Human Perception and Visualization

Moderator: Laurie Hodges Reuter, *Bellcore*

Panelists: Paul Tukey, *Bellcore*
Laurence T. Maloney, *New York University*
John R. Pani, *Emory University*
Stuart Smith, *University of Lowell*

Introduction

Laurie Reuter, *Bellcore*

As scientific data visualization emerges from its infancy, pretty pictures are not enough. If we consider that the ultimate goal of data visualization is to aid scientists with the exploration of their data and discovery of new facts, then certain types of visually dazzling presentations may actually be counterproductive scientifically. Even less complex visualizations sometimes use colors or symbols that make it difficult to accurately interpret the data. Increasingly, creators of data visualizations need to be sensitive to what types of presentations will enhance and complement the perceptual skills of the scientists who are the intended users of data visualization.

The power and usefulness of scientific data visualization is due largely to the strength of human perception. The human visual system is the major player in visualization, but the other senses, in particular the auditory and tactile senses, are starting to be applied to visualization tasks as well. Fortunately, there is a very large body of work in perceptual psychology that has explored the functioning of the human perceptual system. Unfortunately, the typical visualization scientist may have little to no background in perceptual psychology. In such a broad and voluminous literature, it is difficult for the visualization scientist to know where to begin probing for practical information on how to create perceptually effective visualizations.

The purpose of this panel is to introduce some of the fundamentals of human perception to an audience of visualization creators and users. The panel discussion is intended to raise the "perceptual consciousness" of visualization scientists and offer suggestions for where they can learn more about applying the theories of human perception. The emphasis will be on useful tips and guidelines and will include warnings about potential pitfalls, as related to human perception. Users of data visualization will benefit from the panel by learning how to recognize good visualizations that apply principles of perception. The discussion will also identify areas where further research in visualization and perception is needed.

Each of the panelists represents a different area of perception expertise. *Paul Tukey* is a Member of Technical Staff in the Statistics and Economics Research Division of Bellcore. His remarks are on the effective presentation of statistical data. *Laurence Maloney*, an Associate Professor in the Department of Psychology and Center for Neural Science at New York University, discusses data representation using color. *John Pani* is an Assistant Professor in the Department of Psychology at Emory University and he will offer suggestions for the display of orientation and motion of objects in the pictorial display of three-dimensional structures. *Stuart Smith*, an Associate Professor in the Department of Computer Science at the University of Lowell, discusses the use of sound to represent multidimensional data. The moderator, *Laurie Reuter*, is a Member of Technical Staff in the Computer Graphics and Interactive Media Research Group at Bellcore. Position statements and brief bibliographies of references from each of the four panelists follow.

Visual Perception and Scientific Data Display

Paul Tukey, *Bellcore*

When statistical or scientific data is presented in a plot or other graphical display, the viewer is asked to use his or her visual perception to make quantitative judgements about the display, to compare relative sizes, locations, orientations, colors, densities, textures, etc, of the elements of the display. Rarely does the designer of the display take into account the fact that human visual perception is a very complicated and subjective process, and that the effectiveness of the display for conveying objective understanding hinges crucially on a wide range of subtle factors, only some of which are under the control of the person making the display.

As *data visualization* becomes more and more widespread in science, both because today's computers and display hardware make it easy to produce pictures, and because pictures have inherent power to convey complex information, it becomes more and more urgent to study these issues of visual perception in order to choose types of displays that avoid the worst pitfalls and convey the relevant information as objectively and effectively as possible.

Perceptual psychologists and graphic designers have long understood many of the issues involved, but very little of their wisdom appears to have seeped into other fields of science where data is being extensively plotted and made into pictures. This is partly due to the relative isolation of various scientific disciplines (we don't read each other's journals, go to each other's meetings, etc.), but it is also partly due to a problem of communication and language. Each field develops its own language (jargon, perhaps), so that even when we *do* talk to each other about these issues, we fail to communicate fully.

Everyone is familiar with a certain basic set of optical illusions: parallel lines that don't look parallel, equal-length lines that appear to be of different length, squares of identical color that appear as different colors. In each case, one's perception is confounded by other elements of the display. Optical illusions dramatically illustrate how components of a picture can interact with each other to affect perception. What people *see* in a picture may not be what is actually (objectively) there.

To complicate matters further, different people may have different perceptions of the same picture, depending on physiology (color blindness and other visual impairments, for instance) but also depending on prior experience and what they are ''looking for''.

It is unlikely that perceptual psychologists and graphic designers can provide a complete inventory of all the factors and interactions that affect visual perception. And if they could, it would still be a daunting task to determine the extent to which those factors might come into play in a particular display of a particular set of data.

Given this situation, it might seem hopeless to design good data graphics. Luckily, the opposite is true: although we may never fully understand all the subtleties, certain broad principles emerge, and to the extent that we understand them, we can make judicious choices among all the currently available display alternatives, and we can use them as the basis for new kinds of graphical displays. Indeed, many important principles of good data display don't depend on deep understanding of psychophysics, but on just plain common sense (and a little thinking about the problem).

The media abound with examples of poor and misleading graphics. They can be found in newspapers and news magazines, but also, surprisingly, in scientific journals. A typical example is a popular class of displays which includes the ''shrinking dollar'' plot: dollar bills are drawn at various sizes for various years to represent their real (inflation-adjusted) values. (Just for fun, George Washington is replaced by the recent presidents, in turn.) But the person who made the plot has scaled the *linear* dimension of the dollars to be proportional to real value, failing to recognize that people perceiving the *size* of an object are more likely to perceive its area than its linear dimension. Thus a dollar bill ''twice as large'' as another (corresponding to an inflation factor of 2) will be perceived as *four times* as large by the viewer. In the most blatant cases of this kind it is fair to ask whether the designer has deliberately taken advantage of the perceptual confusion.

Another simple example is a line chart showing recent profit figures for a corporation (or stock prices, etc), where we are invited to be shocked by the erratic behavior of the time series, whereas, in fact, the plot fails to show the origin (zero point) of the vertical scale, and thus fails to convey the fact that recent variations in the series are only a tiny fraction of the prevailing level of the series. Put another way, the plot fails to provide a visual ''yardstick'' to judge the importance of the observed variation. This is closely connected to a basic principle of responsible statistical data analysis which says that point estimates should always be accompanied by some assessment of their variability or reliability.

Some other general principles that will be expanded in the presentation are:

- Careful choice and expression of scales is essential.
- Objects are perceived in relation to their surroundings.
- Straight lines are easier to perceive than curves.
- Horizontal lines are easier to perceive than oblique lines.
- Things that are closer together are easier to compare than things far apart.
- Things of equal importance should have roughly equal visual impact.

- Irrelevant material ("chart junk", as Ed Tufte calls it) can seriously interfere with a plot.
- Motion is more effective for conveying 3-dimensional depth than stereopsis or perspective.
- Principles of good graphical display are often in conflict with each other, necessitating trades-offs among them.

In summary, if *Scientific Visualization* is ever to amount to more than just a series of pretty pictures, then issues of good graphic design and human visual perception must be taken into consideration by those who design these displays in order to make them convey the insights that they are meant to convey.

[1] Binns, B., and N. Clark, "The Graphic Designer's Use of Visual Syntax", in *1986 Proceedings of the Section on Statistical Graphics of the American Statistical Association*, pp. 36-41, 1986.

[2] Chambers, J. M., W. S. Cleveland, B. Kleiner, and P. Tukey, *Graphical Methods for Data Analysis*, [esp. Ch. 8: "General Principles and Techniques"], Belmont, CA: Wadsworth Press, 1983.

[3] Cleveland, W. S., *The Elements of Graphing Data*, Monterey, CA: Wadsworth, 1985.

[4] Cleveland, W. S., "Graphs in Scientific Publications", *The American Statistician*, Vol. 38, pp. 261-269, 1984.

[5] Dalal, Siddhartha R., Fowlkes, E. B., and Hoadley, B., "Risk Analysis of the Space Shuttle: Pre-*Challenger* Prediction of Failure", *Journal of the American Statistical Association, Vol. 84, pp. 945-957, 1989.*

[6] Hochberg, J., and D. H. Krantz, "Perceptual Properties of Statistical Graphs", in *1986 Proceedings of the Section on Statistical Graphics of the American Statistical Association*, pp. 29-35, 1986.

[7] Huff, D., *How to Lie with Statistics*, New York: Norton, 1954.

[8] Kruskal, W. H., "Criteria for Judging Statistical Graphics", *Utilitas Mathematica*, Vol. 21B, pp. 283-310, 1982.

[9] Laseau, P., *Graphic Thinking for Architects and Designers*, New York: Van Nostrand Reinhold, 1980.

[10] Tufte, E., *The Visual Display of Quantitative Information*, Cheshire, CT: Graphics Press, 1983.

[11] Tukey, J. W., and P. A. Tukey, in V. Barnett (ed.), *Interpreting Multivariate Data*, pp. 189-275, Chichester, UK: Wiley, 1981.

[12] Tukey, J. W., *The Collected Works of J. W. Tukey*, Monterey, CA: Wadsworth, 1984.

[13] Tukey, P., "A Data Analyst's View of Statistical Plots", in *1986 Proceedings of the Section on Statistical Graphics of the American Statistical Association*, pp. 21-29, 1986.

[14] Wainer, H. and D. Thissen, "Graphical Data Analysis", *Annual Review of Psychology*, Vol. 32, 191-241, 1981.

[15] Wainer, H., "Making Newspaper Graphs Fit to Print", in P. A. Kolers, et al. (eds.) *Processing of Visible Language 2*, p. 139, New York.

Data Representation Using Color

Laurence Maloney, *New York University*

The initial visual information that determines color appearance in human vision depends as much on the lighting in a scene as on the spectral properties of surfaces in the scene. A visual system that bases color appearance on the properties of surface, discounting the contribution of the illuminant, is termed *color constant*. In a color constant visual system, then, color serves to represent certain objective data about surfaces. A better understanding of the mapping from surface properties to color in human vision could suggest novel methods for representing arbitrary data sets effectively using color.

Recently, several authors have succeeded in designing algorithms that allow vision systems to be color constant. These algorithms share common assumptions about the statistical properties of illuminants and surface reflectance functions present in the environment. I describe these assumptions and suggest methods of using color to represent arbitrary data with analogous statistical properties.

[1] Maloney, L. T., and B. A. Wandell, "Color constancy: A computational methods for recovering surface spectral reflectance," *Journal of the Optical Society of America A*, Vol. 3, pp. 29-33, 1986.

[2] Maloney, L. T., "Evaluation of linear models of surface spectral reflectance with small numbers of parameters," *Journal of the Optical Society of America A*, Vol. 3, pp. 29-33, 1986.

[3] Hilbert, D. R., *Color and Color Perception: A Study in Anthropocentric Realism*, Stanford, CA: CSLI, 1987.

Suggestions for the Display of Orientation and Motion of Objects in the Pictorial Display of Three-Dimensional Structures

John Pani, *Emory University*

The problem of using flat displays to give the impression of three-dimensional space is old and has a variety of solutions. Different solutions are better depending on the objects to be portrayed, the need for information, and the resources available for creating the displays. The discussion summarized here addresses the presentation of different views of an object, including the display of rotational motion of the object. In general, I assume that an object is shown in polar projection, so that the impression of depth is maximized. Below are six recommendations for creating effective displays of three-dimensional structure. These recommendations come from perception psychology, architectural drawing, and current practices in computerized display of spatial information.

1. *Assign directions to an object (i.e., top, bottom, front, and back).* That is, give the object a standard orientation with respect to the recognized spatial reference frame. In cases where the object has no intrinsic standard orientation, an orientation should be invented and standardized. Mapmakers did this when they decided that north would be at the top of maps. Without standardization of object orientations, recognition will be inefficient and will fail in many cases [1, 3, 5].

2. *For some displays, provide an explicit spatial reference frame aligned with the directions of the object.* This common practice is particularly helpful when objects are relatively nonrectilinear (i.e., amorphous, curvilinear, or contain many slanted surfaces). The spatial frame appears to serve at least three functions. First, the frame aids in identifying directions of the object (e.g., front and back). Second, because we are familiar with rectilinear solids, the frame tends to disambiguate the slant of surfaces. Specifying the slant of a surface helps to specify the form of the surface [4]. Third, the spatial frame facilitates relating different views of the object to each other.

3. *Make standard perspective and orthographic views immediately available.* Perspective views provide a single look at all three dimensions of an object (e.g., front, top, and one side). Thus, perspective views often are informative about the object as a whole. Orthographic views (i.e., in which the picture plane is parallel to a critical plane of the object) provide a clear look at planar relationships. There are two ways to characterize this function of orthographic views. First orthographic views display planes independent of the effects of linear perspective and foreshortening. Second, orthographic views display surfaces in the frontal plane, and the orientation of the frontal plane is perceived very accurately. When we are certain about the orientation of a displayed surface, we are certain about its form [4].

4. *Provide simulations of rotational motion of the object in depth.* There are two reasons for this. First, even small rotational motions provide very effective depth information [6, 7]. Second, rotational motion is very effective at conveying the coherence of the object across different views. Multiple views of an object permit a look at structure that may have been occluded in other views (e.g., the back of an object). Multiple views also permit displaying both perspective and orthographic views. Rotation can make immediately clear how a front view relates to a side view, or how an orthographic view relates to a perspective view.

5. *Use rotational motions that best provide coherence to multiple views of an object.* Two variables are of prime importance to determining the degree to which people find a rotational motion of an object coherent: The orientation of the object to the axis of rotation, and the orientation of the axis of rotation to the recognized spatial reference frame [2].

A. People comprehend rotational motion best when objects are square to the axis of rotation (e.g., as wagon wheels or paddle wheels are square to their axles). Display programs should permit people to align axes of rotation with edges surfaces, or principle axes of the object. With axes aligned to the object, rotation in depth can relate together a set of perspective views or a set of orthographic views.

B. People best comprehend rotations in which the axes of rotation are square to the recognized spatial reference frame. All else equal, people comprehend rotations about the vertical axis most easily.

C. If an object must be slanted to the axis of rotation (e.g., a radar dish), it is important to have the axes of rotation square to the spatial reference frame. Again, when all else is equal, people comprehend rotation about the vertical axis best.

D. Except in unusual circumstances, avoid using rotations in which an object is slanted to the axis of rotation and the axis of rotation is slanted to the recognized spatial reference frame. People find such motions confusing. In this case, the use of motion to provide coherence across different views of an object will probably fail. (Ironically, such rotations provide both perspective and orthographic views within single motions.)

6. *Exploration of unfamiliar relationships among parts of an object may be aided by violating the above recommendations.* For example, consider a cube tipped onto a corner, so that a main diagonal through the cube becomes vertical (in violation of 1, above) [1]. If the cube is rotated about that vertical (in violation of 5A), the motion is rather unfamiliar [2]. However, relationships among parts of the cube previously missed by most people become apparent (e.g., opposing triples of diamond-shaped surfaces) [1].

In general, our intuitions guide us well in the design of effective displays of familiar objects. It is when display systems are designed for novel objects with uncertain properties that it is important to be explicit about what will constitute an effective display system.

[1] Hinton, G., "Some Demonstrations of the Effects of Structural Descriptions in Mental Imagery," *Cognitive Science,* Vol. 3, pp. 231-250, 1979.

[2] Pani, J. R., "Axis of Rotation and Orientation of the Object in Mental Representation of Rotation," Poster presented at the thirtieth annual meeting of The Psychonomic Society, Atlanta, GA, 1989.

[3] Rock, I., *Orientation and Form,* New York: Academic Press, 1973.

[4] Rock, I., *An Introduction to Perception,* New York: Macmillan, 1975.

[5] Rock, I., *The Logic of Perception,* Cambridge, MA: MIT Press, 1983.

[6] Ullman, S., *The Interpretation of Visual Motion,* Cambridge, MA: MIT Press, 1979.

[7] Wallach, H., and D. N. O'Connell, "The Kinetic Depth Effect," *Journal of Experimental Psychology,* Vol. 45, pp. 205-217, 1953.

Representing Data with Sound

Stuart Smith, *University of Lowell*

Sound is used in the human-computer interface in two quite different ways. The first way is to provide cues for the state of a system or the occurrence of important events. Bill Gaver's Sonic Finder for the Macintosh is perhaps the preeminent example of this use of sound. Sonic Finder generates auditory cues for such events as the disposal of an object in the Trash Can (it produces a metallic "clunk" when the user clicks on the Trash Can). The second way sound is used is as a medium for representing data. Here, the values of various sound parameters—pitch, loudness, duration, and so on—represent the values of multidimensional data. Interesting examples of this use have been produced by Sara Bly, Steve Frysinger, and Dave Lunney. My own work is also in this area.

These two uses of sound do not overlap to any great extent. Their goals, and the techniques used to achieve them, differ significantly. It appears that they even depend on different perceptual capacities. When sound is used to provide cues, it is important that the user be able to identify the source of the sound rapidly and effortlessly. When sound is used for data representation, it is important that the user be able to hear the data-bearing relationships within and between sounds; the focus is on the sound itself, as when we listen to music. A current school of thought holds that our auditory system is optimized for the sound source identification task, and most people seem to do this task well. Musical skills—which are presumably required to some degree for the task of decoding an auditory data representation—are very unevenly distributed in the population, and it is clear that training makes a big difference in performance.

The use of sound for data representation is the auditory counterpart of data visualization. Bill Buxton suggests that this activity be called "sonification." While this field is attracting increasing attention, it must overcome at least three major obstacles if it is to grow. The first is the prevailing sonification model, which is simply to map data to sound parameters arbitrarily. The resulting sound is typically unpleasant and lacking in any "natural" connection to the data represented (one intuitively feels that medical images, for example, ought to sound somehow different from demographic data or satellite imagery). Models of sonification more sensitive to the kinds of data presented must be developed.

The second major obstacle is the lack of suitable sound generation hardware. Sonification requires a general purpose real-time sound synthesis capability, preferably at an affordable price. The ubiquitous MIDI devices are supposedly inexpensive and very flexible, but the benefits of MIDI equipment are mostly illusory. MIDI devices are designed primarily for the performance of rock and pop music, and they are very difficult to bend to other purposes. Furthermore, MIDI equipment is not really inexpensive. Although individual devices may be cheap, a complete system with which you can do experimental work is not. There is some hope that the new generation of digital signal processing chips will provide the basis for a powerful real-time sound generation capability that can be incorporated into a low-cost workstation.

Finally, the third major obstacle is the nearly total absence of the kinds of models that allow the design of computer graphics software systems that can run successfully on hardware made by many different manufacturers. The principal reasons for this situation are the lack of a satisfying comprehensive theory of timbre perception and the lack of an agreed-upon theory of timbre generation. These translate directly into the situation we observe today: multiple incompatible sound-generation devices, each accompanied by its own suite of non-standard applications packages. As a consequence, workers in this field usually have to start a project by designing and building a one-of-a-kind sound system rather than going straight to work on the interesting research questions. Such systems may not work as hoped, thereby creating the risk of costly and time-consuming dead-ends. There are no easy answers to this problem.

[1] Bly, S. A., "Presenting Information in Sound," Proceedings of the CHI '82 Conference on Human Factors in Computer Systems, pp. 371-375, 1982.

[2] Buxton, W., "The Use of Non-Speech Audio at the Interface," Tutorial #10, CHI '89, pp. 2.1-2.15, 1989.

[3] Gaver, William W., "The SonicFinder: an interface that uses auditory icons," *Human-Computer Interaction,* Vol. 4, No. 1, pp. 67-94, 1989.

[4] Grinstein, Georges, and Stuart Smith, "The Perceptualization of Scientific Data, " Proceedings of the SPIE/SPSE Conference on Electronic Imaging, Santa Clara, CA, 1990, to appear.

[5] Lunney, D. and R.C. Morrison, "High technology laboratory aids for visually handicapped chemistry students," *Journal of Chemical Education,* Vol. 58, No. 3, pp. 228-231, 1981.

[6] Smith, Stuart, "An auditory display for exploratory visualization of multidimensional data" Proceedings of the IFIP Conference on Workstations, Lowell, MA, 1989. Springer-Verlag, to appear.

[7] Smith, Stuart, and Marian G. Williams, "The Use of Sound in an Exploratory Visualization Environment," University of Lowell Computer Science Department Technical Report #R-89-002, 1989.

[8] Smith, Stuart, R. Daniel Bergeron, and Georges Grinstein, "Stereophonic and surface sound generation for exploratory data analysis," CHI '90 Conference Proceedings, Reading, MA: Addison-Wesley, 1990.

[9] Mezrich, J.J., S. Frysinger, and R. Slivjanovski, "Dynamic Representation of Multivariate Time Series Data," *Journal of the American Statistical Association,* Vol. 79, No. 385, pp. 34-40, 1984.

[10] Wenzel, Elizabeth M., Frederic L. Wightman, and Scott H. Foster, "Development of a three-dimensional auditory display system," *SIGCHI Bulletin,* Vol. 20, No. 2, pp. 52-57, 1988.

[11] Williams, Marian G., Stuart Smith, and Giampiero Pecelli, "Computer-Human Interaction Issues in the Design of an Intelligent Workstation for Scientific Data Visualization (Phase 1)," *ACM SIGCHI Bulletin,* April 1990.

[12] Yeung, E.S., "Pattern Recognition by Audio Representation of Multivariate Analytical Data," *Analytical Chemistry,* Vol. 52, No. 7, pp. 1120-1123, 1980.

Mailing Addresses for the Participants

- Laurie Reuter
  Bellcore, MRE 2M-344
  445 South St.
  Morristown, NJ 07960-1910
  reuter@bellcore.com
- Paul Tukey
  Bellcore, MRE 2M-391
  445 South St.
  Morristown, NJ 07960-1910
  paul@bellcore.com
- Laurence Maloney
  Department of Psychology
  Center for Neural Science
  6 Washington Place, 8th floor
  New York, NY 10003
  ltm@xp.psych.nyu.edu
- John Pani
  Department of Psychology
  Emory University
  Atlanta, GA 30322
- Stuart Smith
  Department of Computer Science
  University of Lowell
  Lowell, MA 01854
  stu@ulowell.edu