GERALD M. EDELMAN

Vider Than the Sky

the phenomenal gift of consciousness



YALE NOTA BENE

yale university press new haven and london

Contents

	Preface	xi
	Acknowledgments	xv
1.	The Mind of Man: <i>Completing Darwin's Program</i>	1
2.	Consciousness: The Remembered Present	4
3.	Elements of the Brain	14
4.	Neural Darwinism: A Global Brain Theory	32
5.	The Mechanisms of Consciousness	48
6.	Wider Than the Sky: <i>Qualia</i> , <i>Unity</i> , and <i>Complexity</i>	60
7.	Consciousness and Causation: The Phenomenal Transform	76
8.	The Conscious and the Nonconscious: <i>Automaticity and Attention</i>	87
9.	Higher-Order Consciousness and Representation	97
10.	Theory and the Properties of Consciousness	113

11.	Identity: The Self, Mortality, and Value	131
12.	Mind and Body: Some Consequences	140
	Glossary	149
	Bibliographic Note	181
	Index	187

Chapter i

The Mind of Man

COMPLETING DARWIN'S PROGRAM

n 1869, Charles Darwin found himself vexed with his friend Alfred Wallace, the co-founder of the theory of evolution. They had differed on several issues related to that theory. But the main reason for Darwin's disturbance was a publication by Wallace concerning the origin of the brain and mind of man. Wallace, who by that time had spiritualist leanings, concluded that natural selection could not account for the human mind and brain.

Darwin wrote to him before publication: "I hope you have not murdered too completely your own and

my child," meaning, of course, natural selection. Wallace, in fact, concluded that natural selection could not explain the origin of our higher intellectual and moral faculties. He claimed that savages and prehistoric humans had brains almost as large as those of Englishmen but, in adapting to an environment that did not require abstract thought, they had no use for such structures and therefore their brains could not have resulted from natural selection. Unlike Wallace, Darwin understood that such an adaptationist view, resting only on natural selection, was not cogent. He understood that properties and attributes not necessarily needed at one time could nevertheless be incorporated during the selection of other evolutionary traits. Moreover, he did not believe that mental faculties were independent of one another. As he explained in his book The Descent of Man, for example, the development of language might have contributed to the process of brain development.

This rich work has prevailed, along with Darwin's other views, but the program he established remains to be completed. One of the key tasks in completing that program is to develop a view of consciousness as a product of evolution rather than as a Cartesian substance, or *res cogitans*, a substance not accessible to scientific analysis. A major goal of this book is to develop such a view.

What is required to carry out such a project? Before answering this question, let us consider Darwin's entry in his notebook of 1838: "Origin of man now proved metaphysic must flourish—He who understands baboon will do more towards metaphysics than Locke." These statements point in the direction we must follow. We must have a biological theory of consciousness and provide supporting evidence for that theory. The theory must show how the neural bases for consciousness could have arisen during evolution and how consciousness develops in certain animals.

Two subtle but important issues strongly influence our interpretation of these requirements. The first of these is the question of the causal status of consciousness. Some take the view that consciousness is a mere epiphenomenon with no material consequences. A contrary view is that consciousness is efficacious-that it causes things to happen. We will take the position, which we shall explore in detail later, that it suffices to show that the neural bases of consciousness, not consciousness itself, can cause things to happen. The second major challenge to any scientific account of consciousness is to show how a neural mechanism entails a subjective conscious state, or quale, as it is called. Before we can meet these two challenges, it is necessary to provide a sketch of the properties of consciousness and consider some matters of brain structure and function.

Chapter 2

Consciousness

THE REMEMBERED PRESENT

e all know what consciousness is: it is what you lose when you fall into a deep dreamless sleep and what you regain when you wake up. But this glib statement does not leave us in a comfortable position to examine consciousness scientifically. For that we need to explore the salient properties of consciousness in more detail, as William James did in his *Principles of Psychology*. Before doing so, it will help to clarify the subject if we first point out that consciousness is utterly dependent on the brain. The Greeks and others believed that consciousness resided in the heart, an idea that survives in many of our common metaphors. There is now a vast amount of empirical evidence to support the idea that consciousness emerges from the organization and operation of the brain. When brain function is curtailedin deep anesthesia, after certain forms of brain trauma, after strokes, and in certain limited phases of sleepconsciousness is not present. There is no return of the functions of the body and brain after death, and postmortem experience is simply not possible. Even during life there is no scientific evidence for a free-floating spirit or consciousness outside the body: consciousness is embodied. The question then becomes: What features of the body and brain are necessary and sufficient for consciousness to appear? We can best answer that question by specifying how the properties of conscious experience can emerge from properties of the brain.

Before taking up the properties of consciousness in this chapter, we must address another consequence of embodiment. This concerns the private or personal nature of each person's conscious experience. Here is James on the subject:

In this room—this lecture room, say—there are a multitude of thoughts, yours and mine, some of which cohere mutually, and some not. They are as little each-for-itself and reciprocally independent as they are all-belonging-together. They are neither: no one of them is separate, but each belongs with certain others and with none beside. My thought belongs with my other thoughts and your thought with your other thoughts. Whether anywhere in the room there be a mere thought, which is nobody's thought, we have no means of ascertaining, for we have no experience of the like. The only states of consciousness that we naturally deal with are found in personal consciousness, minds, selves, concrete particular I's and you's.

There is no mystery here. Since consciousness arrives as a result of each individual's brain and bodily functions, there can be no direct or collective sharing of that individual's unique and historical conscious experience. But this does not mean that it is impossible to isolate the salient features of that experience by observation, experiment, and report.

What is the most important statement one can make about consciousness from this point of view? It is that consciousness is a process, not a thing. James made this point trenchantly in his essay "Does Consciousness Exist?" To this day, many category errors have been made as a result of ignoring this point. For example, there are accounts that attribute consciousness specifically to nerve cells (or "consciousness neurons") or to particular layers of the cortical mantle of the brain. The evidence, as we shall see, reveals that the process of consciousness is a dynamic accomplishment of the distributed activities of populations of neurons in many different areas of the brain. That an area may be essential or necessary for consciousness does not mean it is sufficient. Furthermore, a given neuron may contribute to conscious activity at one moment and not at the next.

There are a number of other important aspects of consciousness as a process that may be called Jamesian properties. James pointed out that consciousness occurs only in the individual (that is, it is private or subjective), that it appears to be continuous, albeit continually changing, that it has intentionality (a term referring to the fact that, generally, it is about things), and that it does not exhaust all aspects of the things or events to which it refers. This last property has a connection to the important matter of attention. Attention, particularly focal attention, modulates conscious states and directs them to some extent, but it is not the same as consciousness. I will return to this issue in later chapters.

One outstanding property is that consciousness is unitary or integrated, at least in healthy individuals. When I consider my conscious state at the time of this writing, it appears to be all of a piece. While I am paying attention to the act of writing, I am aware of a ray of sunlight, of a humming sound across the street, of a small discomfort in my legs at the edge of the chair, and even of a "fringe," as James called it, that is of objects and events barely sensed. It is usually not entirely possible to reduce this integrated scene to just one thing, say my pencil. Yet this unitary scene will change and differentiate according to outside stimuli or inner thoughts to yet another scene. The number of such differentiated scenes seems endless, yet each is unitary. The scene is not just wider than the sky, it can contain many disparate elements—sensations, perceptions, images, memories, thoughts, emotions, aches, pains, vague feelings, and so on. Looked at from the inside, consciousness seems continually to change, yet at each moment it is all of piece—what I have called "the remembered present"—reflecting the fact that all my past experience is engaged in forming my integrated awareness of this single moment.

This integrated yet differentiated state looks entirely different to an outside observer, who possesses his or her own such states. If an outside observer tests whether I can consciously carry out more than two tasks simultaneously, he will find that my performance deteriorates. This apparent limitation of conscious capability, which is in contrast to the vast range of different inner conscious states, deserves analysis. I will consider its origins when I discuss the difference between conscious and nonconscious activity.

So far, I have not mentioned a property that is certainly obvious to all humans who are conscious. We are conscious of being conscious. (Indeed, it is just such a form of consciousness that impels the writing of this book.) We have scant evidence that other animals possess this ability; only higher primates show signs of it. In the face of this fact, I believe that we need to make a distinction between primary consciousness and higherorder consciousness. Primary consciousness is the state of being mentally aware of things in the world, of having mental images in the present. It is possessed not only by humans but also by animals lacking semantic or linguistic capabilities whose brain organization is nevertheless similar to ours. Primary consciousness is not accompanied by any sense of a socially defined self with a concept of a past or a future. It exists primarily in the remembered present. In contrast, higher-order consciousness involves the ability to be conscious of being conscious, and it allows the recognition by a thinking subject of his or her own acts and affections. It is accompanied by the ability in the waking state explicitly to recreate past episodes and to form future intentions. At a minimal level, it requires semantic ability, that is, the assignment of meaning to a symbol. In its most developed form, it requires linguistic ability, that is, the mastery of a whole system of symbols and a grammar. Higher primates, to some minimal degree, are assumed to have it, and in its most developed form it is distinctive of humans. Both cases require an internal ability to deal with tokens or symbols. In any event, an animal with higher-order consciousness necessarily must also possess primary consciousness.

There are different levels of consciousness. In rapid eye movement (REM) sleep, for example, dreams are conscious states. In contrast with individuals in the waking state, however, the dreaming individual is often gullible, is generally not conscious of being conscious, is not connected to sensory input, and is not capable of motor output. In deep or slow-wave sleep, short dreamlike episodes may occur, but for long periods there is no evidence of consciousness. In awaking from the unconsciousness induced by trauma or anesthesia, there may be confusion and disorientation. And, of course, there may be diseases of consciousness, such as schizophrenia, in which hallucinations, delusions, and disorientation can occur.

In the normal conscious state, individuals experience qualia. The term "quale" refers to the particular experience of some property-of greenness, for instance, or warmth, or painfulness. Much has been made of the need for providing a theoretical description that will allow us directly to comprehend qualia as experiences. But given that only a being with an individual body and brain can experience qualia, this kind of description is not possible. Qualia are high-order discriminations that constitute consciousness. It is essential to understand that differences in qualia are based on differences in the wiring and activity of parts of the nervous system. It is also valuable to understand that qualia are always experienced as parts of the unitary and integrated conscious scene. Indeed, all conscious events involve a complex of qualia. In general, it is not possible to experience only a single quale-"red," say-in isolation.

I shall elaborate later on the statement that qualia reflect the ability of conscious individuals to make highorder discriminations. How does such an ability reflect the efficacy of the neural states accompanying conscious experience? Imagine an animal with primary consciousness in the jungle. It hears a low growling noise, and at the same time the wind shifts and the light begins to wane. It quickly runs away, to a safer location. A physicist might not be able to detect any necessary causal relation among these events. But to an animal with primary consciousness, just such a set of simultaneous events might have accompanied a previous experience, which included the appearance of a tiger. Consciousness allowed integration of the present scene with the animal's past history of conscious experience, and that integration has survival value whether a tiger is present or not. An animal without primary consciousness might have many of the individual responses that the conscious animal has and might even survive. But, on average, it is more likely to have lower chances of survival-in the same environment it is less able than the conscious animal to discriminate and plan in light of previous and present events.

In succeeding chapters, I will attempt to explain how conscious scenes and qualia arise as a result of brain dynamics and experience. At the outset, though, it is important to understand what a scientific explanation of conscious properties can and cannot do. The issue concerns the so-called explanatory gap that arises from the remarkable differences between brain structure in the material world and the properties of qualia-laden experience. How can the firing of neurons, however complex, give rise to feelings, qualities, thoughts, and emotions? Some observers consider the two realms so widely divergent as to be impossible to reconcile. The key task of a scientific description of consciousness is to give a causal account of the relationship between these domains so that properties in one domain may be understood in terms of events in the other.

What such an explanation cannot and need not do is offer an explanation that replicates or creates any particular quale or experiential state. Science does not do that-indeed, imagine that a gifted scientist, through an understanding of fluid dynamics and meteorology, came up with a powerful theory of a complex world event like a hurricane. Implemented by a sophisticated computer model, this theory makes it possible to understand how hurricanes arise. Furthermore, with the computer model, the scientist could even predict most of the occurrences and properties of individual hurricanes. Would a person from a temperate zone without hurricanes, on hearing and understanding this theory, then expect to experience a hurricane or even get wet? The theory allows one to understand how hurricanes arise or are entailed by certain conditions, but it cannot create the experience of hurricanes. In the same way, a brain-based theory of consciousness should give a causal explanation of its properties but, having done so, it should not be expected to generate qualia "by description."

To develop an adequate theory of consciousness,

one must comprehend enough of how the brain works to understand phenomena, such as perception and memory, that contribute to consciousness. And if these phenomena can be causally linked, one would hope to test their postulated connections to consciousness by experimental means. This means that one must find the neural correlates of consciousness. Before addressing these issues, let us turn first to the brain.

CONSCIOUSNESS

Chapter 3

Elements of the Brain

he human brain is the most complicated material object in the known universe. I have already said that certain processes within the brain provide the necessary mechanisms underlying consciousness. In the past decade or so, many of these processes have been identified. Brain scientists have described an extraordinary layering of brain structures at levels ranging from molecules to neurons (the message-carrying cells of the brain), to entire regions, all affecting behavior. In describing those features of the brain necessary to our exploration I will not go into great detail. To provide a foundation for a biological theory of consciousness, however, we do need to consider certain basic information on brain structure and dynamics. This excursion will require some patience on the reader's part. It will be rewarded when we develop a picture of how the brain works.

This short survey on the brain will cover, in order, a global description of brain regions, some notion of their connectivity, the basics of the activity of neurons and their connections-the synapses-and a bit of the chemistry underlying neuronal activity. All this will be necessary to confront a number of critical questions and principles: Is the brain a computer? How is it built during development? How complex are its transactions? Are there new principles of organization unique to the brain that were selected during evolution? What parts of the brain are necessary and sufficient for consciousness to emerge? In addressing these questions, I shall use the human brain as my central reference. There are, of course, many similarities between our brains and those of other animal species, and when necessary I shall describe these similarities as well as any significant differences.

The human brain weighs about three pounds. Its most prominent feature is the overlying wrinkled and convoluted structure known as the cerebral cortex, which is plainly visible in pictures of the brain (Figure 1). If the cerebral cortex were unfolded (making the gyri, its protrusions, and the sulci, its clefts, disappear) it would have the size and thickness of a large table napkin. It would contain at least 30 billion neurons, or nerve



Figure 1. Relative locations of major parts of the human brain. The cerebral cortical mantle receives projections from the thalamus and sends reciprocal projections back; this constitutes the thalamocortical system. Beneath the mantle are three major cortical appendages—the basal ganglia, the hippocampus, and the cerebellum. Below them is the brainstem, evolutionarily the oldest part of the brain, which contains several diffusely projecting value systems.

cells, and 1 million billion connections, or synapses. If you started counting these synapses right now at a rate of one per second, you would just finish counting them 32 million years from now.

Neurons are connected to each other locally to form a dense network in portions of the brain called gray matter; they communicate over longer distances via fiber tracts called white matter. The cortex itself is a sixlayered structure with different connection patterns in each layer. The cortex is subdivided into regions that mediate different sensory modalities, such as hearing, touch, and sight. There are other cortical regions dedicated to motor functions, the activity of which ultimately drives our muscles. Beyond the sensorimotor portions concerned with input and output, there are regions such as the frontal, parietal, and temporal cortices that are connected only to other parts of the brain and not to the outside world.

Before taking up other portions of the brain, I shall briefly describe in simplified form the structure and function of neurons and synapses. Different neurons can have a number of shapes, and there may be as many as two hundred or more different kinds in the brain. A neuron consists of a cell body with a diameter on the order of thirty microns, or about one ten-thousandth of an inch across (Figure 2). Neurons tend to be polar, with a treelike set of extensions called dendrites, and a long specialized extension called an axon, which connects the neuron to other neurons at synapses. The synapse is a specialized region that links the so-called presynaptic neuron (the neuron that sends a signal across the synapse) to a postsynaptic neuron (the neuron that receives the signal). The presynaptic portion of the synapse contains a special set of minute vesicles within which are chemicals known as neurotransmitters. Neurons possess an electrical charge as a result of their membrane properties, and when a neuron is excited current



Figure 2. A diagram illustrating synaptic connections between two neurons. An action potential traveling down the axon of the presynaptic neuron causes the release of a neurotransmitter into the synaptic cleft. The transmitter molecules bind to receptors in the postsynaptic membrane, changing the probability that the postsynaptic neuron will fire. (Because of the number of different shapes and kinds of neurons, this drawing is of necessity a greatly simplified cartoon.)

flows through channels that open across the membrane. As a result, a wave of electrical potential known as an action potential moves from the cell body down the presynaptic axon and causes the release of neurotransmitter molecules from vesicles into the synaptic cleft. These molecules bind to molecular receptors or channels in the postsynaptic cell that, acting cumulatively, can cause it to fire an action potential of its own. Thus, neuronal communication occurs by a combination of controlled electrical and chemical events.

Now try to imagine the enormous numbers of neurons firing in various areas of the brain. Some firings are coherent (that is, they are simultaneous), others are not. Different brain regions have different neurotransmitters and chemicals whose properties change the timing, amplitude, and sequences of neuronal firing. To achieve and maintain the complex patterns of dynamic activity in healthy brains, some neurons are inhibitory, suppressing the firing of others, which are excitatory. Most excitatory neurons use the substance glutamate as their neurotransmitter, while the inhibitory neurons use GABA (gammaaminobutyric acid). We can ignore the chemical details for now and simply accept that the effects of different chemical structures are different and that their distribution and occurrence together can have significant effects on neural activity.

I started by describing the cortex. With the picture of a polar neuron in mind, we can turn briefly to other key regions of the brain. One of the most important anatomical structures for understanding the origin of consciousness is the thalamus. This structure, which is located at the center of the brain, is essential for conscious function, even though it is only somewhat larger than the last bone in your own thumb. When nerves from different sensory receptors serving different modalities (located in your eyes, ears, skin, and so on) travel to your brain, they each connect in the thalamus with specific clusters of neurons called nuclei. Postsynaptic neurons in each specific thalamic nucleus then project axons that travel and map to particular areas of the cortex. A well-studied example is the projection from the neurons of the retina through the optic nerve to the part of the thalamus called the lateral geniculate nucleus and then to the primary visual cortical area, called V1 (for "visual area 1").

There is one striking feature of the many connections between the thalamus and the cortex: not only does the cortex receive many axons from thalamic neurons but there are also reciprocal axonal fibers going from the cortex back to the thalamus. We speak therefore of thalamocortical projections and corticothalamic projections. Reciprocal connections of this type abound within the cortex itself; such reciprocal connections are called corticocortical tracts. A striking example of these is the fiber bundle called the corpus callosum, which connects the two cortical hemispheres and consists of more than 200 million reciprocal axons. Cutting the corpus callosum leads to a split-brain syndrome, which in some cases can lead to the remarkable appearance of two separate and very different consciousnesses.

Each specific thalamic nucleus (and there are many) does not connect directly to any of the others. Surrounding the periphery of the thalamus, however, there is a layered structure called the reticular nucleus, which connects to the specific nuclei and which can inhibit their activity. The reticular nucleus, it is suspected, acts to switch or "gate" the activities of the specific thalamic nuclei, yielding different patterns of expression of such sensory modalities as sight, hearing, and touch. Another set of thalamic nuclei called intralaminar nuclei receive connections from certain lower structures in the brainstem that are concerned with activation of multiple neurons; these then project to many different areas of the cortex. The activity of these intralaminar nuclei is suspected to be essential for consciousness in that it sets appropriate thresholds or levels of cortical response with too high a threshold, consciousness would be lost.

We may now turn to some other brain structures that are important to our efforts to track down the neural bases of consciousness. These are large subcortical regions that include the hippocampus, the basal ganglia, and the cerebellum. The hippocampus is an evolutionarily ancient cortical structure lined up like a pair of curled sausages along the inner skirt of the temporal cortex, one on the right side and another on the left. In cross section, each sausage looks like a sea horse, hence the name "hippocampus." Studies of the neural properties of the hippocampus provide important examples of some of the synaptic mechanisms underlying memory. One such mechanism, which should *not* be equated with memory itself, is the change in the strength, or efficacy, of hippocampal synapses that occurs with certain patterns of neural stimulation. As a result of this change, which can be either positive for long-term potentiation or negative for long-term depression, certain neural pathways are dynamically favored over others.

The point to be stressed is that, while synaptic change is essential for the function of memory, memory is a system property that also depends on specific neuroanatomical connections.

Increased synaptic strength or efficacy within a pathway leads to a higher likelihood of conduction across that pathway, whereas decreases in synaptic strength diminish that likelihood. Various patterns have been found for the so-called synaptic rules governing these changes, following the initial proposals of Donald Hebb, a psychologist, and Friedrich von Hayek, an economist who, as a young man, thought quite a bit about how the brain works. These scholars suggested that an increase in synaptic efficacy would occur when pre- and postsynaptic neurons fired in close temporal order. Various modifications of this fundamental rule have been seen in different parts of the nervous system. What is particularly striking about the hippocampus, where these rules have been studied in detail, is the fact that bilateral removal of this structure leads to a loss of episodic memory, the memory of specific episodes or experiences in life. A very famous patient, H. M., whose hippocampi were removed to cure epileptic seizures, could not, for example, convert his short-term memory of events into a permanent narrative record, a condition

that was depicted dramatically in the movie *Memento*. It is believed that such a long-term record results when particular synaptic connections between the hippocampus and the cortex are strengthened. When these connections are severed, the corresponding cortical synaptic changes cannot take place and the ability to remember episodes over the long term is lost. Such a patient can remember episodes up to the time of the operation, but loses long-term memory thereafter. It is intriguing that in some animals, such as rodents, hippocampal function is necessary for memories of a sense of place. In the absence of hippocampal functions, the animal cannot remember target places that have been explored.

All of the discussion so far has focused on sensory or cognitive functioning. The brain's motor functions, however, are also critically important, not just for the regulation of movement, but also for forming images and concepts, as we shall see. A key output area is the primary motor cortex, which sends signals down through the spinal cord to the muscles. There are also many other motor areas in the cortex, and there are nuclei in the thalamus related to motor function as well. Another structure related to motor functions is the cerebellum, a prominent bulb at the base of the cortex and above the brainstem (see Figure 1). The cerebellum appears to serve in the coordination and sequencing of motor actions and sensorimotor loops. There is no evidence, however, that it participates directly in conscious activity.

An intriguing set of structures known as the basal ganglia is critically important for motor control and sequencing. Lesions of certain structures in these nuclei lead to a loss of the neurotransmitter dopamine, and thus to the symptoms of Parkinson's syndrome. Patients with this disease have tremors, difficulty in initiating motor activity, rigidity, and even certain mental symptoms. The basal ganglia, as shown in Figure 1, are located in the center of the brain and connect to the cortex via the thalamus. Their neural connectivity, which is radically different from that of the cortex, consists of circuits of successive synapses or polysynaptic loops connecting the various ganglia. For the most part, the reciprocal connection patterns seen in the cortex itself and between the cortex and the thalamus are lacking in the basal ganglia. Moreover, most of the activity of the basal ganglia is through inhibitory neurons using GABA as a neurotransmitter. Nevertheless, since inhibition of inhibition (or disinhibition) can occur in these loops, they can stimulate target neurons as well as suppress their activity.

The basal ganglia are believed to be involved in the initiation and control of motor patterns. It is also likely that much of what is called procedural memory (remembering how to ride a bicycle, for example) and other nonconscious learned activity depends on the functions of the basal ganglia. As we shall see later, the regulatory functions of basal ganglia are also significant for forming categories of perceptions during experience.

There is one final set of structures that is critical in brain activity connected with learning and the maintenance of consciousness. These are the ascending systems, which my colleagues and I have called value systems because their activity is related to rewards and responses necessary for survival. They each have a different neurotransmitter, and from their nuclei of origin they send axons up and down the nervous system in a diffuse spreading pattern. These nuclei include the locus coeruleus, a relatively small number of neurons in the brainstem that release noradrenaline; the raphé nucleus, which releases serotonin; the various cholinergic nuclei, so-called because they release acetylcholine; the dopaminergic nuclei, which release dopamine; and the histaminergic system, which resides in a subcortical region called the hypothalamus, a region that affects many critical body functions.

The striking feature of such value systems is that, by projecting diffusely, each affects large populations of neurons simultaneously by releasing its neurotransmitter in the fashion of a leaky garden hose. By doing so, these systems affect the probability that neurons in the neighborhood of value-system axons will fire after receiving glutamatergic input. These systems bias neuronal responses affecting both learning and memory and controlling bodily responses necessary for survival. It is for this reason that they are termed value systems. In addition, there are other loci in the brain with modulatory functions mediated by substances called neuropeptides. An example is enkephalin, an endogenous opioid that regulates responses to pain. In addition, there are other brain areas, such as the amygdala, which are involved in emotional responses, such as fear. For our purposes, these areas need not be described in detail.

To summarize our account so far, we may say that, in a gross sense, there are three main neuroanatomical motifs in our brains (Figure 3). The first is the thalamocortical motif, with tightly connected groups of neurons connected both locally and across distances by rich reciprocal connections. The second is the polysynaptic loop structure of the inhibitory circuits of the basal

Figure 3. Fundamental arrangements of three kinds of neuroanatomical systems in the brain. The top diagram shows the gross topology of the thalamocortical system, which is a dense meshwork of reentrant connectivity between the cortex and the thalamus and among different cortical areas. The middle diagram shows the long polysynaptic loops connecting the cortex with subcortical structures such as the basal ganglia. In this case, these loops go from the basal ganglia to the thalamus, thence to the cortex and back from the target areas of cortex to the ganglia. These loops are, in general, not reentrant. The bottom diagram shows one of the diffusely projecting value systems, in which the locus coeruleus distributes a "hairnet" of fibers all over the brain. These fibers release the neuromodulator noradrenaline when the locus coeruleus is activated.



ganglia. The third consists of the diffuse ascending projections of the different value systems. Of course, this generalization is a gross oversimplification, given the exquisite detail and individuation of neural circuitry. But as we shall see, it provides a useful simplification; we can dispose of it once we have seen its uses.

So much for simplicity. The picture I have painted so far only hints at the remarkably complex dynamics of the neural structures of the brain. After staring at the gross layout of brain regions in Figure 1 and understanding the synapse pictured in Figure 2, close your eyes and imagine myriad neural firings in millions of pathways. Some of this neural activity would occur at certain frequencies while others would show variable frequencies. Bodily activity and signals from the environment and the brain itself would modify which of the pathways were favored over others as a result of changes in synaptic strength. Although I hardly expect the reader to be able to visualize precisely the hyperastronomical numbers of neural patterns in detail, perhaps this exercise will yield a further appreciation of the brain's complexity.

We are now in a position to address some of the questions posed at the beginning of this chapter. Consider the question of whether the brain is a computer. If we examine how neural circuits are built during animal development, this would seem unlikely. The brain arises during development from a region of the embryo called the neural tube. Progenitor cells (cells that are precursors to neurons and support cells called glia) move in certain patterns to make various layers and patterns. As they differentiate into neurons, many also die. From the very beginning of neuroanatomy, there are rich statistical variations in both cell movement and cell death. As a result, no two individuals, not even identical twins, possess the same anatomical patterns.

In the earliest stages of development, the cellular organization characteristic of a species is controlled by families of genes among which are the so-called Hox genes and Pax genes. But at a certain point, the control of neural connectivity and fate becomes epigenetic; that is, it is not prespecified as "hardwiring," but rather is guided by patterns of neural activity. Neurons that fire together wire together. While, at earlier stages, patterned cell movement and programmed cell death determine anatomical structure, the movement and death of individual neurons are nonetheless statistically variable or stochastic. The same holds for which particular neurons connect to each other at later stages. The result is a pattern of constancy and variation leading to highly individual networks in each animal. This is no way to build a computer, which must execute input algorithms or effective procedures according to a precise prearranged program and with no error in wiring.

There are other, even more trenchant reasons for rejecting the idea of digital computation as a basis for brain action. As we shall see later, what would be lethal noise for a computer is in fact critical for the operation of higher-order brain functions. For the moment, though, let us consider some other aspects of brain complexity and its relation to brain structure and function.

A review of what I have said about the overall arrangement of brain areas might tempt one to conclude that the key to brain function is modularity. Since there are regions that are functionally segregated for vision (even for color, movement, and orientation), for example, or similarly for hearing or touch, we might be tempted to conclude that specific brain action is mainly the result of the specialized functioning of these isolated local parts or modules. If pushed to higher levels, this simple notion results in phrenology, the picture of localized separate brain faculties first proposed by Franz Joseph Gall. We now know that modularity of this kind is indefensible. The alternative picture, that the brain operates only as a whole (the holistic view), will also not stand up to scrutiny.

The notion of modularity is based on an overly simple interpretation of the effects of ablation of parts of the brain, either by animal experiments or as a result of a stroke, or of surgery for epilepsy. It is clear, for example, that ablation of cortical area V1 leads to blindness. It does not follow, however, that all the properties of vision are assured by the functioning of V1, which is the first cortical area in a series making up the visual pathway. Similarly, although modern imaging techniques reveal certain areas of the brain that are active in certain tasks, it does not follow that the activity of such areas is the *sole* cause of particular behaviors. Necessity is not sufficiency. But the contrary or holistic argument is not tenable either—one must account for both integration *and* differentiation of brain activity. This will be one of our main tasks in proposing a global brain theory. As we shall see later, the long-standing argument between localizationists and holists dissolves if one considers how the functionally segregated regions of the brain are connected as a complex system in an intricate but integrated fashion. This integration is essential to the emergence of consciousness.

This reasoning is critical to understanding the relationship between brain function and consciousness. Of course, there are areas of the brain that if damaged or removed will lead to permanent unconsciousness. One such area is the midbrain reticular formation. Another is the region of the thalamus containing the intralaminar nuclei. These structures are not the site of consciousness, however. As a process, consciousness needs their activity, but to account for the Jamesian properties of consciousness requires a much more dynamic picture involving integration of the activities of multiple brain regions. We are now in a position to lay the groundwork for just such a picture by considering a global brain theory that accounts for the evolution, development, and function of this most complex of organs.