

A Comparison of Pan and Zoom and Rubber Sheet Navigation

by

Dmitry Nekrasovski

B.C.S., Carleton University, 2000

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Master of Science

in

The Faculty of Graduate Studies

(Computer Science)

The University Of British Columbia

January 2006

© Dmitry Nekrasovski 2006

Abstract

As information visualization tools are used to visualize datasets of increasing size, there is a growing need for techniques that facilitate efficient navigation. Pan and zoom navigation enables users to display areas of interest at different resolutions. Focus+context techniques aim to overcome the drawbacks of pan and zoom by dynamically integrating areas of interest and context regions. To date, empirical comparisons of these two navigation paradigms have been limited in scope and inconclusive.

In two controlled studies, we evaluated navigation techniques representative of the pan and zoom and focus+context approaches. The particular focus+context technique examined was rubber sheet navigation, implemented in a way that afforded a set of navigation actions similar to pan and zoom navigation. The two techniques were used by 40 subjects in each study to perform a navigation-intensive task in a large tree dataset. Study 1 investigated the effect of the amount of screen real estate devoted to context regions for each navigation technique. Performance with both techniques was not significantly affected by this factor, but was influenced by technique-specific strategies developed by subjects. Study 2 compared the performance of the two techniques. Pan and zoom navigation was found to be faster than rubber sheet navigation and was rated by subjects as easier and less mentally demanding. We discuss the implications of these results, including the relationship between navigation technique, task, and user strategy, and propose directions for future work.

Contents

Abstract	ii
Contents	iii
List of Tables	vi
List of Figures	vii
Acknowledgements	ix
1 Introduction	1
1.1 Motivation	1
1.2 Overview	2
2 Related Work	7
2.1 Pan and Zoom Navigation (PZN)	7
2.2 Rubber Sheet Navigation (RSN)	9
2.3 Guaranteed Visibility	13
2.4 Evaluation of Navigation Techniques	14
2.5 Evaluation of Tree Visualizations	17
3 Task and Dataset	19
3.1 Task	19
3.1.1 Task 1: Determining the lowest common ancestor	20
3.1.2 Task 2: Determining the topological distance between nodes	21
3.1.3 Task 3: Determining whether two subtrees are adjacent	21

3.1.4	Task 4: Determining whether a subtree contains unmarked nodes	22
3.2	Dataset	23
4	Study 1	25
4.1	Hypotheses	26
4.2	Interfaces	26
4.2.1	Navigation	29
4.2.2	Overviews and Foci	30
4.2.3	Guaranteed Visibility and Levels of Context	30
4.2.4	RSN-NoOV Interface	34
4.2.5	PZN-NoOV Interface	34
4.2.6	RSN+OV Interface	35
4.2.7	PZN+OV Interface	35
4.3	Task and Dataset	36
4.4	Apparatus	36
4.5	Participants	36
4.6	Design	37
4.7	Procedure	38
4.8	Measures	39
4.9	Results	40
4.10	Summary of Results	42
4.11	Strategies	45
4.12	Discussion	46
5	Study 2	49
5.1	Hypotheses	49
5.2	Interfaces	50
5.3	Task and Dataset	50
5.4	Apparatus	54
5.5	Participants	54

5.6 Design	54
5.7 Procedure	55
5.8 Measures	57
5.9 Results	57
5.10 Summary of Results	61
5.11 Discussion	61
6 Conclusions and Future Work	64
Bibliography	67
A Study 1 Training Protocol	74
B Study 2 Training Protocol	81
C Study 1 Questionnaires	90
D Study 2 Questionnaires	110

List of Tables

4.1	Interface combinations examined in Studies 1 and 2	25
4.2	Latin square used to counterbalance the order of presentation for level of context in Study 1	37
4.3	Levels of context used in Study 1	38
4.4	Results of ANOVAs of level of context by interface in Study 1 . .	41
4.5	Navigation strategies developed by subjects in Study 1	46
5.1	Means and effects of navigation technique in Study 2	59

List of Figures

1.1	Pan and zoom navigation (PZN)	3
1.2	Rubber sheet navigation (RSN)	3
2.1	Panning metaphors examined by Johnson	8
2.2	Rubber sheet navigation with orthogonal stretching	11
2.3	Rubber sheet navigation with polygonal stretching	11
2.4	Rubber sheet navigation in TreeJuxtaposer	12
2.5	Guaranteed visibility in CityLights [51]	15
2.6	Guaranteed visibility in pan and zoom interfaces	15
3.1	Task 1: Determining the lowest common ancestor (LCA).	20
3.2	Task 2: Determining the topological distance between nodes.	21
3.3	Task 3: Determining whether two subtrees are adjacent.	22
3.4	Task 4: Determining whether a subtree contains unmarked nodes.	22
4.1	Study 1 RSN-NoOV interface	27
4.2	Study 1 PZN-NoOV interface	27
4.3	Study 1 RSN+OV interface	28
4.4	Study 1 PZN+OV interface	28
4.5	Motivating scenario for use of multiple foci in Study 1	31
4.6	Calculation of levels of context in Study 1 RSN interfaces	33
4.7	Calculation of levels of context in Study 1 PZN interfaces	33
4.8	Mean per-trial completion times by interface in Study 1	41

4.9	Mean per-trial completion times and trend line for the RSN-NoOV interface	43
4.10	Mean per-trial completion times and trend line for the PZN-NoOV interface	43
4.11	Mean per-trial completion times and trend line for the RSN+OV interface	44
4.12	Mean per-trial completion times and trend line for the PZN+OV interface	44
5.1	Study 2 RSN-NoOV interface	51
5.2	Study 2 PZN-NoOV interface	51
5.3	Study 2 RSN+OV interface	52
5.4	Study 2 PZN+OV interface	52
5.5	Calculation of levels of context in Study 2 RSN interfaces	53
5.6	Calculation of levels of context in Study 2 PZN interfaces	53
5.7	Mean per-trial completion times by interface in Study 2	58
5.8	Mean per-trial completion times by navigation technique and by block	60

Acknowledgements

Research is always a collaborative process, and this thesis is no exception. I would like to thank all the people who have helped me in making it happen.

My supervisors, Joanna McGrenere and Tamara Munzner, provided copious amounts of advice, guidance, and support throughout this project. It has been a fantastic learning experience and a real pleasure to work with both of them.

Adam Bodnar collaborated with me on most aspects of the work described in this thesis. His hard work, creative ideas, and ready wit have been essential to the success of this project (and the preservation of my sanity!)

François Guimbretière committed a great deal of time to help shape and direct this work. His expert advice and attention to detail have proven invaluable.

Karon Maclean somehow found time to act as second reader for this thesis among all her other commitments. Her helpful input is greatly appreciated.

Many other people contributed to this work in a variety of ways. Motivation for the studies came from discussions with phylogenetic biologists, especially David Hillis, Wayne Maddison, and members of their research groups. Early feedback was provided by Patrick Baudisch, Ben Bederson, Jean-Daniel Fekete, Catherine Plaisant, and Katherine St. John. Later on, study design and analysis was influenced by conversations with Leah Findlater, Heidi Lam, Barry Po, and Melanie Tory, while James Slack assisted with implementation issues. Members of the Imager, MUX, and HCT labs helped out as pilot study subjects and created a stimulating and fun environment for conducting this research.

Last but most certainly not least, Julie Lavoie and my parents, Alex and Sofia Nekrasovski, have given me boundless love and constant encouragement. This is for you.

Chapter 1

Introduction

1.1 Motivation

Information visualization uses computer-supported interactive visual representations of abstract data to aid cognition [13]. Information visualization techniques are now being applied to fields such as phylogenetic biology, which is concerned with discovering evolutionary relationships between species, and requires increasingly sophisticated visualization tools for this purpose. As information visualization tools are used to visualize datasets of increased size and complexity, there is a growing need for techniques that facilitate rapid and efficient navigation in such datasets. Two primary approaches have been proposed in the information visualization literature to enable such navigation. **Pan and zoom navigation** (PZN) [24] relies on a combination of panning and zooming operations to enable users to view discrete portions of the dataset at different resolutions. Pan and zoom navigation techniques are often paired with **overview windows** to provide users with contextual information about areas outside the region of current interest. However, overview windows take up screen real estate and may not provide enough resolution to clearly identify features of interest. **Focus+context** (F+C) navigation techniques [11] combine high-resolution displays of areas of interest to the user (focus regions) and contextual information about the rest of the dataset (context regions) into a single unified view. In order to accomplish this without sacrificing screen real estate, many focus+context interfaces rely on distortion to dynamically integrate focus and context regions as users navigate through the dataset. Some researchers

have suggested that focus+context techniques may also benefit from the use of overview windows [5]. There has been no consensus in the literature as to which approach is superior, and some researchers have suggested that their relative effectiveness is highly dependent on the particular navigation task for which they are used [22, 25].

The relative performance of pan and zoom and focus+context navigation techniques can be influenced by a variety of factors. These include the type of dataset being navigated and its visual representation, the nature and level of difficulty of the navigation task, the interactions afforded by the navigation techniques, and the fraction of screen real estate allotted to context regions, whether integrated with focus regions or presented in a separate overview window. This last factor is referred to in this thesis as **level of context**. The dataset types that most commonly motivate research in the area of navigation in the information visualization literature due to their use in a variety of application domains include textual documents [5, 25], maps [22, 24], and graphs, especially tree structures [30, 43]. A variety of navigation tasks can be performed with each of these dataset types. Examples include visual search, browsing, comparison, and more complex compound tasks comprised of multiple instances of these tasks. One category of task that may be particularly suitable for assessing performance of navigation techniques due to their potential for requiring significant amounts of navigation is tasks that involve understanding the topology of a graph or tree structure. The interactions afforded by variants on the pan and zoom and focus+context navigation metaphors often include some or all of panning, zooming in, and zooming out, all of which can be implemented in a variety of ways depending on the design goals and intended dataset of a particular navigation technique. Level of context affects the amount of contextual information available to users during navigation, and may remain static or change dynamically depending on user interactions.

In recent years, empirical evaluation has gained increased prominence in information visualization literature [16]. Recently published user studies have evaluated commercially available visualization tools [30], compared the usability

of different visualization techniques for specific tasks [25], and examined patterns associated with the use of visualization tools in the field [34]. A number of these studies have either compared the performance of different navigation techniques or examined interfaces with and without an overview based on the same navigation technique. However, these investigations have typically involved a relatively narrow subset of variations on the pan and zoom and focus+context paradigms, as well as differences in visual representation and interactions afforded by interfaces using each of the two types of navigation, making study results difficult to interpret. Furthermore, the datasets and tasks used in most of these studies have been devised specifically for study purposes, without regard to the needs of users in a particular domain, and no study has evaluated navigation technique and presence of overview as orthogonal factors. Finally, to date level of context has not been examined as a factor in empirical studies of navigation techniques. The work described in this thesis represents a first attempt to fill these gaps in the literature.

1.2 Overview

This thesis describes experiments performed to quantitatively evaluate the effect of navigation technique (pan and zoom vs. focus+context), presence or absence of an overview, and level of context on user performance and satisfaction. We chose to specifically examine presence or absence of overview and level of context due to the lack of empirical results concerning the relative influence of these factors on performance with each of the two navigation metaphors. The specific focus+context technique that we chose to evaluate is **rubber sheet navigation** (RSN) [45], which allows users to stretch or squish focus areas as though the dataset was laid out on a rubber sheet with its borders nailed down. The differences between the two navigation techniques used in our experiments are illustrated in Figures 1.1 and 1.2. While a zoom action in PZN causes areas outside the selected region of interest to move off-screen, the equivalent operation in RSN causes these areas to be compressed around the edges of the view




Figure 1.1: Selecting (left) and result of zooming into (right) a rectilinear region with pan and zoom navigation. Areas outside the zoomed region are pushed off-screen.

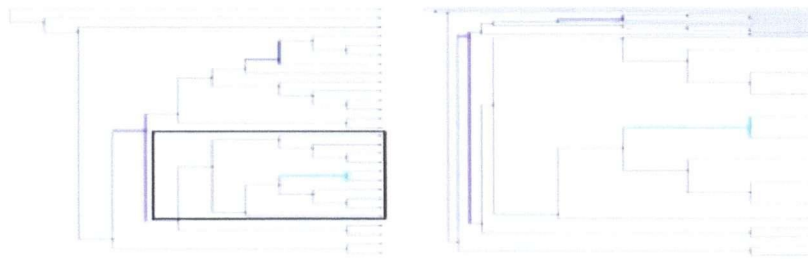


Figure 1.2: Selecting (left) and result of zooming into (right) a rectilinear region with rubber sheet navigation. Areas outside the zoomed region are compressed around the edges of the view.

but remain visible.

We chose RSN as the most appropriate representative technique because of the fact that it is the only distortion-based focus+context technique to date to be combined with **guaranteed visibility** [37], a property that ensures that regions of interest remain visible independent of user navigation actions. While guaranteed visibility is a relatively new concept in the information visualization literature, a recent study [5] suggests that it may provide benefits in terms of both performance and user preference. In an effort to use the best available implementations of each navigation technique for our experiment, in our studies we provided guaranteed visibility in interfaces using both PZN and RSN naviga-

tion. Our use of RSN was also motivated by this technique's similarity to PZN in terms of user interaction. Our implementation of rubber sheet navigation was based on the framework of Slack et al. [47], which, like PZN interfaces, provides multiple levels of magnification and an interaction model based on rectangular selection areas. We were also interested in comparing RSN to PZN because of the lack of empirical evaluation of the former in the literature. A detailed discussion of related work in implementation and evaluation issues for both navigation techniques as well as guaranteed visibility can be found in Chapter 2.

The task used in our study is a topological navigation task motivated by the requirements of phylogenetic biologists, who require sophisticated visualization tools to support their work. Our discussions with phylogenetic biologists lead us to develop a set of compound topological tasks related to their needs, of which we then selected a navigation-intensive task suitable for our comparison. The dataset used in our study is a large tree dataset also derived from phylogenetic biology. We developed abstract versions of the task and the dataset in order to allow us to perform a quantitative study with non-expert users. Further details about the task and dataset used in our studies can be found in Chapter 3.

We conducted two experiments, each involving 40 subjects and the same task and dataset. Each of the experiments also used the same four interfaces, representing all combinations of PZN and RSN with and without overviews, although the design of the interfaces was refined between the studies. The experiments were designed to measure performance by recording the completion times and number of navigation actions required to perform the task with each experimental interface. We also gathered data on self-reported measures such as perceived mental and physical effort, ease of navigation, and ease of use.

Our first study, discussed in Chapter 4, represents the first evaluation of the effect of level of context on performance with different navigation techniques. The results of this study show that level of context did not have a significant impact on performance for any of the interfaces. However, the study's results were strongly affected by differences in navigation strategies developed by subjects to deal with the perceived complexity of the interfaces. This effect motivated us

to simplify the interfaces and develop detailed training strategies for our second study.

Our second study, described in Chapter 5, is the first to evaluate the effects of navigation technique and the presence of an overview as orthogonal factors. We found that subjects performed significantly faster using PZN than RSN regardless of whether an overview was present. Additionally, subjects required fewer navigation interactions and reported a lower mental effort with PZN while completing the task. Our results also indicate that overviews did not appear to improve performance, but were still perceived as beneficial. We discuss the implications of these results, including the relationship between navigation technique and task, and make recommendations for future evaluations of RSN. Chapter 6 outlines the limitations of our experiments, lists some possibilities for future work stemming from our studies, and concludes this thesis.

The research project that comprised the two studies discussed in this thesis was conducted by the author jointly with Adam Bodnar. Within this project, the author was responsible for investigating the effects of navigation technique, while Bodnar investigated the effects of presence or absence of overview. This thesis therefore emphasizes the aspects of the studies related to navigation techniques, while those aspects related to presence of overview are presented in greater detail in Bodnar's master's thesis [12]. As a result, Sections 4.3, 4.4, 4.5, and 4.8 of Chapter 4 and Sections 5.3, 5.4, 5.5, and 5.8 of Chapter 5 are jointly authored with Bodnar, while Chapter 3 is based on a version jointly authored with him. Substantial portions of this thesis also appear in a paper published in the proceedings of the 2006 SIGCHI conference on Human factors in computing systems, which was jointly authored with Bodnar, Joanna McGrenere, François Guimbretière, and Tamara Munzner [38].

To summarize, the remainder of this thesis is organized as follows. Chapter 2 discusses related work and provides a background for our studies. Chapter 3 discusses the task and dataset used in both our studies. The design and results of each experiment are detailed and discussed in Chapters 4 and 5. Finally, Chapter 6 suggests directions for future work and concludes this thesis.

Chapter 2

Related Work

This chapter examines design issues in pan and zoom and rubber sheet navigation interfaces, discusses the concept of guaranteed visibility and its implementations in greater detail, and presents an overview of relevant work in empirical evaluation of navigation techniques and tree visualizations.

2.1 Pan and Zoom Navigation (PZN)

In the context of information visualization, **navigation** can be defined as the traversal of an information structure by selecting parts of the current view of the structure [19]. The dominant metaphor for navigation in information visualization today is **pan and zoom navigation**, which has been used in a variety of experimental systems (see survey in Hornbaek, Bederson, and Plaisant [24]), as well as a number of commercial applications [1, 20]. Pan and zoom navigation combines two classes of navigation techniques: **panning**, which allows users to change the visible region of the dataset through horizontal and vertical translations, and **zooming**, which changes the scale at which the dataset is viewed to allow users to view regions of interest at greater or lesser resolution. In this thesis, pan and zoom navigation is used to mean the combination of these two classes of rigid two-dimensional transformations, as opposed to navigation methods that adapt these techniques for use in distortion-based interfaces [3, 4].

Although panning is one of the most basic techniques for navigating data, there have been relatively few attempts to describe and compare different panning variations. Johnson [27] provides a survey of this literature and describes the following panning metaphors, illustrated in Figure 2.1:




Figure 2.1: Panning metaphors examined by Johnson [27].

1. **Push background:** To view an off-screen region, users pan in the direction opposite to where this region lies, as though moving the background of the visualization.
2. **Push camera:** To view an off-screen region, users pan in the direction this region lies, as though manipulating a camera that is viewing the dataset.
3. **Touch edge camera:** To view an off-screen region, users touch the edge of the view in the direction this region lies.

Johnson compared these three metaphors in a controlled study performed on a touch display. Results showed that Push Background panning was superior to the other kinds of panning in terms of both performance and user preference. Consistent with this finding, we use the Push Background metaphor in our implementation of pan and zoom navigation.

A variety of zooming navigation approaches have been described in the information visualization literature. These primarily differ in terms of how the scale of objects in the dataset is manipulated as users perform zoom in and zoom out actions. Hornbaek et al. [24] describe the following approaches to implementing zooming in terms of scale changes:

1. **Geometric zooming:** The most common approach, where the apparent size of objects increases linearly when users zoom into an area of interest, and decreases at the same rate when they zoom out.

2. **Semantic zooming:** Introduced in Perlin and Fox's Pad system [40], this approach reveals new features in addition to increasing the size of existing ones as users zoom in. This technique is particularly suitable for map datasets, where users are often interested in different kinds of features depending on the scale at which they are viewing the data.
3. **Constant density zooming:** This approach, first used by Woodruff, Landay, and Stonebreaker [50], uses a more complex relationship between scale and appearance, where a constant number of objects is visible regardless of zooming actions.

To date, no study has compared these three approaches to determine whether semantic or constant density zooming offer benefits compared to simple geometric zooming. The implementation of pan and zoom navigation discussed in this thesis therefore relies on geometric zooming for consistency with the majority of the systems documented in the literature.

Two main approaches for implementing scale changes during zoom navigation actions have been described. In **jump zooming** [40], changes of scale occur instantaneously, without intermediate steps, while in **animated zooming** [8], the transition from the old to the new scale is smoothly animated. A study by Bederson and Boltman [7] compared these two methods in a topology recall task. Although no difference in completion time was found, results indicated that the users produced topology reconstructions of higher quality with animated zooming. Based on this result, we use smooth animated transitions during zoom navigation actions in our implementation of pan and zoom.

2.2 Rubber Sheet Navigation (RSN)

As previously mentioned, the major alternative to pan and zoom navigation is the **focus+context** approach, first introduced by Spence and Apperley [48]. Unlike pan and zoom interfaces, which either only present users with regions of current interest or provide contextual information in a separate overview

window, focus+context techniques integrate focus regions and context within a single view [11]. Through this integration, the focus+context approach aims to reduce the cognitive load required for users to maintain a global representation of the dataset and their navigational history [33]. Most focus+context interfaces in the literature integrate focus and context regions using dynamically chosen distortions. Examples of such **distortion-based** techniques include fish-eye views [18, 44], hyperbolic geometry [31], nonlinear magnification [29], and a number of other approaches [35]. Other focus+context approaches that do not rely on distortion include aggregating context regions into glyphs [14, 43] and showing contextual information through layers of lenses [10]. The evaluation discussed in this thesis is intended to be primarily relevant to the literature on distortion-based focus+context interfaces, and in the remainder of this thesis the term “focus+context” is used to refer specifically to these interfaces.

The particular focus+context navigation technique examined in this thesis is **rubber sheet navigation**, originally developed by Sarkar, Snibbe, Tversky, and Reiss [45]. The name of this technique comes from its central metaphor of interacting with the dataset as though it were laid out on a rubber sheet with its borders tacked down. Users can select and stretch or compress arbitrary areas of the rubber sheet, while the rest of the rubber sheet remains visible, though it may be compressed. This approach has the advantage of preserving users’ sense of location in the dataset, which can easily be lost during distortion-based navigation [30]. Sarkar et al. describe two variants of rubber sheet navigation. With **orthogonal stretching**, illustrated in Figure 2.2, users are restricted to selecting vertical or horizontal slices of the dataset areas, which are stretched out without affecting the rest of the dataset. This has the advantage of preserving orthogonal ordering of points in the dataset, but suffers from discontinuity of scale at the boundary between focus and context areas. **Polygonal stretching**, shown in Figure 2.3, enables users to select arbitrary polygons as areas of interest, and smoothly integrates stretched out focus areas with context regions in terms of scale. However, polygonal stretching does not preserve dataset symmetry.




Figure 2.2: Rubber sheet navigation with orthogonal stretching as described in Sarkar et al. [45]. The symmetry of the dataset is preserved at the expense of discontinuities of scale.




Figure 2.3: Rubber sheet navigation with polygonal stretching using a rectangular selection area as described in Sarkar et al. [45]. Focus and context areas are smoothly integrated, but dataset symmetry is not preserved.

The particular version of rubber sheet navigation discussed in this thesis is based on that implemented in the TreeJuxtaposer visualization tool developed by Munzner, Guimbretiere, Taziran, Zhang, and Zhou [37] (see Figure 2.4). The form of rubber sheet navigation used in TreeJuxtaposer represents a middle ground between orthogonal and polygonal stretching as described by Sarkar et al. [45], since it provides smooth integration of focus and context regions while preserving symmetry of tree structures. Similarly to many implementations of pan and zoom navigation, TreeJuxtaposer uses rectilinear regions as selection areas, and provides animated transitions to maintain user context during stretching. These similarities enable an easier comparison between the two types of navigation techniques.

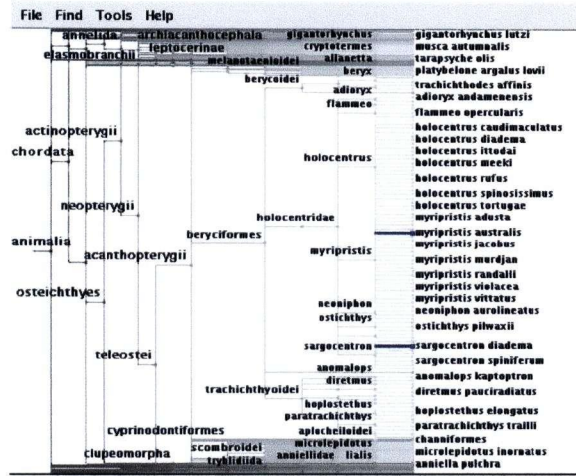


Figure 2.4: Rubber sheet navigation and guaranteed visibility in TreeJuxtaposer [37]. Focus and context regions are smoothly integrated while preserving the symmetry of the tree's topology.

2.3 Guaranteed Visibility

Both pan and zoom and focus+context navigation techniques have been shown to be effective for exploring datasets of up to several hundred items. However, with larger datasets, both classes of navigation techniques often encounter visibility issues, where marked areas of interest to the user, such as landmarks or search results, cannot be seen. In pan and zoom interfaces, this may occur because areas of interest move off-screen due to navigation actions. In its extreme form, the moving of areas of interest off-screen can result in a phenomenon referred to as **desert fog** [28], where the user is faced with a view of the dataset devoid of navigational cues. The marked areas may also become simply too small to be displayed at the set scale and resolution. In focus+context interfaces, areas of interest may be rendered invisible due to the effects of distortion, which can lead them to be culled or aggregated with non-marked areas. A common solution to visibility issues is to augment the visualization with an overview window, which can enable users to see marked areas outside the detail view. However, overview windows have a number of drawbacks. They take up screen real estate, may not provide enough resolution to ensure that marked areas in large datasets are visible to users, and divide users' attention [51].

To address this issue, Munzner et al. [37] introduce the concept of **guaranteed visibility**, the property that marked areas of the dataset are guaranteed to be visible regardless of dataset size or navigation actions taken by users. Munzner et al. differentiate between three cases to consider when guaranteeing visibility of marked areas:

1. **Off-screen:** A marked area may move off-screen due to user navigation actions or restrictions on available screen real estate.
2. **Sub-pixel:** The dimensions of a marked area may shrink to less than a pixel. This situation is particularly likely to occur when the number of items in a dataset is larger than the number of pixels available to the visualization.

3. **Occlusion:** A marked area can be occluded by other parts of the dataset, such as labels in a two-dimensional layout or other items in a three-dimensional visualization.

Munzner et al. [37] implemented all three types of guaranteed visibility in conjunction with rubber sheet navigation in the previously discussed TreeJuxtaposer system, which serves as the basis for the rubber sheet navigation interfaces discussed in this thesis. Guaranteed visibility has also been implemented in several pan and zoom interfaces, including CityLights [51] and Halo [6], illustrated in Figures 2.5 and 2.6, respectively. Based on the Jazz pan and zoom interface toolkit [9], CityLights indicates the direction of off-screen marked areas via compact indicators integrated along view borders. Halo [6] develops the concept of CityLights further to provide an indication of both direction and distance to off-screen marked areas in the context of small-screen devices. Marked areas are surrounded with rings that are just large enough to reach into the border regions of the visualization. The rings enable users to discern the approximate location of the marks, while using a relatively small proportion of the available screen real estate. In a controlled experiment, Halo was found to improve performance on a navigation task compared to an arrow-based technique that, similarly to CityLights, only indicated direction to off-screen marked areas. This result motivated our use of Halo-like arcs to provide off-screen guaranteed visibility in the pan and zoom navigation interfaces examined in our studies.

2.4 Evaluation of Navigation Techniques

To date, no empirical evaluation of rubber sheet navigation has appeared in the information visualization literature. However, a number of evaluations have compared other focus+context navigation techniques to pan and zoom interfaces. One of the first such evaluations was presented by Schaffer et al. [46], who compared a zooming interface to a fisheye interface for performing a navigation and routing task in a hierarchically clustered network dataset. Results

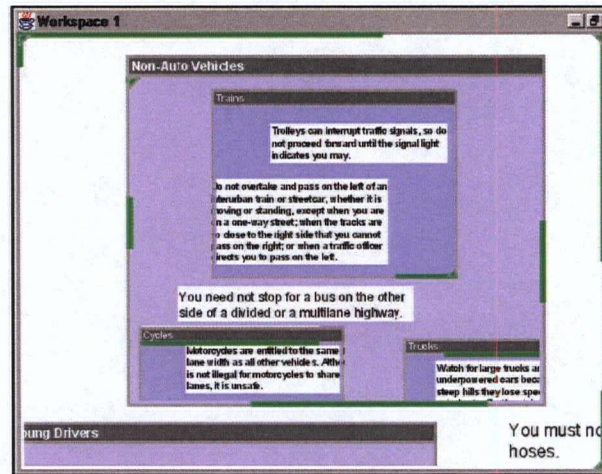


Figure 2.5: Guaranteed visibility in CityLights [51]. Off-screen marked areas are shown by indicators integrated along view borders.



Figure 2.6: Guaranteed visibility in Halo [6]. Direction and distance to off-screen marked areas are indicated by arcs that represent visible portions of rings drawn around the marked areas.

showed that the fisheye interface resulted in significantly lower completion times and was preferred by subjects.

Gutwin and Skopik [22] compared fisheye interfaces using three different distortion models to two panning interfaces with overviews for performing a large steering task. Results indicated that all the fisheye interfaces enabled faster task completion than either of the panning interfaces, although this finding was partially explained by the implementation shortcomings of one of the panning interfaces. Gutwin followed up this study with another investigation [21], which compared fisheye, panning, and two-level zoom interfaces for editing, web navigation, and monitoring tasks. That study found that the fisheye interface was significantly faster for the web navigation task, but the zoom interface performed better on the monitoring task and was strongly preferred by subjects.

Hornbaek and Frokjaer [25] compared panning interfaces with and without overviews to a fisheye interface for reading electronic documents. That study found that subjects read documents faster when using the fisheye interface than when using either of the panning interfaces. However, the panning interface with an overview provided better comprehension and was preferred by subjects. Baudisch, Lee, and Hanna [5] performed a comparative study of three interfaces similar to those examined by Hornbaek and Frokjaer. Both the fisheye interface and the panning interface with an overview provided guaranteed visibility of marked areas, while the panning interface without an overview did not. Although results for both performance and subject preference were highly dependent on the task, the two interfaces with guaranteed visibility were faster than the comparison interface for most tasks and were preferred by all subjects, a finding that motivated our use of guaranteed visibility in the interfaces examined in our experiment. Baudisch et al. also suggested that future studies investigate the potential benefits of combining focus+context navigation with overviews, which corresponds to one of the interfaces examined in our experiment.

The evaluation presented in this thesis attempts to overcome a number of limitations of scope shared by these studies. First, the non-distortion inter-

faces used in these studies provided either panning or zooming capabilities, but not both. Panning and zooming are increasingly used together in both research and commercial interfaces, and some evaluations of interfaces combining these two techniques have appeared in the literature, most notably Hornbaek et al.'s comparison of pan and zoom interfaces with and without overviews [24]. However, combined pan and zoom interfaces have not to date been empirically compared to focus+context interfaces. Second, all the above-mentioned studies used variations on the fisheye interface paradigm. The work presented in this thesis seeks to expand the literature on evaluation of focus+context interfaces to include rubber sheet navigation, which represents one potential alternative to fisheye views. Third, although the fisheye interfaces used in Gutwin and Skopik's study [22] provided context in different ways depending on their distortion models, neither this study nor any of the others discussed in this section examined level of context as a factor. Fourth, none of the studies investigated augmenting focus+context interfaces with overviews to determine whether this would compensate for the drawbacks of distortion, a limitation discussed in more detail in Bodnar's thesis [12]. Finally, apart from the study performed by Schaffer et al. [46], the evaluations discussed above did not rely on tasks or datasets derived from real-world applications, a limitation addressed in our study through the choice of an ecologically valid task and dataset inspired by the needs of phylogenetic biologists.

2.5 Evaluation of Tree Visualizations

The work described in this thesis is related to another branch of the literature on evaluation of information visualizations, namely user studies of tree visualization techniques. In particular, two recent studies have performed controlled experiments involving interfaces for visualizing large tree datasets.

Kobsa [30] compared five tree visualization interfaces, as well as Windows Explorer, which was used as a baseline for comparison. Kobsa's study used a hierarchical tree dataset of more than 5,700 nodes and a variety of tasks related

to both dataset topology and item attributes. Windows Explorer outperformed the comparison interfaces and was also preferred by most subjects, highlighting the difficulty of comparing an interface with which users have experience to those they are encountering for the first time. The comparison was confounded by the fact that some interfaces were missing functionality required to complete some of the tasks.

Plaisant, Grosjean, and Bederson [43] compared their SpaceTree tool, which used a non-distortion-based focus+context interface, to Windows Explorer and a hyperbolic tree browser based on that developed by Lamping, Rao, and Pirolli [31]. The experiment used a large tree dataset of more than 7,000 nodes and a variety of search and topological tasks. The results of the study were mixed, revealing that SpaceTree performed significantly faster for some topological tasks, but not for others, with no significant differences in terms of subject preference.

A common limitation to both these studies is that the interfaces examined used widely different methods of data presentation and interaction, making their results difficult to interpret. The experiments described in this thesis aim to overcome this issue by comparing interfaces that share visual presentation and interaction metaphors and differ only in terms of navigation technique.

Chapter 3

Task and Dataset

In order to lend ecological validity to our experiment, we derived the task and dataset used in it from the domain of phylogenetic biology. Phylogenetic biologists model evolutionary relationships as hierarchical trees in an effort to improve their understanding of how different organisms evolve and co-evolve. The recent flood of molecular data obtained from DNA and protein sequencing has enabled the construction of phylogenetic trees of ever-increasing size. Today, some groups of phylogenetic biologists have constructed trees containing thousands of nodes, and many hope soon to be able to reconstruct the complete Tree of Life, estimated to contain over ten million species [37]. However, a recent survey [15] points out that progress has been hampered by a lack of tools supporting exploration, visual inspection, and structural comparison in such large datasets. This chapter documents the choice of our task and dataset based on the requirements of this domain.

3.1 Task

To gain an understanding of the tasks involved in phylogenetic analysis of large tree datasets using information visualization tools, we conducted interviews with ten phylogenetic biologists from universities in Canada and the United States. We learned that phylogenetic biologists use interactive visualizations of large evolutionary trees to gain a deeper understanding of the relationships between and within groups of organisms. Through the process of topological analysis, these researchers aim to determine how species have evolved and co-evolved, and how characteristics are passed from one species to the next in an evolutionary

lineage.

Based on our discussions, we developed a set of four tasks, described and illustrated below, which were representative of the tree-topological tasks performed by phylogenetic biologists, but did not require specialized knowledge of evolutionary trees. We then validated these tasks with several of the biologists we had previously interviewed to ensure the tasks' ecological validity.

Each of the tasks was composed of several low-level tasks such as find, identify, and compare, as described in the visual task taxonomy of Wehrend and Lewis [49]. In the illustrations below, a colored node represents a species, whereas a colored subtree represents a related group of species.

3.1.1 Task 1: Determining the lowest common ancestor

In a phylogenetic tree, the **lowest common ancestor** of two nodes is an organism that is an ancestor of both the species in question, and that has the greatest depth in the tree, as illustrated in Figure 3.1. Determining the lowest common ancestor is an important task in **phylogenetic taxonomy**, a branch of phylogenetic biology concerned with classifying species based on phylogenetic data.




Figure 3.1: Task 1: Determining the lowest common ancestor. In this case, node A is the lowest common ancestor of nodes B and C.

3.1.2 Task 2: Comparing the topological distances between nodes

Topological distance in a tree is the number of hops between two nodes, and is not the same as geometric distance, which may change with navigation, as illustrated in Figure 3.2. In a phylogenetic tree, the topological distance between two nodes is indicative of the number of evolutionary steps between the species they represent. Measuring and comparing topological distances is one of the primary tasks for which phylogenetic biologists require visualizations of evolutionary trees.

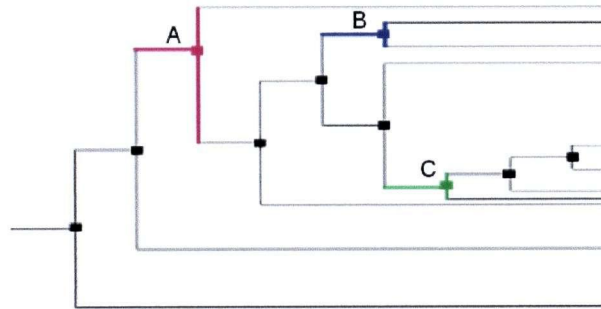


Figure 3.2: Task 2: Comparing the topological distances between nodes. In this case, node A is 2 topological hops from node B and 3 topological hops from node C, making node B topologically closer.

3.1.3 Task 3: Determining whether two subtrees are adjacent

In a tree, two subtrees are adjacent if no other node is between them, as illustrated in Figure 3.3. In phylogenetic biology, this task represents determining whether the groups of species represented by the subtrees are **sister groups**, or groups of organisms who are most closely related to one another in terms of their evolutionary history (for instance, great apes and monkeys).

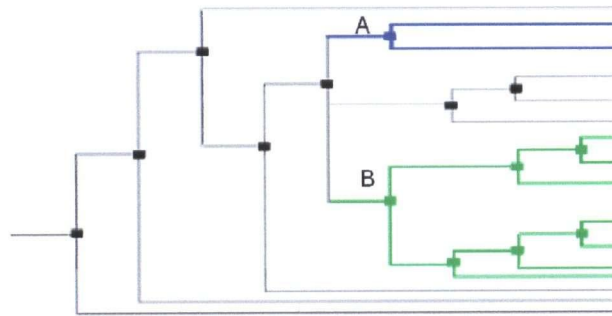


Figure 3.3: Task 3: Determining whether two subtrees are adjacent. In this case, the subtrees labeled A and B are not adjacent.

3.1.4 Task 4: Determining whether a subtree contains unmarked nodes

In a phylogenetic tree, marked nodes may indicate the presence of a unique feature or **character**. The presence of uncolored nodes or subtrees in a marked subtree, shown in Figure 3.4, may therefore indicate a **character reversal**, an event causing the loss of a character formerly present in an evolutionary line (for example, the loss of a tail in great apes and humans).

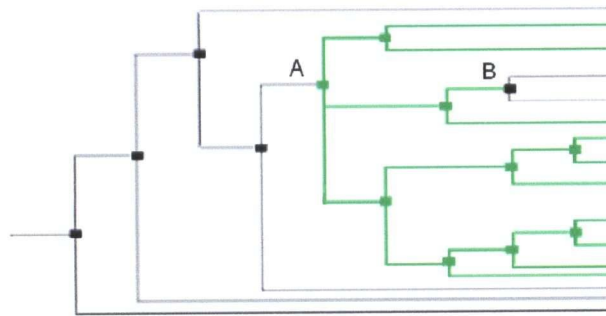


Figure 3.4: Task 4: Determining whether a subtree contains unmarked nodes. In this case, the subtree labeled A contains an unmarked node, B.

After the development of the set of tasks described above, we decided to focus our study on a single task in order to limit the effect of task as a factor. Task 2 was selected for further investigation due to its relative complexity, high importance to phylogenetic analysis, and the fact that it would require subjects to perform multiple navigation actions along well-defined paths, thus reducing performance variability.

Task instances were assessed in pilot studies to ensure that they were isomorphic in difficulty. In particular, topological distances between nodes always fell in a range of 7 to 10, and could not be determined without interacting with the interface for any of the task instances. Also, colored nodes were not located in close proximity to each other in order to ensure that at least one interaction had to be performed to determine each topological distance.

3.2 Dataset

The dataset used for initial piloting for our study was the *animaliaA* dataset from the 2003 Infovis Contest [42], a phylogenetic tree of approximately 190,000 nodes representing a hypothesis about the evolution of organisms in the kingdom *Animalia*. Initial pilot results suggested that this dataset was not an optimal choice for our experiment. Its topology was not sufficiently deep to require subjects to perform a large amount of navigation, while its size necessitated start times of up to 45 seconds for our visualization tools. For this reason, subsequent piloting and the formal experiment used the *phylogenyMatchesTaxonomy* dataset, a binary tree consisting of 5,918 nodes, which also represents evolutionary relationships between species the kingdom *Animalia*. This dataset was used courtesy of David Hillis of the University of Texas, and is available from the Olduvai project website [39]. This dataset allowed for complex topological comparisons requiring a significant amount of navigation while reducing the start times for our tools to under 5 seconds.

Although we had originally assumed that node labels were important to the way biologists interact with phylogenetic tree visualizations, our discussions

with biologists revealed that their typical use of evolutionary trees involved very little label reading. We therefore removed node labels from the dataset for the purpose of our studies. Using no labels enabled us to avoid unnecessary node occlusion and potential confounding of experimental results by subjects' prior knowledge of evolutionary relationships between species.

Chapter 4

Study 1

The goal of Study 1 was to examine the effect of level of context on performance in interfaces with pan and zoom and rubber sheet navigation. The study involved four different interfaces, representing all combinations of the two navigation techniques with and without an overview, as illustrated in Table 4.1. Subjects used these interfaces with varying levels of context to solve a topological task in a large tree dataset. The level of context that resulted in the best performance for each interface was then used to compare the performance of the four interfaces in Study 2, described in Chapter 5. This chapter describes the study, presents the results related to the effects of navigation technique, and discusses their implications and the way they affected Study 2. The results related to the effects of presence or absence of overview and their implications are discussed in detail in Bodnar's thesis [12].

Navigation Overview	Rubber Sheet	Pan and Zoom
Not present	RSN-NoOV	PZN-NoOV
Present	RSN+OV	PZN+OV

Table 4.1: Interfaces representing all combinations of the navigation and presence of overview factors. All four interfaces were used in both Study 1 and Study 2.

4.1 Hypotheses

Our primary hypothesis for this study was that performance in each interface would vary with the level of context according to a U-shaped curve, with very low and very high context levels resulting in poor performance compared to values between these two extremes. We expected that low levels of context would not provide sufficient resolution to enable users to obtain contextual cues, while high levels of context would constrain the amount of screen real estate available to show features of interest to users and therefore adversely impact navigation. Similarly, we expected that, in interfaces with overviews, small overview sizes would provide insufficient resolution, while large overview sizes would detract from the navigation within the detail views. The resulting hypotheses are presented below.

H1: For both pan and zoom and rubber sheet navigation, medium levels of context within the detail view will perform better than either high or low levels of context.

H2: In interfaces with overviews, medium sized overviews will perform better than either small or large overviews.

The values for small, medium, and large levels of context and overview sizes for each interface were chosen based on results of pilot studies, and are listed in Section 4.6. Due to the variation of levels of context and overview sizes within interfaces in this study, no specific hypotheses were developed for the effects of navigation technique or presence of overview on performance, factors that were investigated in more detail in Study 2.

4.2 Interfaces

The interfaces examined in this study are illustrated in Figures 4.1 to 4.4. In order to provide consistent visual representation, drawing performance, and




Figure 4.1: RSN-NoOV interface used in Study 1. A zoom action has stretched a region to fill the top focus region. Nodes outside this region are compressed in the periphery, and marked nodes remain visually salient.

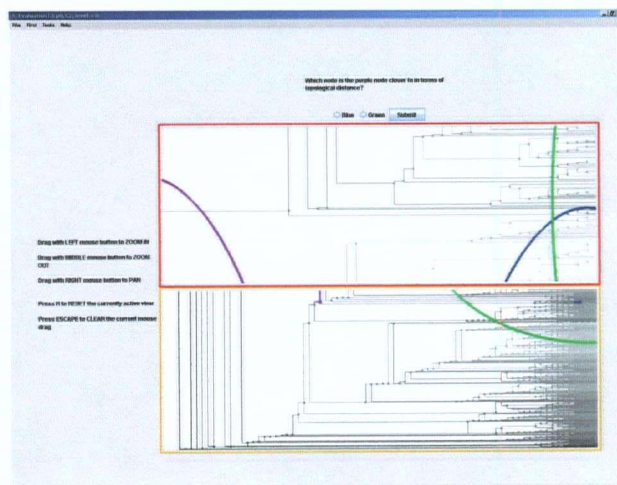


Figure 4.2: PZN-NoOV interface used in Study 1. A zoom action has filled the extent of the top view. Arcs inspired by Halo [6] indicate direction and distance to off-screen marked nodes.

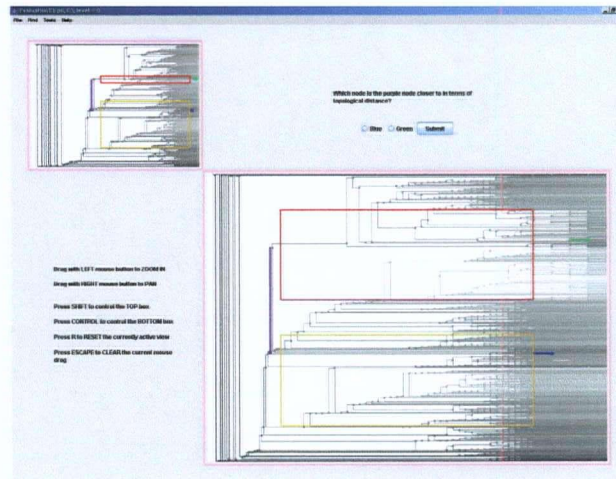


Figure 4.3: RSN+OV interface used in Study 1. A zoom action has stretched the region shown by the field-of-view box in the overview to fill the top focus region of the detail view.

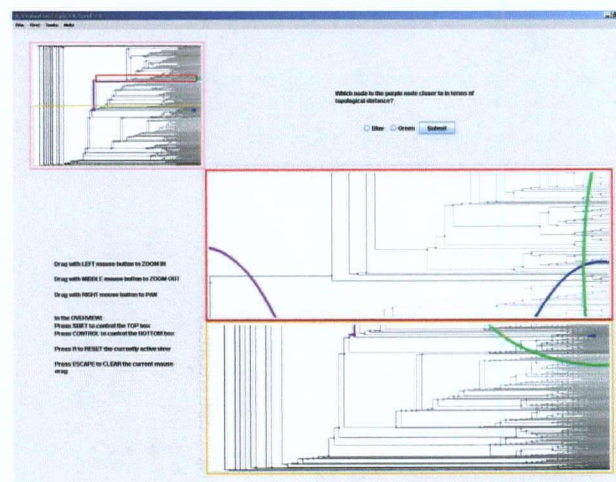


Figure 4.4: PZN+OV interface used in Study 1. A zoom action has filled the extent of the top detail view with the region shown by the field-of-view box in the overview.

interaction model, all interfaces were built on the PRISAD software infrastructure [47], based on the TreeJuxtaposer scalable tree visualization application [37]. The use of the PRISAD infrastructure also helped reduce implementation time for the interfaces. While TreeJuxtaposer was initially developed as a Focus+Context visualization tool using rubber sheet navigation, the inherent similarities between rubber sheet and pan and zoom navigation allowed us to extend its behaviour to support conventional pan and zoom interaction, as well varying levels of context, overviews, and multiple foci. This section discusses the implementation of each of these interface components and then examines the interfaces themselves in detail.

4.2.1 Navigation

The original TreeJuxtaposer application [37] used rubber sheet-style expansions and contractions of arbitrary rectilinear regions for navigation, and included advanced features such as linked navigation between multiple trees. Navigation in TreeJuxtaposer enabled users to select rectangular regions using mouse drags, and resize their selection box to arbitrary size. We replaced this style of navigation with a unified set of navigation actions implemented across all interfaces. All interaction occurred through mouse drags, and in our subsequent analysis, a discrete navigation action refers to a single mouse drag. All transitions were smoothly animated across 20 frames to ensure fluid interaction with the interfaces. In each interface, navigation was controlled using a two button mouse with a scroll wheel, with zoom in mapped to the left mouse button, panning mapped to the right mouse button, and zoom out mapped to the scroll wheel. Each interface also supported a reset function, which was mapped to the R key. Similar to rubber sheet navigation as implemented in TreeJuxtaposer, our implementation of rubber sheet navigation allowed users to select a rectangular region using mouse drags. However, in our implementation, selection boxes were always expanded into focus areas of fixed size and aspect ratio. The dimensions of the focus areas were fixed in order to ensure that the level of context remained

constant regardless of user interactions.

4.2.2 Overviews and Foci

Overviews with movable field of view boxes were present in two of the interfaces. For consistency between interfaces, the view dimensions in each interface were chosen to equalize the total screen real estate across them, with each interface always providing a total of 600,000 pixels of information. Based on the guidelines developed by Ahlberg and Shneiderman [2], we ensured that all navigation actions were tightly coupled between the overview and detail view.

For this study, two foci were implemented in both pan and zoom and rubber sheet navigation interfaces to allow users to simultaneously view and interact with multiple non-adjacent regions of the dataset. In the rubber sheet navigation interfaces, users could select one of two focus regions as the target for rectilinear zooming actions, allowing them to explore two non-adjacent regions of the dataset at different levels of compression. In the pan and zoom interfaces, users could navigate in two separate views, allowing them to explore two different regions of the dataset at different scales. The decision to implement multiple foci was motivated by the scenario where subjects would be required to navigate to features located between the two foci in the process of completing our task. For these instances, we expected rubber sheet navigation to benefit from the context region between the two foci, which would be either not visible or only visible in the overview in interfaces using pan and zoom navigation, as illustrated in Figure 4.5.

4.2.3 Guaranteed Visibility and Levels of Context

Guaranteed visibility of marked areas was provided in both detail views and overviews for both pan and zoom and rubber sheet navigation interfaces, and addressed the three types of guaranteed visibility discussed in Section 2.3. In particular, sub-pixel guaranteed visibility in all interfaces was provided by the underlying PRISAD visualization framework [47], which ensured that items of

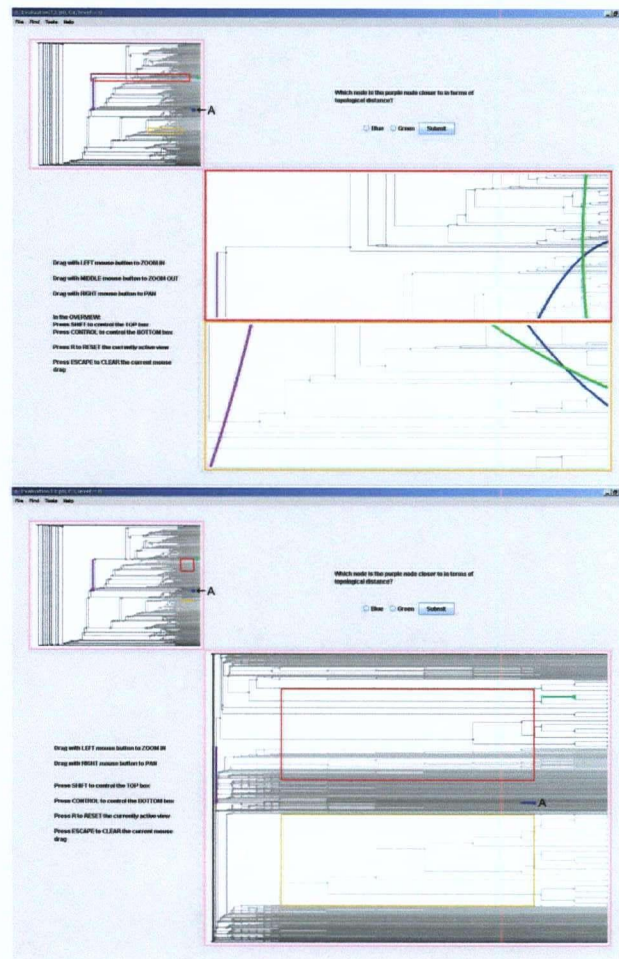


Figure 4.5: Motivating scenario for use of two focus areas in Study 1. Marked node A, located between the two focus areas in each interface, is visible in both views in the RSN+OV interface (above), but only in the overview the PZN+OV interface (below). Rubber sheet navigation was therefore expected to perform better than pan and zoom navigation in this scenario.

interest in all views were visibly marked even when they were compressed to sub-pixel size. Occlusion of marked areas by other parts of the dataset was avoided by using a 2D rather than a 3D spatial layout and removing labels from our dataset.

Off-screen guaranteed visibility was implemented in the experimental interfaces in different ways depending on navigation technique. In rubber sheet interfaces, navigation was constrained so that items outside the focus areas were compressed in context areas along the periphery of the view. In pan and zoom interfaces, direction to and distance from off-screen marked areas were encoded using opaque elliptical arcs similar to those implemented in Baudisch and Rosenholtz's Halo [6]. As in rubber sheet navigation interfaces, these arcs appeared in peripheral context areas along the edges of a view. However, these context areas were not explicitly visually delimited, their resolution did not change, and their shape was oval rather than rectangular. Although the degree of contextual information provided by context areas in both navigation techniques varied with user interaction, we used the total extent of these areas as an approximation for the amount of context within each interface.

In addition to peripheral context areas, contextual information was also provided by overviews in those interfaces that contained them. For the purpose of varying the level of context in this study, we therefore distinguished between two possible levels of context in each interface, illustrated in Figures 4.6 and 4.7:

1. **Level of navigational context:** Fraction of size of navigation-specific context areas C to the total size of focus and context areas in the detail view $F+C$.
2. **Level of overview context:** Fraction of size of the overview O to total size of all views $O+F+C$ (0 for interfaces without an overview).

As previously mentioned, each interface always provided a total of 600,000 pixels of information in all views. In interfaces without overviews, this amount




Figure 4.6: Calculation of levels of context in Study 1 RSN interfaces. Level of navigational context is the fraction of the size of the peripheral context areas C to the total size of the detail view $F+C$. Level of overview context is the fraction of the size of the overview O to the total size of all views $O+F+C$.




Figure 4.7: Calculation of levels of context in Study 1 PZN interfaces. The dotted line indicates the boundary between focus and context regions, which is not visually demarcated in the interfaces. Levels of navigational and overview context are as in Figure 4.6.

was equal to $F+C$, while in interfaces with overviews, it was equal to $O+F+C$. Thus, as shown in Figures 4.1 to 4.4, interfaces without overviews provided larger detail views than their counterparts with overviews, but the level of navigational context was kept constant for each navigation technique, regardless of the presence of an overview.

4.2.4 RSN-NoOV Interface

As illustrated in Figure 4.1, this interface had no overview and allowed users to navigate the dataset using the metaphor of expanding and compressing a rubber sheet with its borders nailed down. Unlike in conventional pan and zoom interfaces, navigation actions did not push context regions off-screen, but compressed them in the periphery of the view, where they remained visually salient. Focus regions were demarcated by colored boxes, which were always located in the center of the view. Users could select a rectangular area of interest for zooming in by dragging out a box with the left mouse button. The contents of the selected area then expanded to fill one of the two focus regions in a smooth transition. By default, the system filled the focus whose centroid was closer to the centroid of the selected area of interest. Users could also specify which focus region was to be filled by pressing a modifier key. The Shift key was used to specify the top focus and the Ctrl key for the bottom focus. These keys were chosen because of their position one above the other on the keyboard, which helped users associate them with the respective foci. An action analogous to panning was accomplished via horizontal and vertical drag motions with the right mouse button, allowing users to fine-tune focus region selections. Users could zoom out by dragging out a rectilinear region larger than the focus region, the contents of which were then compressed to fill the focus region.

4.2.5 PZN-NoOV Interface

As illustrated in Figure 4.2, this interface had no overview and allowed users to navigate using conventional pan and zoom interactions. Just as with the

RSN-NoOV interface, users selected a rectangular area of interest for zooming in with a left mouse drag, resulting in an animated transition that completely filled the focus with the selected area. Similarly to the RSN-NoOV interface, by default the system filled the focus whose centroid was closer to the centroid of the selected area of interest, but users could also specify which focus was to be filled by pressing the Shift or Ctrl modifier key. Users could fine-tune the focus selection by panning with horizontal and vertical right-mouse drags, and gradually zoom out with vertical middle-mouse drags. When marked regions moved off-screen due to navigation actions, colored Halo-like arcs, representing the visible parts of elliptical rings centered on the marked regions, appeared in context areas at the periphery of the interface. The arcs indicated direction and distance to marked regions, and disappeared once the marked regions became visible on-screen. Unlike in the RSN-NoOV interface, the peripheral context areas in this interface were oval rather than rectangular in shape, were not visually delimited, and did not change in terms of resolution.

4.2.6 RSN+OV Interface

Figure 4.3 illustrates that this interface used the same navigation controls as the RSN-NoOV interface. It also had an overview showing two colored field of view boxes corresponding to the extent of each focus in the detail view. The size and location of the field of view boxes were updated dynamically as navigation took place in the detail view. Users could perform the rubber sheet navigation equivalents of panning and zooming, as implemented in the RSN-NoOV interface, directly in the overview by dragging the field of view boxes, which then updated the appropriate foci in the detail view.

4.2.7 PZN+OV Interface

Figure 4.4 illustrates that this interface had the same navigation controls as the PZN-NoOV interface, as well as an overview. Just as with the RSN+OV interface, the field of view boxes in the overview dynamically reflected navigation

in the detail view and could be manipulated directly to control the foci in the detail view.

4.3 Task and Dataset

The task used in the experiment was a tree topological task that required subjects to compare topological distances between colored nodes in a large tree dataset and determine which of the distances was smaller. Task and dataset are both described in detail in Chapter 3.

4.4 Apparatus

The study was conducted on two systems running Windows XP with Pentium 4 processors, 2.0 GB RAM, Nvidia GeForce2 video cards, and 19 inch monitors configured at a resolution of 1280x1024 pixels. The experimental software, including the interfaces, was fully automated and was coded in Java 1.4.2 and OpenGL using the Swing and GL4Java libraries.

4.5 Participants

Forty subjects (15 female) between 18 and 39 years of age successfully completed the study and were each compensated \$10 for their participation. All subjects were right-handed, had normal or corrected to normal vision, and were not color blind. They were recruited through advertisements posted throughout the university campus and through an online participant scheduling system.

Originally, 45 subjects participated in the experiment. One of the subjects was unable to follow the training instructions successfully, while another was not successful in learning to use the mouse scroll wheel, required to perform zooming out in the pan and zoom interfaces. Three others followed the instructions but committed three or more errors (an error rate greater than 10%). These five

Order	Block 1	Block 2	Block 3	Block 4	Block 5
1	3	1	5	2	4
2	4	3	1	5	2
3	2	4	3	1	5
4	5	2	4	3	1
5	1	5	2	4	3

Table 4.2: Latin square used to counterbalance the order of presentation for level of context.

subjects were treated as outliers for the purpose of data analysis, leaving a total of 40 data points.

4.6 Design

The evaluation used a 4 (interface, between subjects) by 5 (level of context, within subjects) design, where each level of context corresponded to a block containing 5 trials. Navigation technique and presence of overview were not examined as separate factors in this study due to the interface-specific variation of levels of context and overview sizes. Interface was chosen to be a between-subjects factor due to the possibility of transfer effects of navigation technique in a within-subjects design, as well as the time required to train subjects on navigation with each technique. Subjects were randomly assigned to each of the four interfaces. Level of context was chosen to be a within-subjects factor to allow comparison between the different levels of context for each interface. To minimize ordering effects, we counterbalanced the order of presentation for level of context using a Latin square, as shown in Table 4.2.

In order to determine the range of levels of context to be used in the study, we piloted each experimental interface with 5 to 9 subjects using the experimental task and dataset. The levels of context used in piloting varied from 10% to 95% for navigational context and from 5% to 50% for overview context. Based on

Interface	Level of Navigational Context	Level of Overview Context
RSN-NoOV	40, 50, 60 , 70, 80	0
PZN-NoOV	30, 40 , 50, 60, 70	0
RSN+OV	60	5, 10 , 15, 20, 25
PZN+OV	50	5, 10 , 15, 20, 25

Table 4.3: Levels of context used for each interface, in percent. Boldface denotes values expected to result in optimal performance for each interface based on pilot results. The level of context in the detail view of interfaces with an overview was set to the middle level of context used for their counterparts without an overview.

pilot results, we selected a range of five levels of context for each interface that performed best in terms of completion times. The levels of context used in each interface are listed in Table 4.3, with the levels of context expected to perform best for each interface based on pilot results shown in bold.

4.7 Procedure

The experiment was designed to fit into a single 60 minute session. The experimenter first instructed subjects on the use of the different navigation techniques afforded by the interface to which they had been randomly assigned. Subjects were then shown the experimental task and instructed that they were to take as much time as necessary to solve it correctly while maximizing their efficiency. The experimenter then trained subjects on the use of long, thin horizontal selection areas to complete the task. This strategy had been found to improve task completion time in all interfaces in piloting. Since many of the paths between colored nodes were horizontal, this strategy enabled subjects to bring them rapidly into focus. Subjects were then given a training block of 5 trials. During these trials, subjects performed the task on their own, and the experi-

menter reminded subjects of the training strategy as needed. For the training block, the middle level of context value of those shown in Table 4.3 was used for each interface. At the end of the training session, subjects were given a one minute break. After the break, the experimenter exited the room where the experiment was conducted, and subjects proceeded with the experiment. The complete training protocol for this study can be found in Appendix A.

In the experiment, subjects were presented with 5 blocks, each containing 5 trials, for a total of 25 trials. All subjects were presented with an identical set of questions, with a predetermined grouping of questions into blocks. The order of blocks was determined using the Latin square shown in Table 4.2. The blocks of questions were verified to be isomorphic in difficulty in piloting. Subjects were given a one minute break between each block of questions.

At the end of the experiment, subjects completed questionnaires, which can be found in Appendix C. The questionnaires were used to collect information about the subjects' demographic background and computer usage. They also included ratings for ease of use, ease of navigation, and interface-specific features on 5-point Likert scales. Space was also provided for subjects to comment on their experiences with the interfaces and provide suggestions for improvement. Short informal interviews were conducted with some of the subjects based on their questionnaire responses.

4.8 Measures

Our performance measures were based on logged data and included task completion times and errors. Additionally, self-reported measures were collected through the post-experiment questionnaire, as described in the previous section. The study was designed to minimize errors, with task completion times used as the primary measure of performance.

4.9 Results

This section presents the experimental results, reporting in detail on the results related to effects of learning and navigation technique. The results for presence of overview are summarized here, and details on these can be found in Bodnar's master's thesis [12].

Prior to analysis, outlier data lying more than 3 standard deviations from the means of each experimental cell were removed from the analysis. The Greenhouse-Geisser adjustment was used for non-spherical data, and the Bonferroni adjustment for post-hoc comparisons. Along with statistical significance, we report partial eta-squared (η^2), a measure of effect size, which is often more informative than statistical significance in applied human-computer interaction research [32]. To interpret this value, .01 is a small effect size, .06 is medium, and .14 is large [17].

The overall results for mean completion times per trial are illustrated in Figure 4.8. A one-way ANOVA was run to understand the effect of interface, level of context, and block on completion time. As expected, performance improved as subjects progressed through the experiment, with a significant main effect of block ($F(4,144) = 12.309$, $p < .001$, $\eta^2 = .255$). There was also a significant main effect of interface on completion time ($F(3,36) = 2.924$, $p < .05$, $\eta^2 = .196$), but post-hoc comparisons revealed no significant pairwise differences between interfaces. No significant interaction effect between block and interface was present.

One-way ANOVAs were performed to investigate the effect of level of context on completion times for each interface. As shown in Table 4.4, there was no significant effect of level of context for any of the interfaces, despite high effect sizes for three of the four interfaces. The best mean performance with the RSN-NoOV and PZN-NoOV interfaces occurred with context levels of 50% and 40%, respectively, while the best mean performance with both the RSN+OV and PZN+OV interfaces occurred with an overview size of 15%. Mean completion times for each interface are shown in Figures 4.9 through 4.12. Separate figures




Figure 4.8: Mean per-trial completion times by interface for blocks 1-5, in seconds ($N=40$).

Interface	df	F	Sig.	Partial η^2
RSN-NoOV	1.826, 16.43	1.272	.303	.124
PZN-NoOV	4,36	.456	.767	.048
RSN+OV	4,36	1.380	.260	.133
PZN+OV	4,36	1.480	.229	.141

Table 4.4: Results of four one-way ANOVAs investigating the effect of level of context on completion time, by interface ($N=10$). Results for the RSN-NoOV interface are adjusted for sphericity.

are used to report these results due to the differences between ranges of levels of context examined for each interface.

Counter to our hypothesis **H1**, there was no significant difference between performance with middle and extreme levels of context for either RSN or PZN. Figures 4.9 through 4.10 illustrate that the expected U-shaped performance trend for this hypothesis was not present in data for either navigation technique. Similarly, Figures 4.11 and 4.12 show that, counter to our hypothesis **H2**, there was no significant difference between performance with middle and extreme overview sizes for either RSN or PZN interfaces with overviews. The implications of this finding are discussed in detail in Bodnar's thesis [12].

On average, subjects committed 0.75 errors over the course of the experiment, for a mean error rate of 3.0%. There were no significant main or interaction effects of interface or level of context on error rate.

4.10 Summary of Results

We summarize our results according to the hypotheses stated in Section 4.1:

- R1:** For both pan and zoom and rubber sheet navigation, medium levels of context did not perform better than low or high levels.
- R2:** For both pan and zoom and rubber sheet navigation, in interfaces with overviews, medium overview sizes did not perform better than low or high ones.

The best performance in terms of mean completion times occurred with a context level of 50% in the RSN-NoOV interface, a context level of 40% in the PZN-NoOV interface, and overview sizes of 15% in both the RSN+OV and PZN+OV interfaces.




Figure 4.9: Mean per-trial completion times (dark line) and quadratic trend line (light line) for the RSN-NoOV interface, in seconds ($N=10$).




Figure 4.10: Mean per-trial completion times (dark line) and quadratic trend line (light line) for the PZN-NoOV interface, in seconds ($N=10$).




Figure 4.11: Mean per-trial completion times (dark line) and quadratic trend line (light line) for the RSN+OV interface, in seconds ($N=10$).




Figure 4.12: Mean per-trial completion times (dark line) and quadratic trend line (light line) for the PZN+OV interface, in seconds ($N=10$).

4.11 Strategies

In order to gain a better understanding of why level of context did not significantly affect completion times, we examined logged navigation and reset action data. Although the logging format used in this study was not conducive to a statistical analysis of these data, a manual analysis of navigation and reset patterns for individual subjects showed that, for each interface, most subjects consistently used a small number of well-defined strategies throughout the experiment. Some of these strategies were related to navigation patterns, while others involved different ways of using overviews or compensating for the absence of overviews by means of using the reset function. The strategies related to navigation are detailed in Table 4.5, while those related to overview and reset usage are discussed in Bodnar's thesis [12].

Based on the manual analysis of log data, we grouped subjects who used the RSN-NoOV, PZN-NoOV, and RSN+OV interfaces according to their primary navigation strategy. One-way ANOVAs were then performed for each interface to understand the effect of navigation strategy on completion time. For both of the RSN interfaces, the "zoom, then pan" strategy resulted in faster mean completion times than the "zoom only" strategy (67.9 seconds vs. 75.7 seconds for RSN-NoOV; 63.8 seconds vs. 75.6 seconds for RSN+OV). Although this difference was not statistically significant for RSN-NoOV ($F(2,47) = 1.616$, $p > .2$, $\eta^2 = .033$), and only borderline significant for the RSN+OV ($F(2,47) = 3.008$, $p < .09$, $\eta^2 = .059$), these results suggested the importance of the use of the rubber sheet equivalent of panning for effective navigation, and led us to emphasize the use of this operation in the training protocol of Study 2.

For the PZN-NoOV interface, the three strategies were significantly different in terms of completion time ($F(2,47) = 9.226$, $p < .001$, $\eta^2 = .282$). Post-hoc pairwise comparisons showed that the "zoom in, then pan" strategy was significantly slower ($p < .01$) than both the "zoom in, then zoom out" strategy and the "overview" strategy (mean completion times: 104.2 seconds, 80.0 seconds, and 75.4 seconds, respectively). This finding motivated us to provide a detailed

Interface	Strategy	Description
RSN-NoOV	Zoom only ($N=6$)	Use RSN zoom only.
	Zoom, then pan ($N=4$)	Use RSN zoom and pan.
PZN-NoOV	Zoom in, then pan ($N=4$)	Zoom into neighbourhood of a marked node, then follow path by panning.
	Zoom in, then zoom out ($N=4$)	Zoom into neighbourhood of a marked node, then zoom out to reveal path.
	Overview ($N=2$)	Zoom in with one focus, then use the other as overview.
RSN+OV	Zoom only ($N=6$)	Use RSN zoom only.
	Zoom, then pan ($N=4$)	Use RSN zoom and pan.
PZN+OV	None	

Table 4.5: Strategies related to navigation developed by subjects in Study 1.

training strategy for the PZN-NoOV interface in Study 2.

4.12 Discussion

Results of this study show that level of context did not significantly affect performance as measured by completion times with either pan and zoom or rubber sheet navigation. Also, the expected U-shape of the completion time data, indicating superior performance on medium rather than high or low levels of context, was not present. This finding seems surprising since, for both navigation techniques, we examined a range of levels of context spanning a significant proportion of the available screen real estate. It is possible that our study design did not provide sufficient power to detect differences within the ranges for each interface. This explanation would account for the medium to high effect sizes of level of context in all interfaces, as shown in Table 4.4.

Another possibility is that the ranges we examined were still not sufficiently large to significantly impact the way subjects navigated. This possibility is supported by qualitative feedback from questionnaires and follow-up interviews with the subjects. Only one subject commented on the difference between small and large levels of context, and a number of subjects stated that they were not even aware of the changes in level of context during the experiment. We therefore hypothesize that, for both pan and zoom and rubber sheet navigation, the range of levels of context conducive to effective navigation is sufficiently broad to encompass the range examined in our experiment, and that differences between levels of context within this range do not significantly impact performance. Thus, for each navigation technique, there may not exist an optimal level of context value for performing a given task, but rather upper and lower bounds beyond which performance deteriorates, similar to the bounds that have been proposed for overview sizes [41]. It is also possible that the effect of level of context on navigation performance can be represented by a discontinuous value function. In particular, a small amount of context may improve performance compared to zero context, with further increases providing no added benefits.

For three of the four interfaces examined in this experiment, navigation strategies used by the subjects had a noticeable effect on performance. In the RSN interfaces, the two strategies differed only in whether subjects chose to use the rubber sheet equivalent of panning or to rely solely on rubber sheet zooming actions. This difference likely resulted from the fact that the navigation strategy used in training (the use of long, thin horizontal selection areas) only involved zooming, and, although subjects were shown how to pan, the use of panning was left to their discretion.

Perhaps more surprising was the development of three distinct navigation strategies with the PZN-NoOV interface. This finding can be explained by noting that this interface contained two focus areas of the same size and resolution, providing subjects with the options of using one view exclusively, alternating between the focus areas depending on the location of the marked nodes, or using one focus as a detail view and another as an overview. Furthermore, the distinc-

tion between the “zoom in, then pan” and “zoom in, then zoom out” strategies can be attributed to the fact that zoom out; like pan, was not specifically covered by the training strategy used in the study. The significantly slower performance of subjects using the “zoom in, then pan” points to the inefficiency of panning at a high magnification level, consistent with results of previous studies [46]. The lack of significant difference in performance between the “overview” and “zoom in, then zoom out” strategies suggests that gradual zooming out may offer an adequate substitute for the lack of an overview in navigation tasks such as the one used in our study, where users must follow a specific navigation path.

We had initially believed that the use of two focus areas in this study would enable users to explore multiple areas of the same dataset simultaneously, and hence provide performance benefits. However, observation data indicated that this feature increased interface complexity, forcing subjects to invest more cognitive effort into coordinating navigation between the multiple focus regions than they did into completing the task. Additionally, the presence of two focus areas was partly responsible for the development of interface-specific strategies, in particular the “overview” strategy in the PZN-NoOV interface, that represented a confounding factor. Finally, the Halo-like arcs as implemented for this study caused some subjects difficulties by occluding regions of interest. These issues were addressed by merging the focus areas in each interface and making the Halo-like arcs translucent for the purposes of Study 2.

Three other issues with this study were noted and addressed in Study 2. First, qualitative data showed that many subjects found it difficult to remember intermediate task results, such as the number of nodes on the first path explored for a given task instance. Second, the absence of the experimenter in the experiment room after the training period hampered our analysis of interface-specific strategies. Subjects in Study 2 were therefore given a pen and paper to record intermediate task results and observed throughout the experiment. Finally, navigation and reset action data from this study were not logged in a format conducive to statistical analysis. This capability was implemented in Study 2, allowing these data to be analyzed as additional dependent variables.

Chapter 5

Study 2

The goal of this study was to evaluate the performance of pan and zoom and rubber sheet navigation techniques with and without an overview. Subjects used interfaces based on those used in Study 1, each providing the best performing level of context as discussed in Chapter 4. This chapter describes the study and presents the aspects of its results related to the effects of navigation technique.

5.1 Hypotheses

Our hypotheses for this study were motivated by findings reported in the literature and the results of Study 1. First, we expected rubber sheet navigation to perform better than pan and zoom navigation because, as discussed in Chapter 2, focus+context approaches have been shown to perform better than pan and zoom interfaces for a variety of navigation tasks. Second, we did not expect an overview to significantly improve the performance of rubber sheet navigation, because focus+context approaches by design attempt to provide the same contextual information as an overview, but in an integrated way. Finally, we expected that an overview would significantly improve the performance of pan and zoom navigation, because most previous studies have shown that overviews decrease navigation time and help the user maintain orientation within a pan and zoom interface. Our hypotheses were therefore as follows:

H3: Rubber sheet navigation will perform better than pan and zoom navigation, independently of the presence or absence of an overview.

H4: For rubber sheet navigation, the presence of an overview will not result in

better performance.

H5: For pan and zoom navigation, the presence of an overview will result in better performance.

5.2 Interfaces

The interfaces examined in this study are illustrated in Figures 5.1 to 5.4. Each of the interfaces was based on an interface used in Study 1, as described in detail in Section 4.2. Based on the results of Study 1, multiple foci were removed from all interfaces, and Halo-like arcs in the PZN interfaces were made translucent, such that they were still visually salient but did not fully occlude areas of the dataset. The design rationale for these changes is discussed in Section 4.12.

The values for level of context and overview size in each interface were derived from the values that provided the best performance in Study 1, as outlined in Section 4.9. In particular, in RSN interfaces, the fraction of the view occupied by the peripheral context area was set to 50%, while in PZN interfaces, the fraction of the view where Halo-like arcs could appear was set to 40%. In interfaces with overviews, overviews and detail views comprised 15% and 85% of the total number of pixels available, respectively. The calculation of levels of context for this study was similar to that in Study 1, but reflected the use of a single focus in each interface, as illustrated in Figures 5.5 and 5.6. As in Study 1, each interface always provided a total of 600,000 pixels of information in all views.

5.3 Task and Dataset

This experiment utilized the same tree topological task and dataset as Study 1. Both the task and the dataset are described in detail in Chapter 3.

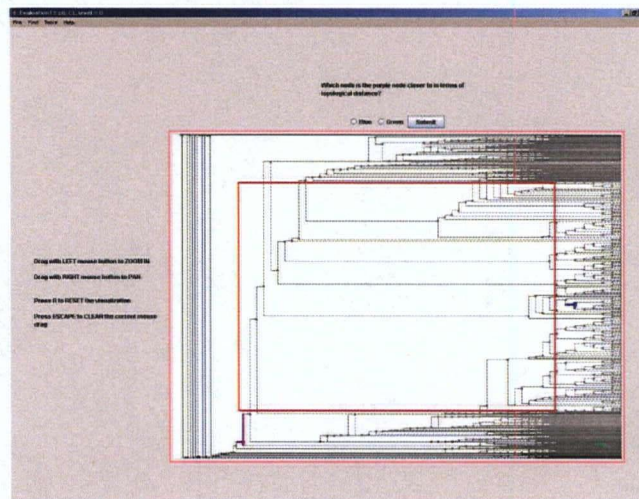


Figure 5.1: RSN-NoOV interface used in Study 2. This interface provided a single focus rather than two foci as in Study 1 (see Figure 4.1). Level of context was set to 50%.

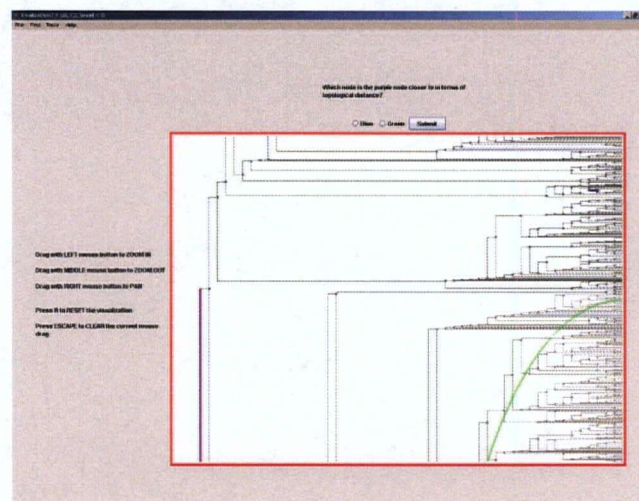


Figure 5.2: PZN-NoOV interface used in Study 2. This interface consisted of a single focus rather than two foci as in Study 1 (see Figure 4.2), and provided translucent rather than opaque arcs. Level of context was set to 40%.

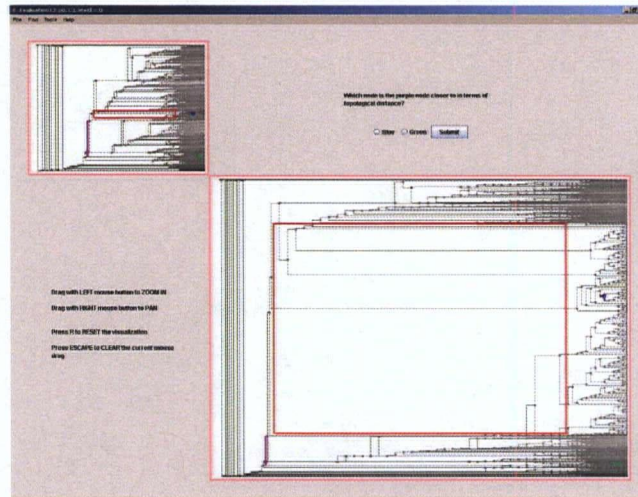


Figure 5.3: RSN+OV interface used in Study 2. This interface provided a single focus rather than two foci as in Study 1 (see Figure 4.3). The overview comprised 15% of total available pixels.

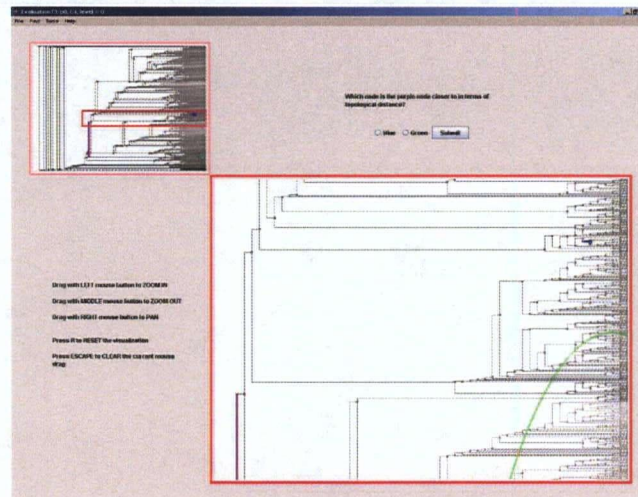


Figure 5.4: PZN+OV interface used in Study 2. This interface provided a single focus rather than two foci as in Study 1 (see Figure 4.4) and translucent rather than opaque arcs. The overview comprised 15% of total available pixels.




Figure 5.5: Calculation of levels of context in Study 2 RSN interfaces. Level of navigational context is the fraction of the size of the peripheral context areas, C , to the total size of the detail view, $F+C$. Level of overview context is the fraction of the size of the overview, O , to the total size of all views, $O+F+C$.




Figure 5.6: Calculation of levels of context in Study 2 PZN interfaces. The dotted line indicates the boundary between focus and context regions, which is not visually demarcated in the interfaces. Levels of navigational and overview context are as in Figure 5.5.

5.4 Apparatus

The experiment was conducted on the two systems used to conduct Study 1, configured as discussed in Section 4.4.

5.5 Participants

Forty subjects (24 female) between 18 and 39 years of age successfully completed the study and were each compensated \$15 for their participation. All subjects were right-handed, had normal or corrected to normal vision, and were not color blind. Participants were recruited through advertisements posted throughout the university campus and through an online participant scheduling system.

Originally, 44 subjects participated in the experiment. Two of the subjects were unable to follow the training instructions successfully, while two others followed the instructions but committed four or more errors (an error rate greater than 10%). These subjects were treated as outliers for the purpose of data analysis, leaving a total of 40 data points.

5.6 Design

The evaluation used a 2 (navigation, between subjects) by 2 (presence of overview, between subjects) by 7 (blocks, within subjects) design, where each block contained 5 trials. Subjects were randomly assigned to each of the four combinations of navigation and presence of overview. As in Study 1, a between-subjects design was chosen for the navigation and presence of overview factors to avoid potential transfer effects, as well due to the time required to train subjects on navigation with each technique. Due to the significant learning effects found in Study 1, the number of blocks was increased from 5 to 7 to ensure that subject performance reached a plateau that would enable an accurate comparison of the effects of navigation and presence of overview. Piloting showed that 7 blocks of 5 trials each were sufficient for performance to plateau.

5.7 Procedure

The experiment was designed to fit into a single 90 minute session. The experimenter first instructed subjects on the use of the navigation functions afforded by the interface to which they had been randomly assigned. Subjects were then shown the experimental task and instructed that they were to take as much time as necessary to solve it correctly while maximizing their efficiency. The experimenter then demonstrated the training strategy to be used by the subject. The training strategies varied depending on the navigation technique and presence or absence of an overview in each interface, and can be found in Appendix B.

All training strategies started with dragging out a long thin selection area along the horizontal path between two of the marked nodes, as described in Section 4.7. For the rubber sheet interfaces, selecting a long thin horizontal area had the effect of stretching the dataset along the vertical axis, as illustrated in Figures 5.1 and 5.3. For the pan and zoom interfaces, selecting a long thin horizontal area had the effect of zooming the contents of the focus box to fill the entire view, as illustrated in Figures 5.2 and 5.4.

For all interfaces, after zooming into the path between the marked nodes, subjects were instructed to count nodes that became visually salient. Following this step, subjects were shown how to drag out long thin horizontal and vertical selection areas to expand other compressed regions along the path. In accordance with the results of Study 1, which showed that using the panning function improved performance, subjects were instructed to use panning in pan and zoom interfaces, or the equivalent of panning in rubber sheet interfaces, to make adjustments to the focus area if needed. In interfaces with overviews, subjects were instructed how to use both the overview and detail views for navigation and counting nodes, but were not explicitly told to navigate in either view.

For the PZN-NoOV, although the “overview” and “zoom in, then zoom out” strategies performed equally well in Study 1, the “overview” strategy was too similar to usage patterns in the PZN+OV interface, and was eliminated as a

possibility through the removal of multiple foci for this study. Subjects were therefore instructed to use the “zoom in, then zoom out” strategy. This strategy involved slowly zooming out and adding nodes as they appeared along the path up the tree, as shown in Figure 5.2.

Following the discovery of one of the two topological distances, subjects were instructed to reset the interface before discovering the second distance. This strategy was motivated by the findings of Study 1 that showed that the use of reset between the discovery of the two topological distances improved performance. This result and its implications are discussed in detail in Bodnar’s thesis [12].

After being shown the strategies, subjects were given a training block of 5 trials. For each of the first 2 trials, the experimenter demonstrated solving the question using the strategies and then asked subjects to repeat this solution. For the last 3 trials of the session, subjects performed the task on their own, and the experimenter reminded them of the trained strategy as needed. At the end of the training session, subjects were given a one minute break before proceeding with the experiment. During both the training block and the experiment, subjects were provided with a pen and paper to record intermediate task results, and the experimenter remained in the experiment room to observe their progress. These changes from Study 1 were implemented to address procedural issues discussed in Section 4.12.

In the experiment, subjects were presented with 7 blocks, each containing 5 trials, for a total of 35 trials. All subjects were presented with an identical set of questions. The grouping of questions to block was predetermined, but the order of blocks was randomly generated for each subject. The blocks of questions were verified to be isomorphic in difficulty in piloting. The experimenter continued to observe the subject throughout the study, but never intervened. Subjects were given a one minute break between each block of questions.

At the end of the experiment, subjects completed questionnaires, which can be found in Appendix D. As in Study 1, the questionnaires were used to collect information about the subjects’ demographic background and computer usage as

well as ratings for ease of use, ease of navigation, and interface-specific features on 5-point Likert scales. For this study, the questionnaires also included the NASA-TLX scales [23], a standardized instrument for assessing various dimensions of workload. This instrument was included to improve our understanding of how demanding each interface was perceived by subjects. Short informal interviews were conducted with some of the subjects based on their questionnaire responses.

5.8 Measures

Our performance measures were based on logged data and included task completion times, errors, navigation actions (pan, zoom in, and zoom out), and reset actions. The inclusion of the two latter measures was motivated by the results of Study 1, where navigation and reset patterns were found to have a strong influence on performance. Self reported measures were collected through the post-experiment questionnaire, as described in the previous section. Like Study 1, this study was designed to minimize errors, with task completion times and number of navigation and reset actions used as the primary measures of performance.

5.9 Results

This section presents the results for both performance and self-reported measures of the experiment. The results for learning effects and navigation technique are reported in detail, while the results for presence of overview are summarized. Detailed reporting of the latter can be found in Bodnar's master's thesis [12].

A series of ANOVAs was run to understand the effect of navigation and overview on the performance and self-reported measures. Prior to these analyses, outlier data lying more than 3 standard deviations from the means of each cell were removed from the analysis. The Greenhouse-Geisser adjustment was used for non-spherical data, and the Bonferroni adjustment for post-hoc com-




Figure 5.7: Mean per-trial completion times by interface for blocks 1-7, in seconds ($N=40$).

parisons. As in the discussion of the results of Study 1 in Section 4.9, we report partial eta-squared (η^2), a measure of effect size. To interpret this value, .01 is a small effect size, .06 is medium, and .14 is large [17].

The overall results for mean completion times per trial are illustrated in Figure 5.7. As expected, performance improved as subjects progressed through the experiment, although the rate of improvement did vary among the interfaces, with a significant main effect of block ($F(3.174, 114.26) = 44.568$, $p < .001$, $\eta^2 = .553$) and a significant interaction between block and navigation ($F(3.176, 114.35) = 3.721$, $p < 0.02$, $\eta^2 = .094$).

Separate one-way repeated measures ANOVAs were run for each of the interfaces to determine performance plateaus. Post-hoc pairwise comparisons showed no differences between blocks 5, 6, and 7 for any of the interfaces, indicating that performance had reached a plateau by the end of the experiment in all interfaces, and therefore that the effect of block was successfully controlled for. Thus, for the analyses of completion times, navigation actions, and reset actions, we focus exclusively on blocks 1 and 7, which represent performance at

Dependent variable	Mean (RSN)	Mean (PZN)	F	Sig.	Partial η^2
Completion time (seconds)	59.58	47.49	13.744	.001	.276
Number of navigation actions	4.55	4.04	3.087	.090	.079
Number of resets	1.228	0.968	4.912	.014	.156

Table 5.1: Means and effects of navigation technique for the completion time, number of navigation actions, and number of resets dependent variables ($N=40$). Degrees of freedom are (1,36) for all dependent variables.

the beginning and end of the experiment. For these analyses, 2 (navigation) by 2 (presence of overview) by 2 (block) ANOVAs were performed.

Counter to our hypothesis **H3**, both our logged and self-reported measures showed that pan and zoom navigation outperformed rubber sheet navigation. The results for completion times by navigation for blocks 1 and 7 are illustrated in Figure 5.8. Both at the beginning and at the end of the experiment, pan and zoom interfaces were significantly faster than interfaces using rubber sheet navigation. Subjects also performed borderline significantly fewer navigation actions and significantly fewer resets using pan and zoom navigation, as shown in Table 5.1.

Two-way (navigation by overview) ANOVAs were conducted on each of the NASA-TLX measures. These analyses showed that mental demand was significantly lower in the pan and zoom interfaces ($F(1,36) = 4.214$, $p < .05$, $\eta^2 = .105$). Subjects also reported that pan and zoom interfaces were significantly easier to navigate ($F(1,36) = 10.385$, $p < .005$, $\eta^2 = .224$). Both of these self-reported measures support the results obtained from the logged performance measures.

Presence of overview had no significant effect on any of the performance measures. This finding supports our hypothesis **H4**, but is counter to our hypothesis **H5**. The self-reported measures did, however, favor an overview. Subjects reported a lower physical demand for interfaces with an overview and




Figure 5.8: Mean per-trial completion times in seconds by navigation technique for blocks 1 and 7, above, and by block, below (N=40). PZN was significantly faster than RSN in both blocks, regardless of presence or absence of overview.

found them more enjoyable to use. Bodnar's thesis [12] reports on and discussed these results in more detail.

On average, subjects committed 1.6 errors over the course of the experiment, for a mean error rate of 4.7%. There were no significant main or interaction effects of navigation or presence of overview on error rate.

5.10 Summary of Results

We summarize our results according to the hypotheses stated in Section 5.1:

- R3:** Pan and zoom navigation interfaces performed better than rubber sheet navigation interfaces in terms of completion times, number of navigation actions, and number of reset actions. Mental demand was also reported as lower in pan and zoom interfaces.
- R4:** For rubber sheet navigation, having an overview made no significant difference in terms of completion times, navigation actions, or resets. Having an overview was, however, reported to reduce physical demand.
- R5:** Similarly, for pan and zoom navigation, having an overview made no significant difference in terms of completion times, navigation actions, or resets. Having an overview was, however, reported to reduce physical demand.

5.11 Discussion

In Study 2, pan and zoom navigation resulted in lower task completion time, number of navigation actions, and number of reset actions than rubber sheet navigation. These results present convergent evidence for pan and zoom navigation outperforming rubber sheet navigation, a finding that can be partially explained by the nature of the task used in our study. Following a topological path in a large tree dataset under the distortion inherent in rubber sheet navigation may have caused a loss of orientation. This explanation is supported by the fact that rubber sheet interfaces required significantly more reset actions.

Our observations showed that, when subjects seemed disoriented in terms of their location in the dataset, they frequently used the reset key to revert to the original state of the visualization, which indicates that the number of reset actions may be a good indicator of loss of orientation.

The finding that pan and zoom navigation performed better than rubber sheet navigation is not consistent with Gutwin and Skopik's finding that focus+context interfaces can perform as well as or better on path navigation tasks as panning interfaces with overviews [22]. To explain this discrepancy, it must be noted that Gutwin and Skopik's non-distortion-based interfaces did not implement zooming and featured an interaction model users found confusing. Furthermore, the task used in their study required users to move the mouse pointer along a path. Gutwin and Skopik suggested that a task that simply requires users to visually navigate a path, which was the case in our study, would be much more amenable to interfaces without integrated focus and context.

Although subjects could be expected to be more familiar with pan and zoom navigation from previous experience, the difference in performance between pan and zoom and rubber sheet interfaces did not decrease significantly as subjects became more adept in the use of the interfaces. We therefore speculate that the experience subjects gained with rubber sheet navigation during the experiment may not have been sufficient to overcome inherent difficulties with this form of navigation. Further investigation is required to determine whether more experience with rubber sheet navigation can enable subjects to improve their performance beyond the plateau observed in this study.

In terms of subjective experience, pan and zoom interfaces were rated by subjects as less mentally demanding and easier to navigate than their rubber sheet equivalents, regardless of the presence of an overview. This result is consistent with results of previous studies [21, 25], which found that non-distortion panning interfaces with overviews were preferred to their distortion-based counterparts. We postulate that, since the pan and zoom interface without an overview used in our study provided an alternative form of contextual information (namely Halo-like arcs), subjects did not consider the lack of an overview to be a liability when

assessing their perceived mental demand and ease of navigation.

Qualitative feedback from questionnaires and interviews highlighted three main areas where subjects encountered difficulty with rubber sheet navigation. First, subjects reported that the global effect of distortions resulting from rubber sheet navigation prevented them from forming a stable mental model of the dataset, even when an overview was present. Second, subjects reported being confused by the effects of the rubber sheet navigation equivalent of panning, which was also implemented using distortion. Some subjects stated that they would have preferred this operation to feel like panning in pan and zoom interfaces they have previously experienced, such as Google Maps [20]. Finally, a number of subjects identified the use of a single fixed focus region in rubber sheet navigation interfaces as a hindrance, and suggested that they would have benefitted from movable, resizable focus areas, as implemented in the original TreeJuxtaposer application [37]. These findings suggest that the constraints imposed on the size and location of focus areas in our studies adversely affected the usability of the rubber sheet navigation interfaces. Our original rationale for imposing these constraints was to control the level of context, a parameter that was shown to have no significant effect on performance in Study 1. For this reason, we recommend that future studies of rubber sheet navigation relax these constraints as necessary to provide usability on a par with comparison navigation techniques.

A possible confounding factor in this study was the lack of adjustment to the best-performing levels of context derived from Study 1 results to account for the merging of the two focus regions used in that study into a single focus region. However, we postulate that the magnitude of such an adjustment would have been minor, and its impact negligible given the lack of significant effect of level of context on the results of Study 1.

Chapter 6

Conclusions and Future Work

This thesis presented the first empirical evaluations in the information visualization literature comparing pan and zoom and rubber sheet navigation techniques. We conducted two experiments, each of which involved interfaces implementing all combinations of these two navigation techniques with and without an overview. Our results indicate that pan and zoom navigation was significantly faster than rubber sheet navigation, required fewer navigation actions, and demanded less mental effort to complete a topological comparison task in a large tree dataset, regardless of the presence of an overview. We also found that level of context did not significantly influence performance with either navigation technique, while interface-specific strategies developed by subjects did.

As controlled laboratory experiments, the studies described in this thesis are necessarily limited in terms of realism and generalizability [36]. The derivation of the task employed in the study from discussions with experts in the domain of phylogenetic biology, combined with the use of a dataset from the same domain, render this work more ecologically valid than most of the studies that have preceded it in the literature. However, the degree of realism provided by these factors is tempered by the use of non-expert subjects and a task that did not require a knowledge of phylogenetic biology. The generalizability of the results of this study is likewise limited by the use of a single, albeit compound, task to represent the variety of possible topological tasks of interest to phylogenetic biologists, as well as the usability issues with our implementation of rubber sheet

navigation identified in the previous section. Given these limitations, this work should be regarded as a starting point for further exploration into the properties of pan and zoom and rubber sheet navigation rather than a definitive statement about their relative performance.

Several possibilities for future studies arise from the results described in this thesis. Comparisons of pan and zoom and rubber sheet navigation methods using other topological tasks such as those described in Chapter 3, as well as more general non-topological tasks and different types of datasets, could yield an improved understanding of the relative strengths and weaknesses of each navigation technique. For instance, a future experiment could use a steering task similar to that used by Gutwin and Skopik [22] to determine whether the benefits of the integration of focus and context regions provided by the fisheye views in that study extend to rubber sheet navigation. The comparison could also be extended to include other variants on the pan and zoom and focus+context navigation metaphors, such as fisheye views and semantic zooming.

The findings from Study 1 suggest that navigation strategy can have a significant impact on user performance, as well as the possibility of an interaction between strategies developed by users and affordances provided by navigation techniques. Therefore, another obvious next step would be to conduct a systematic exploration of the strategies users naturally adopt on their own to complete different tasks with each type of navigation examined in our study. It would be equally valuable to examine possibilities for how the navigation techniques examined in this study might be tuned to better accommodate strategies users would naturally tend to develop when using them. For instance, future studies could examine whether, as suggested by our findings, users can develop zooming strategies that provide performance benefits equivalent to those of an overview, and determine whether zooming enhancements such as speed-dependent automatic zooming [26] can further improve performance with such strategies.

The need for more than one focus region is another issue that affected the results of the studies described in this thesis and deserves further investigation. Our Study 1 showed that the presence of two focus areas in the original versions

of our interfaces increased their complexity and required subjects to invest too much cognitive effort into coordinating navigation between them. However, a number of subjects in Study 2 suggested that multiple focus areas would have proved beneficial, especially in conjunction with rubber sheet navigation. An empirical examination of the combinations of navigation techniques, tasks, and datasets for which multiple focus areas may prove either helpful or a hindrance could provide a valuable addition to the information visualization literature.

Finally, a logical direction for continuing this work is to explore how the navigation methods examined in it compare when used by domain experts in phylogenetic biology in their work. As previously mentioned, although subjects in our study became more proficient in their use of rubber sheet navigation over the course of the experiment, and their learning occurred at a faster rate than that of subjects using pan and zoom navigation, this improvement was not sufficient to significantly decrease the performance differential between them. It would therefore be valuable to investigate whether the use of rubber sheet navigation by domain expert users over an extended period of time would result in this gap decreasing to the point where performance with rubber sheet navigation is not significantly different from that with pan and zoom navigation. Such an investigation could also be extended into other aspects of potential future work we have discussed here, such as task-specific strategies and multiple focus areas, to determine how the usage patterns of domain experts differ from those of non-expert users.

This work represents an initial step in investigating the properties of rubber sheet navigation and how it compares to the more established pan and zoom navigation paradigm. The eventual goal of these investigations would be the development of design guidelines of relevance to a wide range of information visualization applications that could benefit from improved methods for navigating through large amounts of data.

Bibliography

- [1] Adobe Corporation. Adobe Photoshop website. <http://www.adobe.com/products/photoshop>. Retrieved October 15, 2005.
- [2] Ahlberg, C. and Shneiderman, B. (1994). Visual information seeking: tight coupling of dynamic query filters with starfield displays. In *CHI '94: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 313–317, New York, NY, USA. ACM Press.
- [3] Bartram, L., Ho, A., Dill, J., and Henigman, F. (1995). The continuous zoom: a constrained fisheye technique for viewing and navigating large information spaces. In *UIST '95: Proceedings of the 8th annual ACM symposium on User interface and software technology*, pages 207–215, New York, NY, USA. ACM Press.
- [4] Baudisch, P., Good, N., and Stewart, P. (2001). Focus plus context screens: combining display technology with visualization techniques. In *UIST '01: Proceedings of the 14th annual ACM symposium on User interface software and technology*, pages 31–40, New York, NY, USA. ACM Press.
- [5] Baudisch, P., Lee, B., and Hanna, L. (2004). Fishnet, a fisheye web browser with search term popouts: a comparative evaluation with overview and linear view. In *AVI '04: Proceedings of the working conference on Advanced visual interfaces*, pages 133–140, New York, NY, USA. ACM Press.
- [6] Baudisch, P. and Rosenholtz, R. (2003). Halo: a technique for visualizing off-screen objects. In *CHI '03: Proceedings of the SIGCHI conference on*

- Human factors in computing systems*, pages 481–488, New York, NY, USA. ACM Press.
- [7] Bederson, B. B. and Boltman, A. (1999). Does animation help users build mental maps of spatial information? In *InfoVis '99: Proceedings of the 1999 IEEE Symposium on Information Visualization*, page 28, Washington, DC, USA. IEEE Computer Society.
- [8] Bederson, B. B. and Hollan, J. D. (1994). Pad++: a zooming graphical interface for exploring alternate interface physics. In *UIST '94: Proceedings of the 7th annual ACM symposium on User interface software and technology*, pages 17–26, New York, NY, USA. ACM Press.
- [9] Bederson, B. B., Meyer, J., and Good, L. (2000). Jazz: an extensible zoomable user interface graphics toolkit in java. In *UIST '00: Proceedings of the 13th annual ACM symposium on User interface software and technology*, pages 171–180, New York, NY, USA. ACM Press.
- [10] Bier, E. A., Stone, M. C., Pier, K., Buxton, W., and DeRose, T. D. (1993). Toolglass and magic lenses: the see-through interface. In *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 73–80, New York, NY, USA. ACM Press.
- [11] Björk, S., Holmquist, L. E., and Redström, J. (1999). A framework for focus+context visualization. In *InfoVis '99: Proceedings of the 1999 IEEE Symposium on Information Visualization*, page 53, Washington, DC, USA. IEEE Computer Society.
- [12] Bodnar, A. (2005). An evaluation of overviews for large tree navigation. Master's thesis, The University of British Columbia.
- [13] Card, S. K., Mackinlay, J. D., and Shneiderman, B., editors (1999). *Readings in information visualization: using vision to think*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

-
- [14] Card, S. K. and Nation, D. (2002). Degree-of-interest trees: A component of an attention-reactive user interface. In Marsico, M. D., Levialdi, S., and Panizzi, E., editors, *Proc. AVI 2002: Advanced Visual Interfaces.*, Trento, Italy. ISBN=1-58113-537-8.
- [15] Carrizo, S. F. (2004). Phylogenetic trees: an information visualisation perspective. In *CRPIT '04: Proceedings of the second conference on Asia-Pacific bioinformatics*, pages 315–320, Darlinghurst, Australia. Australian Computer Society, Inc.
- [16] Chen, C. and Czerwinski, M. (2000). Empirical evaluation of information visualizations: An introduction. *International Journal of Human-Computer Studies*, 53:631–635.
- [17] Cohen, J. (1973). Eta-squared and partial eta-squared in communication science. *Human Communication Research*, 28:473–490.
- [18] Furnas, G. W. (1986). Generalized fisheye views. In *CHI '86: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 16–23, New York, NY, USA. ACM Press.
- [19] Furnas, G. W. (1997). Effective view navigation. In *CHI '97: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 367–374, New York, NY, USA. ACM Press.
- [20] Google, Inc. Google Maps website. <http://maps.google.com>. Retrieved October 15, 2005.
- [21] Gutwin, C. and Fedak, C. (2004). Interacting with big interfaces on small screens: a comparison of fisheye, zoom, and panning techniques. In *GI '04: Proceedings of the 2004 conference on Graphics interface*, pages 145–152, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada. Canadian Human-Computer Communications Society.

-
- [22] Gutwin, C. and Skopik, A. (2003). Fisheyes are good for large steering tasks. In *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 201–208, New York, NY, USA. ACM Press.
- [23] Hart, S. G. and Staveland, L. E. (1988). Development of NASA-TLX (Task Load Index): results of empirical and theoretical research. In Hancock, P. and Meshkati, N., editors, *Advances in psychology: human mental workload*, pages 139–183. Elsevier Science, Amsterdam.
- [24] Hornbaek, K., Bederson, B. B., and Plaisant, C. (2002). Navigation patterns and usability of zoomable user interfaces with and without an overview. *ACM Trans. Comput.-Hum. Interact.*, 9(4):362–389.
- [25] Hornbaek, K. and Frokjaer, E. (2001). Reading of electronic documents: the usability of linear, fisheye, and overview+detail interfaces. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*.
- [26] Igarashi, T. and Hinckley, K. (2000). Speed-dependent automatic zooming for browsing large documents. In *UIST '00: Proceedings of the 13th annual ACM symposium on User interface software and technology*, pages 139–148, New York, NY, USA. ACM Press.
- [27] Johnson, J. A. (1995). A comparison of user interfaces for panning on a touch-controlled display. In *CHI '95: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 218–225, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co.
- [28] Jul, S. and Furnas, G. W. (1998). Critical zones in desert fog: aids to multiscale navigation. In *UIST '98: Proceedings of the 11th annual ACM symposium on User interface software and technology*, pages 97–106, New York, NY, USA. ACM Press.

-
- [29] Keahey, T. A. and Robertson, E. L. (1997). Nonlinear magnification fields. In *InfoVis '97: Proceedings of the 1997 IEEE Symposium on Information Visualization*, page 51, Washington, DC, USA. IEEE Computer Society.
- [30] Kobsa, A. (2004). User experiments with tree visualization systems. In *InfoVis '04: Proceedings of the IEEE Symposium on Information Visualization*, pages 9–16, Washington, DC, USA. IEEE Computer Society.
- [31] Lamping, J., Rao, R., and Pirolli, P. (1995). A focus+context technique based on hyperbolic geometry for visualizing large hierarchies. In *CHI '95: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 401–408, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co.
- [32] Landauer, T. (1997). *Behavioral research methods in human-computer interaction*, pages 203–227. Elsevier Science, Amsterdam.
- [33] Lau, K., Rensink, R. A., and Munzner, T. (2004). Perceptual invariance of nonlinear Focus+Context transformations. In *APGV '04: Proceedings of the 1st Symposium on Applied perception in graphics and visualization*, pages 65–72, New York, NY, USA. ACM Press.
- [34] Lee, B., Parr, C. S., Campbell, D., and Bederson, B. B. (2004). How users interact with biodiversity information using TaxonTree. In *AVI '04: Proceedings of the working conference on Advanced visual interfaces*, pages 320–327, New York, NY, USA. ACM Press.
- [35] Leung, Y. K. and Apperley, M. D. (1994). A review and taxonomy of distortion-oriented presentation techniques. *ACM Transactions on Computer-Human Interaction*, 1(2):126–160.
- [36] McGrath, J. E. (1995). Methodology matters: doing research in the behavioral and social sciences. pages 152–169.

-
- [37] Munzner, T., Guimbretiere, F., Tasiran, S., Zhang, L., and Zhou, Y. (2003). TreeJuxtaposer: scalable tree comparison using Focus+Context with guaranteed visibility. *ACM Transactions on Graphics*, 22(3):453–462.
- [38] Nekrasovski, D., Bodnar, A., McGrenere, J., Guimbretiere, F., and Munzner, T. (2005). An evaluation of pan & zoom and rubber sheet navigation with and without an overview. In submission.
- [39] Olduvai project website. <http://olduvai.sourceforge.net/tj/index.shtml>. Retrieved October 15, 2005.
- [40] Perlin, K. and Fox, D. (1993). Pad: an alternative approach to the computer interface. In *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 57–64, New York, NY, USA. ACM Press.
- [41] Plaisant, C., Carr, D., and Shneiderman, B. (1995). Image-browser taxonomy and guidelines for designers. *IEEE Software*, 12(2):21–32.
- [42] Plaisant, C. and Fekete, J.-D. Infovis 2003 contest. <http://www.cs.umd.edu/hcil/iv03contest/>. Retrieved October 15, 2005.
- [43] Plaisant, C., Grosjean, J., and Bederson, B. B. (2002). Spacetree: Supporting exploration in large node link tree, design evolution and empirical evaluation. In *InfoVis '02: Proceedings of the IEEE Symposium on Information Visualization*, pages 57–64, Washington, DC, USA. IEEE Computer Society.
- [44] Sarkar, M. and Brown, M. H. (1994). Graphical fisheye views. *Communications of the ACM*, 37(12):73–83.
- [45] Sarkar, M., Snibbe, S. S., Tversky, O. J., and Reiss, S. P. (1993). Stretching the rubber sheet: a metaphor for viewing large layouts on small screens. In *UIST '93: Proceedings of the 6th annual ACM symposium on User interface software and technology*, pages 81–91, New York, NY, USA. ACM Press.

-
- [46] Schaffer, D., Zuo, Z., Greenberg, S., Bartram, L., Dill, J., Dubs, S., and Roseman, M. (1996). Navigating hierarchically clustered networks through fisheye and full-zoom methods. *ACM Trans. Comput.-Hum. Interact.*, 3(2):162–188.
- [47] Slack, J., Hildebrand, K., and Munzner, T. (2005). Prasad: Partitioned rendering infrastructure for stable accordion drawing. In *InfoVis '05: Proceedings of the IEEE Symposium on Information Visualization*, pages 41–48, Washington, DC, USA. IEEE Computer Society.
- [48] Spence, R. and Apperley, M. (1982). Data base navigation: An office environment for the professional. *Behaviour and Information Technology*, 1(1):43–54.
- [49] Wehrend, S. and Lewis, C. (1990). A problem-oriented classification of visualization techniques. In *Vis '90: Proceedings of the 1st conference on Visualization '90*, pages 139–143, Los Alamitos, CA, USA. IEEE Computer Society Press.
- [50] Woodruff, A., Landay, J., and Stonebraker, M. (1998). Constant information density in zoomable interfaces. In *AVI '98: Proceedings of the working conference on Advanced visual interfaces*, pages 57–65, New York, NY, USA. ACM Press.
- [51] Zellweger, P. T., Mackinlay, J. D., Good, L., Stefik, M., and Baudisch, P. (2003). City lights: contextual views in minimal space. In *CHI '03: Extended abstracts of the SIGCHI conference on Human factors in computing systems*, pages 838–839, New York, NY, USA. ACM Press.

Appendix A

Study 1 Training Protocol

All interfaces

Thank you for your willingness to participate in our experiment. You will be helping us evaluate different techniques for visualizing large datasets. You will be asked to complete a series of tasks that involve determining relative distances in large trees. First, let's review some concepts that will help you to complete the tasks.

Present subjects with paper tests.

The task you will perform in this experiment consists of determining the topological distance between a series of marked nodes in the displayed tree, where topological distances are measured by the number of black squares between marked nodes. Remember from the tests that you just completed that topological distance will not equal geometric distance.

We will now explore the features of the interface you will use.

RSN-NoOV

This interface enables you to explore the dataset using a series of zooming and panning actions that use the metaphor of stretching a rubber sheet with its borders tacked down.

The left mouse button will allow you to drag out a box, the contents of which will fill one of the RED or ORANGE focus boxes. The rest of the tree will then be squished around the focus box but will remain visible at all times.

Ask participant to try dragging out a box

As you are dragging out a box, you may hold down the SHIFT key to indicate that you would like this box to be the new RED focus box, or hold down the CTRL key to indicate that you would like this box to be the new ORANGE focus box. Note also that the SHIFT key is above the CTRL key, just like the RED focus box is above the ORANGE focus box.

Ask participant to try dragging out a box using SHIFT and CTRL

If you do not select either the SHIFT or CTRL key, the tool will choose which focus box to place the contents of your new box based on the proximately

of your newly dragged out box to the existing RED and ORANGE focus boxes. You can zoom out by dragging out a box which is larger than either of the colored focus boxes.

Ask participant to try zooming out

The right mouse button will allow you to pan horizontally and vertically within the dataset using either a horizontal or vertical drag motions, which will let you fine tune your selection.

Ask participant to try panning

PZN-NoOV

This interface enables you to explore the dataset using two views which you can navigate through a series of pan and zoom actions.

Show subject paper illustration of two views

The two detail views are independent of one another, so you can navigate in one without affecting the other. It is also possible to overlap the two detail views, and even have one inside the other. The left mouse button will allow you to drag out a box which will become the new extent of your detail view.

Ask participant to try zooming in

Once you are zoomed in, you may hold down the right mouse button and pan in any direction. You cannot pan if you are zoomed out entirely.

Ask participant to try panning

Holding down the middle mouse button and dragging the mouse toward you will allow you to zoom out. As you zoom out, you may also drag the mouse in the opposite direction to zoom back in, but only to the extent that you first began to zoom out.

Ask participant to try zooming out

If a marked node is not currently in view, an arc will appear at the border of the detail view, indicating the direction and distance from your current focus box to the marked node. The arc is part of a circular ring that surrounds one of the nodes which is currently off-screen. This ring is just large enough to reach

the border region of the display. The colour of the arc indicates the colour of the marked node it represents. Once a marked node is visible on screen, the arc will disappear. No marks will appear in the overview window since marked nodes are always visible. Arcs are view dependent.

RSN+OV

This interface enables you to explore the dataset using a series of zooming and panning actions that use the metaphor of stretching a rubber sheet with its borders tacked down. A separate window will provide you with an overview of the dataset, and will not be distorted.

The left mouse button will allow you to drag out a box, the contents of which will fill one of the RED or ORANGE focus boxes. The rest of the tree will then be squished around the focus box but will remain visible at all times.

Ask participant to try dragging out a box

As you are dragging out a box, you may hold down the SHIFT key to indicate that you would like this box to be the new RED focus box, or hold down the CTRL key to indicate that you would like this box to be the new ORANGE focus box. Note also that the SHIFT key is above the CTRL key, just like the RED focus box is above the ORANGE focus box.

Ask participant to try dragging out a box using SHIFT and CTRL

If you do not select either the SHIFT or CTRL key, the tool will choose which focus box to place the contents of your new box based on the proximity of your newly dragged out box to the existing RED and ORANGE focus boxes. You can zoom out by dragging out a box which is larger than either of the colored focus boxes.

Ask participant to try zooming out

The right mouse button will allow you to pan horizontally and vertically within the dataset using either a horizontal or vertical drag motions, which will let you fine tune your selection.

Ask participant to try panning

A separate smaller window will provide you with an overview of the dataset, and indicate where in the dataset your current focus boxes are. In the overview, the left mouse button will allow you to drag out a box, the contents of which will fill one of the RED or ORANGE focus boxes. As you are dragging out a box, you may hold down the SHIFT key to indicate that you would like this box to be the new RED focus box, or hold down the CTRL key to indicate that you would like this box to be the new ORANGE focus box.

Ask participant to try dragging out a box in the overview using SHIFT and CTRL

You may also hold down the right mouse button while inside one of the boxes representing the location of your focus box and move it to wherever you like within the bounds of the overview using a series of drag actions.

Ask participant to try panning in the overview

PZN+OV

This interface enables you to explore the dataset using two detail views which you can navigate through a series of pan and zoom actions.

Show subject paper illustration of two views

The two detail views are independent of one another, so you can navigate in one without affecting the other. It is also possible to overlap the two detail views, and even have one inside the other. The left mouse button will allow you to drag out a box which will become the new extent of your detail view.

Ask participant to try zooming in a detail view

Once you are zoomed in, you may hold down the right mouse button and pan in any direction. You cannot pan if you are zoomed out entirely.

Ask participant to try panning in a detail view

Holding down the middle mouse button and dragging the mouse toward you will allow you to zoom out. As you zoom out, you may also drag the mouse in the opposite direction to zoom back in, but only to the extent that you first began to zoom out.

Ask participant to try zooming out in a detail view

A separate smaller window will provide you with an overview of the dataset, and indicate where in the dataset your current detail views are. In the overview, the left mouse button will allow you to drag out a box which will become the new extent of your detail view. As you are dragging out a box, you may hold down the SHIFT key to indicate that you would like this box to be the new RED detail view, or hold down the CTRL key to indicate that you would like this box to be the new ORANGE detail view.

Ask participant to try zooming in overview using SHIFT and CTRL

Note also that the SHIFT key is above the CTRL key, just like the RED detail view above the ORANGE detail view. You may also hold down the right mouse button while inside one of the boxes representing the location of your detail view and move it to wherever you like within the bounds of the overview using a series of drag actions. The modifier keys only work in the overview window.

Ask participant to try panning in overview using SHIFT and CTRL

If a marked node is not currently in view, an arc will appear at the border of the detail view, indicating the direction and distance from your current focus box to the marked node. The arc is part of a circular ring that surrounds one of the nodes which is currently off-screen. This ring is just large enough to reach the border region of the display. The colour of the arc indicates the colour of the marked node it represents. Once a marked node is visible on screen, the arc will disappear. No marks will appear in the overview window since marked nodes are always visible. Arcs are view dependent.

All interfaces

Do you have any questions about this interface?

The R key can be pressed to reset your current view to its initial startup state.

The ESC key can be pressed during a box drag action to cancel your current

drag.

A question panel at the top of the screen will display a question which will require you to use the interface to solve. The question will ask you to compare the topological distances between marked nodes in the tree. The topological distance between marked nodes will never be equal. The question will never change, but the location of the marked nodes will, thus you will be required to navigate and explore different areas within the large tree to answer the question correctly.

When you have discovered the answer, we ask that you select the appropriate check box and click on the submit button. This will allow you to move onto the next question. An instruction panel at the left of the interface will serve as a reminder of interface specific controls

We will now ask you to complete a series of training tasks using this interface. There is no time limit for completing these tasks - we want you to take as much time as you need to ensure that your answer is correct.

We want to emphasize that we are evaluating the system and not your ability to use it. For this reason, you will receive no feedback as to whether your answers for the tasks were correct.

A good strategy for solving the tasks is to draw out long horizontal thin boxes. This will help you to see the larger tree in more detail.

Appendix B

Study 2 Training Protocol

All interfaces

Thank you for your willingness to participate in our experiment. You will be helping us evaluate different techniques for visualizing large datasets. You will be asked to complete a series of tasks that involve determining relative distances in large trees. First, let's review some concepts that will help you to complete the tasks.

Present subjects with paper tests.

The task you will perform in this experiment consists of determining the topological distance between a series of marked nodes in the displayed tree, where topological distances are measured by the number of black squares between marked nodes. Remember from the tests that you just completed that topological distance will not equal geometric distance.

We will now explore the features of the interface you will use.

RSN-NoOV

This interface enables you to explore the dataset using a view which you can navigate using pan and zoom actions. The view uses the metaphor of stretching and squishing a rubber sheet with its borders tacked down. Note that the colored nodes are visible at all times, even if they are squished to the edges of the view.

The left mouse button will allow you to drag out a box, the contents of which will fill the red box. The rest of the tree will then be squished around the red box but will remain visible at all times.

Ask participant to try dragging out a box.

You can zoom out by dragging out a box which is larger than the red box.

Ask participant to try zooming out.

The right mouse button will allow you to pan horizontally and vertically within the view using either horizontal or vertical drag motions, which will let you fine tune your selection.

Ask participant to try panning.

You can use the colored nodes as visual anchors to help maintain orientation while performing navigation actions. As you zoom or pan, you can monitor the location and size of the colored nodes, which will give you an idea of what path to follow and how much farther you have to go.

PZN-NoOV

This interface enables you to explore the dataset using a view which you can navigate using pan and zoom actions.

The left mouse button will allow you to drag out a box, the contents of which will then zoom to fill the view completely.

Ask participant to try zooming in.

Once you are zoomed in, you may hold down the right mouse button and pan in any direction. You cannot pan if you are zoomed out entirely.

Ask participant to try panning.

Holding down the middle mouse button and dragging the mouse toward you will allow you to zoom out. As you zoom out, you may also drag the mouse in the opposite direction to zoom back in, but only to the extent that you first began to zoom out.

Ask participant to try zooming out.

If a marked node is not currently in view, a colored arc will appear at the border of the detail view, indicating the direction and distance from your current focus box to the marked node. The arc is part of a circular ring that surrounds any marked node which is currently off-screen. The color of the arc indicates the color of the marked node it represents. Once a marked node is visible on screen, the arc will disappear.

You can use the arcs as visual anchors to help maintain orientation of marked nodes while performing navigation actions. As you zoom out or pan, you can monitor the shape and size of the arc, which will give you an idea of what path to follow and how much farther you have to go.

RSN+OV

This interface enables you to explore the dataset using two views which you can navigate through using pan and zoom actions.

The larger view will display detailed information about parts of the dataset. This view uses the metaphor of stretching and squishing a rubber sheet with its borders tacked down. Note that the colored nodes are visible at all times, even if they are squished to the edges of this view.

The smaller view will provide you with an overview of the dataset, and indicate where in the dataset the detail view is at any given time. This view does not use the rubber sheet metaphor.

The left mouse button will allow you to drag out a box, the contents of which will fill the red box. The rest of the tree will then be squished around the red box but will remain visible at all times.

Ask participant to try zooming in in the detail view.

You can zoom out by dragging out a box which is larger than the red box.

Ask participant to try zooming out in the detail view.

The right mouse button will allow you to pan horizontally and vertically within the view using either a horizontal or vertical drag motions, which will let you fine tune your selection.

Ask participant to try panning in the detail view.

In the smaller view, the left mouse button will allow you zoom into an area by dragging out a box which will become the new contents of the red box in the larger view.

Ask participant to try zooming in overview.

You can also hold down the right mouse button while inside the red box in the smaller view, and move it within the view using a series of drag actions.

Ask participant to try panning in overview.

You can use the colored nodes as visual anchors to help maintain orientation while performing navigation actions. As you zoom or pan, you can monitor the location and size of the colored nodes, which will give you an idea of what path

to follow and how much farther you have to go.

PZN+OV

This interface enables you to explore the dataset using two views which you can navigate through using pan and zoom actions. The larger view will display detailed information about parts of the dataset. The smaller view will provide you with an overview of the dataset, and indicate where in the dataset the detail view is at any given time.

The left mouse button will allow you to drag out a box, the contents of which will then zoom to fill the larger view completely.

Ask participant to try zooming in the detail view.

Once you are zoomed in, you may hold down the right mouse button and pan in any direction. You cannot pan if you are zoomed out entirely.

Ask participant to try panning in the detail view.

Holding down the middle mouse button and dragging the mouse toward you will allow you to zoom out. As you zoom out, you may also drag the mouse in the opposite direction to zoom back in, but only to the extent that you first began to zoom out.

Ask participant to try zooming out in the detail view.

In the smaller view, the left mouse button will allow you zoom into an area by dragging out a box which will become the new extent of your detail view.

Ask participant to try zooming in overview.

You can also hold down the right mouse button while inside the red box in the smaller view, and move it within the view using a series of drag actions.

Ask participant to try panning in overview.

If a marked node is not currently in view, a colored arc will appear at the border of the detail view, indicating the direction and distance from your current focus box to the marked node. The arc is part of a circular ring that surrounds any marked node which is currently off-screen. The color of the arc indicates the color of the marked node it represents. Once a marked node is visible on

screen, the arc will disappear.

You can use the arcs as visual anchors to help maintain orientation of marked nodes while performing navigation actions. As you zoom out or pan, you can monitor the shape and size of the arc, which will give you an idea of what path to follow and how much farther you have to go.

All interfaces

Do you have any questions about this interface?

The R key can be pressed to reset your current view to its initial startup state.

The ESC key can be pressed during a box drag action to cancel your current drag.

All the controls I just showed you are also listed at the left of the window in case you need a reminder.

At the top of the window is the task you will perform in this experiment. You will need to determine whether the purple node is topologically closer to the blue node or the green node in the tree. The task will never change, but the location of the marked nodes will with each task. You cannot skip or go back to previously answered questions.

Note that the topological distances to the blue node and the green node will never be equal, but they may be close. If it seems as though they are equal, perform more navigation, and you will discover that they are different from each other.

Note that there is only one path between any two nodes in the tree.

You can use this pen/pencil and sheet of paper to write down topological distances between nodes so that you don't have to remember them as you performing the task.

When you are ready, select the appropriate answer and click on the submit button. This will allow you to move onto the next question.

We want to emphasize that we are evaluating the system and not your ability

to use it. For this reason, you will receive no indication of whether your answer is correct.

There is no time limit for completing these tasks. Take as much time as you need to ensure that your answer is correct, but do work as efficiently as you can.

RSN-NoOV

A good strategy for using this interface is to draw out long thin boxes. This will help you to see the larger tree in more detail. It's often helpful to draw long horizontal boxes to zoom into the details of the dataset, and to draw long vertical boxes to expand areas that are squished vertically.

Demonstrate this, then ask participant to do it.

Another useful strategy is to reset the interface when you have found one of the topological distances before you move onto another distance.

Demonstrate this, then ask participant to do it.

PZN-NoOV

A good strategy for using this interface is to draw out long thin boxes. This will help you to see the larger tree in more detail.

Demonstrate this, then ask participant to do it.

Once you have zoomed in to the area around either the blue or the green node, you can count the number of nodes on the path that are close to it. Then you can slowly zoom out and, as you see more nodes on the path to the purple node, add them to your count.

Demonstrate this, then ask participant to do it.

Additionally, you can reset the interface when you have found one of the topological distances before you move onto another distance.

RSN+OV

A good strategy for using this interface is to draw out long thin boxes. This will help you to see the larger tree in more detail. It's often helpful to draw

long horizontal boxes to zoom into the details of the dataset, and to draw long vertical boxes to expand areas that are squished vertically.

Demonstrate this, then ask participant to do it.

Another useful strategy is to first zoom in to the area around either the blue or the green node using the small view. Then you can use either view to explore the path to the purple node. Note that you can count nodes along the path in either view. If you need to make small adjustments, you can pan; for larger movements, you can zoom in either view.

Demonstrate this, then ask participant to do it.

You can also reset the interface when you have found one of the topological distances before you move onto another distance.

Demonstrate this, then ask participant to do it.

We strongly suggest you use these strategies as you are answering the questions.

PZN+OV

A good strategy for using this interface is to draw out long thin boxes. This will help you to see the larger tree in more detail.

Demonstrate this, then ask participant to do it.

Another useful strategy is to first zoom in to the area around either the blue or the green node using the small view. Then you can use either view to explore the path to the purple node. Note that you can count nodes along the path in either view. If you need to make small adjustments, you can pan; for larger movements, you can zoom in either view.

Demonstrate this, then ask participant to do it.

Additionally, you can reset the interface when you have found one of the topological distances before you move onto another distance.

Demonstrate this, then ask participant to do it.

All interfaces

We strongly suggest you use these strategies as you are answering the questions.

Appendix C

Study 1 Questionnaires



THE UNIVERSITY OF BRITISH COLUMBIA

Experimental Questionnaire

Evaluation of Information Visualization Techniques

Interface # 1

Subject # _____

Part 1**1. Age Group**

- ☐ 19 and under
- ☐ 20 - 29
- ☐ 30 - 39
- ☐ 40 - 49
- ☐ 50 +

2. Gender

- ☐ Male
- ☐ Female

3. Education

- ☐ Some high school
- ☐ Completed high school
- ☐ Some post-secondary education
- ☐ Completed undergraduate degree
- ☐ Some graduate or professional school
- ☐ Completed postgraduate degree

4. Computer Usage (hours per week):

- ☐ 0 - 10
- ☐ 10 - 20
- ☐ 20 - 30
- ☐ 30 - 40
- ☐ 40 - 50
- ☐ 50 +

Part 2

With respect to the visualization you worked with,

a) please indicate the extent to which you agree or disagree with the following statements:

SD = Strongly Disagree

D = Disagree

N = Neutral

A = Agree

SA = Strongly Agree

I found this visualization to be efficient for completing the tasks.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Navigating through the data was easy to do.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Locating coloured nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found this visualization to be frustrating.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Comparing topological distances between nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found it easy to get lost in this visualization.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Using two coloured focus boxes helped me to complete the task.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Being able to see compressed coloured nodes around the edges of the view made the task easier.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I enjoyed using this visualization.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA

b) What particular aspect(s) of this visualization did you *like*?

c) What particular aspect(s) of this visualization did you *dislike*?

d) Please use this space to describe/illustrate any alternative strategies (other than those you were shown at the beginning of the experiment) that you believe would have worked better for you.

e) Please use this space to make any other comments about the experiment or the visualization.

Thank you for your time!



THE UNIVERSITY OF BRITISH COLUMBIA

Experimental Questionnaire

Evaluation of Information Visualization Techniques

Interface # 2

Subject # _____

Part 1**1. Age Group**

- ☐ 19 and under
- ☐ 20 - 29
- ☐ 30 - 39
- ☐ 40 - 49
- ☐ 50 +

2. Gender

- ☐ Male
- ☐ Female

3. Education

- ☐ Some high school
- ☐ Completed high school
- ☐ Some post-secondary education
- ☐ Completed undergraduate degree
- ☐ Some graduate or professional school
- ☐ Completed postgraduate degree

4. Computer Usage (hours per week):

- ☐ 0 - 10
- ☐ 10 - 20
- ☐ 20 - 30
- ☐ 30 - 40
- ☐ 40 - 50
- ☐ 50 +

Part 2

With respect to the visualization you worked with,

a) please indicate the extent to which you agree or disagree with the following statements:

SD = Strongly Disagree

D = Disagree

N = Neutral

A = Agree

SA = Strongly Agree

I found this visualization to be efficient for completing the tasks.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Navigating through the data was easy to do.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Locating coloured nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found this visualization to be frustrating.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Comparing topological distances between nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found it easy to get lost in this visualization.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Using two coloured focus boxes helped me to complete the task.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
The coloured arcs made navigation easier.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I enjoyed using this visualization.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA

b) What particular aspect(s) of this visualization did you *like*?

c) What particular aspect(s) of this visualization did you *dislike*?

d) Please use this space to describe/illustrate any alternative strategies (other than those you were shown at the beginning of the experiment) that you believe would have worked better for you.

e) Please use this space to make any other comments about the experiment or the visualization.

Thank you for your time!



THE UNIVERSITY OF BRITISH COLUMBIA

Experimental Questionnaire

Evaluation of Information Visualization Techniques

Interface # 3

Subject # _____

Part 1**1. Age Group**

- ☐ 19 and under
- ☐ 20 - 29
- ☐ 30 - 39
- ☐ 40 - 49
- ☐ 50 +

2. Gender

- ☐ Male
- ☐ Female

3. Education

- ☐ Some high school
- ☐ Completed high school
- ☐ Some post-secondary education
- ☐ Completed undergraduate degree
- ☐ Some graduate or professional school
- ☐ Completed postgraduate degree

4. Computer Usage (hours per week):

- ☐ 0 - 10
- ☐ 10 - 20
- ☐ 20 - 30
- ☐ 30 - 40
- ☐ 40 - 50
- ☐ 50 +

Part 2

With respect to the visualization you worked with,

a) please indicate the extent to which you agree or disagree with the following statements:

SD = Strongly Disagree

D = Disagree

N = Neutral

A = Agree

SA = Strongly Agree

I found this visualization to be efficient for completing the tasks.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Navigating through the data was easy to do.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Locating coloured nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found this visualization to be frustrating.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Comparing topological distances between nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found it easy to get lost in this visualization.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
The presence of the smaller view made the task easier.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Using two coloured focus boxes helped me to complete the task.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Being able to see compressed coloured nodes around the edges of the view made the task easier.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I enjoyed using this visualization.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA

b) What particular aspect(s) of this visualization did you *like*?

c) What particular aspect(s) of this visualization did you *dislike*?

d) Please use this space to describe/illustrate any alternative strategies (other than those you were shown at the beginning of the experiment) that you believe would have worked better for you.

e) Please use this space to make any other comments about the experiment or the visualization.

Thank you for your time!



THE UNIVERSITY OF BRITISH COLUMBIA

Experimental Questionnaire

Evaluation of Information Visualization Techniques

Interface # 4

Subject # _____

Part 1**1. Age Group**

- ☐ 19 and under
- ☐ 20 - 29
- ☐ 30 - 39
- ☐ 40 - 49
- ☐ 50 +

2. Gender

- ☐ Male
- ☐ Female

3. Education

- ☐ Some high school
- ☐ Completed high school
- ☐ Some post-secondary education
- ☐ Completed undergraduate degree
- ☐ Some graduate or professional school
- ☐ Completed postgraduate degree

4. Computer Usage (hours per week):

- ☐ 0 - 10
- ☐ 10 - 20
- ☐ 20 - 30
- ☐ 30 - 40
- ☐ 40 - 50
- ☐ 50 +

Part 2

With respect to the visualization you worked with,

a) please indicate the extent to which you agree or disagree with the following statements:

SD = Strongly Disagree

D = Disagree

N = Neutral

A = Agree

SA = Strongly Agree

I found this visualization to be efficient for completing the tasks.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Navigating through the data was easy to do.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Locating coloured nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found this visualization to be frustrating.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Comparing topological distances between nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found it easy to get lost in this visualization.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
The presence of the smaller view made the task easier.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Using two coloured focus boxes helped me to complete the task.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
The coloured arcs made navigation easier.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I enjoyed using this visualization.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA

Appendix D

Study 2 Questionnaires



THE UNIVERSITY OF BRITISH COLUMBIA

Experimental Questionnaire

Evaluation of Information Visualization Techniques

Interface # 1

Subject # _____

Part 1**1. Age Group**

- ☐ 19 and under
- ☐ 20 - 29
- ☐ 30 - 39
- ☐ 40 - 49
- ☐ 50 +

2. Gender

- ☐ Male
- ☐ Female

3. Education

- ☐ Some high school
- ☐ Completed high school
- ☐ Some post-secondary education
- ☐ Completed undergraduate degree
- ☐ Some graduate or professional school
- ☐ Completed postgraduate degree

4. Computer Usage (hours per week):

- ☐ 0 - 10
- ☐ 10 - 20
- ☐ 20 - 30
- ☐ 30 - 40
- ☐ 40 - 50
- ☐ 50 +

Part 2

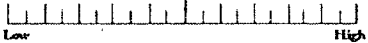
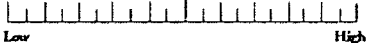
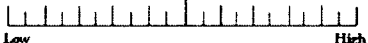
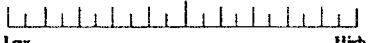
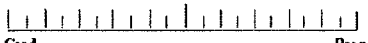
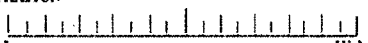
With respect to the visualization you worked with,

a) please indicate the extent to which you agree or disagree with the following statements:

SD = Strongly Disagree
D = Disagree
N = Neutral
A = Agree
SA = Strongly Agree

I found this visualization to be efficient for completing the tasks.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Navigating through the data was easy to do.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Locating coloured nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found this visualization to be frustrating.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Comparing topological distances between nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found it easy to get lost.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Being able to see compressed coloured nodes around the edges of the view made the task easier.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I enjoyed using this visualization.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA

b) With respect to the visualization you worked with, please answer the following questions by marking an 'X' along the scale beside the corresponding question.

How much mental and perceptual activity was required to complete the task (e.g., looking, searching, thinking, deciding, calculating, remembering, etc.)?	MENTAL DEMAND 
How much physical activity was required to complete the task (e.g., moving the mouse, dragging, clicking, pressing keys, etc.)?	PHYSICAL DEMAND 
How much time pressure did you feel due to the rate or pace at which the tasks or task elements occurred?	TEMPORAL DEMAND 
How hard did you have to work (mentally and physically) to accomplish your level of performance?	EFFORT 
How successful do you think you were in accomplishing the goals of the task set by the experimenter (or yourself)?	PERFORMANCE 
How insecure, discouraged, irritated, stressed and annoyed versus secure, gratified, content, relaxed and complacent did you feel during the task?	FRUSTRATION 

c) What particular aspect(s) of this visualization did you *like*?

d) What particular aspect(s) of this visualization did you *dislike*?

e) Please use this space to describe/illustrate any alternative strategies (other than those you were shown at the beginning of the experiment) that you believe would have worked better for you.

f) Please use this space to make any other comments about the experiment or the visualization.

Thank you for your time!



THE UNIVERSITY OF BRITISH COLUMBIA

Experimental Questionnaire

Evaluation of Information Visualization Techniques

Interface # 2

Subject # _____

Part 1**1. Age Group**

- ☐ 19 and under
- ☐ 20 - 29
- ☐ 30 - 39
- ☐ 40 - 49
- ☐ 50 +

2. Gender

- ☐ Male
- ☐ Female

3. Education

- ☐ Some high school
- ☐ Completed high school
- ☐ Some post-secondary education
- ☐ Completed undergraduate degree
- ☐ Some graduate or professional school
- ☐ Completed postgraduate degree

4. Computer Usage (hours per week):

- ☐ 0 - 10
- ☐ 10 - 20
- ☐ 20 - 30
- ☐ 30 - 40
- ☐ 40 - 50
- ☐ 50 +

Part 2

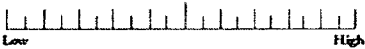
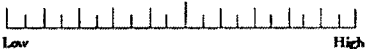
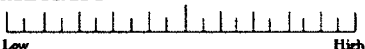
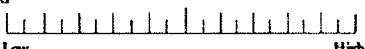
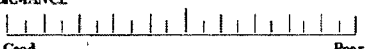
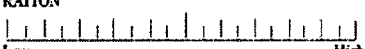
With respect to the visualization you worked with,

a) please indicate the extent to which you agree or disagree with the following statements:

SD = Strongly Disagree
 D = Disagree
 N = Neutral
 A = Agree
 SA = Strongly Agree

I found this visualization to be efficient for completing the tasks.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Navigating through the data was easy to do.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Locating coloured nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found this visualization to be frustrating.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Comparing topological distances between nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found it easy to get lost.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
The coloured arcs made navigation easier.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I enjoyed using this visualization.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA

b) With respect to the visualization you worked with, please answer the following questions by marking an 'X' along the scale beside the corresponding question.

How much mental and perceptual activity was required to complete the task (e.g., looking, searching, thinking, deciding, calculating, remembering, etc.)?	MENTAL DEMAND 
How much physical activity was required to complete the task (e.g., moving the mouse, dragging, clicking, pressing keys, etc.)?	PHYSICAL DEMAND 
How much time pressure did you feel due to the rate or pace at which the tasks or task elements occurred?	TEMPORAL DEMAND 
How hard did you have to work (mentally and physically) to accomplish your level of performance?	EFFORT 
How successful do you think you were in accomplishing the goals of the task set by the experimenter (or yourself)?	PERFORMANCE 
How insecure, discouraged, irritated, stressed and annoyed versus secure, gratified, content, relaxed and complacent did you feel during the task?	FRUSTRATION 

c) What particular aspect(s) of this visualization did you *like*?

d) What particular aspect(s) of this visualization did you *dislike*?

e) Please use this space to describe/illustrate any alternative strategies (other than those you were shown at the beginning of the experiment) that you believe would have worked better for you.

f) Please use this space to make any other comments about the experiment or the visualization.

Thank you for your time!



THE UNIVERSITY OF BRITISH COLUMBIA

Experimental Questionnaire

Evaluation of Information Visualization Techniques

Interface # 3

Subject # _____

Part 1**1. Age Group**

- ☐ 19 and under
- ☐ 20 - 29
- ☐ 30 - 39
- ☐ 40 - 49
- ☐ 50 +

2. Gender

- ☐ Male
- ☐ Female

3. Education

- ☐ Some high school
- ☐ Completed high school
- ☐ Some post-secondary education
- ☐ Completed undergraduate degree
- ☐ Some graduate or professional school
- ☐ Completed postgraduate degree

4. Computer Usage (hours per week):

- ☐ 0 - 10
- ☐ 10 - 20
- ☐ 20 - 30
- ☐ 30 - 40
- ☐ 40 - 50
- ☐ 50 +

Part 2

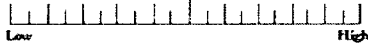
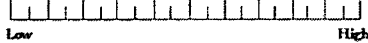
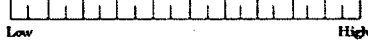
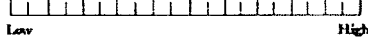
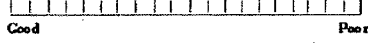
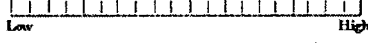
With respect to the visualization you worked with,

a) please indicate the extent to which you agree or disagree with the following statements:

SD = Strongly Disagree
 D = Disagree
 N = Neutral
 A = Agree
 SA = Strongly Agree

I found this visualization to be efficient for completing the tasks.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Navigating through the data was easy to do.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Locating coloured nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found this visualization to be frustrating.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Comparing topological distances between nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found it easy to get lost.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
The presence of the smaller view made the task easier.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Being able to see compressed coloured nodes around the edges of the view made the task easier.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I enjoyed using this visualization.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA

b) With respect to the visualization you worked with, please answer the following questions by marking an 'X' along the scale beside the corresponding question.

How much mental and perceptual activity was required to complete the task (e.g., looking, searching, thinking, deciding, calculating, remembering, etc.)?	MENTAL DEMAND 
How much physical activity was required to complete the task (e.g., moving the mouse, dragging, clicking, pressing keys, etc.)?	PHYSICAL DEMAND 
How much time pressure did you feel due to the rate or pace at which the tasks or task elements occurred?	TEMPORAL DEMAND 
How hard did you have to work (mentally and physically) to accomplish your level of performance?	EFFORT 
How successful do you think you were in accomplishing the goals of the task set by the experimenter (or yourself)?	PERFORMANCE 
How insecure, discouraged, irritated, stressed and annoyed versus secure, gratified, content, relaxed and complacent did you feel during the task?	FRUSTRATION 

c) What particular aspect(s) of this visualization did you *like*?

d) What particular aspect(s) of this visualization did you *dislike*?

e) Please use this space to describe/illustrate any alternative strategies (other than those you were shown at the beginning of the experiment) that you believe would have worked better for you.

f) Please use this space to make any other comments about the experiment or the visualization.

Thank you for your time!



THE UNIVERSITY OF BRITISH COLUMBIA

Experimental Questionnaire

Evaluation of Information Visualization Techniques

Interface # 4

Subject # _____

Part 1**1. Age Group**

- ☐ 19 and under
- ☐ 20 - 29
- ☐ 30 - 39
- ☐ 40 - 49
- ☐ 50 +

2. Gender

- ☐ Male
- ☐ Female

3. Education

- ☐ Some high school
- ☐ Completed high school
- ☐ Some post-secondary education
- ☐ Completed undergraduate degree
- ☐ Some graduate or professional school
- ☐ Completed postgraduate degree

4. Computer Usage (hours per week):

- ☐ 0 - 10
- ☐ 10 - 20
- ☐ 20 - 30
- ☐ 30 - 40
- ☐ 40 - 50
- ☐ 50 +

Part 2

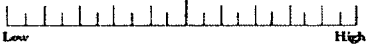
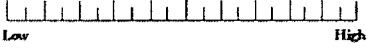
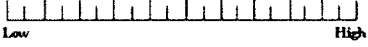
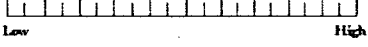
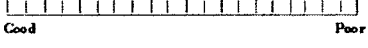
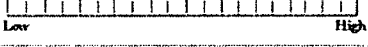
With respect to the visualization you worked with,

a) please indicate the extent to which you agree or disagree with the following statements:

SD = Strongly Disagree
 D = Disagree
 N = Neutral
 A = Agree
 SA = Strongly Agree

I found this visualization to be efficient for completing the tasks.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Navigating through the data was easy to do.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Locating coloured nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found this visualization to be frustrating.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Comparing topological distances between nodes was easy.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I found it easy to get lost.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
The presence of the smaller view made the task easier.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
The coloured arcs made navigation easier.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
I enjoyed using this visualization.	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA

b) With respect to the visualization you worked with, please answer the following questions by marking an 'X' along the scale beside the corresponding question.

How much mental and perceptual activity was required to complete the task (e.g., looking, searching, thinking, deciding, calculating, remembering, etc.)?	MENTAL DEMAND 
How much physical activity was required to complete the task (e.g., moving the mouse, dragging, clicking, pressing keys, etc.)?	PHYSICAL DEMAND 
How much time pressure did you feel due to the rate or pace at which the tasks or task elements occurred?	TEMPORAL DEMAND 
How hard did you have to work (mentally and physically) to accomplish your level of performance?	EFFORT 
How successful do you think you were in accomplishing the goals of the task set by the experimenter (or yourself)?	PERFORMANCE 
How insecure, discouraged, irritated, stressed and annoyed versus secure, gratified, content, relaxed and complacent did you feel during the task?	FRUSTRATION 

c) What particular aspect(s) of this visualization did you *like*?

d) What particular aspect(s) of this visualization did you *dislike*?

e) Please use this space to describe/illustrate any alternative strategies (other than those you were shown at the beginning of the experiment) that you believe would have worked better for you.

f) Please use this space to make any other comments about the experiment or the visualization.

Thank you for your time!