Multiple Realizability Revisited: Linking Cognitive and Neural States

William Bechtel; Jennifer Mundale

Philosophy of Science, Vol. 66, No. 2 (Jun., 1999), 175-207.

Multiple Realizability Revisited: Linking Cognitive and Neural States*

William Bechtel†‡
Philosophy-Neuroscience-Psychology Program
Department of Philosophy, Washington University

Jennifer Mundale
Department of Philosophy, University of Central Florida

The contention that psychological states are multiply realizable in different substrates has been used to support the contention that neuroscience is not likely to be very useful in guiding an understanding how cognition works. But in the context of scientific research, how seriously should we really take this threat of multiple realizability? By examining how brain areas are identified in neuroscience (where the approach is comparative and employs functional criteria), we show that the skepticism about neuroscience's role in understanding cognition is misguided and that the apparent success of multiple realizability is based on methodological error. With respect to the former point, we develop and analyze a representative case in which the organization of the brain provides crucial information about the organization of the cognitive system. Then we show that the initial plausibility of claims to multiple realizability rest on (a) mismatching a broad-grained criterion (to show sameness of psychological states) with a fine-grained criterion (to differentiate brain states), and (b) a failure to attend to the purposes for which taxonomies of brain and psychological states are developed.

---

*Received December 1997; revised October 1998.

†Send requests for reprints to William Bechtel, Philosophy-Neuroscience-Psychology Program, Department of Philosophy, Campus Box 1073, Washington University in St. Louis, St. Louis, MO 63130, USA.

‡Versions of this paper were presented at the Australian Cognitive Science Society Conference, the Conference on Body, Mind, and Brain at the University of Oregon, and at the National Neurobiology Laboratory, National University of Mexico, Querataro, Mexico. We are appreciative for the helpful comments and suggestions made by these audiences. We also received very helpful comments and suggestions from two referees for this journal, Patricia Churchland and Owen Flanagan, for which we are most appreciative.

Philosophy of Science, 66 (June 1999) pp. 175–207. 0031-8248/99/6602-0001$2.00

> Anatomy first, and then physiology; but if physiology first, then not without anatomy. B. von Gudden (cited in Zeki 1993, 72)

1. Introduction. The claim, originating with Hilary Putnam, that psychological states are multiply realizable, has become orthodoxy in the philosophy of mind. It is a fundamental tenet of functionalist accounts of mind, according to which psychological states are to be identified in terms of their interactions with each other and with sensory inputs and behavioral outputs. The claim of multiple realizability is the claim that the same psychological state can be realized by different brain states. Thus, it is claimed that there is a many-to-one mapping from brain states to psychological states. Putnam employed the claim of multiple realizability to argue against the identity theory that was popular in the 1960s, which sought to identify psychological states with specific brain states. In sum, if the relation between brain states and psychological states is many-to-one—the main thrust of multiple realizability—then no identity relation holds between them. One common corollary of this rejection of the identity thesis is the contention that information about the brain is of little or no relevance to understanding psychological processes.

Historically, a major inspiration for the multiple realizability argument was the emergence of artificial intelligence research which attempted to generate computer programs that performed the same cognitive operations as human agents. The guiding assumption was that if mental activities could be characterized in terms of operations in a system, then they should be able to be implemented in different hardware or wetware, thereby providing alternative realizations. Taking this a step further, many philosophers became convinced that the same mental activities could be realized in brains of aliens with radically different composition from ours. The upshot of these speculations about artificial and alien minds is a metaphysical claim that mental processes are the operations themselves, and are not identified with whatever biological or other substances realize them. For the most part, we will have nothing to say about these speculative arguments, nor are we primarily concerned with the metaphysical claim. Our primary concern, rather, is with the implication drawn from the multiple realizability argument that information about the brain is of little or no relevance to understanding psychological processes. This argument generally relies not on the mere logical possibility of multiple realizability, but on the contention that mental or psychological states are multiply realized in existing biological organisms.

The claim that psychological states are in fact multiply realized is usually argued on fairly intuitive grounds. It is widely assumed that different individuals, even of different species, can be in the same psychological state, characterized either at the folk psychological level or at the level of information processing psychology. For the most part, information processing psychology has focused on commonalities in psychological processes across individuals. Yet, it is well known that even within a species brains differ. Even within an individual over time there are differences (neurons die, connections are lost, etc.). Thus, multiple realizability seems to arise within species (including our own) and even within individuals. And it is simply compounded as one moves beyond our species. Putnam (1967), for example, considers pain in mammals, reptiles, and mollusca, and contends that these organisms are not likely to be in the same brain state when they experience pain. Thus, their pains are multiply realized.

This line of reasoning appears plausible, but a proper examination of neurobiological and cognitive neuroscience practice will show that the claim that psychological states are in fact multiply realized is unjustified, and that what is usually taken to be evidence for it, is not. It is precisely on the basis of working assumptions about commonalities in brains across individuals and species that neurobiologists and cognitive neuroscientists have discovered clues to the information processing being performed.

When Putnam characterizes brain states, he treats them as physical-chemical states of the brain. While an appeal to physics and chemistry may be an intuitively plausible way of characterizing brain states, it is not how neuroscientists characterize them. Actually, the notion of a brain state is a philosopher's fiction; a notion closer to what neuroscientists would use is activity in the same brain part or conglomerate of parts. What is then critical to the practice of neuroscientists is how they identify brain parts. In Sections 2 and 3, by considering how neuroscientists map brains, we will demonstrate that the scientifically operative notion of a "brain state" differs from the sort of fine-grained conception of brain states employed in philosophy; it is more coarse-grained and linked to an equally coarse-grained notion of psychological state.

In Sections 2 and 3 we will focus respectively on the neuroanatomical and neurophysiological practice of carving up the brain. Two aspects of these practices will be especially noteworthy: (1) the appeal to function, especially psychological function, is an *essential* part of both the project and its tools, and (2) the cartographic project itself is frequently carried out comparatively—across species. For multiple realization to be a serious option, brain taxonomy would have to be carried out both independently of psychological function, and without comparative evaluation across species. One might think, at first glance, that the ability to make comparisons across species actually depends upon multiple realizability. In fact, it is the very *similarity* (or more precisely, *homology*) of brain structures which permits us to generalize across certain species. So in this latter respect, in the context of neuroscientific research, they are not multiply realized.

If the gloomy implications of multiple realizability were to be taken seriously, one would not expect results based on comparative neuroanatomical and neurophysiological studies to be particularly useful in developing functional accounts of human psychological processing. Some advocates of multiple realizability, such as Fodor (1975), have explicitly contended that, in fact, neural studies have not and will not enhance psychological understanding.[1] But within this decade cognitive neuroscience is rapidly demonstrating the potential for scientific breakthroughs as a result of relating psychological states to brain states. In Section 4 we will develop an example of decomposing the process of visual perception into different psychological processes where neuroscience has indeed been the leading player. This is not the only example that could be chosen; analyses of brain organization also suggest models of psychological processing in such domains as language processing, memory, and attention. In fact, the attempt to use brain organization to guide the analysis of psychological functions is not new; it has roots at least as far back as the 19th century. What is new is that the neuroscientific tools for decomposing functions have matured to such a state where they are able to generate robust results that provide fruitful guides to psychology.

If our case against it is successful, then the appeal of the multiple realizability doctrine in philosophy requires explanation; in Section 5 we focus on two reasons for its appeal. One major factor is that philosophers appeal to different grain sizes in the taxonomies of psychological and brain states, using a coarse-grain in lumping together psychological states and a fine grain in splitting brain states. When a common grain size is insisted on, as it is in scientific practice, the plausibility of multiple realizability evaporates. Another reason that the multiple realizability argument has seemed plausible is that it is presented in a contextual vacuum. Its advocates have typically failed to provide us with a context in which to ask whether two occurrences involve the same or different psychological or brain state. There are different contexts for constructing taxonomies, and these will result in

---

[1] An alternative view of the implications of multiple realizability for relating neuroscientific and psychological explanations, with which we are sympathetic, is developed by Richardson (1979), who argues that the multiple realizability embraced by functionalism is fully compatible with reductionism, properly understood.

different grain sizes for both psychology and neuroscience. In assessing individual differences, researchers will adopt a fine-grained analysis of both psychological and brain states. In developing an evolutionary perspective, researchers will adopt a coarse grain. Context matters to the type of taxonomy one develops, and the philosophical failure to attend to context has resulted in an artificial comparison of psychological and brain states.

2. Neuroanatomical Approaches to Mapping the Brain. The task of mapping the brain has been challenging and controversial one. The brain does not come pre-delineated into parts, no more than does the world's geography. Rather, investigators have to decide what kinds of criteria to use to mark borders. Some of the differentiation of brain areas in higher mammals is rather straightforward. For example, the cerebellum is rather sharply differentiated from the cerebral cortex, and within the cerebral cortex the various gyri and sulci are rather well differentiated and have been assigned labels. While these features are often useful reference points, neuroscientists have not found them adequate. One shortcoming is that there is considerable individual variability within species in the location of the various gyri and sulci. Another is that areas delineated by gyri and sulci have little correspondence to functionally significant areas.

One of the most celebrated maps of the brain was produced at the beginning of the 20th century by Korbinian Brodmann. The neuron doctrine was just gaining acceptance as Brodmann was conducting his research. According to this theory, nerve fibers are seen to be individual cells, as opposed to an uninterrupted syncytium, or network of fibers; furthermore, each neuron is seen to be composed of an axon and dendrites, and separated from adjacent neurons by a gap, or synapse. New stains, such as the Golgi stain, which made it possible to determine that neurons were separate cells, also made it possible to recognize different *types* of neurons in the cortex, and to discover that the cortex generally consisted of six different layers of neurons (which manifested themselves in stain preparations due to different cell types and densities). Not only were there differences between layers of cortex, but it became apparent to Brodmann as well as a number of contemporaries that there were significant differences between regions of cortex, according to such features as the types of cells found and the thickness of different layers. Relying on such criteria, Brodmann (1909) developed his map of 47 brain areas in the human cortex (Figure 1).

Two points are especially noteworthy about Brodmann's map. First, Brodmann's goal of identifying different regions of the brain was to distinguish areas of the brain which performed different functions: "Al-

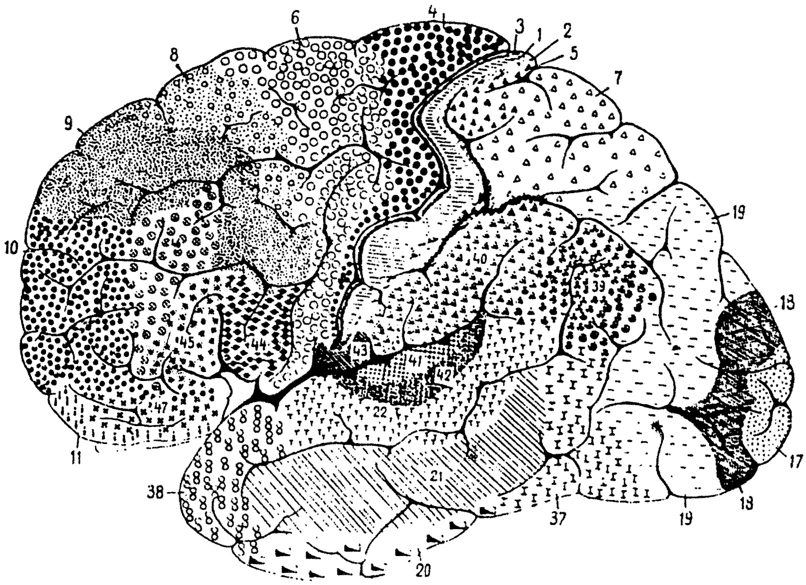

Fig. 1 Korbinian Brodmann's (1909) lateral-view map of the human cortex. He differentiated brain areas primarily in terms of the thickness of different layers of cortex and the type and density of neurons in these layers. Brodmann numbered areas of the cortex in the order in which he studied them. Areas 1, 2, and 3 constitute what we now take to be the sensory cortex, area 4 the primary motor cortex, and area 17 the primary visual cortex.

though my studies of localisation are based on purely anatomical considerations and were initially conceived to resolve only anatomical problems, from the outset my ultimate goal was the advancement of a theory of function and its pathological deviations" (1909/1994, 243). It is because they were likely to be functionally important that Brodmann thought cytoarchitectonic differences should matter in differentiating areas. Second, in identifying brain regions, Brodmann worked comparatively. Thus, in defending the claim of six different layers in cortex, he reports preparations made from 55 species ranging over 11 different orders of mammals, and presents diagrams from the cat, wallaby, kinkajou, and rabbit, as well as the human. When he turns to mapping the cortex, Brodmann presents maps for several other species in addition to humans: two lower monkeys (guenon and marmoset), lemur, flying fox, kinkajou, rabbit, ground squirrel, and hedgehog, using the same numbering systems to identify homologous areas in the different species. (Compare Figure 2, Brodmann's map of the guenon, with Figure 1.) From these maps he argues that there is similarity in

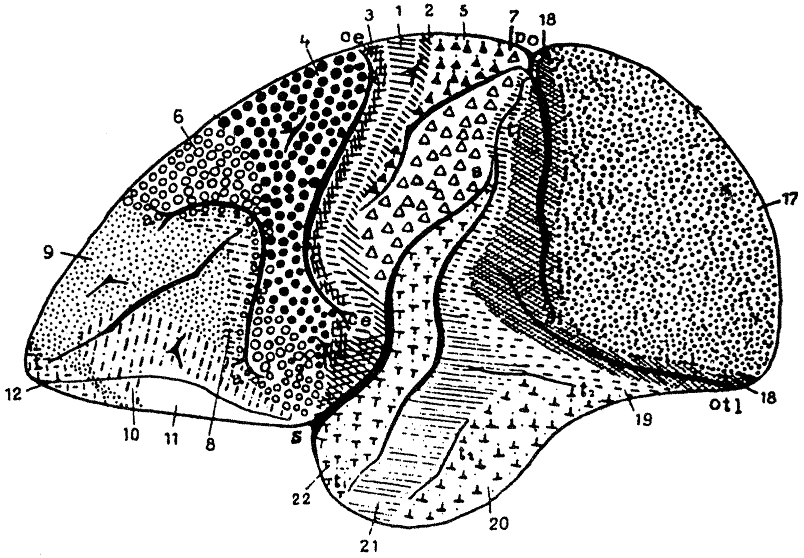

Fig. 2 Brodmann's (1909) lateral-view map of the cortex of the guenon. Brodmann adapted the same numbering system as he used on the human brain for each species that he studied. In some species some areas identified in the human brain were not identifiable, but each region of the cortex that was identified in a given species was associated with a region of the human cortex.

the overall patterns of parcellation, constancy in broader regions (precentral, granular frontal region, etc.) across species, and persistence of individual areas. While Brodmann also drew attention to differences between species, his primary concern was to show commonalities in brain regions, not differences.

Brain maps were controversial, and Brodmann was not the only scientist engaged in this project. He worked in the Institute for Brain Research in Berlin, which was headed by Oskar Vogt, who together with Cécile Vogt, offered a map which divided the brain into many more regions. The ambiguities in identifying regions of the brain are well illustrated by two major critics of brain mapping, Karl Lashley and George Clark. To make their point, they set out to map the cortex of the spider monkey (*Ateles geoffroyi*). Each took a different animal, and developed his own map. Even though they report having relied on the same techniques and on established accounts of primate brain organization, they found significant differences in the maps they generated. Some of these were presumably due to differences between the two animals, while others were due to the ambiguities in developing cortical maps. Despite these difficulties in arriving at a "correct" map, and recognized deficiencies in his map, Brodmann's map was widely accepted, and is still used today as a common basis for referring to brain regions.

Many of the advances beyond Brodmann in the current day involve developing maps of the brain with much finer resolution. For example, working just in the visual areas of the macaque, David van Essen and his collaborators have differentiated 32 different processing areas (Felleman and van Essen 1991). These are shown in Figure 3, with the most prominent areas labeled and shaded. Just as with Brodmann, there remains considerable uncertainty as to exactly how many brain regions there are. Moreover, there are a variety of tools neuroanatomists use

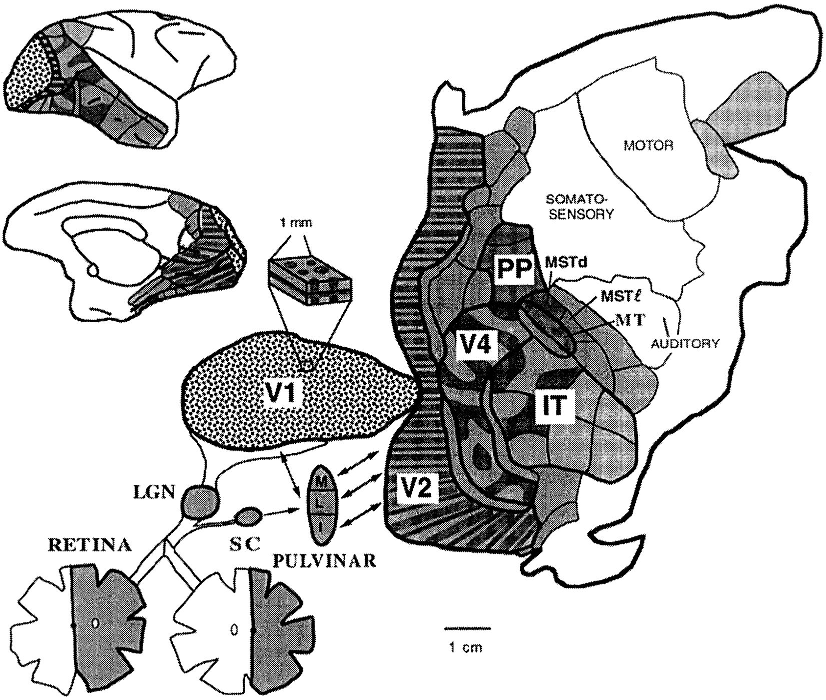

Fig. 3 Flat map of the right hemisphere of the macaque identifying major subcortical and cortical visual processing areas (van Essen and Gallant 1994, 2). The blob and interblob structure as well as layer 4B are differentiated on the expanded representation of VI while the thin, thick, and interstrip regions of V2 are shown. At the upper left are lateral and medial views of the right hemisphere showing where the respective cortical areas are on the three-dimensional cortex. Reprinted with permission of the Cell Press and David C. van Essen.

to identify brain regions. Felleman and van Essen single out three different sets of criteria that figure prominently in their study: architectonics, connectivity, and topographical organization. They note that architectonics, Brodmann's tool, has been useful in identifying only a minority of the areas of visual cortex. Topographical organization, which refers to the orderly projection of the visual field over each area, was useful in distinguishing about half of the areas.[2] Finally, connectivity patterns between putative areas were useful for distinguishing and identifying almost all of the visual areas. Topographical organization and connectivity are, even more clearly than architectonics, features of functional import. Connectivity is important insofar as it provides the vehicle for information to be sent from one processing area to another. Preservation of the topographical arrangement of the visual scene, as projected onto the retina, in other processing areas allows the brain to use spatial relations in the processing area to stand in for spatial relations in the visual scene. Moreover, as we shall see in Section 4, the interest in mapping the visual cortex derives from the assumption that the areas that are distinguished will perform different information processing activities. Lastly, we should emphasize that while van Essen's ultimate interest is in human visual processing, this work has been carried out on the macaque. The clear assumption is that the neural organization in the macaque will provide a defeasible guide to the human brain.

3. Use of Psychological Function in Brain Mapping: Deficits, Stimulation, and Neuroimaging. So far we have focused on attempts to map the brain that rely primarily on neuroanatomical criteria. We have seen that such brain mapping has had as its goal dividing the brain into functionally relevant areas. But these neuroanatomical studies are not the first, nor even the primary tools that have been used in brain mapping. Neurophysiological approaches came first, and some neurophysiological techniques make even more explicit reference to psychological function. Just as with the more neuroanatomical endeavors just discussed, neurophysiological studies are frequently comparative.

One of the first to map psychological function to the brain was Franz Josef Gall, who proposed brain loci for 27 psychological faculties of widely different sorts (e.g., instinct for reproduction, love of offspring,

---

[2] Topographical organization is often identified physiologically through recording from individual cells and comparing the receptive fields of cells in different areas, but it was often first identified neuroanatomically, either through staining axons, tracing pathways of nerve degeneration, or through autoradiography with radioactive tracers. See, e.g., Cragg 1969 and Zeki 1969.

memory of things and facts, memory of words, mathematical ability, and sense of metaphysics). While Gall was a superb neuroanatomist, in this endeavor he relied on the cranium, not the brain itself, on the mistaken assumption that the protrusions and indentations in the cranium corresponded to oversized or undersized brain areas. He further assumed that increased brain volume devoted to a psychological faculty would result in increased performance of its function. The operative tool in his endeavor was to correlate protrusions and indentations of the skull (assuming these reflected underlying brain regions) with increased or diminished activities of the psychological faculties in an individual. Although many of the assumptions underlying Gall's program turned out to be incorrect (and his approach to identifying correlations flawed), Gall set the agenda for many subsequent researchers, who adopted the label *neophrenologists* (after the name Gall's onetime collaborator, Spurzheim, applied to their program).

What is significant for our purposes about Gall is that he demarcated cranial areas in light of psychological functions, insuring that function was built into the characterization of the region. He based his claims on a comparative study of skulls, and contended that 19 of his 27 faculties were common to both humans and animals. The attempt to correlate skulls and psychological faculties was quickly recognized to be a failure, but other researchers soon found other ways to localize psychological faculties in the brain. One approach was to focus on deficits in psychological capacities that followed damage to particular parts of the brain, either damage arising naturally or induced by lesioning the brain. The *locus classicus* for this sort of endeavor was Paul Broca's (1861) research relating damage to the third convolution of the left frontal lobe with loss of articulate speech. Language is perhaps the best known cognitive domain in which deficits have been correlated with brain damage, but in fact deficit studies have been applied to many other domains, including vision (as we will discuss below). In the case of language, deficit studies have been restricted to humans since humans have the best-developed language system. Nonetheless, it is important to note that in interpreting these deficits, researchers implicitly reject multiple realization among human brains and assume that damage to a brain area in anyone will result in a deficit to a particular cognitive function that is performed by that area in undamaged brains. Moreover, cross-species comparisons are sometimes employed even in the study of language processing. To understand what cognitive processing a given area contributes to language processing, researchers may examine the deficits arising in other species with lesions in the same the areas indicated for language performance in humans (Deacon 1997).

While deficit studies have played an important role in attempts to localize psychological function in the brain, there are well-known problems in analyzing deficits. One is that although the correlation of a deficit with a brain region does show that the brain region is necessary for the normal performance, it does not reveal *what* that brain region contributes to normal performance. Its contribution could be quite ancillary to the function involved. (For example, damage to the bus of a computer will interfere with its ability to execute steps in a word processing program, but this would not show that the bus was the locus of the performance of these steps.) An alternative approach is to use mild electrical stimulation to excite the suspected area in normally functioning organisms and determine what psychological activity is thereby increased. The challenge is to detect the increased performance. In humans, where verbal reports can be elicited from stimulated subjects, Wilder Penfield and subsequent investigators used this approach to map somatosensory cortex as well as in more controversial attempts to localize such psychological states as specific memories. The classical use of this approach was to identify regions of the brain related to specific behaviors. In 1870 Gustav Fritsch and Eduard Hitzig developed techniques for stimulating brain areas of a dog with relatively weak currents and found that this generated muscular movements on the opposite side of the body. Subsequently, David Ferrier (1886) greatly extended this approach.

Ferrier's study was explicitly comparative: he reported stimulation experiments on macaque monkeys, dogs, jackals, cats, rabbits, guinea-pigs, rats, pigeons, frogs, and fishes. The studies on pigeons, frogs, and fishes elicited very few locations where stimulations could elicit any responses. But in all of the mammals he found a number of such sites (Figure 4). To make the findings comparable, Ferrier attempted to determine homologous structures according to sulci-defined borders in the brains of monkeys, dogs, jackals, and cats. He then charted the locations at which specific responses could regularly be produced, employing a common numbering system for responses for all species (Table 1). Ferrier could not do similar experiments on humans; nonetheless, in the final chapter he projects his results onto the human brain (Figure 5), again using superficial reference points such as gyri and sulci as his guides.

Ferrier was not just eliciting responses from various areas of cortex; his deeper interest was in understanding the contribution of various areas of cortex. Thus, he contended that some of the responses were not specifically motor, and introduced lesion studies to complement stimulation studies in determining the specific contribution of an area:

> The mere fact of motion following stimulation of a given area does not necessarily signify a motor region. The movements may be the

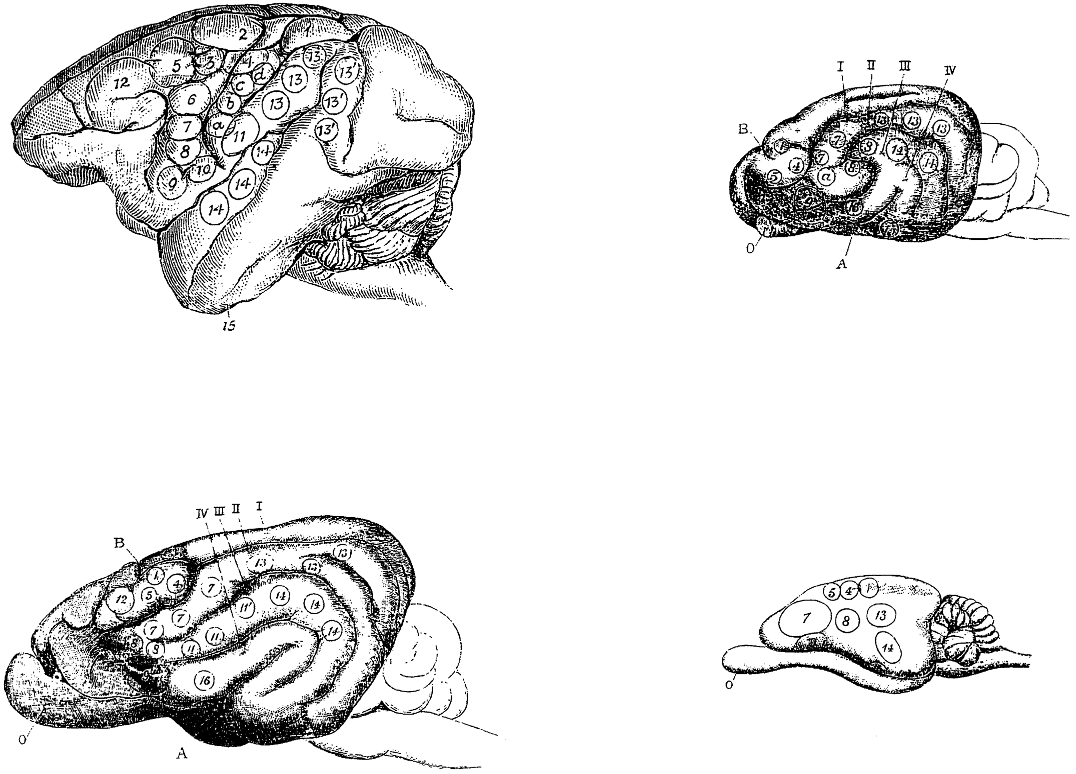

Fig. 4 Ferrier's (1886) maps of sites on the left hemispheres of monkey (upper left), dog (lower left), cat (upper right), and rabbit (lower right) from which motor responses were elicited with mild electrical stimulation. The same numbering pattern is used for each species; Table 1 indicates the response that resulted from stimulating each area. Not all responses that could be elicited in the monkey could be elicited in the other species.

TABLE 1. Areas the Ferrier Found to Respond to Mild Stimulation

1.	Opposite hind limb is advanced as in walking
2.	Flexion with outward rotation of the thigh, rotation inwards of the leg, and flexion of the toes
3.	Movements comparable to 1 and 2, plus movements of the tail
4.	Opposite arm is adducted, extended, and retracted, the hand pronated
5.	Extension forward of the opposite arm (as if reaching or touching something in front)
a,b,c,d	Clenching of the fist
6.	Flexion and supination of the forearm
7.	Retraction and elevation of the angle of the mouth
8.	Elevation of the ala of the nose and upper lip
9.	Opening of the mouth, with protrusion of the tongue
10.	Opening of the mouth, with retraction of the tongue
11.	Retraction of the angle of the mouth
12.	Eyes open widely, pupils dilate, and head and eyes turn to the opposite side
13,13′	Eyes move to the opposite side
14.	Pricking of the opposite ear, head and eyes turn to the opposite side, pupils dilate widely
15.	Torsion of the lip and semiclosure of the nostril on the same side

> result of some modification of consciousness incapable of being expressed in physiological terms, or they may be reflexes; or they may be truly motor in the sense of being due to irritation of a part in direct connection with the motor strands of the crus cerebri and spinal cord. The method of stimulation is by itself incompetent to decide these questions, and requires as a complement the strictly localised destruction of these areas stimulation of which gives rise to definite motor manifestations. (1886, 268)

Ferrier, relying on a combination of lesion studies and stimulation studies, argued that visual processes were localized in the angular gyrus (areas marked 13 in the stimulation maps, where the behavioral response consisted in movement of the eyes) and the occipital lobe,[3] while auditory processes are localized in the temporal cortex (areas marked 14 on the stimulation maps, where the behavioral response consists in the pricking of the opposite ear and head turning). The combination of relatively specific motor responses, and loss of motor responses with

---

[3] Other researchers had already argued for localizing vision in the striate cortex of the occipital lobe; Ferrier initially rejected any role of the occipital lobe in vision, and subsequently allowed it a role, but only in conjunction with the angular gyrus. One factor that probably mislead Ferrier was that his lesions of the angular gyrus also cut the underlying visual pathways from the lateral geniculate nucleus of the thalamus to the occipital lobe (Finger 1993). As well, his lesions of the occipital lobe likely spared some of the areas responsible for foveal vision (Gross 1998), thereby misleading him into thinking that the occipital lobe was not necessary for vision.

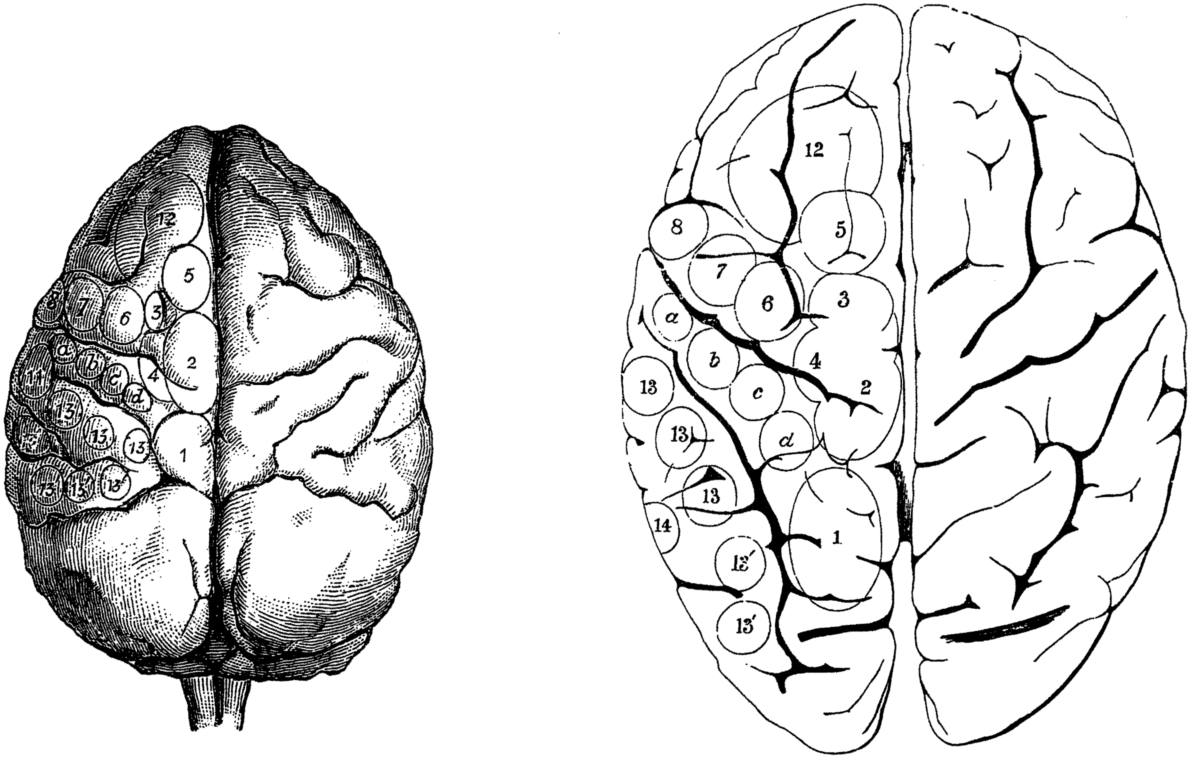

Fig. 5 Ferrier's (1886) attempt to project from the sites on the cortex of the monkey (left) at which electrical stimulation generated motor responses to analogous locations on the human cortex (right). To determine the corresponding sites, Ferrier relied on the projection of superficial features such as major gyri and sulci from the macaque brain to the human.

lesions, led Ferrier to treat the areas on the two sides of the central fissure (Fissure of Rolando) as the primary motor areas.

Recently-developed techniques permit imaging of activity in the brain associated with cognitive activities. First positron emission tomography (PET) and more recently functional magnetic resonance imaging (fMRI), have provided an additional approach to identifying brain areas linked to psychological function. It is important to bear in mind that neither of these techniques measures brain activity (e.g., the firing of neurons) directly, but rather blood flow, which is taken to be an indirect measure of brain activity. There are also a number of complex epistemological issues concerning the interpretation of imaging studies, which we will not go into here (see Bechtel in press). From our vantage point, what is important about these techniques is the central place analysis of psychological processes plays in them. Researchers generally do not do a neuroimaging study to see what areas are involved in performing an overall task, such as reading a word or solving a problem, since huge portions of the brain are active in most psychological activities. Many researchers employ the *subtractive method* to focus on brain activities associated with component psychological processes. The subtractive method was developed by Donders in the mid-19th century for use in purely behavioral studies using reaction time. In this method in reaction time studies researchers choose two tasks that are thought to differ in just one or a few basic operations. They then subtract the time required to perform one task from that required to perform another task that is thought to involve all the same steps plus one or more additional ones. If the steps were performed serially, the difference in the reaction times between the two tasks can be attributed to the extra steps required. In neuroimaging, subtraction is applied spatially: for every pixel in the multiple slices taken through the brain, one subtracts the activation level registered for performing one task from that generated in performing another task which is thought to include all the steps required for the first task, plus additional ones.

The approach is illustrated in a now classic study by Petersen et al. (1989), in which subjects were required in different trials to read nouns or listen to nouns, to say them aloud, and to generate verbs appropriate to them. All of the processing required for reading the nouns was thought to be involved in saying them, and all of the processing involved in saying them was thought to be involved in generating related verbs. Thus, the areas active in the first task were subtracted from the areas active in the second, so that the difference image is designed to indicate areas employed in just one component psychological function (constructing the pronunciation or generating the verb).

Identifying brain areas through neuroimaging depends critically on the cognitive tasks subjects are asked to perform; thus, the possibility of multiple realizability is restricted at the outset. Three other aspects of neuroimaging research are particularly relevant to the question of multiple realizability. First, there are differences between individual brains and researchers have found it necessary to refer to a common atlas (often the Talairach and Tournoux 1988 atlas) in order to designate which areas were active. Since each brain is anatomically different, this requires a mapping of different brains onto the common atlas. (This step is less critical in fMRI since one can map the functional results onto a structural MRI image of the same brain. But if one wants to compare active areas across studies, the preferred basis is in terms of coordinates in the Talairach and Tournoux atlas or proposed successors; see van Essen et al. 1998.) Second, in PET especially, but also in MRI, the signal to noise ratio is relatively low, requiring averaging across subjects. Averaging will cancel out individual differences; thus, the fact that any results at all survive averaging as well as transformation onto a common brain plan indicates a great deal of commonality. This alone suggests much less variability than the multiple realizability arguments would allow.[4] Finally, researchers are generally interested in identifying activation areas with neuroanatomically distinguished areas. However, most neuroimaging to date is performed on humans while the most detailed neuroanatomical and neurophysiological work (using, e.g., single-cell recording) has been done on other species. As a result, researchers often have to try to coordinate the imaging work on humans with neuroanatomy from other primate species (especially the macaque), and thus are assuming that cognitive functions are not differently realized in the two species. Currently, researchers are trying to develop additional bridging techniques by trying to find ways to perform fMRI studies on awake, behaving monkeys.

4. Decomposing Visual Processing: Using the Brain to Guide Cognitive Decomposition. In the previous section our focus was on how the identification of brain parts itself relied on (1) functional considerations that were frequently psychological in nature, and (2) cross-species comparisons. Both of these features undercut the possibility of finding instances of multiple realization of psychological states within terrestrial animals. In this section we turn directly to the task of decomposing psychological function. If the taxonomies of brain states and psycho-

---

[4] Averaging can also conceal small differences in the areas of greatest activation. This can be particularly serious when the activity falls near the boundary of a processing area.

logical states were as independent of each other as the multiple realizability argument suggests, brain decomposition would be a poor guide for psychological decomposition. But we will show how understanding of psychological function is being fostered by appeal to the brain and its organization. We use research on visual processing as our example. The standard view of visual processing in cognitive science is that it is an inferential process that builds up step-by-step from very local information about features of the visual array to a comprehensive perception of the visual scene (Marr 1982, Fodor and Pylyshyn 1981). This portrayal seemed to draw support from Hubel and Wiesel's (1962, 1968) discovery of a hierarchy of feature detectors (simple, complex, and hypercomplex) in primary visual cortex (V1) and the adjacent area V2. Beyond that, however, most research in psychology and AI has been guided principally by information processing considerations, not information about brain organization. Yet, research on the neural structures of the visual system is suggesting a decomposition of visual processing that was not anticipated by psychological or artificial intelligence research but which has, in the 1990s, been adopted by researchers in those disciplines.[5]

Neuroanatomical studies from the late 19th and early 20th century had established that the primary connections between the retina and the cortex in mammals was via a channel passing through the lateral geniculate nucleus of the thalamus in the mid-brain to the rear of the occipital lobe of the cortex (Figure 6).[6] The projection area at the rear of the occipital cortex, located on the banks and lips of the calcarine sulcus, is distinguished by a striation pattern (discovered by Francesco Gennari in 1782) and hence is often referred to as the *striate cortex* (its distinctive cytoarchitecture was recognized by Brodmann and since it was the 17th area of cortex he examined, it has also come to be known as Brodmann's area 17). Around the turn of the century, neuropathologists such as Salomon Eberhard Henschen (1893, 1903) demon-

---

[5] During the 1990s as well, theorists have proposed a radical reconception of what the goal of visual processing is. Rather than involving the construction of a complete representation of the visual world, Ballard (1991) in his program of animate vision and Churchland, Ramachandran, and Sejnowski (1994) in their critique of pure vision, argue that the visual system only needs to support extraction of information about the visual world that is needed to guide different forms of action.

[6] In non-mammalian vertebrates, the primary pathway is from the retina to the optic tectum, known as the *superior colliculus* in mammals; this pathway still exists in mammals and, amongst other things, sends projections to the pulvinar in the thalamus and from there to other visual areas in the cortex. This pathway is generally recognized to play a role in controlling saccadic eye movements, but Milner and Goodale (1995) offer evidence that it figures in mammals' ability to orient towards visual targets.

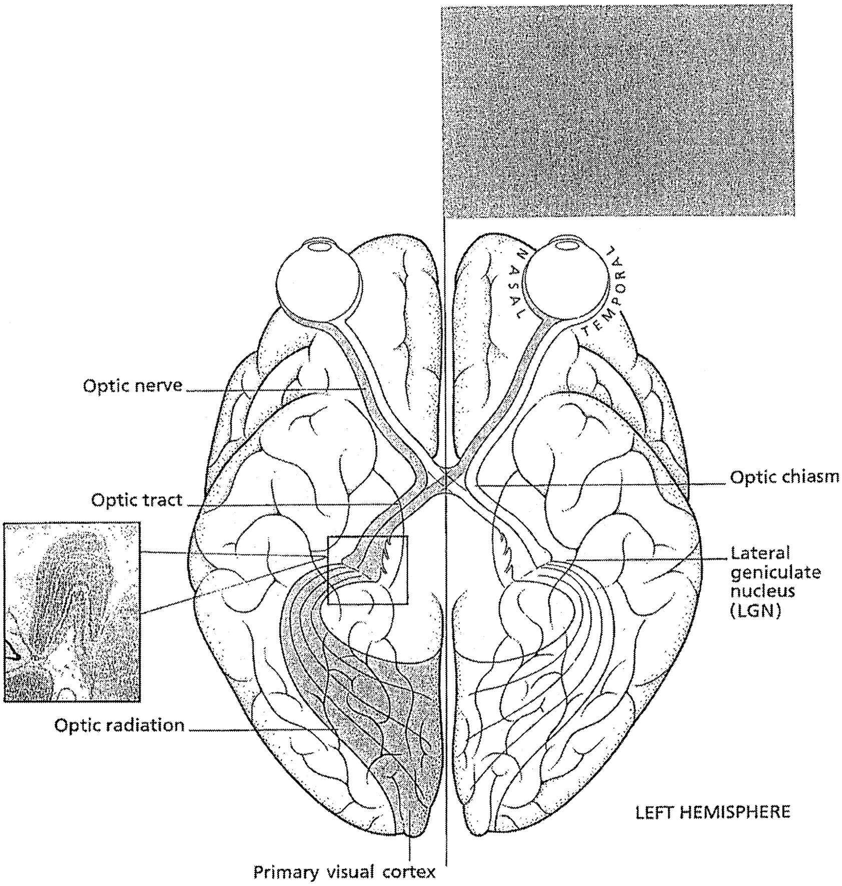

Fig. 6 The primary neural pathway in mammals from the retina to primary visual cortex in the occipital lobe. Nerves project from the retina to the lateral geniculate nucleus of the thalamus, and from there to primary visual cortex. At the optic chiasm, nerves carrying information from the right visual field of each eye (shown in dark) join and project to the left lateral geniculate nucleus and primary visual cortex; nerves carrying information from the left visual field likewise project to the right primary visual cortex. The six-layer structure of the lateral geniculate nucleus is shown in the box on the left. Figure reprinted with permission from Zeki 1993, 25.

strated that various forms of damage to this area result in different forms of blindness: complete destruction of the striate cortex on one side leads to blindness for the contralateral half of the visual field (hemianopia), whereas damage to one lip of the sulcus results in blindness in one quarter of the visual field (quadrantanopia). Damage to a smaller area produces blindness in a small part of the visual field (scotoma). Henschen argued for a highly localized mapping from the retina to the striate cortex; although the mapping that has been accepted reverses that proposed by Henschen (he mapped the fovea of the retina onto the anterior portion of the striate cortex), the idea of topological mapping has been widely supported.

Despite proposals to localize visual functions such as color processing outside the striate cortex, the standard view in the first half of the 20th century was that basic visual processing (perception) occurred in the striate cortex; the areas surrounding the striate cortex, Brodmann's areas 18 and 19, were construed as association areas in which visual impressions could be related and interpreted in terms of previous visual impressions. The account of how these association processes were recognized was left quite vague. One reason for this was that the cytoarchitectonic and myeloarchitectonic tools available in the first part of the 20th century failed to differentiate any areas within areas 18 and 19; even the distinction between the areas was not very sharp. Thus, even though von Bonin and Bailey's (1951) areas OA and OB correspond approximately to Brodmann's areas 18 and 19, they employed merging colors in their map because of uncertainty about the actual boundary. Lashley and Clark (1946) were generally critical of efforts in brain cartography to distinguish areas beyond the primary sensory and motor areas, and they argued that in fact there was no distinction among Brodmann's areas 7 (in the posterior parietal cortex) 18, and 19.

One of the first suggestions that Brodmann's areas 18 and 19 might contain distinct visual processing areas came from Cowey (1964), who demonstrated a second topographically organized area in a part of Brodmann's area 18, immediately adjoining primary visual cortex (V1), which came to be known as V2. Using single-cell recording, Hubel and Wiesel (1965) confirmed the topographical organization of this area and identified yet a third area, V3, in Brodmann's area 18. Cragg (1969) and Zeki (1969) found further systematic projections from V1 to other areas within Brodmann's areas 18 and 19, which came to be known as V4 and either V5 or MT, indicating that the same topographical arrangement was preserved in each of these areas. Areas V2, V3, V4, and MT are deep in the lunate and superior temporal sulci in the macaque; their locations, though, are easily indicated in horizontal slices (Figure 7). Initially there seemed to be few anatomical differences between these areas; the main difference was that as one proceeded further from V1 the size of the receptive fields of the cells increased, indicating a less detailed mapping. After Margaret Wong-Riley's (1979) application to these areas of a stain for cytochrome oxidase activity, which figures in the basic metabolism of the cell, researchers

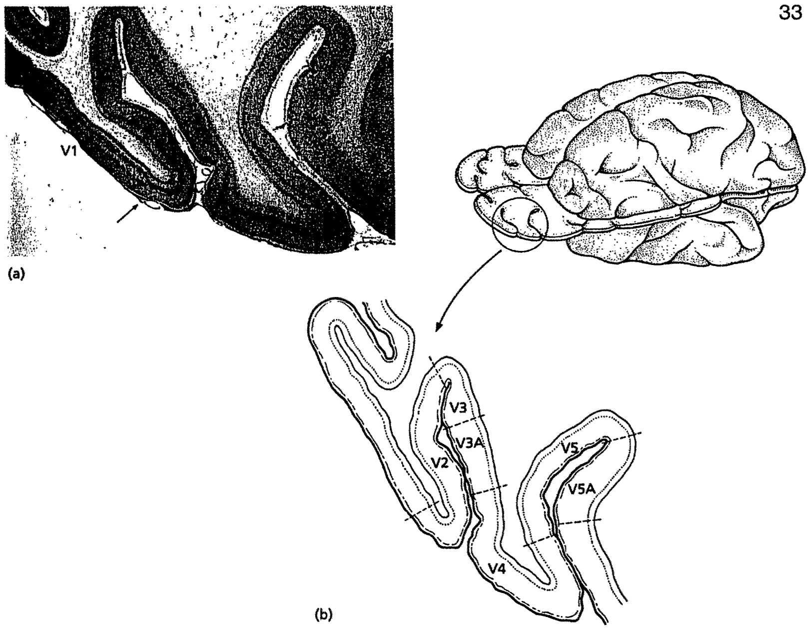

Fig. 7 Horizontal slices through the visual cortex of the macaque. (a) A slice stained to reveal the cytoarchitecture of the cortex. The stain reveals the striated area, which is labeled V1; the arrow marks the boundary between striate cortex and prestriate cortex, which exhibits a uniform cytoarchitecture. (b) A horizontal slice indicating boundaries between areas within prestriate cortex. Figure reprinted with permission from Zeki 1993, 91.

found a pattern of alternating thin and thick strips extending from the cortical surface to the white matter in V2. Application of cytochrome oxidase stain to V1, on the other hand, reveals a pattern of columns that appear as blobs in cortical layers 2 and 3. But cytochrome oxidase stain failed to differentiate areas V3, V4, and V5.

To determine the function of the different visual areas, Semir Zeki turned to single-cell recording so as to identify stimuli that would activate particular cells in these regions. In V5 he discovered cells that were sensitive to motion in different directions, but not affected by the color of the stimulus (Zeki 1974). Since V5 is buried in the superior temporal sulcus, Zeki had to pass his electrodes through V4 on the way to their targets in V5. Accordingly, he sought to identify stimuli that would drive these cells, and discovered that cells in V4 were especially sensitive to the wavelength of the stimulus (Zeki 1977).

These studies were necessarily limited to non-human primates (mostly the macaque and the owl monkey). Initially the only way to extend the results to humans was to analyze patterns of deficits found in patients. Clinical reports of patients with selective deficits in color processing or motion processing extend back to the 19th century and offered some support for the claim that these were processed by separate parallel pathways. But it is difficult to localize the site of the functional area through lesions since, in addition to the shortcomings noted above, lesions typically do not respect anatomical boundaries. The advent of neuroimaging techniques provided a second route for localizing distinct visual functions in humans, and Zeki et al. (1991) used PET to support the claim of distinct processing areas for color and motion in humans. Even with the advent of neuroimaging, permitting localization of processing areas in humans, research on brain visual areas remains fundamentally dependent on monkey research where more detailed neuroanatomical and neurophysiological studies have been possible.

The work we have reviewed so far supported a decomposition of visual processing at a relatively micro-level. In 1982 Leslie Ungerleider and Mortimer Mishkin advanced a decomposition at a more macro-level.[7] Relying largely on lesion studies in monkeys (the first study in this series was Pohl 1973), they differentiated two main routes for processing visual information (Figure 8). From striate cortex, one route proceeds ventrally into posterior temporal cortex along the inferior longitudinal fascilus to areas TEO and TE. Based on the fact that lesions in the posterior temporal area result in loss of pattern discrimination and that lesions in TE in particular result in failure to recognize previously presented objects, they assigned to this pathway the task of analyzing "the physical properties of a visual object (such as its size, color, texture and shape)" (Mishkin et al. 1983, 20). Since there are neural pathways between TE and limbic structures in the temporal lobe and ventral portions of the frontal cortex, Mishkin et al. further claimed that this area figures in assigning meaning to visual stimuli. One feature of neurons in the ventral pathway is that as one proceeds further, the cells have increasingly large receptive fields, suggesting that

---

[7] Ungerleider and Mishkin were not the first to attempt to decompose visual processing into separate pathways, although they were the first to distinguish two pathways *within* the cortex. In mammals, the primary distinction considered was that from the retina to the optic tectum or superior colliculus and that from the retina through the lateral geniculate to the primary visual cortex. Schneider (1969; see also Trevarthen 1968) proposed that the pathway through the superior colliculus figured in the ability to localize stimuli in visual space, whereas the pathway through the lateral geniculate to the visual cortex was employed in identifying the stimulus. Another differentiation between pathways was also proposed for the frog, where the pathway through the lateral geniculate has not yet evolved. Ingle (1973, 1982) differentiated a pathway from the retina to the optic tectum in the midbrain, which figures in catching prey and avoiding predators, and a pretectal system that guides locomotion around barriers.

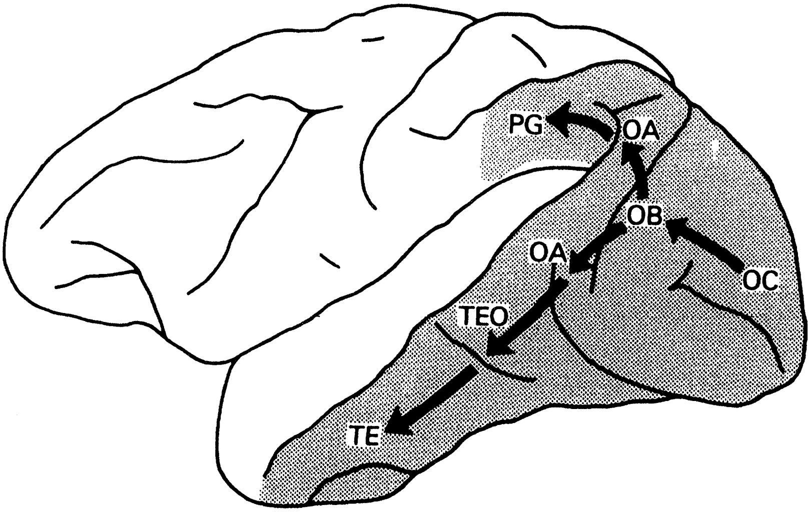

Fig. 8 Mishkin, Ungerleider, and Macko's (1983) portrayal of two pathways originating in primary visual cortex. The dorsal pathway from V1 up to the posterior parietal cortex (area PG) is interpreted as processing information about the location of a stimulus (where), and the pathway from V1 down into superior temporal cortex (areas TEO and TE) is interpreted as processing information about the identity of a stimulus (what).

they become responsible for recognizing objects independently of where they appear in the visual field. Performing a location-independent identification, they contend, requires discarding any location information:

> It is significant that by virtue of the extremely large visual receptive fields of inferior temporal neurons this area seems to provide the neural basis for the phenomenon of stimulus equivalence across retinal translation; i.e. the ability to recognize a stimulus as the same, regardless of its position in the visual field. But a necessary consequence of this mechanism for stimulus equivalence is that within the occipitotemporal pathway itself there is a loss of information about the visual location of the objects being identified. (Mishkin et al. 1983, 21)

The other route proceeds from prestriate cortex dorsally into the posterior parietal cortex. Lesions in the posterior parietal cortex in monkeys result in failure to be able to select a response location on the basis of a visual landmark, suggesting that this pathway figures in "perception of spatial relations among objects, and not in their intrinsic qualities" (ibid.). Further studies revealed that parietal area PG is responsive to tactile input as well as visual, suggesting that it processes multisensory spatial information. Unlike the ventral pathway, this parietal pathway is limited to input from the contralateral part of the visual field. Moreover, whereas the temporal pathway is primarily sensitive to foveal areas, the parietal pathway is equally sensitive to foveal and peripheral areas.

Livingstone and Hubel (1988) expanded Mishkin et al.'s proposal by linking their division into two pathways with an already discovered division between parvocellular and magnocellular pathways from the retina to V1 and the differentiation of areas in V1 made on the basis of stains for cytochrome oxidase. In brief, their comprehensive proposal was that different cells in the retina, which differ in their size, receptive fields, conductance velocities, and duration of response, project to cells in different layers of the lateral geniculate nucleus, and these in turn project to different cells in V1 and V2. In a further section of their paper they address human perception and interpret a wide range of behavioral studies in terms of the differentiation of two visual systems. They then ask the question "*why should the visual system be subdivided*?" Their answer is evolutionary:

> while the magno system is sensitive primarily to moving objects and carries information about the overall organization of the visual world, the parvo system seems to be important for analyzing the scene in much greater and more leisurely detail. These postulated functions would be consistent with the evolutionary relation of the two systems: the magno system seems to be more primitive than the parvo system and is possibly homologous to the entire visual system of nonprimate mammals. If so, it should not be surprising that the magno system is capable of what seem to be the essential functions of vision that an animal uses to navigate in its environment, catch prey, and avoid predators. The parvo system, which is well developed only in primates, seems to have added the ability to scrutinize in much more detail the shape, color, and surface properties of objects, creating the possibility of assigning multiple visual attributes to a single object and correlating its parts. (Livingstone and Hubel 1988, 32)

Although Livingstone and Hubel pointed to behavioral and perceptual studies with humans to support the decomposition, all the research supporting the two pathway analysis to that point had been done with various species of monkeys. Subsequently Haxby et al. (1991) used PET to extend the findings to humans, showing increased activation in parietal areas on a dot matching task that requires spatial processing and increased activation in temporal areas on a face matching task, providing evidence of a similar decomposition in humans. What is more interesting, from our perspective, is how the distinction, once made on the basis of neurobiological research, has been extended to purely behavioral research in psychology and to computational modeling. The computational modeling with neural networks, first by Rueckl, Cave, and Kosslyn (1989) and later by Jacobs, Jordan, and Barto (1991), identified computational efficiency and other beneficial consequences of employing separate processing for *what* and *where* information.

One behavioral domain in which the neurobiological distinction provided valuable heuristic guidance was in the speed-of-processing literature. Most research in that field pointed to a single "global processing-time coefficient that characterizes the speed of all of an individual's basic cognitive processes" (Hale and Jansen 1994, 388). Hale and Jansen, for example, employed a host of different behavioral measures (choice reaction time, letter classification, visual search, abstract matching, line-length discrimination, mental rotation, mental folding) and were unable to find any variable that generated a distinction. Two years later, however, Hale, following the decomposition suggested by the neural anatomy, designed tasks that would selectively require object identity or object location, and now found through a principal-components analysis two factors: a general speed of processing factor and a what/where factor (Hale, Chen, Myerson, and Simon 1996).

The account of two processing systems we have presented, following Mishkin and Ungerleider and Livingstone and Hubel, while offering a novel decomposition of the visual processing system that accounts for a great deal of the data, turns out to be too simple. We described above the research of Felleman and van Essen (1991) identifying 32 different processing areas in the macaque. Their work also shows that the different processing areas are highly interconnected. On average, each of these areas receives inputs from and sends outputs to 10 different areas; as a result, approximately 1/3 of the possible interconnections between areas are realized. Most of these are reciprocal. Moreover, there is a substantial amount of cross-talk between components that Livingstone and Hubel identified as parts of the different processing streams. This is shown, for example, by the fact that cells in the blob region of V1 are driven even in the absence of inputs from the parvocellular pathway through the LGN, suggesting cross-over from the magnocellular system. For another example, cells in V4, usually identified as part of the what system, also receive input from V3, part of the where system (for a review, see Merigan and Maunsell 1993). The evidence seems to suggest that one can identify different primary processing streams along the general lines suggested by Ungerleider and Mishkin, but that the distinctions are not absolute and when one pushes to a further level of resolution, the decomposition of processing mechanisms in visual cortex will be more complex than the simple what/where account suggests. (Figure 9 indicates both the functional specification of areas in the different processing streams and the more important interactions between the streams.)

As neurobiologists produce a more complex account of the processing pathways involved in vision, the question emerges as to whether the complications serve to advance cognitive or psychological theorizing. Milner and Goodale (1995) offer an intriguing idea that may prove useful in further cognitive theorizing. They note that the accounts of different processing pathways to date have focused on the input channels but that one could also focus on the output channels and the question of how processing in different parts of cortex is used. They contend that the output of the dorsal stream (Mishkin and Ungerleider's *where* system) figures rather directly in action whereas the ventral stream is used in higher primates for perception and higher cognitive tasks (Livingstone and Hubel 1988, made a similar point in the passage quoted earlier). Their focus is primarily on the dorsal stream, and they argue that the response of cells in the dorsal processing stream is influenced by behaviors undertaken as well as stimulus input. Milner and Goodale propose that much of the processing in this stream involves the transformation of spatial coordinates of visual inputs into coordinates required for various action systems (e.g., eyes or limbs), although some of it seems to involve coordination of information about visual shape with appropriate responses for grasping objects, etc. Different transformations are segregated in different parts of the parietal pathway. Further, different areas will utilize different components of the input processing which are most useful for the kind of action being controlled.

From this perspective, the two streams should not be construed as two largely separate routes running all the way through the visual system. Rather, the processing components early in the visual system take on responsibility for processing different kinds of information about visual scenes. Later areas in the system are dedicated to solving specific sorts of problems (e.g., coordinating limb movements), and extract from relevant earlier processes information relevant to those respective tasks. This is just the sort of decomposition that can guide further cognitive analysis of visual cognition and the development of comprehensive information processing models of visual processing.

The research on visual processing that we have recounted here is intended to indicate the way in which information about brain organization can be employed to guide development of a functional decom-

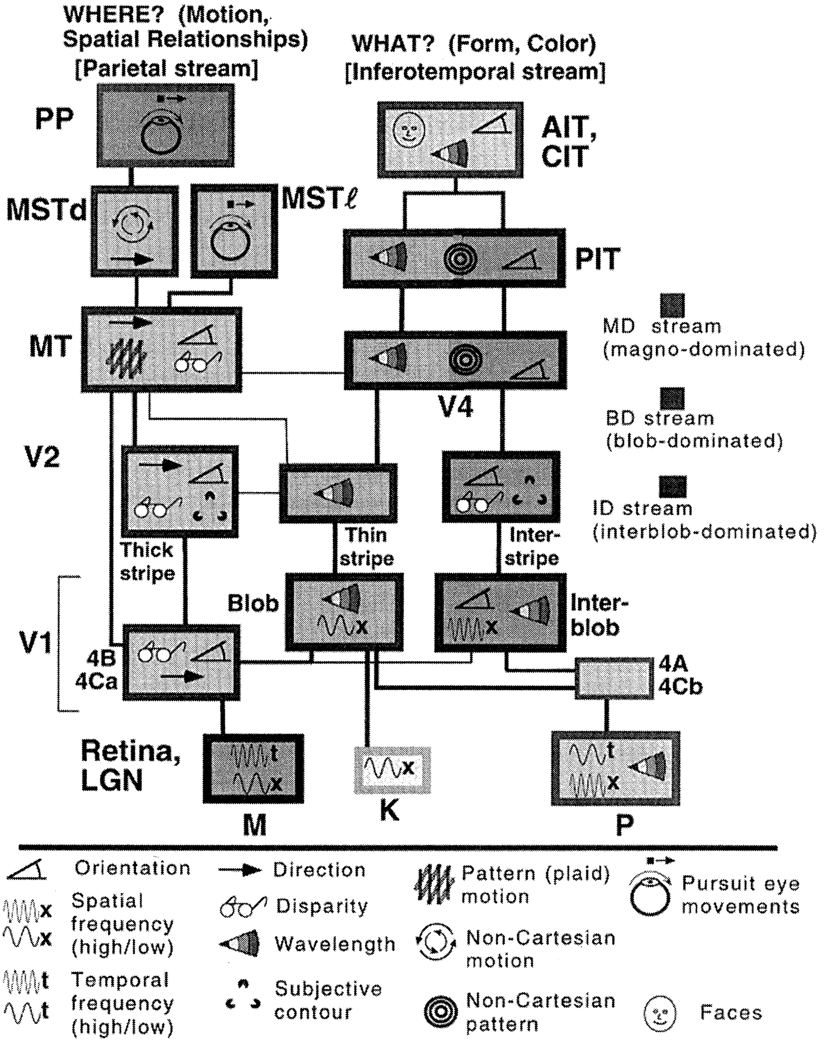

Fig. 9 The organization of multiple processing streams in the macaque, with boxes indicating different processing areas and icons representing the types of information to which cells in each area are responsive (van Essen and Gallant 1994, 3). M, P, and K refer to the magnocellular, parvocellular, and K streams identified in the retina and LGN. AIT, CIT, and PIT refer to the anterior, central, and posterior portions of the IT complex. The thickness of lines indicates the relative degree of connectedness of the various areas. Note that there are connections between the What and Where pathways at a variety of levels. Reprinted with permission of the Cell Press and David C. van Essen.

position of the cognitive system. Since the very manner in which brain parts are identified is in terms of what they do and contribute to cognitive life, they reveal the kinds of component tasks that contribute to overall cognitive performance. This does not depreciate the importance of analysis at the information processing level. It is still necessary to refine behavioral measures to determine precisely what tasks are performed by each component (Churchland and Sejnowski 1989), and to attempt to synthesize the components into an overall model that accounts for the information processing that occurs. Our goal has simply been to show how functional considerations get built into developing the structural taxonomy and how that taxonomy can in turn be a heuristic guide in developing information processing models. This project has not been impaired by multiple realization of psychological states; rather, it relies on the assumption that there is a common realization of mechanisms for processing visual information across species.

5. A Diagnosis of Why Multiple Realizability Looked Plausible. Our efforts have been aimed at discrediting the claim that cognitive states are multiply realized in any way that discounts the relevance of neuroscience to understanding cognitive processes. If we have been successful, then an interesting question arises as to why the claim that psychological states are multiply realized by different biological species has been so compelling in philosophy of mind. In this section we will offer a diagnosis.

One reason is almost certainly that those advancing the claim were not attending to the actual procedures by which neuroscientists identify brain areas, but relying on an intuitive view of what would be the same or different brain areas. Intuitively, a rat brain is sufficiently different from a human brain that identification of brain states in the two brain would be precluded. This is compounded by the fact that much of the philosophy of mind in the era when multiple realizability was first put forward intersected greatly with the philosophy of language, and advanced language ability has traditionally been associated only with human brains. From this perspective, it is natural to construe a brain of a rat as fundamentally different than the brain of a human. Moreover, as brain researchers began to show differences between brains of individuals in the same species, and changes in an individual's brain over time, it seemed clear that, if the brains were different, activity in these different brains would constitute different brain states. Thus, it seems obvious that there are differences in brain states between organisms even though there are many circumstances in which we would attribute the same cognitive state to them. Part of our effort has been to show that these differences between brains do not prevent neuroscientists from identifying common brain areas in different species. Accordingly, neuroscientists will attempt to identify the same brain areas and same brain processing in different members of the same species and across species despite whatever differences there are. In philosophical terms, neuroscientists treat the brain states as instances of the same type.

But there is another way we can look at this same issue. When comparing psychological states across different individuals, psychologists also tend to ignore differences and focus on commonalities. Likewise philosophers such as Putnam, who proposed comparing psychological states such as hunger across species as remote as humans and octopi, have abstracted away from differences. Yet, at anything less than a very abstract level, hunger is different in octopi than in humans. So, just as neuroscientists abstract away from differences between brains in identifying brain areas and brain processes, so do psychologists and philosophers in identifying psychological states. However, philosophers have not noted that the same degree of variability is tolerated by neuroscientists in identifying types of neural processes that they accept in identifying psychological types.

Thus, one diagnosis of what has made the multiple realizability claim as plausible as it has been is that researchers have employed different grains of analysis in identifying psychological states and brain states, using a coarse grain to identify psychological states and a fine grain to differentiate brain states. Having invoked different grains, it is relatively easy to make a case for multiple realization. But if the grain size is kept constant, then the claim that psychological states are in fact multiply realized looks far less plausible. One can adopt either a coarse or a fine grain, but as long as one uses a comparable grain on both the brain and mind side, the mapping between them will be correspondingly systematic. For example, one can adopt a relatively coarse grain, equating psychological states over different individuals or across species. If one employs the same grain, though, one will equate activity in brain areas across species, and one-to-one mapping is preserved (though perhaps further taxonomic refinement and/or delineation may be required). Conversely, one can adopt a very fine grain, and differentiate psychological states between individuals, or even in the same individual over time. If one similarly adopts a fine grain in analyzing the brain, then one is likely to map the psychological differences onto brain differences, and brain differences onto psychological differences. The latter claim may seem contentious, but an analogy from the computer domain may make this plausible. One often speaks of running the same program on different computer architectures, thereby generating a case of multiple realization. But in fact one has surreptitiously adopted a coarse grain in characterizing the program. If one looks carefully enough at performance measures, one will generally be able to identify differences in the way the program runs on different computers. Likewise, if one thinks one has found a case where differences in activity in brain areas do not produce differences in psychological states, the reason may be that one has not used a fine enough grain to analyze psychological states.

The appropriateness of grain size, of course, depends on the reason why one is making the comparison in the first place. This raises our final proposal as to why claims of multiple realizability have been so readily accepted: the lack of context, or frame of reference, in developing taxonomies for both psychology and neuroscience. Whenever one asks whether two items are the same or different, the question makes little sense unless one asks about sameness or difference with respect to some other consideration. Obviously, two items which are the same with respect to color may be different with respect to mass, for example. In fact, as Liebniz and others have pointed out, any two particulars will not only resemble infinitely, but differ infinitely. Why should it be any different if one is talking about brain states and psychological states? A human's psychological state and that of an octopus might well be counted as the same insofar as they are associated with some general feature (such as food-seeking behavior, in the case of hunger). But with respect to other considerations, a human psychological state may be considered different from that of an octopus, even if we limit the scope to mere behavior. Food-seeking behavior for the octopus is different from food-seeking behavior in the human if one is concerned about such things as how one seeks the food, what foods are sought, under what conditions, etc. This much seems simple and apparent, but the assertion that what we broadly call "hunger" is the same psychological state when instanced in humans and octopi has apparently been widely and easily accepted without specifying the context for judging sameness.

We have located the question of whether brain states are the same or not in the context of neuroscientific research, especially neuroanatomical and neurophysiological research devoted to brain mapping. We have tried to show that neuroscientists employ behavior and psychological function in setting the context for the identity of brain states, with the result that they identify areas and processes in brains across species, as well as in different brains within the same species. Not surprisingly, psychologists employ similar contexts for identifying psychological states, resulting in a one-to-one mapping between brain states and psychological states.

So far we have restricted ourselves to intra- and inter-species comparisons. Advocates of multiple realizability often appeal to comparisons of psychological states across systems made out of very different types of components—computers and biological organisms. One might contend that surely a state of a silicon-based computer is different from that of a carbon-based life form, so that if both could exemplify a psychological state, one would have an irrefutable example of multiple realizability. Perhaps, but we would caution that the case is not obvious. One would have to consider how one would analyze the functional components of a computer that genuinely realized psychological states. A computer that could exemplify sufficiently similar behavior to biological organisms to justify the imputation of psychological states is likely to be very different from the ones humans have designed to date, and the characterizations we would have to employ of its physical operations might turn out to be far more similar to those we use of brains than we currently expect. Such machines, for example, would likely have areas devoted to processing different sensory inputs and controlling motor outputs; potentially this might provide a basis for a common taxonomy of the physical processing states underlying psychological function. While it is possible one might discover genuine multiple realization of the same psychological state, this outcome requires argument. There are, after all, cases for which there are no alternative possible realizations for a set of functional properties (e.g., the functional properties of water can only be realized, given physical law as we understand it, by H_2O), and it is possible that mentality is one of these cases.

6. Conclusions. We have tried (i) to demonstrate that the claim that psychological states are multiple realized has not been demonstrated, at least within animal life forms, (ii) to show how denying multiple realization allows fruitful use of neuroscience in guiding the decomposition and understanding of cognitive systems, and (iii) to diagnose why multiple realizability has been so widely accepted. For some theorists, multiple realizability is an essential part of functionalism, the claim that psychological states are to be identified in terms of their functional roles. If multiple realizability were so closely tied to functionalism, a challenge to multiple realizability would also undercut functionalism. We do not, however, see any such serious consequences for functionalism. Even if one can identify psychological states with activity in brain areas, that does not render the functional characterization of psychological states any less important. The pattern of interaction may even be what fixes the identity of psychological states. Moreover, as we have emphasized, it is the functional characterization (the contribution to behavior) that guides the identification of brain areas. A thoroughgoing functionalism uses functional criteria to identify both psychological states and brain states, and can survive even if we jettison the multiple realizability thesis.

REFERENCES

Ballard, Dana H. (1991), "Animate Vision", *Artificial Intelligence* 48: 57–86.

Bechtel, William (in press), "From Imaging to Believing: Epistemic Issues in Generating Biological Data", in J. Maienschien and R. Creath (eds.), *Epistemology and Biology.* Cambridge: Cambridge University Press.

Broca, Paul (1861), "Remarque sur le Siêge de la Faculté Suivies d'une Observation d' Aphémie", *Bulletins de la Société Anatomique de Paris* 6: 343–357.

Brodmann, Korbinian (1909/1994), *Vergleichende Lokalisationslehre der Grosshirnrinde.* Leipzig: J. A. Barth. (English Translation by L. J. Garey (1994), *Brodmann's 'Localisation in the Cerebral Cortex'.* London: Smith-Gordon.)

Churchland, Patricia S., Vilayanur S. Ramachandran, and Terrence Sejnowski (1994), "A Critique of Pure Vision", in C. Koch and J. L. Davis (eds.), *Large-scale Neuronal Theories of the Brain.* Cambridge, MA: MIT Press.

Churchland, Patricia S. and Terrence Sejnowski (1989), "Neural Representation and Neural Computation", in L. Nadel, L. Cooper, P. Culicover, and R. M. Harnish (eds.), *Neural Connections, Mental Computations.* Cambridge, MA: MIT Press, 15–48.

Cowey, Alan (1964), "Projection of the Retina on the Striate and Prestiate Cortex of the Squirrel Monkey *Saimiri Sciureus*", *Journal of Neurophysiology* 27: 366–393.

Cragg, B. G. (1969), "The Topography of the Afferent Projections in the Circumstriate Visual Cortex of the Monkey Studied by the Nauta Method", *Vision Research* 9: 733–747.

Deacon, Terrence W. (1997), *The Symbolic Species: The Co-evolution of Language and the Brain.* New York: Norton.

Felleman, Daniel J. and David C. van Essen (1991), "A Distributed Hierarchical Processing in the Primate Cerebral Cortex", *Cerebral Cortex* 1: 1–47.

Ferrier, David (1886), *The Functions of the Brain.* New York: G. P. Putnam's Sons.

Finger, Stanley (1994), *The Origins of Neuroscience.* New York: Oxford University Press.

Fodor, Jerry A. (1987), *Psychosemantics: The Problem of Meaning in the Philosophy of Mind.* Cambridge, MA: MIT Press.

Fodor, Jerry A. and Zenon W. Pylyshyn (1981), "How Direct is Visual Perception? Some Reflections on Gibson's 'Ecological Approach' ", *Cognition* 9: 136–196.

Gennari, Francesco (1782), *De Peculiari Structura Cerebri Nonnullisque Eius Morbis.* Parma: Ex. Regio Typogratheo.

Gross, Charles G. (1998), *Brain, Vision, and Memory.* Cambridge, MA: MIT Press.

Hale, Sandy, Jing Chen, Joel Myerson, and Adam Simon (1996), *Behavioral Evidence for Brain-based Ability Factors in Visuospatial Information Processing.* Paper presented at the Psychonomics Society, November.

Hale, Sandy and Jennifer Jansen (1994), "Global Processing-Time Coefficients Characterize Individual and Group Differences in Cognitive Speed", *Psychological Science* 5: 384–389.

Haxby, James V., Cheryl L. Grady, Barry Horwitz, Leslie G. Ungerleider, Mortimer Mishkin, Richard E. Carson, Peter Herscovitch, Mark B. Schapiro, and Stanley I. Rapoport (1991), "Dissociation of Object and Spatial Visual Processing Pathways in Human Extrastriate Cortex", *National Academy of Sciences, USA* 88: 1621–1625.

Henschen, Salomon Eberhard (1893), "On the Visual Path and Centre", *Brain* 16: 170–180.

———. (1903), "La projection de la rétine sur la corticalité calcarine", *La Semaine Médicale* 23: 125–127.

Hubel, David H. and Torsten N. Wiesel (1962), "Receptive Fields, Binocular Interaction and Functional Architecture in the Cat's Visual Cortex", *Journal of Physiology (London)* 160: 106–154.

———. (1965), "Receptive Fields and Functional Architecture in Two Non-striate Visual Areas (18 and 19) of the Cat", *Journal of Neurophysiology* 28: 229–289.

———. (1968), "Receptive Fields and Functional Architecture of Monkey Striate Cortex", *Journal of Physiology (London)* 195: 215–243.

Ingle, David J. (1973), "Two Visual Systems in the Frog", *Science* 181: 1053–1055.

———. (1982), "Organization of Visuomotor Behaviors in Vertebrates", in D. J. Ingle, M. A. Goodale, and R. J. W. Mansfield (eds.), *Analysis of Visual Behavior.* Cambridge, MA: MIT Press, 67–109.

Jacobs, Robert A., Michael I. Jordan, and Andrew G. Barto (1991), "Task Decomposition through Competition in a Modular Connectionist Architecture: The What and Where Vision Tasks", *Cognitive Science* 15: 219–250.

Lashley, Karl S. and George Clark, (1946), "The Cytoarchitecture of the Cerebral Cortex of *Ateles:* A Critical Examination of Architectonic Studies", *Journal of Comparative Neurology* 85: 223–305.

Livingstone, Margaret and David Hubel, (1988), "Segregation of Form, Color, Movement, and Depth: Anatomy, Physiology, and Perception", *Science* 240: 740–749.

Marr, David (1982), *Vision.* San Francisco: Freeman.

Merigan, William H. and John H. R. Maunsell (1993), "How Parallel are the Primate Visual Pathways?", *Annual Review of Neuroscience* 16: 369–402.

Milner, A. David and Melvyn G. Goodale (1995), *The Visual Brain in Action.* Oxford: Oxford University Press.

Mishkin, Mortimer, Leslie G. Ungerleider, and Kathleen. A. Macko, (1983), "Object Vision and Spatial Vision: Two Cortical Pathways", *Trends in Neurosciences* 6: 414–417.

Petersen, Stephen E., Peter T. Fox, Michael I. Posner, Mark Mintum, and Marcus E. Raichle, (1989), "Positron Emission Tomographic Studies of the Processing of Single Words", *Journal of Cognitive Neuroscience* 1: 153–170.

Putnam, Hilary (1967), "Psychological Predicates", in W. H. Capitan and D. D. Merrill (eds.), *Art, Mind, and Religion.* Pittsburgh: University of Pittsburgh Press, 37–48.

Pohl, W. (1973), "Dissociation of Spatial Discrimination Deficits Following Frontal and Parietal lesions in Monkeys", *Journal of Comparative Physiological Psychology* 82: 227–239.

Richardson, Robert C. (1979), "Functionalism and Reductionism", *Philosophy of Science* 46: 533–558.

Rueckl, J. G., K. R. Cave, and Stephen M. Kosslyn (1989), "Why are 'What' and 'Where' Processed by Separate Cortical Visual Systems? A Computational Investigation", *Journal of Cognitive Neuroscience* 1: 171–186.

Schneider, Gerald E. (1969), "Two Visual Systems: Brain Mechanisms for Localization and Discrimination are Dissociated by Tectal and Cortical Lesions", *Science* 163: 895–902.

Talairach Jean and Pierre Tournoux (1988), *Co-planar Stereotaxic Atlas of the Human Brain.* New York: Thieme Medical Publishers, Inc.

Trevarthen, Colwyn B. (1968), "Two Mechanisms of Vision in Primates", *Psychologische Forschung* 31: 299–337.

Ungerleider, Leslie G. and Mortimer Mishkin (1982), "Two Cortical Visual Systems", in D. J. Ingle, M. A. Goodale, and R. J. W. Mansfield (eds.), *Analysis of Visual Behavior.* Cambridge, MA: MIT Press, 549–586.

van Essen, David C., Charles H. Anderson, and Daniel J. Felleman (1992), "Information Processing in the Primate Visual System: an Integrated Systems Perspective", *Science* 255: 419–423.

van Essen, David C., Heather A. Drury, Sarang Joshi, and Michael I. Miller (1998), "Functional and Structural Mapping of Human Cerebral Cortex: Solutions Are in the Surfaces", *Proceedings of the National Academy of Sciences, USA* 95: 788–795.

van Essen, David C. and Jack L. Gallant (1994), "Neural Mechanisms of Form and Motion Processing in the Primate Visual System", *Neuron* 13: 1–10.

von Bonin, Gerhardt and Percival Bailey (1951), *The Isocortex of Man.* Urbana: University of Illinois Press.

Wong-Riley, Margaret T. T. (1979), "Changes in the Visual System of Monocularly Sutured or Enucleated Cats Demonstrable with Cytochrome Oxidase Histochemistry", *Brain Research* 171: 11–28.

Zeki, Semir M. (1969), "Representation of Central Visual Fields in Prestriate Cortex of Monkey", *Brain Research* 14: 271–291.

———. (1974), "Functional Organization of a Visual Area in the Posterior Bank of the Superior Temporal Sulcus of the Rhesus Monkey", *Journal of Physiology* 236: 549–573.

———. (1977), "Colour Coding in the Superior Temporal Sulcus of the Rhesus Monkey Visual Cortex", *Proceedings of the Royal Society of London* B197: 195–223.

———. (1993), *A Vision of the Brain.* Oxford: Blackwell Scientific.

Zeki, Semir M., J. D. G. Watson, C. J. Lueck, Karl J. Friston, C. Kennard, Richard S. J. Frackowiak (1991), "A Direct Demonstration of Functional Specialization in Human Visual Cortex", *Journal of Neuroscience* 11: 641–649.