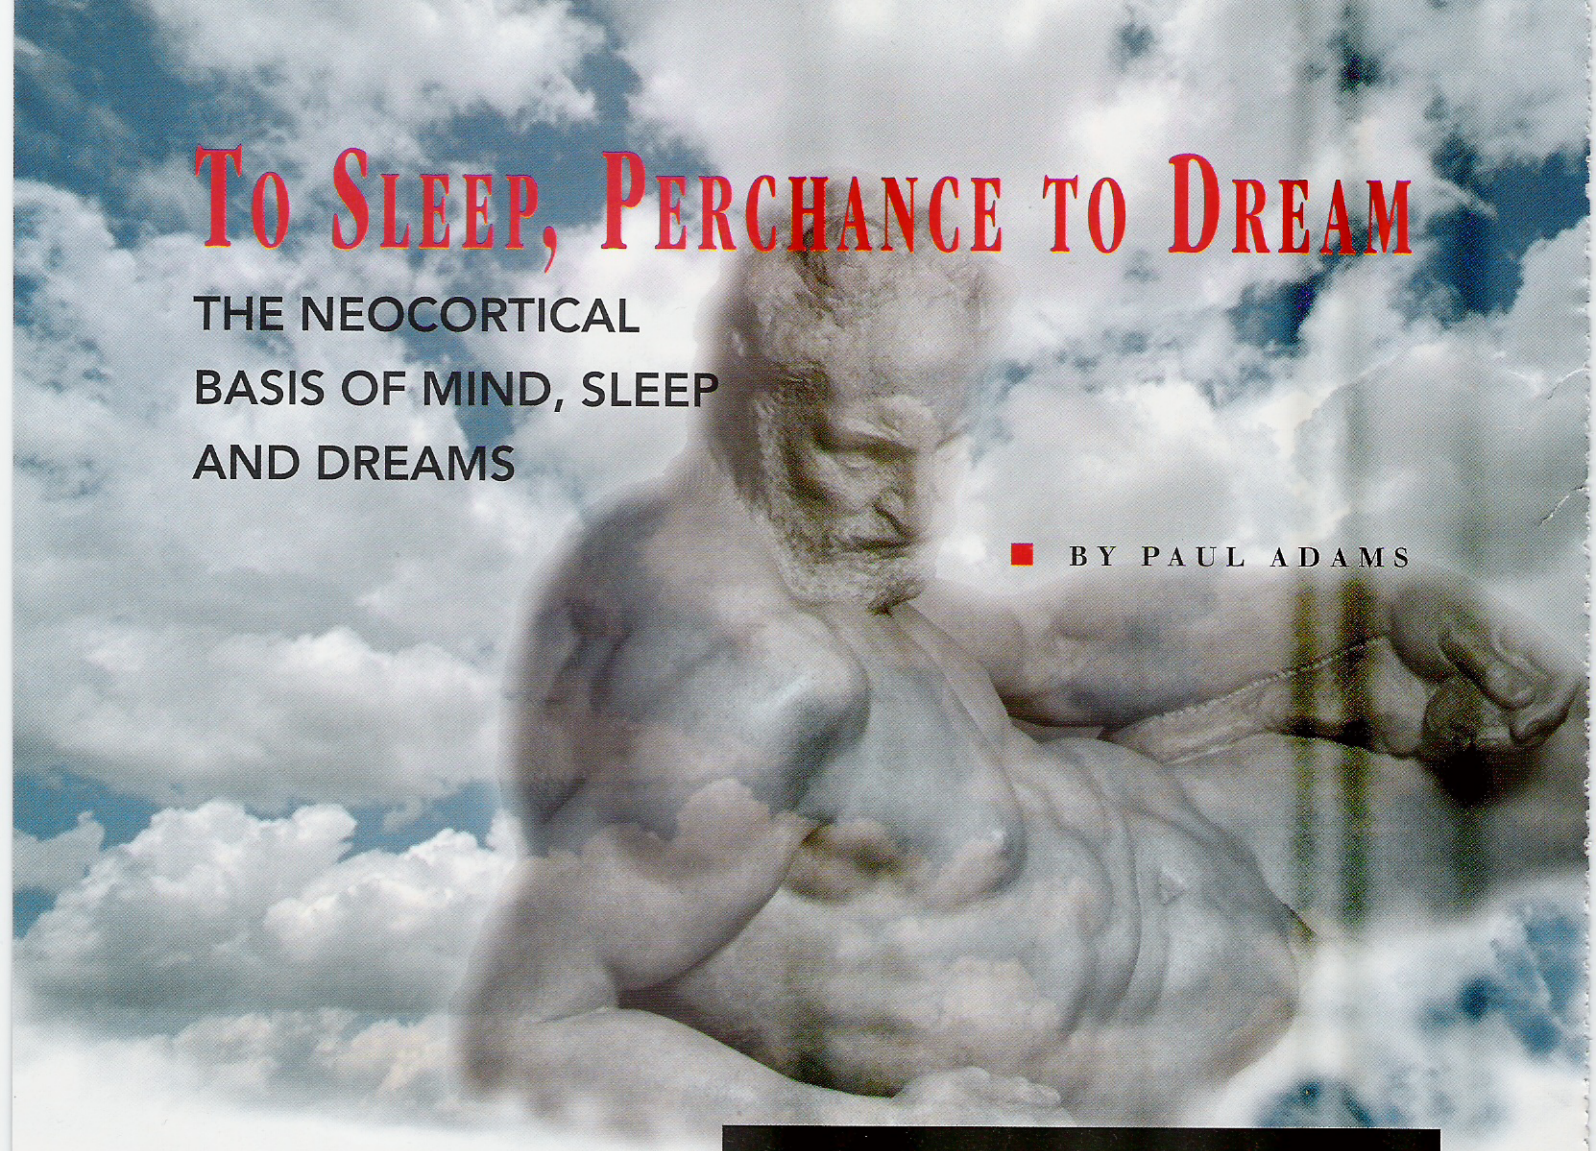

To Sleep, Perchance to Dream

THE NEOCORTICAL BASIS OF MIND, SLEEP AND DREAMS

BY PAUL ADAMS

How does the cerebral cortex work? How are neural connections made? Why do we sleep? New concepts and recent data suggest these three questions may be intimately related. Paul Adams, Professor of Neurobiology at SUNY Stony Brook, shows how these challenging questions in neuroscience are gradually yielding to experimental and theoretical advances.

The cerebral cortex, which occupies about 86 percent of the human brain, is believed to be the structure generating thought. It is composed of a crumpled sheet of neurons about 3 mm thick and 0.25 m^2 in area, together with several thousand km of associated cabling. Although different parts of the sheet are specialized for different tasks—and these areas have subtly different wiring—there appears to be a common blueprint, not only within the human brain, but in all mammals. This plan is sufficiently different from that of various other cortical neural structures (such as the olfactory, hippocampal and cerebellar cortices) that the cerebral cortex is also often called *neocortex*, in keeping with its relatively recent evolutionary appearance.

The computational principles underlying the neocortex must be very powerful to allow its tremendous expansion of functions, culminating in *Homo sapiens*. One principle is gradual, hierarchical processing: each cortical area both solves an immediate practical problem and provides abstract information for use by other cortical areas. For example, the primary visual cortex uses input from the eyes to generate output to subcortical areas involved in eye movement. In the cat visual cortex, the basic operation seems to

be the construction of a local orientation map: locally oriented visual signals identify the edges of objects and suggest useful places to direct the gaze. Progressively higher levels of the cortex then assign particular edges to the possible objects composing the scene. Much work on the cortex concentrates on the detailed local computations performed in each specialized area, together with the interrelations between the many areas. But, as noted earlier, there are strong similarities between all cortical areas and between all mammalian neocortices. The next two sections outline these similarities, with particular reference to the cat visual cortex, the best studied cortical area.

THE NEOCORTICAL PLAN

Cortical circuitry falls into two main categories: the easy half, in which the basic processing of input is fairly straightforward, and a more enigmatic half, perhaps involved not so much in "processing information" as in regulating the behavior of the conventional information-processing circuitry. This is analogous to a sophisticated economy in which much of the workforce is involved not in production but in administration.

Figure 1 shows how the easy half works. An egg-shaped and egg-sized structure at the center of the brain (the *thalamus*) supplies information to the neocortex. The thalamic egg contains numerous yolks, or nuclei, that transmit information to cortex from different subcortical or cortical sites. For example, the lateral geniculate nucleus relays data from the eyes to visual cortex at the back of the brain. Many thalamic signals, the so called "core" signals, arrive in the middle of cortex, typically in layer 4 (in Figure 1, the thalamus itself is designated as layer 0). The signals then pass through connections composed of numerous tiny devices called *synapses* (black dots in Figure 1), where electrical signals are transferred between cells by release and detection of neurotransmitter molecules. The number of synapses contributing to a connection determines its strength. The rather complicated connections from the thalamus to layer 4 cells transform the pattern of activity in a thalamic nucleus into a new activity pattern in the corresponding layer 4 cells in the cortex; the transformed pattern highlights certain features not evident in the original one. For example, when certain lateral geniculate cells receive input from a group of retinal cells responding to a short, oriented light or dark bar, the corresponding layer 4 cells act as local orientation detectors.

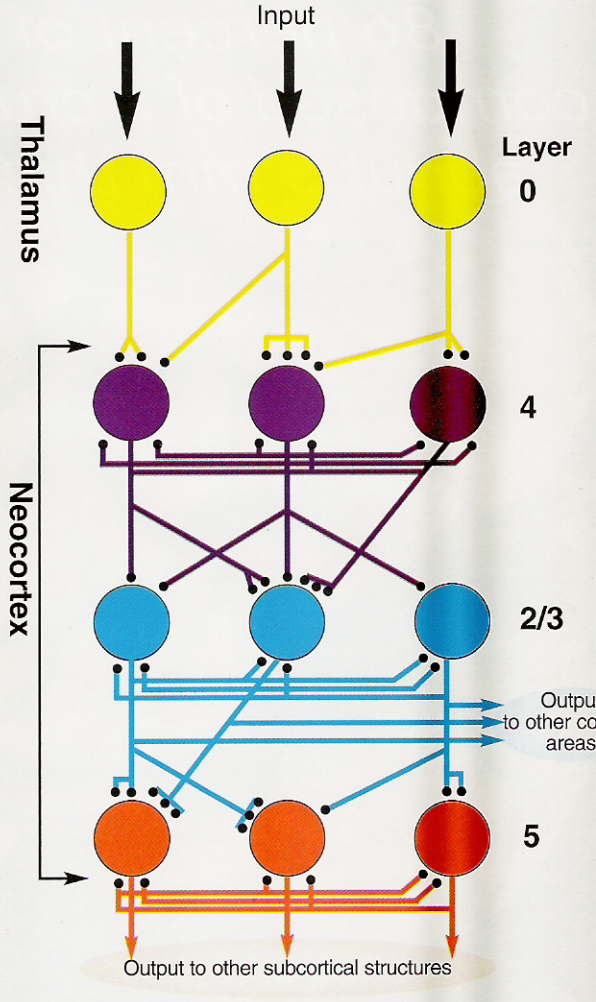

Figure 1: The easy half of the neocortex. Inputs arrive in thalamus (yellow cells, layer 0) from sources throughout the brain, but they undergo little processing there beyond application of "tonic" or "burst" labels that reflect the pattern of firing. Information is then sent to the middle layer of the neocortex (layer 4, purple cells) via synapses (black dots); the number, strength and arrangement of synapses determine the nature of the layer-to-layer pattern transformation. The transformed pattern is then transformed again at layers 2 and 3 (usually lumped together as 2/3) and 5. Within-layer feedback connections further refine the transformations at each cortical layer. Outputs to higher levels of the cortex, to new thalamic nuclei and to the rest of the brain are tapped off at layers 2, 3 and 5. The actual synaptic patterns shown here are for illustration only and have no particular significance.

The pattern of activity in layer 4 is then sent to layers 2 and 3 where it is again transformed via between-layer, feed-forward connections. In the visual cortex, edge detectors reporting a particular orientation in nearby regions of visual space converge onto particular layer 2/3 cells; as a consequence, those cells respond to oriented bars over a larger region of space than do the layer 4

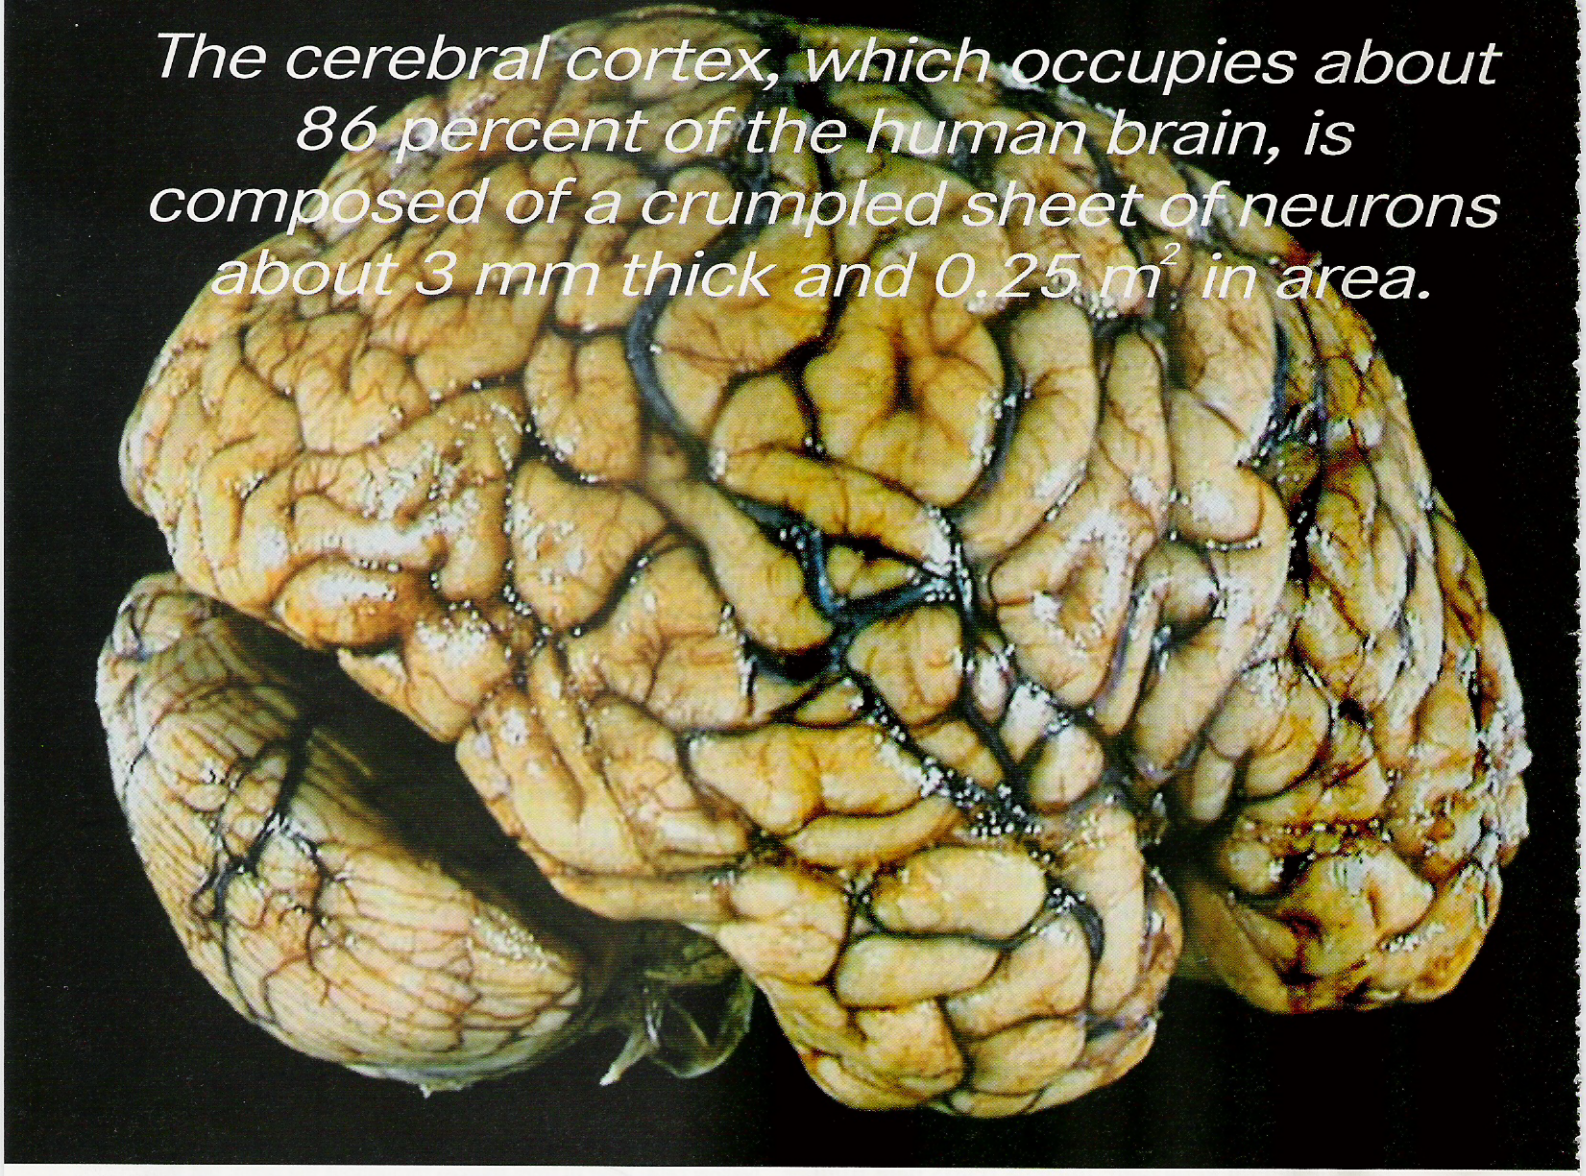

cells. In effect, the layer 2/3 "complex" cells generalize slightly the concept of local orientation initially developed in the earlier "simple" layer 4 cells. This new pattern then goes to layer 5, where it is further transformed and sent to two major destinations: to subcortical targets (such as eye movement centers) and to other thalamic nuclei that can send it on to higher order cortical areas for additional analysis.

This principle of layer-to-layer transformation via feed-forward connections is supplemented by within-layer processes. Cells within a given layer can excite each other, building up activity patterns by positive feedback. For example, cells in layer 2/3 of the visual cortex that respond to the same orientation in nearby regions of visual space tend to reinforce each others' activities. In layer 4, neighboring cells sharing orientation preference tend to excite each other. Thus, weakly oriented signals from previous layers nudge the network to select a single orientation from multiple, perhaps ambiguous, alternatives. Such forced, snap decisions in response to noisy—and possibly conflicting—data may save an animal's life. This within-layer positive feedback process is sometimes referred to as *recurrent* or *attractor dynamics*. It is likely, however, that the cortex can also suspend or revise snap judgments by switching off the lateral feedback and relying solely on the interpretation generated by the feed-forward connections.

The neural machinery itself is probably the main source of noise, hampering rapid cortical judgments. Sequences of electrical pulses called *spikes* send signals from cell to cell. Neural signals are noisy because the spikes are generated by the random opening and closing of ion channels in the nerve cell membrane, and neurons cannot always

precisely control their timing. Noise can be reduced by averaging over either time or space. The lateral feedback process combines signals present simultaneously in a spatial array of neurons. If it is switched off, temporal averaging of signals in individual neurons is required. The cortex could maximize information throughput by making continuous on-line adjustments to the recurrent dynamics on the basis of a preliminary analysis of incoming data.

THE ENIGMATIC HALF OF THE NEOCORTEX

The first puzzling feature of the cortex is already apparent in the scheme shown in Figure 1: the thalamus seems to relay information largely unchanged to synapses on layer 4 cells, i.e. it does little processing itself. Nevertheless, the thalamic relay is vital, not only because almost all information is squeezed through this portal but also because thalamic circuitry is very complex. A shaft of light into this obscurity has come with the realization that thalamic relay cells emit spikes in two quite different ways, called *burst* and *tonic firing*. In the former mode, spikes occur in tight clusters separated by silent periods while in the latter mode spikes occur singly, with intervals that depend on the strength of the input to the relay cells. However, the quantifiable information that the thalamic cells convey about the activity of their inputs appears to be similar in both modes. This suggests the mode of firing conveys a label to the

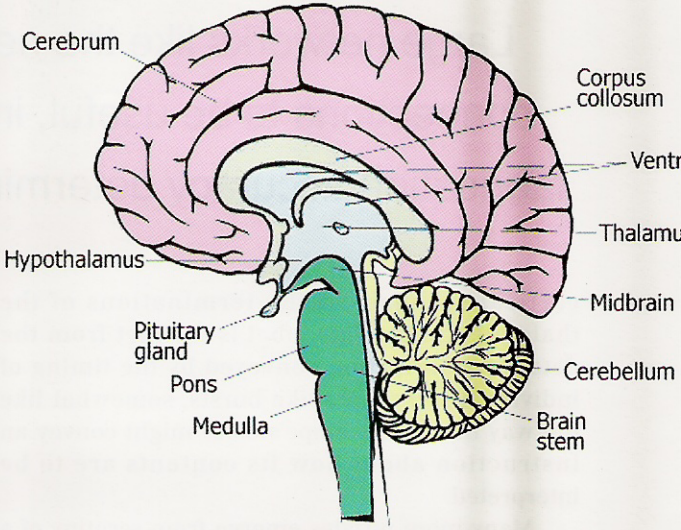

The principle parts of the brain.

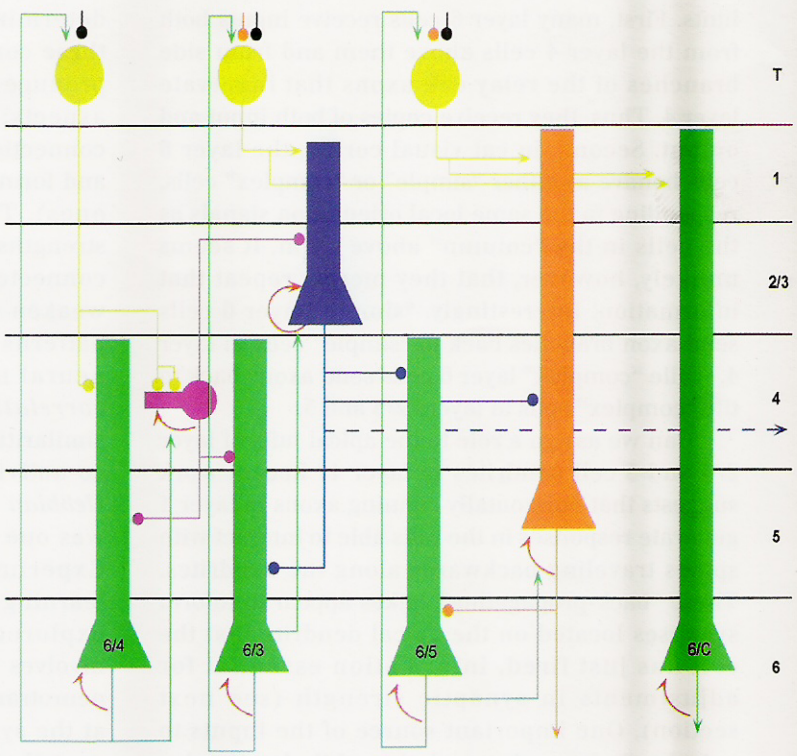

Figure 2: A more complete wiring diagram for the cortex. In this scheme layers 4, 2/3 and 5 have been collapsed to single cells, with red arrows indicating within-layer feedback connections. Cells in layers 2, 3, 5 and 6 have a pyramidal shape, shown here schematically as triangles with long, ascending, cylindrical dendrites. Four types of layer 6 cell are shown: simple-type layer 6 cells (marked 6/4) receive input from layer 4 simple cells; complex layer 6 cells (marked 6/3 or 6/5) receive input from complex cells in layer 2/3 or 5. These inputs are shown on proximal dendrites to represent the correlation-detection role of layer 6 cells. These three types of cell project back to thalamus (layer 0) and do not have apical tufts in layer 1, but a fourth type of layer 6 cell (6/C) projects to claustrum, a structure of unknown function that is associated with the neocortex, and ascends to layer 1. Note that 6/4, 6/3 and 6/5 cells also receive input from the layers that provide input to their defining input layers (i.e. 6/4 cells from thalamus, 6/3 cells from layer 4 and 6/5 cells from layer 2/3). These are placed distally to facilitate computation of correlations. Interactions between layer 6 cells (partly shown as red arrows) allow computation of correlation *ratios*. "Matrix" thalamic relay cells (yellow) provide input to layer 1. This input interacts with spikes that back-propagate along the main apical dendrites of layers 2, 3, 5 and 6/C cells, and may control postsynaptic plasticity.

Large networks like the neocortex require very precise connections to be useful, in much the same manner that fabrication accuracy determines the size of computer chips.

cortical synapses at the terminations of the thalamic axons. This label is distinct from the detailed information conveyed by the timing of individual spikes or spike bursts, somewhat like the way that an envelope's color might convey an instruction about how its contents are to be interpreted.

Many more puzzles emerge from scrutiny of a more complete wiring diagram for the cortex (Figure 2). First, the cortex has two additional layers, 1 and 6, excluded from the main flow path shown in Figure 1. The layer 6 cells have long apical branches, or *dendrites*, extending into layers 2-5 but mostly failing to reach layer 1. Many cells in layers 2/3 and 5 do send apical dendrites up to layer 1, where they terminate as apical tufts. Layer 6 cells typically send axons back to thalamus, where they impinge on the distal dendrites of relay cells. The activity of these layer 6 cells switches the relay cells from burst mode to tonic mode.

What activates layer 6 cells? There are several hints. First, many layer 6 cells receive inputs both from the layer 4 cells above them and from side branches of the relay cell axons that innervate layer 4. Thus, they receive copies of both input and output. Second, in cat visual cortex, the layer 6 cells behave as either "simple" or "complex" cells, responding to the same local orientation signals as the cells in the "column" above them. It seems unlikely, however, that they merely repeat that information. Interestingly, "simple" layer 6 cells send axon branches back to "simple" cells in layer 4, while "complex" layer 6 cells send axons back to the "complex" cells in layers 2/3 and 5.

Can we assign a role to the apical tufts of layer 2/3 and 5 cell dendrites in layer 1? Recent work suggests that horizontally running axons in layer 1 generate responses in the tufts able to interact with spikes traveling backwards along the dendrites. These "back-propagating" spikes appear to inform synapses located on the apical dendrite that the cell has just fired, information essential for adjustments in synaptic strength (see next section). One important source of the inputs to apical tufts is a sub-population of thalamic relay cells (referred to as the "matrix") that projects to layer 1; layers 4 and 6 are the targets of the "core" relay cells. The matrix comprises a minority of relay cells in the lateral geniculate nucleus but some thalamic nuclei are dominated by these superficially terminating matrix cells.

This elaborate additional circuitry, involving various cortex-thalamus-cortex loops, is obviously essential for the correct operation of the neocortex—indeed, it is its most characteristic feature—but there is no clear consensus about the role it serves. One hypothesis is that it is involved in controlling the "plasticity" of the feed-forward synapses that operate in the easy half of the cortex.

HOW ARE NEURAL CONNECTIONS FORMED?

The detailed pattern and strength of the connections between the main processing layers of the cortex, e.g. from thalamus to layer 4, from layer 4 to layers 2/3 and from layers 2/3 to layer 5, determine neocortical computations. How are these connections and synaptic strengths set to produce useful results? This problem has two aspects: setting the strengths of existing connections, the focus of most research to date, and forming new connections (or eliminating old ones). Theoreticians have shown that, if the strengths of connections grow whenever the two connected cells fire spikes simultaneously and weaken when they fire asynchronously, useful patterns of strengths emerge. This is because neural networks are primarily interested in *correlations* between signals, which express similarities between patterns. Correlation-based, co-activity-dependent adjustment is known as *Hebbian learning* in honor of Donald Hebb, who was one of the first to formulate the principle. Experimentalists have shown that Hebbian learning does take place in the brain and have explored its cellular and molecular basis. It involves a simple but elegant trick performed by neurotransmitter receptors, key protein molecules at the synapses. Coincident activity of neurons contributing to a synapse first recruits existing

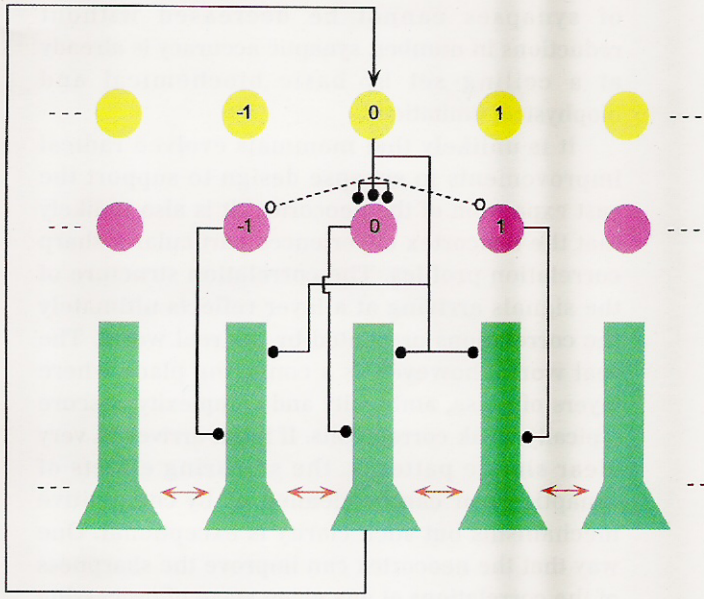

Figure 3: A circuit that guarantees accurate connections. This diagram provides an interpretation of some of the circuitry shown in Figure 2. The solid circles show an existing relay cell-to-layer 4 connection comprised of three synapses. The open circles represent errant or "mutant" synapses that could form on neighboring cells as a result of errors in the strengthening of the existing connection. The cell in the center of layer 6 (bottom row) computes the correlation between the relay cell and its existing target cell in layer 4 (middle row); the distal and proximal placement of the synapses on the layer-6-cell apical dendrite facilitates this computation. The neighbors of this layer 6 cell compute the correlation between the relay cell and the neighbors of the existing target in layer 4. Interactions between layer 6 cells (partly via inhibitory neurons, which are important in every layer, but are not shown) allow the middle layer 6 cell to compute the "sharpness" of the correlation signal at the existing layer 4 connection. If this sharpness is large enough , the layer 6 cell sends spikes back to the appropriate relay cell, shifting it from burst to tonic mode and enabling the plasticity of its layer 4 connections. A similar principle applies to the 6/3 and 6/5 cells, except that they control plasticity postsynaptically, possibly by feeding back to "matrix" relay cells, which, in turn, synapse on the apical tufts of postsynaptic layer 2/3 and 5 neurons.

receptors to the postsynaptic membrane and then stimulates protein synthesis and structural elaboration, culminating in the appearance of new synapses.

Hebbian learning works very well in setting the strength of existing connections and, if each cell in a network is connected to every other cell, setting the strength of those connections is straightforward. However, complete interconnectivity is unlikely in most networks: if every cell in the brain were physically wired to every other, our brains would be bigger than the Earth. So it is usually assumed that existing connections continuously produce local sprouts, which form new connections, and such sprouting has indeed been observed experimentally. If the resulting networks are to be useful, the sprouting should be regulated by Hebbian learning. New sprouts could form either continuously, i.e. independent of neural activity, and then be pruned back by a Hebbian rule, or only in response to coincident activity. In the latter case, one could view the synapses formed at sprouts to be "errors" in the strengthening of existing connections: a new synapse resulting from coincident neural activity is usually added to the co-active connections but occasionally it appears at a neighboring cell instead. Such errors are inevitable in the promiscuous thickets of the brain.

The result is that connections spread out more than they would if there were no errors. This smearing of synaptic connections has both positive and negative consequences. It allows the testing of new connections but degrades the overall performance of the network. In the limit of high error rates, the network is fully connected but unable to perform any useful functions. The spreading of connections is curtailed by Hebbian strengthening or weakening of connections. For example, if a neuron forms new synapses onto neighboring cells, but their activity correlates only weakly with that of the cells involved in the pre-existing connection, the new synapses will tend to be eliminated. Thus the spread of connections depends both on the sharpness of the profile of correlations across the network and on the error rates.

Large networks like the neocortex require very precise connections to be useful, in much the same manner that fabrication accuracy determines the size of computer chips. At first glance there seem to be only two ways to ensure the required accuracy: minimizing error rates and imposing very sharp correlations. However, neither of these solutions is feasible for the cortex. The synapses in the neocortex are essentially like those in other, much less extensive brain regions such as hippocampus, which have many hallmarks of low error rates, e.g. features that ensure chemical signals cannot spread beyond the synapse itself. However, since errors ultimately arise from molecular noise, and since the volume and density

of synapses cannot be decreased without reductions in number, synaptic accuracy is already at a ceiling set by basic biochemical and biophysical limitations.

It is unlikely that mammals evolved radical improvements in synapse design to support the vast expansion of the neocortex. It is also unlikely that the neocortex experiences particularly sharp correlation profiles. The correlation structure of the signals arriving at a layer reflects ultimately the correlations presented by the real world. The real world, however, is a confusing place where layers of noise, ambiguity and complexity obscure typically weak correlations. If input arrives in very clear simple patterns, the smearing effects of synaptic error can be cleaned up by competitive mechanisms but such clarity is exceptional. One way that the neocortex can improve the sharpness of the correlations at any given stage of processing is to improve processing of signals at earlier layers. This effort, though useful, does not guarantee presentation of clear patterns at every step.

Although the neocortex cannot, in the short term, do much to sharpen the arriving correlations, it can measure them since it has access to the signals contributing to those correlations. If its measurements indicate that correlations are sharp, it can allow the existing connections to learn. If, however, the correlations are blunt, it can tell the connections to remain fixed. Learning-enabled connections are termed *plastic* while fixed connections are *implastic*. In the next section we see how this simple idea for limiting the effects of synaptic inaccuracy maps onto the enigmatic half of cortical circuitry.

THE NEOCORTICAL ALGORITHM

Consider the cortical connections formed by a single thalamic relay cell in layer 4. Suppose that the thalamic cell is connected to a single cortical target cell and that these cells have strongly correlated firing. If the synapses comprising this connection are plastic, the correlated firing will add synapses at the connection. Occasionally, however, the new synapses form onto neighbors of the cortical cell; these errant synapses smear out the connection. If the neighboring cell's firing correlates only weakly with that of the thalamic cell, the errant synapses will tend to be eliminated; but if the neighboring cell also correlates strongly with the thalamic cell, the errant synapses can flourish and smear the connection (and perhaps give rise to further errant synapses in flanking cells). This could be avoided by making the entire connection formed by the thalamic cell in layer 4 implastic: implastic connections do not gain synapses so they cannot gain errant synapses. The crucial parameter that should decide whether the connection is plastic or implastic is the relative strength of the correlations between the currently connected cells and between the thalamic cell and the neighbors of the currently connected cell (Figure 3). This *correlation ratio* measures the sharpness of the correlation profile. A layer 6 cell has the right connections to compute this correlation ratio because it receives input from:

- the layer 4 cell immediately above it,
- the thalamic cell(s) that project(s) to that layer 4 cell, and
- the correlation signals of neighboring layer 6 cells.

Note that, in this arrangement, a layer 6 cell fires only when it receives input from both thalamus and layer 4. This is consistent with evidence that layer 6 cells echo the responses of simple or complex cells in the column above them.

Given the prevailing error rates, if the correlation ratio computed by a layer 6 cell is strong enough to eliminate errant synapses, that cell should send a batch of spikes back to the thalamic neuron and make all its synapses in layer 4 plastic. If, on the other hand, the correlation ratio is weak, the layer 6 cell must remain silent so that the corresponding thalamocortical synapses in layer 4 remain implastic.

This arrangement ensures that the accuracy of connection formation remains at any desired level. Although anatomic and physiological observations suggest that this is indeed how layer 6 cells operate, it is unclear how the burst-tonic transition, triggered by the activity of layer 6 cells, actually causes the switch from implasticity to plasticity. Correlation ratios computed by layer 6 cells also can be used for other neocortical tasks such as modulating the operation of the recurrent connections in layers 2, 3 and 5. This may be the function of the layer 6 cell intracortical branches mentioned previously, which form synapses quite different from those of feed-forward and lateral connections and that may well have a modulatory function.

Figure 3 shows the circuitry needed to control the plasticity of relay cells innervating layer 4. Similar arrangements would also be needed for the subsequent, purely intracortical feed-forward synapses. In the latter case, it is likely that plastici-

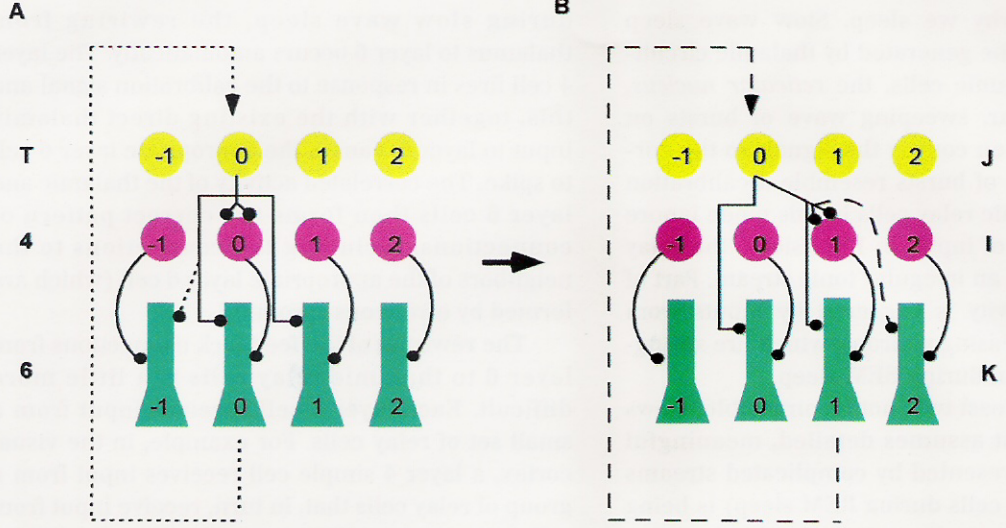

Figure 4: Off-line updating of layer 6 connections. Part A repeats Figure 3 but shows connections that must be broken when the allegiance of the thalamic relay cell shifts to the right as dotted lines. Part B shows the new connection to layer 4 following a shift of allegiance caused by daytime learning and the updated layer 6 connections required to contain further error. The dashed lines indicate new connections that must be formed. The dashed connection from layer 0 to layer 6 can be formed if a strong, bursting calibration signal is played into the thalamic relay cell. The existing feed-forward connection causes firing of the layer 6 cell marked 1. If this connection is plastic, it will strengthen and generate errant synapses onto the correct neighboring cells. The new, dashed feedback connection is trickier. While it can be generated by a localized calibration signal in layer 0 in the case shown, the layer 6 cell marked 1 must innervate all the thalamic cells that innervate the layer 4 cell marked 1. As explained in the text, this requires random, tonic input to thalamus.

ty is controlled postsynaptically, at the level of the recipient cells in layers 2/3 and 5. One possibility is that layer 6 cells influence the firing of "matrix" thalamic cells, which in turn regulate, via their connections in layer 1, the backpropagation of apical dendritic spikes (see Figure 2). Another is that the intracortical terminals of layer 6 cells regulate the plasticity of feedforward connections directly.

Although the scenario shown in Figure 3 guarantees maintenance of accurate connections, it has two important drawbacks. First, the neocortex will learn much more slowly than if there were no errors, connectivity were complete and synapses were always plastic. Second, while this scheme maintains accurate connections, it does not eliminate synaptic error but merely prevents errors from spreading too far. If an errant, or "mutant," synapse actually does form and proves to be useful, it may flourish at the expense of the original, "correct" synapses. This would shift the allegiance of the corresponding thalamic cell (Figure 4) which, in turn, requires updating the connections of the layer 6 cells that compute correlation ratios. These two features—slow learning and connection updating—require that the cortex sleeps.

SLEEP AND THE NEOCORTEX

The difference between sleep and wakefulness is at once the most banal and most perplexing aspect of brain behavior. Sleep is essential for mammals: rats deprived of sleep die more quickly than when starved; dolphins, which must remain awake to breathe, sleep on alternate sides of the cortex; and one species, which must make continuous decisions in muddy estuaries, sleeps in brief micronaps so that it can monitor its environment almost seamlessly. No mammal has successfully occupied the open niche of sleeplessness, although some mammals do function periodically throughout the day and night.

Although sleep is triggered by events deep in the brain, its electrical and subjective manifestations are largely cortical. Two quite distinct phases of sleep alternate through the night. In one, the cortex exhibits rhythmic slow waves, including waxing and waning bursts called *sleep spindles*. A subject awoken from slow wave sleep reports no dreams. This slow wave phase is followed by *rapid eye movement* (REM) sleep, sometimes called *paradoxical* sleep because the brain waves are rapid and small, as in wakefulness. Recent research has illuminated the cellular bases of these states with-

out explaining why we sleep. Slow wave sleep rhythms seem to be generated by thalamic circuitry: a cap of thalamic cells, the *reticular nucleus*, imposes a regular, sweeping wave of bursts on relay cells that then convey the signals to the cortex. These waves of bursts resemble a calibration signal and thalamic relay cells in this mode ignore arriving patterns of input. In REM sleep, the relay cells discharge in an irregular tonic stream. Part of this random activity is triggered by inputs from acetylcholine-releasing neurons, which fire strongly in the brainstem during REM sleep.

There are at least two, not incompatible, views of sleep. The first assumes detailed, meaningful information (represented by complicated streams of spikes in relay cells during REM sleep) is being replayed into the cortex, allowing the brain to catch up with the day's events. Because the cortex learns slowly and is specialized for discovering general relationships rather than storing specific facts, the brain has an auxiliary system, the hippocampus, for rapid learning of recent experiences. Hippocampus can learn rapidly because it makes no attempt to combine experiences into generalizations. In contrast, the neocortex must generate neural representations of experiences that reveal underlying relationships between them. New experiences provoke gradual, but extensive, modifications of existing representations but the latter must be modified very cautiously to avoid destroying existing concepts. Artificial, fully connected neural networks that learn concepts (as opposed to parroting arbitrary facts) must do so slowly, employing information about current performance together with nonlinear response characteristics in order to develop appropriate connection strengths. It seems likely that an analogous process could account for slow learning in sparsely connected networks with synaptic errors.

The second view of sleep is that it represents off-line updating of the layer 6 connections that underlie correlation measurement and plasticity control. If daytime experience produces a flourishing errant synapse, causing a thalamic relay cell to switch its allegiance from one layer 4 cell (Figure 4a) to another (Figure 4b), then some rewiring of the layer 6 circuitry is required to continue to keep synaptic error in check. In particular, the connections shown dotted in Figure 4 must be broken while the connections shown dashed must be created. If a strong calibration signal is played into the rewired thalamic relay cell during slow wave sleep, the rewiring from thalamus to layer 6 occurs automatically. The layer 4 cell fires in response to the calibration signal and this, together with the existing direct thalamic input to layer 6, causes the appropriate layer 6 cell to spike. The correlated activity of the thalamic and layer 6 cells then forms the correct pattern of connections, including the connections to the neighbors of the appropriate layer 6 cell (which are formed by erroneous sprouts).

The rewiring of the feedback connections from layer 6 to thalamic relay cells is a little more difficult. Each layer 4 cell receives input from a small set of relay cells. For example, in the visual cortex, a layer 4 simple cell receives input from a group of relay cells that, in turn, receive input from cells lying along a short oriented bar in the retina. These connections, endowing the cell with its tuned response characteristics, are often referred to as the "receptive field" of the cell. The corresponding layer 6 cell must innervate all thalamic cells innervating a particular layer 4 cell. In other words, the layer 6 cell must find the cells in thalamus constituting the receptive field of its corresponding layer 4 cell, then make or maintain connections to those cells.

How can a layer 6 cell determine this receptive field? Perhaps the same way a neurophysiologist would. The pioneering experiments of Hubel and Wiesel involved waving potentially relevant objects in front of a cat and recording the spikes emitted by specific cells: they stumbled upon the particular relevance of oriented bars, their ticket to Stockholm! But nowadays neurophysiologists use a more objective approach called *reverse correlation analysis*. They confront the cat with computer-generated white noise input (a screen of snowflakes) and the computer records the particular configuration of snowflakes each time the cell spikes. The average of the recorded configurations is the stimulus that best triggers spikes, which is the receptive field.

It is plausible that the brain adopts a similar strategy during paradoxical sleep. Cholinergic brainstem neurons supply random inputs to the thalamus, and layer 6 cells then perform reverse correlation analysis. Whenever the random input to the thalamus happens to fire layer 4 cells, the corresponding layer 6 cells also fire spikes, strengthening their connections with the appropriate relay neurons. This rather elegant solution to the problem of updating corticothalamic feedback connections is possible because the layer 6 cells are wired to act as correlation detectors.

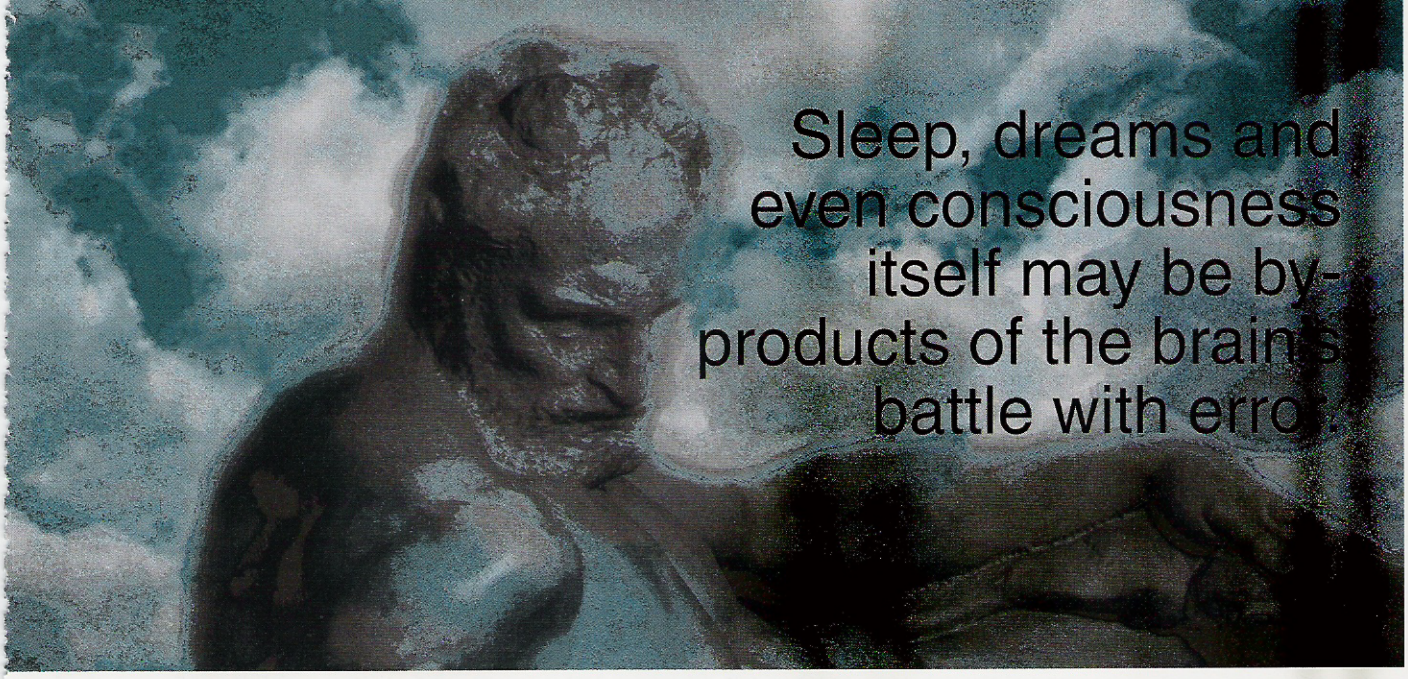

Why is random activation of relay cells experienced consciously, and often vividly, as dreams, while slow wave activation is not? The crucial difference may lie in the firing mode of the thalamic cells: in paradoxical sleep they fire tonically but in slow wave sleep firing is in bursts. Although there may be linkage between tonic-plasticity-remembering and burst-implasticity-forgetting in wakefulness, it is not likely that this accounts for the difference between paradoxical and slow wave sleep. The random thalamic patterns of REM sleep should not be learned since they are meaningless (except for the few thalamic nuclei involved in replaying the important, seemingly random patterns stored in hippocampus during the daytime so that the neocortex can learn during sleep). A more likely explanation is that the feed-forward synapses are mostly implastic during sleep and dreams are remembered only by active processes that occur immediately on waking. The difference in the subjective quality of the two forms of sleep may lie even deeper—we may never be conscious of burst activity simply because consciousness *corresponds* to tonic activity. Indeed, there seems to be little point in being conscious of cerebral activity from which nothing can be learned.

Dreams are characterized not only by vividness but also by their mixture of the bizarre and the compelling. If dreams are built on random inputs, they will be bizarre but not necessarily meaningless: random inputs to higher cortical areas concerned with objects, concepts, etc. will be experienced not as snowflakes but as unusual, even impossible, sequences of familiar events. Moreover, the recurrent mechanisms that try to force every input into a definite interpretation will still be active.

CONCLUSIONS

Recent experimental and theoretical work is producing a new view of the neocortex that complements the older picture of a massively hierarchical, parallel information-processing neural network. In the new paradigm, a large part of cortical machinery and function is devoted to internal housekeeping operations like noise reduction and plasticity control rather than to information-processing in the classical sense. There are, of course, many other operations to ensure that information is correctly sorted, labeled and routed. These are often referred to collectively as "binding," i.e. making sure that neural activity related to particular objects, concepts, etc. is identified as "belonging" and treated differently from activity that does not "belong." The layer 6 arrangements are well suited to these sorting, labeling and routing functions. Although these operations correspond to the brain's bureaucracy—and smack of routine housekeeping—their consequences are far from banal. Sleep, dreams and even consciousness itself may be by-products of the brain's battle with error.

SUGGESTED READING

Adams, P.R. "Hebb and Darwin," *Journal of Theoretical Biology* **195**, 419-438 (1998).

Jones, E.G. "Viewpoint: the core and matrix of thalamic organization," *Neuroscience* **85**, 331-345 (1998).

Mountcastle, V.B. *Perceptual Neuroscience: The Cerebral Cortex*, Harvard University Press, Cambridge, MA (1998).