

Title: Evolution of cognitive function via redeployment of brain areas

Running head: Evolution of cognitive function

Michael L. Anderson
Department of Psychology
Franklin & Marshall College
P.O. Box 3003
Lancaster, PA 17604-3003

Institute for Advanced Computer Studies
Neuroscience and Cognitive Science Program
University of Maryland
College Park, MD 20742

michael.anderson@fandm.edu

Word count (incl. abstract, references and captions): 5197

Keywords: Adaptation, Physiological; Brain; Cortex; Cognition; Evolution

Abstract

The creative re-use of existing neural components may have played a significant role in the evolutionary development of cognition. There are obvious evolutionary advantages to such redeployment, and the data presented here confirm three important empirical predictions of this account of the development of cognition: (1) a typical brain area will be utilized by many cognitive functions in diverse task categories, (2) evolutionarily older brain areas will be deployed in more cognitive functions and (3) more recent cognitive functions will utilize more, and more widely scattered brain areas. These findings have implications not just for our understanding of the evolutionary origins of cognitive function, but also for the practice of both clinical and experimental neuroscience.

Evolution and redeployment

Part of understanding the functional organization of the brain is understanding how it evolved. The current study suggests that while the brain may have originally emerged as an organ with functionally dedicated regions, the creative re-use of these regions has played a significant role in its evolutionary development. This would parallel the evolution of other capabilities wherein existing structures, evolved for other purposes, are re-used and built upon in the course of continuing evolutionary development (“exaptation”: Gould and Vrba 1982). There is psychological support for exaptation in cognition (Cosmides 1989; Cruse 2003; Glenberg and Kashak 2002; Gould 1991; Lakoff and Nuñez 2000; Riegler 2001; Wilson 2001), and neuroanatomic evidence that the brain evolved by preserving, extending, and combining existing network components, rather than by generating complex structures *de novo* (Sporns and Kötter 2004). However, there has been little evidence that integrates these two perspectives, bringing such an account of the evolution of cognitive function into the realm of cognitive neuroscience (although see, e.g., Barsalou 1999).

One recent hypothesis along these lines—that combines a story about the evolution of the brain based on the re-use and extension of existing elements with an exaptive account of cognitive functions—is the massive redeployment hypothesis (Anderson 2006; Anderson in press). The massive redeployment hypothesis proposes that cognitive evolution proceeds in a way analogous to component reuse in software engineering (Heinemann and Council 2001), whereby existing components—originally developed to serve some specific purpose—are used for new purposes and combined to support new capacities,

without disrupting their participation in existing programs. If cognitive functions evolved in this way, then we should be able to make some specific empirical predictions regarding the resulting functional topography of the brain; here I discuss three.

First and most generally, we should expect a typical brain region to support numerous cognitive functions in diverse task categories. If this were not the case, if a typical brain region in fact served a very limited set of cognitive functions, then this would suggest instead that the brain evolved by generating new, dedicated regions for each new purpose.

Second, there should be a correlation between the phylogenetic age of a brain area and the frequency with which it is deployed in various cognitive functions. The longer an area has been around the more likely it will have proved useful to some evolving cognitive capacity, and be incorporated into the functional network of brain regions supporting the new task. Naturally this will not be true for every brain region, since a given area may have evolved to serve a very particular purpose of little use in later developments. But it should be generally the case that the older an area is, the more cognitive functions it supports.

Third and finally, there should be a correlation between the phylogenetic age of a cognitive function and its degree of localization. That is, more recent functions should generally use more, and more widely scattered brain areas than evolutionarily older functions. Again, the reasoning is simple: the more established neural components there

are when a given cognitive capacity is evolving, the more likely that one of them will already serve some purpose useful for the emerging capacity, and there is little reason to suppose that the most useful areas will be grouped together (and less reason to suppose this as evolutionary time passes, making available more functions supported by more areas).

Approach and methods

To evaluate the predictions made by the redeployment hypothesis, I performed some statistical analyses of 135 brain-imaging experiments, collected by Cabeza and Nyberg (2000). They survey 275 fMRI and PET experiments, in ten task categories. Here I focus on only four categories: attention, perception, imagery, and language. The 39 attention tasks included things like tone detection and Stroop tasks (naming colored words); the 42 perception tasks included such things as object identification and facial recognition; the 18 imagery tasks included mental rotation and landmark visualization; and the 36 language tasks included reading out loud and silently, lexical decision tasks (discriminating words from non-words), and the like. I chose these four categories because they are the most dissimilar of the ten categories covered. Since the other six categories all involve various types of memory—episodic memory retrieval, episodic memory encoding, working memory, etc.—it would presumably be much less significant to find evidence of redeployment across these categories.

For each task, Cabeza and Nyberg cataloged the brain areas reported to be activated by that task from a list including 26 numbered Brodmann areas, plus the insula and MT, and three subcortical areas—basal ganglia, thalamus and cerebellum—for each hemisphere.

Each area was divided into a lateral and medial segment, for a total of 124 brain regions. Note that the activations reported by Cabeza and Nyberg do not represent the full network of brain areas activated by a given cognitive task, but only those remaining after the relevant control/comparison tasks have been subtracted out. That is, the areas identified in the studies are understood to be those specifically responsible for the cognitive function under investigation.

Before moving on to the analysis of this data, it is worth pausing to acknowledge two related concerns that might be raised about this particular data set. First, the use of Brodmann areas as the basis for the spatial subdivision of the brain is both anatomically and functionally questionable, and second, the subtractive method has been supplanted in many imaging labs by more advanced analytical methods.

It is of course legitimate to wonder whether the sub-division of the cortex into Brodmann areas will be a feature of our final functional map of the human brain; one rather suspects it will be fully superseded by some yet-to-be developed topographical scheme. Yet, for our purposes—investigating the activations, reactivations and “scatter” of brain areas involved in supporting different functions—virtually any consistent spatial division of the brain will do, and regions approximately half the size of Brodmann’s areas offer adequate spatial resolution for the required analysis.

Likewise, there are indeed other analytical methods that can be used to identify the brain regions supporting cognitive functions, such as principal component analysis, functional

connectivity analysis, coherence analysis, structural equation modeling (a form of multiple regression analysis), dynamic causal modeling, and the like (see, e.g., Harrison and others 2003; Penny and others 2004; Sun and others 2004; Arfanakis and others 2000). And, as with any technique, the subtractive method has its specific limits and dangers (for instance, it may tend to undercount the regions involved in given functions; for discussion, see Lloyd 2000a, 2000b; Anderson in press). However, for the predictions under investigation here, using brain activation data as revealed by subtractive analysis has some specific advantages. First, it is the most widely used, and best understood, of the common analytical techniques. Moreover, insofar as the alternate methodologies mentioned above tend to emphasize the full, distributed network of brain areas activated in a given task, then one would expect to see more overlap in the networks activated by different tasks as compared with those networks revealed using subtractive analysis. That is, if using data produced by the subtractive method skews matters in any direction, it would be toward a more conservative, localizationist interpretation, rather than toward redeployment. Thus, if the massive redeployment hypothesis appears supported even when using a more conservative analytical method, this ought to count in its favor. Note this is not the same as saying that, if an analysis of imaging data based on the subtractive method supports the massive redeployment hypothesis, then it follows that an analysis based on one or another of these different methodologies would necessarily support the same predictions (and perhaps to a greater degree); this would itself be a subject for empirical investigation. Rather, the point is that insofar as the subtractive method is widely used, well understood, and does not in any

obvious way bias the results in favor of the hypothesis under consideration, then Cabeza and Nyberg's review seems a reasonable choice of data set for an initial investigation.

This returns us to the analysis at hand. In order to evaluate the three predictions made by the massive redeployment hypothesis, several values need to be calculated for this data set. First, we need to know how many brain areas are activated by a typical cognitive task, and whether this varies by task category. Second, we need to know how many cognitive tasks a typical brain area supports, and how these tasks are distributed across the four categories. Third, we need to measure the "scatter" of the areas participating in each task, and the variance of this value by task category. Finally, we need to correlate these values with phylogenetic age.

To calculate the first two values was primarily a matter of counting. Cabeza and Nyberg use a coding scheme for activations that forces a decision between lateral and medial activation, such that it is not possible to show both a medial and a lateral activation in a given brain area for a given task. Instead, the possible activations for each brain area are left lateral (LL), right lateral (RL), bilateral lateral (BL); left medial (LM), right medial (RM), bilateral medial (BM). Thus, for instance, they list the following activations for a task involving hearing words vs. a resting condition (Muller 1997): an LL activation in Brodmann area 47, and BL activations in areas 21 and 22. For the purposes of counting areas activated by a task, I treated bilateral activations of an area as two participants, one left and one right (medial or lateral). Thus, the language task above would have five participants, three LL participants (areas 47, 21 and 22) and two RL participants (areas 21 and 22). For the purposes of counting redeployments (areas activated by more than one

task), I matched LL activations in an area to other LL activations of that area, as well as to BL activations, and I matched RL activations in an area to other RL activations of that area, as well as to BL activations. I followed the same procedure for medial activations. I did not match bilateral activations to each other.

To calculate the diversity of activations across task categories I employed a standard measure of population diversity, *Diversity Variability* (DV). DV is calculated using the following equation, a version of standard deviation, where Cat_i is the proportion of activations in category i ; $mean$ is the mean proportion of activations in each category (always 0.25 for 4 categories), and k is the number of categories:

$$DV = \sqrt{\frac{\sum_{i=1}^k (Cat_i - mean)^2}{k}}$$

The category diversity of a given area is just (1-DV). With four categories, category diversity ranges from 0.57 (all items in one category) to 1 (equal numbers in each category). Note that for the purpose of calculating category diversity, the activation counts in each category were normalized to $n=42$.

Finally, to measure the distribution, or “scatter” of areas activated by a given task, I constructed an adjacency graph for the cortex (Figure 3). A graph is a set of objects called points, vertices or nodes connected by links called lines or edges. For constructing a graph of the cortex, I took the nodes to be numbered Brodmann areas (Brodmann 1907) and the edges to indicate adjacency. Adjacency in this context means only that the Brodmann areas share a physical border in the brain.

Graph theory (Diestel 2005) is a branch of mathematics that allows one to explore the topological properties of graphs. Graph-theoretic analyses have been used in neuroscience for such purposes as investigating neural connectivity patterns (Sporns 2002), correcting brain images (Han and others 2002), and analyzing the patterns of neural activations in epilepsy (Suharitdamrong and others 2006).

One of the simplest concepts in graph theory is *minimum graph distance*, which is just the fewest number of edges one must traverse to get from one node to another. Nodes that are adjacent in a graph have a graph distance of 1, nodes not adjacent to each other, but both adjacent to a third have a graph distance of 2, and so on. The minimum graph distance between every pair of nodes in the graph of the cortex was calculated using Dijkstra's algorithm (Dijkstra 1959).

A simple extension of minimum graph distance is *average minimum graph distance*, which is the average of the minimum graph distances between every pair of nodes in some subset of nodes in a graph. Figure 1 illustrates some different graphs, and the average minimum graph distances between all the nodes in the graph.

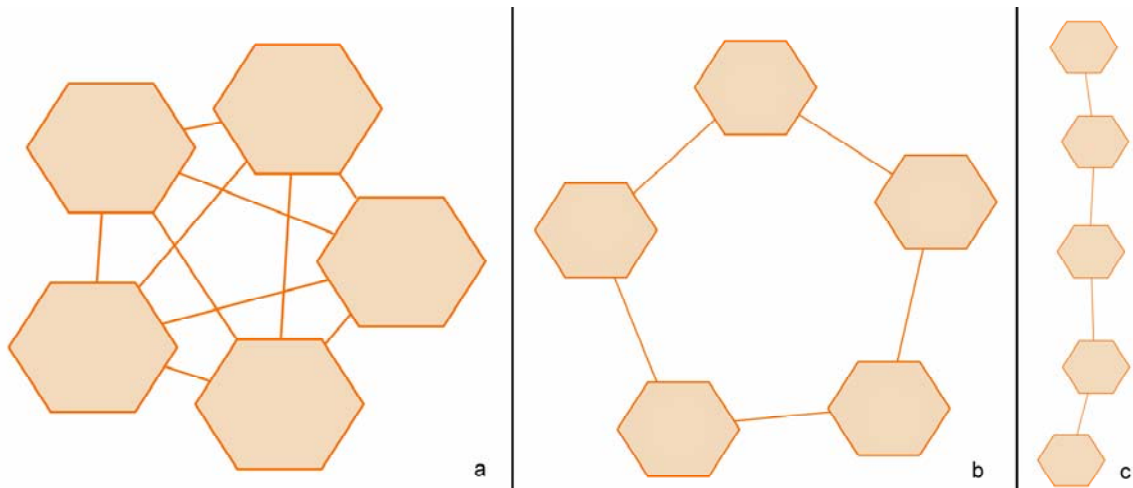


Figure 1: Illustrations of the average minimum graph distance (AMGD). The figure shows some simple graphs, where lines between nodes indicate adjacency. The graph in panel a has $AMGD = 1$, the graph in panel b has $AMGD = 1.5$ and the graph in panel c has $AMGD = 2$. Graphs rendered with aiSee v. 2.2.

Results

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

The activations were not limited to closely related tasks. Of the 86 regions activated in some task, 57 (66.3%) had activations in at least three categories; 28 of these had activations in all four categories. Only 15 regions (17.4%) had activations in just one category. Counting the number of tasks by category that activated each region, and normalizing the count of tasks in each category to $n=42$, shows that an average of 37.8% (SD 21.5) of activations are in categories other than the category with the highest number of activations.

Using the measure of population diversity among categories discussed above shows that the 86 brain regions have a mean category diversity of 0.76 (SD 0.11); ignoring the medial/lateral division gives 0.81 (SD 0.09). As shown in Table 1, an average category diversity of 0.81 suggests a fairly high degree of redeployment throughout the brain.

Area	Normalized proportion of activations by category				Category diversity
	Attention	Imagery	Language	Perception	(1 – DV)
BA46R	0.55	0.24	0.00	0.21	0.80
BA18L	0.26	0.21	0.28	0.24	0.97
BA38L	0.00	0.00	1.00	0.00	0.57

Table 1: Illustrations of the category diversity of some Brodmann areas. The table shows some examples of the diversity of activations across categories for three Brodmann areas.

These results are perhaps even more striking when put in graphical form. Figure 2 represents activations of Brodmann areas in the left hemisphere, in each of the four task domains, by both color and intensity. In panels a-d, the solid color represents the task domain—cyan for language, magenta for attention, yellow for perception and black (grey) for imagery—and the intensity of the color is correlated with the raw number of tasks in the category that activated the area. By using the colors of standard CMYK 4-color printing methods, this data can also be represented by taking the normalized proportion of activations in a Brodmann area for each cognitive domain, and assigning to that area a color composed of cyan, magenta, yellow and black in those proportions. The result is shown in panel e. Where an area has a color close to one of the four “basic” colors, that indicates a significant preference for tasks in the corresponding domain; other

colors reflect the contributions of two or more domains. Thus, for instance, a periwinkle blue indicates a mixture of language (cyan) and attention (magenta), while green suggests a combination of language (cyan) and perception (yellow). Darkness points toward a significant dose of imagery (black), as might be seen in a mixture of attention and imagery (mauve), perception and imagery (olive) or language, attention and imagery (dark purple); a middle gray indicates equal contributions by all task categories. Note the color for Brodmann area 18, which as we saw in Table 1 shows almost perfect category diversity.

Although it is hardly to be expected that anyone, save perhaps the exquisitely color-sensitive, could easily infer the exact mixture of cognitive domains that a given area supports just by looking at its shade, it should nevertheless be clear from the range and mixture of colors in panel e, and the wide distribution of color in panels a-d, that each Brodmann area is a highly diverse instrument. Far from supporting the standard notion that cognitive functions are generally localizable by domain (one specific example of which is illustrated in panel f, compiled by, and reprinted by permission of, Prof. Mark Wm. Dubin, MCD Biology, University of Colorado-Boulder), the data suggest a much more complex and subtle structure, where activity in many (most) brain areas supports multiple tasks in multiple cognitive domains. Rather than large areas of cortex close to the basic four colors, what we see instead is a large array of unique colors, indicating the relative contributions of each Brodmann area to supporting tasks in a given cognitive domain.




Figure 2: Color-coded activations of left cortex. The figure illustrates the activations of Brodmann areas in the left hemisphere according to color and intensity, where color represents the cognitive domain and intensity the raw number of tasks in the domain activating the area. In this figure, cyan represents language, magenta represents attention, yellow represents perception and black represents imagery. Overlaying the single-color images (panels a-d) gives the 4-color image in the bottom center (panel e). This image contrasts sharply with the common notion that cognitive domains are regionally localized, illustrated by panel f (panel f image compiled by, and reprinted by permission of, Prof. Mark Wm. Dubin, MCD Biology, University of Colorado-Boulder).

However, this does not mean that the cortex is in any way *randomly* or *holistically* organized; far from it. In fact, as is illustrated below, we can make (and support) some

specific predictions about the relations between cognitive functions and brain areas based on the phylogenetic age of the function and the brain area.

But first we need to present the data on the “scatter” of brain areas supporting various cognitive functions. The average minimum graph distance between the Brodmann areas activated by each of the 135 tasks is 3.89 (SD 2.00). Broken down by task category, we get attention 3.13 (SD 2.06), perception 3.71 (SD 1.98), imagery 3.97 (SD 1.75), and language 4.82 (SD 1.76). Figure 3 represents the cortex as an adjacency graph, with an attention task (Corbetta and others 1993) superimposed.




Figure 3: Cortex represented as an adjacency graph. The cortex represented as an adjacency graph, showing the Brodmann areas as nodes, with lines between adjacent areas. The yellow nodes are those activated by an attention task reported by Corbetta (1993). That task activated left Brodmann areas 7, 8 and 24, and right Brodmann areas 7 and 32; average minimum graph distance is 4.0, close to the average for all tasks. Graph rendered with aiSee v. 2.2.

With this basic data in front of us, we can look at correlations between these values and phylogenetic age. As noted above, if the evolution of cognition proceeded via the extensive re-use of existing components, then evolutionarily more recent cognitive functions should activate more, and more widely scattered brain areas. Comparing language tasks with perception tasks and attention tasks gives the predicted result. For the mean number of areas activated, language is greater than perception by 2.93 (2-sample Student's-t test, double-sided $p = 0.0165$) and greater than attention by 2.55 ($p =$

0.0475). For average minimum graph distance, language is greater than perception by 1.11 ($p = 0.0121$) and greater than attention by 1.69 ($p = 0.0003$). Differences between other categories are not significant (Table 2).

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

Table 2: Comparisons of cognitive domains. Results for all category comparisons on average number of brain regions activated per task and average minimum graph distance between the activated regions. Note that only the differences between language and perception and language and attention are significant.

The last important prediction of the redeployment hypothesis to be discussed here is that evolutionarily older brain areas should be deployed in more cognitive functions. Figure 4 gives the results of plotting the number of tasks that activate a given Brodmann area versus the Y-coordinate of the area, based on the simplifying assumption that areas in the rear of the cerebral cortex (occipital lobe) are evolutionarily older than those in the front (pre-frontal cortex), *ceteris paribus*. Although the data are highly variable, as expected, there is nevertheless a significant linear correlation.

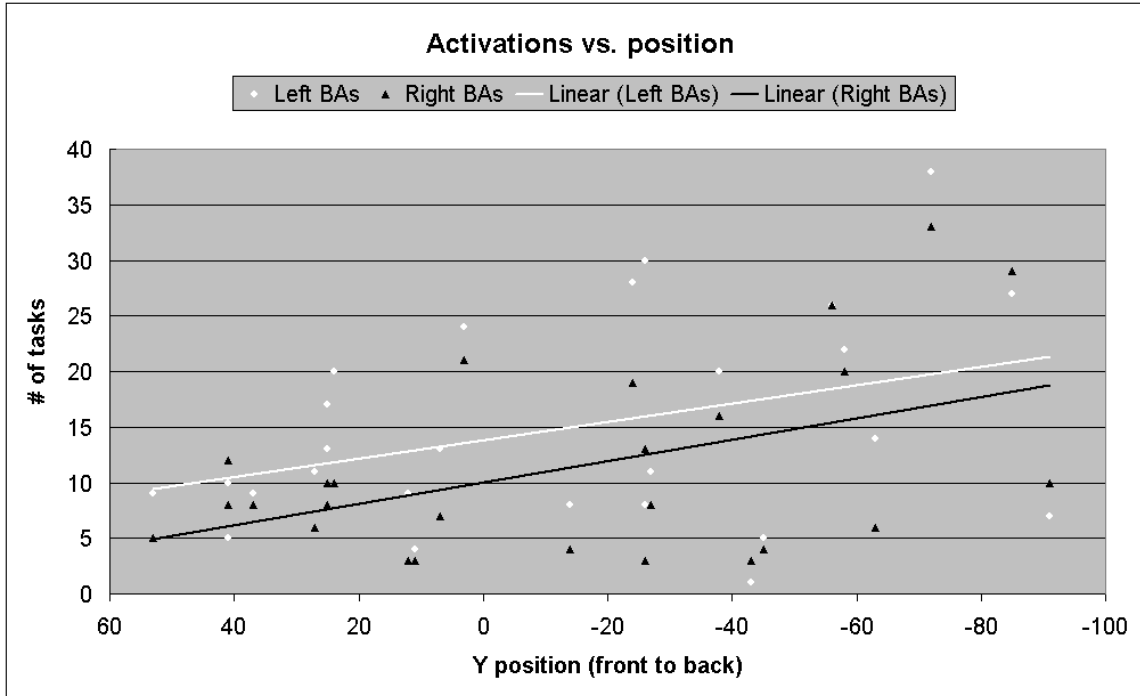


Figure 4: Graph of activations vs. Y-position of activated region. A plot of the number of tasks (out of 135) that activated each Brodmann area vs. the Y coordinate of the area (calculated in Talairach (Talairach and Tournoux, 1988) space using the Brede (Nielsen, 2003) database). The data shows a linear correlation, $R = -0.4121$, $p \leq 0.00244$ ($t = -3.198$, $DF = 50$).

It must of course be admitted that the simplifying assumption used here is relatively crude; for instance, Brodmann areas 4 and 6, associated with motor control, seem likely to be quite old (and yet they are closer to the front of the brain), whereas Brodmann areas 39 and 40, since they do not appear in the macaque cortex, might be much younger (and yet they are near the back of the brain). Further, in interpreting this data in particular, we need to remember that the chance of redeployment for a given neural component is a function of at least two different factors: first, the age of the area, and second the general usefulness of the function(s) it computes. That is, we should not expect (and it is not here claimed) that evolutionary age correlates perfectly with Y-position, nor that the degree of redeployment will correlate perfectly with evolutionary age. Yet, it seems reasonable to expect relatively high correlation in the first case, and this correlation can thus be used to

test the prediction of significant correlation in the second case. If, given these assumptions, we did *not* see significant correlation between redeployment and Y-position, that would certainly count against the hypothesis. Indeed, given the acknowledged limitations of the simplifying assumptions, the fact that there is such a high linear correlation between Y-position and degree of redeployment is all the more striking.

To do better than this would require a more reliable measure of the relative evolutionary age of various brain areas; but in the absence of any consensus on the matter the next best thing is to provide some raw data, allowing researchers to judge for themselves the degree of correlation between redeployment and evolutionary age. Such data is provided in Table 3, which lists each Brodmann area, and its total number of activations, in ascending order.

Area	Left	Right	Total
BA23	1	3	4
BA38	4	3	7
BA31	5	4	9
BA20	8	3	11
BA24	9	3	12
BA4	8	4	12
BA11	5	8	13
BA10	9	5	14
BA17	7	10	17
BA46	9	8	17
BA8	11	6	17
BA42	11	8	19
BA39	14	6	20
BA44	13	7	20
BA32	13	8	21
BA9	10	12	22
BA45	17	10	27
BA47	20	10	30

BA40	20	16	36
BA7	22	20	42
BA21	30	13	43
BA6	24	21	45
BA22	28	19	47
BA37	26	26	52
BA18	27	29	56
BA19	38	33	71

Table 3: Listing of Brodmann areas with total number of activations. Total number of tasks (out of 135) that activated each Brodmann area, listed in ascending order.

Discussion

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

Before concluding, I would like to say a few words about the more theoretical attractions and implications of the massive redeployment hypothesis. First of all, the hypothesis offers the potential for explaining both localization of function (cognitive functions only use limited and specific parts of the brain), and diversity of purpose (a typical brain area is activated by highly diverse cognitive tasks). This may help dissolve the debate between localization and holism (Uttal 2001), which in its typical form offers a false choice between equipotentiality (a given brain area can do many different things when it is activated) and strict localization (each brain area does one and only one thing).

According to the massive redeployment hypothesis, the fact that a brain area is dedicated to some highly specific low-level task is perfectly compatible with its being used to support many different cognitive functions (Anderson 2006, in press). In fact, if brain areas were multi- or equi-potential, and so could easily be recruited to compute substantially different functions, then it is hard to understand why older brain areas are more often recruited than younger ones, and why newer cognitive functions recruit more widely scattered brain areas. It would seem that such a pattern of redeployment would only arise if the low-level (computational) functions of brain areas were relatively fixed, such that developing a new cognitive function requires either developing new capacities *de novo*, or finding areas already performing some required role. If brain areas could instead be easily encouraged to compute many different functions, then considerations of information-processing efficiency would favor recruitment of *nearer* areas over areas already computing some desired function, but further away.

Second, the massive redeployment hypothesis may offer a clearer way of organizing the search for the neurological bases of cognitive function. In particular, it suggests that in order to determine the functional role of a given brain region it is necessary to consider its participation across multiple task categories, and not just focus on one, as has been the typical practice. Making this claim a bit more specific, when modeling a given cognitive function, and attempting to map that model onto brain areas, it will be necessary to consider not just the model of the function under primary consideration, but also the models of other functions recruiting the same brain areas, such that the sub-functional elements of each model attribute the same role to the brain areas where they overlap.

Finding the functional role of a given brain area will be something like finding the right letter to go into a box on a (multidimensional) crossword puzzle, determined not just by the answer to a single clue, but by all the clues whose answers cross that box. This makes the task both harder, because it is multiply constrained, but also easier, because it offers the possibility of leveraging information from several sources to make the attribution.

Third, and closely related to the last point above, as we come to recognize the diverse cognitive functions supported by given brain regions, this should suggest more fine-grained predictions about such matters as priming and cognitive interference, as well as the likely effects (and the localization) of brain injuries. The knowledge that a given brain area is used in multiple tasks and domains opens the possibility of designing experiments leveraging these overlaps, e.g. in cross-domain priming or interference studies, or in the development of cross-domain therapies for brain-injury patients.

Finally, looking at brain organization in this way may eventually offer a different method to assess the relative evolutionary age of cognitive functions, and of brain areas, opening another window on our evolutionary past.

References

Arfanakis K, Cordes D, Haughton VM, Moritz CH, Quigley MA, Meyerand ME. 2000. Combining independent component analysis and correlation analysis to probe interregional connectivity in fMRI task activation datasets. *Magnetic Resonance Imaging* 18:921–30.

Anderson ML. 2006. Evidence for massive redeployment of brain areas in cognitive function. *Proceedings of the Cognitive Science Society* 28:24–29.

- Anderson ML. In press. The massive redeployment hypothesis and the functional topography of the brain. *Philosophical Psychology*.
- Barsalou LW. 1999. Perceptual Symbol Systems. *Behavioral and Brain Sciences* 22:577–660.
- Brodmann K. 1907. Beiträge zur histologischen Lokalisation der Großhirnrinde. *J Psychological Neurology* 10:231–4.
- Cabeza R, Nyberg L. 2000. Imaging cognition II: An empirical review of 275 PET and fMRI studies. *J Cognitive Neuroscience* 12:1–47.
- Corbetta M, Miezin FM, Shulman GL, Petersen SE. 1993. A PET study of visuospatial attention. *J Neuroscience* 13:1202–26.
- Cosmides L. 1989. The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition* 31:187–276.
- Cruse H. 2003. The evolution of cognition—a hypothesis. *Cognitive Science* 27:135–155.
- Diestel R. 2005. *Graph Theory*, 3ed. Heidelberg, Springer-Verlag.
- Dijkstra EW. 1959. A note on two problems in connexion with graphs. *Numerische Mathematik* 1:269–71.
- Glenberg A, Kaschak M. 2002. Grounding language in action. *Psychonomic Bulletin and Review* 9:558–565.
- Gould SJ. 1991. Exaptation: A crucial tool for an evolutionary psychology. *J Social Issues* 3:43–65.
- Gould SJ, Vrba E. 1982. Exaptation: A missing term in the science of form. *Paleobiology* 8:4–15.
- Han X, Xu C, Braga-Neto U, Prince JL. 2002. Topology correction in brain cortex segmentation using a multiscale, graph-based approach. *IEEE Transactions on Medical Imaging* 21:109–21.
- Harrison L, Penny WD, Friston K. 2003. Multivariate autoregressive modeling of fMRI time series. *NeuroImage* 19:1477–1491.
- Heineman GT, Council WT. 2001. *Component-Based Software Engineering: Putting The Pieces Together*. New York, Addison-Wesley.

- Lakoff G, and Núñez R. 2000. *Where Mathematics Comes From*. New York, Basic Books.
- Lloyd D. 2000a. Terra cognita: From functional neuroimaging to the map of the mind. *Brain and Mind* 1:1–24.
- Lloyd D. 2000b. Virtual lesions in the not-so-modular brain, *J International Neuropsychological Society* 6:627–35.
- Muller RA, Rothmel RD, Behen ME, Muzik O, Mangner TJ, Chugani HT. 1997. Receptive and expressive language activations for sentences: a PET study. *Neuroreport* 8:3767–70.
- Nielsen FÅ. 2003. The Brede database: a small database for functional neuroimaging. *NeuroImage* 19:s2879–80.
- Penny WD, Stephan KE, Mechelli A, Friston KJ. 2004. Modelling functional integration: a comparison of structural equation and dynamic causal models. *NeuroImage* 23:s264–74.
- Riegler A. 2001. The cognitive ratchet: The ratchet effect as a fundamental principle in evolution and cognition. *Cybernetics and Systems* 32:411–27.
- Sporns O. 2002. Graph theory methods for the analysis of neural connectivity patterns. Kötter R. (ed.) *Neuroscience Databases: A Practical Guide*. Dordrecht, Kluwer.
- Sporns O, Kötter R. 2004. Motifs in brain networks. *PLoS Biol* 2:e369.
- Suharitdamrong W, Chaovalitwongse A, Pardalos PM. 2006. Graph theory-based data mining techniques to study similarity of epileptic brain network. *Proc. DIMACS Workshop on Data Mining, Systems Analysis, and Optimization in Neuroscience*.
- Sun FT, Miller LM, D'Esposito M. 2004. Measuring interregional functional connectivity using coherence and partial coherence analyses of fMRI data. *Neuroimage* 21:647–58.
- Talairach J, Tournaux P. 1988. *Co-planar stereotaxic atlas of the human brain*. New York, Thieme.
- Uttal WR. 2001. *The New Phrenology: The Limits of Localizing Cognitive Processes in the Brain*. Cambridge, MIT Press.
- Wilson M. 2001. The case for sensorimotor coding in working memory. *Psychonomic Bulletin and Review* 8:44–57.