

Contents

1	The Photoreceptor Mosaic	5
1.1	The S Cone Mosaic	10
1.2	Visual Interferometry	15
1.3	Sampling and Aliasing	18
1.4	The L and M Cone Mosaic	22
1.5	Summary and Discussion	25

List of Figures

1.1	Rods and Cones	6
1.2	Schematic of Rods and Cones	7
1.3	Cone Spectral Sensitivities	8
1.4	Photoreceptor Sampling	9
1.5	Calculating Viewing Angle	10
1.6	Short-wavelength Cone Mosaic: Psychophysics	12
1.7	Short-Wavelength Cone Mosaic: Procion Yellow Stains	13
1.8	Interference and Double Slits	15
1.9	Visual Interferometer	16
1.10	Sinusoidal Interference Pattern	17
1.11	Aliasing Examples	19
1.12	Squarewave aliasing	20
1.13	Drawings of Aliases	23
1.14	Choosing monitor phosphors.	27
1.15	Homework Problem: Sensor sample positions	29

Chapter 1

The Photoreceptor Mosaic

In Chapter ?? we reviewed Campbell and Gubisch's (1967) measurements of the optical linespread function. Their data are presented in Figure ??, as smooth curves, but the actual measurements must have taken place at a series of finely spaced intervals called *sample points*. In designing their experiment, Campbell and Gubisch must have considered carefully how to space their sample points because they wanted to space their measurement samples only finely enough to capture the intensity variations in the measurement plane. Had they positioned their samples too widely, then they would have missed significant variations in the data. On the other hand, spacing the sample positions too closely would have made the measurement process wasteful of time and resources.

Just as Campbell and Gubisch sampled their linespread measurements, so too the retinal image is sampled by the nervous system. Since only those portions of the retinal image that stimulate the visual photoreceptors can influence vision, the sample positions are determined by the positions of the photoreceptors. If the photoreceptors are spaced too widely, the image encoding will miss significant variation present in the retinal image. On the other hand, if the photoreceptors are spaced very close to one another compared to the spatial variation that is possible given the inevitable optical blurring, then the image encoding will be redundant, using more neurons than necessary to do the job. In this chapter we will consider how the spatial arrangement of the photoreceptors, called the *photoreceptor mosaic*, limits our ability to infer the spatial pattern of light intensity present in the retinal image.

We will consider separately the photoreceptor mosaics of each of the different types of photoreceptors. There are two fundamentally different types of photoreceptors in our eye, the *rods* and the *cones*. There are approximately 5 million cones and 100 million rods in each eye. The positions of these two types of photoreceptors differ in many ways across the retina. Figure 1.1 shows how the relative densities of cone

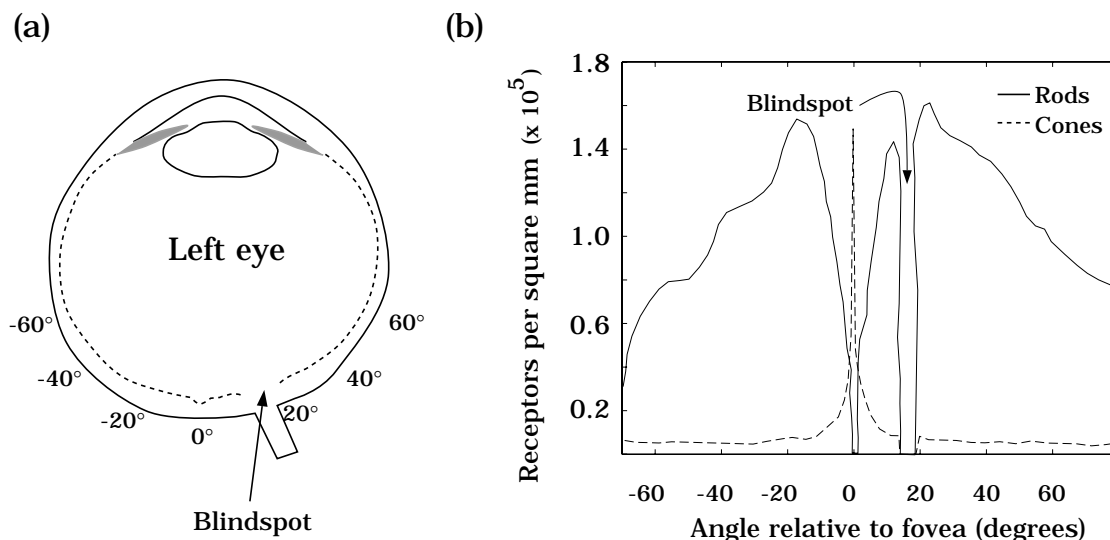


Figure 1.1: *The distribution of rod and cone photoreceptors across the human retina.* (a) The density of the receptors is shown in degrees of visual angle relative to the position of the fovea for the left eye. (b) The cone receptors are concentrated in the fovea. The rod photoreceptors are absent from the fovea and reach their highest density 10 to 20 degrees peripheral to the fovea. No photoreceptors are present in the blindspot.

photoreceptors and rod photoreceptors vary across the retina.

The rods initiate vision under low illumination levels, called *scotopic* light levels, while the cones initiate vision under higher, *photopic* light levels. The range of intensities in which both rods and cones can initiate vision is called *mesopic* intensity levels. At most wavelengths of light, the cones are less sensitive to light than the rods. This sensitivity difference, coupled with the fact that there are no rods in the fovea, explains why we can not see very dim sources, such as weak starlight, when we fixate our fovea directly on them. These sources are too dim to be visible through the all cone fovea. The dim source only becomes visible when it is placed in the periphery and be detected by the rods. Rods are very sensitive light detectors: they generate a detectable photocurrent response when they absorb a single photon of light (Hecht et al., 1942; Schwartz, 1978; Baylor et al. 1987).

The region of highest visual acuity in the human retina is the *fovea*. As Figure 1.1 shows, the fovea contains no rods, but it does contain the highest concentration of cones. There are approximately 50,000 cones in the human fovea. Since there are no photoreceptors at the optic disk, where the ganglion cell axons exit the retina, there is a blindspot in that region of the retina (see Chapter ??).

Figure 1.2 shows schematics of a mammalian rod and a cone photoreceptor. Light imaged by the cornea and lens is shown entering the receptors through the *inner segments*. The light passes into the *outer segment* which contain light absorbing




Figure 1.2: *Mammalian rod and cone photoreceptors* contain the light absorbing pigment that initiates vision. Light enters the photoreceptors through the inner segment and is funneled to the outer segment that contains the photopigment. (After Baylor, 1987)

photopigments. As light passes from the inner to the outer segment of the photoreceptor, it will either be absorbed by one of the photopigment molecules in the outer segment or it will simply continue through the photoreceptor and exit out the other side. Some light imaged by the optics will pass between the photoreceptors. Overall, less than ten percent of the light entering the eye is absorbed by the photoreceptor photopigments (Baylor, 1987).

The rod photoreceptors contain a photopigment called *rhodopsin*. The rods are small, there are many of them, and they sample the retinal image very finely. Yet, visual acuity under scotopic viewing conditions is very poor compared to visual acuity under photopic conditions. The reason for this is that the signals from many rods converge onto a single neuron within the retina, so that there is a many-to-one relationship between rod receptors and neurons in the optic tract. The density of rods and the convergence of their signals onto single neurons improves the sensitivity of rod-initiated vision. Hence, rod-initiated vision does not resolve fine spatial detail.

The foveal cone signals do not converge onto single neurons. Instead, several neurons encode the signal from each cone, so that there is a one-to-many relationship between the foveal cones and optic tract neurons. The dense representation of the foveal cones suggests that the spatial sampling of the cones




Figure 1.3: *Spectral sensitivities of the L, M and S cones in the human eye.* The measurements are based on a light source at the cornea, so that the wavelength loss due to the cornea, lens and other inert pigments of the eye play a role in determining the sensitivity. (Source: Stockman and Macleod, 1993).

must be an important aspect of the visual encoding.

There are three types of cone photoreceptors within the human retina. Each cone can be classified based on the wavelength sensitivity of the photopigment in its outer segment. Estimates of the spectral sensitivity of the three types of cone photoreceptors are shown in Figure 1.3. These curves are measured from the cornea, so they include light loss due to the cornea, lens and inert materials of the eye. In the next chapter we will study how color vision depends upon the differences in wavelength selectivity of the three types of cones. Throughout this book I will refer to the three types of photoreceptors as the L, M and S cones¹.

Because light is absorbed after passing through the inner segment, the position of the inner segment determines the spatial sampling position of the photoreceptor. Figure 1.4 shows cross-sections of the human cone photoreceptors at the level of the inner segment in the human fovea (part a) and just outside the fovea (part b). In the fovea, cross-section shows that the inner segments are very tightly packed and form a regular sampling array. A cross-section just outside the fovea shows that the rod photoreceptors fill the spaces between the cones and disrupt the regular packing arrangement. The scale bar represents $10\mu m$; the cone photoreceptor inner segments

¹The letters refer to Long-wavelength, Middle-wavelength and Short-wavelength peak sensitivity.

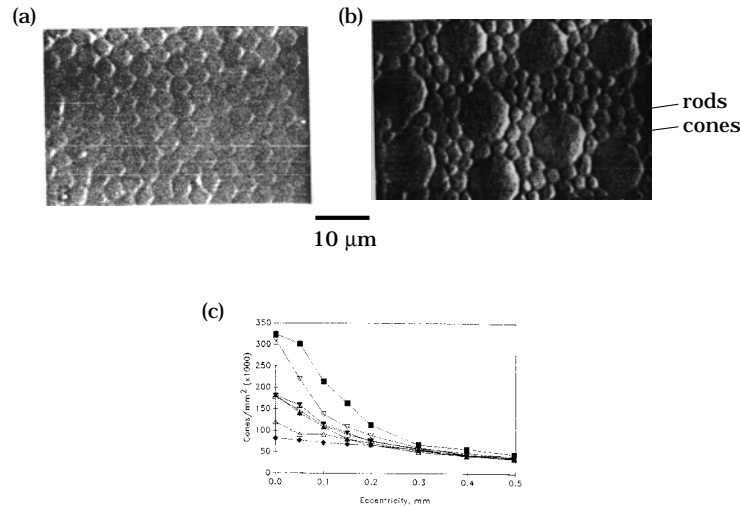


Figure 1.4: *The spatial mosaic of the human cones.* A cross-section of the human retina at the level of the inner segments. Cones in the fovea (a) are smaller than cones in the periphery (b). As the separation between cones grows, the rod receptors fill in the spaces. (c) The cone density varies with distance from the fovea. Cone density is plotted as a function of eccentricity for seven human retinæ (After Curcio et al, 1990).

in the fovea are approximately $2.3\mu\text{m}$ wide with a minimum center to center spacing of about $2.5\mu\text{m}$. Figure 1.4c shows plots of the cone densities from several different human retinæ as a function of the distance from the foveal center. The cone density varies across individuals.

Units of Visual Angle

We can convert these cone sizes and separations into degrees of visual angle as follows. The distance from the effective center of of the eye's optics to the retina is $1.7 \times 10^{-2}\text{m}$ (17 mm). We compute the visual angle spanned by one cone, ϕ , from the trigonometric relationship in Figure 1.5: the tangent of an angle in a right triangle is equal to the ratio of the lengths of the sides opposite and adjacent to the angle. This leads to the following equation:

$$\tan(\phi) = (2.5 \times 10^{-6}\text{m}) / (1.7 \times 10^{-2}\text{m}) = 1.47 \times 10^{-4} \quad (1.1)$$

The width of a cone in degrees of visual angle, ϕ , is approximately 0.0084 degrees, or roughly one-half minute of visual angle. In the center of the eye, then, where the photoreceptors are packed densely, the cone photoreceptors are tightly packed and their centers are separated by one-half minute of visual angle.




Figure 1.5: *Calculating viewing angle.* By trigonometry, the tangent of the viewing angle, ϕ , is equal to the ratio of height to distance in the right triangle shown. Therefore, ϕ is the inverse tangent of that ratio (Equation 1.1).

1.1 The S Cone Mosaic

Behavioral Measurements

Just as the rods and cones have different spatial sampling distributions, so too the three types of cone photoreceptors have different spatial sampling distributions. The sampling distribution of the short-wavelength cones was the first to be measured empirically, and it has been measured both with behavioral and physiological methods. The behavioral experiments were carried out as part of D. Williams dissertation at the University of California in San Diego. Williams, Hayhoe and MacLeod (1981) took advantage of several features of the short-wavelength photoreceptors. As background to their work, we first describe several features of the photoreceptors.

The photopigment in the short-wavelength photoreceptors is significantly different from the photopigment in the other two types of photoreceptors. Notice that the wavelength sensitivity of the L and M photopigments are very nearly the same (Figure 1.3). The sensitivity of the S photopigment is significantly higher in the short-wavelength part of the spectrum than the sensitivity of the other two photopigments. As a result, if we present the visual system with a very weak light, containing energy only in the short-wavelength portion of the spectrum, the S cones will absorb relatively more quanta than the other two classes. Indeed, the discrepancy in the absorptions is so large that it is reasonable to suppose that when short-wavelength light is barely visible, at detection threshold, perception is initiated uniquely from a signal that originates in the short-wavelength receptors.

We can give the short-wavelength receptors an even greater sensitivity advantage by presenting a blue test target on a steady yellow background. As we will discuss in later chapters, steady backgrounds suppress visual sensitivity. By using a yellow background, we can suppress the sensitivity of the L and M cones and the rods and yet spare the sensitivity of the S cones. This improves the relative sensitivity advantage of the short-wavelength receptors in detecting the short-wavelength test light.

A second special feature of the S cones is that they are very rare in the retina. From other experiments described in Chapter ??, it has been suspected for many years that no cones containing short-wavelength photopigment are present in the central fovea. It had been earlier suspected that the number of cones containing the short-wavelength photopigment was quite small compared to the other two classes. If the S cones are widely spaced, and if we can isolate them with these choices of test stimulus and background, then we can measure the mosaic of short-wavelength photoreceptors.

During the experiment, the subjects visually fixated on a small mark. They were then presented with short-wavelength test lights that were likely to be seen with a signal initiated by the S cones. After the eye was perfectly fixated, the subject pressed a button and initiated a stimulus presentation. The test stimulus was a tiny point of light, presented very briefly (10 ms). The test light was presented at different points in the visual field. If light from the short-wavelength test fell upon a region that contained S cones, sensitivity should be relatively high. On the other hand, if that region of the retina contained no S cones, sensitivity should be rather low. Hence, from the spatial pattern of visual sensitivity, Williams, Hayhoe and Macleod inferred the spacing of the S cones.

The sensitivity measurements are shown in Figure 1.6. First, notice that in the very center of the visual field, in the central fovea, there is a large valley of low sensitivity. In this region, there appear to be no short-wavelength cones at all. Second, beginning about half a degree from the center of the visual field there are small, punctate spatial regions of high sensitivity. We interpret these results by assuming that these peaks correspond to the positions of this observer's S cones. The gaps in between, where the observer has rather low sensitivity are likely to be patches of L and M cones. Around the central fovea, the typical separation between the inferred S cones is about 8 to 12 minutes of visual angle. Thus, there are five to seven S cones per degree of visual angle.

Biological Measurements

There have been several biological measurements of the short-wavelength cone mosaic, and we can compare these with the behavioral measurements. Marc and




Figure 1.6: *Psychophysical estimate of the spatial mosaic of the S cones.* The height of the surface represents the observer's threshold sensitivity to a short wavelength test light presented on a yellow background. The test was presented at a series of locations spanning a grid around the fovea (black dot). The peaks in sensitivity probably correspond to the positions of the S cones. (From Williams, Hayhoe, and Macleod, 1981).




Figure 1.7: *Biological estimate of the spatial mosaic of the S cones* in the macaque retina. A small fraction of the cones absorb the procion yellow stain; these are shown as the dark spots in this image. These cones, thought to be the S cones, are shown in a cross-section through the inner segment layer of the retina. (From DeMonasterio, Schein and McCrane, 1985)

Sperling (1977) used a stain that is taken up by cones when they are active. They applied this stain to a baboon retina and then stimulated the retina with short-wavelength light in the hopes of staining only the short-wavelength receptors. They found that only a few cones were stained when the stimulus was a short-wavelength light. The typical separation between the stained cones was about 6 minutes of arc. This value is smaller than the separation that Williams' et al. observed and may be a species-related difference.

F. DeMonasterio, S. Schein, and E. McCrane (1981) discovered that when the dye procion yellow is applied to the retina, the dye is absorbed in the outer segments of all the photoreceptors, but it stains only a small subset of the photoreceptors completely. Figure 1.7 shows a group of stained photoreceptors in cross-section section.

The indirect arguments identifying these special cones as S cones are rather compelling. But, a more certain procedure was developed by C. Curcio and her colleagues. They used a biological marker, developed based on knowledge of the genetic code for the S cone photopigment, to label selectively the S cones in the human retina (Curcio, et al. 1991). Their measurements agree well quantitatively with Williams' psychophysical measurements, namely that the average spacing between the S cones is 10 minutes of visual angle. Curcio and her colleagues could also confirm some early anatomical observations that the size and shape of the S cones differ slightly from the L and M cones. The S cones have a wider inner

segment, and they appear to be inserted within an orderly sampling arrangement of their own between the sampling mosaics of the other two cone types (Ahnelt, Kolb and Pflug, 1987).

Why are the S cones widely spaced?

The spacing between the S cones is much larger than the spacing between the L and M cones. Why should this be? The large spacing between the S cones is consistent with the strong blurring of the short-wavelength component of the image due to the axial chromatic aberration of the lens. Recall that axial chromatic aberration of the lens blurs the short-wavelength portion of the retinal image, the part S cones are particularly sensitive to, more than the middle- and long-wavelength portion of the image (Figure ??). In fact, under normal viewing conditions the retinal image of a fine line at 450 nm falls to one half its peak intensity nearly 10 minutes of visual angle away from the location of its peak intensity. At that wavelength, the retinal image only contains significant contrast at spatial frequency components below 3 cycles per degree of visual angle. The optical defocus force the wavelength components of the retinal image the S cones encode to vary smoothly across space. Consequently, the S cones can sample the image only six times per degree and still recover the spatial variation passed by the cornea and lens.

Interestingly, the *spatial* defocus of the short-wavelength component of the image also implies that signals initiated by the S cones will vary slowly over *time*. In natural scenes, temporal variation occurs mainly because of movement of the observer or an object. When a sharp boundary moves across a cone position, the light intensity changes rapidly at that point. But, if the boundary is blurred, changing gradually over space, then the light intensity changes more slowly. Since the short-wavelength signal is blurred by the optics, and temporal variation is mainly due to motion of objects, the S cones will generally be coding slower temporal variations than the L and M cones.

At the very earliest stages of vision, we see that the properties of different components of the visual pathway fit smoothly together. The optics set an important limit on visual acuity, and the S cone sampling mosaic can be understood as a consequence of the optical limitations. As we shall see, the L and M cone mosaic densities also make sense in terms of the optical quality of the eye.

This explanation of the S cone mosaic flows from our assumption that visual acuity is the main factor governing the photoreceptor mosaic. For the visual streams initiated by the cones, this is a reasonable assumption. There are other important factors, however, that can play a role in the design of a visual pathway. For example, acuity is not the dominant factor in the visual stream initiated by rod vision. In principle the resolution available in the rod encoding is comparable to the acuity




Figure 1.8: *T. Young's double-slit experiment* uses a pair of coherent light sources to create an interference pattern of light. The intensity of the resulting image is nearly sinusoidal, and its spatial frequency depends upon the spacing between the two slits.

available in the cone responses; but, visual acuity using rod-initiated signals is very poor compared to acuity using cone-initiated signals. Hence, we shouldn't think of the rod sampling mosaic in terms of visual acuity. Instead, the high density of the rods and their convergence onto individual neurons suggests that we think of the imperative of rod-initiated vision in terms of improving the signal-to-noise under low light levels. In the rod-initiated signals, the visual system trades visual acuity for an increase in the signal-to-noise ratio. In the earliest stages of the visual pathways, then, we can see structure, function and design criteria coming together.

When we ask why the visual system has a particular property, we need to relate observations from the different disciplines that make up vision science. Questions about anatomy require us to think about the behavior the anatomical structure serves. Similarly, behavior must be explained in terms of algorithms and the anatomical and physiological responses of the visual pathway. By considering the visual pathways from multiple points of view, we piece together a complete picture of how system functions.

1.2 Visual Interferometry

In behavioral experiments, we measure threshold repeatedly through individual L and M using small points of light as we did the S cones. The pointspread function




Figure 1.9: A *visual interferometer* creates an interference pattern as in Young's double-slit experiment. In the device shown here the original beam is split into two paths shown as the solid and dashed lines. (a) When the glass cube is at right angles to the light path, the two beams traverse an equal path and are imaged at the same point after exiting the interferometer. (b) When the glass is rotated, the two beams traverse slightly different paths causing the images of the two coherent beams to be displaced and thus create an interference pattern. (After Macleod, Williams and Makous, 1992).

distributes light over a region containing about twenty cones, so that the visibility of even a small point of light may involve any of the cones from a large pool (see Figures ?? and ??). We can, however, use a method introduced by Y. LeGrand in 1935 to defeat the optical blurring. The technique is called *visual interferometry*, and it is based upon the principle of diffraction.

Thomas Young (1802), the brilliant scientist, physician, and classicist demonstrated to the Royal Society that when two beams of coherent light generate an image on a surface such as the retinal surface, the resulting image is an interference pattern. His experiment is often called the *double-slit* or *double-pinhole* experiment. Using an ordinary light source, Young passed the light through a small pinhole first and then through a pair of slits, as illustrated in Figure 1.8. In the experiment, the first pinhole serves as the source of light; the double pinholes then pass the light from the common original source. Because they share this common source, light emitted from the double pinholes are in a coherent phase relationship and their wavefronts interfere with one another. This interference results in an image that varies nearly sinusoidally in intensity.

We can also achieve this narrow pinhole effect by using a laser as the original source. The key elements of a visual interferometer used by MacLeod et al. (1992) are shown in Figure 1.9. Light from a laser enters the beamsplitter and is divided into one part that continues along a straight path (solid line) and a second path that is reflected




Figure 1.10: *An interference pattern.* The image was created using a double-slit apparatus. The intensity of the pattern is nearly sinusoidal. (From Jenkins and White, 1976.)

along a path to the right (dashed line). These two beams, originating from a common source, will be the pair of sources to create the interference pattern on the retina.

Light from each beam is reflected from a mirror towards a glass cube. By varying the orientation of the glass cube, the experimenter can vary the path of the two beams. When the glass cube is at right angles to the light path, as is shown in part (a), the beams continue in a straight path along opposite directions and emerge from the beamsplitter at the same position. When the glass cube is rotated, as is shown in part (b), the refraction due to the glass cube symmetrically changes the beam paths; they emerge from the beamsplitter at slightly different locations and act as a pair of point sources. This configuration creates two coherent beams that act like the two slits in Thomas Young's experiment, creating an interference pattern. The amount of rotation of the glass cube controls the separation between the two beams.

Each beam passes through only a very small section of the cornea and lens. The usual optical blurring mechanisms do not interfere with the image formation, since the lens does not serve to converge the light (see the section on lenses in Chapter ??). Instead, the pattern that is formed depends upon the diffraction due to the restricted spatial region of the light source.

We can use diffraction to create retinal images with much higher spatial frequencies than are possible through ordinary optical imaging by the cornea and lens. Figure 1.10 is an image of a diffraction pattern created by a pair of two slits. The intensity of the pattern is nearly a sinusoidal function of retinal position. The spatial frequency of the retinal image can be controlled by varying the separation between the focal points; the smaller the separation between the slit, the lower the spatial frequency in the interference pattern. Thus, by rotating the glass cube in the interferometer and changing the separation of the two beams we can control the spatial frequency of the retinal image.

Visual interferometry permits us to image fine spatial patterns at much higher

contrast than when we image these patterns using ordinary optical methods. For example, Figure ?? shows that a 60 cycles per degree sinusoid cannot exceed 10 percent contrast when imaged through the optics. Using a visual interferometer, we can present patterns at frequencies considerably higher than 60 cycles per degree at 100 percent contrast.

But a challenge remains: the interferometric patterns are not fine lines or points, but rather extended patterns (cosinusoids). Therefore, we cannot use the same logic as Williams et al. and map the receptors by carefully positioning the stimulus. We need to think a little bit more about how to use the cosinusoidal interferometric patterns to infer the structure of the cone mosaic.

1.3 Sampling and Aliasing

In this section we consider how the cone mosaic encodes the high spatial frequency patterns created by visual interferometers. The appearance of these high frequency patterns will permit us to deduce the spatial arrangement of the combined L and M cone mosaics. The key concepts that we must understand to deduce the spatial arrangement of the mosaic are *sampling* and *aliasing*. These ideas are illustrated in Figure 1.11.

The most basic observation concerning sampling and aliasing is this: we can measure only that portion of the input signal that falls over the sample positions. Figure 1.11 shows one-dimensional examples of aliasing and sampling. Parts (a) and (b) contain two different cosinusoidal signals (left) and the locations of the sample points. The values of these two cosinusoids at the sample points are shown by the height of the arrows on the right. Although the two continuous cosinusoids are quite different, they have the same values at the sample positions. Hence, if cones are only present at the sample positions, the cone responses will not distinguish between these two inputs. We say that these two continuous signals are an *aliased* pair. Aliased pairs of signals are indistinguishable after sampling. Hence, sampling degrades our ability to discriminate between sinusoidal signals.

Figure 1.11c shows that sampling degrades our ability to discriminate between signals in general, not just between sinusoids. Whenever two signals agree at the sample points, their sampled representations agree. The basic phenomenon of aliasing is this: Signals that only differ between the sample points are indistinguishable after sampling.

The exercises at the end of this chapter include some computer programs that can help you make sampling demonstrations like the one in Figure 1.12. If you print out squarewave patterns and various sampling arrays, using the programs provided, you can print various patterns onto overhead transparencies and explore the effects

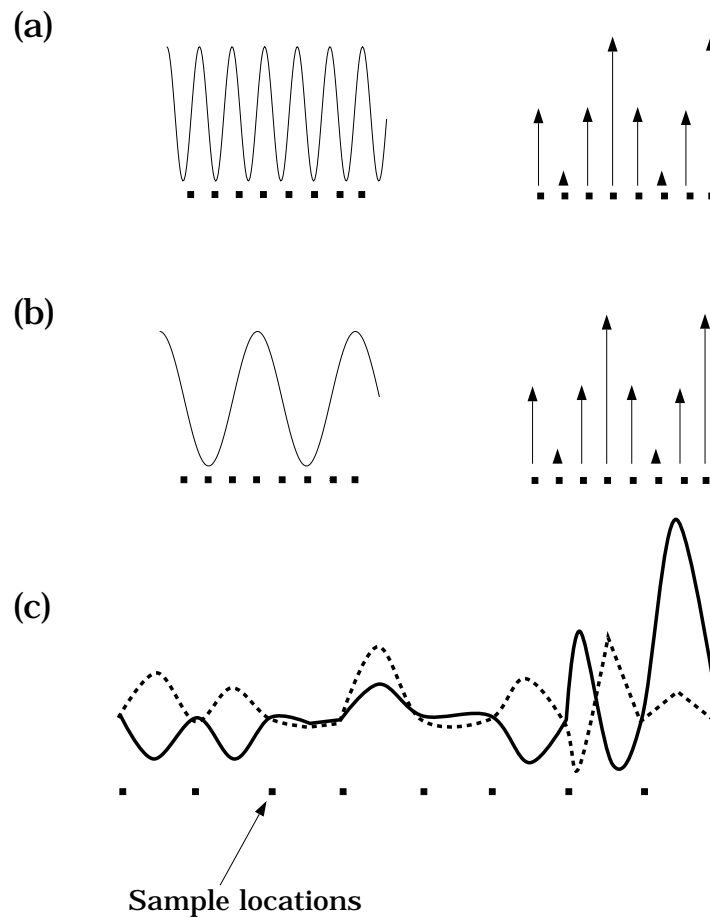


Figure 1.11: *Aliasing* of signals results when sampled values are the same but in-between values are not. (a,b) The continuous sinusoids on the left have the same values at the sample positions indicated by the black squares. The values of the two functions at the sample positions are shown by the height of the stylized arrows on the right. (c) Undersampling may cause us to confuse various functions, not just sinusoids. The two curves at the bottom have the same values at the sampled points, differing only in between the sample positions.




Figure 1.12: *Squarewave aliasing*. The squarewave on top is seen accurately through the grid. The squarewave on the bottom is at a higher spatial frequency than the grid sampling. When seen through the grid, the pattern appears at a lower spatial frequency and rotated.

of sampling. Figure 1.12 shows an example of two squarewave patterns seen through a sampling grid. After sampling, the high frequency pattern appears to be a rotated, low frequency signal.

Sampling is a Linear Operation. The sampling transformation takes the retinal image as input and generates a portion of the retinal image as output. Sampling is a linear operation as the following thought experiment reveals. Suppose we measure the sample values at the cone positions when we present image A ; call the intensities at the sample positions $S(A)$. Now, measure the intensities at the sample positions for a second image, B ; call the sample intensities $S(B)$. If we add together the two images, the new image, $A + B$, contains the sum of the intensities in the original images. The values picked out by sampling will be the sum of the two sample vectors, $S(A) + S(B)$.

Since sampling is a linear transformation, we can express it as a matrix multiplication. In our simple description, each position in the retinal image either falls within a cone inner segment or not. The sampling matrix consists of N rows representing the N sampled values. Each row is all zero except at the entry corresponding to that row's sampling position, where the value is 1.

Aliasing of harmonic functions. For uniform sampling arrays we have already observed that some pairs of sinusoidal stimuli are aliases of one another (part (a) of Figure 1.11). We can analyze precisely which pairs of sinusoids form alias pairs using a little bit of algebra. Suppose that the continuous input signal is $\cos(2\pi f x)$. When we sample the stimulus at regular intervals, the output values will be the value of the cosinusoid at those regularly spaced sample points. Suppose that within a single unit of distance there are N sample points, so that our measurements of the stimulus takes place every $1/N$ units. Then the sampled values will be $S_f(k) = \cos(2\pi f k/N)$. A second cosinusoid, at frequency f' will be an alias if its sample values are equal, that is, if $S_{f'}(k) = S_f(k)$.

With a little trigonometry, we can prove that the sample values for any pair of cosinusoids with frequencies $N/2 - f$ and $N/2 + f$ will be equal. That is,

$$\cos\left(\frac{2\pi(N/2 + f)k}{N}\right) = \cos\left(\frac{2\pi(N/2 - f)k}{N}\right)$$

(To prove this we must use the cosine addition law to expand the right sides of the following equation. The steps in the verification are left as exercise 5 at the end of the chapter.)

The frequency $f = N/2$ is called the *Nyquist frequency* of the uniform sampling array; sometimes it is referred to as the *folding frequency*. Cosinusoidal stimuli whose

frequencies differ by equal amounts above and below the Nyquist frequency of a uniform sampling array will have identical sample responses.

Experimental Implications. The aliasing calculations suggest an experimental method to measure the spacing of the cones in the eye. If the cone spacing is uniform, then pairs of stimuli separated by equal amounts above and below the Nyquist frequency should appear indistinguishable. Specifically, a signal $\cos(2\pi(N/2 + f))$ that is above the Nyquist frequency will appear the same as the signal $\cos(2\pi(N/2 - f))$ that is an equal amount below the Nyquist frequency. Thus, as subjects view interferometric patterns of increasing frequency, as we cross the Nyquist frequency the perceived spatial frequency should begin to decrease even though the physical spatial frequency of the diffraction pattern increases.

Yellott (1982) examined the aliasing prediction in a nice graphical way. He made a sampling grid from Polyak's (1957) anatomical estimate of the cone positions. He simply poked small holes in the paper at the cone positions in one of Polyak's anatomical drawings. We can place any image we like, for example patterns of light and dark bars, behind the grid. The bits of the image that we see are only those that would be seen by the visual system. Any pair of images that differ only in the regions between the holes will be an aliased pair. Yellott introduced the method and proper analysis, but he used Polyak's (1957) data on the outer segment positions rather than on the positions of the inner segments (Miller and Bernard, 1983).

This experiment is relatively straightforward for the S cones. Since these cones are separated by about 10 minutes of visual angle, there are about six S cones per degree of visual angle. Hence, their Nyquist frequency is 3 cycles per degree of visual angle (cpd). It is possible to correct for chromatic aberration and to present spatial patterns at these low frequencies through the lens. Such experiments confirm the basic predictions that we will see aliased patterns (Williams and Collier, 1983).

1.4 The L and M Cone Mosaic

Experiments using a visual interferometer to image a high frequency pattern at high contrast on the retina are a powerful way to analyze the sampling mosaic of L and M cones. But, even before this was technically possible, Helmholtz' (1896) noticed that extremely fine patterns, looked at without any special apparatus, can appear wavy. He attributed this observation to sampling by the cone mosaic. His perception of a fine pattern and his graphical explanation of the waviness in terms of sampling by the cone mosaic are shown in part (a) of Figure 1.13 (boxed drawings).

G. Byram was the first to describe the appearance of high frequency interference gratings (Byram, 1944). His drawings of the appearance of these patterns are shown




Figure 1.13: *Drawings of perceived aliasing patterns* by several different observers. Helmholtz' observed aliasing of fine patterns which he drew in part H1. He offered an explanation of his observations, in terms of cone sampling, in H2. Byram's (1944) drawings of three interference patterns at 40, 85 and 150 cpd are labeled B1, B2, and B3. Drawings W1,W2 and W3 are by subjects in Williams' laboratory who drew their impression of aliasing of an 80 cpd and two patterns at 110 cpd

in part (b) of the figure. The image on the left shows the appearance of a low frequency pattern diffraction pattern. The apparent spatial frequency of this stimulus is faithful to the stimulus. Byram noted that as the spatial frequency increases towards 60 cpd, the pattern still appears to be a set of fine lines, but they are difficult to see (middle drawing). When the pattern significantly exceeds the Nyquist frequency, it becomes visible again but looks like the low-frequency pattern drawn on the right. Further, he reports that the pattern shimmers and is unstable, probably due to the motion of the pattern with respect to the cone mosaic.

Over the last 10 years D. Williams' group has replicated and extended these measurements using an improved visual interferometer. Their fundamental observations are consistent with both Helmholtz and Byram's reports, but greatly extend and quantify the earlier measurements. The two illustrations on the left of part (c) of Figure 1.13 show Williams' drawing of 80 cpd and 110 cpd sinusoidal gratings created on the retina using a visual interferometer. The third figure shows an artist's drawing of a 110 cpd grating. The drawing on the left covers a large portion of the visual field, and the appearance of the patterns varies across the visual field. For example, at 80 cpd the observer sees high contrast stripes at some positions, while the field appears uniform in other parts of the field. The appearance varies, but the stimulus itself is quite uniform. The variation in appearance is due to changes in the sampling density of the cone mosaic. Cone sampling density is lower in the periphery than in the central visual field, so aliasing begins at lower spatial frequencies in the periphery than in the central visual field. If we present a stimulus at a high enough spatial frequency we observe aliasing in the central and peripheral visual field, as the drawings of the 110 cpd patterns in Figure 1.13 show.

There are two extensions of these ideas on aliasing you should consider. First, the cone packing in the fovea occurs in two dimensions, of course, so that we must ask what the appearance of the aliasing will be at different orientations of the sinusoidal stimuli. As the images in Figure 1.12 show, the orientation of the low frequency alias does not correspond with the orientation of the input. By trying the demonstration yourself and rotating the sampling grid, you will see that the direction of motion of the alias does not correspond with the motion of the input stimulus². These kinds of aliasing confusions have also been reported using visual interferometry (Coletta and Williams, 1987).

Second, our analysis of foveal sampling has been based on some rather strict assumptions concerning the cone mosaic. We have assumed that the cones are all of the same type, that their spacing is perfectly uniform, and that they have very narrow sampling apertures. The general model presented in this chapter can be adapted if any one of these assumptions fails to hold true. As an exercise for yourself, a new analysis with altered assumptions might change the properties of

²Use the Postscript program in the appendix section to print out a grid and a fine pattern and try this experiment.

the sampling matrix.

Visual Interferometry: Measurements of Human Optics

There is one last idea you should take away from this chapter: Using interferometry, we can estimate the quality of the optics of the eye.

Suppose we ask an observer to set the contrast of a sinusoidal grating, imaged using normal incoherent light. The observer's sensitivity to the target will depend on the contrast reduction at the optics and the observer's neural sensitivity to the target. Now, suppose that we create the same sinusoidal pattern using an interferometer. The interferometric stimulus bypasses the contrast reduction due to the optics. In this second experiment, then, the observer's sensitivity is limited only by the observer's neural sensitivity. Hence, the sensitivity difference between these two experiments is an estimate of the loss due to the optics.

The visual interferometric method of measuring the quality of the optics has been used on several occasions. While the interferometric estimates are similar to estimates using reflections from the eye, they do differ somewhat. The difference is shown in Figure ?? which includes the Westheimer's estimate of the modulation transfer function, created by fitting data from reflections, along with data and a modulation transfer function obtained from interferometric measurements. The current consensus is that the optical modulation transfer function is somewhat closer to the visual interferometric measurements than the reflection measurements. The reasons for the differences are discussed in several papers (e.g. Campbell and Green, 1965; Williams 1985; Williams et al., 1995).

1.5 Summary and Discussion

The S cones are present at a much lower sampling density, and they are absent in the very center of the fovea. Because they are sparse, we can measure the S cone positions behaviorally using small points of light. The behavioral estimates of the S cones are also consistent with anatomical estimates of the S cone spacing.

The wide spacing of the S cones can be understood in terms of the chromatic aberration of the eye. The eye is ordinarily in focus for the middle-wavelength part of the visual spectrum, and there is very little contrast beyond 2-3 cycles per degree in the short-wavelength part of the spectrum. The sparse S cone spacing is matched to the poor quality of the retinal image in the short-wavelength portion of the spectrum.

The L and M cones are tightly packed in the central fovea, forming a triangular grid

that efficiently samples the retinal image. Ordinarily, optical defocus protects us from aliasing in the fovea. Once aliasing between two signals occurs, the confusion cannot be undone. The two signals have created precisely the same spatial pattern of photopigment absorptions; hence, no subsequent processing, through cone to cone interactions or later neural interpolation, can undo the confusion. The optical defocus prevents high spatial frequencies that might alias from being imaged on the retina.

By creating stimuli with a visual interferometer, we bypass the optical defocus and image patterns at very high spatial frequencies on the cone mosaic. From the aliasing properties of these patterns, we can deduce some of the properties of the L and M cone mosaics. The aliasing demonstrations show that the foveal sampling grid is regular and contains approximately 120 cones per degree of visual angle. These measurements, in the living human eye, are consistent with the anatomical images obtained of the human eye reported by Curcio and her colleagues (Curcio, et al., 1991).

The precise arrangement of L and M cones within the human retina is unknown, though data on this point should arrive shortly (e.g., Bowmaker and Mollon, 1993). Current behavioral estimates of the relative number of L and M cones suggest that there are about twice as many L cones as M cones (Cicerone and Nerger, 1989).

The cone sampling grid becomes more coarse and irregular outside the fovea where rods and other cells enter the spaces between the cones. In these portions of the retina, high frequency patterns presented through interferometry no longer appear as regular low frequency frequency patterns. Rather, because of the disarray in the cone spacing, the high frequency patterns appear to be mottled noise. In the periphery, the cone spacing falls off rapidly enough so that it should be possible to observe aliasing without the use of an interferometer (Yellott, 1982).

In analyzing photoreceptor sampling, we have ignored eye movements. In principle, the variation in receptor intensities during these small eye movements can provide information to permit us to discriminate between the alias pairs. (You can check this effect by studying the images you observe when you experiment with the sampling grids.) The effects of eye movements are often minimized in experiments by flashing the targets briefly. But, even when one examines the interferometric pattern for substantial amounts of time, the aliasing persists. The information available from small eye movements could be very useful; but, the analysis assuming a static eye offers a good account of current empirical measurements. This suggests that the nervous system does not integrate information across minute eye movements to improve visual resolution (Packer and Williams, 1992).

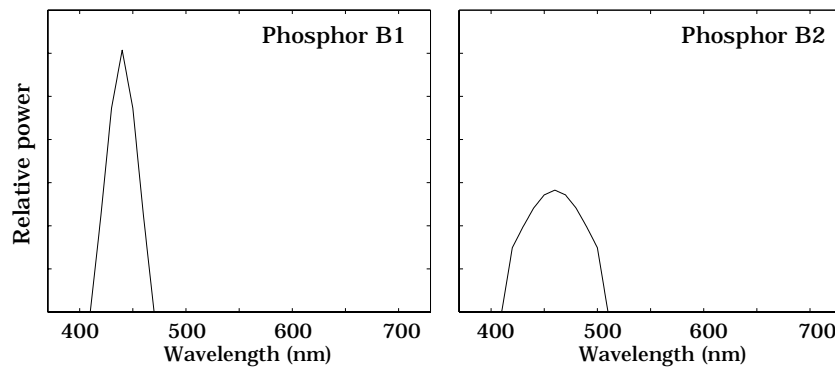


Figure 1.14: Choosing monitor phosphors.

Exercises

1. Answer the following questions related to image properties on the retina.
 - (a) Use a diagram to explain why the retinal image does not change size when the pupil changes size
 - (b) Compute the visual angle swept out by a building that is 200 meters tall seen from a distance of 400 meters.
 - (c) Suppose a lens has a focal length of 100mm. Where will the image plane of a line one meter from the center of the lens be? Suppose the line is 5 mm high. Using a picture, show the size of the image.
 - (d) Use the lensmaker's equation (from Chapter ??) to calculate the actual height on the retina.
 - (e) Good quality printers generate output with 600 dots per inch. How many dots is that per degree of visual angle. (Assume that the usual reading distance is 12 inches.)
 - (f) Good quality monitors have approximately 1000 pixels on a single line. How many pixels is that per degree of visual angle. (Assume that the usual monitor distance is 0.4 meters and the width of a line is 0.2 meters.)
 - (g) Some monitors can only turn individual pixels on or off. It may be fair to compare such monitors with the printed page since most black and white printers can only place a dot or not place one at each location. But, it is not fair to compare printer output with monitors capable of generating different gray scale levels. Explain how gray scale levels can improve the accuracy of reproduction without increasing the number of pixels. Justify your answer using a matrix-tableau argument.

2. A manufacturer is choosing between two different blue phosphors in a display (B1 or B2). The relative energy at different wavelengths of the two phosphors are shown in Figure 1.14. Ordinarily, users will be in focus for the red and green phosphors (not shown in the graph) around 580 nm.
 - (a) Based on chromatic aberration, which of the two blue phosphors will yield a sharper retinal image? Why?
 - (b) If the peak phosphor values are 400 nm and 450 nm, what will be the highest spatial frequency imaged on the retina by each of the two phosphors? (Use the curves in Figure ??.)
 - (c) Given the highest frequency imaged at 450 nm, what is the Nyquist sampling rate required to estimate the blue phosphor image? What is the Nyquist sampling rate for a 400 nm light source?
 - (d) The eye's optics images light at wavelengths above 500 nm much better than wavelengths below that level. Using the curves in Figure 1.3, explain whether you think the S cones will have a problem due to aliasing those longer wavelengths.
 - (e) (Challenge). Suppose the eye is always in focus for 580 nm light. The quality of the image created by the blue phosphor will always be quite poor. Describe how you can design a new layout for the blue phosphor mosaic on the screen to take advantage of the poor short-wavelength resolution of the eye. Remember, you only need to match images after optical defocus.
3. Reason from physiology to behavior and back to answer the following questions.
 - (a) Based purely on the physiological evidence from procion yellow stains, is there any reason to believe that the cones in Figure 1.7 are the S cones?
 - (b) What evidence do we have that the measurements of Williams et al. are due to the positions of the S cones rather than from the spacing of neural units in the visual pathways that are sensitive to short-wavelength light?
4. Give a drawing or an explanation to each of the following questions on aliasing.
 - (a) Draw an example of aliasing for a set of sampling points that are evenly spaced, but do not use a sinusoidal input pattern.
 - (b) Consider the sensor sample positions in Figure 1.15. with the positions unevenly spaced, as shown. Draw the response of this system to a constant valued input signal.




Figure 1.15: Sample positions of a set of sensors.

- (c) Now, draw a picture of a stimulus that is non-uniform and that yields the same response as in the previous question.
 - (d) What rule do you use to make sure the stimuli yield equivalent responses?
 - (e) Suppose that we put a lens that strongly defocuses the stimuli prior to their arrival at the sensor positions. This defocus means that it will be impossible to generate patterns that vary rapidly across space. If this blur is introduced into the optical path, will you be able to deliver your stimulus to the sensor array? Explain.
 - (f) Suppose that somebody asks you to invest in a company. The main product is a convolution operation that is applied to the output of a digital discrete sensor array built into a still camera. The purpose of the filter is to eliminate aliasing due to the sensors spatial sampling. How much would you be willing to invest in the company?
5. Perform the following aliasing calculations
- (a) In this chapter I asserted that $\cos(2\pi(N/2 + f)/N) = \cos(2\pi(N/2 - f)/N)$. Multiply out the arguments of the functions and write them both in the form of $\cos(i + j)$.
 - (b) Use the trigonometric identity $\cos(i - j) = \cos(i)\cos(j) + \sin(i)\sin(j)$ to expand the two functions.
 - (c) What is the value of $\sin(\pi)$? What is the value of $\cos(\pi)$? Use these values to obtain the final equality.
 - (d) Suppose that we represent a signal using a vector with ten entries. Suppose the signal is sampled at five locations, and we describe the sampling operation using a sampling matrix consisting of zeros and ones. How many rows and columns would the sampling matrix have?
 - (e) Write out the sampling matrix for a one-dimensional sampling pattern whose sample positions are at 1,3,5,7,9.
 - (f) Write out the sampling matrix for a non-uniform, one-dimensional pattern in which the sample positions are spaced at locations 1,2,4,and 8.
6. Answer each of the following questions about the relationship between the sampling mosaic and optics of the eye.

- (a) From time to time, some investigators have thought that the long-wavelength photopigment peak was near 620 nm, not 580 nm. Using Figure ??, discuss what implication such a peak wavelength would have for the Nyquist sampling rate required of these receptors.
 - (b) In fact, as you can see from Figure 1.3, the M and L cones both have peak sensitivities in the range near 550 nm to 580 nm. What is required of their spacing in order to accurately capture the retinal image?
 - (c) We have been assuming that the sensors in our array are equally sensitive to the incoming signal. Suppose that we have a sensor array that consists of alternating S and L cones. Draw the response of this array to an uniform field consisting of 450 nm light. Now, draw the intensity pattern that would have the same affect when the light is 650 nm.
7. Here are two Postscript programs, written by Arturo Puente, to create squarewave patterns and sampling patterns. Use the programs to print out the grids and patterns, and then copy the printouts onto an overhead transparency. View the patterns through the grids to see the effects of aliasing.

```

%!PS-Adobe-1.0
% Description:
% Parameters:
%
/widthx1 10 def % <- Value to change
/widthx2 2 def % <- Value to change

/x1 90 def
/y1 180 def
/x2 { x1 widthx1 add } def
/y2 612 def
/numx {432 widthx1 widthx2 add idiv} def

/form
{
  newpath
  x1 y1 moveto
  x1 y2 lineto
  x2 y2 lineto
  x2 y1 lineto
  closepath
  fill
  /x1 x1 widthx1 widthx2 add add def
} def

```

```

numx {form} repeat

/Helvetica findfont 10 scalefont setfont
72 84 moveto
(Width solid lines =) show
widthx1 dup 3 string cvs show
72 72 moveto
(Width white lines =) show
widthx2 dup 3 string cvs show
showpage

%!PS-Adobe-1.0
% Description:
% Parameters:
%
/widthx 20 def % <- Value to change
/widthy 10 def % <- Value to change
/wall 1 def % <- Value to change

/x1 90 def
/y1 180 def
/x2 { x1 widthx add } def
/y2 { y1 widthx add } def
/numx {432 widthx idiv} def
/numy {432 widthy idiv} def

/form {
  newpath
  x1 y1 moveto
  x2 y1 lineto
  x2 y2 lineto
  x1 y2 lineto
  x1 y1 lineto
  x1 wall add y1 wall add lineto
  x1 wall add y2 wall sub lineto
  x2 wall sub y2 wall sub lineto
  x2 wall sub y1 wall add lineto
  x1 wall add y1 wall add lineto
  closepath
  fill
/x1 x1 widthx add def
} def

```

```
/fileform
{
  numx {form} repeat
  /x1 90 def
  /y1 y1 widthy add def
} def

numy {fileform} repeat

%
% Write out parameters
%
/Helvetica findfont 10 scalefont setfont
72 96 moveto
(Width of the rectangle in the x-axis = ) show
widthx dup 3 string cvs show
72 84 moveto
(Width of the rectangle in the y-axis = ) show
widthy dup 3 string cvs show
72 72 moveto
(Thickness of the wall = ) show
wall dup 3 string cvs show

showpage
%
```