

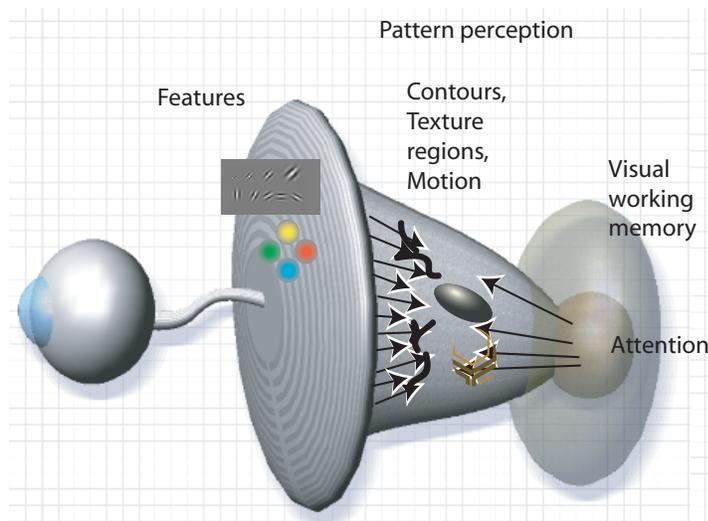
# Static and Moving Patterns

---

Data mining is about finding patterns that were previously unknown or that depart from the norm. The stock-market analyst looks for any pattern of variables that may predict a future change in price or earnings. The marketing analyst is interested in perceiving trends and patterns in a customer database. When we look for patterns, we are making visual queries that are key to visual thinking. Sometimes the queries are vague; we are on the lookout for a variety of structures in the data. Sometimes they are precise, as when we look for a positive trend in a graph. In data exploration, seeing a pattern can often lead to a key insight, and this is the most compelling reason for visualization.

What does it take for us to see a group? How can 2D space be divided into perceptually distinct regions? Under what conditions are two patterns recognized as similar? What constitutes a visual connection between objects? These are some of the perceptual questions addressed in this chapter. The answers are central to visualization, because most data displays are two-dimensional and pattern perception deals with the extraction of structure from 2D space.

Consider again our three-stage model of perception (illustrated in Figure 6.1). At the early stages of feature abstraction, the visual image is analyzed in terms of primitive elements of form, motion, color, and stereoscopic depth. At the next 2D pattern perception stage, the contours are discovered and the visual world is segmented into distinct regions, based on texture, color, motion, and contour. Next, the structures of objects and scenes are discovered, using information about the connections between component parts, shape-from-shading information, and so on. Pattern perception can be thought of as a set of mostly 2D processes occurring between feature analysis and full object perception, although aspects of 3D space perception, such as stereoscopic depth and structure-from-motion, can be considered particular kinds of pattern perception. Finally, objects and significant patterns are pulled out by attentional processes to meet the needs of the task at hand.



**Figure 6.1** Pattern perception forms a middle ground where the bottom-up processes of feature processing meet the requirements of active attention.

There are radical changes in the kinds of processing that occur at the different stages. In the early stages, massively parallel processing of the entire image occurs. This drives perception from the bottom up. But object and visual search recognition is driven from the top down through active attention, meeting the requirements of visual thinking. At the top level, only three to five objects (or patterns) are held in visual working memory. Pattern perception is the flexible middle ground where objects are extracted from patterns of features. Active processes of attention reach down into the pattern space to keep track of those objects and to analyze them for particular tasks; the essentially bottom-up processing of feature primitives meets the top-down processes of cognitive perception. Rensink (2000) calls the middle ground a “proto-object flux.”

Understanding pattern perception provides abstract design rules that can tell us much about how we should organize data so that important structures will be perceived. If we can map information structures to readily perceived patterns, then those structures will be more easily interpreted.

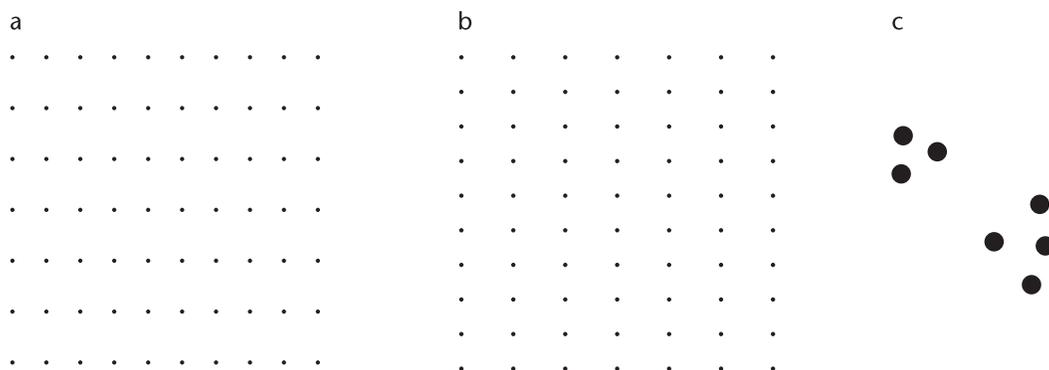
Learning is important in the pattern mechanism. It occurs in the short term through visual priming and in the long term as a kind of skill learning. *Priming* refers to the fact that once we have seen a pattern, it becomes much easier to identify on subsequent appearance. Long-term learning of patterns occurs over hundreds or thousands of trials, but some patterns are much easier to learn than others (Fine and Jacobs, 2002). In this chapter, we consider 2D-pattern perception and what this tells us about information display. In the next two chapters, we consider 3D-space perception, much of which is a form of advanced pattern perception.

## Gestalt Laws

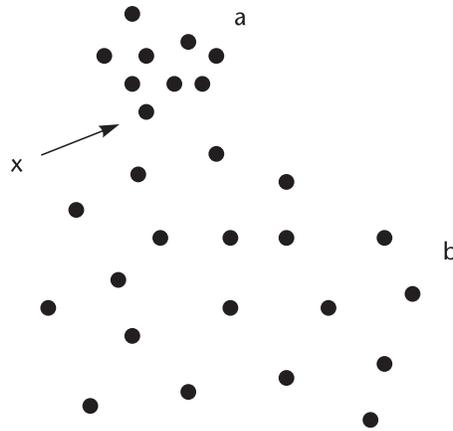
The first serious attempt to understand pattern perception was undertaken by a group of German psychologists who, in 1912, founded what is known as the Gestalt school of psychology. The group consisted principally of Max Westheimer, Kurt Koffka, and Wolfgang Kohler (see Koffka, 1935, for an original text). The word *gestalt* simply means *pattern* in German. The work of the Gestalt psychologists is still valued today because they provided a clear description of many basic perceptual phenomena. They produced a set of *Gestalt laws* of pattern perception. These are robust rules that describe the way we see patterns in visual displays, and although the neural mechanisms proposed by these researchers to explain the laws have not withstood the test of time, the laws themselves have proved to be of enduring value. The Gestalt laws easily translate into a set of design principles for information displays. Eight Gestalt laws are discussed here: proximity, similarity, connectedness, continuity, symmetry, closure, relative size, and common fate (the last concerns motion perception and appears later in the chapter).

### Proximity

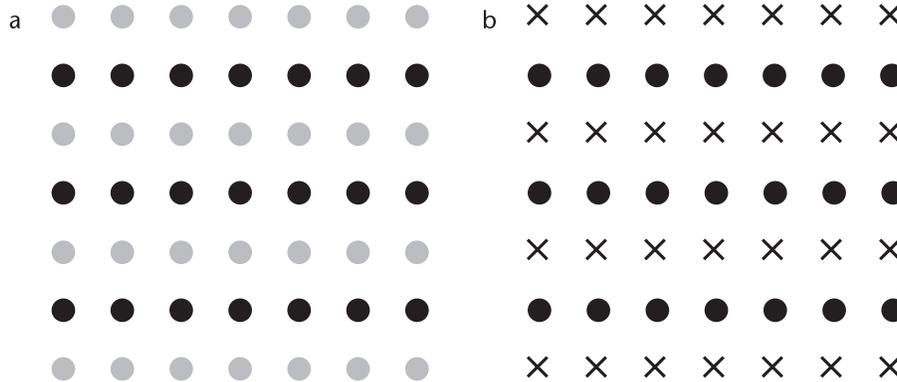
Spatial proximity is a powerful perceptual organizing principle and one of the most useful in design. Things that are close together are perceptually grouped together. Figure 6.2 shows two arrays of dots that illustrate the proximity principle. Only a small change in spacing causes us to change what is perceived from a set of rows, in Figure 6.2(a), to a set of columns, in Figure 6.2(b). In Figure 6.2(c), the existence of two groups is perceptually inescapable. Proximity is not the only factor in predicting perceived groups. In Figure 6.3, the dot labeled *x* is perceived to be part of cluster a rather than cluster b, even though it is as close to the other points in cluster b



**Figure 6.2** Spatial proximity is a powerful cue for perceptual organization. A matrix of dots is perceived as rows on the left (a) and columns on the right (b). In (c), because of proximity relationships, we perceive two groupings of dots.



**Figure 6.3** The principle of spatial concentration. The dot labeled x is perceived as part of group a rather than group b.



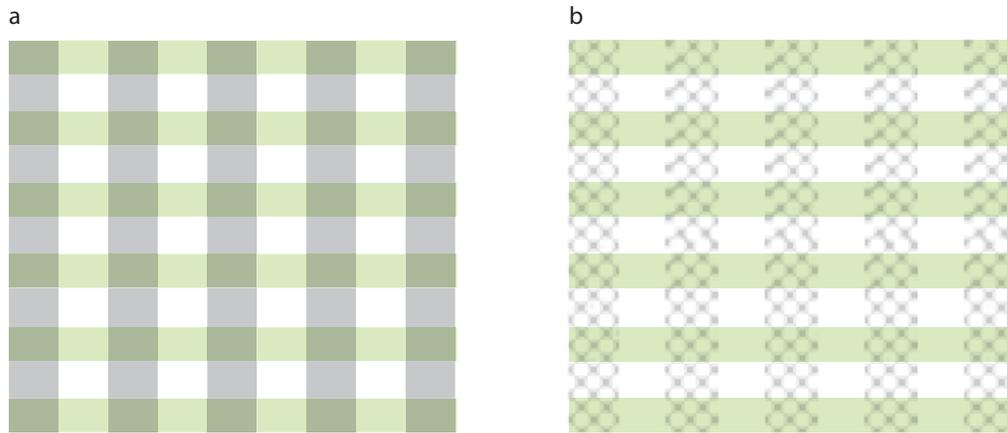
**Figure 6.4** According to the Gestalt psychologists, similarity between the elements in alternate rows causes the row percept to dominate.

as they are to each other. Slocum (1983) called this the *spatial concentration principle*. According to this principle, we perceptually group regions of similar element density.

The application of the proximity law in display design is straightforward: the simplest and most powerful way to emphasize the relationships between different data entities is to place them in proximity in a display.

## Similarity

The shapes of individual pattern elements can also determine how they are grouped. Similar elements tend to be grouped together. In both Figure 6.4(a) and (b), the similarity of the elements causes us to see the rows most clearly.



**Figure 6.5** (a) Integral dimensions are used to delineate rows and columns. (b) When separable dimensions (color and texture) are used, it is easier to attend separately to either the rows or the columns.

We can also apply lessons from the concept of integral and separable dimensions that was discussed in Chapter 5. Figure 6.5 shows two different ways of visually separating row and column information. In 6.5(a), integral color and gray-scale coding is used. In Figure 6.5(b), green color is used to delineate rows and texture is used to delineate columns. Color and texture are separable dimensions, and the result is a pattern that can be visually segmented either by rows or by columns. This technique can be useful if we are designing so that users can easily attend to either one pattern or the other.

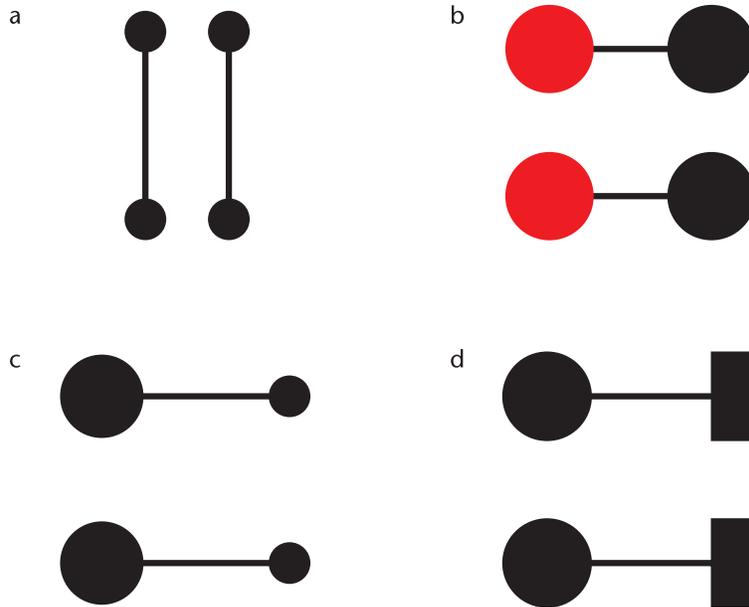
## Connectedness

Palmer and Rock (1994) argue that connectedness is a fundamental Gestalt organizing principle that the Gestalt psychologists overlooked. The demonstrations in Figure 6.6 show that connectedness can be a more powerful grouping principle than proximity, color, size, or shape. Connecting different graphical objects by lines is a very powerful way of expressing that there is some relationship between them. Indeed, this is fundamental to the node–link diagram, one of the most common methods of representing relationships between concepts.

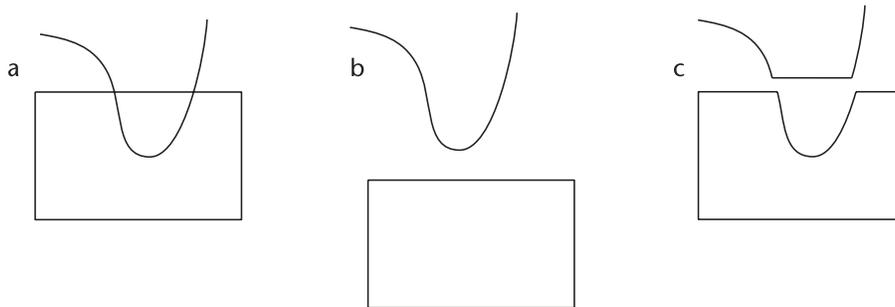
## Continuity

The Gestalt principle of continuity states that we are more likely to construct visual entities out of visual elements that are smooth and continuous, rather than ones that contain abrupt changes in direction. (See Figure 6.7.)

The principle of good continuity can be applied to the problem of drawing diagrams consisting of networks of nodes and the links between them. It should be easier to identify the sources and destinations of connecting lines if they are smooth and continuous. This point is illustrated in Figure 6.8.



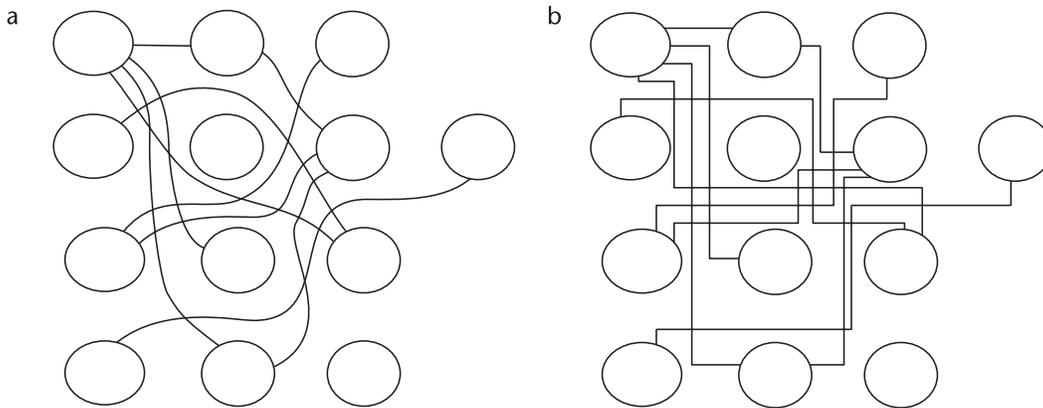
**Figure 6.6** Connectedness is a powerful grouping principle that is stronger than (a) proximity, (b) color, (c) size, or (d) shape.



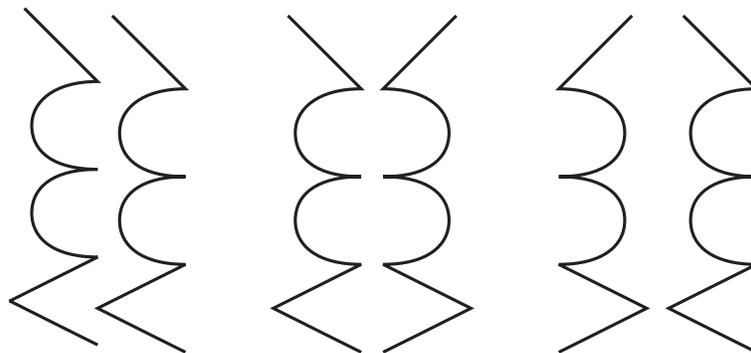
**Figure 6.7** The pattern on the left (a) is perceived as a curved line overlapping a rectangle (b) rather than as the more angular components shown in (c).

## Symmetry

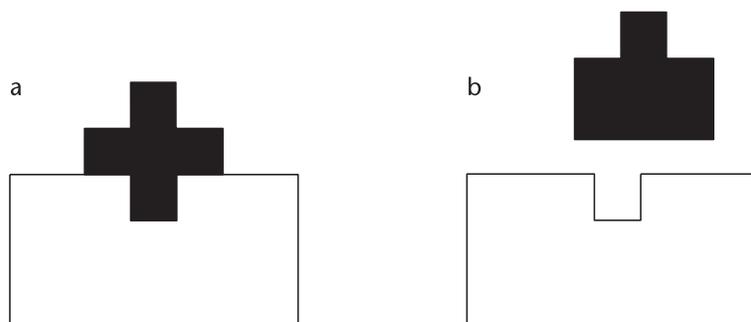
Symmetry can provide a powerful organizing principle. Figures 6.9 and 6.10 provide two examples. The symmetrically arranged pairs of lines in Figure 6.9 are perceived much more strongly as forming a visual whole than the pair of parallel lines. In Figure 6.10(a), symmetry may be the reason why the cross shape is perceived, as opposed to shapes in 6.10(b), even though the second option is not more complicated. A possible application of symmetry is in tasks in which data analysts are looking for similarities between two different sets of time-series data. It may be easier



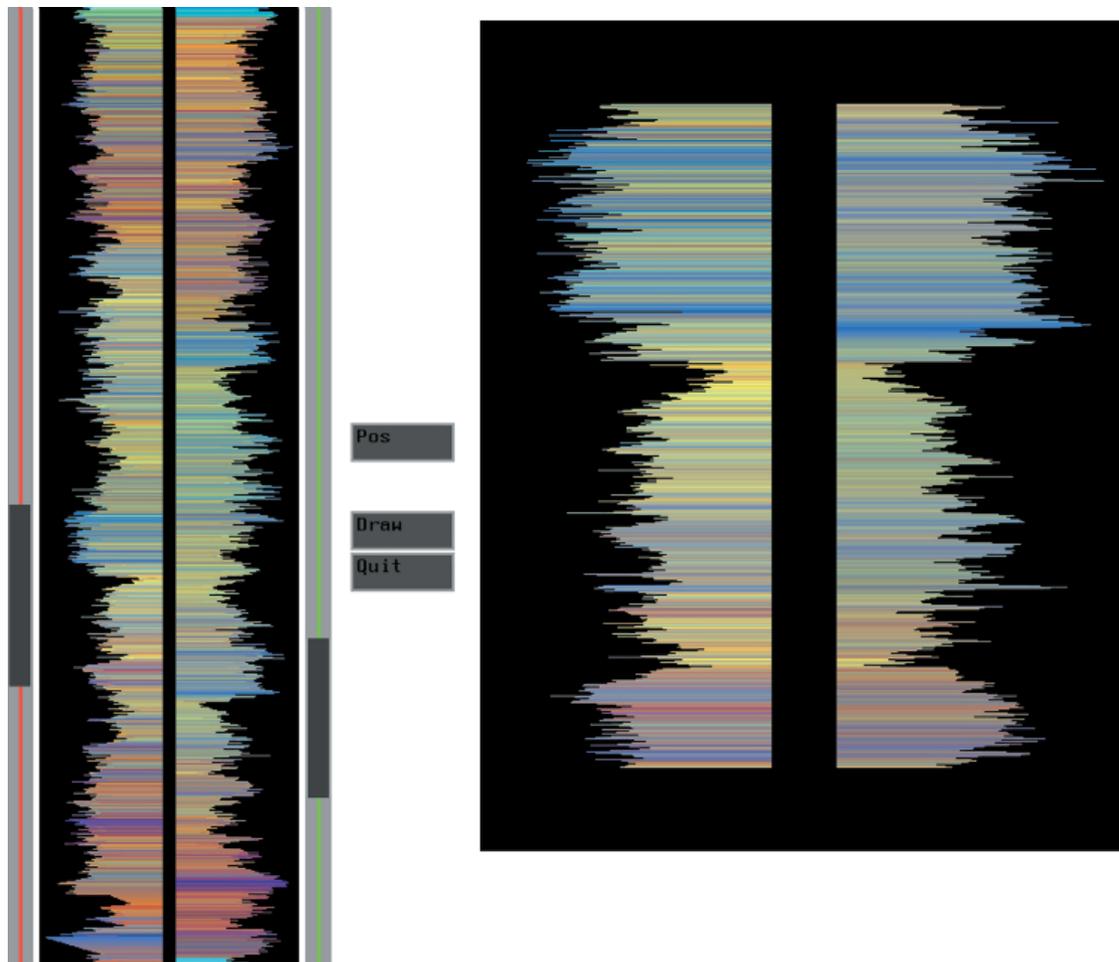
**Figure 6.8** In (a), smooth continuous contours are used to connect the elements, whereas in (b), lines with abrupt changes in direction are used. It is much easier to perceive connections when contours connect smoothly.



**Figure 6.9** The pattern on the left consists of two identical parallel contours. In each of the other two patterns, one of the contours has been reflected about a vertical axis, producing bilateral symmetry. The result is a much stronger sense of a holistic figure.



**Figure 6.10** We interpret pattern (a) as a cross in front of a rectangle. An alternative, two objects shown in (b) are not perceived, even though the black shape behind the white shape would be an equally simple interpretation. The cross on the rectangle interpretation has greater symmetry (about horizontal axes) for both of the components.

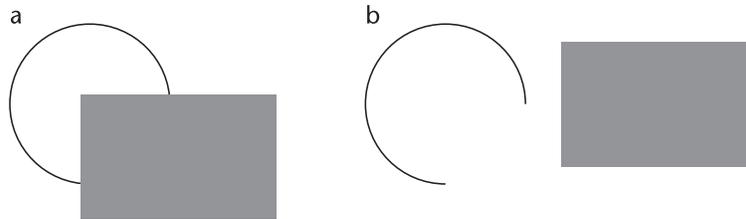


**Figure 6.11** An application designed to allow users to recognize similar patterns in different time-series plots. The data represents a sequence of measurements made on deep ocean drilling cores. Two subsets of the extended sequences are shown on the right.

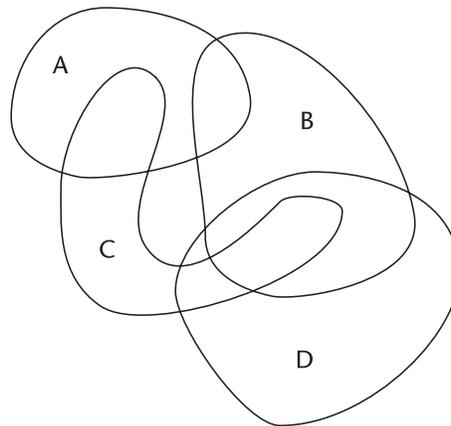
to perceive similarities if these time series are arranged using vertical symmetry, as shown in Figure 6.11, rather than using the more conventional parallel plots.

## Closure

A closed contour tends to be seen as an object. The Gestalt psychologists argued that there is a perceptual tendency to close contours that have gaps in them. This can help explain why we see Figure 6.12(a) as a complete circle and a rectangle rather than as a circle with a gap in it as in Figure 6.12(b).



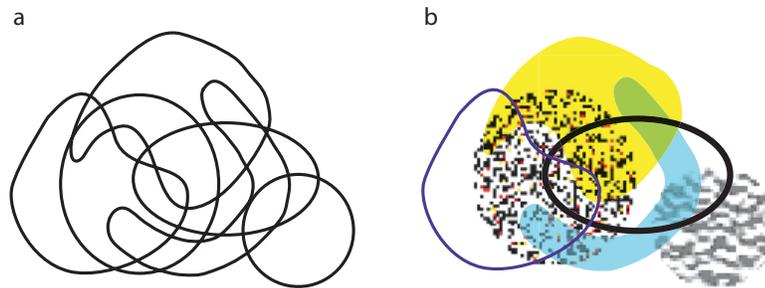
**Figure 6.12** The Gestalt principle of closure holds that neural mechanisms operate to find perceptual solutions involving closed contours. Hence in (a), we see a circle behind a rectangle, not a broken ring as in (b).



**Figure 6.13** An Euler diagram. This diagram tells us (among other things) that entities can simultaneously be members of sets A and C but not of A, B, and C. Also, anything that is a member of both B and C is also a member of D. These rather difficult concepts are clearly expressed and understood by means of closed contours.

Wherever a closed contour is seen, there is a very strong perceptual tendency to divide regions of space into “inside” or “outside” the contour. A region enclosed by a contour becomes a *common region* in the terminology of Palmer (1992). He showed common region to be a much stronger organizing principle than simple proximity. This, presumably, is the reason why Venn-Euler diagrams are such a powerful device for displaying the interrelationships among sets of data. In an Euler diagram, we interpret the region inside a closed contour as defining a set of elements. Multiple closed contours are used to delineate the overlapping relationships among different sets. A Venn diagram is a more restricted form of Euler diagram containing all possible regions of overlap. The two most important perceptual factors in this kind of diagram are closure and continuity.

A fairly complex structure of overlapping sets is illustrated in Figure 6.13, using an Euler diagram. This kind of diagram is almost always used in teaching introductory set theory, and this in itself is evidence for its effectiveness. Students easily understand the diagrams, and they



**Figure 6.14** An Euler diagram enhanced using texture and color can convey a more complex set of relations than a conventional Euler diagram using only closed contour.

can transfer this understanding to the more difficult formal notation. Stenning and Oberlander (1994) theorize that the ease with which Euler diagrams can be understood results specifically from the fact that they have limited expressive power, unlike fully abstract formal notation.

Although simple contours are generally used in Euler diagrams to show set membership, we can effectively define regions using color and texture as well, as discussed in Chapters 4 and 5. Indeed, by using both we should be able to create Euler diagrams that are considerably more complex and still readily understandable. Figure 6.14 illustrates.

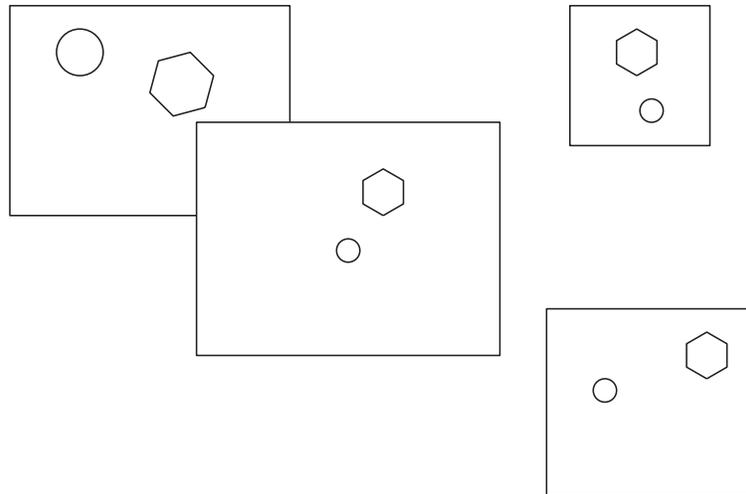
Closed contours are extremely important in segmenting the monitor screen in windows-based interfaces. The rectangular overlapping boxes provide a strong segmentation cue, dividing the display into different regions. In addition, rectangular frames provide frames of reference: the position of every object within the frame tends to be judged relative to the enclosing frame. (See Figure 6.15.)

## Relative Size

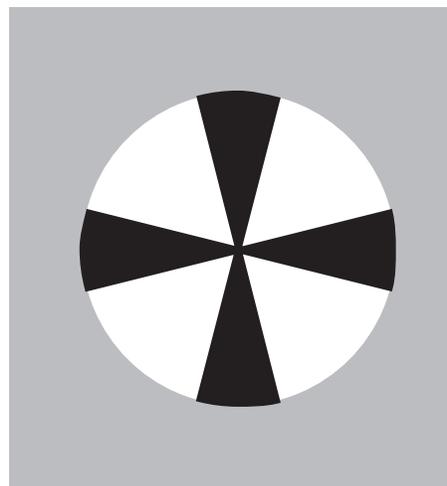
In general, smaller components of a pattern tend to be perceived as objects. In Figure 6.16, a black propeller is seen on a white background, as opposed to the white areas being perceived as objects.

## Figure and Ground

Gestalt psychologists were also interested in what they called *figure-ground* effects. A *figure* is something objectlike that is perceived as being in the foreground. The *ground* is whatever lies behind the figure. The perception of figure as opposed to ground can be thought of as the fundamental perceptual act of identifying objects. All the Gestalt laws contribute to creating a figure, along with other factors that the Gestalt psychologists did not consider, such as texture segmentation (see Chapter 5). Closed contour, symmetry, and the surrounding white area all con-



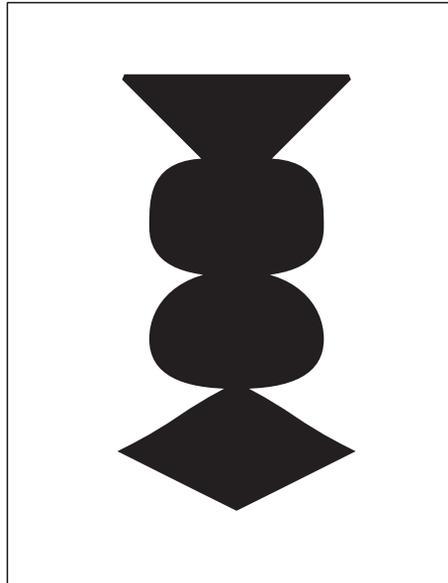
**Figure 6.15** Closed rectangular contours strongly segment the visual field. They also provide reference frames. Both the positions and the sizes of enclosed objects are, to some extent, interpreted with respect to the surrounding frame.



**Figure 6.16** The black areas are smaller, and therefore more likely to be perceived as an object. It is also easier to perceive patterns that are oriented horizontally and vertically as objects.

tribute to the perception of the shape in Figure 6.17 as figure, as opposed to a cut-out hole, for example.

Figure 6.18 shows the classic Rubin’s Vase figure, in which it is possible to perceive either two faces, nose to nose, or a black vase centered in the display. The fact that the two percepts tend to alternate suggests that competing active processes may be involved in trying to construct figures from the pattern.



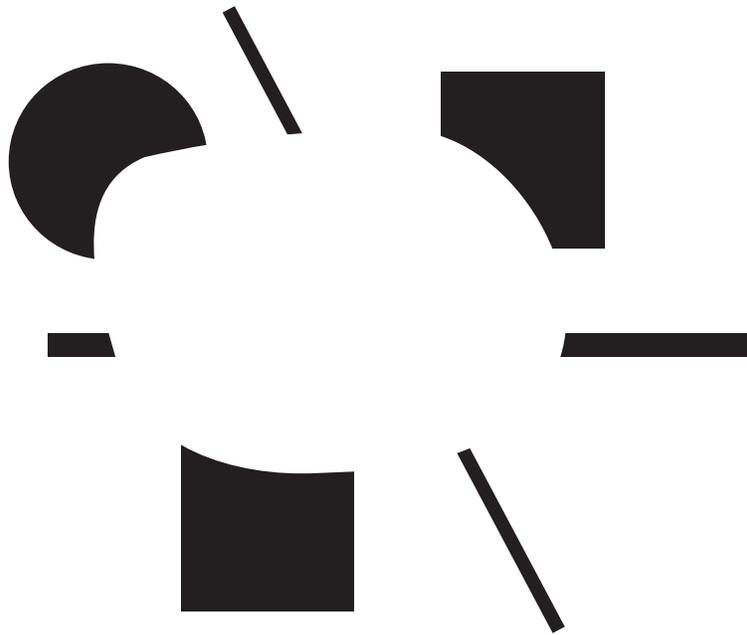
**Figure 6.17** Symmetry, surrounding white space, and a closed contour all contribute to the strong sense that this shape is figure, rather than ground.



**Figure 6.18** Rubin's Vase. The cues for figure and ground are roughly equally balanced, resulting in a bistable percept of either two faces or a vase.

## More on Contours

A contour is a continuous perceived boundary between regions of a visual image. A contour can be defined by a line, by a boundary between regions of different color, by stereoscopic depth, by motion patterns, or by texture. Contours can even be perceived where there are none. Figure 6.19 illustrates an *illusory contour*; a ghostly boundary of a blobby shape is seen even where



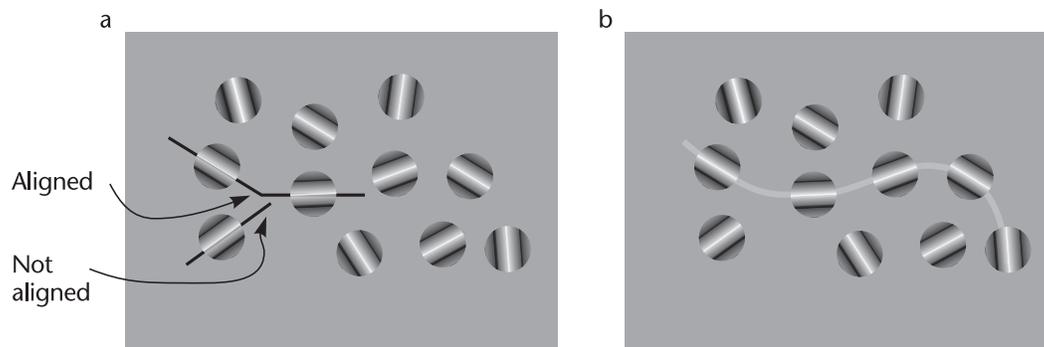
**Figure 6.19** Illusory contour.

none is physically present. There is extensive literature on illusory contours (see Kanizsa, 1976, for an early review).

Because the process that leads to the identification of contours is seen as fundamental to object perception, contour detection has received considerable attention from vision researchers. There are a number of detailed neurophysiological models designed to explain how contours can be extracted from the visual image, based on what is known about early visual processing. See Marr (1982), for example.

Higher-order neurophysiological mechanisms of contour perception are not well understood. However, one result is intriguing. Gray et al. (1989) found that cells with collinear receptive fields tend to fire in synchrony. Thus, we do not need to propose higher-order feature detectors, responding to more and more complex curves, to understand the neural encoding of contour information. Instead, it may be that groups of cells firing in synchrony is the way that the brain holds related pattern elements in mind. Theorists have suggested a fast enabling link, a kind of rapid feedback system, to achieve the firing of cells in synchrony. For a review, see Singer and Gray (1995).

Fortunately, because a theoretical understanding is only just emerging, the exact mechanisms involved in contour detection are less relevant to the purpose of designing visualizations than are the circumstances under which we perceive contours. A set of experiments by Field et al. (1993) places the Gestalt notion of *good continuation* on a firmer scientific basis. In these experiments, subjects had to detect the presence of a continuous path in a field of 256 randomly oriented Gabor patches (see Chapter 5 for a discussion of Gabor functions). The setup is illustrated



**Figure 6.20** A schematic diagram illustrating the experiments conducted by Field et al. (1993). If the elements were aligned as shown in (a) so that a smooth curve could be drawn through some of them, the curve shown in (b) was perceived. In the actual experiments, Gabor patches were used.

schematically in Figure 6.20. The results show that subjects were very good at perceiving a smooth path through a sequence of patches. As one might expect, continuity between Gabor patches oriented in straight lines was the easiest to perceive. More interesting, even quite wiggly paths were readily seen if the Gabor elements were aligned as shown in Figure 6.20(b).

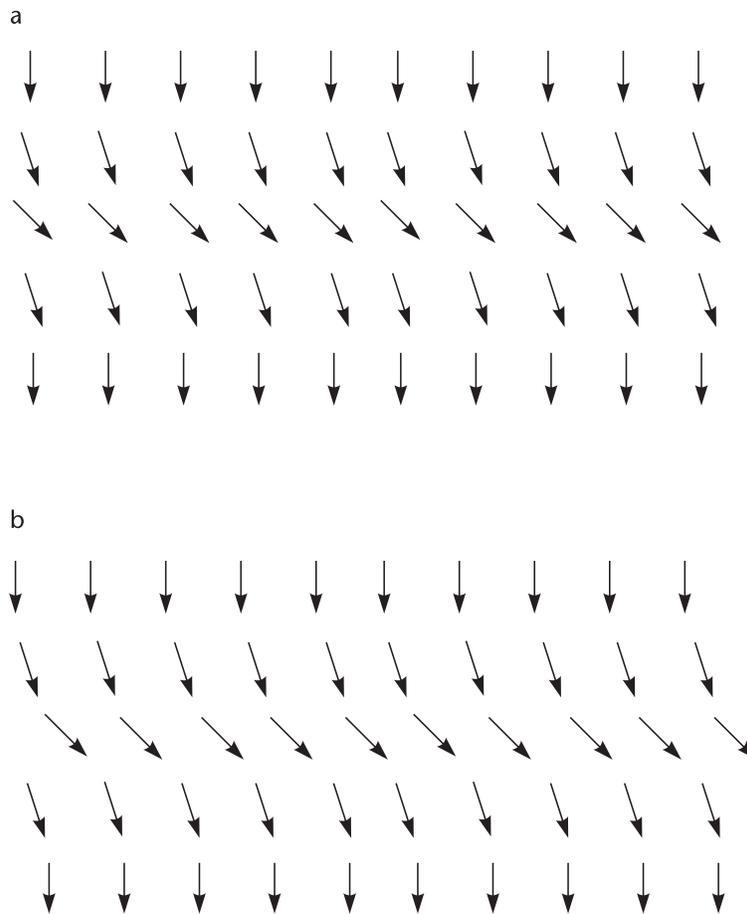
There are direct applications of this result in displaying vector field data. A common technique is to create a regular grid of oriented arrows, such as the one shown in Figure 6.21. When the arrows are displaced so that smooth contours can be drawn between them, the flow pattern is much easier to see.

## Perceiving Direction: Representing Vector Fields

The perception of contour leads us naturally to the perceptual problem of representing vector fields. This problem can be broken down into two components: the representation of orientation and the representation of magnitude. Some techniques display one component but not both.

Instead of using little arrows, one obvious and effective way of representing vector fields is through the use of continuous contours; a number of effective algorithms exist for this purpose. Figure 6.22 shows an example from Turk and Banks (1996). This effectively illustrates the direction of the vector field, although it is ambiguous in the sense that for a given contour there can be two directions of flow. Conventional arrowheads can be added, as in Figure 6.21, but the result is visual clutter. In addition, in Figure 6.22 the magnitudes of the vectors are given by line density and inverse line width, and this is not easy to read.

An interesting way to resolve the flow direction ambiguity is provided in a seventeenth-century vector field map of North Atlantic wind patterns by Edmund Halley (discussed in Tufte, 1983). Halley's elegant pen strokes, illustrated in Figure 6.23, are shaped like long, narrow airfoils oriented to the flow, with the wind direction given by the blunt end. Interestingly, Halley also arranges his strokes along streamlines. We verified experimentally that strokes like Halley's are unambiguously interpreted with regard to direction (Fowler and Ware, 1989).

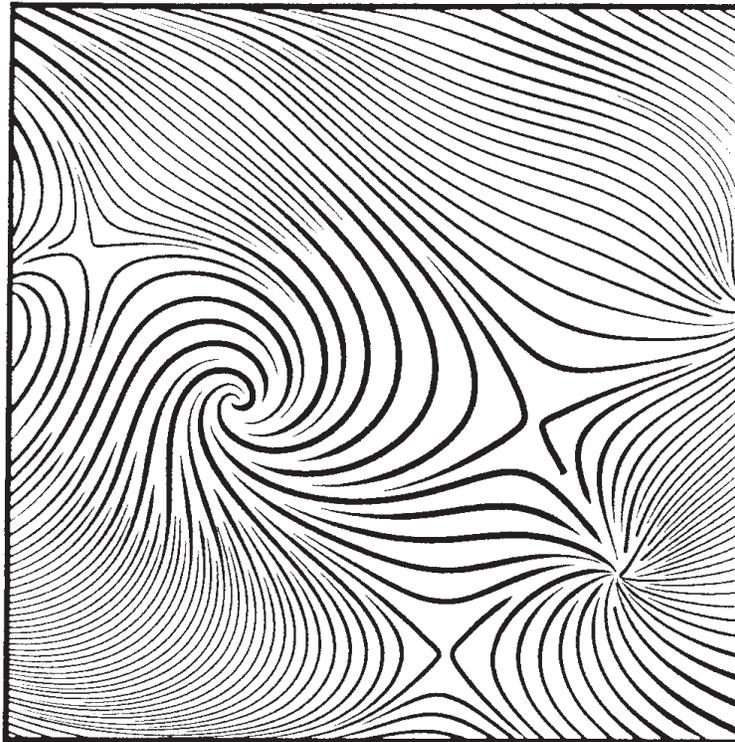


**Figure 6.21** The results of Field et al. (1993) suggest that vector fields should be easier to perceive if smooth contours can be drawn through the arrows. (a) A regular grid is used to determine arrow layout. (b) The arrows have been shifted so that smooth contours can be drawn through the arrows. As theory predicts, the latter is more effective.

We also developed a new method for creating an unambiguous sense of vector field direction that involves varying the color along the length of a stroke. This is illustrated in Figure 6.24. There was a strong interaction between the direction of color change and the background color. If one end of the stroke was given the background color, the stroke direction was perceived to be in the direction of color change away from the background color. In our experiments, the impression of direction produced by color change completely dominated that given by shape.

## Comparing 2D Flow Visualization Techniques

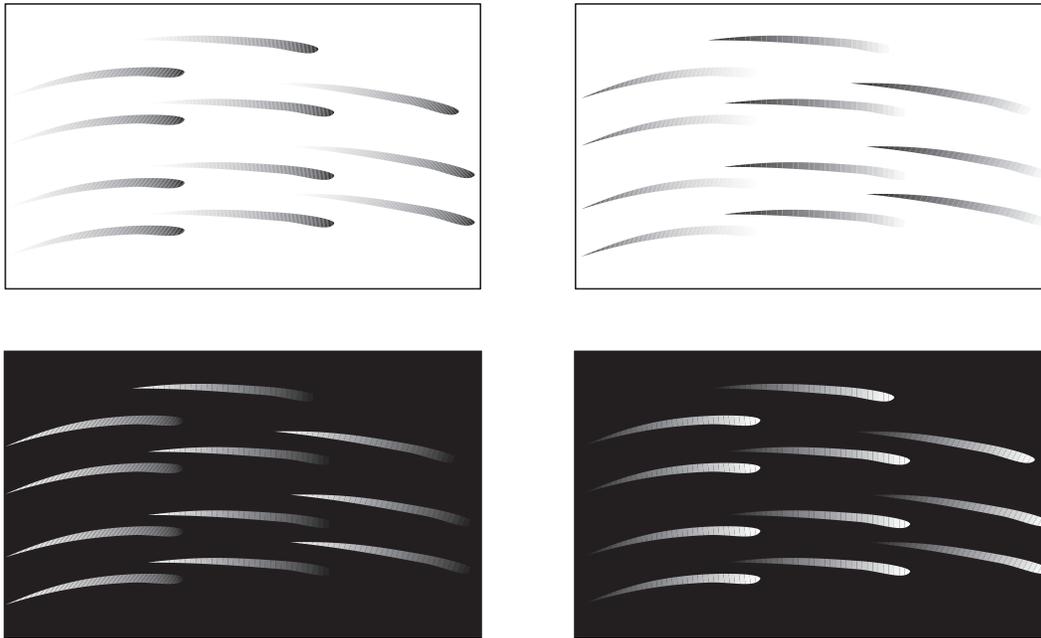
Laidlaw et al. (2001) carried out an experimental comparison of the six different flow visualization methods illustrated in Figure 6.25 and briefly described as follows.



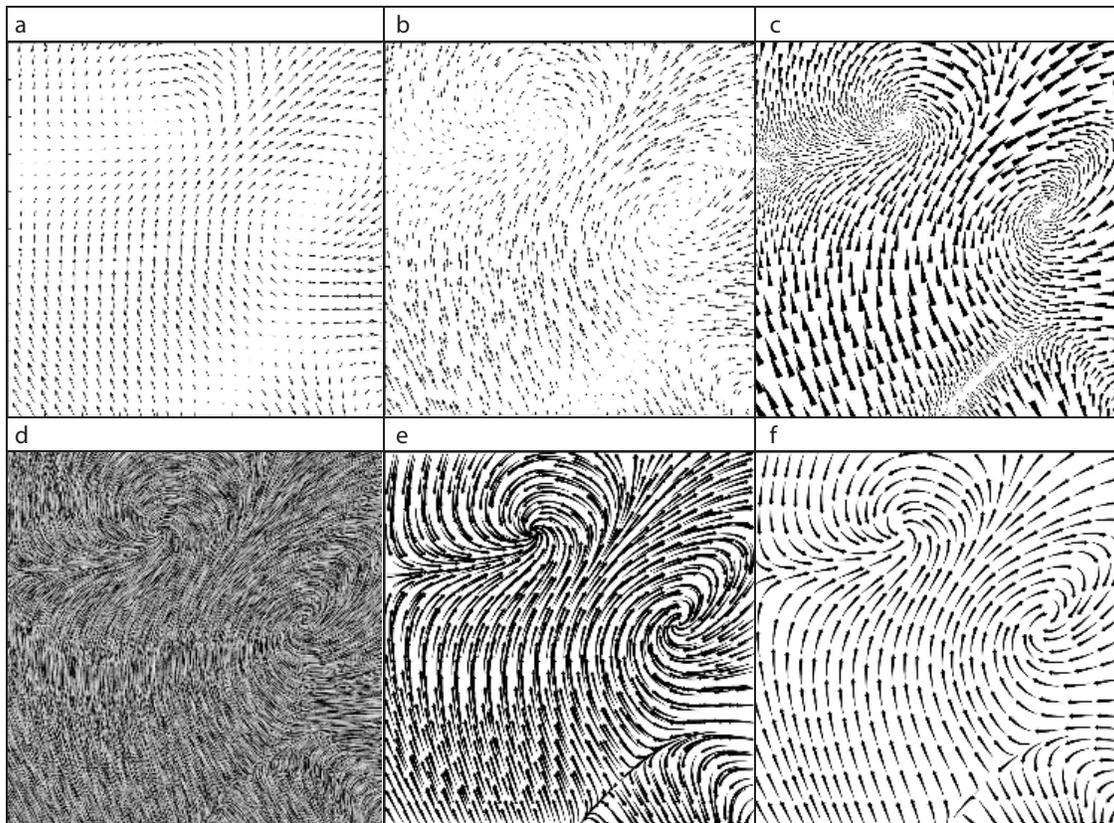
**Figure 6.22** Vector field streamlines are an effective way to represent vector field or flow field data. However, the direction is ambiguous and the magnitude is not clearly expressed (Turk and Banks, 1996).



**Figure 6.23** Drawing in a style based on the pen strokes used by Edmund Halley (1696), discussed in Tufte (1983), to represent the trade winds of the North Atlantic. Halley described the wind direction as being given by “the sharp end of each little stroak pointing out that part of the horizon, from whence the wind continually comes.”



**Figure 6.24** Vector direction can be unambiguously given by means of color change relative to the background.



**Figure 6.25** Six different flow visualization techniques evaluated by Laidlaw et al., 2001. *Used by permission.*

- (a) Arrows on a regular grid. Fixed length.
- (b) Arrows on a jittered grid to reduce perceptual aliasing effects. Fixed length.
- (c) Triangle icons. Icon size proportional to field strength density inversely related to icon size (Kirby et al., 1999).
- (d) Line integral convolution (Cabral and Leedom, 1993).
- (e) Large-head arrows along a streamline using a regular grid (Turk and Banks, 1996).
- (f) Large-head arrows along streamlines using constant spacing algorithm. (Turk and Banks, 1996).

In order to evaluate any visualization, it is necessary to specify a set of tasks. Laidlaw et al. (2001) had subjects identify critical points as one task. These are points in a vector or flow field where the vectors have zero magnitude. The results showed the arrow-based methods illustrated in Figure 6.25(a) and (b) to be the least effective for identifying the locations of these points. A second task involved perceiving advection trajectories. An *advection trajectory* is the path taken by a particle dropped in a flow. The streamline methods of Turk and Banks proved best for showing advection, especially the method shown in Figure 6.25(f). The line integral convolution method, shown in Figure 6.25(d), was by far the worst for advection, probably because it does not unambiguously identify direction.

Although the study done by Laidlaw et al. (2001) is the first serious comparative evaluation of the effectiveness of vector field visualization methods, it is by no means exhaustive. There are alternative visualizations, and those shown have many possible variations: longer and shorter line segments, color variations, and so on. In addition, the tasks studied by Laidlaw et al. do not include all of the important visualization tasks that are likely to be carried out with flow visualizations. Here is a more complete list:

- Identifying the location and nature of critical points
- Judging an advection trajectory
- Perceiving patterns of high and low velocity
- Perceiving patterns of high and low vorticity (sometimes called *curl*)
- Perceiving patterns of high and low turbulence

Both the kinds and the scale of patterns that are important will vary from one application to another; small-scale detailed patterns, such as eddies, will be important to one researcher, whereas large-scale patterns will interest another.

The problem of optimizing flow display may not be quite so complex and multifaceted as it would first seem. If we ignore the diverse algorithms and think of the problem in purely visual terms, then the various display methods illustrated in Figures 6.22 through 6.25 have many characteristics in common. They all consist principally of contours oriented in the flow direction,

although these contours have different characteristics in terms of length, width, and shape. The line integral convolution method illustrated in Figure 6.25(d) produces a very different-looking, blurry result; however, something similar could be computed using blurred contours. Contours that vary in shape and gray value along their lengths could be expressed with two or three parameters. The different degrees of randomness in the placement of contours could be parameterized. Thus, we might consider the various 2D flow visualization methods as part of a family of related methods—different kinds of flow oriented contours. Considered in this way, the display problem becomes one of optimizing the various parameters to reveal important aspects of the data for a particular set of tasks and not so much a problem of developing new algorithms.

## Perception of Transparency: Overlapping Data

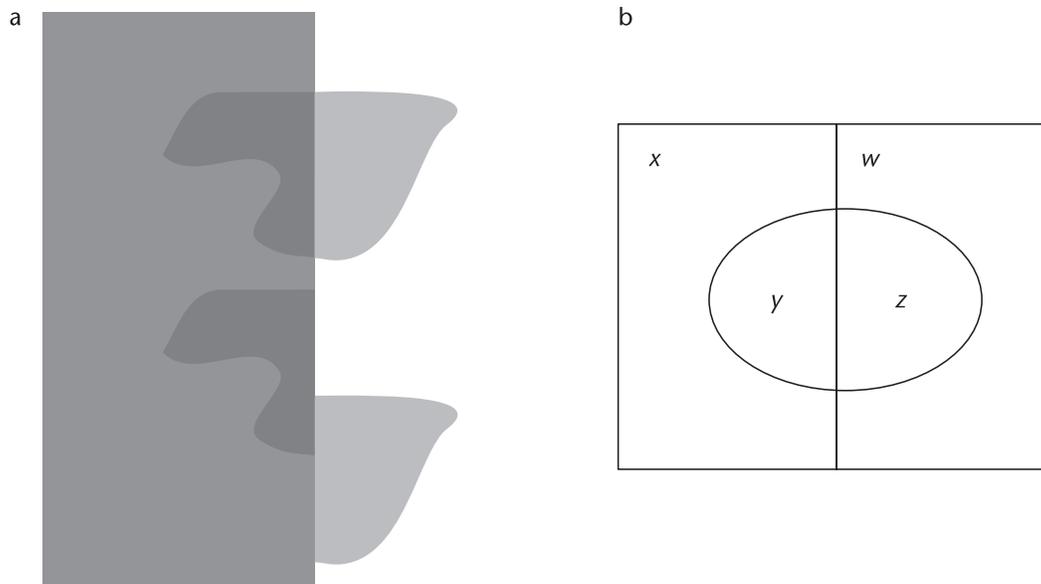
In many visualization problems, it is desirable to present data in a layered form. This is especially common in geographic information systems (GISs). Sometimes, a useful technique is to present one layer of data as if it were a transparent layer over another. However, there are many perceptual pitfalls in doing this. The contents of the different layers will always interfere with each other to some extent, and sometimes the two layers will fuse perceptually so that it is not possible to determine to which layer a given object belongs.

In simple displays, as in Figure 6.26(a), the two main determinants of perceived transparency are good continuity (Beck and Ivry, 1988) and the ratio of colors or gray values in the different pattern elements. A reasonably robust rule for transparency to be perceived is  $x < y < z$  or  $x > y > z$  or  $y < z < w$  or  $y > z > w$ , where  $x$ ,  $y$ ,  $z$ , and  $w$  refer to gray values arranged in the pattern shown in Figure 6.26(b) (Masin, 1997). Readers who are interested in perceptual rules of transparency should consult Metelli (1974).

Another way to represent layers of data is to show each layer as a see-through texture or screen pattern (Figure 6.27). Watanabe and Cavanaugh (1996) explored the conditions under which people perceive two distinct overlapping layers, as opposed to a single fused composite texture. They called the effect *laciness*. In Figure 6.27(a) and (b), two different overlapping rectangles are clearly seen, but in (c), only a single textured patch is perceived. In (d), the percept is bistable. Sometimes it looks like two overlapping squares containing patterns of “–” elements; sometimes a central square containing a pattern of “+” elements seems to stand out as a distinct region.

In general, when we present layered data, we can expect the basic rules of perceptual interference, discussed in Chapter 5, to apply. Similar patterns interfere with one another. Graphical patterns that are similar in terms of color, spatial frequency, motion, and so on, tend to interfere more with one another than do those with dissimilar components.

One possible application of transparency in user interfaces is to make pop-up menus transparent so that they do not interfere with information located behind them. Harrison and Vicente (1996) investigated the interference between background patterns and foreground trans-



**Figure 6.26** In (a), transparency is perceived only when good continuity is present and when the correct relationship of the colors is present. See text for an explanation of (b).

parent menus. They found that it took longer to read from the menu with text or wireframe drawings in the background than with continuously shaded images in the background. This is exactly what would be expected from an interference model. Because a continuously shaded image lacks the high-frequency detail of a wireframe image or text, there will be less interference between the two. The advantages of transparent layered displays must be weighed against the perceptual interference between the layers. For the designer to minimize visual interference, layers must be maximally separated in the different visual channels. Color, texture, motion, and stereoscopic depth channels can all be used in any combination, depending on the design requirements. The more channels used, the better the separation will be.

## Pattern Learning

If pattern perception is, as claimed, fundamental to extraction of meaning from visualizations, then an important question arises. Can we learn to see patterns better? Artists talk about seeing things that the rest of us cannot see, and ace detectives presumably spot visual clues that are invisible to the beat officer.

What is the scientific evidence that people can learn to see patterns better? The results are mixed. There have been some studies of pattern learning where almost no learning occurred. An often-cited example is the visual search for the simple conjunction of features such as color and shape (Treisman and Gelade, 1980). But other studies have found learning for certain patterns (Logan, 1994). A plausible explanation is that pattern learning occurs least for simple, basic