Chapter 5

Visual Perception

5.1 Visual Illusions

The human eye-brain system is arguably the most sophisticated computing system which we have access to. It can easily handle complex visual processing and pattern recognition tasks which would be impossible to attempt on even the most powerful supercomputer.

If we are going to use our visual skills to assist us in data analysis, it is important to remember that they have evolved to handle quite different tasks from those encountered in a typical data analysis. The tasks which our visual system excels at are those which were useful to our hunter-gatherer forbears. These tasks include recognising shapes, discerning colour, judging sizes and distances, and tracing and extrapolating motion in three dimensions. Some of these skills seem directly useful in data analysis, but it is very important that we understand both the strengths and weaknesses of the visual system when used in this way.

Despite its awesome power, we tend to take our visual system for granted; perhaps because we make use of it in virtually every task we perform. One assumption that we make is that we can trust what we see. After all, “seeing is believing”. In fact this is not always the case, and sometimes we can fooled. This is revealed by the existence of visual illusions and we’ll look at a number of these now.

Most of the time our eyes give us a good sense of the way the things are. Sometimes, however, we can be quite misled into seeing effects which are not really present. Images which produce this kind of phenomenon are called *visual illusions*.

Figure 5.1 is a spectacular example of a visual illusion. It is hard to believe, but all the lines in this figure are either horizontal or vertical, and all the black and white polygons are squares. To check this, you should try lying a piece of paper along the horizontal lines. The illusion was first discovered by accident when workmen used the pattern to decorate the outside wall of a cafe with a black and white tiles.

The existence of visual illusions indicates that we need to be careful when using graphics in data analysis. If an apparent feature in a graph is due to a visual illusion rather than a real effect, then we may draw the wrong conclusions from our analysis.

There are many kinds of illusion, some are purely geometric, and others related to colour perception. In this section we will examine a few of the more famous geometric illusions and see what consequences the existence of these illusions may have for data analysis.

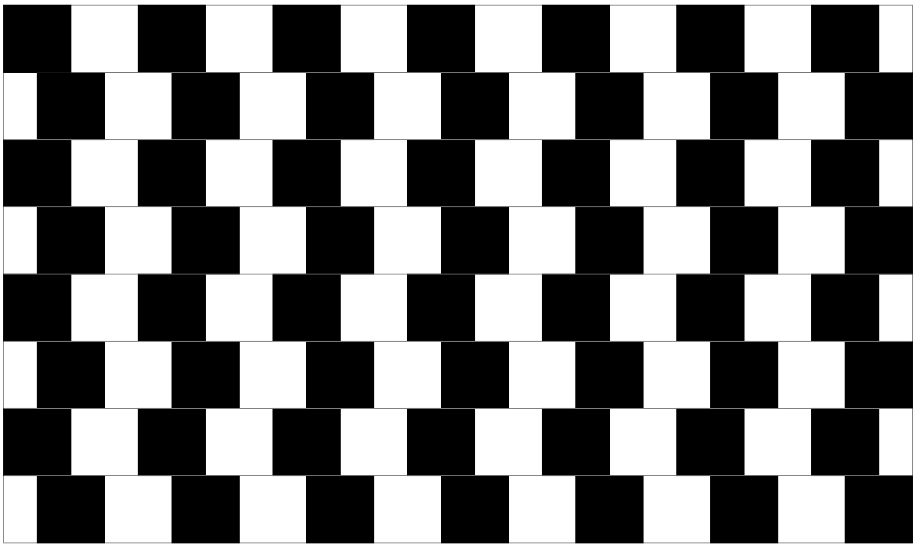

Figure 5.1: The cafe wall illusion.

5.2 Geometric Illusions

Perspective Illusions

We saw in section sec:perspective that the human eye acts as a “pinhole camera” and that this results in the image of an object which appears on our retina as being larger when that object is close than when it is far away. Our brain however corrects our interpretation of the size of objects using any additional knowledge about how far away the object is. When we are mislead about how far an object is away, we can also be mislead about its size.

The Ames room leads us to misjudge the relative size of objects it contains because it deliberately gives false cues about distance. We can also be lead to misjudge the size of things in two dimensional pictures because we use depth cues based on apparent perspective to “correct” the size of objects.

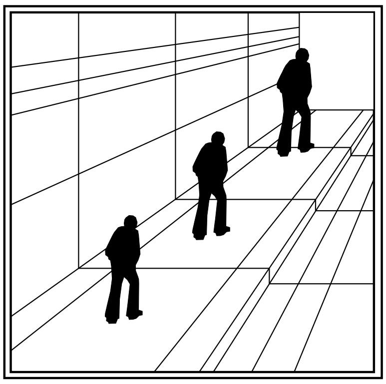

Figure 5.2: A perspective illusion

Figure 5.2 shows such a perspective illusion. The perspective is emphasized because many of the lines in the picture converge to a common vanishing point. The “figures” in the picture are all the same size, but because the appear to be progressively further away, we see them as progressively larger.

In this figure there are obvious features which suggest perspective. Sometimes the cues are more subtle, and we are not always aware that we are being mislead.

The Ponzo Illusion

Figure 5.2 has a clear interpretation as a three dimensional scene, so it is no surprise that perspective influences our perception of the size of objects in it. The effects of perspective an be much more subtle however.

The Ponzo illusion (named after the Italian psychologist Mario Ponzo) is a famous example of how equal length lines can be perceived as having different lengths. The lower of the two horizontal line segments appears to be shorter than the upper one. One explanation for this is that the sloping lines create the same impression of depth as, for example, railway lines. A line of a certain length is perceived as being longer the further away you think it is. The phenomenon shows how the visual system tends to treat figures as three dimensional.

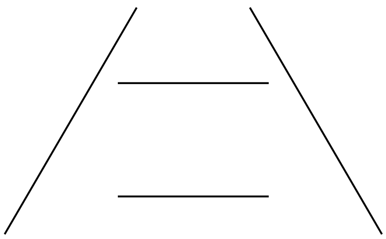

Figure 5.3: The Ponzo illusion

The Müller-Lyer Illusion

The Müller-Lyer illusion takes its name from Franz Carl Müller-Lyer (1857–1916), who studied medicine in Strasbourg and served as assistant director of the city’s psychiatric clinic.

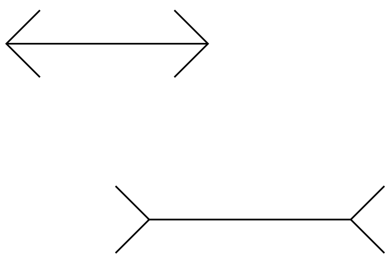

Figure 5.4: The Müller-Lyer Illusion.

The illusion shows that we can be deceived in our perception of the relative lengths of two line segments, even when those line segments are aligned. The simplest variant of the illusion is presented in figure 5.4. Despite the fact that the line segments are of equal length, the lower line segment appears to be longer.

There are a number of theories which attempt to explain this illusion, but there is no definitive explanation. One theory is that we interpret the figures in three dimensions. The upper image is seen as the outside edge of a box, while the lower image is seen as the inside edge of a box. Since it is assumed that the two “solid objects” rest on the same surface, the outside edge appears closer to the viewer than the inside edge. The more “distant” line is then interpreted as larger, to correct for the difference in apparent distance.

The Poggendorff Illusion

This illusion was invented by J. C. Poggendorff in 1860. It shows the incorrect perception of the continuation of a line.

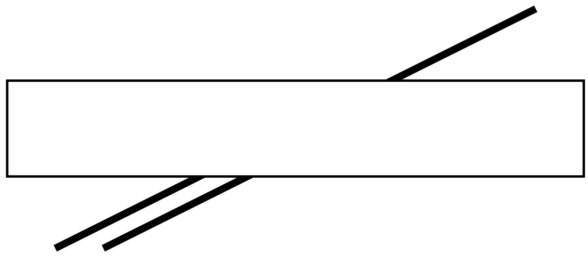

Figure 5.5: The Poggendorff illusion.

In figure 5.5, it appears that it is the rightmost line below the rectangle which continues and protrudes from the top of the rectangle. In fact, it is the leftmost line.

One explanation for the illusion is based on the observation that people tend to overestimate the size of sharp angles formed where the sloping and horizontal lines meet. Figure 4.8 provides experimental backing for this view.

The Zöllner Illusion

The Zöllner illusion is related to the Poggendorff illusion because both are the result of the misjudgement of angles. In the Zöllner illusion, the result is that the diagonal lines through the plot seem to curve in a fashion similar to the horizontal lines in the cafe wall illusion.

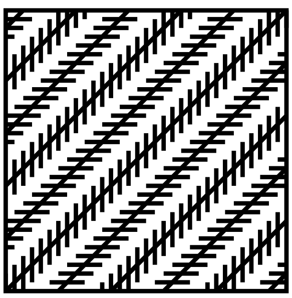

Figure 5.6: The Zöllner illusion.

The Horizontal-Vertical Illusion

The horizontal-vertical (or upside-down T) illusion provides a third example of how equal lengths can be perceived as different. Figure 5.7 shows two different arrangements of two line segments. When most people look at figure 5.7 (a), they perceive the horizontal line as being shorter than the vertical one. In experiments where subjects are required to adjust the length of the vertical line until the lines appear to have equal lengths, most people leave the vertical line at least 10% shorter, often more.

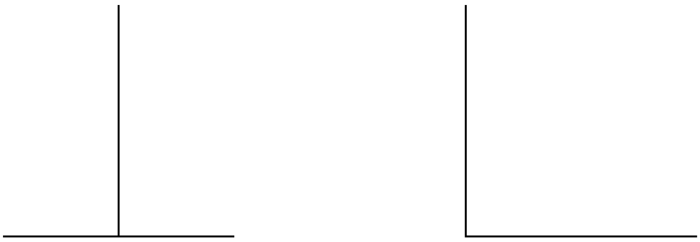

(a) (b)

Figure 5.7: The horizontal-vertical illusion

It might seem that we just perceive vertical and horizontal lengths in different ways, but changing the placement of the lines, as shown in 5.7 (b), destroys the illusion.

5.3 Colour and Irradiation Illusions

The geometric illusions of the last section are based on the distortion of lines and angles. They are created purely by arrangeing line segments on the page. The illusions in this section result from the use of solid areas of gray or colour.

Irradiation Illusions

The irradiation phenomenon causes bright objects on a dark background to appear bigger than the same objects displayed in darker colours on a brighter background. In figure 5.8, the small black and white square are the same size.

This phenomenon is quite obvious in everyday life. When a sporting event takes place with one team in white and one in black (such as when England plays New Zealand at rugby), it is easy to start believing that there is an extreme physical mismatch on the playing field. The players in the white team invariably look bigger.

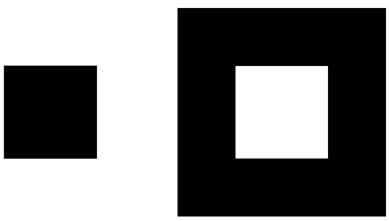

Figure 5.8: The Irradiation Illusion.

It is quite likely that the existence of this illusion explains the why black is a perennially fashionable colour. Light colours make the wearer look larger, and as observed by Wallis Simpson “you can never be too rich, or too thin”.

Mach Bands

The phenomenon of Mach banding occurs when strips of similar shades of gray (or colour) are placed next to each other. The strips do not appear to be of constant colour, but become lighter as they approach darker neighbours and darker as they approach lighter neighbours.

Figure 5.9: Mach bands.

This is almost certainly a result of edge detection being performed somewhere in the visual pathway, and may well be a result of lateral inhibition as discussed in section 4.2.4.

You might think that using a finer partition and a gentler transition between shades of gray would make the banding effect diminish. In fact, doing this makes no difference. The visual system just works harder to make the edges stand out.

One simple way to eliminate the appearance of Mach banding is to replace the straight line boundries between the grays with irregular wandering ones. In this cases the eye does not detect the edges and a smooth blending of the grays takes place.

The Hermann Grid

The Hermann grid consists of a regular array of black squares, separated by thick white lines as shown in figure 5.10. If you examine this figure closely you become aware of “phantom” dark dots which appear and disappear at the intersection of the horizontal and vertical white lines.

Once again, these dots are probably caused by lateral inhibition. Notice that, if you fix your eyes on on one particular intersection, there is no dark dot there. The flashing in and out of existence is caused by the saccades of the eyes as they try (without success) to fixate on the dots.

5.4 Graphical Perception

There are two different ways in which we extract information from graphs. The first of these occurs when we take a quick glance at a graph. Without any apparent conscious effort it is possible to extract a good deal of information. Impressions such as “the

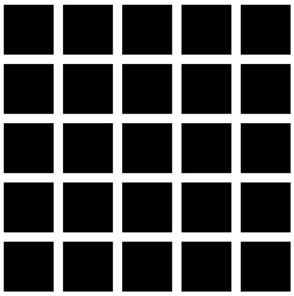

Figure 5.10: The Hermann grid illusion.

graph slopes upwards" are obtained in this way. Because there is no conscious effort involved, this kind of visual processing has been called *pre-attentive vision*. Since our interest here is specifically with statistical graphics, we will use the term *graphical perception* suggested by Cleveland [1].

The second kind of viewing is a more extended process where we consciously think about particular aspects of the graph. It is this kind of viewing that enables us to draw conclusions such as "the tallest peak in the graph occurs very close to $x = 4$" or "a straight line through the points intersects the y axis at about $y = 5$". We will use the term *graphical cognition* for this kind of viewing.

Both these ways of looking at graphs are important, but we will concentrate more on the first of them, because it is the one which makes the use of graphics attractive. In presentation graphics, an understanding of graphical perception can help us to provide a graph with what has been called *inter-occular traumatic impact*[1]. In exploratory work, such an understanding can help us develop and use techniques with the best chance of revealing hidden data features.

Statistical graphs almost always *encode* one or more sets of numbers so that the brain's built-in graphical perception facilities can be used to process them. If we are to have confidence in what we perceive in a plot we need to know that we can effectively *decode* the information in a plot. We will refer to the decoding process as a graphical perception task.

Cleveland [1] lists gives ten elementary graphical perception tasks which are of possible interest in graphics applications.

1. *Angle*
2. *Area*
3. *Colour hue*
4. *Colour saturation*
5. *Density (amount of black)*
6. *Length (distance)*
7. *Position along a common scale*
8. *Position along identical, nonaligned scales*
9. *Slope*
10. *Volume*

[1] It hits the viewer between the eyes.

For the moment, we will ignore the colour related tasks 3, 4, and 5 and consider just the remaining (geometric) tasks.

Angle Judgements

There are a number of ways in which angle judgements are important in statistical graphics. The most familiar example of the use of angles is the commonly used “pie chart”. Figure 5.11 shows a pie chart used to to display five numbers. The figure shows that pie charts do a bad job of displaying numbers. It is not immediately clear from the figure how to order the values.

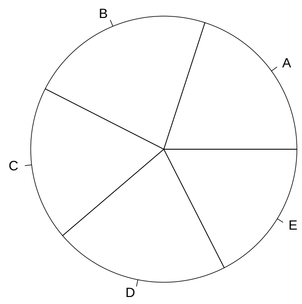

Figure 5.11: A pie chart, showing numbers represented by angles.

It was recognised as far back as last century that there are problems making judgements about angles. We tend to underestimate acute angles and overestimate obtuse ones [6]. This is the source of the Poggendorff illusion. Another problem is that angle judgements depend on orientations. Angles whose bisectors are horizontal tend to be seen as larger than whose bisectors are vertical [8].

Recent experimentation has shown that the best judgements are made about angles of 45° (half-way between acute and obtuse) and line slopes which are close to one (this is a point we will return to).

Area Judgements

Area is often used to encode information in graphs. Perceptual experiments have shown that such information is generally not decoded very well, even when optimally presented. Figure 5.12 shows a square and circle. It is not easy to tell which figure has the

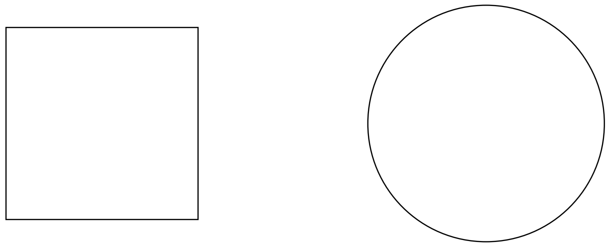

Figure 5.12: Which figure has the larger area?.

Figure 5.13: A map of the south-eastern United States produced using the Albers equal area projection.

larger area. (In fact, the areas are equal.)

The shapes in 5.12 are relatively simple, convex ones. The problem of decoding areas becomes even more difficult with general shapes. In general, if an area is displayed in a long thin figure, it is perceived as larger than if it is displayed in a compact convex one. There are also distortions in area decoding which can occur as a result colour and shading, as demonstrated by the irradiation illusion shown in figure 5.8.

Cleveland [1] points out that these biases in area decoding probably have an influence on how we judge the area of countries, states and provinces. He gives the example of the areas of the states Florida and Georgia in the USA. Figure 5.13 shows a map of the contiguous Unites States constructed using a projection which preserves areas (the Albers equal area projection). The states of Florida and Georgia are labelled on the plot. It appears that Florida is rather larger than Georgia, but in fact, it is slightly smaller. It appears larger because it is more extended.

Length Judgements

Length is an obvious geometric quality which can be used to encode numerical values. The Ponzo, Müller-Lyer, and horizontal-vertical illusion show that we need to take some care when making length judgements, but despite this, we might expect that we might be able to decode lengths fairly well so long as we keep the figure simple. Figure 5.14 shows that this is not the case. It is quite difficult to tell that the *A* and *B* lengths

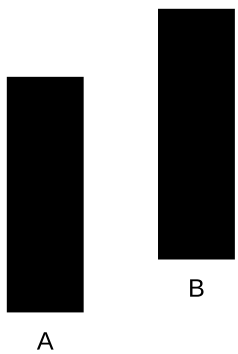

Figure 5.14: Length decoding.

are different.

It is possible to to rather better than figure 5.14 using a the simple device shown in figure 5.15. This amounts to adding a simple benchmark “scale” to the plot.

The device works because rather than comparing the black portions of the figures, we compare the lengths of the white lengths.

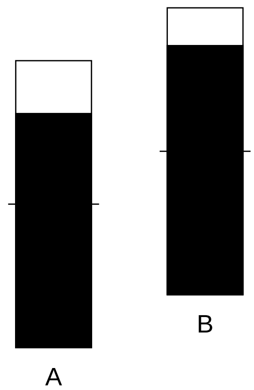

Figure 5.15: Length decoding using a scale benchmark.

Note that the lengths of the white portions of the bars differ by the same absolute amount as the lengths of the black portions. The fact that the white bars are clearly of different lengths while the black ones are not, makes it clear that it is the *relative* difference which we perceive rather than the absolute one.

Position Along a Common Scale

Position along a common scale is the basis for many of the standard plots used in statistics; this includes histograms, bar plots and scatterplots. Decoding values by judging position along a scale is one of the perceptual tasks which we seem to perform best. To see that it is superior to decoding by comparing areas, let us reexamine the values presented in the pie chart of figure 5.11, but this time using a barchart.

The barchart makes it clear that the values are ordered as

$$E < C < A < D < B,$$

something which is only discerned with great difficulty from the pie chart.

Position Along Non-Aligned Scales

Our ability to decode information from positions along identical non-aligned scales is not as good as our ability to decode information from positions on a single scale, but it is still quite good.

Figure 5.15 is an example of decoding along common, non-aligned scales. Clearly we are better at this than we are at direct length judgements.

Slope Judgements

Making slope judgements can be of importance when looking at curves and straight lines. Making assessments of slope has a good deal in common with making assessments of angles as shown in figure 5.18. Indeed, there is evidence that when we make

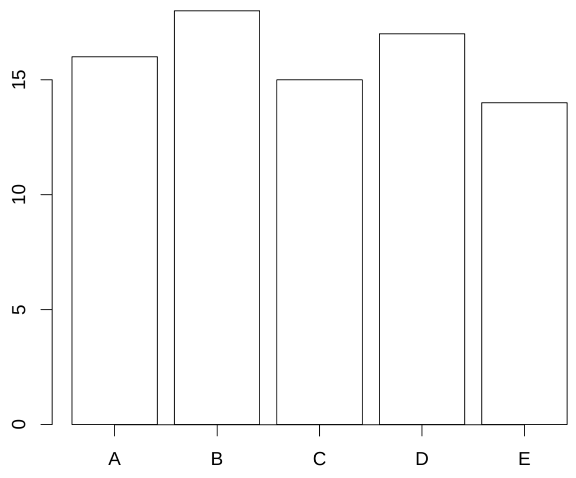

Figure 5.16: Decoding using position along a scale.

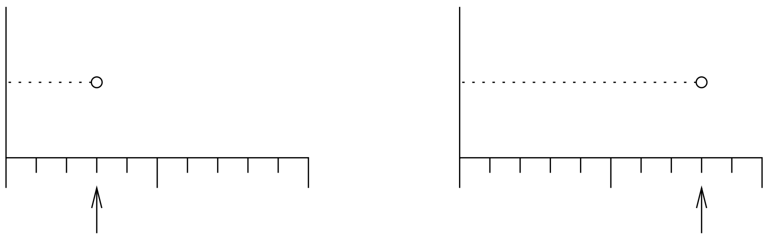

Figure 5.17: Decoding using positions on non-aligned scales.

assessments about slopes and changes of slope, it **is** angles which we look at. These means that making slope judgements suffers from the same problems as making angle judgements. In particular, the best decoding of slopes happens when the slopes are about 1.

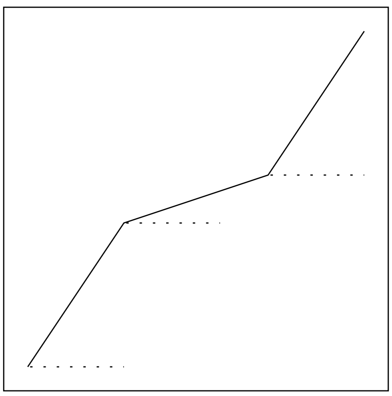

Figure 5.18: Slope decoding using angles.

Volume Judgements

Our ability to visually assess volumes is not good as our ability to assess areas. In data analysis it is very dangerous to try to encode values as volumes because we cannot actually represent three-dimensional objects on screen or paper. Because we work with two dimensional renderings, it is not clear whether the viewer is seeing a two dimensional or a three-dimensional object (or perhaps even something between).

5.5 Perception "Laws"

The way we decode information from graphs and pictures is of great interest to perceptual psychologists, and has been the subject of a great deal of study and speculation. In the course of this study, a number of "laws of perception" have been formulated. These laws are based on empirical study, but give some quite useful guidelines which can be applied in graphics.

Weber's Law

Weber's law, formulated by the 19th century psychophysicist E. H. Weber, is one of the most important of the laws of human perception. Suppose that $w_p(x)$ is a positive number such that a line of length $x + w_p(x)$ is detected as longer than a line of length x with probability p. Weber's law states that for a fixed value of p,

$$w_p(x) = k_p x,$$

where the value of k_p does not depend on x. In simple terms, Weber's law states that it is relative rather than absolute differences which are perceived when we compare lengths.

Figures 5.14 and 5.15 show an example of Weber's law in action. It is easy see the difference of lengths in the second of these figures because the relative difference in the white part of the bars is quite large. Detecting the difference in the first figure is harder because the relative difference in the lengths of the dark portions of the bars is small.

The same example suggests that comparisons along identical but non-aligned scales will be more accurate than pure length comparisons.

Stevens' Law

Suppose that we have chosen a way of encoding numerical values in a graph. How do the values decoded from the graph relate to the original values? Stevens' law, formulated and extensively investigated by psychologist S. S. Stevens, give a general answer to this question. It says that the for a given person and a given perceptual task the perceived values follow the general law

$$p(x) = cx^{\beta}$$

where c and β depend on the particular person and perceptual task.

Suppose that we had encoded values x_1 and x_2 in a graph. Then the perceived values for x_1 and x_2 will be $p(x_1)$ and $p(x_2)$. If we compare the relative sizes of the perceived values we see

$$\frac{p(x_1)}{p(x_2)} = \left(\frac{x_1}{x_2}\right)^{\beta}.$$

This shows that it is β is the key parameter in Steven's law.

Stevens' law provides and excellent description of how we perceive values through a number of encodings, including length, area and volume. Because of the importance of β many experiments have been conducted to try to estimate its value for particular ways of encoding graphical information.

The following ranges have been determined for the value of β (averaged across experiments and subjects).

Attribute	*Range for* β
Length	0.9 – 1.1
Area	0.6 – 0.9
Volume	0.5 – 0.8

Because the value of β is close to 1 for length judgements, the ratio of the perceived values will be close to the ratio of the actual values.

$$\frac{p(x_1)}{p(x_2)} \approx \frac{x_1}{x_2}.$$

This means that judgements of relative sizes of values encoded by length will show little bias.

On the other hand, if we take the middle of the range of β values for area judgements, we see that the ratio of the perceived values is

$$\frac{p(x_1)}{p(x_2)} = \left(\frac{x_1}{x_2}\right)^{0.7}.$$

If we compare an area of $1/2$ with an area of 1, the perceived ratio of areas is 0.62, which is rather larger than the actual ratio of 0.5. If we compare an area of 2 with an area of 1, the perceived ratio of areas is 1.62, which is rather smaller than the actual ratio of 2. In general, large areas are perceived as smaller than they actually are and small areas are perceived as larger than they really are.

Since the range of β values for volume judgement is even further removed from 1, representing values with volumes leads to even greater distortions.

5.6 Perception Experiments

Weber's and Stevens' laws give some guidance in choosing graphical encodings for numerical values. We know for example that length provides a less-biased way of encoding values than area and that area is less biased than volume. We also know that positions along non-aligned scales will provide a better encoding than length.

Although these facts are useful, it would be more useful how all the ten perception tasks listed the start of section 5.4 relate to each other. Ideally, we would like to order the tasks from best to worst. One way to go about this is to conduct experiments to determine which tasks are easier to carry out than others.

Such experiments were carried out in the early 1980s by researchers Bill Cleveland and Robert McGill at Bell Laboratories [3], [2], [4]. In one experiment, subjects were asked to decode values which had been encoded in the following ways.

1. Vertical distances of dots from a common baseline

2. Vertical distances of dots above non-aligned baselines
3. Lengths of lines
4. Slopes of lines
5. Sizes of angles
6. Areas of circles
7. Areas of irregular blobs

Figure 5.19 shows the kind of display used in this experiment. In each of the seven panels, the object on the left is a standard and the subject is asked to give values for the other values in the panel as a percentage of the standard. Thus for each test sheet a subject was asked to make 3×7 judgements and each subject was given 10 such sheets. In all there were 127 subjects in the experiment. The results of the experiment were assessed by considering

$$error = |judged\ percent - true\ percent|.$$

A graphical presentation of the results is given in figure 5.20. We can clearly see that the encoding of values by angle, slope or area produces significantly larger errors than position along aligned or non-aligned scales or length.

Using additional experiments, Cleveland and McGill were able to produce an ordering of the perception tasks from easiest to hardest. Some of the ordering (that connected with colour perception tasks) is based on conjecture, but it seems to be in accord with what is found in practice. The ordering proposed by Cleveland and McGill is as follows.

1. *Position along a common scale*
2. *Position along identical, nonaligned scales*
3. *Length*
4. *Angle — slope*
5. *Area*
6. *Volume*
7. *Colour hue — Colour saturation — Density*

5.7 Applying Graphical Perception Ideas

When constructing a graphical display, it is important to take into account the ideas discussed in this chapter. Taking account of how your display can help you to produce much more effective displays.

- Use an encoding scheme which leads to a decoding task as high up the Cleveland McGill ordering as possible.
- Avoid the use of area, and volume. Stevens' law shows that there will be significant biases which occur as a result of using these encodings.

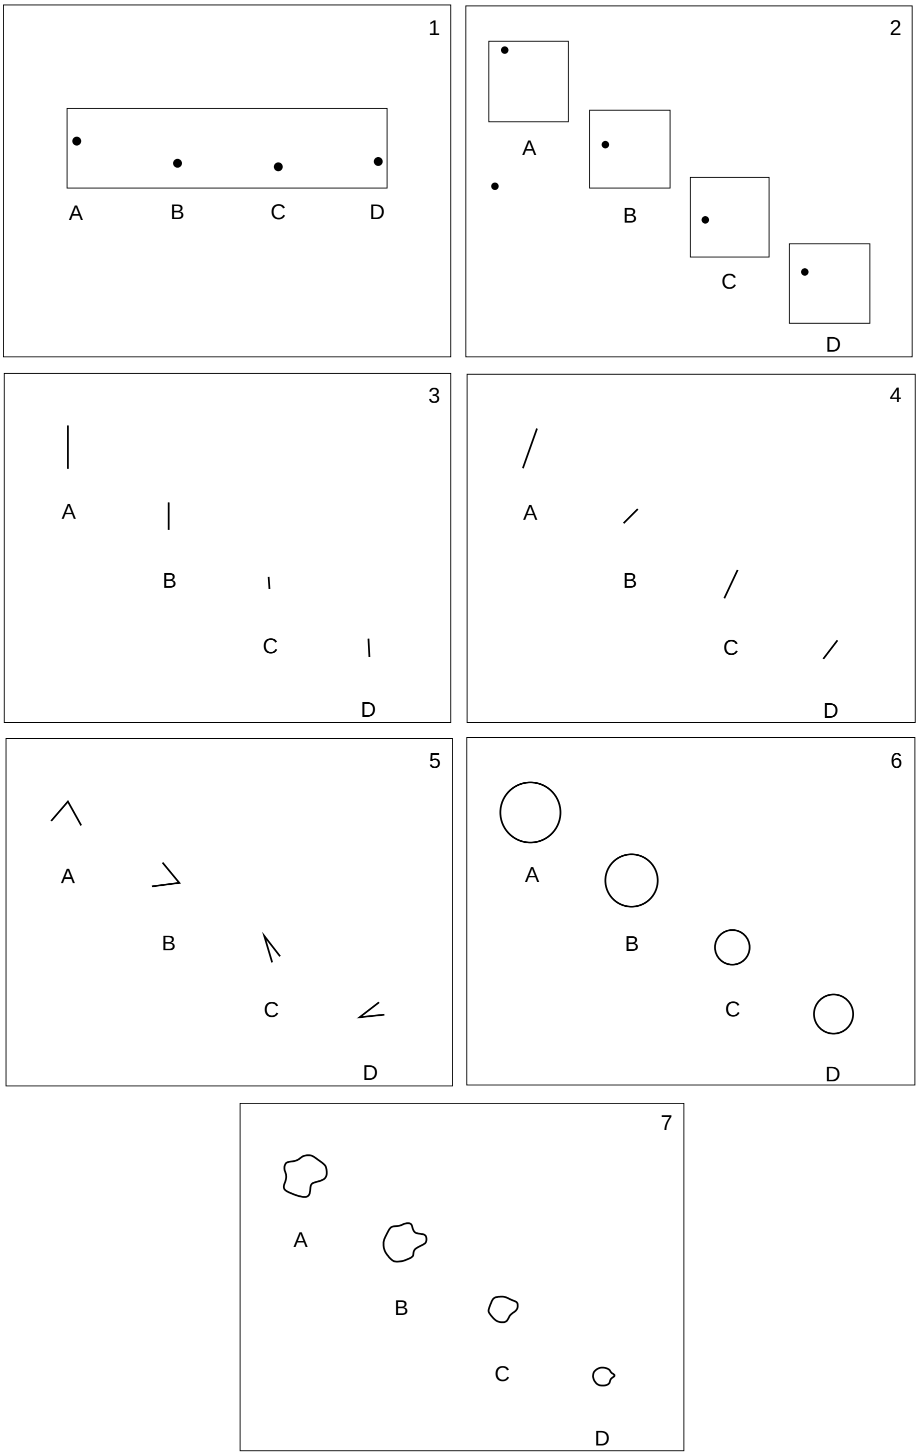

Figure 5.19: A sheet containing seven perception tasks.

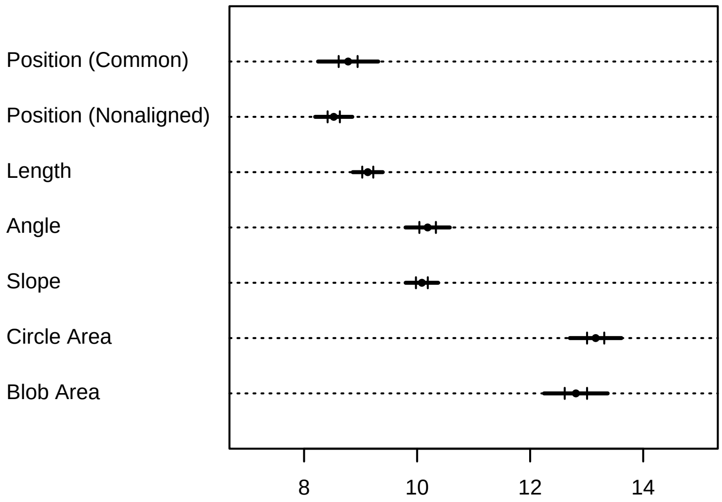

Figure 5.20: Measures of the errors in the elementary perception tasks. The two-tiered error bars are 50% and 95% confidence intervals for the measures.

- Avoid the use of colour and shading to encode numerical values. The decoding task is significantly harder for these encodings and will make your displays harder to use.
- Be careful that you do not create visual illusions either through the layout of your plots or through the use of noisy patterns. Such illusions make it significantly harder to see what the graph is truly saying.