

Information Visualization and the Next Generation Workspace

Ramana Rao <rao@parc.xerox.com>

Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304

Draft of January 17, 1996

Pictures have gotten a bad rap for a long time. Arnheim points out that in the Hebrew tradition, the story of the long hostility toward images begins with the destruction of a piece of sculpture, the golden calf which Abraham burnt in the fire, ground to powder, strew across the water and made the children of Israel drink of. In modern times, many disciplines rely on pictures and diagrams. Indeed, much of modern math and science relies on the invention of the Cartesian plane. A typical example is provided by many economics textbooks which are littered with supply, demand, and utility curves. Much of the intuitive reasoning and explanation is expressed in terms of moving curves or intersection points around. Moreover, stories of the role of visual thinking underly many modern scientific discoveries, for example there are stories of Einstein's riding of particles in his mind's eye.

Visual representations of complex phenomena can enhance or extend our ability to understand these phenomena. This insight has fueled the application of interactive 3d graphics to the visualization of phenomena studied by scientist of many stripes. As with the microscope and the telescope, and the stroboscopic camera and time-lapsed photography, the scientist is able to see phenomena of scales miniscule to galactic, and of durations instantaneous to glacial. Thus over the last few years scientific visualization has been applied to molecular modeling, atmospheric simulations, medical imaging, and the study of fluids and flows, surfaces and volumes. Yet in all of these cases, there is an essential connection to the underlying natural structure of universe.

More recently, graphical techniques have been increasingly applied to problems of visualizing volumes of information of a more artificial or human-made nature: the data and documents of corporations and countries, of stock markets and sports statistics, of goods and services, of computer file systems and the global information network. The work of Card, Mackinlay, and Robertson of Xerox PARC in particular has stimulated a new area of human-computer interaction called information visualization, which is about the building of interactive, visual representations of large amounts of abstract information. To be sure, the bound-

ary between scientific and information visualization is not clearly demarcated—since for example, coloring of a terrain according to temperature seems hardly different from marking the points of a map according to the population or pollution levels—but rather a continuum in which the dimensions of the physical world gradually let go of the dimensions of the picture.

During the last few years, researchers have designed a number of so-called "information visualizations" which can be thought of as new kinds of sensemaking devices. They are what Don Norman calls "cognitive artifacts" — artificial devices designed to maintain, display, or operate upon information. As such, a visualization is a *representation* not just a presentation. They aren't just pretty pictures with which we persuade our managers or our funders. They are tools to extend our ability to think and make sense of the universe.

Three aspects of visualizations are particularly important. They are, well of course, visual representations. The power of such representations partially lies in (as we often chant in our group at Xerox PARC) "offloading the cognitive onto the perceptual." If the equipment that exists between the surface of the eye and the depth of the mind can be shanghaied into the service of extracting sense or meaning from the arrangements of information in space, then more of our conscious effort can be applied to higher order processes of thinking about meaning and action. Second, recent work has focussed on increasing the amount of information that can be examined and manipulated without increasing the burden on the user. It is more and more common to encounter datasets and document collections in office work, in education and commerce, and even in the home that are unmanageable with the prevailing graphical user interface paradigm. The third important aspect of visualizations is that they are interactive, dynamic structures. The soul of computation has been breathed into the formerly inert forms of visual representations that have already played a major role in the history of science and thought. Lets consider each of these aspects in more detail.




Figure 1: These two figures exemplify scientific visualizations in which natural phenomena are represented in a 3d space roughly corresponding to the real world at some scale. Visual techniques are used to represent associated values or properties. The right figure [Courtesy of R. Crawfis and M. Allison, Lawrence Livermore National Laboratory, CA, USA; Modeling data provided by Gerald Potter, Lawrence Livermore National Laboratory, CA, USA] shows wind velocity with texture and heat leaving the earth with color. Two hurricanes are visible off the coast of Central America. The right figure [Still trying to obtain Molecule figure] shows a Nutrasweet molecule with colored clouds representing charge distribution.

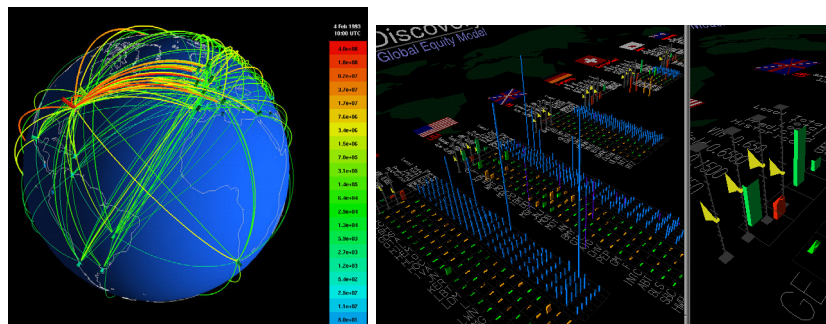


Figure 2: These two figures exemplify information visualization which apply visual techniques to information from the artificial world. The left figure [Courtesy of Steve Eich, AT&T Bell Laboratories] shows a frame from an animation of worldwide Internet traffic. The right figure [Courtesy of Visible Decisions, Inc] shows the yields, performances, risks, etc. of various financial instruments in different countries.

The Power of Graphical Representation

Why is a Diagram (Sometimes) Worth Ten Thousand Words? A good question that is the subject of a 1987 paper by cognitive scientists Larkin and Simon. Their answer is expressed in computational terms. Most computer scientists probably remember the first time they were exposed to the notion of data structures, the core elements of all computer programs. A certain wonder arises from the simple idea that different arrangement of data can behave the same functionally while offering radically different space and time performance profiles for different operations. For computer scientists, it is an early brush with the notion of a tradeoff which is always close at hand during the act of design.

Essentially the thrust of the Larkin/Simon argument lies in the explanation that diagrammatic and sentential representations are different data structures with different computational capabilities. Using a production rule model, a formalism common in artificial intelligence and cognitive science, they show how a diagram can lead to faster computations than a set of sentences. The diagrammatic representations benefits greatly from perceptual recognition apparatus: humans are highly sensitive to the exact form (representation) in which information is presented to senses. Certain elements and arrangements like smooth curves, maxima, discontinuities are perceptually salient and can thus be detected rapidly in diagrams.

[picture of pulley problem from Larkin/Simon]

Larkin/Simon model "perceptual enhancement" by adding formal elements for everything a human can readily perceive. In the sentential description, all of these elements must be painfully deduced and constructed by explicit mental rules. While the diagram along with the eye provides these elements trivially. This is certainly a close accounting of the process of offloading of cognitive onto the perceptual. Thus we can arrange scratches in dirt, cuts in wood, balls in space, marks on paper, and bits on screens to "represent" information to our recognition apparatus to increase the speed and the reach of our thinking. In philosopher Daniel Dennett's words, "we make graphs and maps and all manner of color-coded plottings so that the sought-for regularities and saliencies will just "pop out" at us, thanks to our visual system."

Yet, we really only win, and win big, when these external representations map important concepts onto salient percepts in useful ways. It isn't simply a matter of "seeing" what is meant ("see" often shows up in scare quotes in the discussion of vision and cognition). Although the notion of concepts "popping out" is something that many of us appreciate intuitively, if the representations aren't designed well, then nothing, or perhaps even the wrong things, will pop out. Not any old graph or diagram will do. Tufte has published a couple of elegant books pointing this out.

Another property of graphical representation is that use and understanding of them must be learned, just like lan-

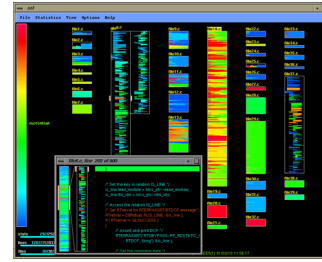


Figure 3: [Courtesy of Steve Eich, AT&T Bell Laboratories] This SeeSoft display shows over 12,000 lines of software. The color scale indicates the recency of modification with red indicating the most recent and blue the oldest. A separate window provides a more detailed view of a portion of the data.

guage or arithmetic. What is being learned is, not for example, how to see edges or smooth curves or discontinuities, but rather how such perceptual objects map onto meaning. Some humorless German philosopher said percepts without concepts are blind. In good diagrams, perceptual objects align with conceptual objects. In fact, so much so that perceptual objects become stand-ins for the underlying conceptual objects. We conveniently traffic in, even confuse, them as the thing referred to by them. Just as transparently as the Fitts law mouse becomes the pointer on the screen. Just as transparently as synecdoche, metonymy, and all their crazy linguistic But for this to be a useful process, the graphical representations must be well-designed, and how to "see" them must be learned. The burden is shared between the writer/designer and the reader/user.

Dealing With More Information

The modern reality is that even the simplest of tasks involves much more information than can be easily explored without visualizations. Say you want to buy a car or a laptop computer. There are hundreds of choices and certainly one problem is that information which may influence your decision may be quite scattered. This is why specialized magazines provide tables and services like consumer reports regularizes information about different products. The problem remains how to make sense of all the numbers and text that you have entered into the tables. With hundreds of cases and dozens of variables, the average person is left with no recourse other than to make some quick simplifying decisions that limit the amount of information that they consider.

I've been asked before, "Why do you want to deal with more information? I would prefer less information, just the right information." Well this person actually wants the same thing as us: we all want more information to make a difference so that we can focus on the information that really matters. This can happen naturally with visual rep-

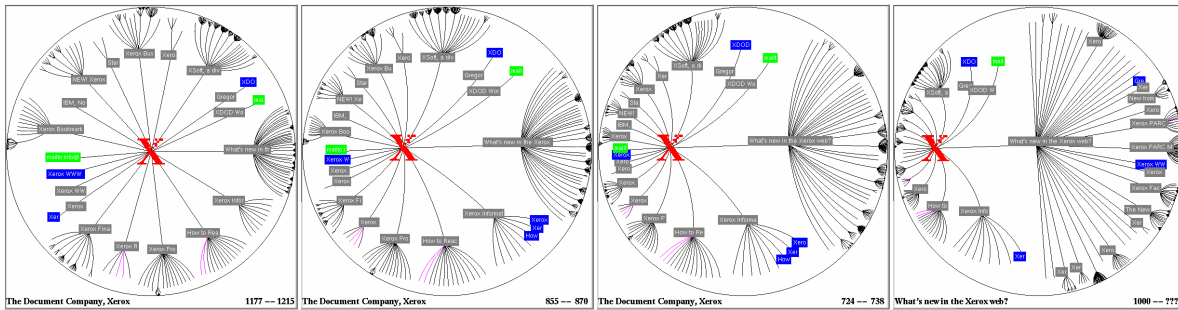


Figure 4: [Courtesy of Xerox PARC] This Hyperbolic Browser display shows the hierarchical link structure of Web pages served by the Xerox WWW Server. It uses a Focus+Context technique that gives more space to some nodes while embedding it amidst many more nodes given less space. Selecting a different node causes the system to respond with an animated transition in which a new node is brought to the focus area in the center.

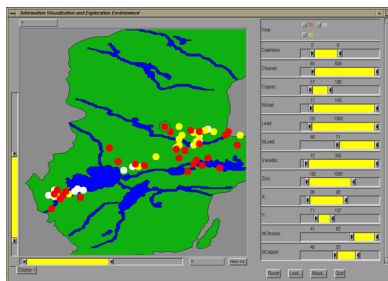


Figure 6: [Courtesy of C. Ahlberg, Chalmers University, Sweden] In this view of Sweden, using the Information Visualization and Exploration Environment, only chemical readings that exhibit high chrome reading are viewed. The points cluster near lakes and rivers where there are factories or cities. This system provides a wide range of controls for manipulating viewing parameters, data selection, and the mapping between data and graphical properties. For example, manipulating range sliders causes data points to blink in or out.




Figure 7: [Courtesy of M. Chuah and S. Roth, Carnegie Mellon University, Pittsburg, USA] This 3d data landscape is from the SDM system which provides a number of interactive operators for controlling the view, data, or mapping between. For example, selected groups of graphical symbols can be scaled as a group (the red group) or can be dragged to a different part of the space for within-group comparison or manipulation.

on the task. Not just information density, how many pixels show how many values, but activity flux, how many inferences based on how much data over time. This property depends on the design of available interactive operators for manipulating the view (percepts), the data (concepts), and the mapping between the two (wherin lies the design and learning). The challenge is to devise a few operations that can be combined in a rich variety of ways to accomplish tasks. Such combinatorics diminishes the costs of learning and remembering how to use the visualization.

Interactive controls break down the strong division between designing and using visualizations, and, to some extent, they start to destroy the concept of a visualization as a distinct object. Writing and reading is blurred in an pro-

cess of exploratory selection of appropriate representations. Yet, there is often a simplifying value to preserving the metaphor of distinct objects that provide the basic racks for information despite rearrangements offered by the controls. Animated transitions across the big form-change operations, say for example, from the Hyperbolic Browser to the Cone Tree, can help preserve the illusion of objectness. The ultimate objective is to build protean artifacts, with controls and transitions, both visual and conceptual. Mighty-morphin power rangers, not procrustean beds.

All For Eda And Eda For All

Our group has been designing visualizations for different kinds of information structures, hierarchies, events across time, document collections, and documents. A few years ago, Stu Card and I turned to the problem of generating a visualization for tabular data with two kinds of things in our head. Certainly, there were the "vague principles" described above—graphical representations, focus+context, animated transitions, combinatorics from a few good operations. Such things are never enough to guarantee that interesting and useful designs will be generated, and certainly not enough for those who would formalize the process into a structured workflow, but they are enough to start a design engagement. The other necessary ingredient is some sense of the kind of process or task that the visualization wants to support. The inspiration in this area came from an area of statistics called exploratory data analysis (EDA).

EDA, fathered by John Tukey, has grown in importance since the late 60s, and is in Tukey's words, "about looking at data to see what it seems to say." Two aspects of this field are particularly relevant. First, EDA is about exploration of data. As opposed to utilizing the data to confirm a carried-in hypothesis or belief, it is about examining the data and attempting to explain what is seen. Thus, EDA is a game of search through a space of models or representations which characterize the data with as little residue as possible. An important motto of EDA is: DATA = MODEL + RESIDUE. The exploration proceeds by performing a series of operations and when things look right you infer some model from the the operations you've performed. The second aspect of EDA is the increased use of graphical displays to guide the process. The "things look right" is based on using various graphical representations that expose properties of the data. It took Tukey, a statistician of great stature even then, to combat the prevailing negative image of graphics in statistics at the time.

In the Table Lens visualization, we have attempted to support a rough and ready style of exploratory data analysis. Instead of supporting the entire repertoire of EDA methods, we wanted a simple artifact with a few good operations, one that might be learned in five minutes or so. The basic operations of data analysis—looking for corre-

lations, spotting outliers or peculiar values, forming and comparing groups of items—can support a wide variety of activities, if the costs of learning and using them can be made insignificant. Thus, our goal was to mainstream and mainline the basics of EDA.

The Table Lens is based on a particular kind of table which has cases or objects as its rows and variables or properties as its columns. Such cases-by-variables tables are widespread in the scientific, commercial, and social arenas. Examples include kinds of cars with their physical or performance characteristics, baseball players with their playing statistics, and countries with their geographical, economic, and demographic properties. The Table Lens supports the facile exploration of such tables with hundreds of cases and tens of variables by using a discrete focus+context warping of space

Rows and columns with less space contain graphical representations of the values, while rows and columns with more space contain both graphical and textual representations of the value. So, for example, quantitative variables are represented by graphical bars with their lengths proportional to the relative size of the represented value. The use of graphical representations not only provides a scale advantage—since the bars can be scaled to one pixel wide without disturbing comparisons—but also an exploration advantage, since large numbers of tiny bars can be scanned much more quickly than a bunch of textually represented numbers.

One important operation, sorting a column, can enable many basic EDA tasks. Many properties of the batch of values in a sorted column are apparent by examining the graphical marks and the shape of the curve in the column. Sorting is also the first step of looking for correlations among variables. After one variable has been sorted, if another variable is correlated then its values will also appear to be sorted. Thus looking for correlated variable is a matter of scanning across the columns to identify other columns which seem to exhibit a descending trend.

Next Generation Workspaces

Documents and other objects that populate our workspaces have attributes based on which you find, select, filter, group, compare, relate and so on. These operations are analogous to the operations of exploratory data analysis. Thus lightweight data analysis techniques can be used to manage workspace objects. For example, you can find things by rearrangement if you understand the ways in which the space means and there are operations to rearrange meaning. Consider a collection of addresses on a palmtop computer displayed with a Table-Lens-like view. You could sort the particular columns to force a sought-for entry to a location which you can focus on to read. Alternatively, imagine dealing with collections found on the Internet or returned from a search. By manipulating external, graphical representation of the col-

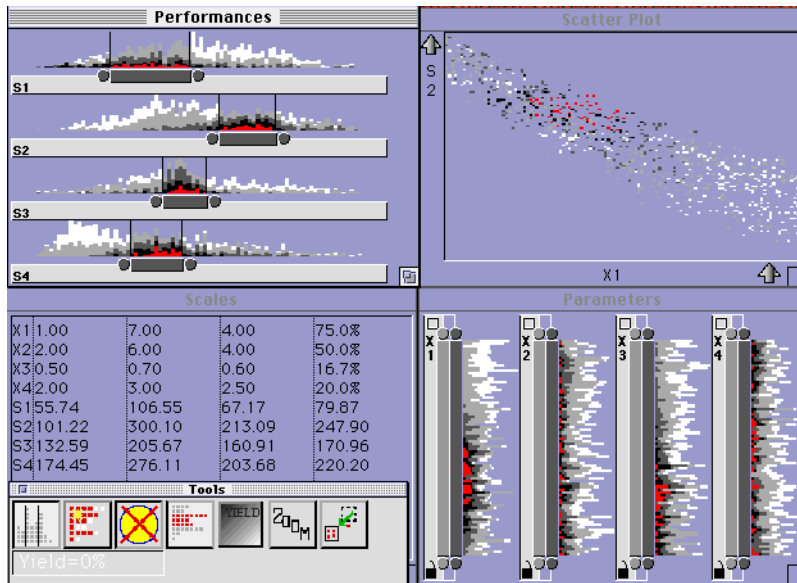


Figure 8: [Courtesy of L.Tweedie and B.Spence, Imperial College, London, England] The Influence Explorer integrates a number of representations from data analysis including histograms, scatterplots, and parallel coordinate plots. Manipulation of controls in different views cause linked actions in other views. For example, limiting the values of interests for the 4 output variables in the upper left view causes coloring of each data point in all views according to how many of the criteria it matches.

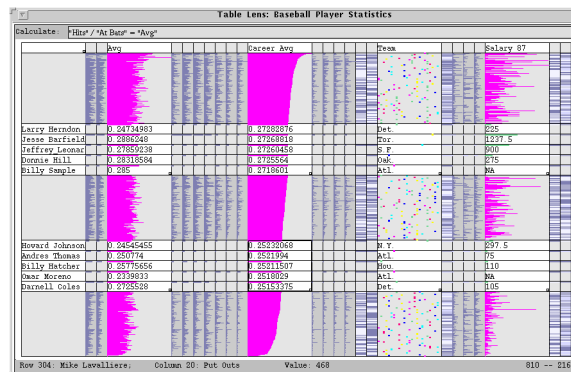


Figure 9: [Courtesy of Xerox PARC] The Table Lens integrates focus and context as well as graphical and textual representations to support looking for unusual cases as well as overall patterns and correlations. This table shows baseball statistic including 25 variables in columns and 323 baseball players in rows. Using a spreadsheet would require 13 vertical and 3 horizontal screen scrolls of a similar sized window. Sorting the "Career Avg" column near the center reveals correlations with many other columns including "Salary".

lection, the user can begin to assimilate overall properties of the collection as well as of elements within.

[Picture of document space visualization]

A major challenge, somewhat exposed in the discussion of EDA above, remains for future work. Currently, the full-blooded process of exploration isn't supported by existing visualization work. Though exploration involves rearranging the data and drawing inferences, it equally involves keeping track of what has been done, so that it can be interpreted as a model of the data. For example, perhaps one line of exploration was set aside to pursue another, but now you want to return to it. Now in rearranging the data suddenly a straight line (or some other striking regularity) "pops out" and you want to know how you got it. In such an exploratory process, one wants tools to manage the interactions in or across sessions or even maybe over a long period of time. Visualizations with their interactive operators are fine during analysis, but can they help manage the process across time? Certainly we could design visualizations that depict actions over times, but there is another important possibility.

Spaces and arrangement can be used in less structured ways than in visualizations. Structuring can be delayed so that tentative results can be dealt with flexibly. For example, think about your office and its contents. The arrangement of documents in space is mostly about interaction. Offices reflect the history and activities of their occupants. There are a number of interesting studies that have established that piles are as much about organizing work as they are about organizing content. You leave trails in space so you don't have to remember constantly what you are doing. And besides reminding you of where you are in a process and how you got there, visual elements are powerful memory cues. For example, thumbnails images of a document are highly evocative in the same way as are quick glimpses of a magazine pages as you riffle through it.

The Web Forager is a prototype workspace which uses space as a resource for keeping track of the state of an activity and for organizing objects. In particular, it focuses on the problem of collecting and organizing pages from the World-Wide-Web. Web pages can be brought into the foreground to read them or pushed into various regions of the room according to their relevance to a current activity or their similarity to other pages. As collections are gathered, they can be grouped into WebBooks and placed on a virtual shelf. This kind of computer workspace with visual depictions like thumbnails and flexible uses of space offers the promise of recapturing the flexibility and power of our physical spaces in next generation computer workspaces.

As important as visualizations may be to making sense of the increasingly larger data sets we deal with in modern life, they may play an even bigger part as elements in the next generation of computer workspaces. Our computer workspaces are becoming more and more dense with documents and we have facile access to more and more docu-

ments spread throughout our organizations and the world-wide information network. Whereas in the past the focus of computer environments was primarily on creating new documents, now it is shifting to accessing documents that already exist. In Stu Card's words, "We've been putting things into computers for a long time, it's time to start taking them out."



Figure 10: [Courtesy of Xerox PARC] The Rooms metaphor provides a set of distinct workspaces which can be used to collect applications, documents, and windows associated with a particular task or project. The overview allows finding and navigating between rooms. This Rooms overview shows a number of minaturized, but live, 3d workspaces containing Xerox PARC visualizations.




Figure 11: [Courtesy of Xerox PARC] The Web Forager utilizes space to allow users to flexibly and informally organize and collect Web pages. In early stages of a task, Web pages can be collected into piles or in regions of the room by "flinging" them to those locations. Later, related pages can be bound into Web Books and placed onto shelves as well as be electronically mailed to collaborators. The 3d physical space could possibly allow organization of millions of pages in a more effective manner than the current 2d folder and icon paradigms.

Suggested Readings

Rudolf Arnheim, *Visual Thinking*, University of California Press, Berkeley, CA, 1969. [A psychologist of art contends that all thinking is perceptual in nature, drawing on material from a broad range of perspectives including the philosophical, scientific, and historical.]

Jacque Bertin, *Semiology of Graphics*, University of Wisconsin Press, 1983. (this one was originally published in French as: *Sémiologie graphique*, 1967, Gauthier-Villars, Paris; Editions Mouton & Cie, Paris-La Haye; and Ecole Pratique des Hautes Etudes, Paris) [A cartographer outlines in detail how graphical symbols and spaces can represent meaning such that we can "see" it.]

Daniel C. Dennett, *Consciousness Explained*, Little, Brown, and Co., Boston, 1991. [A philosopher of mind offers an account of how consciousness might emerge from a material mind in which representation, mental and external, figures important.]

Peter R Keller and Mary M Keller, *Visual Cues: Practical Data Visualization*, IEEE Computer Society Press, 1993 [Experienced visualization practitioners offer over 150 examples of visualizations organized to facilitate effective application of existing techniques.]

Jill H. Larkin and Herbert A. Simon. *Why a Diagram is (Sometimes) Worth Ten Thousand Words*. *Cognitive Science* 11, pp65-99, 1987.

George G. Robertson, Stuart K. Card, and Jock D. Mackinlay, *Information Visualization Using 3D Interactive Animation*, *Communications of the ACM*, 36(4), 1993

Edward R. Tufte, *The Visual Display of Quantitative Information*, 1983, and *Envisioning Information*, 1990, Graphics Press, Cheshire CN, 1983 [A statistician collects and comments on many examples of effective and ineffective uses of graphics to represent quantitative and qualitative information.]