




coeruleus in the pons; the serotonergic raphe nuclei; and GABAergic neurons in the VLPO of the hypothalamus. All of these nuclei exert both direct and indirect effects on the overall cortical activity that determines sleep and wakefulness. Figure 28.12 summarizes the relationships among the various sleep–wake states.

Consciousness

Although being awake is clearly a prerequisite to being conscious in the sense of being normally aware of the world and the self, these functions are not equivalent. Most definitions of consciousness thus refer to two different aspects or meanings of the term: a physiological meaning that describes consciousness in terms of the brain state we think of as wakefulness, and a more abstract meaning that refers to a subjective awareness of the world—a brain state that must have a more subtle signature than wakefulness, since one can be awake and yet be unaware of some or even most aspects of the external and internal environments.

This latter meaning of consciousness refers not to the neuroanatomical and physiological bases of wakefulness, but to the deeply puzzling ability we all have to be subjectively aware of the world and of ourselves as actors in it. Whereas consciousness as wakefulness and its physiological basis fit easily in the conventional framework of neurobiology, consciousness as awareness, or self-awareness, raises more difficult and contentious philosophical issues. It has thus been far more difficult to investigate consciousness with the techniques used in cognitive neuroscience or, for that matter, with any other methods.

Investigating consciousness is further complicated by the fact that being *aware* is not the same as being *self-aware*: one can imagine an animal that is aware of sensory input without having the integrated sense of a separate self that humans possess. Whereas laboratory animals such as rats


FIGURE 28.12 A summary scheme of sleep–wake states. In the waking state, activation is high, modulation is aminergic, and the information source is external. In REM sleep, activation is also high, the modulation is cholinergic, and the information source is internal. The other states can likewise be remembered in terms of this general diagram. (After Hobson, 1989.)

and pigeons seem aware of the world, whether they are self-aware is debatable. However, some non-human primates seem to be both aware of the world and (at least to some degree) aware of themselves as actors in it, although the level, or even existence, of such awareness in controversial. Going further down the scale of brain (or nervous system) complexity, very simple animals presumably lack awareness altogether, at least in any conventional sense. Indeed, the most neurally simple animals would appear to operate as automatons, although it is not clear at what taxonomic level this occurs. Moreover, it is important to note that even in the human nervous system most neural processing is largely automatic, operating below the threshold of awareness. Think, for example, of all the homeostatic neural mechanisms that insure your well-being in innumerable ways even while you are pondering the meaning of this sentence.

These daunting issues notwithstanding, several cognitive neuroscientists have sought to address the basis of consciousness by ferreting out some signature of the neural processing that occurs when we are aware of something (e.g., a visual stimulus) compared with when we are not (e.g., presentation of the same stimulus under circumstances in which it does not elicit a reportable percept). Studies of both normal subjects and patients with various pathological conditions have shed some light on this challenge. In normal subjects, the approach most commonly taken has been to assess the nature and location of neural activity while a particular sensory percept moves in and out of awareness. By asking the person (or experimental animal) to report these perceptual transitions (typically verbally or by a button press), the investigator can compare neural activity during awareness in response to a stimulus with that when the subject is unaware of the stimulus. In such paradigms, the physical stimulus remains unchanged and thus serves as its own control. One paradigm for this purpose has been binocular rivalry. Binocular rivalry refers to the fact that when a particular stimulus is presented to one eye while a discordant stimulus is presented to the other, the visual percept is of either one stimulus or the other, and alternates back and forth every few seconds, rather than being a combination or blending of the views coming from the two eyes (Figure 28.13).

Using humans or monkeys trained to report what they are seeing at any given moment, electrophysiological methods and fMRI can assess changes in brain activity

FIGURE 28.13 Binocular rivalry. The perceptual phenomenon of binocular rivalry is illustrated here by the presentation of vertical lines to the left eye and horizontal lines to the right eye. A grid pattern is not seen, indicating that the views of the two eyes are not simply brought together in V1 by the activity of binocular neurons in the visual cortex.



that occur when there is a switch of conscious content. For example, when the monocular inputs are faces and houses, recordings of fMRI activity show increases in the fusiform face area of the temporal lobe when a face is perceived, and in the parahippocampal place area when houses are seen. Using similar paradigms, single neurons in at least some regions of the visual cortex in monkeys tend to show increased activity when the animal is perceiving the view with one eye but not the other.

A different fMRI paradigm with the same aim of probing activity changes in the low-level visual cortical regions during perception versus its absence takes advantage of visual aftereffects. For example, orientation aftereffects can be induced by exposing participants to a series of lines (referred to as *gratings*) in a particular orientation (say, 45°) for a minute or two; following this exposure, “neutral” line stimuli (vertical lines) are perceived for several seconds as being slightly tilted in the direction opposite the angle of the inducing exposure. The same inducing stimulus can be presented without awareness by masking it during the presentation. One can then ask whether lack of awareness of the inducing stimulus abolishes the aftereffect. It does not, implying that the low-level visual cortical neurons sensitive to orientation are about as active when subjects are aware of the inducing stimulus as when they are not. This lack of effect has been further confirmed in non-human primates by examining the activity of depth-sensitive (disparity-tuned) neurons in the primary visual cortex (V1) that are active in response to stereoscopic stimuli. Again, little or no difference in activity is apparent when the monkeys indicate behaviorally that they experienced a perception of depth compared with when they do not. In short, studies of activity in at least the earliest stages of visual cortical processing have not led to any clear hallmark of visual awareness.

It may help in rationalizing these somewhat confusing results to distinguish lower-order visual processing from higher-order visual processing in the visual association cortices. Recall from Chapter 12 that the higher-order visual areas are broadly divisible into a ventral stream running to the temporal lobe that is concerned with the recognition of visual stimuli (the “what” pathway), and a dorsal stream that is more concerned with stimulus location and spatial

relations (the “where” pathway). Correlations of neural activity with stimulus awareness in monkeys tend to increase as recording sites move away from V1 along the ventral pathway, toward the temporal lobe. This interpretation accords with the majority of the studies mentioned above, showing that awareness of objects such as houses and faces activate the relevant cortical regions.

A further point is that the activity of some cortical regions not generally thought of as visual are also correlated with visual awareness. For example, neuroimaging studies in humans have shown that perceptual changes in binocular rivalry and other bistable-image paradigms are associated with activity in frontal and parietal cortical regions (i.e., there is activity in these cortical regions that is time-locked to the subjects’ reports of the perceptual changes). Consistent with this evidence, other perceptual effects such as “pop-out” of a visual stimulus, or the awareness of a previously missed stimulus feature, are also correlated with activity changes in frontal and parietal areas. Transient functional disruption of processing in these areas by transcranial magnetic stimulation (TMS) also disturbs perception. Disturbances in perception following damage to these higher-level brain regions further confirm that they are involved in perceptual awareness and consciousness, although their roles in these phenomena remain unclear.

Yet another idea concerning possible neural mechanisms underlying awareness involves a longer-latency return of activity back to the relevant sensory processing regions, but much later in time than the initial feedforward cascade. This longer-latency activity, termed **reentrant** or **recurrent neural activation**, has specifically been proposed as a mechanism leading to perceptual awareness. Because of the sluggishness of the fMRI BOLD signal (see Chapter 1), this method is unable to delineate such recurrent activity. EEG studies, however, have suggested that recurrent activity in visual processing regions correlates with reported awareness of an attended visual stimulus, although the tightness of the link between the activity and awareness is still unsettled.

Considered together, this evidence is consistent with the generally accepted idea that awareness of responses to stimuli (i.e., perception and thinking) is based on altered

activity of populations of cortical neurons in the regions of the association cortices that process the relevant stimuli, integrating the result with information arising from other modalities and contextual influences. With respect to vision, the observed correlations described here suggest that activity in the visual association cortices is *necessary* for visual awareness, and that continued activity in these areas may be important. However, such activity does not appear to be a *sufficient* cause of awareness, and no defining neural signature of awareness has been discerned.

Clinical evidence has also contributed to understanding the neural basis of awareness and consciousness. A pathological phenomenon of particular interest is **blindsight**. As described in Chapter 12, patients with damage to the primary visual cortex are blind in the affected area of the contralateral visual field (recall that as a result of the visual system anatomy, V1 in the right occipital lobe processes information arising from the left visual field, and V1 in the left occipital lobe processes input from the right visual field). Objects presented within the area of blindness in the visual field (i.e., the scotoma) are not seen, at least according to the patient's verbal reports. Nevertheless, when some blindsight patients are forced by the experimenter to make a response to simple stimuli presented within their scotoma, the responses are often significantly above chance performance. For instance, if patients are asked to guess whether a stimulus line within their scotoma is vertical or horizontal, they answer correctly much of the time even though they claim to have seen nothing. Blindsight can also be simulated in normal individuals by transient inactivation of V1 by TMS applied over the occipital lobes. TMS creates a temporary scotoma for a specific region of the visual field, and again the features of simple stimuli presented in the unseen region tend to be guessed correctly at levels well above chance.

Functional neuroimaging and electrophysiological studies of patients with blindsight show that the unseen stimuli elicit some activity in extrastriate regions beyond V1, implying that these cortical areas are needed for successful behavior in the absence of awareness. This evidence underscores the conclusion above that extrastriate activity may be necessary for awareness, but is not sufficient. One proposed explanation of blindsight is that subcortical visual processing of the stimulus, or coarse subcortical input to extrastriate cortex that bypasses V1, influences the patient's guesses. This interpretation accords with other evidence that subliminal (unconscious) information processing influences behavior of all sorts.

Another example of neural pathology pertinent to understanding awareness is the experience of *split-brain patients* (see Chapter 33). When the corpus callosum is cut as a treatment for otherwise intractable epileptic seizures, direct communication between the right and left hemispheres is no longer possible. The perceptual

consequences of this surgery, first studied by Roger Sperry and Michael Gazzaniga in the 1960s, showed that the divided hemispheres in these patients function relatively independently, and that awareness generated by neural processing in one hemisphere is largely unavailable to the other. For example, when simple written instructions such as “laugh” or “walk” are presented visually to the left visual field—and thus to the right brain—of a split-brain patient, many subjects have enough rudimentary verbal understanding in the right hemisphere to execute the commanded action. However, when asked to report *why* they laughed or walked, they typically confabulate a response using the superior language skills in the left hemisphere, saying, for instance, that something the experimenter said struck them as funny, or that they were tired of sitting and needed to walk a bit. Thus, the same individual would appear, under these circumstances, to harbor two relatively independent domains of awareness. This evidence raises the provocative question of whether awareness is really the unified function we generally take it to be, as well as what the role of the corpus callosum is in engendering such unity.

Whereas blindsight and hemispheric division lead to circumstances in which patients are unaware of stimulus processing that nonetheless influences their behavior, it is also possible to be aware (or to believe one is aware) of something that doesn't actually exist. Perhaps the most striking illustration of this sort of phenomenon is the *phantom limb* experiences of amputees described in Clinical Applications, Chapter 10. Recall that a common experience following amputation is the patient's subjective awareness of the missing arm or leg, despite the fact that the physical limb and its peripheral sensory input are absent. Although the interpretation of this bizarre condition is debated, the awareness of the missing limb and sensations arising from it seem quite real to the patient (pain is especially problematic for some amputees), emphasizing that processing of peripheral information is an active process in which the cortex constructs the percepts we experience. Hallucinations make the same point.

The Neural Correlates of Consciousness

Thus, efforts to identify a neural basis for awareness and consciousness have not provided the sort of answers sought by many investigators in this field. The neurological aspects of consciousness described in the previous section are insufficient to satisfy philosophers, theologians, and many neuroscientists interested in the broader issues that the phenomenon of consciousness raises. Surely a phenomenon as important to us as consciousness should have a prominent signature. A prominent correlate would answer a number of questions such as whether other animals

are conscious in the same way we are, and whether machines could ever be self-aware.

With respect to the first of these issues, despite a long-standing debate about consciousness in other animals, it would be unwise to assert that humans are alone in possessing this obviously useful biological attribute. However, from a purely logical vantage it is impossible, strictly speaking, to know whether any being other than one's own self is conscious; as philosophers have long pointed out, we must inevitably take the consciousness of others on faith (or at least on the basis of common sense). Nonetheless, it is reasonable to assume that animals with brains structured much like ours (other primates and, to a considerable degree, mammals in general) have, in some measure, the same ability to be self-aware as we do. The ability to reflect on the past and plan for the future, made possible by self-awareness, is surely an advantage that evolution would have inculcated. A reasonable supposition would be that consciousness is present in animals in proportion to the complexity of their brains and behaviors—particularly those behaviors that are sophisticated enough to benefit from reflecting on past outcomes and future eventualities.

The question of whether machines can ever be conscious is a much more contentious issue. If one rejects dualism—the Cartesian proposition that consciousness, or “mind,” is an entity beyond the ken of physics, chemistry, and biology and therefore is not subject to the rules of these disciplines—it follows that a structure could be built that either mimicked human consciousness (by being effectively isomorphic with brains) or achieved consciousness using physically different elements (e.g., computer elements). There is no logical reason to assume that such constructions could not be conscious,

Although some contemporary scientists continue to believe that neurobiology may soon reveal the *basis* of consciousness, such revelations are not likely. A more plausible scenario is that as information grows about the nature of other animals, about computers, and indeed about the brain, the question “What is consciousness?” will fade from center stage in much the same way that the question “What is life?” (which stirred up a similar debate for much of the twentieth century) was asked less and less frequently as biologists and others recognized it to be an ill-posed question that admits no definite answer. We can already create many aspects of living organisms in the laboratory, and there seems to be no reason in principle why artificial neural networks capable of self-awareness should not someday be possible.

Comatose States

An important pathological condition pertinent to consciousness is *coma*. **Coma** (Greek for “deep sleep”) is the term applied to the brain state of individuals who have suffered brain injury that leaves them in a deeply unconscious state defined by apparent unresponsiveness to sensory

stimuli. The condition typically involves compromised function of the brainstem and other deep brain structures, such that the normal interaction of these structures with the cerebral cortex is interrupted.

Coma can arise from varying degrees of brain damage, and the prognosis is often uncertain. Most comatose individuals recover consciousness within a few days or weeks as the compromised neurons and the circuits they contribute to gradually regain their functions. Impairment can persist for much longer, however, if neural damage is more profound. Some patients recover consciousness after months, although typically with residual effects, and extremely rare cases have been reported in which consciousness is regained after some years. This variability has led to social, religious, and political controversy over the point at which an unresponsive patient should be considered to be in a **persistent vegetative state**, a diagnosis that raises ethical issues surrounding decisions about withholding life-support. This sensitive issue means there has been, and will continue to be, interest in techniques that could contribute to a better understanding of a given patient's prognosis. Electroencephalography has been fundamental in diagnosing *irreversible brain death*, which occurs when brain trauma is so severe that no EEG activity can be recorded (i.e., a flat electrical trace). More recently, functional neuroimaging has also been used to evaluate persistent vegetative state, sometimes with surprising results. For example, when the brain of a 23-year-old woman who had been uncommunicative for 5 months following a traffic accident was functionally imaged, fMRI showed that her brain responded to simple verbal commands by generating appropriate patterns of activity.

The “Default State” of the Brain

A final question in this chapter concerns the purpose of high levels of activity in many brain regions when people are awake and aware, but without pursuing any particular goal, thought, or task. On the basis of neuroimaging studies, neuroscientists have come to call this the *default* state of brain activity. Some brain regions, including the posterior cingulate cortex, the ventral anterior cingulate cortex, the medial prefrontal cortex, and the cortex at the junction of the temporal and parietal lobes, show consistently *greater* activity during resting states than during the processing of specific cognitive tasks. When the activity during a baseline or rest condition is subtracted from the activity during cognitive tasks, relative decreases of activity (generally called *deactivations*) are observed in these areas. Because cognitive tasks induce activation increases in other brain areas carrying out the relevant processing, this finding led to the proposal that these regions constitute a network supporting a default mode of brain function that is engaged in the *absence* of any particular cognitive




FIGURE 28.14 Default-mode network activity is anticorrelated with frontal-parietal attentional control regions.

(A) fMRI data were collected from a subject at rest, and one part of the default-mode network, the posterior cingulate cortex (PCC), was chosen as a seed region to help identify other parts of the network. The time course of activity in the seed region was correlated with every other part of the brain, and those correlation coefficients are shown plotted on an inflated representation of the brain surface. Using this simple method, the rest of the default-mode network shows up as warm-colored, indicating a similar activity profile as that of the PCC, whereas areas of the frontal-parietal attentional control network show up in cooler colors, indicating an anticorrelation with the PCC. (B) fMRI activity time courses over one run are plotted from two regions of the default network (PCC and medial prefrontal cortex, or MPF) and from one region of the attentional network (intraparietal sulcus, or IPS). The time courses clearly show that when activity in one network goes up, activity in the other goes down. (From Fox et al., 2005.)

task (Figure 28.14). Further supporting the hypothesis that these areas constitute a network, analysis of the functional connectivity of these areas (i.e., how much their activity covaries across time) has shown strong coupling during the resting state.

Single-unit recording in monkeys has supported the inverse activation pattern of the default-mode network. For example, neurons in the posterior cingulate cortex, a posterior brain region considered to be part of the default-mode network, show changes in firing rates that closely match the fMRI activation patterns observed in humans. In particular, their fluctuations are anticorrelated with activity in dorsal parietal regions associated with the attentional control network and predict lapses in attention and the ability to switch from one task to another. Moreover, neuroimaging studies also have shown abnormal activity patterns in the default-mode network in several major neurological and psychiatric disorders, including being less active in autism and more so in schizophrenia.

The obvious question is what purpose neural activity in a default-mode network serves—that is, why should these regions be active if and when the brain is doing nothing in particular? Although the default-mode network activity

might be related to mental “idling,” another possibility is that this network is activated when attention is inwardly focused, the “standard” attentional control system being activated primarily when a person is focused on events and stimuli in the external environment (see Chapter 29). Whatever this network does, its activation pattern also occurs in monkeys, so it presumably evolved to carry out some relatively basic yet important function. The inverse relationship of activity in the default network and in the dorsal frontal-parietal attentional control network during focused attention suggests that these complementary systems may play an interactive role in system-wide brain function related to engaging and disengaging attention and other cognitive functions.

Summary

The most obvious cortical states are sleep and wakefulness. All animals exhibit a restorative cycle of rest following daily activity, but only mammals divide the period of rest into distinct phases of non-REM and REM sleep. Why humans (and many other animals) need a restorative phase of suspended consciousness accompanied by decreased metabolism and lowered body temperature is not known. Even