

Understanding Charts and Graphs

STEPHEN M. KOSSLYN

Harvard University

SUMMARY

Many charts and graphs do not convey information effectively. This article develops a way of analysing the information in charts and graphs that reveals the design flaws in the display. The analytic scheme requires isolating four types of constituents in a display, and specifying their structure and interrelations at a syntactic, semantic, and pragmatic level of analysis. As the description is constructed, one checks for violations of 'acceptability principles', which are derived from facts about human visual information processing and from an analysis of the nature of symbols. Violations of these principles reveal the source of potential difficulties in using a display.

Most people have had the experience of opening a well-known national news magazine and puzzling over a chart or graph, trying to figure out what it is about and what it is supposed to be telling the reader. Often one can point to some aspect of the offending bit of graphics and say that *those* lines are too close together or *that* mislabelled axis is the root of the problem. But often one is not sure exactly what is wrong, and would be unable to tell the artist how to improve his or her work. In this article I develop a systematic way of characterizing what is right, and wrong, about any given chart or graph.

The present system hinges on describing a display at multiple levels of analysis. Because of the way the scheme was designed, it should be easily used to describe any unambiguous chart or graph in a straightforward way. When the system cannot be easily applied, this indicates that something is wrong. A set of principles is described that should be adhered to if a chart or graph is to be effective, and one or more principles has been violated when the analytic system cannot be applied easily to a display. Similarly, if one uses the system as one is constructing a display, the result should be unambiguous and easily apprehended.

TYPES OF VISUAL DISPLAYS

There are numerous and varied ways in which people illustrate ideas or concepts. We are concerned here with how quantitative information and relations among discrete entities are communicated graphically. These displays necessarily use symbols, which are marks that are interpreted in accordance with convention. There are four common types of 'symbolic' displays, which differ in terms of what information is communicated and how information is communicated. These types can be combined to form hybrids (e.g. a map with pie charts over the states), but

the hybrids will not have emergent properties; they can be understood as simple amalgams of the types delineated here.

Graphs are the most constrained form, with at least two scales always being required and values being associated via a 'paired with' relation that is always symmetrical. Graphs represent greater quantities of the measured substance by greater area, longer lines, or more of some other visual dimension; more along a visual continuum represents more of the symbolized entity. Thus, by this usage a 'pie chart' is actually a graph.

Charts specify discrete relations among discrete entities. Charts have an internal structure in which entities must be visibly related to other entities by lines or relative positions that serve as links. These links can be labelled or unlabelled, directed or undirected, and need not simply pair entities; a very wide variety of types of relations are possible (not simply the paired-with relation illustrated in graphs).

Maps are unlike charts and graphs in that they are not entirely conventional: a part of a map corresponds non-arbitrarily to a part of the pictured territory. The internal relations among parts of a map are determined by the spatial relations of what is pictured. However, maps usually include a symbolic component (e.g. different colours representing different populations), and labels are paired with locations.

Finally, *diagrams* are schematic pictures of objects or entities. These can incorporate objects, such as parts of a machine, or abstract concepts, such as forces acting on the parts. A diagram is symbolic in that special symbols (e.g. cross-hatching to illustrate curvature) are used, which are interpreted (at least in part) by convention; a photograph is not symbolic because no conventional means of representation are exploited. Unlike charts and graphs, the parts of a diagram correspond to parts of some actual object or entity; and unlike maps, parts of diagrams do not represent locations of a territory.

In this paper we will concentrate on a detailed treatment of charts and graphs. In many ways graphs are the most interesting because they are perhaps the most varied form while, at the same time, are among the most constrained. That is, there are numerous different types of graphs—line, bar, surface, divided bar, pictograph—and yet the way they function to communicate information is well-structured. Although some types of displays, such as maps, are more constrained (the shapes must resemble those of regions being represented), they are also less varied. The insight gained from understanding the way charts and graphs convey information applies to all of types of displays; we can simply apply the various strictures as appropriate when making a map or diagram.

Thus, the approach will be first to understand the most structured and demanding cases, where graphs are used to communicate detailed information clearly and concisely. We then will be in a position to consider special cases, where only some subset of the complete information need be conveyed.

This paper has three major parts. The first begins by outlining the key elements of the analytic scheme. Following this, the bases of the evaluative procedure are introduced. Finally, the 'acceptability principles' are developed and are used to diagnose problems with displays. A relatively simple example is used to illustrate these ideas, and a very complex example is presented in the Appendix.

ELEMENTS OF THE ANALYTIC SYSTEM

The analytic system allows one to diagnose problems with a display. The system can be applied easily only to a perfect display; when there is any difficulty in using the systems, this alerts one to a problem. A problem is revealed when one has to puzzle over a display to decide how to use the descriptive procedure; the interpretation of a good display should be unambiguous and immediately obvious.

A chart or graph is described in terms of a set of components and relations among them. This description is generated at three levels of analysis, as is developed below.

'Basic level' graphic constituents

Displays are described as sets of components with specific relations among them. The components can be most usefully described at a 'basic level'. The notion of 'basic level' used here is directly analogous to Rosch, Mervis, Gray, Johnson, and Boyes-Braem's (1976) conception of a 'basic level' in categorization hierarchies. In categorization, the basic level is the one that is as general as possible while still having as many similar members as possible. For example, 'apple', and not 'fruit' or 'delicious apple', is the basic level because that category captures the most exemplars that are still very similar; going down the hierarchy results in fewer exemplars in the category, say delicious apples, whereas going up the hierarchy, to fruit, results in the exemplars, not being very similar to each other (Rosch *et al.* provide additional criteria, but this is adequate for present purposes).

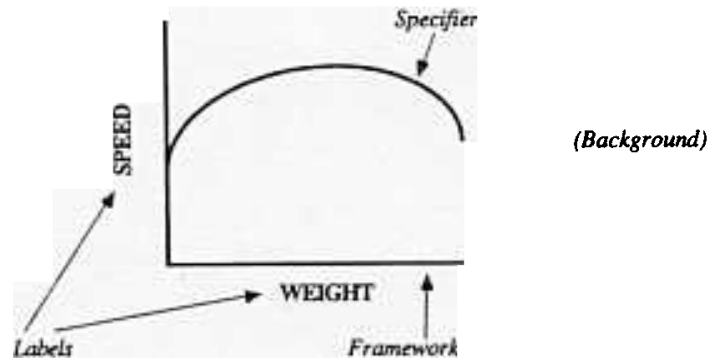
Similarly, the basic-level graphic constituents posited here seem to be the most general way of classifying the components of a chart or graph that still have a high degree of similarity among the different instances of the class. However, in the present case the similarity is not in appearance, but in function, in the role a constituent plays in how information is represented in a display. The four constituents used here are called the *background*, the *framework*, the *specifier*, and the *labels*. These constituents are defined at the level of semantics (described below), in terms of the information directly conveyed. Figure 1 serves to illustrate these basic-level constituents for a typical chart and graph.

Background


The background is a pattern over which the other components of the display are presented. The background serves no essential role in communicating the particular information conveyed by a chart or graph; if the background were removed, the chart or graph would still convey the same information. Although a given background is not a necessary part of a chart or graph (often the background is blank), occasionally a patterned background, such as a photograph, can serve to reinforce the information in a chart or graph (e.g. dead soldiers in a graph about the horrors of war); a patterned background can also interfere with ability to read a display, as will be discussed shortly.

Framework

The framework represents the kinds of entities being related (e.g. year and oil production), but does not specify the particular information about them conveyed



A. GRAPH



B. CHART

Figure 1. The basic-level constituents of charts and graphs. The upper panel displays a graph, and the lower panel displays a chart. Note the differences in the kind of information conveyed by the graphic constituents

by the display (e.g. the amount of oil per year). The framework often has two parts: the *outer framework* extends to the edges of the display and serves the role just described; the *inner framework* is nested within the outer one and often intersects elements of the 'specifier' that indicates specific values of the outer framework (e.g. the function in a line graph). The inner framework (often a grid or regular pattern of lines) usually serves to map points on the outer framework to points on the specifier. In most cases the framework also serves to organize the display into a meaningful whole; in some charts, however, this is not true (e.g. see Fig. 1), although the framework still functions as described above. In some cases a component of the framework may be partially deleted or missing altogether (e.g. the x axis of a graph), forcing the reader to fill in a 'virtual line' upon recognizing the type of display.

Specifier

The specifier conveys the particular information about the entities represented by the framework. The specifier serves to map parts of the framework (actually present or inferred by the reader) to other parts of the framework. In graphs the specifier is often a line or bars which pair values on the x and y axes delineated by the framework. In charts the specifier material often consists of directed arrows connecting boxes or nodes.

Labels

The labels are composed of letters, words, numbers, or pictures and provide an interpretation of a line or region (which is a component of either the framework or the specifier).

In addition to identifying these basic-level constituents, the analytic procedure leads us to describe them in terms of their subcomponents. The individual elements are described in terms of simple 'Gestalt wholes' (such as line segments).^{*} We also note the organization of the elements. For example, the framework of the graph illustrated in Fig. 1 is easily described as a horizontal line and a vertical line (the elements) that abut at the bottom tip of the vertical line and left tip of the horizontal line (the organization). One examines not only the relations among elements belonging to the same basic-level constituent, but also the interrelations among elements belonging to different constituents. This procedure helps one to discover whether the basic-level constituents are easily identified and interpreted, and also leads one to evaluate the overall organization of the display.

Levels of analysis: syntax, semantics, and pragmatics

The basic-level constituents and their interrelations are described at three levels of analysis. Although we identify the constituents by observing their function, we begin the description proper with a detailed *syntactic* analysis. This analysis focuses on properties of the lines and regions themselves; here the lines and regions are not interpreted in terms of what they represent, but are treated as entities in their own right. We describe the individual elements and their organization. As will be discussed in detail shortly, this analysis examines questions such as the detectability of variations in lines, the ways in which they perceptually group into units and so on.

The *semantic* analysis focuses on the meanings of the configurations of lines, what they depict or signify (e.g. axes, labels, etc.). The semantic analysis is the literal reading of each of the components of a chart or graph and the literal meaning that arises from the relations among these components.

Finally, the *pragmatic* analysis focuses on the ways in which meaningful symbols convey information above and beyond the direct semantic interpretation of the symbols. At the level of pragmatics in language, for example, the question 'Can you open the door?' is not comprehended as a question; rather, it is a request to open the door. The conveyed meaning in this case is quite different from the literal semantic interpretation. In addition, pragmatic considerations govern the relationship between the information in a display and the reader's needs. For example, pragmatic considerations dictate that one not provide more or less information than is needed by the purpose of the display; too much information is as bad as too little information.

^{*} At one time I considered introducing a set of 'primitive elements' and relations that would provide a fixed 'alphabet of shapes' to be used in all analyses. This proved very difficult to do, however, and proved to be totally unnecessary for present purposes. The wide variety of charts and graphs seems to preclude specification of a reasonably small set of discrete elements from which all charts and graphs can be constructed, but even if this were possible, the important variations seem to occur at the 'basic level' of organization into the graphic constituents noted above.

FOUNDATIONS OF THE EVALUATIVE PROCEDURE

The analytic scheme hinges on specifying the basic-level constituents and their relations at the three levels of analysis. In so doing, one evaluates whether the display violates 'acceptability principles'. These principles have two roots. The first is the literature on how humans process visual information, and the second is the 'theory of symbols' developed by Goodman (1968).

Visual information processing

A wide range of activities is interposed between that instant when one first fixates one's gaze upon a visual display and the point at which one has successfully extracted relevant information from it. The explosion of interest in cognitive psychology over the past three decades has given us a general framework for conceptualizing these activities, and has given us a rich body of literature concerning their operation. This information is useful for understanding some kinds of problems that beset visual displays—namely problems that arise when a display requires use of mental operations we do not have at our disposal. Thus, it will behoove us to consider briefly now (but in more detail shortly) the basic underpinnings of visual information processing, and then to consider how to use this information to diagnose bad displays and guide in the construction of good ones.

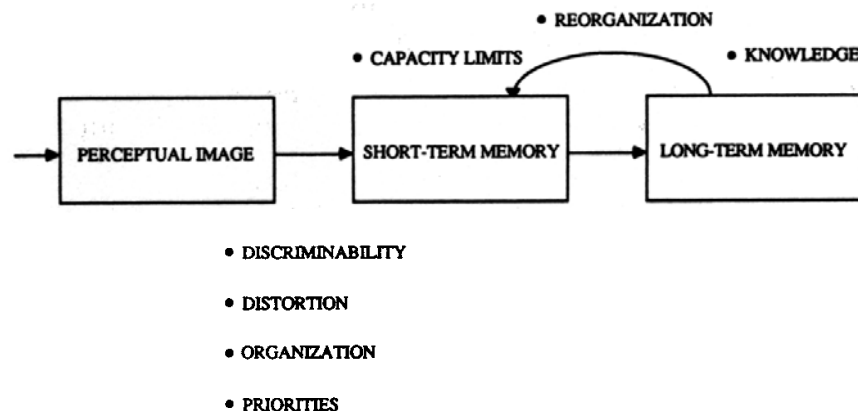


Figure 2. Three stages of visual information processing, with important characteristics noted; these characteristics are described in the text

Figure 2 illustrates the 'standard' overview of visual information processing (see Spoehr and Lehmkuhle, 1982). The boxes in the chart correspond to three types of visual processing. The leftmost box represents the output from a host of perceptual processes, which organize the input into edges and regions. This representation stores the most basic sensory aspects of the input, but does so only as long as one is gazing at the display. The middle box corresponds to 'short-term memory'. As one gazes around a display, some of the material must be held in mind while other material is apprehended. The information stored in short-term memory is usually accompanied by a conscious experience (such as of seeing a word). The capacity of short-term memory is notoriously limited: information can be held in

short-term memory for only a few seconds once one has shifted one's gaze from the stimulus, and only a small amount of information (about four groups of items) can be held in this store at the same time (see Ericsson, Chase, and Faloan, 1980; Norman, 1978; Posner, 1978). In addition to registering input from the eyes, short-term memory also receives input from long-term memory; this input consists of previously learned relevant information, which confers 'meaning' to a stimulus. Short-term memory is also important because it is the locus at which conscious reorganization and reinterpretation takes place, and its limitations severely affect what kinds of reorganization and reinterpretation can occur (as will be discussed shortly). Finally, the rightmost box of Fig. 2 represents 'long-term memory'. This memory stores a huge amount of information for an indefinite amount of time; one's childhood memories, telephone number, and the name of a favourite book are all stored here, as well as one's knowledge of arithmetic and how various types of graphs (e.g. line vs bar) serve to communicate information.

In Fig. 2 are schematized a number of properties of our visual information processing systems that affect reading charts and graphs (along with all other visual stimuli). Four of these properties pertain to how information is transferred from the perceptual input to short-term memory (and hence into awareness). These properties affect our ability to perceive syntactic properties of displays, in the following ways. First, if the stimulus is too small or does not contrast enough with a background, one will simply fail to see it. The *discriminability* limits of the system must be respected if any further processing is going to take place. Second, there are well-known systematic *distortions* in perceiving size and other properties of objects. For example, if one estimates the relative areas of two circles, one is very likely to underestimate the size difference. These distortions are reasonably well understood and can be avoided or compensated for in a display (e.g. see Stevens, 1974). Third, some aspects of a stimulus are given *priority* over others; we pay attention first to abrupt changes of any sort (e.g. heavier marks, brighter colours). Fourth, stimuli are *organized* into coherent groups and units by the time we become aware of them. Much of this organization is 'automatic', not under voluntary control, and is determined by reasonably well-understood properties of stimuli (e.g. proximity of elements). The grouping imposed by these automatic processes must be respected if a chart or graph is to be seen the way a designer intends.

Given that information is represented in short-term memory, the next constraint we must consider is the *capacity limit* of short-term memory. This constraint affects our ability both to integrate syntactic information and to hold semantic information in mind. That is, if too much information must be held in mind at once, a person will be unable to perform a task. Thus, the complexity of a display will be a major factor in determining its comprehensibility. Once information in a display is in short-term memory, it can be *encoded* into long-term memory. That is, it can be compared against previously stored information and categorized. Once categorized, one knows more about the stimulus than is apparent in the input itself. Factors that affect this process affect our ability to extract meaning from the display.

Finally, in long-term memory the major constraint is a person's *knowledge*. The way a display will be interpreted, both at the level of semantics and pragmatics, depends on which stored information is most closely associated with the way the stimulus properties of a display are categorized. If a person does not know the meaning of a word, or of a pattern of lines forming the framework of a display, he

or she will have trouble associating the display with the correct stored information. In addition, if a line is categorized as 'steep' it will be taken to represent a 'sharp rise' in prices or whatever; if it is categorized as 'shallow' it will be taken to represent a 'slow rise' (even if it is the same information, just graphed on different-shaped axes). Furthermore, knowledge of the task at hand can have important consequences: if the initial organization of the display does not help one to interpret the display, knowledge of the task can lead one to consciously *reorganize* the pattern, using information in long-term memory to reconstrue the stimulus—which in turn leads to a new attempt to interpret the pattern against other stored information. For example, if one sees a Star of David, one will probably organize it as two overlapping triangles. If asked whether there is a hexagon in the pattern, one will have to reorganize the pattern before seeing the hexagon in the middle.

All of these activities are relevant whenever one is trying to interpret what one sees. The details of the representations and processes have yet to be specified, but the outlines of these basic forms of processing now seem relatively clear. We certainly know enough about each type of operation to apply this knowledge to the design of visual displays. The basic message for display design is straightforward: an effective display must be easily encoded and comprehended by the visual information-processing system. The analytic scheme to be introduced shortly is in part a systematic way of discovering whether a given display has violated this constraint. If so the scheme is designed to reveal how a display offends our processing abilities and which abilities have been compromised.

Symbol systems

The second foundation of the 'acceptability principles' is Goodman's (1968) theory of symbols. Some aspects of charts and graphs have little to do with the operation of the information-processing system. They have to do with the very nature of how marks serve as meaningful symbols. In the ideal case a chart or graph will be absolutely unambiguous, with its intended interpretation being transparent. One way to think about this sort of unambiguity is in terms of mappings between symbols and concepts. If the display is treated as a complex symbol, then we want a unique mapping between it and one's interpretation of it. Goodman (1968) has characterized systems that have the property of unique *bidirectional* mapping between a symbol and concept as being 'notational'. English, then, is obviously not a notational system because ambiguous words or sentences are possible, whereas a music score is notational. Notational systems are stronger than we need here. In such a system there is not only a single way of interpreting a given symbol, but there is only one symbol that can be used to specify any given piece of information. Our requirement here is less stringent: given a symbol, there should be only one way to interpret it. Thus, for present purposes, there are two important uses of the basic ideas underlying notational systems.

First, we are concerned with the *within-level mappings*, which specify how marks in a chart or graph are paired with other marks, composing a complex symbol; this is especially important when indicating how labels should be paired with different lines. Second, we are concerned with *between-level mapping* between the marks on a page and the interpretation of their meaning. It is important that the lines on the page be read as intended and have the intended effect on the reader.

In Goodman's scheme the first distinction of importance for present purposes is between a 'mark' (also called an 'inscription'), a 'character class', and a 'compliance class'. A mark is a configuration of lines, such as 'A'. A character class defines which groups of marks will be classed as equivalent, such as 'A' and 'a'. A compliance class is the referent, the semantic interpretation, of the character class, such as 'first letter of the alphabet'.

The other useful (for present purposes) ideas from Goodman stem from his formal requirements for a notational system. We shall draw on these to formulate some of the acceptability principles, and thus it is worth considering them briefly. Goodman specifies five requirements for a notational system. Two of these requirements are syntactic, concerning only the properties of marks and characters, and the other three are semantic, pertaining to the interpretation of marks. For present purposes, four of the five requirements are most useful.

The two syntactic requirements are simply put. First, a given mark should map into only one character class. Goodman calls this property 'syntactic disjointness'. Second, one should be able to decide into which character class a given mark falls. Goodman calls this property 'syntactic finite differentiation'. It is important for present purposes to note that this second requirement can be easily violated. Consider an example where lines of different lengths are used as marks and where any difference in length, *no matter how tiny*, affects the character class into which the mark is mapped. Now, in this case between any two marks an infinite number of other marks exist, and so too with any two characters. Given that no physical measuring instrument is infinitely precise, this kind of situation violates the requirement of 'syntactic finite differentiation'; one cannot decide precisely which character class a given mark signifies. In this case the representational system would be called 'syntactically dense'. An example of a syntactically differentiated system is a digital clock, where every reading on the clock (i.e. every mark) is distinctly identifiable. Provided that a.m. and p.m. are distinguished, each mark maps into only one character class, and hence the system is also syntactically disjoint; if A.M. and P.M. are not distinguished, the system is syntactically differentiated but not syntactically disjoint. An example of a syntactically disjoint but syntactically dense system is a 24-hour dial clock with no tick marks. Now every position of a hand is a different mark, which signifies a different — although not uniquely decidable — character (time).

The semantic properties of a perfect 'notational system' are concerned with the way in which one interprets the meaning of marks; in Goodman's terms they are concerned with the way in which character classes are mapped into compliance classes. The two semantic properties of interest parallel the syntactic ones. First, each character class should map into one and only one compliance class. A notational system has the property of 'semantic disjointness'. The compliance classes in such a system do not overlap; they do not share characters. Second, one should be able to identify the compliance class into which a given mark should be placed. That is, a notational system has 'semantic finite differentiation'. If one cannot decide which interpretation a mark should be given, the system is 'semantically dense'. So, for example, a digital clock with a.m./p.m. indicated is semantically disjoint and semantically differentiated because each reading has only one interpretation and the interpretation is identifiable, respectively. In contrast, an axis representing

assign a precise meaning to any given reading (because between every two readings are an infinite number of other ones, precluding precise assessment of an individual reading).

In summary, marks intended to signify different symbols should be distinguishable, every mark should have an interpretation, and the interpretations should be distinct. We are interested in identifying cases in which there is a failure to have an unambiguous mapping between marks and meanings in a chart or graph. One reason for such a failure is ambiguity about the composition of multipart marks, which involve part-to-part correspondences ('within-level mappings'). But even when the syntactic structure of a mark is clear, it can still fail to map into a single concept. In all cases, when such a problem has been identified it can be ameliorated by adding to or changing the marks used in the display: even when a label is ambiguous, a new word or two can be substituted.

DIAGNOSING VIOLATIONS OF ACCEPTABILITY PRINCIPLES. I: THE SYNTACTIC ANALYSIS

In each of the levels of analysis we ask a number of questions that should be easily answered if the graph is unambiguous. These questions are designed to reveal violations of 'acceptability principles'. Because the system is set up to reveal violations of these principles, we will begin the discussion of each level of analysis with a brief overview of the relevant principles themselves. Following this we will consider the actual mechanics of generating a description of a chart or graph.

Acceptability principles

The syntactic principles describe constraints on how lines and regions are detected and organized; these principles grow out of discoveries about human visual information processing. A syntactic problem is not tied to the lines' having a specific meaning, but hinges on problems with extracting any meaning from them. If these principles are violated, one either will not be able to extract some information from a chart or graph (without, perhaps, the aid of a magnifying glass); will systematically distort information when reading it; will tend to have difficulty organizing it correctly; or will find it difficult to hold the number of relevant lines in mind at once.

Three broad classes of acceptability principles are posited at the syntactic level that cannot be violated if a chart or graph is to be completely effective. Each of these classes contains a number of specific principles which themselves have specific aspects.

Perceptual apprehension

The visual system imposes numerous constraints on how marks can be used to convey information in charts and graphs. The first set of principles bears on how lines, colours, and regions are accurately discriminated—which is a necessary

prerequisite for further processing. Two principles bear on the process of discriminating marks:

- (a) *Adequate discriminability.* Variations in marks must be great enough to be easily noted. This principle has two aspects. (i) *Relative discriminability:* two or more marks must differ by a minimal proportion to be discriminated. The laws governing the size of this difference have been worked out for many types of marks and these laws comprise this principle (see Stevens, 1974). (ii) *Absolute discriminability:* a mark must have a minimal magnitude to be detected. This 'absolute threshold' has been computed for many types of marks (e.g. see Smith, 1979).
- (b) *Perceptual distortion.* Marks should be used that are perceived veridically. In addition to producing illusions (see Dodwell, 1975; Frisby, 1980), the visual system often systematically distorts the magnitude of marks along various dimensions (such as area and intensity). This distortion is described by the value of an exponent in a power law, as developed by S. S. Stevens and his colleagues (e.g. Stevens, 1974). For some purposes (e.g. conveying the general impression), it may be appropriate to alter marks intentionally to compensate for the distorting properties of the visual system (which, for example, make increases in area seem smaller than they are; see Cleveland, 1985; Teghtsoonian, 1965).

Principles of perceptual organization

Marks should be drawn so that they are perceptually grouped appropriately. That is, marks are rarely seen as isolated dots on a page, and well-known principles describe the main factors that determine which marks will be grouped together into a single perceptual unit. If these principles operate to group together elements of a display inappropriately, the display must be changed. In addition, good displays should make use of easily seen patterns, exploiting our ability to apprehend changes in slope, groupings and the like.

Gestalt laws. The Gestalt psychologists discovered almost 120 distinct laws that describe how forms are perceptually organized (see Hochberg, 1964; Kaufman, 1974). The more important laws (for present purposes) can be summarized by four general principles:

1. *Good continuity.* Marks that suggest a continuous line will tend to be grouped together. So, a series of marks such as '-----' are seen as forming a single line, not as a series of isolated dashes.
2. *Proximity.* Marks near each other will tend to be grouped together. So, '××× ×××' is seen as two units, whereas '×× ×× ××' is seen as three.
3. *Similarity.* Similar marks will tend to be grouped together. So, '!!!○○○' is seen as two units.
4. *Good form.* Regular enclosed shapes will be seen as single units. So, '[' is seen as a unit, whereas '[' is not.

Dimensional structure. For some types of visual dimensions, values on one dimension are grouped with values on another. That is, marks vary along a number

of dimensions, such as hue, size, height, and some of these dimensions cannot be processed independently of others. For example, it is impossible to see the hue of a mark (i.e. its shade of colour, roughly) without seeing its saturation (i.e. the deepness of the colour, roughly). Thus, some sets of dimensions are organized into single units (including height and width of a rectangle) whereas others (such as hue and height) are not. Thus, the width of a bar will influence one's assessment of its size, even when only the height is relevant. The dimensions that are 'stuck together' in processing are called *integral* dimensions, and the ones that are processed independently are called *separable* dimensions (see Garner, 1974).

Principles of processing priorities and limitations

The visual processing system has quantitative and qualitative limitations. Partly because only a limited amount of information can be held in short-term memory at once, some marks will be given priority over others. The information conveyed by these marks should be central to the display's message. Further, some kinds of comparisons are difficult for the visual system to perform, and hence a display should not require use of them (cf. Cleveland, 1985). These facts are the basis for two kinds of principles:

Processing priorities. Marks should be chosen to be noticed in accordance with their importance in the display. Some colours, weights of line, and sizes are noticed before others. For the most part we do not have formal rules for determining which these are, but instead rely primarily on a general principle: the visual system is 'a difference detector'. Any sharp contrast will draw attention. In addition, some stimulus properties have been determined empirically to be 'salient' (e.g. all other things being equal, a yellowish-orange is noticed before a deep blue; see Frisby, 1984). Physical dimensions of marks should be used to emphasize the message; they should not distract from it (e.g. by making the background too prominent).

Processing limitations. These principles fall into two categories. (a) *Finite capacity:* displays should contain no more than four to seven perceptual groups. At the most, only about seven perceptual groups can be seen at a single glance, and only about four can be held in mind at once (e.g. Ericsson, Chase, and Falcon, 1980). (b) *Unit binding:* perceptual units should be chosen so that separate perceptual units correspond to pieces of information to be sought by a reader. That is, displays should not require readers to decompose perceptual units in order to extract specific information, as occurs if single points along a line must be interpreted. It is more difficult to see and compare parts of a single perceptual unit than it is to see and compare entire units (see Clement, 1978). For example, a triangle is easy to see in a Star of David, whereas a parallelogram is not (Reed, 1974).

The within-level mapping principle. Portions of the chart or graph that are meant to correspond to other portions of the chart or graph should do so unambiguously. Much of the meaning of a display arises from how the marks are organized into a complex symbol. The key, for example, should clearly indicate how labels are paired with different components of the specifier. This principle is distinguished from the principles of perceptual organization in the following way: when a

corrected by rearranging marks already in the display (by repositioning lines and the like). In contrast, when the within-level mapping principle has been violated, new marks must be added (e.g. lines or arrows connecting parts); a necessary ingredient is missing when this mapping principle is violated.

Applying the analytic scheme

The analytic scheme is designed to reveal violations of the principles that impair the effectiveness of the chart or graph. In order to do so, however, one must generate a description of exactly what is out there, exactly how a chart or graph is composed. Thus the scheme requires one to engage in two distinct activities. First, one actually describes the chart or graph. This is especially the case at the syntactic level. Second, one asks questions about the description, checking to ensure that it is unambiguous and immediately obvious. If not, one or more of the principles has been violated. When a difficulty has been encountered, one simply checks the various principles, considering whether each has been violated. The level of detail of the description proper is motivated by the kind of information one will need later on to assign a semantic interpretation, and then a pragmatic evaluation, of the chart or graph—again with an eye towards discovering violations of the respective types of acceptability principles.

At the outset we ask whether the chart or graph is composed of a number of subcharts and graphs. That is, we ask whether there is more than one chart or graph present and whether there are systematic relations among the information in each. If so, the scheme is applied to each one separately and then to the set of charts and graphs together. The following is a description of how the scheme is applied to a single-panel chart or graph. In order to make this description more concrete, we will work through a simple example; the Appendix presents a complex multi-panel example. The example considered here is illustrated in Fig. 3, which was included in a Request for Proposals from the National Institute for Education issued in 1979 (it was described as being included in a Department of Transportation booklet). In each section of the following discussion we first will consider the descriptive procedure and then will turn to the results of applying it to Fig. 3.

Before beginning the analysis it is important to point out the distinction between a 'violation' of a principle and a 'flaw' in a display. Not all violations will necessarily impair reading the chart or graph at the level of detail intended by the designer. Violations reveal difficulties in extracting *all* of the information potentially available in a display, but this may be far in excess of that required to use the display as intended. We cannot distinguish between a violation and a flaw on the basis of the display alone; the distinction depends on knowing the purpose and context of the display (pragmatic considerations).

Background

We first ask whether there is a background and, if so, we describe it. A background extends beyond the framework and does not actually help to convey the information in the display; removing the background would not impair how the chart or graph functions to represent information. We ask whether it is easy to detect the lines and regions delineating the other constituents over the background (if not, it violates the principle of adequate discriminability).




Figure 3. A graph used to illustrate the analytic scheme, analysed in the text

Example

The background in Fig. 3 is blank white

Framework

We ask whether there is an outer and inner framework, and ask the following questions (as appropriate).

Outer framework. We define the outer framework as the set of lines that serve to define the general entities that are addressed in the display. We ask: What are the elements? Are they lines? Are they regions? What is their shape, weight, and colour? Are they clearly discernible? If lines function as axes, are they dense or differentiated? We note whether the framework and its individual component parts are easily identified.

We next ask: How are the elements organized? Are the relations among the different parts clear? Does the juxtaposition and composition of the elements result in an appropriate perceptual organization? How many elements must be held in mind at once in order to organize them into the entire framework?

Inner framework. We next apply the same questions to the inner framework, if one is present. In addition, we ask whether the inner framework masks or distracts attention from the elements of other constituents of the display.

Organization of inner and outer frameworks. Next we consider the organization of the two frameworks, if both types are present. In particular we ask how the similarity, proximity, and continuity of framework elements imply organization.

Following this, we ask a number of general questions about the entire framework: Does the framework represent two-dimensional or three-dimensional space? Are quantities undistorted? Is colour employed in the framework; if so, what is emphasized? If line weights are varied, what is emphasized? (This will be important later in our pragmatic analysis.) Which axis is longer?

Example

Outer framework

Elements. The other framework is composed of two vertical straight lines, which are syntactically dense, and two horizontal straight lines, which are syntactically dense. The lines are black, medium weight.

Organization. The elements of the outer framework are connected to form a rectangle, with the vertical axis being longer.

Inner framework

Elements. The elements are seven straight vertical lines, which are syntactically dense and light black.

Organization. The lines are spaced evenly.

Organization of inner and outer frameworks

The parallel lines of the inner framework are connected to the top and bottom horizontal lines of the outer framework, terminating at those lines.

Overall

The display is in two dimensions, with no apparent distortions; no part of the framework is emphasized. The vertical axis is longer.

Specifier

We begin by isolating the class of visual continua used to represent information. We then describe how shape, size variations, colour and texture are used. We go on to ask: What are the elements used to compose the specifier? Is it clear whether parts are overlapping or contiguous? Are there few enough elements to keep them in mind at once? Are variations used to convey information clearly distinguishable? Do all attention-drawing variations convey information?

We then ask: How are the elements organized? Is the organization clear? Does the specifier clearly imply a two- or three-dimensional shape?

Example*Elements*

There are five rectangles, divided into black and white portions by a vertical line, with the left side being black; *or*, the elements are five black rectangles and five white rectangles. It is not clear whether parts are overlapping or contiguous. This is our first violation of an acceptability principle; violations will be indicated as follows:

VIOLATION. Principles of perceptual organization. It is not immediately clear whether there are bars divided into two segments, or separate white and black bars juxtaposed. That is, it is not clear at first glance whether the white rectangles are to be viewed as representing just their own width, or their width plus the width of the black rectangles.

Organization

The bars are spaced one above the other with the leftmost ends aligned. Each bar is divided into a black part and a white part. Alternatively, each white rectangle abuts the rightmost end of a black one; these pairs of rectangles are spaced vertically, with the leftmost ends of the black rectangles being aligned.

Labels

We next consider three kinds of labels independently, and then turn to an analysis of the relations among them. The following questions are asked about all labels, including those in the title and key.

Alpha labels. Are letters or words used as labels? If so: Are they clearly readable? How many are present? What is the size, colour, and weight of the typeface used for each of the labels? If colour variation is used, are variations clearly discriminable? How do labels group together? Is the perceptual grouping congruent with the intended interpretation?

Numeric labels. Are digits used as labels? If so, ask of them the same questions asked of the alpha labels.

Depictive labels. Are pictures used as labels? If so, are they clearly identifiable? How many are present? Are they all the same size? How do these labels group together? Is the perceptual grouping congruent with the intended interpretation?

Title. We pay special attention to the title, asking first whether there is one. If so, we ask: Is the title clearly discriminable as a title? What is the relation of the title to other elements of the chart or graph? Does it perceptually group correctly, or does it appear to label only a local part of the chart or graph?

Key or legend. Next, we consider whether there is a remote legend or key. If so, we ask: Is the information clearly readable? Does the legend clearly separate itself from one element of the chart or graph? Is there little enough material to be easily held in memory?

Organization of the different types of labels. We next ask about the organization of the labels, examining the relations among each type of label. We ask: Do

perceptual organization principles result in a correct organization of the labels? In particular, does the appearance of the typeface cause the reader to group labels properly, or does a dissimilar typeface cause the reader to separate labels that should be grouped together? Does proximity of labels cause one to group them properly? Are labels ordered in such a way that one groups them properly?

Example

Alpha labels

The words are readable. Three sizes of typefaces are used, they will be referred to as 'large', 'medium', or 'small'. The same-sized font is used to label the horizontal axis, vertical axis, and key.

Numeric labels

There are five numbers in a vertical column at the left; five in vertical column at the right; and six in a horizontal row at the bottom. The font used for the row at the bottom is smaller than the font used elsewhere.

VIOLATION. Principle of processing priorities. The numbers at the bottom are small, leading one to notice them only after the other numbers.

Depictive labels

Pictures are not used, but the key is depictive, consisting of a black rectangle and a white rectangle. The two rectangles in the key are adjacent to each other.

Title

A very large font is used. The title is at the top of the display and centred. Its large size makes it distinct from the other labels and provides it with sufficient scope to encompass the entire display.

Key or legend

It is clearly separated, and it can easily be held in short-term memory.

Organization of the different types of labels

Alpha and numeric. At the left is the label 'Miles Per Hour', above the column of numbers. At the right is the label 'Total Distance', above the column of numbers. At the bottom is the label 'Feet', to left of row of numbers.

VIOLATION. Principle of perceptual organization. The size of the marks used as labels on the vertical axis and the size of the marks used as numbers are different, with the larger font used for the numbers, making it difficult to see them as grouped with the label.

Numeric and depictive. No cases.

Alpha and depictive. The labels are to the right of white and black bars in the key.

Organization of the framework, specifier and labels

We now turn to the overall organization of the display. A badly organized display can be unreadable because the mappings between information-bearing components are opaque.

Organization of framework and specifier. What is the relation of the framework to the specifier? Is the specifier complete?

lines of the inner framework violate boundaries of the specifier? Do perceptual organizational principles cause the reader to group the framework and specifier correctly? Is the dimensionality of the space (two or three dimensions) consistent between the framework and specifier? These questions are asked separately for the outer and inner frameworks (if appropriate).

Organization of framework and labels. The organization of the framework and each type of label is considered separately, with the following questions being considered (as appropriate): How are the labels associated with the framework and parts thereof? Are value markings indicated along the framework? If so, do the labels clearly indicate the correct values corresponding to the associated portion of the framework? Do perceptual organization principles result in a correct organization of the framework and labels?

Organization of labels and specifier. How are the labels and specifier associated? Is all specifier material labelled? If the label is remote, in a key, is the mapping from elements in the key to the specifier clear? Do perceptual organization principles result in a correct organization of the labels and specifier?

Organization of labels, framework, and specifier. Is little enough material present in a small area that it is easy to discriminate the relations among the elements of the different constituents? Do perceptual organizational principles reveal the correct relations among the constituents?

Example

Organization of framework and specifier

Outer framework and specifier. The bars abut the left vertical line, extending to right; the bars are enclosed in the frame.

Inner framework and specifier. The bars are superimposed over the inner grid. The vertical internal lines of the inner framework do not violate the boundaries of the rectangles, and they group properly with the specifier.

Organization of framework and labels

Alpha. The title is at the absolute top of the display; the labels in the key are directly above the highest horizontal line of the outer framework. The left vertical and bottom line of the outer framework are labelled. There is a label at the top left corner of the display, and another label at the bottom, with the first letter directly under extreme left end of the bottom horizontal line.

VIOLATION. Principles of perceptual organization. The 'Miles Per Hour' label is not immediately associated with the vertical scale, being equally close to the horizontal line at the top.

The 'Total Distance' label is at the upper right, above the top horizontal line; it is centred within the segment defined by the rightmost vertical line of the outer framework and first line of the inner framework to its left.

VIOLATION. Principles of perceptual organization. The 'Total Distance' label is closer to the top horizontal line than the column of numbers, and is not immediately grouped with the proper elements of the display.

Numeric. There is a column on the left, regularly spaced outside and to left of leftmost vertical line of the outer framework. There is a row on the bottom, under the horizontal lower line of the outer framework; a number appears under each line of the inner framework, but there is no number under the last internal line on the right. There is a column on the right, evenly spaced, centred between the rightmost outer line and the first internal line to its left.

Depictive. The depictions in the key are above the horizontal line defining the top of the framework.

Organization of labels and specifier

Alpha. No cases.

Numeric. There is a 1:1 alignment between numbers in the column at the left and the bars; the grouping here is clear. In contrast, the 1:1 association between numbers in the column at the right and bars is not immediately clear.

VIOLATION. Principles of perceptual organization. Numbers are not clearly grouped perceptually with the appropriate bars.

Depictive. There are black and white key labels in same order as black and white portions of bars.

Organization of labels, framework, and specifier

Description of the pair-wise relations among the constituents is sufficient; no special problems emerge from the relations among the constituents taken as a whole.

DIAGNOSING VIOLATIONS OF ACCEPTABILITY PRINCIPLES II: THE SEMANTIC ANALYSIS

We again begin by briefly outlining acceptability principles, and then turn to the method of describing and analyzing the semantic information in a display. In addition, it is at this level that the differences between charts and graphs *per se* become important, requiring us to develop two different sorts of semantic interpretations, one based on qualitative relations and other based on quantitative relations.

Acceptability principles

Let us consider four classes of semantic principles. These principles describe constraints on the ways patterns of marks are interpreted. A semantic principle is tied to how a specific meaning can be extracted from a configuration of marks. These principles were derived partly through a review of the literature on how people spontaneously describe visual displays, and partly by generalizing from psycholinguistic phenomena (Clark and Clark, 1977). If the semantic principles are violated, a reader will find the display confusing.

Principle of representativeness

The intended meaning of a mark should correspond to the spontaneous interpretation of it (see Jolicoeur, Gluck, and Kosslyn, 1984). All marks have a preferred interpretation. Thus, labels should name words that are indicative of the class (including the correct connotations) and pictures should depict appropriate objects (a picture of a penguin-like bird should not be used to label birds in general).

Principles of congruence

This principle has three aspects.

Surface compatibility. The appearance of the lines and regions themselves should be compatible with their meanings. For example, people have trouble reporting the colour of the ink used to print a word if the words themselves name other colours (e.g. if the word 'red' is printed in blue ink; this is known as the 'Stroop effect'). Similarly, if 'left' and 'right' are labels along the x axis, the 'left' label should be physically to the left (this constraint is sometimes violated in the literature on the specializations of the cerebral hemispheres). In general, larger areas in the display (mentally categorized as 'larger') should represent larger quantities, faster-rising lines should represent sharper increases, and so on. Colour presents a special case of surface compatibility: although the physical spectrum is ordered in terms of wavelength, colours are psychologically arranged around a circle, not a line (Kaufman, 1974). Thus it is difficult to align colours with quantities. If colour is used, intensity or saturation should covary with hue so that lighter colours correspond to higher values.

Ordering. Pairs of words should be ordered in accordance with natural usage. We say 'bread and butter', not 'butter and bread'. The shorter, less stressed word goes first; if it does not, the phrase will not immediately correspond to a stored memory (see Pinker and Birdsong, 1979).

Markedness. The 'unmarked' term should be used to label the dimension. That is, some words not only name a pole of a dimension but the dimension itself. We say 'how high is that?' without implying necessarily that it is high; but if we say 'how low is that?' we imply it is low. The term that implies a specific value is called the *marked* term, and should not be used to label the dimension itself—if it is, it will mislead the reader (see Clark and Clark, 1977).

Principles of schema availability

In order for a chart or graph to be comprehensible, a reader must have the requisite concepts. That is, a 'compliance class' is in fact something in a reader's head. The reader must know both the individual concepts and the general idea of how a particular graphic design conveys information. We can distinguish among three aspects of these requirements.

Concept availability. A chart or graph should make use of concepts that are likely to be possessed by the intended readership.

Cultural convention. The conventions of a reader's culture should be obeyed when drawing an effective graphic display. So, for example, in the Western world the colour red should not be used to represent 'safe' areas, and green should not be used to signify 'danger'. Similarly, time should increase going left to right or bottom to top.

Familiarity. Information should be presented in a graph type that is familiar to a given readership.

The between-level mapping principle

In the course of describing the semantics of a display we are faced with describing how the marks map into semantic classes. Thus, it is at the point of formulating the semantic description that it is important to ensure that (a) every meaningful difference in the value of a variable is represented by detectable differences in marks, and (b) every mark has one and only one meaning. Ambiguous or missing marks violate this principle and require an alteration at the level of syntax.

Types of semantics

The specific information conveyed by a chart or graph depends on the interrelations among different components of the display. Thus, in evaluating how well a display conveys meaning we will decide which of two general kinds of rules of combination is used, and then will consider whether one can derive the appropriate information. One class of rules is appropriate for graphs. These rules specify quantitative relationships between two or more values on two or more scales. The other class of semantic rules of combination is appropriate for charts. These rules specify the qualitative structure or organization of entities.

Quantitative relational information

Perhaps the best way to present the formal properties of this aspect of graphic semantics is in tabular form. Thus, Table 1 relates values on two scales to each other. We consider all possible combinations on nominal, ordinal, interval and ratio scales except the nominal–nominal combination (which falls into the second class of rules). Nominal scales are not ordered, with numbers being used as names (as on the backs of football players); ordinal scales are rank-ordered according to quantity, but the actual magnitudes of differences are not specified (as in the first, second and third place winners of a race); interval scales are ordered so that the magnitudes of differences mean something, but ratios of numbers do not (as in Fahrenheit degrees, which have no 'natural' zero); finally, ratio scales have numbers that are ordered so that the magnitudes of differences are important and ratios can be computed (as in degrees Kelvin, where 10 degrees is twice as hot as 5 degrees—which is not true with Fahrenheit degrees). In addition to providing an example for each in the table, there are examples of the kinds of information available in each case. Extensions to n -dimensional cases follow in a straightforward manner from the simple two-dimensional cases considered here.

The information content of a graph can then be assessed by interpreting the individual axes, noting the scale types and how points are paired via the specifier(s), and then using the taxonomy in the table to derive the appropriate

Table 1

Scale I	Scale II	Example	Examples of information available
Nominal	Internal	States by rank in coal production	Allow $N(N-1)/2$ inequalities statement to be made.
Nominal	Ratio	States by coal production	Map N things into an infinite number of classes. Comparison of differences.
Ordinal	Ordinal	Ranked oil production by ranked coal production	Nonarbitrary zero point. Ratio comparison of items possible.
Ordinal	Internal	Rank in classes by achievement test score	Relative ranks. Comparison of disparities in ranks
Ordinal	Ratio	Ranked coal production by oil production	Comparison of differences for different ranks.
Interval	Interval	Math achievement score by English achievement score	Relative difference or ratios for different ranks.
Interval	Ratio	Achievement scores by hours/week of TV watching	Differences on both dimensions. Difference in one dimension as a function of difference in other.
Ratio	Ratio	Coal production by oil production	Ratio amounts along one dimension as a function of differences in the others.
			Absolute amounts, differences in amounts, ratios, for both dimensions; absolute amount of one dimension as a function of absolute amount of the other.

kind of information one should be able to extract. If this information is not clear, there may be a failure of within-level mapping (the specifier is not clearly serving to pair points on the framework) or a failure of between-level mapping (a part may be missing). (Violations of many other principles can also distort the relationship, depending on problems in seeing the specifier or organizing parts of it correctly.)

Structural/organizational information

A computer flowchart, an organizational chart for a government agency, and a family tree do not relate values on quantitative dimensions. Rather, they specify the relationships among discrete members of a set. This sort of information can be described using the following three general criteria. These criteria are orthogonal.

The first criterion is whether the links between entities are *directed* or *non-directed*. Elements of the framework (i.e. marks indicating an individual member of the set) can be related together either by symmetrical or by asymmetrical relations. For example, in a kinship diagram the vertical links of the tree are directed, indicating who is the parent of whom (an asymmetrical relation). The horizontal links, such as 'sibling of' (a symmetrical relation), are non-directed.

The second criterion is how many *types of links* are used. More than one kind of relation may be used in a graph. In a kinship diagram, for example, 'cousin-of' and 'brother-of' may both be present. In a computer flowchart, only a single type of arrow—indicating which operations follow one another—occurs.

The third criterion concerns the *type of mapping* used. There are three classes of mappings: One:One, Many:One (or One:Many) and Many:Many which we will consider in turn:

One:One mappings. In this case links in a chart might indicate how husband and wife pairings occur by drawing lines connecting points representing the location of each individual at a cocktail party.

Many:One or One:Many mappings. In this case it is important to consider separately directed and non-directed links. With directed links, inclusion relations may be indicated by a Many:One mapping such as occurs in a hierarchy where many objects are organized under a superset. With non-directed links, collateral relations are indicated. If all diplomatic relations were symmetrical, links on a map illustrating the diplomatic relations of any one country would represent this sort of mapping.

Many:Many mappings. In this case the multiple affiliations of a number of different objects can be represented. For example, a chart might represent different social classes by a drawing of a typical number of each, and might represent different social institutions by drawings of typical buildings (e.g. a church or a bank). Lines could connect the people to the institutions to which a majority of the represented class belong.

In charts, then, the nature of the mapping must be clearly indicated by the specifiers. Too many arrows can obscure mappings among elements, as sometimes happens in tangled organizational charts. Directionality and specific meaning (achieved via labels) should be clear. In actually describing a chart we are careful to consider what kind of information is being conveyed (hierarchical, relational, etc.).

We then consider whether the specifiers effectively convey the meanings of the

Applying the analytic scheme

As in the syntactic analysis, we decompose the problem of describing the semantic content (the literal meaning) of a graphic display into four parts: characterizing the background, the framework, the specifier and the labels. In this case, however, we do not examine the interrelations among the basic-level constituents as distinct from the constituents themselves; as was just noted, in large part the meaning of the display derives from the interrelations among the constituents, and thus we consider the interrelations as we consider the individual constituents. As before, we look for violations of the acceptability principles that come to light when the display is being analyzed. It is of interest that many of the problems noted in the syntactic analysis come home to roost here in a different guise, which is not surprising given that the semantic interpretation rests on the nature of the marks themselves.

Background

If the background is patterned, the meaning of the pattern should be consistent with the information presented in the chart or graph. If background figures are present, do they reinforce the meaning of the chart or graph? Are the elements of the background unambiguous? Are all parts of the rest of the display readily apparent, or do parts of the background occlude or mask parts of the framework such that information is lost?

Example

There are no background figures in Fig. 3.

Framework

As before, we ask about the outer and inner framework separately

Outer framework. We begin by asking whether the display is a chart or graph. In so doing we consider whether the meaning of the elements of the framework is unambiguous. We note whether any part is missing and not clearly implied. For graphs, we next assess the scale type along each axis and note whether the semantic scale is clearly indicated syntactically. For instance, if the scale used on the axes of a graph is syntactically dense, the semantics—the actual scale being represented—should be semantically dense. We next note the extent of the dependent measure scale, attending not only its range, but to the baseline. This may prove important in the subsequent analysis of the pragmatics of the chart or graph.

In addition to these questions, we check whether the lines that compose the framework depict an object (this is quite common in many popular magazines). If so, we ask: What is depicted? Is the meaning clearly evident, and is the depicted object clearly representative of the class of objects being depicted?

Inner framework. We ask what is the purpose of the inner framework, and whether the increments it defines are clearly evident.

Example

Outer framework

Figure 3 is a graph, with the outer framework representing a Cartesian coordinate space, with the dependent variables along the y axis and the independent variable along the x axis.

Vertical axis. The vertical axis represents a ratio scale, with the origin at the top of the line. Although this scale is semantically dense, it has been differentiated into five discrete values with values increasing as one descends down the line. The scale ranges from 0 to 300. We cannot decide whether or not the division into discrete values is a violation without knowing the purpose of the display; however, the ordering of the values is problematical at a semantic level:

VIOATION. Principles of congruence. Larger values should be indicated by physically higher marks.

Horizontal axis. The horizontal axis is a ratio scale, with the origin at the left and values increasing as one moves to the right.

VIOATION. Principles of schema availability. The origin of the two axes in a Cartesian space is usually the same point (the lower left intersection of the axes), which is not true here.

Inner framework

The vertical lines serve to help one assess amounts along the bars; they clearly mark off increments of 50 feet.

Specifier

The meaning of the specifier is derived from how it maps one part of the framework onto another. For a chart this consists of grouping discrete entities. For a graph this consists of pairing values on one scale with values on the other, making available the appropriate types of information noted in Table 1. We ask the following questions about the meaning of the specifier. First, what information is specified? Is it clear what kinds of relations are specified? Is every distinction in the specifiers informative? Does every conceptual distinction correspond to a visible distinction among marks? Do marks used to represent different things look more different than marks used to represent the same thing? Is the literal interpretation of the marks compatible with the role they play?

A special case of this last concern occurs when the specifier is a depiction (e.g. a graph of rising prices could have a jet plane taking off, with the contrail coming out of its exhaust serving as the function); if so, we ask: Are the depictions clearly representative of the compliance class in question? One would not want a picture of a potato to stand for 'plant life', for example (since potatoes are hardly typical—in Rosch's (1978) sense—plants).

Example

The specifier is serving to map discrete values on the vertical axis to continuous values on the horizontal axis, although both are ratio scales. Two functions are plotted: the length of the white rectangle represents average braking distance; the length of the black rectangle represents average reaction distance. Each pair of rectangles represents a discrete and different speed. The relationship between average braking and reaction distance is implicit in the relationship between the length of the black and white portions of the bars.

Labels

In the same vein, the labels along the axes should be compatible with the actual scale being used and with the markings along the axes; the numbers spaced along

the axis should suggest the correct scale type. For each type of label, we ask the following questions:

- (1) *Alpha labels*. What are the labels? Are the words ambiguous? Are the meanings of all the words representative of the class being indicated?
- (2) *Numeric labels*. What are the units? Are they clearly specified? Are the units familiar?
- (3) *Depictive labels*. Are pictures used as labels easily identified? Are they familiar to the intended readers? Are the marks used to depict clearly representative of the concept represented?

Organization of labels. We next ask whether it is clear how the words label numerical values and depictions, and whether it is clear how depictions label numerical values (if relevant).

Example

Alpha labels

There are English words labelling the values of units on the axis, words labelling two small bars in the key, and a label for the total distance column. English words also label the graph as a whole.

VIOLATION. The between-levels mapping principle. The failure to include the word 'distance' on the alpha label associated with the left bar in the key is misleading; no contrast is intended with the right label.

VIOLATION. The between-levels mapping principle. The word 'braking' is ambiguous; it could refer to total stopping distance.

Numeric labels

Distances are specified in feet, and speed is specified in miles per hour. Total braking distance is also provided. The units are clear and likely to be familiar to U.S. audiences.

Depictive labels

No pictures are used as labels. The colour of the bars in the key has no intrinsic meaning.

Organization of labels

Alpha and numeric. The words clearly label scales, and the numbers index values on the scales.

Alpha and depictive. The words clearly serve to label the meaning of the bars in the key.

Numeric and depictive. No cases.

DIAGNOSING VIOLATIONS OF ACCEPTABILITY PRINCIPLES: III: THE PRAGMATIC ANALYSIS

In a typical graph, such as that illustrated in Fig. 1, the line serving as a function is not syntactically differentiated. Indeed, the axes of Fig. 1 themselves are not syntactically differentiated. Thus, this would seem to preclude the graph's being

unambiguous; recall that syntactic differentiation is a prerequisite for unique mapping from mark to compliance class. However, one must consider what is the intended use of the chart or graph. For many purposes only a rough approximation is intended, especially when graphs are idealizations intended to illustrate some general point (such as Fig. 1). In these cases providing additional precision would distract from the point of the display. Thus the prior analyses must be regarded as producing candidates for flaws in a display, and the purpose and context must be taken into account to discover whether these amount to real problems.

In addition, as in language, not all the information humans garner from charts and graphs is dictated by the literal interpretation of the marks on the page. If the number of AIDS victims were indicated in a bar graph by increasingly higher piles of bodies, to take a grisly example, the reader would probably not simply register the literal information conveyed by the height of the column. This aspect of communication with charts and graphs has been discussed at some length by Huff (1954) in his classic book, *How to Lie with Statistics*.

The acceptability principles introduced in this section are concerned with these kinds of 'pragmatic' aspects of communication, focusing on the importance of the purpose of a display, its context, and the indirect meaning conveyed. These acceptability principles were formulated in part by constructing demonstrations in which visual properties were manipulated to produce misleading descriptions of the specifier. We also considered the rather direct correspondences that exist between graphs and language, adapting similar principles for language offered by Grice (1975).

Acceptability principles

Three classes of principles capture the pragmatic considerations we use in evaluating charts and graphs. The classes contain numerous individual principles, however, but we will not develop them here. The classes are:

Principles of purpose-compatibility

The display format should make explicit the information a reader will need to extract. This principle has three separate aspects. (a) *Appropriate amount*: A cardinal rule here is that no more or less information should be provided than is needed by the reader (cf. Grice, 1975). (b) *Appropriate content*: Displays are often used for two kinds of purposes, communication and analysis (e.g. see Bertin, 1983; Chambers, Cleveland, Kleiner and Tukey, 1983; Fisher, 1982; Kosslyn, 1985; Schmid, 1983; Tufte, 1983). Depending on the purpose and intended users, one would or would not place a premium on aesthetics, ease of interpretation, completeness, and so on. (c) *Appropriate format*: In addition, depending on the purpose, different graph formats are more or less appropriate. For example, if one wants a reader to see an interaction, a line graph is better than a bar graph: not only is a line a single perceptual unit, and hence there will be less to hold in short-term memory, but we are good at detecting slope differences—which convey the relevant information (see Pinker, 1988). Indeed, experts learn to recognize patterns formed by lines, which come to function as symbols in their own right (e.g. signifying a crossover interaction). In contrast, if one wants the reader to compare specific point values, a bar graph is better than a line graph: not only are the

discrete values indicated by the tops of the bars, but one does not have to disrupt the single unit formed by the line (as was discussed earlier).

Principles of invited inference

Although a chart or graph may convey the correct information on the semantic level, it may invite us to misread it anyway. This can be done in numerous ways: truncating scales so that small differences appear larger; varying the type of scale used (interval vs. logarithmic, for example); using inferred three-dimensional properties of a display so that we see bars as bigger than they are, and so on. Some of these principles are illustrated by Huff's (1954) observations about how to lie with statistics.

Principles of contextual compatibility

Most displays are embedded in a context, either in text or in an oral presentation. The context and the semantic interpretation of the display must be compatible or comprehension of the display will be impaired. The terminology used in the display should be the same as that in the text or presentation, and the display should be presented at the point when the reader will want to access the information.

Applying the analytic scheme

We again consider first each of the four basic-level constituents, and then turn to questions about the way they are organised. This analysis differs from the earlier ones in an important respect: the syntactic analysis resulted in a rather rich description of the chart or graph itself. This was necessary because elements of the syntax feed into the semantic properties, and hence we needed to have the chart or graph described in a way that would allow us to consider each of the semantic principles. At the level of the semantic analysis there was much less description *per se*, and only some of the semantic aspects of the chart or graph are relevant for this later pragmatic analysis. The pragmatic analysis itself, then, produces very little in the way of a further description of the chart or graph. Rather, the existing description is now rich enough, from the level at which the thickness and colour of the lines is noted to the level at which the elements are interpreted, to allow us simply to ask questions that probe for violations of specific principles. Thus, this analysis consists entirely of questions, as indicated below. These are 'leading questions' in that the answers reveal violations of the acceptability principles described above.

Background

Does the background imply information not explicitly stated in the display (e.g. as might occur if the background was a photo of war dead)? If so, are the implications of background material consistent with the message and the content?

Example

There is no background in Fig. 3.

Framework

We ask questions to determine whether the form of the framework leads the reader to extract the intended message easily. Is there a truncated axis? Does this mislead

in ways not intended by the designer? Are scales transformed? Is this compatible with the point of the display? Are value markings sufficient for the intended purposes? If there is an inner framework, are the distinctions demarked fine enough to access the necessary information, or are they too fine, distracting from the relevant divisions? If the framework is also serving to depict, does the meaning of the depiction help or hinder understanding the content of the chart or graph?

Example

The axis is not truncated, nor are the scales transformed. Without knowing the intended purposes of the display we cannot decide whether the value markings indicated are sufficient for intended purposes. If the display were intended to instruct drivers about the time to brake, it seems adequate.

Specifier

We ask: Is the most appropriate type of specifier element used? For example, are lines used to convey trends, and bars to convey point values? Is more or less information present than is necessary to answer the relevant questions? Are some equivalent elements made to appear more important than others (by colour, width of lines and so on)? If so, is this appropriate given the point of the chart or graph? If the specifier depicts information, does the meaning of the depiction help or hinder understanding the content of the display? If a three-dimensional system is used, do perspective effects distort the apparent sizes of different elements?

Example

Most of the questions cannot be answered without knowing the purpose of the display. For example, the 'average reaction time' is darker and more salient; we cannot know whether this is appropriate without knowing the purpose of the display. If the display is to inform drivers about the relation between reaction time and braking distance, and these drivers are unlikely to have considered the relevance of reaction time, it would be appropriate to emphasize these data.

Labels

Are words consistent with the terminology of the text? Are labels complete? Are labels cryptic or verbose?

Example

It is difficult to answer these questions without knowing the purpose of the display. However, given the inclusion of the inner grid lines, the reader presumably was intended to ascertain more than simply the 'Total Distance' values; if this was not intended, the inner framework should be eliminated.

General context

Does adjacent material on the page distract from or enhance the graph, and vice-versa? Does redundancy, if present, help or hinder understanding of the graph?

Example

(The text was not provided.)

CONCLUSIONS

The present analytic scheme is designed to reveal features of displays that result in their being difficult to understand. This scheme depends on looking very carefully at a display, and generating a detailed description of it at three levels of analysis. At each level we consider the nature and role of the background, framework, specifier, and labels, and we consider the interrelations among these constituents. Whenever we have difficulty generating the description, we consider whether one of the acceptability principles has been violated. This exercise proves to be a systematic and exhaustive way of characterizing potential problems with a display. The Appendix illustrates a particularly compelling example of the power of this system.

The present analysis is intended to be at the most fine-grained, 'picky' level possible. It is important to emphasize that many of the 'violations' we identified would in fact be irrelevant, but this cannot be known without knowing the purpose of the display. That is, charts and graphs are created with a specific purpose in mind; they are intended to allow a reader to answer certain questions and not others. Thus, although an acceptability principle may be violated, the chart or graph may not be flawed—it may still be able to serve its purpose adequately. For example, the sample graph illustrated in Fig. 1 violates what we called the 'within-level mapping principle' because the points on the function do not correspond unambiguously to points on the axes. But this is not an impairment in the graph given its purpose. In fact, when graphs are used as idealizations to present a general principle, the additional information necessary to totally disambiguate the display may distract from the purpose—violating the pragmatic principles of purpose compatibility by providing more than is needed.

Thus, although the analytic scheme faithfully exposes little details that violate an acceptability principle, not all of these violations may be important; whether a violation of a principle renders a display ineffective depends on the purpose to which the display will be put. The scheme errs on the side of being too conservative, leaving it up to the human user to discount particular violations as appropriate. This was the only real option, given that all other alternatives run the risk of not exposing potential problems with the chart or graph.

The present analytic scheme has been designed to analyse existing displays. However, it can also be of use in the design process itself. One could begin with a draft of a display, apply the scheme, and then correct violations of the principles. Better yet, one could keep the principles in mind from the outset, and check each component as the display is being created. Although the present system does not guarantee that one will produce the best possible display, it does virtually guarantee that one will produce a good display. A good display, after all, is one easily usable by human beings.

Acknowledgements

Requests for reprints should be sent to S. M. Kosslyn, 1236 William James Hall, 33 Kirkland Street, Cambridge, MA 02138, U.S.A. The material in this paper was originally written in 1978, and was revised in 1980, to be included as a chapter in a

book coauthored with Steven Pinker. This book was never completed. Steven Pinker played a valuable role in the conception and development of the present work, and I wish to thank him for lending his insights, knowledge, and good sense. I also wish to thank Susan Chipman for her persistent support and encouragement, without which this article would never have been written. Finally, Carolyn Backer Cave, John Gabrieli, Lynn Hillger, Leon Parkin and Bill Simcox not only provided valuable comments on earlier drafts of this paper, but helped analyse the figures and devise the recommendations for improvements offered in the Appendix. The work described here was supported by ONR Contract N00014-K-0291.

APPENDIX

Figure 4 illustrates a very complex, multipanelled display. At one time this format was being considered for use on labels of foodstuffs. Unlike Fig. 3, which is relatively easy to understand, many people have difficulty deciphering the meaning of Fig. 4. Thus, this display is a good test of the utility of the present analytic scheme. It is of interest that the present system not only illuminates what is wrong with this display, but also provides insight into how to improve its design.

NUTRITIONAL INFORMATION PER SERVING




Figure 4. A very complex three-panel graph used to illustrate the analytic scheme, analysed in the Appendix

Syntactic analysis

This display requires an initial analysis into panels, and a later analysis of the interrelations among these displays.

Analysis into panels

The chart is divided into two graphs (left and middle) and a cluster of alpha and numeric material (hereafter referred to as the right table). The rightmost boundary of the left graph is defined by right justification of seven circles and blank space to the right of the circles. The rightmost boundary of the centre graph is defined by an

annular white space between the small radial marks in the centre of the page and the circular justification of the alpha material on the right.

Left graph (LG)

LG Outer framework

Elements. A horizontal axis, indicated by the bracket on the bottom. The axis is syntactically differentiated.

Organization. Except for the presence of the bracket, the outer framework is only implied (by the regular alignment of the inner framework elements).

LG Inner framework

Elements. Twenty-eight closed curved lines, forming circles. These are syntactically dense, and are printed in medium-weight black ink.

Organization. The circles are aligned into columns via proximity.

VIOLATION. Principles of perceptual organization. Proximity results in an organization into columns when an organization into rows is required. The vertical spaces should be larger than the horizontal ones.

Organization of inner and outer frameworks

The bracket encompasses the inner framework elements

LG Specifier

Elements. Black quadrants of circles (i.e. subtending 90 degrees of arc).

Organization. The filled areas are contained within LG inner framework elements. When one of these elements appears in a frame it is positioned in the upper left quadrant. As additional elements are added to a frame, they are placed contiguous to prior elements and fill the frame in a counterclockwise manner. Frames are filled from left to right in rows.

VIOLATION. Principles of perceptual organization. At first glance the inner framework leads one to divide the quantities into fourths, which is incorrect.

LG Labels

Only alpha and numeric labels appear; there are no depictive labels. Because alphas and numerics appear in the same perceptual units, separate syntactic discussions are inappropriate. One typefont (medium weight, black) is used within this subgraph, and alphas may be upper or lower case.

Elements. Subgraph title, 'Nutritional cont.'. The first letter is upper case, remaining letters are lower case, a period appears last.

Organization. Letters have upright orientation and are arranged in two groups in a closely packed horizontal string.

Elements. Seven vertical axis (row) labels are mixed upper and lower case with periods and numerics intermixed.

Organization. Letters have upright orientation and are arranged in one or two groups in closely packed horizontal strings. Labels are left justified at the same column.

Elements. Horizontal axis label 'needed per day' is composed of lower case letters.

Organization. Letters have upright orientation and are arranged in three groups in a closely packed horizontal string.

Organization among different LG labels

LG title is left-justified in the same column as the vertical axis labels. The space between the title and the top vertical axis label is only slightly greater than the space between the various vertical axis labels.

VIOLATION. Principles of perceptual organization. Both the position of and use of the same typefont for all labels impairs identifying the subtitle as distinct.

Organization of the framework, specifier and labels

Framework and specifier. The dark quadrants of circles are contained within inner frame elements, as mentioned above.

Framework and labels. The title is just above, and commences to the left of, the array of circles. The horizontal axis label is below the bracket.

VIOLATION. Principles of perceptual organization. The title of the subgraph does not group clearly with the implied outer framework.

Labels and specifier. The specifier material is not labelled

Middle graph (MG)

MG Outer framework

Elements. 40 short lines, approximately equal in length. The frame comprised of these elements is syntactically differentiated.

Organization. The lines project outward from a common centre and extend from a common distance from the centre to a slightly greater common distance from the centre. The lines are separated by approximately equal angles, but the separating angles are discriminably different.

VIOLATION. Principles of processing limitations. It is not immediately apparent that there are exactly 40 short marks in this frame, but this is important to note in order to understand the graph.

MG Inner framework

There is no inner framework.

VIOLATION. Between-level mapping principle. The specifier looks as if it could be within an inner framework.

MG Specifier

Elements. Two 'pie-slice' wedges; one black, one white. The black wedge is slightly larger than the white.

VIOLATION. Principles of perceptual organization. Failure to include the bottom rim of the white wedge impairs seeing it as a wedge.

Organization. The vertex of the black wedge points straight down while the vertex of the white wedge appears to point straight up. The vertices are joined.

VIOLATION. Principles of perceptual organization. The alignment of the two wedges results in their being organized into a single perceptual unit.

MG Labels

No labels are present within the subgraph.

Organization of the framework, specifier and labels

Framework and specifier. Both wedges have vertices that coincide with the centre of the circle defined by the frame. Both wedges obscure the short radial lines that define the frame.

Framework and labels. Not applicable.

Labels and specifier. Not applicable

Right table (RT)

RT Framework

There is no explicit framework, outer or inner.

RT Specifier

There is no specifier in this table.

RT Labels

There are both alpha and numeric labels in this table. No depictive elements appear. Two typefonts are used: one is small light upper case, the other is large

bold lower case. All letters and numbers in the same cluster have the same typefont. Right justification is used for the table, with the exception of the digit 8's.

VIOLATION. Principles of perceptual organization. The '8's being out of line in the top cluster leads one to focus one's attention on it, for no reason.

Alpha labels

Elements. Small, upper case type appears at the top of the table. Large, lower case type appears at the bottom of the table.

Organization. The small upper case labels at the top are organized into three rows, organized into two columns; only the first row has alpha labels in both columns. Beneath these are three rows in bold type. Spacing again produces two columns. Beneath these elements is one row in bold lower case type ('kilocalorie'). The final line is separated from the rest of the table by a large gap.

VIOLATION. Principles of perceptual organization. The large gap separating the bottom line of the table impairs one from realizing that it belongs to the table.

VIOLATION. Principle of processing priorities. The difference in font size between the upper and middle clusters directs one's attention to the middle cluster first, instead of to the top one.

Numeric labels

Elements. Numbers appear in each cluster.

Organization. When more than one numeral appears in a string, they follow one another in sequence. They appear to the right of rows 3–6, and to the left in row 7.

Organization of labels

Alpha and numeric. Numerals, when present, are intermixed in the same perceptual units with letters.

Organization of the panels

Having discussed the syntax of the various subgraphs, we return to overall structure of the three.

Macroframework

Elements. A macroframework serving to integrate the displays consists of two heavy black lines, each composed of a short vertical segment and a longer horizontal segment ending in an arrowhead.

Organization. One line originates at the centre of the rim of the black wedge in the MG and terminates at the left in an arrowhead, which points at the rightmost part of the title of the LG. The lower line originates at the centre of the rim of the white wedge in the MG and points at the leftmost end of the bottom line of the RT.

Labels

The title is composed of words printed in very large upper case

Organization of macroframework and labels

The title is centred above the macroframework

Overall syntactic properties

VIOLATION. Principles of processing limitations. There is too much information to process at once.

Semantic analysis

Left Graph (LG)

LG Outer framework

The vertical axis (implied by white space to the left of the leftmost column of circles) constitutes a nominal scale. This scale is semantically differentiated (although differentiation is de-emphasized perceptually by narrower spacing between rows than columns, as noted earlier). The horizontal axis constitutes a ratio scale and is semantically differentiated. The extent of this scale represents daily nutritional requirements of given nutrients. The bracket functions as a way of indicating the scope of the label on the bottom, as will be discussed shortly.

LG Inner framework

Each circle in a row may contain as much as one-quarter of the daily requirement for a given nutrient. The circles are thus ratio scales and are semantically differentiated.

VIOLATION. Between-levels mapping principle. The perceptual representation of these circles falsely suggests a dense scale by the lack of differentiation marks on the circle. The semantic differentiation is apparent only by studying the relationship between the specifier and the inner framework.

LG specifier

The basic specifier unit (a black quadrant of a circle) represents one-sixteenth of the daily requirement for a given nutrient. Basic specifier units can be combined to indicate integral multiples of one-sixteenth of the daily requirement.

VIOLATION. Principles of schema availability. The nesting of quadrants within each of the four circles is a novel way of specifying the information, and hence must be clearly noted.

LG Labels

Alpha. English words are used in the title to inform the reader that the subgraph provides information on nutritional contents. They are also used to name the various nutritional components represented as rows of circles and to inform the reader of the meaning of the horizontal axis. Periods inform the reader that a sequence of letters is an abbreviation of an English word.

VIOLATION. Between-levels mapping principle. The graph does not specify all of the nutritional content in the display; only the vitamin and mineral content. It should be labelled more appropriately.

Numeric. Numerals appear as characters which, in part, form the names of the nutritional components.

Organization among labels

Alpha and numeric. Together they comprise names

Organization of framework, specifier, and labels

Outer and inner framework. The bracket can be interpreted as unifying the collection of four circles into one dimension (along the horizontal axis of the inner framework).

Framework and specifier. The basic specifier units (black quadrants of a circle) act in conjunction with the four circles in each row to indicate the extent to which one serving of the food item satisfies the daily requirement for the nutritional component associated with the row.

Framework and labels: Alpha. The alpha labels define the meaning of the axis. The bracket indicates that the horizontal axis is defined by the English words immediately beneath it. This is a one:one mapping. *Numeric.* The numbers and alpha labels name nutritional components represented by rows.

Labels and specifier

The specifier is not labelled directly.

Middle graph (MG)*MG Framework*

This framework is ambiguous. The only interpretation that is consistent with the other subgraphs in the display is that this framework represents two distinct entities. One entity (the top part) is the total daily nutritional requirement for a person. The second (the bottom part) is the total daily caloric requirements for a person.

VIOLATION. Between-levels mapping principle. The ambiguity is due to faulty mapping from syntax to semantics.

Accepting the above interpretation, the framework would constitute a ratio scale. Although the framework appears syntactically differentiated, on the semantic level the issue of denseness and differentiation appears completely indeterminate in the context of all information present or derivable.

VIOLATION. Between-levels mapping principle. The variation in spacing between the marks of the frame seems to have no meaning.

MG Specifier

Black wedge. This represents the proportion of the total daily nutritional requirements supplied by a serving of the food in question. (This interpretation is the only one consistent with the connective relation between the black wedge and the left subgraph.)

White wedge. This represents the proportion of the total daily caloric requirements supplied by a serving of the food in question. This interpretation is uncertain, however, but is suggested by the fact that the arrow from it points to the bottom line of the table on the right.

VIOLATION. Between-levels mapping principle. The meaning of the wedge simply is not clearly defined, allowing one to interpret the meaning of the syntax in more than one way.

VIOLATION. Principles of schema availability. A circle or 'pie' graph is usually used in our culture to show how a whole is divided into parts. The MG, on the other hand, does not use wedges to divide a single entity into parts, but rather treats the two wedges as independent.

MG Labels

No labels of any sort are wholly within the MG.

VIOLATION. Between-levels mapping principle. Missing labels on both the framework and the specifier make this graph very difficult to understand.

Organization of framework and specifier

According to the most consistent reading, the specifier elements represent two distinct entities: (1) proportion of daily nutritional requirement supplied per serving (black wedge), and (2) proportion of daily caloric requirement supplied by a serving (white wedge). The frame represents the whole daily requirement of these two entities (nutrition and calories) and, therefore, supplies ratio scales in which both specifier elements are measured. The different sizes of the two wedges are thereby explained.

VIOLATION. Between-levels mapping principle. If this interpretation is correct, the scale is different things to different objects, and therefore violates the distinctness property required for symbol systems to be unambiguous.

VIOLATION. Within-levels mapping principle. The wedge-shaped specifier elements obscure the hash marks that comprise the outer framework. This prevents any quantitative mapping from the specifier to the framework.

Framework and labels

The framework is not labelled in the MG. If it had been, two different labels would have been required for the same framework, or the framework would have to be divided into two semicircular frameworks, each separately labelled.

Labels and specifier

The specifier in this graph is not labelled within the subgraph. Specifier elements within the graph are connected to labels in the other panels and derive meaning thereby, as will be discussed shortly.

Right table (RT)

RT Framework

There is no actual framework; the framework is only loosely implied by the arrangement on the page of the material.

VIOLATION. Principles of perceptual organization. The organization of the table is not evident from the proximity relations of its components.

RT Specifier

There is no specifier in a table.

RT Labels

Alpha. The labels in the upper cluster are English words that specify quantities of food. The labels in the middle cluster are English words or abbreviations that name of nutritional components of food. The symbol 'g' indicates 'grams'. The lower label is an English word meaning a unit of heat (in this context, the heat equivalent of a serving of food).

Numeric. The numerals specify quantities.

Organization of labels

Alpha and numerics. Alphas and numerics appearing in the same perceptual units together specify a quantity (e.g. '4 grams'). These units in turn specify how much of the named substance is associated with the quantity in a serving.

Macroframework

One arrow associates the white wedge with the '170 kilocalories' label. This, in fact, allowed us to infer the meaning of the white wedge. The other arrow associates the black wedge with the entire LG, which provides an analysis of the total daily requirements of the nutritional components.

VIOLATION. Between-level mapping principle. The lack of labels on the arrows hampers one from realizing that they symbolize different relations, 'decomposes into' (top) and 'corresponds to' (bottom).

VIOLATION. Principles of schema availability. Arrows point from the specifier elements to the labels, instead of the other way around—which is the conventional direction.

Labels of macroframework

Alpha and numeric. The title labels the information provided by the entire set of graphs.

Overall semantic properties

VIOLATION. Within-level mapping principle. Labels are missing that are necessary to coordinate the panels into a single cohesive display.

VIOLATION. Within-level mapping principle. One cannot easily relate the information about protein in the RT to the information about protein in the LG, partly because of the use of 'prot.' and 'protein' in the different subgraphs. In general, use of different notations or abbreviations leads one to infer that different things are being referred to.

Pragmatic analysis

There are no clear cases in which the display has been slanted to lead us to draw incorrect inferences. All we know about the purpose and context is that the display was intended to be part of a label for foods.

VIOLATION. Principles of contextual compatibility. No person in the midst of shopping would want to take the time to decipher this overly complex display.

Recommendations for improvements

The analysis reveals many weaknesses in the display. It would be improved markedly by the following changes:

1. The left graph should be replaced by a bar graph, labelled 'Vitamins and Minerals'. The labels of individual vitamins and minerals should be placed along the *x* axis, and a scale labelled 'Percentage of Daily Requirement' should be placed on the *y* axis; the *y* axis should have 10 demarcations, with every two labelled by a number. An inner framework should consist of horizontal grid lines at each of the numbered values.

2. The information in the black wedge should be eliminated. The undifferentiated total percentage of 'nutritional content' is misleading; the breakdown into specific vitamins and minerals is more useful.
3. The remaining information should be in a table, with all labels aligned in a column on the left and numbers aligned in a column on the right.

REFERENCES

- Bertin, J. (1983). *Semiology of graphs* (trans. W. J. Berg). Madison, WI: University of Wisconsin Press.
- Chambers, J. M., Cleveland, W. S., Kleiner, B. and Tukey, P. A. (1983). *Graphical methods for data analysis*. Belmont, CA: Wadsworth.
- Clark, H. H. and Clark, E. V. (1977). *Psychology and language*. New York: Harcourt Brace Jovanovich.
- Clement, D. E. (1978). Perceptual structure and selection. In E. C. Carterette and M. P. Friedman (Eds), *Handbook of perception*, Vol. IX. New York: Academic Press, pp. 49–84.
- Cleveland, W. S. (1985). *The elements of graphing data*. Monterey, CA: Wadsworth.
- Dodwell, P. C. (1975). Pattern and object perception. In E. C. Carterette and M. P. Friedman (Eds), *Handbook of perception*, Vol. V. New York: Academic Press, pp. 267–300.
- Ericsson, K. A., Chase, W. G. and Faloon, S. (1980). Acquisition of a memory skill. *Science*, **208**, 1181–1182.
- Fisher, H. T. (1982). *Mapping information*. Cambridge, MA: Abt Books.
- Frisby, J. P. (1980). *Seeing: illusion, brain and mind*. New York: Oxford University Press.
- Garner, W. P. (1974). *The processing of information and structure*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Goodman, N. (1968). *Languages of art: an approach to a theory of symbols*. Indianapolis, IN: Bobbs-Merrill.
- Grice, H. P. (1975). Logic and conversation. In P. Cole and J. L. Morgan (Eds), *Syntax and semantics, Vol. 3: Speech Acts*, New York: Seminar Press, pp. 41–58.
- Hochberg, J. E. (1964). *Perception*. Engelwood Cliffs, NJ: Prentice-Hall.
- Huff, D. (1954). *How to lie with statistics*. New York: W. W. Norton.
- Jolicoeur, P., Gluck, M. A. and Kosslyn, S. M. (1984). Pictures and names: making the connection. *Cognitive Psychology*, **16**, 243–275.
- Kaufman, L. D. (1974). *Sight and mind: an introduction to visual perception*. New York: Oxford University Press.
- Kosslyn, S. M. (1985). Graphics and human information processing: a review of five books. *Journal of the American Statistical Association*, **80**, 499–512.
- Norman, D. A. (1978). *Attention and memory*. San Francisco: W. H. Freeman.
- Pinker, S. (in press). A theory of graph comprehension. *Cognitive Science*.
- Pinker, S. and Birdsong, D. (1979). Speakers' sensitivity to rules of frozen word order. *Journal of Verbal Learning and Verbal Behavior*, **18**, 497–508.
- Posner, M. I. (1978). *Chronometric explorations of mind*, Hillsdale, NJ: Lawrence Erlbaum Associates.
- Reed, S. K. (1974). Structural descriptions and the limitations of visual images. *Memory and Cognition*, **2**, 329–336.
- Rosch, E., Mervis, C. B., Gray, W., Johnson, D. and Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, **8**, 382–439.
- Schmid, C. F. (1983). *Statistical graphics*. New York: John Wiley.
- Smith, S. L. (1979). Letter size and legibility. *Human Factors*, **21**(b), 661–670.
- Spoehr, K. T. and Lehmkuhle, S. W. (1982). *Visual information processing*. San Francisco: W. H. Freeman.
- Teghtsoonian, J. (1965). The judgment of size. *American Journal of Psychology*, **78**, 392–402.
- Tufte, E. R. (1983). *The visual display of quantitative information*. Cheshire, CT: Graphics Press.