

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

Michael L. Anderson

Received: ■ / Accepted: 1 June 2007
© Springer Science+Business Media B.V. 2007

1 **Abstract** The massive redeployment hypothesis (MRH) is a theory about the
2 functional topography of the human brain, offering a middle course between strict
3 localization on the one hand, and holism on the other. Central to MRH is the claim
4 that cognitive evolution proceeded in a way analogous to component reuse in software
5 engineering, whereby existing components—originally developed to serve some spe-
6 cific purpose—were used for new purposes and combined to support new capacities,
7 without disrupting their participation in existing programs. If the evolution of cognition
8 was indeed driven by such exaptation, then we should be able to make some specific
9 empirical predictions regarding the resulting functional topography of the brain. This
10 essay discusses three such predictions, and some of the evidence supporting them.
11 Then, using this account as a background, the essay considers the implications of
12 these findings for an account of the functional integration of cognitive operations. For
13 instance, MRH suggests that in order to determine the functional role of a given brain
14 area it is necessary to consider its participation across multiple task categories, and
15 not just focus on one, as has been the typical practice in cognitive neuroscience. This
16 change of methodology will motivate (even perhaps necessitate) the development of a
17 new, domain-neutral vocabulary for characterizing the contribution of individual brain
18 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



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

26 **Keywords** ■

27 **1 Introduction**

28 The massive redeployment hypothesis (Anderson 2006, 2007a,c) is both a theory about
29 the functional topography of the human brain, and also an account of how and why it
30 got that way. As a theory of the functional topography of the brain, it is perhaps most
31 easily understood as a middle ground between strict localization, on the one hand, and
32 holism on the other. The localization-holism debate has generally been presented in
33 terms of a choice between whether cognitive functions are typically instantiated by a
34 few and closely grouped neural circuits, or by many and widely distributed ones. Yet
35 as Mundale (2002) persuasively argues, the belief that cognitive functions typically
36 utilize many and widely distributed neural circuits is perfectly compatible with local-
37 ization. Just so long as one can (more or less clearly) distinguish the neural tissues
38 that support a given cognitive function from those that do not, one might claim to have
39 thereby “localized” the function.

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

51 In contrast to both localization and holism, a redeployment hypothesis splits the
52 difference, answering “no” to question 1, and “yes” to question 2. That is, a rede-
53 ployment hypothesis claims that parts of the brain are specialized, in that they do the
54 same thing each time they are activated. However, the thing that they do—the func-
55 tion they compute or transformation they effect—does not line up with any specific
56 cognitive function. Rather, brain areas must work in concert with other areas to do any-
57 thing interesting, and are therefore not generally deployed in support of only a single

58 function, but are instead *redeployed*¹ in many different functional complexes, which
59 do many different (interesting) things.

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

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

82 To understand these two commitments returns us to the second aspect of MRH
83 mentioned above, its account of how and why the brain came to be organized the
84 way it is. MRH proposes that cognitive evolution proceeded in a way analogous to
85 component reuse in software engineering (Heineman and Council 2001), whereby
86 existing components—originally developed to serve some specific purpose—are used
87 for new purposes and combined to support new capacities, without disrupting their
88 participation in existing programs. This would parallel the evolution of other capa-
89 bilities wherein existing structures, evolved for other purposes, are re-used and built
90 upon in the course of continuing evolutionary development (“exaptation”: Gould and
91 Vrba 1982). There is psychological support for exaptation in cognition (Cosmides
92 1989; Cruse 2003; Glenberg and Kaschak 2002; Gould 1991; Lakoff and Nuñez 2000;
93 Riegler 2001; Wilson 2001), theoretical reason to expect it (Anderson 2003, 2007b,
94 in press), and neuroanatomic evidence that the brain evolved by preserving, extend-
95 ing, and combining existing network components, rather than by generating complex
96 structures de novo (Sporns and Kötter 2004). MRH integrates these various perspec-
97 tives, bringing exaptive accounts of the evolution of the brain at both the functional
98 and neuroanatomic levels together in the realm of cognitive neuroscience.

¹ 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).

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

121 2 Evidence in support of MRH

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

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

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

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

141 a large set of cognitive functions in different domains (perception, language, etc.) for
142 which we had already determined, via functional imaging, the network of brain areas
143 supporting each function. In order to evaluate the empirical predictions suggested by
144 MRH, we would need to extract the following information from that set: For prediction
145 (1) we'd want to know (a) the number of individual functions a typical brain region
146 supports (i.e., how often it is activated), and (b) the distribution of those activations
147 across the cognitive domains; for prediction (2) we'd need to know (c) the number
148 of brain regions activated for a typical cognitive function, along with any significant
149 differences in this number between typical functions in different cognitive domains,
150 and (d) the degree of scatter in the network of brain regions supporting a typical
151 cognitive function, along with any significant differences in scatter between typical
152 functions in different cognitive domains, and for predictions (3) and (4) we'd need to
153 know (e) how this data—and especially data relating to differences in these measures
154 between different brain regions and functional domains—relates to information about
155 the phylogenetic age of the region and cognitive domain.

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

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

184 So, how can we use this list of tasks and activation to get the information we need?
185 To get data on (a) and (c)—the number of tasks activating each brain region, and
186 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).² 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

² 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.

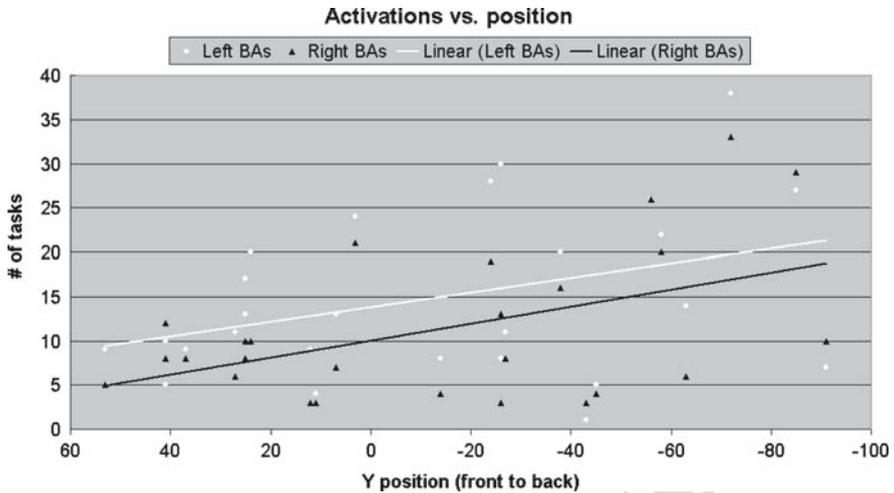


Fig. 1 Plot of number of activations for each Brodmann area versus the Y-position of the area

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

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

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

⁴ 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.

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

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

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

297 Such considerations follow directly from reflection on the evidence for MRH pre-
298 sented above. If it *were* the case that neural circuits, by whatever mechanism, could
299 be easily and quickly made to do different things for different functional complexes,
300 then metabolic considerations (e.g. the cost of establishing, maintaining, and using
301 long-distance connections between brain areas) would tend to favor recruitment of
302 local neural circuits to support developing functionality. Under these conditions, we
303 should expect to observe a consistent degree of localization over evolutionary time.
304 But the evidence presented above suggests that this is not the case. Although alternate
305 hypotheses may be developed to explain the observed decreasing localization of func-
306 tional complexes over evolutionary time, MRH offers a very natural fit with the data:
307 because of the time scales over which configurational changes are possible, neural circuits
308 can't be reconfigured quickly enough to support multiple arbitrary functions. Thus,
309 successful recruitment of these circuits into emerging functional complexes requires
310 that the circuit *already* have a functional role answering some need in the emerging
311 complex. There is no reason to suppose that all such circuits will be near one another,
312 and indeed, as the options grow, so too will the chance that the required circuits are
313 quite scattered. Seen in this light, the belief that neural circuits do the same thing in
314 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.

⁵ 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.

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

327 **3 MRH and the functional integration of cognitive operations**

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

340 First, in order to determine what any neural circuit does—what it contributes to
341 the functional complexes of which it is a part (I call this its “area role”)—it will be
342 necessary to consider the brain area's support of cognitive functions across multi-
343 ple domains, rather than focusing on a single domain or class of tasks as has been
344 the typical practice in cognitive neuroscience. To consider a particular example,
345 Broca's area has long been associated with language processing, but what has
346 recently begun to emerge is its functional complexity (Hagoort 2005). For instance, it
347 has been shown that Broca's area is involved in many different action- and
348 imagery-related tasks, including movement preparation (Thoenissen et al. 2002),
349 action sequencing (Nishitani et al. 2005), action recognition (Procyk et al. 1997;
350 Hamzei et al. 2003; Nishitani et al. 2005), imagery of human motion (Binkofski
351 et al. 2000), and action imitation (Nishitani et al. 2005). This is a rather rich array
352 of functions, and should suggest the poverty of approaching the task of functionally
353 modeling a brain region from within only one cognitive domain.

354 To attribute an area role to Broca's area, then, it will be necessary to consider not
355 just the cognitive models of one or a few language-related functions, and the vari-
356 ous options for mapping the functional boxes in *that* model to specific brain areas,

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

370 The second implication follows directly from the first. Insofar as our approach to
371 discovering the specific functional role of a given brain area involves modeling its
372 activity across different cognitive domains, then it makes little sense to try to charac-
373 terize the contribution of the area using domain-specific terms. This is an important
374 lesson to absorb, for as Cabeza and Nyberg note, researchers tend to offer functional
375 interpretations of observed brain activations using terms drawn from the domain within
376 which they work: "Area 7 activations, for instance, were usually attributed to atten-
377 tional processes in attention studies, to perceptual processes in perception studies,
378 to working memory processes in working memory studies, and so on." (Cabeza and
379 Nyberg 2000: 31). Instead of continuing on this path, we will need to develop a
380 domain-independent functional vocabulary to specify the contribution of each area.
381 Ideally, the vocabulary will also be such as to facilitate an understanding of how to
382 build complex functionality out of simple area roles. Currently, the most promising
383 candidates for this vocabulary use terms drawn from the theory of computation and
384 information processing, but we should be open to different developments, e.g. genetics
385 (Bickle 2003; Marcus 2004) and dynamical systems theory (Kelso et al. 1998; Skarda
386 and Freeman 1987; Thompson and Varela 2001).

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

396 The third implication, which has no doubt already occurred to many readers, is that
397 MRH sits somewhat uneasily with the oft-cited claim from classical cognitive sci-
398 ence that cognitive modules are domain specific (Fodor 1983). Central to the attrac-
399 tion of the modularity hypothesis, and one of the motivations for insisting on the
400 domain-specificity of cognitive modules, was an underlying architecture able to sup-
401 port massively parallel processing, explaining our (apparent) ability to avoid com-
402 putational bottlenecks. Now, if we are content to identify a module with whatever

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

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

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

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

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

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

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

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

⁶ 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.

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

491 Or consider another example of the apparent redeployment of motor-control
492 resources in language, this time to support language *understanding*, rather than lan-
493 guage production: the action-sentence compatibility effect (Glenberg and Kaschak
494 2002). To demonstrate this interesting interaction between comprehension and motor
495 control, Glenberg and Kaschak asked subjects to indicate whether a given sentence
496 (e.g., “Close the drawer”) made sense or not by reaching either toward or away from
497 their bodies to press the appropriate button. They found that response times were
498 longer in cases where the required movement ran counter to the movement suggested
499 by the sentence itself, and that this was true even when the “movement” indicated by
500 the sentence was abstract, as in the transfer of information from one party to another
501 (e.g., “John told you about the party.”).

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

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

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

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

545 4 Conclusion

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

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

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

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

568 References

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

- 574 Anderson, M. L. (2007b). How to study the mind: An introduction to embodied cognition. In F. Santoianni
575 & C. Sabatana (Eds.), *Brain development in learning environments: Embodied and perceptual advance-*
576 *ments*. Cambridge: Cambridge Scholars Press.
- 577 Anderson, M. L. (2007c). The massive redeployment hypothesis and the functional topography of the brain.
578 *Philosophical Psychology*, 21(2), 143–174.
- 579 Anderson, M. L. (in press). Action-grounded cognition: Evolution, embodiment and the nature of the mind.
580 Keynote address, Cognition 2006. In B. Hardy-Vallee & N. Payette (Eds.), *Beyond the brain: embodied,*
581 *situated & distributed cognition*. Cambridge: Cambridge Scholars Press.
- 582 Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577–660.
- 583 Bickle, J. (2003). *Philosophy and neuroscience, a ruthlessly reductive account*. Dordrecht: Kluwer
584 Academic Publishers.
- 585 Binkofski, F., Amunts, K., Stephan, K. M., Posse, S., Schormann, T., Freund, H.-J., Zilles, K., & Seitz,
586 R. J. (2000). Broca's region subserves imagery of motion: A combined cytoarchitectonic and fMRI
587 study. *Human Brain Mapping*, 11, 273–285.
- 588 Cabeza, R., & Nyberg, L. (2000). Imaging cognition II: An empirical review of 275 PET and fMRI studies.
589 *Journal of Cognitive Neuroscience*, 12, 1–47.
- 590 Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason?
591 Studies with the Wason selection task. *Cognition*, 31, 187–276.
- 592 Cruse, H. (2003). The evolution of cognition—a hypothesis. *Cognitive Science*, 27, 135–155.
- 593 Damasio, A. R., & Tranel, D. (1993). Nouns and verbs are retrieved with differently distributed neural
594 systems. *Proceedings of the National Academy of Sciences of the USA*, 90, 4957–4960.
- 595 Decety, J., Grézes, J., Costes, N., Perani, D., Jeannerod, M., Procyk, E., Grassi, F., & Fazio, F. (1997).
596 Brain activity during observation of actions. Influence of action content and subject's strategy. *Brain*,
597 120, 1763–1777.
- 598 Fauconnier, G., & Turner, M. (2002). *The way we think: Conceptual blending and the mind's hidden*
599 *complexities*. New York: Basic Books.
- 600 Fodor, J. (1983). *The modularity of mind*. Cambridge, MA: Bradford Books.
- 601 Fowler, C. A., Rubin, P., Remez, R. E., & Turvey, M. T. (1980). Implications for speech production of a
602 general theory of action. In B. Butterworth (Ed.), *Language production, Volume 1: Speech and talk*
603 (pp. 373–420). London: Academic Press.
- 604 Gibson, J. J. (1979/1987). *The ecological approach to visual perception*. New York: Lawrence Erlbaum
605 Associates.
- 606 Gil-da-Costa, R., Braun, A., Lopes, M., Hauser, M. D., Carson, R. E., Herscovitch, P., & Martin, A. (2004).
607 Conceptual representations in a nonhuman primate: Species-specific vocalizations activate visual and
608 affective processing systems in the macaque *Proceedings of the National Academy of Sciences, USA*,
609 101, 17516–17521.
- 610 Glenberg, A., & Kaschak, M. (2002). Grounding language in action. *Psychonomic Bulletin and Review*, 9,
611 558–565.
- 612 Gorniak, P., & Roy, D. (2006). Perceived affordances as a substrate for linguistic concepts. *Proceedings of*
613 *the twenty-eighth annual meeting of the Cognitive Science Society*.
- 614 Gould, S. J. (1991). Exaptation: A crucial tool for an evolutionary psychology. *Journal of Social Issues*, 3,
615 43–65.
- 616 Gould, S. J., & Vrba, E. (1982). Exaptation: A missing term in the science of form. *Paleobiology*, 8, 4–15.
- 617 Graziano, M. S. A., Taylor, C. S. R., & Moore, T. (2002a). Complex movements evoked by microstimulation
618 of precentral cortex. *Neuron*, 34, 841–851.
- 619 Graziano, M. S. A., Taylor, C. S. R., Moore, T., & Cooke, D. F. (2002b). The cortical control of movement
620 revisited. *Neuron*, 36, 349–362.
- 621 Hagoort, P. (2005). On Broca, brain and binding. In Y. Grodinsky & K. Amunts (Eds.), *Broca's region*.
622 Oxford: Oxford University Press.
- 623 Hamzei, F., Rijntjes, M., Dettmers, C., Glauche, V., Weiller, C., & Büchel, C. (2003). The human action
624 recognition system and its relationship to Broca's area: An fMRI study. *Neuroimage*, 19, 637–644.
- 625 Heineman, G. T., & Councill, W. T. (2001). *Component-based software engineering: Putting the pieces*
626 *together*. New York: Addison-Wesley.
- 627 Kelso, J. A. S., Fuchs, R., Lancaster, T., Holroyd, D., Cheyne, H., Weinberg, H. (1998). Dynamic cortical
628 activity in the human brain reveals motor equivalence. *Nature*, 23, 814–818.
- 629 Koch, C., & Segev, I. (2000). The role of single neurons in information processing. *Nature Neuroscience*,
630 3, 1171–1177.

- 631 Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.
- 632 Lakoff, G., & Johnson, M. (1999). *Philosophy in the flesh: The embodied mind and its challenge to Western*
633 *thought*. New York: Harper Collins.
- 634 Lakoff, G., & Núñez, R. (2000). *Where mathematics comes from*. New York: Basic Books.
- 635 Marcus, G. (2004). *The birth of the mind*. New York: Basic Books.
- 636 Martin, A., Wiggs, C. L., Ungerleider, L. G., & Haxby, J. V. (1996). Neural correlates of category-specific
637 knowledge. *Nature*, 379, 649–652.
- 638 Mundale, J. (2002). Concepts of localization: Balkanization in the brain. *Brain and Mind*, 3(3), 313–330.
- 639 Nishitani, N., Schürmann, M., Amunts K., & Hari, R. (2005). Broca's region: From action to language.
640 *Physiology*, 20, 60–69.
- 641 Prinz, J. (2005). Is the mind really modular? In R. Stainton (Ed.), *Contemporary debates in cognitive*
642 *science*. New York: Blackwell.
- 643 Riegler, A. (2001). The cognitive ratchet: The ratchet effect as a fundamental principle in evolution and
644 cognition. *Cybernetics and Systems*, 32, 411–427.
- 645 Rizzolatti, G., & Arbib, M. A. (1998). Language within our grasp. *Trends in Neurosciences*, 21, 188–194.
- 646 Sandler, W., & Lillo-Martin, D. (2006). *Sign language and linguistic universals*. Cambridge: Cambridge
647 University Press.
- 648 Skarda, C., & Freeman, W. (1987). How the brain makes chaos to make sense of the world. *Behavioral and*
649 *Brain Sciences*, 10, 161–195.
- 650 Sporns, O., & Kötter, R. (2004). Motifs in brain networks. *PLoS Biology*, 2, e369.
- 651 Thoenissen, D., Zilles, K., & Toni, I. (2002). Differential involvement of parietal and precentral regions in
652 movement preparation and motor intention. *Journal of Neuroscience*, 22, 9024–9034.
- 653 Thompson, E., & Varela, F. (2001). Radical embodiment: Neural dynamics and consciousness. *Trends in*
654 *Cognitive Sciences*, 5, 418–425.
- 655 Uttal, W. (2001). *The new phrenology*. Cambridge: MIT Press.
- 656 Wilson, M. (2001). The case for sensorimotor coding in working memory. *Psychonomic Bulletin and Review*,
657 8, 44–57.