

The Visual Scalability of Integrated and Multiple View Visualizations for Large, High Resolution Displays

Beth Ann Yost

Dissertation submitted to the faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
In
Computer Science and Applications

Chris North, Chair
Doug Bowman
Bill Carstensen
Roger Ehrich
Francis Quek

April 11, 2007
Blacksburg, VA

Keywords: Information Visualization, Large Displays, Visual Scalability,
Multiple Views, Usability Testing and Evaluation

The Visual Scalability of Integrated and Multiple View Visualizations for Large, High Resolution Displays

Beth Ann Yost

(ABSTRACT)

Geospatial intelligence analysts, epidemiologists, sociologists, and biologists are all faced with trying to understand massive datasets that require integrating spatial and multidimensional data. Information visualizations are often used to aid these scientists, but designing the visualizations is challenging. One aspect of the visualization design space is a choice of when to use a single complex integrated view and when to use multiple simple views. Because of the many tradeoffs involved with this decision it is not always clear which design to use. Additionally, as the cost of display technologies continues to decrease, large, high resolution displays are gradually becoming a more viable option for single users. These large displays offer new opportunities for scaling up visualization to very large datasets. Visualizations that are visually scalable are able to effectively display large datasets in terms of both graphical scalability (the number of pixels required) and perceptual scalability (the effectiveness of a visualization, measured in terms of user performance, as the amount of data being visualized is scaled-up).

The purpose of this research was to compare information visualization designs for integrating spatial and multidimensional data in terms of their visual scalability for large, high resolution displays. Toward that goal a hierarchical design space was articulated and a series of user experiments were performed. A baseline was established by comparing user performance with opposing visualizations on a desktop monitor. Then, visualizations were compared as more information was added using the additional pixels available with a large, high resolution display. Results showed that integrated views were more visually scalable than multiple view visualizations. The visualizations tested were even scalable beyond the limits of visual acuity. User performance on certain tasks improved due to the additional information that was visualized even on a display with enough pixels to require physical navigation to visually distinguish all elements. The reasons for the benefits of integrated views on large, high resolution displays include a reduction in navigation due to spatial grouping and visual aggregation resulting in the emergence of patterns. These findings can help with the design of information visualizations for large, high resolution displays.

Grant Information

This research was partially supported by the National Science Foundation grants #CNS-04-23611 and #IIS-0617975. This research was also supported and monitored by the Advanced Research and Development Activity (ARDA) and the National Geospatial-Intelligence Agency (NGA) under Contract Number HM1582-05-1-2001. The views, opinions, and findings contained in this report are those of the author and should not be construed as an official US Department of Defense position, policy, or decision, unless so designated by other official documentation.

Acknowledgements

This work would not have been possible without the support and encouragement of numerous people. First, thanks to my advisor, Chris North, for countless hours of thoughtful discussion. You managed to keep it all fun even during the stressful pre-proposal stage. Thanks also to each one of my committee members. You have all contributed something unique to this work through your feedback and guidance along the way. This work benefited significantly from your experience and expertise.

Several people helped with the experiments presented in this document, particularly the early background experiments in Chapter 4. Thanks go to Gregorio Convertino, Jian Chen, and Young-Sam Ryu for your work on the experiment presented in Section 4.2. Your attitude and aptitude provided an encouraging start to my graduate career. Thanks also to Ying Qiao, Bing Fang, and Yamin Wu for your work on the experiment presented in Section 4.1 and to Yogita Bhardwaj, Pak Chung, and Xiaoyan Yu for your work on the experiment presented in Section 4.3. Another invaluable assistant in this process was Yonca Haciahmetoglu, who helped considerably with many aspects of the experiment presented in Section 5.3. Thanks also to Nicholas Polys and Purvi Saraiya for your discussions about and contributions to the design space.

Many other students at Virginia Tech and within the LIVE research group and CHCI have helped throughout the process. Thanks to Lauren Shupp for lending your graphic design expertise and for your encouragement. Thanks to John Booker, Mike Dellanoche, Bob Ball, and Christopher Andrews for your assistance with and maintenance of the displays and to other members of the LIVE research group including: Sarah Peck, Glenn Fink, Sujatha Krishnamoorthy, Qing Li, and Nathan Conklin.

Most of all, to my family and friends, for your loving support throughout this process – thank you!

The Visual Scalability of Integrated and Multiple View Visualizations for Large, High Resolution Displays

Table of Contents

Chapter 1. Introduction.....	1
1.1. Motivation.....	1
1.2. Visualizations.....	2
1.3. Data Types	3
1.3.1. <i>Spatial Data</i>	3
1.3.2. <i>Multidimensional Data</i>	3
1.4. Research Questions.....	4
1.5. Significance.....	5
Chapter 2. Literature Review	6
2.1. Human Perception.....	6
2.1.1. <i>Extracting Low-Level Properties</i>	6
2.1.2. <i>Pattern Processing and User Goals</i>	8
2.2. Large, High Resolution Displays.....	10
2.2.1. <i>Overview of the Technology</i>	10
2.2.2. <i>User Studies with Large, High Resolution Displays</i>	12
2.3. Information Visualizations.....	14
2.3.1. <i>Scalability</i>	14
2.3.2. <i>Integrated and Multiple Views</i>	14
2.3.3. <i>Linked Visualizations</i>	15
2.4. Information Visualization Design Spaces.....	16
2.5. Summary	17
Chapter 3. Design Space	18
3.1. Hierarchical Design Space.....	18
3.2. Extending the Design Space	20
3.3. Definitions.....	22
Chapter 4. Desktop Display Baseline	23
4.1. Experiment 1: Multiple Simple vs. Single Complex Glyphs.....	23
4.1.1. <i>Method</i>	23
4.1.2. <i>Results</i>	26
4.1.3. <i>Discussion</i>	28
4.1.4. <i>Summary</i>	29
4.1.5. <i>Conclusion</i>	29
4.2. Experiment 2: Context Switching and Cognitive Factors.....	30
4.2.1. <i>Method</i>	30
4.2.2. <i>Results</i>	34
4.2.3. <i>Discussion</i>	40
4.2.4. <i>Summary</i>	42
4.2.5. <i>Conclusion</i>	42
4.3. Experiment 3: Visual vs. Interactive Linking	42
4.3.1. <i>Method</i>	42
4.3.2. <i>Results</i>	46

4.3.3.	<i>Discussion</i>	48
4.3.4.	<i>Conclusion</i>	50
4.4.	Summary and Synthesis of Results.....	51
Chapter 5.	Visual Scalability for Large Displays	53
5.1.	Graphical Scalability.....	53
5.2.	Experiment 4: The Perceptual Scalability of Visualization.....	54
5.2.1.	<i>Method</i>	55
5.2.2.	<i>Results</i>	61
5.2.3.	<i>Discussion</i>	68
5.2.4.	<i>Conclusion</i>	68
5.3.	Experiment 5: Scaling Beyond the Limits of Visual Acuity.....	69
5.3.1.	<i>Method</i>	70
5.3.2.	<i>Results</i>	75
5.3.3.	<i>Discussion</i>	81
5.3.4.	<i>Conclusion</i>	82
5.4.	Summary and Synthesis of Results.....	82
Chapter 6.	Information Visualization Design Issues on Large, High Resolution Displays	85
6.1.	Perceptual Design Considerations.....	85
6.1.1.	<i>Visual Encodings</i>	85
6.1.2.	<i>Visual Acuity and Physical Navigation</i>	90
6.2.	Interaction Design Considerations.....	92
6.2.1.	<i>Navigation Techniques</i>	93
6.2.2.	<i>Brushing and Linking</i>	94
6.2.3.	<i>Selection</i>	94
6.2.4.	<i>Control Panels</i>	95
6.3.	Summary and Open Issues.....	95
Chapter 7.	Conclusion	98
7.1.	Research Summary.....	98
7.1.1.	<i>Design Space</i>	98
7.1.2.	<i>Baseline Comparison</i>	98
7.1.3.	<i>Perceptual Scalability</i>	99
7.1.4.	<i>Design Issues</i>	100
7.2.	Research Evolution.....	101
7.3.	Visualization Design and User Task Performance.....	102
7.3.1.	<i>Basic Concept</i>	102
7.3.2.	<i>Tasks</i>	102
7.3.3.	<i>Link between Tasks and Design Space</i>	103
7.3.4.	<i>Empirical Support for the Link</i>	104
7.4.	Large Display Visualization Design Guidelines.....	105
7.5.	Contributions.....	106
7.6.	Future Work.....	107
References		108

List of Figures

Figure 1.1. Large, high resolution display	1
Figure 1.2. Two basic visualization approaches for integrating spatial and multidimensional data.....	2
Figure 1.3. Tabular form for spatially-referenced multidimensional data.....	4
Figure 1.4. Tabular form with temporal data.....	4
Figure 2.1. Digit in red is a pre-attentive visual encoding.....	8
Figure 2.2. Examples of different display technologies	11
Figure 3.1. Space-centric visualizations with different spatial contexts.....	19
Figure 3.2. Attribute-centric visualizations having different spatial contexts	19
Figure 3.3. Space-centric visualization with 3 data types (time-series graphs at each location)	20
Figure 3.4. Attribute/Time-centric design with 3 data types: 2D map for each time/attribute pair.....	20
Figure 3.5. Attribute-centric design with 3 data types: Separate map for every attribute with value-flow maps ([7]) overlaid	21
Figure 3.6. Time-centric design with 3 data types: Separate map for every time point with parallel coordinates ([65]) overlaid.....	21
Figure 3.7. Multiple linked hierarchies correspond to linked visualizations (image shows Snap software in use at the Census Bureau [96])	21
Figure 4.1. One view (space-centric), data attributes mapped to color, size, density, and orientation	24
Figure 4.2. Two views (mixed design), data attributes mapped to color and size on two maps	24
Figure 4.3. Four views (attribute-centric), data attributes on their own map with color used in all	24
Figure 4.4. Completion times for detecting two attributes	27
Figure 4.5. Completion times for one-attribute trend task.....	27
Figure 4.6. Combinations of linked visualizations	31
Figure 4.7. Pattern recognition tasks.....	32
Figure 4.8. Task completion time by task type and combination of visualizations.....	34
Figure 4.9. Completion time and picture completion subtest scores	37
Figure 4.10. Percent of time spent looking at different elements	38
Figure 4.11. Visual paths for SG visualization, finding a pattern within the map.....	38
Figure 4.12. Visual paths for PG visualization, finding a pattern across views	39
Figure 4.13. Visual switches between views	39
Figure 4.14. Integrated 3D view	44
Figure 4.15. Two views with linked scatter plots	44
Figure 4.16. Multiple linked 1D plots.....	45
Figure 4.17. Performance times in seconds for each task and interface, dotted lines indicate significant differences	46
Figure 4.18. Accuracy for each task and interface, the dotted lines indicate significant differences.....	47
Figure 4.19. Satisfaction ratings, the dotted lines indicate significant differences.....	47
Figure 5.1. Visible data by number of available display pixels.....	54

Figure 5.2. Attribute-centric visualization (MULTS) with 9 attributes on the 2Mp display	56
Figure 5.3. Space-centric (embedded) visualization using bar matrices [BARS]	57
Figure 5.4. Space-centric (embedded) visualization using time series graphs [GRAPHS]	57
Figure 5.5. Bar matrices visualization on a 32 Mp display	58
Figure 5.6. Attribute-centric (MULTS) in two display conditions.....	59
Figure 5.7. Space-centric (GRAPHS) in two display conditions.....	59
Figure 5.8. Scaled-up embedded visualizations (BARS, GRAPHS) for a single geographic location.....	59
Figure 5.9. Task completion times.....	62
Figure 5.10. 2 Mp differences in time per attribute. Bars are 95% confidence intervals. Tasks with significant differences are marked with a ‘*’, and have non-overlapping bars.....	62
Figure 5.11. 32 Mp differences in time per attribute. Bars are 95% confidence intervals. Tasks with significant differences are marked with a ‘*’, and have non-overlapping bars.....	63
Figure 5.12. Interaction between display and visualization for DS task	63
Figure 5.13. Interaction between display and visualization for OT task	64
Figure 5.14. Interaction between display and visualization for OS task	64
Figure 5.15. Workload on display sizes. Significant differences have a ‘*’	66
Figure 5.16. Frequency of user visualization preference for 2 Mp display (rank of 1=best, 3=worst)	67
Figure 5.17. Frequency of user visualization preference for 32 Mp display	67
Figure 5.18: Does an optimal display match visual acuity?	69
Figure 5.19. Visualization designs used in visual acuity experiment	70
Figure 5.20. Medium display condition (equal visual acuity)	72
Figure 5.21. Space-centric embedded visualization on different display sizes	73
Figure 5.22. Factor increases in task completion times.....	76
Figure 5.23. Speed and accuracy for the OA task. Small is significantly faster than large and significantly more accurate than both medium and large.	77
Figure 5.24. Speed and accuracy for the OT task. Accuracy was significantly better with medium and large than small. Space-centric was significantly faster than attribute-centric in the large condition.....	78
Figure 5.25. Speed and accuracy for the OST task. Medium is more accurate than small.	79
Figure 5.26. Speed and accuracy for the CA task. Small and medium were significantly faster than large, and medium was the most accurate display size.	80
Figure 6.1. The number of glyphs can be increased using additional pixels, the number of perceptually distinct colors cannot.....	86
Figure 6.2. Both local and global legends should be considered, a user may lose physical context when viewing a global legend during a local tasks.....	88
Figure 6.3: Differences between tiles are exaggerated by the viewing angle.....	89
Figure 6.4: Distortions caused by bezels	90
Figure 6.5. Visual aggregation occurs as a user physically moves away from the display	91

Figure 6.6. Physically adaptive visualizations.....	92
Figure 6.7: Navigation strategy issues.....	93
Figure 6.8: Brushing and linking issues.....	94

List of Tables

Table 1.1. Differences between basic visualization approaches.....	2
Table 2.1. Display specifications for different technologies	11
Table 4.1. Visual encodings used in each visualization condition	25
Table 4.2. Tasks and visual encodings.....	25
Table 4.3. Practice tasks and experimental tasks.....	45
Table 5.1. Calculating graphical scalability with opposing visualizations.....	53
Table 5.2. Task types and examples	60
Table 5.3. Mean familiarity ratings per display condition.....	61
Table 5.4. Task completion time means (seconds).....	62
Table 5.5. Significant normalized time differences	65
Table 5.6. Total correct answers (12 max)	65
Table 5.7. Task workload means	66
Table 5.8. Display size conditions	72
Table 5.9. Tasks (D: detail, O: overview, and C: complex; A: attribute, T: time, and S: space)	74
Table 5.10. User familiarity ratings (1=strongly disagreed, 5=strongly agreed).....	74
Table 5.11. Average task completion times (seconds)	75
Table 5.12. Accuracy (percent correct).....	75
Table 6.1: Examples of graphically scalable and non-scalable representations	85
Table 6.2: Display property impacts on graphical encodings.....	87
Table 7.1. Overview of basic tasks (examples provided in each cell).....	103
Table 7.2. Task hierarchies	103

List of Equations

Equation 1.1. General form for spatial locations	3
Equation 1.2. General form for a single spatially-referenced dimension	3
Equation 1.3. General form for spatially-referenced multidimensional data.....	3

Chapter 1. Introduction

1.1. Motivation

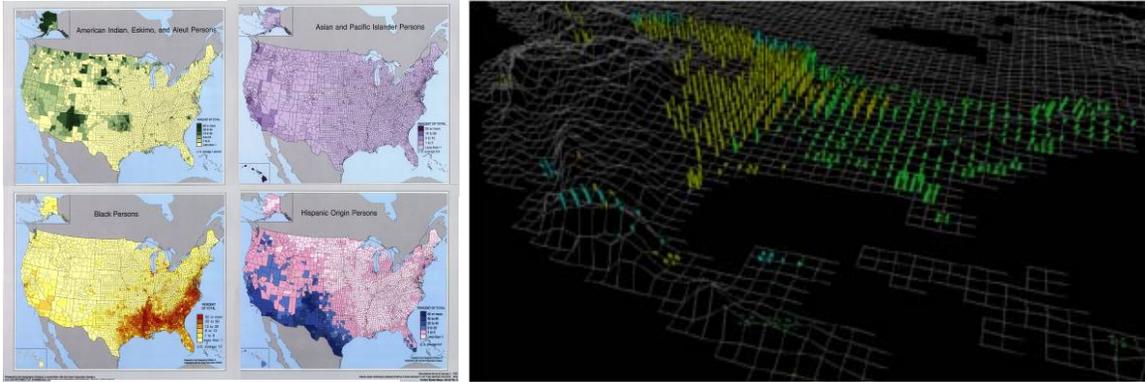
Geospatial intelligence analysts, epidemiologists, and biologists all share a common problem. They are all faced with trying to understand potentially massive datasets that involve integrating spatial and multidimensional data. The intelligence analyst [10] must integrate information about physical structures such as buildings and bridges, location and numbers of enemy forces, and various information related to specific geographic locations. The epidemiologist tries to find the cause of outbreaks of diseases using weather fronts, medical data, and absenteeism across a geographic location [79]. The biologist must consider the relationship between a biological pathway and various experimental results [118].

Information visualizations can be used to help gain insight when integrating the types of information described, but designing these visualizations is a challenge. One aspect of the design space for integrating spatial and multidimensional data, and a difficult design decision, is the question of how many distinct views of the information to present. The integrated view approach [60] emphasizes the advantage of overlaying information, while the multiple views [136] approach emphasizes the advantage of simplifying complex views. There are a number of tradeoffs involved with this design decision [11] that make it unclear which design to apply.

As the cost of display technologies decrease and the size of datasets increase, information visualizations are increasingly being viewed using large, high resolution displays (Figure 1.1) [66, 116]. The implication of the adoption of this technology is that the scalability limits imposed by the number of pixels is now much less of a limit on visualization design. However, as more information can be displayed consideration of human perceptual and cognitive abilities becomes increasingly important. The purpose of this research was to determine the visual scalability [49] of integrated and multiple view visualizations for integrating spatial and multidimensional data types as we move to large, high resolution displays.



Figure 1.1. Large, high resolution display



(a) Attribute-centric [1]

(b) Space-centric using complex glyphs [59]

Figure 1.2. Two basic visualization approaches for integrating spatial and multidimensional data

1.2. Visualizations

An important first step toward understanding the tradeoffs when integrating spatial and multidimensional data was to lay out the design space. The design space evolved during the course of this work, but two basic opposing visualizations remained. We have named these visualizations space-centric and attribute-centric. A **space-centric** visualization has all of the multidimensional data overlaid onto a single spatial structure, typically using complex glyphs. A space-centric visualization is a type of integrated view visualization. An **attribute-centric** visualization has each attribute of the multidimensional data on a separate spatial structure. This visualization is a type of multiple view visualization. Examples of these opposing visualization approaches can be seen in Figure 1.2. A more in-depth explanation of the exact terminology used throughout this work is given alongside the design space in Chapter 3.

There are many tradeoffs between the attribute and space-centric visualizations. Some of the basic differences between the approaches include the location of the complexity in the visualization, the size of the individual views, and tradeoffs in the visual and mental operations required with each visualization. A summary of these issues and tradeoffs is shown in Table 1.1. The basic differences in the visualizations result in different strengths and weaknesses that will impact different types of tasks.

Table 1.1. Differences between basic visualization approaches

	Attribute-centric	Space-centric
Approach	Each attribute displayed on a separate spatial structure	All attributes displayed on a single spatial structure
Location of complexity	Number of views	Glyphs/visual encoding
Size of view	Display size/number of views	Full display size
Tradeoffs	For a single attribute, all locations are grouped	For a single location, all attributes are grouped
	For a single location, must visually integrate across structures	For a single attribute, must visually integrate across locations
	Multiple structures containing irrelevant attributes	A single location containing irrelevant attributes

1.3. Data Types

Central to this work is the concept of integrating spatial and multidimensional data. For the sake of clarity regarding the type of data, the data types under consideration are introduced.

1.3.1. Spatial Data

Although the focus is on geospatial data, by spatial data we mean any data type that has an inherent structure to it (the reason for including other spatial data types will be briefly discussed later). This group includes 2D structures such as geographic data which has latitude (y) and longitude (x), 3D structures which can be specified using an x , y , and z location, and graph structures which can be represented by depth in a graph (y) and the order in which a node appears on that level of the graph (x). Note that the spatially-referenced data we are referring to is discrete and not continuous. In other words, by x and y we are referring to a point within a polygon such as a point within the middle of a U.S. state, a point on a 3D object, or a point within a node in a graph. We refer to these spatially-referenced points as locations and give the general form:

$$l_i = (x_i, y_i)$$

Equation 1.1. General form for spatial locations

For graphs, the graph layout defines the space. The layouts can either be defined by the user or algorithmically. The graph layout is similar to the 2D and 3D layouts because the particular layout of the space is not considered, each node simply has an x,y or x,y,z position. We view this as a reasonable generalization since we are not considering the cases where space is abandoned or manipulated, only cases where the space is displayed as is. Thus by spatial we mean anything where the space is predefined.

1.3.2. Multidimensional Data

We refer to multidimensional data in a similar manner as Muller and Schumann, as the dimensionality of the independent variables [92]. Using that definition, a single spatially-referenced dimension (d) is a quantitative independent variable for which a specific value is given for all locations (l). This can take the general form:

$$d_i = f(l_i)$$

Equation 1.2. General form for a single spatially-referenced dimension

Multiple spatially-referenced dimensions, or spatially-referenced multidimensional data (M), is then defined as a dataset that contains multiple dimensions for each location:

$$M = \{d_1, d_2, \dots, d_n\}$$

Equation 1.3. General form for spatially-referenced multidimensional data

Another way to think of this data is in tabular form (Figure 1.3). In a table, each row might represent a single location. Then, the first two columns might represent the spatial coordinates and each remaining column might represent a different dimension (independent variable).

	X	Y	d_1	d_2	...	d_n
l_1						
l_2						
...						
l_n						

Figure 1.3. Tabular form for spatially-referenced multidimensional data

Temporal data that is spatially-referenced is roughly equivalent to multidimensional data that is spatially-referenced. However, if the dataset is a combination of multidimensional and temporal data, then time is an orthogonal dimension. By this we mean that adding a single time point results in multiple values being added (i.e. there are values for each dimension at that time point and therefore a new row is added to the table). In a tabular format, the data may be arranged by having a separate table for each location (see Figure 1.4), in the same way that Monmonier describes geographic time-series data [91].

$l_1(x_1, y_1)$	d_1	d_2	...	d_j
t_1				
t_2				
...				
t_n				

$l_k(x_k, y_k)$	d_1	d_2	...	d_j
t_1				
t_2				
...				
t_n				

Figure 1.4. Tabular form with temporal data

These definitions lead to the restrictions of a pre-defined spatial structure and a complete dataset. Aggregation and filtering strategies are not considered. The implications of these restrictions are that uncertainty visualization, aggregation, and spatial layout are not discussed.

1.4. Research Questions

The overall research goal was to *compare the visual scalability of integrated and multiple view visualizations for large, high resolution displays* in the context of integrating spatial and multidimensional data. **Visual scalability** includes human perceptual capabilities as well as display capabilities such as resolution. Towards that goal, the main research questions were:

1. What is the visualization design space for integrating spatial and multidimensional data?
2. As a baseline, which visualization results in the best performance for specific tasks on a desktop display?

3. How visually scalable are the visualizations for large, high resolution displays?
 - a. Display Issues: Pixel Count
 - b. Human Issues: Perception and Cognition

4. What characteristics of a visualization determine how visually scalable it is for large, high resolution displays and what design issues remain?

The research approach was to develop a design space for integrating spatial and multidimensional data and then run a series of experiments comparing these visualizations. As the experiments progressed and more was learned the design space evolved into the form presented in this document. Three experiments compared visualizations on a desktop monitor. These comparisons included attribute vs. space-centric, an exploration of cognitive issues with linked visualizations, and a comparison of visual versus interactive integration.

The heart of this work lies in the visual scalability of designs and is found in Chapter 5 and Chapter 6. Two experiments compared the perceptual scalability of visualizations for large, high resolution displays. The first experiment explored general cognitive and perceptual issues. The results of this experiment led to the second experiment which asked: what effect does visual acuity have on the perceptual scalability of visualizations for large, high resolution displays? Based on what we learned during the design of information visualizations used in these experiments we further explored characteristics that make a particular visualization more perceptually scalable.

1.5. Significance

This work helps visualization designers understand the relationship between human capabilities and their design choice as we move from desktop displays to large, high resolution displays. Specific contributions of this work include:

- A hierarchical design space for integrating spatial and multidimensional data
- Analysis of the tradeoffs involved with opposing visualization designs
- Demonstrated that even displays that are large enough and include enough pixels to exceed visual acuity can still be beneficial because of the additional data that can be displayed
- Provided initial design guidelines and suggestions as well as outlining a number of open research questions in the area of information visualization design for large, high resolution displays

This research aids our understanding of the visual scalability of different information visualization design approaches for integrating spatial and multidimensional data as we transition to large, high resolution displays.

Chapter 2. Literature Review

Information visualizations are useful for providing insight into complex datasets. As the size of these complex datasets increases, visualization designs must be able to scale appropriately. Eick and Karr discuss and define visual scalability as, “the capability of visualization tools effectively to display large data sets, in terms of either the number or the dimension of individual data elements” [49]. Beginning with the human and moving away, they include the following as aspects of visual scalability:

- Human perception – perception and cognition of visual patterns
- Monitor resolution – physical size and pixel density (DPI) of the display
- Visual metaphors – choice of metaphor and mapping of data to visual attributes
- Interactivity
- Data structures and algorithms
- Computational infrastructure

Using this as a framework, the work was focused on the first three bullets above – the factors closest to the user. Therefore, this section begins with a discussion of relevant human perceptual and cognitive issues. Next, research related to using large, high resolution displays is discussed. That is followed by a discussion of the visual metaphors (integrated and multiple view visualizations). After covering those aspects of visual scalability, information visualizations design spaces are surveyed to provide background on the ways the design space builds on and extends previous work.

2.1. Human Perception

“On a daily basis, we are fooled about the extent to which we constantly make eye movements and the extent to which these eye movements contribute to the illusion that a whole scene is simultaneously in view. Some scientists refer to this phenomenon as the *grand illusion of complete perception* to make the point that it is the most pervasive and fundamental of all the visual illusions that have been discovered so far.” –(Enns [52] p. 176)

While the grand illusion of complete perception is debatable, it demonstrates the need to better understand how we visually process information. Colin Ware provided a simplified version of a visual information processing (VIP) model [139]. This model will be used to organize research related to perception and cognition of visual patterns. The three stages in his model are 1) extracting low-level properties, 2) pattern processing, and 3) goal directed processing.

2.1.1. Extracting Low-Level Properties

In Stage 1 in Ware’s VIP model (extracting low-level properties) neurons in the eyes are stimulated and basic visual properties are extracted from the visual scene [139]. Three particular aspects of this stage are important to this work. First, the relationship between the resolution of the human eye and a monitor was used by Ware to theorize an “optimal display”. The second and third important aspects of this stage impact glyph design and include the concepts of pre-attentive processing and integral and separable dimensions.

The “Optimal Display”

We have approximately 125 million photoreceptors in each of our eyes. Of these, approximately 120 million are rods that are mainly responsible for our peripheral vision and 5 million are cones located mostly in or near the fovea or center region of our eyes. These cones provide high visual acuity and color vision. However, there is not a 1-to-1 mapping from photoreceptors to the neurons that carry these signals to our brain [139]. Ware uses the term “brain pixel” to refer to the ganglion nerves available to carry signals from the rods and cones to our brains. These brain pixels are not equally distributed. There are significantly more brain pixels per photoreceptor in the fovea and less as the distance from the fovea increases.

Ware states, “In light of the extreme non uniformity of brain pixels, we can talk about the visual efficiency of a display screen by asking what screen size provides the best match of screen pixels to brain pixels?” [139]. Because there are more brain pixels in the fovea, higher-resolution screens help foveal vision. In our periphery, there are fewer brain pixels than screen pixels so additional information displayed there is not conveyed as efficiently. He uses this fact to suggest that current monitor size is about optimal and that “a 4000x4000-pixel resolution monitor should be adequate for any conceivable visual task.” He cites collaboration as an exception to this.

If a small high resolution display is adequate for any task, why might a large high resolution display be beneficial to a single user? One reason is that there may not be enough pixels to display all of the information. In this case, a *high resolution display can be used to show more information simultaneously while still maintaining context*. If aggregation or elimination techniques are used [76], the details are lost. A variety of other techniques such as zooming and panning [19], overview + detail, or fisheye views may be used. These techniques introduce a variety of problems. In the case of zooming and panning context is lost because the surrounding area is not shown. With the overview + detail approach there is a need to switch between views to gain context. With the fisheye approach contextual information is distorted. A high resolution display avoids these problems to a greater extent.

Glyph Design

Glyphs, objects whose visual attributes (size, shape, color, etc.) represent different data attributes, are often used to represent data in the integrated views approach. In stage 1 in the VIP model, the concepts of pre-attentive processing and integral and separable dimensions are useful when designing these glyphs. After information makes it through the photoreceptors in our eyes and to our brain, information is held in a type of memory called iconic memory for less than 2 seconds [143]. Rapid bottom-up parallel processing of visual information from the environment occurs in this stage of visual processing.

Information that is pre-attentive “pops out” from the surrounding information (Figure 2.1). Because eye movements take approximately 200ms to initiate, any task that can be accomplished in less than 250ms is considered pre-attentive. This is useful for target detection and area estimation tasks and a visualization tool exists for mapping data

attributes to pre-attentive visual features [58]. As the target to detect becomes less distinct or when conjunctive encodings are used these features become more difficult to detect and are no longer pre-attentive. Therefore, the scalability of this approach is limited to perhaps two data attributes.

42879542039847542

Figure 2.1. Digit in red is a pre-attentive visual encoding

Related to the concepts of glyph design, and particularly to pre-attentive processing is the work of Bela Julesz. In 1962 he made the famous conjecture that textures could not be discriminated if their first (brightness) and second (granularity, slope) order statistics were identical [69]. When this was later proven incorrect, he modified his thinking and created a theory based on textons [26, 71]. In his 1995 book [70] he went on to explain that textons (local and non-linear features) are rather hard to define because they depend on both the actual glyphs as well as the white space between those glyphs. In this work he explained that rather than being able to discriminate certain textures based on statistics, features such as color, orientation, and motion (i.e. textons) were the determinants.

When glyphs are used to encode multiple attributes, the concept of integral and separable dimensions adds to glyph design by considering whether two visual encodings are perceived holistically or separately [139]. Representing one attribute using red-green and another using yellow-blue is integral, while representing one attribute using color and another using shape is separable. This interaction between features has led to Wickens's Proximity Compatibility Principle (PCP) [33, 142]. Integral dimensions cause interference when a single dimension is needed, but are useful as a redundant encoding when multiple dimensions are considered simultaneously. For separable dimensions there may not be interference, but there is also no redundancy gain. In general, *when trying to combine visual encodings for design, there is a reasonable limit of being able to represent about 8 distinct dimensions* [139]. A variety of rules exist comparing the relative visual salience of encodings for different types of data [38, 83].

2.1.2. Pattern Processing and User Goals

The second and third stages in Ware's VIP model are the pattern processing and goal-directed processing stages. The three aspects of these stages that are most relevant to this research are the capacity of visual working memory, attentional resources, and mental models. The relevance of each will very briefly be introduced.

Working Memory

While there are many different working memory models [90], most include a central executive that controls the flow of information, verbal working memory, and visual working memory. The often cited limit on working memory is 7 +/- 2 items [89]. However, it is now known that this number is more related to verbal working memory than visual working memory. Research on visual working memory [67, 102, 138] suggests a capacity of 3-5 items, where conjunctions of encodings can be considered a

single item. The capacity limit of working memory imposes a human scalability limit on visualization designs, specifically for tasks that involve comparison of individual objects.

Attention

We have already cited being able to display all information without distortion or elimination as a reason to use large high resolution displays. *An additional reason to use large, high resolution displays is that we often shift our attention and visually scan the display rather than always focusing on a single location.*

There are three different types of attention: selective, focused, and divided [139]. Selective attention involves selecting the element that we wish to attend to, for multiple views this requires deciding which of the views may be most important. It has been shown that people monitoring several instruments were most likely to visually scan horizontally and resisted scanning diagonally [44]. Additionally, data that is often of interest should be placed centrally, and information that is typically viewed sequentially should be placed close together [50, 144]. Even if the seemingly most critical views are selected, because memory is not perfect, users may forget to look at a particular source of information if there are many views [143]. If users do not scan all of the views when many are presented in the multiple views visualization, this is likely to result in decreased accuracy. Wickens and Hollands [16] state that although selective attention can occur without a change in direction of gaze, most of the time it holds true that, “our gaze is driven by our need to attend”.

Focused attention allows users to focus on one aspect of the data without being distracted by extraneous information. If a user becomes too focused, they may have trouble dividing their attention or they may not select appropriate information on which to focus. A potential double-edged sword with large, high resolution displays is that people are more likely to look at local details when images are large and global detail when images are small [9, 72]. If users remain close to the display, this may increase the chances of users seeing small scale patterns, but decrease the chance of users seeing global patterns.

Divided attention is important when more than one task is being done simultaneously. Typically, the more attention must be focused on a single task, the worse people will be with dividing their attention between tasks. Focus + context screens consider these shifts in attention and corresponding eye movements yet reduce the overall cost of the display [17]. This type of display provides detailed information in a smaller higher resolution area, and context using lower resolution in the periphery. While this and related techniques [16] can certainly be less expensive than a high resolution display, some type of virtual interaction (mouse movement, eye tracking) is needed whenever a user looks at a new area of the screen. For any visualization task that involves scanning the environment this is likely to be quite inefficient.

Mental Models

A user’s mental model is also likely to play a role in what they choose to attend to and what information they extract from the visualization. Trafton has showed that when

expert meteorologists are presented a series of visualizations they form a qualitative mental model of the data and then extract quantitative information from it [135]. Research on graph comprehension and how both novices and experts perform more complex tasks is also ongoing [108]. This research is necessary since most previous research has focused on very simple tasks and people appear to use different strategies for accomplishing simple compared to complex tasks. Additionally, little is understood about cognitive issues in exploratory data analysis [123].

2.2. Large, High Resolution Displays

2.2.1. Overview of the Technology

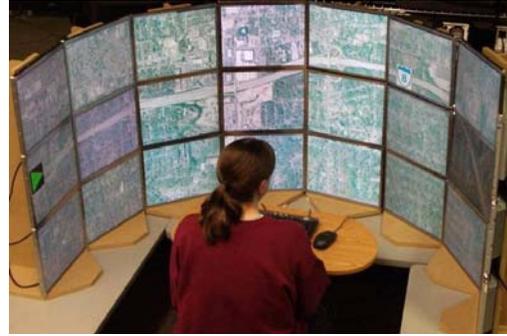
Recent years have seen a dramatic decrease in the cost of LCD displays. According to Odawara's Law, "each doubling in the cumulative area of flat panels produced results in a cost reduction of 22 to 23 percent" [2]. As larger, higher-resolution displays decrease in cost they are becoming accessible to a greater number of people. For example, a twenty-four monitor tiled display of LCD's currently can be built for around US\$10,000. Additionally, improvements in both hardware and software that allow programs to run across tiled displays make this technology accessible to a greater audience since users do not need custom built software to distribute the graphics across displays [61, 64, 75]. Many places such as NASA and AT&T already have large display walls [116, 141].

A raw measure of how much can be shown on a particular display is the number of pixels. In general, the number of pixels in a display can be computed from the DPI and the physical size of the display. Displays can increase the pixel count either by increasing the number of pixels per inch (pixel density), or by maintaining a given number of pixels per inch and tiling projectors or monitors to create a large display (increasing size while maintaining pixel density). For example, IBM's T221 (Big Bertha) has a high pixel density of 204 dots per inch (DPI) in a small 1.5×1 foot area for a total of more than 9 million pixels, while NASA's Hyperwall [66, 116] has a DPI of approximately 90 but in a larger 9×12 foot area for a total of more than 64 million pixels. A number of other display characteristics are also important, including the presence of bezels, viewing angle, color, and brightness uniformity.

Consider five different types of displays: a standard desktop monitor, a display comparable in size but with more pixels per inch (IBM's Big Bertha), a large tiled rear-projected display (VisBlocks), a reconfigurable tiled display of LCD monitors, and a tiled display of fifty LCD touch screens (Figure 2.2). The displays are ordered in terms of the total number of pixels, but also differ in a variety of other ways including their physical size, DPI, and the presence of bezels. The tiled LCD display offers higher DPI than the rear-projection display, but currently requires the presence of bezels between tiles. The rear-projection display offers near-seamless tiling, but is primarily limited by cost and maintenance requirements. The displays can also differ in their physical configuration. Figure 2.2d is configured as a collaborative space, whereas Figure 2.2b is physically arranged for a single user workspace. The specifications for the displays are given in Table 2.1.



(a) IBM's Big Bertha – 204 DPI



(b) Tiled LCD display - curved



(c) VisBlocks – rear-projected



(d) Tiled LCD display – 50 touch screens

Figure 2.2. Examples of different display technologies

Table 2.1. Display specifications for different technologies

Display	Diagonal (inches)	DPI	Number of Pixels
Desktop Monitor	17 [13.3×10.6 in]	96	[1280×1024] = 1,310,720
Big Bertha	22.2 [18.8×11.8 in]	204	[3840×2400] = 9,216,000
VisBlocks (18 blocks)	22 per block [17.6×13.2 in]	58.2	[1024×768] × 18 = 14,155,776
Curved Display (24 monitors)	17 per monitor [13.3×10.6 in]	96	[1280×1024] × 24 = 31,457,280
Touch Screens (50 monitors)	20 per monitor [16×12 in]	100	[1600×1200] × 50 = 96,000,000

Note that the scalability of single monitor displays is currently limited by engineering and manufacturing constraints on DPI and physical size. Presently, the monitor with the most pixels is the IBM Big Bertha (T221-DG5), at approximately 9 Mp. Therefore tiled arrays of monitors are typically used because they are more scalable. Theoretically, one can tile as many monitors as desired, given adequate CPU, graphics processing unit (GPU), and networking infrastructure, but communications become the bottleneck in such configurations (computational performance is beyond the scope of this work). Also note that we refer to the tiled display of LCD monitors as high resolution because the DPI is high relative to other displays of that size (such as a single projector).

There are many other types of displays such as the CAVE, volumetric, and hybrid displays such as focus + context screens [17]. In fact, most research on large, high resolution displays has been about the technology used to create them. Various papers have reported on techniques used to build the displays [75, 125] and software such as Chromium and DMX that can be used to distribute graphics and create a single large desktop across multiple monitors [61, 64, 124]. A survey of these technologies can be found in Ni et al. [95].

Although tiling monitors or projectors is currently the most scalable technique, the seams between the tiles become an issue. Rear-projection arrays have very small seams created by tiling projectors compared to tiled LCD monitors which have larger discontinuities created by physical bezels surrounding each monitor. Each of these situations results in problems aligning images. Some techniques for addressing this have been presented by others [84, 85] and include alternative methods of aligning images and visualizations that provide a representation of the information that is hidden behind the bezels.

The basic issue with the seams is that the spatial position can be affected by bezels which cause spatial distortion if no information is hidden behind the bezels (the default with current software for distributing graphics across large displays) by causing misalignments or geographical areas to appear larger than they actually are. Alternatively, if information is hidden behind the bezel then there is no distortion but the user may be unaware that information is there. There is one study showing that physical discontinuities between information (such as bezels) only impact performance when they are combined with an offset in depth [131]. However in a different study, users expected important information to be hidden behind the bezels and therefore spent time panning images just to determine what might be behind the bezels [80]. Throughout this work we avoid this issue by strategically placing information so that it does not straddle monitor boundaries.

While the technology behind the display is important for assuring a usable display in terms of how much delay is introduced during interaction, more relevant to our work are large display user studies.

2.2.2. User Studies with Large, High Resolution Displays

Benefits of Display Size, DPI, and Field of View

Large displays naturally lead to collaboration research because of size, cost, and privacy concerns. While various papers have dealt with the use of large displays for collaboration [51, 68, 126], in this work we focus on a single user. A variety of benefits have been demonstrated based on the display size, DPI, and field of view. Research on the use of large displays has shown a variety of cognitive and user performance benefits even for basic desktop use [41, 42].

Some of the benefits are likely due to the increased field of view resulting in the creation of better cognitive maps. This is evident by large displays resulting in better performance for 3D navigation tasks [129] and leading to less reliance on way finding aids [94]. This also applies to 2D information, where map navigation is faster and more accurate when

using large displays [13]. Additionally, there is research showing that spatial ability differences due to gender are narrowed when using larger displays because of better optical flow cues [43]. Even when high resolution is not maintained in the periphery of a display, specifically with focus + context screens, there are performance benefits due to the increased field of view [17].

In addition to the advantages of a wider field of view, there also seem to be advantages to information being physically larger and at a higher DPI. When the visual angle is held constant, simply having a larger display compared to a smaller display improved performance on spatial tasks despite increased viewing distance [130, 131]. When the display size was increased and the DPI remained the same, using a larger display improved 3D task performance by making abstract information more visible [94] and allowing for an increased software field of view [104]. While there appear to be many advantages to using large, high resolution displays, none of these studies considered a display of sufficient size and resolution to go beyond the limits of visual acuity - the largest display was less than 5,000 pixels in width. In the portion of this work on perceptual scalability, we go beyond this mark.

Physical Navigation

Typically, information visualizations deal with the graphical scalability limits of desktop monitors by using techniques such as aggregation, elimination, or virtual navigation. Given an unlimited number of pixels, these techniques can be traded for physical navigation. One study showed that for a basic visualization task, physically navigating was faster on a 10 Mp display than panning and zooming on a smaller display [12]. Performance can further be improved if less strenuous physical navigation is required. This can be accomplished by curving the display to bring the outermost pixels closer to the user [122]. That study found performance benefits for a basic route tracing task but not for a basic search task.

Human Vision and the Optimal Display

While brightness, color, and misalignment of images are important issues [5], the most debated issue with large, high resolution displays deals with visual acuity [120]. The question is if pixels are wasted when using higher resolutions in the periphery on large displays [127]. As previously mentioned, Ware asserts that a 4000x4000 display is the optimal display because it is most efficient at matching screen pixels to the “brain pixels” that interpret the signals sent by photoreceptors in our eyes [139].

There is some research that supports the idea that there may be performance costs if visualizations are scaled-up using larger displays. The most relevant research deals with the eccentricity effect, which shows that performance gradually degrades as a target gets farther from our point of visual fixation and that the extent of this effect increases with larger set sizes [32]. With larger displays more details will be farther away and therefore performance may decrease due to the eccentricity effect. Using the additional pixels to scale-up the amount of simultaneously visible information will increase the set size and may thereby exacerbate the problem. Additionally, our visual field is reduced to 92%

under medium mental workload and to 86% under heavy mental workload [106]. This may mean that additional data could lead to more mental workload and thereby a user being able to see less information. However, one experiment evaluating user performance on text-based tasks found that visual separation only hurt performance when combined with an offset in depth [131]. That study maintained a constant amount of information between conditions and consisted of two monitors placed beside each other.

2.3. Information Visualizations

2.3.1. Scalability

Information visualization seems to be one of the best applications for new display technologies since the size of many datasets make scalability an important issue. The number of display pixels places a fundamental upper bound on the amount of information simultaneously visible in a visualization. Most current visualizations can only handle around 10,000 items when shown on a 1600×1200 desktop monitor [53]. According to Huber’s taxonomy of large datasets [63], 10,000 bits of information is classified as a small dataset. This means that most existing information visualization techniques must use aggregation, elimination, or some type of virtual navigation strategy (i.e. zooming, panning, etc.) when visualizing larger datasets.

Despite the importance of scalability to information visualization, analysis of the visual scalability of techniques is sparse, and the research on increasing scalability mainly deals with scalability for a desktop monitor. While Eick and Karr defined the problem of visual scalability for information visualizations, their emphasis was on increasing visual scalability using improved visual representations and interaction when using a single monitor [49]. Likewise, research into visual representations for displaying a million items on standard PC monitors is also ongoing [53]. However, little research has been conducted on how effective different information visualization techniques are when they are scaled-up using large, high resolution displays.

2.3.2. Integrated and Multiple Views

As mentioned in the introduction, the attribute and space-centric visualizations are specific types of multiple and integrated view visualizations respectively. These designs were briefly introduced in Section 1.2 and will be explained further in Chapter 3.

In general, the integrated view strategy seeks to represent all of the information in a single integrated view. It emphasizes the advantages of layering information and integrating information [28, 30, 136]. For example, geographic information systems often overlay spatially correlated information, enabling users to quickly see spatial relationships [146]. Navigation in large and complex data is typically accomplished with zooming [19], fisheye, or other distortion-based techniques [74].

The multiple view strategy splits information into multiple simplified views. It emphasizes the advantages of segmenting complex information into simpler parts [88, 109]. Basic low-level theories and experimental findings in data graphics and visual

perception [22, 38, 139] have enabled the development of high-level guidelines and even automated systems (e.g. [23, 60, 83, 112, 113]) for designing basic integrated views on desktop displays. The focus of this work is on the comparison between integrated and multiple view visualizations across display sizes and for multiple data types, areas which have not been explored to the same extent.

Recall that the data of interest is a combination of spatial and multidimensional data. This combination of data includes a natural spatial structure. However, even without a spatial context, multidimensional data visualization alone is challenging and techniques have been presented for visualizing this type of data. Examples of multidimensional visualization techniques include Parallel Coordinate Plots [65], Pixel Bar Charts [55], and tools such as Table Lens [107]. These techniques must be considered because as mentioned earlier, after approximately 8 attributes the complex glyphs approach fails to create perceptually salient visual encodings. Therefore, these multidimensional visualizations can be used either as a separate linked visualization or as embedded visualizations with a space-centric visualization (again, one type of integrated view visualization).

2.3.3. Linked Visualizations

Another approach, mentioned because of the consideration it was given in the early stages of this research, is to link together distinct visualizations that each provides an alternative perspective of the data. We will refer to this approach simply as linked visualizations. If linked visualizations are used, then the visualizations can be coordinated in a variety of ways to reveal relationships and support navigation [97, 98]. As an example, DataMaps [77] was created for the U.S. Census Bureau. It displays U.S. Census data in several different visualizations including maps, scatter plots, and tables. A variety of coordination strategies are employed, including brushing, dynamic queries, and overview + detail. There are few specific design guidelines and little empirical research on linked visualizations.

Relevant to linked visualizations and also to the comparison between integrated and multiple view visualizations is the work of Baldonado et al. [11]. They provide some heuristics for deciding when it is advantageous to split data between linked visualizations. Their guidelines suggest using linked visualizations only when the *screen space* and computation resource costs are justified and if diversity exists, correlations or disparities can be brought out with visual comparison, and for decomposing a complex view into less cognitively demanding pieces. The major cited downsides are the increased screen space, the cost of learning, and cognitive overhead. However, research has shown effective user performance with different methods of linking visualizations ([77]), and different combinations of visualizations for the exploration of multivariate health data ([48]) on desktop monitors. There is also various related work in geographic information systems suggesting the real world usefulness of linking visualizations on desktop displays [46-48, 56, 82].

If a designer chooses to use linked visualizations, they then must make the decision of whether to present them simultaneously or sequentially. Trafton et al. have studied

distinct visualizations presented to Navy meteorologists sequentially and suggested that there may be advantages to presenting the information concurrently [135]. This leads one to believe that more pixels would prove to be quite advantageous for linked visualizations.

In general, using dynamic interactive graphics has much potential for the data of interest [46]. Researchers have outlined issues related to geospatially-referenced time-series data [6, 7, 81] and also discussed the importance of understanding cognitive issues with geographic visualizations [78, 86, 123]. Examples of well known tools in this area include CommonGIS [8], which seeks to make visualizations available online to the common person, GeoVISTA Studio [128] which allows for the complex linking of visualizations, and space-time cubes [73] which use a 3D cube and represent time as the third dimension.

2.4. Information Visualization Design Spaces

The structuring of information visualization knowledge into a common framework allows for knowledge reuse and generates new design ideas. The importance of this structuring has resulted in various information visualization design spaces and taxonomies over the last two decades. Among the earliest of this work was work done by Bertin who outlined seven visual variables (value, hue, texture, shape, position, orientation, size) and how those visual variables were linked to data attributes (i.e. shape should be used for nominal attributes and not ordinal attributes) [21]. This coincided with the work of Tufte [136] and Cleveland and McGill [38]. Card and Mackinlay extended Bertin's work on the semiotics of graphics to structure their design space [29]. In the same spirit as Bertin, they mapped data variables to visual variables. The results were organized based on domain (Scientific Visualization, GIS, multi-dimensional plots, multi-dimensional tables, information landscapes and spaces, node and link, trees, text) and demonstrated the differences in mappings.

Design spaces also emerged that accounted for interactive information visualizations. Shneiderman's well known task by data type taxonomy outlined high level tasks (overview, zoom, filter, details-on-demand, relate, history, and extracts) and data types (1D, 2D, 3D, temporal, multidimensional, tree, and network) [121]. Others characterized systems using the decomposition of basic visualization interactions [37], or using a combination of data, representation, and interactivity [137]. A taxonomy [99] and a model [25] have also been created that explicitly consider the interaction between windows within multiple view visualizations.

In 1998, Chi and Riedl introduced a framework based on the data state reference model to aid in the implementation and reuse of visualizations [35]. This model consisted of data stages (values, analytical abstraction, visualization abstraction, view) and transformation operations (data transformation, visualization transformation, visual mapping transformation). Chi later created an extensive taxonomy of visualizations using this model [36].

Tory and Moeller's model-based visualization taxonomy uses a model rather than a data-centric approach [134]. This approach clarifies the distinction between scientific visualization and information visualization by classifying a visualization based on the way it is being used (continuous vs. discrete). A unified taxonomic framework has also been proposed that dissects information visualization into five design factors (data, task, interactivity, skill level, and context), input (tool, device), visualization approach (display: dimensions, dynamics, animation), and dimensional overload [101].

Our design space is different because we specifically consider the visual integration of heterogeneous data types (spatial and multidimensional). Other design spaces in this area have been for associating time-series data with geography (single static maps, multiple static maps, single dynamic maps, and multiple dynamic maps) [91], time-series data with graphs (with the major designs being animation, small multiples, and nested visualization) [117], and associating annotations (abstract data) with referents (spatial locations) in information rich virtual environments [105]. As is evident by the given examples, this meeting of heterogeneous data types occurs across domains. Based on the fact that each structure provides a spatial context, we combine Shneiderman's 2D, 3D, tree, and network data types into a single spatial data type. The multidimensional data is then associated with that spatial context.

2.5. Summary

This work builds on and extends previous work in a variety of ways. This is the first work to explore the connection between information visualization and the display. Previous work either focused on comparing visualizations or on aspects of the display, but not a combination of the two. Additionally, in considering the link between the visualization and the display, this work also explores the relationship between human abilities and using more pixels to scale-up visualizations, hence exploring the limits of human perceptual and cognitive abilities with respect to scaling visualizations. The design space fills a gap by specifically considering visualizations that integrate heterogeneous data types. The tradeoffs within the design space are explored and are used to demonstrate the link between visualization and user task performance.

Chapter 3. Design Space

The first research question was: what is the visualization design space for integrating spatial and multidimensional data? In this section, a unifying hierarchical design space is presented as a framework for visually integrating heterogeneous data types. Examples of visualization techniques in different fields are provided to show the commonality across fields. After discussing the basic premise, different ways that the design space can be extended are described. These extensions include adding an additional data type such as temporal data (adding a level to the hierarchy) and constructing linked visualizations (linked hierarchies).

3.1. Hierarchical Design Space

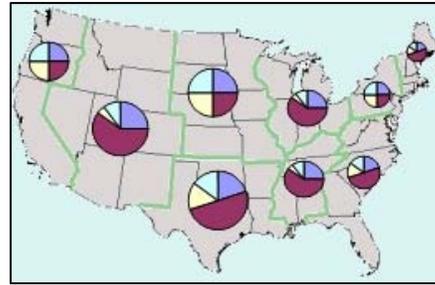
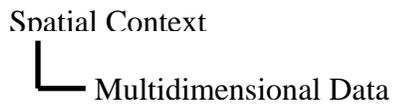
The building blocks of a hierarchical representation are data types and levels. The basic premise is that the closer the visual proximity of a data type's values, the higher that data type is in the hierarchy. A hierarchy can be determined visually or by using the size of constants from the algorithm that determines the spatial layout of the displayed values. The hierarchical representation concept can be demonstrated using the visual integration of two data types: spatial and multidimensional. Assuming for the moment that each data type resides at a different level, there are two possibilities: 1) space is at the top level of the hierarchy or 2) multidimensional data is at the top level of the hierarchy.

If space is at the top level of the hierarchy, we refer to the design as **space-centric**. In this design the values for all dimensions are visually grouped by location. This means that all values for a specific location are in close visual proximity. Examples of the space-centric visualization in different domains and with different spatial contexts can be seen in Figure 3.1.

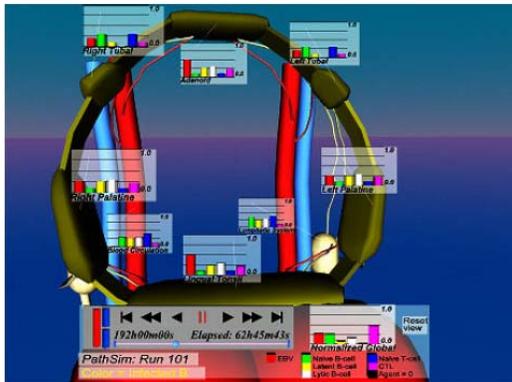
If multidimensional data is at the top level of the hierarchy, then we refer to the design as **attribute-centric**. Attributes (dimensions) are central because the values are visually grouped by attribute. The result is that all values for a specific attribute are in close visual proximity, but multiple attributes referencing the same spatial location are visually distant. This design often corresponds with Tufte's concept of small multiples [136]. Examples of the attribute-centric design are shown in Figure 3.2.

So far, only two basic possibilities have been discussed: the space-centric visualization and the attribute-centric visualization. In practice there is actually a continuum between these approaches based on the distribution of attributes. This continuum can be clarified by asking the question, *how many attributes are overlaid on a single spatial structure?* The answer is all attributes for a space-centric visualization and one attribute for an attribute-centric visualization. If the answer is somewhere in between then we refer to this as a **mixed design**. The distribution of attributes amongst views does not necessarily have to be even. Four attributes can be split two per structure or three and one. The tradeoffs with this design are a mixture of the tradeoffs with the space-centric and attribute-centric designs. In the next subsection we will discuss how the design space expands when an additional data type is added and how the design space accounts for linked visualizations.

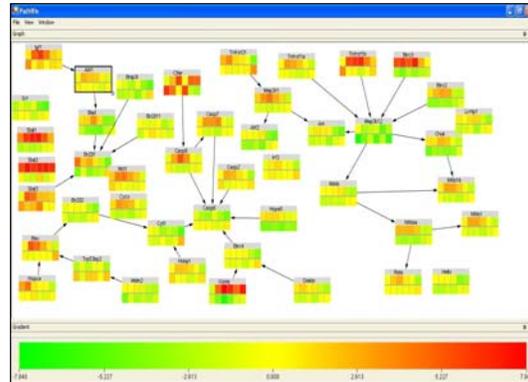
SPACE-CENTRIC DESIGN HIERARCHY



(a) 2D spatial context (pie charts on map)



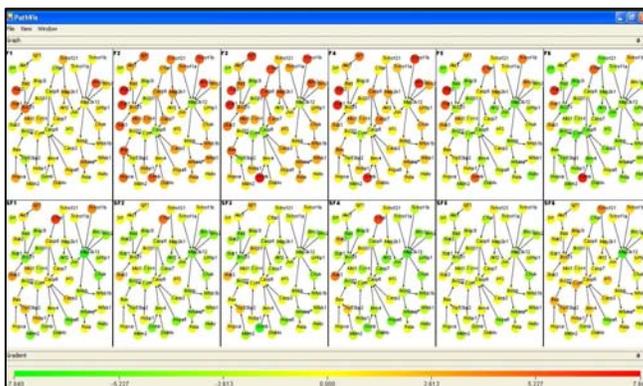
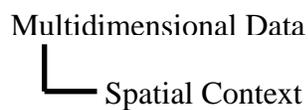
(b) 3D spatial context (pathogens in mouth) [103]



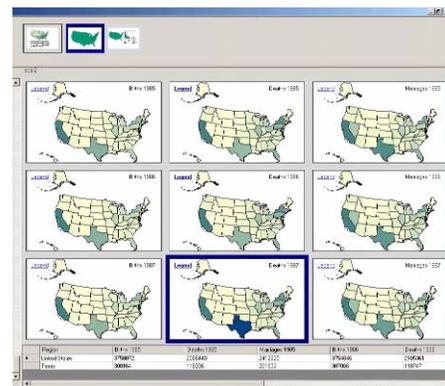
(c) Graph context (gene expression levels and biological pathway) [119]

Figure 3.1. Space-centric visualizations with different spatial contexts

ATTRIBUTE-CENTRIC DESIGN HIERARCHY



(a) Graph context (gene expression levels and biological pathway) [119]



(b) 2D spatial context (demographic data on maps, US Census Bureau)

Figure 3.2. Attribute-centric visualizations having different spatial contexts

3.2. Extending the Design Space

Shneiderman’s data types included 1D, 2D, 3D, tree, network, multidimensional, and temporal data [121]. If 1D, 2D, 3D, tree, and network data are grouped into a larger spatial data type then only multidimensional and temporal data remain. While this design space was originally created as a basic model for visually integrating spatial and multidimensional data, it can easily be expanded to spatial, multidimensional, and temporal data.

Including a third data type such as time requires integrating three data types rather than two. The design space accommodates this situation by including an additional level in the hierarchy. To simplify the design space we do not specifically consider the visual integration of multidimensional and temporal data alone. Instead, we focus on the location of the spatial data type in the hierarchy. This creates four possibilities:

1. **Space-centric:** Space is at the highest level (Figure 3.3)
2. **Attribute/Time-centric:** Space is at the lowest level (Figure 3.4)
3. **Attribute-centric:** Space is in the middle with multidimensional at the highest level (Figure 3.5)
4. **Time-centric:** Space is in the middle with temporal data at the highest level (Figure 3.6)

SPACE-CENTRIC

Spatial Context

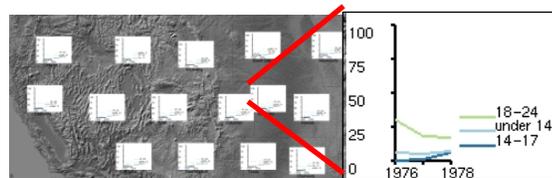
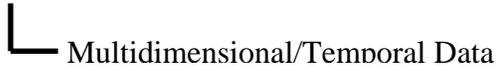


Figure 3.3. Space-centric visualization with 3 data types (time-series graphs at each location)

ATTRIBUTE/TIME-CENTRIC

Multidimensional/Temporal Data

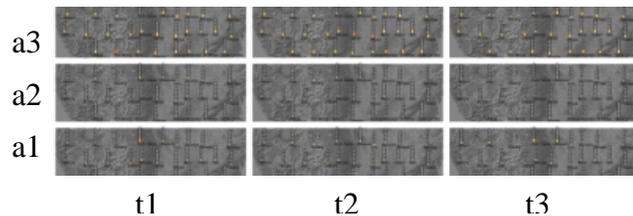
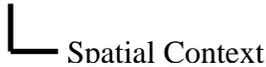


Figure 3.4. Attribute/Time-centric design with 3 data types: 2D map for each time/attribute pair

ATTRIBUTE-CENTRIC

Multidimensional Data

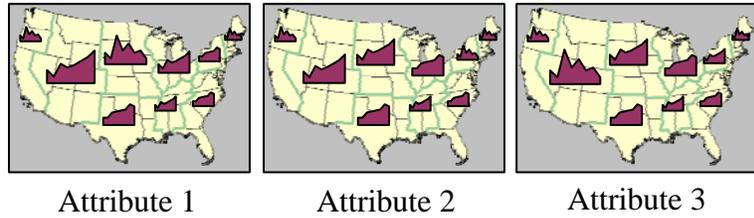
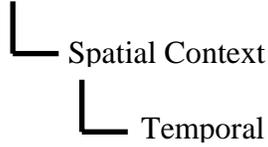


Figure 3.5. Attribute-centric design with 3 data types: Separate map for every attribute with value-flow maps ([7]) overlaid

TIME-CENTRIC

Temporal Data

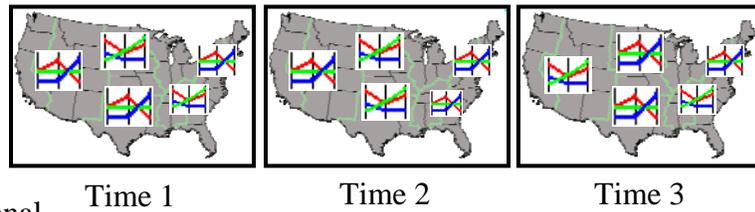
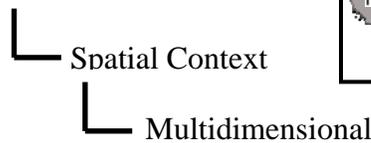


Figure 3.6. Time-centric design with 3 data types: Separate map for every time point with parallel coordinates ([65]) overlaid

One might wonder why animation is not included in these representations since it is often used for visualizing temporal data. For the sake of this work we consider animation a method of filtering data. With animation the time points not shown have essentially been filtered out using time. Because we do not consider filtering or aggregation and because we only consider complete datasets, animation has no distinct representation in this design space. Next we will briefly discuss expanding this design space to linked visualizations.

Linked visualizations create sibling relationships between linked but distinct visualizations. The positive aspects of one visualization then offset the negative aspects of another. The design space can be expanded for linked visualizations by considering each visualization first as an independent hierarchy. Linked visualizations are then simply multiple linked hierarchies (Figure 3.7). Visualizations can then be combined based on the order of data types within the hierarchy, thereby matching the weakness of one visualization with the strengths of another to create a better total result.

LINKED HIERARCHIES

Spatial Context + Multidimensional Data

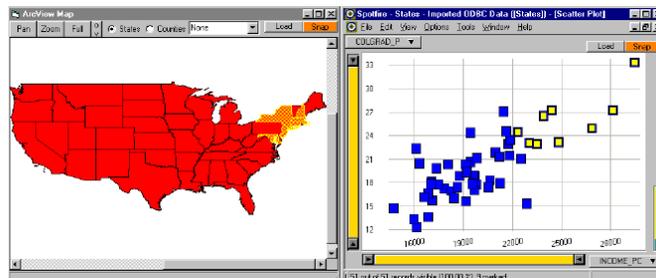


Figure 3.7. Multiple linked hierarchies correspond to linked visualizations (image shows Snap software in use at the Census Bureau [96])

3.3. Definitions

Following are definitions for some of the main terminology that is used throughout this document.

Space-centric visualization: A visualization where space is at the top level of the hierarchy. The integration of spatial and multidimensional data distinguishes this design from other integrated view visualizations that only display a single type of data (or some combination that is not spatial and multidimensional).

Space-centric using complex glyphs: A space-centric visualization where the encoding of information within a spatial location is accomplished using complex glyphs. A complex glyph uses multiple visual attributes such as size, shape, and color to represent data attributes.

Space-centric using embedded visualizations: A space-centric visualization where the encoding of information within a spatial location is accomplished using embedded visualizations. This approach allows for consistency in using the same encoding for all the attribute values. Although embedded visualizations could be considered a type of complex glyph, they are distinguished from complex glyphs because each can exist independently.

Attribute-centric visualization: A visualization where space is at the bottom level of the hierarchy. This language is slightly different than that used in the design space in that it includes the attribute-centric design with two data types and also the attribute/time-centric design with three data types. The integration of spatial and multidimensional data distinguishes this design from those that decompose information into simpler pieces but only display a single type of data (or some combination that is not spatial and multidimensional).

Linked visualizations: Linked visualizations correspond to multiple linked hierarchies where each hierarchy is a distinct visualization. This approach provides multiple alternative perspectives of the data using multiple linked visualizations. An individual visualization does not necessarily include heterogeneous data types. For example, a multidimensional visualization alone can be an independent hierarchy.

Integrated views: This is a loaded term that is used mostly for the sake of readability. This is a term that is commonly used in literature but is vague and therefore the more specific terms above were created. It includes all space-centric visualizations. It also includes all visualizations that visually integrated all data. For example, a 3D scatter plot is visually integrated.

Multiple views: Like integrated views, this is a vague term commonly mentioned in literature that is used mostly for the sake of readability. It includes all attribute-centric visualizations, linked visualizations, and visualizations that are visually decomposed into basic parts.

Chapter 4. Desktop Display Baseline

The second research question was: as a baseline, which visualization results in the best performance for specific tasks on a desktop display? Three experiments were conducted related to this question (published as [40, 114, 148, 150]). These experiments compared complex glyphs in a single view (space-centric visualization using complex glyphs) and simple glyphs in multiple views (attribute-centric visualization), probed the cognitive issues with using linked visualizations, and considered visual versus interactive integration of information.

4.1. Experiment 1: Multiple Simple vs. Single Complex Glyphs

The goal of this experiment was to compare complex glyphs in a single view (space-centric visualization) and simple glyphs in multiple views (attribute-centric visualization). This comparison was done in the context of geospatially-referenced multidimensional data. Simple glyph-based encodings [58] were used for the space-centric visualization. Data points were mapped to visual glyphs and overlaid on a geographic map. The primary tradeoff when comparing the attribute-centric approach to the space-centric with complex glyphs approach is as follows: In a larger space-centric visualization multiple data attributes are encoded using complex glyphs with multiple visual features (such as color, size, and orientation) which could potentially interfere with one another. With an attribute-centric visualization each data attribute is represented in a separate smaller view, using simple glyphs with one most-salient visual feature, as in Tufte’s “small multiples” [10]. The complexity tradeoff is between the number of visual features of glyphs and the number of views. Presumably, both should be minimized to reduce perceptual and cognitive complexity. We wanted to know which visualization approach was best on a desktop display for different types of tasks.

4.1.1. Method

The experiment consisted of a 3x2x4 design with 3 independent variables: number of views, task attributes’ visual encoding, and task type. It was a mixed design. Number of views was between-subjects, while the others were within-subject. The experiment measured three dependent variables. First, user performance time was measured for each task. Second, answer correctness was recorded. The last measure was taken on a post-questionnaire and was a subjective measure of how well the participant believed the interface supported the different types of tasks.

Number of Views (Visualizations)

The number-of-views variable had 3 conditions: one (space-centric), two (mixed), and four (attribute-centric) views were compared (see Figure 4.1, Figure 4.2, and Figure 4.3). The two-view case was an intermediate solution between the attribute and space-centric visualizations (i.e. a mixed design). In each condition, four abstract data attributes were shown: attributes A, B, C, and D. The actual data used was discrete and a modified version of data obtained from the US Census.

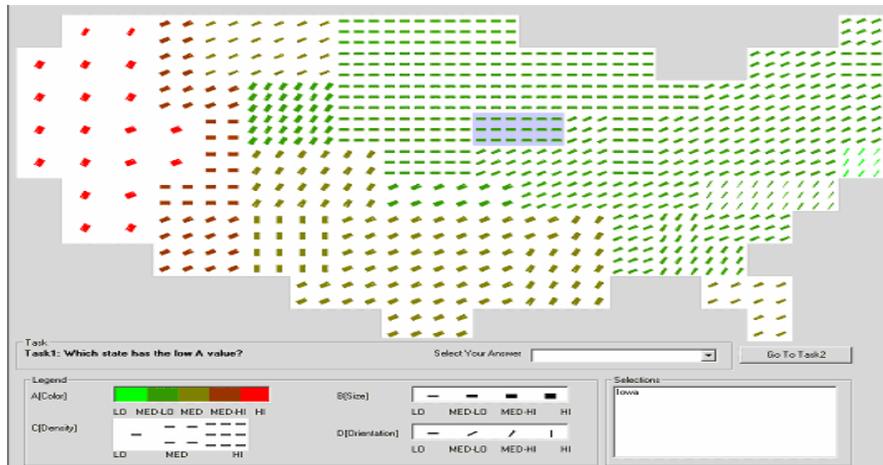


Figure 4.1. One view (space-centric), data attributes mapped to color, size, density, and orientation



Figure 4.2. Two views (mixed design), data attributes mapped to color and size on two maps

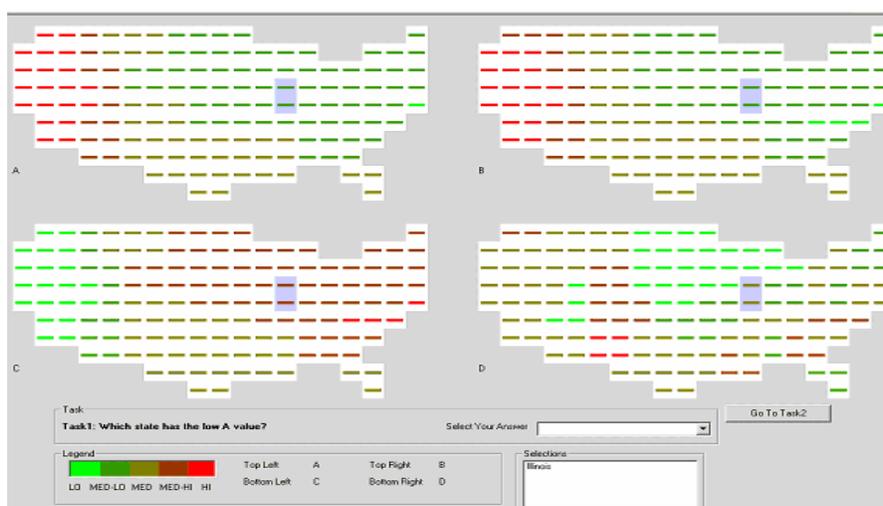


Figure 4.3. Four views (attribute-centric), data attributes on their own map with color used in all

Visual Encoding

Data attributes were distributed equally among views, and each visual encoding was used only once per view as shown in Table 4.1. The visual mappings used for representing these attributes had already been proven effective in [60]. Attribute A was considered the most important and mapped to the best encoding (color [4]), and so on with D the least important. For simplicity in reporting results, we refer to A as having a ‘better encoding’ than B. In all conditions, total screen space used was held constant.

Table 4.1. Visual encodings used in each visualization condition

		Data Attribute			
#		A	B	C	D
V i e w s	1	Color	Size	Density	Orientation
	2	Color (left)	Color (right)	Size (left)	Size (right)
	4	Color (top left)	Color (top right)	Color (low left)	Color (low right)

Tasks

Four different tasks were completed: target detection with 1 attribute and with 2 attributes, and trend finding with 1 attribute and with 2 attributes. These tasks can be seen in the rows of Table 4.2. Each task also had 2 variations (visual-encoding variable, columns of Table 4.2): the first asks about the most important attributes (A and B), while the second includes a less important attribute (C or D). The purpose of this variable was to explore the effect of the tradeoff between glyphs features and views. The most likely effected data attributes are those that are of less importance, and hence mapped to less effective visual encodings in the case of space-centric visualization.

Table 4.2. Tasks and visual encodings

	Most important attribute, Best visual encoding	Less important attribute, Not best visual encoding
Detect: One attribute	Which state has the lowest A value?	Which state has the highest D value?
Detect: Two attributes	Which state has the medium low A value and the low B value?	Which state has the medium low A value and the medium C value?
Trend: One attribute	What’s the trend from West to East in terms of A value?	What’s the trend from North to South in terms of D value?
Trend: Two attributes	What’s the relationship between A and B?	What’s the relationship between A and C?

Procedure

The fifty-seven participants in this study were engineering students from a large public university. Participants performed practice tasks for detecting one and two targets.

During the practice, the administrator explained the various features in the interface. There was no time limit for each task. At the end of the experiment, participants were given a post-questionnaire to rate the interface.

4.1.2. Results

Since it was not of interest to compare between the 4 tasks, each was treated as a separate 3x2 analysis. Two-way ANOVAs were performed for each task (each row in Table 4.2) on the other 2 factors: number of views, visual encoding.

Single Attribute Detection Tasks

For the tasks that involved detecting one data attribute there was a significant effect of number of views, $F(2, 54) = 12.06, p < .01$, visual encoding, $F(1, 54) = 14.06, p < .01$, and an interaction effect between them, $F(2, 54) = 11.85, p < .01$. Detecting the attributes that had the better visual encoding was significantly faster.

Tukey's HSD for the number of views indicated that participants that used one view took significantly less time than those using two views, and those using four views took significantly less time than those using two views (both $p < .01$). There was not a significant difference between one and four views.

For the interaction effect between number of views and visual encoding, task completion times were significantly faster when using two views with the best visual encoding than when using two views with the less perceptually salient encoding ($p < .01$). When detecting a single data attribute with the less perceptually salient visual encoding, task completion times were significantly faster when using one view compared to two views, and also when using four views compared to two views ($p < .01$).

There was no significant difference in correctness. In terms of satisfaction, ANOVA indicated that there was a significant effect of the number of views, $F(2, 54) = 3.57, p < .05$ for these tasks. Tukey's HSD for the number of views indicated that participants that used two views were significantly less satisfied than those using four views ($p < .05$).

Dual Attribute Detection Tasks

For detecting two data attributes there was a significant effect of visual encoding, $F(1, 54) = 4.83, p < .05$. Detecting the two attributes with better visual encodings took significantly less time. Figure 4.4 shows the mean completion times for tasks that involved detecting two attributes. There was not a significant effect for the number of views or their interaction. There was no significant difference in terms of correctness or satisfaction.

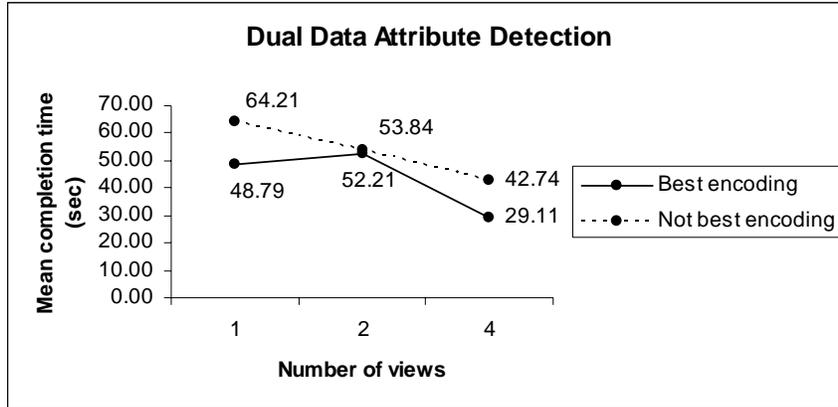


Figure 4.4. Completion times for detecting two attributes

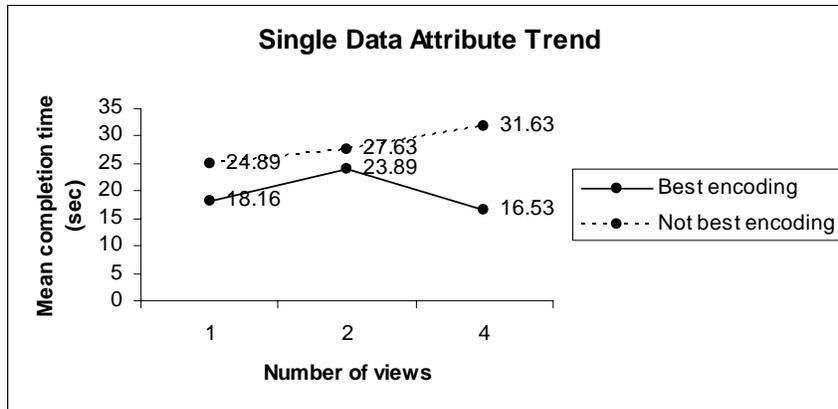


Figure 4.5. Completion times for one-attribute trend task

Trend Finding Tasks

For the tasks that involved finding a trend for a single attribute there was not a significant effect by number of views, but there was by visual encoding $F(1, 54) = 23.94, p < .01$ and for their interaction $F(2, 54) = 3.81, p < .05$. Finding the trend for the attribute that was less perceptually salient took significantly longer. Tukey's HSD for the interaction indicated that, when using one view, finding the trend with the best visual encoding was significantly faster than finding the trend with a less perceptually salient visual encoding. The same was true for four views ($p < .05$). The mean completion times for this task can be seen in Figure 4.5. There was no statistically significant difference between views in terms of correctness or satisfaction for this task.

For the task that involved finding a relationship trend between two data attributes, there was no significant difference in terms of time, correctness, or satisfaction.

4.1.3. Discussion

The decision of how many views to use when designing a visualization involves tradeoffs. Using more views increases the complexity of the interface, but allows more perceptually salient visual encodings to be reused. When using the space-centric visualization with complex glyphs the problem may also arise that all reasonable visual encodings have been exhausted.

Target Detection Tasks

The results of both the single and dual attribute detection tasks indicate that it is faster to find the attribute with the best visual encoding. What is initially surprising is that, in the single attribute detection task, the mixed design seemed to perform worse than the space and attribute-centric visualizations. This suggests that the visual encoding of size may have been harder to detect than either color or, in particular, orientation. It is also possible that the mixed design condition performed worse because it was harder for participants to understand the legend. In the space-centric visualization all attributes are represented by different encodings in a single map. In the attribute-centric visualization every map represented a different attribute. The mixed design was a combination of both strategies and hence required users to decode both methods. Particularly with the mixed design, legends must be carefully designed.

There was no significant difference when color was used in each, but there was a significant difference when different visual encodings were used for the same task. Therefore, it appears that visual encoding has more of an impact on user performance than the number of views and any visual interference due to extra visual encodings. The lesson learned from this is that it is faster to use the best visual encoding for the data attributes of interest. Either the visualization can be designed with the ability to change the visual encodings of attributes as the user's interest changes, or the attribute-centric visualization can be used because of the ability to reuse the best visual encoding for several attributes.

Trend Finding Tasks

The results of the trend finding task on one attribute indicated that using a better visual encoding was faster. Although this is expected when using a single view with different encodings, it is not with the attribute-centric visualization. Aside from the trend direction and data, the only difference was the location of the target view in the 4-view grid. It is possible that this task took longer because participants became confused with which map to consider. This was only true when the trend for one attribute needed to be found. The lesson learned from this is that if the attribute-centric visualization is used there is the potential for confusion as more views are added, particularly when the task becomes more cognitively demanding. To avoid this problem the interface could be designed so that views can be rearranged or highlighted. At a minimum, the interface should be clearly labeled regarding which attributes are represented in each view.

Exploratory Results

In the dual attribute detection task, the best visual encoding condition involved finding color and size in a single view (space-centric), or color in two different views (mixed and attribute-centric). The trend in the results indicates that using the attribute-centric visualization took less time than using the mixed design. At first it seems these should be close to equal considering they both involve finding color in two different views. The difference between these is that in the mixed design there are extra visual encodings, and in the attribute-centric visualization there are extra views. Therefore, it appears that it is easier to identify colors in two views if no extra visual encodings are present in those same views.

Finding the attributes with the less perceptually salient encodings involved finding color and density in a single view (integrated), color and size in a single view (dual), or color in two different views. The trend indicates that the attribute-centric visualization was faster than the mixed design which was faster than the space-centric visualization. This means that using color in two different views was faster than using either color and orientation or color and size in a single view. It appears that being able to use the two best visual encodings in different views was faster than showing both combined when one of the data attributes had a visual encoding that was less salient.

4.1.4. Summary

This experiment was an empirical comparison of complex glyphs in a single view (space-centric visualization) and simple glyphs in multiple views (attribute-centric visualization) [148]. Four attributes in a single view were compared to simple glyphs representing four attributes in four different views. A dual-view (mixed design) situation was also considered. Participants performed target detection and trend finding tasks in the context of multidimensional glyphs overlaid onto geographic maps. Results from the target detection tasks suggest that visual encoding is a more important factor when detecting a single attribute than the number of views. Additionally, for detecting two attributes, the trend indicates that reusing the most perceptually salient visual feature with the attribute-centric visualization provides faster performance than the space-centric visualization which must map one of the attributes to a less salient feature.

4.1.5. Conclusion

The main finding was that visual encoding is the most important factor, regardless of number of views. Furthermore, the trend indicates that it is faster to reuse the most perceptually salient visual feature with an attribute-centric visualization than it is to use a space-centric visualization with complex glyphs that requires a less perceptually salient feature, even for two-attribute tasks.

4.2. Experiment 2: Context Switching and Cognitive Factors

A second study related to establishing a baseline comparison on a desktop display probed the cognitive issues with using linked visualizations. We studied the ways in which people integrate data from multiple linked visualizations by investigating how users identify the relationship between data points in different views. The time cost of context switching was measured and the cognitive processes involved in linked visualizations were explored. Because little empirical research has been done on multiple linked visualizations, we began with an exploratory study. The ways that multiple linked visualizations are used, the cognitive abilities involved with their use, and the effect of context switching on performance was observed. The decision was made to collect as much information during the experiment as possible because of its exploratory nature.

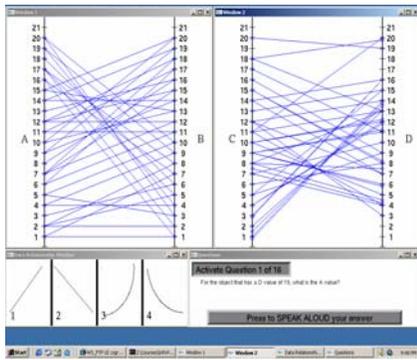
Of particular importance were the decisions to perform cognitive abilities tests and to use the eye-tracking equipment. Wickens and Hollands [143] wrote that even though it's possible for a change in attention to occur without a change in the direction of our visual path, most of the time our visual behaviour is driven by our desire to attend. Studying visual scanning behaviour can provide feedback as to selective attention [111]. Relating this back to linked visualizations, Baldonado et al. [11] point out that there is an increased demand on cognitive attention when using multiple linked visualizations. This suggests that attention is a cognitive function that will influence a user's performance.

Understanding cognitive processes can provide insight into important aspects of designing multiple linked visualization systems since finding relationships in data distributed across visualizations can be a difficult and challenging task. The use of multiple visualizations increases the demand on cognitive attention since a user must make use of numerous perceptual cues simultaneously, thus increasing cognitive load. Aside from increasing cognitive load, linked visualizations occupy more space and require the learning of additional constraints. Under such circumstances, context switching becomes an important concern. As a result, we were interested in measuring the cognitive aspects, especially the time it takes for context switching, involved in such systems.

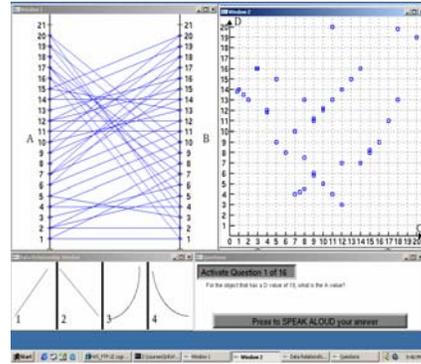
4.2.1. Method

Design

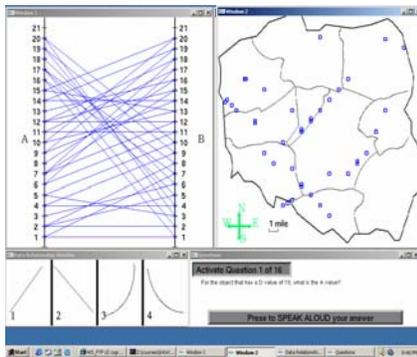
We used a 4x4 factorial, within-subject design. The independent variables in the experiment were the level of context switching and the task difficulty. The level of context switching was varied using increasingly different combinations of visualizations. The given combinations were two parallel coordinate plots (PP), a parallel coordinate plot and a scatter plot (PS), a parallel coordinate plot and a geographic map (PG), and a scatter plot and map (SG). Examples of these combinations can be seen in Figure 4.6.



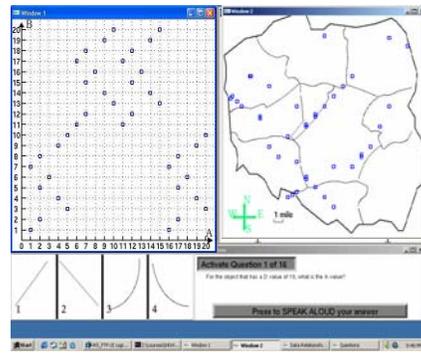
(a) Two parallel coordinate plots (PP)



(b) Parallel coordinates and scatter plot (PS)



(c) Parallel coordinates and geographic map (PG)

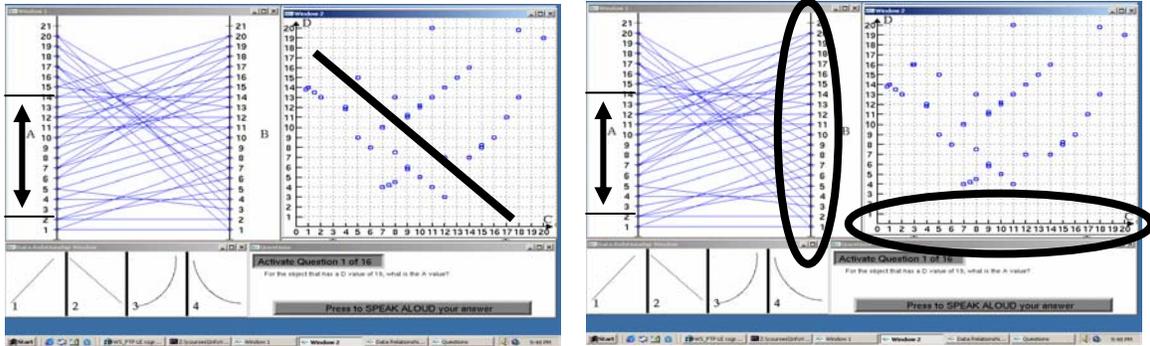


(d) Scatter plot and geographic map (SG)

Figure 4.6. Combinations of linked visualizations

A parallel coordinate plot, a scatter plot, and a geographic map were chosen because of their use in commercial software systems, and their proven usefulness in the field of information visualization. The advantages of using parallel coordinate plots are outlined in [65], and some commercial tools such as Spotfire [4] use an interactive scatter plot as the basic visualization. In a parallel coordinate plot, the axes of a multidimensional space are defined as parallel vertical lines separated by a specified distance. A point in Cartesian coordinates corresponds to a line in parallel coordinates. This visualization was kept consistent throughout three of the combinations in order to reduce influence due to lack of previous experience.

The data consisted of forty data points with four attributes (A, B, C, and D). Abstract data was used to avoid the interference of previous knowledge. A and B attributes were shown in the visualization in the left window and C and D in the right window. When asked a pattern recognition question, the relationships were always between two attributes and involved ten data points with one outlier.



(a) Easier pattern recognition task: select a subset in one view; see a pattern in the other

(b) Harder pattern recognition task: select a subset in one view; find a pattern between views

Figure 4.7. Pattern recognition tasks

Task difficulty changed based on whether the task was a search task or a pattern recognition task, and the minimum number of visual switches between the views a user would have to make to answer. For the search tasks a user either had to find one value or compare two values. In the pattern recognition tasks users either had to find a pattern between a visualization or across visualizations. Figure 4.7 highlights the difference between the pattern recognition tasks. To avoid a learning effect based on the order of questions, the order for each subject was balanced using a Latin-square design.

Dependent variables included user answer performance, cognitive ability test scores, eye-tracking measures, and think-aloud results. With respect to user performance, task completion time and task accuracy were measured, perceived difficulty and user preference were obtained, and logged user-interaction measures such as number of times a data point was selected was collected.

One of the design decisions was whether to use abstract data or real data. Burns used data specific to power plant operators and her results suggest a single integrated view supports performance better than multiple views for that domain [28]. Her explanation for this is that integrated view visualizations are able to show the data in a meaningful way. In that scenario, showing the data in a meaningful way is possible based on domain knowledge. In contraposition with that logic, the explicit purpose of our study was to avoid the influence of factors related to specific domains and focus on more general information visualization issues. The result of that decision is that our results do not apply specifically to any domain. However, they are important during the process of designing a system for any domain, and particularly for situations where linked visualizations are used for exploratory visualization.

User Interface

Instead of using any specific research system or commercial tool, in an attempt to create a situation with maximum experimental control and avoid unchangeable user interface factors, a system was created specifically for this study. The four main features of the system design were generic visualizations (parallel coordinates, scatter plot, and geographic map), abstract data, brushing and linking used for interaction, and an

interaction log kept by the system. In addition to the visualizations, choices of patterns for pattern recognition tasks were displayed in the lower left part of the screen, and the question was shown in the lower right (these features can be seen in Figure 4.6).

The visualization system was implemented in OpenGL and supported the ability to select a single point or to drag and select multiple points. Pairs of visualizations displaying different attributes of the same data set were presented. Using the brushing and linking interaction technique, if a point was selected in one view, the corresponding point was also highlighted in the other view. The system was able to record the history of the users' actions and the corresponding completion times.

Participants

The design was a 4x4 factorial, within-subject design involving sixteen undergraduate computer science students from a human-computer interaction class. Fourteen males and two females between the ages of 18 and 25 participated. All participants received extra credit for their time. This group was chosen because during their previous class work they had been exposed to all types of visualizations used in the experiment.

Although it can be argued that students are not a diverse enough group to be representative of the population and that they are not the true end users, we believe they were sufficient for this study. Because users of exploratory visualization tools deal with complex data sets, it is likely they have a college education. The fact that all participants were motivated enough to be enrolled at a university means they are a group of potential future users. Although the testing of tools designed for specific domains should be tested by users with domain knowledge, the data used in this study was abstract. Therefore the lack of domain knowledge was irrelevant because there was no specified domain.

Procedure

After arriving and signing the consent forms, subjects provided demographic information which included their familiarity with different visualizations, their experience using computers, and other factors. Then, the Raven cognitive ability test was performed. The Raven progressive matrices test measures the ability to form perceptual relations and to reason by analogy [31]. Three subtests of the Weschler Adult Intelligence Scale Revised (WAIS-R), including the Digit Symbol, the Picture Completion, and the Digit Span Verbal subtests were used [132]. Each of the subtests gives an index about some specific cognitive functions. The Digit Symbol performance subtest involves visual-motor coordination and speed. The Picture Completion performance subtest involves visual recognition, general information, and focusing attention. Finally, the Digit Span verbal subtest involves auditory attention, concentration, and short-term memory.

During the second session participants were given instructions, completed two practice tasks, and then performed the sixteen experimental tasks. The practice tasks allowed the subjects to familiarize themselves with the system and the different visualizations. After the practice session concluded, the ISCAN RK-464 eye movement monitoring system was calibrated. This system is used to track eye movements and was recording the eye

position at a rate of 60Hz. Following the calibration, participants proceeded through the experiment, eventually answering a total of 16 questions. To verify that the eye was being tracked appropriately one of the experimenters sat in the observation room watching the output. If the system lost calibration, the system was recalibrated before the start of the next question. Throughout the experiment the participants were videotaped and asked to think aloud. After completing the experimental tasks, they were asked to rate the difficulty on Likert scales of each question type and of using the four pairs of visualizations. Finally, they ranked their preference of combinations and questions and discussed the experiment with the researchers.

4.2.2. Results

This section compiles and expands on results presented in [40] and [114]. It is divided up into user performance results, cognitive ability results, and eye-tracking measures. It is important to note that although the information is divided into these areas, there is overlap throughout.

User Performance

Task Completion Time and Accuracy

Participants took the most time for the situation involving coordinated parallel coordinates. To account for various individual paces, the completion time for each question was divided by the average completion time for a participant. Figure 4.8 shows completion time per question type and task type. Analysis of variance (ANOVA) was conducted and the result showed that the situation with two parallel coordinate plots took significantly longer than the others, $F(3, 45) = 11.32, p < 0.001$.

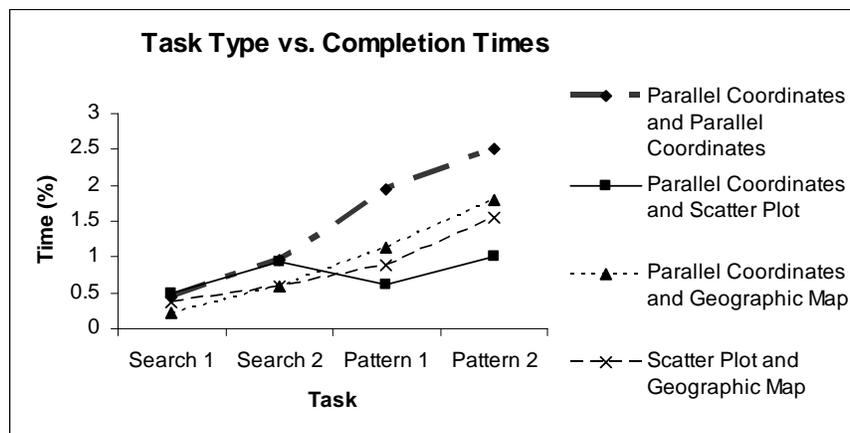


Figure 4.8. Task completion time by task type and combination of visualizations

Participants also were least accurate in the situation involving two parallel coordinate plots. Subjects answered the most questions incorrectly with PP with 28 wrong answers out of 64 questions (16 subjects x 4 questions involving PP), and subjectively rated PP as the most difficult combination to use, but the difference was not significant using an ANOVA, $F(3, 45) = 2.08, p = 0.117$. PS was the best combination with the most correct

answers with 14 wrong answers out of 64 questions and the lowest completion times as shown in Figure 4.8.

After noticing that parallel coordinates coordinated with another parallel coordinates plot took the most time and was the least accurate combination, we thought that it might be a familiarity issue. However, there was no correlation between familiarity with parallel coordinates and completion time ($r = -0.13$). These results were opposite the hypothesis and literature since PP took longer than those presenting data in different manners (or those with greater differences in context).

One confounding factor to these results is the effect of the pattern on participants' ability to recognize it correctly. Participants were more likely to choose the correct pattern if it was a positive relationship than a negative relationship and more likely to choose correctly if it was a linear relationship instead of a non-linear relationship. When factoring out the linearity of the pattern by considering answers wrong only if they were in the wrong direction (e.g. positive instead of negative), the order of combinations with respect to accuracy remained the same, but there was less of a difference. This doesn't affect the significance of the ordering since accuracy was not statistically significantly different in either situation.

We predicted that situations involving two parallel coordinate plots (PP) would take the least amount of time because there would be no cognitive load from context switching. The results were contrary to our hypothesis: questions involving the use of two parallel coordinate plots (PP) actually took significantly longer to complete. This suggests that in some situations, cognitive integration may actually be more difficult when a person is presented with two identical visualizations. However, further research is needed to determine if this is simply a phenomenon that exists when using two parallel coordinate plots.

Perceived Difficulty and User Preference

There was no significant difference between preference rank and difficulty rating based on the combination of visualizations. However, subjects did prefer and therefore ranked the combinations according to their perceived level of difficulty. The difficulty rating was the lowest for the search task that asked the participant to find a value, and the highest for pattern recognition task across visualizations (this confirmed the ordering of task difficulty). Although subjects preferred the questions involving search tasks (finding a single value and comparing two values) to the pattern recognition tasks, they also preferred the questions that required multiple switches between views (comparing two values or finding a pattern across visualizations) to those which required a fewer visual switch (finding a single value or recognizing a pattern within a view) despite their increased level of difficulty.

On the post-experiment questionnaire users were asked their opinions of the combinations of visualizations. Their responses confirmed that they found parallel coordinates the most difficult visualization to use. Five of the subjects commented that the map was the easiest type of visualization to use, while two commented that it was

harder. Four of the subjects commented that relationships were harder to discover using parallel coordinates and that in general parallel coordinates were harder to use. Three subjects felt the opposite about parallel coordinates, preferring them to the other types of visualizations and ranking them the easiest to use. One usability issue with parallel coordinates is that because the lines cross, it can be hard to select multiple edges. Another issue three subjects mentioned was the difficulty in keeping their heads still for the duration of the experiment (a necessity because of the eye-tracking equipment).

From the think-aloud data, it appears that when users were finding patterns across visualizations and had to select a range of A values and find a relationship between B and C values, much of the user's focus was directed toward selected the range of A values. Some users commented that finding patterns between attributes presented in different views was very challenging, and one said she did not know how to do it. As can often be the problem with the think aloud method many users would only talk when prompted by the experimenter then stop after a short time.

Logged User-Interaction Measures

There were three types of mouse actions recorded by the system: the number of clicks selecting data points, the number of clicks on empty space (to clear all selections), and the number of clicks initiating multiple selections. Using the number of edges selected as a dependent measure, both visualization type and question type are significant using an ANOVA for analysis, $F(3, 45)=3.49, p=0.023$ and $F(3, 45)=7.28, p<0.001$, respectively. It is unsurprising that users selected more data points for more difficult tasks. Also unsurprising was the correlation between clicks and raw completion time ($r = 0.66$). The combination with the largest total number of data points selected was the parallel coordinate plots combination (860 total), followed by parallel coordinate plots and geographic map (595), parallel coordinate plot and scatter plot (408), and scatter plot and geographic map (337).

For the number of times a user clicked empty space (meaning they either wished to clear their selected edges or they accidentally missed selecting an edge), both visualization type and question type are significant ($F(3, 45)=3.01, p=0.040$ and $F(3, 45)=11.00, p<0.001$ respectively). However, when considering the number of times multiple selection was used, question type was the only significant factor, $F(3, 45)=8.03, p<0.001$. One reason multiple selection may not have been significant across combinations is that the total number of times it was initiated overall was very low (once initialized users can drag up and down and don't have the need to restart the process).

Cognitive Abilities

From the WAIS-R test, the picture completion subtest score measures focusing attention and visual recognition. This score correlated with task completion times ($r=-0.46$) and accuracy ($r=-0.4$). The correlation with task completion time can be seen in Figure 4.9 (note that higher picture completion scores are better). This means that subjects with better visual recognition and focusing attention completed the tasks faster.

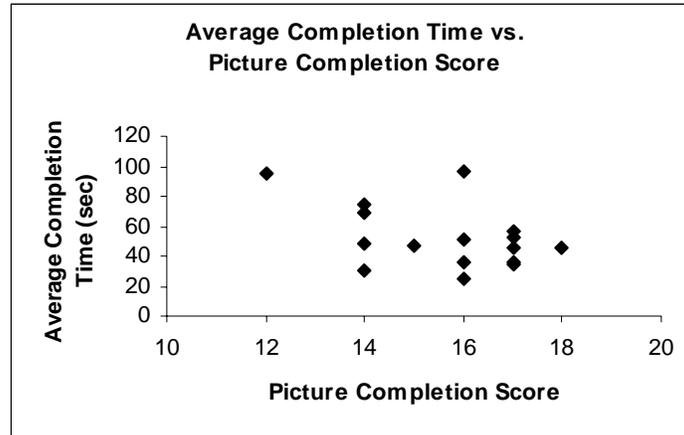


Figure 4.9. Completion time and picture completion subtest scores

The Raven Progressive Matrices score measures the ability to form perceptual relations and reasoning by analogy skills. During the test each subject has to detect a relationship and then reapply the same relationship to other stimuli by analogy. The subjects’ performance on the Raven test correlated with accuracy on tasks with fewer visual switches (the pattern recognition within a view, and searching for a value) ($r = 0.5$).

The correlation between the picture completion subtest and completion time suggests that this visualization strategy requires consistent involvement of visual recognition and focusing attention. In addition, the correlation between the Raven progressive matrices scores and accuracy suggests that analogical reasoning is required to recognize relationships between two attributes in a single visualization, but is not as important when finding patterns across visualizations. These cognitive abilities should be considered when designing multiple view visualizations.

Eye-tracking Measures

The purpose of using the eye tracking data was twofold. On one side, we were able to validate initial hypotheses by integrating eye tracking data with logging data following the approach that other authors have used [62]. The other purpose was to use this tool in a heuristic way to discover typical strategies used for accomplishing the given tasks.

The highest level result from the eye-tracking data was proof that a task can change the way that a visualization is used. In the combination with two parallel coordinate plots the tasks were typically set up so that the left view was used for selection and the right view for finding patterns. In the left parallel coordinate plot users focused on the axes in order to select data. In the right parallel coordinate plot users focused on the slope of the lines in between the axes to find patterns. Figure 4.10 shows the percentage of total time in each view that was spent on different elements of that view.

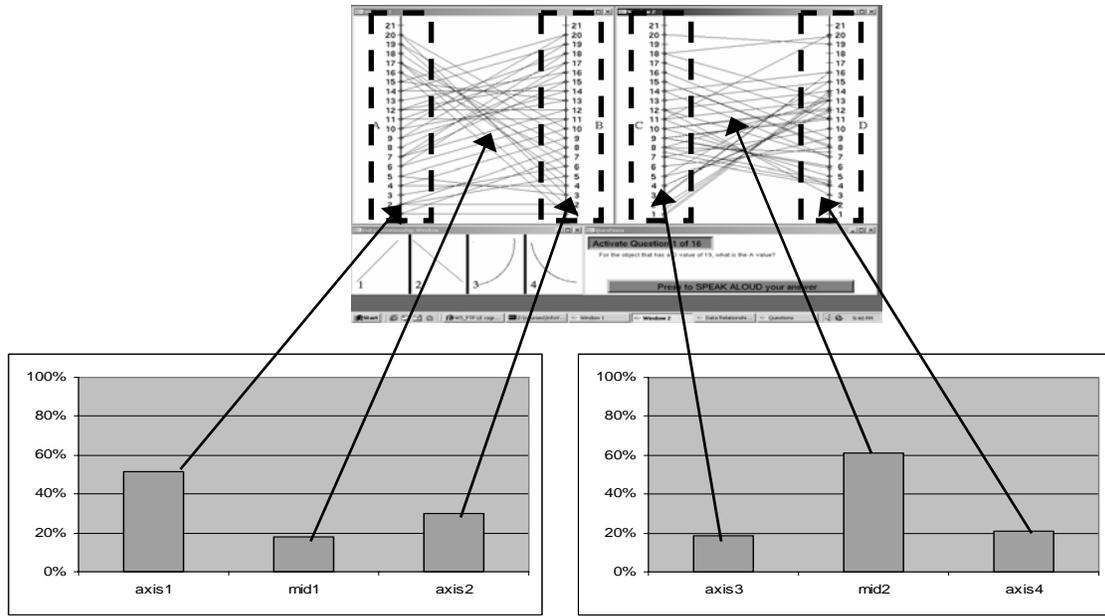


Figure 4.10 Percent of time spent looking at different elements

To get more in depth we chose to look at the extremes. From the qualitative analysis of eye-movement traces [145] of the users we considered the differences between the two subjects that missed the largest number of questions and the two that missed the smallest number of questions. The subjects with the worst performance had fewer fixations, less total area covered by their eye movements (scan path length), and longer fixation time spent at a specific location. A fixation is each instance that the eye remained in an area that was 5 pixels vertically and 3 horizontally for at least 40 milliseconds. Moreover, analyzing the fixations suggests that subjects with shorter fixation times may have performed better in general. Examples of the patterns seen when comparing the most and least accurate subjects are shown in Figure 4.11 and Figure 4.12. Note that the size of the circles represents the amount of time a user spent fixated on the location at the circle's center.

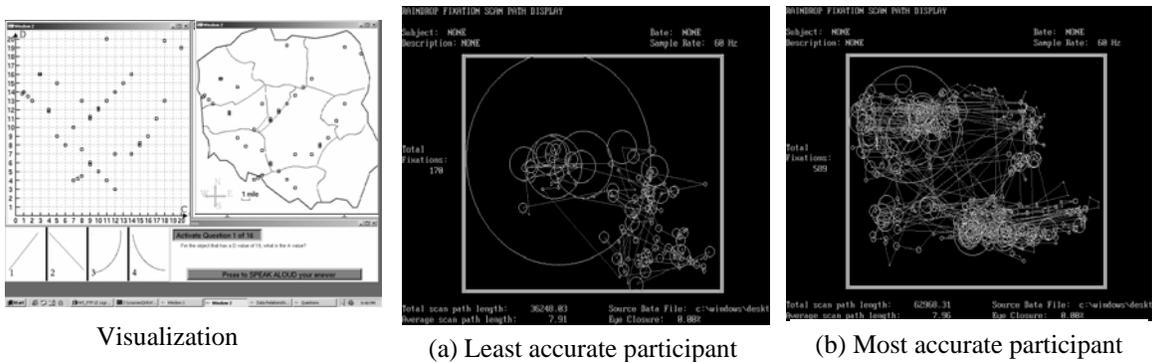


Figure 4.11. Visual paths for SG visualization, finding a pattern within the map

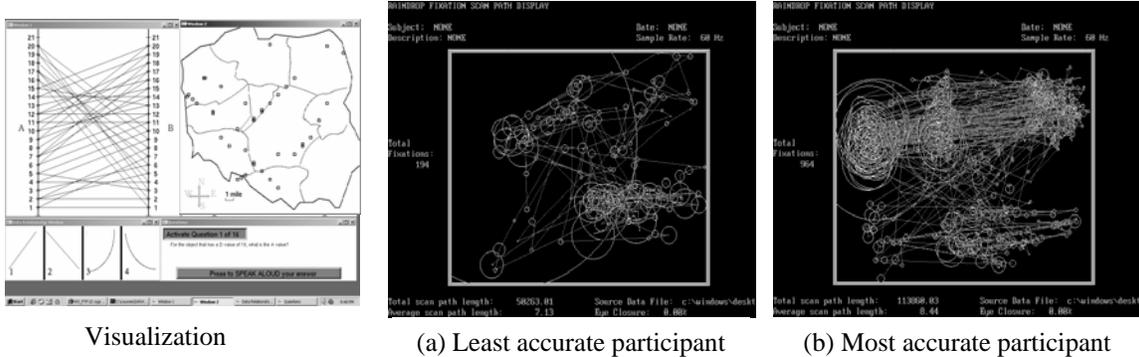


Figure 4.12. Visual paths for PG visualization, finding a pattern across views

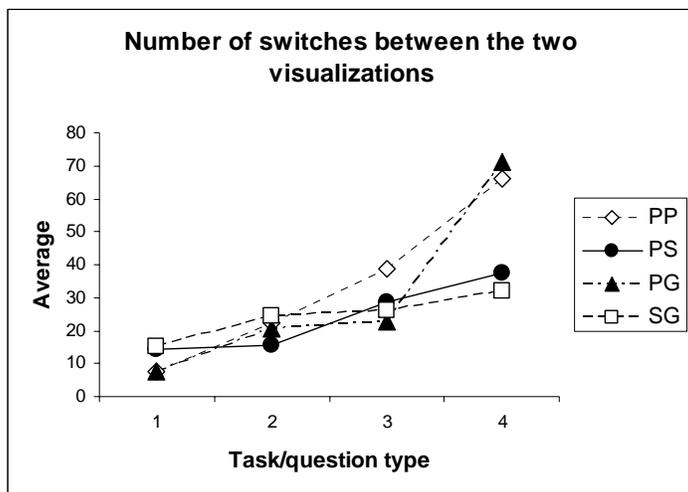


Figure 4.13. Visual switches between views

From the results of the eye-tracking data we were also able to calculate the number of times participants visually switched between the two visualizations (Figure 4.13). This appears to be consistent with the trend for completion times in the sense that both time and the total number of switches between views increased based on task difficulty, and combination of visualizations. This was also consistent with the selection of edges in the sense that as the task difficulty increased so did the number of edges participants selected. The two main factors influencing the number of switches were the type of task and the combination of visualizations.

Although the only analysis to this point was of fixations, analysis of saccades could also be done. It has been suggested that although there is saccadic overhead from visually switching between spatially distributed views (as opposed to views presented sequentially), that this overhead decreases as the task becomes more complex [87]. Because multiple view systems are only used for complex tasks and not basic tasks, the saccadic overhead should typically be low.

4.2.3. Discussion

Useful Results for Designers

The eye-tracking equipment was used to theorize about how people use multiple linked visualizations. A noteworthy observation was that the same visualization can be used in multiple ways depending on the task. For example, users focused on the axes when selecting a subset of data points from a parallel coordinate plot, and focused on the space in between axes when searching for a pattern. This finding is important for designers because they must keep in mind that seemingly less important aspects of a visualization can become important if the task changes.

As is stated in Slocum *et al.* [123], novel methods for visualization (specifically geographic) information will not be very useful unless they are based on a theoretical cognitive framework and created with usability engineering principles in mind. From the results of cognitive abilities tests, focusing attention was important in all tasks, and analogical and spatial reasoning are important in pattern recognition tasks. This provides empirical support to the heuristic walkthrough guidelines in Baldonado *et al* [11] for managing the user's attention and also might explain why in this experiment and in [48] users with worse performance visually explored less. These findings are the beginning of a theoretical framework that extends beyond single visualizations.

User performance

We predicted that situations involving two parallel coordinate plots (PP) would take the least amount of time because there would be no cognitive load from context switching. The results were contrary to our hypothesis: questions involving the use of two parallel coordinate plots (PP) actually took significantly longer to complete. This suggests that in some situations, cognitive integration may actually be more difficult when a person is presented with two identical visualizations.

Subjective rating

There was no significant difference between preference rank and difficulty rating based on the combination of visualizations. However, subjects did prefer and therefore ranked the combinations according to their perceived level of difficulty. The difficulty rating was the lowest for question type 1, and the highest for question type 4 as intended. Although subjects preferred the questions involving search tasks (1 and 2) to the pattern recognition tasks, they also preferred the questions that required multiple switches between views (2 and 4) to those which required a single visual switch (1 and 3) despite their increased level of difficulty.

Cognitive abilities

The relatively strong correlation between the picture completion subtest and completion time suggests that this visualization strategy requires consistent involvement of visual recognition and focusing attention. In addition, the subjects' performance on the Raven test is correlated negatively with the number of errors on question types 1 and 3 (single

switch search and single switch pattern recognition tasks). Since the Raven test measures reasoning by analogy using a set of visual stimuli, this suggests that analogical reasoning is required to recognize relationships between two attributes in a single visualization, but is not as important when finding patterns across visualizations.

Eye-tracking data

Results from the eye-tracking data show that when the parallel coordinate plot is presented on the left the subjects spent most of their time looking at the axes. This is to be expected since most questions require the subjects to use a subset of points from one of these axes. More interesting is the situation of PP combinations in which the subjects have to find patterns within the parallel coordinate plot on the right. Subjects actually spent most of their time looking at the area in between the two axes. This suggests that most of the subjects were looking at the slope of the lines, and not the values on the axes in order to find the pattern.

Selection between combinations

Different combinations of visualizations allow for different types of interaction. One area that differentiates combinations is the type of selection that is afforded. If the view from which the user wishes to select a subset of data is presented in an orthogonal manner, such as a scatter plot, this allows 2D selection. In other words, a user can drag along the x-axis to select a specified range of x values, and then drag upward in order to select a subset of y values for those x values. This type of interaction allows users to easily perform tasks such as those presented in question type 4. In question type 4 tasks users were given a specified range of values along the x-axis and were asked to find the relationship between y values and an attribute displayed within the other visualization. Hence, by maintaining the correct x interval the user could drag the multiple selection box up and down to view the result of increasing the y value for a given x range on an attribute in the other visualization.

While the use of a scatter plot allows for 2D selection, trying to find patterns when selecting two specified ranges of values from parallel coordinate plots can be a difficult task because of the 1D selection it affords. If a user wished to answer a task associated with question type 4, they could attempt to select the specified range on the first axis and drag the multiple selection box up and down. The problem with this strategy is that, because the x and y dimensions are not orthogonal, as they drag it up and down the corresponding points on the second axis typically will not increase and decrease in a standard manner. This type of 2D versus 1D selection provides any combination of visualizations that include a scatter plot or other orthogonal combination a definite advantage. Orthogonal axes provide the possibility of working on an intersection of a range of data on one axis and a range of data on another axis. An additional issue with using parallel coordinates, based on observation of participants' performance, was the rectangular selection box made it difficult not to select more edges than desired, especially with increased numbers of edge crossings.

4.2.4. Summary

The purpose of this experiment was to probe the cognitive issues with using multiple linked visualizations, and to specifically consider context switching [114]. Baldonado et al. [11] cite context switching as a reason not to use multiple views repeatedly in their paper on design heuristics. To explore these issues, a controlled experiment was performed using different combinations of dual view visualizations for search and pattern recognition tasks. To collect the data psychological tests [31], logs of the participants' interaction, eye-tracking equipment, and video recordings were used. Main findings include context switching not being as detrimental as it may first appear, focusing attention correlating with user performance and analogical reasoning being important for finding patterns within a single view.

4.2.5. Conclusion

We have explored the cognitive abilities involved in working with two linked visualizations and the effects of context switching. Using cognitive ability pretests, we were able to find correlations between focusing attention, analogical reasoning, and performance. Additionally, this study showed that context switching may not increase the difficulty of cognitive integration. Similar visualizations may cause interference resulting in decreased performance. An alternate explanation is that subjects have to mentally transform patterns in parallel coordinate plots to an orthogonal representation and this additional step reduces performance.

4.3. Experiment 3: Visual vs. Interactive Linking

When designing visualizations that integrate multiple data attributes, a fundamental tradeoff is whether the interface complexity is in the visual representation or the interactive linking. In integrated view visualizations, the interface complexity exists in the visual representation. In multiple view visualizations, the interface complexity is in the interactive linking. This is particularly true in scenarios where there is no spatial data type. With the attribute-centric visualization information can be mentally integrated without brushing and linking because the spatial context provides a visual connection. However, it is impossible to know the connection between data in different visualizations without interactive linking if only multidimensional data is involved and not spatial data. Therefore, the goal of this experiment was to consider the impact of the basic tradeoff between visual and linked integration of quantitative data attributes in terms of performance time, accuracy, and satisfaction for novice users on simple tasks. In this section the terms integrated and multiple views will be used because without the spatial context there cannot be space and attribute-centric visualizations.

4.3.1. Method

In this study, one, two, and four views were compared, with four data attributes distributed equally amongst the views. Each condition used a spatial encoding for all attributes. The only exception was that the one view condition used color to represent the fourth data attribute. The integrated view condition used a 3D scatter plot and color, the

dual view condition used two 2D scatter plots, and the multiple view condition used four separate 1D plots. This allowed for a comparison between placing all attributes in a single view, distributing them evenly between two views, or displaying them each in their own view. There were two independent variables: number of views and task. Because tasks each had very different requirements, each task was analyzed separately for all interfaces.

Participants

The design was a 3x10 mixed design with visualization a between subjects factor and task a within subjects. There were thirty participants, twenty five undergraduate engineering students and five graduate students in various majors from a large public university.

Materials

Three different visualization interfaces were built for this experiment 1) integrated, 2) dual, and 3) multiple view visualizations. Each displayed four abstract data attributes and the data came from the U.S. Census Bureau. Abstract attribute names (A, B, C, and D) were used to avoid interference from previous knowledge. However, state names were used as a way of labeling data points. The data was quantitative and consisted of a single tuple for each state. Each of the visualizations presented the data in a slightly different manner. The methods of visualizing data are as follows (see Figure 4.14, Figure 4.15, and Figure 4.16):

- **Integrated View:** The four attributes were displayed in a single 3D view. X, Y, and Z axes were used to visualize attributes A, B and C respectively. The color gray [7],[12] was used to visualize attribute D. In terms of interaction, the user was able to grab and rotate the 3D scatter plot using the left mouse button and display a state name and the attribute values by selection using the right mouse button.
- **Dual views:** Attributes A and B were displayed on a 2D scatter plot on the left side of the screen, and attributes C and D on a 2D scatter plot on the right side of the screen. Brushing [3] and linking were used. As a user selected points in one scatter plot, the corresponding points in the other scatter plot were highlighted.
- **Multiple views:** Each attribute was displayed on a separate 1D x-axis. As with dual views, brushing and linking was used to connect plots.

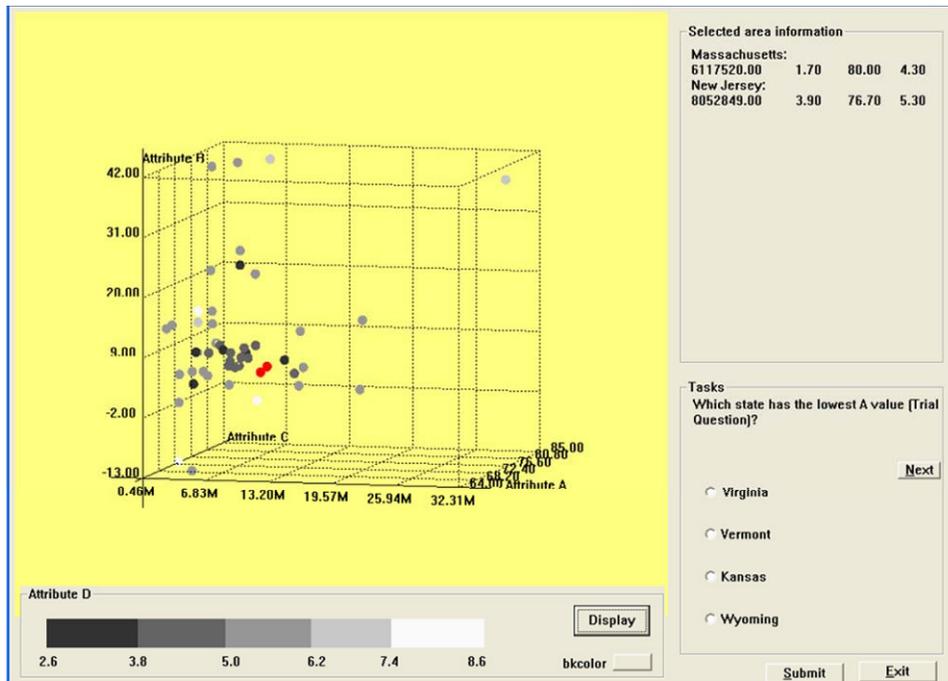


Figure 4.14. Integrated 3D view

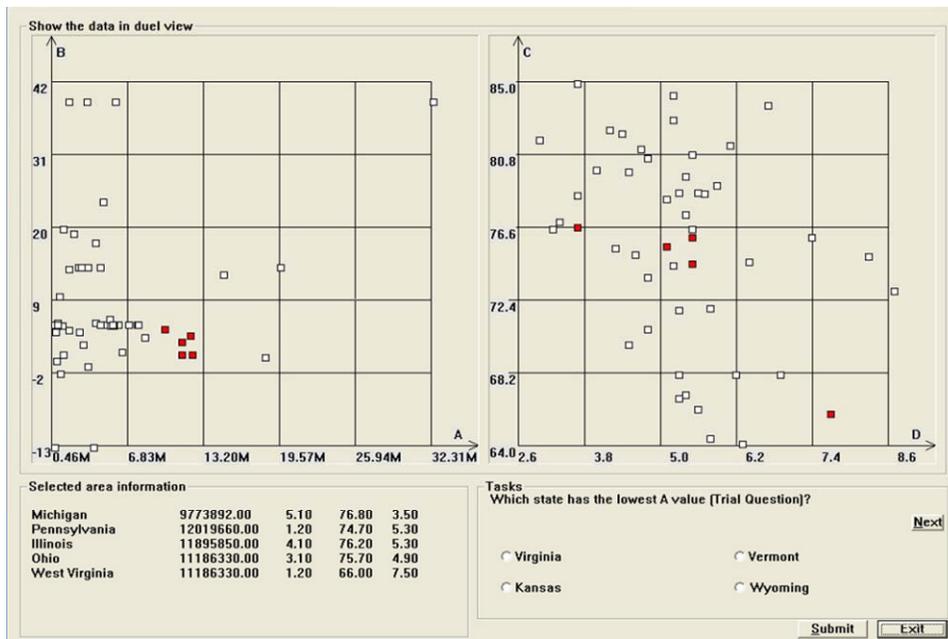


Figure 4.15. Two views with linked scatter plots

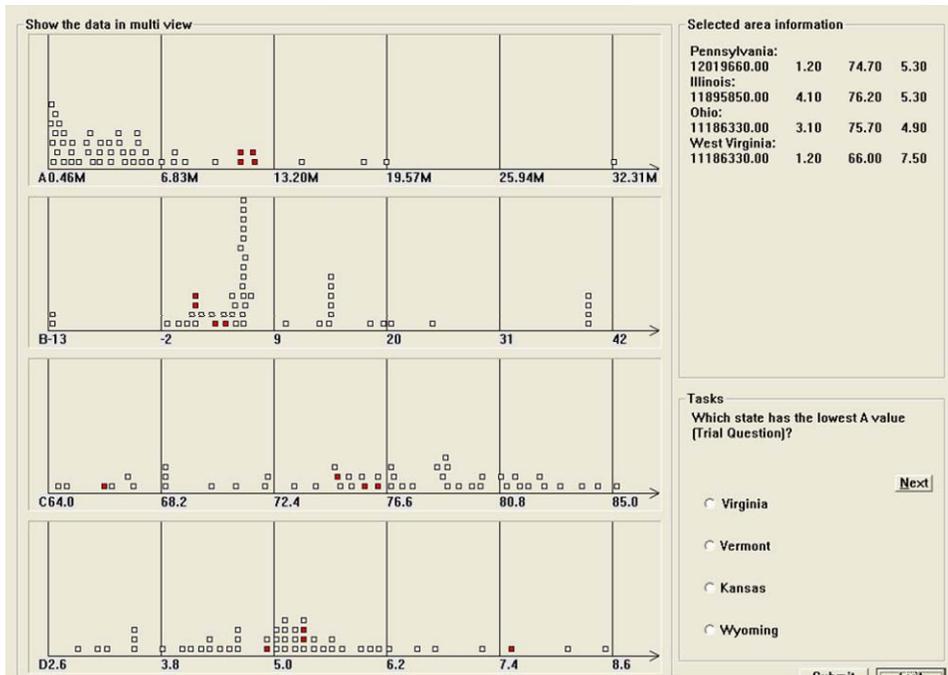


Figure 4.16. Multiple linked 1D plots

Procedure

Before starting the tasks each participant was asked to fill out a demographic survey and then given a demonstration of the interface. Following the demonstration each participant was asked to complete two practice tasks followed by ten experimental tasks. The specific tasks can be seen in Table 4.3.

Table 4.3. Practice tasks and experimental tasks

Task	Question
Practice 1	Which state has the lowest A value?
Practice 2	What kind of state is Virginia in terms of A, B, C, and D?
1	Which state has the highest D value?
2	Which state has B value in range 9.0 ~ 20.0?
3	Which state has A value in range 6.83M ~ 13.20M and B value in range -13.0 ~ -2.0?
4	Which state has A value in range 6.83M ~ 13.20M and C value in range 72.4 ~ 76.6?
5	Which state has A value in range 6.38~13.20M, B value in range -2.0~9.0, C value in range 76.6~80.8 and D value in range 5.0~6.2
6	What's the relationship between A and B?
7	What's the relationship between A and C?
8	How many states have A value in range 6.83M ~ 13.20M and D value in range 5.0 ~ 6.2?
9	How many states have A value in range 13.20M ~ 19.57M and B value in range 9.0 ~ 20.0?
10	Which range of A value do the majority of states fall into?

The tasks consisted of finding a state with a specified value, finding a state with multiple specified values, identifying relationships, and determining how many states fit into a specified range of values. There was no time limit to complete the tasks. The software recorded performance time and accuracy. After completing the tasks students were asked to rate their satisfaction for each task. A rating of 1 indicated that the user was “extremely unsatisfied” and a rating of 4 indicated that they were “most satisfied”.

4.3.2. Results

ANOVAs were performed for each task. For each significant ANOVA, a t-test was used for post-hoc analysis.

Time

Figure 4.17 shows performance times for all tasks. ANOVAs on performance times indicated significant differences for task 1 ($F(2, 27) = 6.72, p < 0.01$), task 3 ($p < 0.01$), task 4 ($p < 0.01$), and task 7 ($p < 0.02$). Post-hoc analysis of task 1, which required the user to find the highest D value, showed that the integrated view condition (72.9 seconds) took significantly longer than both the dual ($p < 0.05$, 17.4 seconds) and multiple view ($p < 0.01$, 13.5 seconds) conditions. The integrated view also (99.8 seconds) took significantly longer than both the dual ($p < 0.01$, 46.5 seconds) and multiple view ($p < 0.01$, 42.3 seconds) visualizations for task 3, where the user had to find a value within a specified A and B range.

The integrated view visualizations took longer for tasks 4 and 7. For task 4, identifying a point within a specified A and C range, the integrated view condition (109.3 seconds) took significantly longer than multiple views (43.9 seconds, $p < 0.01$). For task 7, finding a relationship between A and C, using the integrated view (57.2 seconds) took longer than using dual views (24.6 seconds, $p < 0.05$).

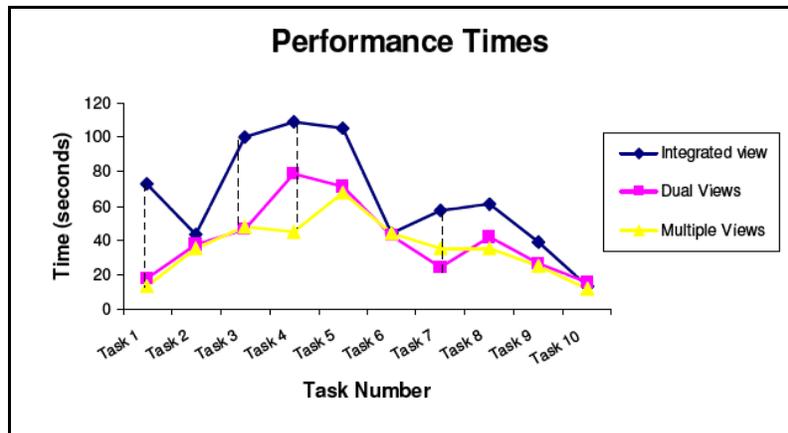


Figure 4.17. Performance times in seconds for each task and interface, dotted lines indicate significant differences

Accuracy

Figure 4.18 shows accuracy for all tasks. Answer correctness was significantly different for task 6 ($p < 0.01$) and task 8 ($p < 0.0001$). For task 6, finding a relationship between A and B, dual view accuracy (50%) was significantly worse than both multiple (100%, $p < 0.05$) and integrated view (90%, $p < 0.05$) accuracy. The same held true on task 8, counting states within a given A and D range of values. The dual view visualization (20%) resulted in significantly less accuracy than both the multiple (100%, $p < 0.01$) and integrated view (90%, $p < 0.01$) visualizations.

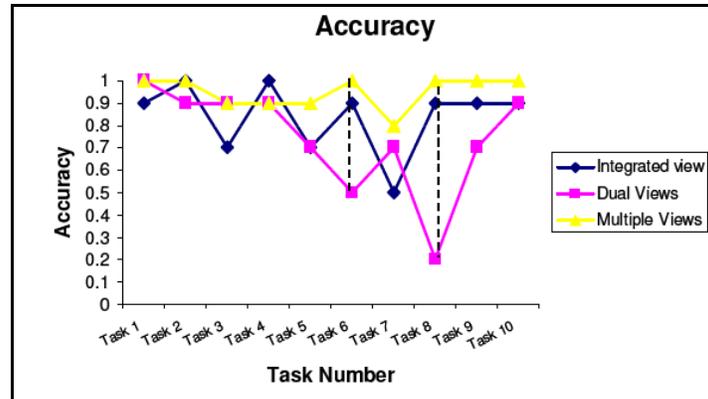


Figure 4.18. Accuracy for each task and interface, the dotted lines indicate significant differences

Satisfaction

Figure 4.19 shows satisfaction ratings for each task. Satisfaction scores were significantly different on task 1 ($p < 0.02$), finding the state with the highest D value, and task 3 ($p < 0.04$), finding a value within a given A and B range. For task 1 participants using the integrated view (3.1) reported significantly lower satisfaction scores than those using either dual (3.9, $p < 0.05$) or multiple views (3.9, $p < 0.05$). For task 3, multiple view users reported significantly lower satisfaction ratings (2.4) than those using dual views (3.4, $p < 0.05$).



Figure 4.19. Satisfaction ratings, the dotted lines indicate significant differences

4.3.3. Discussion

Because each task is unique and has its own requirements, each task that resulted in significant differences is discussed separately below.

Task 1: Finding an Extreme Value of a Single Attribute

Task 1 asked the user to find the highest attribute D value. Using an integrated view in which D was represented as color resulted in significantly longer completion times. The problem with integrated views in this situation was likely the use of a limited number of colors and multiple data points with different D values falling into the same color categorization. This leaves the user requiring details before an answer can be given. Color is also a nominal visual variable and not well suited for this task. Comparatively, dual and multiple view users could simply compare spatial positions in one of the dimensions. Another potential factor adding to the poor performance of integrated views for this task might have been the difference in encoding. A user may have searched axes to determine which axis D was mapped onto, and been slowed down from any 3D rotation involved in this exploration.

Task 3: Discovering That No State Falls within the Specified Range

Task 3 asked the user to find the state with an A and B value in a specified range. The real test in this condition was that the correct answer was 'none'. For this task, using an integrated view took significantly longer than both the dual and multiple view conditions. The longer integrated view times could be due to difficulties in rotating the scatter plot. It is also possible that the added complexity of having four attributes represented in a single view caused the user to believe they were missing something and therefore they searched longer. The latter seems likely to be the cause since there was no significant difference in accuracy across the conditions.

For this task, dual views received higher satisfaction scores than multiple views. While dual views represented A and B in the same scatter plot, multiple views used the x-axis once for each variable. This finding is consistent with Wickens's proximity compatibility principle, suggesting A and B should have close display proximity [14]. In addition to considering whether the attributes were integrated or separated, the other difference between the conditions is that interactive linking (brushing) had to be used to find the relationship between A and B for multiple views. It is possible users preferred a direct visual integration to an indirect interactive link. In other words, one of the scatter plots in the dual view visualization was specifically suited for this task.

Task 4: Finding a State with the Specified Attribute Values

Task 4 involved finding the state that had given attribute A and C values. For integrated views, A and C were represented on the x-axis and z-axis respectively. Here the user would need to rotate the view to see both attributes. When using dual views A was the x-axis in the left hand view and C was on the y-axis in the right hand view. To accomplish this task, a user had to either select all data points within the given A range of values or C

range of values first. The data points in the other scatter plot would then be highlighted and the user could select the correct data point. Using multiple views, A was represented on the top line and C on the third line down. Here the user could select all A or C values in range, and the matching values of the other attribute would be displayed. This intuitively seemed as if it would be more difficult because the points had to be selected individually to find which data points matched across attributes.

For this task using integrated views took significantly longer than using multiple views. Although users could not do a visual comparison without remembering which data attributes matched, the complexity of rotating a 3D interface to find the correct attributes appeared to be much more difficult.

Task 6: Discovering the Relationship between Two Attributes

Task 6 required the user to find the relationship between A and B. For this task the relationship was not one of the three choices provided (there was a 'none of the above' choice). In the integrated view, users had to rotate the 3D scatter plot to find the x (for attribute A) and y-axis (for attribute B). Using dual views both A and B were represented in the left hand scatter plot. For multiple views a user could gradually highlight either A or B to see which corresponded. They were also able to see a general distribution of the values in both A and B separately. For this task the dual view accuracy was significantly worse than integrated and multiple. It is possible that the distribution of A and B in multiple views and the extra encoding of color in integrated views led participants to choose the right answer for the wrong reasons. It is also possible that the overlapping of data points in the dual view scatter plot made identifying the relationship more difficult. Users may have thought that more data points existed than actually did exist.

Task 7: Finding an Inverse Relationship

Task 7 asked the user to find the relationship between A and C. Here, as one attribute increased in value, the other decreased. In an integrated view, A was represented along the x-axis and C along the z-axis. For dual views, A was the x-axis on the left and C was the y-axis on the left. For this task performance time using an integrated view took significantly longer than when using a dual view interface. It appears that having to interactively rotate the scatter plot and mentally filter information in the integrated view took longer than finding the relationship when using interactive brushing.

Task 8: Counting the Number of States Matching Given Criteria

Task 8 provided users with an A value range and a D value range and asked them how many states matched those criteria. For this task, dual view had significantly worse accuracy than both the multiple and integrated view conditions. The real issue here was that users mistook a data point that was near the range as being within the range in the 2D scatter plot. Only by looking at the specific values did users determine that there were 3 and not 4 states matching the criteria. Considering the accuracy rate, it appears users did not double check the values even for very few matching data points. It is surprising that only 20% of users appeared to have checked the actual values in the dual view condition.

Dual view users may have checked values less than the other groups because 2D scatter plots are more familiar and therefore resulted in more confidence.

General Observations

The multiple views condition typically resulted in the best performance. It appears that interactive brushing-and-linking works well at aiding the users in cognitively integrating the data compared to visually integrated views. Overall distributing data attributes across views and using brushing and linking provided more reliable performance across a variety of user tasks than integrating four attributes into a single 3D scatter plot. Projecting the third dimension onto the other two dimensions involves distortion of all three dimensions and provided slower and less accurate estimates.

While satisfaction scores were relatively consistent across view types, dual view accuracy was generally slightly lower and integrated view performance time was generally slower. Rotation in the 3D scatter plot adds complexity to the user interface, and hence adds time without an increase in accuracy. This supports known issues with 3D visualization – that 3D visualization introduces additional visual and interactive complexities. Displaying all attributes in the same view also leads to problems with visual encodings. As the number of attributes is increased, designers must generally resort to less effective encodings such as color. This also increases the potential for visual interference between encodings. The additional complexity may result in users searching longer for relationships that do not exist.

Although the dual view condition received significantly higher satisfaction scores for one task, the use of this interface generally resulted in less accuracy. It is possible the occlusion problems in 2D scatter plots can lead to inaccurate perceptions. Using multiple views, there was no occlusion. Using integrated views and rotation allowed the user to see around some of the occlusion. However, it is difficult to determine if the users were more accurate in the other conditions for potentially wrong reasons.

4.3.4. Conclusion

This experiment began to consider the different types of interaction that integrated and multiple view visualizations afford. In this study the goal was to compare visually linked attributes (an integrated view) to interactively linked attributes (multiple views with brushing) for multidimensional data without a spatial component. As in the experiment comparing multiple simple and single complex glyphs, a dual view situation was also considered. The three different methods of visualizing four data attributes were compared across various tasks. The integrated view interface was a 3D scatter plot that encoded the fourth attribute using gray values. The dual view interface used two linked 2D scatter plots. The multiple views interface used four separate linked 1D axes. While the complexity of the 3D interface resulted in slower task completion times, using two linked views resulted in less accuracy. Overall, it was a bit surprising that the multiple linked 1D views typically resulted in the best performance. Because distributing data attributes across views and using brushing and linking provides more reliable performance across a variety of user tasks than integrating four attributes into a single 3D scatter plot, it

appears that interactive brushing and linking works well at aiding the users in cognitively integrating the data.

4.4. Summary and Synthesis of Results

The overall goal of the three previously described experiments was to determine which visualization resulted in the best user performance on a typical desktop display. Toward answering that question, a basic comparison of the space (using complex glyphs) and attribute-centric visualizations was conducted, cognitive issues with distinct linked visualizations were explored, and interactive versus visual linking was compared. The main findings are as follows.

Research Question 2:

As a baseline, which visualization results in the best user performance for specific tasks on a desktop display?

- Performance with space-centric (using complex glyphs) and attribute-centric visualizations was similar for most tasks, although various trends were indicated.
- The perceptual salience of visual encodings was more important than the number of views. Perceptually salient visual encodings resulted in better user performance than less salient encodings. This gives the attribute-centric approach an advantage compared to the single complex glyphs often used in space-centric visualizations.
- Mixed designs with data split between visualizations almost always resulted in the worst user performance.

While performance was similar, general observations can be made based on non-significant differences:

- Multiple views (including attribute-centric): Not significant, but trends indicate that user performance is better for attribute ranges/values/attribute detection tasks.
- Mixed design: Best when it matches the tasks, but interference and mental integration problems results in the worst user performance.
- Integrated view (including space-centric): Interference and occlusion problems in 3D; not significant, but trend indications that it might be slightly better for trend and relationship tasks in 3D and worse when using complex glyphs (with the space-centric design).

The multiple view approach was generally better for tasks that involved finding an attribute within a given range, a specific value, or for any type of basic search and compare task. The integrated view approach seemed to be slightly better for trend and relationship tasks in 3D with only multidimensional data, but not when employing complex glyphs with the space-centric design. Based on trends in the relationship finding tasks with complex glyphs, it appears easier to mentally integrate information from different structures each displaying a single attribute than to attend to a select number of features in complex glyphs. These results are generally consistent with previous research on visual encodings. Because conjunctive encodings were used in the space-centric

visualization the attributes could not be processed pre-attentively and the other visual encodings used likely interfered with detecting both single and multiple targets. In addition to these differences, it appears brushing helps users focus their attention and also increases the ability to compare values across maps.

The generalization of these results is limited by the small size of the datasets and relatively simple tasks. Therefore, in the next part of the work the dataset size was increased along with the screen size and more complex tasks were completed, this allowed us to determine if the results scaled.

The experiments in this chapter were mostly exploratory in nature and background work which lead into the next chapter, which is central to this work.

Chapter 5. Visual Scalability for Large Displays

The third research question was: how visually scalable are the visualizations for large, high resolution displays? Human perception (perceptual scalability), monitor resolution (the display), and visual metaphors (graphical scalability) are the first three components of visual scalability. In this chapter the graphical scalability of the space and attribute-centric visualizations is first discussed. That is followed by two experiments conducted on the perceptual scalability of visualizations (and published as [147, 149]).

5.1. Graphical Scalability

Graphical scalability is the number of pixels required to display a visualization. In general the scalability of the attribute-centric approach is limited by the number of individual views. The space-centric approach is limited to about eight distinct dimensions with complex glyphs because of the types of encodings available [139]. An alternative and more scalable space-centric approach is to use embedded visualizations in place of complex glyphs. With this approach, the space-centric visualization is limited by over plotting across and within spatial locations.

Again, the data of interest integrates both spatial and multidimensional data types. In this chapter the attributes are actually a combination of multidimensional and temporal data. An example of this type of data is US Census data which has values for each county (spatial), for various demographic groups (multidimensional), over a number of years (temporal). Throughout this work we scaled-up the data by increasing the number of time points and demographic groups while maintaining a constant number of geospatially-referenced locations.

Table 5.1. Calculating graphical scalability with opposing visualizations

Attribute-centric Visualization	Space-centric (embedded) Visualizations
<i>Number of locations</i> = 3100 (US counties)	<i>Number of locations</i> = 3100
<i>Pixels per location</i> = 10×10 = 100 pixels (size of a single county containing a value bar)	<i>Pixels per location</i> = 30×30 = 900 pixels (labels, legend, etc. for a single embedded visualization)
	<i>Initial setup</i> = pixels per location × number of locations = 900 × 3,100 = 2,790,000
<i>Pixels per attribute</i> = pixels per location × number of locations = 100×3,100 = 310,000 (size of 1 map)	<i>Pixels per attribute in a single embedded visualization</i> = 56 (7×8 pixel area – size of single grid position)
	<i>Pixels per attribute</i> = pixels per attribute in a single embedded visualization × number of locations = 56 × 3,100 = 173,600
<i>Pixels available</i> = total pixels – 15% of total pixels	<i>Pixels available</i> = total pixels – 15% of total pixels
Maximum attributes = Pixels available/Pixels per attribute = Pixels available/310,000	Maximum attributes = (Pixels available – Initial setup)/Pixels per attribute = (Pixels available – 2,790,000)/173,600
Maximum attributes on Gigapixel display = 2,741	Maximum attributes on Gigapixel display = 4,880

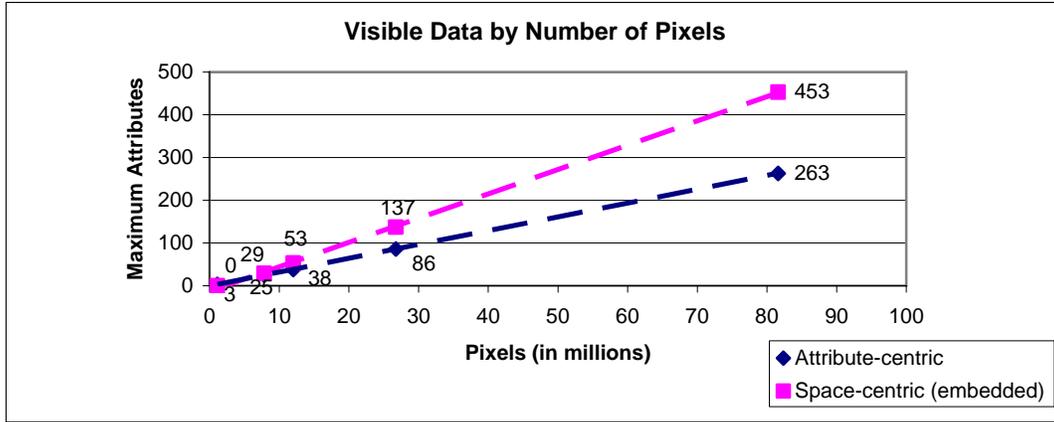


Figure 5.1. Visible data by number of available display pixels

To demonstrate just how much the amount of visible information could be scaled-up with attribute-centric and space-centric using embedded visualizations we created the generic equations shown in Table 5.1. The number of locations was fixed at 3,100 to match the number of US counties. To further simplify the equations we subtracted 15% from the total available pixels to account for overall legends and labels. Although this could be kept constant, larger displays need legends and labels that are viewable from greater distances. We define the number of attributes to be the number of time points \times the number of demographic groups.

Using these equations, the space-centric approach using embedded visualizations cannot be used on a standard desktop display (assuming 3100 locations) due to the initial setup costs and the resulting overlap in embedded visualizations between spatial locations. However, as more information is added this approach graphically scales better than the attribute-centric approach. The reason for this is that when increasing the number of attributes, adding a map requires more space than adding a grid location to all grids. Figure 5.1 shows the maximum number of simultaneously visible attributes for each display in Figure 2.2.

When moving from a desktop display with the attribute-centric approach ($3 \times 3,100 = 9,300$) to space-centric using embedded visualizations on the 50 monitor tiled display ($453 \times 3,100 = 1,404,300$), there is a greater than 150 times increase in the total number of simultaneously visible data entities.

5.2. Experiment 4: The Perceptual Scalability of Visualization

Different visualizations are better able to graphically scale (require fewer pixels), which is especially important on typical desktop displays [49]. However, even designs that may not graphically scale well for a desktop display can be scaled-up to a greater extent using displays that are larger and/or have a higher resolution (DPI). Theoretically, any dataset could be visualized, regardless of the visualization, on an infinite size display.

Therefore, as larger displays are used for visualization, the scalability limit may shift away from the graphical scalability limits imposed by the number of pixels and toward

human limits. The most obvious examples of this occur when the display exceeds a resolution such that the human eye cannot perceive the pixels regardless of distance from the display, and when the display size gets to be so large that significant physical movement would be required by a user (as an example, consider the Vietnam Veterans Memorial Wall in Washington, D.C. which is 493.5 feet (150.42 meters) wide with more than 50,000 names inscribed that are each 0.53 inches (1.35 cm) high [3]). In these cases, the limit is created by human abilities rather than caused by the display technology.

This leads to the question of perceptual scalability of visualizations for large displays. When the screen isn't the limiting factor, just how much data can a person effectively perceive? As more data is shown with increasingly larger displays, do we hit a breaking point, the limits of visualization? And how will visualizations for large displays need to fundamentally differ from visualizations on desktop displays? Are basic visualization design principles different? In this experiment we compared three different visualizations across two different display sizes: a 2 Mp display and a 32 Mp display.

The goal of this study was to examine:

1. *How perceptually scalable are data visualizations for large displays?* In other words, what happens to time/accuracy as both amount of data and number of pixels are increased?
2. *Are some designs more perceptually scalable than others?* In other words, are relative comparisons between designs the same at different screen sizes?

A visualization that is perceptually scalable should not result in an increase in task completion times when time is normalized to the amount of data. It also should not result in decreasing accuracy.

5.2.1. Method

A 2x3x7 mixed design was used. The independent variables were display size (with a proportional increase in data size), visualization design, and task respectively. Display size was treated as a between subjects variable while visualization and task were within subjects. Task completion time, accuracy, and subjective workload were recorded. Each independent variable is described in further detail throughout this section.

Data

The data used in this study was based on the U.S. Department of Justice online Crime & Justice database (<http://bjsdata.ojp.usdoj.gov/dataonline/>). It consisted of percentages of age, race, and gender demographics of homicide victims from 1976-1989. This made it a geospatially-referenced multidimensional time-series dataset in that geospatially-referenced (different states) attributes (demographic groups) were reported over time (1976-1989). To prevent influence of previous knowledge and expectations, participants were not aware of where the data was from or that it was homicide related until completion of the study. During the study they were only aware of the years, the demographic groups, and that each was associated with a value from 0-100. The data was

modified so that a singleton was the answer to every detail/find task and there was always a single definite answer. For temporal overview tasks, the trends generally involved a 5-7 point increase or decrease each year with random variance added. Because the display had a 10 to 3 aspect ratio, the top and bottom portions of the map were cropped to fit, leaving 28 visible states. State boundaries and state names were displayed.

Visualizations

Three different visualizations were used. Of these three, two were space-centric visualizations using the embedded visualization approach and one was an attribute-centric visualization. The three visualizations were as follows:

1. Attribute-Centric (abbreviated as MULTS because it is one type of multiple view visualization): A separate United States map was shown for each attribute (where an attribute is a time and demographic group). On each map a single colored bar was displayed at each state location representing that state's population value for that demographic group and year (see Figure 5.2). A global legend matching the bar color and height to the value was shown in the bottom left corner.

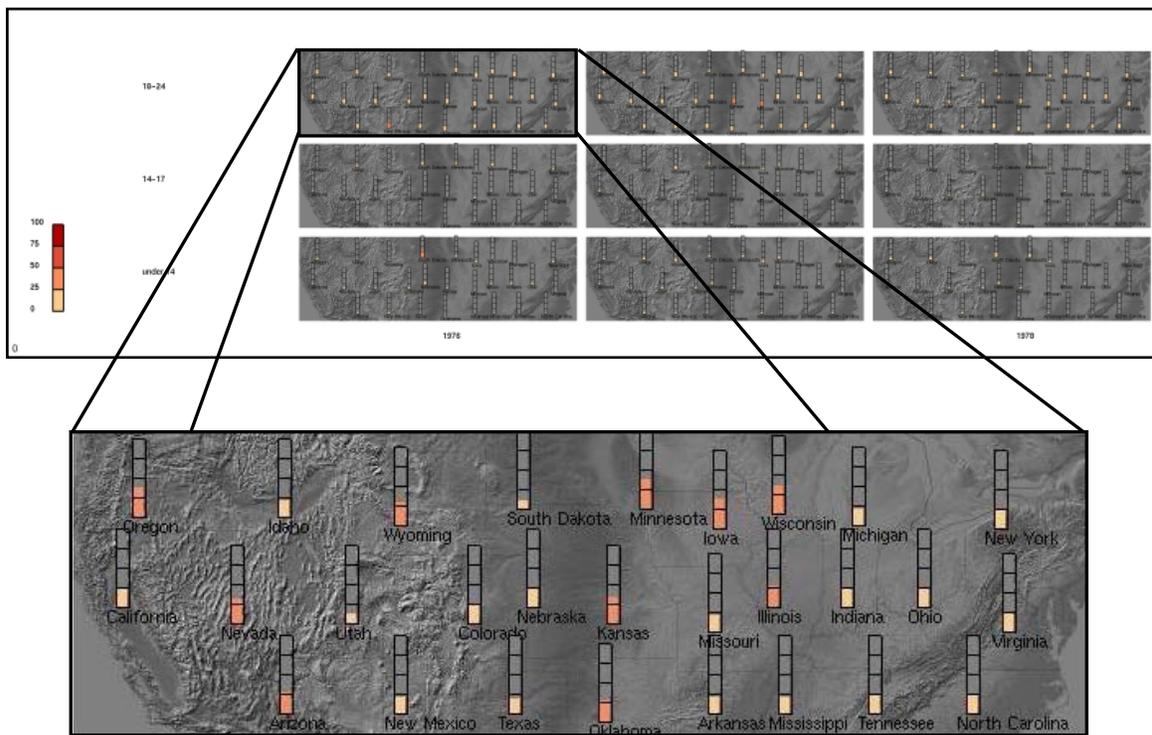


Figure 5.2. Attribute-centric visualization (MULTS) with 9 attributes on the 2Mp display

Space-Centric (integrated view visualizations using the embedded visualization approach): For the space-centric visualizations a single large map was displayed and each state had an embedded visualization containing all of the attributes related to that state. Two different methods of encoding the data were used: bar matrices and time series graphs. Bar matrices were used instead of a technique such as parallel coordinate plots so we could compare the space and attribute-centric visualizations while keeping the visual encoding consistent.

2. Bar encoding (abbreviated as BARS, a space-centric visualization): This visualization represented values as a matrix of colored bars at each geographic location. It was created so the encoding was consistent with the MULTS visualization, but the bars were grouped by their geographic location rather than by their space/time pair (see Figure 5.2). A global legend matching bar colors to values was shown in the bottom left.
3. Line encoding (abbreviated as GRAPHS, a second type of space-centric visualization): This was an attempt to improve on BARS by combining the y-axes and also increasing familiarity since many people have seen time series graphs before (see Figure 5.4). Because of the difficulty matching colors, a local legend was included with each embedded time series graph. The legend was created so that it matched the order of the last time point for that location. However, while the legends were in a different order based on the data at each location, the colors used for each demographic group remained the same across locations.



Figure 5.3. Space-centric (embedded) visualization using bar matrices [BARS]

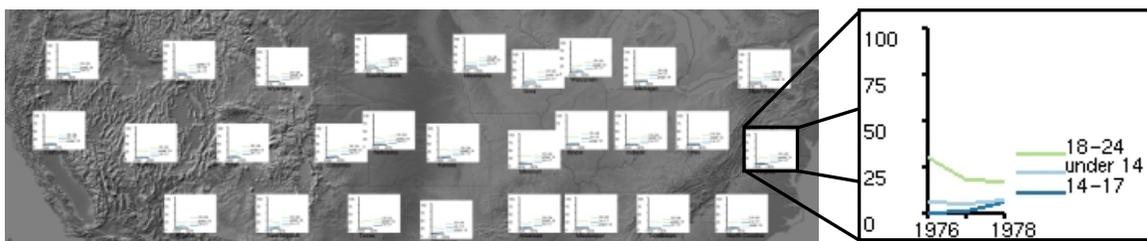


Figure 5.4. Space-centric (embedded) visualization using time series graphs [GRAPHS]

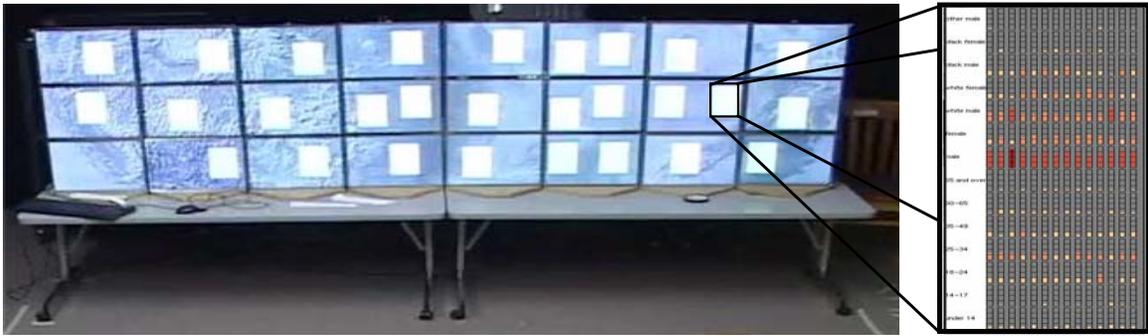


Figure 5.5. Bar matrices visualization on a 32 Mp display

Display Size

We used a 24 monitor tiled display (Figure 5.5). It was arranged in 8 columns that were 3 monitors high. Each monitor was a 17-inch diagonal LCD, 1280x1024 (~96 DPI) for a total resolution of 10,240 x 3,072 or approximately 31.5 million pixels. The total size of the display was roughly 9 feet wide and 3.5 feet high.

Two different display conditions were used:

1. An approximately 2 Mp 2-monitor portion of the display with 3 time points by 3 demographic groups (9 attributes) and 252 total data points
2. An approximately 32 Mp 24-monitor display with 14 time points by 14 demographic groups (196 attributes) and 5488 total data points.

The data to screen size ratio was not a perfect match because some of the display was left blank to maintain an equal number of time points and demographic groups in the MULTS condition. However, the size of each individual map in the MULTS condition was constant between screen sizes as well as the size of each of the bars and the text in the BARS condition. Example of the scaled-up version of the attribute-centric visualization (MULTS, multiple views) and one of the space-centric visualizations (GRAPHS) are shown in Figure 5.6 and Figure 5.7. The attribute-centric technique scaled up the number of individual maps (ex. Figure 5.2). The space-centric (embedded) visualizations (BARS and GRAPHS) scaled-up the number of attributes shown in the embedded visualizations at each geographic location (Figure 5.8).

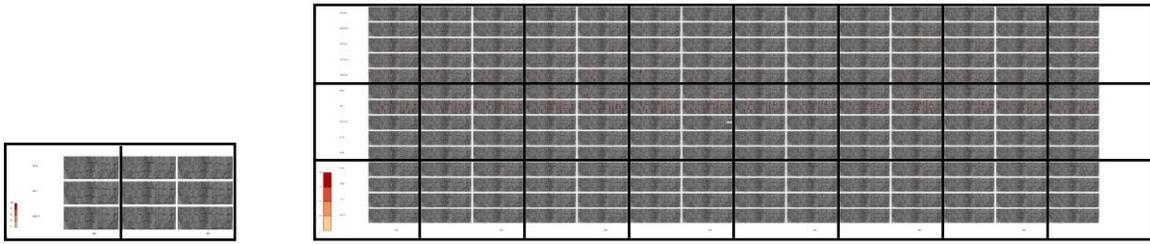


Figure 5.6. Attribute-centric (MULTS) in two display conditions

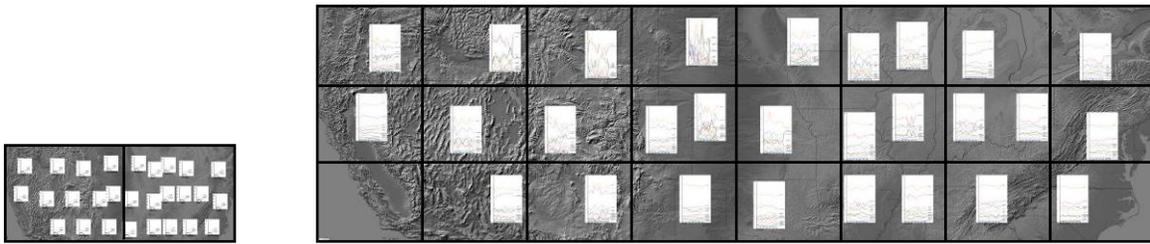


Figure 5.7. Space-centric (GRAPHS) in two display conditions

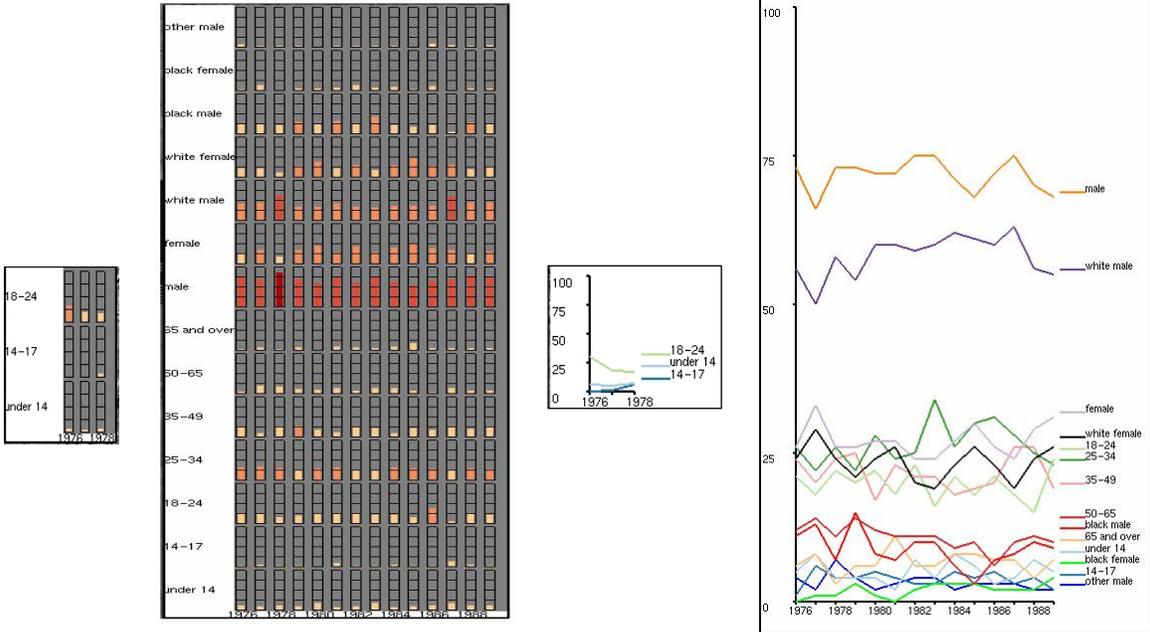


Figure 5.8. Scaled-up embedded visualizations (BARS, GRAPHS) for a single geographic location

Tasks

There were 3 detail tasks and 4 overview tasks for a total of 7 different tasks. For each of these there was a task related to time, attributes, and space. Additionally, there was a spatiotemporal overview task. The tasks are shown in Table 5.2. A modified version of the task was used for each of the different visualizations and each type of task was performed twice with a given visualization. This meant participants completed 3 visualization x 7 tasks x 2 trials = 42 tasks.

Table 5.2. Task types and examples

(D = detail, O = overview, T = time, A = attribute, S = space)

	Task Structure	Example Task
DT	Find a year, given an attribute and location.	Which year was the population of 14-17 the highest in Kansas?
DA	Find an attribute, given a year and location.	Which population was highest in South Dakota in 1976?
DS	Find a location, given a year and an attribute.	Which state had the highest population of under 14 in 1977?
OT	Identify a trend across time for all attributes and locations.	In general, populations have gone [up then down, down then up, up, down]?
OA	Identify a trend in attribute values for all years and locations.	In general, most populations values are in the range [<25, 25-50, 50-75, >75]?
OS	Identify a trend in location for all years and attributes.	In general, populations are the highest in the [North, South, East, West]?
OST	Identify a relationship between locations and times for all attributes.	In general, populations increase the fastest in the [North, South, East, West]?

Procedure

Participants were randomly assigned to either the 2 or 32 Mp display condition. After signing a consent form, participants filled out a demographic form. Participants in the 32 Mp condition were given a stool that could easily be moved, and in the 2 Mp condition they were given a regular office chair. For the 2 Mp condition the second and third displays from the right in the middle row were used. There was also a video camera placed overhead recording physical interaction with the display.

Participants answered all questions with one of the visualizations before moving on to the next design. There were three sets of tasks and datasets that were isomorphic. The tasks were presented in the same order to each participant, but each used a different ordering of visualizations. Detail tasks were asked before overview tasks and a modified version of the NASA TLX was used after each type of task. This asked users to rate the mental demand, physical demand, overall effort, perceived performance, and frustration for the previous tasks. Upon completion of the experiment the users were asked to subjectively compare visualizations.

Participants

There were 9 male and 3 female participants in this study. Most (10) were undergraduate and graduate computer science majors. Participants were recruited from a graduate level information visualization class; hence they were relatively experienced visualization users. No reimbursement was given for their voluntary participation. Participants were randomly assigned to either the 2 or 32 Mp condition, with one female in the 32 Mp condition and two females in the 2 Mp condition.

On a pre-experiment questionnaire participants were asked to rate their familiarity with computers, large displays, information visualization, and geographic information systems on a scale from 1-5 with 1 being strongly disagree and 5 being strongly agree. The mean familiarity ratings for the two groups of subjects are shown in Table 5.3. None of the reported differences in familiarity between the 2 and 32 Mp groups were statistically significant.

Table 5.3. Mean familiarity ratings per display condition

(1 = strongly disagree, 5 = strongly agree)

	Computer	Large Displays	Information Visualization	Geographic Information Systems
2 Mp	4.83	4.25	4.42	3.75
32 Mp	4.7	4.2	4.5	4.0

5.2.2. Results

Performance Time

The task completion times were first normalized by the number of data attributes so that a meaningful comparison between display sizes could be made. This meant that the times in the 2 Mp condition were divided by 9 and the times in the 32 Mp condition were divided by 196 (the respective numbers of attributes). Without normalizing the data it would be expected and unsurprising to see that all 32 Mp times were significantly longer than all 2 Mp times. Normalizing the data allowed for a fair comparison. While the rest of the statistics are done using the normalized times, the actual mean task completion times are shown in Table 5.4.

A 3-way mixed model ANOVA with visualization and task being within subjects factors and display size being a between subjects factor was used to analyze normalized task completion times per attribute. There was a significant 3-way interaction between visualization, display size, and task. Also relevant was a significant main effect by display size $F(1,10)=54.67$, $p<0.001$, such that the normalized time per attribute on the 32 Mp display (0.11s) was faster than on the 2 Mp display (0.84s). Note that although the normalized time per attribute was faster, the actual task completion time was longer (see Figure 5.9). While there was more than a 20x increase in data size (from 9 to 196 attributes), there was less than a 3x increase in task completion times (from 7.54s to 21.25s).

Table 5.4. Task completion time means (seconds)

	2 Mp			32 Mp		
	Mults	Bars	Graphs	Mults	Bars	Graphs
DT	6.85	5.31	4.64	33.10	11.31	22.64
DA	5.59	3.96	4.13	18.84	15.84	11.50
DS	3.34	18.75	21.35	16.04	63.27	53.08
OT	14.17	8.40	4.86	27.96	11.39	6.07
OA	6.78	5.89	6.22	21.15	17.59	18.76
OS	3.99	2.05	7.44	18.94	7.82	17.36
OST	12.79	6.39	5.43	29.04	12.96	11.54

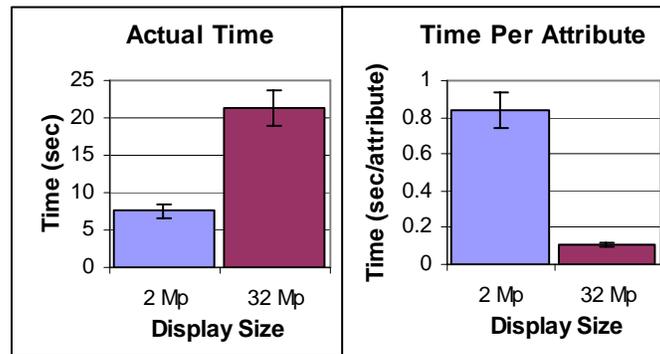


Figure 5.9. Task completion times

Because of the significant 3-way interaction, we first looked for display size by visualization interactions for each task. We next considered, for each display size, the visualization by task interactions. Post-hoc analysis was done using Tukey’s HSD. Visual representations of all results are shown in Figure 5.10 and Figure 5.11.

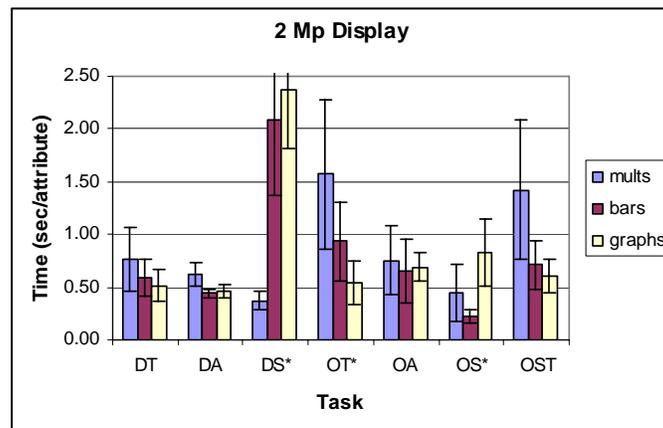


Figure 5.10. 2 Mp differences in time per attribute. Bars are 95% confidence intervals. Tasks with significant differences are marked with a ‘*’, and have non-overlapping bars.

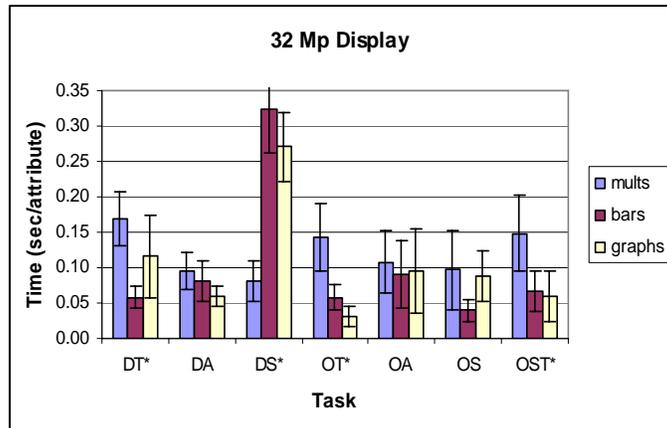


Figure 5.11. 32 Mp differences in time per attribute. Bars are 95% confidence intervals. Tasks with significant differences are marked with a “*”, and have non-overlapping bars.

There were significant display size by visualization interactions for the DS task $F(2,20)=11.49$, $p<0.001$, the OT task $F(2,20)=4.11$, $p=0.032$, and the OS task $F(2,20)=5.19$, $p=0.015$. These interactions are shown in Figure 5.12, Figure 5.13, and Figure 5.14. The space-centric/embedded visualizations saw the most improvement on the DS task with the large display. Graphs saw an 8.76x improvement and bars had a 6.45x improvement while mults only had a 4.54x improvement in time per attribute when moving from the 2 Mp to 32 Mp display. This is likely because finding a location when given a year and attribute was already perfectly suited for mults so there was less room for improvement. The decrease in time for the embedded visualizations may be a result of users mentally filtering a greater portion of the information by remembering where the year x-coordinate was with respect to an embedded visualization. Some users reported remembering the relative location of an attribute and using that to filter as they scanned the embedded visualization for each state. However, while graphs and bars did see a greater percentage of improvement, mults was still significantly faster than both bars and graphs for the DS task, regardless of display size (see Table 5.5 for p values).

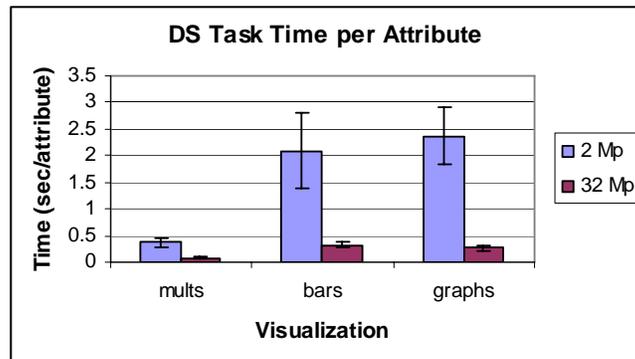


Figure 5.12. Interaction between display and visualization for DS task

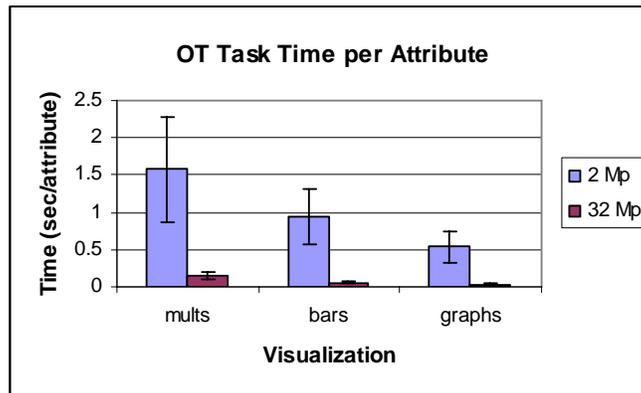


Figure 5.13. Interaction between display and visualization for OT task

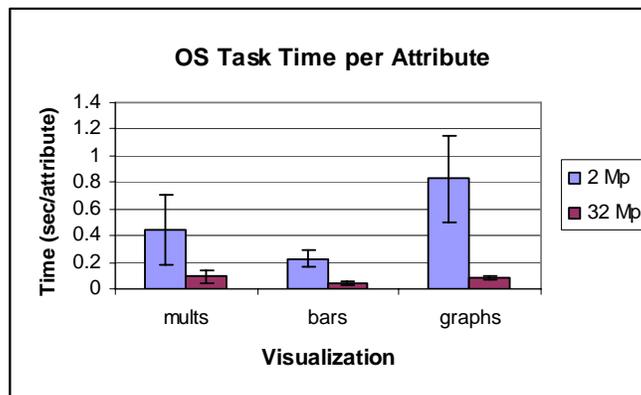


Figure 5.14. Interaction between display and visualization for OS task

For identifying trends across time (OT), mults saw a great improvement on the large display (11x), but improved less than bars (16x) and graphs (17x). However, while mults was only significantly slower than graphs on the small display, it was significantly slower than both graphs and bars on the large display. Again, the relative comparisons still held.

For the OS task, which required identifying a spatial trend, graphs improved the most (9.34x compared to 5.72x for bars and 4.59x for mults). In the 2 Mp condition, bars were significantly faster than graphs for the OS task. This was the only significant difference between the bar encoding and the line encoding. However, this difference was no longer significant in the 32 Mp condition. This is also the only significant difference that appeared on the small display but not on the large display. This suggests that the encoding may play a greater role on the smaller display while the spatial grouping is of greater importance on the larger display.

An additional observation from Table 5.5 is that for all three tasks involving time (DT, OT, and OST), the advantages of bars over mults did not show up on the small display but did appear on the large display. This also suggests that the spatial grouping is of

greater importance on the larger display and encoding on the smaller display since grouping was the only difference between these designs.

Table 5.5. Significant normalized time differences

	2 Mp	32 Mp
DT	-----	MULTS(0.17s) >BARS(0.06s,p=0.0206)
DA	-----	-----
DS	MULTS (0.37s) <BARS (2.08s, p=0.0042) <GRAPHS (2.37s, p=0.0014)	MULTS(0.08s) <BARS (0.32s, p<0.0001) <GRAPHS (0.27s, p=0.0001)
OT	MULTS (1.57s) >GRAPHS (0.54s), p=0.0237)	MULTS (0.14s) >BARS (0.06s, p=0.0035) >GRAPHS (0.03s, p=0.0005)
OA	-----	-----
OS	BARS (0.23s) <GRAPHS (0.83s, p=0.0153)	-----
OST	-----	MULTS(0.15s) >BARS (0.07s,p=0.0063) >GRAPHS (0.06s,p=0.0036)

Accuracy

In general, users were able to correctly answer almost all of the questions; the actual numbers are shown in Table 5.6. For accuracy there was a visualization by task interaction (p=0.001) along with a main effect by visualization and main effect by task, but no other interactions. The difference between display sizes was not significant (p=0.312). The only significant differences occurred in task DS such that graphs (67%) were significantly less accurate than both bars (92%, p=0.0192) and mults (100%, p=0.0019). This was likely the result of the integration of y-axes and perhaps difficulty in distinguishing colors.

Table 5.6. Total correct answers (12 max)

	2 Mp			32 Mp		
	Mults	Bars	Graphs	Mults	Bars	Graphs
DT	12	12	10	9	12	8
DA	12	12	11	11	12	12
DS	12	10	7	12	12	9
OT	11	11	12	10	12	11
OA	12	12	12	11	11	11
OS	12	12	12	12	12	10
OST	12	12	12	10	12	12

Task Workload

Users reported mental demand, physical demand, effort, perceived performance, and frustration on rating scales from 1 to 10 after all detail tasks with each visualization and again after all overview tasks with each visualization. A MANOVA showed all

interactions and main effects as significant. Therefore, univariate analysis was done followed by post-hoc analysis using Tukey’s HSD for individual comparisons. The mean scores for each condition are shown in Table 5.7.

Table 5.7. Task workload means

DETAIL TASKS						
	2 Mp			32 Mp		
	Mults	Bars	Graphs	Mults	Bars	Graphs
Mental	3.08	3.83	7.00	6.08	5.42	7.50
Physical	3.00	3.00	4.08	8.00	6.67	5.92
Effort	3.42	3.42	6.08	7.08	6.33	7.42
Performance	8.17	7.50	6.67	8.25	8.25	7.75
Frustration	0.83	1.00	3.75	5.67	3.92	5.08
OVERVIEW TASKS						
	2 Mp			32 Mp		
	Mults	Bars	Graphs	Mults	Bars	Graphs
Mental	4.75	3.17	3.42	7.25	4.67	4.42
Physical	3.08	2.25	3.33	5.75	2.92	2.67
Effort	4.67	2.75	3.17	7.08	3.92	3.58
Performance	7.50	8.00	7.75	6.67	8.92	8.25
Frustration	1.50	0.75	1.08	5.58	2.25	2.83

Overall, the large display users were significantly more frustrated (4.22 vs. 1.49, $p=0.019$) and reported more physical demand (5.32 vs. 3.13, $p=0.055$) than the small display users (see Figure 5.15). The only visualization specific differences that were seen were between mults on the different display sizes. Users in the large display condition reported more mental demand (6.08 vs. 3.08, $p=0.043$) and physical demand (6.88 vs. 3.04, $p=0.007$) for mults than users in the small display condition. Within the small display conditions, users reported more physical demand with graphs (3.71) than with bars (2.63, $p=0.0024$). Within the large display condition, users reported more physical demand from mults (6.88) than graphs (4.29, $p=0.0272$).

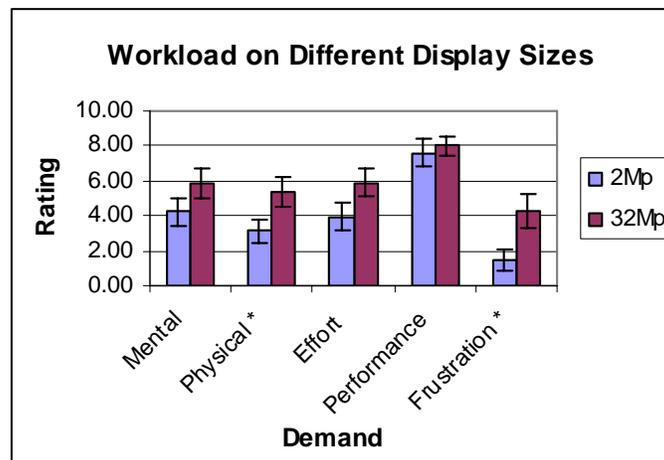


Figure 5.15. Workload on display sizes. Significant differences have a ‘*’.

There were also some differences between visualizations based on the task. For detail tasks, users reported more mental demand with graphs (7.25) than both bars (4.63, $p < 0.001$) and mults (4.85, < 0.001) as well as more effort with graphs (6.75) compared to both bars (4.88, $p = 0.0116$) and mults (5.25, $p = 0.0468$). For overview tasks, mults required more effort (5.88 vs 3.33 for bars and 3.38 for graphs, $p < 0.001$), had a lower perceived performance compared to bars (7.08 vs. 8.46 for bars, $p = 0.0256$) and also had higher frustration levels (3.45 vs. 1.5 for bars, $p = 0.03$). These findings match the proximity compatibility principle [142] in that the embedded time series graphs required more mental demand and effort for detail tasks (where mental separation would need to occur), and for overview tasks mults required more effort, resulted in more frustration, and lower perceived performance (tasks where integration was necessary).

User Preference

User preference was different based on the display size (see Figure 5.16 and Figure 5.17). In the 2 Mp condition, four of five users preferred mults followed by bars and then graphs (the 6th user gave a circular answer). However, in the 32 Mp condition four of five users preferred bars first followed by graphs and then mults. Users always preferred bars to graphs, but mults was the most preferred on the smaller display and the least preferred on the large display. Therefore user preference for visualizations was different based on the display size/amount of data.

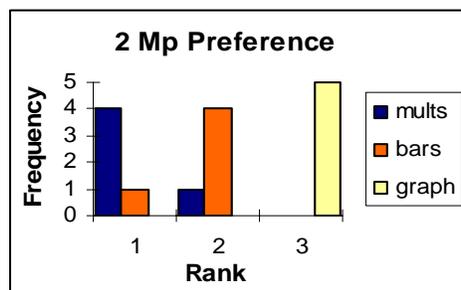


Figure 5.16. Frequency of user visualization preference for 2 Mp display (rank of 1=best, 3=worst)

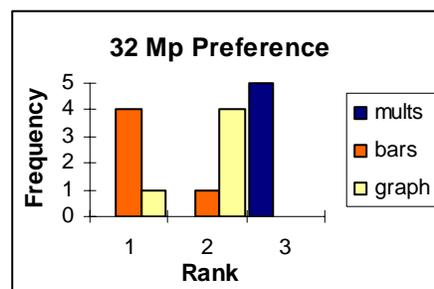


Figure 5.17. Frequency of user visualization preference for 32 Mp display

5.2.3. Discussion

As large displays continue to decrease in cost, we must continue to explore how various visualizations scale for these large displays. In this study, although overall task time increased, the time per number of attributes actually decreased. There was more than a 20x increase in data size for the large display, but less than a 3x increase in task completion times. This means that overall, across tasks and visualizations, these designs were perceptually scalable in terms of time. We also did not find a significant decrease in accuracy on the large display, with a change only from 95% to 92%, again suggesting these designs are perceptually scalable. Despite having 32 Mp, we did not hit the limits of visualization. Users were able to successfully physically navigate to complete detail tasks and were also able to gain a large overview and perceptually integrate information across a display almost 9 feet in width.

It also appears that in general, based on time and accuracy, the relative comparison of these three designs was consistent across display sizes. However, there were some differences. The fact that the only time difference between the two space-centric embedded visualizations (bars and graphs) disappeared on the large display, and that the difference between the attribute-centric approach (mults) and the structure-centric approach with the same visual encoding (bars) showed up on almost every task on the large display suggests that the line vs. bar encoding was most important on the small display. This matches our previous study in suggesting that visual encoding can be more important than grouping on small displays [148]. However, as the visualization was scaled-up using the display, the spatial grouping became much more important. This spatial grouping likely increased visual aggregation and reduced the amount of physical navigation. This was echoed both by the shift in user preference from the attribute-centric visualization (mults) on the 2 Mp display to space-centric embedded visualization with the same encoding (bars) on the 32 Mp display and also by the significant increase with the attribute-centric design in physical and mental effort. Results also support the proximity compatibility principle [142] and go against the idea that the attribute-centric approach might show the most benefit when more than 16 views are displayed [81]. This shows great promise for using space-centric embedded visualizations for geospatially-referenced data on large displays.

5.2.4. Conclusion

In this experiment we compared three different visualizations on both a small and large display. Results showed that the designs used were perceptually scalable – not resulting in an increase in normalized performance time or a significant decrease in accuracy. Accuracy only decreased from 95% to 92%, and a 20x increase in data resulted in only a 3x increase in task completion times. Using a combination of perceptual abilities and physical navigation people were able to effectively use a 32 Mp display (2x Ware’s proposed 16 Mp display [139]).

In terms of tasks, the space-centric embedded visualization using the bar matrices was significantly better than the attribute-centric design for temporal overview (OT), spatiotemporal overview (OST), and temporal detail (DT) finding tasks on the large

display. The attribute-centric design resulted in better performance for the spatial detail task (DS) on both display sizes.

Results also showed that relative comparisons between visualizations with respect to time and accuracy were typically the same regardless of the display size. However, based on user preference and workload, graphical encoding seems to be more important with less data on a small display whereas spatial grouping seems to be more important with more data on a large display. On the large display, both space-centric embedded visualization designs were generally significantly faster than the attribute-centric visualization. User preferences also switched on large displays, with most users preferring both of the space-centric embedded visualizations.

5.3. Experiment 5: Scaling Beyond the Limits of Visual Acuity

Using large, high resolution displays means that graphical scalability (the number of pixels required to visualize a dataset) is less of a limitation. In the last experiment we saw that the chosen visualizations were indeed perceptually scalable. However, one possibility is that visual acuity imposes a perceptual scalability limit. If the dots per inch (DPI) of a display is so high that an individual pixel cannot be seen, regardless of how close a user is to a display, then an even higher DPI is unlikely to be beneficial. More interesting is the situation in which pixels are of a sufficient size, the DPI is constant, and the number of pixels is increased using a larger display. This can result in a total display resolution such that the user cannot see all of the pixels from any particular location without walking. In this case we refer to the display as being beyond visual acuity. This has led to the argument that a display with a resolution equal to visual acuity should be adequate for any single user visualization task (Figure 5.18) [139, 140].

In this experiment we explored the influence of using a display with enough pixels to exceed visual acuity when scaling up information visualizations. This experiment compared two different visualization designs that were scaled-up to larger datasets (with data density kept constant) using display sizes with enough pixels to be within, roughly equal to, and beyond visual acuity.

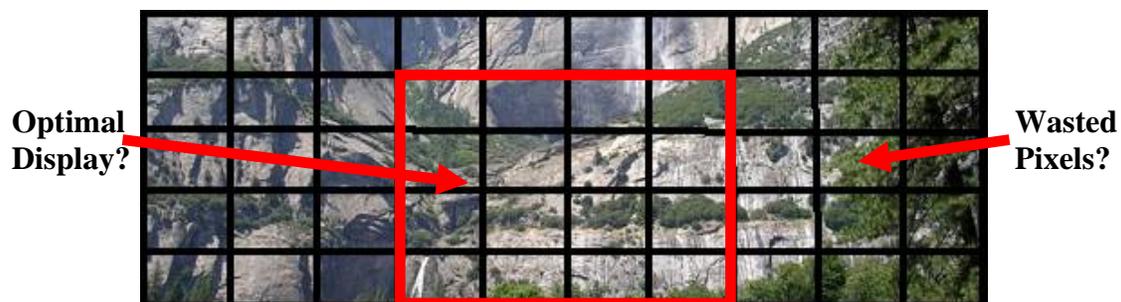


Figure 5.18: Does an optimal display match visual acuity?

The main goal of this study was to explore the effects of using displays that exceed the limits of visual acuity for scaling up information visualizations. In general, we wanted to know *how scaling up visualizations using a display with a resolution exceeding visual acuity would affect user performance*. As a secondary issue we wanted to confirm that the selected visualization designs were still perceptually scalable with significantly more data.

5.3.1. Method

Visualizations

Before conducting the study we first had to design the visualizations for the large display. Again, we were interested in visualizing datasets with a combination of spatial, temporal, and multidimensional data because of their prevalence in many application areas as diverse as sociology, epidemiology, and bioinformatics. The exact data used was artificially generated and consisted of spatial (U.S. states), multidimensional (multiple demographics groups), and temporal (multiple years) data.

The two visualization strategies will once again be referred to using the design space terminology. The attribute-centric visualization corresponds with multiple views and the space-centric visualization corresponds with an integrated view visualization (Figure 5.19). The attribute-centric design had a separate small map for each year/group combination. The space-centric visualization consisted of a single large map with embedded visualization overlaid at each location. In order to prevent the effect of previous knowledge, we used only group numbers instead of meaningful attribute names and abbreviated state labels according to their region. For example, West state labels included w1 and w2. Each individual bar was 2 pixels wide and between 0 and 10 pixels high. A continuous color scale corresponding to the bar height was used as a redundant encoding.

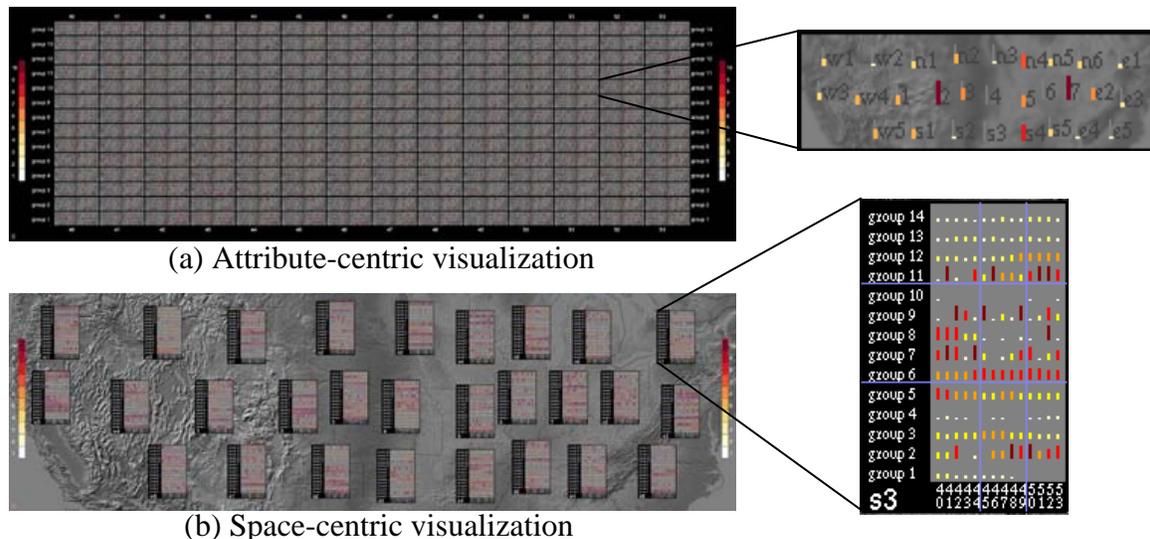


Figure 5.19. Visualization designs used in visual acuity experiment

In the previous experiment [149] we saw that the spatial grouping with the space-centric visualization scaled better. In this experiment we set the amount of data in the smallest display condition equal to the amount of data in the previous experiment's largest display condition. This allowed us to hypothesize as to whether the advantage was due to the size of the display or the amount of data.

Participants

Eighteen volunteers (14 male, 4 female) participated in this study. Most were computer science students at a large public university recruited from HCI related classes.

Equipment

An 8x3 tiled display of LCD monitors was used, the same display used in the previous experiment. Each monitor was a 17 inch LCD with a resolution of 1280x1024 (~96DPI) and a dot pitch of 0.264mm. The total resolution of the display was 10,240x3,072 (approximately 32 Mp) and was 9 feet wide by 3.5 feet high. The bottom of the display was 36.75 inches off the floor. A wider display allows users to access pixels by walking. We avoided issues with the bezels between tiled monitors by ensuring that information never straddled monitor boundaries.

Design

We used a 3x2x9 design with the following independent variables: display size (small, medium, large corresponding to within, equal to, and beyond visual acuity - data density was constant), visualization design (space-centric, attribute-centric), and task (3 detail, 4 overview, and 2 complex). The visualizations were described in a previous section. Display size was a between subjects variable while visualization and task were within subject variables. Task completion time, accuracy, subjective workload, and user preference were recorded.

Display Size and Visual Acuity

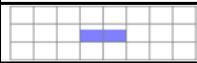
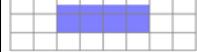
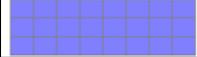
Various interdependent perceptual factors are at play when using large displays. Eye and head movements are used to bring objects into foveal vision. The effect of these movements on the distance and viewing angle parameters are relatively minor compared to the effect of walking. When walking is required to bring objects into foveal vision then the distance from and viewing angle for other objects changes more dramatically. The field of view also increases or decreases based on a user's physical distance from the display.

By *visual acuity* we mean the definition commonly used when referring to 20/20 vision - the ability to distinguish two points (point acuity) [139]. This definition is based on both distance from and the size of a target. Those two parameters along with the viewing angle are then used to calculate the visual angle subtended by a target. For a fixed target size this means that targets that are either too far away or viewed from too great a viewing angle will be "beyond visual acuity" because the visual angle subtended will be too small (< 1 minute of arc, 1/60 degree). In the large display condition, there was no spot where

everything could be distinguished with only eye and head movements because either the distance or the viewing angle would be too great. For example, if users stepped back for an overview, they would not be able to see individual elements. If they stepped forward for a detailed view, they could see a small region of elements, but not the elements on the far side of the display.

Specifications for each display size are shown in Table 5.8. The aspect ratio was kept constant across display sizes by only using part of the display height for the small and medium display conditions. The display size condition that was roughly equal to visual acuity (medium size) is shown in Figure 5.20. In the medium display condition the user started at 711.2 mm (90 degree field of view/half the width of the illuminated surface) away from the center of the display. The starting distance was not the same for all three display size conditions.

Table 5.8. Display size conditions

Display Size	Data		Display	
	Groups x Years	Data Points	Pixels	Area
Small	14 x 14	5,488	2560x768	
Medium	29 x 29	23,548	5120x1536	
Large	58 x 58	94,192	10240x3072	

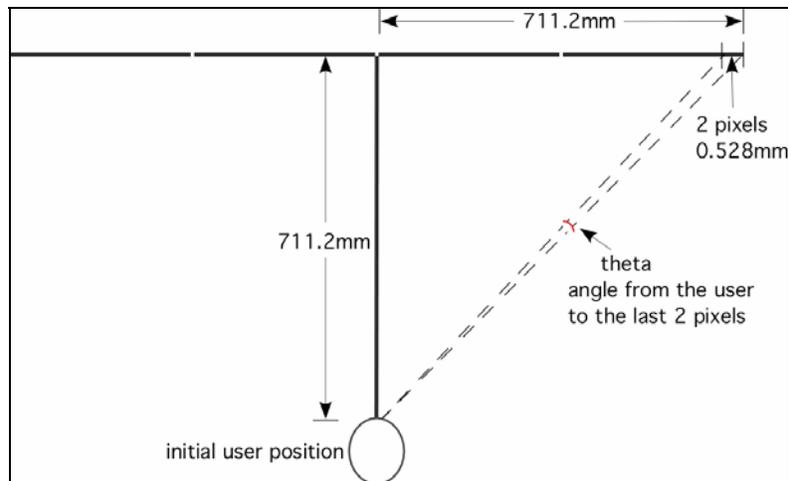


Figure 5.20. Medium display condition (equal visual acuity)

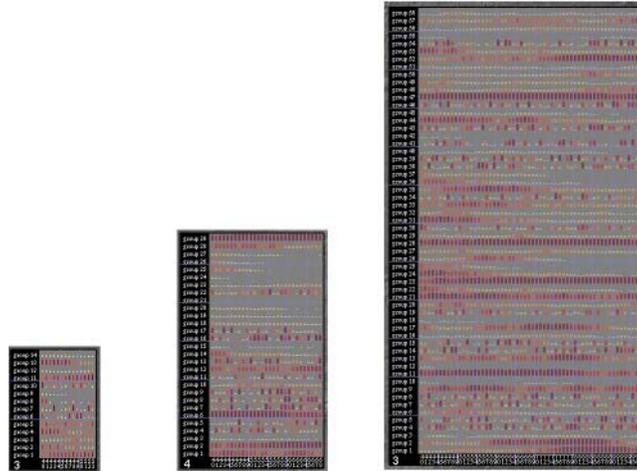


Figure 5.21. Space-centric embedded visualization on different display sizes

For the medium condition the angle from the user to the last 2 pixels, which was the size of visual elements, was 0.0212 degrees ($1.272/60$). This is slightly greater than visual acuity ($1/60 = 0.0166$ degree). This was to compensate for the fact that the last two pixels were viewed at an angle and may be harder to see. As the user moved closer, the angle was too great to see the outermost pixels. As a user moved farther away, the pixels became too small to distinguish.

The data density and size of individual graphical elements were kept constant across display sizes, while the amount of data increased proportionally with display size (Figure 5.21).

Tasks

The tasks are described in Table 5.9. An artificial dataset was generated consisting of 58 time points and 58 attributes for 28 states. This choice of data was modeled after typical demographics datasets available from sources such as the Census Bureau. Data generation provided for an additional level of experimental control.

The datasets had to be created carefully as to result in a fair comparison. Detail task datasets for different types were generated similarly; there was a single definite highest value that was different from the next highest value by either 2 or 4 pixels. For OA tasks, at least 60% of the values were within the specified range. Generation of OT datasets was a little more complex; 4/28 states did not match the pattern and for the remaining 24 states 80% of the groups matched the pattern. Up then down (or down then up): states matching the pattern were split into 4 groups with different inflection points. For OST tasks, states in the fastest increasing area had initial values between 4.5 and 6.5 and increased by a number between -0.3 and 1.2 each year, whereas the rest of the state values changed by only one pixel. For OS tasks, states in the highest area had values between 7 and 10, with some noise added. A correlation was created between two of the groups for CA tasks. An increase/decrease in one group caused an increase/decrease in another group. For CS tasks, matching states were exactly the same, but participants were not told.

Table 5.9. Tasks (D: detail, O: overview, and C: complex; A: attribute, T: time, and S: space)

	Type
Detail	DA: Find a group, given a year and location DT: Find a year, given a group and location DS: Find a location, given a group and year
Overview	OA: Populations are in the range [$<4,4-7,>7$] OT: Population have gone [up,up/down, down/up, down] OS: Populations are highest in [N, S, W, E] OST: Populations increase fastest in [N,S,W,E]
Complex	CA: Cause/effect relationship CS: Find a location that is similar to a given location

Procedure

Participants were assigned to a display size condition based on their self reported familiarity with computers, large displays, information visualization, and geographic information systems. The groups were kept as even as possible in terms of these factors. The average familiarity ratings for each group can be seen in Table 5.10. None of the differences between groups was statistically significantly. The eye sight of participants was not tested and they were only asked if they had corrected vision. We assumed that participants had approximately 20/20 vision and considered that to be a reasonable assumption for the sake of generalizing the results to the overall population.

Table 5.10. User familiarity ratings (1=strongly disagreed, 5=strongly agreed).

Display Size (visual acuity)	Computer	Large Display	Information Visualization	Geographic Information Systems
Small	4.8	3.5	2.7	2.0
Medium	5.0	3.3	3.2	3.3
Beyond	5.0	3.2	2.8	3.0

Participants stood and started at a viewing distance equal to half the width of the illuminated surface (different for each display size). This meant each participant started with a 90 degree field of view. A mark was placed on the floor where users were asked to stand when they started a task. After beginning the task, users were free to move about.

The participants went through a training session to get familiar with the visualizations before beginning. Because visualization was a within subjects factor, half of the participants in each condition used the space-centric first. Before each type of detail task they did a practice task. Each type of task was then performed twice with a given visualization. This meant that each participant performed 2 (visualizations) x 9 (task types) x 2 (trials) = 36 experimental tasks.

After each group of tasks (DA and DT, DS, Overview, CA, and CS) participants answered five questions which were a modified version of the NASA Task Load Index. The questions asked the participants to rate the level of mental demand, physical demand,

effort, perceived performance, and frustration on a scale from 0 to 10. At the end of the experiment they filled out a post-hoc questionnaire indicating their visualization preference for each task.

5.3.2. Results

Time and accuracy averages can be found in Table 5.11 and Table 5.12. Time, accuracy, and workload data was first analyzed using a three-way mixed model ANOVA with display as a between subjects factor and visualization and task as within subjects factors. Post-hoc two-way ANOVAs were then done for each task and significant factors were further compared using Tukey’s HSD. Because task had such a strong influence, we only report the results by task rather than presenting overall results.

Table 5.11. Average task completion times (seconds)

Task	Attribute-Centric			Space-Centric		
	Small	Medium	Large	Small	Medium	Large
DA	10.98	28.75	56.73	7.68	14.43	21.25
DT	15.73	39.53	81.04	8.29	8.48	14.88
DS	6.15	8.74	9.76	44.69	63.68	81.60
OA	6.10	11.24	27.54	9.07	10.15	20.00
OT	25.20	40.48	57.04	27.22	27.87	26.61
OS	10.84	28.75	67.48	3.59	15.37	19.21
OST	22.95	31.21	40.81	26.42	22.93	12.70
CA	52.78	55.16	73.92	60.68	69.68	166.50
CS	52.71	56.34	61.14	29.57	46.75	74.83

Table 5.12. Accuracy (percent correct)

Task	Attribute-Centric			Space-Centric		
	Small	Medium	Large	Small	Medium	Large
DA	100.0	91.7	75.0	91.7	91.7	100.0
DT	91.7	91.7	91.7	83.3	100.0	91.7
DS	91.7	100.0	91.7	91.7	100.0	100.0
OA	100.0	83.3	83.3	100.0	83.3	75.0
OT	0.0	66.7	58.3	33.3	75.0	100.0
OS	100.0	83.3	100.0	100.0	100.0	100.0
OST	83.3	100.0	91.7	75.0	100.0	100.0
CA	16.7	83.3	58.3	25.0	83.3	33.3
CS	58.3	75.0	91.7	83.3	91.7	83.3

Increase Factor for Time

The increases in time between the small and medium, and the medium and large display conditions are shown in Figure 5.22. This is a comparison between how quickly time is increasing as the display gets larger. The display size (and amount of data) increased by a factor of four between conditions (small x 4 = medium, medium x 4 = large). As the

display got larger, almost every task (DA, DS, OT, OS, OST) resulted in a less than proportional increase in time (<4x). The OA and CA tasks were the only tasks where users were less efficient in terms of time as the display got larger (medium/small < large/medium).

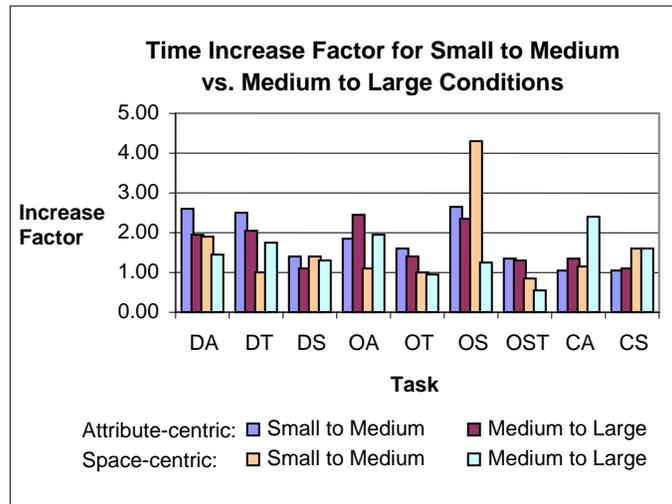


Figure 5.22. Factor increases in task completion times

Detail Tasks

The DA (given 1 year, comparing groups) and DT (given 1 group, comparing years) tasks were similar because the number of groups and years increased with display size. We normalized the task completion times for these tasks by dividing the actual task completions times by the number of groups (14, 29, or 58). The DS task required finding a location with the highest value when given a group and a year. Because the number of locations was fixed, the data for this task was not normalized.

DA and DT Tasks

Overall the space-centric visualization resulted in better user performance than the attribute-centric visualization for the DA and DT tasks. It was significantly faster (DA task: $F(1,15)=30.46$, $p<0.01$; DT task: $F(1,15)=67.44$, $p<0.01$), physical workload was lower (3.11 vs. 5.33, $p<0.01$), perceived performance was better (8.53 vs. 7.6, $p=0.043$), and frustration was less (1.56 vs. 3.39, $p<0.01$). Users also preferred the space-centric visualization for these tasks (DA task: small-3/6, medium-4/6, large-5/6; DT task: small-5/6, medium-5/6, large-5/6). There were no significant differences in terms of accuracy.

The only significant difference related to the display size was that the space-centric visualization was less mentally demanding than the attribute-centric visualization (2.83 vs. 7.08, $p<0.01$) in the large display condition.

DS Task

The attribute-centric visualization resulted in better user performance than space-centric for this task – opposite of the DA and DT tasks. Users were significantly faster with the attribute-centric visualization (significant display size x visualization interaction $F(2,15)=8.59$, $p<0.01$, main effect by visualization $p<0.01$, main effect by size $p<0.01$, faster for every display size $p<0.01$). The attribute-centric visualization also resulted in significantly less mental workload (1.94 vs. 6.54), less physical workload (1.89 vs. 5.89), less effort (1.54 vs. 6.43), better perceived performance (9.39 vs. 7.80), and less frustration (1 vs. 2.8) than space-centric (all $p<0.01$), and users always preferred it (small: 5/6, medium: 4/6, large: 5/6). There were no significant differences in terms of accuracy.

In terms of display size, the only significant difference was not surprising; The space-centric visualization was significantly faster in the small display condition (44.69s vs 81.6s) compared to the large display condition ($p<0.01$).

Overview Tasks

OA Task

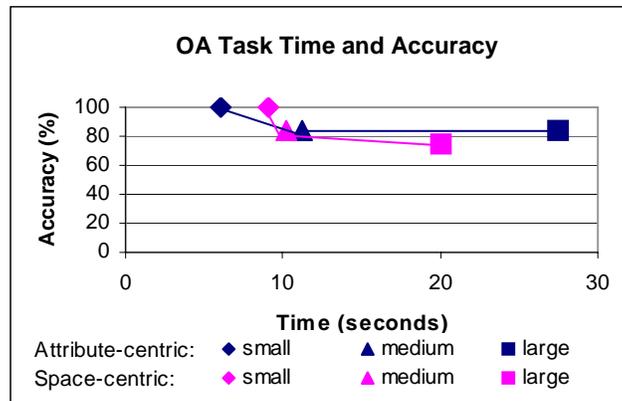


Figure 5.23. Speed and accuracy for the OA task. Small is significantly faster than large and significantly more accurate than both medium and large.

The OA task required participants to see the general range of values. The visualizations were significantly different only in the large display condition where the space-centric visualization was significantly faster than the attribute-centric visualization (20s vs. 27.54s, $p=0.05$). User preference was nearly equally split between the two visualizations (small: 2/6, medium: 3/6, large: 3/6 preferred space-centric).

Increasing the display size seemed to have a negative effect on user performance for this task (Figure 5.23). Not only were the small and medium sizes faster than the large (display x visualization interaction $F(2,15)=7.09$, $p<0.01$, main effect by display size $F(2,15)=11.45$, $p<0.01$, small vs. large and medium vs. large $p<0.01$), but time increased faster between the medium and large displays than between the small and medium displays (1.12x to 1.97x, attribute-centric: 1.84x to 2.45x). The small display condition

(100%) was also significantly more accurate (main effect by size $F(2,15)=4.2$, $p=0.036$) than both the medium (83.3%, $p=0.01$) and large conditions (79.17%, $p=0.022$).

OT Task

The OT task required the user to find a temporal trend. The space-centric visualization resulted in better user performance than the attribute-centric visualization for this task (Figure 5.24). The space-centric visualization was significantly faster in the large display condition (display x visualization interaction $F(2,15)=3.91$, $p=0.043$, main effect by visualization $F(1,15)=8.31$, $p=0.011$, large – visualization comparison $p=0.0247$), and significantly more accurate overall (main effect by visualization $F(1,15)=5.56$, $p=0.032$). Somewhat surprisingly, no participant correctly answered this question using the attribute-centric visualization on the small display. However, the results were better than chance (33%) with the space-centric approach on the small display. This suggests that users have a particularly difficult time identifying temporal patterns with the attribute-centric visualization, but that additional data makes the task easier. Users also preferred the space-centric visualization for this task (small:5/6,medium:6/6, large:5/6).

In terms of display size, both the medium (70.83%) and large (79.17%) sizes were significantly more accurate than the small display condition (16.67%, $p<0.001$).

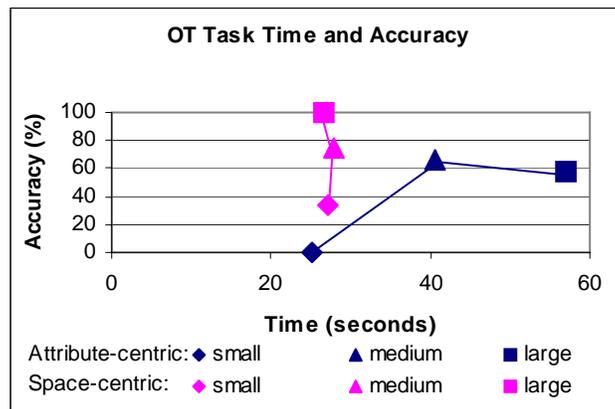


Figure 5.24. Speed and accuracy for the OT task. Accuracy was significantly better with medium and large than small. Space-centric was significantly faster than attribute-centric in the large condition.

OS Task

The OS task required participants to find a spatial trend in the data. The space-centric visualization was better than the attribute-centric visualization for this task. It was significantly faster in every display condition (visualization x display interaction $F(2,15)=3.91$, $p=0.043$, visualization main effect $F(1,15)=8.31$, $p=0.01$, comparison was $p<0.01$).

Not surprisingly, the small condition was significantly faster than the large condition for both visualizations, and the medium condition was faster than the large condition for the space-centric visualization.

OST Task

The OST task was about finding a spatiotemporal pattern. The space-centric visualization appears to be preferable for this task. It resulted in a non-significant decrease in task completion times on the larger displays (Figure 9), and was preferred over the attribute-centric visualization (small: 5/6, medium: 6/6, large: 5/6).

In terms of display size, the medium display (100%) was significantly more accurate than the small display (79.2%, main effect by size $F(2,15)=4.77$, $p=0.025$, $p=0.022$).

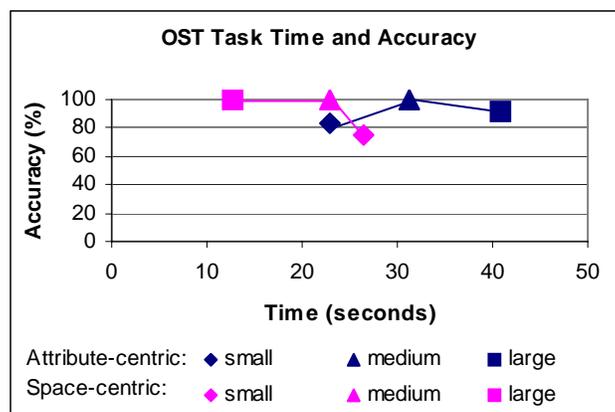


Figure 5.25. Speed and accuracy for the OST task. Medium is more accurate than small.

Workload

Workload measurements were recorded for all overview tasks combined. Overall, the space-centric visualization was better than the attribute-centric visualization for overview tasks. It resulted in less mental demand ($F(1,15)=10.17$, $p<0.01$, 5.18 vs. 7.11, $p<0.01$) and less effort ($F(1,15)=8.9$, $p<0.01$, 4.94 vs. 6.78, $p<0.01$).

In the large display condition the space-centric visualization resulted in less physical demand than the attribute-centric visualization ($F(2,15)=6.64$, $p<0.01$, $p<0.01$). Performance was also perceived to be better ($F(2,15)=4.64$, $p=0.048$) with the space-centric visualization in the medium ($p=0.011$) and large conditions ($p=0.026$).

Complex Tasks

The complex tasks were harder than any of the other tasks because they required complex pattern matching between states or finding correlations between groups across time.

CA Task

The CA task required finding a cause/effect relationship. The attribute-centric visualization was significantly faster than the space-centric visualization (visualization main effect $F(1,15)=7.26$, $p=0.017$). Despite attribute-centric being faster and more accurate, in every display condition 4 of 6 users preferred the space-centric visualization.

In terms of display size, showing more data using the medium size display improved performance, but performance dropped-off with the large display (Figure 5.26). Even though the small and medium displays were significantly faster than the large display (display size main effect $F(2,15)=4.03$, $p=0.04$, both $p<0.01$), this was the only task, in addition to OA, where time increased faster from medium to large than from small to medium. The medium display condition was significantly more accurate than both large and small (size main effect $F(2,15)=7.5$, $p<0.01$, medium vs. large $p=0.012$, medium vs. small $p<0.01$).

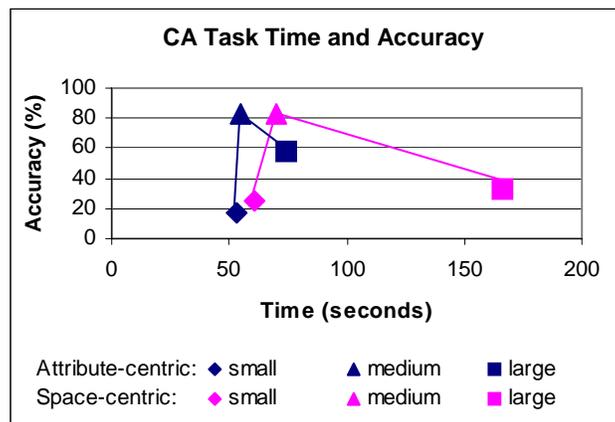


Figure 5.26. Speed and accuracy for the CA task. Small and medium were significantly faster than large, and medium was the most accurate display size.

CS Task

The CS task was a pattern matching task. There were no significant time or accuracy differences for this task. The only significant workload measurement for the CS task was for frustration. The space-centric visualization was less frustrating than the attribute-centric visualization overall (visualization main effect $F(1,15)=5.5$, $p=0.033$, 2.1 vs. 3.36) and particularly when using the small display (interaction $F(2,15)=6.03$, $p=0.012$, comparison $p=0.028$). Most users preferred the space-centric visualization for this task (small: 5/6, medium: 3/3, large: 4/6). For the attribute-centric visualization the large display was significantly less frustrating than the small display ($p<0.01$).

Summary of Results

Detail Tasks

The results from the detail tasks suggest that using a display beyond visual acuity has very little influence on these tasks. There were no significant differences in accuracy or normalized completion times for the DA or DT tasks between display sizes. The space-centric visualization was ideally suited for the DA and DT tasks while the attribute-centric visualization was ideally suited the DS task.

Overview Tasks

Results from the overview tasks suggest that the effect of using a display that exceeds visual acuity depends on the task. The additional data that could be visualized using the larger display resulted in improved accuracy for the overview tasks involving time and space. On the other hand, for the OA task, there was a decrease in accuracy with the larger display. The space-centric visualization was particularly good for overview tasks on the large display.

Complex Tasks

The complex task results revealed a situation where the attribute-centric visualization resulted in better user performance. The CA task, which required finding the strongest correlation, resulted in faster task completion times with the attribute-centric visualization than with the space-centric visualization. Despite this, users still preferred the space-centric visualization.

5.3.3. Discussion

The main goal of this study was to explore visual acuity as a factor in the perceptual scalability of information visualizations for large, high resolution displays. We wanted to know what would happen to user performance as the data was increased using a display with enough pixels that physical navigation was required to see everything. The results suggest that for some tasks performance can actually become more efficient and accurate when showing more data using larger displays. This is supported by the findings that almost all tasks resulted in less than proportional time increases and using the large display resulted in accuracy increases as much as 20% for temporal pattern tasks. These results agree with the previous study [149] in suggesting that visualizations can be perceptually scalable. It adds to that study by showing that this still holds with a larger dataset, when 19x more data is packed on the same size display.

A secondary aim for this study was to confirm the perceptual scalability of the visualizations. As in our previous work, we found that in general the space-centric visualization resulted in better user performance than the attribute-centric visualization. This was particularly true in the largest display condition. In three of four overview tasks the attribute-centric visualization was faster than the space-centric visualization in the small display condition, but in the medium and large display conditions the space-centric

embedded visualization was faster. The task also played a role in this scaling with the advantages of the space-centric being most evident on spatial and temporal overview tasks.

Overall, the role of visual acuity seemed to be outweighed by the advantages of additional data. In five out of nine tasks time increased faster between the smaller displays than between the larger display sizes. In general users were actually becoming more efficient using the additional data, despite the physical navigation that was required.

One might wonder why the space-centric visualization has an advantage. Recall that there was the same amount of data in the small display condition as in the large display condition from the previous experiment. Although the amount of data was the same, the advantages only became apparent on the large display. This suggests that the advantages of the space-centric visualization were more related to the display size than the amount of data. Therefore space-centric embedded visualizations seem to be even more useful on large displays than they are on smaller displays, regardless of the amount of data.

Spatially grouping information and visual aggregation are both important factors in the performance of space-centric visualizations on large displays. Visual aggregation resulted in less physical navigation and thereby allowed for less mental demand and effort. Another observation is related to information overload. Throughout the experiment, some participants stated that it felt like a lot more data was being shown with the attribute-centric visualization (when in fact data was equal). Although performance was better with the attribute-centric visualization on the CA task, more users preferred the space-centric visualization for this task. This may be because they were less intimidated by the space-centric visualization.

All of these factors support the need for additional research into visualizations that are specifically designed for large, high resolution displays. If visualizations are created that consider the display during the design phase then users can take advantage of the additional data that can be displayed on large, high resolution displays despite having to physically navigate to access all of the information. These results provide some insight into designing visualizations that scale-up for large displays.

5.3.4. Conclusion

The scalability of information visualizations has typically been limited by the number of available display pixels. This work explored the effect of using large, high resolution displays to scale-up information visualizations beyond potential visual acuity limitations. Results showed that performance on most tasks was more efficient and sometimes more accurate because of the additional data that could be displayed, despite the physical navigation that was required. Hence, visualization is not limited by visual acuity.

5.4. Summary and Synthesis of Results

The overall goal of the work presented in this chapter was to determine how visually scalable the visualizations were for large, high resolution displays. First, the graphical

scalability of attribute and space-centric (using embedded) visualizations was analyzed. Then, two experiments were performed. The first experiment compared user performance with visualization across display sizes and the second took that a step further by explicitly considering visual acuity as a potential limiting factor and increasing dataset sizes. The main findings are as follows.

Research Question 3:

How visually scalable are the visualizations for large, high resolution displays?

- a. Display Issues: Pixel Count
- b. Human Issues: Perception and Cognition

- The space-centric approach using embedded visualizations was more graphically scalable than the attribute-centric approach. It requires fewer pixels.
- Both visualizations were perceptually scalable. An increase in data using more pixels resulted in a less than proportional increase in time along with maintained accuracy.
- Visual acuity is not the limit for information visualization on large displays. User performance benefits existed for some tasks because of the additional data that could be displayed despite the physical navigation that was required.
- Space-centric embedded visualizations were more perceptually scalable than attribute-centric visualizations. User performance diverged with the increasing display size and the embedded visualization approach typically resulted in better user performance than attribute-centric visualizations on the large display.

The space-centric embedded visualization approach is not only more graphically scalable, but also more perceptually scalable. By perceptually scalable we mean that it resulted in more effective user performance in terms of time and accuracy as the amount of data being visualized was increased using more pixels. The first experiment demonstrated that neither design failed when scaled-up. Accuracy was maintained and the increase in time was less than proportional to the increase in data. The time and accuracy differences between the space and attribute-centric visualizations were similar in the first experiment, but the mental workload increased significantly only with the attribute-centric approach and user preference switched to space-centric embedded visualizations on the large display.

The second experiment demonstrated that there are user performance benefits even beyond visual acuity. Time increased at a slower rate between the larger display sizes than between the smaller display sizes for five tasks. Accuracy also generally improved or was consistent. The space-centric embedded visualization approach also diverged and became better on the large display in this experiment. Based on a cross-experiment comparison we can hypothesize that the advantages of this visualization are more related to display size than the amount of data.

Also of note is task based differences between visualizations. Detail task results were as one might expect. The space-centric embedded visualizations resulted in better user performance on spatial and temporal overview tasks. The attribute-centric visualization

was actually better for finding correlations between groups on the large display, yet users still preferred the space-centric embedded visualization for this task.

The obvious next question was why was the space-centric embedded visualization approach better on large displays? And how can we use this work to improve visualization design? In the next chapter the reasons that the space-centric embedded visualization resulted in better user performance on larger displays are discussed within the larger context of visualization design issues when using large displays.

Chapter 6. Information Visualization Design Issues on Large, High Resolution Displays

The fourth research question was: What characteristics of a visualization determine how visually scalable it is for large, high resolution displays and what design issues remain? The ultimate purpose of determining the scalability of information visualizations for large, high resolution displays is to help the designer create visualizations that are beneficial to the end user. To do this requires understanding various design issues that will be faced when creating visualizations for large, high resolution displays. In the first part of this chapter on perceptual design considerations, various design ideas that arose from the experiments are discussed. Some of the likely reasons for the advantages of the space-centric embedded visualization approach are included throughout. The second part of this chapter focuses on the scalability of interaction. It briefly covers some of the issues that are likely to be faced by information visualization designers as information visualization specific interactions are added to the interface. The chapter concludes by outlining a few remaining open questions with respect to information visualization design on large, high resolution displays.

6.1. Perceptual Design Considerations

6.1.1. Visual Encodings

Graphically Scalable Encodings

The graphical scalability of a visual representation can be used as a first step in determining how useful a large, high resolution display will be for scaling-up a particular visualization design. Understanding this issue may also lead to insights in how to create information visualizations that are particularly well suited for large displays. Graphically scalable representations are those whose limitations are solved by having more pixels. Non-scalable representations are those that are generally not helped by additional pixels (Table 6.1, Figure 6.1).

Table 6.1: Examples of graphically scalable and non-scalable representations

Scalable	Not Scalable
Number of glyphs	Number of perceivable colors
Glyph size	Glyph orientation
Over plotting, spatial position	3D occlusion
Distortion techniques	Network connectivity edge crossing

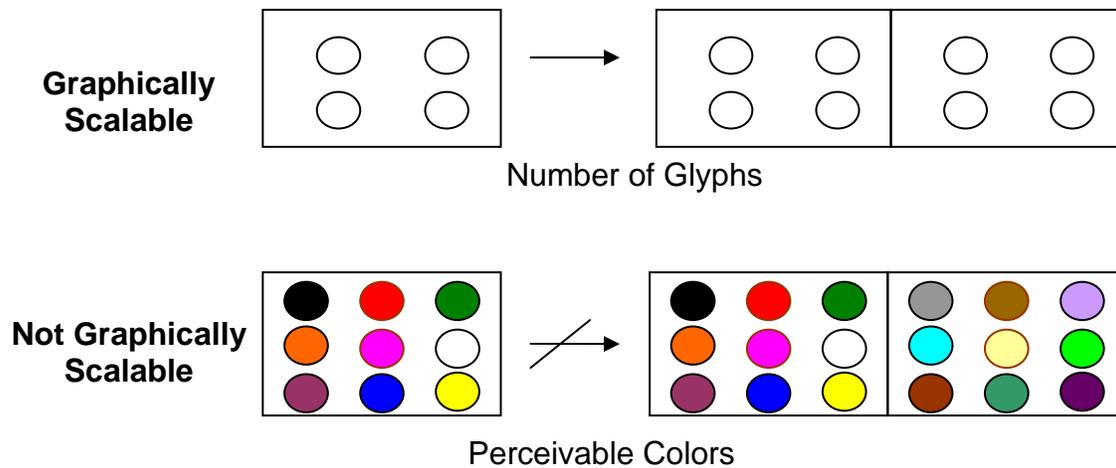


Figure 6.1. The number of glyphs can be increased using additional pixels, the number of perceptually distinct colors cannot

As an example, a bar chart representation is graphically scalable because it can display any amount of data with a large enough display. Similarly, glyph size encoding is scalable because more pixels can enable a greater range or finer granularity of glyph sizes. Over plotting can be reduced with more pixels, making more room for proximal glyphs. Distortion techniques such as focus + context could scale to greater zoom factors or multiple focus regions. Also, particularly with reconfigurable displays, aspect ratio issues could be reduced because the display could be reconfigured to match the best aspect ratio for the visualization. The designer would no longer be constrained by a fixed display aspect ratio.

In contrast, some representations are not graphically scalable. Color encoding is non-scalable because more pixels do not improve the human limitations on the number of distinguishable colors [57]. Similarly, limitations of glyph orientation encoding, 3D occlusion, and network connectivity edge crossings are not solved by more pixels. Therefore, visualizations using these techniques are unlikely to see as much benefit from larger displays.

Recall that the first perceptual scalability experiment used time-series graphs as one type of space-centric embedded visualization. The color of lines was used to represent categorical groups. However, because color is not a graphically scalable encoding (adding more pixels does not allow us to distinguish more colors), this visualization design was no longer usable when we wanted to scale-up the number of attributes. On the other hand, the number of spatial locations was scalable. This meant that the space-centric embedded visualization using bar matrices could be scaled-up.

Visual Encodings and Display Characteristics

We can also look at the interaction between a visual encoding and various display characteristics to help determine how effective a visualization design will be on a given display. Psychophysical requirements for tiled displays, including issues with the brightness, differences in color, and misalignments of images, have been discussed by

others [5, 120]. In this section we explore a different issue related to perception, at the heart of information visualization: Is the ordering of effectiveness of various graphical encodings different on large, high resolution displays?

The order of effectiveness of graphical encodings [38, 83] could be different based on characteristics of the display technology. This idea is suggested by the results that orderings may be different for secondary-tasks that appear in peripheral vision [133]. To provide context for this question, we discuss four different graphical encodings of glyphs: spatial position, size, color, and orientation, and the way they are affected by the display properties (Table 6.2).

Table 6.2: Display property impacts on graphical encodings

<i>Display Property</i>					
	Display Size	DPI	Brightness/ Contrast	Viewing Angle	Bezels
Spatial Position	<ul style="list-style-type: none"> • Outer spatial locations are in peripheral vision • Need to maintain physical context 	<ul style="list-style-type: none"> • Presses limits of visual acuity 		<ul style="list-style-type: none"> • Masks information at edges of display when close 	<ul style="list-style-type: none"> • Can cause spatial distortion • Information may be hidden behind bezels
Glyph Size	<ul style="list-style-type: none"> • Hard to compare when far apart 	<ul style="list-style-type: none"> • If based on number of pixels, higher DPIs lead to less perceptible differences 			<ul style="list-style-type: none"> • Could cause an object to look either bigger or smaller
Color Encoding	<ul style="list-style-type: none"> • Not many color photoreceptors in peripheral vision 		<ul style="list-style-type: none"> • Colors look different across tiled displays 	<ul style="list-style-type: none"> • Low quality display viewing angle can affect the color encoding 	
Orientation	<ul style="list-style-type: none"> • The visual angle can change the apparent orientation 				<ul style="list-style-type: none"> • Glyphs crossing bezels may appear distorted

Display Size

The fact that a display is physically large impacts almost every graphical encoding. Spatial position is affected by size because parts of the visualization that are on the sides of the display are now in peripheral vision. This means that users are more likely to miss those aspects of the visualization and that it will be harder to attract attention to those

areas. Comparing the vertical or horizontal relationship between distant points (as in a scatter plot) may become more difficult.

Using a large display also means that users have to maintain physical context when physically navigating. Even if legends and labels are larger, as a user physically navigates the display they often lose sight of that information. Therefore, based on our experience, it is important to consider having both local and global legends on a large display (or taking a dynamic approach as will be discussed in the section on interaction design). When using space-centric embedded visualizations, only having a local legend that was found within each embedded location meant that users had to walk close to the display to read the legend before moving further back to get an overview. If they had a global legend then that would not have been necessary. On the other hand, the attribute-centric and space-centric embedded using bar matrices visualizations only had global legends, which meant sometimes users had to step back to see the global legend in the middle of doing a detail task (Figure 6.2). Additionally, the group and year labels are often useful in maintaining physical context. Therefore, it would have been useful to place labels at multiple strategic locations. Carefully placing labels and legends may help users maintain physical context.

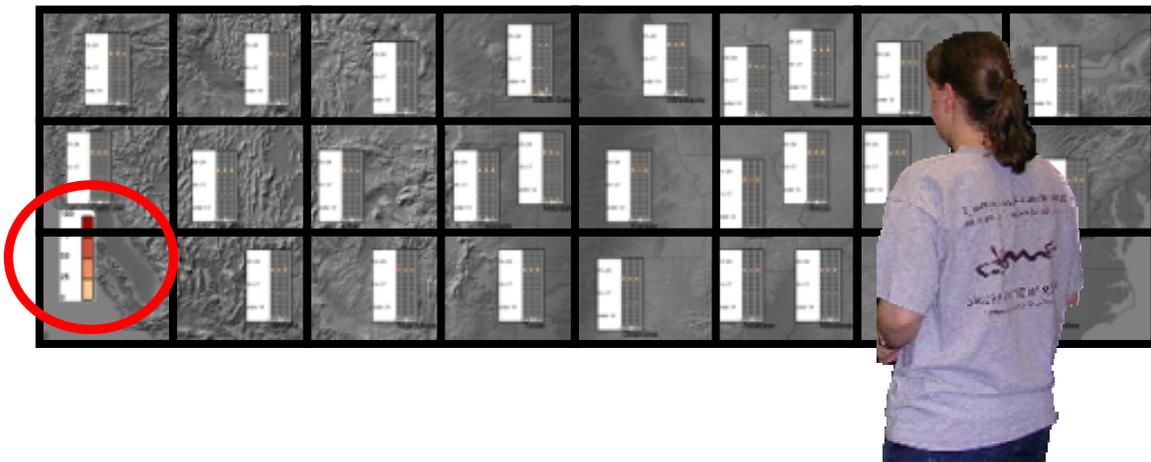


Figure 6.2. Both local and global legends should be considered, a user may lose physical context when viewing a global legend during a local tasks

Graphical encodings of size, color, and orientation are also affected by the physical size of the display. Large displays are likely to require comparing sizes that are much farther apart, which may be more difficult as the distance between two sizes increases. Color is likely to move down in the order of encodings because of the lack of color photoreceptors (cones) in peripheral vision. Line orientation may be affected by the physical size of the display because the angle at which a line is viewed has the potential to change the apparent orientation of a line.

Although intuitively color appears likely to be less effective on large displays, it can still be a useful encoding despite the lack of color receptors in peripheral vision. In a pilot study, two different versions of the embedded visualizations using the bar matrices visualizations were compared. The first version used bar height alone to encode values,

the second version had a redundant encoding of both bar height and color. We found that colors were preferred to bar heights alone because distant colors were easier to compare than distant sizes.

DPI

The pixel density or dots per inch (DPI) of a display is likely to have a significant impact on visualization design and the order of graphical encodings. If the DPI is too high a user simply will not be able to distinguish small differences and combining a large physical size display with a high DPI will press the limits of visual acuity, regardless of physical navigation. If size is one of the graphical encodings being used then the DPI may have a drastic effect because if the size is based on the number of pixels then a higher DPI will lead to a smaller and less perceptible size difference. This may mean that size is a better encoding on lower DPI displays.

Brightness/Contrast and Viewing Angle

Maintaining consistent coloring across tiled displays is a challenging issue because displays age at slightly different rates and can easily end up with slightly different settings. Inconsistencies in color across different tiles in a tiled display (Figure 6.3) can lead a user to believe there is a pattern when there actually is none.



Figure 6.3: Differences between tiles are exaggerated by the viewing angle

The viewing angle rating of the display can also cause problems. When the user is close to a display they may not be able to see any object that is spatially placed on the outermost portions of the display. If color encoding is being used the viewing angle can also affect the color being perceived, particularly with lower quality LCD displays.

Bezels

Another issue is the presence of seams in tiled displays. While rear-projection arrays have very small seams created by tiling projectors, tiled LCD monitors have larger discontinuities created by physical bezels surrounding each monitor. Each of these situations results in problems aligning images. Some techniques for addressing this have

been presented by others [84, 85] and include alternative methods of aligning images and visualizations that provide a representation of the information that is hidden behind the bezels.

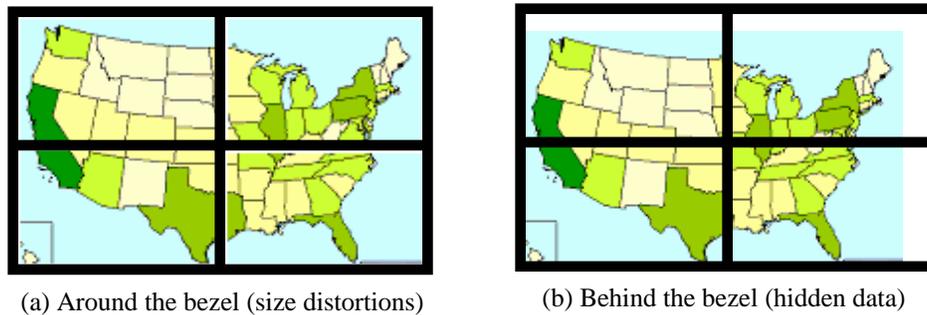


Figure 6.4: Distortions caused by bezels

The basic issue is that the spatial position can be affected by bezels which cause spatial distortion if no information is hidden behind the bezels (the default with current software for distributing graphics across large displays) by causing misalignments or geographical areas to appear larger than they actually are. Alternatively, if information is hidden behind the bezel then there is no distortion but the user may be unaware that information is there (Figure 6.4).

There is one study showing that physical discontinuities between information (such as bezels) only impact performance when they are combined with an offset in depth [131]. However in a different study, users expected important information to be hidden behind the bezels and therefore spent time panning images just to determine what might be behind the bezels [80]. Potentially, new interaction techniques could be developed to allow users to efficiently “peek” behind bezels and distortion techniques that offset the distortion caused by bezels could be applied.

In our studies this issue was avoided by placing multidimensional information such that it never crossed boundaries. Technically it is possible that there was an effect due to the bezels on the space-centric visualization since it distorted the size of the map while individual maps in the attribute-centric visualization never crossed bezels. However, there were no user performance results that indicated that this was a problem for any of the tasks tested.

6.1.2. Visual Acuity and Physical Navigation

Visual Aggregation

It is important to note that overview techniques will still be needed on large, high resolution displays because large enough datasets may still require more pixels than available to the user. Furthermore, highly aggregated overviews designed for small displays may still prove to be useful on large displays.

This illuminates the distinction between computational aggregation and *visual aggregation*. Computational aggregation is seen when an algorithm groups data to

compute a new smaller dataset, and occurs in the graphical representation stage for the purpose of fitting an overview on small displays. Visual aggregation is seen when the human visual system lumps small visual stimuli together, and occurs at the visual perception stage (Figure 6.5). With large displays, users can physically navigate (step back) to gain a visually aggregated overview. It is possible that by being able to see all of the data at once, visual aggregation will become more prevalent and be used to aid computational aggregation techniques.

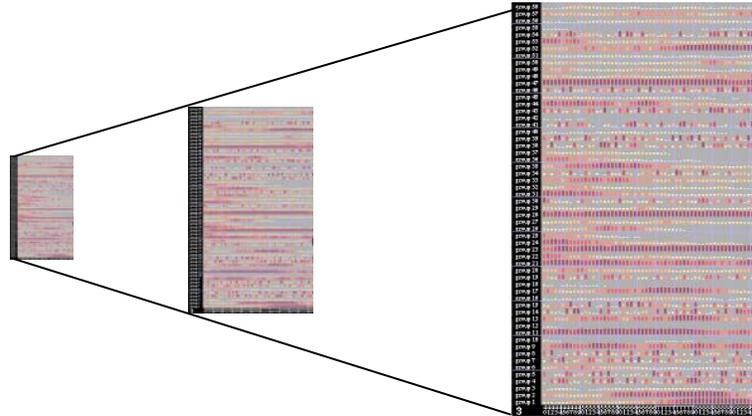


Figure 6.5. Visual aggregation occurs as a user physically moves away from the display

Physical navigation combined with visual aggregation has a somewhat similar effect as geometric zooming in computational aggregation. However, semantic zooming in computational aggregation can offer an entirely different effect, in which the representation changes at different levels of detail.

Visual aggregation is likely the strongest reason that the space-centric embedded visualization approach proved to be advantageous on the large, high resolution display. The proximity of individual bars in the grid resulted in visual aggregation. This meant that although individual bars were not distinguishable from a distance, general patterns emerged. Visual aggregation does not occur in the attribute-centric approach because the individual bars representing data values are always separated on each small map.

Recall that mental demand and physical demand increased significantly between display sizes with the attribute-centric design but not with the space-centric embedded visualization. The visual aggregation of values with the bar matrices visualization and corresponding emergence of patterns was likely the reason that mental demand did not increase with this visualization. Hypothetically an increase in reported physical demand is an indication of an increase in physical navigation. Therefore, visual aggregation also likely helped with performance times because of the reduction in required physical navigation. This is supported by qualitative observation of users' physical interaction with the visualizations on the large display.

Adaptive Visualizations

The physical size of the display coupled with the lack of virtual interaction created certain design considerations for physically navigating the visualizations. The most

simplistic of these considerations was that labels needed to be placed at multiple strategic locations. For the attribute-centric visualization this meant placing group labels at both the left and right sides of the display. While we could have also placed them in the middle, we felt this would create a meaningless grouping and interfere with overview tasks.

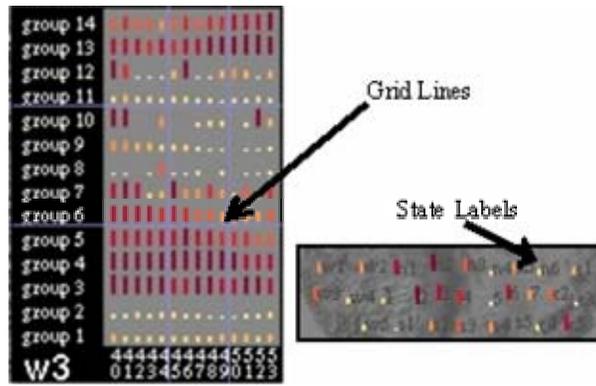


Figure 6.6. Physically adaptive visualizations

We also used what we have termed physically adaptive visualizations. Physically adaptive visualizations create the visual illusion that some details disappear as the user moves away from the display. If carefully designed, visualizations can be created for large, high resolution displays that more fully take advantage of visual aggregation and human perceptual abilities. By using light colors that blend with the background, users have access to details when close without being overwhelmed by details when standing back to get an overview (Figure 6.6). In the space-centric visualization we used light blue gridlines that could only be seen when users were close enough to make use of them. Likewise, in the attribute-centric visualization we used light gray state labels that could only be seen when close to the display. If we had simply made these labels black they would have quickly overwhelmed the visualizations. This approach is likely to work better with wider displays than with taller displays where the highest pixels cannot be accessed by walking. This technique will work even in multi-user scenarios, since each user will perceive the visualization appropriately regardless of their relative positions.

Alternatively, semantic zooming concepts can be made more explicit by tracking the user's physical position with respect to the display, and automatically adjusting the visual representation accordingly. An example might be to enlarge the font of the most important labels as the user steps back. This can be further customized to the perceptual resolution of human focal and peripheral vision. By tracking the user's visual focus (e.g. eye, head tracking), visual information in the focal area can be made more detailed while semantic zooming techniques operate in the periphery.

6.2. Interaction Design Considerations

So far we have covered design issues related to perceptual scalability. In this section we move toward interactive visualizations and begin to explore future extensions of this work by asking: How will user interaction with information visualizations need to change

when using large, high resolution displays? The focus is on interaction scalability issues specific to information visualization, including navigation techniques, brushing and linking, and interaction with separate controls or widgets. For general large display user interface challenges such as reaching distant objects, tracking the cursor, managing space layout, and alternative input devices, see Ni et al. [95]. Many innovative ideas for dealing with general interaction problems have already been implemented [15, 54, 110].

6.2.1. Navigation Techniques

The most basic difference in interaction between display sizes is related to navigation techniques: there is a trade-off between virtual navigation and physical navigation. While physical navigation (moving eyes, head, and body) has advantages in speed and maintaining context, virtual navigation (panning, zooming) requires less strenuous effort from users. While smaller displays emphasize virtual navigation, larger displays offer users the choice of both. One study has shown that larger displays and the corresponding increase in physical navigation can be beneficial for map navigation [13].

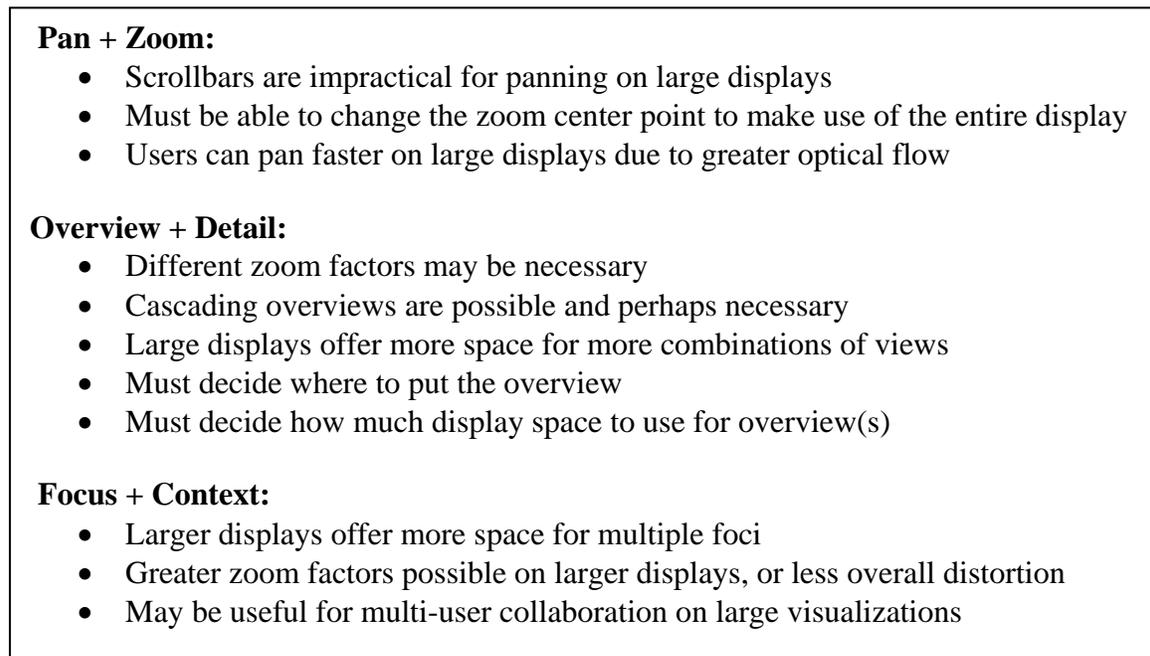


Figure 6.7: Navigation strategy issues

Basic navigation techniques typically used in information visualization (overview + detail, focus + context, and pan + zoom) will likely need to change on large, high resolution displays. Some of these issues are outlined in Figure 6.7. For overview + detail techniques, different zoom factors may be necessary, and cascading overviews are possible and perhaps necessary. Designers must decide where on the display to put the overview and how much display space to use for it. Focus + context techniques may be useful for collaboration on large displays, where each user can have their own focus region, and foci could combine when they get near each other.

6.2.2. Brushing and Linking

Different types of brushing and linking techniques may also be necessary. This is particularly relevant to the design of linked visualizations. It is likely that change blindness [100] will become a major problem on large displays, and users may not see the linked highlight changes in the periphery. Scanning very large complex visualizations for highlighted elements is cumbersome. Some potential solutions are to use motion [14] (which may become too distracting) or temporary flashing, to update views based on eye movement and when a user looks at a view, or to simply change the speed that views update either to a slower speed or to varying speeds based on the distance between the user and the view. Another possibility is to use some form of afterglow [18] or ripple effects for showing change. These issues are outlined in Figure 6.8.

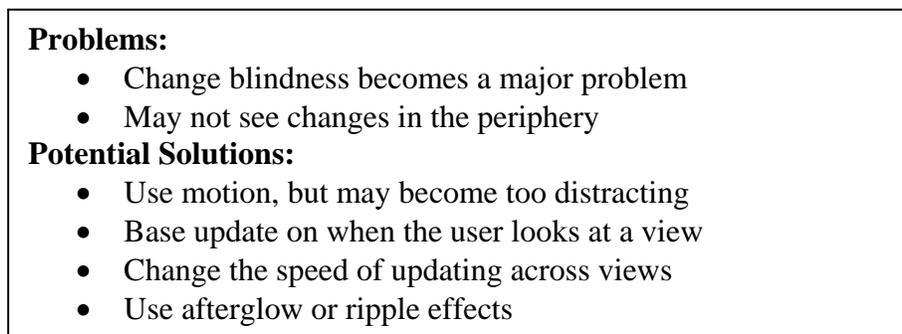


Figure 6.8: Brushing and linking issues

6.2.3. Selection

Interactively selecting data glyphs, as required when brushing and linking, may require new solutions on large high resolution displays due to expanded extremes:

- Very small selections: Selecting single small glyphs on high-DPI displays will require precise pointing.
- Very large selections: Selecting large groups of glyphs in complex patterns across multiple display tiles will require broad scale pointing, fine control, and advanced lasso techniques [54].

New input devices and techniques [34] proposed for large displays may not work well for both extremes. For example, gestures or touch screens may be too coarse for small selections without special interaction techniques [20]. On the other hand, stylus/touch sensitive displays may result in difficulty making large selections, especially in the presence of bezels between monitors. Special adaptations such as virtually extending the hand of the user, bringing objects closer [24], or having a small overview of data from which to select a region can help with these issues [39]. Multi-scale pointing methods are required, with coarse and fine modes or different input devices for each mode. A

thorough review of existing techniques [93] is beyond the scope of this work, but is worth mentioning here because it is likely that techniques currently proposed in the literature will need to be adapted for information visualization specific tasks.

6.2.4. Control Panels

One final important design element in information visualization is the control panel, such as dynamic query sliders, that operate on other visual representations. Large displays afford ample space for scaling up to larger numbers of controls or more advanced controls. However, the standard approach used in small displays, docking the panels around the edges of the display, is problematic on large displays. Controls can become physically distant from the user making them hard to access [115], and distant from the visual elements they operate on making it hard to see effects (e.g. dynamic query filtering). Note that these problems and potential solutions are also applicable to legend placement. By default, controls should be near-centrally located, such as at the bottom-center of the display, to minimize average distance while not impeding the center. Potential solutions for bringing controls nearby include:

- Moveable controls: users move the control panel to where they are currently working on the large display [27]. Optionally, controls could automatically follow the user or cursor [115].
- Pop-up controls: user input causes controls to appear at the user's location, at the correct scale, and disappear when done. These could be similar to context menus and only provide options relevant to the context.
- Hand-held controls: controls are off-loaded onto a separate mobile display device, such as a palm or tablet that users carry with them. This takes better advantage of embodiment principles [45], but requires users to carry a separate device and may have limited screen space. The hand-held should be non-tethered to preserve mobility for physical navigation [122].
- Gestures: visual controls are replaced by gestures, such as a two-handed lateral gesture to specify a dynamic query filter bounds.

In general, from previous research we know that controls should be consolidated onto a small number of control panels so that controls do not get scattered and lost on a large display. In collaborative multi-user scenarios, localized controls such as lenses may be more appropriate than global controls.

6.3. Summary and Open Issues

The purpose of this chapter was to discuss information visualization design issues on large, high resolution displays and to also consider the reasons that the space-centric embedded visualization design was advantageous on the large display. The main findings are as follows.

Research Question 4:

What characteristics of a visualization determine how visually scalable it is for large, high resolution displays and what design issues remain?

- Graphically scalable encodings, or visual encodings that can be scaled simply using more pixels, result in scalable visualizations.
- The interaction between visual encoding and display characteristics must be considered when trying to create scalable visualizations.
- Visual aggregation is beneficial because of the decreased navigation that is required and the reduction in mental demand. This is part of the reason that space-centric embedded visualizations become better than attribute-centric visualizations on large displays.
- Physically adaptive visualizations can be used in a way that allows users to see details when close but without the details overwhelming the visualization from a distance.

Information visualization design can be considered from a perceptual perspective and also from an interaction perspective. The perceptual perspective looks at the visual encodings. To create visualizations that are scalable for large, high resolution displays, graphically scalable encodings such as spatial location should be used. The interaction between the visual encoding and different characteristics of the display should also be considered. A variety of factors from DPI to the presence of bezels can impact the choice of encoding.

The interactive perspective has been explored by many researchers. However, most of that work has been on general interaction with large displays as opposed to information visualization specific interactions. New ideas for dealing with some of these issues were presented. These ideas have not been tested and that experimentation will be left as future work.

Also explored in this chapter were the reasons for the space-centric embedded visualization being advantageous on the large display. Visual aggregation was one major reason for that. The proximity of the bars encoding values resulted in the emergence of visual patterns. Because the patterns visually emerged, the users did not need to physically navigate and then mentally integrate. Instead, they simply stood back and visually scanned the display. The mental and physical demand increases that were seen with the attribute-centric design were likely not seen with the space-centric embedded visualization for this reason. The spatial grouping used in the space-centric visualizations is beneficial since less total navigation is required.

Based on the discussion in this chapter, we can now outline a few of the remaining open issues with respect to information visualization design on large, high resolution displays.

Question 1: Other than increased data scalability, how might information visualizations benefit from larger, higher-resolution displays? Answering this question will require carefully considering the particular display characteristics and how we can exploit each

visually. Creating visualizations that do more with the pixels than simply showing more data is an open issue. Some potential possibilities include showing multiple levels of scale, showing more complex and heterogeneous datasets, and providing more contextual information for tasks.

Question 2: Which types of datasets and user tasks benefit most from the increased data scalability provided by additional pixels? In this work we have seen some benefits for a variety of tasks with geospatially-referenced data. There are many other types of visualizations that will likely benefit from the additional pixels, with new visualization design issues yet to be revealed.

Question 3: Are comparisons between navigation and interaction techniques similar on various display sizes? Focus + context techniques have been compared to overview + detail techniques on desktop displays, but multiple foci and cascading overviews provoke interesting new questions in this realm.

Question 4: Is the ordering of effectiveness of various graphical encodings different on large, high resolution displays? A variety of differences based on display characteristics have been discussed, but these are merely hypotheses and need to be tested. Effective visual encodings will be necessary to exploit the available pixels.

Question 5: When is computationally aggregating data advantageous to using additional pixels and showing more information? This is perhaps the most intriguing of the questions as it really asks what is lost in the aggregation process and what might be gained by not aggregating. One possibility is that more complex multi-scale insights will be gleaned by users of large displays.

This work has only begun to scratch the surface for exploring information visualization design on large, high resolution displays. Knowing that users benefit from additional data, that visualization can still be effective even beyond visual acuity, and exploring the reasons for the advantages of embedded visualizations on large displays, is only a start. Much is left to be explored.

In the last chapter the design space will be tied back in to the performance comparisons between visualizations, and some initial guidelines for information visualization design on large, high resolution displays will be presented.

Chapter 7. Conclusion

7.1. Research Summary

The overall research goal was to compare the visual scalability of integrated and multiple view visualizations for large, high resolution displays in the context of integrating spatial and multidimensional data. This goal was split into four distinct smaller research questions. These questions necessitated outlining the visualization design space, comparing visualizations on a desktop display to establish a baseline, comparing visualizations as data was added using more pixels, and exploring information visualization design issues on large, high resolution displays.

7.1.1. Design Space

First, a hierarchical design space was created for integrating spatial and multidimensional data types. This design space evolved during the course of the research, but the integrated and multiple views approaches remained central. The specific type of integrated views approach was termed space-centric because this design overlays all of the multidimensional data on a single spatial structure. Using the hierarchical representation, space is at the top level of the hierarchy with multidimensional data below. The specific type of multiple views approach was termed attribute-centric because each attribute from the multidimensional data is overlaid on a separate spatial structure. Using the hierarchical representation, the multidimensional data is at the top level of the hierarchy with the spatial data below. Visualizations that use a mixture of these approaches are represented by placing data types on the same level of the hierarchy. We also discussed expanding the hierarchy to include temporal data as a third type. Using the hierarchical notation, the multiple linked visualizations approach is simply multiple linked hierarchies.

7.1.2. Baseline Comparison

After creating an initial version of the design space, we wanted to determine which visualization approach resulted in the best user performance for a variety of tasks using a desktop display. This comparison acted as a baseline before comparing designs as they were scaled-up using more pixels. It is worthwhile to note that this question was mostly exploratory and acted as background work for the next research question.

Three experiments were conducted. The first experiment compared the space-centric complex glyphs approach to the attribute-centric approach using multiple simple glyphs. Although the scalability of the complex glyphs approach is limited, it is a reasonable comparison to make on a desktop display. We found that the perceptual salience of visual encodings was a more important factor than the number of views. This gives the attribute-centric approach an advantage over the space-centric complex glyphs visualization on a desktop display since complex glyphs use less perceptually salient visual encodings. We also saw that mixed designs, or designs where the data was split between the space and attribute-centric approaches almost always resulted in the worst

user performance unless it was specifically suited to a particular task. This visualization seemed to suffer from the worst of both designs in that users had to mentally integrate and visually separate data. Although generally performance was similar between visualizations, various trends were indicated.

The trends indicated in the first experiment can be compared with the findings in the third experiment. That experiment compared a 3D scatter plot where information was visually integrated to four separate 1D plots where multidimensional data was interactively linked. In general, the multiple views visualization was better for tasks that required finding an attribute within a given range, a specific value, or for basic search and comparison tasks. The integrated view visualization was slightly better for trend and relationship tasks in 3D, but not when spatial data was added and a comparison was made using the space-centric visualization with complex glyphs. It appears that it was easier to mentally integrate information from separate spatial structures than to attend to only a select number of features when using complex glyphs.

The findings from the second and third experiment also are connected in their consideration of interaction between linked visualizations. The third experiment demonstrated that brushing can be an effective method of linking views. The second experiment explored cognitive issues with respect to multiple linked visualizations and specifically considering the effects of context switching. The results indicated that context switching may not be as detrimental as thought. Focusing attention and analogical reasoning are important cognitive abilities when finding patterns. This is indicated by a correlation between user performance and cognitive abilities tests. In addition to these findings, brushing helps the user focus their attention and also increases the ability to compare values across maps. This connects to future work on interaction when we wish to determine if brushing will also be an effective method of linking views on large displays.

7.1.3. Perceptual Scalability

This was the key research question and the heart of this work. This part of the work was focused on perceptual scalability, the effectiveness in terms of user performance when information visualizations were scaled-up. This scaling-up of data was accomplished using the additional pixels available with a large, high resolution display. Two experiments compared the perceptual scalability of visualizations for large, high resolution displays. The first experiment explored general cognitive and perceptual issues. The results of this experiment led to the second experiment which looked at visual acuity as a potential perceptual scalability limit. In these studies the dataset sizes were increased along with the total number of pixels.

Before the experiments were conducted we first analyzed the graphical scalability of the space-centric embedded visualization compared to the attribute-centric approach. This analysis demonstrated the space-centric embedded visualization approach was more graphically scalable because it requires fewer pixels to add another encoding in a grid than it does to add another map in the attribute-centric visualization. Although the space-centric embedded visualization approach was more graphically scalable, it was not clear

which would be more perceptually scalable. Some researchers had speculated that the simplicity of the attribute-centric approach would result in it being more beneficial than the complex space-centric approach.

The first experiment demonstrated that neither visualization failed when scaled-up, both were perceptually scalable. By perceptually scalable we mean that the visualizations were effective in terms of user performance time and accuracy as the amount of data being visualized was increased using more pixels. Accuracy was maintained and the increase in time was less than proportional to the increase in data. The time and accuracy differences between the space and attribute-centric approaches were similar in the first experiment, but the mental workload increased significantly only with the attribute-centric approach and user preference switched to space-centric embedded visualizations on the large display.

The second experiment demonstrated that there are user performance benefits even beyond visual acuity. Visual acuity is not the limit for information visualization on large displays. For five of the tasks, increases in time occurred at a slower rate between the larger display sizes than between the smaller display sizes. Accuracy also generally improved or was consistent. User performance benefits existed for some tasks because of the additional data that could be displayed despite the physical navigation that was required to visually distinguish all elements of the visualization.

The space-centric embedded visualization approach also diverged and became better on the large display in this experiment. By considering the results that were seen across the two experiments, we can hypothesize that the advantages of this visualization are more related to display size than the amount of data. This is because when the amount of data shown was equal, there were significant differences between designs on the large display, but not the small display. When the amount of information was different, advantages were only seen on the larger display sizes. The space-centric embedded visualization approach is not only more graphically scalable, but also more perceptually scalable.

These two experiments also allowed us to continue to explore the link between visualization and user performance on a variety of tasks. The results of detail tasks were as one might expect, and could be explained by less navigation. Examples of this will be provided in the following section. In terms of temporal and spatial overview tasks, the space-centric embedded visualizations resulted in better user performance than the attribute-centric approach. The only pattern finding task where the attribute-centric visualization was better was for finding correlations between groups on the large display. Despite performance time and accuracy being better with the attribute-centric visualization, users still preferred the space-centric embedded visualization for this task. The obvious next question was why was the space-centric embedded visualization approach was better on the large display. We also wanted to explore the lessons learned in terms of improving information visualization design on large, high resolution displays.

7.1.4. Design Issues

The next logical progression in the research was to explore information visualization design issues with respect to large, high resolution displays. Within this framework it also

proved useful to answer the question of why the space-centric embedded visualization resulted in better user performance on the larger display.

Information visualization design can be considered from a perceptual and interaction perspective. To create visualizations that are scalable for large, high resolution displays, graphically scalable encodings such as spatial location should be used. Graphically scalable encodings were distinguished from non-scalable encodings as one way of determining how well a visualization would scale for large, high resolution displays. The interaction between visual encodings and various display characteristics were explored. The display technology being used may help determine how effective a particular visualization is.

Other design approaches took advantage of the physical navigation that is inherent with large displays. In particular, visual aggregation was quite beneficial. It resulted in a decrease in navigation and also a reduction in mental demand. Taking this concept a bit further, physically adaptive visualizations can be used in a way that allows users to see details when close but without perceptually overwhelming the overview from a distance. These techniques were used when creating visualizations for the experiments.

The interactive perspective on large display interfaces had already been explored by many researchers. However, that work has not been specific to information visualization specific interactions. Issues with navigation, overviews, brushing and linking, selection, and control panels were presented. Some of the ideas have not been tested and that will be left as future work.

Throughout the chapter the reasons for space-centric embedded visualization advantages on the large display were discussed. Visual aggregation was one major reason for the advantage. The proximity of the bars encoding values resulted in the emergence of visual patterns. This meant that physical navigation was reduced. It also resulted in less mental demand. Users simply stood back and visually scanned the display, rather than having to walk and mentally integrate as much as was necessary with the attribute-centric visualization. The mental and physical demand increases that were seen with the attribute-centric design were likely not seen with the space-centric embedded visualization for this reason. The spatial grouping used in the space-centric visualizations also is beneficial since less total navigation is required. An additional observation is that the spatial grouping meant that users were not overwhelmed. Many users commented that it seemed as if a lot more information was being show with the attribute-centric visualization than the space-centric visualization, although the amount of data was actually equal.

The way that this research evolved will briefly be presented in the following section.

7.2. Research Evolution

Many significant changes occurred as this work progressed. Early on we aimed toward a cognitive framework for the creation of multiple linked visualizations and for understanding the differences between simplified alternative perspectives and single

complex integrated visualizations. Because this was so exploratory in nature and complex, we had to start at a more basic point. This point was the design space and a comparison of single complex and multiple simple glyphs. As the design space slowly evolved, new displays were built and this provided an exciting opportunity to switch the focus of the research to the comparative scalability of integrated and multiple views for large displays. Interaction had to be temporarily left behind at this point because current interaction techniques for large displays and computational performance were likely to have a dramatic influence on the results. We wanted to be sure to focus on the perceptual factors closest to the user. As we scaled-up the amount of data we found the complex glyphs approach no longer worked and had to switch to an embedded visualization approach. When we scaled up again, embedded time series graphs were also left behind.

Perhaps the most unexpectedly difficult aspect of the evolution has been in the terminology. As the design space changed, so did the terminology. Integrated and multiple view visualizations are such loaded terms that one must be careful to describe exactly what is meant.

In the next section the experimental comparisons between visualizations will be tied back to the design space offered at the beginning of this work. Then, in the following section, some initial guidelines that came from the experiments related to information visualization design on large, high resolution displays will be presented.

7.3. Visualization Design and User Task Performance

At the beginning of this work a design space was offered that was used for describing information visualizations that integrate spatial and multidimensional data types. This design space was also expanded to account for the addition of a third, temporal data type. The experiments conducted compared these visualizations. This section reports on the link between the design space for visually integrated data types and user performance on different tasks.

7.3.1. Basic Concept

The basic connection between the design space and user task performance is: *The higher a data type is in the hierarchy, the better user performance is on tasks related to that data type.* This is due to the increased visual proximity of graphical encodings for data types at higher levels in the hierarchy. It also roughly corresponds to the proximity compatibility principle which states “Displays relevant to a common task or mental operation (close task or mental proximity) should be rendered close together in perceptual space (close display proximity)” [142]. It takes this principle further by accounting for multiple data types being visually integrated, accounting for linked views, and providing a framework for spatial data that can extend beyond just 2D data.

7.3.2. Tasks

Central to this discussion is a basic organization of user tasks. In the experiments we split tasks using the basic information visualization overview and detail tasks. Within those

categories tasks were systematically designed to fix certain data types while others varied. For example, if the user was asked to find the year with the highest population of white males in Kansas, the state and demographic group were fixed, and the year varied. An alternative method of organizing tasks when visually integrating spatial and multidimensional data is demonstrated in Table 7.1. This approach is used because of the ability to easily determine which data type a task is most related to. The implication is that we can then match the hierarchical representations with tasks resulting in the best user performance.

Table 7.1. Overview of basic tasks (examples provided in each cell)

	One attribute	Many attributes
One location	Read Value	Search (given location, find attribute) Compare (attributes at one location) Overview (one location, all attributes)
Many locations	Search (given attribute, find location) Compare (locations for one attribute) Overview (one attribute, all locations)	Overview (of all data)

Table 7.2. Task hierarchies

	One attribute	Many attributes
One location	Spatial/Multidimensional	Spatial └ Multidimensional
Many locations	Multidimensional └ Spatial	Spatial/Multidimensional

7.3.3. Link between Tasks and Design Space

Visualization design hierarchies can be matched to task hierarchies as shown in Table 7.2. Two of the cells in Table 7.1 have a clear dominant data type. When a location is fixed (i.e. the one location, many attributes cell) these tasks are most strongly dependent on the spatial data type. This is because space is central to these tasks while the attributes may vary. An example of an overview task is asking about the patterns of all attributes seen at a specified location. For these tasks the space-centric visualization should result in the best user performance. On the other hand, when an attribute is fixed (i.e. the one attribute, many locations cell) these tasks are more strongly dependent on the multidimensional data type. An example of an overview task is asking about the pattern seen for one specific attribute across all locations. For these tasks, the attribute-centric visualization should result in the best user performance.

The remaining two cells are not necessarily related to a specific data type, but rather rely on both. For the read value task we would expect that user performance would depend on the number of attributes vs. the number of locations to sift through. The final cell (many attributes, many locations) may be the most interesting as it involves an overview of all of the data and therefore can result in many different types of insights. In this case, the dominant type of insights is likely to correlate with a data type's position in the hierarchy. A space-centric visualization will result in more insights relating to the spatial data type (i.e. comparing spatial locations with respect to all attributes) while the attribute-centric visualization will result in insights related to the multidimensional data (i.e. comparing attributes with respect to all spatial locations).

7.3.4. Empirical Support for the Link

In general, the results of our empirical studies supported this link between the hierarchical representation of a visualization and user task performance. As previously mentioned, the results of detail tasks were as one might expect. When location was given as part of the task, the space-centric visualization resulted in better user performance. An example of such a detail task would be to find the year with the highest population of males in New York. When an attribute was given as part of the task, the attribute-centric visualization resulted in better user performance. An example of such a detail task would be to find the state with the highest population of males in 1989.

The overview tasks were a bit more interesting and complex. When space-centric complex glyphs were used on the desktop display, the attribute-centric visualization had an advantage for finding correlations when the data was visually separated as opposed to visually integrated. This highlighted the advantages of perceptually salient visual encodings and goes against the general principles outlined here. However, these predictions held on the large display when the visual encoding between the space and attribute-centric visualizations was held constant by using the colored bars as the encoding.

The link between task and hierarchy becomes slightly more complex when discussing the perceptual scalability experiments because three different data types were used rather than two (temporal was included). In these experiments the space-centric embedded visualization resulted in better user performance for spatial and temporal overview tasks on the large display. This matches with the prediction because spatial data was at the top level of the hierarchy for the embedded visualization design. The only pattern finding task where the attribute-centric visualization was better on the large display was for finding correlations between groups.

Although beyond the scope of this document, these predictions can also be extended to other spatial data types and potentially to multiple linked visualizations. Comparisons of the basic visualization approaches have been conducted with biological graphs and 3D structures as the spatial context, with similar results [104, 105, 117, 118]. This connection between visualization and task can also be used when designing multiple linked visualizations by specifying the strengths of each type of visualization and matching the strengths and weakness of different visualizations (i.e. using visualizations that have

different data types at the top level of the hierarchy). The design space provides a common terminology and conceptual framework for integrating data types and allows for the unification of work being done across scientific fields within the information visualization community.

Of note is that relative comparisons between visualizations were similar with a small dataset on a desktop display, but became more apparent with a large dataset and on the large, high resolution display. And, as previously mentioned, the space-centric embedded visualization approach diverged and became significantly better than the attribute-centric approach on the large display. The advantages were due to the spatial grouping of information reducing the amount of navigation, visual aggregation reducing the amount of mental effort that was required, and appearing to be less overwhelming to users. In the next section we leave the comparison between visualizations, and look at guidelines that can be extracted from this work related to information visualization design on large, high resolution displays.

7.4. Large Display Visualization Design Guidelines

A few guidelines for design can be taken from the experiments and the process of designing the visualizations used in these experiments. Following are guidelines along with brief explanations that summarize our initial experiences with information visualization design on large, high resolution displays.

1. Do not be afraid to display a *large amount of information*.

The usefulness of large displays for information visualization had been questioned because of the potential to overwhelm the user with data. However, the experiments demonstrated that designs are perceptually scalable and that users can benefit from the additional data that is shown. This is true even when enough information is shown on a display of sufficient size and resolution to be beyond visual acuity (i.e. physical navigation is required to distinguish all elements).

2. Use *embedded visualizations* on large displays.

Embedded visualizations are one method of encoding data at each location with space-centric visualizations. The spatial grouping of information likely reduces the amount of physical navigation and also leads users to perceive the visualization as being less overwhelming. While there are tasks for which attribute-centric visualizations are preferable, space-centric embedded visualizations have the overall advantage for large displays for the datasets and tasks tested.

3. Use designs that *visually aggregate*.

Visual aggregation can act as an aid to visual acuity and increase the clarity of patterns. The emergence of patterns when a user steps back from the display results in a reduction in the mental effort required to discern patterns. A space-centric embedded visualization that visually aggregate was the best visualization on the large display.

4. Carefully consider the effects of *physical navigation* when designing the visual interface.

The switch from virtual to physical navigation with large display creates a variety of design considerations with respect to information visualization. Having both local and global legends can aid in multi-scale analysis and help users maintain physical context. Placing labels at multiple strategic locations also aids users in maintaining physical context.

Designs can also be created that take advantage of the limits of human perception. We presented what we have termed physically adaptive visualizations. The visual illusion is created that details disappear as a user moves away from a display. This was implemented by including light blue grid lines on a gray background. Users could follow the grid lines to get details when they were close to the display without being overwhelmed by them when they stood back to get an overview. The same concept was used when labeling states in the attribute-centric visualizations. Certain characteristics of the visualization that will only be needed for detail tasks can be presented in this way.

5. Choose representations based on the *graphical scalability* of the visual encodings.

Spatial and size encodings are graphically scalable encodings. The difference between graphically scalable and non-scalable encodings was described. Visualizations that employ scalable encodings are able to be scaled-up using more pixels. Also related to this issue is consideration of how various visual encodings will be affected by display characteristics. Properties such as DPI and size can influence various encodings in different ways.

We hope that these guidelines will aid designers in creating initial visualization prototypes for large displays, and that the issues discussed throughout this paper will motivate future research.

7.5. Contributions

This research can help visualization designers with creating information visualizations for large, high resolution displays. It also helps in understanding human abilities and the relationship between user performance on a variety of tasks and the choice of visualization. The main contributions of this work are:

- The creation of a hierarchical design space for integrating spatial and multidimensional data.
- Design guidelines for choosing a visualization in the design space based on task. This includes an analysis of the tradeoffs involved with opposing visualization approaches.
- Comparisons of visualizations across display sizes. Space-centric embedded visualizations (one type of integrated view) proved to diverge and be

advantageous to attribute-centric visualizations (one type of multiple view visualization) on the larger display.

- Demonstration of the perceptual scalability of visualizations. Visual acuity is not the limit for information visualization and there are benefits to the additional data that can be displayed using more pixels.
- Outline of design issues for information visualization on large displays. Initial design guidelines were suggested and open research questions were presented.

This research increases our understanding of the perceptual scalability of different visualizations for integrating spatial and multidimensional data as we transition to large, high resolution displays. Understanding visual design issues involved with this transition will help us create visualizations that provide users with new and exciting insights into their data.

7.6. Future Work

A variety of remaining open issues were outlined earlier. The next progression in this work is to add interaction back into the visual interface and test the scalability of interactive information visualizations. The major concern becomes change blindness and awareness of updates made to linked views in peripheral vision. Some ideas for addressing this, including the potential to use a ripple update effect, were presented in this document. Many new design issues are likely to arise simply from observations and experience. Information visualization techniques with cascading overviews and multiple foci could be implemented on large displays, but have not been used before now because of pixel limitations.

In addition to adding interaction into the equation, many other significant questions remain. One of these is determining how information visualizations might benefit from large displays other than by simply adding more data. There are likely to be design ideas that have not even been dreamed of yet because of the typical single monitor limitation. Another open area is to determine how computation aggregation compares to visual aggregation. Could a visualization using aggregation on a small display be better than visual aggregation on a large display? Or do users find more complex multi-scale patterns than ever possible before? Another question is in regard to the order of graphical encodings. A variety of reasons for potential differences in effectiveness were discussed, but those were hypotheses and need to be tested. For example, is color truly a less effective encoding on large displays? And is size a more effective encoding on low DPI than high DPI displays?

This work has only begun to explore how information visualization design changes as visualizations are scaled-up using large, high resolution displays. We now know that users can benefit from the data that is displayed, that visual acuity is not the limit of information visualization, and that embedded visualizations are an effective technique on large displays. Still, this is only a start and much is left to be explored.

References

1. U.S. Census Bureau; "Race and Hispanic Origin Population Density of the United States: 1990;" last revised July 06, 2001;
<<http://www.census.gov/geo/www/mapGallery/RHOriginPD-1990.html>>
2. *Global FPD Industry 2003 Executive Summary*. [web site] 2003 [cited 2006 5/24/2006]; Available from:
http://www.usdc.org/technical/USDCroadmap_ExcutiveSummary.htm.
3. *Stones & Mortar*. [web site] 2004; Available from:
<http://www.nps.gov/mall/mortar/mortar.htm#vive>.
4. Ahlberg, C. and Wistrand, E. IVEE: an Information Visualization and Exploration Environment. In *Proc. InfoVis 1995*, IEEE (1995), 66-73, 142-143.
5. Alphonse, G.A. and Lubin, J. Psychophysical Requirements for Tiled Large Screen Displays. In *Proc. High-Resolution Displays and Projection Systems*, SPIE (1992), 230-240.
6. Andrienko, G.L. and Andrienko, N.V. Interactive Maps for Visual Data Exploration. *International Journal of Geographical Information Science*, 13,4 (1999), 355-374.
7. Andrienko, N. and Andrienko, G. Interactive Visual Tools to Explore Spatio-Temporal Variation. In *Proc. AVI*, ACM (2004), 417-420.
8. Andrienko, N., Andrienko, G., Peckham, R., and Voss, H. CommonGIS: Common Access to Geographically Referenced Data. In *Proc. Abstracts of 6th EC-GIS Workshop*, European Commission Joint Research Centre (2000), 59-60.
9. Antes, J.R. and Mann, S.W. Global-local Precedence in Picture Processing. *Psychological Research*, 46,(1984), 247-259.
10. Ashdown, M., Oka, K., and Sato, Y., *Combining Head Tracking and Mouse Input for a GUI on Multiple Monitors*, in *Extended Abstracts on Human Factors in Computing Systems (CHI)*. 2005, ACM: Portland, OR, USA. p. 1188-1191.
11. Baldonado, M.Q., Woodruff, A., and Kuchinsky, A. Guidelines for Using Multiple Views in Information Visualization. In *Proc. AVI2000*, (2000), 110-119.
12. Ball, R. and North, C. Effects of Tiled High-Resolution Displays on Basic Visualization and Navigation Tasks. In *Proc. Ext. Abstracts CHI 2005*, ACM Press (2005), 1196-1199.
13. Ball, R., Varghese, M., Sabri, A., Cox, E.D., Fierer, C., Peterson, M., Carstensen, B., and North, C. Evaluating the Benefits of Tiled Displays for Navigating Maps. In *Proc. IASTED-HCI '05*, ACTA Press (2005), 66-71.

14. Bartram, L., Ware, C., and Calvert, T. Filtering and Integrating Visual Information with Motion. *Information Visualization*, 1,1 (2002), 66-79.
15. Baudisch, P. Interacting with Large Displays. *IEEE Computer*, 39,3 (2006), 96-97.
16. Baudisch, P., DeCarlo, D., Duchowski, A.T., and Geisler, W.S. Focusing on the Essentials: Considering Attention in Display Design. *Communications of the ACM*, 46,3 (2003), 60-66.
17. Baudisch, P., Good, N., Bellotti, V., and Schraedley, P. Keeping Things in Context: A Comparative Evaluation of Focus Plus Context Screens, Overviews, and Zooming. In *Proc. CHI '02*, ACM (2002), 259-266.
18. Baudisch, P., Tan, D., Collomb, M., Robbins, D., Hinckley, K., Agrawala, M., Zhao, S., and Ramos, G. Phosphor: Explaining Transitions in the User Interface Using Afterglow Effects. In *Proc. UIST '06*, (2006), 169-178
19. Bederson, B.B., Hollan, J.D., Perlin, K., Meyer, J., Bacon, D., and Furnas, G.W. Pad++: A Zoomable Graphical Sketchpad for Exploring Alternate Interface Physics. *Journal of Visual Languages and Computing*, 7,(1996), 3-31.
20. Benko, H., Wilson, A., and Baudisch, P. Precise Selection Techniques for Multi-Touch Screens. In *Proc. CHI '06*, ACM (2006), 1263-1272.
21. Bertin, J. *Graphics and Graphic Information Processing*. Walter de Gruyter, Berlin,1981.
22. Bertin, J., *Graphics and Graphic Information Processing*, in *Readings in Information Visualization: Using Vision to Think*, S.K. Card, J.D. Mackinlay, and B. Shneiderman, Editors. 1999, Morgan Kaufmann Publishers: San Francisco, California. p. 62-65.
23. Beshers, C. and Feiner, S. AutoVisual: Rule-based Design of Interactive Multivariate Visualizations. *IEEE Computer Graphics and Applications*, 13,4 (1993), 41-49.
24. Bezerianos, A. and Balakrishnan, R. The Vacuum: Facilitating the Manipulation of Distant Objects. In *Proc. CHI '05*, ACM (2005), 361-370.
25. Boukhelifa, N., Roberts, J.C., and Rodgers, P.J. A Coordination Model for Exploratory Multi-View Visualization. In *Proc. Coordinated and Multiple Views in Exploratory Visualization*, IEEE (2003), 76-85.
26. Bruce, V., Green, P.R., and Georgeson, M.A. *Visual Perception: Physiology, Psychology, and Ecology*. Psychology Press, Hove, East Sussex U.K.,1996.

27. Buennemeyer, T., Booker, J., Sabri, A., and North, C. Intelligence Analysis Using High Resolution Displays. In *Proc. IEEE ISI '06*, (2006), 652-653.
28. Burns, C.M. Putting It All Together: Improving Display Integration in Ecological Displays. *Human Factors*, 42,2 (2000), 226-241.
29. Card, S.K. and Mackinlay, J.D. The Structure of the Information Visualization Design Space. In *Proc. Symposium on Information Visualization*, IEEE Computer Society (1997), 92-99.
30. Card, S.K., Robertson, G.G., and Mackinlay, J.D. The Information Visualizer, an Information Workspace. In *Proc. CHI '91*, ACM (1991), 181-188.
31. Carpenter, P.A., Just, M.A., and Shell, P. What One Intelligence Tests Measures: A Theoretical Account of the Processing in the Raven Progressive Matrices Test. *Psychological Review*, 97,(1990), 404-431.
32. Carrasco, M., Evert, D., Chang, I., and Katz, S. The Eccentricity Effect: Target Eccentricity Affects Performance on Conjunction Searches. *Perception & Psychophysics*, 57,8 (1995), 1241-1261.
33. Carswell, C.M. and Wickens, C.D. Mixing and Matching Lower-Level Codes for Object Displays: Evidence for Two Sources of Proximity Compatibility. *Human Factors*, 38,1 (1996), 1-22.
34. Cheng, K. and Pulo, K. Direct Interaction with Large-Scale Display Systems Using Infrared Laser Tracking Devices. In *Proc. Asia-Pacific Symposium on Information Visualization*, Australian Computer Society (2003), 67-74.
35. Chi, E.H.-H. and Riedl, J.T. An Operator Interaction Framework for Visualization Systems. In *Proc. Symposium on Information Visualization*, IEEE (1998), 63-70.
36. Chi, E.H. A Taxonomy of Visualization Techniques Using the Data State Reference Model. In *Proc. Information Visualization*, IEEE Computer Society (2000), 69-75.
37. Chuah, M.C. and Roth, S.F. On the Semantics of Interactive Visualizations. In *Proc. Information Visualization*, IEEE (1996), 29-36.
38. Cleveland, W.S. and McGill, R. Graphical Perception: Theory, Experimentation and Application to the Development of Graphical Methods. *Journal of the American Statistical Association*, 79,387 (1984), 531-554.
39. Collomb, M., Hascoet, M., Baudisch, P., and Lee, B. Improving Drag-and-Drop on Wall-Size Displays. In *Proc. GI '05*, (2005), 25-32.
40. Convertino, G., Chen, J., Yost, B., Ryu, Y., and North, C. Exploring Context Switching and Cognition in Dual-View Coordinated Visualizations. In *Proc.*

- IEEE Conference on Coordinated and Multiple Views in Exploratory Visualization (CMV)*, (2003), 55-62.
41. Czerwinski, M., Robertson, G., Meyers, B., Smith, G., Robbins, D., and Tan, D. Large Display Research Overview. In *Proc. CHI 2006*, ACM Press (2006), 69-74.
 42. Czerwinski, M., Smith, G., Regan, T., Meyers, B., Robertson, G., and Starkweather, G. Toward Characterizing the Productivity Benefits of Very Large Displays. In *Proc. INTERACT '03*, IOS (2003), 9-16.
 43. Czerwinski, M., Tan, D., and Robertson, G. Women Take a Wider View. In *Proc. CHI 2002*, ACM Press (2002), 195-202.
 44. Donk, M. Human Monitoring Behavior in a Multiple-Instrument Setting: Independent Sampling, Sequential Sampling or Arrangement-Dependent Sampling. *Acta Psychologica*, 86,(1994), 31-55.
 45. Dourish, P. *Where the Action Is: The Foundations of Embodied Interaction*. MIT Press, Cambridge,2001.
 46. Dykes, J.A. Exploring Spatial Data Representation with Dynamic Graphics. *Computers & Geosciences*, 23,4 (1997), 345-370.
 47. Edsall, R.M., Kraak, M.J., MacEachren, A.M., and Peuquet, D.J. Assessing the Effectiveness of Temporal Legends in Environmental Visualization. In *Proc. GIS/LIS '97*, (1997), 677-685.
 48. Edsall, R.M., MacEachren, A.M., and Pickle, L. Case Study: Design and Assessment of an Enhanced Geographic Information System for Exploration of Multivariate Health Statistics. In *Proc. IEEE Symposium on Information Visualization*, IEEE (2001), 159-162.
 49. Eick, S.G. and Karr, A.F. Visual Scalability. *Journal of Computational & Graphical Statistics*, 11,1 (2002), 22-43.
 50. Elkind, J.I., Card, S.K., Hochberg, J., and Huey, B.M. *Human Performance Models for Computer-Aided Engineering*. Academic Press, Orlando, FL,1990.
 51. Elrod, S., Bruce, R., Gold, R., Goldberg, D., Halasz, F., Janssen, W., Lee, D., McCall, K., Pedersen, E., Pier, K., Tank, J., and Welch, B. Liveboard: A Large Interactive Display Supporting Group Meetings, Presentations, and Remote Collaboration. In *Proc. Human Factors in Computing Systems (CHI)*, ACM (1992), 599-607.
 52. Enns, J.T. *The Thinking Eye, the Seeing Brain: Explorations in Visual Cognition*. W.W. Norton & Company, 2004.

53. Fekete, J.D. and Plaisant, C. Interactive Information Visualization of a Million Items. In *Proc. IEEE Symposium on Information Visualization*, (2002), 117-124.
54. Guimbretiere, F., Stone, M., and Winograd, T. Fluid Interaction with High-Resolution Wall Size Displays. In *Proc. UIST '01*, ACM (2001), 11-14.
55. Hao, M.C., Keim, D.A., Dayal, U., Schneidewing, J., and Wright, P. Geo Pixel Bar Charts. In *Proc. Visualization*, IEEE (2003), 89.
56. Harrower, M., MacEachren, A.M., and Griffin, A.L. Developing a Geographic Visualization Tool to Support Earth Science Learning. *Cartography and Geographic Information Society*, 27,4 (2000), 279-293.
57. Healey, C.G. Choosing Effective Colours for Data Visualization. In *Proc. IEEE Visualization*, IEEE (1996), 263-270.
58. Healey, C.G. Fundamental Issues of Visual Perception for Effective Image Generation. In *Proc. SIGGRAPH 99 Course 6*, (1999), 1-42.
59. Healey, C.G. *Pexels: Visualization with Perceptual Textures*. 2003 May 24, 2003 [cited 2007 March 15]; Available from: http://www.csc.ncsu.edu/faculty/healey/HTML_papers/pexel/pexel.html.
60. Healey, C.G., St. Amant, R., and Elhaddad, M. ViA: A Perceptual Visualization Assistant. In *Proc. 28th Applied Imagery Pattern Recognition Workshop*, (1999), 1-11.
61. Hereld, M., Judson, I.R., and Stevens, R.L. Tutorial: Introduction to Building Projection-based Tiled Display Systems. *IEEE Computer Graphics and Applications*, 20,4 (2000), 22-28.
62. Hornof, A.J. and Halverson, T. Cognitive Strategies and Eye Movements for Searching Hierarchical Computer Displays. In *Proc. CHI 2003*, ACM (2003), 249-264.
63. Huber, P. Huge Data Sets. In *Proc. Compstat '94*, Physica (1994), 3-13.
64. Humphreys, G., Houston, M., Ng, R., Frank, R., Ahern, S., Kirchner, P.D., and Klosowski, J.T. Chromium: A Stream-Processing Framework for Interactive Rendering on Clusters. *ACM Transactions on Graphics*, 21,3 (2002), 693-702.
65. Inselberg, A. Multidimensional Detective. In *Proc. InfoVis '97*, IEEE (1997), 100-107.
66. Jedrysik, P.A., Moore, J.A., Stedman, T.A., and Sweed, R.H. Interactive Displays for Command and Control. In *Proc. Aerospace*, IEEE (2000), 341-351.

67. Jiang, Y., Olson, I.R., and Chun, M.M. Organization of Visual Short-Term Memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26,3 (2000), 683-702.
68. Johanson, B., Fox, A., and Winograd, T. The Interactive Workspaces Project: Experiences with Ubiquitous Computing Rooms. *IEEE Pervasive Computing*, 1,2 (2002), 67-74.
69. Julesz, B. Visual Pattern Discrimination. *IRE Transactions on Information Theory*, 8,2 (1962), 84-92.
70. Julesz, B. *Dialogues on Perception*. MIT Press, Cambridge, MA,1995.
71. Julesz, B., Gilbert, E.N., and Victor, J. Visual Discrimination of Textures with Identical Third-order Statistics. *Biological Cybernetics*, 31,(1978), 137-140.
72. Kinchla, R.A. and Wolfe, J.M. The Order of Visual Processing: Top-down, Bottom-up, or Middle-out. *Perception & Psychophysics*, 25,225-231 (1979).
73. Kraak, M. The Space-Time Cube Revisited From a Geovisualization Perspective. In *Proc. International Cartographic Conference*, (2003), 1988-1997.
74. Leung, Y. and Apperley, M. A Review and Taxonomy of Distortion-Oriented Presentation Techniques. *ACM ToCHI*, 1,2 (1994), 126-160.
75. Li, K., Chen, H., Chen, Y., Clark, D.W., Cook, P., Damianakis, S., Essl, G., Finkelstein, A., Funkhouser, T., Housel, T., Klein, A., Liu, Z., Praun, E., Samanta, R., Shedd, B., Singh, J.P., Tzanetakis, G., and Zheng, J. Building and Using a Scalable Display Wall System. *IEEE Computer Graphics and Applications*, 20,4 (2000), 29-37.
76. Li, Q., *Ph.D. Proposal*. 2005, Virginia Tech: Blacksburg, VA.
77. Li, Q. and North, C. Empirical Comparison of Dynamic Query Sliders vs Brushing Histograms. In *Proc. IEEE Symposium on Information Visualization 2003*, IEEE (2003), 147-154.
78. Lloyd, R. *Spatial Cognition, Geographic Environments*. Kluwer Academic Publishers, Dordrecht,1997.
79. Lombardo, J., Burkom, H., Elbert, E., Magruder, S., Lewis, S.H., Loschen, W., Sari, J., Sniegoski, C., Wojcik, R., and Pavlin, J. A Systems Overview of the Electronic Surveillance System for the Early Notification of Community-Based Epidemics (ESSENCE II). *Journal of Urban Health*, 80,(2003), 32-42.
80. Luebbering, C. and Carstensen, L.W. Map Reading Efficiency on Variable Format High-Resolution Displays. In *Proc. Presented at The Association of American Geographers (AAG) 2006 Annual Meeting*, (2006),

81. MacEachren, A.M. *How Maps Work: Representation, Visualization, and Design*. The Guilford Press, New York, 1995.
82. MacEachren, A.M., Boscoe, F.P., Haug, D., and Pickle, L.W. Geographic Visualization: Designing Manipulable Maps for Exploring Temporally Varying Georeferenced Statistics. In *Proc. Information Visualization '98*, IEEE Computer Society (1998), 87-94.
83. Mackinlay, J. Automating the Design of Graphical Presentations. *ACM Transactions on Graphics*, 5,2 (1986), 110-141.
84. Mackinlay, J.D., Good, L., Zellweger, P., Stefik, M., and Baudisch, P. City Lights: Contextual Views in Minimal Space. In *Proc. CHI '03*, (2003), 838-839.
85. Mackinlay, J.D. and Heer, J. Wideband Displays: Mitigating Multiple Monitor Seams. In *Proc. CHI '04*, ACM (2004), 1521-1524.
86. Mark, D.M., Freska, C., Hirtle, S.C., Lloyd, R., and Tversky, B. Cognitive Models of Geographical Space. *International Journal of Geographical Information Science*, 13,8 (1999), 747-774.
87. Matin, E., Shao, K.C., and Boff, K.R. Saccadic Overhead: Information-Processing Time With and Without Saccades. *Perception & Psychophysics*, 53,4 (1993), 372-380.
88. McDonald, J.A., Stuetzle, W., and Buja, A. Painting Multiple Views of Complex Objects. In *Proc. ACM ECOOP/OOPSLA*, ACM (1990), 245-257.
89. Miller, G.A. The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information. *Psychological Review*, 63,(1956), 81-97.
90. Miyake, A. and Shah, P. *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control*. Cambridge University Press, New York, 1999.
91. Monmonier, M. Strategies for the Visualization of Geographic Time-Series Data. *Cartographica*, 27,1 (1990), 30-45.
92. Muller, W. and Schumann, H. Visualization Methods for Time-Dependent Data - An Overview. In *Proc. Winter Simulation Conference*, (2003), 737-745.
93. Myers, B., Bhatnagar, R., Nichols, J., Peck, C., Kong, D., Miller, R., and Long, C. Interacting at a Distance: Measuring the Performance of Laser Pointers and Other Devices. In *Proc. CHI '02*, ACM (2002), 33-40.

94. Ni, T., Bowman, D., and Chen, J. Increased Display Size and Resolution Improve Task Performance in Information-Rich Virtual Environments. In *Proc. GI '06*, (2006), 139-146.
95. Ni, T., Schmidt, G.S., Stadt, O.G., Livingston, M.A., Ball, R., and May, R. A Survey of Large High-Resolution Display Technologies, Techniques, and Applications. In *Proc. IEEE Virtual Reality*, (2006), 223-234.
96. North, C. *A User Interface for Coordinating Visualizations Based on Relational Schemata: Snap-Together Visualization*, Ph.D. Dissertation, University of Maryland, (May 2000).
97. North, C. Multiple Views and Tight Coupling in Visualization: A Language, Taxonomy, and System. In *Proc. CSREA CISST 2001 Workshop on Fundamental Issues in Visualization*, (2001), 626-632.
98. North, C., Conklin, N., Idukuri, K., and Saini, V. Visualization Schemas and a Web-based Architecture for Custom Multiple-View Visualization of Multiple-Table Databases. *Information Visualization*, 1,3-4 (2002), 211-228.
99. North, C. and Shneiderman, B., *A Taxonomy of Multiple Window Coordination*. 1997, University of Maryland, College Park.
100. Nowell, L., Netzler, E., and Tanasse, T. Change Blindness in Information Visualization: A Case Study. In *Proc. IEEE Symposium on Information Visualization*, IEEE (2001), 15-22.
101. Pfitzner, D., Hobbs, V., and Powers, D. A Unified Taxonomic Framework for Information Visualization. In *Proc. Asia-Pacific Symposium on Information Visualisation*, ACM (2003), 57-66.
102. Phillips, W.A. On the Distinction Between Sensory Storage and Short-Term Visual Memory. *Perception & Psychophysics*, 16,(1974), 283-290.
103. Polys, N., Bowman, D., North, C., Laubenbacher, R., Duca, K.,. PathSim Visualizer: An Information-Rich Virtual Environment for Systems Biology. In *Proc. Web3D Symposium*, ACM Press (2004d), 7-14.
104. Polys, N., Kim, S., and Bowman, D. Effects of Information Layout, Screen Size, and Field of View on User Performance in Information-Rich Virtual Environments. In *Proc. VRST '05*, ACM (2005), 46-55.
105. Polys, N.F. *Display Techniques in Information-Rich Virtual Environments*, Ph.D. Dissertation, Virginia Polytechnic Institute and State University,
106. Rantanen, E. and Goldberg, J. The Effect of Mental Workload on the Visual Field Size and Shape. *Ergonomics*, 42,6 (1999), 816-834.

107. Rao, R. and Card, S.K. The Table Lens: Merging Graphical and Symbolic Representation in an Interactive Focus+Context Visualization for Tabular Information. In *Proc. CHI*, ACM (1994), 318-322.
108. Ratwani, R.M., Trafton, J.G., and Boehm-Davis, D.A. From Specific Information Extraction to Inference: A Hierarchical Framework of Graph Comprehension. In *Proc. HFES*, (2004),
109. Roberts, J.C. On Encouraging Multiple Views for Visualization. In *Proc. IV'98*, (1998), 8-14.
110. Robertson, G., Czerwinski, M., Baudisch, P., Meyers, B., Robbins, D., Smith, G., and Tan, D. The Large-Display User Experience. *IEEE Computer Graphics and Applications*, 25,4 (2005), 44-51.
111. Rock, I. and Mack, A., *Attention and perceptual organization*, in *Cognitive approaches to human perception*, S. Ballesteros, Editor. 1994, Lawrence Erlbaum Associates: Hillsdale, NJ. p. 23-41.
112. Roth, S.F., Kolojejchick, J., Mattis, J., and Goldstein, J. Interactive Graphic Design Using Automatic Presentation Knowledge. In *Proc. Conference on Human Factors in Computing Systems (SIGCHI '94)*, (1994), 112-117.
113. Roth, S.F. and Mattis, J. Automating the Presentation of Information. In *Proc. IEEE Conference on Artificial Intelligence Applications*, IEEE (1991), 90-97.
114. Ryu, Y., Yost, B., Convertino, G., Chen, J., and North, C. Exploring Cognitive Strategies for Integrating Multiple-View Visualizations. In *Proc. Human Factors and Ergonomics Society (HFES) 47th Annual Meeting*, (2003), 591-595.
115. Sabri, A., Ball, R., Bhatia, S., Fabian, A., and North, C. High-Resolution Gaming: Interfaces, Notifications and the User Experience. *Interacting with Computers Journal*, 19,2 (2007), (To appear).
116. Sandstrom, T.A., Henze, C., and Levit, C. The Hyperwall. In *Proc. Coordinated and Multiple Views in Exploratory Visualization (CMV '03)*, IEEE (2003), 124-133.
117. Saraiya, P., Lee, P., and North, C. Visualization of Graphs with Associated Timeseries Data. In *Proc. Information Visualization*, IEEE (2005), 225-232.
118. Saraiya, P., North, C., and Duca, K. Evaluation of Microarray Visualization Tools for Biological Insight. In *Proc. IEEE Symposium on Information Visualization*, IEEE (2004), 1-8.
119. Saraiya, P.B. *Insight-Based Studies for Pathway and Microarray Visualization Tools*, Ph.D. Dissertation, Virginia Polytechnic Institute and State University, (July 2006).

120. Sawant, A.P. and Healey, C.G., *A Survey of Display Device Properties and Visual Acuity for Visualization*. 2005, Department of Computer Science, North Carolina State University.
121. Shneiderman, B. The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *Proc. IEEE Visual Languages*, (1996), 336-343.
122. Shupp, L., Ball, R., Yost, B., Booker, J., and North, C. Evaluation of Viewport Size and Curvature of Large, High-Resolution Displays. In *Proc. Graphics Interface (GI)*, (2006), 123-130.
123. Slocum, T.A., Blok, C., Jiang, B., Koussoulakou, A., Montello, D.R., Fuhrmann, S., and Hedley, N.R. Cognitive and Usability Issues in Geovisualization. *Cartography and Geographic Information Society Journal*, 28,1 (2001).
124. Staadt, O.G., Walker, J., Nuber, C., and Hamann, B., *A Survey and Performance Analysis of Software Platforms for Interactive Cluster-Based Multi-Screen Rendering*, in *Workshop on Virtual Environments*. 2003, ACM: Zurich, Switzerland.
125. Starkweather, G.K. DSHARP-A Wide-Screen Multi-Projector Display. *Journal of Optics: Pure and Applied Optics*, 5,5 (2003), S136-S139.
126. Streitz, N.A., Geibler, J., Holmer, T., Konomi, S.-i., Muller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., and Steinmetz, R. i-LAND: An Interactive Landscape for Creativity and Innovation. In *Proc. Human Factors in Computing Systems (CHI)*, ACM (1999), 120-127.
127. Swaminathan, K. and Sato, S. Interaction Design for Large Displays. *Interactions*, 4,1 (1997), 15-24.
128. Takatsuka, M. and Gahegan, M. GeoVISTA Studio: A Codeless Visual Programming Environment For Geoscientific Data Analysis and Visualization. *The Journal of Computers and Geosciences*, (2002).
129. Tan, D., Czerwinski, M., and Robertson, G. Women Go with the (Optical) Flow. In *Proc. CHI 2003*, ACM (2003), 209-215.
130. Tan, D., Gergle, D., Scupelli, P., and Pausch, R. With Similar Visual Angles, Larger Display Improve Spatial Performance. In *Proc. CHI 2003*, ACM Press (2003), 217-224.
131. Tan, D.S. and Czerwinski, M. Effects of Visual Separation and Physical Discontinuities When Distributing Information Across Multiple Displays. In *Proc. Computer-Human Interaction Special Interest Group of the Ergonomics Society of Australia (OZCHI '03)*, (2003), 184-191.

132. Taub, G.E. A Confirmatory Analysis of the Wechsler Adult Intelligence Scale-Third Edition: Is The Verbal/Performance Discrepancy Justified? *Practical Assessment, Research & Evaluation*, 7,22 (2001), Electronic journal; article has 9 pages.
133. Tessendorf, D., Chewar, C.M., Ndiwalana, A., Pryor, J., McCrickard, D.S., and North, C. An Ordering of Secondary Task Display Attributes. In *Proc. CHI '02*, ACM (2002), 600-601.
134. Tory, M. and Moller, T., *A Model-Based Visualization Taxonomy* 2002, Simon Fraser University.
135. Trafton, J.G., Kirschenbaum, S.S., Tsu, T.L., Miyamoto, R.T., Ballas, J.A., and Raymond, P.D. Turning Pictures Into Numbers: Extracting and Generating Information from Complex Visualizations. *International Journal of Human-Computer Studies*, 53,5 (2000), 827-850.
136. Tufte, E.R. *Visual Display of Quantitative Information*. Graphics Press, Cheshire, CT,1983.
137. Tweedie, L. Characterizing Interactive Externalizations. In *Proc. Human Factors in Computing Systemes (CHI)*, ACM SIGCHI (1997), 375-382.
138. Vogel, E.K., Woodman, G.F., and Luck, S.J. Storage of Features, Conjunctions, and Objects in Visual Working Memory. *Journal of Experimental Psychology: Human Perception and Performance*, 27,1 (2001), 92-114.
139. Ware, C. *Information Visualization: Perception for Design*. Morgan Kaufman, 2000.
140. Wegman, E. Huge Datasets and the Frontiers of Computational Feasibility. *Computational and Graphical Statistics*, 4,4 (1995), 281-295.
141. Wei, B., Silva, C., Koutsofios, E., Krishnan, S., and North, S. Visualization Research with Large Displays. *IEEE Computer Graphics and Applications*, 20,4 (2000), 50-54.
142. Wickens, C.D. and Carswell, C.M. The Proximity Compatibility Principle: Its Psychological Foundation and Relevance to Display Design. *Human Factors*, 37,3 (1995), 473-494.
143. Wickens, C.D. and Hollands, J. *Engineering Psychology and Human Performance*. Prentice Hall, Upper Saddle River, New Jersey,2000.
144. Wickens, C.D., Vincow, M.A., Schopper, A.W., and Lincoln, J.E., *Computational Models of Human Performance in the Design and Layout of Controls and Displays*. 1997, Wright-Patterson AFB, Crew Systems Ergonomics Information Analysis Center: Dayton, OH.

145. Wooding, D.S. Fixation Maps: Quantifying Eye-Movement Traces. In *Proc. ETRA'2002*, (2002), 31-36.
146. Woodruff, Landay, J., and Stonebraker, M. Constant Information Density in Zoomable Interfaces. In *Proc. AVI '98*, (1998), 57-65.
147. Yost, B., Haciaahmetoglu, Y., and North, C. Beyond Visual Acuity: The Perceptual Scalability of Information Visualizations for Large Displays. In *Proc. Human Factors in Computing Systems (CHI)*, ACM (2007), to appear.
148. Yost, B. and North, C. Single Complex Glyphs Versus Multiple Simple Glyphs. In *Proc. Extended Abstracts of Human Factors in Computing Systems (CHI)*, ACM (2005), 1889-1892.
149. Yost, B. and North, C. The Perceptual Scalability of Visualization. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of Visualization/Information Visualization 2006)*, 12,5 (2006), 837-844.
150. Yost, B., Qiao, Y., Fang, B., Wu, Y., and North, C., *Visual vs. Linked Integration of Quantitative Attributes in Information Visualization*. 2005, Virginia Tech. p. 8.