Video: Human Vision: Primary Visual Cortex

[00:01] [slide 1] In this video, we'll explore the first stages of processing in visual cortex in the brain, in an area that's often referred to as primary visual cortex, or area V1, and we'll connect this processing to what you learned about edge detection in computer vision systems. This diagram shows the visual pathways from the eyes to visual cortex, which are replicated on the two sides of the brain. Information from the right side of our visual field, shown in green, passes through the lenses to the retina in both eyes, on the left side, and then passes up to the brain on the left side of the head. Information from the left side of the visual field is projected onto the retina on the right side and passes up to the right side of visual cortex. So each side of the brain analyzes half the visual scene, combining information from both eyes - this is a point that we'll come back to when we talk about stereo vision.

[01:12] [slide 2] Any introduction to the early stages of processing in visual cortex begins with the seminal work of David Hubel and Torsten Wiesel who shared the 1981 Nobel Prize in Physiology or Medicine. [slide 3] They recorded neural signals in the cortex of cats and monkeys, and when they first began these experiments in cats, they'd flash spots of light on a screen like we described in the study of retinal ganglion cells, in search of the receptive fields of cortical cells, as they were recording from these cells with a microelectrode. But they weren't able to get much response out of the cells with small spots of light. At the time, they were using a slide projector to flash spots of light on the screen, and after one very long and frustrating day of not being able to get much rise out of these cells, they were pulling a slide out of the projector and heard a vigorous burst of activity from a cell. Probing the cell further, they discovered that it was the dark edge of the slide moving across the screen as they moved it out of the projector that was eliciting this strong response. That discovery led to the systematic study of properties of cortical cells for which they received the Nobel Prize. In this work, they identified three basic cell types in this area, that they called simple, complex, and hypercomplex cells. We'll focus on the simple and complex cells.

[02:49] [slide 4] Some simple cells respond best to bars or lines of a particular sign of contrast, orientation in the image, and position in their receptive field. One cell, for example, might respond to a bright vertical bar on a dark background that's roughly centered in their receptive field, while another cell might respond best to a dark line on a bright background. [slide 5] Other simple cells respond best to edges of a particular sign of contrast, orientation, and position. One cell might respond best to an edge that's dark on the left and bright on the right, or vice versa. Hubel and Wiesel characterized the role of these cells as detecting features like edges and lines in the visual image. Cells have receptive fields of different size, both within a particular region of the visual field, and as we saw in the case of retinal ganglion cells, receptive field size increases as you move away from the center of the eye.

[03:54] [slide 6] Neurons in this area each have a preferred orientation that elicits the strongest response from the neuron, but they give some response within a range of orientations. Suppose we have a neuron that responds to a dark bar on a bright background, at a particular orientation shown on the left here, and imagine that the extent of the receptive field is represented by the dashed rectangle in each picture on the left. To the right of each picture is a trace of neural responses as a line at different orientations is moved back and forth through the receptive field. The orientation in the middle elicits the strongest response, and the response weakens as the line is rotated away from its preferred orientation. On the right is a plot of the firing rate of a sample neuron as a function of orientation, with the angle 0 representing the preferred orientation for the neuron that gives the maximum response. The neuron still has some response to bars that are 20 deg away from the preferred orientation, but by the time you reach about 40 deg away, there's no response. This kind of plot is referred to as a tuning curve, and it conveys how sensitive the neuron is to the orientation of the bar. A population of neurons, each with a different preferred orientation, can together detect bars or edges at all orientations in the image.

[05:33] [slide 7] I'd also like to touch on the spatial structure of simple cell receptive fields in a bit more detail. Earlier I showed simple diagrams of the general pattern of light that some cells prefer - the diagram on the far left depicts a neuron that prefers a vertical bright bar on a dark background, and the diagram on the far right depicts a neuron that prefers a vertical edge that's darker on the right. Similar to what we talked about with retinal ganglion cells, the response of a cell to light varies throughout its receptive field, and I'm displaying the shape of that response as an image here. Imagine that light in the brightest parts of these images excites the neuron the most, and darker regions are areas where light inhibits the neuron the most. The response pattern along a horizontal cross-section through the middle of each receptive field is shaped like the blue curves at the top, and a cross-section in the vertical direction is shaped like a Gaussian in both cases, as shown by the green curve in the middle. This pattern in two dimensions is modeled well by a type of function called a Gabor function that has a few parameters that control the shape of the function in 2D. One of the parameters is its orientation, and along the bottom, I show Gabor functions that approximate receptive fields for neurons with four different preferred orientations. I'm not going to get into the mathematical details of the Gabor function, but I wanted to mention it, as you may see references to Gabor functions in reading that you do later in the course.

[07:24] [slide 8] What about the cells that Hubel and Wiesel called complex cells? Their behavior is similar to simple cells in some ways - they're also selective for bars or edges at preferred orientations. But they generally have larger receptive fields and they're more tolerant to position changes. This will be important later in the course when we talk about visual recognition, but here, let me just clarify what this means. On the left is a typical response of a simple cell. It likes bright bars at this oblique orientation. The cell is very sensitive to the position of the bar in its receptive field. It responds most strongly to a bar right in the center of the receptive field here, but much less to bars that are further to the left or the right edge of their receptive field. In contrast, a complex cell that also likes bright bars at this orientation, will respond vigorously to the bar at any position in its receptive field. We know that complex cells in a particular region of the visual field have larger receptive fields than the simple cells in that region, and Hubel and Wiesel proposed that simple cells responding to bars at particular locations, might provide input to the larger complex cells, so that, for example, if any of these three simple cells respond to a bar at a particular location, then they'll be firing away and their inputs could then lead to firing by the complex cell for any position in the receptive field.

[09:18] Another important property of complex cells is that some are also sensitive to the depth of a bar or edge, its 3D distance from the eye. Remember that the visual cortex on each side of the brain combines information from the left and right eyes and this information can be used to infer the depth of edges in the scene, using stereo vision. Some complex cells are also sensitive to the direction of motion of an edge or bar. We'll come back to both of these observations in the next couple weeks as we talk about stereo and motion.

[09:55] [slide 9] So in the first stage of processing in visual cortex, we find neurons that appear to detect features in the image like edges and lines. The inputs to these cells come from the LGN and represent the convolution of the image with circularly symmetric difference-of-gaussians operators. Where we have an edge or line of a particular orientation in the image, there'll be lines of active on-center and off-center cells. Hubel and Wiesel also proposed that simple and complex cells use this pattern of activity to infer the presence of edges or bars in the image.

[10:36] [slide 10] On a final note, sometimes we perceive edges or bars in the image in places where they don't actually exist - we call these illusory, or subjective contours, and you'll explore one of these illusions in the first assignment, the sun illusion in the bottom left picture here, where you perceive a bright white disk in the center that isn't really there in the image. There are neurons in the next stage of cortical processing, in an area referred to as area V2, that detect real edges and bars in the image, but they're also active in places such as those circled in green, where we perceive illusory contours where there's no intensity change in the image. This phenomenon has also been the focus of many computational studies that attempt to construct these contours and understand why we have this ability. [slide 11] It may, for example, just be an extension of our natural ability to perceive complete objects such as the moon, the horse, the dalmatian, the penguins, and Pendleton Hall, in situations where objects are partially occluded or they blend into the background, like the sides of the head of the penguin. So there may be good reasons why we have this ability that also underlies these visual illusions of subjective or illusory contours.