

Fundamentals of Computer Vision
308-558B

Biological Vision

Prepared By
Louis Simard

1. Optical system

1.1 Overview

The ocular optical system of a human is seen to produce a transformation of the light energy of the visual input stimulus impinging on the eye to an output which is similarly a high-energy signal. A horizontal cross section of the human eye is shown in Figure 1. The input light pattern enters the cornea and then passes in sequence through the anterior chamber, the pupil opening of the iris, the lens, the vitreous humor before impinging on the layer of photoreceptors which constitutes the retina at the rear. The latter is responsible for the actual transduction from light energy into electrical energy in the form of a train of frequency-modulated pulses.

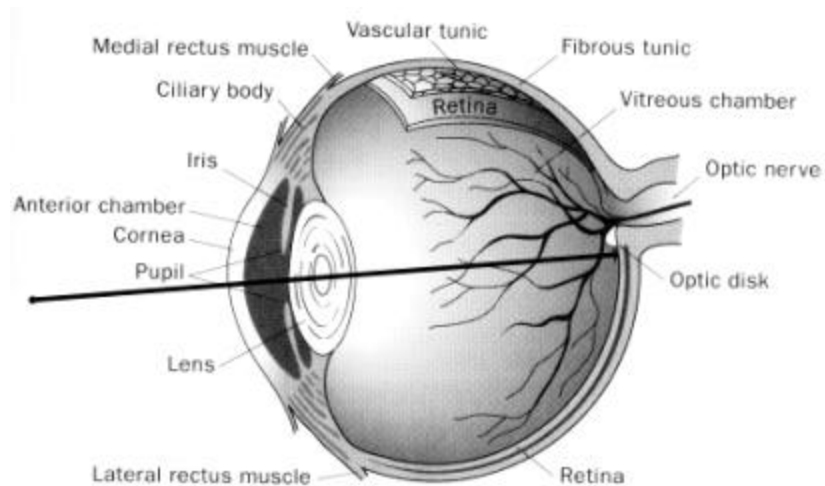


Figure 1

1.2 Retina

The retina resembles a very thin, fragile meshwork. Although no thicker than a postage stamp, this delicate structure has a complex, layered organization. Figure 2 shows how a section of the retina would look if viewed from the side. The arrows represent the direction taken by incoming light. The incoming light must pass through a complex of neural elements before reaching the photoreceptors which are actually responsible for converting light energy into neural signals. Those neural signals, in turn, pass through a network of diverse cells we lump together under the rubric collector cells. These so-called collector cells consist of three major classes of neurons: bipolar cells, amacrine cells, and horizontal cells. Together these gather and modify information registered by the receptors. The output from the network of collector cells provides the input to the retinal ganglion cells whose axons form the optic nerve.

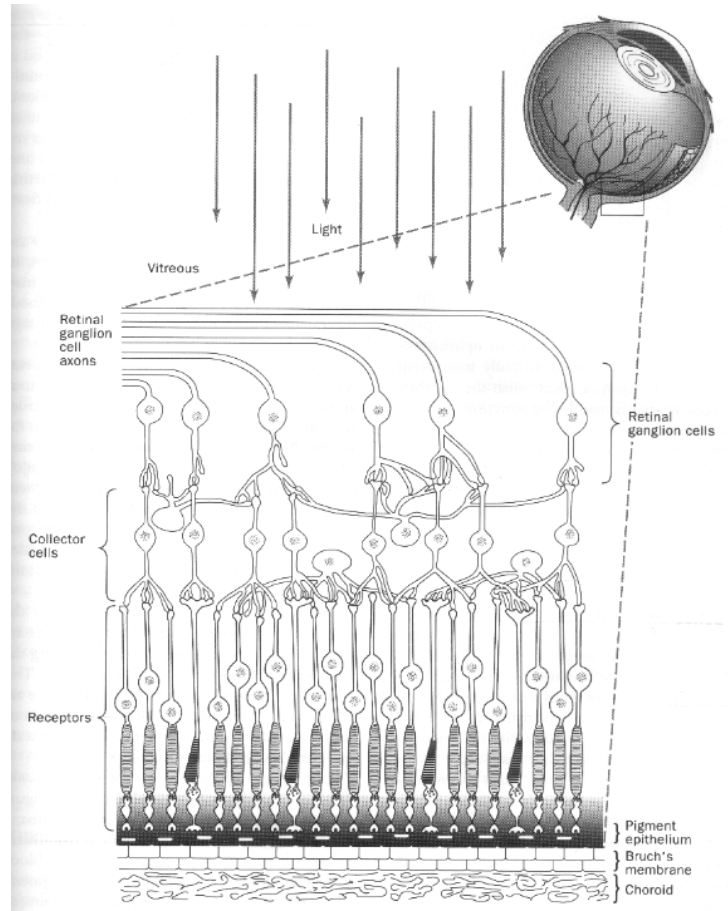


Figure 2

1.2.1 Photoreceptors

The human eye contains two major classes of photoreceptors: rods and cones. The characteristics of both are described in Table X.

| Rod Vision (scotopic) | Cone Vision (photopic) |
|--------------------------------------------|-------------------------------|
| 120 million rods | 8 million cones |
| Peripheral vision | Central vision |
| Massive convergence | Little convergence |
| Black and white vision | Colour vision |
| Low resolution | High resolution |
| High sensitivity | Low sensitivity |
| Night vision | Day vision |
| Slow recovery after bleaching ¹ | Fast recovery after bleaching |
| Maximum sensitivity @ 500nm | Maximum sensitivity @ 550nm |

¹ *Bleaching*: excess light causes a dysfunction in the photopigment such that it is incapable of absorbing any light.

The distribution of rods and cones over the extent of the retina is shown in Figure 3.

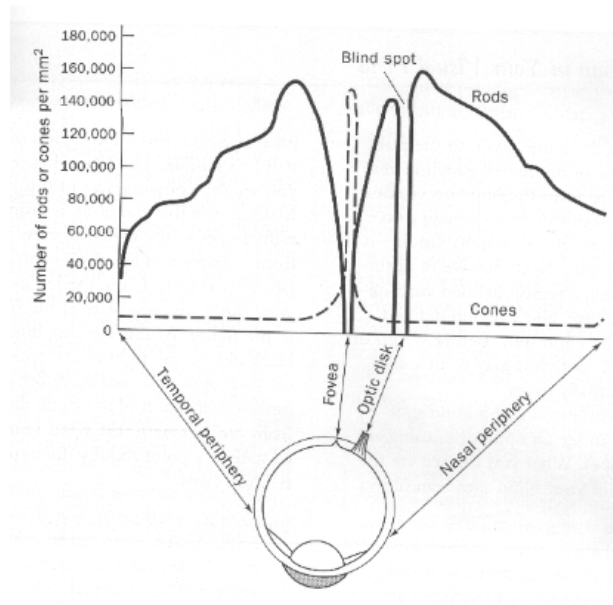


Figure 3

Observe that we find the maximum number of cones in the fovea. On the other hand, we find the maximum number of rods in the periphery. Note the complete absence of rods within the fovea.

1.2.2 Retinal Ganglion Cells - Receptive Fields

The human eye contains roughly 1 million retinal ganglion cells. Comparing this figure to the 128 million receptors in the eyes, you know from the outset that ganglion cells must be condensing the raw messages from the receptors. Therefore, the retinal ganglion cells must collate messages from the more numerous photoreceptors and summarize those messages in a biologically relevant way: this is what we call *convergence*.

If you were to record from a photoreceptor, you would find that it was "ON" (hyperpolarized, paradoxically) whenever light shone on it. If you recorded from a ganglion cell instead, you would find that diffuse light did little to the cell. However, the cell would respond well to a small spot of light, a small ring of light, or a light-dark edge. We say that this cell has a center-surround *receptive field* - the center must be mainly light and the surround mainly dark, or vice versa. What happens between the outer segment and the ganglion cell? This complex receptive field is created by the interneurons of the retina: the bipolar cells and the horizontal cells, primarily.

Let's trace a signal through:

1. Light hyperpolarizes the cone (or rod). For simplicity's sake, we will just say that turns ON the cone, and thereby excites the bipolar cell directly underneath.

That bipolar cell then excites its ganglion cell. The same thing is happening to neighbor cells.

2. However, here's the trick. The neighbor cones also excite horizontal cells. The horizontal cells send processes laterally and inhibit the center bipolar cell.

So, what does diffuse light do? It excites the central bipolar cell, but also inhibits it via the neighbors. Result - the ganglion cell does not get excited. It continues to tick along at its normal, tonic rate.

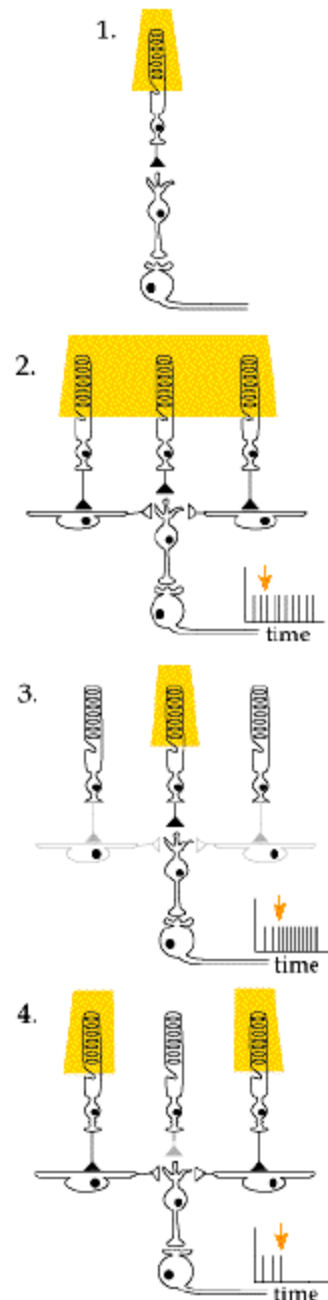
3. A small spot of light, however, excites the bipolar cell but not its neighbors. There is no inhibition, so it is free to get really excited and excite the ganglion cell, which fires like crazy.

4. A ring of light excites only the neighbors. Now, the bipolar cell is strongly inhibited, with no excitation. In response to this strong silencing of the bipolar cell, the ganglion cell shuts down as well. It will not turn on again until the light is turned off, at which time you will see a rebound "off-response".

This is an *ON-center cell*.

The reverse of this entire scenario can be created by reversing all the signals (which we can do with different receptors to the same neurotransmitter) - you then have an *OFF-center cell*.

This unique center-surround receptive field is also a property of lateral geniculate neurons (that will be described later). Things get even more complicated up in the cortex.



What is the point? Well, our entire visual system exists to see borders and contours. We see the world as a pattern of lines, even things as complex as a face. We judge colors and brightness by comparison, not by any absolute scale. This system of *lateral inhibition* in the retina is the first step towards sharpening contours and picking up on borders between light and dark. Diffuse light is ignored by the ganglion cell, but a sharp dot will really turn it on. Higher up in the cortex, all these dots will be combined into lines, which will be combined into curves, etc.

1.2.2.1 Receptive Field Layout

Nearly all ganglion cells have concentrically arranged receptive fields (illustrated in Figure 4), composed of a center and a surround that responds in an antagonistic fashion.

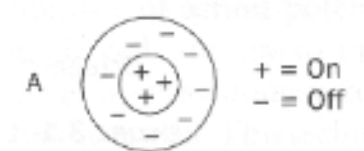


Figure 4

Some of those cells have ON centers and OFF surrounds, while other cells have just the opposite layout, with an OFF center and an ON surround. There are about as many ON-center cells as there are OFF-center cells; both types respond best to light/dark boundaries. The receptive fields on the ON-center cells and OFF-center cells overlap to make up a mosaic covering the entire retina. As a result, whenever light falls on a limited patch of the retina, that light is bound to affect a number of retinal ganglion cells of both types, producing opposite effects in the two. Don't imagine, though, that these opposite effects cancel one another; at higher stages of the visual system, information from ON-center and OFF-center cells remains segregated, allowing information from both types to be used.

1.2.2.2 Receptive Field Size

Cells with small receptive fields will respond best to small objects, while those with large receptive fields will prefer larger objects. This principle suggests that analysis of object size may be inaugurated in the retina.

2. Central Visual Pathways

2.1 The Pathway

Vision is generated by photoreceptors in the retina, a layer of cells at the back of the eye. The information leaves the eye by way of the optic nerve, and there is a partial crossing of axons at the optic chiasm (see Figure 5).

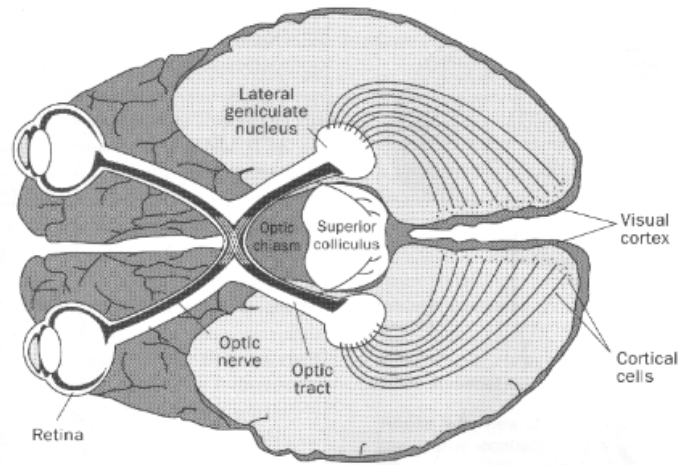


Figure 5

After the chiasm, the axons are called the optic tract. The optic tract wraps around the midbrain to get to the lateral geniculate nucleus (LGN), where all the axons must synapse. From there, the LGN axons fan out through the deep white matter of the brain as the optic radiations, which will ultimately travel to primary visual cortex, at the back of the brain (see Figure 6).

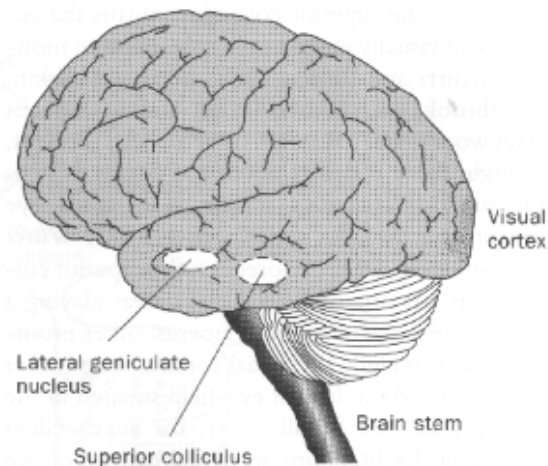
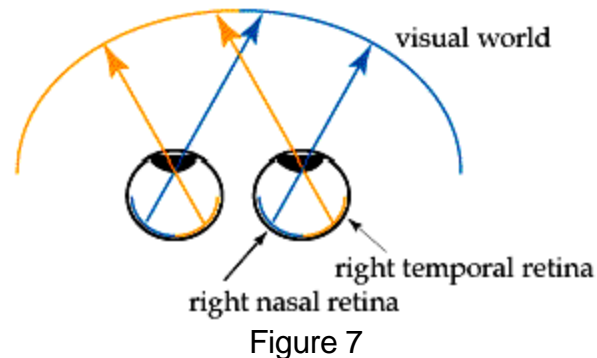


Figure 6

2.2 Visual Fields

Information about the world enters both eyes with a great deal of overlap. Try closing one eye, and you will find that your range of vision in the remaining eye is mainly limited by your nose. The image projected onto your retina can be cut down the middle, with the fovea defining the center (see Figure 7). Now you have essentially two halves of the retina, a left half and a right half. Generally, the halves are referred to as a temporal half (next to your temple) and a nasal half (next to your nose).



Visual images are inverted as they pass through the lens. Therefore, in your right eye, the nasal retina sees the right half of the world, while the temporal retina sees the left half of the world. Notice also that the right nasal retina and the left temporal retina see pretty much the same thing. If you drew a line through the world at your nose, they would see everything to the right of that line. That field of view is called the right hemifield.

So, what you see is divided into right and left hemifields. Each eye gets information from both hemifields. For every object that you can see, both eyes are actually seeing it - this is crucial for depth perception - but the image will be falling on one nasal retina and one temporal retina.

Why bother to divide the retinas at all? Recall that the brain works on a crossed wires system. The left half of the brain controls the right side of the body, and vice versa. Therefore the left half of the brain is only interested in visual input from the right side of the world. To insure that the brain doesn't get extraneous information, the fibers from the retina sort themselves out to separate right hemifield from left hemifield. Specifically, fibers from the nasal retinas cross over at the optic chiasm (these are called contralateral fibers), whereas the temporal retinas, already positioned to see the opposite side of the world, do not cross (these are called ipsilateral fibers). Figure 8 shows what it looks like.

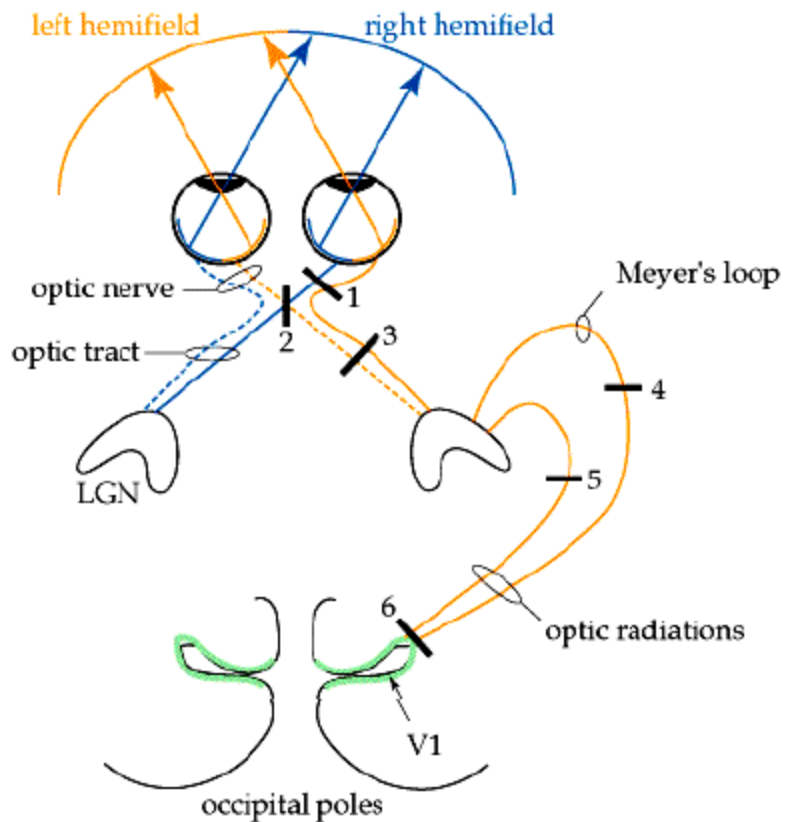


Figure 8

The practical consequences of this crossing are that damaging the visual system before the chiasm will affect one eye, both hemifields - analogous to closing one eye. Damaging the pathway after the chiasm, though, will damage parts of both eyes, and only one hemifield. There is no easy way to imagine what this would look like. Your field of view would be only 90°, from straight ahead to one side.

2.3 After The Retina

Once the ganglion cell axons leave the retina, they travel through the optic nerve to the optic chiasm, a partial crossing of the axons. At the optic chiasm the left and right visual worlds are separated. After the chiasm, the fibers are called the optic tract. The optic tract wraps around the cerebral peduncles of the midbrain to get to the lateral geniculate nucleus (LGN). The LGN is really a part of the thalamus, and remember that nothing gets up to cortex without synapsing in thalamus first (if the cortex is the boss, the thalamus is an excellent secretary). Almost all of the optic tract axons (approximately 80%), therefore, synapse in the LGN. The remaining few (20%) branch off to synapse in the superior colliculus, a neighboring structure in the midbrain.

2.4 The Superior Colliculus

Tucked away at the top of the brain stem, the superior colliculus is a phylogenetically older, more primitive visual area than the visual cortex. In many lower animals, such as frogs and fish, the superior colliculus represents *the* major brain center for visual processing. In human beings and other higher animals, however, the phylogenetically newer visual cortex has supplanted the superior colliculus as the more important visual area. Nonetheless, even in these higher animals, the superior colliculus plays a prominent role in visual orienting reflexes, including the initiation and guidance of eye movements.

Cells in the superior colliculus have receptive fields with rather ill-defined ON and OFF regions. These cells will respond to just about any visual stimulus – edges, bars of any orientation, light flashes – falling within their receptive fields, regardless of the shape or orientation of that stimulus. These properties suggest that the colliculus is not concerned with precisely “what” the stimulus is but rather with “where” it is. For another thing, cells in the superior colliculus clearly are involved in controlling eye movements. Many cells in the superior colliculus emit a vigorous burst of activity just before the eyes begin to move. This burst of activity occurs, however, only if the eyes move in order to fixate a light flashed in the visual periphery; eye movements made in darkness evoke no activity in these cells. So the intention to move the eyes toward some objects seems to be critical for the cells’ responsiveness. Besides initiating eye movements, the superior colliculus also plays some role in guiding the direction and extent of those eye movements.

At present, though, there is debate about the details of how the superior colliculus actually does guide eye movements. Moreover, the superior colliculus is just one of several brain areas playing a role in guidance of eye movements; other prominent areas include the visual cortex and the frontal eye-field situated in the frontal lobe.

In summary, the superior colliculus seems designed to detect objects located away from the point of fixation and to guide orienting movements of the eyes and the head toward those objects. The colliculus does not, however, contain the machinery necessary for a detailed visual analysis of such objects. That job, instead, belongs to the other branch of the optic tract, the one projecting to the lateral geniculate nucleus and then on the visual cortex. It is to those sites that we turn next.

2.5 The Lateral Geniculate Nucleus

2.5.1 Structure of the LGN

The LGN has a very distinct, layered appearance, as you can see in panel A of Figure 9. The layers consist of cell bodies. In humans, the LGN on each side of the brain contains six layers stacked on top of one another and bent in the middle. For the purposes of discussion, let's number the successive layers 1, 2, 3, 4, 5, and 6, going from the bottom to the top of the structure. Note also in Figure 9 that cells in layers 1 and 2 are larger than those in layers 3 through 6. These two large-cell layers are termed the magnocellular layers, and the four small-cell layers are called the parvocellular layers.

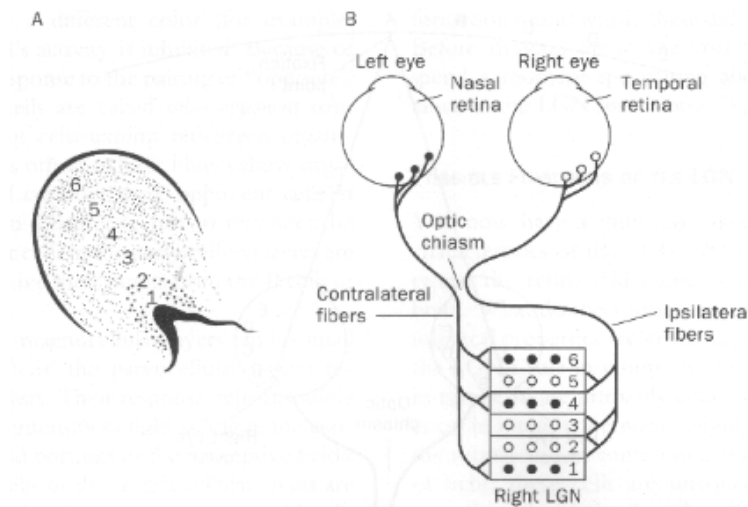


Figure 9

Let's consider some of the anatomical properties of this well-organized nucleus. To begin, consider how fibers of the optic tract make contact with cells in the different layers of the lateral geniculate nucleus. Remember that each optic tract contains a mixture of fibers from both eyes, the temporal retina of the ipsilateral eye and the nasal retina of the contralateral eye. These fibers become very strictly segregated when they arrive at the LGN: the contralateral fibers contact only the cells in layers 1, 4, and 6, whereas the ipsilateral fibers contact only the cells in layers 2, 3, and 5 (see Figure 9, panel B). Keep in mind that the brain contains a pair of LGN, one residing in the left hemisphere and the other in the right hemisphere. This paired arrangement means that a single eye provides the contralateral input to one LGN and the ipsilateral input to the other.

2.5.2 Retinal Maps in the LGN

Each layer of the LGN contains an orderly representation, or map, of the retina. This point is illustrated in Figure 10.

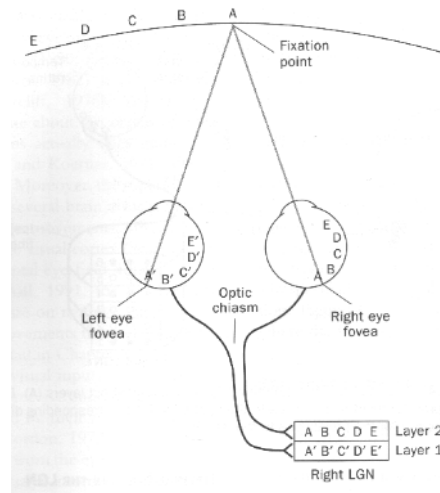


Figure 10

As you can see, axons that carry information from neighboring regions of the retina connect with geniculate cells that are themselves neighbors. This way of distributing retinal information within any layer of the LGN creates a map of the retina within that layer. Since such a map preserves the topography of the retina, it is known as a retinotopic map. Each of the six layers contains its own complete retinotopic map, and these layers are stacked in such a way that comparable regions of the separate maps are aligned with each other. For instance, the foveal parts of adjacent layers are situated on top of one another.

2.5.3 Receptive Field Properties of LGN Cells

Just like their retinal counterparts, LGN cells have circular receptive fields, subdivided into concentric center and surround components. As in the retina, these two components interact antagonistically. There is an important difference, though, between LGN and retina: the surround of an LGN receptive fields exerts a stronger inhibitory effect on its center than does the surround of a retinal ganglion cell. This means the LGN cells amplify, or accentuate, differences in illumination between retinal regions even more than do retinal ganglion cells. Incidentally, because their fields are circular in shape, LGN cells, like their retinal counterparts, will respond to borders or contours of any orientation.

2.5.4 Possible Functions of the LGN

Cells in the LGN are primarily concerned with differences in intensity between neighboring regions of the retina, and in some instances with the color of light;

these cells are unconcerned with the overall level of light. In other words, LGN cells are designed to register the presence of edges in the visual environment. This alone, however, does not represent a unique contribution to vision; retinal cells have already accomplished this job of signaling the presence of edges. So we are led to ask, does the LGN seem uniquely suited for some other role, or does it merely relay messages broadcast from the retina straight on to higher brain sites with no editing or censorship? There are several reasons for thinking the LGN is considerably more than just a relay station.

For one thing, the LGN receives input not only from the retina but from the reticular activating system as well. Buried in the brain stem, the reticular activating system governs an animal's general level of arousal. Hence the flow of visual information from the LGN to higher visual centers could be modified by the animal's arousal.

In addition to input from the reticular activating system, the LGN receives a large input from the visual cortex. You may find this intriguing, since the visual cortex itself receives a substantial amount of input from the LGN. In effect, the LGN formulates a message and transmits it to the cortex; the cortex, in turn, sends back a reply to the LGN, thereby allowing that message to be altered: this is what we call feedback. At present, the biological purpose of the visual feedback between the cortex and the LGN is not known.

Finally, the orderly arrangement of information within the LGN could help the next stage of visual processing to begin piecing the puzzle together. This next stage takes place in the primary visual cortex.

2.6 Visual Cortex

2.6.1 Generalities

The visual cortex is located at the very back of the cerebral hemispheres (see Figure 6), in a region called the occipital lobe. The visual cortex receives input from the LGN, but it is anatomically more complex than the LGN. This anatomical complexity is paralleled by an increase in physiological complexity – particularly in the variety of receptive fields exhibited by cortical cells.

The visual field (and, by extension, the retina) is represented within the cortex in a very orderly, topographic fashion. The center of the visual field maps onto the posterior region of the occipital lobe, and the periphery of the field maps onto the anterior portion. The upper portion of the visual field is represented in the lower portion of the occipital lobe, and vice versa.

The amount of cortical tissue devoted to the central portion of the field far exceeds the amount devoted to the periphery. This distortion, in which

representation of the center of the field is highly exaggerated, has come to be known as cortical magnification means that a large portion of the cortex is devoted to a very small area of the retina. That area of the retina is the fovea, on which the central portion of the visual field is imaged.

2.6.2 Structure of the Visual Cortex

The primary visual cortex is that portion of the occipital lobe receiving input from the lateral geniculate nucleus. It is sometimes referred to as Area 17. Other names for this area include V1 and striate cortex but they all refer to the same portion of the brain.

Like the rest of the cerebral cortex, the visual area consists of a layered array of cells about 2 millimeters thick. In all, there are approximately 100 million cells in the Area 17 of each hemisphere. Figure 11 is a magnified picture of a section of brain tissue from Area 17. Note that various layers differ in thickness and that the concentration of cells varies from layer to layer. The million or so axons arriving from the six layers of the LGN connect with cortical cells within layer 4 of the visual cortex. From here, connections within the cortex carry information to neurons in other layers above and below layer 4, eventually, to other areas within the brain.

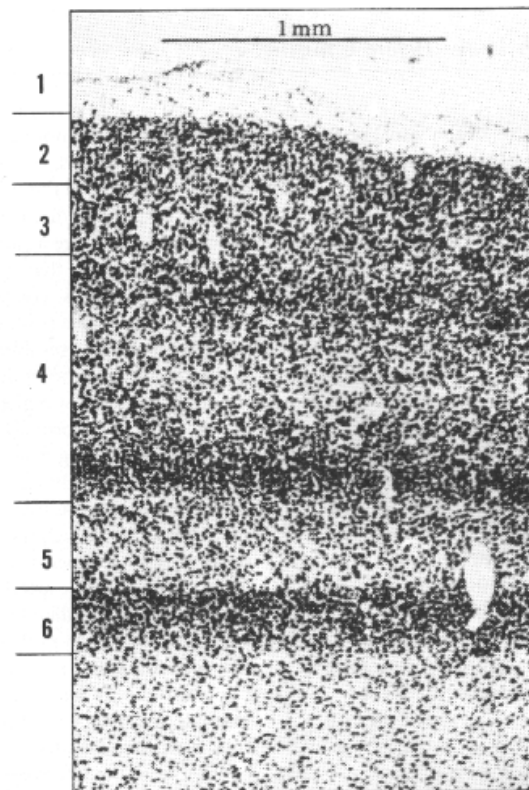


Figure 11

2.6.3 Retinal Maps in the Cortex

The first notable property of cortical cells is that each cell responds to stimulation of a restricted area of the retina – this area constitutes that cell’s receptive field. Considered as a group, the receptive fields of cells in each hemisphere of Area 17 form a topographic map of the contralateral visual field. This cortical map, like its counterpart in the LGN, is said to be topographic because adjacent regions of visual space are mapped onto adjacent groups of cells in the visual cortex. The map in each hemisphere is said to represent the contralateral visual field because the left half of the visual field is represented in the right hemisphere, and vice versa.

2.6.4 Functional Properties of Cortical Cells

One of the most striking characteristics of cortical cells is their orientation selectivity. Recall that retinal and geniculate cells, because of their circular-shaped receptive fields, respond indiscriminately to all orientations. In marked contrast, most cortical cells are very fussy about the orientation of a stimulus. As shown in the left hand portion of Figure 12, any particular cell will respond only if the orientation of an edge or line falls somewhere within a rather narrow range. Each cortical cell has a “preferred” orientation, one to which it is maximally responsive. If a line tilts a little away from this optimum, the cell’s response decreases markedly; if the line tilts even more, the cell no longer responds at all. This behavior can be summarized by graphing a cell’s response strength as a function of orientation; the resulting “tuning curve” resembles a tepee with its peak defining the cell’s preferred orientation (see the right-hand portion of Figure 12).

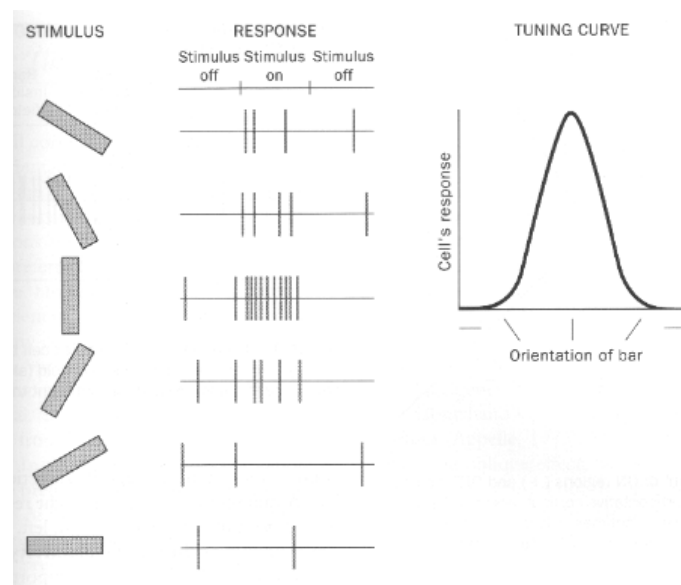


Figure 12

The preferred orientation varies from cell to cell, such that within an ensemble of cells all orientations are represented. Usually the most effective orientation for a cell can be sharply defined. In many cells, a tilt of no more than 15 degrees away from the optimum orientation is sufficient to abolish that cell's response completely. Just about all cortical cells are selective for orientation, with two notable exceptions. These unusual, nonoriented cells are found in the lower subdivision of cortical layer 4 and in regularly spaced patches of cells found near the cortical surface. Outside these restricted areas, in all other regions of visual cortex, orientation matters very greatly.

Let's take a look at the layout of ON regions and OFF regions constituting the receptive fields of some cortical cells. Some representative cortical receptive fields are shown in Figure 13.

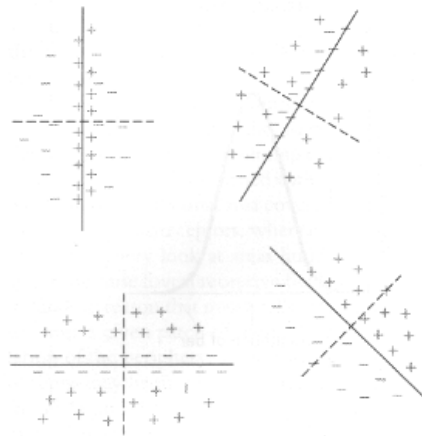


Figure 13

There you can see that receptive fields are elongated, not circular as in the retinal and geniculate cells. Notice, too, that the discrete zones of ON and OFF activity are arranged differently depending on the particular cell. Simply by inspecting these various arrangements, you should be able to picture what kind of oriented stimulus would evoke the best response from a given cell. For this reason, cortical cells within these well-delineated zones are called simple cells: there is a simple relation between their receptive field layout and their preferred stimulus.

Other cortical cells do not have such well-defined ON and OFF zones, yet they also exhibit a preference for a particular orientation. These are called complex cells, because it is more complicated to predict just what stimulus will optimally activate them. Unlike simple cells, complex ones will respond to an appropriately oriented stimulus falling anywhere within the boundaries of the receptive field – precise location is not really so important for a complex cell, so long as the stimulus falls somewhere within its receptive field and remains properly oriented.

Incidentally, because a complex cell lacks clearly defined ON and OFF regions, it will respond quite vigorously to a bar or an edge that moves through the receptive field. In fact, most complex cells strongly prefer moving contours to stationary ones, even if the contour travels rapidly across the receptive field. Simple cells, on the other hand, respond only to stationary or slowly moving contours.

Besides preferring contours that move, complex cells often respond to one direction of motion only. For instance, such a cell (like the one illustrated in Figure 14) might give a burst of activity when a vertical contour moved from the left to right, but would remain unresponsive when that same contour moved in the opposite direction, right to left. A different complex cell might respond only to a contour moving from the right to left. This property – known as direction selectivity – suggests that this class of cells plays an important role in motion perception.

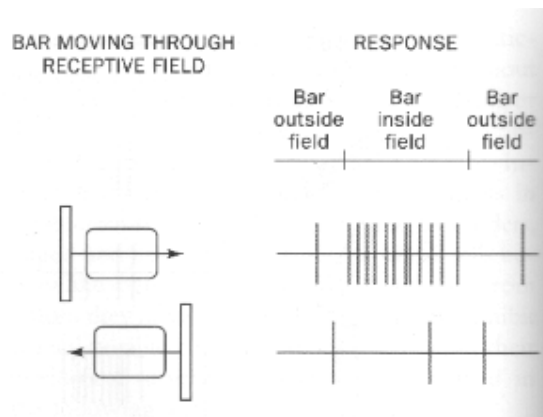


Figure 14

Recall that information from the two eyes is distributed in separate layers of each lateral geniculate nucleus, with individual cells receiving input from either one eye or the other. Once information passes from the geniculate to visual cortex, however, monocular segregation gives way to binocular integration. At the cortical level, individual cells, with few exceptions, are innervated from both eyes. A cell responds moderately well to a line presented to either eye alone, but its response is even stronger when both eyes are stimulated simultaneously. Some cells respond more vigorously to stimulation of the left eye, whereas other cells favor right-eye stimulation. This variation in the relative strength of the connection with the two eyes is called ocular dominance. Any cell that can be excited through both eyes, regardless of its ocular dominance, is called a binocular cell.

2.6.5 Columns and Hypercolumns

Having summarized the major response properties of cortical cells, let's look at how two of these properties – orientation and binocularity – are arranged within the cortex. Earlier we mentioned that cortical cells vary in their preferred orientations, some responding best to vertical, others to horizontal, and still others to orientations in between. These orientation-selective cells are not haphazardly arranged; they are grouped in a very orderly fashion within the cortex. If we penetrate perpendicularly all six cortical layers, we will discover that all cells along the path of this penetration have identical preferred orientations (except for cells in layer 4, which respond to all orientations). Along the way, there may be variation in the size of these cells' receptive fields, and some are likely to prefer edges while others prefer bars. But all cells will exhibit the same orientation. If we repeat the same procedure, this time beginning at a position just a fraction of a millimeter (0.05 millimeter) away from our initial penetration, we will find that all cells prefer an orientation slightly different (about 10 to 15 degrees) from the one encountered in the first sample. If we repeat this procedure over and over, we will uncover a regular sequence of preferred orientations. As we make these repeated samplings, another intriguing emerges: within a given penetration perpendicular to the cortical surface, all cells have the same ocular dominance. If the first cell encountered responds strongest to stimulation of the right eye, all cells in that sequence will be right-eye dominant (with the exception of the monocular cells in layer 4). Moreover, if we now make another sampling penetration right next to the first one, the cells encountered will exhibit a different pattern of ocular dominance, perhaps responding equally well to stimulation of either eye. This is illustrated in Figure 15.

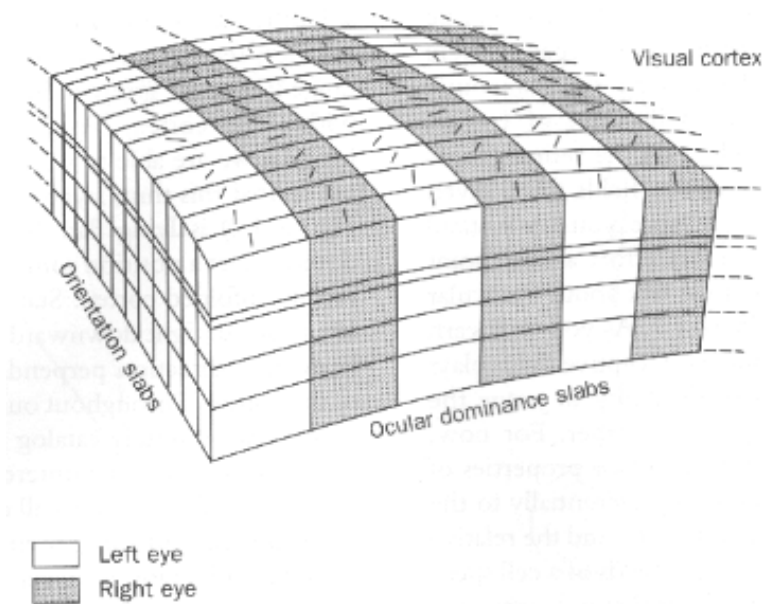


Figure 15

At first, all these facts may seem bewildering, but taken together, they produce an interesting and coherent picture. The visual cortex appears to be composed of columns of cells, with each column consisting of a stack of cells all preferring the same orientation and exhibiting the same ocular dominance. Altogether, it takes roughly eighteen to twenty neighboring columns to cover a complete range of stimulus orientations and ocular dominance. This aggregation of adjacent columns is collectively known as a hypercolumn.