Neural reuse: A fundamental organizational principle of the brain

Michael L. Anderson

*Department of Psychology, Franklin & Marshall College, Lancaster, PA 17604,
and Institute for Advanced Computer Studies, Neuroscience and Cognitive
Science Program, University of Maryland, College Park, MD 20742*
michael.anderson@fandm.edu
http://www.agcognition.org

Abstract: An emerging class of theories concerning the functional structure of the brain takes the reuse of neural circuitry for various cognitive purposes to be a central organizational principle. According to these theories, it is quite common for neural circuits established for one purpose to be exapted (exploited, recycled, redeployed) during evolution or normal development, and be put to different uses, often without losing their original functions. Neural reuse theories thus differ from the usual understanding of the role of neural *plasticity* (which is, after all, a kind of reuse) in brain organization along the following lines: According to neural reuse, circuits can continue to acquire new uses after an initial or original function is established; the acquisition of new uses need not involve unusual circumstances such as injury or loss of established function; and the acquisition of a new use need not involve (much) local change to circuit structure (e.g., it might involve only the establishment of functional connections to new neural partners). Thus, neural reuse theories offer a distinct perspective on several topics of general interest, such as: the evolution and development of the brain, including (for instance) the evolutionary-developmental pathway supporting primate tool use and human language; the degree of modularity in brain organization; the degree of localization of cognitive function; and the cortical parcellation problem and the prospects (and proper methods to employ) for function to structure mapping. The idea also has some practical implications in the areas of rehabilitative medicine and machine interface design.

Keywords: brain; development; evolution; exaptation; functional architecture; localization; modularity

Although an organ may not have been originally formed for some special purpose, if it now serves for this end we are justified in saying that it is specially contrived for it. On the same principle, if a man were to make a machine for some special purpose, but were to use old wheels, springs, and pulleys, only slightly altered, the whole machine, with all its parts, might be said to be specially contrived for that purpose. Thus throughout nature almost every part of each living being has probably served, in a slightly modified condition, for diverse purposes, and has acted in the living machinery of many ancient and distinct specific forms.
— Charles Darwin (1862), p. 348.

1. Introduction and background

Research in the cognitive neurosciences has long been guided by the idealization that brain regions are highly selective and specialized, and that function can be mapped to local structure in a relatively straightforward way. But the degree of actual selectivity in neural structures is increasingly a focus of debate in cognitive science (Poldrack 2006). It appears that many structures are activated by different tasks across different task categories and cognitive domains. For instance, although Broca's area has been strongly associated with language processing, it turns out to also be involved in many different action- and imagery-related tasks, including movement preparation (Thoenissen et al. 2002), action sequencing (Nishitani et al. 2005), action recognition (Decety et al. 1997; Hamzei et al. 2003; Nishitani et al. 2005), imagery of human motion (Binkofski et al. 2000), and action imitation (Nishitani et al. 2005; for reviews, see Hagoort 2005; Tettamanti & Weniger 2006). Similarly, visual and motor areas – long presumed to be among the most highly specialized in the brain – have been shown to be active in various sorts of language processing and other higher cognitive tasks (Damasio & Tranel 1993; Damasio et al. 1996; Glenberg & Kaschak 2002; Hanakawa et al. 2002; Martin et al. 1995; 1996; 2000; Pulvermüller 2005; see sect. 4 for a discussion). Excitement over the discovery of the Fusiform Face Area (Kanwisher et al. 1997) was quickly tempered when it was discovered that the area also responded to cars, birds, and other stimuli (Gauthier et al. 2000; Grill-Spector et al. 2006; Rhodes et al. 2004).

MICHAEL L. ANDERSON, Assistant Professor of Cognitive Science in the Department of Psychology at Franklin & Marshall College, is author or co-author of more than sixty scholarly and scientific publications in cognitive science, artificial intelligence, and philosophy of mind. His papers include: "Evolution of cognitive function via redeployment of brain areas," "Circuit sharing and the implementation of intelligent systems," "Investigating functional cooperation in the human brain using simple graph-theoretic methods," "A self-help guide for autonomous systems," and "Embodied cognition: A field guide." Anderson was recently nominated for the Stanton Prize, recognized as an "emerging leader under 40" by the Renaissance Weekend, and was an invited participant in the McDonnell Project in Philosophy and the Neurosciences workshop for early career researchers.

The ensuing debates over the "real" function of these areas have still not been resolved.

This is just a short list of some highly-studied regions for which the prospect of a clear-cut mapping of function to structure appears dim. In this target article, I will review a great deal more evidence that points in a similar direction. But if selectivity and localization are not in fact central features of the functional organization of the brain, how shall we think about the function-structure relationship? This target article reviews an emerging class of theories that suggest neural circuits established for one purpose are commonly exapted (exploited, recycled, redeployed) during evolution or normal development, and put to different uses, often without losing their original functions. That is, rather than posit a functional architecture for the brain whereby individual regions are dedicated to large-scale cognitive domains like vision, audition, language, and the like, neural reuse theories suggest instead that low-level neural circuits are used and reused for various purposes in different cognitive and task domains.

In just the past five years, at least four different, specific, and empirically supported general theories of neural reuse have appeared. Two of these theories build on the core notion of the sensorimotor grounding of conceptual content to show how it could implicate many more aspects of human cognitive life: Vittorio Gallese's "neural exploitation" hypothesis (Gallese 2008; Gallese & Lakoff 2005) and Susan Hurley's "shared circuits model" (Hurley 2005; 2008). Two other theories suggest that reuse could be based on even more universal foundations: Dehaene's "neuronal recycling" theory (Dehaene 2005; 2009; Dehaene & Cohen 2007) and my own "massive redeployment" hypothesis (M. L. Anderson 2007a; 2007c).[1] These latter two suggest reuse might in fact constitute a fundamental developmental (Dehaene's recycling theory) or evolutionary (my redeployment hypothesis) strategy for realizing cognitive functions. Others are clearly thinking along similar lines, for example, Luiz Pessoa (2008), Gary Marcus (2004; 2008), Steven Scher (2004), William Bechtel (2003), and Dan Lloyd (2000). These models have some interesting similarities and equally interesting differences, but taken together they offer a new research-guiding idealization of brain organization, and the potential to significantly impact the ongoing search for the brain basis of cognition.

I discuss each model, and what these models might collectively mean for cognitive science, in sections 6 and 7, after reviewing some of the broad-based evidence for neural reuse in the brain (sects. 4 and 5). In order to better appreciate that evidence and its implications, however, it will be useful to have before us a more concrete example of a theory of neural reuse, and some sense of where such theories fit in the landscape of cognitive science. To this end, the next subsection briefly details one of the theories of reuse – the massive redeployment hypothesis – and sections 2 through 5 serve to situate reuse with respect to some other well-known accounts of the functional structure of the brain.

1.1. The massive redeployment hypothesis

The core of the massive redeployment hypothesis is the simple observation that evolutionary considerations might often favor reusing existing components for new tasks over developing new circuits de novo. At least three predictions follow from this premise. 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 support the localist story 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 reuse 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 a greater number of and more widely scattered brain areas than evolutionarily older functions, because the later a function is developed, the more likely it is 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 happen to reside in neighboring brain regions. A more localist account of the evolution of the brain would instead expect the continual development of new, largely dedicated neural circuits, and 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 among them.

In a number of recent publications (M. L. Anderson 2007a; 2007c; 2008a) I report evidence for all of these predictions. Consider, for instance, some data demonstrating the first prediction, that a typical brain region serves tasks across multiple task categories. An empirical review of 1,469 subtraction-based fMRI experiments in eleven task domains reveals that a typical cortical region[2] is activated by tasks in *nine* different domains. The domains investigated were various – action execution, action inhibition, action observation, vision, audition, attention, emotion, language, mathematics, memory, and reasoning – so this observation cannot be explained by the similarity of the task domains. And because the activations were post-subtraction activations, the finding is not explained by the fact that most experimental tasks have multiple cognitive aspects (e.g., viewing stimuli, recalling information, making responses). Control tasks would (mostly) ensure that the reported brain activity was supporting the particular cognitive function under investigation. Finally, the observation is not explained by the size of the regions studied. As recounted in more detail in section 5, below, one gets the same pattern of results even when dividing the cortex into nearly 1,000 small regions.[3]

In evaluating the second prediction, one is immediately faced with the trouble that there is little consensus on which areas of the brain are older. I therefore employed the following oversimplification: All things being equal, areas in the back of the brain are older than areas in the front of the brain (M. L. Anderson 2007a). Thus, the prediction is for a relationship between the position of a brain region along the Y-axis in Talairach space (Talairach & Tournaux 1988) and the frequency with which it is used in cognitive functions. The study reports the expected negative correlation[4] between the Y-position and the number of tasks in which it is active ($r = -0.412$, $p = .003$, $t = -3.198$, $df = 50$). A similar analysis using

the data set mentioned above reveals a negative correlation between the number of *domains* in which an anatomical region is activated and the Y-position of the region ($r = -0.312$, $p = 0.011$, $t = -2.632$, $df = 65$). Although the amount of variance explained in these cases is not especially high, the findings are nevertheless striking, at least in part because a more traditional theory of functional topography would predict the *opposite* relation, if there were any relation at all. According to traditional theories, older areas – especially those visual areas at the back of the brain – are expected to be the *most* domain dedicated. But that is not what the results show.

As for the last prediction, that more recently evolved functions will be supported by more broadly scattered regions of activation, in (M. L. Anderson 2007a), I reported that language tasks activate more and more broadly scattered regions than do visual perception and attention. This finding was corroborated by a larger study (M. L. Anderson 2008a), which found that language was the most widely scattered domain of those tested, followed (in descending order) by reasoning, memory, emotion, mental imagery, visual perception, action, and attention. The significant differences in the degree of scatter were observed between attention and each of the following domains: language, reasoning, memory, emotion, and mental imagery; and between language and each of the following domains: visual perception, action, and attention. No other pair-wise comparisons showed significant differences.

Note that, in addition to supporting the main contentions of the massive redeployment hypothesis, this last finding also corroborates one of the main assumptions behind most theories of neural reuse: that cortical regions have specific biases that limit the uses to which they can be put without extensive rewiring. If neural circuits could be easily put to almost any use (that is, if small neural regions were locally poly-functional, as some advocates of connectionist models suggest), then given the increased metabolic costs of maintaining long-distance connections, we would expect the circuits implementing functions to remain relatively localized. That this is not the observed pattern suggests that some functionally relevant aspect of local circuits is relatively fixed. The massive redeployment hypothesis explains this with the suggestion that local circuits may have low-level computational "workings" that can be put to many different higher-level cognitive uses.[5]

If this is the right sort of story, it follows that the functional differences between task domains cannot be accounted for primarily by differences in which brain regions get utilized – as they are reused across domains. And naturally, if one puts together the same parts in the same way, one will get the same functional outcomes. So, the functional differences between cognitive domains should reveal themselves in the (different) ways in which the (shared) parts are assembled. I explored this possibility using a co-activation analysis – seeing which brain regions were statistically likely to be co-active under what task conditions. The results indicated that although different domains do indeed tend to be supported by overlapping neural regions, each task domain was characterized by a distinctive pattern of co-activation among the regions (M. L. Anderson 2008a). This suggests an overall functional architecture for the brain that is quite different from that proposed by anatomical modularity and functional localization (see Fig. 1).

Keeping this substantive introduction to the concept of neural reuse in view, I will devote the next three sections to situating neural reuse with respect to three relevant classes of theory in cognitive science, and return to both neural reuse theory and supporting data in sections 5 and 6. For the purposes of this review, it is important to note that neural reuse theories are *not* full-fledged theories of how the brain (or mind) works. Rather, they are theories of how neural resources are (typically) deployed in support of cognitive functions and processes. Given this, there are at least three relevant comparison classes for neural reuse, each of which I discuss in turn in the sections that follow.

First, in section 2, I briefly discuss some *other* theories – anatomical modularity and global wiring optimization theory – for how neural resources are typically deployed in support of the brain's function. Then, in section 3, I turn to some theories of overall cognitive architecture – ACT-R, massive modularity, and both classic and contemporary parallel distributed processing models – and what they may imply for neural reuse and vice versa. And finally, in section 4, I examine at some length some other theories that predict neural reuse, notably concept empiricism and conceptual metaphor theory, as part of an argument that these established theories are not adequate to account for the full range of neural reuse that can be observed in the brain.

2. How are neural resources deployed in the brain?

There are two prominent theories for how neural resources are deployed in the function and structure of the brain: anatomical modularity and global wiring optimization theory. We will see that neural reuse is deeply
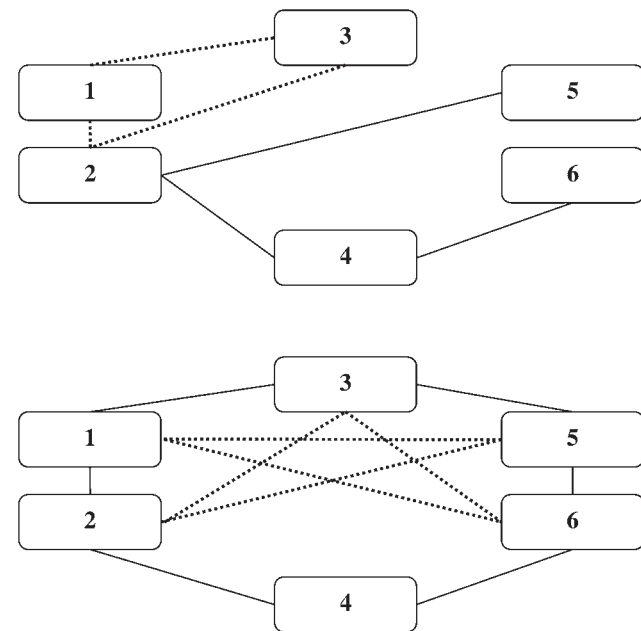

Figure 1. Expected patterns of co-activation in a simple six-region brain for two cognitive functions (solid vs. dashed lines). Anatomical modularity and localization (top) predicts largely non-overlapping sets of regions will contribute to each function, whereas reuse (bottom) suggests that many of the same cortical regions will be activated in support of both functions, but that they will co-activate (cooperate) in different patterns.

incompatible with anatomical modularity, but compatible with wiring optimization theory. In fact, in combination neural reuse and wiring optimization theory make some novel predictions for cortical layout.

2.1. Anatomical modularity

Anatomical modularity is functional modularity plus a strong thesis about how the functional modules are implemented in the brain. Functional modularity is (minimally) the thesis that our cognitive systems are composed of separately modifiable (or "nearly decomposable"; Simon 1962/1969) subsystems, each typically dedicated to specific, specialized functions (see sect. 3.1 for a discussion). Anatomical modularity is the additional thesis that each functional module is implemented in a dedicated, relatively small, and fairly circumscribed piece of neural hardware (Bergeron 2007).

Simply put, neural reuse theories suggest anatomical modularity is false. According to the picture painted by reuse, even if there is functional modularity (see sect. 3.1), individual regions of the brain will turn out to be part of multiple functional modules. That is, brain regions will not be dedicated to single high-level tasks ("uses"), and different modules will not be implemented in separate, small, circumscribed regions. Instead, different cognitive functions are supported by putting many of the same neural circuits together in different arrangements (M. L. Anderson 2008a). In each of these arrangements, an individual brain region may perform a similar information-processing operation (a single "working"), but will not be dedicated to that one high-level use.

Although there are few defenders of a strong anatomical modularity hypothesis, Max Coltheart (2001) goes so far as to include it as one of the fundamental assumptions guiding cognitive neuropsychology. The idea is that the success of neuropsychological research – relying as it does on patients with specific neurological deficits, and the discovery of double-dissociations between tasks – both requires and, in turn, supports the assumption that the brain is organized into anatomical modules. For if it were not, we wouldn't observe the focal deficits characteristic of some brain injuries, and nor would we be able to gather evidentiary support for double-dissociations between tasks.

If this argument were sound, then the success of neuropsychology as a discipline would itself be prima facie evidence against neural reuse. In fact, the inference is fairly weak. First, it is possible for focal lesions to cause specific functional deficits in non-modular systems (Plaut 1995), and double-dissociations do not by themselves support any inference about the underlying functional architecture of the brain (Van Orden et al. 2001). In any event, such deficits are the exception rather than the rule in human brain injuries. Even some of the patients most celebrated for having specific behavioral deficits often have multiple problems, even when one problem is the most obvious or debilitating (see Bergeron 2007; Prinz 2006 for discussions). The evidence coming from neuropsychology, then, is quite compatible with the truth of neural reuse. But is neural reuse compatible with the methodological assumptions of cognitive neuropsychology? Section 7 will discuss some of the specific methodological changes that will be needed in the cognitive neurosciences in light of widespread neural reuse.

2.2. Optimal wiring hypotheses

The layout of neurons in the brain is determined by multiple constraints, including biomorphic and metabolic limitations on how big the brain can be and how much energy it can consume. A series of studies by Christopher Cherniak and others has reported that the layout of the nervous system of C. elegans, the shape of typical mammalian neuron arbors, and the placement of large-scale components in mammalian cortex are all nearly optimal for minimizing the total length of neurons required to achieve the structure (Cherniak et al. 2004; see also Wen & Chklovskii 2008). The last finding is of the highest relevance here. Cherniak et al. examined the 57 Brodmann areas of cat cortex. Given the known connections between these regions, it turns out that the Brodmann areas are spatially arranged so as to (nearly) minimize the total wiring length of those connections.

This is a striking finding; and even though this study examined physical and not functional connectivity, the two are undoubtedly related – at least insofar as the rule that "neurons that fire together wire together" holds for higher-level brain organization. In fact, Cherniak et al. (2004) predict that brain areas that are causally related – that co-activate, for instance – will tend to be physically adjacent. The data reviewed above did not exactly conform to this pattern. In particular, it seems that the neural regions supporting more recent cognitive functions tended to be *less* adjacent – farther apart in the brain – than those supporting older cognitive functions. Nevertheless, neural reuse and the global optimization of component layout appear broadly compatible, for four reasons. First, wiring length can hardly be considered (and Cherniak et al. do not claim that it is) the only constraint on cortical structure. The total neural mass required to achieve the brain's function should also be kept minimal, and reuse would tend to serve that purpose. Second, it should be kept in mind that Cherniak et al. (2004) predict *global* optimization in component layout, and this is not just compatible with, but also positively predicts that subsets of components will be less optimal than the whole. Third, there is no reason to expect that all subsets will be equally suboptimal; global optimality is compatible with differences in the optimality of specific subsets of components. Fourth, when there *is* a difference in the optimality of component subsets, neural reuse would predict that these differences would track the evolutionary age of the supported function. That is, functionally connected components supporting recently evolved functions should tend to be less optimally laid out than those supporting older functions. More specifically, one would expect layout optimality to correlate with the ratio of the age of the cognitive function to the total evolutionary age of the organism. When functional cooperation emerged early in the evolution of the cortex, there is a greater chance that the components involved will have arrived at their optimal locations, and less chance of that for lower ratios, as overall brain morphology will not have had the same evolutionary opportunity to adjust.

This notion is not at all incompatible with the thesis of global (near-) optimality and indeed might be considered a refinement of its predictions. Certainly, this is a research direction worth pursuing, perhaps by merging the anatomical connectivity data-sets from Hagmann et al. (2008) with

functional databases like BrainMap (Laird et al. 2005) and the NICAM database (M. L. Anderson et al. 2010). In fact, I am currently pursuing a related project, to see whether co-activation strength between regions predicts the existence of anatomical connections.

3. Cognitive architectures

In this section, I review four of the most commonly adopted approaches to understanding how the mind is functionally structured, and the implications of these approaches for the functional structure of the brain: massive modularity; ACT-R; and classic and contemporary parallel distributed processing models. Neural reuse appears to undermine the main motivation for positing massive modularity, and although reuse is broadly compatible with the other three theories, it seems likely to somewhat modify the direction of each research program.

3.1. Massive modularity

As noted above, functional modularity is minimally the thesis that the mind can be functionally decomposed into specialized, separately modifiable subsystems – individual components charged with handling one or another aspect of our mental lives. Carruthers (2006) follows this formulation:

> In the weakest sense, a module can just be something like: a dissociable functional component. This is pretty much the everyday sense in which one can speak of buying a hi-fi system on a modular basis, for example. The hi-fi is modular if one can purchase the speakers independently of the tape-deck, say, or substitute one set of speakers for another with the same tape deck. (Carruthers 2006, p. 2)

Massive modularity, which grows largely out of the modularity movement in evolutionary psychology (Pinker 1997; Sperber 1996; Tooby & Cosmides 1992) is the additional thesis that the mind is mostly, if not entirely, composed of modules like this – largely dissociable components that vary independently from one another. Is such a vision for the mind's architecture compatible with widespread neural reuse? Carruthers (2006) certainly thinks so:

> If minimizing energetic costs were the major design criterion, then one would expect that the fewer brain systems that there are, the better. But on the other hand the evolution of multiple functionality requires that those functions should be underlain by separately modifiable systems, as we have seen. As a result, what we should predict is that while there will be many modules, those modules should *share parts* wherever this can be achieved without losing too much processing efficiency (and subject to other constraints: see below). And, indeed, there is now a great deal of evidence supporting what Anderson [2007c] calls "the massive redeployment hypothesis". This is the view that the components of brain systems are frequently deployed in the service of multiple functions. (Carruthers 2006, pp. 23–24; emphasis his)

As much as I appreciate Carruthers' swift adoption of the redeployment hypothesis, I am troubled by some aspects of this argument. First, it appears to contain a false premise: Energetic constraints predict *more compact or localized*, not necessarily *fewer* brain systems. Second, it may be logically invalid, because if functions must be underlain by separately modifiable systems, then they *cannot* be built from shared parts. That is, it appears that this apparently small concession to neural reuse in fact undermines the case for massive modularity.

Consider Carruthers' hi-fi system analogy. There it is true that the various components might share the amplifier and the speakers, the way many different biological functions – eating, breathing, communicating – "share" the mouth. But if neural reuse is the norm, then circuit sharing in the brain goes far beyond such intercommunication and integration of parts. The evidence instead points to the equivalent of sharing knobs and transistors and processing chips. A stereo system designed like this would be more like a boom-box, and its functional components would therefore *not* be separately modifiable. Changing a chip to improve the radio might well also change the performance of the tape player.[6]

To preview some of the evidence that will be reviewed in more detail in section 4, the brain may well be more boom-box than hi-fi. For instance, Glenberg et al. (2008a) report that use-induced motor plasticity also affects language processing, and Glenberg et al. (2008b) report that language processing modulates activity in the motor system. This connection is confirmed by the highly practical finding that one can improve reading comprehension by having children manipulate objects (Glenberg et al. 2007). And of course there are many other such examples of cognitive interference between different systems that are routinely exploited by cognitive scientists in the lab.

This does not mean that all forms of functional modularity are necessarily false – if only because of the myriad different uses of that term (see Barrett & Kurzban 2006 for a discussion). But it does suggest that modularity advocates are guided by an idealization of functional structure that is significantly at odds with the actual nature of the system. Instead of the decompose-and-localize approach to cognitive science that is advocated and exemplified by most modular accounts of the brain, neural reuse encourages "network thinking" (Mitchell 2006). Rather than approach a complex system by breaking functions into subfunctions and assigning functions to proper parts – a heuristic that has been quite successful across a broad range of sciences (Bechtel & Richardson 1993; 2010) – network thinking suggests one should look for higher-order features or patterns in the behavior of complex systems, and advert to these in explaining the functioning of the system. The paradigm exemplars for this sort of approach come from the discovery of common, functionally relevant topological structures in various kinds of networks, ranging from human and insect social networks to the phone grid and the Internet, and from foraging behaviors to the functioning of the immune system (Barabási & Albert 1999; Barabási et al. 2000; Boyer et al. 2004; Brown et al. 2007; Jeong et al. 2000; Newman et al. 2006). Although it is hardly the case that functional decomposition is an ineffective strategy in cognitive science, the evidence outlined above that patterns of neural co-activation distinguish between cognitive outcomes more than the cortical regions involved do by themselves suggests the need for a supplement to business as usual.

Even so, there are (at least) two objections that any advocate of modularity will raise against the picture of brain organization that is being painted here: Such a brain could not have evolved, because (1) the structure

would be too complex, and (2) it would be subject to too much processing interference and inefficiency.

Carruthers (2006) follows Simon (1962/1969) in making the first argument:

> Simon [1962/1969] uses the famous analogy of the two watchmakers to illustrate the point. One watchmaker assembles one watch at a time, attempting to construct the whole finished product at once from a given set of micro components. This makes it easy for him to forget the proper ordering of parts, and if he is interrupted he may have to start again from the beginning. The second watchmaker first builds a set of sub-components out of given micro component parts and then combines those into larger sub-component assemblies, until eventually the watches are complete . . . . Simon's argument is really an argument from *design*, then, whether the designer is natural selection (in the case of biological systems) or human engineers (in the case of computer programs). It predicts that, in general, each element added incrementally to the design should be realized in a functionally distinct sub-system, whose properties can be varied independently of the others (to a significant degree, modulated by the extent to which component parts are shared between them). It should be possible for these elements to be added to the design without necessitating changes within the other systems, and their functionality might be lost altogether without destroying the functioning of the whole arrangement. (Carruthers 2006, pp. 13, 25; emphasis in original)

The argument from design set forth here is more convincing when it is applied to the original emergence of a complex system than when it is applied to its subsequent evolutionary development. What the argument says is that it must be possible for development to be *gradual*, with functional milestones, rather than all-or-nothing; but neural reuse hardly weakens the prospect of a gradual emergence of new functions. And the possibility that new functionality can be achieved by combining existing parts in new ways – which undermines independent variation and separate modifiability, as Carruthers (2006) admits, here – suggests that a modular architecture is only one possible outcome from such gradualism.

Moreover, the strong analogy between natural selection and a designer may not be the most helpful conceptual tool in this case. When one thinks about the brain the way a human designer would, the problem that neural reuse presents is one of taking a given concrete circuit with a known function and imagining novel uses for it. That this process can be *very* difficult appears to place a heavy burden on reuse theories: How could such new uses ever be successfully designed? But suppose instead that, in building a given capacity, one is offered a plethora of components with *unknown* functions. Now the task is quite different: Find a few components that do something useful and can be arranged so as to support the current task – *whatever* their original purpose. Thus is a problem of *design imagination* turned into a problem of *search*. Evolution is known to be quite good at solving problems of the latter sort (Newell & Simon 1976), and it is useful to keep this alternate analogy for the evolutionary process in mind here.

This brings us to the second objection, that non-modular systems would suffer from disabling degrees of interference and processing inefficiency. Here, it may be useful to recall some of the main findings of the situated/embodied cognition movement (M. L. Anderson 2003; Chemero 2009; Clark 1997; 1998). Central to the picture of cognition offered there is the simple point that organisms evolve in a particular environment to meet the particular survival challenges that their environment poses. Situated/embodied cognition emphasizes that the solutions to these problems often rely in part on features of the environments themselves; for example, by adopting heuristics and learning biases that reflect some of the environments' structural invariants (Gigerenzer et al. 1999; Gilovitch et al. 2002). One such useful feature of most environments is that they don't pose all their problems all at once – inclement weather rarely comes along with predator abundance, pressing mating opportunities, and food shortages, for instance. And often when there *are* competing opportunities or challenges, there will be a clear priority. Thus, an organism with massively redeployed circuitry can generally rely on the temporal structure of events in its environment to minimize interference. Were this environment-organism relationship different – or if it were to change – then neural reuse *does* predict that increased interference will be one likely result.

Interestingly, contemporary humans encounter just such a changed organism-environment relationship in at least two arenas, and the effect of reused circuitry can often be seen as a result: First, in the labs of some cognitive scientists, who carefully engineer their experiments to exploit cognitive interference of various sorts; and, second, at the controls of sophisticated machinery, where the overwhelming attentional demands have been observed to cause massive processing bottlenecks, often with dangerous or even deadly results (Fries 2006; Hopkin 1995). It is no coincidence that, in addition to designing better human-machine interfaces, one important way of minimizing the problems caused by processing bottlenecks is to engineer the *environment*, including, especially, changing its task configuration and social structure, for instance by designing more efficient teams (Hutchins 1995).

3.2. ACT-R

At the core of ACT-R is the notion of a cognitive architecture, "a specification of the structure of the brain at a level of abstraction that explains how it achieves the function of the mind" (J. R. Anderson 2007, p. 7). ACT-R is explicitly modular. As of ACT-R 6.0, it consisted of eight functionally specialized, domain-specific, relatively encapsulated, independently operating, and separately modifiable components. Given J. R. Anderson's definition of a cognitive architecture, it might seem to directly follow that ACT-R is committed to the notion that the brain, too, consists of functionally specialized, domain-specific, relatively encapsulated, independently operating, and separately modifiable regions that implement the functional modules of the ACT-R model. Certainly, recent experiments meant to associate ACT-R components with specific brain regions encourage this impression (J. R. Anderson 2007; J. R. Anderson et al. 2007). As he argues:

> As discussed above, modular organization is the solution to a set of structural and functional constraints. The mind needs to achieve certain functions, and the brain must devote local regions to achieving these functions. This implies that if these modules reflect the correct division of the functions of the mind, it should be possible to find brain regions that reflect their activity. Our lab has developed a mapping of the eight modules . . . onto specific brain regions . . . (J. R. Anderson 2007, p. 74)

Given that neural reuse implies that anatomical modularity is false (see sect. 2.1), success in assigning ACT-R modules to specific brain regions would seem to be a problem for neural reuse, and evidence for neural reuse would appear to create problems for ACT-R. But the conclusion does not follow quite so easily as it seems. First, ACT-R does not strictly imply anatomical modularity. ACT-R is committed to the existence of *functional* modules, and to the existence of elements of the brain that implement them. If it turned out that activity in the ACT-R goal module was a better fit to the coordinated activity of some non-contiguous set of small brain regions than it was to the anterior cingulate (to which they currently have the goal module mapped), then this would count as *progress* for ACT-R, and not a theoretical setback. Similarly, if it turned out that some of the brain regions that help implement the goal module *also* help implement the imaginal module, this would pose no direct challenge to ACT-R theory.[7]

Therefore, although J. R. Anderson is at pains to deny he is a functionalist – not just any possible mapping of function to structure will count as a success for ACT-R – there is a good deal of room here for alternatives to the simple 1 : 1 mapping that he and other ACT-R theorists are currently exploring. For its part, neural reuse predicts that the best fit for ACT-R modules, or any other high-level functional components, is much more likely to be some cooperating complex of multiple brain regions than it is a single area, and that brain regions involved in implementing one ACT-R function are likely to be involved in implementing others as well. Interestingly, this is more or less what J. R. Anderson et al. (2007) found. For every task manipulation in their study, they found several brain regions that appeared to be implicated. And every one of their regions of interest was affected by more than one factor manipulated in their experiment. Thus, despite their methodological commitment to a 1:1 mapping between modules and brain regions, J. R. Anderson et al. (2007) are quite aware of the limitations of that approach:

> Some qualifications need to be made to make it clear that we are not proposing a one-to-one mapping between the eight regions at the eight functions. First, other regions also serve these functions. Many areas are involved in vision and the fusiform gyrus has just proven to be the most useful to monitor. Similarly, many regions have been shown to be involved in retrieval, particularly the hippocampus. The prefrontal region is just the easiest to identify and seems to afford the best signal-to-noise ratio. Equally, we are not claiming these regions only serve one function. This paper has found some evidence for multiple functions. For instance, the motor regions are involved in rehearsal as well as external action. (J. R. Anderson et al. 2007, pp. 213–14)

Here, the regulative idealization promoted by decomposition and localization may have unduly limited the sorts of methodological and inferential tools that they initially brought to bear on the project. As noted already in section 1, one of the contributions neural reuse may be able to make to cognitive science is an alternate idealization that can help guide both experimental design and the interpretation of results (M. L. Anderson et al. 2010).

Going forward there is at least one other area where we can expect theories of neural reuse and modular theories like ACT-R to have significant, bidirectional critical contact. Right now, ACT-R is not just theoretically, but also literally modular: it is implemented as a set of independent and separately modifiable software components. It does not appear, however, that separate modifiability is theoretically essential to ACT-R (although it is no doubt a programming convenience). Therefore, implementing overlaps in ACT-R components in light of the evidence from neuroimaging and other studies of the sort recounted here is likely to offer scientific opportunities to both research communities (see Stewart & West 2007 for one such effort). For example, overlaps in implementation might offer a natural explanation and a convenient model for certain observed instances of cognitive interference, such as that between language and motor control (Glenberg & Kaschak 2002) or between memory and audition (Baddeley & Hitch 1974), helping to refine current hypotheses regarding the causes of the interference.

The ACT-R community is already investigating similar cases, where different concurrent tasks (dialing the phone while driving) require the use of the same ACT-R module, and thus induce performance losses (Salvucci 2005). Altering ACT-R so that different modules share component parts might enable it to model some cognitive phenomena that would otherwise prove more difficult or perhaps impossible in the current system, such as the observation that object manipulation can improve reading comprehension (Glenberg et al. 2007). Finally, observations of interference in a modified ACT-R but *not* in human data, might suggest that the ACT-R modules did not yet reflect the correct division of the mind's functions. Such conflicts between model and data could be leveraged to help ACT-R better approximate the high-level functional structure of the mind.

3.3. Classic parallel distributed processing

It is of course true that from a sufficiently abstract perspective, the idea of neural reuse in cognitive functioning is nothing new. It has been a staple of debates on brain architecture at least since the advent of parallel distributed processing (PDP) models of computation (Rummelhart & McClelland 1986). For one widely cited example, consider the following from Mesulam (1990). He writes:

> A central feature of networks is the absence of a one-to-one correspondence among anatomical site, neural computation and complex behavior . . . Figure [2] implies that each behavior is represented in multiple sites and that each site subserves multiple behaviors, leading to a distributed and interactive but also coarse and degenerate (one-to-many and many-to-one) mapping of anatomical substrate onto neural computation and computation onto behavior. This distributed and degenerate mapping may provide an advantage for computing complex and rapid cognitive operations and sets the network approach sharply apart from theories that postulate a nondegenerate one-to-one relationship between behavior and anatomical site. (Mesulam 1990, pp. 601–602)

Broadly speaking, neural reuse theories are one of a family of network approaches to understanding the operation of the brain. They share with these an emphasis on cooperative interactions, and an insistence on a non-modular, many-to-many relationship between neural-anatomical sites and complex cognitive functions/behaviors. But there are also some important differences that set neural reuse apart.
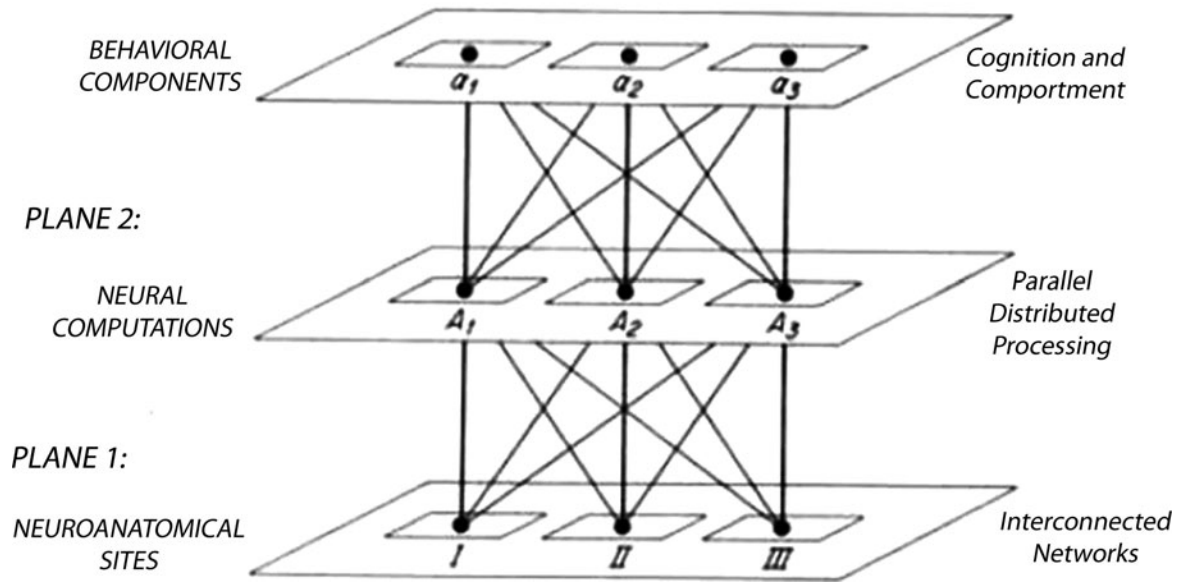
Figure 2. Detail of Figure 3 from Mesulam (1990). Reprinted with permission of the author.

First is a better appreciation of the computational work that can be done by very small groups of, or even individual, neurons (Koch & Segev 2000). Neural reuse theories all agree that most of the interesting cognitive work is done at higher levels of organization, but they also emphasize that local circuits have specific and identifiable functional biases. In general, these models make a strong distinction between a "working" – whatever specific computational contribution local anatomical circuits make to overall function – and a "use," the cognitive purpose to which the working is put in any individual case. For neural reuse theories, anatomical sites have a fixed working, but many different uses.

In contrast, note that in Figure 2 "neural computations" are located at Plane 2, parallel distributed processing. This reflects the belief that computational work can only be done by fairly large numbers of neurons, and that responsibility for this work can only be assigned to the network as a whole. Put differently, on PDP models there are no local workings. Classic PDP models are indeed a powerful way to understand the flexibility of the brain, given its reliance on relatively simple, relatively similar, individual elements. But the trouble for PDP models in this particular case is that there is no natural explanation for the data on increasing scatter of recently evolved functions, nor for the data on the cross-cultural invariance in the anatomical locations of acquired practices (see sect. 6.3). Indeed, on PDP models, investigating such matters is not even a natural empirical avenue to take. This represents a significant distinction between PDP and neural reuse.

Other important differences between neural reuse and classic PDP models flow from the above considerations, including the way neural reuse integrates the story about the cognitive architecture of the brain into a natural story about the *evolution* and *development* of the brain. In a sense, neural reuse theories make some more *specific* claims than generalized PDP – not just that the brain is a kind of network, but that it is a kind of network with functional organization at more levels than previously thought.

As can be seen already in the evidence outlined above, and will be seen in greater detail in sections 5 and 6, this specificity has led to some interesting and empirically testable implications for the brain's overall functional organization.

3.4. Contemporary parallel distributed processing models

More contemporary versions of network models, such as Leabra (O'Reilly 1998; O'Reilly & Munakata 2000) tend to be composed of densely connected, locally specialized networks that are sparsely connected to one another (see Fig. 3).
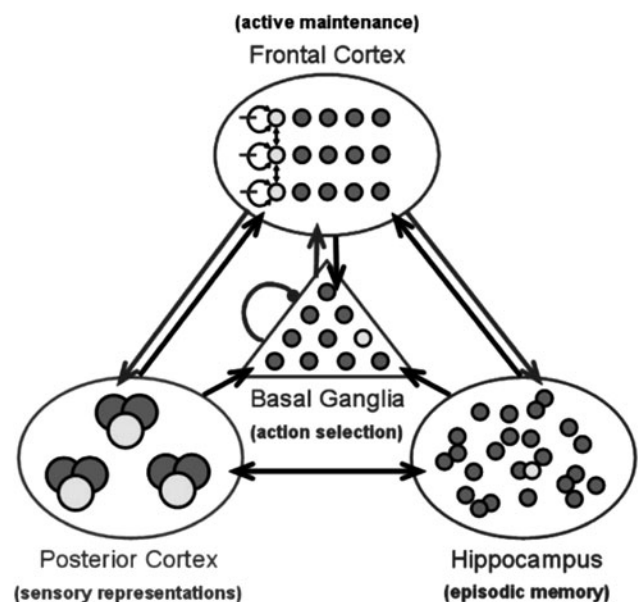

Figure 3. Overview of the Leabra architectural organization. Reprinted from Jilk et al. (2008) with permission of the authors.

In one sense, Leabra appears to be more compatible with neural reuse than classic PDP models are, as Leabra explicitly allows for regional functional biases. But insofar as this new architecture reflects the influence of the selectivity assumption, and represents a more modularist approach to understanding the brain, then there are potentially the same points of conflict with Leabra as there are with those theories. Consider the following, from a recent paper describing Leabra:

> The brain is not a homogenous organ: different brain areas clearly have some degree of specialized function. There have been many attempts to specify what these functions are, based on a variety of theoretical approaches and data. In this paper, we summarize our approach to this problem, which is based on the logic of *computational tradeoffs* in neural network models of brain areas. The core idea behind this approach is that different brain areas are specialized to satisfy fundamental tradeoffs in the way that neural systems perform different kinds of learning and memory tasks. (Atallah et al. 2004, p. 253)

There is nothing here that explicitly commits the authors to the idea that large brain regions are dedicated to specific tasks or cognitive domains – something the data presented here throw into question – although that is certainly one possible reading of the passage. Moreover, O'Reilly (1998) tends to focus more on modeling *processes* over modeling parts, an approach that need not commit one to a specific story about how and where such processes are implemented in the brain – it needn't be the case that individual brain regions implement the processes being modeled, for instance.

And yet, O'Reilly and his collaborators *have* assigned these processes to specific regions:

> The large-scale architectural organization of Leabra includes three major brain systems: the posterior cortex, specialized for perceptual and semantic processing using slow, integrative learning; the hippocampus, specialized for rapid encoding of novel information using fast arbitrary learning; and the frontal cortex/basal ganglia complex, specialized for active and flexible maintenance of goals and other context information, which serves to control or bias processing throughout the system. (Jilk et al. 2008, p. 204)

And, in fact, the Leabra team has gone further than this by recently integrating Leabra with ACT-R to form the SAL architecture:

> When the ACT-R and Leabra research teams started to work together in 2006, they came to a startling realization: the two theories, despite their origins in virtually opposite paradigms (the symbolic and connectionist traditions, respectively) and widely different levels of abstraction, were remarkably similar in their view of the overall architecture of the brain. (Jilk et al. 2008, p. 205)

So it is not clear just what commitments Leabra has to modularity and localization. As with ACT-R, there doesn't seem to be anything *essential* to Leabra that would prevent it from explicitly incorporating neural reuse as one of its organizing principles. In particular, the functional specializations being ascribed to the brain regions mentioned are general enough to plausibly have many different cognitive uses, as predicted by neural reuse theories. But, as with ACT-R, more research will be needed before it becomes clear just how compatible these visions for the functional organization of the brain in fact are. The notion of neural reuse cuts across some old divisions – localization versus holism; modular versus connectionist – and whether theories falling on one or another side of each dichotomy are compatible with the notion of neural reuse will ultimately depend on how their advocates interpret the theories, and how flexible their implementations turn out to be.

4. Other theories predicting forms of neural reuse

One of the most successful theoretical paradigms in cognitive science has been the conceptual metaphor theories originating with Lakoff and Johnson (1980; 1999) and extended by many others, perhaps most notably Fauconnier and Turner (2002).[8] As is well known, conceptual metaphor theories suggest that cognition is dominated by metaphor-based thinking, whereby the structure and logical protocols of one or more domains, combined in various ways, guide or structure thinking in another. For a simple case, consider the *Love Is War* mapping taken from Lakoff and Johnson (1980; 1999). When employing this metaphorical mapping, people use their understanding of war – of how to interpret events and how to respond to them – to guide their thinking about love: One fights for a partner, makes advances, fends off suitors, or embarks on a series of conquests. Similarly, the *Life Is a Journey* mapping allows people to leverage their extensive experience and competence in navigating the physical world in order to facilitate planning for life more generally: We plan a route, overcome obstacles, set goals, and reach milestones. The theory has been widely discussed and tested, and enjoys a raft of supporting evidence in linguistics and cognitive psychology.

A natural question that arises for such theories, however, is how the structured inheritance from one domain to another is actually achieved by the brain. Is it done abstractly, such that *mental models* (Gentner & Stevens 1983; Johnson-Laird 1983) of war or navigation are used as prototypes for building other models of love or life? Or is there a more basic biological grounding, such that the very neural substrates used in supporting cognition in one domain are reused to support cognition in the other? Although some researchers favor the first possibility – notably Lera Boroditsky (e.g., Boroditsky & Ramscar 2002) – it seems fair to say that the greater effort has been focused on investigating the second.

This is at least in part because the debate over the biological basis of conceptual metaphors dovetails with another over the nature and content of cognitive representations – symbols, concepts, and (other) vehicles of thought – that has also played out over the last twenty years or so. At issue here is the degree to which the vehicles of thought – our mental carriers of meaning – are tied to sensory experience (Barsalou 2008; 1999). Concept empiricists (as they are called in philosophy) or supporters of modal theories of content (as they are called in psychology) are generally committed to some version of the thesis that "the vehicles of thought are re-activated perceptual representations" (Weiskopf 2007, p. 156). As one of the core statements of the modal position puts it, perceptual symbols, which "constitute the representations that underlie cognition," are "record[s] of the neural activation that arises during perception" (Barsalou 1999, pp. 578, 583; see Prinz 2002 for a general discussion). This position is meant to contrast

with a rationalist or amodal one in which the vehicles of thought are inherently nonperceptual, abstract, logical, linguistic, or computational structures for which (as the classic semiotics line goes) the relation between signifier and signified is established arbitrarily (see, e.g., Fodor 1975; Fodor & Pylyshyn 1988).

In the case of both debates, it looked as if information about what neural resources were actually deployed to support cognitive tasks could provide evidence favoring one side or another. If planning a task used brain regions different from those used in planning (or imagining) a journey, then this would be prima facie evidence against the notion that the two were related via direct neural grounding. Similarly, if perceptual tasks and cognitive tasks appeared to be handled by distinct brain regions, this would appear to favor the amodal view.

In the event, a series of early findings bolstered the case for modal concepts, on the one hand, and for the idea that direct neural substrates supported metaphorical mappings, on the other. For example, a series of papers from the labs of Antonio Damasio and Alex Martin offered evidence that verb retrieval tasks activated brain areas involved in motor control functions, and naming colors and animals (that is, processing nouns) activated brain regions associated with visual processing (Damasio & Tranel 1993; Damasio et al. 1996; Martin et al. 1995; 1996; 2000). Similarly, it was discovered that perceiving manipulable artifacts, or even just seeing their names, activates brain regions associated with grasping (Chao & Martin 2000). All this suggested that class concepts like HAMMER, RED, and DOG might be stored using a sensory and/or motor code, and, more generally, that high-level, conceptual-linguistic understanding might involve the reactivation of perceptuomotor experiences. This dovetailed nicely with the general idea behind direct neural support for metaphorical mappings, whereby understanding in one domain would involve the reactivation of neural structures used for another. Thus, the finding that mental planning can activate motor areas even when the task to be planned itself involves no motor activity (Dagher et al. 1999) has long been taken to support the case that mappings like *Life Is a Journey* are mediated by the direct sharing of neural resources by both domains.[9]

It seems fair to say that these early discoveries prompted a much larger effort to uncover the neural underpinnings of high-level cognitive functions, one specifically focused on revealing the ways in which these underpinnings were shared with those of the sensorimotor system. The result is literally hundreds of studies detailing the various ways in which neural substrates are shared between various cognitive functions. A representative sample of these studies will be reviewed further on in sections 4.1 through 4.6, but to presage the argument to follow: The effort to uncover instances of neural reuse has been *so* successful that even a cursory examination of the breadth and frequency of reuse suggests that there is much more reuse than can be accounted for by modal concepts or conceptual metaphor theory. Any explanation of the phenomenon must therefore articulate a broader framework within which the prevalence of reuse naturally fits, and which in turn can explain such individual cognitive phenomena.[10] We will review some of the evidence for this claim in the next subsections.

4.1. Reuse of motor control circuits for language

A great deal of the effort to discover the specific neural underpinnings of higher cognitive functions has focused on the involvement of circuits long associated with motor control functions. In a typical example of this sort of investigation, Pulvermüller (2005) reports that listening to the words "lick," "pick," and "kick" activates successively more dorsal regions of primary motor cortex (M1). The finding is consistent both with the idea that comprehending these verbs relies on this motor activation, insofar as the concepts are stored in a motoric code, and also with the idea that understanding these verbs might involve (partial) simulations of the related actions. Either interpretation could easily be used as part of the case for concept empiricism.

Similarly, Glenberg and Kaschak (2002) uncover an interesting instance of the entanglement of language and action that they call the "action-sentence compatibility effect" (ACE). Participants are asked to judge whether a sentence makes sense or not and to respond by pressing a button, which requires a move either toward or away from their body. In one condition "yes" is away and "no" is toward; another condition reverses this. The sentences of interest describe various actions that would *also* require movement toward or away, as in "put a grape in your mouth," "close the drawer," or "you gave the paper to him." The main finding is of an interaction between the two conditions, such that it takes longer to respond that the sentence makes sense when the action described runs counter to the required response motion. More striking, this was true even when the sentences described abstract transfers, such as "he sold his house to you," which imply a direction without describing a directional motor action. Following the reasoning originally laid out by Sternberg (1969), an interaction between two manipulated factors implies at least one shared component between these two different processes – movement and comprehension. A likely candidate for this component would be a neural circuit involved in motor control, a supposition confirmed by Glenberg (2008b).[11] Thus, this seems another clear case in which motor control circuits are involved in, and perhaps even required for, language comprehension, whether via simulation (e.g., in the concrete transfer cases), metaphorical mapping (e.g., in the abstract transfer cases), or by some other mechanism. Glenberg has suggested both that the effect could be explained by the activation of relevant action schemas (Glenberg et al. 2008b) and by the activation and combination of appropriate affordances (Glenberg & Kaschak 2002; Glenberg et al. 2009). Whatever the precise mechanism involved, the finding has been widely interpreted as support for both concept empiricism and for conceptual metaphor theory (although see M. L. Anderson 2008c for a dissent).

4.2. Reuse of motor control circuits for memory

Another interesting description of the motor system's involvement in a different cognitive domain comes from Casasanto and Dijkstra (2010), who describe bidirectional influence between motor control and autobiographical memory. In their experiment, participants were asked to retell memories with either positive or negative valence, while moving marbles either upward or downward from

one container to another. Casasanto and Dijkstra found that participants retrieved more memories and moved marbles more quickly when the direction of movement was congruent with the valence of the memory (upward for positive memories, downward for negative memories). Similarly, when participants were asked simply to relate some memories, without prompting for valence, they retrieved more positive memories when instructed to move marbles up, and more negative memories when instructed to move them down. Because the effect is mediated by a mapping of emotional valence on a spatial schema, the finding seems to most naturally support conceptual metaphor theory. The fact that the effect was bidirectional – recounting memories affected movement *and* movement affected memory retrieval – is a striking detail that seems to suggest direct neural support for the mapping.[12]

4.3. Reuse of circuits mediated by spatial cognition

Many of the apparent overlaps between higher-order cognition and sensorimotor systems appear to be mediated by spatial schemas in this way. For example, Richardson et al. (2003) report that verbs are associated with meaning-specific spatial schemas. Verbs like "hope" and "respect" activate vertical schemas, whereas verbs like "push" and "argue" activate horizontal ones. As the authors put it, "language recruits spatial representations during real-time comprehension." In a similar vein, Casasanto and Boroditsky (2008) suggest that our mental representations of time are built upon the foundations of our experience with space. These findings appear to provide strong and relatively unproblematic support for conceptual metaphor theory, and perhaps also for a generic theory of concept empiricism, according to which the content of our concepts is grounded in (but does not necessarily constitute a simulation or reactivation of) sensorimotor experiences.

On the other hand, even when simulation *is* an important aspect of the reuse of resources between different domains, it does not always play the functional role assigned it by concept empiricism or conceptual metaphor theory. For some time, there has been growing evidence that doing actions, imagining actions, and watching actions done by others all activated similar networks of brain regions (Decety et al. 1990; Decety et al. 1997; Jeannerod 1994). This has suggested to many that social cognition – understanding the actions and intentions of others – could involve simulating our own behaviors, a notion that attracted even more widespread interest after the discovery of mirror neurons (Decety & Grèzes 1999; Gallese et al. 1996; Gallese & Goldman 1998; Rizzolati et al. 1996). The trouble for concept empiricism and conceptual metaphor theory is that the logic governing the reuse of resources for multiple purposes is quite different in this case. Here, the idea is that circuits associated with behavioral control can be used to build predictive models of others, by inputting information about another agent into the system that would normally be used to guide one's own actions (and reactions). Although it could be argued that using simulation in support of such "mind-reading" (Gallese & Goldman 1998) requires a kind of metaphorical mapping (he is *like me* in relevant ways), in fact this is simply a necessary assumption to make the strategy sensible, and does *not* play the role of a domain-structuring inheritance.

Even some of the evidence for the reuse of *spatial* operations in other cognitive domains – which has been a mainstay of research into concept empiricism and conceptual metaphor theory – suggests the existence of more kinds of reuse than can be accounted for by these theoretical frameworks. Consider just a few of the various manifestations of the spatial-numerical association of response codes (SNARC) effect (Dehaene et al. 1993): (1) When participants are asked to judge whether numbers are even or odd, responses are quicker for large numbers when made on the right side of space (canonically with the right hand, although the effect remains if responses are made while hands are crossed) and quicker for smaller numbers when responses are made on the left side of space. (2) Participants can accurately indicate the midpoint of a line segment when it is composed of neutral stimuli (e.g., *XXXXX*), but are biased to the left when the line is composed of small numbers (e.g., *22222* or *twotwotwo*) and to the right when the line is composed of large numbers (e.g., *99999* or *nineninenine*). (3) The presentation of a number at the fixation point prior to a target detection task will speed detection on the right for large numbers and to the left for small numbers. Hubbard et al. (2005) hypothesize that the SNARC effect can be accounted for by the observed reuse in numeric cognition of a particular circuit in left inferior parietal sulcus that plays a role in shifting spatial attention. Briefly, the idea is that among the representational formats we make use of in numerical cognition there is a mental "number line," on which magnitudes are arrayed from left to right in order of increasing size. Once numerals are arrayed in this format, it is natural to reuse the circuit responsible for shifting spatial attention for the purpose of shifting attention between positions on this line. The resulting magnitude-influenced attentional bias can explain the SNARC effect.

This redeployment of visuo-spatial resources in support of alternate cognitive uses is somewhat difficult to explain from the standpoint of either concept empiricism or conceptual metaphor theory. In these examples, the effects would not be accounted for by the fact that numbers might be grounded in or involve simulations of basic sensorimotor experience, nor is it immediately obvious what metaphorical mapping might be implicated here. In fact, if the reuse of spatial schemas *were* in support of some semantically grounding structural inheritance from one domain to the other, we would expect the numbers to be arrayed *vertically*, with magnitude increasing with height. Instead, the reuse in this case appears driven by more abstract functional considerations. When doing certain numerical tasks, a number line is a useful representational format, and something like the visuo-spatial sketchpad (Baddeley 1986) offers a convenient and functionally adequate storage medium. Similarly, reusing the spatial shifting mechanism is a sensible choice for meeting the functional requirements of the task, and need not ground any semantic or structural inheritance between the domains.

4.4. Reuse of circuits for numerical cognition

In fact, several such examples can be found in the domain of numerical cognition. Zago et al. (2001) found increased activation in the premotor strip in a region implicated in

finger representation during multiplication performance compared to a digit reading condition. Similar findings were reported by Andres et al. (2007), who found that hand motor circuits were activated during adults' number processing in a dot counting task. That these activations play a functional role in both domains was confirmed by Roux et al. (2003), who found that direct cortical stimulation of a site in the left angular gyrus produced both acalculia and finger agnosia (a disruption of finger awareness), and by Rusconi et al. (2005), who found that repetitive Transcranial Magnetic Stimulation (rTMS) over the left angular gyrus disrupted both magnitude comparison and finger gnosis in adults.

Here again, this reuse of a basic sensorimotor function in an alternate cognitive domain does not seem to follow the logic of conceptual metaphor theory or concept empiricism. These theories are not making the claim that magnitudes inherit their meanings from finger representations, nor is any mathematical metaphor built in any straightforward way on our finger sense. Rather, the idea is that this neural circuit, originally developed to support finger awareness, is offering some functionally relevant resource in the domain of numerical cognition. For instance, Butterworth (1999c) suggests that the fingers provide children a useful *physical* resource for counting, with the neural result that the supporting circuits now overlap, while Penner-Wilger and Anderson (2008; submitted) suggest instead that the circuit in question might *itself* offer useful representational resources (such as a storage array).[13] This is not to question the notion that mathematical concepts and procedures are in some way grounded in sensorimotor experience (Lakoff & Núñez 2000), but this specific overlap in neural circuitry isn't straightforward to explain in the context of such grounding, nor is it anything that would have been *predicted* on the basis of either conceptual metaphor theory or concept empiricism. In fact, proponents of conceptual metaphor theory in mathematics tend to focus on relatively higher-level concepts like sets and investigate how our understanding of them is informed by such things as our experience with physical containers.

A similar argument can be made when considering the interrelations of speech and gesture, and the cognitive importance of the latter (see, e.g., Goldin-Meadow 2003). According to Goldin-Meadow (2003), gesture is typically used not just to signal different moments in the learning process (e.g., to index moments of decision or reconsideration in a problem-solving routine), but also appears to have utility in *advancing* the learning process by providing another representational format that might facilitate the expression of ideas currently unsuited (for whatever reason) to verbal expression. The motor control system is here being used for a specific cognitive purpose *not* because it is performing semantic grounding or providing metaphorically guided domain structuring, but because it offers an appropriate physical (and spatio-temporal) resource for the task.

4.5. Reuse of perceptual circuits to support higher-order cognition

There are examples of the reuse of circuits typically associated with perception that also make the same point. Although there have certainly been studies that appear

to unproblematically support concept empiricism – for example, Simmons et al. (2007) report the discovery of a common neural substrate for seeing colors, and for knowing about (having concepts for) color – other studies suggest that such cases represent only a small subset of a much broader phenomenon. Consider one of the earliest and most discussed cases of the reuse of neural circuits for a new purpose, the Baddeley and Hitch model of working memory (Baddeley & Hitch 1974; 1994; Baddeley 1986; 1995). One strategy for remembering the items on a grocery list or the individual numbers in a phone number involves (silently) saying them to one's self (producing a "phonological loop"), which engages brain areas typically used both in speech production and in audition.

A pattern of findings supports the existence of a phonological loop, and the engagement of both inner "speaking" and inner "hearing" to support working memory (see Wilson 2001 for a review). First, there is poor recall of similar sounding terms; second, there is poor recall of longer words; third, there is poor recall if the subject is made to speak during the maintenance period; and fourth, there is poor recall when the subject is exposed to irrelevant speech during the maintenance period. Moreover, imaging studies have found that such memory tasks cause activation in areas typically involved in speech production (Broca's area, left premotor cortex, left supplementary motor cortex, and right cerebellum) and in phonological storage (left posterior parietal cortex) (Awh et al. 1996).

In this interesting and complicated case, we have something of a triple borrowing of resources. First is the use of a culturally specific, acquired representational system – language – as a coding resource, and second is the application of a particular learned skill – silent inner speech – as a storage medium. These two together imply the third borrowing – of the neural resources used to support the first two functions. And note that all of this borrowing is done in support of what is likely an enhancement of a basic evolved function for storing small amounts of information over short periods. This raises the obvious question of whether and to what degree evolutionary pressures might have shaped the language system so that it was capable of just this sort of more general cognitive enhancement (Carruthers 2002). In any case, it seems clear that this sort of borrowing is very hard to explain in terms of concept empiricism or conceptual metaphor theory. In the case of sensorimotor coding in working memory, the phonological loop is not metaphorically *like* speech; rather, it *is* a form of speech. In this, it is another instance of a straightforward functional redeployment – the reuse of a system for something other than its (apparent) primary purpose because it happens to have an appropriate functional structure.

4.6. Reuse is not always explained by conceptual metaphor theory or concept empiricism

These various examples suggest something along the following lines: One of the fundamental principles guiding reuse is the presence of a sufficient degree of functional relatedness between existing and newly developing purposes. When these functional matches result in the reuse of resources for both purposes, this history

sometimes – but not always – reveals itself in the form of a metaphorical mapping between the two task domains, and sometimes, but not always, results in the inheritance or grounding of some semantic content. This way of thinking makes conceptual metaphors and "grounded" symbols into two possible side-effects of the larger process of reuse in cognition. It also muddies the causal story a bit: Planning is like locomotion because it inherits the structure of the existing domain via neural overlap; but planning also overlaps with the neural implementation base of locomotion to the degree that it is like locomotion.

The suggestion here is not that planning or communication or any other cognitive function has some predetermined Platonic structure that entirely reverses the causal direction typically supposed by conceptual metaphor theory. Rather, the idea is to point out the need to be open to a more iterative story, whereby a cognitive function finds its "neural niche" (Iriki & Sakura 2008) in a process codetermined by the functional characteristics of existing resources, and the unfolding functional requirements of the emerging capacity (Deacon 1997).

Consider, in this regard, 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 2cm 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, p. 355)

There are certainly functional characteristics that a unit of acoustic communication must have in order to adequately perform its communicative purpose, and not just any neural substrate would have had the required characteristics. But there remain degrees of freedom in how those characteristics are implemented. Speech production, then, developed its *specific* phonemic character as the result of the circuits on which it was built. Had the motor control system been oriented instead around (for example) simple, repeatable contractions of individual muscles – or had there been some other system with these functional characteristics available for reuse as acoustic communication was evolving – the result of the inheritance might have been a communication code built of more purely temporal elements, something closer to Morse code.[14]

Finally, consider what may be a case *not* of the reuse of a basic sensorimotor area for higher cognitive functions, but rather the reverse. Broca's area has long been associated with language processing, responsible for phonological processing and language production, but what has recently begun to emerge is its functional complexity (Hagoort

2005; Tettamanti & Weniger 2006). 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 (Decety et al. 1997; Hamzei et al. 2003; Nishitani et al. 2005), imagery of human motion (Binkofski et al. 2000), and action imitation (Nishitani et al. 2005). Note that Müller and Basho (2004) suggest that these functional overlaps should not be understood as the later reuse of a *linguistic* area for other purposes, but are rather evidence that Broca's area already performed some sensorimotor functions that were prerequisites for language acquisition, and which made it a candidate for one of the neural building blocks of language when it emerged. That seems reasonable; but on the other hand, Broca's area is also activated in domains such as music perception (Tettamanti & Weniger 2006). While it is possible that this is because processing music requires some of the same basic sensorimotor capacities as processing language, it seems also possible that this reuse was driven by functional features that Broca's acquired as the *result* of its reuse in the language system, and thus by some more specific structural similarity between language and music (Fedorenko et al. 2009). Whatever the right history, this clearly represents another set of cases of functional reuse *not* explained by conceptual metaphor theory or concept empiricism.

Assuming the foregoing is sufficient to establish the existence of at least some cases of neural reuse that cannot be accounted for by these theoretical frameworks alone, the question naturally arises as to whether these anomalous cases should be dealt with by post-hoc elaborations of these theories (and/or by generating one or a few similarly specific theories), or whether this is a situation that calls for a global theory of reuse that supersedes and at least partially subsumes these existing frameworks. Far be it from me to argue a priori that one tack *must* be the correct one to take – science works best when we pursue multiple competing research paths – but one thing it might be useful to know when deciding how to spend one's research time is exactly how widespread neural reuse *is*. That is, the more widespread reuse appears, and the more instances of reuse that can be identified that do not involve the sensorimotor system, the stronger the justification would seem for trying to formulate a more global theory of neural reuse.

5. Further evidence that neural reuse is a pervasive feature of brain organization

Given the success of the theoretical frameworks just mentioned, as well as the growing interest in embodied cognition (M. L. Anderson 2003; Chemero 2009; Clark 1997; 1998), it is quite easy to find studies reporting that the neural implementations of higher cognitive functions overlap with those of the sensorimotor system. Indeed, this was the theme of a recent Attention and Performance Symposium, culminating in the 27-essay volume *Sensorimotor Foundations of Higher Cognition* (Haggard et al. 2008). In contrast, there are only a few examples of reuse not involving the sensorimotor system that are reported as such in the literature. This fact would seem to favor the post-hoc elaboration approach to explaining

the sorts of cases outlined above. On the other hand, the lack of such reports could simply be because people are not looking in the right place, or looking in the right way; after all, nobody is trying to establish a theory of attention-grounded, mathematics-grounded, or music-grounded cognition (as interesting as that sounds!). Absence of evidence of these cases, this is to say, is not evidence of absence. A typical literature search, then, will not help answer our question.

The literature can, however, be used in a somewhat different way. There are many, many thousands of studies in the neuroimaging literature that purport to uncover the neural underpinnings of various cognitive functions. If one were to compile a number of these studies in various task domains, one could ask, for each region of the brain, whether it supported functions in multiple domains, and whether such reuse was typically limited to regions of the brain implicated in supporting sensorimotor tasks.

The NICAM database (M. L. Anderson et al. 2010) currently contains information from 2,603 fMRI studies reported in 824 journal articles. All the studies involve healthy adults and use a within-subjects, subtraction-based, whole-brain design. That is, for all the studies in the database, brain activity during an experimental task was observed over the whole brain (not just a region of interest), and then compared to and subtracted from activity observed in the same participant during a control task. The logic of subtraction method is such that it should uncover *only* the regions of activation that support the specific mental function that best captures the difference between the experimental and control task. The neural activations supporting the mental operation that the two tasks have in common – the visual process allowing one to see the stimuli in a language task, for example – should be subtracted out. The database lists, among other things, the locations of the 21,553 post-subtraction fMRI activations observed during those 2,603 studies – that is, the regions of activation that are purported to specifically support those 2,603 mental operations. These features make the database ideal for investigating whether and to what degree specific brain regions support multiple functions across various task domains.

The general methodology for this sort of study is simple and straightforward. First, choose a spatial subdivision of the brain, then see which experiments, in which (and how many) domains, showed activity in each of the regions. To get the results reported in the next paragraph, below, I used the same 998 anatomical regions of interest (ROIs) used by Hagmann et al (2008).[15] The study was restricted to the following eleven task domains: three action domains – execution, inhibition, and observation – two perceptual domains – vision and audition – and six "cognitive" domains – attention, emotion, language, mathematics, memory, and reasoning.[16] Any study that was assigned to more than one domain was excluded. Activations were assigned to the ROI with the closest center; any activation that was more than 13mm from the center of one of the ROIs was excluded. This left 1,469 experiments collectively reporting 10,701 eligible activations.[17]

There were 968 regions that were active in at least one experiment (and thus in one domain). Of these, 889 (91.8%) were active in at least two domains – that is, were reused at least once. On average, these 968 regions were active in 4.32 different domains (SD 1.99), and 555

of the regions were active in action tasks, with 535 of these "action" areas also active in an average of 3.97 (SD 1.58) non-action domains, and 530 active in an average of 3.16 (SD 1.23) cognitive domains. There were 565 regions active in perception tasks; 555 of these "perception" regions were also active in an average of 4.00 (SD 1.61) non-perception domains, and 550 were active in an average of 3.20 (SD 1.24) cognitive domains. There were 348 regions active in *both* action and perception tasks. On average, these were reused in 3.33 (SD 1.22) cognitive domains. There were also 196 regions not active in either perception or action tasks; 143 of these (72.96%) were active in two or more domains and averaged 2.97 (SD 0.95) domains. With all 196 regions included, the average is 2.43 (SD 1.19) of the six cognitive domains.[18]

Naturally, if one uses larger regions – for instance, the 66 cortical ROIs[19] used by Hagmann et al (2008) – the average amount of reuse increases accordingly. All 66 regions were active in at least one domain; 65 (98.5%) were active in two or more domains.[20] As noted already above, the 66 regions were active in an average of 9.09 (SD 2.27) different domains. The 60 regions active in action tasks were also active in an average 7.38 (SD 0.98) non-action domains and 5.5 (SD 0.81) cognitive domains. The 64 regions active in perception tasks were also active in 7.39 (SD 1.87) non-perceptual domains and 5.34 cognitive domains. The 59 regions active in both perception and action tasks were also active in an average of 5.53 (SD 0.80) other domains, and the 7 regions not active in both perception and action tasks were active in an average of 3.00 (SD 1.41) of the cognitive domains. Only one region was active in only cognitive tasks, and that region was active only in memory.

These data appear to support the following claims: (1) Regions of the brain – even fairly small regions – are typically reused in multiple domains. (2) If a region is involved in perception tasks, action tasks, or both, it is more likely to be reused than if it is not involved in such tasks.[21] (3) Regions not involved in such tasks are nevertheless more likely than not to be reused in multiple domains. Note that the way of counting adopted above makes the best possible case for the "action and perception are special" position, by classifying as an "action" or "perception" region every region that is active in any such task. But it seems unlikely that there are 60 large cortical "action areas" and 64 "perception areas" in the way this term is usually understood. If instead some of these regions in fact contain instances of the reuse of "cognitive" circuits[22] for action or perception tasks, then this way of counting likely overestimates the relatively higher reuse frequency of action and perception circuits. That is, neural reuse appears to be a pervasive feature of the functional organization of the brain, and although circuits that support action and perception may be favored targets for reuse, reuse is by no means restricted to sensorimotor circuits. Therefore, the situation appears to call for an assimilative, global theory, rather than the elaboration of existing theoretical frameworks.

6. Global theories of neural reuse

As mentioned at the outset, there are currently four candidates for a broad, general theory of neural reuse (or for the

core around which such a theory could be built): Gallese's neural exploitation hypothesis, Hurley's shared circuit model, Dehaene's neuronal recycling hypothesis, and my massive redeployment hypothesis (already outlined in sect. 1.1 of this article). In this section, I will discuss each theory in turn and explore some of their similarities and differences.

6.1. Neural exploitation hypothesis

The neural exploitation hypothesis is a direct outgrowth of conceptual metaphor theory and embodied cognition, and largely sits at the intersection of these two frameworks. The main claim of the framework is that "a key aspect of human cognition is ... the adaptation of sensory-motor brain mechanisms to serve new roles in reason and language, while retaining their original function as well." (Gallese & Lakoff 2005, p. 456) This claim is the conclusion of an argument about the requirements of understanding that runs roughly as follows:

1. *Understanding requires imagination*. In the example most extensively developed by Gallese and Lakoff (2005), understanding a sentence like "He grasped the cup" requires the capacity to imagine its constituent *parameters*, which include the agent, the object, the action, its manner, and so on.

2. *Imagination is simulation*. Here, the neural exploitation hypothesis dovetails with concept empiricism in arguing that calling to mind individuals, objects, actions, and the like involves reactivating the traces left by perceiving, doing, or otherwise experiencing instances of the thing in question.

3. *Simulation is therefore neural reuse*. Simulation involves reuse of the same *functional clusters* of cooperating neural circuits used in the original experience(s).

As much of the evidence for these claims has been laid out already in earlier sections, it won't be recounted here. The reader will of course notice that the theory as stated is limited to the adaptation of sensorimotor circuits, and we have already seen that reuse in the brain is much more broad-based than this. This is indeed a drawback of the theory, but it is nevertheless included here for two reasons: first, because it has been expanded to include not just the case of concept understanding, but also of human social understanding (Gallese 2008); and, second, because it incorporates a detailed computational model for how the reuse of circuitry might actually occur, based on work by Feldman and Narayanan (2004). This model has broader applicability than is evidenced in the two main statements of the neural exploitation hypothesis (Gallese 2008; Gallese & Lakoff 2005).

The core of the computational model is a set of *schemas*, which are essentially collections of features in two layers: *descriptions* of objects and events and *instructions* regarding them. These two layers are systematically related to one another and to the sensorimotor system, such that event schemas can be used both to recognize events and to guide their execution, and object schemas can be used both to recognize objects and also to guide actions with respect to them.[23.] The schemas are also connected to the conceptual system, such that the contents of our concepts are built from the same features that form the schemas. The general idea is that the features' connections to the sensorimotor system give semantic substance to the concepts, as well as a natural model for understanding as the activation of neurally (or, in the current case, neural-network-ly) instantiated features and schemas. Like Gallese and Lakoff (2005), Feldman and Narayanan (2004) focus primarily on cases of understanding that can be directly ("He grabbed the cup") or metaphorically ("He grabbed the opportunity") mapped to basic perception-action domains.

But there is no reason in principle that the model need be limited in that way. As the authors note, by adding layers of abstraction, one can move from concrete action execution plans to abstract recipes like mathematical algorithms. Given this flexibility, it seems that action schemas need not be limited to providing guidance for the manipulation of independent objects (whether concrete or abstract) but could presumably also become control systems for the manipulation of neural circuits. That is, the same action schema that might normally be used to control rhythmic speech production, could be reused to guide silent memory rehearsal, and more abstract schemas might form the basis of control systems for predictive modeling or other applications.[24]

Of course, this emendation would constitute a significant departure from the model as originally formulated.[25.] In particular, it would turn a system in which neural reuse was driven by *grounding* – the inheritance of semantic content from one level to another – into one in which reuse was driven by the need to create control systems for functionally relevant outcomes. Although it is far from clear that this switch *precludes* the possibility that grounding plays a role in driving neural reuse, it certainly moves it from center stage, which may have undesirable theoretical consequences for the theory as a whole, and for the way it interfaces with related ideas in linguistics, philosophy, and psychology. On the other hand, without some emendation that significantly broadens the kinds of cases that it can cover, the neural exploitation hypothesis risks being inadequate to the full range of available empirical evidence. We will return to these issues when we come to our general discussion of the four candidate theories.

6.2. The shared circuits model

The shared circuits model (Hurley 2005; 2008) is organized around five control layers of similar structure, which are differentiated by the increasing abstraction of inputs and outputs. Each layer consists of an adaptive feedback loop that takes state information as input and generates control information as output. The first, lowest layer is a simple perception-action feedback loop that monitors progress toward action goals (reaching a target) and adjusts motor output in light of perceptually generated state information. It is, in this sense, a model of the simplest sort of thermostat; and the idea is that behavioral control systems might consist, at the most basic level, of multiple such control systems – or circuits. Layer 2 takes input from the external world, but also from layer 1, and becomes in essence an adaptive feedback loop monitoring the original layer. That is, layer 2 is in essence a forward model of layer 1. As is well known, incorporating such models into adaptive control systems tightens overall control by allowing for the prediction of state information, so appropriate action can be taken without waiting for the (typically slower) external feedback signal.[26.] The more

hysteresis in the system – the longer it takes control interventions to produce expected results – the more improvement forward models can offer.

Circuit *sharing* really begins with layer 3, in which the same control circuits described by layers 1 and 2 take as input observations of the actions (or situations) of other agents. Hurley's suggestion is that the mirror system (Decety & Grèzes 1999; Gallese et al. 1996; Rizzolati et al. 1996) should be modeled this way, as the activation of basic control circuits by state information relevant to the situations of other agents. Layer 3 also implements output inhibition, so agents don't automatically act as if they were in another agent's situation whenever they observe another agent doing something. Layer 4 incorporates monitoring of the output inhibition, supporting a self-other distinction; and layer 5 allows the whole system to be decoupled from actual inputs and outputs, to allow for counter-factual reasoning about possible goals and states and about what actions might follow from those assumptions. The idea is that the same circuits normally used to guide action in light of actual observations can *also* be fed *hypothetical* observations to see what actions might result; this can be the basis of predictive models. By the time we achieve the top layer, then, we have the outline for a model both of deliberation about possible actions, and also of multi-agent planning, which could serve as the basis for high-level social awareness and intelligence.

Like the neural exploitation hypothesis, one of the main explanatory targets of the shared circuits model is the possibility of mindreading and intelligent social interaction. And like the neural exploitation hypothesis, it is built entirely on the foundation of sensorimotor circuits. However, *unlike* the neural exploitation hypothesis, the shared circuits model does not revolve around the inheritance of semantic content from one level to another, but rather around the inheritance of function. The core capacities of the higher layers are based on exploiting the functional properties of the lower layers; all the layers are essentially control loops containing predictive models because they are reusing the basic implementation of the lowest levels. This is an advantage in that it is easier to see how the shared circuits model could be used to explain some of the specific instances of function-driven inheritance canvassed above; for, although Hurley models layer 1 on low-level sensorimotor circuits, there seems no reason in principle that the general approach couldn't allow for other kinds of basic circuits, on which other functional layers could be built.[27] It is also a potential weakness, in that it is less easy to see how it could be used to account for the central findings of concept empiricism or conceptual metaphor theory; can the sort of functional inheritance allowed by this model *also* allow for semantic inheritance? The inheritance of a basic feedback structure does not seem to lend itself to any obvious examples of this sort. This is not a criticism of the model as it stands – it was meant only to account for our understanding of instrumental actions; but it suggests that there is no very simple way to generalize the model to a wider set of cases. On the other hand, there seems no *fundamental* conflict between inheriting a function and thereby inheriting semantic content or domain structure.

I mentioned at the outset that the hierarchy of levels was characterized by an increasing abstraction of input and output. Consider layer 3 in this regard – at this level, input will be both impoverished and abstract as compared with lower layers. It will be impoverished because it will be missing a great deal of the richness of embodied experience – tactile experience, proprioceptive feedback, and efference copy are all absent when observing as opposed to acting. One is left with the *visual* experience of an action. And note that an action viewed from the first-person perspective looks different from the same action viewed from the third-person perspective. This points to one reason that the information must be *abstract*: since the visual experience of another agent's action will differ in most, if not all, of its low-level particulars, the system must be sensitive not to these, but to high-level *features* of the action that are common to the two situations.[28] Moreover, by hypothesis, layer 3 responds not just to actions, but to situations in which actions are possible – not just to another agent *reaching* for a banana, but to the banana being *within the reach* of another agent. This requires imputing possible goals to the observed agent, as well as encoding the high-level features of situations (relations between other agents, their capacities, and the objects in a scene). Here, the shared circuits model may need to be supplemented with something like the feature schemas from the neural exploitation model, itself expanded to allow for *situation* schemas, and not just object-action ones.

Similarly, if layer 4 is to appropriately and selectively inhibit the control outputs, it must take as input information about the relationships among the actions, agents, goals, and situations – who is in which situation doing what – which requires at least a rudimentary self/other distinction. And if layer 5 is going to be useful at all, the predictions it provides as output must be abstract, high-level action descriptions, not low-level motor commands.

These facts might seem to be nothing more than interesting and functionally useful features of the model, but in fact the requirement for abstraction at higher levels raises a puzzle: If low-level circuits respond to high-level features as inputs, and can produce high-level commands as outputs, might this not imply that layers 1 and 2 are more abstract than the model assumes? The trouble this raises is not with the coherence of the model, but with the evidence for it: All the evidence for layer 1 and 2 type controllers comes from on-line control systems dealing with real-time effector-specific, low-level feedback and control information, and *not* with abstract, feature-based information.

One obvious way to address this puzzle is to say that each layer is in fact a separate control structure that takes input from and delivers output to the layer below it, but this would undercut the entire premise of the model, since it would no longer be clear in what sense circuits were being "shared." That high-level control systems are structurally *like* low-level ones is a perfectly reasonable hypothesis, but this is not the hypothesis put forward by this model, nor is it one for which there is a great deal of biological evidence.

A different approach would be to retain the central hypothesis that control circuits are shared among layers – that layer 3 reuses the control circuit defined by layers 1 and 2, and layer 5 reuses the control circuit defined by layers 1–4 – but suggest that the inputs between layers must be mediated by translators of various kinds. That is, layer 3 takes high-level feature information and translates this into the low-level information favored by layers 1 and 2 before passing it on. Indeed, one might hypothesize it does this by reusing *other* circuits, such as those that

translate abstract plans into successive low-level motor actions. Similarly, layer 5 accepts the low-level motor commands natively output by layer 1, but translates them into high-level action descriptions. This picture is pretty plausible in the case of layer 3 observations of abstract action features, but it is much less clear how *situations* might get translated appropriately; and it is *especially* unclear how the reverse inference from low-level motor commands to high-level action descriptions might work. Just as a high-level action might be implemented any number of ways, a specific motor movement might be put in the service of innumerable high-level actions. The fact that part of the sensory information used to retroduct the action/intention from motor movement is the observed *effect* of the motor movement will help somewhat, but the basic problem still remains: There is a many-to-many relationship between movement and actions, so the valid deduction of a movement from an intention, and the valid retroduction of an intention from a movement need not follow the same paths in opposite directions.

These are hard problems to address; and because they originate from the fact that the shared circuits model requires that different kinds of inputs be fed to the same neural circuits, they may be problems that will surface for *any* theory of neural reuse (see discussion in section 6.4). Hence, it seems that the best approach to this puzzle may be to bite the bullet and say that, in at least some cases, circuit reuse is arranged such that different data – both information pertaining to different targets, as well as information about the same targets but at different levels of abstraction – can be fed *without translation* to the same circuits and still produce useful outputs.[29] Many sorting algorithms can just as easily sort letters as numbers; and if you feed a given algorithm pictures instead, it will do *something* with them. Naturally, this raises some pressing questions that seem ready-made for an enterprising theorist of neural computation: Under what conditions might useful things be done by circuits working with non-standard data? What kinds of implementations increase the chances of functionally beneficial outcomes given the fact of reuse? We will return to these issues in sections 6.4 and 7.

At its core, the shared circuits model offers an approach to understanding how high-level function could possibly be enabled by low-level circuits – and specifically by the reuse of low-level circuits for various purposes. Unfortunately, it is left fairly unclear exactly how they might *actually* be so enabled, given the different input-output requirements for each level; I have tried to sketch a solution that does the least damage to the intentions of the model, but I have to admit that some deep puzzles potentially remain.

Nevertheless, the model is interesting as an example of what might come from adopting a fairly classical "boxological" approach to cognitive modeling – understanding information processes via decomposition and interrelation – but without the underlying assumption of anatomical modularity.[30] If neural reuse is indeed a pervasive feature of the functional organization of the brain – as the current article is arguing – we will need to see more such work in the future.

6.3. The neuronal recycling hypothesis

The neuronal recycling hypothesis (Dehaene 2005; Dehaene & Cohen 2007) originates from a set of considerations rather different from those motivating the two theories just discussed (i.e., the neural exploitation hypothesis and the shared circuits model). While those are neutral on the question of how and over what time-scales the brain organization they propose came about, Dehaene is interested specifically in those cognitive capacities – such as reading and mathematics – that have emerged too recently for evolution to have generated cortical circuits specialized for these purposes. Such cultural practices must be learned, and the brain structures that support them must therefore be assigned and/or shaped during development.

There are two major ways to explain how recent cultural acquisitions, which emerge and are maintained in a population only by learning and not via genetic unfolding, can be supported by neural structures, as of course they must partly be. One way is to take our capacity to acquire such practices as reading and arithmetic as evidence for domain-general learning mechanisms (Barkow et al. 1992) and fairly unconstrained neural plasticity (Quartz & Sejnowski 1997). The other way is to suggest that cultural acquisitions must find a "neuronal niche" – a network of neural structures that already have (most of) the structure necessary to support the novel set of cognitive and physical procedures that characterize the practice. The neuronal recycling hypothesis is of the latter sort.

Note the interesting implication that the space of possible cultural acquisitions is partly constrained by cortical biases. The phrase "neuronal niche" is clearly meant to echo the idea of an ecological niche, and suggests both that acquired cognitive abilities "belong" in specific neural locations (i.e., can only survive where the neural climate is appropriate) and that the neural ecology may partly determine the characteristics that these cultural acquisitions possess, by limiting what is even possible to learn (and therefore which cognitive animals survive). Assuming the set of evolutionarily determined cortical biases is consistent across the species, we should expect to find evidence of at least three things: First, the neural manifestations of acquired abilities should be relatively consistent across individuals and even cultures; second, these practices should have some common cross-cultural characteristics; and third, the same sorts of cortical biases, as well as some of the same possibilities for learning, should be present in nonhuman primates.

As evidence for the first expectation, Dehaene and Cohen (2007) note that the visual word form area, functionally defined as a region specifically involved in the recognition and processing of written words, appears in the same location in the brain across participants, whether the participants in question are using the same language and writing system or using different ones. Similarly, the intraparietal sulcus has been implicated in numeric tasks, regardless of the culture or number representation system used by the participants. As evidence for the second expectation, they point to work by Changizi and colleagues (Changizi & Shimojo 2005; Changizi et al. 2006) that writing systems are characterized by two cross-cultural invariants: an average of three strokes per written letter; and a consistent frequency distribution for the types of contour intersections among the parts of those letters (T, Y, Z, etc.). Finally, the third expectation has been supported by some interesting and groundbreaking work by Atsushi Iriki and colleagues (Iriki 2005; Iriki &

Sakura 2008) who uncover evidence for real-time neural niche construction in primate brains (specifically *Macaca fuscata*) as the result of learning to use simple tools. The location of the observed neuro-morphological changes following tool training is roughly homologous to the regions associated with tool-use in the human brain (Culham & Valyear 2006). Thus, the theory suggests a novel pathway by which *Homo sapiens* may have achieved its current high-level cognitive capacities.

The neuronal recycling hypothesis outlines a universal developmental *process* that, although illustrated with specific examples, is meant to describe the way any acquired ability would come to have a neural instantiation. In this sense, it is broader in conception than the neural exploitation and shared circuits theories described in sections 6.1 and 6.2, respectively (although as noted their scope might well be increased with a few modifications). How much neural plasticity would be required in any given case will vary with the specifics of the acquisition, but one strength of the neuronal recycling theory is it makes clear some of the limits and costs that would be involved. The greater the distance between the function(s) required by a given practice, and the existing cortical biases, the harder the learning process will be, and the more likely that the learning process will disrupt whatever *other* functions the affected brain regions support.

On the other hand, the more the requirements of the acquisition match what is already possible, the less novel and potentially less valuable the cultural practice is likely to be – unless, that is, it is possible to combine existing capacities in new ways, to use old wheels, springs, and pulleys to form new machines. It is interesting to note in this regard that while the neural exploitation and shared circuits theories outlined earlier tend to envision neural circuits being put to fairly similar new uses – for example, forward models in motor control being used to support forward models in social interaction – the neuronal recycling hypothesis suggests that neural circuits might be put to uses quite other than the ones for which they were originally developed. As already noted above, this notion is central to the massive redeployment hypothesis, which we will briefly review next.

6.4. The massive redeployment hypothesis

Since the massive redeployment hypothesis has already been discussed in section 1.1, I will only review the main idea here. The primary distinction between massive redeployment and neuronal recycling is the time course over which each is supposed to operate. Massive redeployment is a theory about the *evolutionary* emergence of the functional organization of the brain, whereas neuronal recycling focuses on cognitive abilities for which there has been insufficient time for specialized neural circuits to have evolved. Both, however, suggest that the functional topography of the brain is such that individual circuits are put to various cognitive uses, across different task domains, in a process that is constrained in part by the intrinsic functional capacities (the "workings" or "cortical biases") of local circuitry.

It is worth noting that the concepts of a "working" and of a "cortical bias" are not identical. The workings envisioned by the massive redeployment hypothesis commit that theory to the existence of cortical biases – that is,

limitations on the set of functions it is possible for the circuit to perform in its present configuration. However, Dehaene is not committed to the notion of a local working in virtue of belief in cortical biases. Although it would be natural to understand cortical biases as the result of fixed local workings, a region could in fact perform more than one working and still have a cortical bias. However, the more flexible regions are, the less their individual biases will differ, and the harder it will be to explain the findings that recently evolved cognitive functions use more and more widely scattered neural components. On the other hand, as noted already above, the data are consistent with a number of functionally relevant constraints on local operation. For example, it could be that the dynamic response properties of local circuits are fixed, and that cognitive *function* is a matter of tying together circuits with the right (relative) dynamic response properties (for a discussion, see Anderson & Silberstein, submitted). In this sense, "cortical bias" is perhaps useful as a more generic term for denoting the functional limitations of neural regions.

In any event, both theories are committed to the notion that putting together the same neural bits in different ways can lead to different – in some cases very different – functional outcomes. In the discussion of the shared circuits model (sect. 6.2), I raised the issue of whether and how a single circuit could be expected to deal with various different kinds of data, as reuse theories seem to require. The question arises here as well: Exactly *how* is such reuse possible? It must be considered a weakness of both the massive redeployment and the neuronal recycling hypotheses that they lack any semblance of a functional model. In describing my theory (M. L. Anderson 2007a; 2007b; 2007c), I have used the metaphor of component reuse in software engineering, which may be useful as a conceptual heuristic for understanding the proposed functional architecture but cannot be taken as a model for the actual *implementation* of the envisioned reuse. In software systems, objects are reused by making virtual copies of them at run-time, so that there can be multiple, separately manipulable tokens of each object type. With wetware systems, no such process is possible. What is reused is the *actual circuit*.

In general, how such reuse is actually effected must be considered an open question for the field. Going forward, supporters of recycling and redeployment need to provide at least three things: specific models of how information could flow between redeployed circuits; particular examples of how different configurations of the same parts can result in different computations; and a more complete discussion of how (and when and whether) multiple uses of the same circuit can be coordinated. Penner-Wilger and Anderson (2008; submitted) have taken some tentative steps in this direction, but much more such work is needed. It is to the credit of both Hurley and Gallese that they each offer a (more or less concrete) proposal in this regard (see Gallese 1996; 2008; Hurley 2005; 2008). That neither seems wholly adequate to the task should not be surprising nor overemphasized; the neurosciences are replete with what must be considered, at best, partial models of the implementation of function by neural structure. More important by far is that neural reuse offers a unique guide to discovery – a sense of what to look for in understanding brain function, and

how to put the pieces together into a coherent whole. If neural circuits are put to many different uses, then the focus on explaining cognitive outcomes should shift from determining local circuit activity and single-voxel effects to uncovering the complex and context-dependent web of relations between the circuits that store, manipulate, or otherwise *process and produce* information and the functional complexes that *consume* that information, putting it to diverse purposes.[31]

One way this effort might be abetted is via the formulation of an even more universal theory of neural reuse than is offered by any of the four theories canvassed above. As should be clear from the discussion, none of the four proposals can explain *all* the kinds of reuse in evidence: reuse supporting functional inheritance, reuse supporting semantic inheritance, reuse that occurs during development, and reuse that occurs during evolution. In fact, each is strongest in one of these areas, and weaker in the others. This opens the obvious possibility that the four theories could be simply combined into one.[32] While it is true that there seems no *obvious* barrier to doing so, in that none of the theories clearly contradicts any of the others, this lack of conflict is in part an artifact of the very under-specification of the theories that leaves them covering distinct parts of the phenomenon. As mentioned already, it may turn out that the kind of functional inheritance required by the shared circuits model precludes the kinds of semantic inheritance required by the neural exploitation hypothesis, or that the schemas envisioned by neural exploitation cannot be modified and expanded along the necessary lines. Likewise, it could turn out that the processes driving massive redeployment are in tension with those driving neuronal recycling; or that one, or the other, but not both can explain semantic and/or functional inheritance.

Should such problems and conflicts arise, no doubt solutions can be found. The point here is simply: We don't yet even know if there will be problems, because no one has yet even tried to find a solution. I would encourage all those interested in the general topic of brain organization to ponder these issues – how does the fact of reuse change our perspective on the organization, evolution, development, and function of the brain? Within what framework should findings in neuroscience ultimately be placed? There is enough work here for many hands over many years.

7. Implications

Although the question of how neural reuse is actually effected must be considered open, the question of whether there *is* significant, widespread, and functionally relevant reuse must be considered closed. In light of all the evidence discussed above, it is clear that there is neural reuse, and there is a lot of it. Neural reuse is a real feature of brain organization, but it is also a novel concept – something about the brain that we are just now beginning to notice. What might it mean? What is the way forward? I close the article with a few thoughts on these topics.

First, and most obviously, the fact of widespread neural reuse seems to favor modal and "embodied" accounts of cognition – and of representational content, in particular – over amodal or more abstract accounts. On the other hand, the neuroscientific evidence for these theories has generally been over-read (M. L. Anderson 2008c). Especially in light of the many different *kinds* of reuse, and the many potential mechanisms by which it may have come about, the claims made on behalf of concept empiricism and embodied cognition need close examination. Although a lack of neural reuse would have been evidence against embodied cognition, concept empiricism, and conceptual metaphor theory, the fact that it is even more widespread than these theories predicted means that neural overlaps are not by themselves evidence for these theories, and do not fully explain the relationships between cognitive domains that are at the heart of these ideas. In particular, it needs to be asked what kinds of reuse will, and will not, support the kinds of inheritance of structure and content these theories require; and whether the evidence actually points specifically to that sort of reuse. In fact, this is one of the main open areas of research for neural reuse: *How* is functional inheritance possible, and *what kinds* of implementations of reuse can lead to *semantic* inheritance of the sort described in concept empiricism, conceptual metaphor theory, and other theories of cognitive grounding? Providing this sort of story would offer the possibility of unifying these different theories of grounding with one another, under the umbrella of general neural reuse. In the absence of such a story, general neural reuse instead threatens to undermine some of the justification for these accounts.

If regions of the cortex are indeed put to many different cognitive uses, this suggests that cortical parcellation and function-to-structure mapping should be approached via multiple- or cross-domain investigations (Penner-Wilger & Anderson 2008; submitted). One way to move forward on this task is via the increased use of effect location meta-analysis, in which multiple imaging studies, each reporting significant effects, are analyzed together to get more accurate information about the brain locations of mental operations (Fox et al. 1998). Although such studies are increasingly common, they are also typically limited to one task domain. There is nothing intrinsic to effect-location meta-analysis or cognitive modeling in general that militates against cross-domain modeling, but in practice it is very rarely done. This is, I presume, because there remains a very strong, and perhaps typically unconscious, assumption that brain regions are both unifunctional and domain dedicated.[33] Widespread neural reuse suggests that this assumption must be given up.

Neural reuse offers an alternative to these assumptions, as well as to the more general selectivity and localization assumptions that have long been the guiding idealization for research in the cognitive neurosciences. In their place, neural reuse offers the strong distinction between working (or local cortical bias) and cognitive use, which can help guide the (re-)interpretation of experimental results, especially those based on single brain-imaging experiments. It also offers the suggestion that attention paid to the *interactions* of multiple regions over the activity of single ones will be well rewarded. Methodological tools that take us beyond single-voxel effects – such as functional connectivity analysis and multi-voxel pattern analysis – may have an important role to play in supporting these efforts (Anderson & Oates 2010; M. L. Anderson et al. 2010; Honey et al. 2007; Pereira et al. 2009; Sporns et al. 2000; 2004).

Once we give up these assumptions, our vocabulary of cognitive function might need specific revision to include fewer domain-specific concepts. In current practice, cortical regions are assigned visual functions by vision researchers, memory functions by memory researchers, attention functions by attention researchers, and so on (Cabeza & Nyberg 2000). But if cortical circuits contribute to multiple task domains, then this practice will not lead to the accurate attribution of workings to these circuits. In light of neural reuse, it appears that this practice can *at best* reveal one of the uses to which a region is put, but is unlikely to hit upon the actual local working (see M. L. Anderson 2007b; Bergeron 2008 for discussions). This best-case scenario requires that the process models are themselves accurate, but it seems implausible to suppose that these models – also typically generated on the basis of domain-focused experimentation – will themselves survive widespread acceptance of neural reuse without significant revision. In this sense neural reuse is a potentially disruptive finding, although hopefully in the service of increased theoretical fertility.

Widespread neural reuse makes it quite clear that there *is not* and cannot be anatomical modularity in the brain. Whether this means there is no *functional* modularity is an open question. Can cognitive functions be independent when they have overlapping neural implementations? Questions about what functional modularity requires are vexed, and different researchers have come to many different conclusions on the matter (Barrett & Kurzban 2006; Carruthers 2006). Whether and precisely how neural reuse constrains this debate is a matter that deserves careful attention.

There are some practical upshots as well. Maps of the overlaps among the circuits supporting cognitive function will support robust predictions regarding cognitive processes and tasks that are likely to interfere with one another. Not only does this offer leverage to the experimentalist in designing inquiries into brain function, it also offers advice to the system designer in designing work flows and machine interfaces. As consumer devices, medical instruments, and heavy machinery become more sophisticated and powerful, increasing attention will need to be paid to the cognitive demands of operating them, and information about neural overlaps will be one important tool in the designers' toolbox (Rasmussen & Vicente 1989; Ritter & Young 2001), especially as leading cognitive models start incorporating information about reuse into their systems (Stewart & West 2007).

Similarly, knowledge of neural overlaps might suggest novel therapies for brain injury. Many therapies for traumatic brain injury are based on the "use it or lose it" principle – the more tasks that stimulate a brain region, the more likely patients are to recover function. Knowledge about the *range* of different tasks that potentially stimulate each region may serve as the basis for unexpected therapeutic interventions, ways of *indirectly* recovering function in one domain by exercising capacities in another. Indeed, there is evidence from healthy subjects that such indirect approaches to strengthening neural function can in fact work – for example, the finding that object manipulation can increase reading comprehension in school-age children (Glenberg et al. 2007).

Finally, given that brain functions are apparently supported by multiuse components, there are possible implications for how cognition might be engineered and reproduced in robotic artificial intelligence (AI) (M. L. Anderson 2008a). That is, neural reuse might recommend a shift from building intelligent systems out of separate, specialized modules dedicated to language, motor-control, vision, and such, to engineering low-level multi-use components that offer services to many different high-level functions. There has been some theoretical and practical work in this direction (Hall 2009; Stewart & West 2007), but much more is needed. Such work is probably the necessary precursor to any satisfactory theory of how it is that component reuse can engender both functional and semantic inheritance. I hope the present article gives some sense that such efforts will be rewarded.

NOTES
1. It is perhaps worth mentioning that, although the first publications on the massive redeployment hypothesis did not appear in print until 2007, the original article detailing the theory was received by *Philosophical Psychology* in 2005. It hence appears likely that all the neural reuse theories of cognition discussed here were independently developed in the very same year.

2. The cortical regions studied were the same as those used in Hagmann et al (2008): "The 66 cortical regions are labeled as follows: each label consists of two parts, a prefix for the cortical hemisphere (r = right hemisphere, l = left hemisphere) and one of 33 designators: BSTS = bank of the superior temporal sulcus, CAC = caudal anterior cingulate cortex, CMF = caudal middle frontal cortex, CUN = cuneus, ENT = entorhinal cortex, FP = frontal pole, FUS = fusiform gyrus, IP = inferior parietal cortex, IT = inferior temporal cortex, ISTC = isthmus of the cingulate cortex, LOCC = lateral occipital cortex, LOF = lateral orbitofrontal cortex, LING = lingual gyrus, MOF = medial orbitofrontal cortex, MT = middle temporal cortex, PARC = paracentral lobule, PARH = parahippocampal cortex, POPE = pars opercularis, PORB = pars orbitalis, PTRI = pars triangularis, PCAL = pericalcarine cortex, PSTS = postcentral gyrus, PC = posterior cingulate cortex, PREC = precentral gyrus, PCUN = precuneus, RAC = rostral anterior cingulate cortex, RMF = rostral middle frontal cortex, SF = superior frontal cortex, SP = superior parietal cortex, ST = superior

temporal cortex, SMAR = supramarginal gyrus, TP = temporal pole, and TT = transverse temporal cortex."

3. If cognitive scientists are very bad at categorizing their experiments – at knowing what cognitive domains or tasks their experiments in fact explore – that *could* explain the simple finding that regions are activated by multiple tasks, because some experiments that belonged in one category would have instead been placed in another. I don't doubt we *are* pretty bad at this. But this fact alone would not explain the specific patterns of findings reported in support of the other predictions of redeployment. Moreover, Tony Chemero and I have performed a clustering analysis on the data to see if there is a way of dividing experiments into groups so that the neural activations do not overlap. There does not seem to be any clustering that avoids overlaps (unpublished data). We have not yet determined whether and to what degree it is possible to *minimize* overlap with alternate clusterings of the experiments.

4. In Talairach space, the origin is located deep in the center of the brain, and regions anterior of that are increasingly positive, and posterior to that are increasingly negative.

5. The terms "working" and "use" are adopted from Bergeron (2008). That brain regions have fixed low-level functions ("workings") that are put to many high-level "uses" is the assumption followed by most work on the massive redeployment hypothesis (M. L. Anderson 2007a; 2007b; 2007c; 2008a; Penner-Wilger & Anderson 2008), but it should be noted that there are other possibilities consistent with the data. For example, it could be that the dynamic response properties of local circuits are fixed, and that cognitive *function* is a matter of tying together circuits with the right (relative) dynamic response properties. See Anderson and Silberstein (submitted) for a discussion.

6. Terry Stewart (personal communication) suggests that an even better analogy might be modern Graphics Processing Units (GPUs). GPUs were initially intended as specialized devices to offload computationally intensive graphics rendering from the main CPU, but it has turned out they are useful for many other tasks. He writes: "it's turning out that they're extremely useful for general parallel processing, and lots of people (including us) are using them to run neural simulations. And, it's an interesting balancing task for the GPU developers to support this new use of the same working while maintaining the graphics use as well." (See, e.g., Ho et al. 2008; Nvidia 2007, sect. 1.1.)

7. ACT-R modules are separately modifiable, and, if neural reuse is true, the functional components of the brain will often *not* be. But separate modifiability does not appear to be an essential aspect of ACT-R theory, the way it is at the core of massive modularity (see sect. 3.1).

8. Some proponents of blending have argued to me that Conceptual Blending Theory (CBT) and Conceptual Metaphor Theory (CMT) are much more different than this brief description allows. For instance, Cristóbal Pagán Cánovas (personal communication) writes that:

Fauconnier and Turner argue that double-scope blending is a defining capacity of our species, of which metaphor is just a surface product, emergent from complex integration network that cannot be described by binary unidirectional mappings. I think that: a) this makes CBT and CMT hardly compatible; b) CBT (unlike CMT) accounts for frame shifting, bidirectional or multidirectional conceptual mappings, emergence of new meanings not present in their inputs, opportunistic re-use of conceptual materials, etc. and thus constitutes a change of paradigm; c) CBT is much more compatible with the massive redeployment hypothesis; d) a deeper debate about CMT and CBT is necessary. (For more on this, see Pagán Cánovas 2009.)

This is certainly a very interesting issue, and I would be especially pleased if Conceptual Blending turned out to be more compatible with the observed extent of neural reuse than CMT appears to be (although whether it could account

for *all* of it is a different matter), but space constraints dictate that we leave the matter for future discussion.

9. Note the reuse of the same neural circuits to support abstract planning wouldn't necessarily mean that one simulates motor experience as part of the planning process. Rather, for conceptual metaphor theory, the neural overlap would support the inheritance of elements of one domain (e.g., its inferential structure) by the other. The discovery of such circuit reuse therefore does offer support for both theories – although, as I have complained elsewhere (M. L. Anderson 2008c), little attention has been paid to the fact that concept empiricists and conceptual metaphor theorists in fact need to interpret this evidence in quite different ways for it to support their specific claims.

10. Apropos of which it should be noted that this approach is broadly compatible with the developmental theories of Piaget, according to which abstract thought depends on the acquisition of sensorimotor skills and concrete operations (e.g., Piaget 1952).

11. Glenberg et al. (2008b) confirmed that motor regions were involved by applying TMS over the motor areas and measuring a motor-evoked potential (MEP) at the hand while having a subject judge both action sentences, describing concrete and abstract transfers, and neutral sentences. A larger MEP response was seen during transfer sentences as compared with non-transfer sentences, consistent with the notion that the motor areas are specifically activated by action sentences.

12. If it were the case that emotional valence was metaphorically mapped to movement in space *without* direct neural sharing, we would be more likely to see that emotions affected movement, but not the reverse, for presumably movement is not metaphorically mapped to anything. The fact that the effect is bidirectional suggests that it is mediated by the activation of something shared by and necessary to both systems, and a shared neural circuit seems a likely (although naturally not the only) possibility.

13. Note that on both views the neural overlaps could remain even if numbers were entirely amodally represented. A complete review of the evidence for and the various theories regarding the nature of this resource would take us too far afield to include here. For a discussion, see (Penner-Wilger 2009; Penner-Wilger & Anderson, submitted).

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

15. The advantages of using this subdivision are that it ensures a neutral choice of ROIs, and lays the groundwork for future studies in which the domain-related topology of the cortex can be directly compared to the cortical connection matrix reported in that study. Thanks to the authors for sharing their ROI data.

16. The domains follow the standards defined by the Brain-Map database (Fox & Lancaster 2002; Laird et al. 2005), and are generally determined by the authors of the study. Where available, we adopted the classification entered into the Brain-Map database itself.

17. The disadvantage of using this set of ROIs is that it is based on 1.5cm^2 regions of the cortical *surface*; hence, many activations deeper in the brain are not captured by this subdivision. One can mitigate this problem by defining a set of cubes of roughly the same size as those from Hagmann et al. (2008) – 12mm on a side – but distributed equally through the entire brain. This brings the eligible total of 12,279 activations in 1,486 experiments. For the sort of counting we are presenting here, this addition of only 17 new experiments does not materially change the results.

18. These are averages of the raw counts. If the averages are normalized to 11 (the number of possible domains in the overall average), the numbers are as follows: Action areas are active in the equivalent of 5.46 (SD 2.17) nonaction domains and 5.79

(SD 2.26) cognitive domains; perception areas are active in 4.90 (SD 1.97) non-perception domains and 5.87 (SD 2.28) cognitive domains; perception-action areas are active in the equivalent of 6.11 (SD 2.23) cognitive domains; and cognitive areas are active in 4.46 (SD 2.18) cognitive domains.

19. See Note 2.

20. The one region active in only one domain was left Frontal Pole, which was active only in memory.

21. The differences are indeed significant, 2-tailed student's t-test, $p \ll 0.01$, whether one uses the raw or normalized counts. Note that the massive redeployment hypothesis would explain this finding in terms of the relative age of the brain regions involved. Perceptual and motor circuits are more frequently reused because they are older, and not necessarily because they are functionally special.

22. Note that for the purposes of this article, the term "circuit" is more-or-less interchangeable with "small neural region." I take the evidence of this section to indicate that small neural regions are activated in multiple tasks across multiple domains, which for current purposes is interpreted to indicate that local neural structures – that is, neural circuits – are reused in these tasks and domains. One certainly *could* reserve the term "circuit" for larger neural structures, such as might be revealed by combining fMRI results with Diffusion Tensor Imaging data that can reveal the physical connectivity underlying function (see, e.g., Behrens & Johansen-Berg 2005; Honey et al. 2009; Sporns et al. 2000), but this lexical preference would not materially alter the claims of this section. And although I do believe that one of the upshots of this article as a whole is that much more attention should be paid to functional connectivity and other measures of the cooperation *between* cortical regions, rather than making functional attributions primarily on the basis of differential activation, following out this implication in detail will have to wait for some future paper (but see, e.g., M. L. Anderson 2008a).

23. The authors explicitly relate this latter aspect to the concept of affordances (Gibson 1979), the perceived availability of objects for certain kinds of uses or other interactions.

24. Hurford (2003) suggested something like this when he hypothesized that the division between ventral stream and dorsal stream vision provide the biological basis for predicate-argument structure.

25. There is the further drawback that, in contrast to the model actually built by Feldman and Narayanan (2004), there is no proof it is actually possible to build such a control system.

26. In fact, most household electronic thermostats contain such forward models, one reason they are more efficient than the older mercury-switch models.

27. This might push the architecture in the direction of something like the "servo stacks" concept (Hall 2009), which imagines building diverse high-level cognitive components from the iterative combination of simple, relatively homogenous, low-level building blocks.

28. The problem remains even given (1) observed actions will be associated with motor commands and those commands may be simulated by the observer, and (2) part of the observation is not just the movements of an agent, but also the effects of the agent's actions. Even if motor simulations kick in, enriching our observational experience, one must *begin* with the visual experience of the action – it is that which drives the initial categorization. And the sensory effects of an action will still differ for actor and observer, so that abstraction – attention to high-level features – will still be required.

29. A somewhat different approach to this problem is offered by Wolpert et al. (2003). In this model, there are multiple predictors and multiple controllers arranged in an abstraction hierarchy. Actions and observations activate different controllers and predictors to different degrees, and the ones that generate the fewest errors (of prediction or of movement) over time are the ones that come to dominate. That is, they are the ones that

come to drive action, or action understanding. Miali (2003) describes how such a model might be instantiated in a large brain circuit involving F5 mirror neurons cooperating with cerebellum and cortical motor areas. In this model, there is no need for translation between levels because there are multiple specialist modules, each corresponding to some (class of) actions or situations, already arranged in an appropriate hierarchy; but there is also little need for reuse.

30. Hurley appears to accept functional modularity, but explicitly denies anatomical modularity.

31. The producer/consumer distinction is based on Millikan's (1984) as it pertains to the content of representations. It will surely need to be part of the model for how circuit reuse is possible. The "same" representation can have different content depending on the characteristics of the representation consumer. Similarly, the same neural activity or output can have different functional significance depending on the nature of the neural partners.

32. My colleague Tony Chemero and I are developing one such model, by adapting and combining insights from the literature on niche construction (Odling-Smee et al. 2005) and the evolution of food-webs (Quince et al. 2002), but the space of illuminating models of this process is surely quite large.

33. Consider the titles of some recent meta-analyses of imaging data: "Functional neuroanatomy of emotion: A meta-analysis of emotion activation studies in PET and fMRI" (Phan et al. 2002); "Meta-analysis of the functional neuroanatomy of single-word reading: Method and validation" (Turkeltaub et al. 2002); "Functional neuroanatomy of emotions: A meta-analysis" (Murphy et al. 2003); "The functional neuroanatomy of autobiographical memory: A meta-analysis" (Svoboda et al. 2006); "A systematic review and quantitative appraisal of fMRI studies of verbal fluency: Role of the left inferior frontal gyrus" (Costafreda et al. 2006). In fact, of the 51 papers that cite Fox et al. (1998), the only one to consider studies in more than one task domain was a paper analyzing the functional connectivity of the basal ganglia (Postuma & Dagher 2006).

Open Peer Commentary

Reuse or re-function?

Daniela Aisenberg and Avishai Henik
Department of Psychology, Ben-Gurion University of the Negev, Beer-Sheva, 84105, Israel.
aisenber@bgu.ac.il
henik@bgu.ac.il
http://www.bgu.ac.il/~henik

Abstract: Simple specialization cannot account for brain functioning. Yet, we believe Anderson's reuse can be better explained by re-function. We suggest that functional demands shape brain changes and are the driving force behind reuse. For example, we suggest that the prefrontal cortex (PFC) is built as an infrastructure for multi-functions rather than as a module for reuse.

Anderson is impressed by reuse; namely, by the fact that the same brain structures are used in different tasks and contexts. He points out that "in combination neural reuse and wiring optimization theory make some novel predictions for cortical layout" (sect. 2, para. 1). We agree that theories assuming simple structural specialization cannot account for all brain functioning. Yet, we suggest that functional demands drive reuse. More than thirty years ago, Paul Rozin suggested that the evolution of intelligence is marked by exploiting routines designed

for a special task or goal, to achieve other goals (Rozin 1976). Namely, routines (programs, served by specific brain tissue) that were designed to provide specific solutions to unique problems become accessible to other systems through evolution and within the individual lifetime. Such routines are also examples for reuse, but they are better described as a change or expansion of function, rather than reuse, because we "make these (adaptive specializations) more generally available or accessible. This would have adaptive value when an area of behavioral function could profit from programs initially developed for another purpose" (Rozin 1976, p. 256).

Rozin connects such changes in accessibility to genetic programs in which "A specialization [circuit] could be extended by releasing (or depressing) the appropriate genetic program at the appropriate time in appropriate neural context. Such extensions have probably occurred many times in the evolution of organisms" (Rozin 1976, p. 260). Dehaene's neuronal recycling hypothesis (Dehaene 2005; Dehaene & Cohen 2007) fits with this conceptualization; " 'neuronal recycling'... refer[s] to the putative mechanism by which a novel cultural object encroaches onto a pre-existing brain system . . . (which) occurs during the life span as a result of brain plasticity" (Dehaene & Cohen 2007, p. 384). We suggest that functional demands are the driving force behind reuse and that these demands shape brain changes.

Frontal control and brain connectivity. Anderson's second assumption in his massive redeployment hypothesis (MRH) is that older areas in the brain would be more subjective to reuse (sect. 1.1, para. 1). In contrast, the frontal lobes are able to perform more functions (or are more reused, in Anderson's words) than lower and older areas (Miller 2000). The assumption that higher and more novel areas in the brain perform more functions can be explained by their connectivity. Specifically, it has been suggested that the prefrontal cortex (PFC) is "built for control," because it is composed of several interconnected areas that are linked to cortical sensory and motor systems and to a wide range of subcortical structures, so that it is provided with the ability to synthesize a wide range of information. Miller (2000) and Duncan (2001) suggested that the characteristics of the system and its connections allow flexibility that enables the system to adjust and control different situations. In addition, the PFC has widespread projections back to lower systems, which allow for a top-down influence. These features make it reasonable to assume that the PFC is built as an infrastructure for multi-functions, rather than as a module to be reused.

Attention. In visuo-spatial attention, responding is commonly faster and more efficient at cued (valid) than non-cued (invalid) locations. In exogenous-reflexive orienting of attention this validity effect is replaced, after 300 msec from cue onset, by faster responding to non-cued locations. This was described as inhibition of return (IOR), which helps to avoid automatic returning to already searched locations and is dependent on involvement of the midbrain superior colliculus (Klein 2000; Posner & Cohen 1984; Sapir et al. 1999). It has been suggested that the evolutionarily older retinotectal visual system developed a mechanism (IOR) which, through connections with higher brain structures (e.g., parietal lobe; Sapir et al. 2004), enabled coordination of reflexive and voluntary attentional systems (Sapir et al. 1999). Connectivity with other brain areas helped to transfer control to higher brain centers.

Anterior cingulate cortex (ACC) – "Dedicated to one high-level use." In his target article, Anderson (suggests that "an individual brain region . . . will not be dedicated to . . . one high-level use" (sect. 2.1, para. 2). Anterior cingulate cortex (ACC) function is of interest here. There is wide agreement that the ACC is involved in conflict monitoring (Botvinick et al. 2004; Kerns et al. 2004). However, recent reports indicate that the ACC and close structures are also involved in outcome evaluation and in reward-based action (Botvinick et al. 2004; Ridderinkhof et al. 2004). Such results suggest that conflict monitoring may be a manifestation of a more general function of the ACC.

Specifically, the ACC is involved in monitoring and evaluating the outcomes of actions, and, in turn, serves to mold goal-directed behavior and achievement of planned behavior.

Numerical cognition. In the area of numerical cognition, many assume that the ability to grasp the number of displayed objects (e.g., counting) is an essential part of the core system that enables the development of the number sense and arithmetic skills. However, there are clear indications for a connection between numerical processing and size perception and judgment (Ashkenazi et al. 2008; Henik & Tzelgov 1982). Accordingly, it is possible that another system, heavily dependent on the processing of size, is the antecedent for the human numerical system. Namely, routines and neural structures built for size judgments were made available, through evolution, due to the need to develop an exact numerical system. Cantlon and colleagues (Cantlon et al. 2009) presented a similar idea: "a system that once computed one magnitude (e.g., size) could have been hijacked to perform judgments along a new dimension (e.g., number)" (p. 89).

Summary. We suggest that functional demands shape brain changes and are the driving force behind reuse. This is a different point of view, rather than just a terminology change.

From the physical to the psychological: Mundane experiences influence social judgment and interpersonal behavior

John A. Bargh,[a] Lawrence E. Williams,[b] Julie Y. Huang,[a] Hyunjin Song,[a] and Joshua M. Ackerman[c]

[a]*Department of Psychology, Yale University, New Haven, CT 06520;* [b]*Leeds School of Business, University of Colorado at Boulder, Boulder, CO 80309-0419;* [c]*Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02142.*

john.bargh@yale.edu   lawrence.williams@colorado.edu
julie.huang@yale.edu   hyunjin.song@yale.edu
joshack@mit.edu
www.yale.edu/acmelab   leeds-faculty.colorado.edu/lw/
web.mit.edu/joshack/www/

Abstract: Mere physical experiences of warmth, distance, hardness, and roughness are found to activate the more abstract psychological concepts that are analogically related to them, such as interpersonal *warmth* and emotional *distance*, thereby influencing social judgments and interpersonal behavior without the individual's awareness. These findings further support the principle of neural reuse in the development and operation of higher mental processes.

The principle of neural reuse and the various competing theories regarding its underlying mechanisms are of great value to the understanding of a growing body of findings within social psychology – those in which concrete physical sensations are shown to influence higher-order processes involved in trust, interpersonal and situational evaluation, and interpersonal behavior. For example, briefly holding a cup of hot (versus iced) coffee just before an impression formation task involving the identical set of information about a given target person changes that impression (Williams & Bargh 2008a): those who had contact with the warm cup subsequently judged the person as warmer (more prosocial, generous, helpful; see Fiske et al. 2007) than did those in the cold-coffee condition. (The effect was specific to variables related to interpersonal warmth, and not an overall positivity effect, as the coffee-temperature manipulation did not affect impression judgments on dimensions unrelated to prosocial behavior.) In a second study, those in the warm-coffee condition were more likely to give their compensation for being in the experiment to a friend (in the form of a

gift certificate), whereas those in the cold-coffee condition were more likely to keep it for themselves. Thus, physical experiences of warmth directly influence perceptions of psychological warmth in another person, as well as the participant's own behavioral warmth towards others (see also IJzerman & Semin 2009; Zhong & Leonardelli 2008).

Similarly, perceptions of physical distance produce corresponding analogous influences on perceptions of psychological and emotional distance. Merely plotting two points on Cartesian graph paper that are relatively far versus close together on the page causes participants to feel more *psychologically* distant from their friends and family, and, in further studies, to show less physiological reactivity to emotionally laden photographs (i.e., more *emotionally* distant; see Williams & Bargh 2008b; Williams et al. 2009a).

In both cases, these effects were predicted in part from the observed ubiquity of priming effects in social psychology in which incidental stimuli are shown to influence higher-order cognitive and behavioral outcomes without the individual's awareness or appreciation of this influence (see, e.g., Dijksterhuis et al. 2007). These priming effects have become so prevalent that the prevalence itself requires an explanation (Bargh 2006). Ours (Bargh & Morsella 2008; Williams et al. 2009b) involved the notion of *scaffolding*, in which the development of more abstract concepts is said to be grounded in earlier-formed concrete concepts (such as spatial concepts that form in infancy and young childhood out of the comprehension of the physical world; Clark 1973; Mandler 1992), or exapted from pre-existing innate structures such as evolved motivations for reproduction and survival (Huang & Bargh 2008). In this manner, associative connections develop between the original physical and the analogous later psychological versions of the concept (warmth, distance), creating multiple physical avenues for psychological priming effects in adults.

It is also possible that such warmth and distance effects have an innate basis. The attachment theorist John Bowlby (1969) notably argued that distance information was of survival relevance to many, if not all, organisms, because it facilitates both keeping close to caretakers when young and vulnerable, as well as the dispersal of conspecifics to reduce competition for scarce resources, as in territoriality behavior. And, at least in the case of primates, Harlow's (1958) pioneering studies of monkeys raised alone showed the importance of early warmth experiences in infancy for successful social functioning as adults; those raised with a cloth mother, with a 100-watt light bulb behind the cloth, adapted much better than did the other parent-less monkeys.

The physical-to-psychological effects are not limited to warmth and distance, and may instead represent a general phenomenon involving many forms of sensory experience. For example, six experiments reported recently by Ackerman et al. (2010) reveal how the sense of touch influences analogously related psychological variables. Holding a relatively heavy (versus light) clipboard on which to evaluate a job candidate causes evaluators to see the candidate as more serious (heavy = serious) about his or her work and also causes the evaluators to take their own judgment task more seriously (they spend significantly longer on it). Working on a jigsaw puzzle with a rough versus smooth surface causes participants to subsequently rate an interpersonal interaction as going less (versus more) smoothly. Likewise, sitting on a hardwood versus cushioned chair produced greater rigidity (less attempt to compromise) in an interpersonal negotiation task.

Taken together, these demonstrations suggest a cognitive architecture in which social-psychological concepts metaphorically related to physical-sensory concepts – such as a warm person, a close relationship, and a hard negotiator – are grounded in those physical concepts, such that activation of the physical version also activates (primes) the more abstract psychological concept. Again, as in most priming research involving these social-psychological concepts and variables, the experimental participants are unaware of these potentially biasing

influences on their social judgments and behavior and so do not correct or adjust for them (Wilson & Brekke 1994).

The principle of neural reuse – specifically, that "local circuits may have low-level computational 'workings' that can be put to many different higher-level cognitive uses" (sect. 1.1, para. 5) – also helps to explain how activation of presumably evolved motivations, such as the mating (reproduction) goal, can exert influences outside of its focal domain of mating – effects that are difficult to understand under the principles of anatomical modularity or functional localization. For example, priming the mating goal influences the evaluation of other living kinds (flowers, fruits), as well in terms of "prime" life stages (Huang & Bargh 2008). Viewed in terms of the principle of reuse, this finding suggests that the mating goal makes use of a "prime lifestage" appraisal circuit, which is activated when the mating goal is primed and is thus influential in other domains as well, not exclusively mate selection.

Overall, these findings are in harmony with Anderson's central point that our mental carriers of meaning are tied to sensory experience to such an extent that one's physical state exerts a pervasive and often unconscious influence over the workings of the mind.

Neural reuse and cognitive homology

Vincent Bergeron

Department of Philosophy, University of Ottawa, Ottawa, ON K1N 6N5, Canada.
vincent.bergeron@uottawa.ca

Abstract: Neural reuse theories suggest that, in the course of evolution, a brain structure may acquire or lose a number of cognitive *uses* while maintaining its cognitive *workings* (or low-level operations) fixed. This, in turn, suggests that *homologous* structures may have very different cognitive uses, while sharing the same workings. And this, essentially, is homology thinking applied to brain function.

The study of human cognition is, in many ways, linked to the study of animal cognition. This is perhaps most apparent if one considers the large number of animal models of human cognitive functions developed in the past few decades. In memory research, for example, various forms of memory or memory systems have been modeled extensively in other species – for example, long-term memory in rats, working memory in nonhuman primates. Vision research provides another good example, where a great deal of our current knowledge of the human visual system comes from an extensive mapping of the macaque monkey's visual system. A less obvious candidate is the study of language. Despite it being a uniquely human cognitive capacity, there is mounting evidence that experimental work in nonhuman primates may illuminate various aspects of language processing (Petrides et al. 2005; Rauschecker & Scott 2009; Schubotz & Fiebach 2006).

In using animal data to explain human cognitive functions, one must assume that there is sufficient evolutionary continuity between the human brain and that of other species. Not all animal data are equally relevant, of course, and whether a piece of data in a given species appears to be relevant to human studies depends on the interplay between several different factors, such as the kind of cognitive systems involved, the evolutionary distance between the two species, and the particular experimental methods used. For example, basic neurobiological mechanisms like long-term potentiation can be studied in evolutionarily distant animals such as Aplysia and rats, whereas higher cognitive functions like executive functions are best studied in evolutionarily closer species such as nonhuman primates.

In its simplest form, this evolutionary continuity assumption is uncontroversial. The human brain shares many of its principles

and functions with that of other species; and for any human cognitive function, we can expect that (at least) some component(s) of it could be found in the cognitive repertoire of another species. What is less clear, however, is how best to exploit this evolutionary continuity in building models of human cognition. This is the challenge of finding precisely which components of human cognitive functions can be successfully studied in other species. Anderson's target article, and neural reuse theories in general, provide a unique perspective on how to accomplish this task.

Central to the concept of neural reuse is a distinction between two concepts of function, namely, "working" and "use." The cognitive *workings* of a brain structure (e.g., Broca's area) are the low-level operations that it performs, whereas the cognitive *uses* of that structure are the higher-level operations (or capacities) to which it contributes. What neural reuse theories suggest is that, in the course of evolution, a brain structure may acquire or lose a number of cognitive uses while maintaining its cognitive workings fixed. This, in turn, suggests that *homologous* structures may contribute to very different cognitive capacities, and thus have very different cognitive uses, while sharing essentially the same low-level internal operations or workings. And this, one might think, is homology thinking applied to brain function.

The idea of functional homology may seem confused at first (Love 2007). After all, the concept of homology was originally defined as "the same organ in different animals under every variety of form and function" (Owen 1843, p. 379), where sameness is defined by common phylogenetic origin. And in fact, homologous brain structures will often have very different functions. For example, Broca's area, unlike its homologue in the macaque monkey (Petrides et al. 2005), is heavily involved in language and music processing (Patel 2003). However, as we have just seen, the fact that these two structures appear functionally dissimilar based on a comparison of their cognitive uses obscures the fact that they may share the same workings. By specifying the workings of the two structures independently of their specific uses, as neural reuse theories suggest we do, one could test whether this is in fact the case. Recent models of Broca's area's workings (Schubotz & Fiebach 2006) provide a first step. For example, Fiebach and Schubotz (2006) propose that Broca's area may function as a hypersequential processor that performs the "detection, extraction, and/or representation of regular, rule-based patterns in temporally extended events" (p. 501). As the model attempts to explain Broca's area's contribution to complex, behaviorally relevant sequences that are also present in nonhuman primates (e.g., action sequencing and the manipulation of objects), and because there is a homologue of the area in the macaque monkey, Fiebach and Schubotz's account of Broca's area's workings appears be a good candidate for a *cognitive homology* – that is, the same workings in different animals regardless of cognitive use, where sameness is defined by the common phylogenetic origin of the associated structures (see also Love 2007 for a similar proposal regarding "homology of function").

Anderson's discussion of the spatial-numerical association of response codes (SNARC) effect (Dehaene et al. 1993) provides another illustration of how homology thinking might apply to cognitive function. When subjects are asked to classify numbers as even or odd by making their responses on either the right or the left side of space, their responses to larger numbers are faster when made on the right side of space, whereas responses to smaller numbers are faster when made on the left side of space. Hubbard et al. (2005) review several lines of evidence in monkeys and humans that point to a region in the intraparietal sulcus as the site of this interaction between numerical and spatial cognition. They hypothesize that the interaction arises because of the common involvement, in both attention to external space and internal representations of numbers, of a particular circuit in this region. Here again, we can think of their account of the workings of this brain structure in both monkeys and humans as a cognitive homology.

Homology thinking applied to brain structures is already an integral part of cognitive neuroscience. The perspective offered by neural reuse theories allows us to extend homology thinking to brain function.

Neural reuse implies distributed coding

Bruce Bridgeman

Department of Psychology, University of California at Santa Cruz, Santa Cruz, CA 95064.
bruceb@ucsc.edu
http://people.ucsc.edu/~bruceb/

Abstract: Both distributed coding, with its implication of neural reuse, and more specialized function have been recognized since the beginning of brain science. A controversy over imageless thought threw introspection into disrepute as a scientific method, making more objective methods dominate. It is known in information science that one element, such as a bit in a computer, can participate in coding many independent states; in this commentary, an example is given.

The tension between interpreting the brain as a collection of specialized areas and as a distributed network is as old as brain research itself. Rejecting the medieval idea that the brain was unitary because the soul was indivisible, nineteenth-century phrenologists emphasized modularity. Although these phrenologists got the details wrong because of inadequate methods, they introduced the idea of different functions being handled by distinct cortical areas. The idea was made concrete by neurologists such as Fritsch and Hitzig (1870/1960) (sensory and motor areas), Broca (1861) and Wernicke (1874) (language areas), and many others. Distributed coding and the neurological evidence for it came from Karl Lashley's (1929) mass action, through his student Karl Pribram's (1971) distributed coding, to present-day parallel distributed processing.

The contrast between the "concept empiricists" and the rational or amodal concept also has a long history, far more than "the last twenty years or so" (as Anderson writes in sect. 4, para. 3) and unknown to most philosophers. The idea that "the vehicles of thought are re-activated perceptual representations" (Weiskopf 2007, p. 156) – which Anderson refers to in this section (same paragraph) – was championed at Cornell by Titchner, a student of Wilhelm Wundt. He fought a long battle with the followers of Külpe at Würzburg, who saw mental life as built of a great hierarchy of ideas. The controversy was defined as an evaluation of the role of imageless thought. Külpe insisted that some ideas had no associated images, and that Titchner just hadn't found those ideas yet. Titchner, in turn, held that all ideas included images, and that Külpe hadn't found the images yet. Each founded a school to press his introspection, and the battle raged around the turn of the twentieth century. Eventually, the whole controversy blew up in their faces, as it became clear that the introspective method could not resolve the issue. Objective data, not private opinions, were necessary for psychology to become scientific. Tragically, philosophers of mind continue to use the discredited method, disguised in phrases such as "it is obvious that" or "a moment's thought will reveal that." Introspection is a good starting point for an investigation, but it can never be an ending point.

The essential features of the "action-sentence compatibility effect" are also older than Glenberg and Kaschak (2002) (referred to in section 4.1, para. 2, of the target article). An obvious example is the Stroop effect (Stroop 1935), well known to cognitive psychologists. Color naming is easy when a printed color name is in the corresponding color, but difficult when the color and name are incompatible, such as the word "blue" printed in red ink.

There are really two parts to the reuse hypothesis: First, a given brain area can be involved in processing functions of more than one kind; and second, a brain area that evolves to perform one function can later be pressed into service to participate in performing other related functions as well.

Models like Hurley's (2008, p. 41) may be too specific in assigning hardwired logic to each problem. It is like tracing the functions of my word processor, or my spreadsheet, through the hardware of my computer. Reuse implies that the hardware can be more flexible, like the general-purpose hardware in my computer that supports a variety of software in the same array of logic elements.

In this light, can cognitive functions be independent when they have overlapping neural implementations? Of course. For example, numbers in a computer's register do not gain their meaning from any particular bit. Rather, it is the combination of bits, 16 or 32 at a time, that determines what is represented. With 16 bits operating as independent detectors, a brain could store 16 different events. But when combined as a binary number, the same 16 bits can code more than 64,000 distinct states. As the number of elements available increases, the combinatoric advantage of this distributed coding becomes overwhelming. Given the large swaths of brain involved in almost any mental operation, neural reuse becomes inevitable.

Sensorimotor grounding and reused cognitive domains

Maria Brincker

Department of Philosophy, Graduate Center, City University of New York, New York, NY 10016.
mbrincker@gc.cuny.edu
sites.google.com/site/mariabrincker/

Abstract: Anderson suggests that theories of sensorimotor grounding are too narrow to account for his findings of widespread "reuse" supporting multiple different cognitive "task domains." I call some of the methodological assumptions underlying this conclusion into question, and suggest that his examples reaffirm rather than undermine the special status of sensorimotor processes in cognitive evolution.

Anderson's massive redeployment hypothesis (MRH) proposes that "reuse" of local cognitive circuits is a general evolutionary principle. "Reuse" is understood as the exaptation of cognitive circuits to new cognitive uses, while retaining prior but separate functions. The evidence for widespread reuses is based on statistical analyses of overlapping activations across predetermined task domains in a wide array of fMRI studies. On this basis, Anderson raises a two-pronged objection to theories of sensorimotor grounding: (1) That they cannot explain all his findings of reuse, and (2) that the functional properties of sensorimotor circuits are not special in regard to evolutionary reuse, nor in grounding higher cognition; these are simply older circuits and hence reused more in evolution.

While I am deeply sympathetic to the project of questioning modularity and investigating neural co-activations and overlaps, I am puzzled by Anderson's approach and suspicious of his conclusions. I propose that his assumptions about "reuse" and "task domains" seem implausible from such a sensorimotor grounding point of view – and hence that his arguments against such theories lose their bite.

Anderson analyzes findings of fMRI activation overlaps in terms of predefined "task domains" (such as visual perception, action execution, inhibition, emotion, memory, attention, language, etc.); and given this methodology, he finds significant activation overlaps in regions beyond typical perceptual or motor areas for multiple, typically "cognitive" tasks domains.

He concludes that sensorimotor theories are too narrow to accommodate such findings of "reuse."

In spite of many admittedly ambiguous expressions, the idea of sensorimotor grounding is *not* that *all* cognitive processes are localized in areas supporting typical action output or perception input. Rather, generally the core claim is that brains develop and have evolved on the basis of and in support of sensorimotor engagements between animal and environment (Clark 1997; Glenberg 2010; Haggard et al. 2008; Hurley 1998; Nunez & Freeman 2000). In short, it is not simply about location, but also about evolution and development.

But how can we tell whether fMRI activation overlaps are due to evolutionary "reuse," rather than simply repeated use of the same functional circuit? Anderson's answer seems to be that, "For neural reuse theories, anatomical sites have a fixed working, but many different uses" (sect. 3.3, para. 3). That is, exaptation does not imply an evolutionary *change* in the local circuit, but simply a reuse of this very circuit to form a new combination with other circuits to support a new cognitive function. This sort of atomistic combinatorial idea features prominently in Anderson's methodological equation between fMRI activation overlaps and evolutionary reuse: "Reuse" simply *is* repeated use of the same anatomical circuit across task domains (sect. 4.4., paras. 4–5).

Anderson himself notes his theory does not address how the brain circuits have expanded and changed over the course of evolution. This is, however, a central issue for sensorimotor grounding theories, and such a perspective precisely calls Anderson's notion of reuse and methodology of counting task domains into question. First, primitive cognitive circuits might be multifunctional at the outset – that is, supporting not only action and perception, but also other of Anderson's "task domains" such as, for example, primitive attention, emotion, and memory functions. Secondly, differentiation from within, in concert with newer cognitive circuits, could form cognitive support systems for increasingly more complex organism-environment engagements. Accordingly, cognitive exadaptions could involve both old and new anatomical regions, and local activation overlaps might be the result of either "repeated use" of already evolved processes, or of evolutionary "reuse."

Anderson's key assumptions that (1) neural activation overlaps = evolutionary "reuse" and (2) his statistical use of predefined "task domains" are therefore questionable. And, given a sensorimotor grounding of reuse and "task domains," there is no obvious incompatibility between findings of areas outside the sensorimotor system, say, medial prefrontal regions, being involved in multiple higher cognitive tasks such as memory, imagery, or motivation – or that other additional "cognitive domains" such as attention would interact with these "default network" processes (Bruckner et al. 2008).

Anderson uses the phonemic character of human speech as an example of a reuse exadaption. His discussion is illustrative in that it shows how he assumes that certain abilities or "task domains" as functionally independent and to a certain extent reified by their cognitive purpose independently of the actual neurobiological instantiation that they happened to get. He describes (via Graziano et al. 2002b) how the evolution of phonemic speech piggybacked on the specifics of the preexisting motor control mechanism organized around endpoint postures. So far so good. But then he writes: "Had the motor control system been oriented instead around (for example) 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"(sect. 4.6, para. 4). Anderson here assumes that complex symbolic and structured language could have evolved absent a motor system organized around perceptual end-goals in abstraction from the precise physical vectors of the kinetic movements. Maybe so, but he makes the tacit assumption that one can separate the sophisticated cognitive function of language not only from its

phonetic character and the concrete physical constraints of the vocal system, but also from what might be a core organizing principle of motor control, namely, sensorimotor goal or end-state representations (Gallese 2003; Hommel et al. 2001; Rizzolatti et al.1988). In my work on mirror neurons and sensorimotor integration (Brincker, forthcoming), I argue that this organization of the motor system exactly presents a seed for abstraction that can be exploited for higher cognitive processes, including language. Accordingly, one might think that sign language could have evolved without our specific vocal system but probably not without sensorimotor end-state organizations.

In summary, Anderson's assumptions differ significantly from the essential ideas of sensorimotor grounding, namely, that there is something about the basic biological acting and perceiving organism that structures the evolution and development of higher cognition. His findings of neural activation overlaps are not incompatible with sensorimotor grounding per se, as these statistical findings simply suggest that neural regions are *used* independently of sensorimotor engagements and say nothing about whether their *evolution* and *primary function* can be understood independently of such.

The importance of ontogenetic change in typical and atypical development

Tessa M. Dekker and Annette Karmiloff-Smith
Centre for Brain and Cognitive Development, Birkbeck College, University of London, London WC1 7HX, United Kingdom.
tdekke01@students.bbk.ac.uk   a.karmiloff-smith@bbk.ac.uk
http://www.psyc.bbk.ac.uk/research/DNL/personalpages/tessa.html
http://www.bbk.ac.uk/psyc/staff/academic/annettekarmilofsmith

Abstract: The compelling case that Anderson makes for neural reuse and against modularity as organizing principle of the brain is further supported by evidence from developmental disorders. However, to provide a full evolutionary-developmental theory of neural reuse that encompasses both typical and atypical development, Anderson's "massive redeployment hypothesis" (MRH) could be further constrained by considering brain development across ontogeny.

Neural reuse is the notion that new cognitive skills are comprised of recombined and reused neural solutions, rather than independently evolved modules. In Anderson's version of such theories, the "massive redeployment hypothesis" (MRH), he predicts that newer cognitive functions will be more scattered across the brain. His reasoning is that local neural circuits have fixed internal workings across evolutionary time, which enables solutions to newer evolutionary problems to draw upon a more widely spread out set of neural building blocks. By providing evidence that all cognitive domains overlap and are distributed across the brain, Anderson convincingly negates the need for implementation of cognitive functions as sets of independently evolved, localized modules in the brain, and, at the same time, makes a compelling case for neural reuse. In our view, however, the MRH falls short of providing a full evolutionary-developmental explanation of brain organization because the roles of ontogenetic change and plasticity across the life span are overlooked.

In fact, one of the strongest lines of evidence against modular organization in the brain comes from in-depth analyses of developmental disorders across ontogeny. Although impairments in developmental disorders seem to be specific to particular cognitive domains and are often taken as evidence for innately specified modularity, this turns out not to be the case. On closer inspection, claims about intact and impaired cognitive modules have consistently overlooked subtle deficits in "intact" domains and have failed to trace cognitive-level impairments in the phenotypic outcome back to their basic-level origins in infancy; that is, they do not account for the full atypical cognitive spectrum over developmental time (see discussions in Karmiloff-Smith 1998; 2009; Southgate & Hamilton 2008).

Take, for example, the case of Williams syndrome (WS), caused by a hemizygous deletion of genes on chromosome 7, resulting in decreased expression of affected gene products throughout the brain from conception onwards. Although the effects of the deletion may be superficially more apparent in certain cognitive domains, in fact they turn out to be widespread across the multiple cortical regions where the genes are expressed and are therefore highly unlikely to be specific to single domain-specific modules. Indeed, in WS, impairments across several domains such as face processing, number, auditory and spatial perception (Brown et al. 2003; Elsabbagh et al., in press; Paterson et al. 1999; Van Herwegen et al. 2008) can be traced to a featural processing bias in infancy (Karmiloff-Smith et al. 2004), which itself is likely to be due to very early atypical saccadic eye movement planning (Karmiloff-Smith 2009). Theories that explain WS in terms of intact and impaired, innately specified modules are based on static descriptions of the phenotypic end state (Bellugi et al. 1999; Pinker 1994; Rossen et al. 1996), ignoring the complex dynamics of development. In contrast to modular theories of the brain, theories of neural reuse are far more likely to explain why pure cognitive deficits in specific brain regions have been so difficult to identify.

How the massive redeployment theory of neural reuse could give rise to adult-like brain organization across the life span needs to be specified further, however. Firstly, it remains unclear whether Anderson considers locally fixed internal workings to be already present at birth – in which case one innately specified mechanism (modules) is simply being replaced by another (fixed internal neuronal workings) – or whether his approach encompasses the development of such neural functions over ontogeny. On the one hand, aspects of neuronal differentiation may indeed emerge early in development through intrinsic factors that determine cortical connections, causing cortically localized functions to be highly preserved across individuals, cultures, and even species (but see Han & Northoff 2008; Orban et al. 2004). On the other hand, research on brain plasticity shows that developmental pressures can dramatically reshape the inner workings of neurons. Most strikingly, this is illustrated by classic studies in which developing patches of cortex received abnormal sensory input. For example, when ferret auditory cortex neurons were rewired to receive visual input, and visual cortex neurons to receive auditory input, the inner workings of both types of neurons changed. The auditory cortex took on characteristics and assumed functions of the visual cortex and vice versa (von Melchner et al. 2000). A neuroconstructivist approach to brain development reconciles these two apparently contradicting sets of findings by suggesting that early differentiation may render certain parts of the cortex more relevant to performing certain functions. However, these initial systems are coarsely coded, and competition between regions gradually settles which regions with domain-relevant biases *become* domain-specific over time, ultimately giving rise to the structured adult brain (e.g., Johnson 2001; Karmiloff-Smith 1998; 2009).

A second issue that remains unclear is whether recombination of connections between specialized regions is the only mechanism that Anderson considers relevant, leaving no role for localized plasticity of neural computation in response to newly learnt tasks such as mathematics and reading. Dehaene's neuronal recycling hypothesis (2005) proposes that such culturally transmitted skills invade neural systems that are already present and that lend themselves well to performing these new tasks. If there is any difference between functions, optimizing a neural circuit with an existing function for a new task will consequently affect tasks that already relied on the same circuit. It remains unclear whether Anderson accepts this possibility or whether he maintains that inner neuronal workings are truly

fixed, which would imply that learning a new task (e.g., reading) should never adversely affect other tasks that depend on the shared neuronal circuitry (e.g., object processing).

To summarize, we maintain that the consideration of ontogenetic change and developmental disorders can provide vital evidence for the organizational principles of the brain, principles that run counter to modular views. We agree with Anderson that neural reuse is a promising organizing principle of the brain, as opposed to the notion that the brain has evolved into a Swiss army knife with innately specified modules uniquely designed for each new cognitive function. However, we suggest that Anderson's massive redeployment hypothesis could be further constrained by considering brain development across ontogeny in order to provide a full evolutionary-developmental theory of neural reuse that encompasses both typical and atypical development.

How and over what timescales does neural reuse actually occur?

Francesco Donnarumma, Roberto Prevete, and Giuseppe Trautteur

Dipartimento di Scienze Fisiche, Università di Napoli Federico II, Complesso Universitario Monte Sant'Angelo, I-80126 Napoli, Italy.

donnarumma@na.infn.it   prevete@na.infn.it
trau@na.infn.it   http://vinelab.na.infn.it

Abstract: We isolate some critical aspects of the reuse notion in Anderson's massive redeployment hypothesis (MRH). We notice that the actual rearranging of local neural circuits at a timescale comparable with the reactivity timescale of the organism is left open. We propose the concept of programmable neural network as a solution.

Reuse, working, function. *Merriam-Webster's Collegiate Dictionary*, 11th edition, gives the definition of *reuse* as: "to use again, especially in a different way or after reclaiming or reprocessing." Thus, for example, the well-known evolutionary sequence from jaw bones of reptiles to the ossicles of mammalian ears may be taken as an instance of an acoustic reuse of the manducatory reptilian jaw bones after an (extensive and) exaptive "reprocessing." Is this the use of "reuse" (no pun intended) in the target article?

Notice that, in the above example, reuse completely obliterates original use. On the contrary, the overwhelming connotation of the term one gleans from an overview of the target article is: "new use or uses, without losing the original function." In the article's Note 5, Anderson clarifies the meaning of *working*: "brain regions have fixed low-level functions ('workings') that are put to many high-level 'uses'."

"Function" or "functionalities" occur in contexts in which it is difficult to separate their meanings from working, or cortical bias, except on the basis of the granularity of the neural circuits considered. "Working" is used for local circuits; "function," for overall cortical, cognitive behavior.

Drawing on numerous excerpts of the article, we summarize the gist of the reuse idea in the massive redeployment hypothesis (MRH), and stress the timescale aspect, as follows: The brain – at least, but not exclusively, in sensorimotor tasks – obtains its enormously diversified functional capabilities by rearranging in different ways (i.e., putting to different uses) local, probably small, neural circuits endowed with essentially fixed mini-functionalities, identified as "workings," *and does so on a timescale comparable with the reactivity timescale of the organism.*

There is one exception where reuse seems to originate in the circuit itself – as contrasted with the empirical rejection of "small neural regions [were] locally polyfunctional" (sect. 1.1, para. 5) – and not in the putting together of circuits: "in at least some cases, circuit reuse is arranged such that different data – both information pertaining to different targets, as well as information about the same targets but at different levels of abstraction – can be fed *without translation* to the same circuits and still produce useful outputs" (sect. 6.2, para. 9; emphasis Anderson's).

A surprising disconnection occurs, though, with respect to timescales. Indeed, Anderson states: "Massive redeployment is a theory about the *evolutionary* emergence of the functional organization of the brain" (sect. 6.4, para. 1). But the actual reuse of neural circuits must occur at the timescale of the organism's intercourse with the environment, as we stressed above. Any "*evolutionary* emergence" does not explain *how* the mechanism of reuse is deployed at real time.

Synaptic plasticity is of no use here, both because of its slower timescale with respect to the reactivity timescale and because the synaptic structure of the neural tissue gets altered and the previous function is lost. Indeed, plasticity is very aptly distinguished, in the target article's Abstract, from reuse and, therefore, from learning.

Need of programming. The conundrum implicit in the MRH, succinctly stated in the quote we chose for our title, is as follows: Evolutionary or exaptive processes have determined a structure of synaptic connections, which must be considered as *fixed* over current reactivity timescales, bringing about all possible *use*ful "arrangements" of local circuits which give rise to the multiplicity of cognitive functions. But how can a fixed structure deploy at reactivity time selectivity over the specific pre-wired arrangements? How can a specific routing of connections be selectively enabled at reactivity time, if the connections are fixed?

The answer is, by programming. Anderson almost says so: "I have used the *metaphor* of component reuse in software engineering" (he writes in sect. 6.4, para. 3; our emphasis) – but then he argues against assuming the metaphor as literal.

Fixed-weight programmable networks. We propose a model that allows real-time programmability in fixed-weight networks, thus solving the conundrum. The model is realized in the Continuous Time Recurrent Neural Networks (CTRNNs) environment. CTRNNs are well known, neurobiologically plausible, modeling tools – as attested, for instance, by Dunn et al. (2004). The architecture we developed sustains a programming capability which is usually associated with algorithmic, symbolic systems only. By means of this architecture one can design either local circuits or networks of local circuits having the capability of exhibiting on-the-fly qualitative changes of behavior (function) caused and controlled by *auxiliary* (*programming*) inputs, without changing either connectivity and weights associated with the connections.

The main idea underlying this approach is as follows: The post-synaptic input to biological neurons is usually modeled in artificial neural networks – and it is so in CTRNNs – as sums of products between pre-synaptic signals originating from other neurons, and the weights associated with the synapses. So, the behavior of a network is grounded into sums of products between pre-synaptic signals and weights. In the proposed architecture, we "pull out" the multiplication operation by using auxiliary (*interpreting*) CTRNN sub-networks providing the outcome of the multiplication operation between the output of the pre-synaptic neuron and the synaptic weight. In this way, one obtains a Programmable Neural Network (PNN) architecture with two kinds of input lines: *programming* input lines fed to the interpreting CTRNN subnetworks, in addition to *standard data* input lines. As a consequence, a PNN changes on the fly the mapping (working/function) it is performing on standard input data, on the basis of what is being fed into its programming input lines. Notice that a PNN is strictly fixed-weight. More importantly, notice that the two kinds of input signals are different only on a contextual basis. If input signals are

fed to the appropriate lines, then they will be interpreted as *code*, but – as in programming practice – they have the nature of *data*, and, as such, can be processed or originated by other parts of a complex network.

The proposed solution. By using PNNs, one can develop an artificial neural network composed of fixed, that is, non-programmable, local neural circuits which can be rearranged in different ways at "run-time" by programmable, and still fixed-weight, routing networks. The local circuits will be thus reused in different arrangements, giving rise to different overall functions and cognitive tasks. But PNNs are also hypothetical models for fully programmable local networks, thus suggesting an answer to the "exception" we mentioned above. There we take *"without translation"* to mean that those data are fed to the programming inputs – an enticing possibility.

Bibliographical notice. The seminal motivations for programming neural networks, in a more general setting than that of reuse, are expounded in Tamburrini and Trautteur (2007) and Garzillo and Trautteur (2009); some toy realizations were presented in Donnarumma et al. (2007), and a full implementation of the concept is reported in Donnarumma (2010).

Sleep, neural reuse, and memory consolidation processes

William Fishbein, Hiuyan Lau, Rafael DeJesús, and Sara Elizabeth Alger

Laboratory of Cognitive Neuroscience and Sleep, The City College and Graduate Center, The City University of New York, New York, NY 10031.
wfishbein@ccny.cuny.edu
cuhiuyan@yahoo.com   ralphdj01@aol.com   sealger78@gmail.com

Abstract: Neural reuse posits development of functional overlap in brain system circuits to accommodate complex evolutionary functions. Evolutionary adaptation evolved neural circuits that have been exploited for many uses. One such use is engaging cognitive processes in memory consolidation during the neurobiological states of sleep. Neural reuse, therefore, should not be limited to neural circuitry, but be extended to include sleep-state associated memory processes.

Anderson's neural reuse hypothesis posits the development of functional overlap in brain system circuits to accommodate increasingly complex and evolutionarily more advanced functions. The notion of reuse is also consistent with many researchers' thinking regarding multiple functions of brain circuitry. The work in our laboratory centers around the ongoing processes of cognitive functions during sleep, and its various stages, that might be associated with different forms of memory, including implicit, explicit, and emotional salient memories with specific yet overlapping or instantiated neural circuits associated with each. Yet we operate with the implied assumption that memory is not the sole function of sleep, but an evolutionary epiphenomena that has played a central role in the development and retention of complex and advanced cognitive abilities.

Species adaptation of the basic rest-activity cycle seen in plants and animals suggest an evolutionary adaptive aspect to this universal behavior. The development of this universal process appears to fulfill a myriad of ancillary activities. There has certainly been much debate about the functional purpose of sleep – the rest aspect of the rest-activity cycle. While there is no debate about the functional importance of eating, drinking, and engaging in sexual behavior, a clear conclusion regarding the biological state of sleep has yet to be determined. Many theories abound, and memory consolidation is one such theory. Although sleep is more likely to be an adaptive state for more vital purposes such as energy conservation, the sleeping state depends upon essential neural circuitry that hosts neurophysiological and neurochemical dynamics important for memory processing. One such example, the cortical cycle of synchronized and desynchronized neural firing during slow wave sleep (SWS) may serve to globally reduce and restrict unsustainable synaptic growth resultant of learning experiences in wakefulness (Tononi & Cirelli 2003; 2006). At the same time, the reduction of weak synaptic connection may inadvertently enhance the signal-to-noise ratio for more significant connections that are strong enough to survive this global downscaling. Another example might be the neurophysiological and neurochemical dynamics occurring during the various stages of sleep, involving brainstem activation or hippocampal to cortical off-line activation (Buszáki 1998; Hasselmo 1999), acting upon newly acquired information, thereby facilitating long-term memory consolidation.

Several laboratories, including our own (Alger et al. 2010; Tucker et al. 2006), have provided evidence demonstrating that the neurobiological state of sleep plays an essential role in facilitating the formation of direct associative and non-associative memories. We (Lau et al. 2010), along with Wagner et al. (2004), Ellenbogen et al. (2007), and Payne et al. (2009), have extended these findings demonstrating that sleep also facilitates the formation of relational memories – the flexible representation and expression of items not directly learned. The mechanisms underlying the processing of direct associative and relational memory appear related to the physiological events occurring during SWS (Lau et al. 2010). Besides temporally coordinated physiological activities specific to the hippocampal-neocortical circuitry (Buszáki 1998), SWS is also characterized by global synchronized oscillatory activities (Tononi & Cirelli 2003; 2006) and depressed acetylcholine level (Hasselmo 1999). Perhaps associations between items learned before sleep are strengthened and reorganized inadvertently through these widespread activities during sleep to form more energy-efficient and functionally flexible networks among existing neural substrates.

Similarly, once treated as distinct and separate from that of cognition, emotions involve neural circuitry that host neurophysiological and neurochemical dynamics. The traditional limbic system theory supports the idea that neural resources (e.g., physiological or somatic) were carried out by the evolutionarily old cortex (i.e., the so-called reptilian brain), whereas cognitive processes (i.e., higher-order functions) were subserved by the neocortex. The present view, however, integrates both the limbic system and the neocortex as separate but interacting brain systems functioning in parallel. The processes of long-term potentiation (LTP), long-term depression (LTD), and neural plasticity are just some of the ways that the brain can reorganize and change its pattern of activity across cortical regions (and across modalities) in response to experiences.

Following this logic, one can imagine that neural reuse, whether evolutionary old or new, also follows a similar trend whereby mental functions are mediated by separate but interdependent brain processes. In the context of emotional arousal, the domain highly implicated for such processing is the amygdala, interacting with the hippocampus, thereby playing a role in supporting the formation and storage of emotionally salient forms of declarative memories. The brain state supporting such a process appears to occur primarily during the low-voltage, fast activity of rapid eye movement (REM) and stage II sleep (DeJesús et al., in preparation).

Therefore, the processes ongoing during the different sleep stages, stage II, SWS and REM sleep, might serve to consolidate distinct aspects of emotionally salient declarative memories. Whether it is one process or mechanism or another, it would appear that evolutionary adaptation evolved neural circuits that may have been exploited for different uses, and one such use may be the cognitive processes engaged in memory consolidation that occur during the neurobiological states of sleep.

Therefore, the notion of neural reuse should not be limited to recycling of neural circuitry, but should extend to recycling of neurobiological processes that may have well served the evolutionary advancement in mammalian intelligence.

Reuse (neural, bodily, and environmental) as a fundamental organizational principle of human cognition

Lucia Foglia[a] and Rick Grush[b]

[a]*Dipartimento di Studi Storico, Sociali e Filosofici, Università degli Studi di Siena, 52100 Arezzo, Italy;* [b]*Philosophy Department, University of California – San Diego, La Jolla, CA 92093-0119.*

foglialucia@libero.it   rick@mind.ucsd.edu
http://mind.ucsd.edu

Abstract: We taxonomize the varieties of representational reuse and point out that all the sorts of reuse that the brain engages in (1) involve something like a model (or schema or simulator), and (2) are effected in bodily and external media, as well as neural media. This suggests that the real *fundamental* organizational principle is not *neural* reuse, but *model* reuse.

The target article discusses a number of proposals concerning the reuse of neural mechanisms, and these fall broadly into two categories: those which are motivated primarily by representational considerations, and those which are motivated by purely neurophysiological considerations (e.g., cortical areas determined to be active during a variety of tasks). We won't discuss the latter sort of proposals, but will focus on the former. These all involve the reuse of something like a model of some domain. They differ on how the model is reused. In one sort of case, a model of some domain D1 is used to represent, or model, some distinct domain D2. An example would be using models of *space*, or movement through space, to represent *time*. Call this *domain reuse*. The other sort of case is where a model of D1 is still representing the same domain D1, but serves a different function. For example, a model used for *perceptual processing* of environmental scenes is used to generate *imagery* of those same scenes. In this case, the domain represented is the same, but the function (perception, imagery, memory, planning, language comprehension) may be different. Call this *functional reuse*. It isn't obvious what other sort of reuse there could be.

We want to point out that, remarkably, both these sorts of reuse are not limited to neural models. Domain reuse is evident in using physical lines (or circles on clocks) to represent time, or using parts of one's body, such as fingers, to represent numbers. Functional reuse occurs, for instance, when one uses a chess-board to not only *play* a game, but to *plan moves* by physically implementing mock sequences of moves on the board. Another example would be cultural rituals where important prior events are *remembered*, as opposed to *performed*, through re-enactment (re-enactments of Civil War battles are not battles, any more than a memory of a birthday party is itself a birthday party).

This suggests that what is *most* interesting about the human brain is not neural reuse per se, but the fact that the brain is able to use things as models, and then submit those models to both domain and functional reuse. The deep interesting principle here is *model reuse*. That some of these models are implemented neurally is obviously interesting, but it may not be the crucial organizational principle.

Domain reuse includes, among other things, the examples falling under the heading of conceptual metaphor theory. Most generally, a representation of the source domain is used to model the target domain. Familiar examples are space being used to represent time (or money, or state transitions). But the entity reused need not be neural: Fingers can be used to model numbers; drawn circles on the ground to represent logical inclusion relations, or possible state transitions. Interestingly, the latter is a strategy widely used in cognitive-behavioral therapy where drawn diagrams can represent emotional state transitions to help patients understand their situation and possible remedies.

Functional reuse includes the examples of so-called concept empiricism, among others. In concept empiricism, the idea is that some model or scheme that is used in perception, say,

perceiving a spatial relationship such as spatial inclusion, is reused for a different function, such as imagery, information retrieval, or language comprehension (e.g., the word "in"). A related view is Grush's emulation theory of representation (Grush 2004), which describes in detail how models used for perceptual functions can be reused for visual imagery, motor planning, and many others. Other examples include making a sensibility judgment (whether the sentences such as "kick a ball" or "kick a cloud" convey a feasible body movement), which, as the target article discusses, requires the activation of the motor circuits usually involved with modeling the body for planning and guidance of real actions. Here, a model of the body does not serve one of its primary functions, like motor planning, but is reused for a totally different purpose: language comprehension.

This ability, however, seems to transcend neural models. We can take a chess-board, from its primary use of an arena in which to make moves and play a game, and reuse it to plan moves, or even to help understand why someone might have made a certain move. Of course, we could also use a neural model for the purpose.

In some cases, it is not obvious whether functional or domain reuse is the best analysis. Mirror neurons, for example, could be analyzed either way. If one takes it that their proper domain is the agent's own behavior, then using mirror neurons to model or understand another agent's behavior would be domain reuse. On the other hand, if one takes their proper domain to be motor behavior generally, then using mirror neurons to execute behavior versus to understand another agent's motor behavior would be functional reuse. And sometimes there are combinations of both kinds. We can use an external spatial arrangement, like a calendar, to represent time, but we can also use it for different functions: to keep a record of what actually happened at various times, to plan what we might do at different times, to communicate to someone what we want them to do at some time, and so forth.

We can imagine that some might quibble with our use of the expression "model," but in the relevant sense, what others call schemas or simulators are models, or degenerate cases of models. It also might be maintained that other kinds of reuse do not involve anything like a model – for example, some have claimed that merely reusing non-model-involving motor areas is sufficient for generating imagery. We have explained elsewhere (Foglia & Grush, in preparation) that such "simulation" or "enactive" accounts require a model in order to do their job, and we won't rehash those points here.

Our present points are that while we agree that neural reuse is interesting, it seems to us that (1) they are all examples of the reuse of one or another kind of a model, and (2) the human brain is not limited to neural models. Accordingly, we suggest that investigations into the architectural requirements for constructing, using, and reusing models, whether in neural or non-neural media, will teach us much about the brain.

Understanding brain circuits and their dynamics

Antoni Gomila[a] and Paco Calvo[b]

[a]*Department of Psychology, University of the Balearic Islands, 070XX Palma, Spain;* [b]*Department of Philosophy, University of Murcia, 30003 Murcia, Spain.*

toni.gomila@uib.cat   fjcalvo@um.es

Abstract: We argue that Anderson's "massive redeployment hypothesis" (MRH) needs further development in several directions. First, a thoroughgoing criticism of the several "embodied cognition" alternatives is required. Second, the course between the Scylla of full holism and the Charybdis of structural-functional modularism must be plotted more distinctly. Third, methodologies better suited to reveal brain circuits must be brought in. Finally, the constraints that naturalistic settings provide should be considered.

In his target article, Anderson points to the fact that currently available fMRI neuroimaging data clearly show that "neural reuse" or, more precisely, anatomical polyfunctionality, is a pervasive feature of brain organization. He further argues that this polyfunctionality makes it impossible to distinguish which of the various versions of cognitive embodiment proposed so far is more plausible. His main point is that the evidence just shows that multiple cortical regions are involved in multiple tasks, whereas the different theories of embodied cognition conceive in different ways the functional import of such "reuse" or polyfunctionality: as semantic grounding, as simulation of experience, or as anticipation of feedback. However, Anderson does not develop a sustained criticism of such approaches; rather, he insists in their shortcomings as general approaches to brain function. In this regard, much more could be said, precisely on the grounds of the neurophysiological evidence he discusses. Thus, for instance, "simulationist" accounts that appeal to internal models in the brain as grounding for higher cognitive functions, ought to consider the evidence that the "efference copies" are fed to a distinct brain region, as in the case of motor control, where modelling appears to take place in the cerebellum (Kawato et al. 2003), not in the motor cortex. Conversely, both simulationist and conceptual metaphor theories should explain how it is possible for the activation of the same circuits to correspond to different tasks and levels of cognitive abstraction (Gomila 2008).

Second, Anderson's approach has the potential to avoid the Scylla of full holism and the Charybdis of structural-functional modularism, but maybe not as it stands. In the article, full holism is represented by connectionist neural networks, although it refers to the more general idea that function emerges out of the interaction of basic equipotent units. Modularism, by contrast, views the brain as an aggregate of independent, decomposable, functional units with their own proprietary anatomic (maybe even genetic) structure. Anderson's proposal requires that basic units of brain circuitry be identifiable, both structurally (say, in terms of cell assemblies) and functionally, in order to look for the different "higher-level" circuits in which they can be "used," again both structurally (how this basic functional units can be multiply connected with many others) and functionally (what they do depending on which connectivity gets activated). The problem here is whether the requirement of a basic, independent "functionality" – in Anderson's terminology, the "work" of the circuit, distinguishable from the "uses" to which it is put through its "redeployment" – makes any neuronal sense. In principle, it could even happen that it is not the whole basic unit that gets reused, but rather that different components are differentially involved across functions. In other words, the challenge resides in the very possibility of specifying such elementary components in the brain, given that the individuation of circuits cannot be made in the abstract, but always within a functional setting.

Moreover, third, although Anderson widely uses the expression "brain circuit," standard fMRI-based methodologies simply uncover differential, regional, metabolic activity, and are therefore inadequate to unearth brain connectivity as such. Towards the end of the target article, Anderson calls for new methodological approaches, such as multiple- or cross-domain studies; but these should avoid the limitations of subtractive methodologies. An alternative methodology in this regard is to look for common patterns of activity through different tasks. Inspired by a complex systems approach to the brain, this approach applies the analytical techniques of network analysis to find out which nodes are shared by multiple tasks (Eguiluz et al. 2005; Sporns et al. 2004). This approach initially confirms a hierarchy of levels of structural organization, suggesting that neural reuse does not characterize equally all network nodes: brain connectivity takes the structure of scale free networks. Another interesting option is tensor diffusion, a method based on structural magnetic resonance, which uncovers white matter interregional connectivity, and whose functional import has

already been shown (Behrens & Johansen-Berg 2005; Fuentemilla et al. 2009).

Lastly, fourth, one may wonder how we can discover whether neural reuse does constitute an "evolutionary [. . .] strategy for realizing cognitive functions" (sect. 1, para. 3), when the data reported in support of Anderson's framework is not ecological after all. It is noteworthy that in order to enhance neural specificity, experimental designs require a high degree of sophistication; a form of manipulation that, although needed, prevents us from knowing whether the results thus obtained still hold true under naturalistic settings. For example, we do not know if, say, the Fusiform Face Area responds to "ecological" faces in the wild (Spiers & Maguire 2007). Hence, in our view, beyond the exploitation of methodologies other than fMRI in order to be able to properly speak of "brain circuits," the explanation of how structure relates to function requires paying closer attention to the way the environment and the body constrain the sensory and cognitive structure and function in naturalistic, non-task-evoked, settings. In fact, task-evoked responses promote a static interpretation of brain function, which is orthogonal to the spirit of the version of embodiment that underlies Anderson's MRH.

Anderson presents his MRH as a general account of how structure relates to function in the brain. His view somehow reminds us of Simon's (1962/1982) monograph, "The architecture of complexity." Even if Anderson does not mention the term "hierarchy" in his proposal, it seems to be implicit in his conception of a set of basic anatomical circuits, with distinctive functionalities, that constitute several "second-order" circuits by re-wiring, thus giving rise to new functionalities, and so on and so forth. Each new level of organization inherits the capabilities of the circuits involved. In addition, the same basic circuitry can become part of multiple higher-level circuits/functions. In Anderson's proposal, this process of amplifying capabilities by re-wiring of circuits is thought to be characteristic of evolution. However, it doesn't need to be so restricted. It could also account for the possibility of new circuits appearing in phylogenesis (that is, new circuits, not just reuse of the ones available, as it was the case in human brain evolution), as well as of functional reorganization in ontogenetic development (Casey et al. 2005), in learning, and in cases of brain plasticity after stroke, for instance. But, if neuroimaging data are to help us choose among competing views of cognition, the set of issues raised in this commentary must be addressed with an eye to furthering Anderson's project.

Neural reuse in the social and emotional brain

Mary Helen Immordino-Yang,[a] Joan Y. Chiao,[b] and Alan P. Fiske[c]

[a]Brain and Creativity Institute and Rossier School of Education, University of Southern California, Los Angeles, CA 90089; [b]Psychology Department, Northwestern University, Evanston, IL 60208; [c]Anthropology Department, University of California Los Angeles, Los Angeles, CA 90095.

mhimmordino-yang@post.harvard.edu   jchiao@northwestern.edu
afiske@anthro.ucla.edu

http://rossier.usc.edu/faculty/mary_helen_immordinoyang.html
http://culturalneuro.psych.northwestern.edu
http://www.sscnet.ucla.edu/anthro/faculty/fiske/

Abstract: Presenting evidence from the social brain, we argue that neural reuse is a dynamic, socially organized process that is influenced ontogenetically and evolutionarily by the cultural transmission of

mental techniques, values, and modes of thought. Anderson's theory should be broadened to accommodate cultural effects on the functioning of architecturally similar neural systems, and the implications of these differences for reuse.

Reuse of tissues, organs, and systems is a key adaptive strategy in all phyla across evolution and through development. Neural systems are reused in the evolution and development of complex human behaviors, including social emotion and the representation of social status. Research shows: (1) evolutionary and developmental reciprocal reuse between social and nonsocial neural systems; (2) the importance of cultural transmission as a mode for learning evolutionarily and ontogenetically new uses and combinations of neural systems; and (3) the possibility that socially mediated reuse may affect the original, primitive function of a neural system, either developmentally or evolutionarily. In short, although Anderson's approach maps distinct cognitive functions to unique networks, neural reuse within and between networks is a dynamic process involving culture and sociality.

Compassion and admiration: Neural reuse between a social and a somatosensory system. A growing body of evidence points to developmental and evolutionary reuse between a social and a somatosensory system in the feeling of social emotions. Brain systems involved in the direct sensation of physical pain in the gut and viscera (e.g., during stomach ache), are also involved in the feeling of one's own social or psychological pain (Decety & Chaminade 2003; Eisenberger & Lieberman 2004; Panksepp 2005). These systems are also involved in the feeling of late-developing social emotions about another person's psychologically or physically painful, or admirable, circumstances (Immordino-Yang et al. 2009). These systems most notably involve the anterior insula, anterior middle cingulate, and ascending somatosensory systems in the dorsal midbrain, most directly associated with regulation of arousal and homeostasis.

Comparative social status: Neural reuse between a social and a cognitive system. The intraparietal sulcus (IPS) is important in representing comparative numerosity, quantity, magnitude, extent, and intensity (Cohen et al. 2008; Dehaene et al. 2003); it is also involved in representing social status hierarchy (Chiao et al. 2009b). Particularly when comparisons are close, neural activations observed within the IPS for numerical and social status comparisons parallel behavioral distance effects in reaction time and error rates, and are thought to reflect a domain-independent spatial representation of magnitude, including the "magnitude" of social rank.

All animals are responsive to magnitudes, distances, temporal intervals, and intensities (Gallistel 1993). The neurocognitive systems that support this seem to have been reused in evolution to represent the linear dominance hierarchies that are ubiquitous in both vertebrates and invertebrates. Social dominance hierarchies existed long before the invention of symbols to mediate mathematical calculation, so it is likely that the neural systems modern humans use for analog processing of numerical symbols reflect this phylogenetic history.

The social chicken or the useful egg? Learning cognitive skills through cultural transmission. In addition to demonstrating neural reuse in the social brain, the juxtaposition of these examples demonstrates the importance of considering the social sources and functions of the complex skills underlain by neural reuse. Many of modern humans' complex mental functions, both social and nonsocial, are learned through cultural transmission of practices and cognitive techniques, and are further shaped by social values, emotional relevance, and cultural modes of thought. For example, the use of numeral symbols to represent, remember, and communicate magnitude depends on the cultural invention and transmission of such symbols. Learning to use a number board or abacus allows the reuse of systems in the motor and visual cortices to calculate and remember quantities. Similarly, the cultural invention and transmission of calendars and later digital PDAs entails the reuse of perceptual object recognition and spatial relations systems, in conjunction with fine motor control skills, for temporal mnemonics. Similar

processes operate in neurochemistry. For example, oxytocin, whose original functions were to mediate birth and lactation, was evolutionarily reused to bond infants and mothers, then further reused in a small proportion of mammals for parental pair-bonding (Lee et al. 2009). Subsequently, oxytocin systems were culturally reused in diverse social bonding rituals and recently exploited in recreational ingestion of MDMA (ecstasy).

The function of culture in shaping the use of neural systems is demonstrated by cultural variation in the neural correlates of visual attention (Lin et al. 2008) and self-representation (Chiao et al. 2009a), including differential activation patterns within the same neural systems, which can be manipulated by cultural priming in bicultural individuals (Chiao et al. 2010). Together, these findings suggest that Anderson's assertion that putting "together the same parts in the same way [will lead to] the same functional outcomes" (sect. 1.1, para. 6) may not adequately account for the dynamic effects of socialization on neural reuse.

Conversely, the reuse of a neural system for a more complex, culturally organized task apparently can affect its recruitment for a phylogenetically or ontogenetically earlier use. Cross-cultural psychiatric research shows that various Asian populations tend to manifest psychosocial distress somatically, in medically unexplained bodily symptoms, whereas Westerners tend to express depression psychologically (Parker et al. 2001). Cross-cultural work in progress by Immordino-Yang and colleagues suggests that such tendencies may be associated with cultural differences in the recruitment of neural systems for somatosensation in the cortex and brain stem during social processing, extending even into midbrain nuclei that regulate basic bodily functions.

From use to reuse and back: Toward a dynamic, sociocultural theory of reuse. Anderson's theory proposes that neural reuse is mainly a process of organizing low-level circuits with relatively fixed functions into interconnected networks, and that functional differences between cognitive domains correspond to differences in the architecture or organization of these networks. Here, we argue that Anderson's model should be expanded to account for the possibilities that social learning produces distinct culturally informed operations within architecturally similar complex networks, and that the reuse of a low-level neural circuit may, in turn, influence its original, primary function. Future research should investigate how socioculturally shaped ontogenetic processes interact with the constraints and potentials of neural subsystems, connectivity, and chemistry. Are there (as Anderson assumes) fundamental components of neurocognition that are not decomposable – or how modifiable are the functions of such basic components? What biologically and culturally transmitted processes, and what social and nonsocial experiences at what stages of development, determine how neurocognitive components are combined? In humans, neural reuse involves dynamic interplay among social and nonsocial (re)uses over developmental, cultural-historical, and evolutionary timescales.

Neural reuse: A polysemous and redundant biological system subserving niche-construction

Atsushi Iriki

Laboratory for Symbolic Cognitive Development, RIKEN Brain Science Institute, Wako 351-0198, Japan.

iriki@brain.riken.jp

http://www.brain.riken.jp/en/a_iriki.html

Abstract: Novel functions, which emerge by reusing existing resources formerly adapted to other original usages, cannot be anticipated before the need eventually arises. Simple reuse must be accidental. However, to survive the evolutionary race, one cannot merely keep hoping for a

string of good fortune. So, successful species might be gifted with "rational" and "purposeful" biological mechanisms to prepare for future reuse. Neural reuse must be extrapolated from such mechanisms.

Anderson thoroughly reviews neural reuse as a common brain mechanism for human cognitive functions to emerge. During evolution, whenever organisms were faced with a novel unforeseen environment, they had no other means to overcome any immediate problems than by reusing any materials at hand. So, neural reuse appears to be a truly "fundamental organization principle" (target article title). However, it remains an open question as to how human higher-cognitive functions appear as though they are "ensured" to "evolve" much quicker than via ordinary biological evolutionary processes (sect. 6.3, para. 1). To bridge this gap, I try to propose here a "more universal theory of neural reuse" (sect. 6.4, para. 5) grounded in a broader evolutionary framework.

Anderson's "massive redeployment hypothesis" (MRH) stands on two major observations – selectivity and localization are not central features of the brain (sect. 1, para. 1), and newer brain networks of cognitive functions tend to involve more brain areas than older ones (sect. 1.1, para. 1). Four other premises could be recognized further: (1) Biological systems are never ultimately efficient – systems require some redundancy to be stable, adaptable, and sustainable, leading extreme (over-adapted) efficiency to risk flexibility to survive novel situations. (2) A somewhat redundant brain structure would allow representational bistablility, for both the original and adapted functions. Such bistability, or "polysemy," would support the use of metaphor in conceptual structure (sect. 4, para. 1). In addition, gains of further redundancy to stabilize this adapted alternative usage, perhaps by rapid brain expansion, would allow rapid construction of new neural-niche (sect. 4.6, para. 2). (3) Humans have attained unusually long post-reproductive life spans, particularly for females. Reuse-based acquisition of cognitive functions, and resulting accumulation of knowledge, continues over the whole lifespan, tending to peak in middle to old age. Hence, for semantic inheritance (sect. 7, para. 2) over generations to happen, some extra-genetic mechanisms are indispensable. Finally, (4) a "novel concept" (sect. 7, para. 1) that realizes neural reuse should not be found only in *Homo sapiens*, but precursors must exist in nonhuman primates (sect. 6.3, para. 3) and are perhaps also present in other extant taxa.

Evolution ushers diversity and complexity (or, adaptive radiation), perhaps through two major different paths: Species with short life spans and mass-reproduction adapt to environmental changes through variations in their numerous offspring, expecting at least a few to survive. Species with long life spans and low birth rates do so through an individual capacity to adapt. This is readily carried out through expansion of an organ to control adaptive behaviors – the primate brain, and that of humans in particular, is the extreme example. The latter evolutionary process may not be naïve mutation and selection, but rather like the Baldwin effect that initially induced modification, within the range of preprogrammed adaptation, stands by for later mutations to optimize it – modular structures and their exploratory behaviors are proposed to be essential to realize such a phenomenon (Kirshner & Gerhart 2005). The concept of reuse would reinforce this path. That is, slightly excessive redundancy of the brain, initially designed to stabilize a system against unexpected environmental noise, occasionally allowed the system to be polysemous. This newly acquired bistable state enabled systems to be reused for completely different functions in the future, maybe in combination with other parts of the brain. Such novel networks could wait for optimization through later genetic changes, perhaps induced by an emergent epigenetic factor, and become embedded in the environment as a result of the function of the network itself – thus, enabling post-reproductive inheritance. This hypothetical mechanism, referred to as "real-time neural niche-construction" (sect. 6.3, para. 4), seems to be supported by recently discovered concrete biological phenomena, which are described below.

Monkey intraparietal neurons normally coding body image could be trained to code a tool in a way equivalent with the hand holding it (sect. 6.3, para. 4; Iriki et al. 1996) – thus, bistable or polysemous for the hand or the tool. This functional plasticity might range within a fringe of the system prepared for body growth, which came across adaptable to "sudden elongation" by the tool. This accidentally established equivalence between body parts (hands) and tools, in turn demonstrated additional polysemic interpretations, that is, hands were extended towards tools (externalization of innate body), or alternatively, tools were assimilated into the body schema (internalization of external objects). This "self-objectification" process happened to adapt further for the mind and the intention to emerge (Iriki 2006). However, if this new function stays limited within existing neural machinery, it is merely plasticity, or a learning process. But the evidence suggests this is not the case – monkeys exhibited substantial expansion (detectable in each individual monkey) of the grey matter, including above cortical areas, during only a two-week tool-use training period (Quallo et al. 2009). This directly proves the phenomena previously suggested (and detected statistically in groups) that humans who are experts in certain cognitive domains tend to have slightly thicker grey matter in the part corresponding to the area subserving such mental functions. Concrete biological and genetic mechanisms realizing this expansion could be studied using the monkey paradigm in the near future.

Once a novel, alternative, bistable state was found to be useful, additional resources will be invested to stabilize the system, perhaps allowing further redundancy. Humans can induce such expansion intentionally, to create a better, more comfortable environmental niche. Subsequently, triggered by (extra-genetic, or epigenetic) factors embedded in such an environment, the corresponding neural niche in the brain could be reinforced further – thus, comprising recursive intentional niche construction (Iriki & Sakura 2008). Indeed, human-specific cognitive characteristics (or polysemous bias) seem to be subserved mainly by these "expanded" brain areas (Ogagwa et al. 2010; in press). Some aspects of recently evolved cognitive functions resulting from such neural reuse could be the mind (as described above; Iriki 2006), language, or culture, all of which contribute remarkably to semantic inheritance of the benefits acquired during the unusually elongated human post-reproduction period. "Thus, the theory suggests a novel pathway by which *Homo sapiens* may have achieved its current high-level cognitive capacities" (target article, sect. 6.3, para. 4).

Multi-use and constraints from original use

Justin A. Jungé and Daniel C. Dennett
Center for Cognitive Studies, Tufts University, Medford, MA 02155.
ddennett@tufts.edu   junge@wjh.harvard.edu
http://www.tufts.edu

Abstract: Anderson's theory is plausible and largely consistent with the data. However, it remains underspecified on several fronts, and we highlight areas for potential improvement. Reuse is described as duplicating a functional component, preserving one function and tinkering to add another function. This is a promising model, but Anderson neglects other reasonable alternatives and we highlight several. Evidence cited in support of reuse fails to uniquely support it among a broader set of multi-use theories. We suggest that a more stringent criterion for direct support of reuse may be satisfied by focusing on previous adaptive functions (original use).

On the whole, Anderson's theoretical framework appears plausible and advances a flexible computational architecture for brains. Although this framework works well in the abstract, there are several points for further refinement and investigation.

Our first suggestion is to better constrain the concept of reuse in order to set clear criteria for evidential support. One way to do this is by focusing on previous adaptive functions, original use. Until we have some sense of the functions that specific parts were optimized to perform in the past, it remains unclear how such parts might (or might not) be reused. Reuse promises (among other things) to go beyond original use. But how do former functions of neural components constrain the possibilities for reuse, if at all? Anderson is largely silent on this account, perhaps advantageously at this stage. Casting the theory abstractly leaves plenty of room for it to be generally accurate, and avoids objections to uncertain particulars. However, filling in more details will eventually be required for the theory to gain explanatory and predictive traction.

Anderson's discussion of modularity could benefit from additional examples, narrowing in the specific thesis of reuse. Modularity is a versatile – perhaps too versatile – concept. Is "massive modularity" a thesis about the size (crudely analogous to mass) or scale of the modules, the large number of modules (whatever their size), or the ubiquity of modular architecture in brains? Carruthers' (2006) comparison with hi-fi components may have misled Anderson. A better parallel might be the random number generator and the graphics processing card in a laptop, which can vary independently, and interact in many different applications. However, probably any parallel with technological modules is of very limited utility, since no such module exhibits the sorts of plasticity that neural tissue is known to enjoy. Sperber (1996), for instance, is a proponent of modularity, but he insists that modules are there to be exploited to meet new demands. Anderson might categorize Sperber's (1996; 2001) views as more closely aligned with reuse than massive modularity, but this suggests fuzzy boundaries between modularity and potential alternatives. A software theory of massive modularity – programs evolved to serve particular adaptive functions within brains – without commitments about implementation (unlike anatomical modularity), could survive Anderson's critique largely untouched. The grain and level of analysis where modularity is applied can make all the difference.

An important point for clarification concerns Anderson's occasional conflation of two partially overlapping (classes of) hypotheses. Reuse and multi-use should be better distinguished. Reuse theories form a set of related hypotheses. Multi-use is a larger set, including cases where original function is lost, as well as cases where original function is preserved (preservation is a defining attribute of Anderson's reuse theory). The term "reuse" strongly suggests exaptation, and Anderson is explicit that his reuse differs from typical exaptation by proposing that components continue to serve some previous adaptive function while *also* becoming available to "time share" new functions (though he doesn't put it in exactly those terms). Anderson takes the multiplicity of functions – a brain area being activated by multiple different tasks – as evidence for reuse. However, if multi-use is an available move in design space, what reason do we have to assume that original function is preserved? Without preserving original function, reuse is an inaccurate account, and adaptation to multi-use is more accurate. The case for multi-use is strong, but all of the evidence cited implicating multi-use, while consistent with the reuse hypothesis, is not *evidence* for the more specific hypotheses of reuse. This ties in with our first point. Until the original use of components is specified, along with examples, Anderson hasn't yet made the strong case for reuse.

To illustrate our suggestion that Anderson's theory should be fleshed out with details, we conclude with a specific example. As mentioned above, the picture of reuse that Anderson offers appears analogous to a time-sharing model: (1) At any given time, one high-level process uses the "workings" of multiple lower-level areas, and (2) numerous high-level processes are hypothesized to alternately access a common pool of specialized lower-level resources. While this account may be accurate, we wish to highlight an alternative that focuses on a finer mechanical

grain, such as individual neurons (or perhaps small collections of neurons, such as minicolumns). It is possible that specialized brain areas contain a large amount of structural/computational redundancy (i.e., many neurons or collections of neurons that can potentially perform the same class of functions). Rather than a single neuron or small neural tract playing roles in many high-level processes, it is possible that distinct subsets of neurons within a specialized area have similar competences, and hence are redundant, but as a result are available to be assigned individually to specific uses (similar to the way that redundancies due to gene duplication provide available competences for reassignment, leaving one copy to perform the original function). Over development or training, subsets of neurons in a specialized brain area could then be recruited for involvement in distinct high-level processes. This model emphasizes multi-potential of neurons, but single-use of individual neurons, as determined in the course of development and learning.

In a coarse enough grain, this neural model would look exactly like multi-use (or reuse). However, on close inspection the mechanism would be importantly different. In an adult brain, a given neuron would be aligned with only a single high-level function, whereas each area of neurons would be aligned with very many different functions. This model of multi-potential and single-use may account for all the data that Anderson cites in support of reuse, and it also avoids time-sharing for specific neurons. Whether or not the model sketched here is accurate, it illustrates the kind of refinement that could make Anderson's abstract theoretical proposal more concrete, and perhaps subtly improved.

Comparative studies provide evidence for neural reuse

Paul S. Katz
Neuroscience Institute, Georgia State University, Atlanta, GA 30302-5030.
pkatz@gsu.edu
http://neuroscience.gsu.edu/pkatz.html

Abstract: Comparative studies demonstrate that homologous neural structures differ in function and that neural mechanisms underlying behavior evolved independently. A neural structure does not serve a particular function so much as it executes an algorithm on its inputs though its dynamics. Neural dynamics are altered by a neuromodulation, and species-differences in neuromodulation can account for behavioral differences.

Anderson begins his article by quoting one of Darwin's explanations about how homologous structures can differ in function across species. Such a realization was clear even to Richard Owen who, although not accepting Darwin's theory of evolution, defined homology as "the same organ in different animals under every variety of form and function" (Owen 1843). It is therefore surprising that Anderson uses very little comparative data to support his theory of neural reuse through "massive redeployment." Comparative research examining neural circuitry across species, which has led to important insights into the evolution of neural circuits, needs to be included in any global theory about the evolution of human cognitive abilities. By concentrating solely on humans and extending analogies only to primates, one misses the strength of the comparative approach. Evolutionary principles can be generalized across species; humans are not more special for their cognitive abilities than bats are for their sonar abilities or song birds are for vocal learning abilities. Even the more distantly related invertebrates can provide lessons about how nervous systems evolved.

As a structure, the cortex is very adaptable; similar circuitry can be used for different functions. For example, in the

absence of auditory afferents, primary auditory cortex in ferrets can be experimentally induced to process visual information (Sur et al. 1988), and the ferrets respond to visual stimuli as being visual in nature (Von Melchner et al. 2000). Such a situation may occur naturally in congenitally blind humans; primary visual cortex, which lacks visual input, is instead responsive to somatosensory input and is necessary for reading Braille (Cohen et al. 1997). Therefore, the "function" of cortex is very much determined by the inputs that it receives. It may be better to refer to the algorithm that cortex performs on its inputs than on its innate function.

Because of cortical plasticity, it can be problematic to call one area of cortex "homologous" to a region in other species based on its function (Kaas 2005). Evidence suggests independent evolution of higher-order cortical areas, indicating that there may be innate directions for evolutionary change (Catania 2000; Krubitzer 2007; 2009; Padberg et al. 2007). In discussing the "neuronal recycling hypothesis," Anderson refers to changes following tool training in an area of the macaque brain that is "roughly homologous to the regions associated with tool-use in the human brain" (sect. 6.3, para. 4). It is difficult to develop any theory about the evolution of a structure without being able to unambiguously identify homologous structures in other species.

Homology of neural structures can be more precisely determined in invertebrates, where individual neurons are uniquely identifiable and can be recognized as homologous across species (Comer & Robertson 2001; Croll 1987; Meier et al. 1991). This allows the role of homologous neurons across species exhibiting different behaviors to be assessed. For example, homologous neurons in nudibranch molluscs have different effects and are involved differently in the production of different types of swimming behavior (Newcomb & Katz 2007; 2008). There is also evidence to suggest that homologous neurons have independently been incorporated into circuits that perform analogous swimming behaviors (Katz & Newcomb 2007). This is reminiscent of the reuse of cortical areas across mammals for similar tasks (Catania 2000; Krubitzer 2007; 2009; Padberg et al. 2007). Thus, a corollary of neuronal reuse may be that constraints on neuronal structure preclude some potential avenues and allow evolution to proceed in only particular directions, which leads to reuse.

Work on invertebrates has established the existence of multifunctional neural circuits, in which the same set of neurons in a single animal produces different types of behaviors at different times (Briggman & Kristan 2008). One mechanism for shifting activity of neurons is neuromodulatory inputs, which alter cellular and synaptic properties (Calabrese 1998; Katz 1999; Katz & Calin-Jageman 2008; Marder & Thirumalai 2002). This has been particularly well studied in circuits that produce rhythmic motor patterns. Cortex has been likened to such a circuit in that it can exhibit different dynamic activity states depending upon its neuromodulatory input (Yuste et al. 2005). It has been proposed that phylogenetic differences in neuromodulation could be a mechanism by which neural circuits exhibit different behaviors across species (Arbas et al. 1991; Katz & Harris-Warrick 1999; Meyrand et al. 2000; Wright et al. 1996). This would allow core functions of a neural circuit to remain intact, while enabling the circuit to produce different dynamic states, corresponding to the neural exploitation theory.

A nice example of changes in neural modulation that leads to large changes in behavior has been documented in the social behavior of voles (Donaldson & Young 2008; McGraw & Young 2010). Prairie voles pair-bond after mating, whereas meadow voles do not. In addition to displaying partner preference, pair-bonding involves a number of complex behavioral traits, including increased territoriality and male parental care. The difference in the behavior of male voles can largely be accounted for by the neural expression pattern of vasopressin V1a receptors. These receptors are highly expressed in the ventral pallidum of prairie voles, but not in non-monogamous species. Using viral

gene expression to express the V1a receptor in the ventral forebrain of the meadow vole substantially increased its partner-preference behavior (Lim et al. 2004).

The evolutionary mechanism for differences in gene expression patterns in voles has been traced to an unstable stretch of repetitive microsatellite domains upstream from the coding region of the V1a receptor gene (Hammock & Young 2005). Although similar genetic mechanisms do not play a role in the expression pattern in primates (Donaldson et al. 2008), monogamous primate species such as the common marmoset display high levels of V1a receptor expression in ventral forebrain regions, whereas non-monogamous species such as rhesus macaques do not (Young 1999). This suggests that similar social behaviors have arisen independently through changes in the expression of V1a receptors in the ventral forebrains of rodents and primates. Once again, this supports the neural exploitation model: The basic connectivity of the brain has not been altered; instead, there is change in the expression of a particular receptor, which can modulate the dynamics of the activity through that area. The ventral forebrain areas are involved in more than pair-bonding; they also play a role in addiction and reward-based learning (Kalivas & Volkow 2005; Schultz et al. 1997). Pair-bonding results from these types of reward-learning processes being applied to a mate. This further supports the neural exploitation theory.

Anderson expresses several ideas relating to the "age" of a particular brain area influencing its ability to undergo evolutionary change. This notion smacks of *Scala Natura* because it assigns youngest age to structures that are found in humans and not in other animals. The fallacy of this line of thinking can be seen with the above example. By all accounts, the ventral forebrain areas predate mammals. Yet, even closely related voles can exhibit behavioral differences caused by evolutionary change to this "older" region of the forebrain. Furthermore, the ventral forebrain area is also involved in learning in birds (Jarvis et al. 2005; Perkel 2004).

In summary, comparative studies offer important insights into how brains evolved. There are surely many mechanisms that can be found. It is clear, however, that assigning a function to a particular brain structure is a gross simplification and can lead to false conclusions about its evolution. Neural circuitry is multifunctional and dynamic. Anything that changes the dynamics of the circuit will alter the behavioral output.

No bootstrapping without semantic inheritance

Julian Kiverstein

School of Philosophy, Psychology, and Language Sciences, University of Edinburgh, Edinburgh EH8 7PU, Scotland, United Kingdom.
j.kiverstein@ed.ac.uk
http://www.artandphilosophy.com/philosophy.html

Abstract: Anderson's massive redeployment hypothesis (MRH) takes the grounding of meaning in sensorimotor behaviour to be a side effect of neural reuse. I suggest this grounding may play a much more fundamental role in accounting for the bootstrapping of higher-level cognition from sensorimotor behaviour. Thus, the question of when neural reuse delivers semantic inheritance is a pressing one for MRH.

Evolution has devoted by far and away the largest part of its history to building organisms that can move around in a dynamic environment, sensing their environments in ways conducive to their own survival and reproduction (Brooks 1991). The challenge to cognitive scientists is to explain how the strategies organisms use to solve these basic problems of perception

and action scale up to the strategies humans use in solving abstract higher-level problems. I call this the "bootstrapping challenge." Embodied cognitive science offers a programmatic response to the bootstrapping challenge that attempts to show how high-level problem solving might have been built upon the foundation of a substrate of perception and sensorimotor control. The ascent from sensing and moving to thinking, planning, and language understanding is an incremental and gradual one, and a key strategy may have been the redeployment of sensorimotor capacities to perform high-level cognitive tasks.

Anderson has done the embodied cognition community the enormous service of framing a global hypothesis about how these incremental changes might have taken place in our brains over the course of evolution. The central claim of his massive redeployment hypothesis (MRH) is that more recent cognitive functions such as those involved in abstract problem solving might have their origin in the reuse of evolutionarily older neural circuits that served biologically basic functions. In this commentary, I want to take up Anderson's claim that the principle guiding reuse is "functional inheritance" and not "semantic inheritance." By "semantic inheritance," I mean the kind of relation that concept empiricists and conceptual metaphor theories take to hold between concepts and sensorimotor representations. What connects both theories is the use of our experience and competence in one domain to guide our thinking in a distinct domain. Anderson describes very many instances of neural reuse that do not obviously involve the sensorimotor system, and hence do not involve semantic inheritance. He takes this to show that semantic inheritance may be a "side effect" (see sect. 4.6) of neural reuse. I will argue that it is only when reuse is accompanied by semantic inheritance that you find any bootstrapping from low-level cognitive functions to high-level cognitive functions. This follows from an argument Anderson himself makes against Susan Hurley's (2008) shared circuits model (SCM). Therefore, the question of what kinds of reuse support semantic inheritance (a question Anderson himself raises in sect. 7) becomes a particularly pressing issue for the embodied cognition research programme. I will finish up by suggesting that neural reuse and semantic inheritance may actually be much more closely tied than Anderson suggests.

We can see how semantic inheritance is required for bootstrapping by considering Anderson's discussion of Susan Hurley's (2008) shared circuits model (SCM). The model is complex, and I shall restrict my discussion to layer 3 of SCM, which describes how functional mechanisms used to predict sensory feedback in the control of motor behaviour might be reused to "simulate" the motor processes that stand behind the observed behaviour of another. This simulation is hypothesised to take the form of "mirroring" that can underwrite the copying of instrumental behaviour either in the form of priming, emulation, or imitation. Anderson worries that the inputs and outputs required for mirroring are "impoverished" and "abstract" when compared to those inherited from layer 2. When I perform an action myself, for instance, the action is represented from my own point of view. Anderson supposes that when I observe another's action, I must represent the other's action from a third-person point of view. Hence, the progression from layer 2 to layer 3 would seem to require a translation of a first-person representation of action into a third-person representation. Without some explanation of how this translation gets effected, we will not have shown how high-level cognitive abilities like imitative learning can have their basis in the reuse of low-level sensorimotor representation.

This problem Anderson has identified for SCM would however seem to apply equally to MRH. What allows the control mechanisms found at layer 2 to be reused at layer 3 are the functional properties of those control mechanisms. According to MRH, it is a neural region's functional properties that allow a region used in one domain to get reused in a distinct domain. The inheritance of functional properties falls some way short of

guaranteeing semantic inheritance. Functional inheritance doesn't on its own explain the abstraction and informational impoverishment you find as you move from lower-level sensorimotor behaviour to higher-level cognition. If this is right, it seems to follow that neural reuse won't suffice for bootstrapping.

Hurley's SCM may however have resources for responding to this problem that are different from those outlined by Anderson in his target article. What is missing from Anderson's framing of the problem is any mention of the sensorimotor associations that drive the predictions at layers 2 and 3 of SCM. Predictions of the sensory effects of movement are possible at layer 2 only because the motor system has learned that movements of a given type are correlated with certain sensory effects. Hurley followed Cecilia Heyes in thinking of this learning as arising in development through associations that wire sensory neurons (in superior temporal sulcus, for example) together with motor neurons (in premotor and parietal cortices; see Heyes [2010] for more on a recent presentation of this hypothesis). Crucially, Hurley is assuming that the sensory inputs from one's own movement and from the movement of others are similar enough for sensory neurons to respond to both without distinguishing them. Thus, sensorimotor associations can underwrite an "inference" from the sensory effects of observed behaviour to the motor processes that tend to cause behaviour. In this way, sensorimotor associations can be used both to control the sensory effects of movement and to simulate the movements that have similar sensory effects when carried out by others. For SCM then, it is associative learning that delivers the kind of semantic inheritance required for bootstrapping.

I finish by drawing a tentative moral for MRH. The functional inheritance that underpins neural reuse might bear cognitive fruit only when it is accompanied by semantic inheritance. Reuse of functional mechanisms in SCM is understood as simulation that establishes a space of shared meaning. Semantic inheritance, as appealed to in concept empiricism and conceptual metaphor theories, is also naturally understood as a form of simulation which opens up a space of shared meaning. While neural reuse could well turn out to be a "fundamental organisational principle" of the brain, the pressing question that remains is how neural reuse could deliver a shared space of meaning of a kind that supports bootstrapping.

Redeployed functions versus spreading activation: A potential confound

Colin Klein

Department of Philosophy, University of Illinois at Chicago, Chicago, IL 60607.
cvklein@uic.edu
http://tigger.uic.edu/~cvklein/links.html

Abstract: Anderson's meta-analysis of fMRI data is subject to a potential confound. Areas identified as active may make no functional contribution to the task being studied, or may indicate regions involved in the coordination of functional networks rather than information processing per se. I suggest a way in which fMRI adaptation studies might provide a useful test between these alternatives.

That there is a many-to-one mapping between cognitive functions and brain areas should now be beyond dispute. The tricky part is figuring out what to say about it. Anderson's massive redeployment hypothesis (MRH) is a plausible position in the debate. Good engineers often find new uses for old tricks; we should expect nature to be no less clever.

A crucial piece of evidence for the MRH is Anderson's impressive meta-analyses of fMRI experiments (Anderson 2007b; 2007c). These show that phylogenetically older areas

tend to be more active, across a variety of tasks, than phylogenetically newer ones. Crucially, Anderson assumes that the areas identified as active make a functional contribution to the experimental tasks being studied. That is often assumed in fMRI experiments, and so may seem unproblematic. This assumption is subject to a potential confound, however, and one that becomes especially troublesome when doing large-scale meta-analyses.

The BOLD response on which fMRI depends is a measure of physiological change. *Which* physiological change fMRI tracks is a matter of considerable debate. There is increasing evidence that the BOLD response better tracks regional increases in synaptic activity, rather than increased output of action potentials (Logothetis et al. 2001; Nair 2005, sect. 2.2 reviews; Viswanathan & Freeman 2007). Crucially, this means that observed BOLD activity may represent a mix of both excitatory and inhibitory inputs. A region which receives subthreshold excitatory input, or one which is both excited and inhibited enough to suppress further activation, may nevertheless show a measurable – even strong – BOLD response (Logothetis 2008). However, these "active" regions would make no *functional* contribution to the experimental task.

Hence the potential confound. The fact that phylogenetically older areas are more often active may be explained by redeployment. It may also be explained by assuming that older areas simply receive more input than do newer ones. This potential confound may be manageable in individual fMRI experiments. Meta-analyses increase statistical power, however, making even small effects more likely to be noticed. Further, meta-analyses necessarily lack the fine-grained detail that might normally allow these functional by-products to be explained away.

This is not a merely academic worry. To give one example: Mahon and Caramazza (2008) recently reviewed the fMRI evidence for the sensorimotor account of conceptual grounding (including many of the studies reviewed by Anderson in sect. 4). They conclude that the evidence is consistent with a view on which the semantic analysis of a sentence activates motor areas as an inevitable consequence of spreading activation within a complex neural system. Hence, although the motor system may often be activated during semantic analysis tasks, this activation need not represent a functional contribution to semantic analysis itself. It would instead be the natural consequence of a system in which the typical *consumers* of representations were primed for action, but inhibited (or simply under-excited) if their further, functionally specific, contribution was unnecessary. Note that a reliance on subtraction-based imaging does not obviate this problem: distinct semantic terms may well prime distinct motor regions.

Spreading activation and massive redeployment are not mutually exclusive hypotheses. Indeed, it seems to me that the redeployment model should accept some version of the former. If the brain does consist of pluripotent regions that flexibly combine into functional networks, problems of coordination – and especially the necessity of inhibiting preponent but contextually inappropriate dispositions – become paramount. Further, phylogenetically newer areas evolved in the context of organisms which already had well-functioning brains. We should expect newer areas to project heavily to older areas, both because the information they provide might be relevant to these older adaptive repertoires and because those older functions will need to be coordinated in light of newer capacities.

The crucial question, then, is how we might get experimental evidence that favors redeployment over the alternatives. Anderson suggests several plausible possibilities for testing his hypothesis. I suggest a further possibility: the use of fMRI adaptation. This technique exploits the fact that recently active neurons tend to show a decreased response to further stimulation; a decreased BOLD response across experimental conditions thus provides evidence that a region is making the same contribution to both tasks. Adaptation would allow one to distinguish areas

which are truly redeployed from those which have simply parcellated into functionally specific areas that are smaller than the resolution of fMRI (an open evolutionary possibility; Streidter 2005, Ch. 7 reviews). Further, adaptation would allow us to distinguish areas that are truly reused from areas that are involved in the coordination of complex networks.

Crinion et al. (2006) used this technique to distinguish the contribution of various cortical and subcortical areas in language processing. Proficient bilingual speakers showed both within- and cross-language priming in the left anterior temporal lobe, suggesting a shared substrate for semantic information (and thus supporting a form of reuse). Activation in the left caudate, in contrast, did not show a priming effect. This supports a hypothesized role for the caudate in language *control*: Plausibly, the caudate helps inhibit contextually inappropriate responses, a real problem when distinct languages partially share the same substrate. fMRI adaptation might thus allow us to disentangle the contribution of frequently activated areas in a variety of tasks, and so provide a further test of Anderson's intriguing hypothesis.

Implications of neural reuse for brain injury therapy: Historical note on the work of Kurt Goldstein

Barry Lia

Dizziness and Balance Center, Otolaryngology/Head and Neck Surgery, University of Washington Medical Center, Seattle, WA 98195-6161.
barrylia@uw.edu

Abstract: This commentary suggests how the target article raises new implications for brain injury therapies, which may have been anticipated by the neurologist Kurt Goldstein, though he worked in and earlier era of fervent localization of brain function.

I first took interest in Anderson's article by dint of the notion that neural circuits established for one purpose may be exapted for new functions during evolution and development. In a previous BBS commentary (Lia 1992), I had proposed an exaptation of the peripheral visual system for the adaptive evolution of enactive focal vision and praxic use of the forelimbs in primates, a crucial feature of our cognitive niche. I applaud Anderson's discussions of the co-determination of organism and environment and of the idea of "neural niche" within the organism itself as most welcome for cognitive science.

But it was implications for therapies for brain injury – which this article raises in closing – which brought to mind the work of Kurt Goldstein (1963) for comment. Anderson refers to "network thinking" which "suggests one should look for higher-order features or patterns in the behavior of complex systems, and advert to these in explaining the functioning of the system" (sect. 3.1, para. 6). Writing in 1939, Goldstein was a neurologist primarily occupied with patients' recovery from brain injury. Similar to Anderson, Goldstein was also concerned with method in biological research and ways of conceptualizing the empirical material. The influence of the Gestalt school of psychology upon Goldstein is reflected in the following passage from Goldstein (1963), which refers to the "figure" of a performance:

> Localization of a performance no longer means to us an excitation in a certain place, but a dynamic process which occurs in the entire nervous system, even in the whole organism, and which has a definite configuration for each performance. This excitation configuration has, in a certain locality, a special formation ("elevation") corresponding to the figure process. This elevation finds its expression in the figure of the performance. A specific location is characterized by the influence which a particular structure of that area exerts on the total process, i.e., by the contribution which the excitation of that area, by virtue of its structure, makes to the total process. (pp. 260–61)

This foreshadows the dynamic view of functional recruitment and brain organization which neural reuse theories present. Goldstein would likely have appreciated Anderson's hope that "Knowledge about the *range* of different tasks that potentially stimulate each region [akin to Goldstein's notion of 'excitation configuration'] may serve as the basis for unexpected therapeutic interventions, ways of *indirectly* recovering function in one domain by exercising capacities in another" (sect. 7, para. 8, emphasis Anderson's). Such specific knowledge of the "excitation configuration" was unknown and unavailable to Goldstein; he could only infer it. But by taking a holistic, organismal perspective, somewhat akin to Anderson's "network thinking," Goldstein intuited such an understanding and postulated such an indirect recovery of function in his work with rehabilitation of brain injury. Goldstein's outlook echoes Anderson's "call for an assimilative, global theory, rather than the elaboration of existing theoretical frameworks" (sect. 5, para. 7). This target article may point toward advances which a Goldstein would be striving toward today had he had our modern tools for studying the brain and cognitive function.

Reuse in the brain and elsewhere

Björn Lindblom
Department of Linguistics, Stockholm University, 10691 Stockholm, Sweden.
lindblom@ling.su.se   **http://www.ling.su.se**

Abstract: Chemistry, genetics, physics, and linguistics all present instances of reuse. I use the example of how behavioral constraints may have contributed to the emergence of phonemic reuse. Arising from specific facts about speech production, perception, and learning, such constraints suggest that combinatorial reuse is domain-specific. This implies that it would be more prudent to view instances of neural reuse not as reflecting a "fundamental organizational principle," but as a fortuitous set of converging phenomena.

Hallmark of true phonology. It is easy to forget that the words we use everyday are built rather "ingeniously." They code meanings in a combinatorial way, arranging a limited number of phonetic properties in various combinations (phonetic segments, phonemes) and permutations (syllables, morphemes, words). This method of reuse provides tremendous expressive power and creates the means for developing large and open-ended vocabularies. In the realm of animal communication, it appears to be unique to humankind.

How did it come about? Combinatorial structure is hardly a product of humankind's ingenuity, a cultural invention. It is more likely to have evolved. But how? Is it an idiosyncrasy pre-specified in our genetic endowment for language? Or did performance factors drive language towards phonemically structured signals? If, as Anderson claims, neural reuse is a general principle of brain organization, did perhaps this process play a role in the emergence of linguistic reuse?

On-line speaking. Assuming that phonetic reuse evolved from existing capacities, we are led to ask: What were those capacities? Recent work (Lindblom et al., in press) suggests three factors. The first two identify general characteristics of motor control (not specific to speech). The third highlights expressive needs arising from the growth of humankind's cognitive capacity.

1. Positional control (targets > discrete units).
2. Motor equivalence (movement trajectories > recombination).
3. Cognitive processes (expressive needs > sound-meaning link > vocabulary size).

Voluntary non-speech motions are output-oriented, that is, organized to produce desired results in the subject's external environment. So is speech. Experiments indicate that speech movements are controlled by commands specifying a series of positions (targets) in articulatory space. Goals can be attained from arbitrary initial conditions, and the system compensates in response to obstacles and perturbations. Transitions between targets are typically smooth and show stable velocity profiles reminiscent of point-to-point reaching motions. We conclude that speech is in no way special. Both speech and non-speech show positional (target-based) control and motor equivalence (mechanisms for deriving trajectories from arbitrary initial to arbitrary final locations within the work space).

A difference worth pointing out is that, since we speak to be understood, perceptual factors play a role in determining the extent to which targets are reached. But, although information dynamics may modulate the speaker's performance (cf. clear/casual speech), its motor organization is basically the same. Significantly, "target" is a context-independent notion, whereas its associated articulatory movements are highly context-sensitive.

Evo/devo implications. The above account implies that the end-state of phonetic learning is a mastery of targets and motor equivalent trajectory formation. What the learner does in imitating ambient speech is to find the sparsest way of activating the dynamics of the speech effectors. Using a least-action strategy, the child residually ends up with targets.

The context-free nature of target implies that once a target is learned in one context, it can immediately be recruited in another. There lies the key to reuse in the present account. Learning targets speeds up the acquisition process, compared with learning contextually variable movements. For evolution, this means that lexical inventories that are phonemically coded are easier to learn than systems consisting of Gestalt (holistic) sound patterns. Seen in this light, phonetic reuse appears to be an adaptation linked to ease of acquisition. If discrete units are to be derived from positional control and recombination from motor equivalence – two general cross-species motor characteristics – we must ask why other animals do not end up speaking.

This is where the third factor comes in. Humankind's cognitive capacity has developed dramatically from skills not unlike those of present-day apes. It makes it possible to use language to encode a virtually infinite set of meanings. For an account of how that may have happened, see Donald's (1991) synthesis of a broad range of evidence. Donald assumes that, as gestural messages grew more elaborate, they eventually reached a complexity that favored faster and more precise ways of communicating. The vocal/auditory modality offered an independent, omnidirectional channel useful at a distance and in the dark. It did not impede locomotion, gestures, or manual work. The vocal system came to be exploited more and more as the growing cognitive system pushing for lexical inventions and sound-meaning pairs. The reuse capability implicit in discrete targets and motor equivalence conveniently provided the expressive means for these growing semantic abilities to interact in a process of mutual reinforcement. Accordingly, the reason why no other species has extensive reuse lies in the felicitous convergence of all three factors. According to the present account, one would expect reuse *not* to be limited to the vocal/auditory modality. The formal organization of sign language corroborates that prediction.

Neural reuse: Organizational principle or widespread phenomenon? While it may be the case that true phonology is uniquely human, combinatorial reuse is known to occur in other domains. Studdert-Kennedy (2005) draws attention to the work of Abler (1989) who "recognized that a combinatorial and hierarchical principle is a mathematically necessary condition of all natural systems that 'make infinite use of finite means', including physics, chemistry, genetics, and language. He dubbed it 'the particulate principle'." (Studdert-Kennedy 2005, p. 52).

I take the word "principle" here to be used descriptively, rather than as referring to the possibility that there is a hidden abstract formal condition to be discovered which can be used for explaining all instances of combinatorial and hierarchical

coding. In other words, each case of reuse is likely to have its own history.

Which takes us back to neural reuse. If the central nervous system (CNS) exhibits massive reuse of neural circuitry, we may, as Anderson does, choose to talk about a fundamental organizational principle of the brain. Or we might prefer saying that massive reuse of neural circuitry is a widespread phenomenon, bearing in mind that every example of reuse may have its own story.

Let us redeploy attention to sensorimotor experience

Nicolas Michaux,[a] Mauro Pesenti,[a] Arnaud Badets,[b] Samuel Di Luca,[a] and Michael Andres[a]

[a]*Institut de Recherche en Sciences Psychologiques, Université catholique de Louvain, 1348 Louvain-la-Neuve, Belgium;* [b]*Centre de Recherches sur la Cognition et l'Apprentissage, CNRS UMR 6234, France.*

nicolas.michaux@uclouvain.be   **mauro.pesenti@uclouvain.be**
arnaud.badets@univ-poitiers.fr   **samuel.diluca@uclouvain.be**
michael.andres@uclouvain.be
http://www.uclouvain.be/315041.html
http://cerca.labo.univ-poitiers.fr

Abstract: With his massive redeployment hypothesis (MRH), Anderson claims that novel cognitive functions are likely to rely on pre-existing circuits already possessing suitable resources. Here, we put forward recent findings from studies in numerical cognition in order to show that the role of sensorimotor experience in the ontogenetical development of a new function has been largely underestimated in Anderson's proposal.

With his massive redeployment hypothesis (MRH), Anderson proposes an attractive view of neural reuse by claiming that neural circuits initially dedicated to a specific function can be reused in the course of human evolution to support novel cognitive functions. Because this is meant to occur whenever a pre-existing circuit already possesses useful mechanisms for a novel function, Anderson's proposal challenges the assumption of concept empiricism that neural reuse is causally related to sensorimotor experience. Here, we question the idea that the mere availability of neural resources is sufficient to explain how new functions emerge from neural reuse, and we highlight the role of sensorimotor experience during the ontogenetical development of a new function by reviewing recent findings from studies in numerical cognition.

In the past few years, finger control and numerical cognition have been shown to share common areas in the parietal and premotor cortices (Andres et al. 2007; Pesenti et al. 2000; Zago et al. 2001). This common ground for finger movements and number processing may be a developmental trace of the use of fingers when learning to count (Butterworth 1999a). In contrast, Anderson and colleagues (see Penner-Wilger & Anderson 2008) propose that the neural network originally evolved for finger representation has been redeployed to serve numerical cognition only because it offers suitable resources to represent numbers, such as a register made of switches that can be independently activated. Accordingly, sensorimotor experience would play no role in the development of numerical cognition. However, a growing body of empirical data makes this perspective untenable. Indeed, finger use was found to deeply impact the acquisition of numerical skills in at least four different ways.

First, developmental studies indicate not only that children with poor abilities to discriminate their fingers are more likely to experience difficulties in mathematical tests (Fayol et al. 1998; Noel 2005), but also that an extensive training in finger differentiation, via sensorimotor exercises, improves both finger

gnosis and numerical abilities (Garcia-Bafalluy & Noël 2008). This shows that sensorimotor experience critically contributes to reaching an optimal performance during the acquisition of new numerical skills and, more generally, to making neural reuse effective in supporting new functions.

Second, a cross-cultural brain imaging study with participants from Eastern and Western cultures showed that cultural and educational habits can shape neural resources (Tang et al. 2006). Various numerical tasks activated similar networks in occipito-parietal, perisylvian, and premotor areas in both cultures, but English participants showed higher activity in the perisylvian areas, whereas Chinese participants showed higher activity in premotor areas, a finding difficult to explain unless one considers their high level of practice in calculations with an abacus which requires a fine control of finger movements (Cantlon & Brannon 2007; Tang et al. 2006). The cerebral network underlying numerical cognition can thus be shaped by the constraints that culture and/or education exert on the way individuals physically represent and manipulate numbers, thereby providing key evidence against the deterministic view conveyed by the MRH.

Third, even if Anderson's proposal makes it clear why pre-existing neural resources may underlie new representations, such as numbers, it remains confusing how these representations acquire their conceptual meanings. The idea that number semantics could also pre-exist in the brain is still disputed (see Rips et al. 2008; and our comment, Andres et al. 2008). We argue that the use of finger counting can account for conceptual properties of numbers that are left undefined in the initial redeployment of pre-existing neural resources. For instance, the stable sequence of finger movements performed by children while counting, presumably under the combined influence of motor constraints and cultural habits, may lead them to understand that natural numbers include a unique first element, and that each number in a sequence has a unique immediate successor and a unique immediate predecessor, except the first (Wiese 2003). This suggests that neural reuse involves domain-structuring inheritance, as predicted by concept empiricism, but not by a strong version of the MRH.

Furthermore, the recurrent use of a stable finger-counting strategy during childhood keeps on influencing the way numbers are represented and processed in adults. Indeed, we recently showed that, when participants are asked to identify Arabic digits by pressing keys with their ten fingers, a finger-digit mapping congruent with their prototypical finger-counting strategy leads to a better performance than any other mapping, suggesting that number semantics of educated adults is grounded in their personal experience of finger counting (Di Luca et al. 2006). The finding that, in long-term memory, the structure of newly acquired concepts reflects idiosyncratic aspects of sensorimotor experience challenges Anderson's proposal that neural reuse anticipates concept formation. One may argue that neural redeployment may constrain or predispose individuals to count the way they do. However, this alternative explanation cannot account for the multiplicity of finger-counting strategies observed across individuals and cultures (Butterworth 1999b; Wiese 2003). It is also incompatible with the results of an unconscious priming study showing that number semantics are linked not only to finger-counting, but also to finger-monitoring configurations (i.e., finger configurations used to show numerosities to other people; Di Luca & Pesenti 2008).

Finally, recent findings show that object-directed actions mediate some aspects of the functional relationship between fingers and numbers. For example, observing grip closure movements interferes with numerical magnitude processing, suggesting the automatic activation of a magnitude code shared by numbers and finger movements (Badets & Pesenti 2010). Critically, this interference is not observed when viewing non-biological closure movements, which suggests that it does not result from a general system for processing movement amplitude. This finding rather underlines the need to postulate a

grounding mechanism, as predicted by empiricist accounts only.

In conclusion, although pre-existing circuits might be reused to provide representational resources for novel functions, we propose that these resources remain insufficient, and possibly unspecified, without the involvement of sensorimotor experience. In order to obtain a more universal theory of neural reuse, future studies now have to clarify how representational resources are shaped by cultural and educational constraints and how they interact with the functions they support.

Neural reuse as a source of developmental homology

David S. Moore[a] and Chris Moore[b]

[a]Department of Psychology, Pitzer College and Claremont Graduate University, Claremont, CA 91711; [b]Department of Psychology, Dalhousie University, Halifax, NS B3H4J1, Canada.
dmoore@pitzer.edu   Chris.Moore@Dal.Ca
http://pzacad.pitzer.edu/~dmoore/
http://myweb.dal.ca/moorec/index.html

Abstract: Neural reuse theories should interest developmental psychologists because these theories can potentially illuminate the developmental relations among psychological characteristics observed across the lifespan. Characteristics that develop by exploiting pre-existing neural circuits can be thought of as developmental homologues. And, understood in this way, the homology concept that has proven valuable for evolutionary biologists can be used productively to study psychological/behavioral development.

Conventional wisdom in the neurosciences has long held that specific brain regions have specific functions. However, several recent studies have undermined the claim that cognitive functions can typically be mapped in straightforward ways to highly specialized brain areas, leading Anderson (2007c) to propose his massive redeployment hypothesis (MRH). In the target article, Anderson has considered his theory, along with others that posit similarly, that existing neural structures are normally reused/recycled/redeployed as new brain functions develop. This new approach has enormous potential for helping neuroscientists rethink the relationship between brain structures and their functions, as well as for helping those interested in the development and/or evolution of behavioral organization to understand changes in that organization across ontogeny and phylogeny.

Anderson uses the MRH to predict that a brain area's phylogenetic age should correlate with how often that area is deployed for various cognitive functions, and that a cognitive function's phylogenetic age should correlate with how localized that function is in the brain. However, although Anderson recognizes that neural reuse theories bear on questions of development, his article focuses on phylogeny to the virtual exclusion of ontogeny. Brief mentions of development are made, and a note points out that neural reuse "is broadly compatible with the developmental theories of Piaget" (target article, Note 10); but, in fact, neural reuse should interest all developmental psychologists because the approach is compatible with most current theories of development and could contribute to theoretical progress in the field in general. Anderson cites Dehaene's "neuronal recycling" theory as having potentially identified a "fundamental developmental . . . strategy for realizing cognitive functions" (sect. 1, para. 3); but, like other promissory notes in Anderson's text, this one is never fully redeemed. Neither Anderson nor Dehaene and Cohen (2007) fully consider the implications of neural reuse theories for understanding development.

The idea of neural reuse could have profound and general implications for the understanding of behavioral development.

In particular, we believe that neural reuse produces a type of developmental homology, and that just as evolutionary biology has profited from the discovery and analysis of evolutionary homologies (Hall 2003), so developmental psychology may profit from the identification of developmental homologies, some of which likely arise as a result of neural reuse. Because two or more psychological characteristics present at a given point in development might both (re)use neural circuits formed much earlier in development, thinking about such characteristics in terms of developmental homology could well illuminate their relationship to each other (as well as to other psychological characteristics present earlier in development that also depend on these circuits). Consequently, we believe that importing the concept of homology into developmental psychology has the potential to help behavioral scientists understand when, how, and why specific traits have common developmental origins.

Within biology, several types of homology have been identified, including among others (1) taxic homology (Griffiths 2007), in which characteristics in different species (e.g., bat wings and human forearms) have derived from a characteristic present in a common ancestor; (2) serial homology (Rutishauser & Moline 2005), in which parts of an individual organism are of the same type (e.g., the corresponding bones in a person's right hand and right foot, or any two vertebrae in mammalian spinal columns); and (3) ontogenetic homology (Hoßfeld & Olsson 2005), in which distinct individuals of the same species have differing features that nonetheless derive from common embryonic tissues (e.g., human ovaries and testes). Developmental homologies arising from neural reuse would be most similar to the kinds of homologies identified by Bertalanffy in 1934 (described in Hoßfeld & Olsson 2005), and would include pairs of psychological characteristics, both of which emerged from a common characteristic present earlier in development. In addition, much as human forearms are homologous to the forearms of extinct Australopithecines, psychological characteristics of adults could be recognized as homologues of psychological characteristics present in juveniles in various developmental stages. Such homologues *could* arise in ways that would not require neural reuse – after all, "a structure that is homologous across species can develop based on non-homologous genes and/or developmental processes, and vice-versa" (Brigandt & Griffiths 2007, p. 634) – but any characteristics known to emerge following the redeployment of a specific neural circuit would seem prima facie to be homologous, at least structurally if not functionally.

Several examples of possible developmental homologies may be identified. Temporal cognition in the form of episodic thinking develops later than spatial cognition and makes use of related conceptual structures (Clayton & Russell 2009). The discovery that these mental processes also make use of certain shared neural circuits would indicate that they are homologous, thereby shedding light on the nature of their developmental relationship. Linguistic structures, likewise, may well depend upon earlier-developing social interactive communicative structures. Tomasello (2003), for example, argues that syntax can be understood as a form of joint attention, a conceptualization that implies that these are homologous psychological characteristics, their different appearances notwithstanding. Still other psychological characteristics that appear similar across age have been assumed to be homologues, such as the neonatal imitation reported by Meltzoff and Moore (1977) and later-developing forms of imitation observed in older children and adults. Even so, studies of the neural circuits that contribute to neonatal and later imitation might or might not support this conclusion; a finding that adult imitation normally recruits neural circuits previously used during neonatal imitation would certainly qualify as support for the contention that these behaviors are homologous.

As Anderson suggests, neural reuse might be a fundamental organizational principle of the brain; and just as this idea can be used to formulate testable hypotheses about the evolution of

both the brain and its function, we think it could also influence the study of psychological development in significant ways. Similarly, importing the idea of homology from evolutionary biology into developmental psychology could help researchers conceptualize behavioral development in new and potentially informative ways. Taken together, the concepts of neural reuse and developmental homology could be used to further our understanding of brain development, psychological development, and the relationships between these phenomena.

Reuse of identified neurons in multiple neural circuits

Jeremy E. Niven[a] and Lars Chittka[b]

[a]Department of Zoology, University of Cambridge, Cambridge, CB2 3EJ, United Kingdom; [b]Research Centre for Psychology, School of Biological and Chemical Sciences, Queen Mary University of London, London E1 4NS, United Kingdom.

jen22@cam.ac.uk   l.chittka@qmul.ac.uk
http://www.neuroscience.cam.ac.uk/directory/profile.php?jen22
http://chittkalab.sbcs.qmul.ac.uk/

Abstract: The growing recognition by cognitive neuroscientists that areas of vertebrate brains may be reused for multiple purposes either functionally during development or during evolution echoes a similar realization made by neuroscientists working on invertebrates. Because of these animals' relatively more accessible nervous systems, neuronal reuse can be examined at the level of individual identified neurons and fully characterized neural circuits.

The principle of neural reuse is widespread within peripheral sensory and motor circuits in both vertebrates and invertebrates. Peripheral sensory circuits, such as those in the retina, extract and process information that is used in many behaviors. Indeed, the coding of visual scenes or odors requires that overlapping sets of sensory neurons are activated in response to different scenes or odors. Likewise, overlapping sets of premotor and motor neurons may be activated in disparate behaviors that require activation of overlapping sets of muscles.

The detailed characterization of invertebrate neurons and neural circuits has demonstrated that neurons can be reused to form neural circuits that perform multiple functions. One striking example comes from the stomatogastric ganglion (STG) of the crab *Cancer borealis*. The ~30 neurons of the STG control rhythmic muscle activity involved in chewing and digestion of food – the gastric mill and pyloric rhythms, respectively. Individual identified neurons may contribute to the production of more than one rhythm. The VD neuron, for example, is involved in the generation of both the gastric mill and pyloric rhythms (Weimann & Marder 1994). Thus, the dynamic restructuring of neural circuits within the STG provides a clear example of the reuse of neurons for the production of different behaviors.

Reuse may also be found in neurons involved in learning and memory. In the pond snail (*Lymnea stagnalis*), the breathing rhythm is generated by three synaptically connected neurons that form a central pattern generator. One of these neurons, RPeD1, is also necessary for many aspects of learning and memory; and removing the RPeD1 cell body can prevent the formation or reconsolidation of long-term memories (Sangha et al. 2003). In honeybees (*Apis mellifera*), a single identified neuron (VUMmx1) in the suboesophageal ganglion mediates the reward pathway in associative olfactory learning, but this neuron has also been implicated in learning phenomena as diverse as second-order conditioning and blocking (Menzel 2009).

The above examples emphasize that within the adult nervous system neurons are reused for different functions; but as Anderson points out, neurons may also be reused during development.

One such example is the reuse of larval motor neurons in the adult nervous system of the tobacco hornworm moth (*Manduca sexta*). *Manduca* caterpillars, like those of all moths and butterflies, undergo a metamorphosis that involves restructuring of the nervous system. Motor neurons that innervate leg muscles in the caterpillar have been shown to remodel their axons and dendrites during metamorphosis before innervating newly developed leg muscles (Kent & Levine 1993). Memories can also be retained between larval and adult forms of insects, despite the remodeling of neural networks during metamorphosis. For example, adult fruit flies (*Drosophila melanogaster*) retain memories of odors associated with aversive stimuli formed as third instar larvae (Tully et al. 1994). Memory retention between developmental stages suggests that those elements of neural circuits that are the loci of these stored memories are reused in adult animals.

Anderson also suggests that neurons may be reused during evolution, acquiring novel functions and possibly losing their original function. Again, invertebrate neural networks provide examples of such reuse during evolution. In the desert locust (*Schistocerca gregaria*), more than 20 interneurons have been identified from the neural networks controlling the flight muscles. Some of these interneurons have homologues in abdominal neuromeres, which innervate segments that do not bare wings or contain motor neurons innervating flight muscles (Robertson et al. 1982). Yet, these interneurons can reset the flight rhythm in the locust, showing that despite their location they are components of the flight control machinery. Indeed, their role in the flight control circuitry may have influenced the structure of the insect ventral nerve cord (Niven et al. 2006). Robertson et al. (1982) have suggested that these interneurons are remnants of control circuits for ancestral appendages that have been lost.

Neural reuse may be more prevalent in invertebrate brains, especially those of insects, which contain relatively few neurons compared to those of many mammals. Many insects possess small brains that have been miniaturized during evolution (Beuthel et al. 2005). Their small size means that insects are under selective pressure to reduce energetic costs and brain size (Chittka & Niven 2009). Anderson suggests that energy minimization in the absence of behavioral constraints would promote the reduction of neural structures and, thereby, the reuse of neural substrates. The possibility of reusing neurons for different behaviors through the dynamic restructuring of neural circuits means that the consequences of miniaturization may not be as severe as is often assumed.

Anatomical modularity is clear within invertebrate nervous systems (e.g., Niven et al. 2006) but, as Anderson mentions, neural reuse may blur the boundaries between anatomical modules. Indeed, most behaviors involve sensory and motor circuits that are overlapping anatomically, and it seems unlikely that the majority of behaviors are localized entirely within specific anatomical modules. As discussed above, the locust neurons involved in wing control, which include examples of evolutionary reuse, are spread across six neuromeres although only two segments bear wings (Robertson et al. 1982). Indeed, even reflex arcs confined to a single neuromere can be modified by descending and local control, allowing the neurons to be reused in different behaviors (Burrows 1996). Anatomical modularity has been suggested to reduce the energy consumption of neural processing by reducing the length of relatively common local connections and increasing the length of relatively rare long-distance connections. Thus, although modularity may be beneficial for efficiency, it may be opposed by neural reuse, which may not minimize the lengths of connections within neural circuits. In small brains, the low number of neurons and the short distances of most connections may promote further functional reuse, even when some components of neural circuits are in different anatomical segments. Thus, in small brains there may be an increased prevalence of neural reuse.

The Leabra architecture: Specialization without modularity

Alexander A. Petrov,[a] David J. Jilk,[b] and Randall C. O'Reilly[c]

[a]*Department of Psychology, Ohio State University, Columbus, OH 43210;*
[b]*eCortex, Inc., Boulder, CO 80301;* [c]*Department of Psychology and Neuroscience, University of Colorado, Boulder, CO 80309.*
apetrov@alexpetrov.com   david.jilk@e-cortex.com
Randy.OReilly@colorado.edu
http://alexpetrov.com   http://www.e-cortex.com
http://psych.colorado.edu/∼oreilly

Abstract: The posterior cortex, hippocampus, and prefrontal cortex in the Leabra architecture are specialized in terms of various neural parameters, and thus are predilections for learning and processing, but domain-general in terms of cognitive functions such as face recognition. Also, these areas are not encapsulated and violate Fodorian criteria for modularity. Anderson's terminology obscures these important points, but we applaud his overall message.

Anderson's target article adds to a growing literature (e.g., Mesulam 1990; Prinz 2006; Uttal 2001) that criticizes the recurring tendency to partition the brain into localized modules (e.g., Carruthers 2006; Tooby & Cosmides 1992). Ironically, Anderson's critique of modularity is steeped in modularist terms such as *redeployment*. We are sympathetic with the general thrust of Anderson's theory and find it very compatible with the Leabra tripartite architecture (O'Reilly 1998; O'Reilly & Munakata 2000). It seems that much of the controversy can be traced back to terminological confusion and false dichotomies. Our goal in this commentary is to dispel some of the confusion and clarify Leabra's position on modularity.

The target article is vague about the key term *function*. In his earlier work, Anderson follows Fodor (2000) in "the pragmatic definition of a (cognitive) function as whatever appears in one of the boxes in a psychologist's diagram of cognitive processing" (Anderson 2007c, p. 144). Although convenient for a meta-review of 1,469 fMRI experiments (Anderson 2007a; 2007c), this definition contributes little to terminological clarity. In particular, when we (Atallah et al. 2004, p. 253) wrote that "different brain areas clearly have some degree of specialized function," we did *not* mean cognitive functions such as face recognition. What we meant is closest to what Anderson calls "cortical biases" or, following Bergeron (2007), "working."

Specifically, the posterior cortex in Leabra specializes in slow interleaved learning that tends to develop overlapping distributed representations, which in turn promote similarity-based generalization. This computational capability can be used in a myriad of cognitive functions (O'Reilly & Munakata 2000). The hippocampus and the surrounding structures in the medial temporal lobe (MTL) specialize in rapid learning of sparse conjunctive
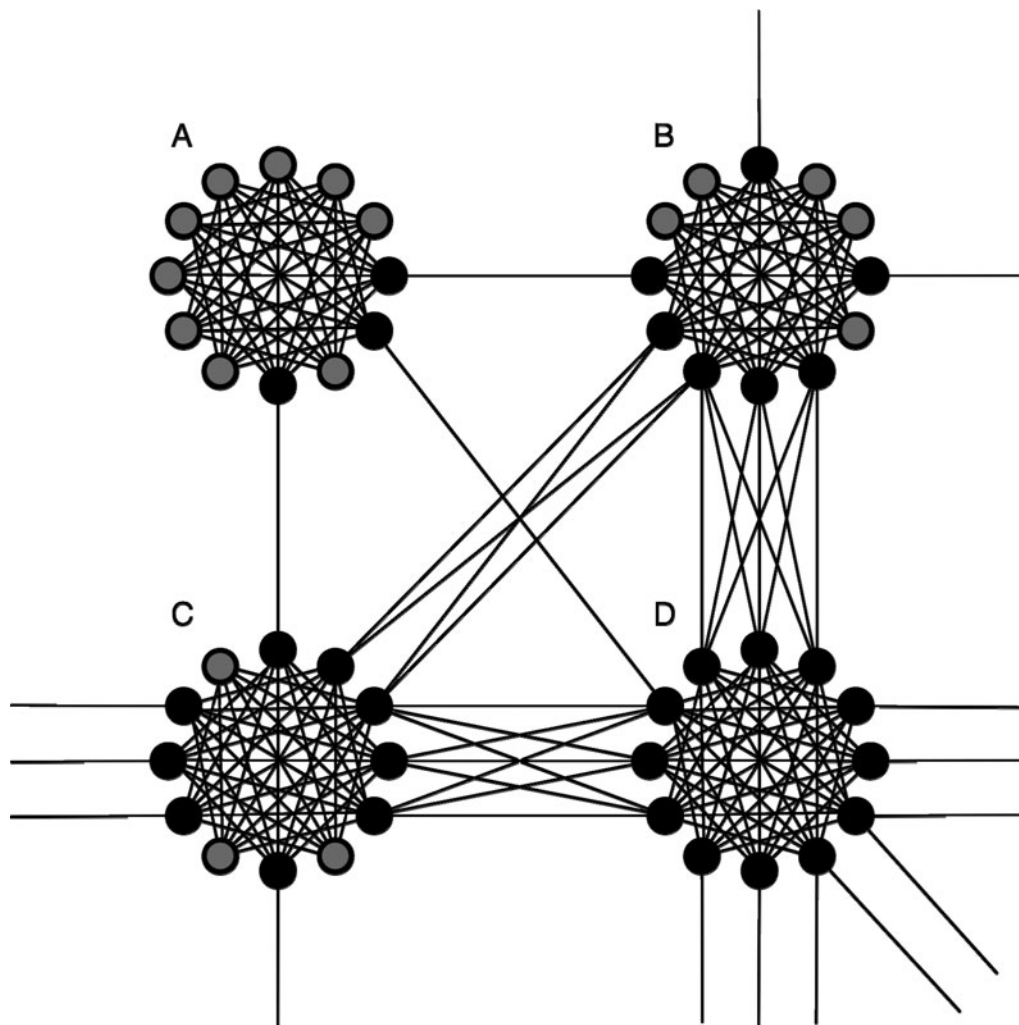

Figure 1 (Petrov et al.) Information encapsulation is a matter of degree. Four neuronal clusters are shown, of which A is the most and D the least encapsulated. Black circles depict exposed (input/output) units that make distal connections to other cluster(s); grey circles depict hidden units that make local connections only.

representations that minimize interference (e.g., McClelland et al. 1995). The prefrontal cortex (PFC) specializes in sustained neural firing (e.g., Miller & Cohen 2001; O'Reilly 2006) and relies on dynamic gating from the basal ganglia (BG) to satisfy the conflicting demands of rapid updating of (relevant) information, on one hand, and robust maintenance in the face of new (and distracting) information, on the other (e.g., Atallah et al. 2004; O'Reilly & Frank 2006). Importantly, most[1] of this specialization arises from parametric variation of the same underlying substrate. The components of the Leabra architecture differ in their learning rates, the amount of lateral inhibition, and so on, but not in the nature of their processing units. Also, they are in constant, intensive interaction. Each high-level task engages all three components (O'Reilly et al. 1999; O'Reilly & Munakata 2000).

We now turn to the question of modularity. Here the terminology is relatively clear (e.g., Carruthers 2006; Fodor 1983; 2000; Prinz 2006; Samuels 2006). Fodor's (1983) foundational book identified nine criteria for modularity. We have space to discuss only *domain specificity* and *encapsulation*. These two are widely regarded as most central (Fodor 2000; Samuels 2006).

A system is *domain-specific* (as opposed to *domain-general*) when it only receives inputs concerning a certain subject matter. All three Leabra components are domain-general in this sense. Both MTL and PFC/BG receive convergent inputs from multiple and variegated brain areas. The posterior cortex is an interactive multitude of cortical areas whose specificity is a matter of degree and varies considerably.

The central claim of Anderson's massive redeployment hypothesis (MRH) is that most brain areas are much closer to the general than the specific end of the spectrum. This claim is hardly original, but it is worth repeating because the subtractive fMRI methodology tends to obscure it (Uttal 2001). fMRI is a wonderful tool, but it should be interpreted with care (Poldrack 2006). Any stimulus provokes a large response throughout the brain, and a typical fMRI study reports tiny differences[2] between conditions – typically less than 1% (Huettel et al. 2008). The importance of Anderson's (2007a; 2007c) meta-analyses is that, even if we grant the (generous) assumption that fMRI can reliably index specificity, one still finds widespread evidence for generality.

MRH also predicts a correlation between the degree of generality and phylogenetic age. We are skeptical of the use of the posterior-anterior axis as a proxy for age because it is confounded with many other factors. Also, the emphasis on age encourages terms such as *reuse*, *redeployment*, and *recycling*, that misleadingly suggest that each area was deployed for one primordial and specific function in the evolutionary past and was later *redeployed* for additional functions. Such inferences must be based on comparative data from multiple species. As the target article is confined to human fMRI, the situation is quite different. Given a fixed evolutionary endowment and relatively stable environment, each human child develops and/or learns many cognitive functions simultaneously. This seems to leave no room for *redeployment* but only for deployment for multiple uses.

Anderson's critique of modularity neglects one of its central features – *information encapsulation*. We wonder what predictions MRH makes about this important issue. A system is encapsulated when it exchanges[3] relatively little information with other systems. Again, this is a matter of degree, as our Figure 1 illustrates. The degree of encapsulation depends on factors such as the number of exposed (input/output) units relative to the total number of units in the cluster, and the density and strength of distal connections relative to local ones. Even when all units are exposed (as cluster D illustrates), the connections to and from each individual unit are still predominantly local because the units share the burden of distal communication. Long-range connections are a limited resource (Cherniak et al. 2004) but are critical for integrating the components into a coherent whole. The Leabra components are in constant, high-bandwidth interaction, and parallel constraint satisfaction among them is

a fundamental implicit processing mechanism. Hence, we eschew the terms *module* and *encapsulation* in our theorizing. This is a source of creative tension in our (Jilk et al. 2008) collaboration to integrate Leabra with the ACT-R architecture, whose proponents make the opposite emphasis (J. R. Anderson 2007; J. R. Anderson et al. 2004). Much of this tension is defused by the realization that the modularist terminology forces a binary distinction on what is fundamentally a continuum.

NOTES
1. There are exceptions, such as the use of a separate neurotransmitter (dopamine) in the basal ganglia.
2. Event-related designs do not escape this criticism because they too, via multiple regression, track contingent variation around a common mean.
3. *Encapsulation* on the input side is usually distinguished from *inaccessibility* on the output side. We discuss them jointly here because of space limitations. Also, the reciprocal connectivity and the task-driven learning in Leabra blur the input/output distinction.

Neural reuse and human individual differences

Cristina D. Rabaglia and Gary F. Marcus
Department of Psychology, New York University, New York, NY 10003.
rabaglia@nyu.edu
gary.marcus@nyu.edu

Abstract: We find the theory of neural reuse to be highly plausible, and suggest that human individual differences provide an additional line of argument in its favor, focusing on the well-replicated finding of "positive manifold," in which individual differences are highly correlated across domains. We also suggest that the theory of neural reuse may be an important contributor to the phenomenon of positive manifold itself.

Anderson's compelling case for neural reuse is well motivated by empirical results and evolutionary considerations and dovetails nicely with the "descent with modification" perspective put forward by our lab (Marcus 2006; Marcus & Rabagliati 2006). An important additional line of support comes from the study of human individual differences.

In an entirely modular brain, one might predict that individual differences in specific cognitive domains would be largely separate and uncorrelated, but the opposite is in fact true: An extensive literature has shown that performance on separate cognitive tasks tends to be correlated within individuals. This "positive manifold," first noted by Spearman (1904), is arguably one of the most replicated findings in all of psychology (e.g., Deary et al. 2006). At first glance, such correlations might appear to be a statistical by-product of the fact that any individual cognitive task draws on multiple underlying processes. However, even when the impurity of individual tasks is taken into account, using more sophisticated structural equation models that form latent cognitive constructs (representing a cognitive ability, such as short-term memory, by the shared variance among performance on diverse tasks with different specific task demands), clear correlations between cognitive capacities within individuals remain. Positive manifold is not an artifact, but a fact of human cognitive life. (Our point here is reminiscent of Anderson's observation that patterns of co-activation in fMRI remain even after subtraction, and are therefore not attributable solely to mechanistic impurities at the task level.)

These correlations between cognitive domains have now been shown in hundreds of separate data sets, and at many levels, ranging from parts of standardized tests such as SAT math and SAT verbal, to broad ability domains such as memory and spatial visualization (see Carroll 1993), to more specific links

such as susceptibility to memory interference and sentence processing (Rabaglia & Marcus, in preparation). Recently, it has been pointed out that "the existence of g creates a complicated situation for neuroscience" (Deary et al. 2010). Adequate theories of brain organization and functioning will have to be consistent with the robust finding of positive manifold, and Anderson's theory of neural reuse is one of the few that is. Strictly modular theories would not predict such between-domain correlations, and nor would theories that are driven purely by experience (since experience is likely to differ heavily between domains).

At the same time, the concept of neural reuse (or decent with modification) may help to shed some light on the interpretation of positive manifold itself. Despite being noted for more than 100 years, there is not yet a consensus on how to explain the phenomenon. Spearman's view was that positive manifold reflected the operation of a general intelligence factor, referred to as "g." Since then, proposed causes range from biological factors such as overall mental speed (Jensen 1998) or myelination (Chiang et al. 2009), to some special rudimentary cognitive ability influencing the operation of others, such as the optimal allocation of resources or a limited central memory capacity (e.g., Kyllonen & Christal 1990); but each of these individually only accounts for (at most) a portion of the variance. If neural reuse characterizes brain functioning for most of human cognition, overlap in the neural substrates recruited by separate cognitive capacities could, in fact, be another factor contributing to positive manifold. One finding that could lend support to this notion is the fact that Ravens Progressive Matrices – arguably the gold standard for tapping into "g" – is an abstract reasoning task, and, as Anderson points out, reasoning tasks are among the most widely distributed in terms of neural areas of activation. Indeed, the most heavily "g-loaded" tasks, or, in other words, the tasks that seem *most* related to what a range of cognitive abilities tend to share, are usually those involving frontal-lobe type abilities (see, for example, Jung & Haier 2007) – the very same abilities that are presumably latest-evolving and thus perhaps most characterized by reuse.

Reuse of molecules and of neural circuits

Mark Reimers

Department of Biostatistics, Virginia Commonwealth University, Richmond, VA 23284.

mreimers@vcu.edu   **http://www.people.vcu.edu/~mreimers**

Abstract: Reuse is well established in molecular evolution; some analogies from this better understood field may help suggest specific aspects of reuse of neural circuits.

Reuse is a settled issue in molecular evolution: most functions in modern cells reuse proteins, or parts of proteins, which previously evolved under different selective pressures. This commentary on Anderson's target article draws analogies between specific aspects of molecular evolution and the ideas he presents about neural reuse, and suggests how the better understood field of molecular evolution may illuminate aspects of and inform hypotheses about neural reuse.

1. Analogy between protein domains and local neural architecture. A protein domain is a chain of typically 20 to 70 amino acids, which folds consistently into a specific compact shape (under normal cellular conditions). Early in protein evolution, a set of useful folds emerged; these domains are essential components of almost all known proteins (Caetano-Anolles et al. 2009a, Finn et al. 2010). Most proteins contain several domains, many contain dozens, and some large proteins contain hundreds of domains. These domains typically perform similar physical functions, such as binding specific protein partners or catalyzing specific reactions, in most of the proteins in which the domains occur. However, the role of each of these domains in the overall economy of the cell has diverged over evolutionary time. Thus domains are prime examples of molecular reuse, reflecting the general evolutionary principle that it is hard to invent something new.

We may think of specific types of neural circuitry as analogous to protein domains. For example, the six-layer local circuit is broadly similar across the cortex, and yet this relatively narrow range of circuit architectures has become involved with almost every aspect of behavior. The typical striatal circuit with inhibitory output cells has also been reused in regions such as the central nucleus of the amygdala (Ehrlich et al. 2009). As the phylogeny of neurodevelopment is uncovered, we might expect to find more examples of newer brain regions borrowing (and mixing) developmental programs from older brain regions.

2. Analogy between metabolic networks and functional circuits. A metabolic network is a set of metabolites, each of which may be readily inter-converted with one of several neighboring metabolites (by gain or loss of some atoms) through the catalytic action of a specific enzyme. The principal metabolic reactions have evolved with the enzymes that catalyze them. Early enzymes catalyzed a set of analogous chemical reactions inefficiently on a wide variety of chemically similar substrates. During the course of early evolution, these enzymes were duplicated by DNA copying errors, and each of the descendant enzymes came to act much more effectively on a narrower range of substrates (Caetano-Anolles et al. 2009b, Yamada & Bork 2009).

There was for some years a controversy over how novel metabolic pathways are assembled, which is analogous to the controversy in cognitive science between dedicated modules and ad hoc neural reuse. An early theory suggested that when genes for enzymes duplicated, they acted on the same kinds of substrates, but catalyzed novel reactions. The major alternative theory, now widely accepted, is that novel metabolic pathways are assembled by duplication of genes for enzymes that perform relevant biochemistry with different substrates; these enzymes then adapt to the substrates of the novel pathway (Caetano-Anolles et al. 2009b; Yamada & Bork 2009). The enzymes that structure novel metabolic functions or pathways are therefore a "patchwork" of adapted pieces from previously existing functions.

Thus, many important pathways of recent vintage are constructed mostly of unrelated proteins. Some of these pathways are crucial to most current forms of life. For example, many of the proteins of the Krebs cycle are distant cousins of proteins that catalyze amino acid metabolism, which evolved earlier in the history of life (Gest 1987; Melendez-Hevia et al. 1996). This "patchwork" model is analogous to Anderson's prediction that more recently evolved pathways invoke more distal brain regions.

These themes in metabolic evolution suggest by analogy that during development many brain regions both become more specialized – dealing with a subset of functions performed previously – and also paradoxically acquire novel functions in the expanding repertoire of behavior. Although a particular behavior may elicit broad brain activity in early life, the same behavior would recruit only a subset of those early regions in later life. However, each individual region active in the original behavior would later become active in many related behavioral functions, in which the region was not originally active. This kind of idea could be tested using chronic implants in animals or using fMRI at several points during child development.

3. Analogy comparing neural reuse to proteins that acquire novel functions very different from their earlier functions. A well-known example concerns cell adhesion molecules (CAMs), whose sticky domains were crucial in forming multi-cellular bodies early in animal life. These same sticky domains have been reused in a variety of other contexts, notably as immunoglo-

bulins in the adaptive immune system, a key to the evolutionary success of vertebrates (Edelman 1987). By analogy, we would expect that during development some brain regions would become important in functions unrelated to those for which they had been used. This idea could be tested as described in the previous paragraph.

4. Analogy between neural circuits and signaling proteins. The majority of proteins in mammals are not enzymes catalyzing reactions, nor even structural components of our diverse tissues, but rather regulators and organizers of small molecules or other proteins. Some of these proteins sequester or transport small molecules, while others modify other proteins (often by adding a small molecular group such as phosphate or methyl), or regulate access to DNA. These "classical" signaling pathways are well-studied because they are reused in many situations. For example, the Wnt signaling pathway is widely reused throughout animal development (Croce & McClay 2008). (Wnt is an example of a protein with three unrelated mechanisms of action, which seem to have been acquired independently.) The fibroblast growth factor (FGF) family of proteins plays a crucial role in the emergence of limb buds, and individual members of the family are reused at several points in mammalian development (Popovici et al. 2005). In all these cases, the specific protein interactions have been preserved while being adapted to a novel function. By analogy then, we might expect different brain regions to preserve the dynamics of their interactions as these regions become co-opted to new functions. This idea might be tested by identifying pairs of brain regions with multiple behavioral functions and recording from these regions simultaneously during several types of behavior in which both regions are active.

Several families of DNA-binding proteins regulate transcription of genes by attracting or blocking the transcription machinery (RNA polymerase complex) at the locations on DNA where they bind. Reuse of these proteins is at the core of some of the most exciting current work in molecular biology: evolutionary developmental biology ("evo-devo") (Carroll et al. 2005). The homeobox genes are famous for their role in early patterning of the front-to-back axis of the embryos of vertebrates and many invertebrates, and these functions are believed to date to the original bilaterian ancestor. However, most of these proteins have lesser-known roles in patterning limbs or digits or epithelia of organs, using the same mechanisms but responding to different signals. By analogy, we might expect that brain regions involved in early aspects of planning actions may also play a role in the fine-tuning of a subset of actions. This suggestion might be tested by recording from "executive" regions of the prefrontal cortex (PFC) during a variety of tasks.

Molecular evolution provides many specific examples of reuse, of which I have only scratched the surface. By analogy, these may provide some concrete inspiration for further research in the evolution and development of mental function.

Massive modularity is consistent with most forms of neural reuse

J. Brendan Ritchie and Peter Carruthers
Department of Philosophy, University of Maryland, College Park, MD 20742.
britchie@umd.edu   pcarruth@umd.edu
https://sites.google.com/site/jbrendanritchie/Home
http://www.philosophy.umd.edu/Faculty/pcarruthers/

Abstract: Anderson claims that the hypothesis of massive neural reuse is inconsistent with massive mental modularity. But much depends upon how each thesis is understood. We suggest that the thesis of massive

modularity presented in Carruthers (2006) is consistent with the forms of neural reuse that are actually supported by the data cited, while being inconsistent with a stronger version of reuse that Anderson seems to support.

Carruthers (2006) characterizes the mind as composed out of the interactions of a large set of mental modules, utilizing the "global workspace" provided by perception and working memory to recruit the resources of multiple specialized systems in the service of cognition and behavior. The sense of "module" in question is quite weak, however. Modules are functionally dissociable, intentionally characterized processing systems, each with its own neural realization. Modules need not be encapsulated, domain-specific, or innate (although many probably are). And the neural systems that realize them certainly need not be anatomically localized. On the contrary, modules can be realized in spatially distributed interconnected networks of brain regions. Moreover, many modules are constructed out of, and share parts with, other modules. Hence, the distinctness of different modules and their neural realizers will only be partial.

Anderson claims that the thesis that modules can share parts is inconsistent with the idea that modules are functionally dissociable and separately modifiable, committing Carruthers to a strong version of anatomical modularity. But this is a mistake. Provided that two modules sharing a part differ from one another in other respects, it will be possible to disrupt the operations of one without having any impact on the other (by disrupting only the parts of the former that are *not* shared), and it will be possible for natural selection to make modifications in the one without changing the other (again, by making improvements in the parts that are not shared). Indeed, at the limit, two modules could share *all* of their processing parts while still remaining dissociable and separately modifiable. For the differences might lie entirely in the patterns of connectivity among the parts, in such a way that those connections could be separately disrupted or improved. In short, the functional dissociation and separate modifiability of modules do not preclude the possibility of neural reuse.

The shared-parts doctrine provides a clear sense of neural reuse that is consistent with massive modularity. Moreover, each shared part can be given a dual functional characterization. Its function can either be described in univocal local semantic terms, or it can be said to be multi-functional, characterized in terms of the different longer-range uses for which its outputs are employed. (This seems to correspond to *one* way of understanding Anderson's distinction between "workings" and "functions," respectively, although he himself characterizes the former in "low-level computational" rather than intentional terms [sect. 1.1, para. 5].) Consider, for example, the region of fusiform gyrus that is often characterized as a face-recognition area (Coltheart 1999; Kanwisher et al. 1997). At one level of description, this is a module that recognizes faces. But it will contribute to, and be a part of, a number of larger systems. One is a person-file system, which uses face-recognition to help collect and store information about the individuals in one's community (especially their personality traits and mental states). Another is an affiliative, social-bond-building, module which uses face-recognition as part of the process of creating and activating positive affective reactions to specific others. And a third is a Westermarck-style incest avoidance module (Fessler & Navarrete 2004), which uses the face-recognition module during childhood to track the extent to which other children are co-present in the home, and then above a certain threshold of cohabitation produces sexual disgust at the prospect of intercourse with those individuals post-adolescence. We can then say that the fusiform gyrus is a module with one local function (face-recognition) which is part of at least three other larger-scale modules (and hence is at the same time multi-functional).

Notice that nothing much needs to change in this account if one thinks that the fusiform gyrus isn't a face area, but is rather a holistic shape-processing area, which can be used for

recognizing any type of object that requires a combination of local detail and overall form (Gauthier et al. 2000; 2003). For we can now characterize its local function in just such semantic terms; and yet on this account, there will be an even larger set of systems of which it constitutes a modular part.

However, the massive modularity hypothesis is inconsistent with a distinct, stronger, doctrine of neural reuse. This would claim that a neural region can be implicated in multiple long-range functions *without* there being a single semantic characterization of its local function. Perhaps Anderson endorses this stronger view. He emphasizes, for example, how the same brain regions can be involved in very different tasks like reading comprehension and manual object-manipulation (sect. 3.1, para. 5). And he thinks that local functions (or "workings") are "low-level" and computational rather than intentional. But nothing in the evidence that Anderson presents actually supports such a view over the weaker account sketched above. Moreover, it strikes us as quite implausible. It is hard to see how the same set of computations could realize distinct representational properties on different occasions of use. For the consumer, systems for those computations would have no way of knowing which representational properties are involved on a given occasion, and hence no way of determining how the outputs should be used.

Anderson might accept a more modest position with which the data are equally consistent: Under such a view, the neural region of interest subdivides into a number of more fine-grained areas (too fine-grained to show up in fMRI data, for example), each of which has a specialized semantically characterizable function. Furthermore, for all that the data show, distinct local modules might spatially interpenetrate one another, with the neurons involved in one being interspersed among neurons involved in the other, in something like the way that mirror neurons are interspersed among purely motor-related neurons in premotor regions of macaque cortex (Rizzolatti & Craighero 2004). However, such a position would also be consistent with the thesis of massive modularity.

We conclude that to the extent that the data reviewed by Anderson support a thesis of massive neural reuse, the resulting thesis is fully consistent with the hypothesis of massive mental modularity, as characterized by Carruthers (2006).

More than modularity and metaphor: The power of preadaptation and access

Paul Rozin

Department of Psychology, University of Pennsylvania, Philadelphia, PA 19104-6241.

rozin@psych.upenn.edu

Abstract: Neural reuse demonstrates preadaptation. In accord with Rozin (1976), the process is an increase in accessibility of an originally dedicated module. Access is a dimension that can vary from sharing by two systems to availability to all systems (conscious access). An alternate manifestation is to reproduce the genetic blueprint of a program. The major challenge is how to get a preadaptation into a "position" so that it can be selected for a new function.

For more than ten years, I have intended to submit an article to *Behavioral and Brain Sciences* on the power of preadaptation and access. The excellent article by Anderson on neural reuse provides strong evidence for preadaptation and access, and I jump at this opportunity.

Preadaptation is a basic principle in twentieth-century evolutionary biology (Bock 1959; Mayr 1960). As Ernst Mayr points out: "The emergence of new structures is normally due to the acquisition of a new function by an existing structure . . . the resulting 'new' structure is merely a modification of a preceding

structure" (Mayr 1960, p. 377). The basic idea is that something that evolved for one function is used for another. Occasionally the original structure is not itself an adapted entity, falling under the broader category of exaptation (Buss et al. 1998; Gould 1991; Gould & Vrba 1982). The human brain is surely a preadaptation: a very large processing system selected to solve a wide range of problems, then adapted to solve (or create!) problems other than those for which it was originally selected.

In 1976, in response to the view that learning was accomplished by a few general-purpose and domain-insensitive mechanisms, I put forth some ideas in a paper entitled "The Evolution of Intelligence and Access to the Cognitive Unconscious," ideas that were related to preadaptation and to the issues raised by Anderson (Rozin 1976). Below, I list a few points made in this 1976 paper and in some subsequent work (Rozin 1999; 2006) that anticipate some of the later findings and/or suggest directions for future work.

1. The building blocks for innovations in evolution, and particularly the brain, are adaptive specializations (called modules by Fodor) which are circuits or structures specifically dedicated to performing a specific function. These can be considered preadaptations.

2. In the course of evolution, these modules may be accessed by other systems, and thus acquire a new function. The original function may remain (e.g., shared circuitry – neural reuse), or the original function may disappear.

3. Accessibility is a dimension, varying from a dedicated module at one extreme to attainment of consciousness, which usually means total access to all systems. The brain (mind) is neither totally modular nor totally a general processor. It is both and everything in between.

4. A parallel process of increasing access occurs in development (e.g., Piaget's décalage), and an inversion of this process is associated with degeneration of the nervous system.

5. Alphabetic writing and reading presumes some level of access (or even "insight") into the fact that "bat" has three sounds. This can be framed as gaining access to the phonological processing "module."

6. In addition to the idea of reuse, there is an alternate preadaptive pathway (Rozin 1976): that is, to reproduce the genetic/ developmental plan for a particular neural circuitry in another part of the brain. This presumably happened, for example, with multiple topographic representations of space in different visual areas of the brain.

The impressive recent data supporting the idea of a literally embodied mind are an instance of preadaptation and access, in the use of sensory and motor cortical structures to represent "higher" functions.

The framework I present highlights the critical developmental-evolutionary problem with this whole class of models. As formulated by Mayr, the problem is: "How can an entirely new structure be gradually acquired when the incipient structure has no selective advantage until it has reached a considerable size and complexity?" (Mayr 1960, p. 350). How do we get from a photosensitive cell to an eye, from a fin to a limb, from a jaw articulation to middle ear bones? Many of the imaginable intermediate stages are not adaptive. In terms of the reuse (as opposed to duplicate circuitry) model, *physical contact* is necessary between a brain area whose function could be improved and the other circuitry that could enhance its function, in order for selection pressure to operate. *Getting closer is not more adaptive; it is contact that is needed.* One must identify the selective force that leads to contact, as demonstrated beautifully by Bock (1959) in his analysis of how an enlarging muscle insertion point on the mandible of a particular bird species becomes a jaw articulation after it contacts the skull. There is no doubt that some type of contact has been established in many important examples of preadaptation in evolution, as, for example, the invasion of land by aquatic vertebrates. There are examples of preadaptation where the new adaptation replaces the old (reptile jaw

articulation to middle ear bones) and others more like reuse, where a structure maintains its old function and acquires a new one (such as the human tongue functioning both in ingestion and in the articulation of speech sounds).

Brain anatomy, and developmental constraints, probably make it difficult to physically co-opt circuits that are not in close proximity. One possibility, implied by some of the work of Dehaene and Cohen (2007), is that expansion of a particular area of the brain brings it into contact with neural tissue that can improve its function by integrating this circuitry.

Natural selection is powerful when there is transmission. But it can only act on the available variants, and it can be trapped by local optima and the necessity to bridge maladaptive intermediate phases. And here is where something wonderful comes in to speed up and expand the process immensely. Culture! Preadaptation, however impressive in biological evolution, is massively important in cultural evolution, because the variants can be generated purposively, and there is tolerance for maladaptive intermediate stages, motivated by the desire to reach a goal. The extraordinary power and speed of cultural evolution is well documented (Girifalco1991; Newson et al. 2007). Natural selection can work without constraints! The results are computers, memory storage systems that evolve by the year, Mozart symphonies, and the like. I am astonished that evolutionary psychologists are not excited by the application of the principle of natural selection to the study of cultural evolution, given that they can watch it happen (Rozin, in press). I was excited to learn from Anderson that Dehaene and Cohen (2007) have been examining how processes like access can occur in the developing brain under the selective guidance of cultural selection. I think this is what I was talking about in 1976 as accessibility in development and in cultural evolution (Rozin 2006). But we still have to figure out how Mother Nature built such an extraordinary creature as the human before intentional cultural actions made abilities and artifacts available as preadaptations.

Optical holography as an analogue for a neural reuse mechanism[1]

Ann Speed, Stephen J. Verzi, John S. Wagner, and Christina Warrender

Sandia National Laboratories,[2] Albuquerque, NM 87185-1188.
aespeed@sandia.gov   www.sandia.gov
sjverzi@sandia.gov   jswagne@sandia.gov   cewarr@sandia.gov

Abstract: We propose an analogy between optical holography and neural behavior as a hypothesis about the physical mechanisms of neural reuse. Specifically, parameters in optical holography (frequency, amplitude, and phase of the reference beam) may provide useful analogues for understanding the role of different parameters in determining the behavior of neurons (e.g., frequency, amplitude, and phase of spiking behavior).

Optical holography hypothesis. In this commentary, we highlight a possible physical mechanism for neural reuse. Importantly, how reuse is implemented in neural tissue is one of the primary open questions, as the author states in section 6.4, paragraph 4, of the target article. Specifically, we wonder if there might be utility in a theory of reuse (i.e., recruitment of the same cortical area for multiple cognitive functions) based on an analogy to optical holography. This analogy has been proposed by several authors as early as the late 1960s and early 1970s (e.g., Westlake 1970) and as recently as 2008 (Wess 2008). It has influenced work in distributed associative memory, which involves neural reuse in the form of individual processors

contributing to multiple distributed representations (Plate 1995; Sutherland 1992). However, the full potential of the analogy does not appear to have been realized. Therefore, we describe optical holography and the neural analogy, state some predictions about neural function based on this analogy, and propose means for testing these predictions. (See our Fig. 1.)

Optical holography was developed by Dennis Gabor, for which he won the 1971 Nobel Prize in Physics. It is a method for encoding, and then retrieving, multiple images onto a single (color-sensitive) photographic plate using different wavelengths of light emitted by a laser. Illustrated in Figure 1, laser light is split into two equally coherent beams of light by a beam splitter. One path goes through the beam splitter and reflects off of the target (real three-dimensional) object; some of this reflected light hits the storage media (photographic film). The other path is reflected by the beam splitter directly towards the storage media. The difference in path length of the two coherent beams of light from the beam splitter to the storage media creates a phase difference that exposes the photographic film with an intoferogram image (inset, Fig. 1). To retrieve the stored image, the real object can be removed and the laser light is sent through the beam splitter, all of which is reflected to the photographic film. After passing through the photographic film, the optical holographic image is reconstructed and visible to the eye. Importantly, if lasers of different wavelength are used, different holograms can be encoded on the same photographic film, essentially allowing reuse of that film.

Reuse. That multiple images can be encoded in a distributed manner on a single plate at different wavelengths is the foundation of the applicability to the neural reuse hypothesis, although we imagine it would apply to more than just storage of memories. Specifically, optical holography has fundamentally three parameters that can be varied to encode each unique hologram onto a photographic medium: (1) frequency of the laser, (2) amplitude of the laser, and (3) phase relationships to other stored representations. On the surface, these three variables might be analogous to frequency, amplitude, and phase relationships in firing of individual neurons and neural circuits or ensembles. However, there are additional variables affecting neural behavior, including: (i) involvement of various neurotransmitters; (ii) afferent, lateral, and feedback connectivity; and (iii) temporal relationships between thousands of inputs. This implies a large variety of ways in which an individual neuron, circuit, or area can be recruited for different functions.

One prediction that follows from this part of the analogy is that one should be able to elicit fundamentally different behaviors from a neuron, circuit, or even a cortical region by changing the input or the larger context in which the input occurs. This could take the form of electrical stimulation with different properties or the presentation of different neurotransmitters to the same neuron or circuit and measuring the resulting outputs. If the electrical stimulation or neurotransmitter acts as an analogue to the wavelength of the reference beam, different behaviors should result. Such testing could be done in slice preparations, in engineered networks, or in simple animal models such as *Drosophila*.

Interference. Harmonic interference, or holographic aliasing, leads to errors that may have analogues in human memory and learning. Specifically, aliasing in holography may be analogous to confabulation or abstraction into a schema representation. Aliasing can occur as a result of two objects being encoded onto the same plate using lasers with similar wavelengths, and results in retrieval of an image that is an amalgamation of the two original objects. An analogue in human behavior would be two skills learned under similar contextual conditions. When those skills are functionally similar, they are considered to positively transfer to one another, and can result in a generalized representation of that particular class of problems. When they are functionally dissimilar, this is considered to be an example of negative transfer (e.g., Cormier 1987; Gick & Holyoak 1983;
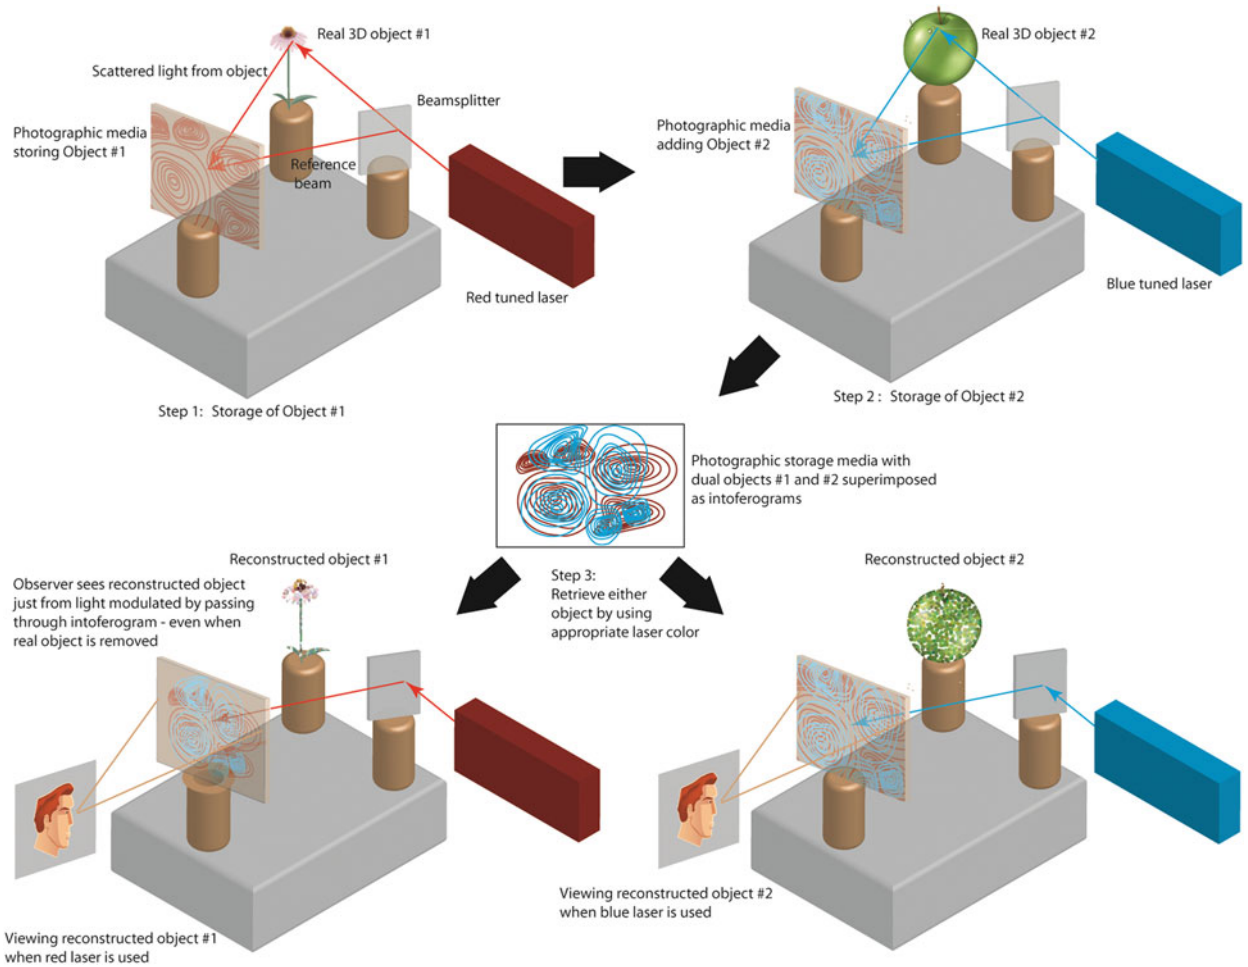
Figure 1 (Speed et al.) Holographic data storage and retrieval. Inset illustrates an interference pattern on film that is the physical storage of the holographic image.

Novick 1988). This might also cause enhancement of a particular memory (e.g., the von Restorff effect; Hunt 1995).

Additional implications of the holographic analogy include:

1. The fact that a single beam of a particular frequency recalls the entire image may be analogous to redintegration (Roodenrys & Miller 2008).

2. The capacity for storage and reuse increases with the number of variables used in defining neural and circuit behavior (Kalman et al. 2004; Plate 1995; Psaltis & Burr 1998; Sutherland 1992).

3. The number of parameters defining neural and circuit behavior in a given organism should predict behavioral/cognitive complexity, and such complexity should scale similarly to the predicted capacity.

As indicated above, tests of predicted capacity and interference can be done using computational simulations, experiments with networks in preparation, or engineered networks of neurons.

In the past, the optical holography analogy has been criticized (e.g., Wilshaw et al. 1969). Certainly, the analogy does break down in certain places – for example, the fact that any piece of the photographic plate encodes the entire image, thus destroying some parts of the plate, merely degrades image quality rather than creating an analogue to aphasias seen in humans. However, using the holographic analogy as a starting point for hypothesis development might provide a foundation from which the physical mechanisms of neural reuse might be identified.

Massive redeployment or distributed modularity?

Alexia Toskos Dils and Stephen J. Flusberg
Department of Psychology, Stanford University, Stanford, CA 94305.
atoskos@stanford.edu   **sflus@stanford.edu**

Abstract: In distinguishing itself from other distributed approaches to cognition, Anderson's theory of neural reuse is susceptible to some of the same criticisms that have been leveled at modular approaches. Specifically, neural reuse theories state that: (1) the "working" of a given brain circuit is fixed, rather than shaped by its input, and (2) that high-level cognitive behaviors can be cleanly mapped onto a specific set of brain circuits in a non-contextualized manner.

The target article does an excellent job of exploring the behavioral, neural, and theoretical evidence supporting the idea that brain regions are reused in the service of many different cognitive functions and that traditional, modular approaches to neural architecture may be misguided. This viewpoint echoes other recent critics of contemporary cognitive neuroscience (e.g., Uttal 2001) and fits well alongside related distributed, emergent approaches to cognitive functioning (Rumelhart & McClelland 1986; Thelen & Smith 1994; Varela et al. 1991). A distinguishing feature of Anderson's neural reuse framework is that it highlights how local neural circuits with *fixed* "workings" may be combined in evolutionary (or developmental) time to support new cognitive "uses." However, we are concerned that some of the same criticisms that have been leveled at modular approaches to the mind may also pose problems for the current formulation of the neural reuse theory.

First, much like classical modular views of mind, Anderson's theory of neural reuse de-emphasizes the role that the immediate environment plays in the development of the functional properties of a particular neural circuit (Fodor 1983; Pinker 1997). In fact, the target article explicitly claims that the working of any given anatomical brain site is fixed, in stark contrast to classical PDP (parallel distributed processing) models. However, there is evidence that the function of a given neural circuit may be largely shaped by the structure of its input. For example, Sur and colleagues (Sharma et al. 2000; von Melchner et al. 2000) surgically rewired the optic tract of a ferret so that primary auditory cortex received visual input from the eyes of the animal. Not only did the ferret seem to develop normal visual (and auditory) behavior, but also the circuitry in auditory cortex exhibited many of the properties traditionally associated with visual cortex, such as orientation selective cortical columns. This suggests that the working of circuits even in the most evolutionarily ancient cortical regions is not restricted to any particular modality, let alone any specific function. Such flexibility provides evidence in favor of computational mechanisms that derive their function based in part on the statistical structure of the input (Rumelhart & McClelland 1986).

Second, while Anderson's theory of neural reuse rejects the idea that high-level cognitive functions (e.g., "language comprehension") can ultimately be mapped onto any single brain module, the approach still calls for the one-to-one mapping between these high-level functions and a specific, distributed set of neural circuits. However, it may be the case that distinct instances of what we would label as the *same* cognitive behavior might actually emerge from the distributed activation of *different*, contextually variable sets of neural circuits. For example, although visual object recognition has been shown to automatically activate motor brain regions (Chao & Martin 2000; Tucker & Ellis 1998), very different motor circuitry might be recruited to *recognize* a chair when you are tired and want to sit down than when you need to reach something on a high shelf. There may also be individual differences across a population in what neural resources are recruited for a particular cognitive task. For example, some people seem to readily recruit direction-selective neurons when listening to stories describing both literal and metaphorical motion, whereas others do not, even though both groups *comprehend* the story (Toskos Dils & Boroditsky, forthcoming). Thus very different neural representations might subserve the very same high-level cognitive behavior (i.e., "object perception" and "language comprehension") both within and across individuals. This suggests that it may be a *category mistake* to try to reduce complex, person-level cognitive phenomena to a unique set of neural circuits (Ryle 1949). Rather, these mental operations are always a contextually bound, emergent function of the history of the organism, the immediate environment, and the bodily state of the organism (Thelen & Smith 1994).

In sum, while Anderson's theories of neural reuse offer a much-needed counterpoint to traditional, modular views of neural architecture, they still suffer from some of the same difficulties these modular views have in accounting for complex cognitive behaviors that develop over the course of learning and experience. Dynamic models of cognitive function preserve many features of the neural reuse framework that account for data unexplained by massive modularity models. They should be preferred because, unlike neural reuse models, they also predict that the function of a given circuit should change as the structure of its input changes, and they do not require that high-level cognitive functions cleanly map onto specific cortical circuits. These approaches currently provide the additional benefit of computational models that can be used to make precise predictions about the development of cognition function. Proponents of neural reuse should point to specific ways in which they can accommodate the limitations of the current formulation of neural reuse theory.

Belling the cat: Why reuse theory is not enough

Oscar Vilarroya
Unitat de Recerca en Neurociència Cognitiva, Departament de Psiquiatria i Medicina Legal, Universitat Autònoma de Barcelona, and Fundació IMIM, Barcelona 08193, Spain.
oscar.vilarroya@uab.cat

Abstract: I agree with Anderson's approach to reuse theories. My main concern is twofold. Anderson assumes certain nomological regularities in reuse phenomena that are simply conjectures supported by thin evidence. On the other hand, a biological theory of reuse is insufficient, in and of itself, to address the evaluation of particular models of cognition, such as concept empiricism or conceptual metaphor.

I would first like to welcome Anderson's target article. Extant cognitive neuroscience and neuroimaging studies, as well as the growing importance of biological analyses in cognitive science, increasingly show the unsuitability of a modular approach to cognition. In this situation, a new framework is required to model the functional architecture of cognitive processes in the nervous system. Anderson's article is a remarkable effort in this direction. I agree with his general approach to the issue. My main concern, though, is twofold. On the one hand, Anderson assumes certain nomological regularities in reuse phenomena that are simply conjectures supported by shaky evidence. On the other hand, a biological theory of reuse by itself is inadequate for the task of evaluating particular models of cognition, such as concept empiricism or conceptual metaphor. We need an independent characterization of cognitive phenomena, a model that we currently lack.

First, extracting biological regularities from evolutionary phenomena is not a straightforward issue. Elsewhere (Vilarroya 2001), I have suggested that cognitive systems are constrained by what I called "bounded functionality," which accounts for the dynamics of the functional paths leading to solutions to adaptive problems. One of the bounded functionality constraints is what I call the "bricoleur constraint," defined as the fact that natural selection favors the shortest design path. In other words, the solutions to adaptive problems have to take into account the resources that were available to the system before the adaptive problem appeared. The bricoleur constraint is the evolutionary characterization of the reuse approach. However, the bricoleur constraint can be realized in many ways for any evolutionary phenomenon. For instance, Anderson's principle, that "older areas, having been available for reuse for longer, are ceteris paribus more likely to have been integrated into later-developing functions" (sect. 1.1, para. 1), can be a good starting point, but it cannot be taken as an evolutionary law. Evolutionary biology is full of surprises; older areas can serve a small

range of functions at the same time that an intermediately incorporated area which proved more useful in later functions results in more pervasive implications. Evolutionary tinkering is, in itself, not susceptible to lawlike regularities (see, e.g., Jacob 1977). Additionally, the evidence by which Anderson tries to sanction the abovementioned principle is based on the hypothesis that "the older the areas, the more back in the brain they are" (see sect. 1.1, para. 3), which is, to say the least, highly contentious. The foundation of his entire argument is therefore a shaky one.

Second, in order to address the evaluation of particular models of cognition, we require, apart from reuse theory, a characterization of the cognitive processes the nervous system actually carries out; and the jury is still out on nearly all the available hypotheses. Indeed, Anderson examines cognitive models while taking for granted some functional attributions, for example, of fMRI studies, to form the basis of his argumentation, but such characterizations are under discussion precisely in part because of reuse theories. For example, in section 4.4, Anderson uses neuroimaging studies to argue against conceptual metaphor. However, the functional interpretation of such studies (e.g., "finger representation") are prone to criticism, as is any other neuroimaging study, precisely on account of reuse theories, and therefore cannot be used as arguments against conceptual metaphor or any other hypotheses. Neuroimaging studies are task-oriented, and the interpretations are reverse-engineering biased. Previously (Vilarroya 2001), I addressed the issue of "functional mesh," that is, the assumed tight fit between a cognitive trait's design and the adaptive problem it is supposed to solve. It is now widely assumed, even by Anderson, that the "optimality move" that creeps in behind functional mesh is misplaced – namely, that cognitive mechanisms need not be specially designed to solve the adaptive problems for which they were selected. Even if Anderson seems to agree with such an approach, my impression is that he eventually falls into the functional mesh trap, by assuming the functions of certain areas.

I have also defended (Vilarroya 2002) a dual reverse-engineering and biological analysis to characterize cognitive functioning. However, biological analyses in cognitive science are of a particular type. Usually, biological explanations are teleonomic explanations that first identify the trait that is likely to be under selection, and then identify the adaptive problem that the trait is supposed to solve. Yet, certain aspects of cognitive science force a change in this methodology. In trying to explain the cognitive mechanisms of a biological organism, the researcher can identify the adaptive problem that the brain is supposed to solve, but in reality it is difficult to identify the actual trait itself, because the trait is not as self-evident as, say, an eye, a liver, or a wing. Moreover, the explanatory strategy of cognitive science cannot simply be an inversion of the first steps of the teleonomic explanation. It is not enough to identify the adaptive problem and then infer the mechanism. Rather, we need to complement an initial assumption about a trait's design with a characterization of how the adaptation might have appeared over evolutionary time – first characterizing the adaptive problem that the organism is supposed to solve, then the fitness-maximization process, as well as showing that the trait is specialized for solving the adaptive problem, unlikely to have arisen by chance alone, and not better explained as the byproduct of mechanisms designed to solve some alternative adaptive problem.

In summary, functional attribution in cognitive science is not a straightforward operation, but rather, requires an independent characterization from the functional mesh assumption; reuse theory alone cannot provide this type of tool. Hence, in my opinion, Anderson lacks the basis to apply his functional characterizations as arguments against specific models of cognition. Once we have the necessary tools to account for functional characterization in cognitive science, of course, reuse theory will prove extremely useful.

Author's Response

Cortex in context: Response to commentaries on neural reuse

Michael L. Anderson

Department of Psychology, Franklin & Marshall College, Lancaster, PA 17603, and Institute for Advanced Computer Studies, Neuroscience and Cognitive Science Program, University of Maryland, College Park, MD 20742.
michael.anderson@fandm.edu
http://www.agcognition.org

Abstract: In this response, I offer some specific examples of neural workings, discuss the uncertainty of reverse inference, place neural reuse in developmental and cultural context, further differentiate reuse from plasticity, and clarify my position on embodied cognition. The concept of local neural workings is further refined, and some different varieties of reuse are identified. Finally, I lay out some opportunities for future research, and discuss some of the clinical implications of reuse in more detail.

Behavioral and Brain Sciences (*BBS*) is a unique and extremely valuable resource, and so I would like to begin this response by thanking the editors for their continued service to our field. *BBS* has been an important part of my intellectual life since I was an undergraduate. I vividly remember my first encounter with the journal in the library stacks. Its debates were deeply helpful to me in preparing my senior thesis, and have remained crucial to my intellectual development since. I know many of us in the cognitive sciences are similarly indebted. Naturally, this arena for discussion would serve no purpose without willing participants, who spend their time and energy to help improve the ideas of others. For this gift from my commentators, I am truly grateful.

The commentaries cover an astonishingly broad range of issues – from history to holograms, modularity to memory consolidation – and I will do my best to at least touch on all of the many ideas they contain. Many commentators are especially concerned about the core notion of cortical "workings," and about my emphasis on neural context as the main determiner of cognitive function, to the apparent exclusion of the social, environmental, and developmental contexts that also help determine functional outcomes. A few commentators take issue with my stance on embodied/grounded cognition. Some commentators have concerns about the general adequacy of the theory; others, about the adequacy of the data; and a few offer some alternate hypotheses to account for the data I review. Very many commentators offered specific advice for ways to improve the theory – proposals for better integrating neural reuse with evolutionary theory, for specifying the mechanisms driving reuse, and for some experimental approaches that could further elucidate the functional organization of the brain. I try to treat each of these topics in the following sections.

R1. What neural reuse is

Before getting to those specific responses, let me begin with a short section in which I discuss two specific

examples of what I take a "working" to be, as it might help clarify the theory of neural reuse more generally. As was hopefully made clear in the target article, the basic idea behind neural reuse is that neural circuits, which may have initially served one specific purpose, have come over evolutionary and developmental time to play many different roles. Because terms like *role, purpose*, and *function* have many different meanings, and were in fact being used in conflicting ways in the literature on function-to-structure mapping, Bergeron (2008) introduced the terms *use* and *working*. Neural reuse holds that the "workings" of local neural circuits are put to many different higher-level "uses," and that the flexibility and variety of our cognitive repertoire results in part from the ability to put together the same parts in different configurations to achieve different behavioral outcomes.

From the perspective of neural reuse, it appears that the field of psychology has typically concerned itself with investigating *uses*, which is of course a necessary part of any investigation of the mind. Nevertheless, given the apparent many-to-many mapping between uses and brain regions, it behooves the cognitive scientist interested in the neural basis of cognition to think about *workings*, as well.

What, then, is a working? Abstractly, it is whatever single, relatively simple thing a local neural circuit *does for* or *offers to* all of the functional complexes of which the circuit is a part. Concretely, consider two examples:

In Penner-Wilger and Anderson (2008), we suggested that a brain circuit known to subserve both finger and number representations might be offering to both a kind of ordered storage. The idea was that a register – an ordered set of containers in which to store specific values – offered a functional structure useful to both finger and number representation, and so that structure might have been deployed for both uses.[1]

Somewhat more speculatively, consider the ability to fixate the eye on a specific region of the visual field. This is known as *foveation*, because its purpose is to move the eye so that the fovea (the retinal region offering the greatest acuity) is receiving the desired input. Foveation is important to many visual capacities, including the visual grasp reflex, smooth ocular pursuit, and reading. One component of the foveation function might be the ability to map any arbitrary element in a matrix (a two-dimensional grid that could represent the retina) onto the center of that matrix, that is, the ability to re-center the matrix around any of its elements.[2] Such a working could play a functional role not just in the visual uses mentioned above, but also in such things as shifting spatial attention and Braille reading – and even in the "tactile foveation" exhibited by the star-nosed mole (Catania & Remple 2004). Hence, we should not be surprised to find that parts of the foveation circuit are deployed not just in visual tasks, but in these other tasks as well.

These are, of course, just examples of the kinds of thing that workings could be. As noted in the target article, the discovery and definition of specific neural workings can only come at the end of a long process of investigation and testing. Nevertheless, I hope these examples – however speculative or provisional – can serve to clarify the basic notion, and improve understanding of the theory as a whole.

R2. Context, context, context

One of the central implications of neural reuse that did not come out as clearly in the target article as I would have liked is the deep uncertainty of reverse inference as a strategy of functional explanation in cognitive neuroscience (Poldrack 2006). If brain regions contribute to multiple uses – if, that is, they fail to be "use-selective" – then the mere observation of activity in a given region provides very little information about what the brain is up to. Certainly, one cannot assume that a region is being put to the same use in an attention task as it was in a language task or an emotion task. This goes also, and perhaps especially, for inferences based on seeing more (or less) activation in a region under different task conditions. If one doesn't know what the brain is doing just in virtue of observing regional activity, then one cannot know it is doing *more* of some specific thing (more attention, more control, more calculation) in virtue of any observed increase in that activity. Differences in activity level could equally well be a sign of being in a different information state.

R2.1. The importance of neural context

For many, these observations will simply add fuel to the skeptical fire being built under the use of neuroimaging in cognitive science (see Coltheart 2006; Klein 2010; Roskies 2007 for discussions of the general issue). Certainly there is reason to be cautious, but the potential value of neuroimaging is so vast that it would be foolish to forego its use. So how should we address this issue? The target article emphasizes that although neural regions are not use-selective, they may be "working selective," and so clearly one promising empirical project is to begin to define local workings with various sorts of cross-domain modeling. What was less clear in the target article, but helpfully raised by **Gomila & Calvo** and **Reimers,** is that there is another, complementary empirical project: Although local regions are apparently not use-selective, *networks* of regions may well be use-selective. That is, it might be possible to recover selectivity by attending to larger scale patterns of regional co-activation and coherence. Cognitive tasks and task categories may turn out to have characteristic signatures at the network level (for discussion, see Anderson et al. 2010; Poldrack 2006).

Whether and to what degree specific identifiable networks of interacting regions turn out to be use-selective is an empirical question, one that is only now starting to be investigated. But it seems clear that this is currently the most promising, and perhaps the only viable way to uncover use selectivity in the brain. Note the implication that knowing what activity in a given region means for the tasks engaging the brain will require careful attention to the *neural context* of that activation – to what *else* the brain is doing when it exhibits the activation of interest. Seeing activation in Broca's area may give one confidence that it is doing *something* of importance (although see **Klein**'s commentary and Klein [2010] on the uncertainty of this inference), but knowing what *else* is active along with Broca's may tell us what that something is, the use to which Broca's is being put.

R2.2. Bodily, environmental, social, and cultural context

This point about the value of attending to neural context was somewhat eclipsed by my attention to the importance of local neural workings. Moreover, my near exclusive attention to *neural* facts appeared to **Toskos Dils & Flusberg**; **Gomila & Calvo**; **Immordino-Yang, Chiao, & Fiske** [**Immordino-Yang et al.**]; and **Rozin** to eclipse the value of attending to even broader contexts. I certainly endorse the general point that broader contexts – bodily, environmental, social, cultural – matter to our ascription of function. If we don't understand what an organism is doing, we can hardly hope to understand what its brain is up to; and figuring out the best way to describe or characterize a behavior certainly involves detailed attention to context. On the other hand, this is where semantics – and, specifically, the imprecision of the term *function* – can sometimes get in the way of science. It is an obvious point, and barely worthy of scientific attention, that a single mechanism can be construed as serving multiple functions; the alarm system detects motion, and protects the house from intrusion. This is less an observation about the fundamental nature of alarm systems than about the potential variety of our epistemic interests when we ask "What is it doing?" The cases of scientific interest are those where a single mechanism is put to genuinely different uses, the way many people (myself included) use their e-mail inbox also as a "to-do" list. Note the implied converse, that I thereby put different mechanisms – my inbox, my task list—to the *same* use.

So, is it the case, as **Toskos Dils & Flusberg** hypothesize, that the same cognitive behaviors can emerge from different, contextually variable sets of neural circuits? It is an interesting question and worth pursuing. But here is where attention to context, and its role in defining the conditions under which we will call one behavior the "same" as another, becomes crucial to the scientific enterprise. There can be no doubt that the same behaviors mechanically defined (writing on a piece of paper, say, or calculating exchange rates) can involve different neural circuits. But of course writing out a confession, or a love letter, or an essay are vastly different enterprises, and we should expect correspondingly different neural involvement. These would be cases where the neural context *tracks* the task context in ways that are interesting to discover, but also unsurprising.

What would be somewhat surprising is the discovery of different unique sets of circuits for the very same function, where there is no discoverable contextual or other difference to explain the apparent redundancy. Here would be a failure of the neural context to track task context because of the surfeit of neural mechanisms for the task in question. The discovery of such an example would be very interesting and illuminating, although it would not be a specific counterexample to neural reuse. Nothing about the theory suggests that there is only a single neural token instantiating any given type of neural working, much less a single, unique high-level neural network for every identified use. *Some* redundancy, especially in workings, is to be expected; it will sometimes be the case, because energetic constraints favored the outcome, or perhaps just as the result of chance, that different neural circuits developed the same working or came to subserve a similar use. And the discovery that

such redundancy was extensive at the use/network level, and, more importantly, that differences in which networks subserved specified uses did not track context in some identifiable way, would be very puzzling and would affect far more than the theory of neural reuse, including the dynamic models that **Toskos Dils & Flusberg** favor.

Immordino-Yang et al. ask a similarly interesting and challenging series of questions. Can the same network of neural circuits in fact serve quite different uses, a difference that would only become apparent once cultural context was considered? Here again, it is not the least surprising (although not for that reason uninteresting) that cultural context affects which neural resources are brought to bear on cognitive tasks; for, after all, the context may well change (if only in subtle ways) the nature of the task. One would *expect* the neural context to track the environmental/cultural context in this way. What would be harder to assimilate is if it were often the case that the same network subserved *different* uses at the cultural level – genuinely different uses not arising from shift in epistemic perspective – without there being a detectable neural difference. It would be a bigger challenge because this would imply the existence of many cases where neural context does not track environmental context, and this would leave a large explanatory gap that behavioral science has not yet discovered a way to fill. Here again, this would not be a challenge specific to neural reuse; the discovery of radical context dependence in behavior would not undermine the discovery that neural resources are deployed in support of multiple uses across task categories and that differences in uses are better explained by patterns of inter-neural cooperation than by differences of activity in individual brain regions. But it certainly *would* suggest that this was only a part – perhaps a very small part – of the explanation of behavior. There are perhaps some hard-core neuroscientists who think that neural facts are the only relevant inputs to behavioral science, but I am not a member of that tribe, and the implications of neural reuse seem largely orthogonal to that debate.

R2.3. Context and complexity

Still, there is an interesting quasi-dilemma that is illuminated by the advocates of context. Insofar as neural context tracks broader context, then although *initial* attention to context would be necessary to identify the nature of the cognitive task or behavior in question, the lab-minded neuroscientist could then (safely?) focus on the brain, for the broader context would be reflected therein. This somewhat blunts the force of the argument for the necessity of contextualization. On the other hand, although the discovery of cases where neural context did *not* track broader context would appear to strengthen the case for the necessity of contextualization in science, the attendant increase in the complexity of the problem space could encourage the community to simply ignore such cases as beyond the current reach of scientific method. If there is no neural difference despite relevant differences in context and behavior, to what degree are subjects aware of the difference, and controlling their behavior with respect to context? How do the non-neural aspects of intention and control manifest themselves? It is incumbent on advocates

of context to go beyond gestures to dynamic modeling or genetic mechanisms; they must both identify examples of this sort *and* describe an approach to understanding them that promises more illumination than mystification (Chemero 2009). I should be clear that I am not faulting the commentators for not doing so here; this is an ongoing challenge for all of us in the neuro-genetic-social-behavioral sciences.

R2.4. Niche construction

For a final word on the topic of context, it is important to keep in mind the facts that context is itself malleable, and that we are among the most important agents of that change. **Iriki** makes the case not just for the importance of neural niche construction, but also for the possibility that the changing neural context influences our evolutionary pathway, by inducing stable *cultural* changes and thereby changing the environment within which selection operates. Both **Rozin** and **Lindblom** make similar points. Culture (and especially language) can speed up preadaptation, both by increasing the degree and frequency of innovation and by buffering group members against selection pressures that might otherwise tend to weed out initially maladaptive exaptations. There is an extremely interesting literature emerging at the intersection of brains, genes, and culture (Boyd & Richerson 2009; Hawks et al. 2009; Kitayama & Cohen 2007; Richerson et al. 2010; Suomi 2004), and I would be pleased if neural reuse turned out to be a more amenable perspective for coming to grips with these interdependencies than competing proposals on brain organization, such as modularity and localization (something **Reimers** suggests in noting the many parallels between neural reuse and molecular and epigenetic reuse).

R3. Workings 9 to 85

In addition to worrying about my apparent lack of attention to context, **Toskos Dils & Flusberg** and **Immordino-Yang et al.** also question whether the notion of fixed local workings really gives an adequate picture of the functioning of the brain, since it appears to underplay the importance of development and plasticity, a sentiment echoed also by **Aisenberg & Henik**; **Dekker & Karmiloff-Smith**; and **Katz**. I certainly do not want to deny the importance of plasticity to the brain and its functions. But plasticity is a change in use as a result of a change in working. Neural reuse is the acquisition of a new use without a change in working. The target article reviews evidence for importance of the latter process in the functional organization of the brain; it is not an argument against the importance of the former.

R3.1. Workings versus plasticity

Still, neural reuse does suggest that these two processes will be mutually constraining, not to say antagonistic, and that opens some very interesting avenues for further exploration. I think the matter should be framed in the following way. The regions of the developing brain are likely to (and the massive redeployment hypothesis positively predicts that they will) have some innate functional biases, the strength and specificity of which undoubtedly

vary from region to region. Where the nature of experiential input and the characteristics of the task demands being placed on the organism are consistent with these neural biases, then plasticity and reuse can act happily in concert. Neural plasticity generates local workings that reuse can arrange into various circuits that subserve the uses required for the organism's flourishing. (Apropos of which, it should be noted contra **Michaux, Pesenti, Badets, Di Luca, & Andres** [**Michaux et al.**] that nothing in the theory of neural reuse requires denying the necessity of experience to shaping local circuitry; more on this issue in section R4, para. 3.) However, where the nature of the input or the characteristics of the task are *inconsistent* with existing cortical biases or established workings, then these processes can easily come into conflict.

The experiments reported by Sur et al. (1988) and cited by many commentators here offer an excellent paradigm to further explore these sorts of conflicts. As is well known, Sur et al. (1988) rewired the ferret cortex by redirecting right visual field outputs to auditory rather than visual cortex. The result of this manipulation was the induction of neural circuitry in auditory cortex resembling that typically found in normally developing visual cortex – the classic "pinwheel" organization, for instance. The rewired cortex apparently instantiated workings typically associated with visual processing, such as orientation and direction selectivity, and subserved typical visual uses, such as orienting toward a visual stimulus (von Melchner et al. 2000). Plasticity is clearly a powerful force in the development of the brain. It is not, however, omnipotent; visual acuity in the right visual field was significantly lower than in the left visual field. This finding is consistent with the existence of congenital cortical biases potentially in conflict with the nature of sensory inputs, which had to be overcome to accommodate visual stimuli.

From the perspective of neural reuse, it would be interesting to have a more complete behavioral inventory of these animals. Although in this particular case behavioral evidence would have to be treated with caution, given the multiple neural ablations these experiments involved, such information could nevertheless offer some clues as to what other uses the now missing auditory workings might have served. What performance impact did the induction of "visual" circuitry into auditory areas have on *other* functions relying on the same neural region? Were the neural complexes underlying these other behaviors systematically altered by the rewiring? If so, how? Certainly, this is a paradigm that could be used to systematically investigate such questions for various regions of the brain.

Other opportunities for investigating the potential conflicts between plasticity and neural reuse come in the form of manipulations not of neural wiring, but of the task environment, and in particular manipulations of the order in which tasks are learned. Before local neural circuits have acquired very specific workings, and before these workings have been incorporated into multiple functional complexes subserving different uses, it may well be that the most efficient way to acquire novel capacities is inducing plasticity in local circuitry. But later in development such plasticity could prove costly, and learning may favor neural reuse as a strategy. If it is the case that different tasks induce different local workings when

learned early, then it might be possible to systematically investigate the conflicts between plasticity and reuse by looking for specific order effects in learning. For instance, it might be easier to learn task A after task B than after task C, even when A is learned at the same stage of development in each case. (Naturally, choosing appropriate tasks would take some ingenuity; that it will be harder to learn calculus after learning Spanish than after learning algebra is obvious and uninteresting. That it might be easier to learn simple arithmetic after manipulating objects in gloves than after manipulating objects in mittens looks potentially more interesting.) Similarly, it may be that learning task D after A and B disrupts A, but does not do so when learned after tasks A and C, because in the former case the local workings needed to allow for neural reuse as a learning strategy have not developed, leaving plasticity as the only option. **Reimers** suggests some similar developmental studies, and I know that the entire community eagerly awaits the release of the first analyses from the various longitudinal fMRI studies currently underway (Paus 2010).

R3.2. Evolution or development? Both!

In short, I think that the neural reuse model is much friendlier to the developmental perspective than it might have appeared in the target article (**Dekker & Karmiloff-Smith** and **Moore & Moore** rightly point out that development was under-discussed there) and that the two perspectives together suggest some novel avenues for further exploration. I think this account may also shed some light on the issue of *how fixed* I take neural workings to be (**Aisenberg & Henik, Immordino-Yang et al., Toskos Dils & Flusberg**) and how I take them to *be fixed* (**Michaux et al**.). While I think innate cortical biases are a likely feature of our neural organization, workings emerge over time, driven by experience and task demands. Although I think the brain eventually achieves a maturity that is marked in part by the existence of strong and stable workings, plasticity always remains a possibility, whether in response to extraordinary task demands or to physical injury. In this light, I think one can understand neural reuse as a learning strategy that greatly increases the flexibility of the brain while avoiding some of the potentially disruptive effects of local plasticity (especially plasticity that occurs late in development).

This may make it sound like I am giving up on the phylogenetic claims of the massive redeployment hypothesis. I am not. Instead, I am suggesting that the evolutionary and developmental aspects of learning – especially when considered in the neural context – are typically complementary, mutually influencing, and extremely difficult to disentangle. Developmental trajectories, even those highly sensitive to local context, may nevertheless depend on specific evolutionary inheritances. Genetic effects can be influenced by environmental factors such as resource availability (Devlin et al. 2004), and even social factors such as parenting styles (Suomi 2004), and may themselves rely on typical developmental trajectories which, although not themselves hard-coded, have been driven long enough by stable environmental factors to have become established among the dependencies of the genetic pathway.

R3.3. Workings versus polyfunctionality

This may assuage some of the concerns that my workings were too fixed to account for the dynamic nature of the brain, but several commentators question the very notion of local workings. **Aisenberg & Henik**; **Brincker**; **Gomila & Calvo**; **Jungé & Dennett**; **Katz**; **Petrov, Jilk, & O'Reilly** [**Petrov et al.**]; and **Toskos Dils & Flusberg** all argue that local regions might well be natively polyfunctional, obviating the need for any explanation based on reuse. It is true that much of my imaging data is consistent with this possibility, as they show at most that neural regions subserve multiple uses, and multi-use could certainly result from the polyfunctionality of these regions. Moreover, as **Jungé & Dennett**, **Klein**, and **Ritchie & Carruthers** point out, the imaging data are *also* consistent with there being multiple local workings in close proximity, such that the multiple uses only *appear* to use overlapping circuitry due to the poor spatial resolution of current functional imaging techniques. And, as I noted in the target article, these data are even consistent with there being no local functions at all. If brain functions are primarily established not by the structure of local circuitry but by the relative response patterns of neurons or neural assemblies (if, that is, functions are defined by the relational rather than the local properties of neural assemblies), then multi-use could result when these assemblies cooperate with different partners, thereby establishing different relational – and therefore different functional – properties.

But imaging data demonstrating neural overlaps are not the only data I cited, and I think the broader picture sits uneasily with these possibilities. First, there are the data suggesting that more recently evolved uses are subserved by more broadly scattered neural circuits than are older uses. If we may presume an increase in the metabolic cost of establishing and maintaining more broadly scattered functional complexes, then, if polyfunctional local circuits were an option, one would expect uses to be consistently subserved by localized functional complexes. These data seem to favor the existence of local and relatively defined cortical biases. Second, there are the data on cognitive interference and cross-domain interactions. These data would appear to weigh against the possibility of multiple local workings, and favor actually shared neural components. Third – and most telling in my view – are the cases where there appears to be functional or semantic inheritance that results from the sharing of components. This suggests that the functional contributions of the shared local neural circuits are stable and detectible across multiple uses, a possibility apparently inconsistent with both relationally defined function and polyfunctionality. I recognize, of course, that these arguments are more suggestive than definitive, and will be more or less compelling depending on one's other intellectual commitments. In the end, the only evidence that can truly establish the case is the consistent discovery of local workings that can explain the multiple uses to which the local circuit is put. I am laying an empirical bet that this will prove possible, but I must recognize that the evidence may not break my way.

To counter the worries that apparent neural overlaps might be a side effect of the relatively poor spatial resolution of fMRI, **Klein** suggests that experiments leveraging

neural adaptation effects might be in order. Functional magnetic resonance imaging adaptation, fMRIa (Krekelberg et al. 2006), exploits the fact that neurons exposed to the same stimulus attenuate their responses to that stimulus over time, resulting in a decreased BOLD signal. Klein's idea is roughly this: Choose two tasks that appear to involve some overlapping neural circuits, attenuate neural responses by repeatedly engaging in one task, and then switch tasks. If the attenuation disappears in the overlapping region, this would be evidence that "fresh" neurons from a distinct neural subpopulation were responsible for the second task; whereas if the attenuation remained, this would be evidence that the very same neurons were responsible for processing in both tasks.

Let me first of all heartily endorse any and all calls for converging evidence from multiple techniques, and for multiple alternate uses of MRI in particular, for example, Diffusion Tensor Imaging (DTI) (Hagmann et al. 2008); fMRI coherence analysis (Muller et al. 2001; 2003; Sun et al. 2004); Multi-Voxel Pattern Analysis (Periera et al. 2009), and so forth. Nevertheless, I do have some concerns about this particular suggestion. First, although there is good evidence for neural adaptation as the result of repeated stimuli, that is, as the result of being asked to *represent* the same thing, there is less evidence for *task* adaptation, that is, for the idea that there is a reduction in neural response as a result of being asked to *do* the same thing. This matters, because in most cases of neural reuse, the hypothesis is that the region of overlap is *not* being used to represent the same thing in both tasks, so any inherited neural response suppression between tasks would have to result from the region being asked to *do* the same thing, that is, express the same working in both tasks. Second, even if one were to observe neural response suppression as the result of repeating a task, it would remain difficult to interpret the outcome of any experimental manipulation. For, consider the case where the BOLD signal in a region continued to be attenuated during the second task. This could be because they use the same brain regions, and there is some inherited response suppression (apparent evidence for reuse); or it could be because practice at the first task makes the second task easier, or changes the strategy participants use to engage in it (evidence compatible with multiple views). Similarly, if the attenuation disappears, this could be because a distinct neural subpopulation in the same region was being recruited (apparent evidence against reuse); because in the two tasks the same neural populations are being asked to represent different things (compatible with reuse); or because the first task induced neural *changes* that interfere with performance of the second task (apparent evidence for reuse; see Glenberg et al. [2008a] for one such example, and Grill-Specter et al. [2006] for a discussion of various ways to interpret repetition suppression).

For these reasons, I think that fMRIa might be of limited usefulness in teasing apart "real" neural reuse from the possibility that neighboring neural populations are responsible for the different uses to which individual regions are apparently put. As noted above, better techniques for this include cross-domain interference and use-induced plasticity paradigms (Glenberg & Kashak 2002; Glenberg et al. 2008a), and I certainly hope to see more such work in the near future.

Of course, the possibility that the limited spatial resolution of fMRI might be hiding some functional segregation isn't **Klein**'s only worry about that data set. He also wonders whether fMRI activations are particularly good at identifying which brain regions are making genuinely functional contributions to a task in the first place. Rather, activation may spread around the brain network, leading to false positives: regions that are activated only as a side effect of their connectivity, and not because they are making a functional contribution to the task under investigation. He is right, of course; this is a possibility (although not one that cuts against neural reuse in particular). That is why it is crucial to have not just imaging data, but also behavioral data and, where possible, results from techniques like Transcranial Magnetic Stimulation (TMS). If applying TMS over a region thought to be functionally involved in two different tasks in fact disrupts both of those tasks, that is evidence that the activation there is not just a side effect of connectivity, but is making a genuine functional contribution. The target article purposely included data from all of these sources, but here again I would encourage and welcome more studies along these lines.

R3.4. How do workings work?

Even those willing to entertain the possibility that the brain might actually *have* local workings had some questions about how best to understand what they are. **Ritchie & Carruthers,** for instance, ask whether they ought to be understood in computational or intentional terms, and express some skepticism that they could be both computational and multi-use, since it is hard to see how the same computations could realize distinct representational properties on different occasions of use. Rather than repeat or expand upon the arguments from the target article on this matter, I would like instead to refer the reader to the very interesting suggestions made by **Bridgeman**; **Speed, Verzi, Wagner, & Warrender** [**Speed et al.**]; and **Donnarumma, Prevete, & Trautteur** [**Donnarumma et al.**] Bridgeman quite succinctly describes the representational power and flexibility of even fairly simple computational elements, and Donnarumma et al. offer a specific proposal for how this representational flexibility might be harnessed for multiple uses via programmable neural networks. One especially noteworthy aspect of their proposal is a solution to one apparent problem for neural reuse, mentioned also by **Klein** and Ritchie & Carruthers, that reused brain components might send their outputs to all their consumer complexes all the time, which would presumably result in a great deal of noise and behavioral confusion. That this does not appear to be the outcome means either that there is little neural reuse, or that the brain has managed a solution to this problem. Donnarumma et al. offer one model for how selective routing of outputs could be achieved even in a network with fixed connections.

Equally interesting is the proposal made by **Speed et al.** that reuse might be enabled by a mechanism similar to that employed in optical holography. Here, it is somewhat harder to understand what form local workings would take (as these commentators note, in optical holography every piece of the plate encodes the entire image, and nothing like this appears to obtain in the

brain), but that massive reuse is possible on this model is quite clear; and the proposal is notable for detailing the high-level functional predictions that emerge from taking a holographic perspective. Whether these solutions resemble the ones implemented by the brain for managing reuse remains to be seen, of course, but knowing that there exist *possible* solutions is certainly a positive step forward. I look forward to seeing how these diverse research programs evolve.

R4. Embodied cognition: Still standing

Perhaps the most vehement commentators were those objecting to my criticism of embodied cognition, including **Brincker, Michaux et al.,** and **Vilarroya**. Let me be clear: I was an early proponent of embodied cognition (O'Donovan-Anderson 1996; 1997), continue to be a staunch advocate (M. L. Anderson 2003; 2008b; 2009; M. L. Anderson & Rosenberg 2008), and do not think that any of the arguments made in the target article falsify any of the claims made on behalf of the embodied perspective. What I do think is that *embodied* cognition is only one form of a much larger phenomenon of *borrowed* cognition, driven largely by neural reuse. This most certainly does not mean that the kinds of semantic inheritance from sensorimotor to higher-order cognitive systems so important to embodied accounts can be explained away. Quite the contrary: they urgently need to be explained. The worst effect my arguments will have on advocates of embodied cognition (a limitation apparently lamented by **Gomila & Calvo,** who wish I had pressed the critique further) is to strip away the illusion that the semantic inheritance observed in so many domains was ever actually explained by the discovery of neural overlaps. But such disillusionment should be welcomed by any scientist, as it lays down the direction of future research.

Therefore, although **Vilarroya** is right in his argument that we need better methods for attributing functions to brain regions, he is wrong to think that without the ability to specify local workings, it is not possible to establish the limitations of extant models of cognition such as concept empiricism and conceptual metaphor theory. First of all, I do not criticize these theories per se; I argue that not all of the evidence taken to support the theories in fact does so without further assumptions, including especially the assumption that neural overlaps imply semantic inheritance. My evidence shows that this assumption is unwarranted. For this argument, one does not need to know what the workings of any given region in fact are. Rather, one needs to know what some of the *uses* are. It is apparent that in some cases of overlap – as between spatial schemas and evaluative concepts – the working underlying the (presumably earlier) spatial use exerts a structural and semantic influence on the (presumably later) conceptual/linguistic use (hence "better" is also conceptually "higher" or "above"). But in other cases, there seems no obvious evidence for such inheritance. The borrowing of spatial resources for number storage revealed by the SNARC effect (Dehaene et al. 1993), the use of gesturing in learning (Goldin-Meadow 2003), and the use of a common brain region for both finger and magnitude representation (Penner-Wilger & Anderson 2008; submitted; Zago et al. 2001), can all be explained

by positing that both later and earlier use share some functional requirements, such that one or more of the workings underlying the earlier use can also be of service in supporting the later use. In such cases, there need not be, and we in fact see no evidence in these particular cases for, any conceptual grounding or semantic inheritance between these domains as a result of these overlaps.

Michaux et al. object to this line of reasoning in the specific case of the overlap between finger and number representations, but in fact all the evidence they cite is compatible with the functional inheritance account (see Penner-Wilger & Anderson, submitted, for a more detailed account). As noted already above, we argue that the observed overlap results from the fact that one of the functional needs in both domains is for a specific sort of ordered storage. If this is the case, any activity that tended to improve the functionality of the shared working would be expected to improve performance in both cognitive domains. Hence, one doesn't need to posit semantic inheritance to explain the finding that finger differentiation exercises improve math performance (Gracia-Bafalluy & Noël 2008). In fact, this same finding suggests that although sensorimotor experience can be crucial to the development of numerical cognition, insofar as it helps establish the functional structure of the brain regions used in both domains, the crucial experience needn't be of using the fingers to do mathematics. Exercises that improve the sensory acuity of finger representations could be expected to improve performance on certain numerical tasks, without the further requirement that the actual fingers be used in a mathematical context.

Similarly, whenever there is a shared resource, the overlapping uses would have the potential to interfere with one another. That there is indeed such interference between finger and number representations (e.g., Badets & Pesenti 2010) is therefore not surprising. More specifically, Penner-Wilger and Anderson (2008) predicted that there should be a set of self-interfering counting procedures, just in virtue of the fact that on some procedures the representations of which fingers had been touched or otherwise attended to would be incompatible with the representations of which number the fingers were standing in for (that is, the respective representation consumers would systematically misinterpret the content of the representations that resulted from the procedure). Once again, this explains the differences in performance (the number and kinds of mistakes, for example) observed when using nonstandard counting procedures (Di Luca et al. 2006) without needing to posit any semantic inheritance between the domains.

Note that this at least begins to answer the question about encapsulation raised by **Petrov et al.** Neural reuse predicts that encapsulation will *not* be a prominent feature of functional complexes, precisely because in sharing parts each will have access to the information stored and manipulated by the others. Naturally, it is not the case that everything overlaps or communicates with everything else; there is definite and detectable structure to the functional arrangements. Hence, as Petrov et al. correctly describe, the degree of relative encapsulation between functional complexes will depend on the specifics of the physical overlaps and functional connections between them.

Finally, I think the evidence from the cross-cultural imaging study (Tang et al. 2006) raised by **Michaux**

et al. favors neither view. Anyone would predict differences in the relative contributions of some brain regions to otherwise similar cognitive tasks if the methods by which the tasks are taught, and the tools used to achieve them, were significantly different. The evidence simply does not bear on the question of the nature of the inheritance in this case. Nevertheless, I certainly agree with Michaux et al. that it remains to be explained how number concepts acquire their meanings. It may well be that the practice of finger counting *can* play a role in establishing number semantics, but it seems equally clear that there must be other avenues, because not all children who are proficient in math can or do use their fingers in this way. Much more research along these lines is needed.

Brincker lays down a broader and more radical objection to my critique of embodied cognition. She questions whether the evidence that significant neural overlaps occur between regions that are not canonically associated with sensorimotor functions actually shows that neural reuse is a broader phenomenon than can be accounted for by embodied cognition. After all, if higher functions like language are grounded in and by sensorimotor engagement, then reuse of language areas is simply further evidence for the reuse of (fundamentally) sensorimotor circuitry. One problem with this objection is that it ignores the other sources of evidence in support of my claim. But the main trouble is that the critique comes dangerously close to the following argument: All observations of neural overlap – all repeated uses of neural circuitry – are compatible with embodied cognition, because all task domains are ultimately grounded in the sensorimotor system. That argument would indeed undermine my claim that neural reuse is a broader scope phenomenon that can be accounted for by embodied cognition, concept empiricism, and conceptual metaphor theory, but it equally undermines the claim that any specific observation of reuse is evidence *for* these theories. That this is not the way such discoveries have generally been interpreted suggests that this more radical view of the scope of the embodied cognition hypothesis is not widely shared. Moreover, the constraint that *all* task domains (and all aspects of all tasks) must be grounded in sensorimotor systems requires that we read prototypes of all the functional aspects of higher-order cognitive systems into the grounding system.

The case in point here is language, which **Brincker**'s view *requires* that motor control systems have a means-end intentional structure, because language has that structure and language is built upon motor control. As it happens, I am a fan of this particular hypothesis in the case of language (M. L. Anderson 2007b), and so I look forward to the detailed exposition I expect will be offered by Brincker (forthcoming). But the requirement seems far too stringent to apply more generally. Must all the semantic and functional characteristics of recent systems be inherited directly from sensorimotor analogues? Can nothing novel emerge? I am doubtful that evolution has strictly observed this constraint.

A somewhat more subtle challenge along similar lines is offered by **Kiverstein**. He suggests that although semantic inheritance may not be an outcome of every functional borrowing, the cases where there *is* such inheritance play a particularly crucial role in our intellectual evolution,

because only in these cases is there the possibility of bootstrapping from lower- to higher-order functions. The idea is that one hallmark of higher-order cognition is its use of symbols and abstraction, but when these representations are not grounded in sensorimotor systems, they remain semantically empty. Thus, bootstrapping useful higher-order systems out of older parts will require semantic inheritance. Kiverstein is right that the symbol grounding problem is a crucial one for cognitive science (Harnad 1990); that neural reuse does not by itself solve the problem; and that the embodied cognition movement offers some of the more promising approaches to it (M. L. Anderson 2003).

But I think there are at least two things to explain in bootstrapping. One is indeed the origins of any contentful representations manipulated in these systems; but the other is the specific functional character of the system itself. Although I agree that neural reuse alone doesn't address the content issue, I think it goes further toward addressing the functional one than does embodied cognition alone, because it leverages the full power of combinatorix (**Lindblom**) in undergirding new functional arrangements. Moreover, I think that the discovery of neural reuse shows that the embodied cognition movement actually hasn't got as close to solving the content issue as has often been supposed, precisely *because* mere reuse doesn't explain semantic inheritance. I see neural reuse as a strong ally to embodied cognition – and here I think **Kiverstein** agrees – but one that in the near term will be taking up the role of Socratic gadfly.

R5. Reuse, reuse everywhere

One of the more striking aspects of the set of commentaries was the number of *additional* examples of reuse they discuss. **Katz** cites evidence for the reuse of neural circuits across species; **Immordino-Yang et al.** offer discussion of the reuse of the somatosensory system in the processing of social emotions, numerical circuits in recognizing social hierarchy, and the oxytocin system in mother-infant bonding and parental pair bonding; **Niven & Chittka** review many examples of the redeployment of individual neurons for multiple uses in invertebrates; **Bargh, Williams, Huang, Song, & Ackerman** [**Bargh et al.**] discuss various physical-to-psychological effects that suggest reuse in the underlying control systems; **Reimers** reviews some striking analogies between neural reuse and reuse in genetics; **Fishbein, Lau, DeJesús, & Alger** [**Fishbein et al.**] suggest that the sleep cycle may have been redeployed for various purposes; **Rozin** notes that there can be reuse not just of actual circuits, but of the developmental processes or plan that generated them; **Foglia & Grush** note that physical objects like fingers, drawings, and chessboards are reused in multiple representational and task contexts; and **Michaux et al.** review further evidence for the overlaps between motor control and mathematical processing. This range of examples is evidence for the power of the concept of reuse as an organizational frame, but at the same time it greatly complicates the task of specifying a theory adequate to the variety. Perhaps, as **Lindblom** suggests, its reach must necessarily exceed its grasp, and there can be no universal theory of reuse, but only a group of

explanations individually applying to sub-classes of a more general phenomenon.

Although I agree that the theory is *not* fully adequate as it stands – and said as much in the target article – I am not quite ready to abandon the project of specifying a unified theory of neural reuse. And it is perhaps worth noting that **Bargh et al.** found the theory helpful in interpreting some of the fascinating findings coming out of their lab, demonstrating the influence on one's social judgments of others of apparently unrelated physical stimuli such as the warmth of a coffee cup or the weight of a clipboard; **Rabaglia & Marcus** suggest it may help explain the positive manifold – inter-individual performance correlations observed across disparate cognitive tasks; and **Kiverstein** avers that it offers a useful frame for understanding the evolutionary mechanisms for bootstrapping. Although the theory *is* currently underspecified, it is nevertheless still useful.

R5.1. Some additional classes of reuse

Foglia & Grush suggest that one way to further specify the theory is to distinguish between neurophysiological reuse – the physical use of neural circuits to support multiple tasks – and representational reuse – the reuse of physical and mental stand-ins for multiple purposes. They further divide the latter category into "domain" and "functional" reuse: the reuse of a model in one domain (space) to support tasks in another domain (time) versus the reuse of a model in a single domain for multiple purposes (visual representations for both on-line processing and imagining). From their perspective, what is most striking is the brain's capacity to use and reuse models across diverse contexts. Foglia & Grush suggest that *this* may be the more general principle driving instances of neurophysiological reuse, and that the latter occurs when the models in question are neutrally instantiated.

I think the distinction between neurophysiological and representational (model) reuse is a good one, and I agree that our ability to reuse models across contexts is a crucial cognitive ability (Landy et al., in press; Landy & Goldstone 2007a; 2007b). However, I don't think it is right that the latter category falls under the former. Instead, I think these are largely but not entirely overlapping sets: There is model reuse without neural reuse (using a chessboard to play chess, and as a calendar); model reuse with neural reuse (using perceptual systems for imagination); and, I would argue, neural reuse *without* model reuse. For an example of the last category, consider again the case of using the fingers to represent numbers, raised by **Foglia & Grush** and **Michaux et al.** One way to use the fingers to support mathematical reasoning is to use them as models of numbers, and Butterworth (1999c) argues that this results in and explains the observed neural overlap between neural circuits involved in finger-representation and number-representation. But I think the evidence points to a different explanation for the observed overlap: *infrastructural* reuse without model reuse (Penner-Wilger & Anderson 2008; submitted; Penner-Wilger 2009). Here, the idea is that part of the finger representation circuit just happens to have a functional structure that lends itself to supporting certain aspects of number representation. It is not because the fingers are or can be used as models (although

they certainly are and can), nor is the neural circuit being used as a model of anything; it is simply being *used* because it can serve the purpose. Although this example involved reuse of *neural* infrastructure, one imagines there could be important cases of non-neural infrastructural reuse – the use of the hands as an alternative representational resource that aids learning (Goldin-Meadow 2003) may be one such case.

Thus, it seems there are three large classes of cognitively relevant reuse: neural reuse, model reuse, and infrastructural reuse. None of these classes entirely reduces to the others. **Foglia & Grush** further divide model reuse into cross-domain and intra-domain reuse (I drop their term "functional" here, since all of this reuse seems functional to me), and, following **Rozin** and **Fishbein et al.,** we can divide infrastructural reuse into structural token reuse, physiological process reuse, and developmental plan reuse. A cross-cutting category is reuse resulting in semantic inheritance, which **Kiverstein** has suggested has an especially important role in cognitive bootstrapping. Naturally, the target article was centrally concerned with neural reuse, and with understanding when such reuse is (merely) infrastructural, when it involves reuse of models, and when it results in semantic inheritance. But Foglia & Grush are quite right to draw our attention to the cognitive importance of non-neural reuse as well.

R5.2. Does reuse rule out modularity?

One side-effect of the apparent underspecification of neural reuse theory is that it seemed to several commentators – including **Toskos Dils & Flusberg, Jungé & Dennett,** and **Ritchie & Carruthers** – to leave neural reuse closer to modularity than my own rhetoric on the matter would indicate. For example, Jungé & Dennett suggest that a software theory of modularity that posits a modular organization at the abstract level, with no commitment about neural implementation, could survive the critique offered by neural reuse. And Ritchie & Carruthers argue that their *specific* version of massive modularity is in fact compatible with neural reuse. Indeed, one might argue that *no* largely functionalist account of the mind, insofar as it leaves open the many possibilities for implementation, would have any intellectual friction with an account of how neural resources are deployed in the service of cognitive functions. Although I can see the attraction of this position, I think it doesn't apply to the case of modularity. Any theory of modularity worthy of the name must have *modules*, of course, and these modules need to have some specific functional characteristics, such as relative encapsulation, or functional separability. These characteristics in fact place limits on the way such modules can be implemented and, in my view of how our brain is organized, this means they cannot be implemented there.

Ritchie & Carruthers try two (apparently incompatible) tacks to avoid this conclusion: First, they suggest that it would be possible to functionally separate two part-sharing modules via double-dissociation, just so long as one did this by disrupting parts that were not shared between them; and, second, they suggest that maybe modules don't share parts after all, since my imaging evidence is indeed compatible with there being distinct neural regions, smaller than the spatial granularity of

fMRI, dedicated to individual modules. I have already discussed why I think this second possibility is not likely to obtain, but will note here that even if the brain were like that, this argument would fail to demonstrate the compatibility of massive modularity and neural reuse. For what it would show is *not* that these two theories were compatible, but that neural reuse was false.

Unfortunately for *one* of our two theories, I think the first argument fares no better. It is certainly true (and not in conflict with neural reuse) that there will be pairs of functional complexes that are functionally dissociable because they share no parts. And it is also true that even functional complexes that share *some* parts can be distinguished by their reactions to disruptions of the parts not shared. But although the claim that *for all modules A and B it will be possible to functionally dissociate them by disrupting any of their respective parts X and Y* may be too strong for even the most dedicated modularist, surely a claim of the following logical form is too weak to distinguish between any competing theories of brain organization: that *there exist some modules A and B that can be distinguished by disrupting some of their respective parts X and Y*. Wouldn't that be true on just about *anyone's* theory of brain organization? So the fact that we both accept that statement doesn't make for a particularly strong alliance. Yet, I don't see that **Ritchie & Carruthers** have anything stronger to offer here. And if it is the case, as reuse predicts, that in disrupting region X one might not disrupt functional complex B, but *would* disrupt some complex C (and often many complexes C, D, . . .), then even though specific pairs of functional complexes will be functionally separable, it would appear that functional separability will not be a *general* characteristic of the brain. But this is exactly the core claim of massive modularity. I am forced to conclude once again that the two theories are incompatible and, as noted in the target article, in fact point cognitive science in very different empirical directions.

This being said, it is important to note that the term *module* is used in many different ways in many different disciplines, and many of these senses of module are compatible with neural reuse. For instance, in graph theory the term *module* is often used to refer to a set of nodes that are highly connected to one another, and less connected with other parts of the graph. Note that this is a *structural* rather than a functional characterization. Modules are defined in terms of features of the abstract topology of the representational vehicle: the graph. Nevertheless, one of the reasons graphs have proven a useful representational format is that these abstract structures often identify functionally relevant features of the underlying system. In molecular and developmental biology, for instance, a "module" is a set of interacting elements – genes, gene networks, proteins, brain regions, and so forth – that make a specific, relatively context-insensitive contribution to some developmental process (Rives & Galitski 2003; Spirin & Mirny 2003; Tong et al. 2004; von Dassow & Munro 1999) wherever it is instantiated. This sense of module is roughly equivalent to what I have been calling a functional complex, and is perfectly compatible with the notion that the elements of the functional "module" cooperate with different sets of partners to support other outcomes in other circumstances. And, indeed, we know that developmental modules share parts and are often

nested as components of larger modules (Jablonka & Lamb 2006; Schlosser & Wagner 2004).

This is a perfectly viable use of the term *module*, but note that these modules are individuated in ways quite distinct from the *mental* modules posited by evolutionary psychology (e.g., Pinker 1997). Mental modules are entities with abstract *functional* characteristics (encapsulation, domain specificity, functional separability, etc.) and flexible structural characteristics. In contrast, typical biological modules (as found in gene networks or neural co-activation graphs, for example) are entities with well-defined abstract *structural* characteristics but flexible functional characteristics. As tidy as it would be for neuroscience if the modules in neural co-activation graphs identified brain structures with the functional features of mental modules, that is not the way the brain is organized. Therefore, it is important in debates about brain organization and function to try to keep the different senses of "module" distinct; it is all too easy to let them blur into one another.

R5.3. Reuse and evolutionary theory

One place where neural reuse theory is somewhat short on specifics involves its precise fit with existing evolutionary theory. The massive redeployment hypothesis, for instance, is based in part on an overly simplified, armchair-evolutionary story. I think we can and should do better than this. Thankfully, several of the commentators point the way to a better integration of reuse and evolution.

Moore & Moore and **Bergeron** both suggest that the concept of homology can serve as an organizing framework for the further exploration of reuse in evolutionary and developmental context. Bergeron argues that we ought to search for cross-species *cognitive homologies* – workings with the same phylogenetic origins serving different uses in different animals – and he offers some evidence that the search will prove fruitful. Such discoveries would not only help further specify the evolutionary mechanisms behind neural reuse, but could also offer some unique insights into the cognitive relatedness of various species.

Naturally, such a project would involve a great deal of comparative work. **Katz** and **Niven & Chittka** are rightly dismayed by the dearth of comparative data (in my defense, I plead lack of both expertise and space). These authors offer many examples of reuse *in* other species, and Katz in particular offers evidence for just the sorts of cognitive homologies that **Bergeron** suspects should exist. One interesting upshot from the commentaries of both **Katz** and **Niven & Chittka** is that invertebrates may prove the most promising class of animals for initial investigations. All I can say is I think that sounds like a great idea, and hope that someone – if not these authors, then some enterprising young graduate students – will successfully take up the challenge.

Moore & Moore have a somewhat different take on the same general idea. They argue that the concept of homology can *also* be applied in developmental context, whenever two (or more) psychological traits or processes share a neural circuit that has been put to different uses. In this case, we may have identified a *developmental homology*, a shared ontogenetic "ancestor" circuit being used in

different ways. I agree that this perspective offers the possibility of leveraging useful analogies from the evolutionary literature to form hypotheses about features of developmental change, and think that it can help continue the important process of integrating these two perspectives. There are some disanalogies to be careful of as well, however. Chief among these is the fact that traditional evolutionary and cognitive homologies in different species are far less functionally entangled than developmental homologies in a single animal. Whatever limitations imposed by the nature of the inheritance, the use one species subsequently makes of that inheritance does not affect the uses made by others. This is not the case with developmental homologies, where subsequent use can affect the functional properties of other uses, if only because of increased processing demand on the shared circuit. Thus, in cross-species homologies it is more possible to alter the properties of the underlying inheritance, whereas in a developmental homology changing the nature of the shared circuit could have deleterious consequences. Nevertheless, when properly noted, I think both the analogies *and* the disanalogies will prove a fruitful source of developmental hypotheses for future investigations.

Exploring the parallels with reuse in genetics offers another very promising avenue both for hypothesis generation and for coming to better understand the mechanisms of neural reuse. As **Reimers** details, there are many examples of reuse in molecular evolution: Protein domains can be used for multiple purposes; novel metabolic pathways are often assembled by reusing and adapting parts from existing pathways; and signaling pathways are widely reused throughout development. That there is a such a pattern of structure-to-function mapping at the molecular level suggests, among other things, that the neural overlaps I uncovered by reviewing fMRI experiments are not going to go away with higher-resolution imaging techniques. There is too much to be gained functionally by taking advantage of reuse and recombination for this strategy, evident at the micro level, to be absent from the macro level.

R5.4. A history of reuse

Bridgeman and **Lia** both do the community a service by placing the neural reuse hypothesis in historical context, pointing out some intellectual forebears of the idea in addition to those identified in the target article. Awareness of history is important to scientific progress, for while intellectual cycles are an inevitable by-product of the social structure of science, we can at least try to notice whether we are on a spiral staircase or a high-school track. Right now, the cognitive sciences feel to me more like the former than the latter – and neural reuse theory seems a genuine advance – but others should of course make their own judgments.

R6. Where do we go from here?

By this point, I hope it will be evident to the reader that, with a lot of help from the commentaries, neural reuse offers a useful, well-specified, and potentially research-guiding perspective in the cognitive sciences. Several commentaries offer specific techniques and other suggestions for future research. Many have been discussed already, but in this last section I would like to briefly acknowledge a few that have not yet been mentioned. As I noted in the target article, there is work here for many labs; I hope that at least a few of them are inspired to take some of it up.

In the case of reuse that emerges in the course of development, **Reimers** and **Gomila & Calvo** suggest that developmental brain studies of various sorts would be very useful, perhaps especially those that focus on identifying the properties of the networks responsible for high-level cognitive function. I couldn't agree more. The next several years should see the release of data from longitudinal studies tracking both the change in structural connectivity (DTI) and functional connectivity over the course of human brain development (Paus 2010). The opportunity to see how these networks change over time and achieve their adult configuration will be an incredible boon not just for research on neural reuse, but across the whole spectrum of neuroscience.

Possible guides for research in this area are suggested by **Fishbein et al., Foglia & Grush,** and **Rozin**. Perhaps one can use observations of the reuse of physiological processes, of models, and of developmental plans to guide the search for neural reuse. Clearly, not every instance of such reuse will involve the reuse of neural circuitry, but many probably will.

And last, but certainly not least, **Lia** suggests that we should start thinking seriously about the potential clinical applications of both neural reuse and of the scientific techniques that, in light of widespread reuse, ought achieve greater prominence. Perhaps most important is the use of data-mining and meta-analysis of large imaging databases (Anderson et al. 2010; Fox et al. 1998). The amount of information we have about the functional organization of the brain is astounding but, as I have been arguing, we have too often been looking at that information through the wrong lens. I hope to have provided a better one, and with it – or, more likely, with a refined version of it – it should be possible for a dedicated group of researchers equipped with ample computational resources and expertise in data extraction (are you listening, Google?) to mine the many thousands of existing imaging studies to give an accurate model of the functional behavior of the brain under many different circumstances. Will the outcome of such a project be clinically relevant? It seems to me that such an exercise could begin to lay the foundations for baseline expectations of normal brain function across tasks – the identification of typical use-selective networks – which can be as necessary a part of improving our understanding neurological disorders as the discovery of healthy cholesterol ratios was to improving our understanding of heart disease.

Having a good measure of "normal" baseline function, one can begin to catalog the various deviations from these expectations, and their relations to psychiatric diagnoses. Of course, it may not prove possible to do so, but the payoff for success could be quite significant. The ability to define neural signatures for certain disorders can play a role in their diagnosis, of course, but it may also help with our ongoing attempts to categorize and understand them; the discovery that two distinct disorders appear to result from quantitatively similar deviations

from baseline expectations (e.g., increased coherence between regions generally only loosely coupled; or the substitution of one region for another in a known functional complex) might lead to a reassessment of the similarity of the disorders; likewise, the finding of two distinct signatures for a single identified disorder could be part of the argument for splitting the designation.

As goes our understanding, so follows our clinical recommendations. In the ideal case, the features of the neural signatures could themselves suggest treatment options (e.g., can rTMS or deep brain stimulation be used to entrain brain regions to one another? Do neural overlaps suggest which indirect behavioral therapies might be effective?). But even without suggesting particular therapies, knowing a patient's neural signature could sharpen the clinical picture, if one can discover a relation between features of that signature and the range of therapies to which a patient is likely to respond. Perhaps neural signatures will turn out to be as important a part of providing personalized medical care as gene sequencing (Ginsberg & McCarthy 2001; Westin & Hood 2004). Such possibilities are of course quite distant. But we have the technology and the ingenuity. Why not put them to work?

NOTES
1. This attribution was refined after discovering multiple other uses for this circuit (see Penner-Wilger 2009; Penner-Wilger & Anderson, submitted).

2. Thanks to the Editor for this particular suggestion.

References

[The letters "a" and "r" before author's initials stand for target article and response references, respectively.]

Abler, W. L. (1989) On the particulate principle of self-diversifying systems. *Journal of Social and Biological Structures* 12:1–13. [BLin]

Ackerman, J. A., Nocera, C. C. & Bargh, J. A. (2010) Incidental haptic sensations influence social judgments and decisions. *Science* 328:1712–15. [JAB]

Alger, S. E., Lau, H. & Fishbein, W. (2010) Delayed onset of a daytime map facilitates retention of declarative memory. *PLoS ONE* 5(8);e12131. doi:10.1371/journal.pone.001213. [WF]

Anderson, J. R. (2007) *How can the human mind occur in the physical universe?* Oxford University Press. [aMLA, AAP]

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C. & Qin, Y. (2004) An integrated theory of mind. *Psychological Review* 111:1036–60. [AAP]

Anderson, J. R., Qin, Y., Junk, K. -J. & Carter, C. S. (2007) Information processing modules and their relative modality specificity. *Cognitive Psychology* 57:185–217. [aMLA]

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

Anderson, M. L. (2007a) Evolution of cognitive function via redeployment of brain areas. *The Neuroscientist* 13:13–21. [aMLA, AAP]

Anderson, M. L. (2007b) Massive redeployment, exaptation, and the functional integration of cognitive operations. *Synthese* 159(3):329–45. [arMLA, CK]

Anderson, M. L. (2007c) The massive redeployment hypothesis and the functional topography of the brain. *Philosophical Psychology* 21(2):143–74. [aMLA, CK, DSM, AAP]

Anderson, M. L. (2008a) Circuit sharing and the implementation of intelligent systems. *Connection Science* 20(4):239–51. [aMLA]

Anderson, M. L. (2008b) Evolution, embodiment and the nature of the mind. In: *Beyond the brain: Embodied, situated and distributed cognition*, ed. B. Hardy-Vallee & N. Payette, pp. 15–28. Cambridge Scholar's Press. [rMLA]

Anderson, M. L. (2008c) On the grounds of x-grounded cognition. In: *The Elsevier handbook of cognitive science: An embodied approach*, ed. P. Calvo & T. Gomila, pp. 423–35. Elsevier. [aMLA]

Anderson, M. L. (2009) What mindedness is. *Europe's Journal of Psychology* 5(3). Available at: http://www.ejop.org/archives/2009/11/what_mindedness.html. [rMLA]

Anderson, M. L., Brumbaugh, J. & Şuben, A. (2010) Investigating functional cooperation in the human brain using simple graph-theoretic methods. In: *Computational neuroscience*, ed. A. Chaovalitwongse, P. M. Pardalos & P. Xanthopoulos, pp. 31–42. Springer. [arMLA]

Anderson, M. L. & Oates, T. (2010) A critique of multi-voxel pattern analysis. *Proceedings of the 32nd Annual Meeting of the Cognitive Science Society*, ed. S. Ohlsson and R. Catrambone, pp. 1511–16. Cognitive Science Society. [aMLA]

Anderson, M. L. & Rosenberg, G. (2008) Content and action: The guidance theory of representation. *Journal of Mind and Behavior* 29(1–2):55–86. [rMLA]

Anderson, M. L. & Silberstein, M. D. (submitted) Constraints on localization as an explanatory strategy in the biological sciences. [aMLA]

Andres, M., Di Luca, S. & Pesenti, M. (2008) Finger counting: The missing tool? *Behavioral and Brain Sciences* 31:642–43. [NM]

Andres, M., Seron, X. & Oliver, E. (2007) Contribution of hand motor circuits to counting. *Journal of Cognitive Neuroscience* 19:563–76. [aMLA, NM]

Arbas, E. A., Meinertzhagen, I. A. & Shaw, S. R. (1991) Evolution in nervous systems. *Annual Review of Neuroscience* 14:9–38. [PSK]

Ashkenazi, S., Henik, A., Ifergane, G. & Shelef, I. (2008) Basic numerical processing in left intraparietal sulcus (IPS) acalculia. *Cortex* 44:439–48. [DA]

Atallah, H. E., Frank, M. J. & O'Reilly, R. C. (2004) Hippocampus, cortex, and basal ganglia: Insights from computational models of complementary learning systems. *Neurobiology of Learning and Memory* 82(3):253–67. [aMLA, AAP]

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

Baddeley, A. D. (1986) *Working memory*. Oxford University Press. [aMLA]

Baddeley, A. D. (1995) Working memory. In: *The cognitive neurosciences*, ed. M. S. Gazzaniga, pp. 755–64. MIT Press. [aMLA]

Baddeley, A. D. & Hitch, G. (1974) Working memory. In: *The psychology of learning and motivation*, ed. G. H. Bower, pp. 647–67. Erlbaum. [aMLA]

Baddeley, A. D. & Hitch, G. (1994) Developments in the concept of working memory. *Neuropsychology* 8:485–93. [aMLA]

Badets, A. & Pesenti, M. (2010) Creating number semantics through finger movement perception. *Cognition* 115:46–53. [rMLA, NM]

Barabási, A.-L. & Albert, R. (1999) Emergence of scaling in random networks. *Science* 286:509–12. [aMLA]

Barabási, A.-L., Albert, R. & Jeong, H. (2000) Scale-free characteristics of random networks: The topology of the World Wide Web. *Physica A* 281:69–77. [aMLA]

Bargh, J. A. (2006) What have we been priming all these years? On the development, mechanisms, and ecology of nonconscious social behavior. *European Journal of Social Psychology* 36:147–68. [JAB]

Bargh, J. A. & Morsella, E. (2008) The unconscious mind. *Perspectives on Psychological Science* 3:73–79. [JAB]

Barkow, J. H., Cosmides, L. & Tooby, J., eds. (1992) *The adapted mind: Evolutionary psychology and the generation of culture*. Oxford University Press. [aMLA]

Barrett, H. C. & Kurzban, R. (2006) Modularity in cognition: Framing the debate. *Psychological Review* 113(3):628–47. [aMLA]

Barsalou, L. W. (1999) Perceptual symbol systems. *Behavioral and Brain Sciences* 22:577–660. [aMLA]

Barsalou, L. W. (2008) Grounded cognition. *Annual Review of Psychology* 59:617–45. [aMLA]

Bechtel, W. (2003) Modules, brain parts, and evolutionary psychology. In: *Evolutionary psychology: Alternative approaches*, ed. S. J. Scher & F. Rauscher, pp. 211–27. Kluwer. [aMLA]

Bechtel, W. & Richardson, R. C. (1993) *Discovering complexity: Decomposition and localization as strategies in scientific research*. Princeton University Press. [aMLA]

Bechtel, W. & Richardson, R. C. (2010) *Discovering complexity: Decomposition and localization as strategies in scientific research*, 2nd edition. MIT Press/ Bradford Books. [aMLA]

Behrens, T. E. & Johansen-Berg, H. (2005) Relating connectional architecture to grey matter function using diffusion imaging. *Philosophical Transactions of the Royal Society of London, B: Biological Sciences* 360:903–11. [aMLA, AG]

Bellugi, U., Lichtenberger, L., Mills, D., Galaburda, A. & Korenberg, J. R. (1999) Bridging cognition, the brain and modular genetics: Evidence from Williams syndrome. *Trends in Neuroscience* 22:197–207. [TMD]

Bergeron, V. (2007) Anatomical and functional modularity in cognitive science: Shifting the focus. *Philosophical Psychology* 20(2):175–95. [aMLA, AAP]

Bergeron, V. (2008) *Cognitive architecture and the brain: Beyond domain-specific functional specification*. Unpublished doctoral dissertation, Department of

Philosophy, University of British Columbia. Available at: http://circle.ubc.ca/handle/2429/2711. [arMLA]

Beuthel, R. G., Pohl, H. & Hunefeld, F. (2005) Strepsipteran brains and effects of miniaturization (Insecta). *Arthropod Structure and Development* 34:301–13. [JEN]

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* 11:273–85. [aMLA]

Bock, W. J. (1959) Preadaptation and multiple evolutionary pathways. *Evolution* 13:194–211. [PR]

Boroditsky, L. & Ramscar, M. (2002) The roles of body and mind in abstract thought. *Psychological Science* 13(2):185–88. [aMLA]

Botvinick, M. M., Cohen, J. D. & Carter, C. S. (2004) Conflict monitoring and anterior cingulate cortex: An update. *Trends in Cognitive Sciences* 8:539–46. [DA]

Bowlby, J. (1969) *Attachment and loss.* Hogarth Press. [JAB]

Boyd, R. & Richerson, P. J. (2009) Culture and the evolution of human cooperation. *Philosophical Transactions of the Royal Society of London, B: Biological Sciences* 364:3281–88. [rMLA]

Boyer, D., Miramontes, O., Ramos-Fernández, G., Mateos, J. L. & Cocho, G. (2004) Modeling the searching behavior of social monkeys. *Physica A* 342:329–35. [aMLA]

Brigandt, I. & Griffiths, P. E. (2007) The importance of homology for biology and philosophy. *Biology and Philosophy* 22:633–41. [DSM]

Briggman, K. L. & Kristan, W. B. (2008) Multifunctional pattern-generating circuits. *Annual Review of Neuroscience* 31:271–94. [PSK]

Brincker, M. (forthcoming) Moving beyond mirroring – A social affordance model of sensorimotor integration during action perception. Doctoral dissertation, Department of Philosophy, Graduate Center, City University of New York. (forthcoming in September 2010). [rMLA, MB]

Broca, P. (1861) Remarques sur le siège de la faculté du langage articulé, suivies d'une observation d'aphémie (perte de la parole). *Bulletin de la Société Anatomique* 6:330–57. [BB]

Brooks, R. (1991) Intelligence without representation. *Artificial Intelligence* 47:139–60. [JK]

Brown, C. T., Larry S. Liebovitch, L. S. & Glendon, R. (2007) Lévy flights in dobe ju/'hoansi foraging patterns. *Human Ecology* 35:129–38. [aMLA]

Brown, J., Johnson, M. H., Paterson, S., Gilmore, R., Gsödl, M., Longhi, E. & Karmiloff-Smith, A. (2003) Spatial representation and attention in toddlers with Williams syndrome and Down syndrome. *Neuropsychologia* 41:1037–46. [TMD]

Buckner, R. L, Andrews-Hanna, J. R. & Schacter, D.L. (2008) The brain's default network: Anatomy, function and relevance to disease. *Annals of the New York Academy of Sciences* 1124:1–38. [MB]

Burrows, M. (1996) *The neurobiology of an insect brain.* Oxford University Press. [JEN]

Buss, D., Haselton, M. G., Shackleford, T. K., Bleske, A. L. & Wakefield, J. C. (1998) Adaptations, exaptations and spandrels. *American Psychologist* 53:533–48. [PR]

Buszáki, G. (1998) Memory consolidation during sleep: A neurophysiological perspective. *Journal of Sleep Research* 7(Suppl. 1):17–23. [WF]

Butterworth, B. (1999a) A head for figures. *Science* 284:928–29. [NM]

Butterworth, B. (1999b) *The mathematical brain.* Macmillan. [NM]

Butterworth, B. (1999c) *What counts – How every brain is hardwired for math.* The Free Press. [arMLA]

Cabeza, R. & Nyberg, L. (2000) Imaging cognition II: An empirical review of 275 PET and fMRI studies. *Journal of Cognitive Neuroscience* 12:1–47. [aMLA]

Caetano-Anolles, G., Wang, M., Caetano-Anolles, D. & Mittenthal, J. E. (2009a) The origin, evolution and structure of the protein world. *Biochemical Journal* 417:621–37. [MR]

Caetano-Anolles, G., Yafremava, L. S., Gee, H., Caetano-Anolles, D., Kim, H. S. & Mittenthal, J. E. (2009b) The origin and evolution of modern metabolism. *International Journal of Biochemistry and Cell Biology* 41:285–97. [MR]

Calabrese, R. L. (1998) Cellular, synaptic, network, and modulatory mechanisms involved in rhythm generation. *Current Opinion in Neurobiology* 8:710–17. [PSK]

Cantlon, J. F. & Brannon, E. M. (2007) Adding up the effects of cultural experience on the brain. *Trends in Cognitive Sciences* 11(1):1–4. [NM]

Cantlon, J. F., Platt, M. L. & Brannon, E. M. (2009) Beyond the number domain. *Trends in Cognitive Sciences* 13:83–91. [DA]

Carroll, J. B. (1993) *Human cognitive abilities: A survey of factor analytic studies.* Cambridge University Press. [CDR]

Carroll, S. B., Grenier, J. K. & Weatherbee, S. D. (2005) *From DNA to diversity: Molecular genetics and the evolution of animal design.* Blackwell. [MR]

Carruthers, P. (2002) The cognitive functions of language. *Behavioral and Brain Sciences* 25(6):657–74. [aMLA]

Carruthers, P. (2006) *The architecture of the mind: Massive modularity and the flexibility of thought.* Clarendon Press/Oxford University Press. [aMLA, JAJ, AAP, JBR]

Casasanto, D. & Boroditsky, L. (2008) Time in the mind: Using space to think about time. *Cognition* 106:579–93. [aMLA]

Casasanto, D. & Dijkstra, K. (2010) Motor action and emotional memory. *Cognition* 115(1):179–85. [aMLA]

Casey, B. J., Tottenham, N., Liston, C. & Durston, S. (2005) Imaging the developing brain: What have we learned about cognitive development? *Trends in Cognitive Sciences* 9(3):104–10. [AG]

Catania, K. C. (2000) Cortical organization in insectivora: The parallel evolution of the sensory periphery and the brain. *Brain, Behavior, and Evolution* 55:311–21. [PSK]

Catania, K. C. & Remple, F. E. (2004) Tactile foveation in the star-nosed mole. *Brain, Behavior, and Evolution* 63:1–12. [rMLA]

Changizi, M. A. & Shimojo, S. (2005) Character complexity and redundancy in writing systems over human history. *Proceedings of the Royal Society of London B: Biological Sciences* 272:267–75. [aMLA]

Changizi, M. A., Zhang, Q., Ye, H. & Shimojo, S. (2006) The structures of letters and symbols throughout human history are selected to match those found in objects in natural scenes. *American Naturalist* 167:E117–39. [aMLA]

Chao, L. L. & Martin A. (2000) Representation of manipulable man-made objects in the dorsal stream. *NeuroImage* 12:478–84. [aMLA, ATD]

Chemero, A. (2009) *Radical embodied cognitive science.* MIT Press. [arMLA]

Cherniak, C., Mokhtarzada, Z., Rodrigues-Esteban, R. & Changizi, K. (2004) Global optimization of cerebral cortex layout. *Proceedings of the National Academy of Sciences USA* 101:1081–86. [aMLA, AAP]

Chiang, M. C., Barysheva, M., Shattuck, D. W., Lee, A. D., Madsen, S. K., Avedissian, C., Klunder, A. D., Toga, A. W., McMahon, K. L., de Zubicaray, G. I., Wright, M. J., Srivastava, A., Balov, N. & Thompson, P. M. (2009) Genetics of brain fiber architecture and intellectual performance. *Journal of Neuroscience* 29:2212–24. [CDR]

Chiao, J. Y., Harada, T., Komeda, H., Li, Z., Mano, Y., Saito, D. N., Parrish, T. B., Sadato, N. & Iidaka, T. (2009a) Neural basis of individualistic and collectivistic views of self. *Human Brain Mapping* 30(9):2813–20. [MHI-Y]

Chiao, J. Y., Harada, T., Komeda, H., Li, Z., Mano, Y., Saito, D. N., Parrish, T. B., Sadato, N. & Iidaka, T. (2010) Dynamic cultural influences on neural representations of the self. *Journal of Cognitive Neuroscience* 22(1):1–11. [MHI-Y]

Chiao, J. Y., Harada, T., Oby, E. R., Li, Z., Parrish, T. & Bridge, D. J. (2009b) Neural representations of social status hierarchy in human inferior parietal cortex. *Neuropsychologia* 47(2):354–63. [MHI-Y]

Chittka, L. & Niven, J. (2009) Are bigger brains better? *Current Biology* 19:R995–1008. [JEN]

Clark, A. (1997) *Being there: Putting brain, body, and world together again.* MIT Press. [aMLA, MB]

Clark, A. (1998) Embodied, situated, and distributed cognition. In: *A companion to cognitive science,* ed. W. Bechtel & G. Graham, pp. 506–17. Blackwell. [aMLA]

Clark, H. H. (1973) Space, time, semantics, and the child. In: *Cognitive development and the acquisition of language,* ed. T. E. Moore, pp. 27–63. Academic Press. [JAB]

Clayton, N. S. & Russell, J. (2009) Looking for episodic memory in animals and young children: Prospects for a new minimalism. *Neuropsychologia* 47:2330–40. [DSM]

Cohen Kadosh, R., Lammertyn, J. & Izard, V. (2008) Are numbers special? An overview of chronometric, neuroimaging, developmental and comparative studies of magnitude representation. *Progress in Neurobiology* 84(2):132–47. [MHI-Y]

Cohen, L. G., Celnik, P., Pascual-Leone, A., Corwell, B., Falz, L., Dambrosia, J., Honda, M., Sadato, N., Gerloff, C., Catala, M. D. & Hallett, M. (1997) Functional relevance of cross-modal plasticity in blind humans. *Nature (London)* 389:180–83. [PSK]

Coltheart, M. (1999) Modularity and cognition. *Trends in Cognitive Sciences* 3:115–20. [JBR]

Coltheart, M. (2001) Assumptions and methods in cognitive neuropsychology. In: *The handbook of cognitive neuropsychology,* ed. B. Rapp, pp. 3–21. Psychology Press. [aMLA]

Coltheart, M. (2006) What has functional neuroimaging told us about the mind (so far)? *Cortex* 42(3):323–31. [rMLA]

Comer, C. M. & Robertson, R. M. (2001) Identified nerve cells and insect behavior. *Progress in Neurobiology* 63:409–39. [PSK]

Cormier, S. M. (1987) The structural processes underlying transfer of training. In: *Transfer of learning: Contemporary research and applications,* ed. S. M. Cormier & J. D. Hagman, pp. 152–82. Academic Press. [AS]

Costafreda, S. G., Fu, C. H. Y., Lee, L., Everitt, B., Brammer, M. J. & David, A. S. (2006) A systematic review and quantitative appraisal of fMRI studies of verbal fluency: Role of the left inferior frontal gyrus. *Human Brain Mapping* 27(10):799–810. [aMLA]

Crinion, J., Turner, R., Grogan, A., Hanakawa, T., Noppeney, U., Devlin, J. T., Aso, T., Urayama, A., Fukuyama, H., Stockton, K., Usui, K., Green, D. W. & Price, C. J. (2006) Language control in the bilingual brain. *Science* 312(5779):1537. [CK]

Croce, J. C. & McClay, D. R. (2008) Evolution of the Wnt pathways. *Methods in Molecular Biology* 469:3–18. [MR]

Croll, R. P. (1987) Identified neurons and cellular homologies. In: *Nervous systems in invertebrates*, ed. M. A. Ali, pp. 41–59. Plenum Press. [PSK]

Culham, J. C. & Valyear, K. F. (2006) Human parietal cortex in action. *Current Opinion in Neurobiology* 16:205–12. [aMLA]

Dagher, A., Owen, A., Boecker, H. & Brooks, D. (1999) Mapping the network for planning. *Brain* 122:1973–87. [aMLA]

Damasio, A. & Tranel, D. (1993) Nouns and verbs are retrieved with differently distributed neural systems. *Proceedings of the National Academy of Sciences USA* 90:4957–60. [aMLA]

Damasio, H., Grabowski,T. J., Tranel, D., Hichwa, R. D. & Damasio, A. R. (1996) A neural basis for lexical retrieval. *Nature* 380:499–505. [aMLA]

Darwin, C. (1862) *On the various contrivances by which British and foreign orchids are fertilised by insects, and on the good effects of intercrossing.* John Murray. [aMLA]

Deacon, T. (1997). *The symbolic species.* Norton. [aMLA]

Deary, I. J., Penke, L. & Johnson, W. (2010) The neuroscience of human intelligence differences. *Nature: Neuroscience* 11:201–11. [CDR]

Deary, I. J., Spinath, F. M. & Bates, T. C. (2006) Genetics of intelligence. *European Journal of Human Genetics* 14:690–700. [CDR]

Decety, J. & Chaminade, T. (2003) Neural correlates of feeling sympathy. *Neuropsychologia* 41(2):127–38. [MHI-Y]

Decety, J. & Grèzes, J. (1999) Neural mechanisms subserving the perception of human actions. *Trends in Cognitive Sciences* 3:172–78. [aMLA]

Decety, J., Grezes, 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* 120:1763–77. [aMLA]

Decety, J., Sjoholm, H., Ryding, E., Stenberg, G. & Ingvar, D. (1990) The cerebellum participates in cognitive activity: Tomographic measurements of regional cerebral blood flow. *Brain Research* 535:313–17. [aMLA]

Dehaene, S. (2005) Evolution of human cortical circuits for reading and arithmetic: The "neuronal recycling" hypothesis. In: *From monkey brain to human brain*, ed. S. Dehaene, J.-R. Duhamel, M. D. Hauser & G. Rizolatti, pp. 133–57. MIT Press. [DA, aMLA, TMD]

Dehaene, S. (2009) *Reading in the brain.* Viking. [aMLA]

Dehaene, S., Bossini, S. & Giraux, P. (1993) The mental representation of parity and numerical magnitude. *Journal of Experimental Psychology: General* 122:371–96. [aMLA, VB]

Dehaene, S. & Cohen, L. (2007) Cultural recycling of cortical maps. *Neuron* 56:384–98. [DA, aMLA, DSM, PR]

Dehaene, S., Piazza, M., Pinel, P. & Cohen, L. (2003) Three parietal circuits for number processing. *Cognitive Neuropsychology* 20(3):487–506. [MHI-Y]

DeJesús, R., Lau, H., Alger, S. & Fishbein, W. (in preparation) Nocturnal sleep enhances retention of emotional memories: Total sleep deprivation, and to a greater extent REM and stage II sleep deprivation impedes the retention enhancement. *Abstract, Society for Neuroscience.* [WF]

Devlin, R. H., D'Andrade, M., Uh, M. & Biagi, C. A. (2004) Population effects of growth hormone transgenic coho salmon depend on food availability and genotype by environment interactions. *Proceedings of the National Academy of Sciences USA* 101(25):9303–308. [rMLA]

Di Luca, S. & Pesenti, M. (2008) Masked priming effect with canonical finger numeral configurations. *Experimental Brain Research* 185(1):27–39. [NM]

Di Luca, S., Graná, A., Semenza, C., Seron, X. & Pesenti, M. (2006) Finger-digit compatibility in Arabic numerical processing. *The Quarterly Journal of Experimental Psychology* 59(9):1648–63. [rMLA, NM]

Dijksterhuis, A., Chartrand, T. L. & Aarts, H. (2007) Effects of priming and perception on social behavior and goal pursuit. In: *Social psychology and the unconscious: The automaticity of higher mental processes*, ed. J. A. Bargh, pp. 51–132. Psychology Press. [JAB]

Donald, M. (1991) *Origins of the modern mind.* Harvard University Press. [BLin]

Donaldson, Z. R., Kondrashov, F. A., Putnam, A., Bai, Y., Stoinski, T. L., Hammock, E. A. & Young, L. J. (2008) Evolution of a behavior-linked microsatellite-containing element in the 5′ flanking region of the primate AVPR1A gene. *BioMed Central Evolutionary Biology* 8:180. [PSK]

Donaldson, Z. R. & Young, L. J. (2008) Oxytocin, vasopressin, and the neurogenetics of sociality. *Science* 322:900–904. [PSK]

Donnarumma, F. (2010) *A model for programmability and virtuality in dynamical neural networks.* Doctoral dissertation in Scienze Computazionali ed Informatiche (Computational and Information Sciences), Dipartimento di Matematica e Applicazioni "R. Caccioppoli," Università di Napoli Federico II. Available at: http://people.na.infn.it/~donnarumma/files/donnarumma09model.pdf. [FD]

Donnarumma, F., Prevete, R. & Trautteur, G. (2007) Virtuality in neural dynamical systems. Poster presented at the International Conference on Morphological Computation, ECLT, Venice, Italy, March 26–28, 2007. Available at: http://vinelab.na.infn.it/research/pubs/donnarumma07virtuality.pdf. [FD]

Duncan, J. (2001) An adaptive coding model of neural function in prefrontal cortex. *Nature Reviews Neuroscience* 2:820–29. [DA]

Dunn, N. A., Lockery, S. R., Pierce-Shimomura, J. T. & Conery, J. S. (2004) A neural network model of chemotaxis predicts functions of synaptic connections in the nematode *Caenorhabditis elegans. Journal of Computational Neuroscience.* 17(2):137–47. [FD]

Edelman, G. M. (1987) CAMs and Igs: Cell adhesion and the evolutionary origins of immunity. *Immunology Review* 100:11–45. [MR]

Eguiluz, V. M., Chialvo, D. R., Cecchi, G., Baliki, M. & Apkarian, A. V. (2005) Scale-free brain functional networks. *Physical Review Letters* 94:18102. [AG]

Ehrlich, I., Humeau, Y., Grenier, F., Ciocchi, S., Herry, C. & Luthi, A. (2009) Amygdala inhibitory circuits and the control of fear memory. *Neuron* 62:757–71. [MR]

Eisenberger, N. I. & Lieberman, M. D. (2004) Why rejection hurts: A common neural alarm system for physical and social pain. *Trends in Cognitive Sciences* 8(7):294–300. [MHI-Y]

Ellenbogen, J. M., Hu, P. T., Payne, J. D., Titone, D. & Walker, M. P. (2007) Human relational memory requires time and sleep. *Proceedings of National Academy of Science USA* 104(18):7723–28. [WF]

Elsabbagh, M., Cohen, H., Cohen, M., Rosen, S. & Karmiloff-Smith, A. (in press) Severity of hyperacusis predicts individual differences in speech perception in Williams syndrome. *Journal of Intellectual Disability Research.* [TMD]

Fauconnier, G. & Turner, M. (2002) *The way we think: Conceptual blending and the mind's hidden complexities.* Basic Books. [aMLA]

Fayol, M, Barrouillet, P. & Marinthe, C. (1998) Predicting arithmetical achievement from neuropsychological performance: A longitudinal study. *Cognition* 68:63–70. [NM]

Fedorenko, E., Patel, A., Casasanto, D., Winawer, J. & Gibson, T. (2009) Structural integration in language and music: Evidence for a shared system. *Memory and Cognition* 37(1):1–9. [aMLA]

Feldman, J. & Narayanan, S. (2004) Embodied meaning in a neural theory of language. *Brain and Language* 89:385–92. [aMLA]

Fessler, D. & Navarrete, D. (2004) Third-party attitudes towards sibling incest: Evidence for the Westermarck hypothesis. *Evolution and Human Behavior* 24:277–94. [JBR]

Fiebach, C. J. & Schubotz, R. I. (2006) Dynamic anticipatory processing of hierarchical sequential events: A common role for Broca's area and ventral premotor cortex across domains? *Cortex* 42(4):499–502. [VB]

Finn, R. D., Mistry, J., Tate, J., Coggill, P., Heger, A., Pollington, J. E., Gavin, O. L., Gunasekaran, P., Ceric, G., Forslund, K., Holm, L., Sonnhammer, E. L., Eddy, S. R. & Bateman, A. (2010) The Pfam protein families database. *Nucleic Acids Research* 38:D211–22. [MR]

Fiske, S. T., Cuddy, A. J. C. & Glick, P. (2007) Universal dimensions of social cognition: Warmth and competence. *Trends in Cognitive Sciences* 11:77–83. [JAB]

Fodor, J. (1975) *The language of thought.* Harvard University Press. [aMLA]

Fodor, J. (1983) *The modularity of mind.* MIT Press. [ATD, AAP]

Fodor, J. (2000) *The mind doesn't work that way.* MIT Press. [AAP]

Fodor, J. & Pylyshyn, Z. W. (1988) Connectionism and cognitive architecture: A critical analysis. *Cognition* 28:3–71. [aMLA]

Foglia, L. & Grush, R. (in preparation) The limitations of a purely enactive (nonrepresentational) account of imagery. *Journal of Consciousness Studies.* [LF]

Fowler, C. A., Rubin, P., Remez, R. E. & Turvey, M. T. (1980) Implications for speech production of a general theory of action. In: *Language production, vol. 1: Speech and talk*, ed. B. Butterworth, pp. 373–420. Academic Press. [aMLA]

Fox, P. T. & Lancaster, J. L. (2002) Mapping context and content: The BrainMap model. *Nature Reviews Neuroscience* 3:319–21. [aMLA]

Fox, P. T., Parsons, L. M. & Lancaster, J. L. (1998) Beyond the single study: Function-location meta-analysis in cognitive neuroimaging. *Current Opinions in Neurobiology* 8:178–87. [arMLA]

Fries, R. C. (2006) *Reliable design of medical devices.* CRC Press. [aMLA]

Fritsch, G. & Hitzig, E. (1870/1960) On the electrical excitability of the cerebrum. In: *Brain and behaviour: Vol. 2. Perception and action*, ed. K. H. Pribram. Penguin. (Original work published in 1870). [BB]

Fuentemilla, L. I., Cámara, E., Münte, Th., Krämer, U., Cunillera, A., Marco-Pallarés, J., Tempelmann, C. & Rodríguez-Fornells, A. (2009) Individual differences in true and false memory retrieval are related to white matter brain microstructure. *Journal of Neuroscience* 29:8698–703. [AG]

Gallese, V. (2003) A neuroscientific grasp of concepts: From control to representation. *Philosophical Transactions of the Royal Society London, B: Biological Sciences* 358(1435):1231–40. [aMLA, MB]

Gallese, V. (2008) Mirror neurons and the social nature of language: The neural exploitation hypothesis. *Social Neuroscience* 3(3–4):317–33. [aMLA]

Gallese, V., Fadiga L., Fogassi L. & Rizzolatti G. (1996) Action recognition in the premotor cortex. *Brain* 119:593–609. [aMLA]

Gallese, V. & Goldman, A. (1998) Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Sciences* 2(12):493–501. [aMLA]

Gallese, V. & Lakoff, G. (2005) The brain's concepts: The role of the sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology* 22(3–4):455–79. [aMLA]

Gallistel, C. R. (1993) *The organization of learning.* MIT Press. [MHI-Y]

Garcia-Bafalluy, M. & Noël, M.-P. (2008) Does finger training increase young children's numerical performance? *Cortex* 44:368–75. [rMLA, NM]

Garzillo, C. & Trautteur, G. (2009) Computational virtuality in biological systems. *Theoretical Computer Science* 410:323–31. [FD]

Gauthier, I., Curran, T., Curby, K. & Collins, D. (2003) Perceptual interference supports a non-modular account of face processing. *Nature Neuroscience* 6:428–32. [JBR]

Gauthier, I., Skudlarski, P., Gore, J. C. & Anderson, A. W. (2000) Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience* 3:(2):191–97. [aMLA, JBR]

Gentner, D. & Stevens, A. L., eds. (1983) *Mental models.* Erlbaum. [aMLA]

Gest, H. (1987) Evolutionary roots of the citric acid cycle in prokaryotes. *Biochemical Society Symposia* 54:3–16. [MR]

Gibson, J. J. (1979) *The ecological approach to visual perception.* Erlbaum. [aMLA]

Gick, M. L. & Holyoak, K. J. (1983) Schema induction and analogical transfer. *Cognitive Psychology* 15:1–38. [AS]

Gigerenzer, G., Todd, P. M. & The ABC Research Group (1999) *Simple heuristics that make us smart.* Oxford University Press. [aMLA]

Gilovich, T., Griffin, D. & Kahneman, D., eds. (2002) *Heuristics and biases: The psychology of intuitive judgment.* Cambridge University Press. [aMLA]

Ginsberg, J. S. & McCarthy, J. J. (2001) Personalized medicine: Revolutionizing drug discovery and patient care. *Trends in Biotechnology* 19(12):491–96. [rMLA]

Girifalco, L. A. (1991) *Dynamics of technological change.* Van Nostrand Reinhold. [PR]

Glenberg, A. (2010) Embodiment as a unifying perspective for psychology. *Wiley Interdisciplinary Reviews: Cognitive Science* 1(4):586–96. [MB]

Glenberg, A. M., Becker, R., Klötzer, S., Kolanko, L., Müller, S. & Rinck, M. (2009) Episodic affordances contribute to language comprehension. *Language and Cognition* 1:113–35. [aMLA]

Glenberg, A. M., Brown, M. & Levin, J. R. (2007) Enhancing comprehension in small reading groups using a manipulation strategy. *Contemporary Educational Psychology* 32:389–99. [aMLA]

Glenberg, A. M. & Kaschak, M. P. (2002) Grounding language in action. *Psychonomic Bulletin and Review* 9:558–65. [arMLA, BB]

Glenberg, A. M., Sato, M. & Cattaneo, L. (2008a) Use-induced motor plasticity affects the processing of abstract and concrete language. *Current Biology* 18:R290–91. [arMLA]

Glenberg, A. M., Sato, M., Cattaneo, L., Riggio, L., Palumbo, D. & Buccino, G. (2008b) Processing abstract language modulates motor system activity. *Quarterly Journal of Experimental Psychology* 61:905–19. [aMLA]

Goldin-Meadow, S. (2003) *Hearing gesture: How our hands help us think.* Belknap Press. [arMLA]

Goldstein, K. (1963) *The organism: A holistic approach to biology derived from pathological data in man.* Beacon Press. [BLia]

Gomila, A. (2008) Mending or abandoning cognitivism? In: *Symbols, embodiment and meaning*, ed. A. Glenberg, M. de Vega & A. Glaesser, pp. 799–834. Oxford University Press. [AG]

Gould, S. J. (1991) Exaptation: A crucial tool for evolutionary psychology. *Journal of Social Issues* 47:43–65. [PR]

Gould, S. J. & Vrba, E. S. (1982) Exaptation: A missing term in the science of form. *Paleobiology* 8:4–15. [PR]

Graziano, M. S. A., Taylor, C. S. R. & Moore, T. (2002a) Complex movements evoked by microstimulation of precentral cortex. *Neuron* 34:841–51. [aMLA]

Graziano, M. S. A., Taylor, C. S. R., Moore, T. & Cooke, D. F. (2002b) The cortical control of movement revisited. *Neuron* 36:349–62. [aMLA, MB]

Griffiths, P. E. (2007) The phenomena of homology. *Biology and Philosophy* 22:643–58. [DSM]

Grill-Specter, K., Henson, R. & Martin, A. (2006) Repetition and the brain: Neural models of stimulus-specific effects. *Trends in Cognitive Sciences* 10(1):14–23. [rMLA]

Grill-Specter, K., Sayres, R. & Ress, D. (2006) High-resolution imaging reveals highly selective nonface clusters in the fusiform face area. *Nature Neuroscience* 9(9):1177–85. [aMLA]

Grush, R. (2004) The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences* 27:377–442. [LF]

Haggard, P., Rossetti, Y. & Kawato, M., eds. (2008) *Sensorimotor foundations of higher cognition.* Oxford University Press. [aMLA, MB]

Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C. J., Wedeen, V. J. & Sporns, O. (2008) Mapping the structural core of human cerebral cortex. *PLoS Biology* 6(7):e159. Available at: http://biology.plosjournals.org/perlserv/?request=get-document. doi:10.1371/journal.pbio.0060159. [arMLA]

Hagoort, P. (2005) On Broca, brain and binding. *Trends in Cognitive Sciences* 9(9):416–23. [aMLA]

Hall, B. K. (2003) Descent with modification: The unity underlying homology and homoplasy as seen through an analysis of development and evolution. *Biological Reviews* 78:409–33. [DSM]

Hall, J. S. (2009) The robotics path to AGI using servo stacks. In: *Proceedings of the Second Conference on Artificial General Intelligence*, ed. B. Goertzel, P. Hitzler & M. Hutter, pp. 49–54. Atlantis Press. doi:10.2991/agi.2009.5. [aMLA]

Hammock, E. A. & Young, L. J. (2005) Microsatellite instability generates diversity in brain and sociobehavioral traits. *Science* 308:1630–34. [PSK]

Hamzei, F., Rijntjes, M., Dettmers, C., Glauche, V., Weiller, C. & Büchel (2003) The human action recognition system and its relationship to Broca's area: An fMRI study. *Neuroimage* 19:637–44. [aMLA]

Han, S. & Northoff, G. (2008) Culture-sensitive neural substrates of human cognition: A transcultural neuroimaging approach. *Nature Reviews: Neuroscience* 9:646–54. [TMD]

Hanakawa, T., Honda, M., Sawamoto, N., Okada, T., Yonekura, Y., Fukuyama, H. & Shibasaki, H. (2002) The role of rostral Brodmann area 6 in mental-operation tasks: An integrative neuroimaging approach. *Cerebral Cortex* 12:1157–70. [aMLA]

Harlow, H. (1958) The nature of love. *American Psychologist* 13:673–85. [JAB]

Harnad, S. (1990) The symbol grounding problem. *Physica D* 42:335–46. [rMLA]

Hasselmo, M. E. (1999) Neuromodulation: Acetylcholine and memory consolidation. *Trends in Cognitive Sciences* 3(9):351–59. [WF]

Hawks, J., Wang, E. T., Cochran, G. M., Harpending, H. C. & Moyzis, R. K. (2009) Recent acceleration of human adaptive evolution. *Proceedings of the National Academy of Sciences USA* 104 (52):20753–58. [rMLA]

Henik, A. & Tzelgov, J. (1982) Is three greater than five: The relation between physical and semantic size in comparison tasks. *Memory and Cognition* 10:389–95. [DA]

Heyes, C. (2010) Where do mirror neurons come from? *Neuroscience and Biobehavioural Reviews* 34 (4):575–83. [JK]

Ho, T. -Y., Lama, P. -M. & Leung, C. -S. (2008) Parallelization of cellular neural networks on GPU. *Pattern Recognition* 41(8):2684–92. [aMLA]

Hommel, B., Müsseler, J., Aschersleben, G. & Prinz, W. (2001) The theory of event coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences* 24(5):849–78. [MB]

Honey, C. J., Kötter, R., Breakspear, M. & Sporns, O. (2007) Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *Proceedings of the National Academy of Sciences USA* 104:10240–45. [aMLA]

Honey, C. J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J. P., Meuli, R. & Hagmann, P. (2009) Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences USA* 106(6):2035–40. [aMLA]

Hopkin, V. D. (1995) *Human factors in air traffic control.* CRC Press. [aMLA]

Hoßfeld, U. & Olsson, L. (2005) The history of the homology concept and the "Phylogenetisches Symposium." *Theory in Biosciences* 124:243–53. [DSM]

Huang, J. Y. & Bargh, J. A. (2008) Peak of desire: Activating the mating goal changes life stage preferences across living kinds. *Psychological Science* 19:573–78. [JAB]

Hubbard, E. M., Piazza, M., Pinel, P. & Dehaene, S. (2005) Interactions between number and space in parietal cortex. *Nature Reviews Neuroscience* 6(6):435–48. [aMLA, VB]

Huettel, S. A., Song, A. W. & McCarthy, G. (2008) *Functional magnetic resonance imaging.* Sinauer. [AAP]

Hunt, R. R. (1995) The subtlety of distinctiveness: What von Restorff really did. *Psychonomic Bulletin and Review* 2:105–12. [AS]

Hurford, J. (2003) The neural basis of predicate-argument structure. *Behavioral and Brain Sciences* 26(3):261–83. [aMLA]

Hurley, S. L. (1998) *Consciousness in action.* Harvard University Press. [MB]

Hurley, S. L. (2005) The shared circuits hypothesis: A unified functional architecture for control, imitation and simulation. In: *Perspectives on imitation: From neuroscience to social science*, ed. S. Hurley & N. Chater, pp. 76–95. MIT Press. [aMLA]

Hurley, S. L. (2008) The shared circuits model (SCM): How control, mirroring, and simulation can enable imitation, deliberation, and mindreading. *Behavioral and Brain Sciences* 31(1):1–58. [aMLA, BB, JK]

Hutchins, E. (1995) *Cognition in the wild.* MIT Press. [aMLA]

IJzerman, H. & Semin, G. R. (2009) The thermometer of social relations: Mapping social proximity on temperature. *Psychological Science* 20:1214–20. [JAB]

Immordino-Yang, M. H., McColl, A., Damasio, H. & Damasio, A. (2009) Neural correlates of admiration and compassion. *Proceedings of the National Academy of Sciences USA* 106(19):8021–26. [MHI-Y]

Iriki, A. (2005). A prototype of homo-faber: A silent precursor of human intelligence in the tool-using monkey brain. In: *From monkey brain to human brain*, ed. S. Dehaene, J. R. Duhamel, M. Hauser & G. Rizzolati, pp. 133–57. MIT Press. [aMLA]

Iriki, A. (2006) The neural origins and implications of imitation, mirror neurons and tool use. *Current Opinion in Neurobiology* 16:660–67. [AI]

Iriki, A. & Sakura, O. (2008) Neuroscience of primate intellectual evolution: Natural selection and passive and intentional niche construction. *Philosophical Transactions of the Royal Society of London, B: Biological Science* 363:2229–41. [aMLA, AI]

Iriki, A., Tanaka, M. & Iwamura, Y. (1996) Coding of modified body schema during tool use by macaque postcentral neurons. *NeuroReport* 7:2325–30. [AI]

Jablonka, E. & Lamb, M. J. (2006) *Evolution in four dimensions*. MIT Press. [rMLA]

Jacob, F. (1977) Evolution and tinkering. *Science* 196(4295):1161–66. [OV]

Jarvis, E. D., Gunturkun, O., Bruce, L., Csillag, A., Karten, H., Kuenzel, W., Medina, L., Paxinos, G., Perkel, D. J., Shimizu, T., Striedter, G., Wild J. M., Ball, G. F., Dugas-Ford, J., Durand, S. E., Hough, G. E., Husband, S., Kubikova, L., Lee, D. W., Mello, C. V., Powers, A., Siang, C., Smulders, T. V., Wada, K., White, S. A., Yamamoto, K., Yu, J., Reiner, A. & Butler, A. B. (2005) Avian brains and a new understanding of vertebrate brain evolution. *Nature Reviews Neuroscience* 6:151–59. [PSK]

Jeannerod, M. (1994) The representing brain: Neural correlates of motor intention and imagery. *Behavioral and Brain Sciences* 17:187–245. [aMLA]

Jensen, A. R. (1998) The suppressed relationship between IQ and the reaction time slope parameter of the Hick function. *Intelligence* 26:43–52. [CDR]

Jeong, H., Tombor, B., Albert, R., Oltvai, Z. N. & Barabási, A. -L. (2000) The large-scale organization of metabolic networks. *Nature* 407:651–54. [aMLA]

Jilk, D. J., Lebiere, C., O'Reilly, R. C. & Anderson, J. R. (2008) SAL: An explicitly pluralistic cognitive architecture. *Journal of Experimental and Theoretical Artificial Intelligence* 20:197–218. [aMLA, AAP]

Johnson, M. H. (2001) Functional brain development in humans. *Nature Reviews Neuroscience* 2:475–83. [TMD]

Johnson-Laird, P. N. (1983) *Mental models: Towards a cognitive science of language, inference, and consciousness*. Harvard University Press. [aMLA]

Jung, R. E. & Haier, R. J. (2007) The parieto-frontal integration theory (P-FIT) of intelligence: Converging neuroimaging evidence. *Behavioral and Brain Sciences* 30(2):135–54. [CDR]

Kaas, J. H. (2005) The future of mapping sensory cortex in primates: Three of many remaining issues. *Philosophical Transactions of the Royal Society of London, B: Biological Sciences* 360:653–64. [PSK]

Kalivas, P. W. & Volkow, N. D. (2005) The neural basis of addiction: A pathology of motivation and choice. *American Journal of Psychiatry* 162:1403–13. [PSK]

Kalman, E., Kobras, S., Grawert, F., Maltezos, G., Hanssen, H., Coufal, H. & Burr, G. W. (2004) Accuracy and scalability in holographic content-addressable storage. Paper presented at the Conference on Lasers and Electro-Optics (CLEO). San Francisco, CA, May 2004. [AS]

Kanwisher, N., McDermott, J. & Chun, M. (1997) The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience* 17(11):4302–11. [aMLA, JBR]

Karmiloff-Smith, A. (1998) Development itself is the key to understanding developmental disorders. *Trends in Cognitive Sciences* 2:389–98. [TMD]

Karmiloff-Smith, A. (2009) Nativism versus neuroconstructivism: Rethinking the study of developmental disorders. *Developmental Psychology* 45(1):56–63. [TMD]

Karmiloff-Smith, A., Thomas, M. S. C., Annaz, D., Humphreys, K., Ewing, S., Brace, N., van Duuren, M., Pike, G., Grice, S. & Campbell, R. (2004) Exploring the Williams syndrome face processing debate: The importance of building developmental trajectories. *Journal of Child Psychology and Psychiatry* 45:1258–74. [TMD]

Katz, P. S. (1999) *Beyond neurotransmission: Neuromodulation and its importance for information processing*. Oxford University Press. [PSK]

Katz, P. S. & Calin-Jageman, R. (2008) Neuromodulation. In: *New encyclopedia of neuroscience*, ed. L. R. Squire, pp. 497–503. Academic Press. [PSK]

Katz, P. S. & Harris-Warrick, R. M. (1999) The evolution of neuronal circuits underlying species-specific behavior. *Current Opinion in Neurobiology* 9:628–33. [PSK]

Katz, P. S. & Newcomb, J. M. (2007) A tale of two CPGs: Phylogenetically polymorphic networks. In: *Evolution of nervous systems*, ed. J. H. Kaas, pp. 367–74. Academic Press. [PSK]

Kawato, M., Kuroda, T., Imamizu, H., Nakano, E., Miyauchi, S. & Yoshioka, T. (2003) Internal forward models in the cerebellum: fMRI study on grip force and load force coupling. *Progress in Brain Research* 142:171–88. [AG]

Kent, K. S. & Levine, R. B. (1993) Dendritic reorganization of an identified during metamorphosis of the moth *Manduca sexta*: The influence of interactions with the periphery. *Journal of Neurobiology* 24:1–22. [JEN]

Kerns, J. G., Cohen, J. D., MacDonald III, A. W., Cho, R. Y., Stenger, V. A. & Carter, C. S. (2004) Anterior cingulate conflict monitoring and adjustments in control. *Science* 303:1023–26. [DA]

Kirschner, M. W. & Gerhart, J. C. (2005) *The plausibility of life – Resolving Darwin's dilemma*. Yale University Press. [AI]

Kitayama, S. & Cohen, D., eds. (2007) *Handbook of cultural psychology*. Guildford Press. [rMLA]

Klein, C. (2010) Images are not the evidence in neuroimaging. *British Journal for the Philosophy of Science* 61:265–78. [rMLA]

Klein, R. M. (2000) Inhibition of return. *Trends in Cognitive Sciences* 4:138–47. [DA]

Koch, C. & Segev, I. (2000) The role of single neurons in information processing. *Nature Neuroscience* 3:1171–77. [aMLA]

Krekelberg, B., Boynton, G. M. & van Wezel, R. J. A. (2006) Adaptation: From single cells to BOLD signals. *Trends in Neurosciences* 29(5):250–56. [rMLA]

Krubitzer, L. (2007) The magnificent compromise: Cortical field evolution in mammals. *Neuron* 56:201–208. [PSK]

Krubitzer, L. (2009) In search of a unifying theory of complex brain evolution. *Annals of the New York Academy of Sciences* 1156:44–67. [PSK]

Kyllonen, P. C. & Christal, R. E. (1990) Reasoning ability is (little more than) working-memory capacity? *Intelligence* 14:389–433. [CDR]

Laird A. R., Lancaster, J. L. & Fox P. T. (2005) BrainMap: The social evolution of a functional neuroimaging database. *Neuroinformatics* 3:65–78. [aMLA]

Lakoff, G. & Johnson, M. (1980) *Metaphors we live by*. University of Chicago Press. [aMLA]

Lakoff, G. & Johnson, M. (1999) *Philosophy in the flesh: The embodied mind and its challenge to western thought*. Basic Books. [aMLA]

Lakoff, G. & Núñez, R. (2000) *Where mathematics comes from: How the embodied mind brings mathematics into being*. Basic Books. [aMLA]

Landy, D., Allen, C. & Anderson, M. L. (in press) Conceptual discontinuity through recycling old processes in new domains. Commentary on Susan Carey: Précis of *The Origin of Concepts. Behavioral and Brain Sciences* 33(6). [rMLA]

Landy, D. & Goldstone, R. L. (2007a) Formal notations are diagrams: Evidence from a production task. *Memory and Cognition* 35(8):203–340. [rMLA]

Landy, D. & Goldstone, R. L. (2007b) How abstract is symbolic thought? *Journal of Experimental Psychology: Learning, Memory, and Cognition* 33(4):720–33. [rMLA]

Lashley, K. S. (1929) *Brain mechanisms and intelligence*. University of Chicago Press. [BB]

Lau, H., Tucker, M. A. & Fishbein, W. (2010) Daytime napping: Effects on human direct associative and relational memory. *Neurobiology of Learning and Memory* 93(2010):554–60. [WF]

Lee, H., Macbeth, A. H., Pagani, J. H. & Young, W. S., III. (2009) Oxytocin: The great facilitator of life. *Progress in Neurobiology* 88:127–51. [MHI-Y]

Lia, B. (1992) Ontogeny and ontology: Ontophyletics and enactive focal vision. *Behavioral and Brain Sciences* 15(1):43–45. [BLia]

Lim, M. M., Wang, Z., Olazabal, D. E., Ren, X., Terwilliger, E. F. & Young, L. J. (2004) Enhanced partner preference in a promiscuous species by manipulating the expression of a single gene. *Nature (London)* 429:754–57. [PSK]

Lin, Z., Lin, Y. & Han, S. (2008) Self-construal priming modulates visual activity underlying global/local perception. *Biological Psychology* 77(1):93–97. [MHI-Y]

Lindblom, B., Diehl, R., Park, S.-H. & Salvi, G. (in press) Sound systems are shaped by their users: The recombination of phonetic substance. In: *Where do features come from? The nature and sources of phonological primitives*, ed. N. Clements & R. Ridouane. John Benjamins. (Publication will appear in March 2011) [BLin]

Lloyd, D. (2000) Terra cognita: From functional neuroimaging to the map of the mind. *Brain and Mind* 1(1):93–116. [aMLA]

Logothetis, N. K. (2008) What we can do and what we cannot do with fMRI. *Nature* 453:869–78. [CK]

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T. & Oeltermann, A. (2001) Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412(6843):150–57. [CK]

Love, A. C. (2007) Functional homology and homology of function: Biological concepts and philosophical consequences. *Biology and Philosophy* 22:691–708. [VB]

Mahon, B. Z. & Caramazza, A. (2008) A critical look at the embodied cognition hypothesis and a new proposal for grounding conceptual content. *Journal of Physiology-Paris* 102(1–3):59–70. [CK]

Mandler, J. M. (1992) How to build a baby: II. Conceptual primitives. *Psychological Review* 99:587–604. [JAB]

Marcus, G. F. (2004) *The birth of the mind: How a tiny number of genes creates the complexities of human thought*. Basic Books. [aMLA]

Marcus, G. F. (2006) Cognitive architecture and descent with modification. *Cognition* 101:43–65. [CDR]

Marcus, G. F. (2008) *Kluge: The haphazard construction of the human mind*. Houghton Mifflin. [aMLA]

Marcus, G. F. & Rabagliati, H. (2006) The nature and origins of language: How studies of developmental disorders could help. *Nature Neuroscience* 10:1226–29. [CDR]

Marder, E. & Thirumalai, V. (2002) Cellular, synaptic and network effects of neuromodulation. *Neural Networks* 15:479–93. [PSK]

Martin, A., Haxby, J. V., Lalonde, F. M., Wiggs, C. L. & Ungerleider, L. G. (1995) Discrete cortical regions associated with knowledge of color and knowledge of action. *Science* 270:102–105. [aMLA]

Martin, A., Ungerleider, L. G. & Haxby, J. V. (2000) Category-specificity and the brain: the sensorimotor model of semantic representations of objects. In: *The new cognitive neurosciences*, 2nd edition, ed. M. S. Gazzaniga, pp. 1023–36. MIT Press. [aMLA]

Martin, A., Wiggs, C. L., Ungerleider, L. G. & Haxby, J. V. (1996) Neural correlates of category-specific knowledge. *Nature* 379:649–52. [aMLA]

Mayr, E. (1960) The emergence of evolutionary novelties. In: *Evolution after Darwin, vol. 1: The evolution of life*, ed. S. Tax, pp. 349–80. University of Chicago Press. [PR]

McClelland, J. L., McNaughton, B. L. & O'Reilly, R. C. (1995) Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review* 102:419–57. [AAP]

McGraw, L. A. & Young, L. J. (2010) The prairie vole: An emerging model organism for understanding the social brain. *Trends in Neurosciences* 32:103–109. [PSK]

Meier, T., Chabaud, F. & Reichert, H. (1991). Homologous patterns in the embryonic development of the peripheral nervous system in the grasshopper *Schistocerca gregaria* and the fly *Drosophila melanogaster*. *Development* 112:241–53. [PSK]

Melendez-Hevia, E., Waddell, T. G. & Cascante, M. (1996) The puzzle of the Krebs citric acid cycle: Assembling the pieces of chemically feasible reactions, and opportunism in the design of metabolic pathways during evolution. *Journal of Molecular Evolution* 43: 293–303. [MR]

Meltzoff, A. N. & Moore, M. K. (1977) Imitation of facial and manual gestures by human neonates. *Science* 198:75–78. [DSM]

Menzel, R. (2009) Conditioning: Simple neural circuits in the honeybee: In: *Encyclopedia of neuroscience, vol. 3*, ed. L. R. Squire, pp. 43–47. Academic Press. [JEN]

Mesulam, M.-M. (1990) Large-scale neurocognitive networks and distributed processing for attention, language and memory. *Annals of Neurology* 28:597–613. [aMLA, AAP]

Meyrand, P., Faumont, S., Simmers, J., Christie, A. E. & Nusbaum, M. P. (2000) Species-specific modulation of pattern-generating circuits. *European Journal of Neuroscience* 12:2585–96. [PSK]

Miali, R. C. (2003) Connecting mirror neurons and forward models. *NeuroReport* 14(17):2135–37. [aMLA]

Miller, E. K. (2000) The prefrontal cortex and cognitive control. *Nature Reviews Neuroscience* 1:59–65. [DA]

Miller, E. K. & Cohen, J. D. (2001) An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience* 24:167–202. [AAP]

Millikan, R. G. (1984) *Language, thought and other biological categories*. MIT Press. [aMLA]

Mitchell, M. (2006). Complex systems: Network thinking. *Artificial Intelligence* 170:1194–212. [aMLA]

Muller, K., Lohmann, G., Bosch, V. & von Cramon, D. Y. (2001) On multivariate spectral analysis of fMRI time series. *NeuroImage* 14 347–56. [rMLA]

Muller, K., Mildner, T., Lohmann, G. & von Cramon, D. Y. (2003) Investigating the stimulus-dependent temporal dynamics of the BOLD signal using spectral methods. *Journal of Magnetic Resonance Imaging* 17:375–82. [rMLA]

Müller, R.-A. & Basho, S. (2004) Are nonlinguistic functions in "Broca's area" prerequisites for language acquisition? fMRI findings from an ontogenetic viewpoint. *Brain and Language* 89(2):329–36. [aMLA]

Murphy, F. C., Nimmo-Smith, I. & Lawrence, A. D. (2003) Functional neuroanatomy of emotions: A meta-analysis. *Cognitive, Affective and Behavioral Neuroscience* 3(3):207–33. [aMLA]

Nair, D. G. (2005) About being BOLD. *Brain Research Reviews* 50:229–43. [CK]

Newcomb, J. M. & Katz, P. S. (2007) Homologues of serotonergic central pattern generator neurons in related nudibranch molluscs with divergent behaviors. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology* 193:425–43. [PSK]

Newcomb, J. M. & Katz, P. S. (2008) Different functions for homologous serotonergic interneurons and serotonin in species-specific rhythmic behaviours. *Proceedings of the Royal Society of London, B: Biological Sciences* 276:99–108. [PSK]

Newell, A. & Simon, H. A. (1976) Computer science as empirical enquiry. *Communications of the ACM* 19(3):113–26. [aMLA]

Newman, M., Barabasi, A.-L. & Watts, D. J. (2006) *The structure and dynamics of networks*. Princeton University Press. [aMLA]

Newson, L., Richerson, P. J. & Boyd, R. (2007) Cultural evolution and the shaping of cultural diversity. In: *Handbook of cultural psychology*, ed. S. Kitayama & D. Cohen, pp. 454–76. Guilford Press. [PR]

Nishitani, N., Schürmann, M., Amunts K. & Hari, R. (2005) Broca's region: From action to language. *Physiology* 20:60–69. [aMLA]

Niven, J. E., Graham, C. M. & Burrows, M. (2006) Diversity and evolution of the insect ventral nerve cord. *Annual Review of Entomology* 53:253–71. [JEN]

Noël, M.-P. (2005) Finger gnosia: A predictor of numerical abilities in children? *Child Neuropsychology* 11(5):413–30. [NM]

Novick, L. R. (1988) Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 14:510–20. [AS]

Nunez, R. & Freeman, W. (2000) *Reclaiming cognition: The primacy of action, intention, and emotion*. Imprint Academic. [MB]

Nvidia Corporation. (2007) *CUDA Programming Guide, version 1.1*. Santa Clara, CA. Available at: http://developer.download.nvidia.com/compute/cuda/1_1/NVIDIA_CUDA_Programming_Guide_1.1.pdf. [aMLA]

Odling-Smee, F. J., laland, K. N. & Geldman, M. W. (2005) *Niche construction: The neglected process in evolution.* Princeton University Press. [aMLA]

O'Donovan-Anderson, M., ed. (1996) *The incorporated self: Interdisciplinary perspectives on embodiment.* Rowman & Littlefield. [rMLA]

O'Donovan-Anderson, M. (1997) *Content and comportment: On embodiment and the epistemic availability of the world*. Rowman & Littlefield. [rMLA]

Ogawa, A., Yamazaki, Y., Ueno, K., Cheng, K. & Iriki, A. (2010) Neural correlates of species-typical illogical cognitive bias in human inference. *Journal of Cognitive Neuroscience* 22:2120–30. [AI]

Ogawa, A., Yamazaki, Y., Ueno, K., Cheng, K. & Iriki, A. (in press) Inferential reasoning by exclusion recruits parietal and prefrontal cortices. *NeuroImage*. doi:10.1016/j.neuroimage.2010.05.040. [AI]

Orban, G. A., Van Essen, D. & Vanduffel, W. (2004) Comparative mapping of higher visual areas in monkeys and humans. *Trends in Cognitive Sciences* 8:315–24. [TMD]

O'Reilly, R. C. (1998) Six principles for biologically based computational models of cortical cognition. *Trends in Cognitive Sciences* 2:455–62. [aMLA, AAP]

O'Reilly, R. C. (2006) Biologically based computational models of high-level cognition. *Science* 314(5796):91–94. [AAP]

O'Reilly, R. C., Braver, T. S. & Cohen, J. D. (1999) A biologically based computational model of working memory. In: *Models of working memory: Mechanisms of active maintenance and executive control*, ed. A Miyake & P. Shah, pp. 375–411. Cambridge University Press. [AAP]

O'Reilly, R. C. & Frank, M. J. (2006) Making working memory work: A computational model of learning in the prefrontal cortex and basal ganglia. *Neural Computation* 18:283–328. [AAP]

O'Reilly, R. C. & Munakata, Y. (2000) *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. MIT Press. [aMLA, AAP]

Owen, R. (1843) *Lectures on the comparative anatomy and physiology of the invertebrate animals, delivered at the Royal College of Surgeons, in 1843.* Longman, Brown, Green, and Longmans. [VB, PSK]

Padberg, J., Franca, J. G., Cooke, D. F., Soares, J. G., Rosa, M. G., Fiorani, M., Jr., Gattass, R. & Krubitzer, L. (2007) Parallel evolution of cortical areas involved in skilled hand use. *Journal of Neuroscience* 27:10106–15. [PSK]

Pagán Cánovas, C. (2009) La emisión erótica en la poesía griega: una familia de redes de integración conceptual desde la Antigüedad hasta el siglo XX. Departmento Filología Clásica, Universidad de Murcia, Spain. http://www.tesisenred.net/TDR-0519110-103532/index.html. [aMLA]

Panksepp, J. (2005) Why does separation distress hurt? Comment on MacDonald and Leary (2005). *Psychological Bulletin* 131(2):224–30. [MHI-Y]

Parker, G., Cheah, Y. C. & Roy, K. (2001) Do the Chinese somatize depression?: A cross-cultural study. *Social Psychiatry and Psychiatric Epidemiology* 36:287–93. [MHI-Y]

Patel, A. D. (2003) Language, music, syntax and the brain. *Nature Reviews Neuroscience* 6(7):674–81. [VB]

Paterson, S. J., Brown, J. H., Gsodl, M. K., Johnson, M. H. & Karmiloff-Smith, A. (1999) Cognitive modularity and genetic disorders. *Science* 286:2355–58. [TMD]

Paus, T. (2010) Population neuroscience: Why and how. *Human Brain Mapping* 31(6):891–903. [rMLA]

Payne, J. D., Schacter, D. L., Propper, R. E., Huang, L. W., Wamsley, E. J., Tucker, M. A., Walker, M. P. & Stickgold, R. (2009) The role of sleep in false memory formation. *Neurobiology of Learning and Memory* 92(3):327–34. [WF]

Penner-Wilger, M. (2009) *Subitizing, finger gnosis, and finger agility as precursors to the representation of number.* Unpublished doctoral dissertation, Department of Cognitive Science, Carleton University, Ottawa, Canada. http://gradworks.umi.com/NR/52/NR52070. [arMLA]

Penner-Wilger, M. & Anderson, M. L. (2008) An alternative view of the relation between finger gnosis and math ability: Redeployment of finger representations for the representation of number. In: *Proceedings of the 30th Annual Meeting of the Cognitive Science Society, Austin, TX, July 23–26, 2008*, ed. B. C. Love, K. McRae & V. M. Sloutsky, pp. 1647–52. Cognitive Science Society. [arMLA, NM]

Penner-Wilger, M. & Anderson, M. L. (submitted) The relation between finger recognition and mathematical ability: Why redeployment of neural circuits best explains the finding. [arMLA]

Pereira, F., Mitchell, T. & Botvinick, M. M. (2009) Machine learning classifiers and fMRI: A tutorial overview. *NeuroImage* 45:S199–209. [arMLA]

Perkel, D. J. (2004) Origin of the anterior forebrain pathway. *Annals of the New York Academy of Sciences* 1016:736–48. [PSK]

Pesenti, M., Thioux, M., Seron, X. & De Volder, A. (2000) Neuroanatomical substrate of Arabic number processing, numerical comparison and simple addition: A PET study. *Journal of Cognitive Neuroscience* 121(3):461–79. [NM]

Pessoa, L. (2008) On the relationship between emotion and cognition. *Nature Reviews Neuroscience* 9:148–58. [aMLA]

Petrides, M., Cadoret, G. V. & Mackey, S. (2005) Orofacial somatomotor responses in the macaque monkey homologue of Broca's area. *Nature* 435(7046):1235–38. [VB]

Phan, K. L, Wager, T., Taylor, S. F. & Liberzon, I. (2002) Functional neuroanatomy of mmotion: A meta-analysis of emotion activation studies in PET and fMRI. *NeuroImage* 16(2):331–48. [aMLA]

Piaget, J. (1952) *The child's conception of number*. Routledge and Kegan Paul. [aMLA]

Pinker, S. (1997) *How the mind works*. Norton. [aMLA, ATD]

Pinker, S. (1999) *Words and rules: The ingredients of language*. Basic Books. [TMD]

Plate, T. A. (1995) Holographic reduced representations. *IEEE Transactions on Neural Networks* 6(3):623–41. [AS]

Plaut, D. C. (1995) Double dissociation without modularity: Evidence from connectionist neuropsychology. *Journal of Clinical and Experimental Neuropsychology* 17:291–321. [aMLA]

Poldrack, R. A. (2006) Can cognitive processes be inferred from neuroimaging data? *Trends in Cognitive Sciences* 10:59–63. [arMLA, AAP]

Popovici, C., Roubin, R., Coulier, F. & Birnbaum, D. (2005) An evolutionary history of the FGF superfamily. *Bioessays* 27:849–57. [MR]

Posner, M. I. & Cohen, Y. (1984) Components of visual orienting. In: *Attention and performance X*, ed. H. Bouma & D. Bouwhuis, pp. 531–56. Erlbaum. [DA]

Postuma, R. B. & Dagher, A. (2006) Basal ganglia functional connectivity based on a meta-analysis of 126 PET and fMRI publications. *Cerebral Cortex* 16(10):1508–21. [aMLA]

Pribram, K. H. (1971) *Languages of the brain.* Prentice-Hall. [BB]

Prinz, J. (2002) *Furnishing the mind: Concepts and their perceptual basis*. MIT Press. [aMLA]

Prinz, J. (2006) Is the mind really modular? In: *Contemporary debates in cognitive science*, ed. R. J. Stainton, pp. 22–36. Blackwell. [aMLA, AAP]

Psaltis, D. & Burr, G. W. (1998) Holographic data storage. *Computer* February 1998:52–60. [AS]

Pulvermüller, F. (2005) Brain mechanisms linking language and action. *Nature Reviews Neuroscience* 6:576–82. [aMLA]

Quallo, M. M., Price, C. J., Ueno, K., Asamizuya, T., Cheng, K., Lemon, R. N. & Iriki, A. (2009) Gray and white matter changes associated with tool-use learning in macaque monkeys. *Proceedings of the National Academy of Sciences USA* 106:18379–84. [AI]

Quartz, S. R. & Sejnowski, T. J. (1997) The neural basis of cognitive development: A constructivist manifesto. *Behavioral and Brain Sciences* 20:537–56. [aMLA]

Quince, C., Higgs, P. G. & McKane, A. J. (2002) Food web structure and the evolution of ecological communities. In: *Biological evolution and statistical physics: Lecture notes in Physics* 585, ed. M. Laessig & A. Valleriani, pp. 281–98. Springer-Verlag. [aMLA]

Rabaglia, C. D. & Marcus, G. F. (in preparation) Individual differences in sentence comprehension: Beyond working memory. [CDR]

Rasmussen, J. & Vicente, K. J. (1989) Coping with human errors through system design: implications for ecological interface design. *International Journal of Man-Machine Studies* 31:517–34. [aMLA]

Rauschecker, J. P. & Scott, S. K. (2009) Maps and streams in the auditory cortex: Nonhuman primates illuminate human speech processing. *Nature Reviews Neuroscience* 12(6):718–24. doi: 10.1038/nn.2331. [VB]

Rhodes, G., Byatt, G., Michie, P. T. & Puce, A. (2004) Is the Fusiform Face Area specialized for faces, individuation, or expert individuation?. *Journal of Cognitive Neuroscience* 16(2):189–203. [aMLA]

Richardson, D., Spivey, M., Barsalou, L. & McRae, K. (2003) Spatial representations activated during real-time comprehension of verbs. *Cognitive Science* 27:767–80. [aMLA]

Richerson, P. J., Boyd, R. & Henrich, J. (2010) Gene-culture coevolution in the age of genomics. *Proceedings of the National Academy of Sciences USA* 107:8985–92. [rMLA]

Ridderinkhof, K. R., Ullsperger, M., Crone, E. A. & Nieuwenhuis, S. (2004) The role of the medial frontal cortex in cognitive control. *Science* 306:443–47. [DA]

Rips, L. J., Bloomfield, A. & Asmuth, J. (2008) From numerical concepts to concepts of number. *Behavioral and Brain Sciences* 31:623–87. [NM]

Ritter, F. E. & Young, R. M., eds. (2001) Using cognitive models to improve interface design. *International Journal of Human-Computer Studies* 55(1):1–107. (Special Issue.) [aMLA]

Rives, A. W. & Galitski, T. (2003) Modular organization of cellular networks. *Proceedings of the National Academy of Sciences USA* 100:1128–33. [rMLA]

Rizzolatti, G., Camarda, R., Fogassi, L., Gentilucci, M., Luppino, G. & Matelli, M. (1998) Functional organization of inferior area 6 in the macaque monkey. II. Area F5 and the control of distal movements. *Experimental Brain Research* 71:491–507. [MB]

Rizzolatti, G. & Craighero, L. (2004) The mirror-neuron system. *Annual Review of Neuroscience* 27:169–92. [JBR]

Rizzolatti, G., Fadiga, L., Gallese, V. & Fogassi, L. (1996) Premotor cortex and the recognition of motor actions. *Cognitive Brain Research* 3:131–41. [aMLA]

Robertson, R. M., Pearson, K. G. & Reichert, H. (1982) Flight interneurons in the locust and the origin of insect wings. *Science* 217:177–79. [JEN]

Roodenrys, S. & Miller, L. M. (2008) A constrained Rasch model of trace redintegration in serial recall. *Memory and Cognition* 36:578–87. [AS]

Roskies, A. L. (2007) Are neuroimages like photographs of the brain? *Philosophy of Science* 74:860–72. [rMLA]

Rossen, M., Klima, E. S., Bellugi, U., Bihrle, A. & Jones, W. (1996) Interaction between language and cognition: Evidence from Williams syndrome. In: *Language, learning, and behavior disorders: Developmental, biological, and clinical perspectives*, ed. J. H. Beitchman, N. Cohen, M. Konstantareas & R. Tannock, pp 367–92. Cambridge University Press. [TMD]

Roux, F. -E., Boetto, S., Sacko, O., Chollet, F. & Tremoulet, M. (2003) Writing, calculating, and finger recognition in the region of the angular gyrus: A cortical stimulation study of Gerstmann syndrome. *Journal of Neurosurgery* 99:716–27. [aMLA]

Rozin, P. (1976) The evolution of intelligence and access to the cognitive unconscious. In: *Progress in psychobiology and physiological psychology, vol. 6*, ed. J. A. Sprague & A. N. Epstein, pp. 245–80. Academic Press. [DA, PR]

Rozin, P. (1999) Preadaptation and the puzzles and properties of pleasure. In: *Well being: The foundations of hedonic psychology*, ed. D. Kahneman, E. Diener & N. Schwarz, pp. 109–33. Russell Sage. [PR]

Rozin, P. (2006) About 17 (+/ − 2) potential principles about links between the innate mind and culture: Preadaptation, predispositions, preferences, pathways and domains. In: *The innate mind, vol. 2: Culture and cognition*, ed. P. Carruthers, S. Laurence & S. Stich, pp. 39–60. Oxford University Press. [PR]

Rozin, P. (in press) Evolutionary and cultural psychology: Complementing each other in the study of culture and cultural evolution. In: *Evolution, culture, and the human mind*, ed. M. Schaller, A. Norenzayan, S. J. Heine, T. Yamagishi & T. Kameda. Psychology Press. [PR]

Rumelhart, D. E. & McClelland, J. L (1986) *Parallel distributed processing: Explorations in the microstructure of cognition*. MIT Press. [aMLA, ATD]

Rusconi, E., Walsh, V. & Butterworth, B. (2005) Dexterity with numbers: rTMS over left angular gyrus disrupts finger gnosis and number processing. *Neuropsychologia* 43:1609–24. [aMLA]

Rutishauser, R. & Moline, P. (2005) Evo-devo and the search for homology ("sameness") in biological systems. *Theory in Biosciences* 124:213–41. [DSM]

Ryle, G. (1949) *The concept of mind*. Hutchinson. [ATD]

Salvucci, D. D. (2005) A multitasking general executive for compound continuous tasks. *Cognitive Science* 29:457–92. [aMLA]

Samuels, R. (2006) Is the human mind massively modular? In: *Contemporary debates in cognitive science*, ed. R. J. Stainton, pp. 37–56. Blackwell. [AAP]

Sandler, W. & Lillo-Martin, D. (2006) *Sign languages and linguistic universals*. Cambridge University Press. [aMLA]

Sangha, S., Scheibenstock, A. & Lukowiak, K. (2003) Reconsolidation of a long-term memory in *Lymnaea* requires new protein and RNA synthesis and the soma of right pedal dorsal 1. *Journal of Neuroscience* 23:8034–40. [JEN]

Sapir, A., Hayes, A., Henik, A., Danziger, S. & Rafal, R. (2004) Parietal lobe lesions disrupt saccadic remapping of inhibitory location tagging. *Journal of Cognitive Neuroscience* 16:503–509. [DA]

Sapir, A., Soroker, N., Berger, A. & Henik, A. (1999) Inhibition of return in spatial attention: Direct evidence for collicular generation. *Nature Neuroscience* 2:1053–54. [DA]

Scher, S. J. (2004) A lego model of the modularity of the mind. *Journal of Cultural and Evolutionary Psychology* 2(21):248–59. [aMLA]

Schlosser, G. & Wagner, G. P., eds. (2004) *Modularity in development and evolution*. University of Chicago Press. [rMLA]

Schubotz, R. I. & Fiebach, C. J. (2006) Integrative models of Broca's area and the ventral premotor cortex. *Cortex* 42:461–63. [VB]

Schultz, W., Dayan, P. & Montague, P. R. (1997) A neural substrate of prediction and reward. *Science* 275:1593–99. [PSK]

Sharma, J., Angelucci, A. & Sur, M. (2000) Induction of visual orientation modules in auditory cortex. *Nature* 404:841–47. [ATD]

Simmons, W. K., Ramjee, V., Beauchamp, M. S., McRae, K., Martin, A. & Barsalou, L. W. (2007) A common neural substrate for perceiving and knowing about color. *Neuropsychologia* 45(12): 2802–10. [aMLA]

Simon, H. A. (1962/1969) The architecture of complexity. *Proceedings of the American Philosophical Association* 106:467–82. Reprinted in: H. Simon,

The sciences of the artificial, 1st edition, pp. 192–229. MIT Press, 1969. [aMLA]

Simon, H. A. (1962/1982) The architecture of complexity: Hierarchical systems. Reprinted in: H. Simon, *The sciences of the artificial*, 2nd edition, pp. 183–216. MIT Press, 1982. [AG]

Southgate, V. & Hamilton A. F. (2008) Unbroken mirrors: Challenging a theory of autism. *Trends in Cognitive Sciences* 12:225–29. [TMD]

Spearman, C. (1904) "General intelligence" objectively determined and measured. *American Journal of Psychology* 15:201–93. [CDR]

Sperber, D. (1996) *Explaining culture*. Blackwell. [aMLA, JAJ]

Sperber, D. (2001) In defense of massive modularity. In *Language, brain, and cognitive development: Essays in honor of Jacques Mehler*. MIT Press. [JAJ]

Spiers, H. J. & Maguire, E. A. (2007) Decoding human brain activity during real-world experiences. *Trends in Cognitive Sciences* 11(8):356–65. [AG]

Spirin, V. & Mirny, L. A. (2003) Protein complexes and functional modules in molecular networks. *Proceedings of the National Academy of Sciences USA* 100:12123–28. [rMLA]

Sporns, O., Chialvo, D. R., Kaiser, M. & Hilgetag, C. C. (2004) Organization, development and function of complex brain networks. *Trends in Cognitive Sciences* 8:418–25. [aMLA, AG]

Sporns, O., Tononi, G. & Edelman, G. M. (2000) Theoretical neuroanatomy: Relating anatomical and functional connectivity in graphs and cortical connection matrices. *Cerebral Cortex* 10:127–41. [aMLA]

Sternberg, S. (1969) The discovery of processing stages: Extensions of Donders' method. *Acta Psychologica* 30: 276–315. [aMLA]

Stewart, T. C. & West, R. L. (2007) Cognitive redeployment in ACT-R: Salience, vision, and memory. Paper presented at the 8th International Conference on Cognitive Modelling, Ann Arbor, MI, July 26–29, 2007. [aMLA]

Striedter, G. F. (2005) *Principles of brain evolution*. Sinauer. [CK]

Stroop, J. (1935) Studies of interference in serial verbal reactions. *Journal of Experimental Psychology* 18:643–62. [BB]

Studdert-Kennedy, M. (2005) How did language go discrete? In: *Language origins: Perspectives on evolution*, ed. M. Tallerman, pp. 48–67. Oxford University Press. [BLin]

Sun, F. T., Miller, L. M. & D'Esposito, M. (2004) Measuring interregional functional connectivity using coherence and partial coherence analyses of fMRI data. *NeuroImage* 21:647–58. [rMLA]

Suomi, S. J. (2004) How gene-environment interactions shape biobehavioral development: Lessons from studies with rhesus monkeys. *Research in Human Development* 1:205–22. [rMLA]

Sur, M., Garraghty, P. E. & Roe, A. W. (1988) Experimentally induced visual projections into auditory thalamus and cortex. *Science* 242:1437–41. [rMLA, PSK]

Sutherland, J. G. (1992) The holographic neural method. In: *Fuzzy, holographic, and parallel intelligence*, ed. B. Soucek & the IRIS Group, pp. 7–92. Wiley. [AS]

Svoboda, E., McKinnon, M. C. & Levine, B. (2006) The functional neuroanatomy of autobiographical memory: A meta-analysis. *Neuropsychologia* 44(12):2189–208. [aMLA]

Talairach, J. & Tournaux, P. (1988) *Co-planar stereotaxic atlas of the human brain*. Thieme. [aMLA]

Tamburrini, G. & Trautteur, G. (2007) A note on discreteness and virtuality in analog computing. *Theoretical Computer Science* 371:106–14. [FD]

Tang, Y., Zhang, W., Chen, K., Feng, S., Ji, Y., Shen, J., Reiman, E. M. & Liu, Y. (2006) Arithmetic processing in the brain shaped by cultures. *Proceedings of the National Academy of Sciences USA* 103:10775–80. [rMLA, NM]

Tettamanti, M. & Weniger, D. (2006) Broca's area: A supramodal hierarchical processor? *Cortex* 42:491–94. [aMLA]

Thelen, E. & Smith, L. B. (1994) *A dynamic systems approach to the development of cognition and action.* MIT Press. [ATD]

Thoenissen, D., Zilles, K. & Toni, I. (2002) Differential involvement of parietal and precentral regions in movement preparation and motor intention. *Journal of Neuroscience* 22:9024–34. [aMLA]

Tomasello, M. (2003) *Constructing a language*. Harvard University Press. [DSM]

Tong, A. H. Y., Lesage, G., Bader, G. D., Ding, H., Xu, H., Xin, X., Young, J., Berriz, G. F., Brost, R. L., Chang, M., Chen, Y., Cheng, X., Chua, G., Friesen, H., Goldberg, D. S., Haynes, J., Humphries, C., He, G., Hussein, S., Ke, L., Krogan, N., Li, Z., Levinson, J. N., Lu, H., Ménard, P., Munyana, C., Parsons, A. B., Ryan, O., Tonikian, R., Roberts, T., Sdicu, A.-M., Shapiro, J., Sheikh, B., Suter, B., Wong, S. L., Zhang, L. V., Zhu, H., Burd, C. G., Munro, S., Sander, C., Rine, J., Greenblatt, J., Peter, M., Bretscher, A., Bell, G., Roth, F. P., Brown, G. W., Andrews, B., Bussey, H. & Boone, C. (2004) Global mapping of the yeast genetic interaction network. *Science* 303:808–13. [rMLA]

Tononi, G. & Cirelli, C. (2003) Sleep and synaptic homeostasis: A hypothesis. *Brain Research Bulletin* 62:143–50. [WF]

Tononi, G. & Cirelli, C. (2006) Sleep function and synaptic homeostasis. *Sleep Medicine Review* 10:49–62. [WF]

Tooby, J. & Cosmides, L. (1992) The psychological foundations of culture. In: *The adapted mind: Evolutionary psychology and the generation of culture*, ed. J.

Barkow, L. Cosmides & J. Tooby, pp. 19–136. Oxford University Press. [aMLA, AAP]

Toskos Dils, A. & Boroditsky, L. (forthcoming) A motion aftereffect from literal and metaphorical motion language: Individual differences. To appear in: *Proceedings of the 32nd Annual Conference of the Cognitive Science Society. Austin, TX*, ed. S. Ohlsson & R. Catrambone. Cognitive Science Society. [ATD]

Tucker, M. & Ellis, R. (1998) On the relations between seen objects and components of potential actions. *Journal of Experimental Psychology: Human Perception and Performance* 24:830–46. [ATD]

Tucker, M. A., Hirota, Y., Wamsley, E. J., Lau, H., Chaklader, A. & Fishbein, W. (2006) A daytime nap containing solely non-REM sleep enhances declarative but not procedural memory. *Neurobiology of Learning and Memory* 86(2):241–47. [WF]

Tully, T., Cambiazo, V. & Kruse, L. (1994) Memory through metamorphosis in normal and mutant *Drosophila. Journal of Neuroscience* 14:68–74. [JEN]

Turkeltaub, P. E., Eden, G. F., Jones, K. M. & Zeffiro, T. A. (2002) Meta-analysis of the functional neuroanatomy of single-word reading: Method and validation. *NeuroImage* 16:765–80. [aMLA]

Uttal, W. R. (2001) *The new phrenology: The limits of localizing cognitive processes in the brain*. MIT Press. [AAP, ATD]

Van Herwegen, J., Ansari, D., Xu, F. & Karmiloff-Smith, A. (2008) Small and large number processing in infants and toddlers with Williams syndrome. *Developmental Science* 11:637–43. [TMD]

Van Orden G. C., Pennington, B. F. & Stone, G. O. (2001) What do double dissociations really prove? *Cognitive Science* 25:111–72. [aMLA]

Varela, F. J., Thompson, E. & Rosch, E. (1991) *The embodied mind: Cognitive science and human experience.* MIT Press. [ATD]

Vilarroya, O. (2001) From functional "mess" to bounded functionality. *Minds and Machines* 11:239–56. [OV]

Vilarroya, O. (2002) "Two" many optimalities. *Biology and Philosophy* 17(2):251–70. [OV]

Viswanathan, A. & Freeman, R. D. (2007) Neurometabolic coupling in cerebral cortex reflects synaptic more than spiking activity. *Nature Neuroscience* 10(10):1308–12. [CK]

von Dassow, G. & Munro, E. (1999) Modularity in animal development and evolution: Elements of a conceptual framework for EvoDevo. *Journal of Experimental Zoology, Part B: Molecular and Developmental Evolution* 285(4):307–25. [rMLA]

von Melchner, L., Pallas, L. L. & Sur, M. (2000) Visual behavior mediated by retinal projections directed to the auditory pathway. *Nature* (*London*) 404:871–76. [rMLA, TMD, PSK, ATD]

Wagner, U., Gais, S., Haider, H., Verleger, R. & Born, J. (2004) Sleep inspires insight. *Nature* 427(6972):352–55. [WF]

Weimann, J. M. & Marder, E. (1994) Switching neurons are integral members of multiple oscillatory networks. *Current Biology* 4:896–902. [JEN]

Weiskopf, D. (2007) Concept empiricism and the vehicles of thought. *The Journal of Consciousness Studies* 14:156–83. [aMLA, BB]

Wen, Q. & Chklovskii, D. B. (2008) A cost–benefit analysis of neuronal morphology. *Journal of Neurophysiology* 99:2320–28. [aMLA]

Wernicke, C. (1874) *Der aphasische Symptomenkomplex*. Weigert. [BB]

Wess, O. J. (2008) A neural model for chronic pain relief by extracorporeal shockwave treatment. *Urological Research* 36:327–34. [AS]

Westin, A. D. & Hood, L. (2004) Systems biology, proteomics, and the future of health care: Toward predictive, preventative, and personalized medicine. *Journal of Proteome Research* 3(2):179–96. [rMLA]

Westlake, P. R. (1970) The possibilities of neural holographic processes within the brain. *Biological Cybernetics* 7(4):129–53. [AS]

Wiese, H. (2003) *Numbers, language, and the human mind*. Cambridge University Press. [NM]

Williams, L. E. & Bargh, J. A. (2008a) Experiencing physical warmth promotes interpersonal warmth. *Science* 322:606–607. [JAB]

Williams, L. E. & Bargh, J. A. (2008b) Keeping one's distance: The influence of spatial distance cues on affect and evaluation. *Psychological Science* 19:302–308. [JAB]

Williams, L. E., Bargh, J. A., Nocera, C. C. & Gray, J. R. (2009a) The unconscious regulation of emotion: Nonconscious reappraisal goals modulate emotional reactivity. *Emotion* 9:847–54 [JAB]

Williams, L. E., Huang, J. Y. & Bargh, J. A. (2009b) The scaffolded mind: Higher mental processes are grounded in early experience of the physical world. *European Journal of Social Psychology* 39:1257–67. [JAB]

Wilshaw, D. J., Buneman, O. P. & Longuet-Higgins, H. C. (1969) Non-holographic associative memory. *Nature* 222:960–62. [AS]

Wilson, M. (2001) The case for sensorimotor coding in working memory. *Psychonomic Bulletin and Review* 8:44–57. [aMLA]

Wilson, M. (2002) Six views of embodied cognition. *Psychonomic Bulletin and Review* 9(4):625–36. [aMLA]

Wilson, T. D. & Brekke, N. (1994) Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychological Bulletin* 116:117–42. [JAB]

Wolpert, D. M., Doya, K. & Kawato, M. (2003) A unifying computational framework for motor control and social interaction. *Philosophical Transactions of the Royal Society B: Biological Sciences* 358:593–602. [aMLA]

Wright, W. G., Kirschman, D., Rozen, D. & Maynard, B. (1996) Phylogenetic analysis of learning-related neuromodulation in molluscan mechanosensory neurons. *Evolution* 50:2248–63. [PSK]

Yamada, T. & Bork, P. (2009) Evolution of biomolecular networks: Lessons from metabolic and protein interactions. *Nature Reviews Molecular Cell Biology* 10:791–803. [MR]

Young, L. J. (1999) Oxytocin and vasopressin receptors and species-typical social behaviors. *Hormones and Behavior* 36:212–21. [PSK]

Yuste, R., MacLean, J. N., Smith, J. & Lansner, A. (2005) The cortex as a central pattern generator. *Nature Reviews Neuroscience* 6:477–83. [PSK]

Zago, L., Pesenti, M., Mellet, E., Crivello, F., Mazoyer, B. & Tzourio-Mazoyer, N. (2001) Neural correlates of simple and complex mental calculation. *NeuroImage* 13:314–27. [arMLA, NM]

Zhong, C. B. & Leonardelli, G. J. (2008) Cold and lonely: Does social exclusion literally feel cold? *Psychological Science* 19:838–42. [JAB]