What Do Brain Data Really Show?

Valerie Gray Hardcastle†‡
Virginia Tech

C. Matthew Stewart
University of Texas Medical Branch-Galveston

There is a bias in neuroscience toward localizing and modularizing brain functions. Single cell recording, imaging studies, and the study of neurological deficits all feed into the Gallian view that different brain areas do different things and the things being done are confined to particular processing streams. At the same time, there is a growing sentiment that brains probably don't work like that after all; it is better to conceive of them as fundamentally distributed units, multi-tasking at every level. This sentiment, however, is much less congenial to the tried-and-true experimental protocols available today and to theorizing about the brain in general. This essay examines the tension between current experimental methods and large-scale views of the brain. We argue that this disconnection between experiment and what really are guiding theoretical metaphors seriously impedes progress in neuroscience.

1. Introduction. There is a burgeoning cottage industry in philosophy dedicated to criticizing the modularity of mind hypotheses in evolutionary psychology (see e.g. Buller 1999; Davies 1996, 1999; Davies et al. 1995; Grantham and Nicols 1999; Griffiths 1997; Richardson 1996; Shaprio 1998, 1999; Shapiro and Epstein 1998; Sober 1997; Sterleny 1995). We know because we are part of it (cf., Buller and Hardcastle, 2001). But the criticisms mustered against evolutionary psychologists' wanton and unwarranted methodological assumptions of mind modules also find a home in cognitive neuroscience. Cognitive neuroscientists assume that they can

†For reprints, contact Valerie Gray Hardcastle, STS/CIS, Virginia Tech, Blacksburg, VA 24061–0227, USA, e-mail: valerie@vt.edu.

‡This paper benefited greatly from discussion with Eric Dietrich, audience members at the PSA symposium, and the McDonnell Neurophilosophy Research Fellows. Work on this project was supported in part by a grant from the McDonnell Foundation.

Philosophy of Science, 69 (September 2002) pp. S72–S82. 0031-8248/2002/69supp-0007$10.00
Copyright 2002 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/341769 Published online by Cambridge University Press

localize brain function; they seek discrete, physically constant brain "modules," a material analogue for the psychologists' set of distinct mental software packages. Indeed, the criticisms should be all the more pointed because, according to some, the brain is where we should be finding these alleged cognitive units. (Mother Nature does build psychologies out of brains, after all.) If anybody has localization data that would support modularity, it should be the neuroscientists.

But they don't have the data. Actually, it is worse than that, because—appearances to the contrary—they don't even have a good way of accessing the appropriate evidence. It is a *bias* in neuroscience to localize and modularize brain functions. Single cell recordings, the study of neurological deficits, and imaging studies all feed into the Gallian view that different brain areas do different things and the things being done are confined to particular processing streams. But these are just prejudices, nothing more. And the underlying assumptions could very well be wrong.

2. Localization and Single Cell Recordings. When scientists do single unit recordings from a set of neurons they assume that they are busy examining a discrete system. Why else record from this particular set, isolated in this particular fashion, in the first place? Many neuroscientists spend their entire careers delineating small functional areas in the brain, and they are wildly successful in their endeavors. They have identified at least 36 different topographical visual processing areas in cortex (DeYoe and Van Essen 1988); they have differentiated the "what" from the "where" object processing streams (Mishkin et al. 1983); and they have distinguished motion detection from contour calculations (Livingstone and Hubel 1988). Our maps of brain function are getting more and more complicated as we learn more and more about the processing capacities of individual cells. And all these projects are founded on the belief that our brain operates in this way, that we have discrete processing streams that feed into one another.

Yet the most neurons we have ever been able to record from simultaneously are around 150; the most cells we can ever see summed local field potential activity over are a few thousand. But brain areas have hundreds of thousands of neurons, several orders of magnitude more than we can access at any given time. And these neurons are of different types, with different response properties and different interconnections with other cells, including other similar neurons, neurons with significantly different response properties, and cells of other types completely (such as glia). Any conclusions we draw about the behavior of whatever cells we are recording from are going to be limited to very basic stimulus-response and correlation analyses of whatever neuronal subtype we are currently examining. Hence, the functionality we ascribe based on these relatively meager sorts

https://doi.org/10.1086/341769 Published online by Cambridge University Press

of experiments might be much more restricted than what the cells are actually doing.

We insert an electrode in or near a cell and then record what it does as we stimulate the animal in some fashion. We record from a cell in a vestibular nucleus and then move the animal's head about to see if that changes the activity of the neuron (cf., Newlands and Perachio 1990a, 1990b). If it does, then we move it some more or we move it differently and see how that changes the neuronal output. If it doesn't, then we either try another nearby cell or we try some other stimulus. But what we can't do is record from all the neurons in some isolated area, even if the area is very small. And what we can't do is test any given cell for all the known functional contributions of brain cells in general. So, what we conclude about any cell will only reflect the cells we or others have actually recorded from using stimuli we or others have actually used. This research strategy systematically underestimates when neurons actually respond and under what conditions.

This sort of unit study attempts to combine scores, hundreds, or even thousands of single-unit recordings together to try to analyze the population. Theoretically, we could, perhaps, in principle, delineate a nervous system region stereotaxically if it had reproducible correlations between afferent and efferent connections such that we could ultimately articulate the neurophysiological function of the defined region. However, the likelihood of success for this type of study decreases as the complexity of the organism increases. We can draw functional conclusions regarding the activities of neurons in the abdominal ganglia of *Aplysia,* or the segmental ganglia of the leech. But the architecture of these organisms' central nervous system is so different from mammals' that the probability of successfully using similar techniques is very low to zero.

In addition, the actual processing of information that goes on in those cells involves lots of different kinds of excitatory and inhibitory inputs from other areas in the brainstem, cerebellum, and cerebral cortex. Our dorsal horn is supposed to integrate afferent nociceptive information from the periphery and pass it onto the motor system (among other things), but it doesn't do that segregated from the rest of the brain and what the brain is trying to do. It is integrating and passing as we are trying to pursue prey or flee from an enemy. Moreover, the brain regions that perform these tasks are often connected to the very area we are recording from. The motor system feeds back down into the dorsal horn, as does the thalamus and significant parts of cortex (cf. Hardcastle 1999).

We sometimes wonder whether we can't play a variant of the parlor game Six Degrees of Kevin Bacon in the brain. (This is the game in which you can connect any actor or actress in a movie to a movie starring Kevin Bacon via common actors in no more than six moves. So, for example,

https://doi.org/10.1086/341769 Published online by Cambridge University Press

Elizabeth Taylor has a Bacon number of 2; she was in *Rhapsody* (1954) with Vittorio Gassman, who was in *Sleepers* (1996) with Kevin Bacon.) We can connect any neuron in the brain to any other neuron with just a few steps, probably fewer than six.

The impact on cognitive processing of such rampant feedback connections in the brain is only just now starting to be explored in neuroscientific research, though exactly how to do this is a difficult question to answer. Of course, neurophysiologists aren't fools; they design their experiments keeping in mind the known anatomic connections between and among the relevant structures. At the same time, any actual experimental observations of all the remote influences on the dorsal horn, for example, are impossible, despite however many individual neurons we record from. We simply don't have any way of conducting such extensive, invasive tests on live animals. At best, the particular influences assumed in any particular recording series are a matter of previously accepted gospel, dogma, and faith.

Let me give you an analogy for what we are talking about. We have all heard the old joke about the four men who examine an elephant blindfolded. The one who feels the legs believes he is before a tree; the one who gets the truck thinks he has a snake on his hands, and so forth. Experiments in neurophysiology using single-cell recordings are much like that, only each investigator examines a different creature in a different pen on different days, and they only get to look at one square inch per examination period. They simply can't holler out to their companions, "See what happens at your end when I kick the tree trunk." Moreover, they have to read about each other's conjectures and descriptions in the local newspaper, which may or may not accurately report what the others think and do, if they report such activities at all. And not only do they have to identify what it is they are examining, but they also have to figure out what the function is for each part they are allowed to touch.

Discerning brain function when limited to touching single cells is no easy feat. No wonder neuroscientists assume functional localization. It makes their task at least conceivable. If they didn't assume some sort of functional specificity in discrete brain regions, then what would the purpose of single-cell recordings be at all? In order to get anything theoretically useful out of single-cell recordings, we have to have already accepted that the data are going to tell us about how a specific brain area functions.

3. Lesion Studies and the Assumption of Brain Constancy. Ideally, neuroscientists try to conjoin their single-cell studies with some sort of lesion experiment. Once scientists construct a general flowchart of the relevant structures based on anatomy experiments, and they have estimated normal unit behavior from a series of single cell studies, they then try to knockout

https://doi.org/10.1086/341769 Published online by Cambridge University Press

the hypothesized functions by placing lesions in otherwise normal animals. They run their experiments based on the assumption that these lesions, placed in regions known to be important, will change the unit behavior of the cells they are studying in a consistent fashion. If they witness such a change, they use that information to explain the relative functional contributions of the lesioned region to the cells under scrutiny. In other words, they are using lesion studies to try to derive a functional boxology for the brain, just as cognitive psychologists use reaction time distributions and error measurements to find one for the mind.

Right off the bat, though, we find technical difficulties. There are many different types of lesions, including electrolytic, incisional, excisional, and chemical. The optimal lesion spares the nerve tracts themselves while eliminating populations of nerve cells. (Chemical lesions are generally the first choice for this type of lesion.) But lesion studies are notorious for highly variable functional damage. Where the lesion actually occurs and how widespread it actually is—facts over which scientists have much less control than one might assume—leads to all sorts of differently ordered remaining populations. As a result, each lesion is unique, and the actual functional deficit after lesioning an animal is highly variable. A genuine replication of a lesion study in neuroscience is a practical impossibility.

But there is a larger theoretical concern. What neuroscientists know, but generally ignore, is that *any* functional change in the central nervous system will lead to compensatory changes elsewhere (see discussion in Hardcastle and Stewart, 2001). If we ablate the semi-circular canals in a rat's ear, then within hours the rat recovers static functioning in its vestibular system (Baarsma and Collewjin 1975; Schaefer and Meyer 1973; Sirkin et al. 1984). The vestibular nuclei of the brainstem recover a level of baseline activity, even though it is receiving *no* information from the periphery anymore. How does it do this? Frankly, we have no idea.

Here is another example closer to home for many of us. Temporomandibular joint syndrome (TMJ), a progressive degeneration of the joint that connects the lower jar to the skull, is often accompanied by tinnitus in the ipsilateral ear. Indeed, ringing in the ear is a diagnostic criterion for TMJ. But we don't know why having one's mandibular joint fall apart would provoke some change in the auditory system. In fact, we didn't even know that there was an auditory nerve that ran alongside the lower jaw until recently. It is hard to derive the functions of various brain areas if we don't even have a sketch of the correct neuroanatomy in the first place, and this task becomes even more difficult if the brain won't hold still for us.

Because it is highly plastic, poking holes in the brain in one place will provoke to it react in some fashion in some other place. Usually these other places are not components in the system or region being studied. But even if they are, neuroscientists ignore plasticity of the brain in favor

https://doi.org/10.1086/341769 Published online by Cambridge University Press

of assuming a consistent functional alteration as caused by the lesion and nothing more. How are investigators supposed to evaluate some observed functional change when the difference they see might have been evoked by the brain's attempt to compensate for its loss and not by any specific deficit induced by the lesion?

The short answer is that they can't if they are restricted to single-cell recordings and lesion studies. To answer this question we need to be able to see the activity of the entire brain at once and over time. But we can't do that.

4. Functional Imaging to the Rescue? We want to claim that neuroscientists' assignments of function to brain regions or areas are not warranted by the data. They are not warranted because their simplifying assumptions of localization of function and functional constancy are radically false. We admit these seem ludicrous conclusions, for neuroscientists give us functional hypotheses all the time with a perfectly straight face, about brain areas both large and small. The frontal lobes allow for planning complex behaviors (cf., Goldman-Rakic 1996; though see Carpenter et al. 2000); the lingual gyrus is sensitive to peripheral visual information (Emiliano et al. 2000); certain cells in superior colliculus are multi-modal and can process both visual and auditory inputs (cf., de Gelder 2000).

Moreover, and more importantly, the excited hoopla over fMRI and other imagining techniques concerns exactly this point: we do have a way of looking at the activity of the whole brain at one time as tied to some cognitive activity or other. At least, that is what the PR claims. But magnetic resonance imagining, the best non-invasive recording device we currently have, only has a spatial resolution of about 0.1 millimeter and each scan samples about five seconds of activity (cf., Churchland and Sejnowski 1988). This imprecision forecloses the possibility of directly connecting single cell activity—which operates three to four orders of magnitude smaller and faster—with larger brain activation patterns. What are we to do? The answer given all too often by neuroscientists is to fudge.

Methodological difficulties with current imagining techniques are now becoming well known (cf., Bechtel 2000; Cabeza and Nyberg 1997; Stufflebeam and Bechtel 1997). Let us just add briefly to this discussion. Here is how most functional imagining studies work. The experimenter picks two experimental conditions that she believes differ along only one dimension; they differ only with respect to the cognitive or perceptual process she wants to investigate. She then compares brain activity recorded under one condition with what happens in the second condition, looking for regions whose activity levels differ significantly across the two. These areas, she believes, comprise the neural substrates of the task under scrutiny.

Let us set aside the fact that this so-called subtraction method has no

https://doi.org/10.1086/341769 Published online by Cambridge University Press

way of determining whether the differences found are actually tied to the cognitive process and not to something else occurring concurrently but coincidentally (cf., Shulman 1996). Notice that how well the subtraction method will work depends upon the sensitivity of the measuring devices and that the *worse* the instrument is the *better* the method seems to be for localization studies. Low signal-to-noise ratios (SNR) means that we will find only a few statistically significant differences across conditions. And these are the sort of results neuroscientists need in order to bolster any claims identifying particular cognitive processes with discrete brain regions (Wandell 1999).

But as the imagining technology improves and the SNR increases, we see more and more sites that differ across trials. The more sites we get, the more it looks as though essentially the entire brain is involved in each cognitive computation. And the more it looks as though the entire brain is involved in each thought, the less it is we can justify any assumption of functional specificity in the brain.

We are but victims of imprecise instrumentation. If we extrapolate from what we might learn with more sensitive measures, we can easily see that there will come a time when this whole approach just won't work anymore. It simply ceases to be interestingly informative to learn that, as Brian Wandell notes, "the activity of one pixel is significant to one part in 10^{10} and the adjacent pixel is significant to one part in 10^{11}" (1999, 170). Put in the harshest terms, brain imaging seems to support localist assumptions because we aren't very good at it yet.

It also seems to support functional specificity because there have been few meta-analyses done that look at activity in single brain regions across a wide variety of experimental conditions (though see Cabeza and Nyberg 1997, 2000; Lloyd 2000). It doesn't take much work to see that there is something seriously wrong with trying to assign a particular function to each statistically significant area of subtracted activity. (For simplicity's sake, in the following discussion, we are ignoring the difficulties associated with summarizing the results from multiple studies, including the fact that different researchers use different standards of significance, different baseline conditions, and different statistical methods for culling their data. But even if we were to take these factors into consideration, we believe the conclusions we draw below would still stand.)

Brainmap is a wonderful database housed at the University of Texas Health Science Center in San Antonio (cf. http://ric.uthscsa.edu/projects/brainmap.html). It currently has archived 225 imaging articles, for a total of 771 experiments with 7,683 significant areas of difference being discussed. A quick search for Brodmann area 6 returns 36 articles that found area 6 significantly active in at least one experimental trial. This in and of itself might not be so interesting, but what is interesting is the wide variety

https://doi.org/10.1086/341769 Published online by Cambridge University Press

of cognitive tasks the area apparently underlies. Area 6 appears significantly active after subtraction in studies of phonetic speech processing, voluntary hand and arm movements, sight-reading music, spatial working memory, recognizing facial emotions, binocular disparity, sequence learning, idiopathic dystonia, pain, itch, delayed response alternation, and category-specific knowledge, to name but a few (Ceballos-Baumann et al. 1995; Colebatch et al. 1991; Derbyshire et al. 1994; George et al. 1993; Gold 1994; Gulyas and Roland 1994; Hsieh et al. 1994; Jonides et al. 1993; Martin et al. 1996; Rauch; Sergent et al. 1992a; Sergent et al. 1992b; Zatorre et al. 1992).

Keep in mind that area 6's actual activity is systematically underreported in published articles because it would only be mentioned in studies in which it was not differentially active during baseline scanning. That is, any study in which both the control condition and the test condition prodded area 6 to light up would subtract out this activity as uninteresting for the task at hand, even though that area might be absolutely crucial for carrying out the cognitive process. (Our quick scan of the database would also miss studies that discussed area 6 under a different rubric, such as premotor cortex.) Now, what should we say the function of area 6 is? It isn't at all clear anymore how the brain is using this structure or for what purpose.

It could be the case that we simply haven't identified the function of area 6 and if we keep on doing the sort of subtraction studies that we are currently, then eventually we will find a unifying and pithy way to describe what premotor cortex is doing for us. In this instance, neuroscience would be on the right track to determining brain function, but we still have a long ways to go yet. But, it could also be the case that how a region functions depends heavily on the "neural context" (McIntosh 1999, 2000). Its functional role in our cognitive economy depends on how it is connected to other areas and how those other areas are responding. (The function of these areas would also be dependent on their particular connectivity and the current patterns of activation. And so it would go.) If this is correct, then searching for "the" function of particular areas is misguided, for different brain regions play different roles depending upon the cognitive tasks at hand.

In addition, as we all know, the underlying scientific principle of things like fMRI and PET is that inferences can be made due to indirect observations of *metabolic* changes. Consequently, they completely overlook all the components of cognition that do not have a corresponding change in metabolism, such as those due to small changes in small areas or volumes. For example, rapid (seconds to minutes) changes in membrane channel protein populations, which can lead to fundamental and distinct changes in the type and quantity of information a synaptic membrane sends or

https://doi.org/10.1086/341769 Published online by Cambridge University Press

receives, are invisible in fMRI and PET. It is likely that the majority of cognitive processing occurs "below the radar" of inducing a metabolic change, and having the change be big enough that it can be detected indirectly.

We liken voxel and pixel analysis to having 1000 people in a large room. If you measure the oxygen consumed in this room you will not differentiate the following states, no matter how accurately you can measure the oxygen consumed: 500 people with red hats sleeping while 500 people with blue hats walk on a treadmill, 500 red hats on the treadmill while 500 blue hats sleep, or 250 red hats and 250 blue hats on the treadmill while 250 red hats and 250 blue hats sleep, and so on. But in order to understand what is actually going on in the room, we need to know which people are sleeping or exercising, and why. Measuring gross oxygen consumption won't tell us that.

The long and the short of it is that our experiments are not designed such that they give us much definitive functional information about any particular brain structure. Neuroscientists' simplifying assumptions of discreteness and constancy of function are simply not justified, nor is it clear that they will be theoretically appropriate any time soon. These assumptions end up causing scientists to abstract over the neurophysiology improperly in that they deny diversity of function or multi-tasking all the way down the line prior to examining the data for exactly these possibilities.

As a result, neuroscientists cannot use the data they get to support their claims of function, for they are assuming local and specific functions prior to gathering appropriate data for that claim. But, gathering the appropriate data is simply beyond our kin at the moment, for we have no way to approach studying the brain except through a modularist lens. Basically, we are stuck theoretically and empirically in a counter-productive circle. What do all the brain data we have amassed tell us about how the brain works? Precious little thus far.

REFERENCES

Baarsma, E. A. and H. Collewjin (1975), "Changes in Compensatory Eye Movements after Unilateral Labyrinthectomy in the Rabbit", *Archives of Otorhinolaryngology* 211: 219–230.

Bechtel, W. (2000), "From Imaging to Believing: Epistemic Issues in Generating Biological Data", in R. Creath and J. Maienschein (eds.), *Epistemology and Biology.* Cambridge, Mass.: Cambridge University Press, 138–163.

Buller, D. and V. G Hardcastle (2001), "Evolutionary Psychology, Meet the Developing Brain: Combating Promiscuous Modularity", *Brain and Mind* 1: 307–325.

Buller, D. J. (1999), "DeFreuding Evolutionary Psychology: Adaptation and Human Motivation", in V. G. Hardcastle (ed.), *Where Philosophy Meets Psychology: Philosophical Essays.* Cambridge, Mass.: MIT Press, 99–114.

Cabeza, R. and L. Nyberg (1997), "Imaging Cognition: An Empirical Review of PET Studies with Normal Subjects", *Journal of Cognitive Neuroscience* 9: 1–26.

https://doi.org/10.1086/341769 Published online by Cambridge University Press

——— (2000),"Imagining and Cognition II; An Empirical Review of 275 PET and fMRI Studies", *Journal of Cognitive Neuroscience* 12: 1–47.
Carpenter, P. A., M. A. Just, and E. D. Reichle (2000), "Working Memory and Executive Function: Evidence from Neuroimaging", *Current Opinions in Neurobiology* 10: 195–199.
Ceballos-Baumann, A. O., R. E Passingham, T.Warner, E. D. Playford, C. D. Marsden, and D. J. Brooks (1995), "Overactive Prefrontal and Underactive Motor Cortical Areas in Idiopathic Dystonia", *Annals of Neurology* 37: 363–372.
Churchland, P. S. and T. J Sejnowski (1988), "Perspectives on Cognitive Neuroscience", *Science* 242: 741–745.
Colebatch, J. G., M. P. Deiber, R. E. Passingham, K. J., Friston, and R. S. Frackowiak (1991), "Regional Cerebral Blood Flow During Voluntary Arm and Hand Movements in Human Subjects", *Journal of Neurophysiology* 65: 1392–1401.
Davies, P. S. (1996), "Discovering the Functional Mesh: On the Methods of Evolutionary Psychology", *Minds and Machines* 6: 559–585.
——— (1999), "The Conflict of Evolutionary Psychology", in V. G. Hardcastle (ed.), *Where Philosophy Meets Psychology: Philosophical Essays.* Cambridge, Mass.: MIT Press, 67–81.
Davies, P. S., J. Fetzer, and T. Foster (1995), "Logical Reasoning and Domain Specificity: A Critique of the Social Exchange Theory of Reasoning", *Biology and Philosophy* 10: 1–38.
De Gelder, B. (2000), "More to Seeing than Meets the Eye", *Science* 289: 1148–1149.
Derbyshire, S. W., A. K. Jones, P. Devani, K. J. Friston, C. Feinmann, M. Harris, S. Pearce, J. D. Watson, and R. S. Frackowiak (1994), "Cerebral Responses to Pain in Patients with Atypical Facial Pain Measured by Positron Emission Tomography", *Journal of Neurology, Neurosurgery, and Psychiatry* 57:1166–1172.
DeYoe, E. A. and D.C. Van Essen (1988), "Concurrent Processing Streams in Monkey Visual Cortex", *Trends in Neuroscience* 11: 219–226.
Emiliano, M., C. D. Frith, and J. Driver (2000), "Modulation of Human Visual Cortex by Crossmodal Spatial Attention", *Science* 289: 1206–1208.
George, M. S., T. A. Ketter, D. S. Gill, J. V Haxby, L. G. Ungerleider, P. Herscovitch, and R. M. Post (1993), "Brain Regions Involved in Recognizing Facial Emotion or Identity: An Oxygen-15 PET Study", *Journal of Neuropsychiatry and Clinical Neuroscience* 5: 384–394.
Gold, J. M. (1995), "PET Validation of a Novel Prefrontal Task: Delayed Response Alternation", *Neuropsychology* 10: 3–10.
Goldman-Rakic, P. S. (1996), "The Prefrontal Landscape: Implications of Functional Architecture for Understanding Human Mentation and the Central Executive", *Philosophical Transactions of the Royal Society of London* (B) 351: 1445–1453.
Grantham, T. and S. Nichols (1999), "Evolutionary Psychology: Ultimate Explanations and Panglossian Predictions", in V. G. Hardcastle (ed.), *Where Philosophy Meets Psychology: Philosophical Essays.* Cambridge, Mass.: MIT Press, 47–66.
Griffiths, P. (1997), *What Emotions Really Are.* Cambridge, Mass.: MIT Press.
Gulyas, B. and P. E. Roland (1994), "Binocular Disparity Discrimination in Human Cerebral Cortex: Functional Anatomy by Positron Emission Tomography", *Proceedings of the National Academy of Sciences, USA* 91: 1239–1243.
Hardcastle, V. G. (1999), *The Myth of Pain.* Cambridge, Mass.: MIT Press.
Hardcastle, V. G. and C. M. Stewart (2001), "The Structure of Neuroscientific Theories", in P. Machamer, R. Grush, and P. McLaughlin (eds.) *Theory and Method in the Neurosciences.* Pittsburgh: University of Pittsburgh Press, 30–44.
Hsieh, J. C., O. Hagermark, M. Stahle-Backdahl, K. Ericson, L. Eriksson, S. Stone-Elander, and M. Ingvar (1994), "Urge to Scratch Represented in the Human Cerebral Cortex during Itch", *Journal of Neurophysiology* 72: 3004–3008.
Jonides, J., E. E. Smith, R. A. Koeppe, E. Awh, S. Minoshima, and M. A. Mintun (1993), "Spatial Working Memory in Humans as Revealed by PET", *Nature* 363: 623–625.
Livingstone, M. and D. Hubel (1988), "Segregation of Form, Color, Movement, and Depth: Anatomy, Physiology, and Perception", *Science* 240: 740–749.

https://doi.org/10.1086/341769 Published online by Cambridge University Press

Lloyd, D. (2000), "*Terra cognita:* From Functional Neuroimaging to the Map of the Mind", *Brain and Mind* 1: 93–116.
Martin, A., C. L. Wiggs, L. G. Ungerleider, and J. V. Haxby (1996), "Neural Correlates of Category-specific Knowledge", *Nature* 379: 649–652.
McIntosh, A. R. (1999), "Mapping Cognition to the Brain through Neural Interactions", *Memory* 7: 523–548.
——— (2000), "Towards a Network Theory of Cognition", *Neural Networks* 13: 861–870.
Mishkin, M., L. G. Ungerleider, and K. A. Macko (1983), "Object Vision and Spatial Vision: Two Cortical Pathways", *Trends in Neuroscience* 6: 414–417.
Newlands, S. D. and A. A. Perachio (1990a), "Compensation of Horizontal Canal Related Activity in the Medial Vestibular Nucleus following Unilateral Labyrinth Ablation in the Decerebrate Gerbil: I. Type I Neurons", *Experimental Brain Research* 82: 359–372.
——— (1990b), "Compensation of Horizontal Canal Related Activity in the Medial Vestibular Nucleus following Unilateral Labyrinth Ablation in the Decerebrate Gerbil: I. Type II Neurons", *Experimental Brain Research* 82: 373–384.
Rauch, S. L. (2001), "A PET Investigation of Implicit and Explicit Sequence Learning", *Human Brain Mapping* 3: 271–286.
Richardson, R. (1996), "The Prospects for an Evolutionary Psychology: Human Language and Human Rationality", *Minds and Machines* 6: 541–557.
Sergent, J., S. Ohta, and B. MacDonald (1992), "Functional Neuroanatomy of Face and Object Processing. A Positron Emission Tomography Study", *Brain.* 1: 15–36.
Sergent, J., E. Zuck, S. Terriah, and B. MacDonald (1992), "Distributed Neural Network Underlying Musical Sight-reading and Keyboard Performance", *Science* 257: 106–109.
Shaefer, K.-P. and D. L. Meyer (1974), "Compensation of Vestibular Lesions", in H. H. Kornhuber (ed.) *Handbook of Sensory Physiology,* Vol. Vi/2. New York: Plenum, 462–490.
Shapiro, L. (1998), "Do's and Don'ts for Darwinizing Psychology", in C. Allen and D. Cummins (eds.) *The Evolution of Mind.* New York: Oxford University Press, 243–259.
——— (1999), "Presence of Mind", in V. G. Hardcastle (ed.), *Where Philosophy Meets Psychology: Philosophical Essays.* Cambridge, Mass.: MIT Press, 83–98.
Shapiro, L. and W. Epstein (1998), "Evolutionary Theory Meets Cognitive Psychology: A More Selective Perspective", *Mind and Language* 13: 171–194.
Shulman, R. G. (1996), "Interview with Robert G. Shulman", *Journal of Cognitive Neuroscience* 8: 474–480.
Sirkin, D. W., W. Precht, and J. H. Courjon (1984), "Initial, Rapid Phase of Recovery from Unilateral Vestibular Lesion in Rat not Dependent on Survival of Central Portion of Vestibular Nerve", *Brain Research* 302: 245–256.
Sober, E. (1997), "Is the Mind an Adaptation for Coping with Environmental Complexity?", *Biology and Philosophy* 12: 539–550.
Sterelny, K. (1995), "The Adapted Mind", *Biology and Philosophy* 10: 365–380.
Stufflebeam, R. S. and W. Bechtel (1997), "PET: Exploring the Myth and the Method", *Philosophy of Science* 64 (Supplement): S95–106.
Wandell, B. A. (1999), "Computational Neuroimaging of Human Visual Cortex", *Annual Review of Neuroscience* 22: 145–173.
Zatorre, R. J., A. C. Evans, E. Meyer, and A. Gjedde (1992), "Lateralization of Phonetic and Pitch Discrimination in Speech Processing", *Science* 256: 846–849.

https://doi.org/10.1086/341769 Published online by Cambridge University Press