

5 Adventures into *Terra Incognita*

Probing the Neural Circuits along the Ventral Visual Stream

Supplementary content at <http://bit.ly/2TpAg3w>

Around the 1950s, a wealth of behavioral experiments had characterized many phenomenological aspects of visual perception that begged for a mechanistic explanation (Chapter 3). Lesion studies had provided a compelling case that damage to circumscribed brain regions led to specific visual processing deficits (Chapter 4). These lesion studies pointed to specific brain areas to investigate visual processing, especially the primary visual cortex in the back of the brain. In addition, the successful use of microelectrode electrical recordings had led to direct insights about the function of neurons within the retinal circuitry (Chapter 2). The time was ripe to open the black box of the brain and begin to think about how vision emerges from the spiking activity of neurons in the cortex.

Retinal ganglion cells project to the lateral geniculate nucleus (LGN) in the thalamus, and the principal output projection from the LGN conveys visual information to primary visual cortex (V1; see Section 2.10), the first stage of cortical processing for visual information. From V1, information is propagated into a large number of visual cortical areas that are responsible for transforming a pixel-like representation of sensory information into rich and complex visual percepts (Chapter 1, Figure 1.5). The exploration and computational modeling of the visual cortex is an ongoing adventure, where courageous conquistadors dare to peek inside the inner workings of the most complex system ever examined by science. Fundamental structural and functional principles of computation are beginning to emerge out of the sometimes seemingly enigmatic *terra incognita* of visual cortex. These basic principles are introduced in this chapter and the next one and form the basis of the computational models of vision discussed in Chapters 7–9.

5.1 About the Neocortex

The neocortex is the outer structure of the neural tissue in the brain and is thought to be responsible for cognition. The prefix “neo” stands for new, which should be understood in evolutionary time scales and contrasts with the older paleocortex, which includes the olfactory system and the hippocampus. The human neocortex is about 2–4 millimeters thick, comprises about 40 percent of the brain mass, and contains on the order of 10^{10} neurons. The cortex shows a large number of folds such that it can fit about 2,600 square

centimeters, into the size of the brain. Because of its extensive surface and relatively shallow depth, many investigators think of neocortex as a quasi-two-dimensional structure. The most prominent fold is the longitudinal fissure separating the right and left hemispheres. The human neocortex has more folds than that of many other mammals; for example, the mouse cortex appears relatively smooth in comparison to the human cortex. Mechanical tension, combined with the strong constraint to save wiring and space, is likely to have been an important factor in determining the shape and folds of cortex throughout evolution.

To a pretty reasonable first-order approximation, *cortex is cortex*: staining of cortical tissue appears at a gross level to be very similar across different parts of the brain. Furthermore, cortical staining also appears quite similar across different species. It takes a connoisseur to distinguish a section of mouse cortical tissue from human cortical tissue. This similarity is perhaps remarkable to some people. Egocentric or anthropomorphic considerations might lead some people to think that the human cortex might be substantially different; after all, mice do not play chess, nor do they read Shakespeare. The coarse similarities in the basic cortical structure suggest that approximately the same pieces of hardware can be combined in different and exciting ways to account for the cognitive capacities of different species. As a rough analogy, similar transistors can be used to build an electronic calculator, a smartphone, and a laptop.

Upon further scrutiny, specialists can distinguish between different species by examining cortical tissue. Furthermore, it is also possible to demarcate different brain regions by examining cortex. The German neuroanatomist Korbinian Brodmann (1868–1918) devised a parcellation of the human and monkey brains – as well as many other species – based on morphological cytoarchitectonic criteria. Many parts of the neocortex are still referred to by their Brodmann area number (Figure 5.1). For example, the primary visual cortex corresponds to Brodmann area 17. Neurophysiological and lesion studies have shown that several of the structural subdivisions proposed by Brodmann, as well as subsequent neuroanatomical work, correlate with functional specialization. Attempts to




Figure 5.1 The cortex can be subdivided into multiple brain areas based on cytoarchitectonic criteria. Brodmann subdivided the neocortex into multiple areas based on cytoarchitectonic criteria. The primary visual cortex corresponds to Brodmann area 17 in this diagram. From Wikipedia

separate cortical regions, particularly combined with attempts to attach cognitive functions to different regions, have a long and rich history that continues to current days.

5.2 Connectivity to and from the Primary Visual Cortex

The primary visual cortex is the first stage where information from the two eyes converges onto individual neurons. Each hemisphere in V1 represents the contralateral visual field. The part of the retina that is closer to the nose is called *nasal*, while the other half of the retina is called *temporal*. The left visual hemifield (left of the center of gaze) is represented by the nasal part of the retina on the left eye and by the temporal part of the retina on the right eye. Information from the nasal retina on the left eye will cross the brain and end up represented in the right hemisphere in the primary visual cortex. Information from the temporal retina on the right eye will turn at the *optic chiasm* and also end up represented in the right hemisphere in the primary visual cortex.

Like most other aspects of neuroanatomy, the first drawings of the primary visual cortex were made by Santiago Ramón y Cajal, who was introduced in Chapter 2. The basic architecture of the primary visual cortex turned out to be approximately similar to that of other parts of the visual neocortex. The neocortical sheet is characterized by six layers that can be distinguished with Nissl staining, a technique used to sparsely introduce a dye into many neurons in a given area for visualization. Sparse staining is important here because the density of neurons in the cortex is so large that it would be hard to see much upon staining all neurons using standard microscopy. The six layers are characterized by a stereotypical connectivity pattern that is often referred to as the *canonical cortical microcircuit*. With some exceptions – it is biology after all – this canonical connectivity pattern is shared across different visual areas and also across different sensory modalities.

Connections among different areas of cortex are often described as “bottom-up,” “top-down,” or “horizontal” connections, a nomenclature that is also used to describe connectivity in artificial neural network architectures (Section 7.3, Figure 7.4). A given individual neuron will only project in a bottom-up manner, or horizontally, or provide top-down signals, but not all of these. These different types of connections are defined based on the specific layer of the pre- and postsynaptic neurons. The connections between and within visual cortical areas follow a stereotypical pattern that has been used to define what area is “upstream” or “downstream” and, therefore, which connections are bottom-up or top-down (Figure 5.2). Bottom-up connections arrive at layer 4 – the LGN projects to pyramidal neurons in layer 4 in primary visual cortex. Layer 1 is the most superficial and contains mostly dendrites and few neuronal cell bodies; the neuron cell bodies for those dendritic arbors are mostly located in layers 2 and 3. Top-down connections from other visual cortical areas typically end in the deep layers 5 and 6 and also, to a lesser degree, in layers 2 and 3. After the LGN input (or input from a “lower area”) arrives onto layer 4, information flows from layer 4 to layers 2/3 and then on to layer 5 and layer 6. Information from layer 6 provides back projections to the LGN (or to a “lower” visual area) and is also fed back to layer 4.

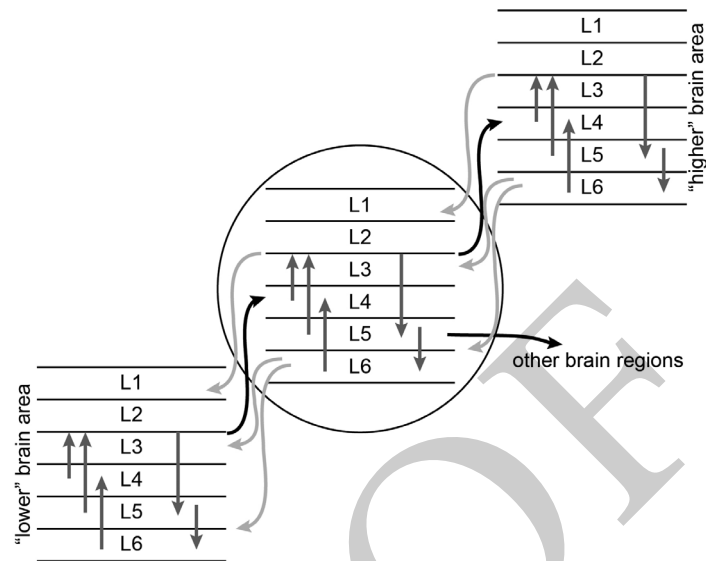


Figure 5.2 Canonical cortical circuits. Cortical connectivity across the visual cortex follows stereotypical connectivity patterns illustrated here. L1 through L6 refer to the six cortical layers. “Bottom-up” connectivity between areas is shown in black, “Top-down” connectivity between areas is shown in light gray, and connections within an area are shown in medium gray.

An important aspect of connectivity in the visual cortex is that connections between areas are almost invariably reciprocal. If area A provides bottom-up input into area B, area B provides top-down inputs to area A. Furthermore, these reciprocal connections are quantitatively comparable: the number of projections from A to B and from B to A are approximately similar.

By scrutinizing the connectivity patterns across layers in multiple brain areas, investigators have come up with an approximate map of the anatomical paths through which different visual areas communicate with each other (Chapter 1, Figure 1.5). Based on the separation of connections into bottom-up and top-down, it is possible to arrange the multiple different visual brain areas into an approximately hierarchical structure. The diagram in Figure 1.5 provides a semi-hierarchical description of the anatomical flow of information in the visual system.

The more we study connectivity in visual cortex, the more we realize that this stereotypical pattern is full of exceptions. There are differences across species, differences between visual cortex and motor cortex, and even differences between different visual cortical areas. To make matters more complicated, these layers can, in turn, be subdivided into sublayer structures, and the connectivity patterns may be different depending on the types of neurons being considered. For example, we started this section by stating that the primary visual cortex is *approximately* similar to other visual cortical areas. Perhaps because of its unique position in receiving more direct thalamic inputs than all other visual areas, V1 is actually thicker, layer 4 has different numbers of sublayers, and the pattern of inputs and outputs is also distinct from other visual areas.

In addition to the variations in the canonical microcircuit across cortical areas and across species, the hierarchical nature of visual cortex should not be interpreted too strictly. For example, numerous “bypass” connections send information from area A to area C without going through the intermediate area B (e.g., information flows from V1 to V2 to V4, but there are also direct connections from V1 to V4). Despite the subdivisions, exceptions, and refinements, the basic principles of connectivity in the visual cortex have played an important role in imposing method to the apparent madness and have inspired the development of the best computational models that we have today (Chapters 7 and 8).

A word of caution about nomenclature is pertinent, particularly for computer scientists used to thinking about neural networks. Biologists talk about different cortical *areas* – such as V1, V2, and V4. Each of these areas has six *layers*, as described earlier. In Chapters 7 and 8, we will discuss computational models of visual processing, which often refer to computational steps instantiated in different “layers.” Those computational layers should not be confounded with the cortical layers described here. A layer in a neural network is not necessarily directly linked to one of the six layers in the neocortex in any given brain area. The exact mapping between computational layers and brain areas is not always well defined by modelers. In fact, in many cases, people think about a layer in a neural network as potentially equivalent to a whole brain area in the cortex. We will come back to the question of making a commitment in the mapping between computational models and biological anatomy. For the moment, here we refer to layers in the biological sense discussed in the previous paragraph and in Figure 5.2. In addition to information flowing from one layer to another layer within a visual area and information flowing between brain areas, there are extensive *horizontal* connections whereby information flows *within* a layer. Some investigators use the term *recurrent* connections to refer both to horizontal and top-down connections, but it is conceptually clearer to keep different terms for these two distinct types of signal paths.

5.3 The Gold Standard to Examine Neural Function

Every problem has an appropriate scale of study: a Goldilocks scale, so to speak – not too coarse, not too fine. For example, it is particularly tedious and challenging to attempt to read the newspaper using a microscope (too fine a resolution). It is also extremely challenging to read a newspaper from a distance of 200 meters away (too coarse). A plethora of methods are available to study the brain, ranging from elucidating the three-dimensional structure of specific types of ion channels, all the way to indirectly measuring signals that show some degree of correlation with blood flow, averaged over coarse spatial scales.

In the case of neocortical circuits, this Goldilocks scale is the activity of individual neurons. Studying the three-dimensional structure of each protein inside a neuron is equivalent to trying to read the newspaper with a microscope – but it can be extremely useful for other questions such as understanding the kinetics and properties of ion channels in the neuronal membrane. Studying the average amount of blood flowing

through half a cubic centimeter of cortex over several seconds is equivalent to attempting to extract ink tones from the newspaper from 200 meters away – but it can be useful for other questions such as differentiating general and coarse properties of a part of cortex.

In addition to this spatial scale, there is also a natural time scale to examine neuronal activity. Most neurons communicate with each other by sending electrical signals called action potentials lasting about two milliseconds. For most purposes, it is sufficient to study neuronal activity at the millisecond level. With a few exceptions (e.g., small differences in timing between signals arriving at the two ears), microsecond resolution time scales do not provide additional information. One day has 1,440 minutes, and, therefore, the analogy for studying brains at the microsecond instead of millisecond scale (a factor of 1,000) would be to reread the same newspaper every minute. At the other end of the spectrum, techniques that average activity over many seconds are too coarse to elucidate cortical computations. The analogy for studying brains at the scale of several seconds instead of milliseconds (a factor of 1,000) would be to average the newspaper over three years.

Studying the activity of neocortical circuits at the neuronal resolution at a scale of milliseconds is not trivial and requires inserting thin microelectrode probes into the areas of interest. Neuronal action potentials lead to changes in the electrical potential in the extracellular milieu. It is possible to amplify and measure this electrical potential in the extracellular space and measure the action potentials emitted by individual neurons. The methodology was established by Edgar Adrian (1889–1977), and we already introduced example measurements of single-neuron activity in the retina in Section 2.7.

5.4 Neurons in the Primary Visual Cortex Respond Selectively to Bars Shown at Specific Orientations

Human primary visual cortex consists of about 280 million neurons arranged in a 2-millimeter-thick sheet that encompasses a few square inches in surface. There are more papers examining the neurophysiology of the primary visual cortex than the rest of the visual cortex combined. Neurons in the primary visual cortex – as well as those in the retina and LGN (Section 2.7) and also neurons in other parts of visual cortex – show spatially restricted receptive fields; that is, they respond only to a specific part of the visual field (Figure 2.9). The ensemble of all the neurons tiles the entire visual field. On average, the receptive field size of neurons in the primary visual cortex is larger than the receptive field sizes in the retina and LGN, typically encompassing about 0.5 to 1 degree of visual angle. A typical neurophysiology experiment often starts by determining the receptive field location of the neuron under study. After determining the location of the receptive field, a battery of stimuli is used to probe the neuron's response preferences.

The initial and paradigm-shifting strides toward describing the neurophysiological responses in the primary visual cortex were introduced by Torsten Wiesel (b. 1924) and David Hubel (1926–2013). The history of visual neuroscience revolves around the history of visual stimuli. Before the Hubel-Wiesel era, investigators had examined the

responses in the primary visual cortex using diffused light or the type of point sources that had successfully elicited activity in the retina and LGN. By a combination of inspiration, perspiration, and careful observation, Hubel and Wiesel realized that neurons in the primary visual cortex responded most strongly when a bar of a particular orientation was presented within the neuron's receptive field. The story of how this discovery came about is particularly fascinating and is recounted in David Hubel's Nobel Lecture. Hubel and Wiesel did not have particularly grandiose hypotheses about the function of neurons in the visual cortex before they embarked on these investigations but rather intuited that compelling principles would emerge by courageously placing electrodes in V1. After a particularly long day recording the activity of a V1 neuron, they were frustrated by how little the neuron seemed to care about the presence of a light or dark annulus inside the receptive field. In those days, they did not have computers to present stimuli; instead, they used slides inserted into a projector. Suddenly, their careful power of observation led them to realize that the neuron would show a burst of activity every time they inserted the slide into the projector. It was the edge of the slide moving in and out of the projector that triggered activation, much more than the content of the slide. Excited by this finding, they went on to discover that the orientation of an edge placed within the receptive field mattered for the neuron: specific orientations led to much larger activation than others.

A typical pattern of responses obtained from V1 recordings is illustrated in Figure 5.3. In this experiment, an oriented bar was moved within the receptive field of the neuron under study. The direction of movement was perpendicular to the bar's orientation. Different orientations elicited drastically distinct numbers of action potentials. While the number of action potentials (or spike count) is not the only variable that can be used to define the neuronal response, it provides a simple and adequate starting point to examine neuronal preferences. When the bar was approximately at a -45° angle (Figure 5.3D), the neuron emitted more spikes than for any other orientation. Moreover, the activity of this neuron was also dependent on the direction of motion. When the bar was moving toward the upper right, the neuron was vigorously active, whereas there was minimal activation in the opposite direction of motion.

Hubel and Wiesel went on to characterize the properties of V1 neurons in terms of their topography, orientation preference, ocular preference, color, direction of motion, and even how those properties arise during development. Their Nobel Prize-winning discovery inspired generations of neurophysiologists to examine neuronal responses throughout the visual cortex.

5.5 Complex Neurons Show Tolerance to Position Changes

In the example shown in Figure 5.3, the V1 neuron responds preferentially to a moving bar. Neurons in V1 also respond to flashes of static stimuli. When flashing a stimulus, how precise does the position of the oriented bar within the neuron's receptive field have to be to trigger a response? A distinction has been observed between two types of neurons in V1 based on how picky they are with respect to stimulus position within the




Figure 5.3 Example responses of a neuron in the monkey primary visual cortex. Physiological responses of a neuron in the primary visual cortex to bars of different orientations. In these examples, the bar was moved in a direction perpendicular to its orientation. The dashed lines on the left indicate the receptive field, the black rectangle is the oriented bar, and the arrows indicate the direction of motion. The neuronal response traces are shown on the right. Reproduced from Hubel and Wiesel 1968

receptive field: *simple* and *complex* V1 neurons. Complex neurons are less sensitive to the exact position of the bar within the receptive field. When using gratings containing multiple oriented bars at a given spatial frequency, complex neurons tolerate larger changes in the spatial frequency than simple cells. Simple and complex neurons are often distinguished by the ratio of the “DC” maintained response to their “AC” response elicited by a moving grating. Complex neurons show a small AC/DC ratio (typically <10), whereas simple neurons have a larger AC/DC ratio (typically >10). In other words, complex neurons show a higher degree of *tolerance* to the exact position of an oriented bar within the receptive field compared to a simple neuron whose response magnitude decreases when the bar is shifted away from the preferred position (Figure 5.4). This progression from a simple neuron to a complex neuron showing increased tolerance has inspired the development of hierarchical computational models of object recognition that concatenate operations reminiscent of simple and complex cells as a way of keeping selectivity while achieving tolerance to transformations in the stimulus (Section 8.5).

Some complex neurons also show “end-stopping,” meaning that their optimum stimulus includes an end within the receptive field, as opposed to very long bars whose ending is outside of the receptive field. The end-stopping phenomenon can be understood as a form of contextual modulation where the patterns in the region surrounding the receptive field (in this case, whether the line continues or stops) influence the responses to the stimulus inside the receptive field. Such influences from

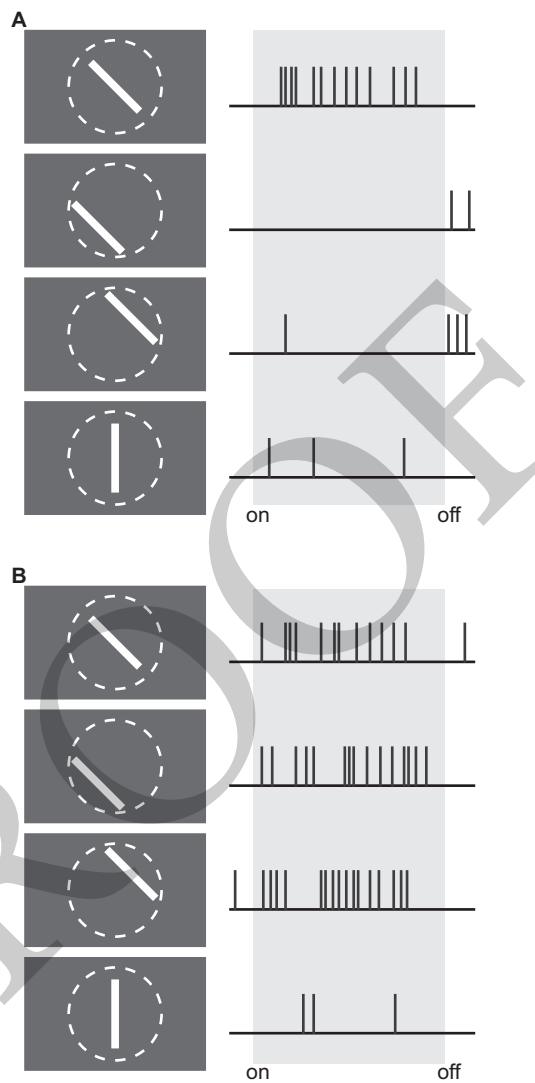


Figure 5.4 Complex neurons show tolerance to position changes. (A) Schematic diagram showing responses from a simple neuron that responds maximally to a -45 -degree oriented line when it is positioned in the center of the receptive field (top) but not when the position is shifted (rows 2, 3) or when the orientation changes (bottom). (B) Schematic diagram showing responses of a complex neuron that shows tolerance to position changes.

outside the receptive field are not restricted to end-stopping. V1 neurons also show surround suppression, similar to the suppressive effects of light around the receptive field center for on-center retinal ganglion cells described in Section 2.8 (Figure 2.10). In sum, V1 neurons are particularly sensitive to spatial changes, detecting edges indicative of a discontinuity in the visual field, and some neurons also detecting where the edge stops.

5.6 Nearby Neurons Show Similar Properties

Neurons in the primary visual cortex are topographically organized, in a similar fashion to the situation described in the retina in Section 2.7. The V1 topography is inherited from the LGN: the connections from the LGN to the primary visual cortex are topographically organized, meaning that nearby neurons in the LGN map onto nearby neurons in the primary visual cortex. V1 neurons cover the visual field, with a much higher density of neurons covering the foveal region. These neurophysiological observations are consistent with the types of scotomas observed in cases of localized V1 lesions (Section 4.4) and also with the locations of phosphenes reported upon stimulation in V1 (Section 4.9).

Hubel and Wiesel discovered another aspect of the topographical arrangement of neurons in V1 by comparing the tuning preferences of different neurons recorded during the same electrode penetration. In addition to sharing properties with their two-dimensional neighbors along the cortical sheet, neurons also share similar response patterns with their neighbors in the third dimension representing cortical depth. Advancing the electrode in a direction approximately tangential to the cortical surface, different neurons along a penetration share similar orientation tuning preferences. This observation led to the notion of a *columnar structure*: neurons within a column have similar preferences; neurons in adjacent columns show a continuous variation in their orientation tuning preferences.

Such topography may be critical for saving wires by virtue of arranging neurons with similar properties that need to be connected near each other. In particular, interneurons that have short dendrites may require having their targets nearby. However, if we keep the neuron-to-neuron connectivity intact, we could, in principle, rearrange the geometry of the neurons in arbitrary ways while keeping the computations intact. Topography may thus be mostly dissociated from function. Therefore, the smooth map of tuning properties within V1 is probably not a requirement for V1 computations. In fact, recent work has shown that this level of organization may not be a universal property. The primary visual cortex in mice does not have such a precise topographical mapping of orientation preferences; the geometrical arrangement of tuning preferences is described as “salt-and-pepper.”

Even if this topography is not strictly required for computational purposes, it may come in quite handy for investigators. For example, recording techniques with a reduced spatial resolution that average the activity of many neurons may depend strongly on topography (because average responses from completely randomly arranged neurons may yield nothing). For similar reasons, as discussed in Section 4.10, stimulation of many neurons via current injection may also be dependent on topography.

5.7 Quantitative Phenomenological Description of the Responses in the Primary Visual Cortex

Let $D(x,y)$ denote the responses of a neuron at position x, y . The receptive field structure $D(x,y)$ of orientation-tuned simple V1 neurons is often mathematically

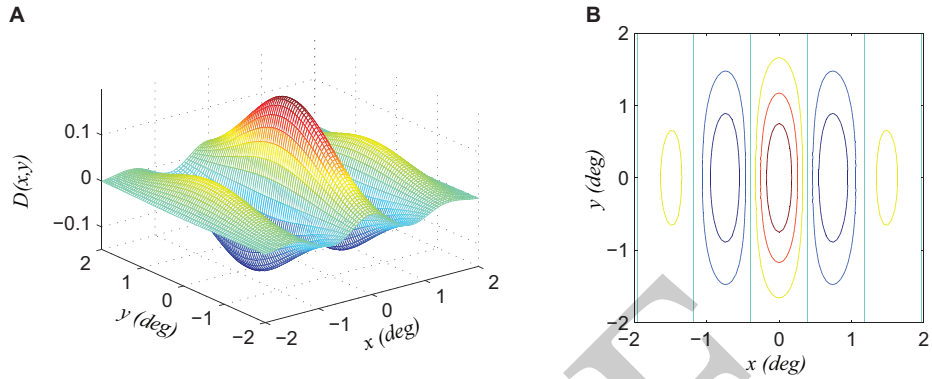


Figure 5.5 The spatial structure of receptive fields of V1 neurons is often described by a Gabor function (Equation (5.1)). (A) Illustration of a Gabor function. (B) Contour plot.

described by a *Gabor* function – that is, the product of an exponential and a cosine function:

$$D(x, y) = \frac{1}{2\pi\sigma_x\sigma_y} e^{-\left[\frac{x^2}{2\sigma_x^2} + \frac{y^2}{2\sigma_y^2}\right]} \cos(kx - \phi), \quad (5.1)$$

where σ_x and σ_y control the spatial spread of the receptive field, k controls the spatial frequency, and ϕ the phase. The Gabor function is characterized by an elongated excitatory region whose angle corresponds to the orientation preference of the V1 neuron, as well as a surrounding inhibitory region. An example illustration of a Gabor function is shown in Figure 5.5.

In addition to the spatial aspects of the receptive field, it is important to characterize the temporal dynamics of responses in V1. In most cases, the spatial and temporal aspects of the receptive fields in V1 can be considered to be approximately independent; that is, they can be separated without considering complex interactions between space and time. The temporal aspects of the receptive field can be fitted by the following phenomenological equation:

$$D(t) = \alpha e^{-\alpha t} [(\alpha t)^5 / 5! - (\alpha t)^7 / 7!] \quad (5.2)$$

for $\tau \geq 0$ and 0 otherwise. This equation is a fancy way of fitting the rapid and transient increase in firing rate upon flashing a stimulus at time 0 (Figure 5.6). The parameter α controls the latency and width of the temporal receptive field.

5.8 A Simple Model of Orientation Selectivity in the Primary Visual Cortex

Equation (5.1) provides a phenomenological description of the receptive field structure. In a remarkable feat of intuition, Hubel and Wiesel proposed a simple and elegant biophysically plausible model of how orientation tuning could arise from




Figure 5.6 The temporal structure of receptive fields of V1 neurons. Equation (5.2) is shown for different values of the parameter α .




Figure 5.7 Building orientation tuning by combining circular center-surround neurons. Schematic diagram showing how multiple LGN neurons with a circular center-surround receptive field structure can be combined to give rise to a V1 simple neuron that shows orientation tuning when those receptive field centers are adequately aligned. Modified from Hubel and Wiesel 1962

the responses of neurons with LGN-type receptive fields (Figure 5.7). In their model, multiple LGN neurons with circularly symmetric center-surround receptive fields (Figure 2.10) arranged along a line project onto a V1 simple neuron. Orientation tuning is thus constructed in a bottom-up fashion by combining the inputs of the right set of LGN neurons.

Subsequent work gave rise to a plethora of other possible models, and there is still an ongoing debate about the extent to which the purely bottom-up Hubel-Wiesel model represents the only mechanism giving rise to orientation selectivity in area V1. Still, this simple and elegant interpretation of the origin of V1 receptive fields constitutes a remarkable example of how experimentalists can provide reasonable and profound models that account for their data. Furthermore, the basic ideas behind this model have

been extended to explain the buildup of neuronal preferences for more complex shapes in other areas (Section 8.5).

In addition to orientation selectivity, there are many other properties of V1 neurons that are also arranged topographically including their spatial receptive fields, their ocular dominance (stronger responses to inputs coming from one or the other eye), their direction selectivity (stronger responses for specific directions of motion), and their retinal disparity (sensitivity to shifted positions between the right and left eyes used for stereopsis). It turns out that all of these other properties can also be mapped onto the specific arrangements of inputs from the LGN.

Extending their model for orientation selectivity in simple neurons by combining the output of LGN neurons (Figure 5.7), Hubel and Wiesel proposed that the responses of complex neurons could originate by the nonlinear combination of responses from multiple simple neurons with similar orientation preferences but slightly shifted receptive fields. These pioneering ideas of a linear filtering operation bringing about the responses of simple neurons in V1 followed up by a nonlinear pooling operation giving rise to complex neurons in V1 has played an influential role in inspiring computational models of visual processing (Section 8.5).

Figure 5.8 summarizes schematically how a V1 simple neuron would respond in a real-world image. This neuron has a receptive field in the upper right part of the visual field (black circle). Two fixations are shown in this figure. In the first fixation (**A1**), the image inside the receptive field is similar to the neuron's preferred orientation (**B**). After a nonlinear activation function (**C**), the neuron shows a strong response (**D1**). When the subject makes a small eye movement to the right, landing on fixation 2 (**A2**), the image inside the receptive field does not resemble the neuron's preferred features anymore, and the response is weak (**D2**).

5.9 Many Surprises Left in V1

Despite significant amounts of work investigating the neuronal properties in primary visual cortex, much remains to be explained. Multiple biases contribute to a partial view of V1 function. First, many of the recording procedures to date tend to focus on neurons that have higher firing rates and that are easier to pick up through extracellular recordings. Interneurons are smaller and harder to record from than the larger pyramidal cells. Additionally, there could be “shy” neurons that may be overlooked. Second, the types of stimuli that we use to probe neuronal responses also have biases (Section 5.11, Figure 5.10). Perhaps there are neurons in V1 that respond strongly to purple triangles with a sunflower on top, but, not surprisingly, nobody has tested this. Why would anyone test such a stimulus? None of our theories suggest that such a stimulus would be particularly relevant for V1 neurons. However, our theories may also be biased. Another important point to keep in mind is that neuronal responses in V1 are often probed in monkeys that are not performing any visual task other than fixating. Spatial context, temporal context, internal expectations, and task demands can modulate the responses of V1 neurons.

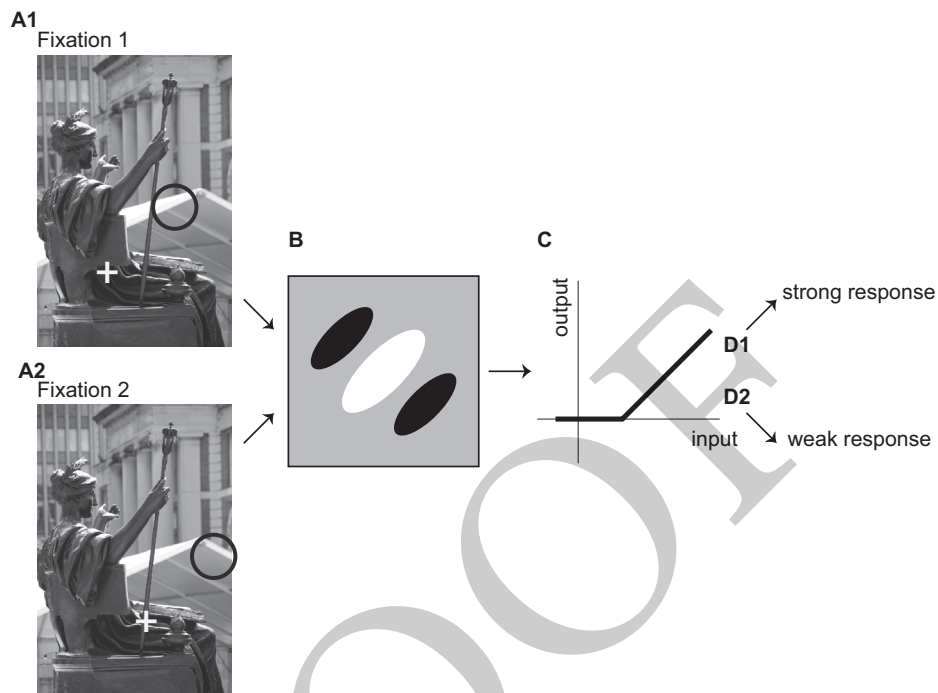


Figure 5.8 Beyond gratings and into the real world. Schematic example of how a V1 simple neuron might respond in the real world. “+” indicates the fixation location, and the black circle indicates the receptive field location. In **A1**, the image inside the receptive field is similar to the neuron’s preferred orientation (**B**), eliciting a high response (**D1**), whereas the reverse is true in the bottom case.

The last decade has seen an exciting increase in studies of mouse V1, many of which have opened our eyes to a world full of surprises, even at the heart of the most studied cortical area. Many of the experiments in mice are performed while the animal is running on a ball, a sort of treadmill exercise, while the mouse is watching a flashed stimulus or a movie. One of the most shocking findings in V1 in the last decade is that the running speed strongly modulates V1 neuronal responses. The same visual stimulus can trigger very distinct responses depending on whether the animal is still, trotting slowly, or sprinting. If this is not astounding enough, the responses of those V1 neurons can also be modulated by running in the dark, in the absence of any visual stimulation. Continuing with the list of intriguing observations in mice, there are direct connections from the primary auditory cortex onto V1, and it is possible to trigger responses in V1 neurons with auditory tones! These responses are weaker than visually triggered ones, but it is an auditory signal driving the most visual part of cortex. Whisker deflections can also modulate V1 neurons. And head movements too.

It remains unclear whether any of these observations extend to monkeys, let alone humans. It is not easy to do neurophysiological recordings in monkeys running around, and it is very challenging to perform neurophysiological recordings in human V1.

To the best of our knowledge, there is no report of auditory stimuli modulating V1 responses in monkeys (after controlling for eye position, attention, and visual stimulus). The rodent brain is much smaller than the macaque monkey brain (the *Mus musculus* and *Macaca mulatta* diverged about 75 million years ago), which is – in turn – smaller than the human brain (*Macaca mulatta* and *Homo sapiens* diverged about 25 million years ago). Introspectively, our visual world does not seem to change when we are walking or running around. However, there could be compensatory mechanisms that account for modulatory responses in V1 during running (remember that we are not even aware of the massive and pervasive visual changes caused by blinks and eye movements, Chapter 2). The auditory cortex, somatosensory cortex, and motor cortex are closer to V1 in mice than in monkeys, and there are more convolutions that could isolate brain areas in the macaque brain, and even more so in the human brain. Of note, this is all speculation, and we will need to evaluate all of these possibilities in neurophysiological recordings in monkeys and humans. We should keep our brains open and expect many exciting surprises ahead.

5.10 Divide and Conquer

Leaving the primary visual cortex and ascending through the hierarchy of cortical computations, we reach the fascinating and bewildering cortical areas that bridge low-level visual features into the building blocks of perception. In the primary visual cortex, there are neurons that respond selectively to lines of different orientations (Figure 5.3). At the other end of the visual hierarchy, there are neurons in the inferior temporal cortex (ITC) that respond selectively to complex shapes and help us identify chairs, faces, and planets (to be discussed in Section 6.2). In between V1 and the representation of complex object shapes, there is a vast expanse of cortex involved in the seemingly magical transformations that convert oriented lines into complex shapes. How do we go from oriented lines to recognizing chairs, faces, and planets (Figure 5.9)?

Despite heroic efforts by a talented cadre of investigators to scrutinize the responses between the primary visual cortex and the highest echelons of the ITC, the ventral visual cortex remains mostly *terra incognita*. Visual information flows along the ventral visual stream from V1 into areas V2, V4, posterior, and anterior parts of the ITC. The cortical real estate between V2 and the ITC constitutes a mysterious, seductive, and controversial ensemble of neurons whose functions remain unclear and are only beginning to be deciphered. Courageous investigators – armed with computational models, electrodes, and intuition – are beginning to describe the neuronal tuning preferences of neurons in areas V2, V3, and V4, in terms of features including curvature, disparity, color, texture, and shapes.

To solve the complex task of interpreting a scene, the visual system seems to have adopted a *divide-and-conquer* strategy. Instead of trying to come up with a single function that will transform lines into complex shapes in one step, the computations underlying visual cognition are implemented by a cascade of multiple approximately sequential computations. Each of these computations may be deceptively simple, and

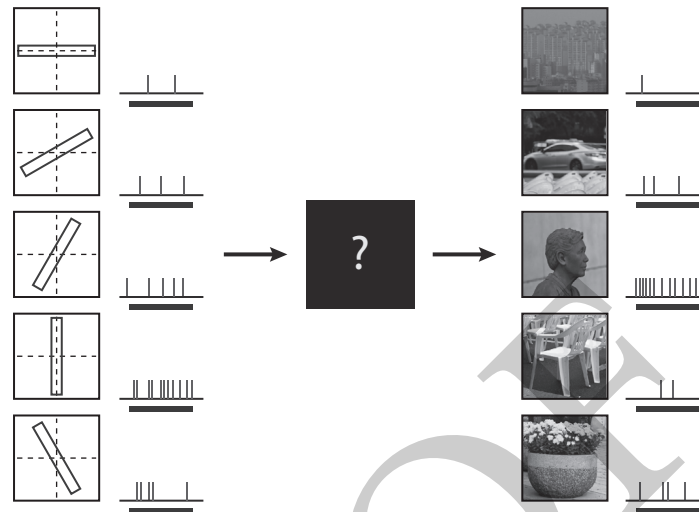


Figure 5.9 How does the cortex convert pixels to percepts? Through the cascade of computations along the ventral visual stream, the brain can convert preferences for simple stimulus properties such as orientation tuning into sophisticated features such as faces.

yet the concatenation of such steps can lead to interesting and complex emergent results. As a rough analogy, consider a factory making cars. There is a long sequence of specialized areas, departments, and tasks. One group of workers may be involved in receiving and ordering different parts; others may be specialized in assembling the carburetor; others in painting the exterior. The car is the result of all of these sequential and parallel steps. To understand the entire mechanistic process by which a car is made, we need to dig deeper into each of those specialized sub-steps without losing touch with the overall objective that each of these sub-steps contributes to – that is, the final product.

5.11 We Cannot Exhaustively Study All Possible Visual Stimuli

It would be nice to be able to describe the tuning preferences of neurons along the ventral visual stream in an analogous way to orientation tuning and Gabor functions for V1 neurons. There have been many empirical attempts to characterize the neuronal preferences of V2, V4, and ITC neurons, yielding exciting insights. As in the famous parable of blind men trying to describe an elephant by touching separate parts, different investigators have come up with several examples of how neurons respond to angles, colors, curvatures, and other shapes.

One of the main challenges to investigate the function and preferences of neurons in cortex is that there are too many possible images and we only have a limited amount of recording time for a given neuron. Given current techniques, it is simply impossible to exhaustively examine the large number of possible combinations of different




Figure 5.10 The curse of dimensionality in vision. With current techniques, we cannot exhaustively sample all possible stimuli. Here we consider a 5×5 grid of possible binary images (top) or possible grayscale images (bottom). Even for such simple stimuli, the number of possibilities is immense.

stimuli that might drive a neuron. Consider a simple scenario where we present image patches of size 5×5 pixels, where each pixel is either black or white (Figure 5.10, top). There are 2^{25} (more than 33 million) such stimuli. If we present each stimulus for 100 milliseconds and we do not allow for any intervening time in between stimuli, it would take more than five weeks to present all possible combinations. There are many more possibilities if we allow each pixel to have gray tones from 0 to 255 (Figure 5.10, bottom): 256^{25} such images (about 10^{60} such images!). Moreover, the problem becomes even worse if we allow three colors (red, green, and blue) and if we allow images larger than a mere 5×5 pixels. Even after restricting our analyses to the ill-defined subset of natural images (Section 2.1), we would still have an astronomically large number of possible images. We can typically hold extracellular recordings with single (non-chronic) electrodes for a couple of hours. Recent extraordinary efforts have managed to track the activity of a given neuron for up to a year. However, even with such chronic electrodes, it is challenging to keep an animal engaged in a visual presentation task for more than a few hours a day. Thus, we cannot record the responses of a neuron to all images.

Because of the severe limitations in the number of stimuli that can be tested, investigators often recur to several astute strategies to decide which stimuli to use in order to investigate the responses of cortical neurons. These strategies typically involve a combination of (i) inspiration from previous studies (past behavior of neurons in other studies is a good predictor of how neurons will behave in a new experiment); (ii) intuitions about what types of images might or might not matter for neurons (for example, many investigators have argued that real-world objects such as faces should be important); (iii) statistics of natural stimuli (as discussed in Section 2.1, it is reasonable to assume that neuronal tuning is sculpted by exposure to images in the natural world); (iv) computational models (to be discussed in more detail in Chapters 7–9); (v) serendipity (the role of rigorous scrutiny and systematic

observation combined with luck should not be underestimated). Combining these approaches, several investigators have probed the neural code for visual shapes along the ventral visual cortex.

5.12 We Live in the Visual Past: Response Latencies Increase along the Ventral Stream

Visual processing is very fast (Section 3.6). Indeed, as we argued in Chapter 1, the speed of vision is likely to have conferred critical advantages to the first species with eyes and may well constitute one of the key reasons why evolution led to the expansion of visual capabilities. However, even though the world seems to materialize in front of us upon opening our eyes, as we noted in Section 2.6, processing in the retina takes time. The intuition that vision is instantaneous is nothing more than an illusion. It takes about 30 to 50 milliseconds for signals to emerge from retinal ganglion cells into the thalamus, and it takes further time for signals to propagate through the cortex.

A small fraction of this time has to do with the speed of propagation along dendrites and axons within a neuron. However, within-neuron delays are relatively short. In particular, action potential signals within axons that are insulated by myelin can propagate at speeds of about 100 meters per second. Thus, signals from a single myelinated axon could, in principle, traverse the entire length of the human brain of approximately 15 centimeters in about 1.5 milliseconds. Dendrites tend to be shorter than axons, and propagation speeds within dendrites are also quite fast. The main reason why vision is far from instantaneous is the multiple computations and integration steps in each neuron combined with the synaptic handoff of information from one neuron to the next executed throughout the multi-synaptic circuitry of cortex.

At each processing stage in the visual system, it is possible to estimate the time it takes for neurons in that area to realize that a flash of light was presented. Response latencies to a stimulus flash within the receptive field of a neuron increase from ~45 milliseconds in the LGN to ~100 milliseconds in the inferior temporal cortex (Table 5.1). There is an increase in the average latency within each area from the retina to the LGN to V1, to V2, to V4, to the ITC. This progression of latencies has further reinforced the notion of the ventral processing stream as an approximately hierarchical and sequential architecture. Each additional processing stage along the ventral stream adds an average of ~15 milliseconds of computation time.

It should be emphasized that these are only coarse values, and there is significant neuron-to-neuron variability within each area. An analysis of neural recordings in anesthetized monkeys by Schmolesky and colleagues showed latencies ranging from 30 milliseconds all the way to 70 milliseconds in the primary visual cortex. Because of this heterogeneity, the distributions of response latencies overlap, and the fastest neurons in a given area (say V2) may fire before the slowest neurons in an earlier area (say V1). Not only is there heterogeneity in response latencies from one neuron to another within a given visual area, but even the same neuron can also show different latencies depending on the nature of the stimulus. For example, response latencies tend

5.13 Receptive Field Sizes Increase along the Ventral Visual Stream

Table 5.1 Response latencies in different areas in the macaque monkey. From Schmolesky et al. 1998.

Area	Mean (ms)	S.D. (ms)
LGNd M layer	33	3.8
LGNd P layer	50	8.7
V1	66	10.7
V2	82	21.1
V4	104	23.4
V3	72	8.6
MT	72	10.3
MST	74	16.1
FEF	75	13

to be inversely proportional to the stimulus contrast. The notion of sequential processing is only a coarse approximation. However, the response latencies constitute an important constraint to the number of possible computational steps along the visual system.

Because of these latencies, we continuously live in the past in terms of vision. The notion that we only see the past events is particularly evident when we consider distant stars. The light signals that reach the Earth left those stars a long time ago. Although much less intuitive, the same idea applies to visual processing in the brain. Of course, the time it takes for light to bounce on a given object and reach the retina is negligible, yet signal propagation in the brain takes on the order of a hundred milliseconds, as discussed earlier. Through learning, the brain might be able to account for these delays by predicting what will happen next. For example, how is it possible for a Ping-Pong player to respond to a smash? The ball may be moving at about 50 kilometers per hour (apparently, the world record is about 112 kilometers per hour), and thus, the ball traverses the ~3-meter length of the table in about 200 milliseconds. By the time the opponent has to hit the ball back, his or her visual cortex is processing sensory inputs from the time when the ball was passing the net in the best-case scenario, not to mention the fact that orchestrating a movement also takes time (signals need to propagate from vision to the decision centers of the brain, and then from there to the muscles; all of these steps cost time). The only way to play Ping-Pong and other sports is to use the visual input combined with predictions learned through experience. Because of these predictions, players not only capitalize on smashing speed but also recur to other strategies such as embedding the ball with spinning effects to confuse the opponent.

5.13 Receptive Field Sizes Increase along the Ventral Visual Stream

Concomitant with the prolonged latencies, as we ascend through the visual hierarchy, receptive fields become larger (Figure 5.11). Receptive fields range from less than one degree in the initial steps (LGN, V1) all the way to several degrees or even tens of degrees in the highest echelons of the cortex. Each area has a complete map of the visual

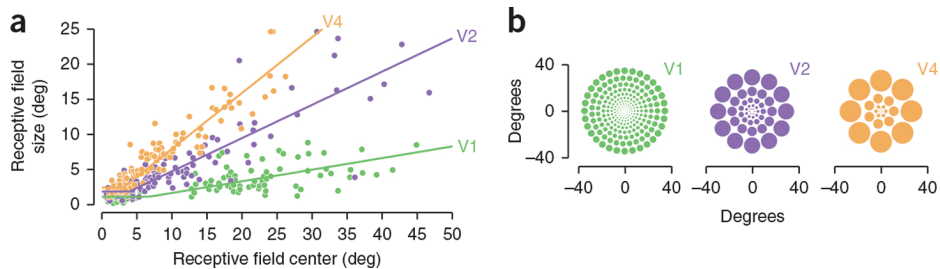


Figure 5.11 Receptive field sizes increase with eccentricity and along the ventral stream. Receptive field size increases with eccentricity for a given area. Additionally, receptive field size increases along the ventral visual stream at a fixed eccentricity. (A) Experimental measurements based on neurophysiological recordings in macaque monkeys. (B) Schematic rendering of receptive field sizes in areas V1, V2, and V4. Reproduced from Freeman and Simoncelli 2011

field; thus, the centers of the receptive fields go from the fovea to the periphery. As discussed for the primary visual cortex, within each area, the size of the receptive field increases as we move farther away from the fovea. There is always better resolution in the fovea, across all visual areas. The range of receptive field sizes within an area also increases with the mean receptive field size. The distributions are relatively narrow in primary visual cortex, but investigators have described a wide range of receptive field sizes in V4 and inferior temporal cortex. The scaling factor between receptive field size and eccentricity is more pronounced in V4 than in V2 and more pronounced in V2 compared to V1.

The increase in receptive field size from one area to the next may be a natural consequence of pooling-like operations in a hierarchical network, as we will discuss in more detail when we introduce computational models of visual cortical processing in Section 8.2. The increase in receptive field size provides several interesting properties: (i) a specific mechanism of discarding precise positional information in favor of (ii) extracting visual features that show progressively larger degrees of invariance to the exact position or scale of relevant visual features, and (iii) the ability to combine shapes from slightly shifted locations to build progressively more complex visual feature descriptors.

5.14 What Do Neurons beyond V1 Prefer?

There have been a few systematic parametric studies of the neuronal preferences in areas V2 and V4. These studies have opened the doors to investigate the complex transformations along the ventral visual stream. Even though multiple interesting studies compared responses in V1, V2, and V4, we do not yet have a clear, unified theory of what neurons “prefer” in these higher visual areas. Of course, the term “prefer” is an anthropomorphism. Neurons do not prefer anything. They fire spikes whenever the integration of their inputs exceeds a given threshold. Investigators often speak about neuronal preferences in terms of what types of images will elicit high firing rates.

The notion that V1 neurons show a preference for orientation tuning is well established, even if this only accounts for part of the variance in V1 responses to natural stimuli. There is significantly less agreement as to the types of shape features that are encoded in V2 and V4. There have been several studies probing responses with stimuli that are more complex than oriented bars and less complex than everyday objects. These stimuli include sinusoidal gratings, hyperbolic gratings, polar gratings, angles formed by intersecting lines, and curvatures with different properties, among others. Simple stimuli such as Cartesian gratings can certainly drive responses in V2 and V4. As a general rule, neurons in V2 and V4 can be driven more strongly by more complex shapes. As discussed earlier in the context of latency, there is a wide distribution of stimulus preferences in V2 and V4.

Perhaps one of the challenges is that investigators seek an explanation of neural coding preferences in terms of colloquial English expressions such as orientation, color, or curvature. An attractive idea that is gaining momentum is the notion that neurons in these higher visual areas filter the inputs from previous stages to produce complex tuning functions that defy language-based descriptions. A neuron may be activated by a patch representing complex shapes and textures that cannot be simply defined as an angle or a convex curve. Ultimately, the language of nature is mathematics, not English or Esperanto. Neuronal tuning properties do not have to map in any direct way to a short language-based description; we will come back to this idea in Chapters 7 and 8 when we discuss computational models of vision.

5.15 Brains Construct Their Interpretation of the World: The Case of Illusory Contours

A pervasive illusion is the notion that our senses contain a veridical representation of precisely what is out there in the world. This notion can be readily debunked through the study of visual illusions. In Section 3.1, we argued that our brains make up stuff by constructing an interpretation of the outside world. Our brains “making up stuff” implies that there should be neurons that explicitly represent those constructs. Let us revisit the Kanizsa triangle (Figure 5.12), where we have the strong illusion of perceiving an equilateral triangle in the midst of the three Pacman icons. The small parts of the sides of the triangle near the vertices are composed of real black contours. However, the center of each side is composed of a line that does not really exist. These lines represent illusory contours – that is, edges created without any change in luminance.

It is relatively easy to “trick the eye,” except that the eye is typically *not* tricked in most visual illusions. Visual illusions represent situations where our brains construct an interpretation of the image that is different from the pixel level content. In most such illusions, the responses of retinal ganglion cells (RGC) follow the pixel-level content in the image relatively well. Consider recording the activity of an RGC whose receptive field center corresponds to position A in Figure 5.12, right along one side of the Pacman. There is a luminance change inside the receptive field, and we expect the neuron to fire vigorously at this location upon flashing the Kanizsa figure. Now consider an RGC with




Figure 5.12 V2 neurons can represent lines that do not exist except in the eyes of the beholder. The figure shows the Kanizsa triangle visual illusion and a schematic rendering of neurophysiological recordings from four neurons: two retinal ganglion cells (RGC) and two V2 neurons. When the receptive fields (gray dotted circles) encompass locations that have a real contour (A), both RGC and V2 neurons fire vigorously. In contrast, when the receptive fields encompass an illusory contour (B), the V2 neuron fires vigorously, but the RGC neuron only fires a few baseline spikes.

a receptive field center located at position B, smack in the middle of the illusory contour. We do not expect this neuron to fire above baseline levels because there is no stimulus inside the receptive field. In other words, the activity of RGCs does not correlate with our perception. If the retina does not reflect perception, then who does? It seems reasonable to conjecture that there must be neurons somewhere that explicitly represent the contents of our perception – in this case, the illusory contours. This explicit representation is a critical postulate that we will discuss again in more depth when we take up the question of the neuronal correlates of consciousness in Section 10.3.

Indeed, neurons in area V2 respond to illusory contours (Figure 5.12). A V2 neuron that prefers horizontal edges would fire strongly if its receptive field is at location A because there is a real horizontal line there. Remarkably, a V2 neuron that prefers leftward edges would also fire if its receptive field is at position B, where there is an illusory edge. V2 neurons respond almost equally well to an illusory line or to a real line. The responses to illusory contours are remarkable because there is no contrast change within the neuron's receptive field. Hence, these responses indicate a form of context modulation that is consistent with the subjective interpretation of borders. There are also neurons in V1 that respond to illusory contours, but there are more such neurons in V2. Interestingly, the responses to illusory contours show a short delay with respect to the responses to real lines. These delays may reflect the need for additional computational steps required to infer the presence of a line when there is none.

5.16 A Colorful V4

Neurons in the retina (cones), LGN (parvocellular neurons), and primary visual cortex (particularly those within so-called blobs in V1) are all sensitive to the color of the

stimulus within their receptive field. Neurons in area V4 demonstrate sensitivity to color properties that are more complex than those in earlier areas. A notable property is that neurons in V4 have been implicated in the phenomenon of color constancy whereby an object's perceived color is relatively insensitive to large changes in the overall illumination, in contrast to the responses earlier in the visual system.

There are many visual illusions based on the phenomenon of color constancy. A banana typically appears to be yellow to our eyes, whether we see it at noon, or in the early evening, or under the kitchen light. The actual spectrum of light reaching the eyes depends quite strongly on the environment illumination, and cones in the retina signal the actual wavelengths reflected off the banana. However, our perception discounts the background illumination and interprets the banana to be yellow. The integration of color signals emanating within the receptive field with those in the surround to perform this type of discounting is thought to take place in V1, and even more clearly in V4 neurons. The responses of V4 neurons better correlate with how primates perceive colors. Furthermore, the rare condition of cortical color blindness known as achromatopsia has been associated with damage to area V4 (Section 4.8).

5.17 Attentional Modulation

As noted earlier in this chapter, neurons along the ventral visual cortex receive numerous top-down signals in addition to their bottom-up inputs. Through these top-down signaling mechanisms, the activity of neurons along ventral visual cortex can be strongly modulated by signals beyond the specific visual content within their receptive fields – including spatial context, temporal context, expectations, and higher-level cognitive influences such as task goals.

Despite keen interest in such top-down signals, there have been many more studies about the role of bottom-up inputs on neuronal responses. At least partly, this imbalance is due to the fact that it is much easier to change what is shown on the screen than to change an animal's internal expectations and goals.

A prime example of the study of top-down modulatory signals in visual processing involves *spatial attention*. One way to allocate attention to one part of the visual field is by moving the eyes. However, spatial attention effects can also be demonstrated outside of the fixation focus. A subject can be looking at one location and paying attention to another place, a phenomenon known as *covert attention* (as opposed to *overt attention*, which is the more common scenario where attention is allocated to the fixation area). Through a series of astute training paradigms, investigators have been able to train animals to deploy covert spatial attention, thus enabling them to investigate the consequences of spatial attention on neurons with receptive fields outside the fovea.

An animal is trained to fixate in the center of the screen, and its eye movements are strictly monitored to ensure that attentional effects are not driven by saccades. In some trials, the animal is rewarded for detecting a visual stimulus in a certain location on the right, and that tells the animal to allocate attention to that region of the visual field without breaking fixation. Compliance can be checked by randomly probing a stimulus

presented at another location and showing that performance is better (faster, more accurate) in the attended area.

Under these experimental conditions, neurons typically show an enhancement in the responses when their receptive field is within the focus of attention, particularly upon presentation of a visual stimulus. In other words, imagine a neuron in V2 with a receptive field location that is right at the center of the attended area in some trials and outside the attended area in other trials. The neuron will respond to an identical visual stimulus with more spikes in those trials when attention encompasses the receptive field. The effect of spatial attention is not all-or-none. Neurons still respond vigorously to a stimulus placed within their receptive field regardless of whether or not the animal is paying attention to that location. Attention leads to about 5 to 30 percent increased firing rates. The magnitude of this attention effect follows the reverse hierarchical order, being significantly stronger in area V4 compared to area V1.

Neuronal responses can also be modulated in a feature-specific manner. Instead of paying attention to a particular location, the animal can be trained to pay attention to a specific stimulus feature such as the color red or vertical lines. When the animal is paying attention to the neuron's preferred features, the neuron shows an enhanced firing rate.

5.18 Summary

- Visual computations transpire in the six-layered neocortical structure.
- The cortex is characterized by stereotypical connectivity patterns from one area to the next, forming approximately canonical microcircuits.
- The gold standard to study cortical function is to scrutinize the activity of individual neurons.
- Neurons in the primary visual cortex detect edges and show orientation tuning, responding more strongly to a bar in a specific orientation within the receptive field.
- Complex neurons in the primary visual cortex show tolerance to the exact position of the preferred stimulus within the receptive field.
- A Gabor function can phenomenologically fit the responses of V1 neurons.
- A mechanistic model posits that V1 simple cell receptive fields can be created by adequately combining the outputs of center-surround neurons from the lateral geniculate nucleus positioned to create the desired orientation.
- A model posits that V1 complex cell receptive fields can be created by adequately combining the outputs of V1 simple cells with the same orientation preferences but slightly shifted receptive fields.
- The visual cortex uses a divide-and-conquer strategy, subdividing visual processing into a sequence of computations in tens of different brain areas arranged into an approximate hierarchy.

- Ascending through the visual hierarchy, neurons show increased receptive field sizes, more complex tuning preferences, and longer latencies.
- Neurons in area V2 respond to illusory contours.
- Spatial context, temporal context, and task demands like attention can modulate neuronal responses along ventral visual cortex.

Further Reading

See <http://bit.ly/2TpAg3w> for more references.

- Carandini, M.; Demb, J. B.; Mante, V.; Tolhurst, D. J.; Dan, Y., et al. (2005). Do we know what the early visual system does? *Journal of Neuroscience* 25: 10577–10597.
- Hubel, D.H.; and Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology* 195: 215–243.
- Kremkow, J.; Jin, J.; Wang, Y.; and Alonso, J. M. (2016). Principles underlying sensory map topography in primary visual cortex. *Nature* 533: 52–57.
- Markov, N.T.; Ercsey-Ravasz, M. M.; Ribeiro Gomes, A. R.; Lamy, C.; Magrou, L., et al. (2014). A weighted and directed interareal connectivity matrix for macaque cerebral cortex. *Cerebral Cortex* 24: 17–36.
- Schmolesky, M.; Wang, Y.; Hanes, D.; Thompson, K.; Leutgeb, S.; et al. (1998). Signal timing across the macaque visual system. *Journal of Neurophysiology* 79: 3272–3278.