

Methods for Visual Mining of Data in Virtual Reality

Henrik R. Nagel, Erik Granum, and Peter Musaeus

Lab. of Computer Vision and Media Technology, Aalborg University, Denmark
{hrn, eg, petermus}@cvmt.dk

Abstract. Recent advances in technology have made it possible to use 3-D Virtual Reality for Visual Data Mining. This paper presents a modular system architecture with a series of tools for explorative analysis of large data sets in Virtual Reality. A 3-D Scatter Plot tool is extended to become an "Object Property Space", where data records are visualized as objects with as many statistical variables as possible represented as object properties like shape, color, etc. A working hypothesis is that the free and real-time navigation of the observer in the immersive virtual space will support the chances of finding interesting data structures and relationships. The system is now ready to be used for experiments to validate the hypothesis.

Keywords: data exploration, visualization, perception

1 Introduction

Visual Data Mining traditionally uses 2-D graphics or very simple 3-D graphics to visualize results on ordinary monitors. Real-time interaction is only used to a limited extent. One of the reasons for this has been the lack of adequate hardware for visualizing complex graphics. However, during the last years graphics cards have doubled their speed every half year. Together with advances in supercomputers, user interface technology, etc., this has made it possible to view large and complex visualizations in immersive 3-D Virtual Reality (VR), with real-time interaction. Thus today it is possible to design new methods for visualizing the often very large and complex databases. These methods enable analysts to perceive data also from the inside out, thereby, hopefully, adding extra opportunities to recognizing patterns, clusters, etc.

Some statisticians have experimented with extending the methods traditionally used in statistical data exploration, to work in VR. An example of this is XGobi/VRGobi [12] with a VR version of "The Grand Tour" [2, 4, 18]. These methods were originally designed for ordinary workstations with standard 2-D monitors and are often used to analyze relative small data sets. It may therefore be possible to improve upon these methods using the enhanced visualization, processing, and interaction facilities available for VR today. Except for systems like TIDE (Tele-Immersive Data Explorer) [15] and, DIVE-ON (Data mining in an Immersed Virtual Environment Over a Network) [1], little scientific research

has been done on how best to use VR in visual data mining. TIDE focuses on the use of collaboration and resource sharing to mine large data sets in VR, while DIVE-ON focuses on interaction with a Virtual Data Warehouse over a network.

In the autumn of 1999 a new Virtual Reality center now called "VR Media Lab" was inaugurated at Aalborg University in Denmark. Among its facilities are a 3-D Power Wall, a 160 degree Panorama, a 6-sided cubic CAVE, and a 16 processor SGI Onyx2 with 6 graphics pipes. A research project called "3-D Visual Data Mining" (3DVDM) was also initiated to study how VR may be used in Visual Data Mining. The project group consists of persons with expertise within the scientific fields: databases, statistics, perceptual psychology, and visualization.

This paper presents work of the visualization group of the 3DVDM project. It contains a discussion of system architecture, data exploration tools, and experience with using natural human perceptual skills for mining data in VR.

1.1 Motivation and Tasks

With the new high-end VR technology, it is possible to let users perceive visual worlds from inside 3-D immersive environments. Users are able to recognize whether objects appear large and far away or small and close by. The scenes look realistic, since close objects stand out in space in front of the users. It gives possibilities for exploiting natural human perceptual skills in finding unspecified patterns in a visual 3-D representation of data.

Tracking of a user's position, orientation, and pointing gestures, allows the computer to calculate the position and orientation for the right visualization. This means that the user in real-time can move around objects or clusters of objects close to him, and see them from all sides. Virtual reality thus allows a whole range of new possibilities for expressing and inspecting information.

Visual Data Mining projects aim at allowing analysts to explore large and complex data sets. To investigate the possibilities available in VR, our tasks are:

1. Developing a modular system architecture suitable for empirical studies in a research context.
2. Investigating state-of-the-art software methodologies for optimally exploiting VR technology, allowing as many degrees of freedom as possible.
3. Developing a visual language suitable for expressing information in VR. This requires investigations on how to artificially generate perceptual sensations in VR, which optimally exploit the excellent faculties of human perception, for analysis of information content and structure of data.
4. Developing new VR data exploration methods that make use of our findings.

2 The Visualization System

3DVDM is also the name of our software system for exploring large databases in VR. The aim of the software system is to accommodate new methods for data

exploration, that makes full use of supercomputing, VR visualization, real-time interaction, human perceptual skills, etc.

In figure 1 is shown the general approach adopted in this research project for visualizing representations of data from databases.




Fig. 1. Data Flow and Interaction Patterns

The system contains different data processing modules in a pipeline with the possibility of feedback from the user to each module.

First database technology extracts a relevant subset of the data in a database, and produces an easily accessible internal database, which is passed on for statistical processing. Data is then transformed into an equivalent symbolic graphical representation. This data format is independent of specific hardware and software requirements. Last step is to transform this data to polygons, which are rendered in a 3-D space.

The following sections contain a description of the project in each of the 3 areas: system architecture data exploration tools, and visual perception

2.1 Design Goals

It is expected that several different transformations of data will be invented during the project's lifetime. The system architecture has been designed in a highly modular way to allow easy addition and substitution of functionality by specifying a set of formal interfaces between the modules. Existing modules are therefore completely compatible with modules designed months or years later, provided that all modules use the same interface. New interfaces can easily be added to the system, if needed.

The primary goals for the 3DVDM system, were:

1. Highly modular programming
2. Automatic handling of data flow between modules
3. Automatic handling of process flow between modules
4. Structural flexibility to implement most kinds of modules.
5. Interface rigidity to ensure compatibility between modules

2.2 System Architecture

The approach taken was to design an object-oriented framework, with graphs and nodes as the basic structuring mechanism. To create an application a number

of nodes are connected in a graph, so that the data generated as output of one node is sent to the input of another node. Each node acts as a tool that contains the implementation of a method devised by participants in the project.

The nodes can be divided into source nodes, mapper nodes, and sink nodes. A source node receives no input, but produces data that is subsequently processed by other nodes. Data is represented internally as objects. Mapper nodes receive an input data object, and produce an output data object. This could be a node that filters data according to user defined criteria. A sink node is a node that receives input data object, but does not produce any output data object. These nodes produce the result of the program by other means, such as visualizing graphics in a VR arena.

Graphs transport data between nodes and execute the nodes in the correct order. They could be considered the skeleton of the program. In the original design of the system architecture, 3 kinds of graphs were included:

- **Sequential graphs** execute their nodes in sequential order, starting with the source nodes in the graph, and ending with the sink nodes.
- **Thread graphs** execute their nodes in parallel using e.g. POSIX threads. Data objects are passed from one node to another using shared memory. This kind of graph is particularly useful for allowing parallel execution on a supercomputer.
- **Process graphs** execute their nodes in parallel processes, using popular message passing protocols such as Message Passing Interface (MPI), or Parallel Virtual Machine (PVM). These kinds of graphs are particularly useful when executing programs on clusters of PC's or workstations.

In the current version of the software only sequential graphs are implemented.

2.3 Data Preparation

One of the main differences between ordinary statistical data exploration and data mining is the amount of data being analyzed. While it is normal in the former to work with databases containing only a few 100 records, the latter concerns much larger databases. It is thus necessary either to use a data exploration tool, which can handle large amounts of data, or perform analysis on a sample of the data. The system supports both methods by defining a parameter called "Coarseness". If its value is n , where $n \geq 1$, then $1/n$ part of the data is extracted. The extracted data is stored in an easily accessible internal database. Random sampling is not yet supported.

The internal database may be passed through a filter to extract a subset of the data using user-defined rules such as $Age \geq 20$ and $Age \leq 80$.

After this, selected statistical variables are mapped into visual properties like: position, color, shape, size, and pose. This allows subsequent data exploration tools to perform a visual analysis of the data based on analyst's choice of variables.

2.4 VR Visualization

Visualization is done by creating lists of object-vectors. To each visual object corresponds an object-vector. Each object vector contains only the minimum necessary information about a visual object. Object-vectors have the advantage of being easy to process, and independent of special software and hardware requirements.

For the rendering part, SGI's OpenGL Performer has been chosen as the basic 3-D graphics toolkit in the 3DVDM system. OpenGL Performer exists on SGI's Irix based computers, and on Linux. It provides real-time 3-D graphics, with automatic multi-processing in the low-level parts of the rendering system.

Performer stores the data that define virtual worlds in scene graphs. A scene graph includes low-level descriptions of object geometry and their appearance, as well as higher-level spatial information, such as positions and object transformations. The object vectors are transformed to polygon data, and stored in a Performer scene graph.

This scene graph can be rendered by Performer itself, in which case one can view the visualizations on ordinary monitors. However, one can also choose the combination of OpenGL Performer and VRCO's CAVELib VR toolkit, in which case the visualizations are done in 3-D virtual reality arenas. The latter gives analysts the benefit of being immersed in data, and thus being able to study local phenomena with a higher degree of detail from any viewpoint and in the context of the full data set.

3 Perception of Visual Cues

The task in Visual Data Mining is to extract and analyze as much interesting information as possible. This entails encoding and processing of visual stimuli by the human perceptual system.

What guidelines can inform our construction of data exploration tools in order to facilitate visual data exploration? One guideline is that VR-displays for visual data exploration are constructed with perceptual cues, which pop-up pre-attentively in order that the encoding happens quickly and reliably [16, 17]. Even though visual data mining is mainly about exploring data, one prerequisite for such exploration can in some cases be that data is read accurately from the VR-displays. Here we can be informed by guidelines concerning perception of 2-D graphs or traditional displays, such as dashboards [6, 11, 3].

As mentioned in section 2.3 the data preparation in the 3DVDM-system is performed in such a way that a single geometric shape represents an observation in a data set. The parameters in the static object property space are: Position, pose, size, shape, color and texture. In the following we will briefly present some thoughts on the potential use of these perceptual parameters in terms of mapping statistical variables.

Position Position is a fundamental parameter, which determines the relative position of objects. Position can be a strong pop-up cue (e.g. close objects tend

to be grouped together according to the gestalt law of proximity). Furthermore, stereoscopic depth - the distance to the object from an observer - is a pop-up cue. When the task is merely to read off data in displays such as in a 2-D histogram, position has been found to be the most efficient way to map data. With regards to a 3-D Scatter Plot, the perceptual system can discriminate fine changes in position and position can be used to map three continuous statistical variables.

Pose The pose - or spatial orientation - of an object is often perceived to be upright in relations to the observer's interpretation of vertical and horizontal in the visual space. Vertical or horizontal visual stimuli are perceived more efficiently than tilted visual stimuli [9]. Co-linearity of line structures is a pop-up phenomenon. Pose can theoretically be used to map up to two continuous variables.

The use of pose requires coordination with the use of the shape property, as orientation characteristics should be maintained for all shape-variations used.

Size The size of an object is potentially important to consider since a large object stands out from a population of smaller ones, and groupings of data points (in 2-D) which occupy less area is perceived as figure whereas regions with bigger area is perceived as ground. In traditional displays, objects should have no more than three different sizes in order to be efficiently encoded. In VR colored objects should not be too small, e.g. if the color-difference is in the yellow-blue direction the smallest size should be larger than half a grade of the visual angle.

In the 3DVDM system we have so far found it useful to have "frozen" object size to be constant by some parameter (e.g. volume). This reserves size for use by the observer for depth perception. Ambiguities regarding statistical information and distance to object should thus be avoided.

Shape Symmetric shapes are often thought to be encoded and processed more efficiently than non-symmetric shapes [7]. The contour of the shape can have pop-up qualities since the length - and width - of a line and curvature are pop-up phenomena. Furthermore, shapes with added marks work as pop-up stimuli. For instance a dot added to a square in a population of squares will make that particular square pop-up. The perceptual system clearly differentiates between topologically different objects, such as a ring with a hole and a ring without a hole, but not between topologically equivalent, such as triangle and square, square and circle, or triangle and circle [5]. Up to 15 different shapes can be distinguished in traditional displays, but no more than five different shapes should be used. This guideline seems to apply for all the 3DVDM data exploration tools, but the project only use 3 shapes so far. In order for the human perceptual system to notice a difference, distortions in the length of a shape, as measured horizontally or vertically, are not perceived if the distortion is less than 1.4% of the original length. Given the limits in number of shapes that can be efficiently encoded, shape can be used to map one categorical variable.

The data exploration tools developed so far use object shape as one or more visual properties to be varied parametrically or via fixed categories (like cube, tetrahedron etc.). Similar shapes tend to be grouped together (gestalt law of similarity).

Color The color (both hue and saturation) of objects can act as particularly strong pop-up feature. E.g. in a cluttered visual space with heterogeneous objects, the observer detects specific objects faster knowing in advance the color of the object as opposed to its size or shape. The pop-up effect is enhanced when objects are colored with a black rim on a white background or a white rim on a black background. Perceiving visual stimuli in VR, more than 6 colors are easily confused. Generally it is advised against using color as a continuous variable for two reasons: First due to limits in the human perceptual system in distinguishing accurately between hue, saturation, or brightness. Second due to the fact it does not unequivocally make sense to say that one hue is more or less than another (e.g. green is not more or less than red). In the 3DVDM system color can in fact be used to map continuous variables.

Texture Texture can e.g. be defined according to granularity, orientation and pattern [19]. The texture of an object aids the observer in determining the object's pose and shape. Texture can theoretically be used to map one or more continuous variable.

Dynamic Object Properties In the current system blinking is also an option as a dynamic object property. Blinking is particularly appropriate for drawing attention to alert signals, but has been shown to tire the observer [7] and must, therefore, be used with caution.

Spatial Distance Metric We are currently developing design rules for constructing VR-visualizations for visual exploration by adopting experience from perceptual psychology [10]. First we establish a spatial distance metric on the basis of maximal object size as the basic spatial unit. This conveniently allows a scaling of the visualization to refer to perceptually relevant measures. We also evaluate all object properties in terms of distance range within which variations of the visualization of the properties can be distinguished for individual objects.

We suggest initially the following upper bounds for the range of three variables:

Texture 25 spatial distance units
Shape 50 spatial distance units
Color 100 spatial distance units

The potential importance of these different upper bounds is that the perceptual grouping and the perceptual pop-up phenomenon on the basis of these properties are correspondingly bounded spatially, relative to the observer. Hence

a statistical variable encoded as texture properties will only work informatively in a relatively close neighbourhood as defined by the three variables currently used as positional variables. If the figures above hold, color may offer a potential for perceptual structuring in a neighbourhood up to 16 times larger. A more thorough investigation into these relationships is required, as they reveal some important relationships for the mapping between statistical variables and object properties.

4 Results

We have designed and implemented a software system that can be used for conducting experimental research on new methods for Visual Data Mining in VR using, e.g. the object property space.

4.1 Data Exploration Tools

We have until now designed and implemented the following data exploration tools:

3-D Histogram The 3-D histogram tool divides the space into cubes in a coordinate system with 3 axes. 4 variables from the data set are used for the visualization. 3 variables are mapped to position. The average value of a 4th variable over the records that fall inside a cube, is mapped to the color attribute of the cube as shown in figure 2.




Fig. 2. Visualizing 4 statistical variables.

The size of each cube is here used as an additional dimension for showing the number of counts in them as illustrated in figure 3.




Fig. 3. Count is mapped to the size of each cube.

This tool is substantially more useful than an ordinary 2-D histogram, since it utilizes 3-D space more effectively. One can take a closer look at a subset of the data by "flying" into the middle of the data to study a phenomena there. These kinds of visualizations have also been explored in, e.g. DIVE-ON [1].

3-D Scatter Plot The 3-D Scatter Plot tool maps each data record as a data point in a 3-D coordinate system of continuous variables, see figure 4.




Fig. 4. Scatter Plot

A data point is illustrated as an object with the minimal number of surface polygons, a tetrahedron with 4 surfaces. To maintain the smooth real time response our system can handle up to 80.000 polygons, which allows it to visualize about 20.000 data points simultaneously and still have smooth visualization when changing viewpoint. The data points may be colored to visualize one more variable.

Initially the 3-D Scatter Plot just gives us spatial resolution compared to the 3-D Histogram. But this higher resolution allows further exploration of the navigation facility in virtual space. One may have a close look at a local configuration of data points and smoothly change to alternative viewing direction and/or viewing distance, and hence gradually obtain a global view and observe local configurations in the large context.

Another possibility is to calculate a surface map of the data. A (very) simple "Kernel density estimate" function is used for the calculations as shown in figure 5. More work has to be done to make this a useful tool.




Fig. 5. Scatter Plot with surface.

It is also possible to "highlight" some of the visual objects. This is done by specifying an expression exactly as when choosing a set of filters. The highlighted records are visualized as visual objects blinking between their original color and white.

3-D Object Property Space In an attempt to map a larger number of statistical variables into the visualized world we have extended the 3-D Scatter Plot tool to become the "Object Property Space". The objects visualizing the data points may be given various visual properties that may illustrate other variables as exemplified in figure 6.




Fig. 6. A look inside Object Property Space, using position, color, shape, and size to represent statistical variables.

In principle a large range of possibilities emerge when taking this approach. Visually perceivable object characteristics are form, size, surface texture, and/or color and object orientation. These are all static properties and we may add animations that make objects vibrate or rotate with different amplitude, frequency and phase.

Taking this line of thinking further one may also let one of the variables drive a temporal development of the visual space. In a simple version a time series of "snapshots" may be visualized and in the more advanced version, for which methods still need to be developed, a continuous temporal development may be visualized. At this stage of the project we have implemented the use of object form, size, orientation and surface color as well as the snapshot series.

3-D Scatter Plot Matrix The 3-D Scatter Plot Matrix tool allows one to view multiple, small 3-D Scatter Plots simultaneously, making it possible to obtain an overview of a data set with alternative combinations of variables used for spatial position, see figure 7.

To maintain smooth real time interaction, fewer data points are used in each of the small Scatter Plot. Since the coordinate systems are relatively small, the tetrahedra are only shown in white. The highlight function also works in this plot. The highlighted tetrahedra are shown in magenta.

This tool is useful, when deciding which variables to use for the bare spatial distribution of data points (objects) in a full Scatter Plot.

3-D Scatter Plot Tour The 3-D Scatter Plot Tour tool is equivalent to the 3-D Scatter Plot Matrix tool, except that it only shows one coordinate system at the time. All possible unique combinations of the selected variables are shown




Fig. 7. Scatter Plot Matrix.

as positional variables in snapshots of 5 seconds each. It is possible to pause the animation at any given snapshot and navigate around before next combination is visualized.

The main advantage of this visualization in comparison to the "3-D Scatter Plot Matrix" is that much more data points can be used in each Scatter Plot. This makes it possible to add more visual cues to each data point.

4.2 Use and Performance

The 3DVDM system uses GNU's General Public License and is publicly available on the Internet [13]. It has been downloaded by users from international research institutions and bug-reports have been received. So far, however, it has not been used to any large extent by the public.

The system has automatic installations scripts for both SGI Irix computers, and Linux computers and it has a tool for loading data. However, configuration of CAVELib has to be done individually for each VR arena to be used.

Parameters are specified through a user-friendly X-Windows menu. The system also allows real-time interaction through VR input devices.

Flexibility One of the main advantages of the system is the ease with which it is possible to extend the system with new modules. Developers can concentrate on working with their modules, without interfering with the work of other developers.

Performance Loading of data from databases is done with MySQL, and is considered to be relative slow compared with other applications.

Currently only one processor in a multiprocessor computer is used for statistical processing.

However, Performer and CAVELib use multiple processes for rendering the 3-D graphics on monitors, as well as in VR arenas. On our Onyx2 supercomputer this allows visualization of 100.000 tetrahedra simultaneously. If smooth real-time interaction is required then the limit is about 20.000 tetrahedra simultaneously.

Immersiveness Our preliminary experiences indicate that in particular the "object property space", where the analyst is immersed in data, gives possibilities of navigating around to find "new" interesting structures and relationships in data. Even the Panorama arena provides well for the immersive experience, and it has the advantage of allowing a larger group of more than 10 people to take part in and discuss the same visual analysis.

5 Discussion and Future Work

As of today, the data handling part of the system performs poorly compared with state-of-the-art, and it is insufficient for very large data sets. An interesting solution to this problem would be to implement parallel data extraction into the system.

The system should use multiprocessing to:

1. Increase performance.
2. Allow different time-consuming visualizations, e.g. snapshots, to be calculated simultaneously.
3. Make it possible for a user of the system to obtain response to a complicated question in real-time.
4. Experiment with dynamic object properties.

3-D menus, buttons, etc. are highly desirable, as well as audio aspects of VR.

It should also be possible to map, e.g. the position of the user in real-time directly to a process performing a statistical calculation based on this information. This would also make it possible for a statistical process to control the users' relative position dynamically in real-time.

Dynamic particle systems have been known since 1983 [14]. An interesting experiment would be to investigate how useful such systems are for "object property space" visualizations concerning, e.g. dynamic object properties.

Visual perception is often made in reference to a solid ground. In particular we suggest that such a perceptually "solid ground" for depth perception in the nearest neighbourhood provides a good reference for exploiting object pose as an informative property. Thus the free navigation and browsing in the 3-D world may provide the opportunity for seeing the data (objects) from arbitrary viewpoints and from arbitrary view-directions. This might aid the analyst in finding

visual events of potential interest, and it is in particular important for exploiting the information encoded in the pose-parameter. Seeing along their longitudinal versus perpendicular view could give cues about perceptual grouping hinting at clusters and structures.

We plan to carry out experiments on perceptual tasks in relation to the mentioned objects properties. Hereby we hope to test our working hypothesis that visual data mining is facilitated in immersive virtual environments.

6 Conclusion

This paper presented work in progress, with emphasis on a description of an approach to Visual Data Mining in 3-D immersive environments. The project aims at investigating and hopefully verifying that current immersive visualization technology in combination with the human perceptual capabilities provide for a new scope of explorative data analysis.

A flexible, maintainable system architecture was presented, as well as several methods for exploring data in VR. Our first findings concerning the encoding of a larger number of statistical variables and the use of human perceptual skills in the field of Visual Data Mining in VR are promising.

Acknowledgements

We gratefully acknowledge the support to the 3DVDM project from the Danish Research Councils, grant no. 9900103.

References

- [1] A. Ammoura. Dive-on: From databases to virtual reality. *ACM Crossroads Database Special Edition*, 7(3), 2001.
- [2] D. Asimov. The grand tour: A tool for viewing multidimensional data. *SIAM. Journal of Science and Statistical Computing.*, 6:128–143, 1985. (original paper, 2-D grand tour).
- [3] K.R. Boff and J.E. Lincoln. *Engineering data compendium: human perception and performance*. Ohio: Harry G. Armstrong Aerospace Medial Research Laboratory, 1988.
- [4] A. Buja and D. Asimov. Grand tour methods: an outline. In *Computing Science and Statistics: Proceedings of the Seventeenth Symposium on the Interface*, pages 63–67, 1985. (2-D grand tour).
- [5] L. Chenk. Topological structure in visual perception. *Science*, 218(4573):699–700, 1982.
- [6] R.E. Christ. Review and analysis of color coding research for visual displays. *Human factors*, 17:542–570, 1975.
- [7] E.G. Davis and R.W. Swezey. Human factors guidelines in computer graphics: A case study. *Man Machine Studies*, 18(2):113–133, 1983.
- [8] George Eckel. *IRIS Performer Getting Started Guide*. Silicon Graphics, Inc., 1997.

- [9] A. Friedman and D.L. Hall. The importance of being upright: Use of environmental and viewer-centered reference frames in shape discriminations of novel three-dimensional objects. *Memory and Cognition.*, 24(3):285–295, 1996.
- [10] E. Granum and P. Musaeus. Constructing virtual worlds for visual explorers. In L. Qvortrup, editor, *Virtual Space Construction: The Spatiality of Virtual Inhabited 3D Worlds*. Springer, Berlin, 2001. (In press).
- [11] E.J. McCormick and M.S. Sanders. *Human factors in engineering and design. (5th ed.)*. Mcgraw-Hill Book Company., 1983.
- [12] L. Nelson, D. Cook, and C. Cruz-Neira. Xgobi vs the c2: Results of an experiment comparing data visualization in a 3-d immersive virtual reality environment with a 2-d workstation display. *Computational Statistics*, 14:39–51, 1999.
- [13] The 3DVDM project group. <http://www.cs.auc.dk/3DVDM/>.
- [14] W. T. Reeves. Particle systems - a technique for modeling a class of fuzzy objects. In *Proc. of SIGGRAPH '83*, 1983.
- [15] N. Sawant, C. Scharver, J. Leigh, A. Johnson, G. Reinhart, E. Creel, S. Batchu, S. Bailey, and R. Grossman. The tele-immersive data explorer: A distributed architecture for collaborative interactive visualization of large data-sets. *Proceedings of 4th International Immersive Projection Technology Workshop, Ames, Iowa*, June 2000.
- [16] A. Treisman. Perceptual grouping and attention in visual search for features and for objects. *Experimental Psychology: Human Perception and Performance.*, 8(2):194–214, 1982.
- [17] C. Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann Interactive Technologies Series, 2000.
- [18] E. J. Wegman. The grand tour in k-dimensions. In *Computing Science and Statistics: Proceedings of the 22nd Symposium on the Interface*, pages 127–136, 1991. (general k-dimensional grand tour).
- [19] L. Wilkinson. *The Grammar of Graphics*. Springer, NY, 1999.