Towards a Synthetic Eye:
Psychological Issues in Data Visualization

by

Susan F. Chipman
Office of Naval Research
November, 1992

Typically, discussions of data visualization extol the wonderful powers of the human eye. Apparently instantaneous perception of structures in visual displays of data is contrasted to the unintelligibility of masses of numbers in traditional computer outputs. Period. Little is said about what the powers of that eye might actually be, or—more importantly—about the limits on those powers. If the full potential of our investments in visualization is to be realized, we need to understand much more both about the perceptual characteristics of the human visual system and about the integration of the perception of data displays into complex scientific thought processes. This paper outlines some of the research issues and questions that must be addressed to gain that improved understanding.

To a psychologist, such as myself, who has studied human sensitivity to various kinds of visual structure (Chipman, 1977, Chipman & Mendelson, 1979), it is striking that the illustrations of visual power almost always show strong symmetry or approximate symmetry. Consider the cover illustration of the NSF report on the Visualization in Scientific Computing (McCormick, DeFanti, & Brown, 1987). Human vision is extremely sensitive to symmetry. There is a large psychological literature demonstrating the impact of symmetry on almost any visual task—from detection to search to visual memory. It is highly probable that the human visual system has special computational mechanisms for the perception of symmetry. In contrast, there are many other kinds of structure that might be considered equivalent from an abstract, information-theoretic point of view that are quite ineffective perceptually.

In my own research I explored human sensitivity to structure in simple patterns of black and white squares, either 6x6 matrices or combinations of two of these in various spatial relationships.

Figure 1 illustrates 10 different kinds of structure that were studied in one of these experiments.

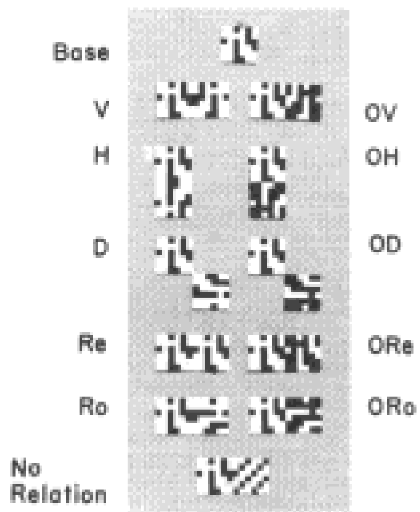

People were asked to judge the complexity of these patterns, using a method called magnitude estimation (Stevens, 1957, 1960). The instructions for these judgments allow the assignment of any number to the complexity of the first pattern seen and request that the next judgment represent the complexity ratio of the next pattern: if the next pattern is twice as complex, give it a number twice as large, and so on. Judgments are averaged using the geometric mean. The results of this experiment are shown in Figure 2.

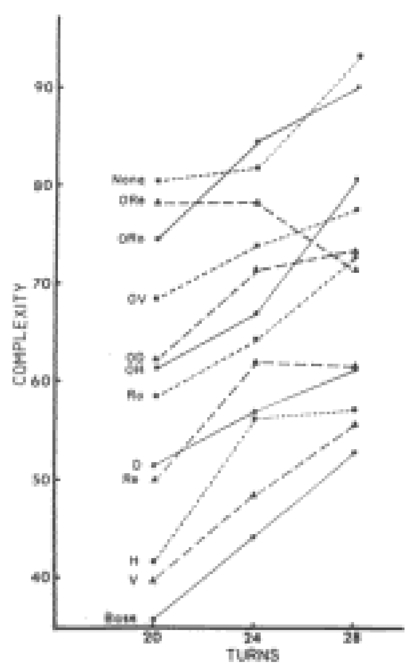

For random patterns with no perceptible structure, the number of corners (turns) on the black figure or the amount of perimeter relative to the area, determines the perceived complexity of the pattern. The plot of the 6x6 "base" patterns in Figure 2 illustrates this effect: these base patterns were chosen on the basis of previous experiments to have minimal perceived internal structure. The plot labeled "none" shows the results for the combination of two such unrelated patterns to make a pattern of double size and approximately double perceived complexity. The position of the other plots in between these two extremes measures the relative perceptual effectiveness of these other kinds of structure.

Notice that symmetry around the vertical axis is highly effective in reducing perceived complexity, closely reflecting the amount of objective information reduction (to one half plus the specification of symmetry) in the pattern, but even symmetry about other axes is less effective perceptually. The structural relationships labeled <u>opposition symmetries</u> are much less effective. They were studied for two reasons. Visual perception is known to rely heavily upon <u>opponent processes</u> for contrast enhancement, feature detection, and color perception. Therefore, it was plausible that our nervous system would be equipped to perceive the opposition symmetries. Also, all of these types of structure can be found in art, particularly American Indian art (Appleton, 1971). This is a kind of evidence for their perceptual effectiveness. An extreme case of opposition symmetry, incidentally, is the familiar checkerboard pattern which nests opposition symmetry at successive levels of organization.

As it turns out, the frequency with which these kinds of structure are found in art seems to roughly approximate their perceptual effectiveness. Obviously, forms of structure which are nearly equivalent informationally differ greatly in their perceptual effectiveness.

The evident lesson of such research is that people will perceive structure in visual displays of data if and only if the type of structure in the data display happens to be well-aligned with the analytic mechanisms of human vision. Thus, visualization is coming to be understood as, "... transformations that convert data into a format amenable to understanding by the human perceptual system, while maintaining data integrity," (Haber, 1989).

Some Limits on Human Perception

What are the "formats amenable to understanding by the human perceptual system"? What are some of the key known or likely limits on the powers of human vision that impact data visualization? As general background, it is important to know that primate visual systems seem to be organized into a number of rather separate parallel processing modules or streams: spatial information, object or pattern recognition, color perception, motion perception. The interaction and ultimate integration of information from these separate streams is not yet well understood, but the possible separateness and the problem of integration has many possible implications for data visualization. Some of these will be commented upon in the discussions of the perceptual topics below. Human visual perception is characterized by:

<u>Dimensionality</u>: Ordinarily, we perceive three dimensional space, although conventional data displays are two dimensional. Often, a two dimensional display will be perceived as a projection of some three dimensional entity, whether or not that is intended. On the other hand, it is known that when a three-dimensional graph is represented as a two-dimensional projection, little quantitative information is conveyed for the third dimension (Shah & Carpenter, 1992).

Newer virtual reality displays attempt to provide a third dimension with binocular disparity in the images presented to the two eyes and by representing the change in view (motion parallax) that occurs as a result of motion in the environment. For data displays, it is important to understand that a significant proportion of the population does not actually see depth as a result of binocular disparity. For all of us, binocular disparity is effective only to a distance of about 20 feet. Workers in data visualization are well aware of the limited dimensionality of human spatial vision because they are often struggling to represent and explore data sets of very high dimensionality. Motion or change in time is sometimes used to represent a fourth dimension, but this takes us over into another perceptual processing system; it is unclear whether it can provide another dimension on a par with the first three. In fact, motion information is typically processed in order to arrive at an understanding of the three dimensional structure of the environment (motion parallax).

Structural sensitivities and insensitivities. This was discussed above;it is very critical since data visualization is essentially about the detection and comprehension of structure. There are not many psychological research investigations of this issue, beyond the many demonstrations of the importance of symmetry.

Rates of movement and change. There is a certain range of values to which we are sensitive. Slow motion and time-lapse photography, for example, function to bring phenomena into our perceptual range. Visualization systems may have very high computational demands in order to bring change into our perceptual range or in order to effectively represent a variable with motion and change. As mentioned above, motion perception is a specialized component of the visual system with a long evolutionary history. It may have many quirks or special features that reflect the functional significance of motion perception during that evolutionary history; that is, it may not always work in the way we expect from a rational, engineering perspective. As mentioned above, it may not integrate well with space perception to provide a smoothly integrated fourth dimension.

Variables affecting the kinetic depth effect. This is a special case of the phenomena discussed above, the perception of depth as a result of motion. If possible, the human visual system makes sense of motion in a 2-dimensional display by interpreting it as a stable array moving in depth. Anecdotally, the perception of motion and kinetic depth seem to be strongly endorsed as the most powerful visualization tool. Martha Geller's discovery (Geller & Huchra, 1991) of the non-uniform distribution of galaxies in the universe, their clustering in approximate surfaces, is an excellent illustration of this. This example shows that this perceptual mechanism is quite powerful and general, not limited to simple or familiar forms.

Preferred sizes of elements in the display. The human visual system has a well-known spatial frequency sensitivity curve.
Furthermore, there is some reason to believe that the computational analysis of visual information may proceed in separate spatial frequency channels.

Spatial extent of the analysis. It is likely that the analysis of even symmetry relations has some capacity limitations, either in spatial extent or in number of elements, or both. Quite possibly, the reason for the ineffectiveness of the opposition symmetries, despite opponent processing in the visual system, is that the spatial extent of opponent processing is very limited. Opponent symmetries are perceived, but apparently by a time-consuming and effortful serial process. More generally, high resolution foveal vision is restricted to a few degrees of visual angle. Although we are not conscious of it, our perception of the world around us or of any large visual display is constructed from the products of many visual fixations, about 4 per second. Consequently, it is important to understand the following topic.

Role of eye movements in constructing perception. Successive eye movements may be driven by features that "automatically" capture attention, in the absence of other determinants of attention. Obviously, our visual systems must be adapted to do a good job of constructing a wholistic representation of the surrounding environment, and we have a great deal of experience doing that before we are ever exposed to visualized data. However, the mechanisms of this process are not understood. Whatever they are, they are likely to have important implications for the visualization of data. For example, it seems likely that the information that patches together the islands of high resolution vision may have quantitative limitations. We can probably process structural relationships within these islands of high resolution vision that are not accessible across a wider visual field. Obviously, the effective use of our eye movements to process structure in a visual delay is something that might require a great deal of learning. Some kind of internal memory organization is also required to support a visual analysis that involves many fixations. Either of these factors may explain why the developmental course of structure perception was found to be so protracted.

Color in pattern vision. In the early stages of visual processing, the processing of color and pattern or form information proceed separately. Color vision may be structure blind. It has been shown that color differences without brightness differences are very ineffective in producing form perception (Livingston & Hubel, 1987). Thus structure in visualized data that depends on color coding may not be well-perceived. Yet extensive use of colors in the attempt to provide extra dimensions for the representation of data variables is common in data visualization. Of course, color also has problems as a dimensional representation because the perceived relations among colors do not provide an analogue to a spatial dimension; instead they form the well-known color circle.

Object-like icons and textures. These are also used in our efforts to represent additional dimensions of data. At the first IEEE Visualization conference, a very interesting paper (Enns, 1990) claimed that some icons may engage special quick and dirty visual mechanisms designed to spot objects. These icons "pop-out" of the display, making themselves conspicuous. Others do not. Even for the icons that do pop out, it is an open empirical question whether we can easily perceive the structure in the visual arrays they form. How does the perception of such an array compare with the perception of an array based on brightness contrast? Texture perception is intimately involved with the notion of spatial frequency channels.

Other possible interactions. It might be, for example, that element size and the nature of the forms interact with other variables in our ability to perceive the kinetic depth effect. In general, whenever one crosses the boundary between one of the low-level visual processing modules and another, whenever one is calling for the integration of information across these systems, there are likely to be important interactions requiring research investigation.

All of these items lay out rather obvious and large sets of research questions that require investigation. The existing psychological literature should be systematically reviewed from the perspective of the needs of data visualization. That job is beyond the scope of this paper, but I have called out the technical terms that are used to label relevant work in the psychological literature. When that review is done, the results will reveal that many questions of importance to data visualization have not been studied by psychologists. In particular, one will find that most of the research has employed very simple and reduced displays of information. The contrast between the patterns used in my own research and the illustrations of data visualization will prove typical. There are two reasons for that fact. One is the desire to start simple and maintain strict experimental control. But many questions which have interested psychologists have not been studied because of the past technical difficulties of doing the studies. Psychologists have not had, and generally still do not have, ready access to the kinds of computational resources and display systems that permit efficient investigation of these issues. The typical size of psychological research budgets does not permit purchasing Silicon Graphics terminals, let alone supercomputers. Given the very substantial investments now being made in visualization facilities at supercomputer centers, psychologists should be incorporated in the research teams working with these facilities. Otherwise, it is unlikely that the full potential value of those investments will be realized.

Understanding the limitations of human vision is critical if one is going to transform data into a form amenable to understanding by the human perceptual system. But there may be a more radical alternative or supplement to that basic visualization goal.

Emulating and Generalizing Perception

Modern theories of visual perception are expressed in computational form. The computations that account for perceptual performance can become a basis for emulation of perception in an artificial system. For example, the experiments on the perception of structure described above were not limited to the clearcut cases graphed in Figure 2. Other experiments probed sensitivity to approximate and partial symmetries, and computer programs were developed to analyze the patterns, detecting and measuring these partial symmetries. Such programs could be used to locate the symmetries and other structures in a data display, to make masses of numbers intelligible in the same way that human vision does.

Furthermore, the mathematical or computational expression of perceptual structures may yield obvious generalizations of them. For example, a change of sign in the comparison process yields opposition symmetry rather than ordinary symmetry. In this way, we can build artificial systems that are sensitive to kinds of structure that the human eye sees dimly, if at all. Having found them computationally, we may be able to convert them to a form to which human vision is sensitive so that we can appreciate them.

There are other ways in which such a computational <u>synthetic eye</u> could help us overcome the limits of both human perception and computational power. The computational equivalents of perception can be carried out in N dimensions, at arbitrary rates. They need not be done fast enough satisfy a human user. Mathematicians and statisticians interested in the examination of high dimensional data sets have spoken about learning how to explore such a dataset by bringing parts of it into 2 or 3 dimensional view, by learning to do some sort of <u>grand tour</u> (Banchoff, 1986). It seems possible that a user of data visualization could have a computer do a thorough exploration, finding the structure that the user would perceive as interesting, and displaying it efficiently.

Beyond Perception to Scientific Cognition

Scientific data visualization involves more than perception. It is part of a process of scientific problem solving or scientific discovery. One seeks structure in the data that will prove theoretically meaningful. One may see structure that implies known mathematical relationships, for example. Understanding how the visualization of scientific data can be productive requires understanding its role within the larger thought processes of scientists. Workers in data visualization are already seeing the need to provide users with guidelines concerning the most effective ways to display and explore their data (Rosenblum, 1989). For the simpler case of conventional statistical graphics, a start in this direction has been made (Chambers, Cleveland, Kleiner & Tukey, 1983).

In the past few years, cognitive scientists have developed a strong interest in the role of imagery and diagrams in thinking, scientific problem solving and scientific discovery. Visualization of data is a sophisticated extension of these more traditional visual tools for thinking. A key contribution to this area of research is a paper by Larkin & Simon (1987), "Why a diagram is (sometimes) worth ten thousand words." Examining a few cases, they did a computational analysis of the value of diagrams in problem solving. They concluded that computational efficiency depends on both the organization of data structures and the processes operating on them. They suggest that the human visual system provides <u>perceptual operators</u> that generate inferences at almost no cost. There are <u>perceptual productions</u> that use visual pattern recognition to link into relevant knowledge in memory. Seeing structure that implies a known mathematical relationship in the data would be an example of that phenomenon.

The precise nature of our available <u>perceptual operators</u> remains to be discovered. Presumably they are closely related to or even identical with the kinds of structure which we perceive in displays, as discussed at the outset of this paper. One conjecture about this repertoire of structural kinds is that it is the set of structural relationships generated by the motion in the world that we are capable of (Chipman, 1977; Palmer, 1983) Quite possibly, if a motion we could make transforms one display into another, or one part of a display into another region, we will be able to perceive the structural relationships between the two.

Some of the perceptual mechanisms that support perceptual operators are undoubtedly built into the genetically determined organization of the visual system (cf Gallant, Braun & Van Essen, 1993). However, there is a complex interplay between basic perceptual mechanisms and learning from experience. Although the perceptual effectiveness of symmetry, especially symmetry about a vertical axis, seems to be present very early in life (Fisher, Ferdinandsen & Bornstein, 1981), sensitivity to the kinds of visual structure discussed in the beginning of this paper shows a surprisingly extended developmental course (Chipman & Mendelson, 1979). When it comes to "seeing" mathematical relationships in a visual display of data, there is no doubt about the role of learning and experience. In discussing visualization, the statistical graphics community places a great deal of emphasis on the notion of learning to see via experience with the visual presentation of known mathematical structure (Banchoff, 1986). But this is not yet a skill for which a systematic training regime has been developed. Similarly, learning to see the graphic representation of mathematical functions and relationships was the first educational application of computers to generate great excitement in the mathematics education community. Despite educators' interest, I do not believe there has been research on the most effective way to develop this knowledge of the mathematical significance of graphical displays of data.

Although our knowledge of perception gives us some hints about where to look for perceptual operators and the pattern recognition end of perceptual productions, research on this kind of scientific thinking is at the most preliminary stage. What are some of the methods that we could use to give it a serious start?

Cognitive science has often relied upon the naturalistic study of expertise, and that seems appropriate here. In a context where visualization techniques are in use, one can begin to study how they are used effectively. The perceptual features that are noticed and used can be recorded, ideally by observing the person at work, collecting "think aloud" protocols. As a still more preliminary approach, or an alternative if the critical episodes of insight are too rare, one can collect reports of past episodes from experienced users of visualization. One can examine the examples of successful use of visualization that have been published. Analysis of these episodes should seek to identify perceptual features or structures that were critical to the episode and to determine how those features were linked to either general mathematical relationships or specific knowledge of the subject matter of the data. One of the anecdotes in the NSF report, for example, involved discovering a display program bug by noticing a conspicuous reflection (McCormick et.al. 1987, p.5) An important category of examples to study would be those in which, as in this example, visualization has been used to steer, modify, or correct an on-going computation.

A catalog of the perceptual features or structures that seem to capture attention and trigger insightful discoveries about the data would have value in two ways. This would generate the information that we would need in order to produce a synthetic eye to do search and discovery work for us. For that purpose, we need a formal specification of the features that we find perceptually interesting. The same information would tell us how to train new users of visualization facilities.

Another category of phenomena that we would wish to garner from our observations is the interplay of perception and knowledge in the episodes of insight or discovery from data displays.

Let me give a simple example of this interplay from my own research on the perception of structure in patterns. The base patterns shown in Figure 1 represent a restricted range of relatively high complexity, but other experiments used a wider range so that it is interesting to look at complexity as a function of the number of turns, and that function is not linear. Usually, judgment data of this type are well fit by some power function. On the other hand, another method of judgment was used to study children's perception. When the data from the experiment with children were plotted across the full range of numbers of turns, the data for random patterns and the data for symmetrical patterns formed two parallel lines. See Figure 3, which shows data from the adult group in the developmental study.

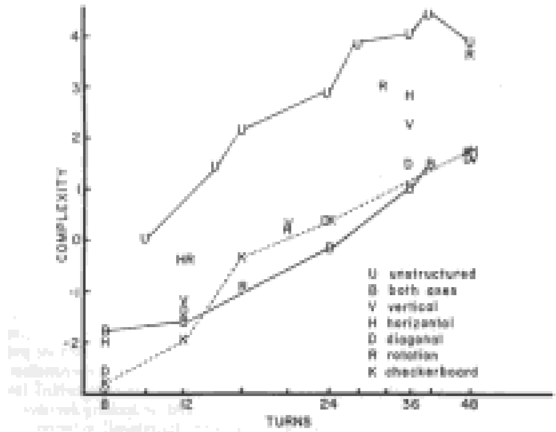

Two parallel lines is an example of the sort of perceptual feature that we are well-equipped to detect. I "saw" immediately that this meant they were related by a translation function, that the structure of the symmetrical pattern could be regarded as equivalent to reducing the number of turns by some constant fraction. This is an example of general mathematical knowledge of the mathematical meaning of graphic representations. My research collaborator, who was less mathematically inclined, not only did not see this; he was quite startled by my claim. Furthermore, knowing that the two methods of judgment are usually related by a log function allowed me to extend this relationship to the other data. This is an example of the role of specific knowledge of the scientific subject matter. (The effects of the different kinds of structure can be expressed concisely by a set of reduction factors that apply to the underlying physical "turns" variable and account for both types of judgment data.) Thus, this minor discovery had its origin in the visualization of data, but it depended heavily upon knowledge. The graph was of no use without the associated knowledge.

After the preliminary observational or retrospective report approach, it is possible to move on to more formal experimental research to test and confirm the hypotheses one is developing. Cognitive scientists have often chosen to compare and contrast the performance of novices and experts. In studying visualization, we would want to determine how what is "seen" in the display depends upon the experience and formal training of the individual viewing it. By contrasting what is seen by persons with varying degrees of knowledge of the scientific area of the data, one can separate the influence of general human perceptual capacities from the influence of special knowledge. For example, the approximate symmetry of the display is what strikes me about the cover illustration of the NSF visualization report, but the caption says that it shows, "the fundamental 'kink' mode of instability in a supersonic gas jet." Even for the expert, it is probably accurate to say that the kink becomes visible against the background of symmetrical organization that enables us to process the display as a whole.

Given some displays that are known to be interesting cases, we can investigate how relevant knowledge affects the units and structures that are perceived to be salient in a display. This can be done by collecting reports as people with varying backgrounds inspect the displays or by requiring people to recall what they can from the display. Psychological studies have shown enormous individual differences in what can be remembered from a circuit diagram as a function of knowledge of electronics (Egan & Schwartz, 1979)

Another interesting area for investigation would be the intuition that there are ways one can learn to explore a complex multidimensional dataset effectively. Is this true? Certainly we can observe experienced users of visualization facilities as they explore new data sets, determining whether and how they control the display in a goal oriented way. Is there in fact evidence of a systematic search strategy, of doing some sort of grand tour of the data? If not, then we probably really do need automated aids to exploration that take a brute force approach to exploration for which we ourselves would not have the patience. But skill can go beyond mere exploration of the data as it is first presented. The likelihood of gaining visual insights into data can be increased if one develops skill in manipulating and transforming displays or data in such a way as to engage our maximum perceptual strengths. For example, workers in statistical graphics (Chambers et. al, 1983) have noted that we are very good at perceiving straight lines and even minor deviations from them. If I had had only data from the first method of judgment, I might have discovered the same relationship by subjecting that data to a log transformation and seeing parallel lines. If I had been more skilled in interpreting graphed data than I was, I might have detected clues in the curves that I had that a log transformation might be worthwhile. (At that time I did not have access to a computer, so the effort of such an exploration was non-trivial.) Determining the circumstances that cue an expert to try such actions is part of the characteristic cognitive science approach: books written by practitioners often tell you what can be done but they rarely tell you how to decide when it might be worth doing. Often that information does not seem to be readily available for conscious report. Once determined, that information can guide a heuristic approach to data exploration, whether implemented by a human or a computer.

Formalization of skill with conventional statistical graphics is just beginning: The far richer world of data visualization must have many such tricks waiting to be discovered.

Conclusion

A systematic program of research that addressed the psychological issues I have raised here could yield these benefits:

-- An understanding of the perceptual factors affecting display designs and effective analytic strategies.

-- A basis for guidelines in carrying out analyses, a need recognized by workers in data visualization, a counterpart to Chambers, Cleveland, Kleiner & Tukey's Graphical Methods for Data Analysis (1983).

-- Approaches to training a new generation of scientists in the effective use of visualization facilities.

-- A basis for developing automated aids to the exploration of complex data sets for potentially interesting phenomena.

Footnotes

Dr. Susan F. Chipman is Program Manager for Cognitive Science, U.S. Office of Naval Research, Code 1142CS, 800 N. Quincy Street, Arlington, VA, 22217-5660. Phone: 703-696-4318. Email: chipman@nprdc.navy.mil. The ideas in this paper were originally developed in an effort to initiate a small targeted research program on the visualization of scientific data. That effort was not successful, and there are no specially designated funds for research on visualization available at the Office of Naval Research.

References

Appleton, L. H. (1971). American Indian Design and Decoration. New York: Dover.

Banchoff, T. F. (1986). Visualizing two-dimensional phenomena in four-dimensional space: A computer graphics approach. In E. J. Wegman & D. J. DePriest (Eds.), Statistical Image Processing and Graphics. New York & Basel: Marcel Dekker, Inc.

Chambers, J. M., Cleveland, W. S., Kleiner, B., & Tukey, P. A. (1983). Graphical Methods for Data Analysis. Boston: Duxbury Press.

Chipman, S. F. (1977). Complexity and structure in visual patterns. Journal of Experimental Psychology: General, 106, 269-301.

Chipman, S. F., & Mendelson, M. J. (1979). Influence of six types of visual structure on complexity judgments in children and adults. Journal of Experimental Psychology: Human Perception and Performance, 5, 365-378.

Egan, Dennis E.; Schwartz, Barry J. (1979) Chunking in recall of symbolic drawings. Memory & Cognition. Vol 7(2) 149-158.

Enns, J. T. (1990). The promise of finding effective geometric codes. Paper presented at Visualization '90 (IEEE, ACM/SIGGRAPH), San Francisco, October, 1990.

Fisher, C. B., Ferdinandsen, K., & Bornstein, M. H. (1981). The role of symmetry in infant form discrimination. Child Development, 52, 457-462.

Gallant, J. L., Braun, J., & VAn Essen, D. C. (1993) Selectivity for polar, hyperbolic, and Cartesian gratings in Macaque visual cortex. Science, 259, 100-103.

Geller, M. J., & Huchra, J. P. (1991, August). Mapping the universe. Sky and Telescope, 82, 134.

Haber, R. B. (1989, August). Scientific visualization and the rivers project at the National Center for Supercomputing Applications. Computer, 22(8), 84-89.

Larkin, J.H. & Simon, H.A. (1987) Why a diagram is (sometimes) worth ten thousand words. Cognitive Science, 11, 65-100.

Livingstone, M. S., & Hubel, D. H. (1987). Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. Journal of Neuroscience, 7, 3416-3468.

McCormick, B. H., DeFanti, T. A., & Brown, M. D. (1987, November). Visualization in Scientific Computing. Computer Graphics, 21(6).

Palmer, S. E. (1983. The psychology of perceptual organization: A transformational approach. In: J. Beck, B. Hope, L.A. Rosenfeld (Eds) Human and Machine Vision, New York: Academic Press.

Rosenblum, L. M. (1989, August). Scientific visualization at research laboratories. Computer, 22(8), 68-70.

Shah, P., & Carpenter, P. A. (1992). Conceptual limitations in comprehending line graphs. Unpublished manuscript, Department of Psychology, Carnegie Mellon University.

Stevens, S. S. (1957). On the psychophysical law. Psychological Review, 64, 153-181.

Stevens, S. S. (1966). A metric for the social consensus. Science, 151, 530-541.