The Psychology of Visualization

by

Andrew Csinger

Department of Computer Science
University of British Columbia
November 1992

Abstract

This document is a review of the literature of three related areas: psychophysical vision research, automatic display generation, and multi-dimensional data visualization. Common threads are explored, and a model of the visualization process is proposed which integrates aspects of these three areas.

In the review of psychophysics, attempts to find a set of primitive perceptual channels are explored. In the literature on automatic generation and visualization, attempts to employ these psychophysical findings are investigated. Finally, the proposed model is a framework which might facilitate this kind of cooperation.

Contents

List of Tables

List of Figures

1 Introduction

You don't have to be Hungarian to be a genius, but it sure helps.

—attributed to Béla Julész

1.1 Motivation

Humans have been engaged in visualization since the day they acquired the capability of abstraction. There is nothing esoteric or arcane about the need for better ways of studying the sometimes subtle relationships among the variables in often large data streams, but surprisingly little is known about the relative merits of the different ways of going about the enterprise. New hope springs from at least two different sources, however: Insights from modern psychophysical studies promise a firm theoretical foundation for future visualization techniques, and feedback from practitioners in the graphic arts and the software industry provide both validation for these early theories and effectiveness criteria of their own.

Visualization is different things to different people. Broadly interpreted, even source listings of programs belong to the domain of visualization [Car79]. Fractal imagery, already of interest by virtue of its graphic beauty and paradoxical complexity, is a powerful visualization tool for exploring the stability of feedback systems [Pic88]. Computer-Aided Software Engineering (CASE) tools are also visualization aids [JLSU87, LJSR89].

Many research areas in science and medicine are concerned with data of many dimensions, including those of the three spatial and the temporal axes. Nonetheless, scientists still plot two-dimensional graphs, doctors still diagnose from two-dimensional X-rays, engineers still pore endlessly through numerical data, and designers still draw two-dimensional plans [Gre88]. The technology now exists to take advantage of new and old insights into human perception and improve the presentation of multi-dimensional data to humans.

In this review, I identify two broad sub-areas of the visualization research endeavor: the use of iconographic displays to explore multi-dimensional data [Fli90, FC90, JLSU87, SBG90, GPW89, PG88, CM84] and the use of graphs to represent multi-dimensional relational data [Mac86a, Mac86b, RM90, BP90].

The concerns of these camps are quite different, but their respective explorations have brought them into close contact, since both have recognized the importance of the human perceptual processor [Ber77, Ber81, Ber83, Enn90, ER90a, ER90b, ER91, KP81, Pom81, WB85, Mur85, SCS90].

1.2 Structure of this paper

Section 2 is a survey of psychophysical research whose goal is the elaboration of the primitive perceptual channels in the human. I discuss the perceptual principles which have been proposed to underlie the human visual mechanism, with particular attention to issues which arise in the design of effective displays for visualization tasks: attention and the ranking of perceptual tasks. I also relate some of the psychophysical tests that have been employed.

Section 3 is a survey of work to date on visualization. Here, I pay special attention to the application of knowledge about perception in the subdomains of automatic display generation for relational information, and multi-dimensional data exploration. I treat color under a separate heading because the results in this area may apply equally well to iconography and automatic generation.

Section 4 is further justification for the way I have structured this document. I propose a model of the visualization process which explicitly shows the roles that knowledge from different fields of study might play. Having refrained this far from serious editorializing, I freely proselytize in this section.

The final section concludes with some broad generalizations and makes some suggestions for steps that might lead to more effective visualization techniques.

2 Attention, Pre-attention, and the Ranking of Perceptual Tasks

Much work in the field of (perceptual) psychology [Nei63, Tre90, Tre86, TG88, TCF$^+$90] has been concerned with elaborating a putative dichotomy between processes which are pre-attentive and those which require attention. Pre-attentive processes are characterized by their speed: they are fast, typically accomplished within 100ms, suggesting that they are performed in parallel by the human perceptual system. Such processes are sometimes referred to as automatic, or early-vision processes. Although such a dichotomy is conceptually attractive, it has been increasingly unable to account for the data, and new models are appearing which refer to a continuous ranking of perceptual difficulty. Pre-attentive tasks are at the extreme 'easy' end of this continuum, while tasks requiring attention are at the other, 'hard' end of the scale. The possibilities afforded by such a ranking are explored in later chapters.

Much work in the areas of automatic presentation and multi-dimensional data display and exploration has been concerned with ranking perceptual tasks in order to provide meaningful *effectiveness* and *expressiveness* criteria for visualization systems. The connections between these apparently disparate fields have not, to my knowledge, been explicitly related.

2.1 Psychophysics

The early work of many researchers (*cf.* Pomerantz, Garner, Prinzmetal, Rock, Julész, etc.) sets the stage for the modern inventory of psychophysical testing techniques. The task with which subjects are faced is usually some variation on a search theme: a target in a display is distinguished from a varying number of distractors by differences on a single dimension. The temporal latencies of these search tasks is recorded for analysis. Followers in this paradigm seek to determine from the pattern of search latencies those features which are coded automatically in early vision[TG88].

In what follows I have singled out a few researchers and their work as representative of work in their field and their era. I will pay particular attention to Pomerantz, as exemplary of early work, to Treisman as a proponent of feature integration theory, and to Enns as an explorer of ecological effects on pre-attention. Montalvo's work is included here to show that similar work is being done outside mainstream psychology.

I superimpose the following terminology on my discussion throughout the remainder of this paper:

Definition 1 *Property: A* property[1] *is a generalization of Treisman's use of the term* dimension *(see Section 2.1.2).*

Definition 2 *Data Space:* Data space *is the mathematical space in which the domain-dependent data to be visualized are defined.*[2]

Definition 3 *Stimulus Space: A* Stimulus space *consists of some subset of the properties referred to above, called* stimulus dimensions. I.e.*, stimulus space is defined by the choice of perceptual (visual) dimensions to which a subset of data dimensions are to be mapped.*

Definition 4 *Stimulus:* Stimuli *are composed of n* variables $x_1, x_2, \cdots, x_n$ *which vary along orthogonal* stimulus dimensions.

Definition 5 *Value: The* values *of all the variables of a stimulus* instantiate *that stimulus.*

Definition 6 *Series: A* series *is a (temporal) sequence Σ of stimuli $\mathcal{S}_1, \mathcal{S}_2, \cdots, \mathcal{S}_s$.*

2.1.1 Pomerantz: Perceptual Grouping

Pomerantz [Pom81] discusses *divided attention versus selective attention*. Lamenting the loose terminology of the Gestaltists, he proposes a metric for the concept of perceptual grouping arrived at via tests for both selective and divided attention. I will describe the tests because they are highly suggestive of the kinds of questions one might want to ask in this paradigm, and even of how one might go about answering some of these.

Selective Attention is measured by presenting to the subject stimuli that vary in two or more of their component parts or dimensions. The task is to classify these stimuli according to one part while ignoring the other part; "if the two parts in question are dissected into separate perceptual units, then selective attention to just one part should be possible. But if the two parts are parsed into the same perceptual unit or group, then the two should not split, and so selective attention should be difficult or impossible."

Filtering search tasks: Tests of this sort are called filtering tasks, and take the following form. A series Σ_F of stimuli $\mathcal{S}_1, \mathcal{S}_2, \cdots, \mathcal{S}_s$ are presented to the subject, who uses only the value of x_k in each stimulus, for some k between 1 and n to decide the membership of each stimulus in pre-assigned categories. Reaction times are recorded and averaged over the series.

In particular, Pomerantz composes stimuli of two dimensions α and β which vary in each case over the set $\{(,)\}$, *i.e.*, over the set consisting of the right and left bracket characters. Four stimuli $\alpha \times \beta$ are possible: $\{\boxed{()}, \boxed{)(}, \boxed{((}, \boxed{))}\}$

The control series Σ_C consisted of 32 stimuli which varied in only a single dimension, for instance, $\boxed{((}$ and $\boxed{)(}$, where x_1 is the only relevant dimension because x_2 is invariant. The

[1] Some researchers refer to these properties as *channels*, but I hesitate to employ this term because of its physiological implications. The reader should feel free to consider my usage equivalent, bearing in mind that I might allow 'channels' which are not primitive in any physiological sense.

[2] The distinction between data and stimulus space is explored in detail in Section 4. Psychophysics, at any rate, is concerned only with what I refer to as stimulus spaces.

dependent variable was the total time required for the discrimination of all the stimuli in the control series.

The filtering or selective attention series Σ_F Pomerantz used consisted of 32 stimuli drawn once again from $\alpha \times \beta$. In this series, however, both the α and β dimensions are varied, while the subject is still required to make discriminations based upon only one of the x_k values ($k = 1$ in the experiment cited by Pomerantz).

In fact, the response times for this series using the dimensions described above showed that the irrelevant variation in the second dimension hindered the filtering task. Pomerantz suggests that this can be explained by the tendency of the brackets to group, leading to a difficulty in selective attention. The same series drawn from different dimensions $\alpha = \{(,)\}$ (as before), and $\beta = \{\smile, \frown\}$ (*i.e.*, as before but rotated through 90 degrees) produce no effect over the control series. Variation in the irrelevant dimension had no effect on response times, indicating that selective attention to these dimensions was possible, and that the components of these stimuli had less tendency to group perceptually.

Divided Attention tasks require that subjects pay attention to at least two aspects of a stimulus and base their responses on both aspects.

Condensation search tasks: Tests of this sort are called condensation tasks, and take the following form. A series Σ_D of stimuli $\mathcal{S}_1, \mathcal{S}_2, \cdots, \mathcal{S}_s$ are presented to the subject, who uses the values of two or more x_k for some k's between 1 and n to decide the membership of each stimulus in pre-assigned categories. Reaction times are recorded and averaged over the series.

As in the filtering tasks described above, Pomerantz composed stimuli from two dimensions which varied in each case over the set $\{(,)\}$,

The condensation or divided attention series Σ_D Pomerantz used consisted of 32 stimuli drawn once again from $\alpha \times \beta$. Both the α and β dimensions were varied, while the subject was required to make discriminations based upon the values of both x_1 and x_2 .

The results revealed that the condensation task for these stimuli was easier than the filtering task (when the parentheses were closely spaced in the stimuli). When the distance between the parentheses was increased, selective attention improved, and divided attention deteriorated. The divided attention test was not conducted with the stimuli composed of rotated parentheses, although Pomerantz assures the reader that results would favor selective, and offend divided attention, once again due to the differing tendencies of these stimuli to group.[3]

To summarize: Results show that increased grouping (stronger emergent properties) leads to increased divided attention and decreased selective attention.

Pomerantz addressed the separability of dimensions. Perceptual dimensions are said to be *separable* when the time to discriminate variations on any one of the dimensions is the same whether the dimension is shown by itself or in combination. Dimensions are *configural* when discrimination of variations on any one of the dimensions will be different when viewing the dimensions in combination. When correlated variations on data dimensions result in decreased discrimination time, the stimulus dimensions to which they are mapped are said to be *integral*.

[3]To rephrase the question which experiments of this type seek to answer: the experimentor wishes to determine whether the subject is using the component or the configural (*i.e.*, emergent) properties of the stimulus [Gar81].

(Garner has called this effect a 'redundancy gain'). Uncorrelated variation on integral dimensions leads to increased discrimination time.

There is a subtle connection between selective and divided attention and integral stimuli and separable dimensions. In general, when a task calls for the division of attention between correlated data dimensions, response times will be faster for a pair of integral than for a pair of separable stimulus dimensions. For a task requiring selective attention to a dimension, integral dimensions yield slower response times than separable dimensions.

2.1.2 Treisman: Feature Integration Theory

Treisman [TG88] defines a *dimension* to be a set of mutually exclusive values for a single stimulus. Thus, a line can have the value **red** on the color dimension, as well as the value **vertical** on the orientation dimension, but can not be both **red** and **green** or **vertical** and **horizontal**, since these are values along the same dimension. Treisman uses the word *feature* to refer to a value on a dimension if that dimension appears to be coded as a distinct and separable entity, and the feature in question is coded independently of any other features on the same dimension that are also present in the field.

Experiments have been performed to gather evidence for the separability of features within a dimension [TG88] as well as the separability of one dimension from another. Treisman hypothesizes a model of visual perception called *Feature Integration Theory* to account for the data. Different feature maps (*e.g.*, for orientation, color, etc. . .) mark the pre-attentively signalled presence or (amount) of a stimulus on a dimension, while a master map of locations must be accessed attentively to determine *where* in the display field a particular object lies. Such attention is also required to detect the presence of stimuli which signal activity on more than one map (*i.e.*, where a stimulus is composed of a conjunction of features, as in a line which is both **vertical** and **green**). Weber's law[4] can be brought to bear on the search time latency data to provide support for the theory. In this paradigm, separability is a context-dependent relation between features.

Any search function that increased substantially in time with display size was interpreted as evidence for a serial scan. Pop-out in visual search tasks occurs when the target has a unique feature, which is coded early in visual processing and which is not shared by the distractors. The following distinct dimensions were provided as candidate primitives in the taxonomy of early vision, and comprise a partial list of separable, parallel, pre-attentive features:

- colors and sufficiently distinct, different levels of contrast
- line curvature
- line tilt or misalignment
- quantitative values, *e.g.*, length, number or proximity
- terminators and closure
- direction of movement and stereoscopic disparity (work by Nakayama and Silverman)

In particular, shape and color were found to be separable, as well as were lightness and size. Chroma and value are integral, as well as spatial location.

[4] See Appendix B.

2.1.3 Enns: Ecological Constraints

A visual search paradigm was employed by both Treisman and Enns. In one of Enns' experiments, subjects searched for targets distinguished from the background of distractors by the luminance relations among three polygons. Search times were low when these relations were consistent with the interpretation of the targets as three-dimensional cubes, but search was difficult when no three-dimensional interpretation was possible.

In another experiment by Enns *et alias*, subjects searched for target items distinguished from a background field of distractors by the spatial relations between lines. When changes in these spatial relations afforded interpretations which differed in three-dimensional orientation, search was easy (*i.e.*, fast); when these relations were varied among items that had no obvious three-dimensional interpretations, search became difficult (*i.e.*, slow).

Sensitivity to systems of line relations was demonstrated empirically [ER90b], and argued to be supportive of the hypothesis that preattentive processes can extract three-dimensional orientation from line drawings [ER91]. These results hint that it may be time to revise assumptions about the role and operation of the human early vision system.[5] Enns suggests that the assumption "that early vision is designed primarily to reduce the pattern of light on the retina into a useful set of visual primitives" for later processing by the attentional system, might give way to a view of the pre-attentive processor as a "high-speed system that performs a 'quick and dirty' description of objects in the three-dimensional world." Enns summarizes recent findings which he says have been guided by this hypothesis:

- *Early vision is sensitive to the direction of light in the scene.* This was shown by Ramachandran [Ram88] with a texture segregation task where the textural elements varied only in the direction of their shading gradient.

- *Three values of luminance are sufficient to determine direction of lighting.* Enns and Rensink showed that smooth, 'natural' shading is not necessary to achieve Ramachandran's effect. In fact, three different levels of grey are sufficient when items have three-dimensional interpretations.

- *Luminance relations between items and their background are critical in early vision.* Enns and Rensink find that these effects are maximal when the luminance of the background is between the extremes of the luminances in the item. They suggest that this background-contrast relativity indicates a sensitivity of the pre-attentive system to the sign of the contrast relation between item and background.

- *Early vision's sensitivity to three-dimensional orientation is premised on important constraints:* Search is easiest, for example, when the line junctions in the items correspond to orthogonal edges in the scene.

2.1.4 Montalvo: Another approach

Montalvo [Mon90] is looking for independent, or orthogonal stimulus dimensions, as are the other researchers cited in this paper. She is motivated, as are the automatic graph generation and

[5]This kind of sensitivity to scene-based properties also has repercussions on the choice of effective geometric codes for multi-dimensional data visualization [Enn90]: see Section 3, this paper.

Figure 1: Bongard Problem number 6.

visualization researchers described in the next section, by a desire to improve communication between human and machine.[6] She measures perceptual ranking according to the *vividness* of visual primitives and the *conciseness* of their symbolic representations.

Montalvo considers the acquisition of what she calls *natural, visual primitives* [Mon90]. The experimental paradigm she employs involves asking a subject to verbally describe the minimal difference between two sets of figures, simultaneously presented on either side of a vertical dividing line. Such presentations are known as Bongard problems;[7] see Figure 1 for an example. The objects on either side of the dividing line vary randomly on all dimensions but the one which defines the target property, so the "solution makes verbally explicit only the one property that distinguishes all the elements on the right from those on the left." She claims in this way to be "capturing some of the visual categories that humans have about the world,"[8] carefully pointing out that this paradigm does not depend at all upon the introspection of the subject.

She goes at the issue of perceptual ranking in terms of *vividness* and *conciseness*. Vividness is a determinant of which visual property will be seen first, as well as a measure of how fast it will be seen. Conciseness is a determinant of which symbolic description is most easily expressed and remembered. She claims that these two criteria are linked, but she has not yet been able to unravel the relationship between them.

Figure 1 is representative. In this Bongard problem, the polygon elements on the left differ from those on the right only in the 'count' property of their *sidedness*. Bongard problems can be defined to discriminate any dimension; problems can also be composed and decomposed.

[6] She is more interested in the interactive requirements of such systems.

[7] Montalvo cites N. Bongard, *Pattern Recognition*, Macmillan, London 1970.

[8] The Bongard paradigm also enables the elaboration of compositional operators for visual primitives.

3 Visualization: The Orchestra Metaphor

My eyes were drawn –dare I say: pre-attentively?– to the patterns emerging from the string section, as the individual cellists, violinists and basists drew their bows over their instruments. While my ears responded to the tones of the music, I gradually grew aware that I was hearing the melody with my eyes *as well.*

—The author's reflections on a concert.

But the silliest feature of all was that if you wanted your company accounts represented as a piece of music, it could do that as well. Well, I thought it was silly. The corporate world went bananas over it.

—Douglas Adams
Dirk Gently's Holistic Detective Agency

When Brahms is the data, one conventionally and justifiably expects to visualize it with one's auditory system. Although much would be lost by foregoing the *sound* in music, I wonder how much appreciation the experienced concert-goer gets from the visual channel. The marvellous visualization tool that is the orchestra[9] is suggestive of some of the extant approaches to data visualization.

"The critical requirement of an effective data display is that it stimulate spontaneous perceptions of structure in data." [SBG90] The preceding remark obviously applies to all display applications, from word-processing to medical imaging. It certainly motivates some of the work in automatic graph layout and presentation, where salient dimensions of relational data are mapped to appropriate graphical techniques. In that paradigm, effectiveness and expressiveness criteria derived from vision research and graphic arts knowledge [Arn74] are used to render a meaningful view of the data. This area is further explored below.

The best known approach to data visualization is the scatterplot [GPW89, WB85]. The success of this technique is due to the ability of the early vision system to group points in space based upon proximity and similarity in color, size and shape. Ware and Beatty have shown that up to five dimensions can be effectively mapped to a full color scatterplot display, and suggest ways in which the visual effect can be maximized (See Section 3.4). Were it not for the need to detect patterns in data of arbitrarily high dimension, efforts might have stopped here.

An increasingly popular approach to increasing the dimensionality of displays is what has been variously referred to as *iconography*, and *geometric coding*. This approach employs a generalization of the traditional graphic primitive, the *pixel*, into a parameterized icon whose features are mapped to distinct dimensions of the data stream. A famous example which proved more useful as a characterization of the method than representative of its success, is the *Chernoff Face* icon family [Che73].

[9]This idea is not as fanciful as it might seem: a colleague of mine has proposed just such a tool.

3.1 Automatic Graph Generation

Some work which has had to progress without the benefit of a mature perceptual theory has implicitly recognized this by trying to establish a ranking of perceptual tasks [CM84, Mac86a, Ber83]. This kind of endeavor has been referred to as a search for a "taxonomy of measurement scales" [Ste46].

Bertin This man's work [Ber83] is nothing short of fascinating. Though highly idiosyncratic[10] and thoroughly at odds with much of the current conceptualization of perception (and notably of color perception),[11] there is a wealth of accumulated experience to be gleaned from his rambling volumes.[12] He provides many examples of thought experiments which, properly conducted, could have been valid psychophysical investigations; Bertin, like many of the researchers cited in this review, is after a system of rankings of perceptual tasks, which he refers to as the 'level of organization' of the visual variables.[13]

The visual variables are composed of the 'planar' and the 'retinal' variables:

- The planar variables: x and y
- The 'retinal variables':
  - Size
  - Value (saturation)
  - Texture
  - Color
  - Orientation
  - Shape

The retinal variables are necessary not only in cartography, Bertin points out, but in all graphics problems involving more than two dimensions, where the planar spatial dimensions are already in use.[14] And most interesting is his analysis of the ways in which the different variables can be combined:

- A variable is *associative* when it promotes the recognition of emergent properties in spite of variation along other dimensions (variables). On the other hand, a variable is *dissociative*

[10]Though his books contain no references, he cites the work of Zipf [Zip35], who elaborated a notion of 'mental cost', along the lines of the concept of perceptual rank referred to in this document.

[11]His usage of the term *value*, for instance, corresponds to the conventional concept of grey-level saturation, and the term *color* to refer to any color variation including color saturation. It is apparent from a footnote on page 73, however, that Bertin is entirely aware of these divergences, and simply *chooses* to be different.

[12]Clarity is not served by reading the original French versions; these merely ramble in French!

[13]Bertin's syncretic terminology supports the thesis that he actually developed these ideas independently, over years of cartographic production work. These are, as of this version of this paper, only conjectures on my part.

[14]He further refines these notions in terms of his concepts of 'implantation' (whether the marks in question take the form of points, lines, or areas), and 'imposition' (the roles of the variables in diagrams, networks, maps, or symbols). Another term he invents is 'length', referring here to 'the number of elements or categories which we are able to identify in a given . . . variable."

Retinal Variable	*Level*			
	Association	Selection	Order	Quantity
Planar Dimensions	√	√	√	√
Size		√	√	√
Value		√	√	
Texture	√	√	√	
Color	√	√		
Orientation	√	√		
Shape	√			

Table 1: Properties of the Retinal Variables

when it "dominates all combinations made with it and prohibits carrying out an immediate visual selection for the other variables." Bertin suggests tests for associativity which are highly suggestive of those performed by Pomerantz (see Section 2, this paper) to decide between the use of configural and component properties of stimuli. [p65]

- A variable is *selective* when the eye can "isolate *all* the elements of this category, disregard all the other signs [marks], and perceive the image formed by the given category. Such perception can be immediate, in which case the variable is selective, and each category forms a family. . . . the perception can necessitate going through sign by sign, in which case the variable is not selective." Bertin suggests tests for separability which are suggestive of psychophysical experimentation to demonstrate separable perceptual dimensions (see Section 2, this paper).

- A variable is *ordered* when it leads to pre-attentive[15] apperception of underlying order.

- A variable is *quantitative* when it gives pre-attentive cues to the numerical ratio of two signs which differ along the dimension (variable) in question. "It is readily apparent that only size variation is quantitative."

- A combination of several variables utilized to represent a single data dimension is a *redundant combination*: 'Redundant combinations increase the separation between the steps of the retinal variables. They are the basis for selective legibility.' [p187]

- When two variables are each associated with a different data dimension, the combination is *meaningful*: 'the visual properties of the combination of variables are obviously applicable to meaningful correlations within the information.' [p189] This definition runs parallel to the earlier discussion of integral dimensions.

Table 1 summarizes the properties of the retinal variables [p96]. Bertin goes on to describe in amazing and overwhelming detail the characteristics of the visual variables, and gives dozens of rules-of-thumb designed to help graphic designers with their chores. The task at hand is one of

[15] I depart here from Bertin's cumbersome, unorthodox terminology, having given the reader an inkling of the difficulties involved with the material.

matching the characteristics of the data to be displayed to the characteristics of the visual variables via which they will be presented.[16] These considerations lead naturally to the issues faced by those involved with automatic display generation, which I discuss next.

Mackinlay explores effectiveness and expressiveness criteria for the automatic generation of graphical presentations, focussing upon explicitly relational data. The importance of this work for our present purposes is Mackinlay's emphasis upon the human role as perceiver of the presentation. He conjectures a theory of human perception that ranks the difficulty of perceptual tasks associated with the interpretation of presentations; a presentation composed of tasks which are perceptually 'easy' is more *effective* than one which is composed of 'harder' tasks.

Expressiveness criteria for a graphical language are derived from conventions about the usage of such languages; for example, the conventional interpretation for bar charts can be captured in a language definition that indicates that bar lengths encode an ordinal relationship between items: 'The bar chart language *cannot* express functions that map to nominal domain sets without encoding additional, incorrect information.' He gives convincing examples of success and failure in presentations along the dimensions of both expressiveness and effectiveness.

Over a dozen pages of Mackinlay's dissertation are devoted to elaborating the ranking of perceptual tasks. Beginning with an analysis of the work of Bertin (see above) and that of Cleveland and McGill [CM84],[17] he goes on to suggest ranking criteria of his own. While Cleveland and McGill performed psychophysical experiments to determine their ranking, Mackinlay uses the work of Julész on texture [Jul81], Kahneman [KH81] and Ware and Beatty on color,[18] as well as that of Bertin, but does not try very hard to justify his rankings with psychophysical studies. Table 2 is a summary of these rankings, where the bracketed terms are not applicable to the domain type in which they appear [Mac86a, p69]. This kind of knowledge has been harnessed in automatic presentation systems to decide how best to graph different relational data.

Part of the contribution of Roth [RM90] has been to refine the taxonomy developed by Mackinlay. He has subdivided ordinal types into *coordinates* and *amounts* to ensure appropriate graphic techniques in the rendering of graphs which represent information of these types. He also advances further subdivisions according to the domain of membership of the information. Thus, his characterization of data recognizes that sets can belong to domains of time, space, temperature or mass. These distinctions, says Roth, help to preserve "subtle stylistic conventions, such as using a horizontal axis for time coordinates and a vertical axis for temperature. This characterization can also be helpful for judging how to group and integrate relations within pictures."[19]

A fair summary of Roth's recent work might be to say that he has been concerned with semantics, where Mackinlay's approach has been more or less syntactic. Others have explored the use of user-specified constraints to improve automatic graph layout [BP90].[20] Extensions to non-standard displays have also been proposed [FM90], and some early implementations have been developed.

[16] Bertin includes much detail about media and printing techniques that need not concern us here.

[17] Cleveland and McGill developed the ranking of quantitative tasks given here as the first column of table 2; the other two rankings in the table were developed during Mackinlay's own doctoral research.

[18] See Section 3.4, this paper.

[19] These observations are suggestive of the information which might properly be channelled along the K_{UM} stream of the diagram in Section 4.

[20] This idea is once again suggestive of the K_{UM} channel in Section 4.

Quantitative:	Ordinal:	Nominal:
Position	Position	Position
Length	Gray Saturation	Color Hue
Angle	Color Saturation	Texture
Slope	Color Hue	Connection
Area	Texture	Containment
Volume	Connection	Gray Saturation
Gray Saturation	Containment	Color Saturation
Color Saturation	Length	Shape
Color Hue	Angle	Length
(Texture)	Slope	Angle
(Connection)	Area	Slope
(Containment)	Volume	Area
(Shape)	(Shape)	Volume

Table 2: Ranking of Elementary Perceptual Tasks

3.2 Iconographic Data Visualization

The *generalized icon* (gicon) is a generalization of the pixel to higher dimensions. The strategy has been to allow the information in different channels of the input data to control corresponding pixels in each gicon. In general, the gicon is an $n \times m$ array of pixels, each mapped to a different input channel. The available display surface is then tiled with these icons. The logic of this and related approaches is that the number of information channels which can be displayed is increased: "Geometric coding allows for further and far reaching extensions [over color] of dimensionality. Observers can utilize shape perceptions to sense the combinations of data at each location and texture perception to sense how those combinations are spatially distributed." This was the rationale behind the Chernoff Face icon family, as well as the stick figure family described by Pickett and Grinstein [Pic88]. The latter is a stick figure consisting of five connected line segments, where the angle of inclination of each limb is controlled by a different dimension of the numerical data to be visualized.

In keeping with the psychological findings related earlier in this document, these and other workers in this area assume that the most effective icons for exploration of multidimensional data will be those composed with features along integral dimensions of visual coding, so that seeing the data dimensions in combination is facilitated (*cf.* The *redundancy gains* referred to earlier).

The search for effective icons is also lead by studies of pre-attention: "Shifts along certain dimensions of color, shape and motion of elements lead to preattentive discrimination, and it is variation in these dimensions that we must seek to bring under data control in our texture displays."

Another factor is ecological potency, as described earlier in this document. See also the work of Rogowitz *et alias* [RV90], who argue that since the spatial coding properties of the human visual system have evolved in response to the statistical properties of the physical environment, similarities between the logarithmic spacing of spatial-frequency "channels" and the logarithmic spacing of frequency components in fractal images is not surprising. This leads them to the observation that the human visual system appears to be particularly well-suited to the perception of fractal objects. They hint that their attempt to characterize Rorschach test images in terms of fractals is meeting with at least some early success.

Some early implementations have been described in the literature. The Exploratory Visualization (Exvis) project [SBG90], for instance:

> is a multi-disciplinary effort to develop new paradigms for the exploration of data with very high dimensionality. The fundamental philosophy behind Exvis is that data representation tools should be driven by the perceptual powers of the human. In addition, the interpretation of data of very high dimensionality will be maximized only when we learn how to capitalize simultaneously on multiple domains of human perceptual capabilities.

This project is in the early stages of exploring the possibilities of iconographic data representation using sound attributes, along with the integration of auditory and visual displays into a single unified data exploration facility. (See also Grinstein *et alias* [GPW89] and Pickett [Pic91]).

3.3 Requirements for Effective Geometric Codes

Enns has summarized some of what is known about perception as it relates to the design of effective icons [Enn90]. I have generalized some of his observations:

- *More than n variables simultaneously displayed.* The value of n might be approached from two, opposite, directions. There is the question of the bandwidth of the human perceptual processor, which is currently unknown, and which effectively places an upper limit on n, the dimensional limit of stimulus space.[21] There is also the practical, application-dependent issue of how many variables need to be viewed in combination, which presents itself as a lower bound on n.[22]

- *Spontaneous region segregation.* Appeal is made here to the correlative abilities of the human pre-attentive processor.

- *Combination of variables for detection and identification of interactions.* The stimulus dimensions chosen should encourage pre-attentive 'pop-out' phenomena when combined.

- *Perceptual separability by early vision system of parameters.* Enns points out that: "...not all logically-independent perceptual properties are perceptually-independent." Not all combinations of stimulus dimensions will be equally *effective*.

3.4 Color Considerations

Color deserves a separate section in this document for several reasons. Our world is, for most of us, a very colorful place. Color is, not surprisingly, one of the most effective stimulus dimensions. Even in the absence of a complete neurophysiological underpinning, a tremendous amount of rule-of-thumb information is available on the use of color to accomplish various communicative tasks. It is not that the use of color in visualization is essential, but that the kind of knowledge we have of color capabilities provokes questions about other human perceptual capacities.

[21] Grinstein *et al.* [GPW89] claim there displays can encompass at least 20 dimensions, but say little of their effectiveness.

[22] At any rate, if fewer than five dimensions are to be displayed, they can be displayed quite effectively as a color scatterplot.

Ware and Beatty show that it is possible for human observers to perceive five data dimensions simultaneously [WB85]. The data they used was characterized by a hyperellipsoidal probability density distribution, but they conclude with respect to the generality of their results that: "colour is likely to be effective in assisting in the perception of correlations in multidimensional space". Although in most cases they found that adding color was expressively equivalent to adding three more spatial dimensions, color is not a completely heterogeneous perceptual space, and "resolution is worse in some directions than in others." In particular, when clusters are separated along dimensions which have been mapped to color, perception suffers. Clusters are perceived as distinct when they are separated by between three and five standard deviations along most of the possible vectors; "much greater cluster separation is necessary before two clusters can be resolved" when they are separated on [only] "a few" specific color vectors. One solution, discussed elsewhere in this paper, is to vary the mapping of variables to dimensions; they call this mapping the *permutation vector*, and I too employ this terminology.

They observe that users require no training to use their color-based five-dimensional visualization tool, but point out the importance of control over the background color, which tends to emphasize particular colors in the display, and consequently particular correlations in the data.

Murch also gives an interesting summary of the use of color, from the point of view of a graphics practitioner [Mur85]. He distinguishes between the qualitative and quantitative uses of color, and provides a stimulating list of guidelines for the effective use of color derived from physiological, perceptual and cognitive studies. I summarize his observations in Appendix A. Murch also provides a list of the sixteen best and worst color combinations.

These guidelines are suggestive of a beginning for a database of default axioms for reasoning about presentations of information.

4 A Model of the Visualization Process

To further motivate the unification of the areas in this literature review, I advance the following model of the visualization enterprise. This model is intended as a general, covering model for a broad class of visualization activities, subsuming all the approaches discussed thus far in this paper.

In particular, I wish to construe the visualization process as composed of the following sub-processes:

1. Determining the *permutation vector*. Three steps are involved in this process.

   - Choose the data dimensions (*i.e.*, the subspace of the data) to project into stimulus space.
   - Choose the perceptual properties to use as the dimensions of stimulus space.
   - Choose the mapping from data to stimulus space.

   What I have thus far been calling the *stimulus* is therefore the result of applying a function to the input data: $\mathcal{S} = f(< x_1, x_2, \cdots, x_n >) =< P^n_{i_1}, P^n_{i_2}, \cdots, P^n_{i_k} >$, where P^n_m is the m-th

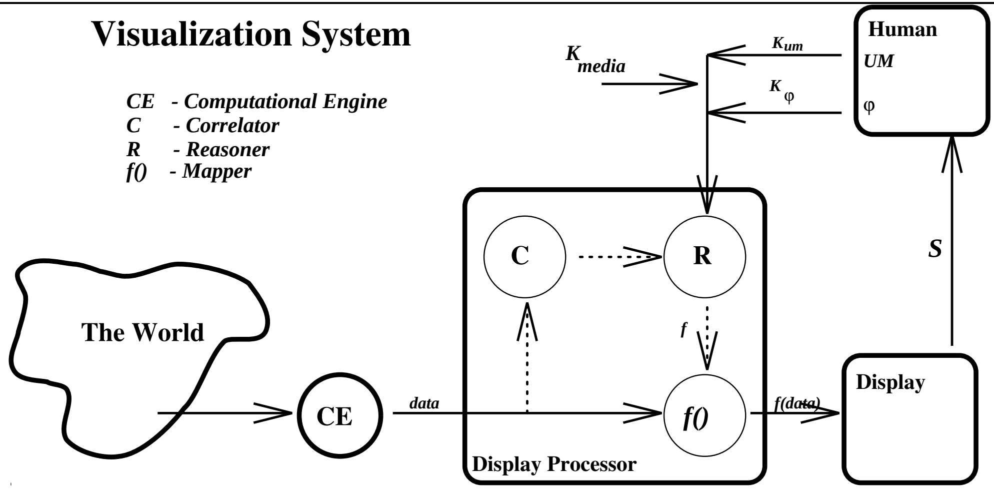

Figure 1: The Visualization System

element of the n-place projection function. The permutation vector for a k-dimensional stimulus space[23] is thus $< i_1, i_2, \cdots, i_k >$.

2. Iterating, possibly interactively, through step 1 in order to explore the data effectively.

This approach to visualization can be characterized as the process of reducing the domain-dependent data to a suitable domain-independent stimulus, upon which can be brought to bear the results of psychological studies. More to the point is that there is a body of domain-independent psychological knowledge which can and should be used to derive the projection function referred to above.

The component descriptions which follow refer to the high-level block diagram in Figure 1. I will also make a few remarks about the channels connecting the various components of the system.

4.1 Components

The Domain: The target of the computation, analysis, discourse, etc., at hand. This might be weather data for meterologists using a weather forecasting system, the Herbrand Universe for logicians searching for a proof of their favorite would-be theorem, AIDS case information for epidemiologists searching for partitions in their samples [GPW89], software components for an engineer engaged in CASE analysis, and so on.

Computational Engine: The process that performs the domain-dependent calculations, manipulations, required by the application.

[23] *I.e.*, for a stimulus space of k properties.

Display Processor: The component that extracts relevant information from the output of the computation, and sends it on to the display in some suitable format. In psychophysical terms, the role of this process is to generate the *stimulus*. The various sub-components are briefly:

- Correlator: Performs mathematical analyses on the data, passing results on to the *reasoner*.
- Reasoner: Generates the permutation vector based upon the available information from the correlator and various other knowledge sources.
- Mapper: Applies the permutation vector to the data, producing the stimulus.

Different dimensions of the input data are mapped to various perceptual dimensions in stimulus space, leading ultimately to differing partitions of the visual field.[24] If information from the different sources (see below) permit the display processor (DP) to uniquely determine this mapping according to some effectiveness and effectiveness criteria (see Section 3), the result is a single *permutation vector* (see Section 3.4). Otherwise, different permutation vectors could be generated simultaneously or in sequence, with the aim of encouraging the powerful pattern matching system of the human user to perceive a 'pop-out' phenomenon. Constraint-satisfaction technologies from work in artificial intelligence could be brought to bear on the generation of these permutation vectors in some order of decreasing utility using the effectiveness criteria as a metric, though in the worst case random vectors might be tried.

Display: The data displayed at this stage should exhibit properties that permit the next stage to extract pertinent relations with minimum effort. The display may or may not be a conventional monitor.[25] In psychophysical terms, the contents of the display comprise the *stimulus*.

Human Processor: This is arguably the most interesting and perhaps the most powerful component of the proposed visualization system. Certainly, for the purposes of this paper, it is central. A crude division of this processor into its pre-attentive and attentional components serves to emphasize the different roles of these units.

4.2 Channels

Data: That which is to be visualized. The output of the computational engine can have an arbitrary number of dimensions, while the display processor will have a certain finite number of output channels. The DP must therefore decide, based upon information it receives along its information channels, how to allocate its remaining input and output channels. The relevant aspect of the interface is that the DP need know nothing about the functionality of the CE; it is licensed to search 'blind' through the output of the CE for correlated dimensions. This means that the visualization system can be connected to any existing code. Such a connection has been called 'symbiotic' [BG90].

[24] There are, in general, $P_k^n = \frac{n!}{k!}$ such mappings, or permutations, for data of n dimensions and a stimulus space of k dimensions.

[25] *I.e.,* I do not specifically exclude any non-traditional display media, *cf.* virtual reality, etc...

Information flows along several pathways in the diagram.

- $\mathcal{K}_{\psi}$ represents knowledge about human psychophysics. Psychophysical research provides a wealth of information about perception, which is provided to the DP along this data channel. This channel is concerned with the pre-attentive capabilities of the human processor.
- $\mathcal{K}_{UM}$ represents knowledge about the human user of the visualization system. Such knowledge has been referred to as a *User Model* [Csi90]. The goals and desires of the user are quite relevant to the functioning of the application; a model of the user is made accessible to the DP along this channel. This channel is concerned with the directed, conscious, attentive aspects of the human processor.
- $\mathcal{K}_{media}$ represents knowledge about the available display media and technology. Such knowledge must be brought to bear upon the process of generating the stimulus, since it is the display surface which mediates the stimulus.

4.3 Interactivity Requirements

In designing a visualization system, it is not enough to consider the perceptual capacities of the human processor: one must perform a task-analysis of the visualization effort itself. One might wish to supply the following modes of interaction:

- Control over the permutation vector.
- Query of a point (or region or volume or hyper-volume) on the display to ask for:
  - Values at that point along dimensions in stimulus space.
  - Values at that point along dimensions in data space.

  The answers to these kinds of queries may involve new displays, and eventually a network of linked displays.
- Highlighting all points mapped to a (conjunction of) stimulus dimensions.
- Highlight all points in data space which have a (conjunction of) values on arbitrary (stimulus or data) dimensions.

Declarative Graphics: It is conceivable that in at least some visualization scenarios, users might want to interact with the data in some non-standard ways. I have in mind here such things as direct manipulation of the display to alter the values of the underlying data; parameterized icons are called *interactive graphical objects* when communication is bi-directional through the icon. Using such icons in a visualization environment may permit direct-manipulation, *what-if* analyses of the data. Users could in this fashion ask 'what-if' questions about the data they are visualizing, and receive perceptual feedback about both their question and its answer.[26]

When the user of a visualization system detects a pattern which he believes to be in some way meaningful in terms of the conceptualization of the domain he is exploring, it makes sense to allow

[26]This would likely involve connections not only to the display processor, as shown in Figure 1, but to the computational engine as well. These issues are peripheral here, however.

him to identify this pattern to the system. We want the user and the system to be able to 'talk' about this new object. This kind of declaration could be made in terms of a conjunction along ranges of either stimulus or data space, in some kind of logical calculus, or by direct graphical interaction with the perceptual phenomena of interest. An example of the former (using stimulus dimensions) might be:

$$\forall_{points}\text{red(point)} \vee \text{x(point, X)} \wedge (\text{X} > 100) \wedge \text{blue(point)}$$

while an example of the latter might be the mere pointing of a locator graphic device at a cluster of points in a scatter-plot.

5 Conclusions and Discussion

Or you can turn your figures into, for instance, a flock of seagulls, and the formation they fly in and the way in which the wings of each gull beat will be determined by the performance of each division of your company. Great for producing animated corporate logos that actually mean *something.*

—Douglas Adams
Dirk Gently's Holistic Detective Agency

From the preceding exploration of the psychophysics of perception, and the overview of the visualization undertaking, one might conclude that matters are well in hand. There is, to be sure, a concerted effort on the parts of psychologists to elaborate the variables which figure in the human perceptual system. There is, as well, a will on the parts of visualization practitioners to make use of these and other results. The connections, however, that would make the psychophysical results apply in some direct or obvious way to the visualization process remain unclear.

Observations on the Contribution of Psychophysics Classical approaches to the study of grouping have serious drawbacks, according to Pomerantz. They are first of all subjective, and are a scientific dead-end. The latter objection is made from the observation that these studies can only establish the existence of a phenomenon, but fall short of 'localizing it somewhere in the chain of events that constitutes perception'. I have two things to say about this:

- This criticism is overly harsh, especially in view of the fact that we are far from having established the causal 'chain of events that constitutes perception'. I admit that the hopeful, naive phenomenology of the classical Gestaltists has not delivered on its early promises, but this surely leaves some maneouvering room for would-be neo-Gestaltists!
- Were the full force of Pomerantz' criticism entirely well-deserved, I would nonetheless console myself with an exhaustive catalogue of perceptual modes, or ways-of-perceiving, along with their relative strengths and weaknesses. With such a compendium, my visualization system would be able to choose the 'best' way to render a set of data to a user.

Feature integration theory [Tre90, Tre86, TG88, TCF+90], in spite of the appealing simplicity of its proposed dichotomy between what is pre-attentive and what is attentive, is unable to account for the continuum of search times in psychophysical experiments (see Section 2, this paper). The suggestion is that perhaps there is a continuum of perceptual difficulty from the very easy, fast, pre-attentive to the very difficult, attention-requiring. Feature integration theory is a useful conceptual model, but it should not be allowed to hinder efforts to arrive at the ranking of perceptual tasks required by visualization practitioners.

Lessons from Psychophysics for visualization When looking for correlations between dimensions of the data, one may wish to map correlated dimensions to perceptually integral channels in human stimulus space, so that the perceptual system benefits from any redundancy gain. When looking for basic differences between dimensions, one might wish to map the data dimensions to perceptually separable channels.[27]

The design of geometric codes might profitably proceed from an analysis of perceptually separable and perceptually integrated dimensions of visual coding.

Flinn says [FC90, p39] that one way to address the issue of transforming a data stream into a smaller and more comprehensible format is to shift the burden of data reduction from the computer to the user. "The hope is that a data stream can, through a . . . change of representation, be converted to a form which exploits the capabilities of the human visual system". By 'shifting the burden' in this fashion, I imagine designers trying to express data using techniques which fall as far to the 'easy' end as possible along the perceptual cost continuum. The elaboration of this continuum is therefore of great interest to practitioner and theorist alike. Psychophysical data of this sort will fuel further investigations of their neurophysiological foundations, as well as provide guidelines for effective representation in data visualization task environments.

Ecological Potency The idea of exploiting real-world statistics (in addition to human perceptual norms) on the premise that human perception evolved in the same statistical framework is a good one, and should be exploited wherever possible. Thus, for instance, if animating a feature on a visualization display, it would come as no surprise to me to find that the human visual system is more sensitive to motion which is gravity-governed[28] than to other arbitrary movements of features along dimensions.

More Questions...

- Do the **spatial dimensions** *need* to be represented in the permutation vector in all cases? Does their *a priori* inclusion yield undesirable (and so far unexplored) perceptual biases?

- Are all **correlations** equal? Specifically, is the human perceptual pattern matcher referred to throughout this paper sensitive to correlations between dimensions of data according to one of the traditional mathematical formulations of this relation? Or is it something quite different? Is there some sort of overlap? What kind? How much? It seems that humans are at least as sensitive to correlation as the mathematical definitions. So sensitive, in fact, that practices have evolved in human societies that take advantage of it. Consider the

[27] The preceding from a conversation with Jim Enns, 1991 March 22.

[28] *I.e.*, motion that can be interpreted as 'falling'.

well-known casino game of craps; inveterate gamblers insist against published theory either that they "have a system" or that they have a "feel" for the die.[29]

- What is the **maximum dimensionality of stimulus space**? (*I.e.*, how many properties are there?) The recent work of Enns has demonstrated that the search for perceptual primitives is not yet over. If there are such late-breaking discoveries in the area of visual perception, which has seen much research activity over more than a century, what then of the non-visual channels, which are much less explored?

5.1 Final Remarks

Psychophysics is the exploration of stimulus space. Visualization is the exploration of data space via a suitable mapping from data to stimulus space.

The goal of psychophysics, from the visualization perspective, is a more expressive visual language with which the construction of meaningful graphical displays is facilitated. The motivation is to maximize the bandwidth of the communications channel between user and data [CCBD87]. This goal is being met from one side by fundamental psychophysical studies of the limits on human perception, and from another angle by continual improvements in the computational methods and technologies used to direct data to humans.

While I think that it is premature to infer neurophysiological structure from the psychophysical studies alluded to, I see much to be gained by a practical marriage of visualization practitioners and psychophysicists.

[29] There is certainly some plausibility in claiming that, due to irregularities in the physical shapes of the die, certain patterns of throws will be more likely than others, but casino earnings testify to the expertise of die-makers!

A Murch's Guidelines on the use of Color

- Physiological Guidelines
  - Avoid the simultaneous display of highly saturated, spectrally extreme colors
  - Avoid pure blue for text, thin lines and small shapes
  - Avoid adjacent colors differencing only in the amount of blue
  - Older viewers need higher brightness levels to distinguish colors
  - Colors change appearance as ambient light level changes
  - The magnitude of a detectable change in color varies across the spectrum
  - Difficulty in focusing results from edges created by color alone
  - Avoid red and green in the periphery of large-scale displays
  - Opponent colors go well together
  - For color-deficient observers, avoid single-color distinctions
- Perceptual Guidelines
  - Not all colors are equally discernable
  - Luminance does not equal brightness
  - Different hues have inherently different saturation levels
  - Lightness and brightness are distinguishable on a printed hard copy, but not on a color display
  - Not all colors are equally readable or legible
  - Hues change with intensity and background color
  - Avoid the need for color discrimination in small areas
- Cognitive Guidelines
  - Do not overuse color (*cf.* the five plus-or-minus rule)
  - Be aware of the non-linear color manipulation in video and hard-copy
  - Group related elements by using a common background color (*set* and *pre-attention* in subjects)
  - Similar colors connote similar meanings
  - Brightness and saturation draw attention
  - Link the degree of color change to event magnitude
  - Order colors by their spectral position
  - Warm and cold colors should indicate action levels

B Weber's Law

Weber's Law describes the amount by which the magnitude of a physical quantity x must be increased so that the difference will be discernable with probability p. This formulation is a small refinement of Weber's Law in terms of the concept of *just noticeable difference* [LN77].

$$w_p(x) = k_p x$$

C Steven's Law

Steven's Law describes the relationship between an actual physical magnitude x, and its perceived magnitude $p(x)$:

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

β is empirically determined for particular physical quantities. For instance, experiments show that the average β for length judgements range from 0.9 to 1.1; for area, from 0.6 to 0.9; for volume, from 0.5 to 0.8. When the stimuli are intended to be interpreted in one-to-one correspondence with physical values, people are therefore more accurate at length judgements than area or volume, and are worse with volume than with area [CM84].

References

[Arn74] Rudolf Arnheim. *Art and Visual Perception: A Psychology of the Creative Eye*. University of California Press, Berkeley, CA, 1974.

[Ber77] Jacques Bertin. *La Graphique et le Traitement Graphique de l'Information*. Flammarion, 1977.

[Ber81] Jacques Bertin. *Graphics and Information Processing*. Walter de Gruyter, 1981.

[Ber83] Jacques Bertin. *Semiology of Graphics: diagrams, networks, maps*. University of Wisconsin Press, 1983.

[BG90] Kellogg S. Booth and W. Morven Gentleman. A symbiotic approach to visualization and the user interface. *Unpublished*, 1990.

[BP90] Karl-Friedrich Bohringer and Frances Newbery Paulisch. Using constraints to achieve stability in automatic graph layout algorithms. In *CHI'90 Proceedings*, pages 43–51, April 1990.

[Car79] T.A. Cargill. A view of source text for diversely configurable software. Technical Report CS-79-28, University of Waterloo Computer Science Department, 1979.

[CCBD87] A. Csinger, H.da Costa, B.Forghani, and D.A.Lowther. Increasing cad throughput with a programmable user interface. In *Official Proceedings of the 3rd Int'l IMS '87, SATECH '87, Part I*, Long Beach, CA., 1987.

[Che73] H. Chernoff. The use of faces to represent points in k-dimensional space graphically. *Journal of the American Statistical Association*, 68:361–368, 1973.

[CM84] William S. Cleveland and Robert McGill. Graphical perception: theory, experimentation and application to the development of graphical methods. *American Statistical Association*, 79:531–554, September 1984.

[Csi90] Andrew Csinger. Implementing a theory of communications in a default reasoning framework. Master's thesis, University of British Columbia, 1990.

[Enn90] James T. Enns. The promise of finding effective geometric codes. *Conference Presentation, SIGGRAPH'90*, 1990.

[ER90a] James T. Enns and Ronald A. Rensink. Influence of scene-based properties on visual search. *Science*, 247:721–723, February 1990.

[ER90b] James T. Enns and Ronald A. Rensink. Sensitivity to three-dimensional orientation in visual search. *Psychological Science*, 1(5):323–326, September 1990.

[ER91] James T. Enns and Ronald A. Rensink. Preattentive recovery of three-dimensional orientation from line-drawings. *Psychological Review*, 1991.

[FC90] Scott Flinn and William Cowan. Visualizing the execution of multi-processor real-time programs. In *Proceedings Graphics Interface 1990*, pages 293–300, Halifax, Nova Scotia, May 1990.

[Fli90] Scott Flinn. Visualizing program execution in real-time: Representing multi-processor real-time program behavior as a statistical view of state. Master's thesis, University of Waterloo, Dept. of Computer Science, Waterloo, Ontario, 1990.

[FM90] Steven K. Feiner and Kathleen R. McKeown. Coordinating text and graphics in explanation generation. In *Proceedings AAAI*, pages 442–449, Boston, MA, July 1990.

[Gar81] W. R. Garner. The analysis of unanalyzed perceptions. In Michael Kubovy and James R. Pomerantz, editors, *Perceptual Organization*, chapter 5. Lawrence Erlbaum Associates, 1981.

[GPW89] Georges Grinstein, Ronald Pickett, and Marian G. Williams. Exvis: An exploratory visualization environment. In *Graphics Interface '89*, London, Canada, 1989.

[Gre88] Donald P. Greenberg. Coons award lecture siggraph87. *Communications of the Association for Computing Machinery*, 31(2), February 1988.

[JLSU87] Jeffrey Joyce, Greg Lomow, Konrad Slind, and Brian Unger. Monitoring distributed systems. *ACM Transactions on Computer Systems*, 5(2):121–150, May 1987.

[Jul81] Béla Julész. Figure and ground perception in briefly presented isodipole textures. In Michael Kubovy and James R. Pomerantz, editors, *Perceptual Organization*, pages 27–54. Lawrence Erlbaum Associates, Hillsdale, NJ, 1981.

[KH81] Daniel Kahneman and Avishai Henik. Perceptual organization and attention. In Michael Kubovy and James R. Pomerantz, editors, *Perceptual Organization*, pages 181–211. Lawrence Erlbaum Associates, Hillsdale, NJ, 1981.

[KP81] Michael Kubovy and James R. Pomerantz, editors. *Perceptual Organization*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1981.

[LJSR89] David Alex Lamb, Nitin Jain, Wendy Sharp, and Arthur Ryman. Knowledge-based and formal methods of software design: A preliminary survey for the 4thought project. Technical Report TR 74.042, IBM Canada Laboratory, September 1989.

[LN77] Peter H. Lindsay and Donald A. Norman. *Human Information Processing: An Introduction to Psychology*. Academic Press, Inc., 111 Fifth Avenue, New York, NY 10003, 1977.

[Mac86a] Jock D. Mackinlay. Automatic Design of Graphical Presentations. Technical Report STAN-CS-86-1138, Stanford University Department of Computer Science, December 1986.

[Mac86b] Jock D. Mackinlay. Automating the Design of Graphical Presentations of Relational Information. *Association for Computing Machinery Transactions on Graphics*, 5(2):110–141, April 1986.

[Mon90] Fanya S. Montalvo. Diagram understanding: The symbolic descriptions behind the scenes. In Tadao Ichikawa, Erland Jungert, and Robert R. Korfhage, editors, *Visual Languages and Applications*. Plenum Publishing Corporation, 1990.

[Mur85] Gerald M. Murch. Color graphics – blessing or ballyhoo? In Ronald M. Baecker and William A. S. Buxton, editors, *Readings in Human-computer Interaction: A Multidisciplinary Approach*, pages 333–341. Morgan Kaufmann, 1985.

[Nei63] U. Neisser. Decision time without reaction time: Experiments in visual scanning. *American Journal of Psychology*, 76:376–385, 1963.

[PG88] Ronald Pickett and Georges Grinstein. Iconographic displays for visualizing multidimensional data. In *Proceedings of the 1988 IEEE Conference on Systems, Man and Cybernetics*, pages 514–519, Beijing and Shenyang, People's Republic of China, 1988.

[Pic88] Clifford A. Pickover. Pattern formation and chaos in networks. *Communications of the Association for Computing Machinery*, 31(2):136–151, February 1988.

[Pic91] Ronald Pickett. Integrated displays of multispectral imagery at air force geophysics laboratory. *Draft*, March 1991.

[Pom81] James R. Pomerantz. Perceptual organization in information processing. In Michael Kubovy and James R. Pomerantz, editors, *Perceptual Organization*, chapter 6. Lawrence Erlbaum Associates, 1981.

[Ram88] Vilayanur S. Ramachandran. Perceiving shape from shading. *Scientific American*, 259:76–83, August 1988.

[RM90] Steven F. Roth and Joe Mattis. Data Characterization for Intelligent Graphics Presentation. In *CHI'90 Proceedings*, pages 193–200, Seattle, WA, April 1990.

[RV90] Bernice E. Rogowitz and Richard F. Voss. Shape perception and low-dimension fractal boundary contours. In Bernice E. Rogowitz and Jan P. Allebach, editors, *Human Vision and Electronic Imaging: Models, Methods, and Applications*, February 1990.

[SBG90] Stuart Smith, R. Daniel Bergeron, and Georges G. Grinstein. Stereophonic and surface sound generation for exploratory data analysis. In *CHI'90 Proceedings*, pages 125–132, April 1990.

[SCS90] Frankie K. Sun, William B. Cowan, and Kellogg S.Booth. Understanding visual effects in a windowed environment. In *Proceedings Graphics Interface 1990*, pages 100–105, Halifax, Nova Scotia, May 1990.

[Ste46] S. S. Stevens. On the theory of scales of measurement. *Science*, 103:677–680, 1946.

[TCF+90] Anne Treisman, Patrick Cavanaugh, Burkhart Fischer, V.S. Ramachandran, and Rudiger von der Heydt. Form perception and attention: Striate cortex and beyond. In Lothar Spillmann and John S. Werner, editors, *Visual Perception: The Neurophysiological Foundations*, pages 273–316. Academic Press, 1990.

[TG88] Anne Treisman and Stephen Gormican. Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, 95:15–48, 1988.

[Tre86] Anne Treisman. Properties, parts, and objects. In Boff, Kaufmann, and Thomas, editors, *Handbook of Perception*, volume II, pages 35–1 to 35–70. 1986.

[Tre90] Anne Treisman. Features and objects in visual processing. In Irvin Rock, editor, *The Perceptual World: Readings from Scientific American*. W. H. Freeman and Company, New York, 1990. Originally November, 1986 Issue of Scientific American.

[WB85] Colin Ware and John C. Beatty. Using color as a tool in discrete data analysis. CS 21, University of Waterloo Computer Science Department, Waterloo, Canada, August 1985.

[Zip35] George Kingsley Zipf. *The Psycho-biology of Language*. Houghton-Mifflin, Boston, MA, 1935.