7 Cognition and the Prefrontal Cortex

We'll discuss two functions central to cognition: working memory and attention. Working memory has to do with how we keep information active while the information is useful for the actions and thoughts that we perform. Attention, as the common usage of the term implies, refers to what we monitor in the world, what really matters. Next, we discuss how these multifaceted functions engage the prefrontal cortex. Neuroscientists like to describe this part of the brain as it's crowning apex, where cognition—including abstract thought—reigns supreme. The chapter then picks this idea apart by tracing its origins and discussing how to think about the brain as a distributed system. We thus start building the case for an entangled view of the brain in terms of interacting parts.

Researchers have always sought to unravel features of the brain that are, presumably, uniquely human, or that at least confer the species with "superior" mental capabilities. The pseudoscientific phrenology movement led by Franz Gall and his disciples, discredited as it is now, galvanized the pursuit of localizing mental functions to brain territories. Phrenology was particularly active from the 1810s to 1840s and sought to find functions in the brain by observing and exploring the skull; proponents would run their fingertips and palms over a person's skull to feel for distinct patterns, such as enlargements or indentations. Perhaps because of its location at the front and top of the brain, Gall placed humans' "highest functions" in parts of the frontal lobe.[1] As neuroscience took shape as an active area of research in the second half of the nineteenth century, the question of localization received much attention and was greatly invigorated by Paul Broca's 1861 report linking language and the prefrontal cortex (chapter 4). His presentation before the *Société d'Anatomie* in Paris, which described his clinical and neurological conclusions about Tan, had a tremendous impact on both clinicians and experimentalists. Broca proposed that the frontal lobe was important for speaking and for higher intellectual functions, including judgment, reflection, and abstraction. Interest in this part of the brain was kindled considerably.

Seventeen years later, Broca would again have a major impact in shaping the prefrontal cortex's importance for intellectual functions. In his monumental theoretical paper of 1878 discussed in chapter 6, Broca developed this theme in a forceful manner, concluding that the primate brain can be distinguished from that of other mammals by "characteristics numerous and very striking." But chief among them was the *predominance* of the frontal lobe. Enamored with this view he threw all caution to the wind:[2]

> The simultaneous appearance of these numerous characteristics leads to major external changes and throws the entire cerebral morphology into such upheaval that one might believe oneself to be seeing a brand new order of things, as if the chain had been broken, nature had smashed its old molds, and the project had been started up again using a completely different set of plans.

In chapter 9, we will discuss brain evolution in more detail and explain how Broca, admittedly without the tools of modern neuroscience, painted an extreme—and incorrect—portrait of the organization of the mammalian brain. Yet, intellect and frontal cortex would be paired for the next 100 years—and still are.

Measuring Electrical Activity from Cortical Neurons

The same type of electrical stimulation studies that uncovered links between the cortex and the body (chapter 6), also sought to uncover the response properties of the cortex. But the discovery of the functions of particular brain parts was advanced enormously by the refinement of the electrophysiological recording technique in the late 1950s and early 1960s. With this method, electrodes were inserted into the cortex to directly record cellular activity, including spiking activity, the all-or-none firing of cells. Initially, recordings were made in anesthetized animals, but the technique was later extended to allow measurements in alert animals, although they still needed to be restrained. (Techniques to record neuronal activity in alert, freely moving animals are expanding rapidly.)

Early studies characterized responses in parts of the cortex that respond to sensory stimulation. Because the animals were restrained and anesthetized, this was a natural place to start. (You might wonder if cell firing can be detected when the animal is anesthetized; in the case of vision, even their eyes were kept open artificially. For sensory responses, basic response properties are similar when the animal is under anesthesia compared to alert.) A second key reason to study sensory regions was that it allowed researchers to choose and control stimuli rather precisely. For example, they could present a visual stimulus of a certain elongation, orientation, and speed at which it moved in front of the eyes. If it was an auditory stimulus, it could be a sound of precise intensity, duration, and wave characteristics.

Altogether, this strategy proved immensely productive and has kept neurophysiologists busy for decades. In the case of the visual system, little by little researchers built a body of knowledge describing how cortical areas process increasingly more complex objects. For example, in the primary visual cortex (area V1), cells respond to fairly specific and basic stimulus properties, such as the orientation of a “bar” (think of a pen). Based on physiological and anatomical properties, researchers gradually assembled a catalog of areas (unimaginatively, some were called V2, V3, and V4). Cell responses in each area are sensitive to different stimulus properties, including object shape, color, and motion. By the middle of the 1980s, around 15 regions in occipital, temporal, and parietal cortex were classified as “visual.” Whereas the responses in some parts of the occipital lobe are fairly basic and tied to the physical quality of objects, responses in parts of temporal and parietal cortex are noticeably more abstract. Some even respond to human faces! Cells in the lower part of the temporal cortex fire intensely in response to pictures of faces and, in some cases, are quite selective in their responding; they might respond to pictures of a certain individual vigorously but much less to other faces or not at all to pictures without a face. One of the cells recorded in a 2005 study of an epilepsy patient with implanted electrodes seemed to have a clear preference for pictures of the actress Jennifer Aniston, leading to infelicitous media reports of the “Jennifer Aniston cell.”[3] Of course, were that cell to die, the patient would not suddenly lose the concept or even the visual image of the actress. The investigators eventually found out that the neuron also fired in response to pictures of her former castmate, Lisa Kudrow, likely reflecting memory associations of the patient in question.

While charting visual areas in the monkey brain, it didn’t take long for researchers to reach the frontal cortex, where some cells are heavily involved in the control of eye movements. Primates, in particular, are attuned to the shape and form of objects in the world, which rely on detailed visual processing at the location where the eye is fixating. As neurophysiology gradually moved to parts of the brain where responses are not directly tied to stimulus properties (it doesn't matter if the visual object is green or yellow, at least insofar as this isn't important for the task), the focus of research shifted to "cognitive variables." How does an animal remember a stimulus, decide what to pay attention to, or switch between tasks?

Holding Information in Mind

In the early 1970s, a major discovery opened the door to the understanding of how cognitive processes are realized in the brain. Psychologists had been interested in how humans and other animals actively maintain information for several decades. If given a list of words to remember, after a brief interval of 20 seconds, say, how many items can people recall? Can they remember both the items and their order? This type of "working memory" (which is different from long-term memory of events in the past) had been extensively investigated in monkeys as early as the 1930s and the importance of the prefrontal cortex already established: lesions of the so-called dorsolateral part of the prefrontal cortex substantially impair working memory abilities.

In a typical working memory experiment, on each trial, a visual stimulus is first shown. A delay period of several seconds then ensues during which the stimulus is absent (in early experiments before the use of computer screens, the monkeys' view was obstructed), followed by the presentation of the so-called test stimulus. The task is to indicate whether or not the test matches the initial stimulus. Successful performance thus necessitates some form of memory trace of the first display, which must be further matched to the test to determine the correct answer.

What type of neural signal could bridge the gap between the two visual items during the delay? Cell responses to visual stimuli are transient and decay back to pre-stimulation levels within half a second or so (in the absence of stimulation, many cells are not completely quiescent but respond at low "baseline" level). Thus, researchers were looking for evidence of *sustained* cell firing during the delay, and that is what they found when recording from neurons in the prefrontal cortex of rhesus monkeys (figure 7.1). The prolonged activity was interpreted to be the neural correlate of remembering

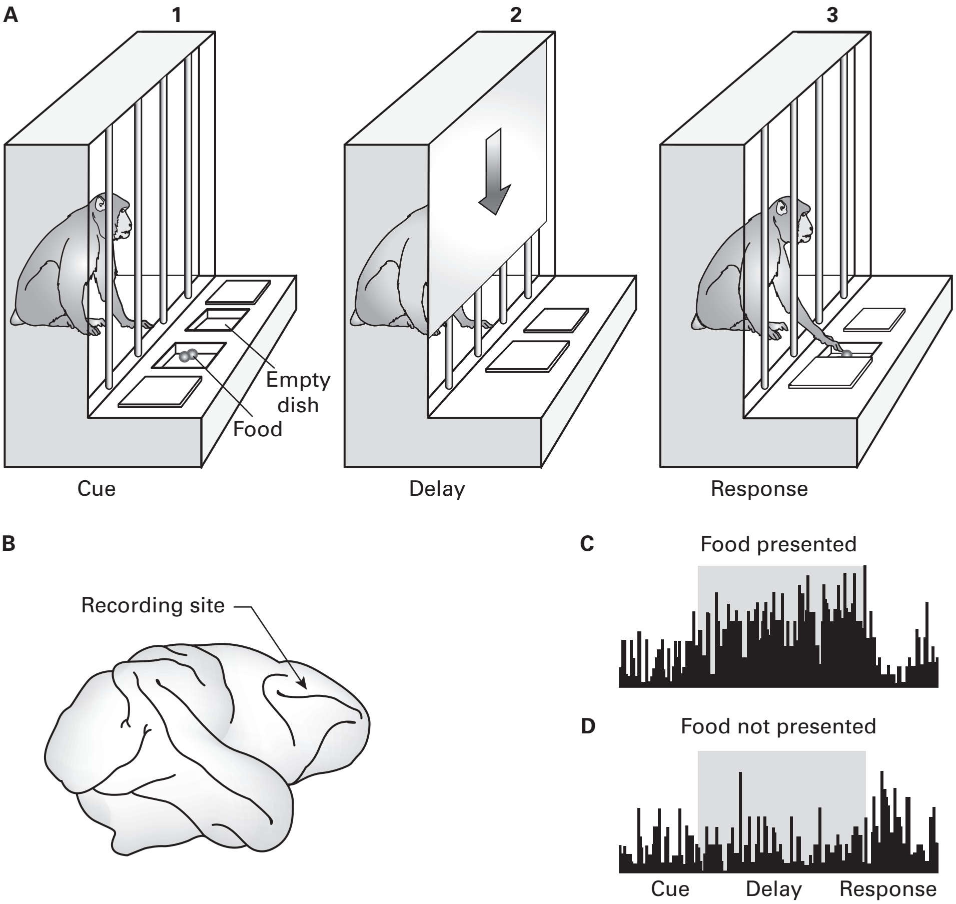

Figure 7.1
Keeping information in mind when it is out of sight. (a) Early experimental setup. (b, c, d) Cell recordings in the prefrontal cortex uncovered cells with sustained firing during the delay interval when the animal does not see the food item.

during the delay and taken to be the *neural signature* of working memory. Two independent papers published in 1971 discovered sustained firing and were enormously influential. For the first time, neuroscientists had shown that an abstract mental operation that wasn't directly linked to a physical stimulus—no stimulus is shown during the delay—could be identified. Studying cognitive tasks was thus viable, and the cellular mechanisms that underlie complex mental functions could be uncovered.

Can we conclude that cells with uninterrupted responses during the delay interval implement short-term remembering? For one, they could be tied to motor preparation, because at the end of the trial the monkey has to indicate whether or not the test stimulus matches the initial one—for example, by touching a response area. So, sustained firing might be related to the motor part of the task, not the mnemonic component. How can we strengthen the evidence in favor of maintenance processes? Cells in the prefrontal cortex prefer certain types of stimuli, of a certain shape or color, say. Multiple studies have shown that sustained responses are tuned to the same attributes of the to-be-remembered stimulus, which is precisely what would be required for a true neural signature (for example, Constantinidis, Franowicz, and Goldman-Rakic 2001). For example, if a prefrontal cell has a preference for a round stimulus when displayed briefly, it will fire vigorously during the delay if that very stimulus needs to be remembered (recall that no stimulus is shown during the delay); otherwise it will respond more weakly.

Should I Pay Attention to You?

Consider three lionesses as they attack a giraffe. They aim to confuse the prey and bounce its attention away; the giraffe cannot keep track of everything that is going on. Animals are not capable of processing all of the information they receive perfectly. How do they *select* some information while ignoring other information sources?

Individual cells in the visual cortex don't respond to stimuli everywhere; they respond only to objects in delimited parts of the field of vision, called the "receptive field." In addition, cells exhibit preference for particular features: Is it moving? Is it colored? And so on. In visual area V4, for example, neurons often exhibit preference for stimuli of certain colors. Researchers capitalized on these properties to investigate the mechanisms of visual attention.[4] How do the responses to something that matters differ from responses to an item that is less relevant? First, the researchers determined an "effective" stimulus that drove a cell's response vigorously, as well as an "ineffective" one that produced weak responses. They then placed both stimuli within the cell's receptive field (so both were capable of eliciting responses from the cell) and taught the animal to pay attention to one of them while ignoring the other. When they rewarded the animal for indicating the elongation (horizontal or vertical) of the relevant one, they found that the locus of attention had a dramatic effect on the cell's response. If the monkey paid attention to the effective stimulus, the cell responded vigorously. Now, when the monkey paid attention to the ineffective stimulus, the output was considerably decreased. This is remarkable because the actual physical stimulation is exactly the same in both conditions (figure 7.2). What changes is the *state* of the animal: In one case the effective stimulus is significant, and in the other the ineffective stimulus is the one that matters to obtain a reward by the experimenter. To a good extent, then, the responses were determined by the attributes of the attended object. It's as if the brain implemented a *filtering* mechanism capable of reducing the influence of the irrelevant information while zooming in on the important information at that time.

Since findings like those in figure 7.2 were first described in the mid-1980s, the literature on the mechanisms of attention has blossomed. We now know that the brain carries out a number of *competition* mechanisms: In the case of vision, the processing of certain objects or parts of the visual field are favored, whereas other objects or spatial locations are de-emphasized. Attention is better understood not as a single process but as a collection of processes that help solve this problem: How does the brain choose information that is relevant while ignoring less pertinent signals?

The Final Frontier: Frontal Cortex

We saw that the frontal cortex plays an important role in working memory—lesions impair a monkey's ability to actively maintain information in the mind. Attention is also believed to involve the frontal cortex. In the example described above, one of the regions influencing visual responses in V4 is an area in the frontal cortex called the "frontal eye field," which was originally known to control eye movements (hence the name) but contributes to attention, too. Anatomical pathways from the frontal eye field to the visual cortex allow the former to direct the latter so as to favor the processing of task-relevant information—the one the animal is instructed to pay attention to.

Before continuing, we need to clarify the distinction between the frontal and prefrontal cortex. The central sulcus is a prominent fissure on the lateral surface of the cortex that separates the parietal and frontal cortices. The first gyrus on the frontal side is where the primary motor cortex is found (figure 7.3); it's the motor part of the homunculus described by Wilder

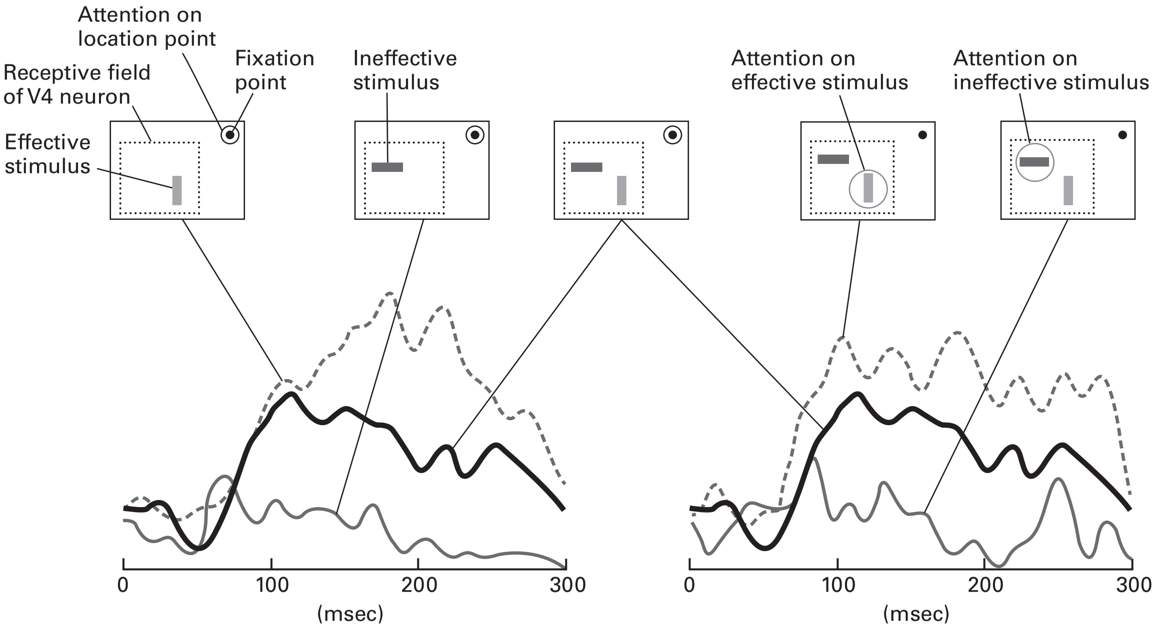

Figure 7.2

Attention can work as a filter helping to select relevant information. In the experimental setup, the animal is trained to keep the eye fixed during multiple stimulation conditions. When both effective and ineffective stimuli are in the receptive field, the response is intermediate (see middle rectangle and solid, dark gray line). When attention is focused on one of the stimuli while another stimulus is simultaneously present (two rightmost rectangles), the response approaches what is observed when a stimulus is presented alone (two leftmost rectangles). In other words, the process of attention appears to “filter out” the unattended stimulus.

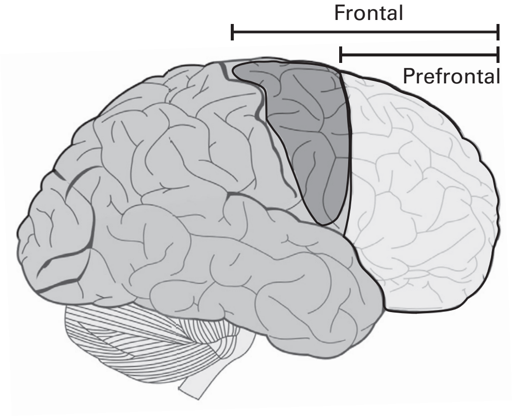

Figure 7.3
Frontal and prefrontal cortex. The prefrontal cortex includes areas that are less directly involved in motor function and are often considered involved in "abstract" processing.

Penfield (chapter 6), and cells there are intimately linked to motor actions. If we move further to the front of the brain, we encounter multiple "premotor" areas, where cell firing reflects motor intentions and planning but don't actually generate movements. To separate these regions of the frontal cortex involved in motor function from the more "noble parts" of the lobe (see discussion in the next section), the latter are described as "prefrontal."

Both working memory and attention are paradigmatic cognitive processes—they involve a prefrontal cortex that matches most neuroscientists' *expectations* of how the brain works.[5] But what is the general function of the prefrontal cortex? What does it do? Does damage to the prefrontal cortex "remove" intelligence from a person? Studies of patients with lesions have yielded somewhat of a puzzle or what's been described as the "riddle of the frontal lobe" (Damasio 1979). Indeed, it is frequently difficult to detect a neurological disorder based on patients' everyday behavior. They do not display obvious impairments in perceptual abilities, and their speech can be fluent and coherent. On conventional tests of intelligence, including standard IQ tests, they perform normally. But with more sensitive and specific tests, it becomes clear that frontal lesions disrupt function.

In some cases, the test administered is not so specific or even sophisticated. Take as an illustration (with no small amount of black humor) how the behavior of a prefrontal patient can be dominated by immediately available information. F. Lhermitte, a neurologist in Paris, directed patients to sit at a desk containing a hammer, a nail, and a picture (Lhermitte 1983). One patient picked up the items to hang the picture on the wall. In another setting, the neurologist approached, placed a hypodermic needle on the desk, dropped his trousers, and turned his back to the patient. What did the patient do? Unfazed, he simply picked up the needle and gave the doctor a healthy jab in the buttocks![6] If you are wondering what the patients were thinking, they would state something like, "You held out objects to me; I thought I had to use them."

As fascinating as they are, these examples don't help much in dissecting what might be happening. Is it a matter of disinhibition, only being able to focus on what is immediately present, loss of social knowledge, or some combination of these factors? Interestingly, in the cases above, the patients were capable of carrying out behaviors which, although relatively simple, were definitely nontrivial (think of the challenges of developing a robot to execute those tasks). So what does the prefrontal cortex do? More developed experimental paradigms converge on the idea that it is involved in what is called *cognitive control*. I'll illustrate this by describing two tasks: the Wisconsin card sort task and the Stroop task.[7]

In the Wisconsin card sort task, participants are asked to sort cards according to the shape, color, or the number of symbols (figure 7.4); they aren't told the rule by which to sort the cards, only whether or not they are correct after responding to each card. After the person obtains 10 correct choices, the rule switches. The catch is that the participant is not informed of this change, so after making a mistake, the person has to change the sorting rule. Once a correct response is generated, the rule is maintained until 10 correct choices are produced, and so on, until all cards on the pile are sorted. In a landmark study, the neuropsychologist Brenda Milner discovered that, in this task, all her frontal patients performed more poorly than her control participants, who had lesions of other parts of the cortex, such as the parietal and occipital cortex. In what way did they struggle? After the surgery, they tended to stick to one rule (sort by color, say) and to perseverate as they sorted the cards, *even* as they received continued negative feedback. Milner's approach was experimentally impeccable. Not only did she examine performance in patients with lesions elsewhere in the cortex, but each of her study groups was tested before and after the operation, allowing her to isolate the *change* in behavior based on lesion location.

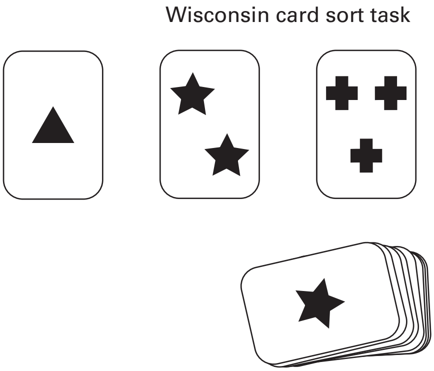

Figure 7.4
The Wisconsin card sort task requires sorting all cards according to shape, color, or number of symbols. Here, the test card at the bottom could match the first or second cards along the row above, based on number or shape, respectively. Participants are never explicitly told the rule, only if they are correct or not.

Nevertheless, interpreting the results of the Wisconsin task is challenging not least because it is difficult, and even some regular participants without a brain lesion perform poorly at it.

The Stroop task (named after its developer, J. Ridley Stroop) is attractive given its combination of simplicity and effectiveness. Imagine a page with words used for colors, such as “red,” “green,” and so on, written in black ink. Now imagine writing the words with pens of different colors, including one that matches the word and one that doesn’t—for example, the word “red” written in red ink (called congruent) or the word “red” written in green ink (called incongruent). Your task is to tell the color of the ink used to write the word. If the word’s name doesn’t match the ink color, the task feels somewhat harder than it should be. In the Stroop task, the participant is asked to either read the word or to name the color in which it is written in different trials. Every stimulus contains two properties (word meaning and color), and the participant must attend selectively to the appropriate attribute to perform adequately. The key comparison is to contrast a person’s reaction time to answer incongruent and congruent trials—one is a little slower in the former. Patients with frontal lobe damage have difficulty with this task, particularly in more challenging versions when the instructions are provided shortly before seeing the word (Dunbar and Sussman 1995).

The Stroop task is called a *conflict task* because during incongruent trials one is asked to focus on a particular dimension (color) in the presence of competing information (word meaning). Both *detecting* the presence of conflict and *resolving* it are important cognitive functions. Versions of conflict occur in many behavioral contexts, in particular when more automatic and controlled behaviors interact. When certain actions are practiced over and over, they become habitual, a process of *habit formation* that's important for making them efficient—that is, more automatic. Now, once the behavior is performed in a fairly automatic fashion, it can be rather difficult to break it: Once triggered, it's executed. But if a habitual action is uncalled for, it should be possible to recalibrate it, modify it, or call it off completely. That's when *cognitive control* comes into play. As a simple example, consider the act of stepping on the gas pedal when the light turns green at an intersection. For those of us who have been driving for many years, this action is completely habitual. But upon spotting a child crossing the street on a bicycle, one should be able to step on the brake pedal immediately (hopefully). Seeing both the green light and the child triggers the conflict (the detection part); releasing the accelerator and rapidly stepping on the brake solves the problem (the resolution part).

The Stroop task illustrates a fundamental aspect of cognitive control and *goal-directed* behaviors: the ability to select a weaker but task-relevant response (or source of information) in the face of competition from an otherwise stronger but task-irrelevant one.[8] Researchers believe that this is the central contribution of the prefrontal cortex, which allows adherence to goals in the presence of competing, stronger actions. It has been proposed that the key function of the prefrontal cortex is to enable and ensure adopting and following the "rules of the game." Think back to Lhermitte's patient: whatever the rules of being examined by a doctor in Paris, they don't include jabbing them in the buttocks with a syringe!

I'll end this section with a historical note. The quest to discover *the* function of the prefrontal cortex conceptualizes this territory as a unit, or at least as a relatively coherent functional entity. But the prefrontal cortex is large, and it is not surprising that the search to find its essence has been beleaguered with difficulties. By the 1980s, it was becoming increasingly clear that the heterogeneity of this large sector would need to be confronted for progress to accelerate. Even at a rather coarse level, we can distinguish three broad segments associated with the lateral, medial, and orbital surfaces.

The lateral aspect includes the dorsolateral prefrontal cortex, which is frequently linked to executive control functions. Both the medial and orbital surfaces have long been linked to emotional and motivational processing, too (chapter 6). Whereas most researchers have by now relinquished the goal of a more unified theory of prefrontal working and accepted its *multifunctional* nature, still, the territory is most frequently linked to cognitive processes, in particular the so-called *executive functions* required during nonroutine, challenging situations. Later, in chapters 10 and 11, we'll explore how emotion and motivation engage the prefrontal cortex together with cognition.

From Sensation to Cognition, One Step at a Time

How are cognitive mechanisms built? In this section, I'll discuss a view that was popular for many decades, until at least the late 1980s. Although researchers currently don't subscribe to stronger incarnations of this model, the main idea lies dormant in the background, and current thinking is still influenced by it. It goes roughly like this.

Parts of the cortex are sensory, others are motor. In the former, cell responses are tied to stimulus and perceptual properties; in the latter, they are linked to movements and actions in a fairly direct manner. We can think of the prefrontal cortex as the part of the brain that's decoupled from immediate sensory and motor variables. In a nutshell, it deals with the *abstract*. On the sensory end, researchers think of the cortex in terms of information flowing from early sensory regions (for instance, visual, auditory) to parts of the prefrontal cortex (figure 7.5). Along the way, intermediate regions have cells whose responses are progressively more independent from sensory variables. Sensory signals provide a steady stream of signals that progress through a *cortical hierarchy* until reaching the prefrontal cortex. At each junction, responses become more refined—namely, less about physical properties. In the prefrontal cortex, they are sufficiently abstract that they can support "symbolic processing" of the type that possibly distinguishes humans from other apes, or maybe apes from other primates.

Consider the case of vision. Cells in the primary visual cortex (area V1) respond to basic stimulus properties, which are elaborated in subsequent areas, including V2, V3, and V4. After subsequent steps along the visual hierarchy, responses recorded along the temporal cortex reflect novel properties;

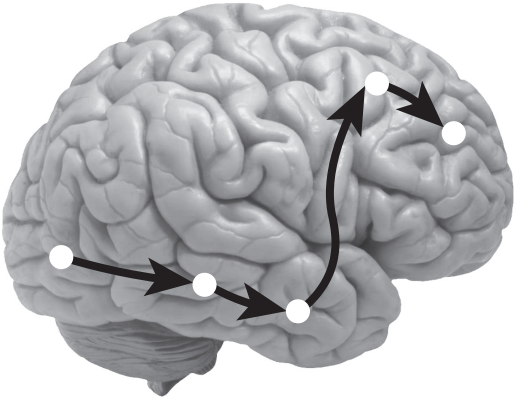

Figure 7.5
Flow of information processing in the brain. In the scheme favored by neuroscientists until at least the late 1980s, sequential processing steps move from the early visual cortex to the prefrontal cortex, where “abstract processing” is assumed to take place.

cells respond to object shapes and complex properties including “face-ness.” Eventually, this processing stream reaches the prefrontal cortex. Within the prefrontal cortex itself multiple areas can be discerned, and as one moves toward the front of the brain, cells integrate multiple dimensions of the sensory world (they may respond to visual, auditory, and tactile properties simultaneously) in a manner that reflects the task at hand. Neurons care about features and objects that are behaviorally pertinent at that moment (for example, the color of a fruit matters when the animal is inspecting it for ripeness, not when navigating around it).

In this way, the information flows from sensory regions, which register the world “out there,” to an *apex of integration* sitting atop the brain. It is thought that this hierarchical scheme protects the low-level stages, the ones doing basic sensory processing, so that they can reflect the environment with verisimilitude. The external world is thus represented with the least distortion possible. Interactions of diverse signals (such as from multiple sensory modalities) and integration with other processes (say, expectations and memory) are kept for intermediate and later stages, thus insulating the initial registering of the world from bias.

A related narrative can be told from the perspective of motor actions, but here we need to reverse the progression of signals. Signals flow from the regions of prefrontal cortex where the most abstract response properties are found, to so-called premotor regions in the frontal cortex engaged by action planning, and eventually to primary motor cortex, still located in the frontal cortex, where motor commands are issued. In other words, the sequence is from thought to musculoskeletal movements.

The Predictive Brain

While the conceptual framework outlined in the previous section was dominant, an alternative view was espoused by a minority of researchers, on and off, since the early days of brain research in the nineteenth century. Initially at least, it wasn't based on anatomical information per se but on conceptualizing the brain as a *prediction* device. In the leading view, action follows sensation in a progression that eventually culminates in a motor act: A monkey sees a fruit, registers its properties and position, and reaches out to grab it if it seems ripe. The predictive framework flips this logic on its head: Perception is directed by action so that effective behaviors can be generated. According to the traditional view, vision is relatively passive, like a camera pointed at the world, clicking away. In the predictive framework, vision is active and guided by endogenous computations that try to anticipate the most valuable future information for the animal.

From the standpoint of the active approach, the flow of information in the brain *can't* be like in figure 7.5. There should exist both connections from a "lower" to a "higher" area *and* the reverse—an idea that receives overwhelming empirical support. For instance, area V1 receives major projections from the visual part of the thalamus and is accordingly termed "primary" visual cortex. We know that V1 projects to area V2, too, but in addition, V2 projects to V1. We can therefore consider the V1 → V2 connections as "feedforward" and the V2 → V1 ones as "feedback." More generally, feedback connections, which are abundant in the cortex, provide a mechanism for predictions to influence earlier processing. In fact, Stephen Grossberg, a theoretical neuroscientist and a pioneer of the field of artificial neural networks (and one of my mentors in graduate school), proposed in the 1980s that feedforward and feedback pathways are arranged in what can be considered a basic building block: Connections from "lower" regions are reciprocated such that what ensues is a "consensus" between bottom-up and top-down signals (figure 7.6). This type of bidirectional architecture has profound implications for our understanding of how the brain works. Signals don't flow just one way but bidirectionally—higher areas project to, and influence, lower areas.

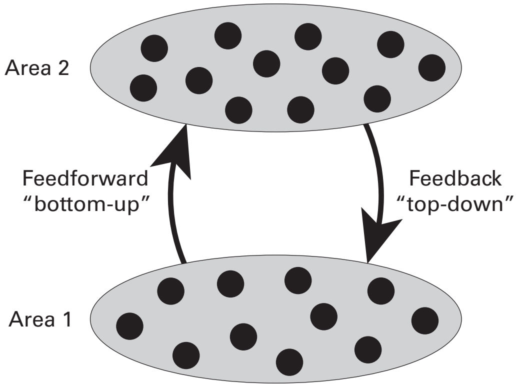

Figure 7.6
Bidirectional pathways between areas. The arrangement allows the two areas to mutually influence each other. One way to think about the pair is that one of them (here Area 2) provides "prediction signals" to the other.

The upshot is that brains are not passive. When a stimulus is processed, it does not encounter a tabula rasa. Instead, it is registered against a host of expectations constructed from prior experience, leading to the idea of a *matching* process between incoming information and feedback "template" signals. The template represents the system's predictions of the input and is updated to reflect the animal's past history.[9] Despite differences, predictive-brain approaches share a key concept: the brain doesn't reconstruct the external world but *constructs* a version of it.

Let's Do This All Together

In a critical assessment of the literature, the neuroanatomist Patricia Goldman-Rakic summarized the state of knowledge by the end of the 1980s as follows:

> The conclusion traditionally reached in virtually all comprehensive studies of cortical connections is that they are organized in a step-wise hierarchical sequence proceeding from relatively raw sensory input at the primary sensory cortices through successive stages of intramodality [that is, specific to one of the senses] elaboration allowing progressively more complex discriminations of the features of a particular stimulus. (Goldman-Rakic 1988, 146)

This statement is, of course, aligned with the traditional feedforward framework described in this chapter. But the time was ripe for different types of organization to be entertained. For one, computer science was evolving rapidly, and the centralized hardware organization of computers was changing to allow distributed processing. Instead of the architecture with a central processing unit, or CPU, parallel computers were designed with a large number of processing elements (some machines had thousands of them). A second impetus came from the field of artificial neural networks, which resurfaced quite strongly at the time. Although neural network models containing large numbers of artificial neurons—which could be thought of as simple processing units—were developed in the 1940s and studied by a small but very active group of investigators, the field of "connectionism," as some called it, only flourished in the mid-1980s. Neuroscientists took notice and started to think in more flexible ways.

Viewing the brain as a distributed system emerged as a guiding principle, as persuasively described by Goldman-Rakic in the paper quoted above. In a centralized architecture, signals converge at a specific element (the CPU, say), where they are further elaborated. Conceptually, this is the most straightforward type of arrangement one can imagine. Given that signals are all present at one location, they can be manipulated and operated on to generate an output. The distributed architecture, instead, carries its operations in a spatially scattered manner—and even the "results" themselves may be decentralized (figure 7.7). To illustrate with an early example, the neurologist Marsel Mesulam proposed that the parietal, frontal, and cingulate cortex work together to implement "attentional functions" (Mesulam

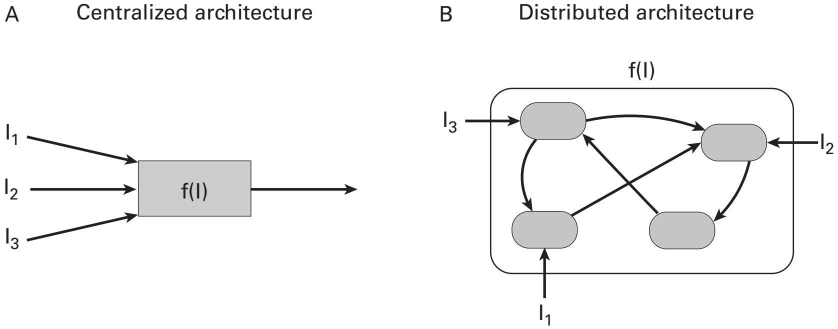

Figure 7.7

Centralized vs. distributed computation. (a) In a centralized system, a function **f()** is computed by a processing element based on its inputs, generating an output. (b) In a distributed system, multiple basic processing elements interact such that the function **f()** is carried out by the system as a whole.

1981). That is to say, mental functions that we call “attention” rely on joint contributions from regions within these three cortical sectors to carry them out. What’s more, although working together, they provide distinct contributions to attentional processing, with the parietal, frontal, and cingulate cortices most important for sensory, motor, and motivation, respectively. Thus, while composed of multiple parts, the three areas of the attention circuit aren’t exchangeable.

Coda

Let’s recap the path taken in this chapter. Our goal was to describe some of the mechanisms that support cognition, with a particular focus on executive functions. Cognition reflects processing that is less tied to the sensory world. Distinct dimensions of the same stimulus can be relevant based on the context at hand, and the system needs to be updated on the fly, always at the ready. Simplistically, we can think of sensory signals as flowing from the periphery to the central headquarters (the prefrontal cortex), where information is manipulated and put together in complex ways. This transformation of perception into cognition is supported by an anatomical architecture that takes signals, step by step, from the sensory cortex to the prefrontal cortex. This view was indeed favored by most neuroscientists until at least the middle of the 1980s.

But another mode of communication, one based on feedback connections, must be considered, too. Abundant *bidirectional* connectivity fosters a view that processing is as much about exogenous as about endogenous signals, leading to an active, predictive system. Furthermore, parallel pathways are capable of conveying information in a distributed manner, creating an elaborate anatomical infrastructure that can support nonsequential and decentralized mechanisms. Thus, cognition is the product of much richer and nuanced mechanisms than a piecewise building-up process. In fact, in chapter 10, we’ll describe how large-scale cortical-subcortical interactions are essential for constructing cognition and melding it together with emotion and motivation. Before doing so, let’s discuss the concept of complex systems in chapter 8. To understand how the central nervous system supports sophisticated mental functions, we need to have a better grasp of *systems thinking*.

The Entangled Brain

How Perception, Cognition, and Emotion Are Woven Together

By: Luiz Pessoa

Citation:

The Entangled Brain: How Perception, Cognition, and Emotion Are Woven Together
By: Luiz Pessoa
DOI: 10.7551/mitpress/14636.001.0001
ISBN (electronic): 9780262372107
Publisher: The MIT Press
Published: 2022

The open access edition of this book was made possible by generous funding and support from MIT Press Direct to Open

The MIT Press

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Pessoa, Luiz, author.
Title: The entangled brain : how perception, cognition, and emotion are woven together / Luiz Pessoa.
Description: Cambridge, Massachusetts : The MIT Press, [2022] | Includes bibliographical references and index.
Identifiers: LCCN 2021061878 (print) | LCCN 2021061879 (ebook) | ISBN 9780262544603 (paperback) | ISBN 9780262372107 (pdf) | ISBN 9780262372114 (epub)
Subjects: LCSH: Perception. | Emotions and cognition. | Brain. | Neuropsychology.
Classification: LCC BF311 .P3767 2022 (print) | LCC BF311 (ebook) | DDC 153—dc23/eng/20220411
LC record available at https://lccn.loc.gov/2021061878
LC ebook record available at https://lccn.loc.gov/2021061879