properties, such as the number, strength and topology of its connections.<sup>3</sup> Then it follows that a given neural circuit, in a given configuration, does some specific functional thing (let’s speak generally for the moment) when it is activated, the way a gear, or a lever, or a circuit board does some specific thing whenever a certain amount and sort of mechanical or electrical energy is applied. It is only when the configuration of these things changes that they can be said to be doing something different. Of course, we have to be sensitive to the conceptual scope here. That is, assuming that some configurally identical circuit is installed in both the iPhone and the Palm,<sup>4</sup> it would of course be true that using a description inspired by global context, the circuit does something different for each device. But it is only because by some more local description the identical circuits do the *same* thing (and it is possible to specify precisely what this thing is) that the same circuit would be functionally suitable for these different applications.

So far so good. So long as the configuration of a neural circuit remains fixed, we should say—using a suitably local scope—that it is doing the same specific thing whenever activated. But mightn’t the configuration of the network change? Indeed we know that neural circuits *do* change, and precisely in terms of the number, strength and topology of its connections, as the result of development, learning, drugs, injury,

<sup>3</sup> Not an exhaustive list by any stretch (see, e.g. Bickle 2003), but enough to motivate the argument.

<sup>4</sup> This is just an example! I have no idea if these two devices share components (although they almost certainly do at some very low level), and it may well be the case that the designers of both would be scandalized by the thought.



and the like. By my own assumption, we should say that the network's functional properties change as the result of such configurational changes, and in fact we can observe such functional differences at the level of behavior (e.g. learning—or losing—some physical ability like walking or talking).

But the question for MRH is not whether neural circuits remain fixed, or change, but *how quickly* they can do this, and *under what circumstances*. MRH holds that a typical neural circuit is part of many functional complexes; the question is whether the circuit does “the same thing”, or something different, for each. If it is right to say that functional difference entails configurational difference, then for the circuits to do different things for different functional complexes would require configurational changes to occur on time scales of seconds or less, as quickly as an agent can switch back and forth between cognitive functions. Thus, the first point to be made in support of the commitment that these circuits do the *same* thing in each case is that there is currently no evidence for configurational changes at these time scales.

Yet, as Donald Rumsfeld infamously noted, absence of evidence is not evidence of absence, and we should of course admit that evidence for rapid configurational change in neural circuitry may be one issue of *Nature* away. Moreover, it could turn out that among the relevant function-determining configurational properties of neural assemblies are elements that *are* known to change quickly, such as the passive resistance of the individual soma.<sup>5</sup> So if the argument here is to have a stronger form than “until we know that circuits change we should assume that they don't”, we need to introduce some further considerations.

Such considerations follow directly from reflection on the evidence for MRH presented above. If it *were* the case that neural circuits, by whatever mechanism, could be easily and quickly made to do different things for different functional complexes, then metabolic considerations (e.g. the cost of establishing, maintaining, and using long-distance connections between brain areas) would tend to favor recruitment of local neural circuits to support developing functionality. Under these conditions, we should expect to observe a consistent degree of localization over evolutionary time. But the evidence presented above suggests that this is not the case. Although alternate hypotheses may be developed to explain the observed decreasing localization of functional complexes over evolutionary time, MRH offers a very natural fit with the data: because of the time scales over which configurational changes are possible, neural circuits can't be reconfigured quickly enough to support multiple arbitrary functions. Thus, successful recruitment of these circuits into emerging functional complexes requires that the circuit *already* have a functional role answering some need in the emerging complex. There is no reason to suppose that all such circuits will be near one another, and indeed, as the options grow, so too will the chance that the required circuits are quite scattered. Seen in this light, the belief that neural circuits do the same thing in each of their functional complexes is not just an arbitrary commitment for MRH, but an important part of the explanation for the empirical observations made so far.

---

<sup>5</sup> The issue of just what matters for understanding the functional properties of neurons and neural circuits is still very much open; see, e.g. (Koch and Segev 2000) for a discussion.

Note that while I have framed the latter argument in terms of the relation between the configural and functional properties of neural circuits, and the time scales over which the former can change, in fact the argument does not depend on these assumptions. It could turn out that neural circuits have rapidly changing configural properties, or that they do not even require configural changes to implement different functions, and the following response would remain available: perhaps so, yet it appears that these circuits, for whatever reason, do *not* change their roles, for if they did metabolic considerations would favor local recruitment, etc. (see paragraph above). The burden to provide an alternate explanation of the reported observations would remain on any theorist who prefers to suppose that neural circuits could do different things for different functional complexes. It may be that such an alternative is rapidly forthcoming, and will displace MRH as the best current explanation of this data, but *c'est la science*.

### 3 MRH and the functional integration of cognitive operations

Even given this truncated account of MRH and the supporting evidence, it should be clear that it has a number of implications for understanding the functional integration of cognitive operations in the (human) brain. In this section, I'll discuss four. The first three I take to be perfectly straightforward implications of MRH, relevant in various ways to understanding functional integration and organization in the brain. The fourth is more speculative, and concerns the ways in which coming to recognize that neural resources are shared between cognitive domains, and reused over evolutionary time, can help shape our understanding of the basic nature of those cognitive domains: are there ways in which language, supported partly by neural circuits also supporting perception and motor control, is somehow *like* perception and motor control, and vice-versa? We will consider such questions shortly, but turn initially to easier things.

First, in order to determine what any neural circuit does—what it contributes to the functional complexes of which it is a part (I call this its “area role”)—it will be necessary to consider the brain area's support of cognitive functions across multiple domains, rather than focusing on a single domain or class of tasks as has been the typical practice in cognitive neuroscience. To consider a particular example, 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- and imagery-related tasks, including movement preparation (Thoenissen et al. 2002), action sequencing (Nishitani et al. 2005), action recognition (Procyk 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). This is a rather rich array of functions, and should suggest the poverty of approaching the task of functionally modeling a brain region from within only one cognitive domain.

To attribute an area role to Broca's area, then, it will be necessary to consider not just the cognitive models of one or a few language-related functions, and the various options for mapping the functional boxes in *that* model to specific brain areas, but also the cognitive models of all the *other* functions recruiting the area, such that

the sub-functional elements of each model attribute the same role to the brain areas where they overlap. Thus, finding the role of a given brain area will be something like finding the right letter to go into a box on a (multidimensional) crossword puzzle, determined not just by the answer to a single clue, but by all the clues whose answers cross that box. This makes the task both harder, because it is multiply constrained, but also easier, because it offers the possibility of leveraging information from several sources to make the attribution. For instance, the overlaps should suggest more fine-grained predictions about such matters as priming and cognitive interference, and this opens the possibility of designing experiments leveraging these overlaps, e.g. in further imaging, cross-domain priming, and interference studies. The hope is that such an approach will allow a more focused and fruitful search for the roles of individual brain areas.

The second implication follows directly from the first. Insofar as our approach to discovering the specific functional role of a given brain area involves modeling its activity across different cognitive domains, then it makes little sense to try to characterize the contribution of the area using domain-specific terms. This is an important lesson to absorb, for as Cabeza and Nyberg note, researchers tend to offer functional interpretations of observed brain activations using terms drawn from the domain within which they work: “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). Instead of continuing on this path, we will need to develop a domain-independent functional vocabulary to specify the contribution of each area. Ideally, the vocabulary will also be such as to facilitate an understanding of how to build complex functionality out of simple area roles. Currently, the most promising candidates for this vocabulary use terms drawn from the theory of computation and information processing, but we should be open to different developments, e.g. genetics (Bickle 2003; Marcus 2004) and dynamical systems theory (Kelso et al. 1998; Skarda and Freeman 1987; Thompson and Varela 2001).

Here it is worth returning to a point made much earlier: the ability to establish empirically what a given area *does*, and whether and when it is *doing the same thing*, awaits not just the adoption of a methodology that is significantly different from the current standard, but the development of a shared, domain-neutral vocabulary within which to formulate and evaluate rival functional attributions. Thus, such results must be considered a goal of ongoing research, and MRH’s commitment to the notion that brain areas have relatively fixed and definable roles (and to the idea that it is useful to talk this way) should be considered a research-guiding principle, useful only to the extent that it continues to prove fruitful and illuminating.

The third implication, which has no doubt already occurred to many readers, is that MRH sits somewhat uneasily with the oft-cited claim from classical cognitive science that cognitive modules are domain specific (Fodor 1983). Central to the attraction of the modularity hypothesis, and one of the motivations for insisting on the domain-specificity of cognitive modules, was an underlying architecture able to support massively parallel processing, explaining our (apparent) ability to avoid computational bottlenecks. Now, if we are content to identify a module with whatever organized network of brain areas in fact supports a given cognitive task (e.g. verb

retrieval), then it seems safe to say that this functional complex will indeed be a domain (and perhaps task) specific entity. But, insofar as MRH is on the right track, few if any of the elements that compose this entity will *themselves* be domain specific. Indeed, to stick with the case of verb retrieval, we know that its supporting network includes areas of left premotor cortex canonically involved in motor control (Damasio and Tranel 1993). The trouble for modularity is that the more that different functional complexes utilize the same brain areas, the less easy it will be, from an architectural standpoint, to avoid competition and interference between different cognitive tasks. In contrast to the modularity-inspired architecture of multiple, largely isolated, domain-specific processing streams, MRH suggests an architecture of organized high-level complexes that are densely interconnected by low-level, shared components.

Now, it seems to me that the various data on our cognitive limitations, the wealth of information we have on cross-domain interference (and the number of experimental approaches that depend upon it), and the clinical evidence that focal brain lesions can often result in cognitive deficits in multiple domains, all suggest that MRH is closer to architectural reality than is the modularity hypothesis. However, marshalling the evidence necessary to gore that particular bull is beyond the scope of the current article (for a recent attempt along these lines, see (Prinz 2005); also (Uttal 2001)). It is worth emphasizing, however, that what is a bane for this particular architectural theory is a boon for research in cognitive science: as noted already above, the more overlap of functional elements there might be, the better chance we have of exploiting that overlap in interference studies that will help us specify the functional contributions of the overlapping elements.

This brings us to the promised reflection on a fourth, and more speculative consequence of MRH for approaching the issue of functional integration. What I have in mind is the following: given the hypothesis that later-developing cognitive domains inherit and redeploy neural resources originally developed for some different cognitive purpose, might it be the case that the functions in these later domains (or the domains as a whole) thereby inherit something of the *character* of the earlier? And if we can identify this character, might it not tell us something interesting and useful about the nature of—the proper way to describe—*both* domains?

This implication situates MRH in a mutually reinforcing relationship with several long-standing research projects in cognitive science that have emphasized the deep interconnections between various cognitive domains. This includes work on metaphorical inheritance and conceptual blending (Lakoff and Johnson 1980; Lakoff and Johnson 1999; Fauconnier and Turner 2002), on the perceptual roots of abstract representations (Barsalou 1999; Martin et al. 1996), and on the relations between language and motor control (Damasio and Tranel 1993; Rizzolati and Arbib 1998). The fact that MRH makes the independently motivated prediction that there is likely to be inheritance (or similarity) of characteristics between disparate domains, a prediction borne out by some celebrated observations in the cognitive sciences, strikes me as a point in its favor. And by way of reinforcing these more established projects, MRH may offer a framework within which to develop some physiologically-grounded causal explanations of the relevant observations.

Let me illustrate this first by analogy.<sup>6</sup> Suppose we were to discover that both Michael Jordan and Bill Russell had the same high-school basketball coach. We might then look to see what these two very different players have in common. It could be some very low-level element, like the mechanics of how they move the ball between their hands on the dribble, which might be relatively uninteresting with respect to understanding the overall play of these men, but would have the advantage of telling us something very specific about the (possible) contribution of the coach. Alternately, we might notice something at a bit higher level—the way they manage the game, the way they see the court, or the way they exploit their own repertoire of skills—the sort of thing people mean when they say “his playing reminds me of . . .”. This commonality in the higher-level *character* of their game might give us less information about the specific contribution of the coach, but will go a long way in helping us understand the two players. In addition to suggesting ways to organize the search for the specific functional roles of individual neural circuits, as outlined above, I think that MRH also holds out the possibility of gaining this latter sort of insight into the nature of cognitive domains.

Consider, for instance, the particular phonemic character of human speech. A phoneme is defined by a certain posture of the vocal apparatus, and is produced by moving the apparatus toward that posture while making some noise (Fowler et al. 1980). Why should speech production be this way? In an article outlining their discoveries regarding the postural organization of the motor-control system, Graziano et al. (2002b) write:

One possibility is that the mechanisms for speech were built on a preexisting mechanism for motor control, one that emphasized the specification of complex, behaviorally useful postures. When we stimulated the ventral part of the precentral gyrus, in the mouth and face representation, we often caused the lips and tongue to move toward specific postures (Graziano et al. 2002a). For example, at one site, stimulation caused the mouth to open about 2 cm and the tongue to move to a particular location in the mouth. Regardless of the starting posture of the tongue or jaw, stimulation evoked a movement toward this final configuration. This type of posture may be useful to a monkey for eating, but could also be an evolutionary precursor to the phoneme. (Graziano et al. 2002b: 355)

Speech production, then, inherited the postural organization of pre-existing motor-control circuits, and thereby also developed a specific phonemic character. Had the motor control system been oriented instead around (for instance) simple, repeatable contractions of individual muscles, the result of the inheritance might have been a communication code built of more purely temporal elements, something closer to Morse code. Interestingly, this inheritance by the language system of the postural organization of motor control circuits also has the potential to help explain why even American Sign Language (ASL) seems to have a phonemic structure, despite

---

<sup>6</sup> I'm playing to *hoi polloi* with a sports analogy. However, Tony Chemero suggests the following more cultured alternative: a cognitive domain may inherit a character via its redeployed neural circuits the way a wine inherits a *gout de terroir* from its region.

differences in modality that might otherwise have predicted a rather different organization (Sandler and Lillo-Martin 2006).

Or consider another example of the apparent redeployment of motor-control resources in language, this time to support language *understanding*, rather than language production: the action-sentence compatibility effect (Glenberg and Kaschak 2002). To demonstrate this interesting interaction between comprehension and motor control, Glenberg and Kaschak asked subjects to indicate whether a given sentence (e.g., “Close the drawer”) made sense or not by reaching either toward or away from their bodies to press the appropriate button. They found that response times were longer in cases where the required movement ran counter to the movement suggested by the sentence itself, 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., “John told you about the party.”).

This apparent interference between language and motor control may have some important implications for our understanding of each. One intriguing possibility raised by Glenberg and Kaschak is that both language and motor control are fundamentally related to affordance processing (Gibson 1979/1987). More particularly, they 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. One general implication of this claim, explored in detail in work by, e.g., Lakoff and Johnson (1980, 1999), Barsalou (1999), Damasio and Tranel (1993) and others, is that the language system is not encapsulated and autonomous but is instead deeply integrated with older systems, with a function inherited from and grounded in our experience of and capacities for perceiving and navigating the world. Part of what’s behind the work of these authors is the idea that understanding nouns, or verbs, or whole sentences, is in some way *like* seeing, or acting, or planning to act—not phenomenologically (although this would be an interesting area of research), but functionally, structurally, logically. Something of the *character* of the perceptual and motor systems is evident in the language understanding system—putting linguistic elements together in a meaningful sentence is like putting motor primitives together in an executable motor plan. Indeed, the things we say, and the way we reason, often bears witness to this inheritance. Consider the notion of an upright person, the head of an organization, facing the future, being on top of things, seeing red. According to these various authors, linguistic meaning, and with it our sense of how concepts connect and flow, has its origin in, and retains the structure of, our perceptual and physical coping with the world. MRH offers the beginnings of one possible causal, *physiological*, account of these long-noted and oft-discussed findings.

This is interesting enough, and ought to help us better describe the nature of language and language processing. However, there is also a reverse implication that is worth considering: what does the fact that language is built in part on motor-control circuits tell us about *motor control*? Since affordances, the perceived availability of objects for certain kinds of interaction, aren’t just motor programs, but features of the environment with specific significance for the organism, this opens the possibility that the motor control system is also, already, a primitive meaning processor (Gorniak and Roy 2006). This would offer one explanation of how it is even possible to leverage

motor control to support and constrain higher-order processes like language understanding. After all, on a more mechanistic understanding of the nature of motor control, it would be nearly impossible to say why a motor-control system would have *any* of the right basic elements for building a language understanding system.

MRH not only offers a possible physiological explanation of these findings in terms of the neural resources shared between systems, but also suggests that the inheritance of cognitive-functional characteristics will be a common phenomenon. Insofar as this proves true, we should be prepared to exploit this phenomenon as an epistemic lever to help us generate better, more accurate, and more fruitful descriptions of high-level cognitive domains and their relations.

## 4 Conclusion

In this essay, I've outlined the basic idea behind the massive redeployment hypothesis, and introduced some of the evidence that supports it. More importantly, I've identified four specific implications of MRH for the project of understanding the functional integration of cognitive operations:

- (1) MRH suggests that we need to develop a domain-neutral vocabulary for characterizing the functional roles of local neural circuits.
- (2) To actually determine what those area roles *are*, we need to look at the participation of each area in a wide variety of cognitive functions, in a number of cognitive domains.
- (3) We should not expect cognitive modules to have domain-specific parts, and we should be prepared to exploit these overlaps in designing studies to help with specific functional attributions.
- (4) Cognitive domains may inherit a certain *character* from their evolutionary predecessors, via shared, redeployed neural circuits. Recognizing these inherited characters can help us better describe the nature (and relations) of the domains themselves.

With a bit of luck, all these elements can work together, in an iterative, mutually-reinforcing process leading to a better, more complete understanding of cognition, whether high-level or low, ancient or new.

**Acknowledgements** Thanks are due to Jakob Hohwy for his patience and persistence in compiling this special collection, and to Tony Chemero and two anonymous reviewers for their helpful comments and suggestions.

## References

- Anderson, M. L. (2003). Embodied cognition: A field guide. *Artificial Intelligence*, 149(1), 90–130.
- Anderson, M. L. (2006). Evidence for massive redeployment of brain areas in cognitive function. *Proceedings of the Cognitive Science Society*, 28, 24–29.
- Anderson, M. L. (2007a). Evolution of cognitive function via redeployment of brain areas. *The Neuroscientist*, 13(1), 13–21.



- Anderson, M. L. (2007b). How to study the mind: An introduction to embodied cognition. In F. Santoianni & C. Sabatana (Eds.), *Brain development in learning environments: Embodied and perceptual advancements*. Cambridge: Cambridge Scholars Press.
- Anderson, M. L. (2007c). The massive redeployment hypothesis and the functional topography of the brain. *Philosophical Psychology*, 21(2), 143–174.
- Anderson, M. L. (in press). Action-grounded cognition: Evolution, embodiment and the nature of the mind. Keynote address, Cognitio 2006. In B. Hardy-Vallee & N. Payette (Eds.), *Beyond the brain: embodied, situated & distributed cognition*. Cambridge: Cambridge Scholars Press.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577–660.
- Bickle, J. (2003). *Philosophy and neuroscience, a ruthlessly reductive account*. Dordrecht: Kluwer Academic Publishers.
- Binkofski, F., Amunts, K., Stephan, K. M., Posse, S., Schormann, T., Freund, H.-J., Zilles, K., & Seitz, R. J. (2000). Broca's region subserves imagery of motion: A combined cytoarchitectonic and fMRI study. *Human Brain Mapping*, 11, 273–285.
- Cabeza, R., & Nyberg, L. (2000). Imaging cognition II: An empirical review of 275 PET and fMRI studies. *Journal of Cognitive Neuroscience*, 12, 1–47.
- Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition*, 31, 187–276.
- Cruse, H. (2003). The evolution of cognition—a hypothesis. *Cognitive Science*, 27, 135–155.
- Damasio, A. R., & Tranel, D. (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.
- Fauconnier, G., & Turner, M. (2002). *The way we think: Conceptual blending and the mind's hidden complexities*. New York: Basic Books.
- Fodor, J. (1983). *The modularity of mind*. Cambridge, MA: Bradford Books.
- Fowler, C. A., Rubin, P., Remez, R. E., & Turvey, M. T. (1980). Implications for speech production of a general theory of action. In B. Butterworth (Ed.), *Language production, Volume 1: Speech and talk* (pp. 373–420). London: Academic Press.
- Gibson, J. J. (1979/1987). *The ecological approach to visual perception*. New York: Lawrence Erlbaum Associates.
- Gil-da-Costa, R., Braun, A., Lopes, M., Hauser, M. D., Carson, R. E., Herscovitch, P., & Martin, A. (2004). Conceptual representations in a nonhuman primate: Species-specific vocalizations activate visual and affective processing systems in the macaque *Proceedings of the National Academy of Sciences, USA*, 101, 17516–17521.
- Glenberg, A., & Kaschak, M. (2002). Grounding language in action. *Psychonomic Bulletin and Review*, 9, 558–565.
- Gorniak, P., & Roy, D. (2006). Perceived affordances as a substrate for linguistic concepts. *Proceedings of the Cognitive Science Society*, 28, 279–284.
- Gould, S. J. (1991). Exaptation: A crucial tool for an evolutionary psychology. *Journal of Social Issues*, 3, 43–65.
- Gould, S. J., & Vrba, E. (1982). Exaptation: A missing term in the science of form. *Paleobiology*, 8, 4–15.
- Graziano, M. S. A., Taylor, C. S. R., & Moore, T. (2002a). Complex movements evoked by microstimulation of precentral cortex. *Neuron*, 34, 841–851.
- Graziano, M. S. A., Taylor, C. S. R., Moore, T., & Cooke, D. F. (2002b). The cortical control of movement revisited. *Neuron*, 36, 349–362.
- Hagoort, P. (2005). On Broca, brain and binding. In Y. Grodinsky & K. Amunts (Eds.), *Broca's region*. Oxford: Oxford University Press.
- Hamzei, F., Rijntjes, M., Dettmers, C., Glauche, V., Weiller, C., & Büchel, C. (2003). The human action recognition system and its relationship to Broca's area: An fMRI study. *Neuroimage*, 19, 637–644.
- Heineman, G. T., & Councill, W. T. (2001). *Component-based software engineering: Putting the pieces together*. New York: Addison-Wesley.
- Kelso, J. A. S., Fuchs, R., Lancaster, T., Holroyd, D., Cheyne, H., Weinberg, H. (1998). Dynamic cortical activity in the human brain reveals motor equivalence. *Nature*, 23, 814–818.
- Koch, C., & Segev, I. (2000). The role of single neurons in information processing. *Nature Neuroscience*, 3, 1171–1177.

- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.
- Lakoff, G., & Johnson, M. (1999). *Philosophy in the flesh: The embodied mind and its challenge to Western thought*. New York: Harper Collins.
- Lakoff, G., & Núñez, R. (2000). *Where mathematics comes from*. New York: Basic Books.
- Marcus, G. (2004). *The birth of the mind*. New York: Basic Books.
- Martin, A., Wiggs, C. L., Ungerleider, L. G., & Haxby, J. V. (1996). Neural correlates of category-specific knowledge. *Nature*, 379, 649–652.
- Mundale, J. (2002). Concepts of localization: Balkanization in the brain. *Brain and Mind*, 3(3), 313–330.
- Nishitani, N., Schürmann, M., Amunts K., & Hari, R. (2005). Broca's region: From action to language. *Physiology*, 20, 60–69.
- Prinz, J. (2005). Is the mind really modular? In R. Stainton (Ed.), *Contemporary debates in cognitive science*. New York: Blackwell.
- Riegler, A. (2001). The cognitive ratchet: The ratchet effect as a fundamental principle in evolution and cognition. *Cybernetics and Systems*, 32, 411–427.
- Rizzolatti, G., & Arbib, M. A. (1998). Language within our grasp. *Trends in Neurosciences*, 21, 188–194.
- Sandler, W., & Lillo-Martin, D. (2006). *Sign language and linguistic universals*. Cambridge: Cambridge University Press.
- Skarda, C., & Freeman, W. (1987). How the brain makes chaos to make sense of the world. *Behavioral and Brain Sciences*, 10, 161–195.
- Sporns, O., & Kötter, R. (2004). Motifs in brain networks. *PLoS Biology*, 2, e369.
- 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.
- Thompson, E., & Varela, F. (2001). Radical embodiment: Neural dynamics and consciousness. *Trends in Cognitive Sciences*, 5, 418–425.
- Uttal, W. (2001). *The new phrenology*. Cambridge: MIT Press.
- Wilson, M. (2001). The case for sensorimotor coding in working memory. *Psychonomic Bulletin and Review*, 8, 44–57.