EXVIS: An Exploratory Visualization Environment

Georges Grinstein, Ph.D.
Ronald M. Pickett, Ph.D.
Marian G. Williams, M.S.

Graphics Research Laboratory
University of Lowell
Lowell, MA 01854

Abstract

The Exploratory Visualization (Exvis) project at the University of Lowell is devoted to the study of new technologies for the visualization of scientific data. Traditional visualization techniques are limited in the number of data dimensions that they can present. Also, they are generally done in batch mode, with the analyst providing input at the beginning of the process and receiving output at the end, but having no opportunity to direct the process itself. We are interested in approaches to scientific visualization that provide for the representation of very high-dimensional data and that support direct interaction between the data and the analyst.

We present here a description of the Exvis project and of its goals. We discuss in detail one particularly promising new technology for interactive visualization: creating iconographic displays that exploit the human observer's innate texture perception abilities. Each record of multidimensional data is represented by an icon whose visual and auditory attributes are under data control. The icons are densely packed into a 2D display, creating a visual texture. The display has an auditory texture, as well, which may be explored by sweeping the mouse cursor across the screen. Gradients and contours in the overall texture of the display reveal information about trends in the underlying data.

We describe our prototype implementation and provide several illustrative uses from diverse data domains including satellite imagery, census data, epidemiological data, and solutions of partial differential equations. Finally, we discuss our approaches to studying the formal foundation of the iconographic display approach and summarize the future directions of our work.

KEYWORDS: Displays, Exploratory Data Analysis, Visualization

I. Introduction

A recent National Science Foundation Workshop Report [1] addresses the need for improved techniques for data visualization. The report stresses the disparity between the powerful technologies that exist for collecting massive quantities of data and the inadequate methods that exist for analyzing them.

There appear to us to be three critically important issues that must be dealt with in providing improved visualization technologies. These are:

1. representing very high-dimensional data;
2. supporting interaction between data and the analyst; and
3. providing a powerful development environment for visualization researchers to use in studying ways to enhance the technology.

In this paper, we describe the recent accomplishments of our visualization project in addressing these issues. Pickett and Grinstein [2] introduced an iconographic display technique for composite displays of multispectral imagery data. The conventional approach is to let each of up to three imagery channels drive one of the color guns in a CRT to create a pseudocolor picture. The iconographic display approach replaces color pixels with visual icons whose features are determined by the data from the various images. Figure 4c (from [2]) shows an integrated display of five channels of weather satellite data. We also present here a generalization of the iconographic display technique to forms of data other than imagery.

A. Goals of the Exvis Project

The Exploratory Visualization (Exvis) project is dedicated to the investigation of innovative technologies for scientific data visualization. The project has three goals:

1. to provide an intelligent, interactive environment in which a domain expert can explore large quantities of data of high dimensionality;
2. to provide a platform for the design, running, and analysis of perception experiments in order to study the efficacy of new visualization technologies; and
3. to provide a powerful development environment for visualization researchers who are developing and extending new technologies.

Traditional visualization techniques lean almost entirely on a single perceptual metaphor. Data are mapped into positions of points in 2D or 3D spaces, and the analyst's visual perceptions of the resulting objects, surfaces or clusters are relied upon to find structure in the data. This basic metaphor, exploited since quite early in the history of science [3], has proved extremely useful in domains that generate data of relatively few dimensions. Modern graphics technology has added to its effectiveness, with such refinements as volume rendering [4] and scatterplot displays that can be "rotated," "flown-through" and "brushed" [5]. But even with these refinements, analysts in many domains find the traditional approaches inadequate for analyzing data of high dimensionality.

Our general aim is to identify and use perceptual abilities that have not yet been exploited for data visualization. Our current focus is to study how data may be turned into visual and auditory textures so as to exploit the domain of texture perception via an innovative visualization technology that uses iconographic displays to represent data of high dimensionality.

B. A Reference Model for Scientific Visualization

A reference model for scientific visualization [6] has arisen out of our work on interactive visualization environments. We summarize the salient points of the model here. We speak of an interactive *visualization pipeline* that takes in data, allows the analyst to manipulate it, and then produces visual and auditory output for exploration (see Figure 1).

While we do not define a database management system (DBMS) as part of the visualization pipeline, we assume the existence of one in the application layer. Data enter the pipeline form the DBMS. The user may then define and apply modelling transformations to the data, in the *logical layer.* The transformed data are then associated with (or in our terminology "mapped onto") the parameters of a physical representation. Since the viewing specifications are dependent on the type of visual and auditory representation being used (for example, icons), and therefore on the workstation being used, we say that the viewing specifications are defined and applied in the *workstation visualization layer.* The resulting data structures are not yet physical representations of the data; rather, they are descriptions of what the physical representation should look like. They become visual and auditory representations of the data in the *physical layer.*

The term "pipeline" implies uni-directional motion of data from its raw or normalized state, through modelling transformations and viewing specifications, to a physical representation. Our intent is entirely different. The user is free to interact with any module of the pipeline at any time in the visualization process in order to undo or edit any previous decisions. Moreover, data produced by the pipeline may be re-introduced as new data. For instance, an iconographic display could itself become input for the pipeline. Without this kind of flexibility, the analyst might as well be working in batch mode. And in fact most traditional visualization is done that way: the user gives input at the beginning of the pipeline and receives output at the end.

We classify data for visualization into two main categories: coherent and non-coherent. Coherent databases have an inherent continuity on some field or fields; they include inherently visual data like medical scans and satellite imagery. Non-coherent databases, which have no inherently continuous fields, include statistical databases such as census and epidemiological databases (unless the statistical database can be projected onto geographic coordinates).

II. A New Approach to Visualization: Iconographic Displays

The critical requirement of an effective data display is that it stimulate spontaneous perceptions of structure in the data. Consider, for example, the scatterplot display, which is immensely effective for analyzing data because it is a potent stimulator of a natural and spontaneous capacity to sense the clustering of points in space. The primary rationale of our iconographic approach is that it lets us exploit another such spontaneous perceptual capacity: the capacity to sense and discriminate texture [7,8].

To stimulate visual texture perception, many small but discriminable elements must be displayed over a relatively small area. Such an image is experienced, for example, when one looks at an expanse of wall-to-wall carpeting. With deliberate visual analysis, one can obtain specific kinds of information about the lengths and materials of individual fibers, but one also receives, without any deliberate effort, an impression of the overall texture of the carpeting. It is just such natural impressions of texture that we seek to exploit.

The icons in an iconographic display serve as the fibers in the carpet. Just as footprints or the sweep of a vacuum cleaner will create variations in the texture of the carpet, changes in the shape, size, and spacing of icons will create gradients or contours in the visual texture of the iconographic display. The challenge, of course, is to design icons that, when varied, can stimulate perceptions of highly distinctive textures.

Pickett and White [9] created such a texture display using arrays of triangles whose shape, orientation, and tilt in a stereographic display graphically conveyed a 7-dimensional database of college entrance examination scores. Chernoff [10] displayed geological data in as many as eleven dimensions using arrays of tiny faces. In both of these research efforts, the particular icons and icon features under data control were chosen arbitrarily, and no claim was made that those particular graphic codes were the most visually effective.

Research on texture perception [11,12,13,14,15] has provided some good leads on the kinds of icon geometries that will work. Variation in the orientation of line segments has been found to be a particularly potent texture coding variable. The texture of a display composed of many small line segments can be made to look clearly different from one region of the display to another simply by shifting the orientation of the lines in the two regions. The differences can be strong enough to support the perception of a fairly sharp contour at the transitions. Indeed, the stick-figure icons used in the illustrations in this paper (see Figures 4a-d) were chosen on this basis.

In addition to their visual attributes (for example, the lengths and orientations of the limbs of a stick-figure), our icons have auditory attributes. Thus, we have the facility for representing data via sound. Studies [16,17,18] suggest that sound can be successfully used in two ways to enhance the visual presentation of data: it can reinforce data that are presented simultaneously through the visual attributes of an icon and it can be used to increase the number of data dimensions that the icon represents. Our icons currently have five sound parameters. We speak of the iconographic display having an "auditory texture" in addition to its visual texture. The auditory texture can be explored by sweeping the mouse across the iconographic display. We speculate that there is a synergy between the auditory and visual textures that will become apparent through experimentation.

The dimensionality of data that a single icon can convey is further increased by the icon's color and animation attributes. In the discussions that follow, it will become clear that we have given our icons more power than we are currently able to explore. Thus, we will in general confine our discussion to their static visual properties and their auditory properties, and leave the use of animation to our section on future work.

In the sections which follow we first discuss the needs of the various Exvis user groups. We then describe our prototype workstation for meeting those needs and discuss the two approaches we are taking for giving a formal foundation to the iconographic display technology. Finally, we provide four examples of displays that represent data from diverse domains, and then discuss the future directions of our work.

III. User Groups of the Exvis Environment

The Exvis environment is designed to support two different classes of users whom we call end users and developer users. The end users are data analysts and psychophysics experimenters who need to use the system without doing any programming. End users will rely heavily on visual languages [19] for interaction with the system. The developer users are computer graphics, imaging, and artificial intelligence researchers whose goal is to enhance the technology available to the end users. They will need an extremely powerful integrated programming environment for implementating and enhancing new technologies. The interaction requirements of the various user groups summarized here are detailed elsewhere [20].

End users of the Exvis environment include scientists with data to analyze and psychophysics experimenters who are studying the perceptibility of iconographic displays. These and other end users will need to be able to construct and modify iconographic displays. They will also want to perform operations:

1. on data (such as defining and applying modelling transformations)

2. on icons (such as selecting and editing icon families, and defining and applying viewing specifications)

3. on displays (such as assigning data parameters to the x- and y-axes, and exploring the auditory texture of the display), and

4. on the system (such as saving the state of the system in order to restore it at a later time).

It seems unlikely that an end user will know a *a priori* which of a multitude of possible iconographic representations will give the most information about his/her data. See, for instance, Figure 3, which shows different iconographic representations of the same data. Above all else, the scientist user must be able to browse a variety of iconographic representations of the same data. The feasibility of this kind of visual exploration depends on two things: (1) a person's spontaneous perception of texture, so that each representation may be viewed only briefly; and (2) the capability of the hardware to compute and display images rapidly. If, for example, the twelve members of the five-limb icon family described in Section IV.C. are used, there are 12x5!=1440 possible pictures. With a user spending even as little as three seconds per picture, this exhaustive search of all possible pictures would take more than an hour. User modelling may make it possible to reduce the search space for data from a given domain [20].

We expect that the psychophysics experimenter will need to generate artificial data, specify testing protocols, run experiments, keep records of how subjects perform, and store experiments for future use.

Image management facilities are clearly of paramount importance for both data analysts and phychophysics experimenters. Iconographic displays must be saved to memory or to disk and retrieved with ease and rapidity.

Developer users include visualization researchers from computer graphics, imaging and artificial intelligence who are studying ways to improve the technology offered to the end users. We have adopted an object-oriented design for the Exvis environment to allow enhancements to be incorporated easily. For example, it is a relatively simple matter to define a new class of icons and add it to Exvis. The icon class must simply be provided with a description of its own topology and with methods for displaying or representing itself, and for interacting with other objects.

IV. Implementation

The current Exvis prototype is implemented on a Symbolics 3600 computer in Flavors, an object-oriented language written over Lisp [21]. It is being re-implemented on a VAX in C++ and X Windows, incorporating what was learned from the experience of rapid prototyping on the Symbolics [22].

The heart of the Exvis environment is called the *kernel.* The design of the kernel is object-oriented. It consists of the objects that contain the user's data, that define and apply modelling transformations to the data, that define and apply viewing specifications, that describe the visual and auditory representation of the data and that define the input sensitivity and interaction rules with the data.

A. Data Transformations

The first major task of the data analyst is to define and apply modelling transformations to the data. Modelling transformations include projections, extensions, and clipping. For imagery data, they also include general image-processing operations such as scaling, filtering, and edge detection. Statistical parameters of the data set are computed when a data set is loaded, and those parameters are available for use in the modelling transformations. These transformations take place in the logical layer of the pipeline.

Exvis extends conventional visualization by allowing the data to be transformed and fed back into earlier stages of the pipeline. Figure 4a could, for example, be used as input back at the beginning of the pipeline; some pattern recognition module could then be invoked to produce a new data display.

B. Data Mappings onto Icon Attributes

The second major task of the data analyst is to associate the transformed data with the visual and auditory attributes of icons. These associations are made via viewing specifications. Each association maps one piece of data onto one attribute of an icon. For example, in a demographic database, a person's age could determine the color of the icon, while the person's income could determine the pitch of the icon's sound. Since they are dependent on the visualization techniques that the workstation supports, the viewing specifications take place in the workstation visualization layer of the pipeline.

In order to place icons on the screen, two data fields are mapped to the x- and y-axes. In the case of two-dimensional coherent databases, the default choice of the data fields is obvious. In other cases, the user must select mappings. Exvis permits the visualization of coherent databases as if they were non-coherent. For example, a satellite image can be considered as a non-coherent database where the x- and y- coordinates are just two other data fields.

C. The Stick-Figure Icon Family

Exvis allows the user to map each data record onto the properties of an icon, but places no restriction on the nature of the icon. In the examples in Section VI, we illustrate the use of stick-figure icons. A stick-figure icon is a collection of line segments or "limbs". It is characterized both by the number of limbs it has and by the topological arrangement of those limbs. Figure 2 shows a typical icon family. Each member of the family has a body and four limbs. With the restriction that limbs are attached only to the ends of the body or to the ends of other limbs and that only two attachments can be made at any point, there are twelve possible different configurations.

The numerical data to be analyzed are mapped onto an icon by associating one data value with each icon parameter. Each limb has length, width, angle, and color. Thus, a five-limb icon can represent 5x4=20 data dimensions via the visual attributes of its limbs. Add to this the icon's five sound parameters, the two data dimensions represented by the x- and y-axes, the possibility of using animation, and the fact that there is no arbitrary limit imposed on the number of limbs that a stick-figure icon can have, and it becomes clear that Exvis icons can represent a large number of data dimensions.

V. Toward a Formal Foundation

The iconographic display approach to scientific visualization is intuitively appealing both to us and to the data analysts with whom we work. However, there is a clear need to study the formal foundations of the technology. We are approaching this task from two directions: (1) working with data domain experts to see what they can learn from iconographic representations of their data and (2) performing perception experiments to learn how known structures in data reveal themselves in different iconographic representations.

A. Collaboration with Data Domain Experts

We are working with experts from domains as diverse as radiology and mathematics to see whether iconographic displays can help such analysts visualize their data. Consider the domain of radiology. Magnetic resonance imaging produces, for each tomographic crosssection of the body organ under study, several gray scale pictures, each conveying different information about the structure and physiology of the tissue. Much of the diagnostic information comes from comparing the pictures, noting similarities and differences in the shapes and shadings of corresponding areas. These comparisons can only be done in very crude fashion by back and forth visual inspection of the separate gray scale pictures. Radiologists are concerned about how they might create an integrated display [23]. We are working to devise an integrated iconographic display like the one we show for the satellite imagery in Figure 4c. This new display may not only facilitate the kinds of crude comparisons the radiologists now make, but it may enable them to detect much more subtle relationships among the pictures of potential diagnostic interest.

B. Perception Experiments

We are conducting perception experiments using artificial data to examine how a subject's performance in discerning gradients and areas of interest in an incongraphic display is affected by such variables as: the type and strength of statistical structure in the data; the size of the data sample; the type of icon used; the particular mappings of data fields to icon attributes; the use of auditory texture in additional to visual texture; and the possible role of perceptual learning. Prelimary studies [24] do verify that some statistical properties are visible in iconographic displays, that the choice of icon type and specific mapping make a difference, and that perceptual learning is an important consideration.

VI. Application Illustrations

We provide here four illustrations[1] to show the generality of the iconographic technique. Our primary interest here is to illustrate the different possible representations of data and the diversity of the data that may be input for iconic display. The specific mappings of data fields to icon attributes are shown with each figure.

A. Census Data on Scientists, Engineers and Technicians

The database displayed in Figure 4a consists of a portion of the Public Use Microsample-A (PUMS-A) of the 1980 United States Census. PUMS-A is a 5 percent representative sample of the entire census. This portion contains information on all individuals classified as scientists, engineers or technicians. It contains thirty fields of data selected for their potential relevance to geographic mobility. Those used in the illustration are: occupation, age, education level, marital status and sex. Each icon in the figure represents a scientist, engineer, or technician and is positioned on the screen on the basis of the person's income (x-axis) and age (y-axis). Icon family member 10 (see Figure 3) has been used.

The clear shift in texture over the screen indicates that one or more of the data fields controlling the shapes of the icons is varying as some function of income and age. The statistical structure underlying that shift might be a very simple one, easily discovered by conventional numerical analysis, provided, of course, that one knows in advance how to partition the sample, or it might be rather complex, and probably tough to find by conventional numerical analysis.

B. AIDS Database

The database displayed in Figure 4b comes from the Center for Disease Control (CDC) and contains data of epidemiological interest on AIDS cases on record at the CDC. Our current version [1987] contains 38,000 cases. The data on each case include age, race, sex, sexual orientation, drug use, geographic location and presenting opportunistic disease. Each icon in the figure represents a single AIDS patient. We have partitioned the sample of patients by race and age. Within quadrants the patients are assigned to locations randomly. Icon family member 10 has been used. Only three data fields are displayed, and the outermost segment that would have represented a fourth data field has been suppressed.

It would be of great interest to epidemiologists to find differences in the epidemiology of AIDS among the different races. We conducted numerous visual explorations and found no differences with the limited set of data fields available in the CDC database. Here we show a display in which we have built an artifactual difference, as a way of illustrating how a

[1]We are indebted to the following persons or institutions for providing the databases for these illustrations: Dr. Ralph Gentile, College of Management Science, University of Lowell—Census Database; The Center for Disease Control, Atlanta—AIDS Database; Dr. Robert P. d'Entremont, Air Force Geophysics Laboratory—Satellite Imagery Database; and Dr. John McKelliget, College of Engineering, University of Lowell—Partial Differential Equation Data.

difference might have looked were it there. The apparent difference between the top and bottom half of this display is purely forced, because age, one of the data dimensions controlling assignment to the top versus bottom half of the display, also controls the angle of limb 1.

C. Imagery Data from Weather Satellite—NOAA-7 AVHRR

The imagery database displayed in Figure 4c was acquired form the Air Force Geophysics Laboratory. It contains geosynchronous satellite data of North America, with approximately 1 nautical mile resolution in five bands. The figure shows how a single image can be created to represent information that would ordinarily have to be viewed in five separate gray scale images. Such an integrated display has previously been produced for up to three spectral lines, using a color composite technique, but ours is the first such integrated display that gets beyond the three channel limit of the color composite. Each icon represents all the pixel values at the same given location in all five images. Icon family member 12 has been used.

This image shows a geographic region covering all of the western end of Lake Ontario and part of the eastern tip of Lake Erie. A rather strikingly well-structured picture emerges in the total absence of gray scale coding. This picture nicely demonstrates how strongly one can differentiate regions in an array of icons on the basis of texture differences. In actual applications, one would typically be looking for regions one did not know were there, regions defined perhaps by some interaction among the channels.

D. Solutions of Partial Differential Equations

The database displayed in Figure 4d was acquired through the mechanical engineering department of the University of Lowell. It represents the flow field, temperature field, electromagnetic field and reaction kinetics inside an inductively coupled plasma chemical reactor. The field variables consist of the Z and R components of the velocity, the stream function, pressure field, turbulence kinetic energy, turbulent energy dissipation, plasma enthalpy, laminar and turbulent viscosity, peak current density, thermal conductivity, joule heating per unit volume, radiation loss per unit volume, electrical conductivity, volumetric production rate of silicon, and the concentration of reactant. Icon family member 12 has been used. The x-axis represents the Z component of the velocity and the y-axis the R component. Five of the field variables control orientation of the icon segments.

The turbulence produced by the flow field is clear in this display. Both Figure 4b and Figure 4c clearly show the natural advantage resulting from the spatial coherence of the original input data.

VII. Summary and Future Work

We have illustrated a potentially powerful new display technique for data visualization. A picture of a database can be created as an aggregation of icons, each representing one of the data records. Texture perception can be harnessed as a way of detecting structure in the data.

We have also described a general purpose system that can facilitate design and experimentation with such displays and that can support both the scientist user and the visualization researcher. We envision several directions for future work:

- Extend the visualization reference model.
- Install Exvis in diverse sites and obtain feedback to understand user interface issues and build evidence about the practical use of the iconographic technique. Super-mini graphics workstations now available can provide great support for browsing through data.
- Codify the mapping process into rules in order to (a) simplify the design of an interactive program to define the mapping and (b) support later experimentation with expert system assistance in generating mappings.
- Design other types and families of icons. The stick-figure icon of the present illustrations is but one of many possible types. Exvis provides for the development of user-defined icons, but base types and families are necessary for the scientist user.
- Incorporate icons on a 3D surface with hidden line and hidden surface removal. This will remove clutter and provide much more perceptible displays.
- Investigate the visualization of large databases. The issues involved are not just performance or hardware ones. Currently we can display up to 10,000 data records. How can we extend this technique to handle large databases where the number of records is orders of magnitude larger? What computational techniques can we develop to support panning?
- Explore dynamic icon encoding. We have incorporated this in our system, and although we do not report details here, it looks very promising. Structure in the data may become more salient as the data driving the icons causes them to move and perhaps "behave" under data control.
- Investigate the use of icons with auditory attributes. Exvis can accept arbitrary icons and mappings, including icons with auditory attributes supplementing the geometric ones. Can we provide tools to split data representation across the visual and aural bands so as to increase the user's perception of data? Some work has been done in that area.
- Investigate projection pursuit algorithms and user modeling to reduce the search space of possible iconographic representations of a given data set.

REFERENCES

[1] B.H. McCormick, T.A. DeFanti, and M.D. Brown (Eds). "Visualization in Scientific Computing—A Synopsis," *Computer Graphics*, **21** (6) November 1987.

[2] R.M. Pickett and G.G. Grinstein, "Iconographic Displays for Visualizing Multidimensional Data," Proceedings of the 1988 IEEE Conference on Systems, Man, and Cybernetics, Beijing and Shenyang, People's Republic of China.

[3] E.R. Tufte, "The Visual Display of Quantitative Information," Graphics Press, Cheshire, Connecticut, 1983.

[4] R. A. Drebin, L. Carpenter, and P. Hanrahan, "Volume Rendering," *Computer Graphics*, **22**(4), 65-74, July, 1988.

[5] R.A. Becker and W.S. Cleveland, "Brushing Scatterplots," *Technometrics*, **29**, 127-142, 1987.

[6] R.D. Bergeron and G.G. Grinstein, "A Reference Model for the Visualization of Multi-dimensional Data," accepted for presentation at Eurographics 89, Hamburg, Germany, September, 1989.

[7] J.J. Gibson, "The Perception of the Visual World," Boston, MA, Houghton-Mifflin Company, 1950.

[8] R.M. Pickett, "Visual Analyses of Texture in the Detection and Recognition of Objects," in B.S. Lipkin and A. Rosenfeld (Eds.), Picture Processing and Psycho-Pictorics, New York, Academic Press, 289-308, 1970.

[9] R.M. Pickett and B.W. White, "Constructing Data Pictures," Proceedings of the Society for Information Display—Seventh National Symposium, 1966, 75-81. *Information Display*, **3**(5), 57, September/October, 1966.

[10] H. Chernoff, "The Use of Faces to Represent Points in k-Dimensional Space Graphically," *Journal of American Statistical Association*, **68**, 361-368, 1973.

[11] R.M. Pickett, "The Perception of a Visual Texture," *Journal of Experimental Psychology*, **68**, 13-20, 1964.

[12] R.M. Pickett, "Response Latency in a Pattern Perception Situation," *Acta Psychologica*, **27**, 160-169, 1967.

[13] J. Beck, "Relations Between Similarity Grouping and Peripheral Discriminability," *Journal of Experimental Psychology*, **102**, 1145-1147, 1974.

[14] B. Julesz, "A Theory of Preattentive Texture Discrimination Based on First-Order Statistics of Textons," *Biological Cybernetics*, **41**, 131-138, 1981.

[15] A. Triesman, "Preattentive Processing in Vision," *Computer Vision*, Graphics and Image Processing, **31**, 156-177, 1985.

[16] E.S. Yeung, "Pattern Recognition by Audio Representation of Multivariate Analytical Data," *Analytical Chemistry*, **52**, 1120-1123, 1980.

[17] S. Bly, "Sound and Computer Information Presentation," Unpublished doctoral Dissertation, University of California, Davis, 1982.

[18] J.J. Mezrich, S. Frysinger, and R. Slivjanovski, "Dynamic Representation of Multivariate Time Series Data," *Journal of the American Statistical Association*, **79**(385), 34-40, 1984.

[19] S. Chang, "Visual Languages: A Tutorial and Survey," *IEEE Software*, **4**(1), 29-39, January, 1987.

[20] M.G. Williams, S. Smith, and G. Pecelli, "An Intelligent Assistant for Interactive Data Visualization (Phase I)," Technical Report, Computer Science Department, University of Lowell, Lowell, MA, 1989.

[21] M.G. Williams, R.D. Bergeron, G.G. Grinstein, G. Pecelli, R.M. Pickett, M.J. Schedlbauer, S. Smith, and S.M. Streeter. "The Architecture of the Exvis Kernel," Technical Report, Computer Science Department, University of Lowell, Lowell, MA, 1989.

[22] M. J. Schedlbauer, M. G. Williams, and G. G. Grinstein, "Issues in the Implementation of an Experimental Data Visualization System," Third Annual X Technical Conference, Massachusetts Institute of Technology, Cambridge, MA, January, 1989.

[23] M.W. Vannier, R.L. Butterfield, D. Jordan, W.A. Murphy, R.E. Levitt and M. Gado, Multispectral Analysis of Magnetic Resonance Images, *Radiology*, **154**, 221-224, 1985.

[24] M.G. Williams, S. Smith, and G. Pecelli, "CHI Issues in the Design of an Intelligent Visualization Workstation," ACM SIGCHI '89, Austin, TX, April 30-May 5, 1989.

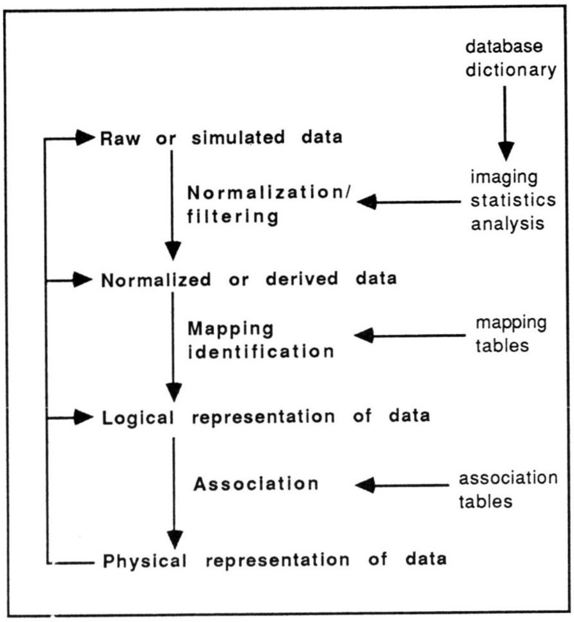

Figure 1. A schematic representation of the visualization pipeline.

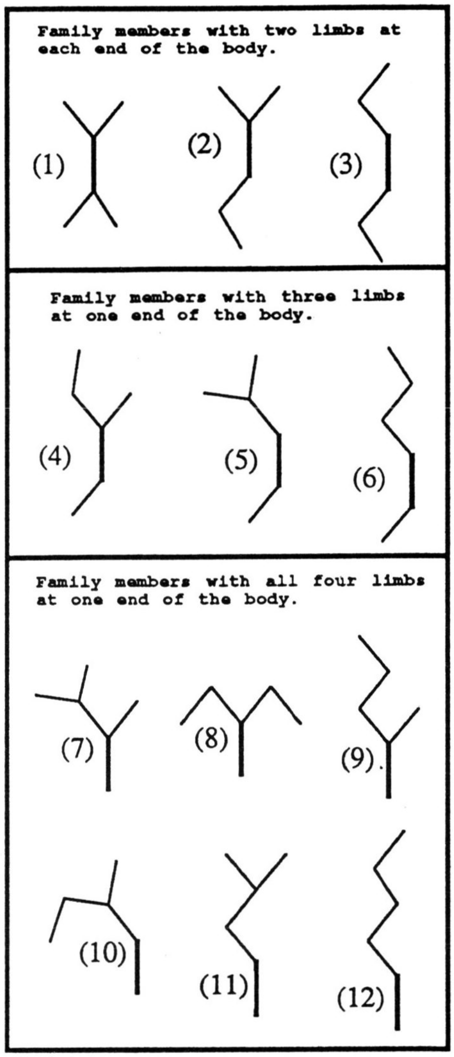

Figure 2. The 12 members of a stick-figure icon family. Each member has a body (bold) and four limbs. A limb must be attached to the body and/or to another limb.

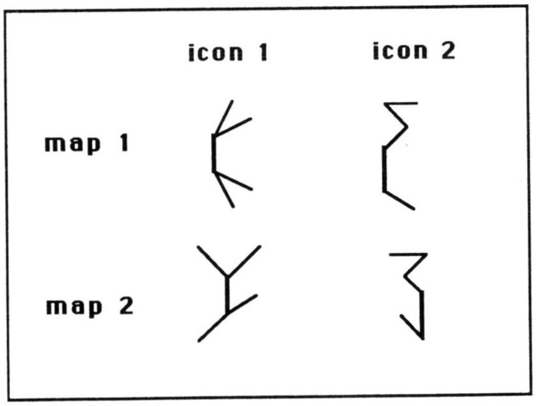

Figure 3. The effect of mapping the same input data onto different icons.

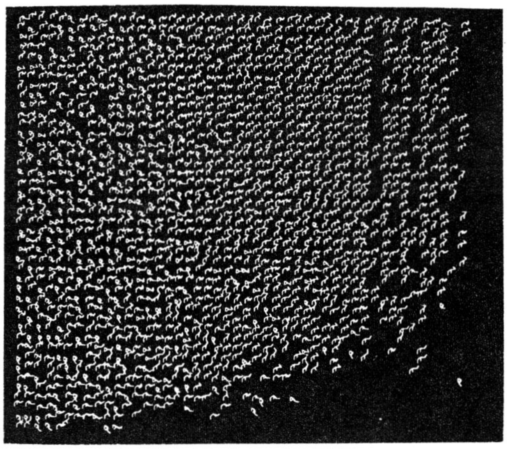

Figure 4a. An iconographic display of the census database (see Section VI). The icons are positioned on the x-axis by income and on the y-axis by age. Icon 10 has been used, with the following mapping:

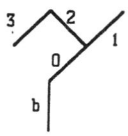

Icon feature	data field
b	fixed
0	sex
1	education
2	occupation
3	marital status

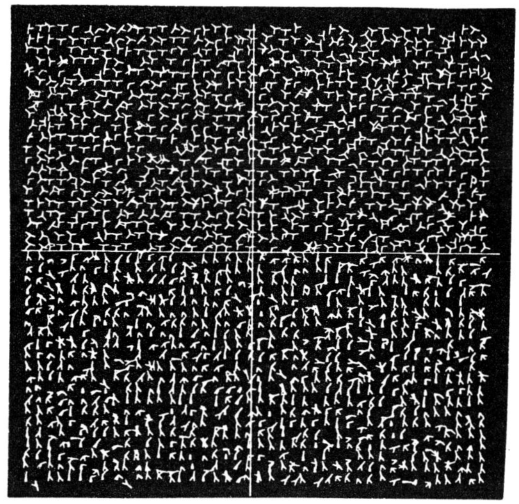

Figure 4b. An iconographic display of the AIDS epidemiological database (see Section IV). The icons are partioned by age (top row-old, bottom row-young) and by race (left column-Black, right column-White) Icon 10 has been used, with the following mapping:

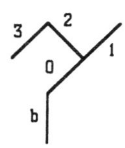

Icon feature	data field
b	suppressed
0	symptom
1	age
2	dx date
3	suppressed

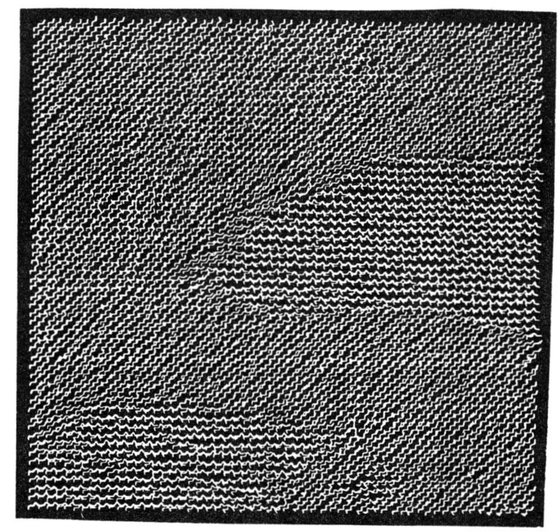

Figure 4c. An iconographic display of the satellite imagery data (see Section VI). An area centered roughly on the western tip of Lake Ontario is depicted. Icon 12 has been used, with the following mappings:

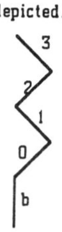

Icon feature	channel (μ m)
b	11.5-12.5
0	0.725-1.1
1	3.55-3.93
2	0.58-0.68
3	10.3-11.3

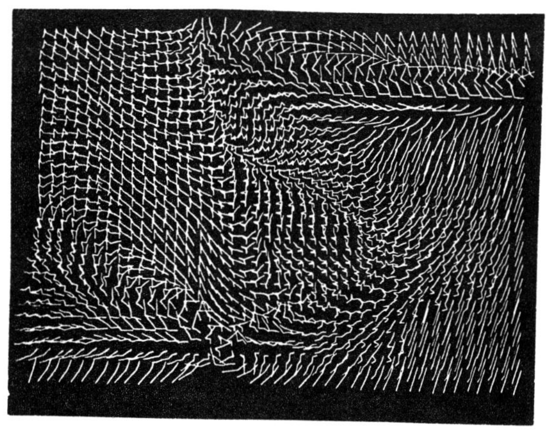

Figure 4d. An iconographic display of the PDE data (see Section VI). A components of the field variable are plotted on the x-axis and 2 components on the y-axis. Icon 12 has been used with the following mappings:

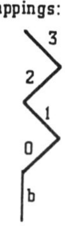

Icon feature	data field
b	6
0	8
1	10
2	12
3	14