

Neural Processing, Graphemes, and Tuned Receptors

We now consider the same problem from a neurological perspective. Triesman and others claim that preattentive processing is due to *early* visual processing. What is the neurological evidence for this?

Visual information leaves the retina, passing up the optic nerve, through the neural junction at the lateral geniculate nucleus (LGN), and on to the much richer world of processing in the cortex. The first areas in the cortex to receive visual inputs are called, simply, visual area 1 (V1) and visual area 2 (V2). Most of the output from area 1 goes on to area 2, and together these two regions make up more than 40% of vision processing (Lennie, 1998). There is plenty of neural processing power, as several billion neurons in areas V1 and V2 are devoted to analyzing the signals from only two million nerve fibers coming from the optic nerves of two eyes. This makes possible the massively parallel simultaneous processing of the incoming signals for color, motion, texture, and the elements of form. It is here that the elementary vocabularies of both vision and data display are defined.

Figure 5.11 is derived from Livingston and Hubel's diagram (1988) that summarizes both the neural architecture and the features processed in this area of the brain. A key concept in understanding this diagram is the *tuned receptive field*. In Chapter 3, we saw how single-cell recordings of cells in the retina and the LGN reveal cells with distinctive concentric receptive fields. Such cells are said to be *tuned* to a particular pattern of a white spot surrounded by black or a black spot surrounded by white. In general, a tuned filter is a device that responds strongly to a certain kind of pattern and responds much less, or not at all, to other patterns. In the early visual cortex, some cells respond only to elongated blobs with a particular position and orientation, others respond most strongly to blobs of a particular position moving in a particular direction at a particular velocity, and still others respond selectively to color.

There are cells in V1 and V2 that are differentially tuned to each of the following properties:

- Orientation and size (with luminance) via the Gabor processor described later in this chapter
- Color (two types of signal) via the opponent processing channel mechanisms discussed in Chapter 4
- Elements of local stereoscopic depth
- Elements of motion

Moreover, all these properties are extracted for each point in the visual field. In V1 and V2 and many other regions of the brain, neurons are arranged in the form of a spatial map of the retina. It is a highly distorted map, because the fovea is given more space than the periphery of vision. The receptive fields are smaller for cells that process information coming from the fovea than for cells that process information from peripheral regions of the visual field. Nevertheless, for every

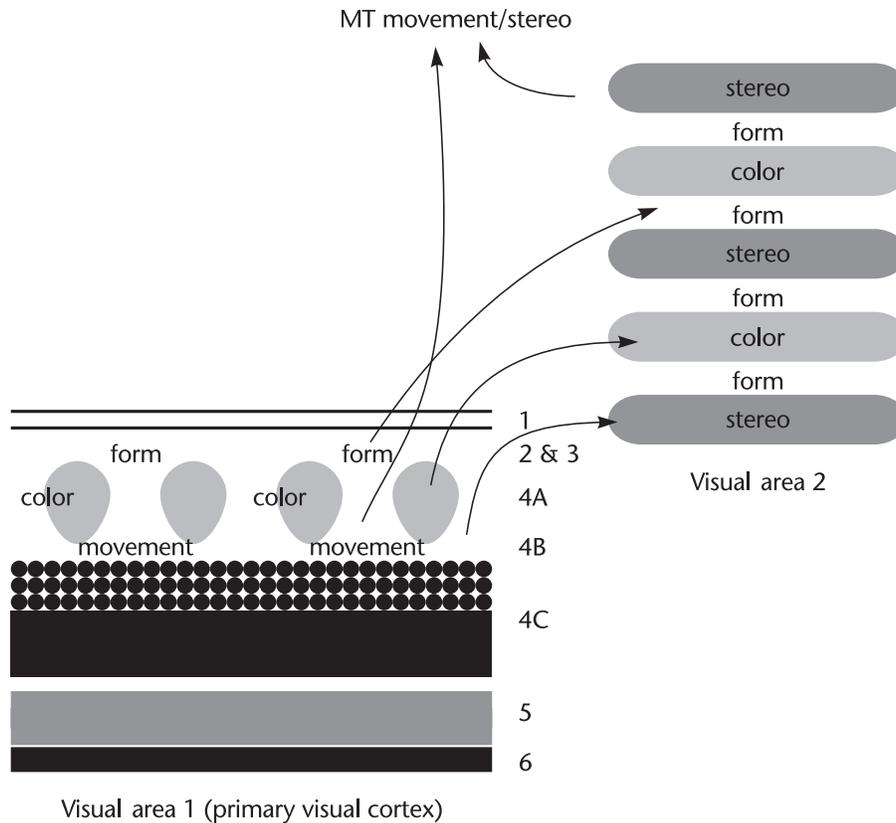


Figure 5.11 Architecture of primary visual areas. Adapted from Livingston and Hubel (1988).

point in V1, there is a corresponding area of the visual field in a topographic relationship (adjacency is preserved between areas). It is a massively parallel system in which, for each point in visual space, there are tuned filters for many different orientations, many different kinds of color information, many different directions and velocities of motion, and many different stereoscopic depths.

The Grapheme

It is useful to think of the things that are extracted by the early neural mechanisms as the “phonemes” of perception. *Phonemes* are the smallest elements in speech recognition, the atomic components from which meaningful words are made. In a similar way, we can think of orientation detectors, color detectors, and so on as “visual phonemes,” the elements from which meaningful perceptual objects are constructed.

We use the term *grapheme* to describe a graphical element that is primitive in *visual* terms, the visual equivalent of a phoneme. The basis of the grapheme concept is that the pattern that most efficiently excites a neuron in the visual system is exactly the pattern that the neuron is

tuned to detect (Ware and Knight, 1995). Thus, the most efficient grapheme is one that matches the receptive field properties of some class of neurons. An orientation detector will be excited most efficiently by a pattern whose light distribution is exactly the same as the sensitivity distribution of the cell. This is simply another way of saying that the detector is tuned to that particular pattern. Once we understand the kinds of patterns the tuned cells of the visual cortex respond to best, we can apply this information to create efficient visual patterns. Patterns based on the receptive field properties of neurons should be rapidly detected and easily distinguished.

A number of assumptions are implicit in this account. They are worth examining critically. One basic assumption is that the rate at which single neurons fire is the key coding variable in terms of human perception. This assumption can certainly be questioned. It may be that what is important is the way in which groups of neurons fire, or perhaps the temporal spacing or synchronization of cell firings. In fact, there is evidence that these alternative information codings may be important, perhaps critical. Nevertheless, few doubt that neurons that are highly sensitive to color differences (in terms of their firing rates) are directly involved in the processing of color and that the same thing is true for motion and shape. Moreover, as we shall see, the behavior of neurons fits well with studies of how people perceive certain kinds of patterns. Thus, there is a convergence of lines of evidence.

We also assume that *early-stage* neurons are particularly important in determining how distinct things seem. We know that at higher levels of processing in the visual cortex, receptive fields are found that are much more complex; they respond to patterns that appear to be composites of the simple receptive field patterns found at earlier stages. The evidence suggests that composite patterns analyzed further up the visual processing chain, are not, in general, processed as rapidly. It seems natural, then, to think of early-stage processing as forming the graphemes, and of later-stage processing as forming the “words,” or objects, of perception.

Much of the preattentive processing work already discussed in this chapter can be regarded as providing experimental evidence of the nature of graphemes. The following sections apply the concept to the perception of visual texture and show how knowledge of early mechanisms enables us to create rules for textures that are visually distinct.

The Gabor Model and Texture in Visualization

A number of electrophysiological and psychophysical experiments show that visual areas 1 and 2 contain large arrays of neurons that filter for orientation and size information at each point in the visual field. These neurons have both a preferred orientation and a preferred size (they are said to have *spatial* and *orientation tuning*). These particular neurons are not color-coded; they respond to luminance changes only.

A simple mathematical model used widely to describe the receptive field properties of these neurons is the *Gabor function*. This function is illustrated in Figure 5.12. It consists of the product of a cosine wave grating and a gaussian. Roughly, this can be thought of as a kind of fuzzy bar detector. It has a clear orientation, and it has an excitatory center, flanked by

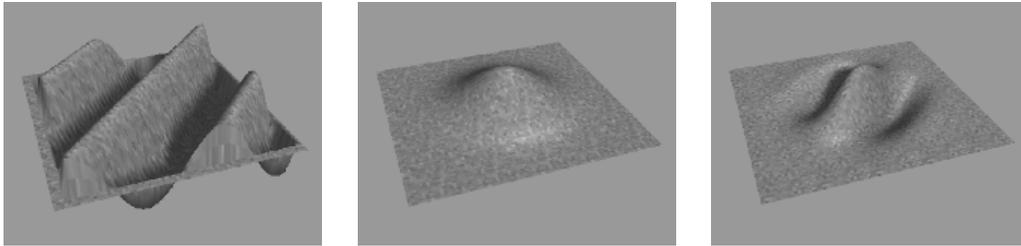


Figure 5.12 Gabor receptive field, composed of cosine and gaussian components. Multiply the cosine wave grating on the left by the gaussian envelope in the center to get the two-dimensional Gabor function shown on the right. This example has an excitatory center flanked by two inhibitory bars.

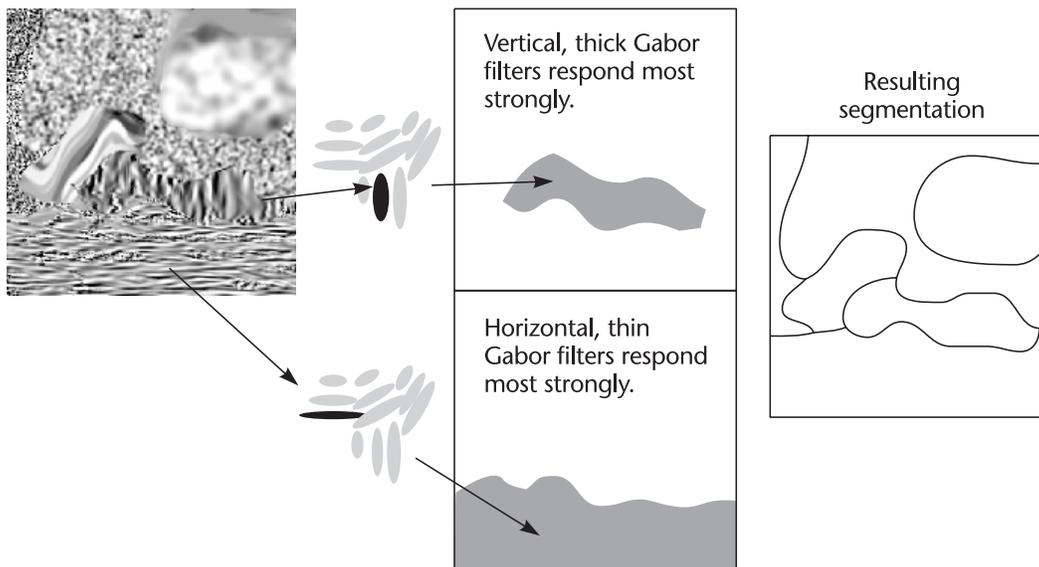


Figure 5.13 The texture segmentation model. Two-dimensional arrays of Gabor detectors filter every part of the image for all possible orientations and sizes. Areas exciting particular classes of detectors form the basis of visually distinct segments of the image.

inhibitory bars. The opposite kind of neuron also exists, with an inhibitory center and an excitatory surround.

Many things about low-level perception can be explained by this model. Gabor-type detectors are used in theories of the detection of contours at the boundaries of objects (form perception), the detection of regions that have different visual textures, stereoscopic vision, and motion perception.

The Gabor function has two components, as illustrated in Figure 5.12: a cosine wave and a gaussian envelope. Multiply them together, and the result is a Gabor function. Mathematically, a Gabor function has the following form (simplified for ease of explanation):

$$\text{Response} = C \cos\left(\frac{Ox}{S}\right) \exp\left(-\frac{(x^2 + y^2)}{S}\right) \quad (5.1)$$

The C parameter gives the amplitude or *contrast* value; S gives the overall *size* of the Gabor function by adjusting both the wavelength of the cosine grating and the rate of decay of the gaussian envelope. O is a rotation matrix that *orients* the cosine wave. Other parameters can be added to position the function at a particular location in space and adjust the ratio of the gaussian size to the sine wavelength; however, orientation, size, and contrast are most significant in modeling human visual processing.

Texture Segmentation

One way to apply the Gabor model is in understanding how the visual system *segments* the visual world into regions of distinct visual texture. Suppose we wish to understand how people perceptually differentiate types of vegetation based on the visual textures in a black-and-white satellite image. A model based on Gabor filters provides a good description of the way people perform this kind of texture segmentation task (Bovik et al., 1990; Malik and Perona, 1990).

The segmentation model is illustrated in Figure 5.13. It has three main stages. In the first stage, banks of Gabor filters respond strongly to regions of texture where particular spatial frequencies and orientations predominate. In a later stage, the output from this early stage is low-pass-filtered. (This is a kind of averaging process that creates regions, each having the same general characteristic. At the final stage, the boundaries are identified between regions with strongly dissimilar characteristics.) This model predicts that we will divide visual space into regions according to the predominant spatial frequency and orientation information. A region with large orientation and size differences will be the most differentiated. Also, regions can be differentiated based on the texture contrast. A low-contrast texture will be differentiated from a high-contrast texture with the same orientation and size components.

Tradeoffs in Information Density: An Uncertainty Principle

A famous vision researcher, Horace Barlow, developed a set of principles that have become influential in guiding our understanding of human perception. The second of these, called “the second dogma” (Barlow, 1972), provides an interesting theoretical background to the Gabor model. In the second dogma, Barlow asserted that the visual system is simultaneously optimized in both the spatial–location and spatial–frequency domains. John Daugman (1984) showed mathematically that Gabor detectors satisfy the requirements of the Barlow dogma. They optimally preserve a combination of spatial information (the location of the information in visual space) and oriented-frequency information. A single Gabor detector can be thought of as being tuned to a little packet of orientation and size information that can be positioned anywhere in space.

Daugman (1985) has also shown that a fundamental uncertainty principle is related to the perception of position, orientation, and size. Given a fixed number of detectors, resolution of size can be traded for resolution of orientation or position. We have shown that same principle applies to the synthesis of texture for data display (Ware and Knight, 1995). A gain in the ability to display orientation information precisely inevitably comes at the expense of precision in displaying size information. Given a constant density of data, orientation or size can be specified precisely, but not both.

Figure 5.14 illustrates this tradeoff, expressed by changing the shape and size of the gaussian multiplier function with the same sinusoidal grating. When the gaussian is large, the spatial frequency is specified quite precisely, as shown by the small image in the Fourier transform. When the gaussian is small, position is well specified but spatial frequency is not, as shown by the large image in the Fourier transform. The lower two rows of Figure 5.14 show how the gaussian envelope can be stretched to specify either the spatial frequency or the orientation more precisely. Although a full mathematical treatment of these effects is beyond the scope of this book, the main point is that there are fundamental limits and tradeoffs related to the ways texture can be used for information display. To restate them simply, large display glyphs can only show position imprecisely; precise orientation can be shown at the expense of precise size information, and both trade off against precision in position.

Texture Coding Information

If texture perception can be modeled and understood using the Gabor function as a model of a detector, the same model should be useful in *producing* easily distinguished textures for information display. The ideal grapheme for generating visual textures will be the Gabor function expressed as a luminance profile, as shown in Figure 5.15. A neuron with a Gabor receptive field will respond most strongly to a Gabor pattern with the same size and orientation. Therefore, textures based on Gabor primitives should be easy to distinguish.

Primary Perceptual Dimensions of Texture

A completely general Gabor model has parameters related to orientation, spatial frequency, contrast, and the size and shape of the gaussian envelope. However, in human neural receptive fields, the gaussian and cosine components tend to be coupled so that low-frequency cosine components have large gaussians and high-frequency cosine components have small gaussians (Caelli and Moraglia, 1985). This allows us to propose a simple three-parameter model for the perception and generation of texture.

Orientation O: The orientation of the cosine component

Scale S: The size = $1/(\text{spatial frequency component})$

Contrast C: An amplitude or contrast component

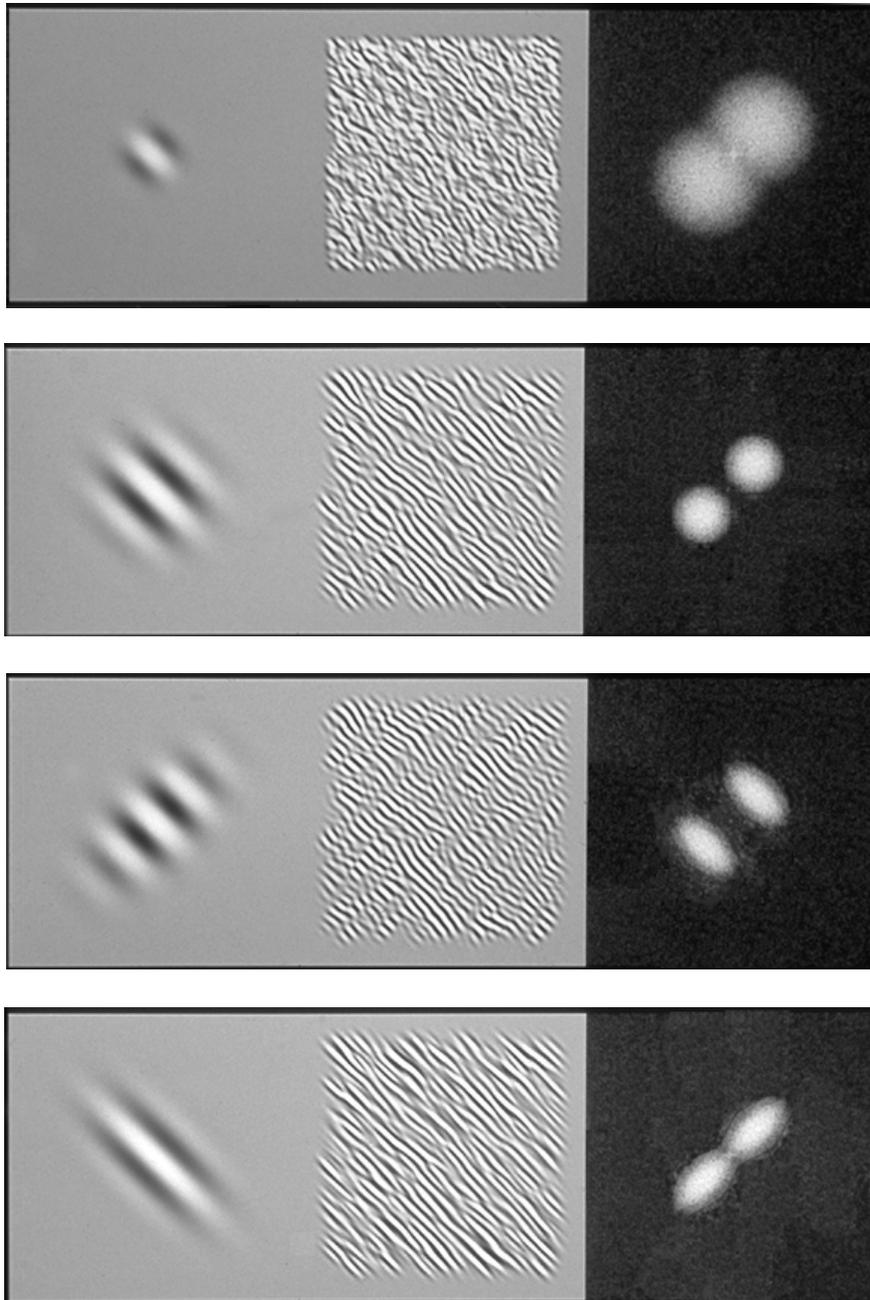


Figure 5.14 In the left-hand column, the same cosine pattern is paired with different gaussian multipliers. In the center column are textures created using each Gabor function by reducing the size by a factor of 5 and spattering it in the field. In the right-hand column are 2D Fourier transforms of the textures.

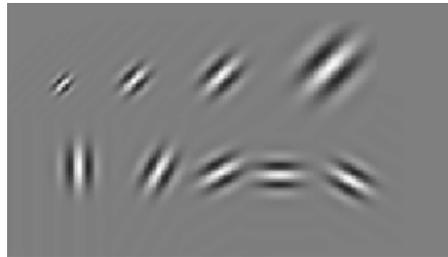


Figure 5.15 Gabor receptive fields shown as gray-scale images. Different sizes and orientations are represented for each part of the visual field.

Generation of Distinct Textures

With this simple model, it is straightforward to generate textures using Gabor functions as primitives. These textures can be varied in orientation, size ($1/\text{frequency}$), or contrast.

One method is to randomly splatter down Gabor functions whose orientation, size, and contrast have been determined by data values for the region in space where each splatter lands (Ware and Knight, 1995). When enough splatters have been accumulated in this way, we will have a continuous map that can represent up to three variables (a trivariate map). We can also map an additional variable to hue, producing a four-variable map.

Data value 1 → Orientation

Data value 2 → Size

Data value 3 → Contrast

Data value 4 → Hue

Figure 5.16 provides an example showing a magnetic field displayed using orientation and size manipulations. Color coding is added to the Gabor textures to illustrate field strength. A word of caution—Figure 5.16 illustrates a direct application of low-level visual theory, but it should not be taken as an optimal display. It is based on a feature-level model; to understand how to better show flow patterns, we need to move up the visual system and consider how patterns are formed from features. A more effective approach to vector field visualization, through pattern perception, is discussed in Chapter 6.

Note that textures need not be made of Gabor patterns for the method or the theory to work. It is only necessary that texture elements have a dominant orientation and spatial frequency. It is also important to note that the fundamental tradeoffs in our ability to represent spatial information using texture are independent of whether or not the Gabor model of texture perception is correct. To take a simple example, if we consider that texture elements, or *textons*, can be made from small graphical shapes representing data, the number of such shapes that can be drawn per unit area is inversely proportional to their size. The location of the packet of

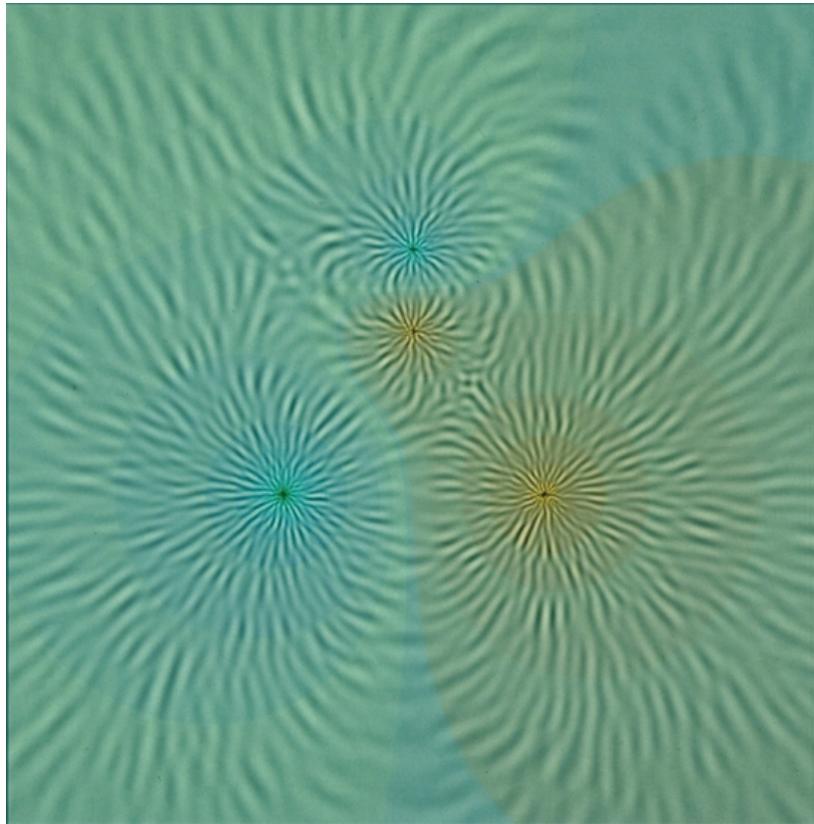


Figure 5.16 Magnetic field shown using Gabor textures.

information can be specified only to a precision determined by the size of the object representing that information.

Spatial-Frequency Channels, Orthogonality, and Maps

Sometimes we may wish to display many different kinds of information in a single map. For example, we might wish to show sea-surface temperature and sea-surface salinity at the same time. Naturally, we would prefer that the different sources of information do not interfere with one another. It would be unfortunate if regions of high salinity appeared to have a greater apparent temperature than they really have, due to visual crosstalk between the way we display temperature and the way we display salinity. Thus, our goal is to create display methods that are *perceptually independent*.

The concept of the *visual processing channel* can be taken directly from vision research and applied to the independence problem. We have already discussed the concept of color channels in Chapter 4. Here, the same idea is applied to spatial information. The idea is that information carried on one channel should not interfere with information displayed on another. It is

probably not the case that any of the perceptual channels we shall discuss are fully independent; nevertheless, it is certainly the case that some kinds of information are processed in ways that are more independent than others. A channel that is independent from another is said to be *orthogonal* to it. Here, the concept is applied to the spatial information carried by Gabor detectors.

A given Gabor-type neuron is broadly tuned with respect to orientation and size. The half-width of the spatial tuning curve is approximately a period change (in the sinusoid) of a factor of 3, and the total number of spatial-frequency channels is about four. Wilson and Bergen (1979) determined these values using a masking technique, which essentially determines the extent to which one type of information interferes with another. The resulting estimation of spatial-frequency channels is illustrated in Figure 5.17.

Orientation tuning-in appears to be about ± 30 degrees (Blake and Holopigan, 1985). Therefore, textures that differ from one another by more than 30 degrees in orientation will be easily distinguished.

These experimental results can be applied to problems in information display. For textured regions to be visually distinct, the dominant spatial frequencies should differ by at least a factor of 3 or 4, and the dominant orientations should differ by more than 30 degrees, all other factors (such as color) being equal. In general, the more displayed information differs in spatial frequency and in orientation, the more distinct that information will be. In practical applications, this means

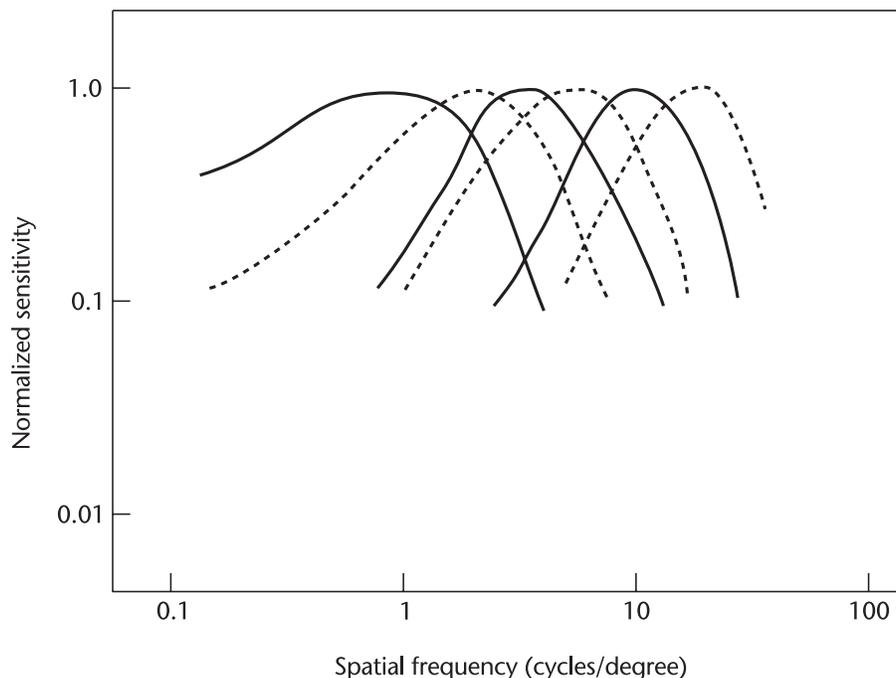


Figure 5.17 Wilson and Bergen spatial channels.

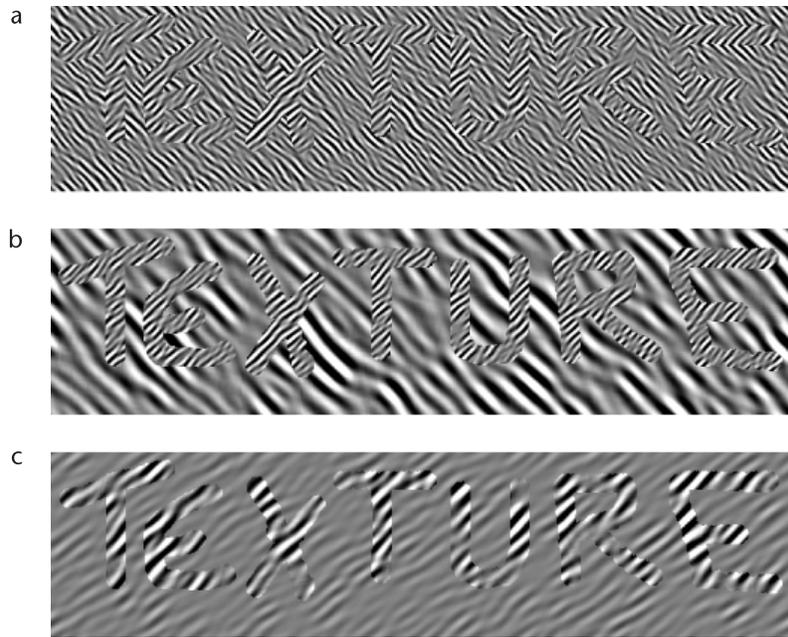


Figure 5.18 The word *TEXTURE* is visible only because of texture differences between the letters and the background; overall luminance is held constant. (a) Only texture orientation is altered. (b) Texture orientation and size are altered. (c) Texture contrast is altered.

that if we want different regions to be distinct because of their texture, the dominant orientations of the patterns should be made as different as possible. In Figure 5.18(a), only orientation is changed between different regions of the display, and although the word *TEXTURE* appears distinct from its background, it is weak. The difference appears much stronger when both the spatial frequency and the orientation differ between the figure and the background, as in Figure 5.18(b). The third way that textures can be made easy to distinguish is by changing the contrast, as illustrated in Figure 5.18(c).

Texture Resolution

The model of texture segmentation described previously predicts performance when people are asked to rapidly classify regions of a display. However, if we ask how small a difference people can *resolve*, we need a different model. When people are allowed to stare at two regions of a display for as long as they like, they can resolve far smaller differences than those perceived in brief presentations.

The *resolvable* size difference for a Gabor pattern is a size change of about 9% (Caelli et al., 1983). The resolvable orientation difference is about 5 degrees (Caelli and Bevan, 1983). These resolutions are much smaller than the channel-tuning functions would predict. This implies that higher-level mechanisms are present to sharpen up the output from individual receptors. The

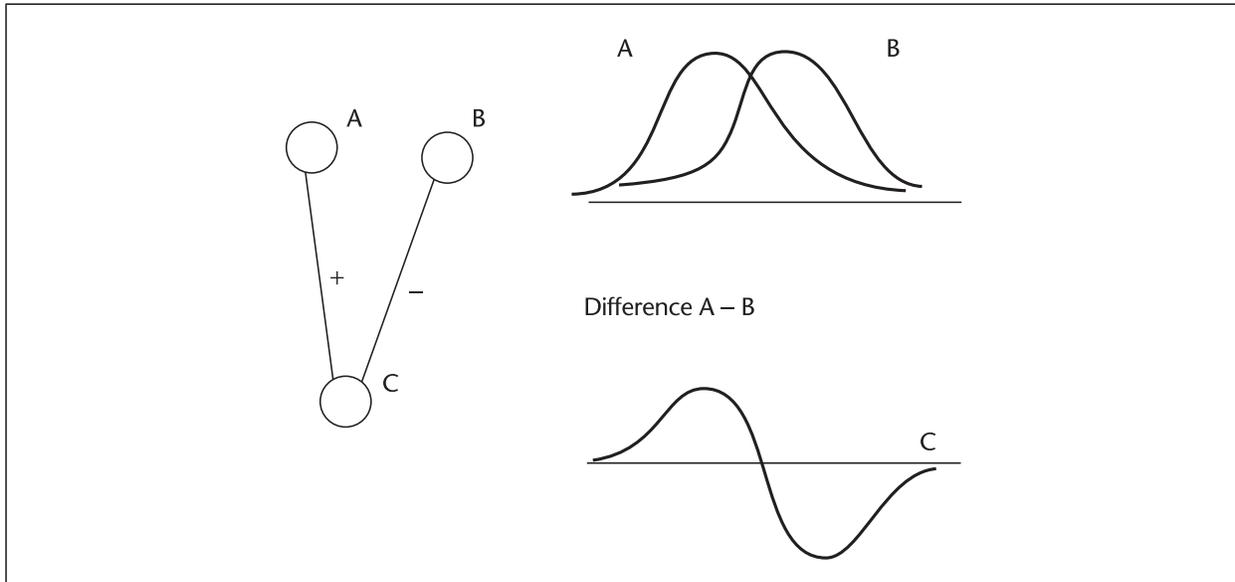


Figure 5.19 Differences between two signals are created by an excitatory and an inhibitory connection.

mechanism is based on inhibition. If a neuron has an excitatory input from one neuron and an inhibitory input from another with a slightly different tuning, the resulting difference signal is much more sensitive to spatial tuning than either of the original signals. This kind of sharpening is common in neural systems; it appears in color systems, edge detection, and heading detection (for navigation). Figure 5.19 illustrates the concept. Neurons A and B both have rather broadly tuned and somewhat overlapping response functions to some input pattern. Neuron C has an excitatory input from A and an inhibitory input from B. The result is that C is highly sensitive to differences between A and B at the crossover point.

Texture Contrast Effects

Textures can appear distorted because of contrast effects, just like the luminance contrast illusions that were described in Chapter 3. Thus, a given texture on a coarsely textured background will appear finer than the same texture on a finely textured background. This phenomenon is illustrated in Figure 5.20. The effect is predicted by higher-order inhibitory connections. It will cause errors in reading data that is mapped to texture element size. Texture orientation can cause contrast illusions in orientation, and this, too, may cause misperception of data. See Figure 5.21.

Other Dimensions of Visual Texture

Although there is considerable evidence to suggest that orientation, size, and contrast are the three dominant dimensions of visual texture, it is clear that the world of texture is much richer

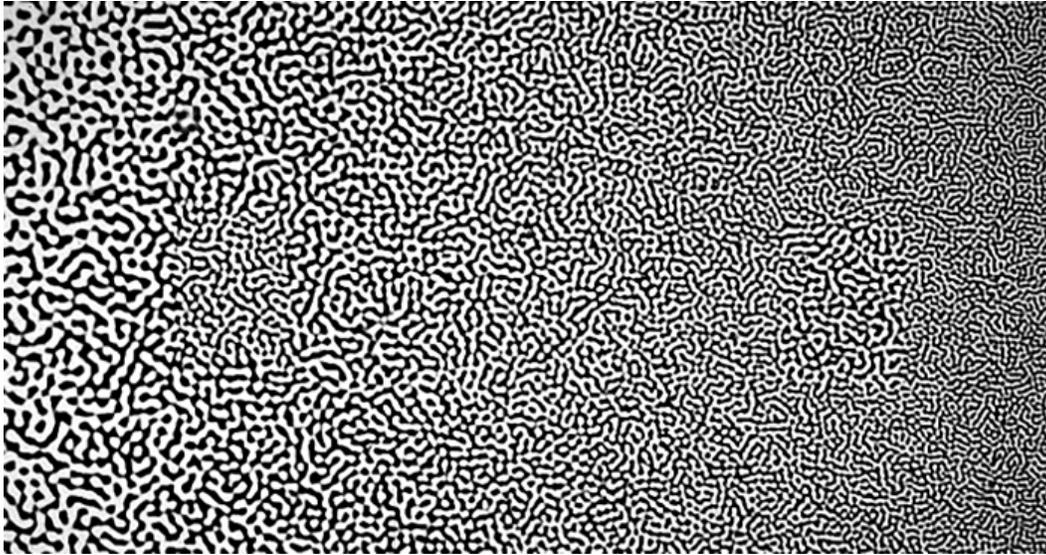


Figure 5.20 Texture contrast effect. The two patches left of center and right of center have the same texture granularity, but texture contrast makes them appear different.

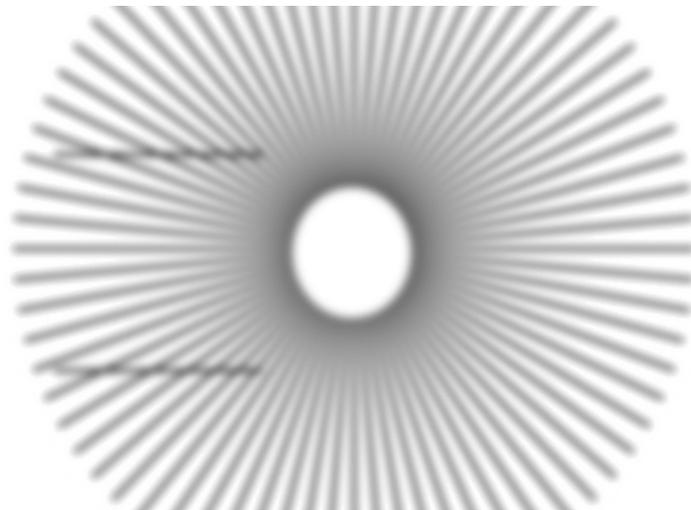


Figure 5.21 The two horizontal lines are parallel, but they appear tilted due to orientation contrast.

than this. The dimensionality of visual texture is very high, as a visual examination of the world around us attests. Think of the textures of wood, brick, stone, fur, leather, and other natural materials. One of the important additional texture dimensions is certainly randomness (Liu and Picard, 1994). Textures that are regular have a very different quality from random ones.

Texture Field Displays

We would do well to learn to use texture more effectively in information displays. The world of visual texture is arguably as rich and expressive as the world of color. Users of GISs commonly require the display of many overlapping variables on the same map, such as geological information, vegetation type, surface topography, and magnetic anomalies. In light of the theory of parallel feature processing, we are now in a position to say something about various solutions that apply visual texture to such problems.

The Exvis tool (Pickett and Grinstein, 1988) mapped data values to various attributes of stick-figure icons such as those shown in Figure 5.22. This package had many display options, including changing the angles of line segments relative to each other, or relative to a reference orientation, and changing the line segment widths. These glyphs could then be displayed in a dense field over a plane producing a visual texture. Although the Exvis developers implemented the capability to map data to icon colors, they worked mostly with angles (Pickett et al., 1995).

What does early visual processing tell us about the Exvis glyph? The theory of visual texture segmentation based on low-level Gabor detectors suggests a problem. With the Exvis glyph, multiple segments of a single glyph can have the same or similar orientations, although each represents a different data dimension. These line segments will be visually confounded when the glyphs are densely displayed, ensuring that unrelated aspects of the underlying data will be visually confounded. Because the orientation tuning of V1 neurons indicates that glyph element orientations should be separated by at least 30 degrees, and because a line-oriented segment will be confused with an identical segment rotated through 180 degrees, fewer than six orientations can be rapidly distinguished.

Weigle et al. (2000) developed a technique called *oriented sliver textures* specifically designed to take advantage of the parallel processing of orientation information. Each variable in a multivariate map was mapped to a 2D array of slivers where all the slivers had the same orientation. Differently oriented 2D sliver arrays were produced for each variable. The values of each scalar map were shown by controlling the amount of contrast between the sliver and the background. Combining all of the sliver fields produced the visualization illustrated in Figure 5.23.

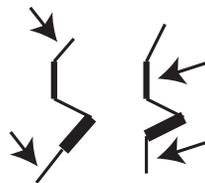


Figure 5.22 The Exvis data glyph used to form visual textures. Different variables are mapped to the angle between line segments and their thickness.

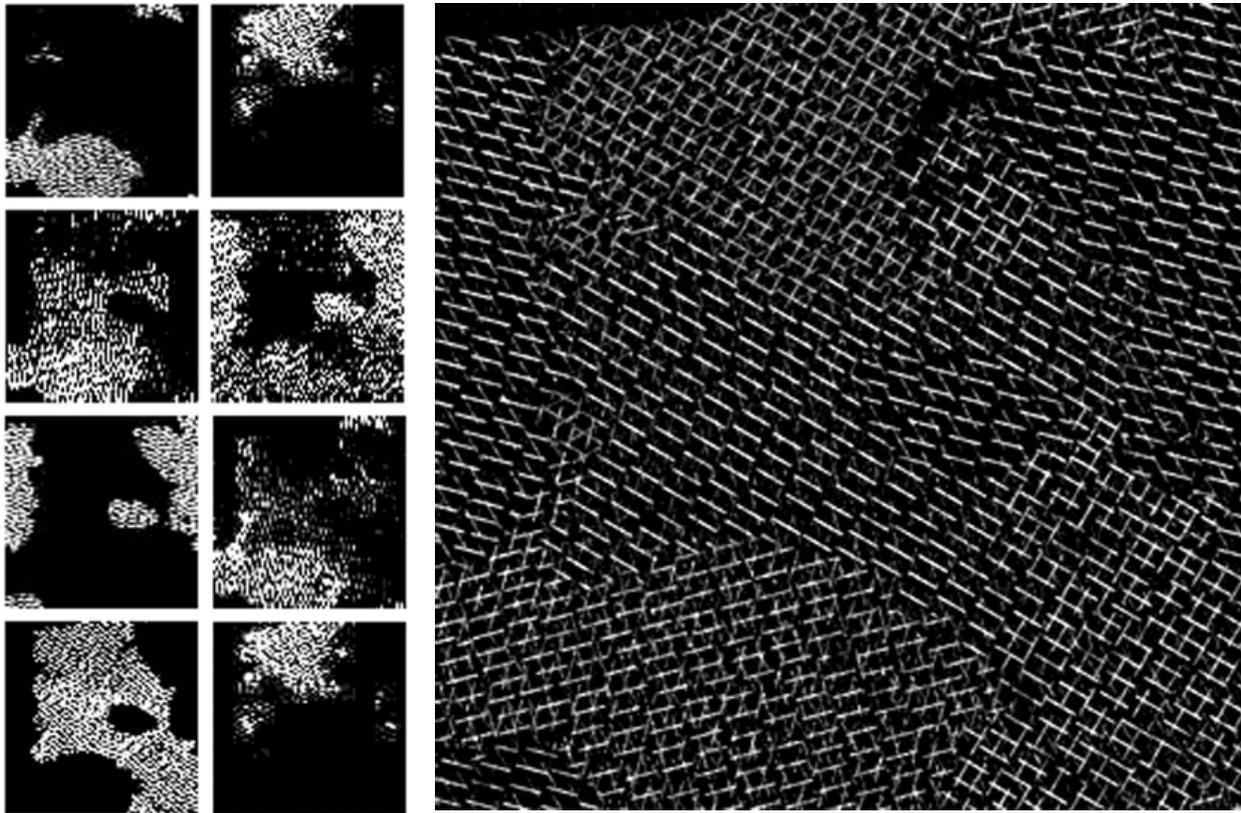


Figure 5.23 The sliver plot of Weigle et al. (2000). Each of the variables shown in the thumbnail patterns in the left part of the above figure is mapped to a differently oriented sliver field. *Courtesy of Chris Weigle.*

The right-hand part of this figure shows the combination of the eight variables illustrated in the thumbnail patterns shown on the left. Weigle et al. conducted a study showing that if slivers were oriented at least 15 degrees from surrounding regions, they stood out clearly. However, the experiment was only carried out with a single sliver at each location (unlike Figure 5.23). To judge the effectiveness of the sliver plot for yourself, try looking for each of the thumbnail patterns in the larger combined plot. The fact that many of the patterns cannot easily be seen suggests that the technique is not effective for so many variables. The tuning of orientation-sensitive cells suggests that slivers should be at least 30 degrees apart to be clearly readable (Blake and Holopigan, 1985), perhaps more, but in Figure 5.23 some differ by only 15 degrees.

Figure 5.24 shows another sliver plot with only three orientations. This adds a colored background and also uses slivers having both positive and negative contrast with the background. It is easier to see the different patterns in this example.

Two other examples of high-dimensional data display from Laidlaw and his collaborators (Laidlaw et al. 1998) (Figures 5.25 and 5.26) were created using a very different design strategy.

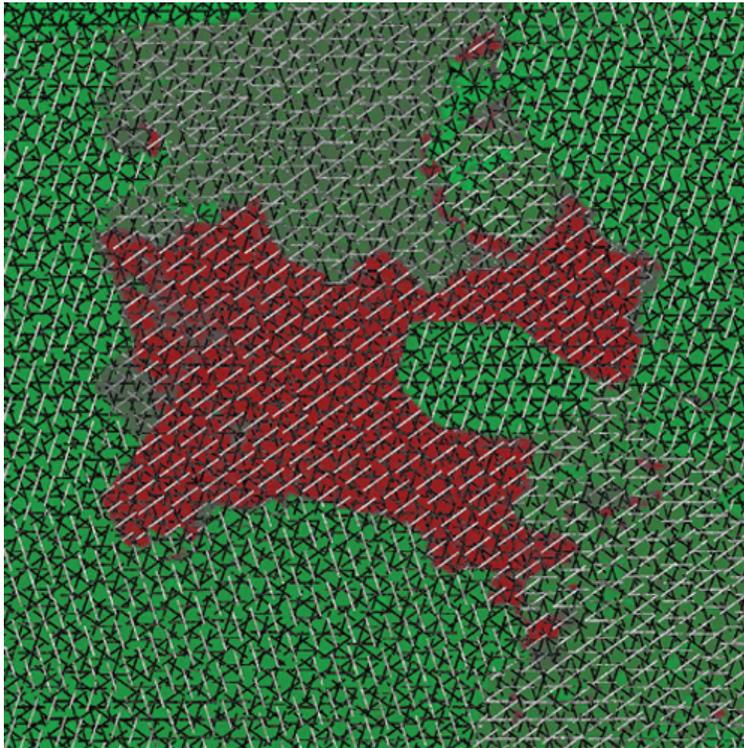


Figure 5.24 Another example of a sliver plot. Three variables are mapped to three differently oriented slivers. A fourth variable is mapped to the background color. *Courtesy of Chris Weigle.*

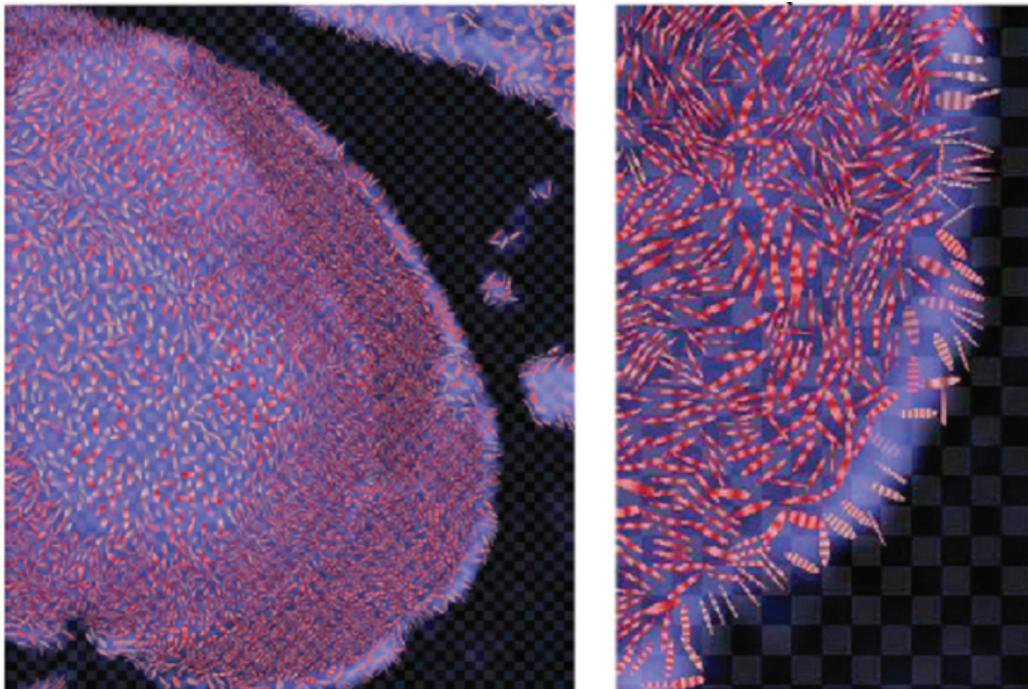


Figure 5.25 A cross section of a mouse spinal column. Seven variables are shown at each location. Part of the image is enlarged on the right. See text for description. *Courtesy of David Laidlaw.*

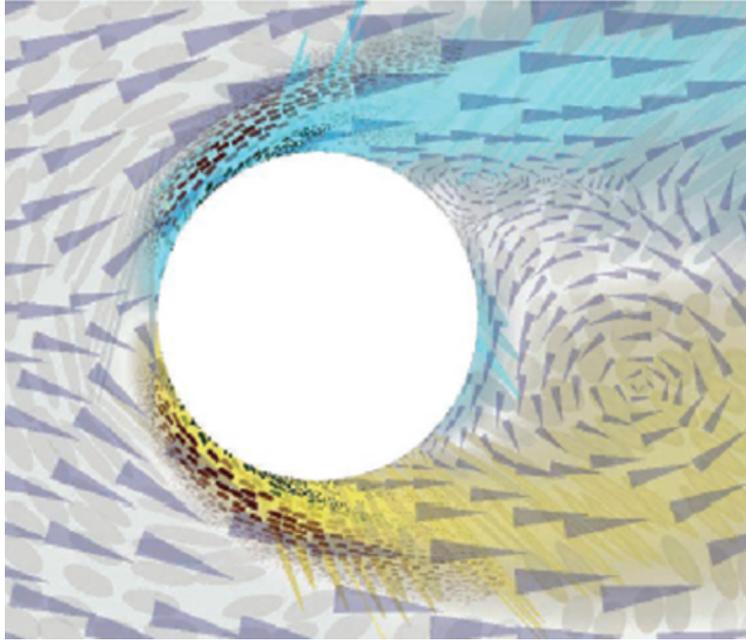


Figure 5.26 A flow visualization showing six variables relating to fluid flow around a cylinder. *Courtesy of David Laidlaw.*

Rather than attempting to create a simple general technique (like slivers), both figures were hand-crafted in a collaboration between the scientist and the designer. Figure 5.25 shows a cross section of a mouse spinal column. The data has seven values at each location in the image. The image is built up in layers: image intensity, sampling rate determines the grid, elliptical shapes show the in-plane component of principal diffusion and anisotropy, texture on the ellipses shows absolute diffusion rate.

The image in Figure 5.26 is a flow visualization. It displays six variables relating to the flow pattern of a fluid around a cylinder. These values are 2D velocity (two values are mapped to arrow direction and area), vorticity (one value is mapped to color and texture on ellipses), and deformation rate tensor (three values are mapped to shape and orientation of ellipses).

Without specific knowledge of mouse physiology or fluid dynamics, it is impossible to judge the success of these examples. Nevertheless, they provide a vivid commentary on the tradeoffs involved in trying to display high-dimensional multivariate maps. The first point to be made is that none of the preceding three examples (Figures 5.24, 5.25, and 5.26) shows much detail, and there is a good reason for this. We only have one luminance channel, and luminance variation is the only way of displaying detailed information. If we choose to use texture (or any kind of glyph field), we inevitably sacrifice the ability to show detail, because to be clear each glyph element must be displayed using luminance contrast. Larger glyphs mean that less detail can be shown.

There are also tradeoffs when displaying orientation. It may be only possible to display a single orientation clearly at each point in space for the purpose of showing flow patterns. Figure 5.26 suggests that if we need to show differently oriented glyphs in the same region, the glyphs must be widely spaced. This reduces the data density further. Also, Figure 5.26 suggests that the colors of different glyph layers must be very carefully chosen to be dissimilar. This, in turn, severely restricts how color-coding can be used on individual glyphs. In Figure 5.25, each of the elliptical glyphs is textured to display an additional variable. However, the texture striations are at right angles to the ellipse major axes. This camouflages the glyphs, making their orientation more difficult to see. The use of texture will inevitably tend to camouflage glyph shape; if the textures are oriented, the problem will be worse. In general, the more similar the spatial frequencies of the different pattern components, the more likely they are to disrupt one another visually.

The complexity of the design tradeoffs suggests that the problem of creating complex visualizations will be more of a craft than a science for quite some time. The problem is too difficult for automatic assignments of data maps to graphical attributes to be successful. Still, the designer needs to be aware of the perceptual tradeoffs in order to make informed decisions about the best choice of glyph size, shape, and color distribution.

It is also worth pointing out that there are some perceptual dimensions that may be used in addition to color, shape, and texture. In some cases, it is helpful to use stereoscopic depth and motion in displaying multidimensional data. Stereoscopic depth, especially if used with a high-resolution display, can undoubtedly help us perceptually segment data layers. So can motion. Making all of the points in a data layer move coherently, even by a small amount, may make it possible to visually attend to either the static layer or the moving layer (as shown by the possibility of preattentive conjunction search with motion).

Glyphs and Multivariate Discrete Data

In the previous section, we saw how texture could be used to represent continuous map data. In Chapter 4, it was shown that color could be used in a similar way. However, sometimes multivariate *discrete* data is the subject of interest. For example, a marketing specialist may have data for every person in a particular geographical area, including estimates of income, educational level, employment category, and location of residence. The marketer would like to see each person on a map in such a way that the concentrations of individuals with particular sets of attributes can easily be seen. In this way, neighborhoods to be blanketed with flyers might be selected most effectively.

To create a glyph, multiple data attributes are mapped in a systematic way to show the different aspects of the appearance of the graphical object. In the aforementioned marketing example, income might be mapped to the glyph's size, education level to its color, employment category to its shape, and geographic location to the x,y location where the glyph is plotted. All the previously discussed results relating to preattentive detection of size, orientation, and color-coding of data apply to the design of glyphs.