

APPENDIX

BRAIN ARCHITECTURE

When you look at three-dimensional views of the human brain, there is an obvious architectural arrangement that you grasp with the naked eye. The overall pattern is similar from brain to brain, and certain components show up in every brain in the same position. Their relationship is like that of the components of our faces—eyes, mouth, nose. Their exact shape and size are somewhat different in each individual, but the range of variation is limited. There are no human faces in which the eyes are square or in which an eye is larger than the nose or mouth, and symmetry is by and large respected. Comparable restrictions apply to the relative positions of the elements. Like our faces, our brains are extremely similar in terms of the grammatical rules according to which the parts are arranged in space. And yet brains are quite individual. Each brain is unique.

Another aspect of the architecture that is relevant to the ideas in this book, however, is invisible to the naked eye. Lying beneath the surface, it consists of a massive cable work made up of *axons*—the fibers that interconnect neurons. The brain has billions of neurons (about 1011), and those neurons make trillions of connections among themselves (about 1015). Nonetheless, the connections are made according to *patterns*, and not every neuron connects to every other neuron. On the contrary, their meshwork is highly selective. Seen from afar it constitutes a wiring diagram, or many wiring diagrams, depending on the sector of the brain.

Understanding the wiring diagrams is one road to understanding what the brain does and how. But it is not easy because the wiring dia-

grams undergo considerable changes during development and beyond. We are born with certain connection patterns, put into place under the instruction of our genes. These connections were already influenced by several environmental factors in the womb. After birth individual experiences in unique environments get to work on that first connection pattern, pruning it away, making certain connections strong and others weak, thickening or thinning the cables in the network, under the influence of our own activities. Learning and creating memory are simply the process of chiseling, modeling, shaping, doing, and redoing our individual brain wiring diagrams. The process that began at birth continues until death makes us part with life, or some time before, if Alzheimer's disease disrupts the process.

How does one uncover the design of the wiring diagrams? Until quite recently, research on this problem required brain specimens, largely postmortem material from either humans or experimental animals. Samples of brain tissue would be fixed and stained with identifiable dyes, and very thin slices of tissue could be analyzed under the microscope. There is a venerable tradition of such studies in experimental neuroanatomy, and they have yielded most of the knowledge we have today about the brain's networking. But our knowledge of neuroanatomy remains embarrassingly incomplete, so there is an urgent need for such studies to continue, making use of considerable progress in the available stains and in the power of modern microscopes.

Recently, new possibilities have opened with the use of magnetic resonance methods in living humans. Noninvasive methods such as diffusion imaging are allowing us a first glimpse of in vivo human connection networks. Although the techniques are still far from satisfactory, they promise to yield fascinating revelations.

How do the billions of neurons inside a human brain and the trillions of synapses they form manage to produce not just the actions that constitute behaviors but also minds—minds of which each owner can be con-

scious and minds that can give rise to cultures? To say that so many neurons and synapses do the job by massive interactivity and by the ensuing complexity is not a good answer. Interactivity and complexity must surely be present, but interactivity and complexity are not amorphous. They derive from the varied designs of local circuit arrangements and the even more varied ways in which such circuits create regions and regions become affiliated in systems. How each region is made, internally, determines its function. A region's location in the overall architecture is important too, because its place in the global plan determines its partners in the system—the regions that talk to a particular region and to which it talks back. To make matters even more complicated, the opposite is also true: to a certain extent the partners that it interacts with determine where its place is going to be. But before we go any further, we should give a brief account of the materials used to construct brain architecture.

BRICKS AND MORTAR

The mind-making brain is made of neural tissue, and neural tissue, like any other living tissue, is made of cells. The principal type of brain cell is the *neuron*, and for reasons that I alluded to in Chapters 1, 2, and 3, the neuron is a distinctive cell in the universe of biology. Neurons and their axons are embedded—*suspended* might be a better term—in a scaffolding made up of another type of brain cell, the *glial cell*. Besides providing neurons with physical support, glial cells also provide part of their nourishment. Neurons cannot survive without glial cells, but everything indicates that neurons are the critical brain unit as far as behavior and mind are concerned.

When neurons use their axons and send messages to muscular fibers, they can produce movements; and when neurons are active within very complex networks of map-making regions, the result is images, the main currency of mental activity. Glial cells, as far as we know, do nothing of the sort, although their full contribution to the operation of

neurons has not been fully elucidated. On a somber note, glial cells are the origin of the most deadly brain tumors, the gliomas, for which there is no cure to date. Even worse, for reasons that are entirely unclear, the incidence of malignant gliomas is rising worldwide, unlike practically all other malignancies. The other common origin of brain tumors is the cells of the meninges—the skinlike membranes that cover brain tissue. Meningiomas tend to be benign, although, by dint of their location and unchecked growth, they can compromise brain function seriously and are anything but innocent.

Each neuron has three main anatomical elements: (1) the *cell body*, which is the cell's powerhouse and includes the cell nucleus and organelles such as mitochondria (the neuron's genome, its complement of governing genes, is located within the nucleus, although DNA is also to be found within mitochondria); (2) the main output fiber, known as the *axon*, which arises from the cell body; and (3) input fibers, known as *dendrites* that stick out from the cell body a bit like antlers. Neurons are connected to one another via a border area called the *synapse*. In most synapses the axon of one neuron makes chemical contact with the dendrites of another.

Neurons can be active (firing) or inactive (not firing), “on” or “off.” The firing consists of producing an electrochemical signal that crosses the border to another neuron, at the synapse, and makes that other neuron fire too, provided the signal meets the requirements of the other neuron to fire. The electrochemical signal travels from the neuron's body down the axon. The synaptic border is located between the end of an axon and the beginning of another neuron, generally at the dendrite. There are several minor variations and exceptions to this standard description, and different kinds of neurons vary in shape and size; but this outline is acceptable as far as the big picture goes. Each neuron is so small that one needs the major amplification of a microscope to see it, and in order to see a synapse one needs an even more powerful

microscope. Still, smallness is relative, entirely in the amplified eye of the beholder. Compared to the molecules that make them up, neurons are truly gigantic creatures.

When neurons “fire,” the electric current known as the action potential is propagated away from the cell body and down the axon. The process is very fast—it takes only a handful of milliseconds, which should give an idea of the remarkably different time scales of brain and mind processes. We need *hundreds* of milliseconds to become conscious of a pattern presented to our eyes. We experience feelings in a time scale of *seconds*, that is *thousands* of milliseconds, and *minutes*.

When the firing current arrives at a synapse, it triggers the release of chemicals known as neurotransmitters (glutamate is an example) in the space between two cells, the synaptic cleft. In an excitatory neuron, the cooperative interaction of many other neurons whose synapses are adjacent and that release (or do not) their own transmitter signals, determines whether the next neuron will fire, that is, whether it will produce its own action potential, which will lead to its own neurotransmitter release, and so forth.

Synapses can be strong or weak. Synaptic strength determines whether and how easily impulses will continue to travel into the next neuron. In an excitatory neuron, a strong synapse facilitates impulse travel, while a weak synapse impedes or blocks it.

One critical aspect of learning is the strengthening of a synapse. Strength is translated into ease of firing and thus ease of activation of the neurons downstream. Memory depends on this operation. Our understanding of the neural basis of memory at neuron level can be traced to the seminal ideas of Donald Hebb, who, in the mid-twentieth century, first raised the possibility that learning depended on the strengthening of synapses and the facilitation of the firing of subsequent neurons. He did so on a purely theoretical basis, but his hypothesis was subsequently proven correct. In the past few decades the understanding of learning has deepened to the level of molecular mechanisms and gene expression.

On average each neuron talks to relatively few others, not to most, and never to all. In fact, many neurons talk only to neurons that are close by, within relatively local circuits; others, even if their axons travel for several centimeters, make contact with only a small number of other neurons. Still, depending on where the neuron sits in the overall architecture, it may have more or fewer partners.

The billions of neurons are organized in circuits. Some are very small microcircuits, truly local operations invisible to the naked eye. When many microcircuits are placed together, however, they form a region, with a certain architecture.

The elementary regional architectures come in two varieties: the *nucleus* variety and the *cerebral cortex patch* variety. In a patch of cerebral cortex, the neurons are displayed on two-dimensional surface sheaths stacked in layers. Many of these layers have a fine topographical organization. This is ideal for detailed mapping. In a nucleus of neurons (not to be confused with the cell nucleus inside each neuron), the neurons are usually displayed like grapes inside a bowl, but there are partial exceptions to this rule. The geniculate nuclei and the collicular nuclei, for example, have two-dimensional, curvy layers. Several nuclei have topographical organization as well, which suggests that they can generate coarse maps.

Nuclei contain “know-how.” Their circuitry embodies knowledge about how to act or what to do when certain messages make the nucleus active. Because of this dispositional know-how, nucleus activity is indispensable for the management of life in species with smaller brains, those with little or no cerebral cortex and limited map-making abilities. But nuclei are also indispensable for managing life in brains such as ours, where they are responsible for basic management—metabolism, visceral responses, emotions, sexual activity, feelings, and aspects of consciousness. The governance of endocrine and immune systems depends on nuclei, and so does affective life. But in humans, a good part of the operation of nuclei is under the influence of the mind, and that means largely, though not entirely, the influence of the cerebral cortex.

Importantly, the separate regions defined by nuclei and by cerebral cortex patches are interconnected. They form, in turn, larger- and larger-scale circuits. Numerous patches of cerebral cortex come to be wired together interactively, but each patch is also wired to subcortical nuclei. Sometimes the patch of cortex is a recipient of signals from a nucleus, or sometimes it is a sender of signals; sometimes it is both recipient and sender. The interactions are especially significant in relation to the myriad nuclei of the thalamus (regarding which the connections to the cerebral cortex tend to be two-way) and in relation to the basal ganglia (regarding which the connections tend to be either downward from cortex or up toward it, but not both).

In sum, neuron circuits constitute cortical regions, if they are arranged in sheaths placed in parallel layers like those of a cake or constitute nuclei, if they are grouped in nonlayered arrangements (but note the exceptions mentioned earlier). Both cortical regions and nuclei are interconnected by axon “projections” to form *systems* and, at gradually higher levels of complexity, *systems of systems*. When bunches of axon projections are large enough to be seen by the naked eye, they are called “pathways.” In terms of scale, all neurons and local circuits are microscopic, while all cortical regions, most nuclei, and all systems of systems are macroscopic.

If neurons are the bricks, what is the brain’s equivalent of mortar? Quite simply, it is the large number of *glial* cells that I introduced as the scaffolding for the neurons everywhere in the brain. The myelin sheaths that wrap around fast-conducting axons are also glial. They provide protection and insulation for those axons, conforming yet again to the role of mortar. Glial cells are very different from neurons in that they do not have axons and dendrites and do not transmit signals over long distances. In other words, glial cells are not about the other cells in an organism, and their role is neither to regulate nor to represent other cells. The imitative role of neurons does not apply to glial cells. But the roles that glial cells play go beyond mere shelving for neurons. Glial cells intervene in the nutrition of neurons by holding and delivering

energy products, for example, and, as suggested earlier, their influence may actually go deeper.

MORE ON THE LARGE-SCALE ARCHITECTURE

The nervous system has central and peripheral divisions. The main component of the *central nervous system* is the *cerebrum*, which is made up of two *cerebral hemispheres*, left and right, joined by the *corpus callosum*. A facetious tale says that the corpus callosum was invented by nature to keep the cerebral hemispheres from sagging. But we know that this thick collection of nerve fibers connects the left and right halves, in both directions, and performs an important integrative role.

The cerebral hemispheres are covered by the cerebral cortex, which is organized in lobes (*occipital, parietal, temporal, and frontal*) and includes a region known as the *cingulate cortex*, visible only on the internal (mesial) surface. Two regions of the cerebral cortex that are not visible at all when one inspects the surface of the cerebellum are the *insular cortex*, buried underneath the frontal and parietal regions; and the *hippocampus*, a special cortical structure hidden in the temporal lobe.

Underneath the cerebral cortex, the central nervous system also includes deep conglomerates of nuclei such as the *basal ganglia*, the *basal forebrain*, the *amygdala*, and the *diencephalon* (a combination of the *thalamus* and the *hypothalamus*). The cerebrum is joined to the spinal cord by the *brain stem*, behind which the *cerebellum* is located with its two hemispheres. Although the hypothalamus is usually mentioned together with the thalamus to constitute the diencephalon, in reality the hypothalamus is functionally closer to the brain stem, with which it shares the most critical aspects of life regulation.

The central nervous system is connected to every point of the body by bundles of axons originating in neurons. (The bundles are known as nerves.) The sum total of all nerves connecting the central nervous system with the periphery and vice versa constitutes the *peripheral nervous system*. Nerves transmit impulses from brain to body and from body to




Figure A.1: The large-scale architecture of the human brain shown in a three-dimensional reconstruction of magnetic resonance data. The lateral (external) views of both right and left cerebral hemispheres are shown on the left panels; the medial (internal) views are shown on the right. The white curved structure in the right panels corresponds to the corpus callosum.

brain. One of the oldest and most important sectors of the peripheral nervous system is the *autonomic nervous system*, so called because its operation is largely outside our volitional control. The components of the autonomic nervous system include the *sympathetic*, *parasympathetic*, and *enteric* systems. The system plays a critical role in life regulation and in emotions and feelings. The brain and the body are also interconnected by chemical molecules such as hormones, which travel in the bloodstream. The ones that travel from brain to body originate in nuclei such as those in the hypothalamus. But chemical molecules also travel in the opposite direction and influence neurons directly at locations such as the area postrema, where the protective blood-brain barrier is missing.




Figure A.2: The panels on the left depict three-dimensional reconstructions of the human brain seen from lateral and medial perspectives (top and bottom, respectively).

The panels on the right depict three sections of the brain volume. The sections were obtained along the lines marked *a*, *b*, and *c*. The sections reveal a number of important brain structures located under the surface: 1 = basal ganglia; 2 = basal forebrain; 3 = claustrum; 4 = insular cortex; 5 = hypothalamus; 6 = thalamus; 7 = amygdala; 8 = hippocampus. The cerebral cortex covers the entire surface of the cerebral hemispheres, including the depth of every sulcus. In the sections, the cerebral cortex appears as a dark rim easily distinguishable from the lighter white matter underneath. The black areas at the center of the sections correspond to the lateral ventricles.

(The blood-brain barrier is a protective shield against certain molecules circulating in the bloodstream.) The area postrema is located in the brain stem, very close to important life-regulating structures such as the parabrachial and periaqueductal nuclei.

When one slices the central nervous system in any direction and looks at the cross-section, one notices a difference between dark and pale sectors. The dark sectors are known as the *gray matter* (although they are more brown than gray), and the pale sectors are known as the *white matter* (which is more tan than white). The gray matter gets its darker hue from the tight packing of many neuron cell bodies; the white matter gets its lighter appearance from the insulating sheaths of the axons that emanate from the cell bodies located in the gray matter. As noted, the insulation is made of myelin and speeds up the conduction of electric current in the axons. Myelin insulation and fast conduction of signals are hallmarks of evolutionarily modern axons. Unmyelinated fibers are quite slow and of older vintage.

The gray matter comes in two varieties. By and large the layered variety is found in the *cerebral cortex*, which envelops the cerebral hemispheres, and in the *cerebellar cortex*, which envelops the cerebellum. The nonlayered variety is made of *nuclei*, the main examples of which were listed earlier: the *basal ganglia* (located in the depth of each cerebral hemisphere and made up of three large nuclei, the caudate, the putamen, and the pallidum); the *amygdala*, a single and sizable lump located in the depth of each temporal lobe; and several aggregations of smaller nuclei that form the *thalamus*, the *hypothalamus*, and the gray sectors of the *brain stem*.

The cerebral cortex is the cerebrum's mantle, covering the surfaces of each cerebral hemisphere, including those that are located in the depth of fissures and sulci, the crevices that give the brain its unique folded appearance. The thickness of the cortex is about three millimeters, and the layers are parallel to one another and to the brain's surface. The evolutionarily modern part of the cerebral cortex is the *neocortex*. The main divisions of the cerebral cortex are designated as lobes:

frontal, temporal, parietal, and occipital. All other gray structures (the various nuclei mentioned earlier and the cerebellum) are subcortical.

In the text I often refer to *early sensory cortices* or to *association cortices* or even to *higher-order association cortices*. The designation *early* has no time connotation at all; it refers to the position occupied by a region in space, along a sensory processing chain. Early sensory cortices are those located near and around the point of entry of peripheral sensory pathways into the cerebral cortex—for example, the point of entry for vision or hearing or touch signals. The early regions tend to be organized concentrically. They play a critical role in producing detailed maps using the signals brought in by the sensory pathways.

The association cortices, as the name implies, interrelate signals arising from the early cortices. They are located everywhere in the cerebral cortex where there are no early sensory cortices or motor cortices. They are organized hierarchically, and the ones higher up in the chain are usually known as higher-order association cortices. The prefrontal cortices and the anterior temporal cortices are examples of higher-order association cortices.

The various regions of the cerebral cortex are traditionally identified by numbers corresponding to the distinctive architectural design of its neuron arrangements, which is known as cytoarchitectonics. The best-known system for numbering the regions was proposed by Brodmann a century ago, and it remains a useful tool today. The Brodmann numbers have nothing whatsoever to do with the area's size or functional importance.

THE IMPORTANCE OF LOCATION

The internal anatomical structure of a brain region is an important determinant of its function. Where a given brain region is located within the three-dimensional volume of a brain is another important determinant. Placement in the global brain volume and internal anatomical structure are largely consequences of evolution, but they are

also influenced by individual development. Individual experience shapes the circuitry, and although this influence is most marked at the microcircuitry level, it is inevitably felt at the macroanatomic level as well.

The evolutionary vintage of nuclei is old, a throwback to a time in the history of life when whole brains were little more than chains of ganglia resembling beads in a rosary. A ganglion is, in essence, an individual nucleus before being evolutionarily incorporated into a brain mass. The brains of the nematodes I mentioned in Chapter 2 consist of chains of ganglia.

The location of nuclei within the brain's whole volume is fairly low, always below the mantle provided by the cerebral cortex. They sit in the brain stem, the hypothalamus and thalamus, the basal ganglia, and the basal forebrain (whose extension includes the collection of nuclei known as the amygdalae). Banished as they are from the prime cortical estate, they still have an evolutionary pecking order. The older they are, historically speaking, the closer they are to the brain's midline. And because everything in the brain has two halves, left and right with a dividing median, it so happens that very old nuclei sit looking at their twin on the other side of the midline. This is the case with the brain-stem nuclei that are so vital for life regulation, and for consciousness. In the case of somewhat more modern nuclei—say, the amygdala—the left and right exemplars are more independent and clearly separate from each other.

The cerebral cortices are evolutionarily more recent than the nuclei. They are all distinguished by their two-dimensional sheathlike structure, which confers upon some of them detailed map-making abilities. But the number of layers in a cortex varies from a mere three (for old-vintage cortices) to six (for more recent vintages). The complexity of the circuitry within and across those layers varies as well. The overall location in the whole brain volume is also functionally telling. In general, very modern cortices occur at and around the point at which the major sensory pathways—e.g., auditory, visual, somatosensory—enter

the cerebral cortex mantle and are thus connected with sensory processing and map-making. In other words, they belong to the “early sensory cortex” club.

Motor cortices also have varied vintages. Some motor cortices are quite old and small, again located at the midline in the anterior cingulate and supplementary motor regions, clearly visible on the internal (or medial) surface of each cerebral hemisphere. Other motor cortices are modern and structurally sophisticated and occupy a sizable territory on the external surface of the brain (the lateral surface).

What a given region ends up contributing to the overall business of the brain depends significantly on its partners: which talks to the region and which is talked back to, specifically, which regions project their neurons to region *X* (thus modifying the state of region *X*) and which regions receive projections from region *X* (thus being modified by its output). A lot depends on where region *X* is located within the network. Whether region *X* has map-making abilities is another important factor in its functional role.

Mind and behavior are the moment-to-moment results of the operation of galaxies of nuclei and cortical parcels articulated by convergent and divergent neural projections. If the galaxies are well organized and work harmoniously, the owner makes poetry. If not, madness ensues.

AT THE INTERFACES BETWEEN THE BRAIN AND THE WORLD

Two kinds of neural structures are located at the border between the brain and the world. One points *inward*, the other *outward*. The first neural structure is made up of the sensory receptors of the body’s periphery—the retina, the cochlea in the inner ear, the nerve terminals in our skin, and so forth. These receptors do not receive neuron projections from the outside, at least not naturally, although neuronlike electrical inputs from prosthetic implants are changing this situation. They receive *physical stimuli* instead—light, vibration, mechanical contact.

Sensory receptors initiate a chain of signals from the body’s border to inside the brain, across multiple hierarchies of neuron circuits that penetrate deeply into the brain territories. But they don’t just move up like water in a pipe system. At every new station they undergo processing and a transformation. In addition, they tend to send signals back to where the inbound projection chains started. These understudied features of brain architecture probably have great significance for certain aspects of consciousness.

The other kind of border point occurs where the *outward* projections from the brain end and the environment begins. The chains of signals arise within the brain but end up either releasing chemical molecules into the atmosphere or connecting to muscular fibers in the body. The latter enables us to move and speak, and that is where the principal outward chains terminate. Beyond the muscle fibers there lies direct movement in space. In earlier stages of evolution, the release of chemical molecules at the membrane or skin border played important roles in the life of an organism. It was an important means of action. In humans, this facet remains understudied, although the release of pheromones is not in doubt.

One may conceptualize the brain as a progressive elaboration of what began as a simple reflex arc: neuron NEU senses object OB and signals to neuron ZADIG, which projects to muscle fiber MUSC and causes movement. Later in evolution a neuron would be added to the reflex circuit, in between NEU and ZADIG. This is an *interneuron*, and let us call it INT; it behaves such that the response of neuron ZADIG is no longer automatic. Neuron ZADIG responds, for example, only if neuron NEU fires all its guns upon it and not if neuron ZADIG receives a weaker message; a critical part of the decision is left in the hands of the interneuron INT.

A major aspect of brain evolution has consisted of adding the equivalent of interneurons at every level of brain circuitry—a slew of such equivalents, in fact. The largest such equivalents, located in the cerebral cortex, might well be called *interregions*. They become sandwiched

between other regions, for the good and obvious purpose of modulating simple responses to varied stimuli and making the responses less simple, less automatic.

On the path to making the modulation more subtle and sophisticated, the brain developed systems that map stimuli in such detail that the ultimate consequence was images and mind. Eventually the brain added a self process to those minds, and that permitted the creation of novel responses. Finally, in humans, when such conscious minds were organized in collectives of like beings, the creation of cultures became possible along with their attending external artifacts. In turn, cultures have influenced the operation of brains over generations and eventually influenced the evolution of the human brain.

The brain is a system of systems. Each system is composed of an elaborate interconnection of small but macroscopic cortical regions and subcortical nuclei, which are made of microscopic local circuits, which are made of neurons, all of which are connected by synapses.

What neurons do depends on the local assembly of neurons to which they belong; what systems end up doing depends on how local assemblies influence other assemblies within an interconnected architecture; finally, whatever each assembly contributes to the function of the system to which it belongs depends on its place in that system.

A NOTE ON THE MIND-BRAIN EQUIVALENCE HYPOTHESIS

The perspective adopted in this book contains a hypothesis that is not universally liked, let alone accepted—namely, the idea that mental states and brain states are essentially equivalent. The reasons for the reluctance in endorsing such a hypothesis deserve a hearing.

In the physical world, of which the brain is unequivocally a part, equivalence and identity are defined by physical attributes such as mass, dimensions, movement, charge, and so forth. Those who reject the identity between physical states and mental states suggest that while a brain map that corresponds to a particular physical object can be dis-

cussed in physical terms, it would be absurd to discuss the respective mental pattern in physical terms. The reason given is that to date science has not been able to determine the physical attributes of mental patterns, and if science cannot do so, then the mental cannot be identified with the physical. I fear, however, that this reasoning may not be sound. Let me explain why I think so.

First, we need to consider how we determine that nonmental states are physical. In the case of objects out in the world, we proceed by perceiving them with our peripheral sensory probes and by using varied instruments to execute measurements. In the case of mental events, however, we cannot do the same. This is not because mental events are not equivalent to neural states but because, given their place of occurrence—the interior of the brain—mental states are simply not available for measurement. In fact, mental events can be perceived only by part of the very same process that includes them—the mind, that is. The situation is unfortunate but says nothing whatsoever about the physicality of the mind or lack thereof. The situation does impose major qualifications on the intuitions that can emerge from it, however, and it is thus prudent to doubt the traditional view that asserts that mental states *cannot* be equivalent to physical states. It is unreasonable to endorse such a view purely on the basis of introspective observations. The personal perspective should be used and enjoyed for what it gives us directly: experience that can be made conscious, and that can help guide our life, provided extensive reflective analysis conducted offline—which includes scientific scrutiny—validates its counsel.

The fact that neural maps and the corresponding images are found *inside* the brain, accessible only to the brain's owner, is a hurdle. But where else would the maps/images be found but within a private, secluded sector of the brain, given that they are formed inside the brain to begin with? What would be surprising would be to find them outside the brain, given that brain anatomy is not designed to externalize them.

For the time being, the mental state/brain state equivalence should be regarded as a useful hypothesis rather than a certainty. It will take a

