


6

The Cortical Representation

In this chapter we will review the representation of information in the visual cortex. There have been many advances in our understanding of the visual cortex over the last 25 years. Even today, our view of the visual cortex is changing rapidly; new results that change our overall view sometimes seem to appear weekly. In the beginning of this chapter, I will review what is commonly accepted concerning the visual cortex. Toward the end, I will introduce some of the broader claims that have been made about the relationship between the visual cortex and perception. We will take up the issue of connecting the cortex, computation, and seeing again in the later chapters.

Overview of the Visual Cortex

The most prominent feature in the lateral view of the brain shown in Figure 6.1 is the cortex. The human cortex is a 2-mm-thick sheet of neurons with a surface area of 1,400 cm². Rather than lining the skull, as the retina lines the eye, the cortex is like a crumpled sheet stuffed into the skull. Each location where the folded cortex forms a ridge visible from the exterior is called a **gyrus**, while each shallow furrow that separates a pair of gyri is called a **sulcus**. The pattern of sulci and gyri differ considerably across species; the human brain contains more sulci than other primate brains. There are also significant differences between human brains, although the broad outlines of the sulcal and gyral patterns are usually present and recognizable across different people. The gyri



6.1 CORTICAL REGIONS. Based on its overall shape, anatomists divide the human brain into four regions called the occipital, temporal, parietal, and frontal lobes. Based on its internal connections, the cortex can be further divided into many anatomically distinct areas. Visual input to the brain arrives in primary visual cortex, or area V1, which is located in the occipital lobe.

and sulci are convenient landmarks, but they probably have no functional significance.

The most visible sulci are used as markers to partition the human brain into four lobes. The lobes are called **frontal**, **parietal**, **temporal**, and **occipital** to describe their relative positions (see Fig. 6.1). Each lobe contains many distinct brain areas, that is, contiguous groups of cortical neurons that appear to function in an interrelated manner. A cortical area is identified in several ways, though perhaps the most significant is by its anatomical connections with other parts of the brain. Each brain area makes a distinctive pattern of anatomical connections with other brain areas. The inputs arriving to one area come from only a few other places in the brain, and the outputs emerging from that area are sent to a specific set of destination areas.

In primates, the great part of the visual signal from the retina and the LGN arrives at a single area within the occipital lobe of the cortex called **area V1**, or the **primary visual cortex**. This is a large cortical area, comprising roughly 1.5×10^8 neurons, many more than the 10^6 neurons in the LGN. Area V1 can be identified by a prominent striation made up of a dense collection of myelinated axons within one of the layers of visual cortex. The striation is coextensive with area V1 and appears as a white band to the naked eye.¹ Because of its promi-

¹ Because area V1 was defined by the presence of this striation, it is sometimes called *striate cortex*.

nence, importance has been the subject with a review of area V1.

In addition, receive a strong physiology, an active study and come. We will be performed chapters concerning functional roles. Van Essen, 1995

Most of what experimental studies in the anatomical species. These cortical manipulations removal of the visual jumps, rudimentary (1974) has still removed. Initially, Over time, he was able to work up small cases of area V1 is essential. In essence, I have and describe it

The Archi

There is a great visual areas. The the two retinas: ture of the primary area V1 first. two retinas is represented in

The Layers of

Like the cortex shows a cross identified each ences in the ferences in the layer 1 has v

nence, important anatomical location, and large size (24 cm²), area V1 has been the subject of intense study, and we will begin this chapter with a review of the anatomical and electrophysiological features of area V1.

In addition to area V1, more than 20 other cortical areas that receive a strong visual input have been discovered. The anatomy, electrophysiology, and computational purpose of these areas are now under active study and will be an important topic for study for many years to come. We will review some of the preliminary experiments that have been performed in these visual areas at the end of this chapter. In later chapters concerning motion and color, we will return to consider the functional roles of these visual areas (Zeki, 1978, 1990; Felleman and Van Essen, 1991).

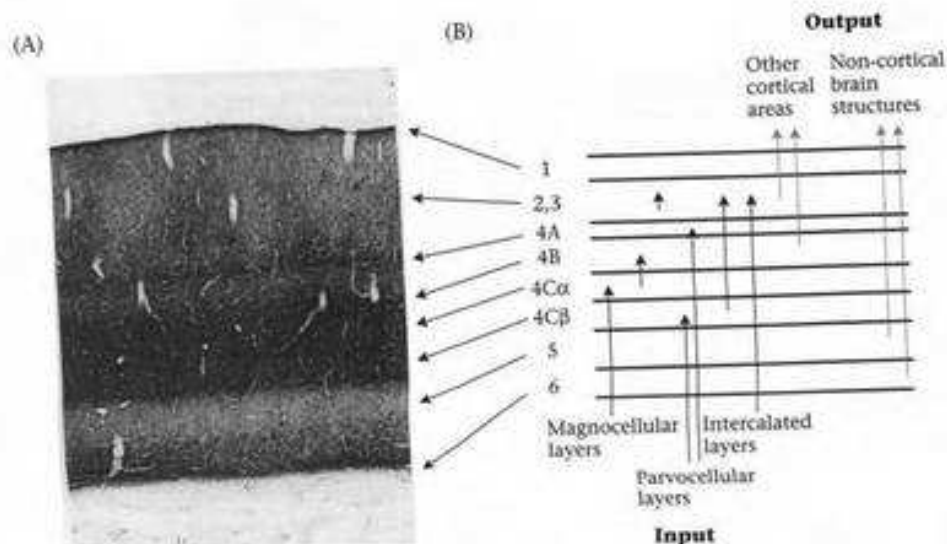
Most of what we know about cortical visual areas comes from experimental studies of cats and monkeys. There are significant differences in the anatomy and functional properties of the cortices of different species. These differences can be demonstrated in simple experimental manipulations. For example, Sprague et al. (1977) have shown that removal of the cat primary visual cortex does not blind the cat: the animal jumps, runs, and appears normal to the casual observer. Humphrey (1974) has studied the behavior of a monkey whose area V1 was removed. Initially the lesion appeared to blind the monkey completely. Over time, however, the monkey recovered some visual function and was able to walk around objects, climb a tree, and even find and pick up small candy pellets in her play area. In humans, however, the loss of area V1 is devastating to all visual function. Because of these differences, I have focused this review on experimental studies of primates, and describe measurements of the human brain whenever possible.

The Architecture of the Primary Visual Cortex

There is a great deal of precision in the interconnections of cortical visual areas. The specific pattern of connections received by area V1 from the two retinas via the LGN results in certain regularities of the architecture of the primary visual cortex. We review the anatomical structure of area V1 first. Then, we review how the pattern of connections from the two retinas imposes an overall organization on the visual information represented in cortical area V1.

The Layers of Area V1

Like the cortex in general, area V1 is a layered structure. Figure 6.2A shows a cross section of the visual cortex. Several major layers can be identified easily. Area V1 is segregated into six layers based on differences in the relative density of neurons, axons, and synapses, and differences in the interconnections to the rest of the brain. The superficial layer I has very few neurons but many axons, dendrites, and synapses,



6.2 AREA V1 IS A LAYERED STRUCTURE. (A) A stained cross section of the visual cortex of a macaque shows the individual layers. Each layer has different proportions of cell bodies, dendrites, and axons and may be distinguished by the density of the staining and other properties. The light areas are blood vessels. (B) The organization of the neural inputs and outputs to area V1 are shown. The parvocellular and magnocellular inputs make connections in layer 4C. The intercalated neurons make connections in the superficial layers. The outputs are sent to other cortical areas, back to the LGN, and to other subcortical nuclei. Photomicrograph courtesy of Jennifer Lund.

which collectively are called **neuropil**. Layers 2 and 3 consist of a dense array of cell bodies and many local dendritic interconnections. These layers appear to receive a direct input from the intercalated layers of the LGN as well (Fitzpatrick et al., 1983; Hendry and Yoshioka, 1994), and the outputs from layers 2 and 3 are sent to other cortical areas. Layers 2 and 3 are hard to distinguish based on simple histological stains of the cortex. Frequently, layers 1–3 are grouped together and simply called the **superficial layers** of the cortex.

Layer 4 has been subdivided into several parts as the interconnections with other brain areas and layers have become clarified. Layer 4C receives the primary input from the parvocellular and magnocellular layers of the LGN. The magnocellular neurons send their output to the upper half of this layer, which is called 4C α while the parvocellular neurons make connections in the lower half, called 4C β . Layer 4B receives a large input from 4C α and sends its output to other cortical areas. Layer 4B can be defined anatomically by the presence of a large striation, called the **stria of Gennari**, which is composed mainly of cortical axons.

Layer 5 contains relatively few cell bodies compared to the surrounding layers. It sends a major output to the superior colliculus, a

stru
outp
solu
the
tion
(Roc
from
the
cent
inte
mar
ical
catic
the

The
The
to t
Amc
sph
only
V1
field

thei
toge
are
the
in tl
side
the
brai
both

bas
resp
spat
trate
and

gin
ent
num
the
the

structure in the midbrain. Layer 6 is dense with cells and sends a large output back to the LGN (Lund et al., 1975). As a general though not absolute rule, forward outputs to new cortical areas tend to come from the superficial layers and terminate in layer 4. The feedback projections tend to come from the deep layers and terminate in layers 1 and 6 (Rockland and Pandya, 1979; Felleman and Van Essen, 1991).

The wiring diagram in Figure 6.2B shows that the signals to and from area V1 are complex and highly specific. One must suppose that the interconnections within area V1 are specific, too. Roughly 25 percent of the neurons in all layers are inhibitory interneurons, and their interconnections must be governed by the presence of biochemical markers that identify which neurons should connect and how. Anatomical classification of the cell types within the visual cortex, and identification of the local circuitry, will provide us with many more clues about the functional significance of this area.

The Pathway to Area V1

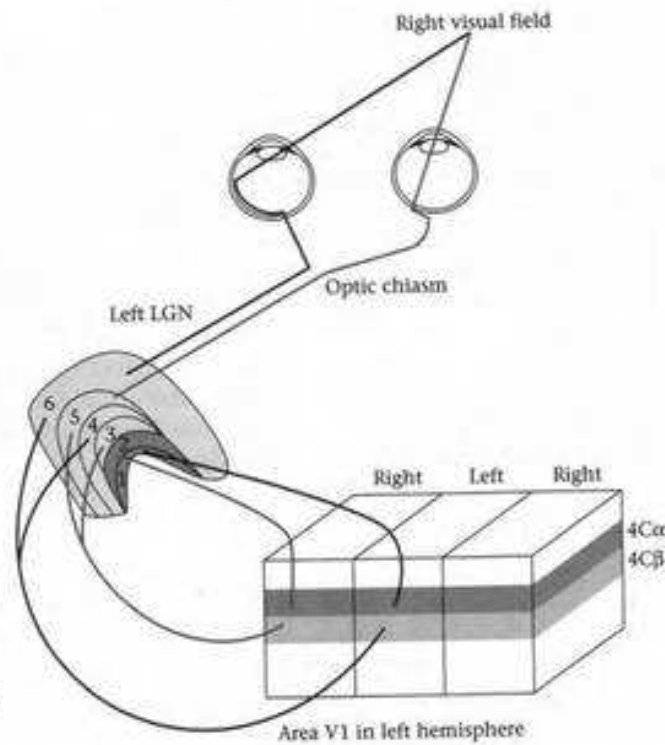
The structure of the anatomical pathways leading from the two retinas to the cortex defines many of the fundamental properties of area V1. Among the most significant properties is that area V1 in each hemisphere has only a restricted field of view: area V1 in the left hemisphere only receives visual input from the right half of the visual field; area V1 in the right hemisphere only receives input from the left visual field.

We can see how this arises by considering how retinal signals make their way to area V1. The optic tract fibers from the two retinas come together at the optic chiasm, as shown in Figure 6.3. There the fibers are sorted into two new groups that each connect to only one side of the brain. Axons from ganglion cells whose receptive fields are located in the left visual field send their outputs towards the LGN on the right side of the brain, while axons of ganglion cells with receptive fields in the right visual field communicate their output to the left side of the brain. Consequently, each LGN receives a retinal signal derived from both eyes, but from only one half of the visual field.

The signals reaching the cortex from the retina respect three other basic organizational principles. The interconnections are organized with respect to (a) the eye of origin, (b) the class of ganglion cell, and (c) the spatial position of the ganglion cell within the retina. Figure 6.3 illustrates the pattern of connections schematically, starting at the retinas and continuing to area V1.

EYE OF ORIGIN. Within the LGN, information about the eye of origin is preserved, since fibers from each eye make connections in different layers of the LGN. The parvocellular and magnocellular layers are numbered as 1-6: layers 2, 3, and 5 receive input from the retina on the same side of the head, while layers 1, 4, and 6 receive input from the retina on the opposite side. The connections of these layers for the

6.3 THE SIGNALS FROM THE TWO RETINAS ARE COMMUNICATED TO AREA V1 via the LGN. Points in the right visual field are imaged on the temporal side of the left eye and the nasal side of the right eye. Axons from ganglion cells in these retinal regions make connections with separate layers in the left LGN. Neurons in the magnocellular and parvocellular layers of the LGN send their outputs to cortical layers 4C α and 4C β , respectively. The signals from each eye are segregated into different bands within area V1. Signals from these bands converge on individual neurons in the superficial layers of the cortex.



left LGN are illustrated in Figure 6.3. Why this particular pattern of ocular connections exists is a mystery. The eye of origin for the intercalated layers, which fall between the parvocellular and magnocellular layers, has not yet been ascertained.

The signals from the two eyes remain segregated as they arrive at the input layers of area V1. One can observe this segregation by measuring the electrophysiological responses of the units in layer 4C. As the recording electrode is moved within layer 4C, there is an abrupt shift as to which eye drives the unit. In layer 4C, the shift from one eye to the other takes place over a distance of less than 50 μm . Above and below layer 4C the signals from the two eyes converge onto single neurons, although there is still a tendency for individual neurons to receive inputs predominantly from one eye or the other, and this pattern is aligned with the input pattern. The transition between eyes of origin is less abrupt in the superficial layers, perhaps extending over 100 μm . The relative segregation of information across the visual cortex with respect to the eye of origin has led researchers to describe its organization in terms of **ocular dominance columns** (Hubel and Wiesel, 1977; Bishop, 1984).

In addition to evidence from electrophysiological measurements, one also can use anatomical methods to visualize the ocular dominance columns. After injection into one eye, the tritiated amino acid

prol
tic c
4C
deve
ence
ligh
the
mor
they
199


from
sup
the
upo
that
fect
neu
cal
for
an
et a
Mit
col
upc
198

of r
tex.
sol
mic

proline will be transported from the retina to the cortex across synaptic connections. By sectioning the visual cortex tangentially at layer 4C and exposing the sections to a photographic emulsion, one can develop a pattern of light and dark stripes that corresponds to the presence and absence of the tritiated proline. Figure 6.4 shows a pattern of light bands that mark regions receiving input from the injected eye; the intervening dark areas receive input from the opposite eye. In the monkey, these bands each span approximately $400\ \mu\text{m}$; in the human they span approximately $800\ \mu\text{m}$ (Hubel et al., 1978; Horton and Hoyt, 1991).

In the superficial layers of area V1 many neurons respond to stimuli from both eyes; in a normal monkey, 80 percent of the neurons in the superficial layers of area V1 are binocularly driven. The development of the interconnections necessary to drive the binocular neurons depends upon experience during maturation. Hubel and Wiesel (1965) showed that artificially closing one eye or cutting an ocular muscle strongly affects the development of neurons in area V1. Specifically, the binocular neurons fail to develop. Behaviorally, if one eye is kept closed for a critical period during development, the animal will remain blind in this eye for the rest of its life. This is quite different from the result of closing an adult eye for a few months, which has no significant effect (Hubel et al., 1977a; Shatz and Stryker, 1978; Movshon and Van Sluyters, 1981; Mitchell, 1988). In the cat, normal development of ocular dominance columns and presumably the binocular interconnections as well depend upon neural activity originating in the two retinas (Stryker and Harris, 1986).

GANGLION CELL CLASSIFICATION. Information from different classes of retinal ganglion cells remains segregated along the path to the cortex. Neurons in the magnocellular layers receive fibers from the parasol cells; neurons in the parvocellular layers receive fibers from the midget ganglion cells. It is uncertain precisely which retinal ganglion



6.4 THE OCULAR DOMINANCE COLUMNS IN AREA V1 can be visualized by using a radioactive marker, tritiated proline. When the marker is injected into one eye it is transported via the LGN nucleus to the cortex. The radioactive uptake is revealed in this dark-field photograph. The light bands in this tangential section show the places where the radioactive marker was located and thus reveal the ocular dominance columns. Source: Hubel et al., 1978.

cells project to the intercalated layers. The segregation of signals continues to the input of area V1. Within layer 4C, the upper half ($4C\alpha$) receives the axons from the magnocellular layers, while the lower half ($4C\beta$) receives the parvocellular input. The neurons in the intercalated layers send their output to the superficial layers 2 and 3 (Hendry and Yoshioka, 1994).

RETINOTOPIC ORGANIZATION. The spatial position of the ganglion cell within the retina is preserved by the spatial organization of the neurons within the LGN layers. The back of the nucleus contains neurons whose receptive fields are near the fovea. Toward the front of the nucleus, the receptive field locations become increasingly peripheral. This spatial layout is called **retinotopic organization** because the topological organization of the receptive fields in the LGN parallels the organization in the retina.

The signals in area V1 are also retinotopically arranged. From electrophysiology in monkeys, one can measure the location of receptive fields with an electrode that penetrates tangentially through layer 4C, traversing through the ocular dominance columns. The receptive-field centers of neurons along this path correspond systematically to locations from the fovea to the periphery. This trend is interrupted locally by small, abrupt jumps at the ocular dominance borders. Within the first ocular dominance column the receptive-field center positions change smoothly; as one passes into the next ocular dominance region there is an abrupt shift of the receptive field positions equal to about half of the space spanned by receptive fields in the first column. Hubel and Wiesel (1977) refer to this organization as "two steps forward and one step back."

In the last 15 years, it has become possible to estimate spatially localized activity in the human brain. Beginning with **positron emission tomography (PET)** studies, and more recently by using **functional magnetic resonance imaging (fMRI)**, researchers have been able to measure activity in volumes of the cortex as small as 10 mm^3 containing a few hundred thousand neurons.²

Human area V1 is located within the **calcarine sulcus** in the occipital lobe. The calcarine sulcus in my brain, and its retinotopic organization, is shown in Figure 6.5. Neurons with receptive fields in the central visual field are located in the posterior calcarine sulcus, while neurons with receptive fields in the periphery are located in the anterior portions of the sulcus. At a given distance along the sulcus, the receptive fields are located along a semicircle in the visual field. Neurons with receptive fields on the upper, middle, and lower sections of

²Both of these methods are based on indirect measures of neural activation. With the PET method, an observer receives a low dose of radiation in his bloodstream, and neural activity is indicated by brain regions showing increased radioactivity. The fMRI signal detects differences in the local concentration of blood oxygen. Both the increased radioactivity and the change in local blood oxygenation are due to vascular responses to the neural activity (Kwong et al., 1992; Ogawa et al., 1992; Posner and Raichle, 1994).


(A)

6.5
it ma
sulcu
neur
appe
sulcu
hemi
upper
the lo

the :
the
Hort

from
The
was
cates
edge
new
of th
tivity
ante

fMR
is ar
tion:
nal
show
fMR
wan




6.5 HUMAN AREA V1 is located mainly in the calcarine sulcus, and in some individuals it may extend onto the occipital pole. (A) Seen in sagittal view, the calcarine is a long sulcus that extends roughly 4 cm. The visual eccentricities of the receptive fields of neurons at different locations are shown. (B) In the coronal plane the calcarine sulcus appears as an indentation of the medial wall of the brain. At a given distance along the sulcus, the receptive fields of neurons fall along a semicircle within the visual field. Each hemisphere represents one half of the visual field. Neurons with receptive fields on the upper, middle, and lower sections of a semicircle of constant eccentricity are found on the lower, middle, and upper portions of the calcarine sulcus respectively.

the semicircle are found on the lower, middle, and upper portions of the calcarine sulcus, respectively (Inouye, 1909; Holmes, 1918, 1945; Horton and Hoyt, 1991).

Engel et al. (1994) measured the human retinotopic organization from fovea to periphery by using the stimulus shown in Figure 6.6A. The stimulus consisted of a series of slowly expanding rings; each ring was a collection of flickering squares. The ring began as a small spot located at the fixation mark, and then it grew until it traveled beyond the edge of the visual field. As a ring faded from view, it was replaced by a new ring starting at the center. Because of the retinotopic organization of the calcarine sulcus, each ring causes a traveling wave of neural activity beginning in the posterior calcarine sulcus and traveling in the anterior direction.

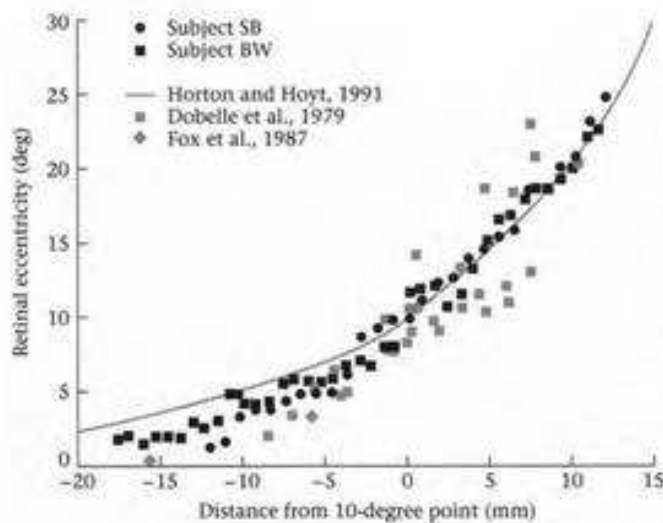
We can detect the traveling wave of activation by measuring the fMRI signal at different points along the calcarine sulcus. Figure 6.6B is an image of the brain within the plane of the calcarine sulcus. Positions within the calcarine sulcus are indicated in black. The fMRI signal at each point within the sulcus, plotted as a function of time, is shown in the mesh plot in Figure 6.6C. Notice that the amplitude of the fMRI signal covaries with the stimulus; the fMRI signal waxes and wanes four times through the four periods of the expanding annulus.



6.6 RECEPTIVE-FIELD LOCATIONS OF NEURONS IN THE HUMAN CALCARINE SULCUS measured by functional magnetic resonance imaging (fMRI). (A) The observer viewed a series of concentric expanding annuli presented on a gray background. Each annulus contained a high-contrast flickering radial checkerboard pattern. As an annulus expanded beyond the edge of the display, a new annulus emerged in the center creating a periodic image sequence. The sequence was repeated four times in a single experiment. (B) An image within the plane of the calcarine sulcus. The dark lines indicate points identified as within the calcarine sulcus. (C) The fMRI temporal signal at different points within the left calcarine sulcus. The fMRI signal follows the time course of the stimulus; the phase of the signal is delayed as one measures from the posterior to the anterior calcarine sulcus. After Engel et al., 1994.

The temporal phase of the fMRI signal varies systematically from the posterior to anterior portions of the sulcus. Activity in the posterior portion of the sulcus is advanced in time compared to activity in the anterior portion. This traveling wave occurs because the stimulus creates activity in the posterior part of the sulcus first, and then later in the anterior part of the sulcus.

In addition to fMRI, there are several other estimates of the mapping from visual-field eccentricity to locations in the calcarine sulcus. These estimates are compared in Figure 6.7. The fMRI measurements from two observers are shown as the black symbols. Estimates from direct electrical stimulation of the cortex are shown as gray squares (Dobelle et al., 1979). In these experiments the volunteer observer's brain was stimulated, and he indicated the location of the perceived visual stimulation within the visual field (see also Brindley and Lewin, 1968). The three gray diamonds show measurements using PET. These data



6.7 METHODS TO ESTIMATE THE RECEPTIVE-FIELD

LOCATION OF NEURONS in the calcarine sulcus. The black symbols show measurements from two observers using the fMRI method (Engel et al., 1994). The shaded squares are from a microstimulation study on a blind volunteer (Dobbelle et al., 1979). The shaded diamonds are measurements averaged from five observers using PET (Fox et al., 1987). The shaded curve is an estimate based on studying the locations of scotoma in stroke patients and single-cell data from nonhuman primates (Horton and Hoyt, 1991). Source: Engel et al., 1994.

represent the average of five different observers, normalized for differences in brain size. The shaded line shows an estimate Horton and Hoyt (1991) made by studying the positions of scotoma in observers with localized brain lesions and by extrapolating from data on monkeys. These estimates are in good agreement, and they all show that considerably more cortical area is allocated to the foveal representation than to the peripheral representation.

The allocation of more cortical area to the foveal representation than to the peripheral representation seems a natural consequence of the fact that more photoreceptors and retinal ganglion cells represent the fovea than the periphery. Wässle et al. (1990; see also Schein, 1988) suggested that the expanded foveal representation can be explained by assuming that every ganglion cell is allocated an equal amount of cortical area. More recently, Azzopardi and Cowey (1993) have suggested that there is a further expansion of the foveal representation, and that foveal ganglion cells are allocated 3–6 times more cortical area than peripheral ganglion cells.

Electrical Stimulation of Human Area V1

Direct electrical stimulation of the visual cortex causes the sensation of vision. When a visual impression is generated by nonphotic stimulation, say by pressing on the eyeball or by electrical stimulation, the resulting perception is called a **visual phosphene**. In order to develop visual prostheses for individuals with incurable retinal diseases, several research groups have studied the visual properties of phosphenes created by electrical stimulation of the visual cortex (Brindley and Lewin, 1968; Dobbelle et al., 1979; Bak et al., 1990).

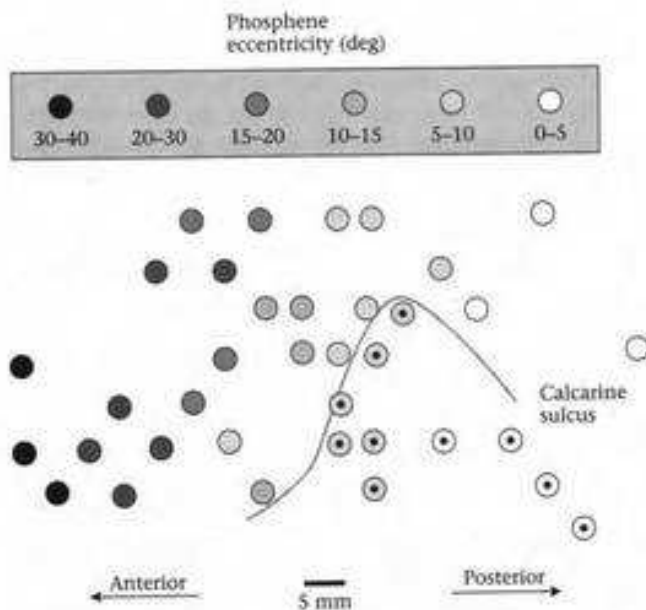
Brindley and Lewin (1968) describe experiments with a human volunteer who was diabetic and suffered from bilateral glaucoma and a right retinal detachment, and who was effectively blind throughout most of the visual field, and completely blind in the central fifteen degrees. With the patient's consent, Brindley and Lewin built and implanted a stimulator that could deliver current to the surface of her brain, near the patient's primary visual cortex. They asked her to describe the appearance of the electrical stimulation following stimulation by the different electrodes, at various positions within her primary visual cortex. She reported that electrical stimulation caused her to perceive a phosphene that appeared to be a point of light or a blob in space. Her description of the visual impression caused by most of the electrodes was "like a grain of rice at arm's length." Occasionally one electrode might cause a slightly larger impression, "like half a matchstick at arm's length."

As might be expected from the retinotopic organization of the visual cortex, the position of the phosphenes varied with the position of the stimulating electrodes. The observer told the experimenters where she perceived the phosphenes to be using a simple procedure. She grasped a knob with her right hand and imagined she was fixating on that hand. She then pointed to the location of the phosphenes relative to the fixation point using her left hand.

Figure 6.8 shows the positions of the electrodes and the corresponding phosphenes. The pattern of results follows the expectations from the retinotopic organization of the calcarine sulcus. Stimulation by electrodes near the back of the brain created phosphenes in the central 5 degrees; stimulation by forward electrodes created phosphenes in more eccentric portions. More cortical area is devoted to the central regions of vision than to peripheral regions.

Brindley and Lewin tested the effects of superposition by stimulating with separate electrodes and then stimulating with both electrodes at once. When electrodes were far apart, the visual phosphene generated by stimulating both electrodes at once could be predicted from the phosphenes generated by stimulating individually. Superposition also held for some closely spaced electrodes, but not all. The test of superposition is particularly important for practical development of a prosthetic device. To build up complex visual patterns from stimulation of area V1, it is necessary to use multiple electrodes. If linearity holds, then we can measure the appearance from single electrode stimulations and predict the appearance corresponding to multiple stimulations. That superposition held approximately suggests that it may be possible to predict the appearance of the multiple electrode stimulation from measurements using individual electrodes. Without superposition, we have no logical basis for creating an image from the intensities perceived at single points.

There have been a few recent reports of stimulation of the human visual cortex. For example, Bak et al. (1990) used very fine (37.5- μm)



6.8 ELECTRICAL STIMULATION OF HUMAN AREA V1 using implanted microelectrodes reveals the retinotopic organization of human cortex. The symbols are plotted at the electrode positions on the medial wall of the brain. The shading of the symbol indicates the visual eccentricity of the phosphene created by electrode stimulation. A dot within the symbol means that the phosphene was perceived in the upper visual field. The shaded curve shows the inferred position of the calcarine sulcus. Source: Brindley and Lewin, 1968.

microelectrodes inserted within the cortex. They experimented on patients who were having epileptic foci removed. These patients were under local anaesthesia and could report on their visual sensations. Bak et al. observed that, when the electrodes were embedded within the visual cortex, visual sensations could be obtained with quite low current levels. Brindley and Lewin (1968) had used about 2 mA of current, but Bak et al. found thresholds about 100 times lower, near 20 μ A. The appearance of the visual phosphene was steady in these patients, and some of the phosphenes appeared to be colored. Time was quite limited in these studies, and only a few experimental manipulations were possible. But, Bak et al. report that when the microelectrodes were separated by more than 0.7 mm, the two phosphenes could be seen as distinct, while separations of 0.3 mm were seen as a single spot. For one subject, nearly all of the phosphenes were reported to be strongly colored, unlike the phosphenes reported by Brindley and Lewin's patient. While the subjects were stimulated, they could also perceive light stimuli. The phosphenes were visible against the backdrop of the normal visual field.

Receptive Fields in the Primary Visual Cortex

The receptive fields of neurons in area V1 are qualitatively different from those in the LGN. For example, LGN neurons have circularly symmetric receptive fields, but most V1 neurons do not. Unlike LGN neurons, some neurons in area V1 respond well to stimuli moving in one

direction but fail to respond to stimuli moving in the opposite direction. Some area V1 neurons are binocular, responding to stimuli from both eyes. These new receptive-field properties must be related to the visual computations performed within the cortex such as the analysis of form and texture, the perception of motion, and the estimation of depth. We might expect that these new receptive-field properties have a functional role in these visual computations.

Much of what we know about cortical receptive fields comes from Hubel and Wiesel's measurements during their 25 years of collaboration. Others had accomplished the difficult feat of recording from cortical neurons first, but the initial experiments used diffuse illumination, say, turning on the room lights, as a source of stimulation. As we have seen, pattern contrast is an important variable in the retinal neural representation; consequently, cortical cells respond poorly to diffuse illumination (von Baumgarten and Jung, 1952). Hubel and Wiesel made rapid progress in elucidating the responses of cortical neurons by using stimuli of great relevance to vision and by being extremely insightful. Hubel and Wiesel's papers chart a remarkable series of advances in our understanding of the visual cortex. Their studies have defined the major ways in which area V1 receptive fields differ from LGN receptive fields. Their qualitative methods for studying the cortex continue to dominate experimental physiology (Hubel and Wiesel, 1959, 1962, 1968, 1977; Hubel, 1982).

Hubel and Wiesel recorded the activity of cortical neurons while displaying patterned stimuli, mainly line segments and spots, on a screen that was imaged through the animal's cornea and lens onto the retina. As the microelectrode penetrated the visual cortex, they presented line segments whose width and length could be adjusted. First, they varied the position of the stimulus on the screen, searching for the neuron's receptive field. Once the receptive-field position was established, they measured the response of the neuron to lines, bars and spots presented individually.

One important goal of their work was to classify the cortical neurons based on their responses to the small collection of stimuli. They sought classifications that represented the neurons' receptive field properties and that also helped to clarify the neurons' function in seeing. Recall that classification of the receptive-field types was also an important theme when we considered the responses of retinal ganglion cells. It is of great current interest to try to understand whether the classifications of cortical neurons and retinal neurons can be brought together to form a clear picture of this entire section of the visual pathways.

A second important aspect of characterizing cortical neurons is to measure the transformation from pattern contrast stimulus to firing activity. We used linear-systems methods to design experiments and create quantitative models of this transformation for retinal ganglion cells. Linearity is an important idea when applied to cortical receptive fields too. The most important application of linearity is Hubel and Wiesel's


classification of cortical neurons into two categories, called **simple cells** and **complex cells**. This classification is based, in large part, on an informal test of linearity (Skottun et al., 1991). As Hubel writes, "For the most part, we can predict the responses of simple cells to complicated shapes from their responses to small-spot stimuli" (Hubel, 1988, p. 72). Complex cells, on the other hand, do not satisfy superposition. The response obtained by sweeping a line across the cell's receptive field cannot be predicted accurately from the responses to individual flashes of a line.

Orientation Selectivity

Because simple cell responses to light approximately satisfy homogeneity and superposition, their receptive fields can be measured using the linear methods described in Chapter 5. Simple-cell receptive fields differ from those of retinal ganglion cells or cells in the LGN. The receptive fields of simple cells consist of adjacent excitatory and inhibitory regions, as illustrated in Figure 6.9. Simple cells have oriented receptive fields, and hence they respond to stimuli in some orientations better than others. This receptive field property is called **orientation selectivity**. The orientation of the stimulus that evokes the most powerful response is called the cell's preferred orientation.

Orientation selectivity of cortical neurons is a new receptive-field property. LGN and retinal neurons have circularly symmetric receptive fields, and they respond almost equally well to all stimulus orientations. Orientation-selective neurons are found throughout layers 2 and 3, though they are relatively rare in the primary inputs within layer 4C.

Figure 6.9 shows several orientation-selective linear receptive fields and how these might be constructed from the outputs of LGN neurons. The simple-cell receptive fields consist of adjacent excitatory and in-




hibitory regions that are longer in one direction than the other. The main axis of the receptive fields defines the preferred orientation; stimuli oriented along the main axis of these receptive fields are more effective at exciting or inhibiting the cell than stimuli in other orientations. The figure shows the excitatory regions as resulting from the combined output of neurons with excitatory centers and the inhibitory regions resulting from the combined output of neurons with inhibitory centers.³

By comparing the three panels in Figure 6.9 you will see that receptive fields sharing a common preferred orientation can differ in a number of other ways. Panels (A) and (B) show two receptive fields with the same preferred orientation but different spatial arrangements of the excitatory and inhibitory regions. Panels (A) and (C) show two receptive fields with the same preferred orientation and arrangement of excitatory and inhibitory regions, but differing in the overall length of the receptive field. The neuron with the longer receptive field will respond well to a narrower range of stimulus orientations than the neuron with the shorter receptive field.

Complex cells also show orientation selectivity. Complex cells are nonlinear, so to explain the behavior of complex cells, including orientation selectivity, will require more complex models than the simple sums of neural outputs used in Figure 6.9.

The preferred orientation of neurons varies in an orderly way that depends on the neuron's position within the cortical sheet. Figure 6.10 shows the preferred orientation of a collection of neurons measured during a single, long, tangential penetration through the cortex. In any small region of layers 2 and 3, the preferred orientation is similar. As the electrode passes tangentially through the cortical sheet, the preferred orientation changes systematically, varying through all angles. Figure 6.10A shows an extensive set of measurements of preferred orientation made during a single tangential penetration (Hubel and Wiesel, 1977). Upon later review, Hubel and Livingstone (1987) noted that during these measurements there were certain intervals during which the receptive-field orientation was ambiguous. Figure 6.10B shows a second penetration in which regions with no preferred receptive-field orientation are identified. As we shall see, Hubel and Livingstone also report that the regions lacking orientation selectivity coincide with locations in layers 2 and 3 of the cortex where an enzyme called **cytochrome oxidase** is present in high density. However, there is some debate whether these measurements represent true differences in the receptive fields of individual neurons, or whether they represent differences in the distribution of activity in local collections of neurons (Leventhal et al., 1993; O'Keefe et al., 1993).

³In principle, one might construct an oriented receptive field from the outputs of a single line of LGN neurons. But recall that the receptive fields of LGN neurons have a weak opposing surround. The inhibitory and excitatory regions of the cortical neurons often are more nearly balanced in their effect. Hence, I have constructed these regions by combining the outputs from separate groups of neurons.




6.10 THE PREFERRED ORIENTATION OF NEURONS IN AREA V1 measured during single tangential penetrations. The horizontal axes show the distance along the tangential penetration, and the vertical axes show the orientation of the receptive field. (A) Data reported by Hubel and Wiesel (1977); (B) Data from a similar experiment by Hubel and Livingstone (1987). In this graph, open rectangles denote locations where responses to all orientations are equal.

The alternative interpretation is based on measurements of the spatial organization of cortical regions with common orientation preference. Obermayer and Blasdel (1993) measured regions with a common orientation preference using a high-resolution optical imaging method. In this method, a voltage-sensitive dye is applied to the cortex. Local neural activity causes reflectance changes in the dye, and these can be visualized by reflecting light from the exposed cortex. By stimulating with visual signals in different orientations and measuring the changes in reflectance, Obermayer and Blasdel (1993) visualized regions with common orientation preference; by stimulating with images originating in different eyes, they could identify ocular dominance columns (see also Hubel and Wiesel, 1977).

Figure 6.11 represents Obermayer and Blasdel's (1993) measurements as a contour plot. Regions with common orientation preference are shown as gray iso-orientation lines, and the boundaries of the ocular dominance columns are shown as dark lines. The figure shows that the variation in preferred orientation corresponds with the variation in ocular dominance. A full range of preferred orientations takes place within about 0.5 mm of the cortex, about equal to one ocular dominance column. Near the edges of the ocular dominance columns, the iso-orientation lines are arranged in linear, parallel strips extending roughly 0.5–1 mm. These linear strips are oriented nearly perpendicular to the edge of the ocular dominance edge. In the middle of the ocular dominance columns, the iso-orientation lines converge toward single points called **singularities**. In these regions, neurons with receptive fields with different preferred orientations are brought close to one another, and they may also be the position of the high density of cytochrome oxidase (Blasdel, 1992). These regions will have high

6.11 REGIONS WITH COMMON ORIENTATION PREFERENCE are shown as gray lines in this contour plot. The dark lines show the boundaries of ocular dominance columns. At the edges of the ocular dominance columns, regions with common orientation are arranged in parallel lines that are nearly perpendicular to the ocular dominance columns. These lines converge to singular points located near the center of the ocular dominance columns. From Obermayer and Blasdel, 1993.



metabolic activity since, for any stimulus orientation some of the neurons in the region will be active. This is an explanation for why regions with high densities of cytochrome oxidase tend to exhibit reduced orientation selectivity of the neural response.


There are a number of broad questions that remain unanswered about the orientation selectivity in the visual cortex. First, we might ask how the receptive field properties of cortical neurons are constructed from the cortical inputs. Figure 6.9 shows that we can explain orientation selectivity theoretically, since combining signals from center-surround neurons with adjacent receptive-field locations results in an oriented receptive field. But there is no empirical counterpart to this theoretical explanation. Second, the regularity of the iso-orientation contours shows that the orientation preferences of neurons are created in a highly regular and organized pattern. What are the rules for making the interconnections that lead to this spatial organization of orientation selectivity? What functional role do they have in perceptual processing? Is this spatial organization essential for neural computations, or is it merely a convenient wiring pattern for an area whose output is communicated to other processing modules?

Direction Selectivity

Hubel and Wiesel (1968) found a second specialization that emerges in the receptive fields of area V1 neurons. Certain cortical neurons in the monkey respond well when a stimulus moves in one direction, and poorly or not at all when the same stimulus moves in the opposite direction. This feature is called **direction selectivity**. Figure 6.12 shows the response of a neuron in monkey area V1 to a line first moving in one direction and then in the opposite direction. Notice that the cell shows



orient
tation
moves
line n
the lo
tic of
wheth
the sti
T
of the



6.12 DIRECTION SELECTIVITY of a cortical neuron's response. The left-hand portion of each panel shows the receptive-field location, the orientation of the line stimulus, and the two motion directions. The action potentials shown on the right are the neuron's response to motion in each of the two opposite directions. The neuron's response depends upon the direction of motion and the orientation of the line. After Hubel and Wiesel, 1968.

orientation selectivity; it only responds well to the line in one orientation. In addition, the cell shows direction selectivity. When the line moves up and to the right the cell responds well, but when the same line moves down and to the left the cell responds poorly. Because of the low spontaneous response rate of this neuron, which is characteristic of many cortical neurons, we cannot tell from these measurements whether the neuron simply fails to respond or is actively inhibited by the stimulus moving in the wrong direction.


The direction-selective neurons are found mainly in certain layers of the cortex and are quite rare or absent from others. The main layers

containing direction-selective neurons are 4A, 4B, 4C α , and 6 (Hawken et al., 1988). These layers receive the main input from the magnocellular pathway and send their outputs to selected brain areas. Hence, these neurons may be part of a visual stream that is specialized to carry information about motion.

Direction selectivity of the receptive-field response may arise from neural connections that are analogous to the connections underlying orientation selectivity. A cell with a direction-selective receptive field can be built by sending the outputs of neurons with spatially displaced receptive fields onto a single cortical neuron and introducing temporal delays into the path of some of the input neurons. The temporal delays of the signal are displacements of the signal in time. As we will review in more detail in Chapter 10, the result of a combined spatial and temporal displacement is to create a cortical neuron that responds better to stimuli moving in one direction, when the delay reinforces the signal, than to stimuli moving in the opposite direction, when the delay works against the two signals. This scheme for connecting neurons is plausible, but like the mechanisms of orientation selectivity, the precise neural wiring used to achieve direction selectivity in primate cortical neurons is unknown.

Contrast Sensitivity of Cortical Cells

Perhaps the most straightforward way to classify simple and complex cells is based on their responses to contrast-reversing sinusoidal patterns. Examples of the response of a simple and a complex cell to a contrast-reversing pattern are shown in Figure 6.13.




6.13 THE TIME COURSE OF RESPONSE OF CORTICAL CELLS to a contrast-reversing spatial frequency pattern at a period of 0.5 seconds. (A) The response of a simple cell is a half-wave rectified sinusoid. (B) The response of the complex cell is full-wave rectified. Consequently, the temporal response is at twice the frequency of the stimulus. Source: DeValois et al., 1982.

Recall from Chapter 5 that contrast-reversing patterns are periodic in both space and time. The stimulus used to create the neural responses shown in Figure 6.13 had a temporal period of 0.5 seconds. Figure 6.13A shows the firing rate of a simple cell averaged over many repetitions of the contrast-reversing stimulus. Were the simple cell perfectly linear, the variation in firing rate would be sinusoidal, and one period of the response would equal one period of the stimulus. This sinusoidal variation is impossible, however, because the spontaneous discharge rate of the neuron is close to zero; hence, the firing rate cannot fall below the spontaneous rate. The response shown in the figure is typical of cortical simple cells because many have a low spontaneous discharge rate. When a signal follows only the positive part of the sinusoid and has a zero response to the negative part, it is called **half-wave rectified**. The response of many simple cells shows this half-wave rectification.

Figure 6.13B shows the average response of a complex cell during one period of the stimulus. Unlike the simple cell, the complex cell's response does not vary at the same frequency as the input stimulus; the cell's response is elevated during both phases of the flickering contrast. This response pattern is called **full-wave rectification**, and the temporal response varies at twice the temporal frequency of the stimulus. This nonlinear **frequency doubling** is typical of complex cells. These cells make up a large proportion of the neurons in area V1.

DeValois et al. (1982) measured the spatial contrast-sensitivity functions of cortical neurons. Figure 6.14 shows a sample of these measure-



6.14 SPATIAL FREQUENCY SELECTIVITY of six neurons cells in area V1 of the monkey. These responses were recorded at nearby locations within the cortex, yet the neurons have different spatial frequency selectivity. Source: DeValois et al., 1982.

ments, for both simple and complex cortical neurons. The contrast-sensitivity functions of these neurons are narrower than those of retinal ganglion cells. Moreover, even though these measurements were made from neurons close to one another in the cortex, there is considerable heterogeneity in the most effective spatial frequency of the stimulus. This variation in spatial tuning is not found in retinal neurons from a single class. This may be due to a new specialization in the cortex, or it may be that we have not yet identified the classes of cortical neurons properly. In either case, the different peak spatial frequencies of the contrast-sensitivity functions raise the question of how the signals from retinal neurons within a small patch are recombined to form cortical neurons with such varied spatial receptive-field properties.


Movshon et al. (1978a, b) and Tolhurst and Dean (1987) tested the linearity of cat simple cells. Taking into account the low spontaneous rate and the resulting half-wave rectification, they found that they could predict quantitatively a range of simple-cell responses from measurements of the contrast-sensitivity function. The predictions work well for stimuli with moderate or weak contrast, that is, stimuli that evoke a response that is less than half of the maximum response rate of the neuron. There have not been extensive tests of linear receptive fields in the monkey cortex, but contrast-sensitivity curves are probably adequate to predict simple cell responses in the monkey, too.

Figure 6.14 also includes contrast-sensitivity functions of nonlinear complex neurons. Recall from our discussion in earlier chapters that when a system is nonlinear, its response to sinusoidal patterns is not a fundamental measurement of the neuron's performance: we cannot use it to predict the response to other stimuli. For these nonlinear neurons, the contrast-sensitivity function defines the response of the cell to an interesting collection of stimuli. These measurements may also help us understand the nature of the nonlinearity. However, the contrast-response function of a nonlinear system is not a complete quantitative measurement of the cell's receptive field.

Contrast Normalization

If one takes into account the low spontaneous firing rate, simple cells are approximately linear for moderate contrast stimuli. As one expands the stimulus range, however, several important response properties of cortical simple cells are nonlinear. One deviation from linearity, called **contrast normalization**, can be demonstrated by measuring the contrast-response function (cf. Figure 5.11).

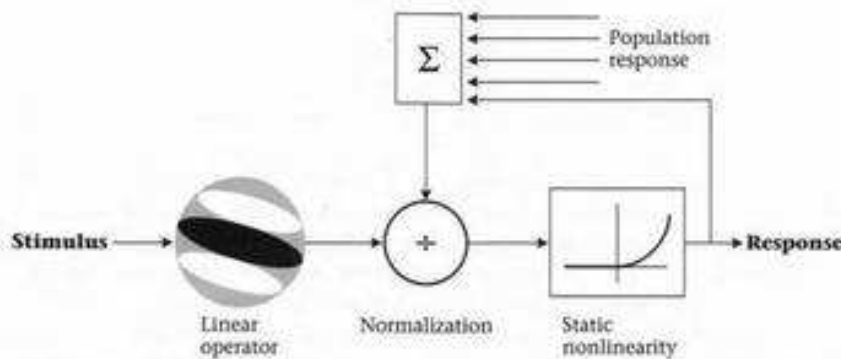
Figure 6.15 shows the contrast-response function of a neuron in area V1 to four different sinusoidal grating patterns. The stimulus contrast and neuronal responses are plotted on logarithmic axes. The rightward displacements of the curves indicate that the neuron is differentially sensitive to the spatial patterns used as test stimuli. This shift is what we expect from a simple linear system followed by a static nonlinearity (see Chapter 4).



6.15 CONTRAST-RESPONSE FUNCTIONS of a neuron in area V1. Each curve shows the responses measured using a different spatial frequency grating. The spatial frequencies of the stimuli are shown at the right. The neuron's sensitivity and maximum response depend on the stimulus spatial frequency. Source: Albrecht and Hamilton, 1982.

The entire set of data is not consistent with such a model, however, because the response saturation level depends on the spatial frequency of the stimulus. Were the nonlinearity static, then the response saturation level would be the same no matter which stimulus we used. Since the saturation level is stimulus-dependent, it cannot be based on the neuron's intrinsic properties. Rather, it must be mediated through an active process (Albrecht and Geisler, 1991; Heeger, 1992).

Heeger (1992) has described a model of this process (Figure 6.16). The model assumes that the neuron's response is initiated by a linear process. This linear signal is divided by a second signal whose value de-



6.16 A MODEL OF CONTRAST NORMALIZATION is shown. According to this model, each neuron's response is derived from an initial linear encoding of the stimulus. The linear response is divided by a factor that depends on the activity of the neural population. Finally, the entire signal is modified by a static nonlinearity. Source: Heeger, 1992.

depends on the pooled activity of the population of cortical neurons. This is a nonlinear term. It is not a static nonlinearity because the divisive term depends on the contrast of the stimulus.

This model explains the data in Figure 6.15 as follows. First, the sensitivity of the neuron varies with the spatial frequency of the stimulus because the initial linear receptive field will respond better to some stimuli than others. This causes the response to be displaced along the horizontal axis in the log-log plot. Second, the response saturation level depends on the ratio of the neuron's intrinsic sensitivity to the stimulus and the neural population's sensitivity to the stimulus. This saturation level is set by the normalization process. If the neuron is relatively insensitive to the stimulus compared to the population as a whole, then the peak response of the neuron will be suppressed by the divisive signal. Finally, the overall shape of the response function is determined by the nature of the static nonlinearity that follows.

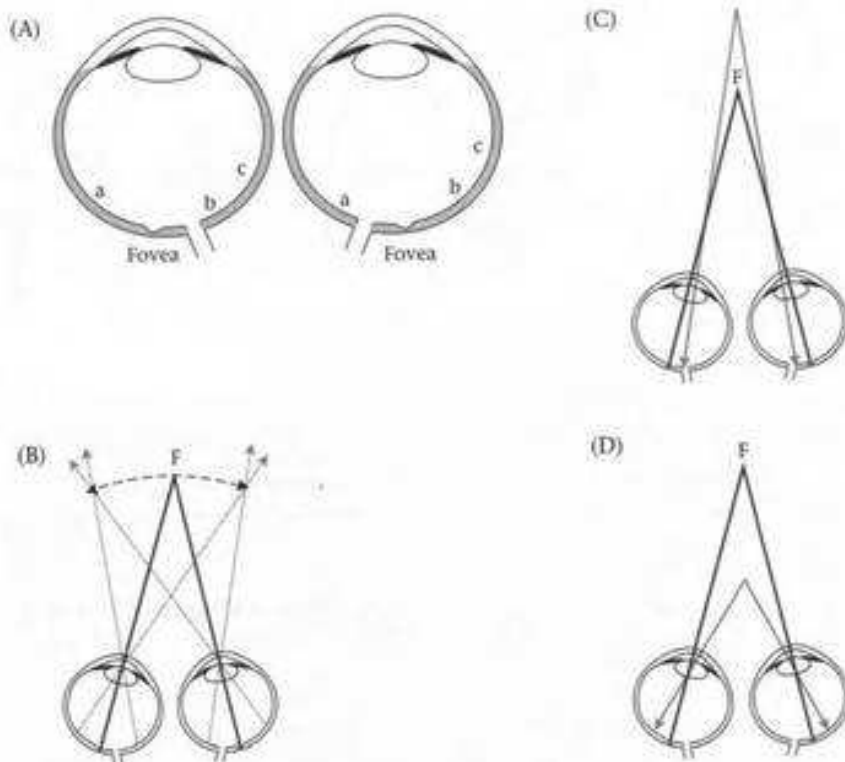
What purpose does the contrast-response nonlinearity serve? From the data in Figure 6.15, notice that the response ratio to different patterns remains approximately constant at all contrast levels. Without the contrast-normalization process, the neuron's response would saturate at the same level, independent of the stimulus. In this case, the response ratios at different contrast levels would vary. For example, at high contrast levels all area V1 neurons would be saturated, and their signals would be nondiscriminative with respect to the input signal. The normalization process adjusts the saturation level so that it depends on the neuron's sensitivity; in this way the ratio of the neuronal responses remains constant across a wide range of contrast levels.

Binocular Receptive Fields

At the input layers of the visual cortex, signals from the two eyes are spatially segregated. Within the superficial layers, however, many neurons respond to light presented to either eye. These neurons have **binocular receptive fields**. Cortical area V1 is the first point in the visual pathways where individual neurons receive binocular input. One might guess that these binocular neurons may play a role in our perception of stereo depth. What binocular information is present that neurons might use to deduce depth?

First, consider the two retinas as illustrated in Figure 6.17A. We can label points on the two retinas with respect to their distance from the fovea. We say that a pair of points on the two retinas fall at corresponding locations if they are displaced from the fovea by the same amount. Otherwise, the two points fall at noncorresponding positions.

Now, suppose that the two eyes are positioned so that a point F casts an image on the two foveas. By definition, then, the images of the point F fall on corresponding retinal locations. By tracing a ray from the corresponding retinal positions back into space, we can find the points in space whose images are cast on corresponding retinal positions (Fig-



6.17 RETINAL DISPARITY AND THE HOROPTER. (A) The fovea and three pairs of points at corresponding retinal locations are shown. (B) When the eyes are fixated on a point *F*, rays originating at corresponding points on the two retinas and passing through the lens center intersect on the horopter (dashed curve). The images of points located farther away (C) or closer (D) than the horopter do not fall at corresponding retinal locations.

ure 6.17B). These points sweep out an arc about the viewer that is called the **horopter**.

The image of a point closer or farther away than the horopter will fall on noncorresponding retinal positions. The difference between the image locations and the corresponding locations is called the **retinal disparity**. Because the main separation between the two eyes is horizontal, the retinal disparities are mainly in the horizontal direction as well. The horopter is the set of points whose images have zero retinal disparity.

Panels C and D of Figure 6.17 show two examples in which image points fall on noncorresponding retinal points. Figure 6.17C shows an example when both images fall on the nasal side of the fovea, and Figure 6.17D shows an example when both images fall on the temporal side of the fovea. These panels show that the size and nature of the horizontal retinal disparity vary with the distance from the visual horopter.

Hence, the horizontal retinal disparity is a binocular clue for estimating the distance to an image point.⁴

Do binocular neurons represent stereo depth information by measuring horizontal disparity? There are two types of experimental measurements we can make to answer this question. First, we can measure the receptive fields of individual binocular neurons. If retinal disparity is used to estimate depth, then the receptive fields of the binocular neurons should show some selectivity for horizontal disparity. Second, we can look at the properties of the population of binocular neurons. While no single neuron alone can encode depth information, the population of binocular neurons should include enough information to permit the population to estimate image depth.

A complete characterization of binocular receptive fields requires many measurements. First, one would like to measure the spatial receptive fields of individual neurons when stimulated by each eye alone. These are called the **monocular receptive fields** of the binocular neurons. Then, we should characterize how binocular neurons respond to simultaneous stimulation of the two eyes. In practice there have been very few complete measurements of binocular neurons' receptive fields. The vast majority of investigations have been limited to localization of the monocular receptive-field centers which are then used to derive the retinal disparities between the monocular field centers.

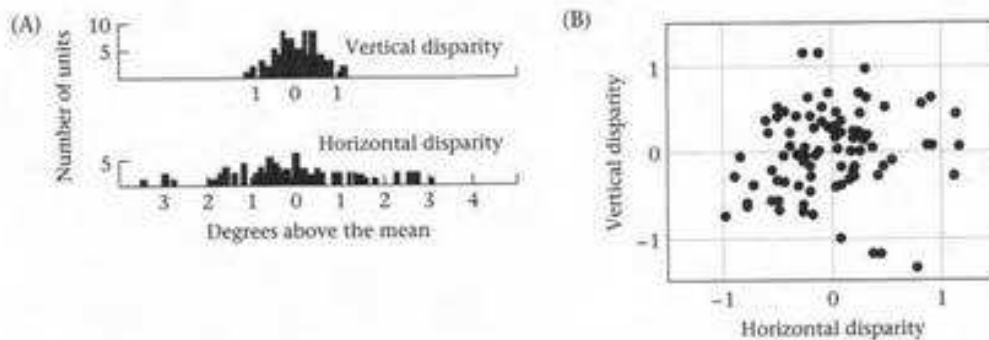
Given the variability inherent in biological systems, the two monocular receptive fields will not be in perfect register. We would like to decide whether the observed horizontal disparities serve a purpose, or whether they are due to unavoidable random variation. To answer this question several groups have measured both the horizontal and the vertical disparities of binocular neurons in the cat cortex (Barlow et al., 1967; Joshua and Bishop, 1970; von der Heydt et al., 1978). The histograms in Figure 6.18A show the initial measurements from Barlow et al. (1967). They observed more variability in the horizontal disparity than in vertical disparity, and they concluded that the horizontal variation was purposeful and used for processing depth. Joshua and Bishop (1970) and von der Heydt et al. (1978) saw no difference in the range of disparities in the horizontal and vertical directions. A scatterplot of the retinal disparities observed by Joshua and Bishop (1970) is shown in Figure 6.18B. While these data do not show any systematic difference between the horizontal and vertical disparity distributions, the authors do not dispute Barlow et al.'s hypothesis that variations in the horizontal disparity are used for stereo depth detection.⁵

⁴You can demonstrate the relative shift in retinal positions to yourself as follows. Focus on a nearby object, say your finger placed in front of your nose. Then, alternately look through one eye and then the other. Although your finger remains focused in the fovea, the relative positions of points nearer or farther away than your finger will change as you look through each eye in turn.

⁵A frequently suggested alternative is that these disparity cues serve to converge the two eyes. Since the same cues are used to converge the eyes and estimate depth, this alternative hypothesis is virtually impossible to rule out.

(A)

6.11
the
of t
dis
besdis
Ho
rorspe
ity
po
pa
Or
is
th
res
de
poof
br.
we
di
wi
th
be
ul
di
reco
cu
el




6.18 THE HORIZONTAL AND VERTICAL DISPARITIES OF BINOCULAR NEURONS in the cat visual cortex are shown. (A) Histograms of the horizontal and vertical disparities of binocular neurons in cat cortex. (B) A scatter diagram of the vertical and horizontal disparities of cells in cat cortex with receptive fields located within 4 degrees of the cat's best region of visual acuity. Sources: A from Barlow et al., 1967; B from Bishop, 1973.

For the moment, accept the premise that the variation in horizontal disparity of these binocular neurons is a neural basis for stereo depth. How might we design the binocular response properties of these neurons to estimate depth?

One possibility is to create a collection of neurons that each respond to only a single disparity. One might estimate the local disparity by identifying the neuron with the largest response. An alternative possibility, suggested by Richards (1971), is that one might measure disparity by creating a few **pools** of neurons with coarse disparity tuning. One pool might consist of neurons that respond when an object feature is beyond the horopter, and a second pool might consist of neurons that respond when the feature is in front of it. The third pool might respond only when the feature is close to the horopter. To estimate depth, one would compare the relative responses in the three neural pools.

Some support for Richards's hypothesis comes from measurements of individual neurons in area V1 and an adjacent area of a monkey brain. Poggio and Talbot (1981; see also Ferster, 1981) measured how well individual neurons respond to stimuli with different amounts of disparity. They used experimental stimuli consisting of bar patterns whose width and velocity were set to generate a strong response from the individual neuron. The experimenters varied the retinal disparity between the two bars presented to the two eyes. They plotted the binocular neuron's response to the moving bars as a function of their retinal disparity. The curves in Figure 6.19, plotting response as a function of retinal disparity, are called **disparity-tuning curves**.

Poggio and Talbot (1981) found that the disparity-tuning curves could be grouped into a small number of categories. Typical tuning curves from each of these categories are illustrated in the separate panels of Figure 6.19. The two neurons illustrated in panels (A) and (B)




6.19 DISPARITY-TUNING CURVES of binocular neurons in areas V1 and V2 in a monkey. Each panel plots the response of a different neuron to moving bar patterns. The independent variable is the retinal disparity of the stimulus. Panels (A) and (B) show the responses of neurons that respond best to stimuli with near zero disparity, that is, near the horopter. Responses of a neuron that responds best to stimuli with positive disparity (C) and a neuron with negative disparity (D) are also shown. The curves represent data measured using binocular stimulation. Source: Poggio and Talbot, 1981.

respond to disparities near the fixation plane; for these neurons, stimuli near the horopter stimulate or inhibit the cell. Panels (C) and (D) illustrate neurons with opponent tuning. One neuron is excited by a bar whose disparity places the object beyond the horopter and is inhibited by bars in front of the horopter. The second neuron shows approximately the complementary excitation pattern: it is excited by objects nearer than the horopter and inhibited by objects farther away. Poggio and his colleagues view their measurements in monkeys as support for Richards's hypothesis that the encoding of binocular depth is based on the response of neurons organized in disparity pools (also see Ferster, 1981).

We have been paying attention mainly to the retinal disparity of the binocular neurons. But disparity tuning is only one measure of the receptive-field properties of these neurons. In addition, the receptive fields must have spatial, temporal, and chromatic selectivities. To fully understand the responses of these neurons we must make some progress in measuring all of these properties.

To obtain a more complete description of binocular neurons, Freeman and Ohzawa (1990; also see DeAngelis et al., 1991) studied the fundamental issue of how signals from the two eyes converge to create responses in the binocular neurons. Specifically, Freeman and Ohzawa measured the response of binocular neurons to signals that originated in each eye alone (i.e., monocular receptive fields) and to stimuli that were presented to both eyes (binocular receptive fields). One of their most striking findings is that the two monocular spatial receptive fields of a single neuron can be quite different from one another, as shown in Figure 6.20, an example of the right and left monocular receptive fields of a single binocular neuron in the visual cortex of the cat. The left-eye spatial receptive field and the right-eye field are displaced relative to one another. If we only concern ourselves with disparity, we will report that this cell's receptive field has significant horizontal disparity. Notice, however, that the spatial receptive fields are different from one another. The spatial receptive field in the left eye is a mirror-reversal of the field in the right eye.

Freeman and Ohzawa suggest that these different monocular spatial receptive fields are important to the way stereo depth is estimated by the nervous system. They hypothesize that stereo depth depends on having neurons with different monocular spatial receptive fields. Perhaps most important, however, their measurements remind us that, in order to understand the biological computation of stereopsis, we must



6.20 MONOCULAR SPATIAL RECEPTIVE FIELDS of two binocular neurons in cat cortex. The panels show examples of left and right monocular receptive fields whose centers are displaced horizontally and thus have nonzero retinal disparity. In addition to the disparity, the left and right monocular spatial receptive fields differ. Source: Freeman and Ohzawa, 1990.

study more than just the center position of the monocular receptive fields.

Visual Streams in the Cortex

We have reviewed two major principles that characterize the flow of information from retina to cortex. First, visual information is organized into separate visual streams. These streams begin in the retina and continue along separate neural pathways into the brain. Second, the receptive-field properties of neurons become progressively more sophisticated. Receptive fields of cortical neurons show selective responses to stimulus properties that are more complex than retinal neurons. The new receptive-field properties are clues about the specialization of the computations performed within the visual cortex.

As we study visual processing within the cortex we should expect to see both of these principles extended. First, we should expect to find new visual streams that play a role in the cortical computations. Some new visual streams will arise in the visual cortex, and some, like the rod pathway in the retina, will have served their purpose and will merge with other streams. Second, as we explore the cortex we should expect to find neurons with new receptive-field properties. We will need to characterize these receptive fields adequately in order to understand their computational role in vision.

Our understanding of cortical visual areas is in an early and exciting phase of scientific study. In this section, we will review some of the basic organizational principles of the cortical areas. In particular, we will review how information from area V1 is distributed to other cortical areas, and we will review the experimental and logical methods that relate activity within these cortical areas to what we see. We will review some of the more recent data and speculative theories in Chapters 9 and 10.

The Fate of the Parvocellular and Magnocellular Pathways

The segregation of visual information into separate streams is an important organizing principle of neural representation. Two of the best-understood streams are the magnocellular and parvocellular pathways, whose axons terminate in layers $4C\alpha$ and $4C\beta$ within area V1. What happens to the signals from these pathways within the visual cortex?

Along one branch, signals from the magnocellular pathway continue from area V1 directly to a distinct cortical area. The magnocellular pathway in layer $4C\alpha$ makes a connection to neurons in layer 4B where there are many direction-selective neurons. These neurons then send a strong projection to cortical area MT (middle temporal). It seems reasonable to suppose, then, that the information contained within the magnocellular stream is of particular relevance for the visual process-

ing
has
cor
in
MT
to
wil

inc
pat
al.
sur
no
me
tio
al.,
an
ins
lul
als
tra
pu
Th
tril
thi
foi

wi
mi
ga
co
ter
m:
fre

Cy
Li
er:
a
m
da
he
th
fo
se
ap


ing in area MT. As we saw in Chapter 5, the magnocellular pathway has particularly good information about the high-temporal-frequency components of the image. Earlier in this chapter we saw that neurons in layer 4B show strong direction selectivity, as do the neurons in area MT (Zeki, 1974, 1990a,b). Taken together, these observations have led to the hypothesis that area MT plays a role in motion perception. We will discuss this point more fully in Chapter 10.

While one branch of the magnocellular stream continues on an independent path, another branch of this stream converges with the parvocellular pathway in the superficial layers of area V1. Malpeli et al. (1981) and Nealey and Maunsell (1994) made physiological measurements demonstrating that signals from the parvocellular and magnocellular streams converge on individual neurons. In these experiments parvocellular or magnocellular signals were blocked by application of either a local anaesthetic, lidocaine hydrochloride (Malpeli et al., 1981), or the neurotransmitter γ -aminobutyric acid (GABA) (Nealey and Maunsell, 1994) to small regions of the LGN. Both studies report instances of neurons whose responses are influenced by both parvocellular and magnocellular blocking. Anatomical paths for this signal have also been identified. Lachica et al. (1992) injected retrograde anatomical tracers (i.e., tracers that are carried from the injection site toward the inputs to the injection site) into the superficial layers of the visual cortex. They concluded that the magnocellular and parvocellular neurons contribute inputs into overlapping regions within the superficial layers of the visual cortex. Hence, these anatomical pathways could be the route for the physiological signals.

Just as the rod pathways are segregated for a time, and then merge with the cone pathways, so too signals from the magnocellular stream merge with parvocellular signals. The purpose of the peripheral segregation of the parvocellular and magnocellular signals, then, may be to communicate rapidly certain types of image information to area MT. After the signal has been efficiently communicated, the same information may be used by other cortical areas, in combination with information from the parvocellular pathways.

Cytochrome Oxidase Staining

Livingstone and Hubel (1984a, b, 1987a, b, 1988) have argued that several new visual streams begin in area V1. Their argument begins with a discovery made by Wong-Riley (1979), who used a histochemical marker to detect the presence of an enzyme called cytochrome oxidase, and found that the enzyme is present mainly in a continuous and heavy pattern in layer 4, but is present mainly as a set of patches in the superficial and deeper layers. These patches can be seen by staining for cytochrome oxidase and viewing the superficial layers in tangential section. In that case, the regions of high cytochrome oxidase density appear as a set of darkened spots (Figure 6.21). These darkened regions



6.21 CYTOCHROME OXIDASE PUFFS ARE LOCATED IN THE MIDDLE OF OCULAR DOMINANCE COLUMNS. The three panels show tangential sections of area V1. (A) The locations of the ocular dominance columns in layer 4C correspond to the light and dark patterns. The columns were identified by a monocular injection of the radioactive tracer tritiated proline. (B) The puffs are the darkened spots identified by a high concentration of cytochrome oxidase. The image is a tangential section in the superficial layers. (C) An overlay of (A) and (B), placed in register by comparing the positions of larger blood vessels, shows that the puffs are located in the center of the ocular dominance columns. Scale bar = 1 mm. Source: Horton and Hubel, 1981.

are called variously "puffs," "blobs," or "CO-rich" areas. These puffs are visible in several, but not all, primate species, including humans. Cytochrome oxidase density is correlated with regions of high neuronal activity, though the enzyme itself is not known to have any specific significance for visual function (Wong-Riley, 1979; Humphrey and Hendrickson, 1980; Livingstone and Hubel, 1982; Hendrickson, 1985).

Several lines of evidence suggest that cytochrome oxidase labeling is correlated with the presence of distinct visual streams in the cortex. First, Horton and Hubel (1981) discovered that the cytochrome oxidase staining pattern is related to at least one aspect of visual function: the ocular dominance columns. Figure 6.21A shows the ocular dominance columns measured by injecting one eye of a monkey with a radioactive tracer, tritiated proline. The section is through layer 4C (and a little bit of layer 5) where the ocular dominance columns are best segregated. The data in this figure replicate the demonstration of ocular dominance columns depicted in Figure 6.4. Figure 6.21B shows the puffs in a tangential section from layers 2 and 3, just above the region shown in Figure 6.21A. Figure 6.21C shows an overlay of the ocular dominance columns and the puffs. The images were placed into spatial register by finding blood vessels common to the two images and aligning their positions. Figure 6.21C demonstrates that the puffs trace a path down the center of the ocular dominance columns.

The second suggestion that the presence of cytochrome oxidase identifies distinct visual streams comes from studying the connectivity between neurons within area V1. Burkhalter et al. (1989) report that the

neurons within the V1 puffs in humans are connected to other neurons in the puffs, while neurons between the puffs are connected to other neurons between the puffs (Rockland and Lund, 1983).

A third piece of evidence is the relationship between cytochrome oxidase puffs and connections to other cortical areas. Like area V1, cytochrome oxidase is distributed unevenly in a second cortical area, V2, which is adjacent to area V1. In area V2, staining for cytochrome oxidase yields a regular three-striped pattern defining regions that stain to different degrees. The three types of stripes are labeled **thick stripes**, **thin stripes**, and **interstripes**. The thick and thin stripes contain more cytochrome oxidase than the interstripe regions. This pattern of stripes is visible in human area V2 as well as in some species of monkey (Tootell et al., 1983).

Livingstone and Hubel (1982, 1987a) demonstrated the specificity of these interconnections. Regions with a high density of cytochrome oxidase in area V1, that is, regions with high metabolic activity, are connected to high-density regions in area V2. Neurons in layer 4B of area V1 send outputs to the thick stripes of area V2. Neurons in the puffs send outputs to the thin stripes. Neurons in between the puffs send outputs to the interstripes.

Finally, the pattern of cytochrome oxidase staining of area V2 correlates with the connections from area V2 to other visual areas. In monkeys, visual area MT receives a strong projection from the thick stripes in area V2. **Area V4** receives signals mainly from the thin stripes and the interstripes (Deyoe and Van Essen, 1988; Merigan and Maunsell, 1993). These measurements suggest that there is considerable organization of the signals communicated within individual cortical areas.

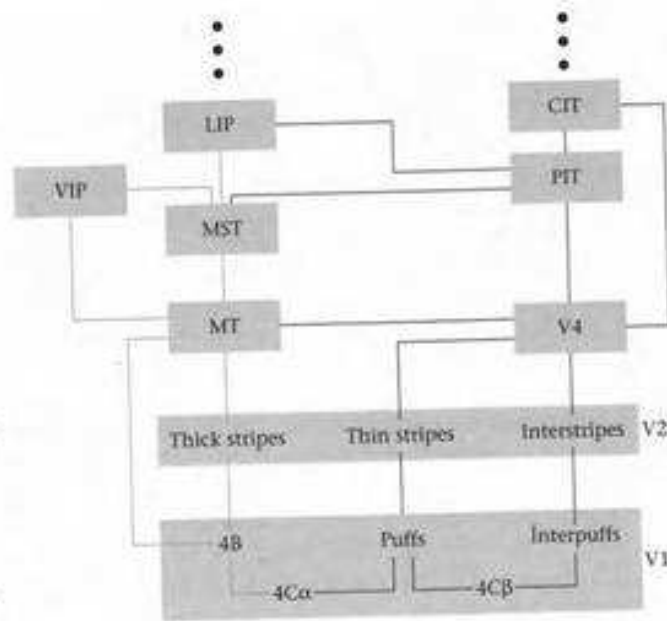
Areas Central to Primary Cortex

Figure 6.22 shows a few of the many cortical visual areas. Often, the connections between areas are reciprocal; ascending axons make connections with the primary input layer, 4C, and descending axons make connections in layers 1 and 6. A single visual area can have connections with several other cortical areas (Rockland and Pandya, 1979; Felleman and Van Essen, 1991).

Figure 6.22 is arranged to emphasize one aspect of the segregation of visual information within the cortex, namely, that information from the visual areas in the occipital lobe separates into two visual streams. One stream sends its outputs mainly to the posterior portion of parietal lobe, while the second stream makes its connection mainly in the inferior portion of the temporal lobe.

There are several known connections between these two streams, but a study by Baizer et al. (1991) shows that the segregation is impressive. These authors injected large amounts of two retrograde tracers into individual monkey brains: one tracer was injected into the posterior parietal lobe and the other into the inferior temporal cortex. They

6.22 AN OVERVIEW OF THE ORGANIZATION OF VISUAL AREAS is shown. The signals from the occipital lobe areas are sent along two major streams. One stream passes into the posterior parietal lobe and a second to the inferior temporal lobe. The visual signals in the temporal and parietal lobe areas appear to arise from different neurons within the occipital lobe. MST = medial superior temporal area, MT = medial temporal area, VIP = ventral intraparietal area, PIT = posterior inferotemporal area, LIP = lateral intraparietal area, CIT = central inferotemporal area. Source: Merigan and Maunsell, 1993.



examined where the two types of tracers could be found in visual areas within the occipital lobe, including areas V1, V2, V4, and MT. They report finding almost no neurons that contained both tracers, suggesting that the signals from individual neurons are communicated mainly to either the parietal or temporal lobes. Baizer et al. also report that neurons in the parietal lobe received information mainly from neurons with receptive fields located in the periphery, while neurons in the temporal lobe received information mainly from neurons with receptive fields located near the fovea.

Baizer et al.'s observations support Ungerleider and Mishkin's (1982) proposal that the parietal and temporal streams serve different visual functions. Ungerleider and Mishkin observed that clinical damage to the parietal stream of one hemisphere causes difficulties in visual and motor orientation. It also causes **hemineglect**, a condition in which the observer appears to be unaware of stimulation arising in the hemifield that projects to that parietal lobe. Such patients also have trouble orienting toward or reaching for objects in the visual periphery. The clinical symptoms associated with damage to the temporal stream are quite different. In this case, patients have impaired form discrimination or recognition. They also have problems with visual memory. These behavioral deficits are very different, so that neurologists have supposed that the visual function of the parietal and temporal lobes are quite different. One brief characterization of the distinction is this: the parietal system defines *where*, while the temporal system defines *what* (Ungerleider and Mishkin, 1982; Merigan and Maunsell, 1993).

occ
etal
it se
mu
fere



The
leap
tion
ma
wa
sen
ing
me

us
spe
the
Cri
nes

abc
ity
for
ha
ne
ava
ha
to
tir
col
for

dis
ey
cor
rie
lay

un
of-
ste
pr

The anatomical segregation of the neural signals coming from the occipital lobe into these two streams shows that the temporal and parietal areas receive different information about the visual image. Hence, it seems likely that the computations in these two portions of the brain must serve different goals, though a more refined analysis of these differences will be helpful.

Cortical Representations and Perception

The brilliant visual scientist William Rushton enjoyed needling his colleagues. On one occasion, he challenged neuroscientists with the assertion that only the hope of understanding perception and consciousness makes neuroscience worth doing. Much of the work I have reviewed was inspired by the desire to understand conscious perception. In this sense, some neuroscientists think of themselves as philosophers studying the mind-body problem. They form a field one might call experimental philosophy; it is the expensive branch of philosophy.

More recently Crick (1993) has taken up Rushton's call and pressed us to consider the question: what aspect of the cortical response corresponds to conscious experience? Since neuroscientists nearly all make the assumption that consciousness is a correlate of cortical activity, Crick points out that identifying the relationship between consciousness and neural activity is properly an experimental question.

What experimental tests might we perform to answer questions about the relationship between our conscious awareness and the activity of our brains? One way to study this question is to compare the information available within a cortical area with the visual experience we have. An interesting example of an experiment concerning consciousness is the comparison of conscious awareness with the information available in area V1. From the anatomy and physiology of area V1, we have learned that there is plenty of information there that we can use to deduce which eye is the source of a visual signal. Within layer 4C, entirely different sets of neurons, confined to different ocular dominance columns, respond depending on which eye sends the signal. Is this information available to us?

We can answer this question experimentally by asking subjects to discriminate between visual signals originating in the right and left eyes. If we can accurately decide on the eye of origin, then we might conclude that the information in area V1 is part of our conscious experience. If we cannot, then we should conclude that information within layer 4C can be lost prior to reaching our conscious experience.

Notice that eye-of-origin information is certainly available to us unconsciously. For example, Helmholtz (1896) pointed out that eye-of-origin information is necessary and used for the computation of stereopsis. The question, therefore, is not whether the information is present, but whether it is accessible to conscious experience. Based on

his own introspections, Helmholtz answered the question in the negative. And, while there have been occasional reports that some discriminations are possible, Ono and Barbeito's (1985) careful experiments suggest that no reliable eye-of-origin discriminations are possible. Hence, the massive information available in the input layers of area V1 (and earlier) about the eye-of-origin is not part of our conscious experience. Through this negative result, we have made a small amount of progress in localizing consciousness.

The Function of the Visual Areas

Even when the computation performed in a visual area is not part of our conscious experience, we would still like to know what the area does. Over the last 15 years, there have been a broad variety of hypotheses concerning the perceptual significance of the cortical areas. Mainly, we have seen a flurry of proposals suggesting that individual visual areas are responsible for the computation of specific perceptual features, such as color, stereo, form, and so forth.

What is the logical and experimental basis for reasoning about the perceptual significance of visual areas? Horace Barlow (1972) has set forth one specific doctrine to relate neurons to perception, the **neuron doctrine**. This doctrine asserts that a neuron's receptive field describes the percept caused by excitation of the neuron. You will see the idea expressed many times as you read through the primary literature and study how investigators interpret the perceptual significance of neural responses.

Our understanding of the peripheral representation lends little support to the neuron doctrine. For example, the principle does not serve us well when analyzing color appearance. In that case, we know with some certainty that a large response from an L photoreceptor does not imply that the observer will perceive red at the corresponding location in the visual field. Rather, the color appearance depends upon stimulation at many adjacent points of the retina. The conditions for a red percept include a pattern of peripheral neural responses, including more L- than M-cone excitation. Data from the periphery are generally more consistent with the notion of a representation in which an experience depends on the response of a collection of neurons.

Oddly, the failure of the neuron doctrine in the periphery is often used to support the neuron doctrine. After all, the argument goes, the periphery is not the site of our conscious awareness, so failures of the doctrine in the periphery are to be expected. The neuron doctrine's significance depends on the idea that there will be a special place, probably located in the cortex, where the receptive field of a neuron predicts conscious experience when that neuron is active. This location in the brain should only exist at a point after the perceptual computations needed to see features we perceive (e.g., color, form, depth) have taken place.

wor
ron
per
con
the
as i
per

cor
Ess
(19

the
on
foi
V4
m
th.
ne
te

tai
TF
ov
eli
di
ce
as
dc
se
in
ta

in
re
at
th

In the past, secondary texts sometimes cited Hubel and Wiesel's work in area V1 as providing an example of a location where the neuron doctrine might hold. The receptive fields in area V1 seem like basic perceptual features; orientation, motion selectivity, binocularity, and complex cells all emerge for the first time in area V1. Consequently, these texts often described our knowledge of receptive fields in area V1 as if it were a theory of vision, with the receptive fields defining salient perceptual features.

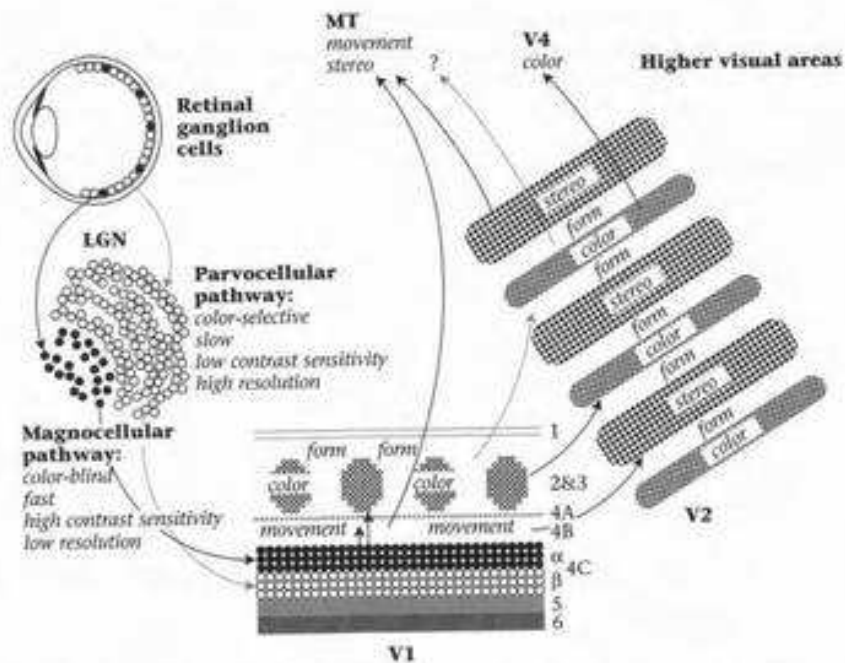
However, by 1979 the significance of the other cortical areas had become undeniable (Zeki, 1974, 1978; Felleman and Van Essen, 1991; Van Essen et al., 1992). In reviewing the visual pathways, Hubel and Wiesel (1979) wrote:

The lateral geniculate cells in turn send their axons directly to the primary visual cortex. From there, after several synapses, the messages are sent to a number of further destinations: neighboring cortical areas and also several targets deep in the brain. One contingent even projects back to the lateral geniculate bodies; the function of this feedback path is not known. The main point for the moment is that the primary visual cortex is in no sense the end of the visual path. It is just one stage, probably an early one in terms of the degree of abstraction of the information it handles.

Acknowledging this point leads one to ask, What is the function of these other cortical areas? The answer to this question has relied mainly on the neuron doctrine. For example, when Zeki (1980, 1983, 1993) found that color contrast was a particularly effective stimulus in area V4, he argued that this area is responsible for color perception. Since movement was particularly effective in stimulating neurons in area MT, that became the motion area (Dubner and Zeki, 1971). The logic of the neuron doctrine permits one to interpret receptive-field properties in terms of perceptual function.

A particularly vigorous application of the neuron doctrine is contained in articles by Livingstone and Hubel (1984a,b, 1987a,b, 1988). They supported Zeki's basic views and added new hypotheses of their own. Their hypothesis, which continues to evolve, is summarized in the elaborate anatomical/perceptual diagram shown in Figure 6.23. In this diagram, anatomical connections in visual cortex are labeled with perceptual tags, including color, motion, and form. The logical basis for associating perceptual tags with these anatomical streams is the neuron doctrine. Receptive fields of neurons in one stream were orientation-selective; hence the stream was tagged with form perception. Neurons in a different stream were motion-selective, and hence the stream was tagged with motion perception.

The perceptual-anatomical hypotheses proposed by Zeki and by Livingstone and Hubel define a new view of cortex. On this view, the relationship between cortical neurons and perception should be made at the level of perceptual features. These investigators did not study the computation within the neural streams, but rather they summa-



6.23 AN ANATOMICAL/PERCEPTUAL MODEL OF THE VISUAL CORTEX. In this speculative model, visual streams within the cortex are identified with specific perceptual features. The anatomical streams are identified using anatomical markers; the perceptual properties are associated with the streams by applying the neuron doctrine. Source: Livingstone and Hubel (1988).

alized the receptive field properties and based their interpretations on the neuron doctrine (see Hubel and Wiesel [1977] for a description of this approach).

The use of the neuron doctrine to interpret brain function is very widespread, but there is very little evidence in direct support of the doctrine (Martin, 1992). The main virtue of the hypothesis is the absence of an articulated alternative. The most frequently cited alternative is the proposal that perceptual experience is represented by the activity of many neurons, so that no individual neuron's response corresponds to a conscious perceptual event. Such models are often called **distributed processing models**; they are not widely used by neurophysiologists since they do not provide the specific guidance for interpreting experimental measurements from single neurons, the neurophysiologist's stock-in-trade. The neuron doctrine, on the other hand, provides an immediate answer.

In my own thinking about brain function, I am more inclined to wonder about the brain's computational methods than the mapping between perceptual features and tentatively identified visual streams. I

find it satisfying to learn that the magnocellular pathway contains the best representation of high temporal frequencies, but less satisfying to summarize the pathway as the motion pathway, since high temporal frequency information may also be used in many other types of performance tasks. The questions I find fundamental concerning computation are "how?" not "where?" How are essential signal-processing tasks, such as multiplication, addition, and signal synchronization, carried out by the cortical circuitry? What means are used to store temporary results, and what means are used to represent the final results of computations? What decision mechanisms are used to route information from one place to another?

Although the neuron doctrine is widely used because it is an easy tool to relate perception and brain function, my own view is that the doctrine distracts us from the most important question about visual function: how do we *compute* perceptual features like color, stereo, and form? Even if it turns out that a neuron's receptive field is predictive of experience, the question we should be asking is how the neuron's receptive-field properties arise. Answering these computational questions will help us most in designing practical applications that range from sensory prostheses to robotics applications. We should view the specific structures within the visual pathways as a means of implementing these principles, rather than as having an intrinsic importance.

Hubel and Wiesel (1979) once expressed something like this view. While reviewing their accomplishments in the study of area V1, they wrote (p. 23):

What happens beyond the primary visual area, and how is the information on orientation exploited at later stages? Is one to imagine ultimately finding a cell that responds specifically to some very particular item? (Usually one's grandmother is selected as the particular item, for reasons that escape us.) Our answer is that we doubt there is such a cell, but we have no good alternative to offer. To speculate broadly on how the brain may work is fortunately not the only course open to investigators. To explore the brain is more fun and seems to be more profitable.

There was a time, not so long ago, when one looked at the millions of neurons in the various layers of the cortex and wondered if anyone would ever have any idea of their function. Did they all work in parallel, like the cells of the liver or the kidney, achieving their objectives by pure bulk, or were they each doing something special? For the visual cortex the answer seems now to be known in broad outline: Particular stimuli turn neurons on or off; groups of neurons do indeed perform particular transformations. It seems reasonable to think that if the secrets of a few regions such as this one can be unlocked, other regions will also in time give up their secrets.

In the remaining chapters, we will see how other areas of vision science, based on behavioral and computational studies, might help us to unlock the secrets of vision.

Exercises

1. Answer these questions about forming cortical receptive fields from the inputs of LGN neurons.
 - (a) What is an orientation column?
 - (b) Suppose the cortical cell response is the sum of two LGN neurons' outputs. Describe the LGN neurons' receptive fields needed for the cortical cell receptive field to be circularly symmetric.
 - (c) Use a matrix representation to describe how the receptive field of a linear cortical neuron depends on the linear receptive fields of a collection of LGN neurons.
 - (d) Use the matrix representation to design a weighted combination of LGN responses that indicates how much each LGN neuron's response should contribute to the cortical cell's response in order to achieve different receptive-field functions.

2. Answer these questions about retinotopic organization.
 - (a) What is a retinotopic map?
 - (b) Certain neurons in the central nervous system have very large receptive fields, spanning 20 degrees of visual angle. Can visual areas containing such neurons have retinotopic maps? Experimentally, how would you convince yourself that such an area had a retinotopic organization?
 - (c) Draw a picture describing the logical organization of an area of visual cortex that is retinotopic and also has orientation columns. Are the two types of organization necessarily linked to one another, or can you imagine that each would follow its own independent layout?
 - (d) How many different kinds of organization can be simultaneously superimposed within a single cortical region?

3. Lesion studies have been an important source of information about neural function. The logical foundation of lesion studies, however, is quite involved. Often, it is difficult to be precise about the conclusions one may draw from a lesion study. The problems of interpreting lesion studies can be illustrated by this old joke:

A scientist once decided to study binocular vision. He removed the right eye of a cat and observed that the cat lost its stereo vision. He then wrote an article describing how stereovision was localized to the right eye.

To follow on his groundbreaking work, this scientist studied the animal's behavior again and noted that the animal could perform monocular tasks. He enucleated the left eye and observed that the animal could no longer perform such tasks. The scientist wrote a second paper describing how monocular vision is localized to the left eye.

Write a paragraph that describes what conclusions the scientist should have drawn. Choose your words carefully.

4. Answer these questions about color coding.
 - (a) How would you measure the wavelength responsivity of a cortical neuron?

- (b) Recall the definition of separability for space and time in the receptive fields of neurons from Chapter 5. Make an analogous definition for wavelength-space separability.
 - (c) Suppose you measure the wavelength encoding of four neurons in the cortex. Recall that the initial encoding of wavelength is based on three types of cones. Do the wavelength encoding properties of any two of the neurons have to be precisely the same?
 - (d) If the four neurons have different wavelength encoding properties, will this have any implications for trichromacy?
 - (e) Suppose the neurons respond linearly to wavelength mixtures. Describe how you might be able to infer the connections between the neurons and the three different cone classes.
5. "Functional specialization" refers to the notion that different brain areas are specialized for certain visual functions, such as the perception of motion or color.
- (a) What interconnections would you anticipate finding between different brain areas if each one is specialized to represent a different perceptual function?
 - (b) What properties do you think a neuron should exhibit before we believe that it is specialized for a visual function?
 - (c) What special role might the LGN play? Could it be connected with eye movements? Saccadic suppression? Attentional gating? Synchronization of the images from the two eyes? How would you study this question?