An Evaluation of Overview+Detail and Focus+Context Interfaces with Guaranteed Visibility

First Author Name [Blank if Blind Review]
Affiliation [Blank if Blind Review]
Address [Blank if Blind Review]
e-mail address [Blank if Blind Review]
Optional phone number [Blank if Blind Review]

Second Author Name [Blank if Blind Review]
Affiliation [Blank if Blind Review]
Address [Blank if Blind Review]
e-mail address [Blank if Blind Review]
Optional phone number [Blank if Blind Review]

ABSTRACT

The datasets grow in size and complexity, new techniques become necessary to efficiently explore and navigate their structure. One domain which can greatly benefit from these techniques is phylogentics, which is the study of the evolutionary relationships between species. Phylogenetic analysis relies heavily on visual inspection and topological analysis of large trees, yet the small number of techniques that have been developed to support these tasks lack evaluation. Based on an ethnographic investigation into the challenges and tasks faced by phylogenetic researchers, we ran a study to examine the navigation patterns and compare performance between Pan and Zoom and Focus+Context interfaces, with and without an overview, on navigating large and complex trees. Each interface implements guaranteed visibility – a recent innovation in Information Visualization which guarantees areas of interest to be visible at all times, independent of navigation action Results indicate that navigation strategy can play a significant role in user performance. Additionally, P&Z navigation appeared significantly faster than RSN, regardless of the presence or absence of an overview. Additionally, we found that RSN was reported to be significantly more mentally demanding than P&Z techniques.

Author Keywords

Evaluation, Information Visualization, Focus+Context, Overview+Detail, Guaranteed Visibility

ACM Classification Keywords

H.5.2 User Interfaces, *Evaluation/methodology*. I.3.6 Methodology and Techniques, *Interaction techniques*

INTRODUCTION

As information grows in size and complexity, conventional methods of information display, such as documents and spreadsheets, are unable to effectively represent the data in a form that facilitates rapid exploration and analysis. Information Visualization aims to overcome the drawbacks of these conventional methods by creating interactive visual representations that can be used to rapidly explore and navigate large and complex sets of data. raditional navigation techniques, such as pan and zoom (P&Z) have been shown to be effective for exploring relatively small data sets [Hornbaek et al., 2002], and in conjunction with an overview window, overcome several of the drawbacks documented in the literature; however, it remains uncertain how well these techniques scale for large, complex data sets.

Focus+Context [F+C] navigation techniques, which commonly use distortion, aim to overcome several of the drawbacks of traditional P&Z methods; however, the effect of introducing distortion has been shown to reveal mixed performance results. One recent F+C navigation technique is rubber-sheet navigation [RSN], which uses the metaphor of a stretching and squishing a rubber-sheet with its borders tacked down, to explore a data space. The effect of adding an overview to F+C techniques, such as RSN, has not been previously investigated.

A recent innovation in Information Visualization is the concept of Guaranteed Visibility [Munzner et al., 2003], which is the property that regions of interest are guaranteed to be visible, regardless of the navigation actions taken by a user. While we believe that this is a desirable property for user interfaces, the effect of adding guaranteed visibility to either P&Z or F+C interfaces has not been previously investigated.

We ran a study to examine navigation patterns and compare performance between P&Z and F+C interfaces with and without an overview, which implement the property of guaranteed visibility. Prior to our study, we conducted an investigation with phylogenetic researchers to understand the requirements and challenges of phylogenetic tree exploration and analysis. Our study is the first to compare different visualization techniques that all share the same underlying method of visual presentation. Our comparison uses a large, real-world dataset, in conjunction with a generalized version of a complex topological comparison task, which is based on informal discussions with phylogenetic researchers. Results indicate that P&Z navigation appeared significantly faster than RSN, regardless of the presence or absence of an overview, when comparing topological characteristics in large trees. Additionally, we found that RSN was reported to be significantly more mentally demanding than P&Z techniques.

Our contributions are as follows:

1. The first evaluation of the effect of navigation techniques (P&Z and F+C) with the presence or absence of an overview as orthogonal factors.
2. The first evaluation of the effect of adding an overview to F+C.
3. The first evaluation of the effect of adding an overview to a visualization with guaranteed visibility.

Our paper is organized as follows. We start by defining several terms, based on information in the literature and our own investigation. We then summarize the related work, with emphases on the evaluation of P&Z and F+C techniques. We then describe our experimental methodology, followed by a detailed analysis of our results. A discussion of our results, along with our conclusion and goals for future work completes this paper.

TERMINOLOGY

A number of information visualization techniques have been proposed to facilitate interaction with large**,** complex sets of data. Many of these techniques use the concept of **focus and context regions** to enable users to selectively identify parts of a dataset that interest them and explore them in detail, while maintaining a mental model of the dataset as a whole. The term **focus** is generally taken to mean a high resolution region of interest to the user, while **context** is understood to be a low resolution region which has been compressed to occupy less screen space, but is still visually salient [Spence and Apperley, 1982]**.** However, to date no rigorous definition of these terms has been presented or agreed upon.

The term **Focus+Context** has been generally used for techniques that visualize focus and context regions in an integrated way. Many Focus+Context techniques use distortion-oriented techniques (see survey in [Leung and Apperley, 1994]) such as fisheye views [Furnas, 1982] to dynamically expand focus regions while contracting surrounding context regions. Other approaches that have been proposed for integrating focus and context regions include **aggregating** context regions into glyphs [Plaisant et al., 2002; Card and Nation, 2002], **pointing** to off-screen context regions [Bederson and Hollan, 1994; Baudisch and Rozenholtz, 2003], and **space filling** to present as much context information as possible within a given number of pixels [Munzner et al., 2003].

An alternative approach to Focus+Context is to provide context in a separate view, commonly referred to as an **overview**. Visualizations that use overviews are usually described as **Overview+Detail** visualizations [Hornbaek et al., 2002]. Many of these rely on a particular type of overview, namely a smaller separate window placed in a corner of the screen, which is referred to as a **radar view** [Smith et al., 1998]. The extent of the main view of the visualization is usually represented in the overview by a moveable **field of view box** [Hornbaek et al., 2002].

Navigation can be defined as a means of performing a set of actions afforded by a visualization. Commonly used techniques for navigation in large datasets include **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 resolution (see survey in [Hornbaek et al., 2002]). The two techniques have been frequently used in combination. While panning and zooming have been used in some distortion-oriented Focus+Context visualizations [reference?], in this paper we use the term Pan and Zoom to denote the combination of non-distorting panning and zooming, as distinguished from distortion-oriented navigation techniques.

While Pan and Zoom may be sufficient for navigating smaller datasets, the drawbacks of this approach become apparent with increases in dataset size. These include inefficient navigation patterns [Hornbaek et al., 2002] and the phenomenon of **desert fog,** where the user loses context due to zooming or panning into an empty area of the dataset [Jul and Furnas, 1998]. A navigation technique that attempts to avoid some of these issues in Focus+Context visualizations is **rubber sheet navigation** [Sarkar et al., 1993], which allows users to select and enlarge or compress different areas of the dataset as though it were laid out on a rubber sheet with its edges tacked down. Accordion drawing [Munzner et al., 2003, Slack et al., 2005] uses rubber sheet navigation to provide **guaranteed visibility**, that is, to ensure that marked areas always remain visible regardless of navigation actions taken by the user

RELATED WORK

While a number of interaction techniques have been proposed in the literature to facilitate the exploration of large datasets, the evaluation of these techniques has been inconclusive.

The literature on Overview+Detail visualizations reveals mixed results. Studies have shown that navigation is more efficient since users are able to navigate in both the overview window and the detail window [Beard and Walker, 1990]. The extra contextual information provided by the overview may also help users maintain orientation within a large data space [Plaisant et al., 1995], and provides them with global information to help them made decisions about future navigation actions [Hornbaek and Frokjaer, 2001] and a feeling of control [Shneiderman 1998]. However, other studies have found that the addition of an overview means that a user must divide their attention between two separate windows to accomplish their task. This extra cognitive load has been shown to strain memory and increases the time for visual search [Card et al. 1999]. Moreover, the addition of an overview window means that less screen space may be devoted to the detail window [Hornbaek et al, 2002].

The literature on Focus+Context visualizations also reveals mixed results. Studies have shown that distortion approaches, such as fisheye techniques, are beneficial for tasks such as steering navigation [Gutwin and Skopik, 2003], hierarchical network navigation [Schaffer et al, 1996], web browsing [Baudisch et al, 2004], spatial collaboration [Schafer and Bowman, 2003], and calendar use [Bederson et al, 2004]. However, other studies found that distortion can negatively impact performance for tasks such as interactive layout [Gutwin and Fedak, 2004a], location recall [Skopik and Gutwin, 2003], and visual scanning [Kobsa, 2004]. It has also been shown that the performance of Focus+Context visualizations can be significantly affected by parameters of the distortion mechanism, such as extent, magnification level, and shape of the distortion [Gutwin and Fedak, 2004a].

There have been few empirical evaluations of navigation methods for Focus+Context and Overview+Detail visualizations. [Johnson, 1995] found panning by dragging the background of the visualization to be superior to other panning methods. A number of studies have shown zooming interfaces to perform better than those providing only basic scrolling or panning navigation (e.g. [Kaptelinin, 1995; Gutwin and Fedak, 2004b]). However, results of evaluations of zooming interfaces have been generally inconclusive due to differences in comparison interfaces, size and type of the datasets used, and implementations of zooming [Hornbaek et al., 2002].

Guaranteed visibility appears to be a desirable property given previously mentioned issues that may arise when navigating a large dataset with traditional panning and zooming techniques. A recent study [Baudisch et al., 2004] explored the benefits of a form of guaranteed visibility in the context of reading short electronic documents. However, further evaluation is needed to determine whether these benefits extend to larger datasets in other domains.

Recent work has explored the benefits and drawbacks of different techniques for visualizing large trees. Kobsa [Kobsa, 2004] performed a comparative experiment with five tree visualization systems, with Windows Explorer as a benchmark. Kobsa's evaluation used a large hierarchical dataset based on a subset of taxonomy of items on eBay, which consisted of 5 levels and a total of 5,799 nodes. Tasks were generated by the experimenters and informed from an early version of the tasks detailed in the InfoVis 2003 contest. The results of his study revealed significant differences between the different tree visualization tools with respect to accuracy, performance, and user satisfaction. These differences were attributed to the fact that each tree visualization system used a different visualization paradigm, and additionally some had interface problems or were missing functionality required to complete the tasks. SpaceTree [Plaisant et al, 2002] was evaluated in a controlled experiment against a hyperbolic tree browser and Windows Explorer. Rather than pit the interfaces against each other, as in [Kobsa, 2004], the experimenters' goal was to understand what features appeared to help users perform certain tasks. The SpaceTree evaluation used a large tree dataset of more than 7,000 nodes from the CHI '97 BrowseOff. Tasks were generated by the experimenters and included questions concerning tree topology. The results of the study were mixed, revealing that SpaceTree performed significantly faster for some classes of topological tasks, such as listing all the ancestors of a given node, while it also performed significantly slower for other classes of topological tasks, such as finding 3 nodes that have more than 10 direct descedents.

While these evaluations demonstrate the utility of tree visualizations, they also illustrate the inherent difficulty of comparing the performance of different tools which each have their own unique method of visually presenting the underlying data. Our evaluation aims to overcome this difficulty by focusing on techniques that share the same underlying method of visual presentation, and only differ in their form of navigation.

EXPERIMENAL METHOD

The goal of our study is to compare the performance of RSN and P&Z interfaces with and without an overview.

Interfaces

We developed 4 different visualization interfaces to evaluate the performance of RSN and P&Z, with and without an overview. Each interface is built on top of a de-featured version of TreeJuxtaposer [Munzner, 2003], which provided us with a stable platform for evaluating different visualization techniques, while preserving a consistent visual presentation of the data and interaction controls across the different conditions, a characteristic that previous evaluations in the literature have been lacking.

Each interface is designed to support a user in exploring and navigating a large tree. While each interface supports fluid and interactive navigation, the method of navigation, and thus corresponding strategy for completing tasks, differed based on the interface. For each interface, navigation was controlled using a 3 button mouse, with rectilinear zooming mapped to the left mouse button, panning mapped to the right mouse button, and zoom out mapped to the middle mouse button. Each interface supported a reset function, which was mapped to the 'R' key on a standard keyboard.

The following is a description of the visualization techniques and interaction metaphors used in each interface:

RSN with GV - Enables users to navigate the data set using pan and zoom actions, which use the metaphor of stretching and squishing a rubber sheet with its boarders tacked down. Focus areas are selected for zoom by dragging out a rectilinear box, the contents of which are stretched to fill the red focus area (see Figure 1). Panning is accomplished via horizontal and vertical drag motions, which allow the user to finely tune their focus selection. Zooming out is accomplished by dragging out a rectilinear focus region that is larger than the red focus area, the contents of which are squished to fill the red focus area.

Colored nodes are guaranteed to be visible at all times, even if they are squished to sub-pixel size due to navigation actions.

P&Z with GV – Enables users to navigate the data set using pan and zoom actions, which follow the traditional PZ metaphor. Focus areas are selected for zoom by dragging out a rectilinear box, the contents of which zoom to fill the view completely. Panning is accomplished via horizontal and vertical drag motions, which allow the user to finely tune their focus selection. Zooming out is accomplished via vertical drag motions, which allow the user to gradually zoom out.

Colored nodes are guaranteed to be visible at all times. In the event that a colored node is off-screen due to navigation action, a colored arc will appear at the border of the screen, indicating the direction and the distance from the current view to the colored node (see Figure 2). The arc is part of a circular ring that surrounds any colored node that is currently off-screen. Once a colored node is visible on-screen, the arc disappears. This technique is inspired by Halo [Baudisch and Rozenholtz, 2003], which uses arcs to visualize off-screen locations on small devices, such as personal digital assistants.

RSN with GV + Overview – Same as *RSN*, with the addition of an overview window with moveable field of view box. In addition to the navigation and visualization techniques described in previously, focus regions could also be selected by dragging out a rectilinear box in the overview window, the contents of which are stretched to fill the red focus area in the detail window. Panning is accomplished in the overview window via a series of mouse drags, which allowed the user to position their field-of-view box anywhere within the data space.

In addition to the guaranteed visibility provided within the larger detail view as previously described, colored nodes are also guaranteed to be visible within the overview window (see Figure 3).

PZ with GV + Overview - Same as *PZ with GV*, with the addition of an overview window with moveable field of view box. In addition to the navigation mechanisms described in Condition 2, focus regions could also be selected by dragging out a rectilinear box in the overview window, the contents of which zoom to full the larger detail view completely. Panning is accomplished in the overview window via a series of mouse drags, which allows the user to position their field-of-view box anywhere within the data space.

In addition to the guaranteed visibility provided within the larger detail view, as previously described, colored nodes are also guaranteed to be visible within the overview window (see Figure 4).

Dataset

The dataset used in our evaluation is a binary tree consisting of almost 6,000 nodes. The tree represents the phylogenetic relationships between species within the kingdom Animalia, and can be downloaded from the Olduvai project website.[1]

Task

Based on discussions with phylogenetic researchers at the University of British Columbia, University of Texas at Austin, and the 2004 Evolution conference, we developed a generalized version of a topological comparison task that required no specialized knowledge of phylogenetics or trees. Our task, as illustrated in figure 5, is a complex task composed of several low-level tasks, such as find, identify, and compare. The task required a subject to perform several navigation actions to be successfully completed, which was important since we wanted to exercise the different navigation mechanisms for each interface.

The task required subjects to compare the topological distances between 2 sets of colored nodes in a large tree, and answer which of the distances was smaller. The task was always the same, though the location of the colored nodes in the tree varied from question to question. Questions were assessed prior to the experiment to ensure that each was isomorphic in difficulty. This was done by ensuring that topological distances always fell in a range of 7 to 10, and that the complexity of the local topology surrounding each colored node was consistent from question to question.

Measures

We collected both quantitative and qualitative data from subjects. Quantitative data was logged as subjects interacted with the interface and completed tasks. We logged performance data, including task completion time,

[1] http://olduvai.sourceforge.net/tj/data.shtml

number of navigation actions such as pans and zooms performed with the interface, and number of times subjects reset the interface to its initial state, as well as errors.

Qualitative data was collected by having subjects complete a questionnaire and participating in an exit interview following the experiment. The questionnaire asked subjects to report on usability issues with the interface they used, including perceived ease of use, efficiency for completing the task, and overall enjoyment. Subjects were also asked to report on the perceived usefulness of the specific visualization techniques they used, such as the value of seeing compressed nodes in the RSN interfaces, the value of seeing arcs in the P&Z interfaces, and the value of having an overview if they used an overview interface. Additionally, we collected responses on the NASA-TLX workload questionnaire [Hart and Staveland, 1988] to gauge the workload placed on the subject by the interface. Finally, subjects were asked to comment on what they liked and disliked about the interface and to provide any alternative strategies they believe may have worked better for them other than the ones they were trained on.

Apparatus

Our study was conducted on 2 systems running Windows XP with Pentium 4 processors, 2.0 GB RAM, and Nvidia GeForce2 video cards. The experimental software, including the visualization conditions, was fully automated and was coded in Java 1.4.2 and OpenGL.

Pilot

We initially ran a pilot study to explore navigation strategies, context levels, and the benefits of guaranteed visibility. Results indicated that performance was significantly better for subjects who used guaranteed visibility to complete the task. Additionally, results showed that having multiple foci, in combination with minimal training, led to highly variable strategies, and in general made the interfaces more complex to use.

The lessons learned from our pilot were used to inform the design of our controlled study. These included providing guaranteed visibility of marked areas across all interfaces, restricting the number of foci to one, and providing subjects with a comprehensive training session which instructed them on the best strategies as observed from our pilot.

Training

Subjects were instructed on the use of the different navigation and interaction techniques afforded by each interface. Based on the results of our pilot experiment, subjects were provided with and strongly encouraged to use strategies to complete questions.

For each interface, the strategy was first demonstrated by the experimenter, who then asked the subject to repeat the strategy. All strategies started with dragging out a long thin selection area along the horizontal path between the nodes in question. Following the discovery of one of the two topological distances, subjects were instructed to reset the interface and continue using the same strategy. The strategies provided for each condition were as follows:

RSN with GV – Dragging out a long thin selection area along a horizontal path between the two nodes in question has the effect of stretching the dataset along the horizontal axis (see Figure 1). Using this strategy, subjects were shown how to count nodes which were now visually salient. Following this step, long thin vertical selection areas could be dragged out to expand areas that were compressed vertically.

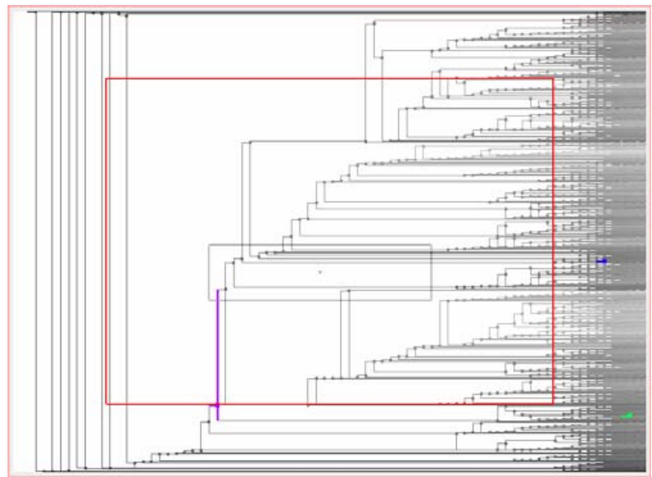

Figure 1: RSN with GV condition. The use of the "select a long thin horizontal area" strategy is shown.

PZ with GV – Dragging out a long thin selection area along a horizontal path between the two nodes in question has the effect of zooming the contents of the focus box to fill the entire view. Using this strategy, subjects were shown how to count nodes which were now visually salient. Following this step, the subjects could slowly zoom out and add nodes as they appeared along the path up the tree (see Figure 2).

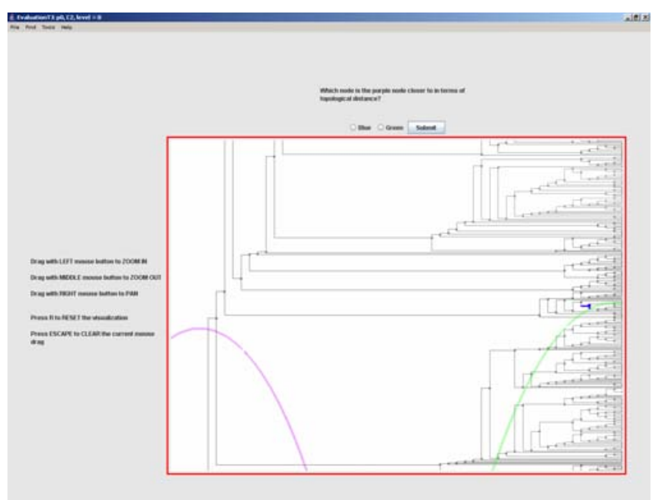

Figure 2: PZ with GV condition. Results of the "zoom out to reveal the path up the tree" strategy are shown.

RSN with GV + Overview – Dragging out a long thin selection area along a horizontal path between the two nodes in question has the effect of stretching the dataset along the horizontal axis. Using this strategy, subjects were shown how to count nodes which were now visually salient (see Figure 3). Following this step, long thin vertical selection areas could be dragged out to expand areas that were compressed vertically. Subjects were also instructed to use both the overview and detail views for navigation and counting nodes, but were not explicitly told that they had to use either the overview or detail views.

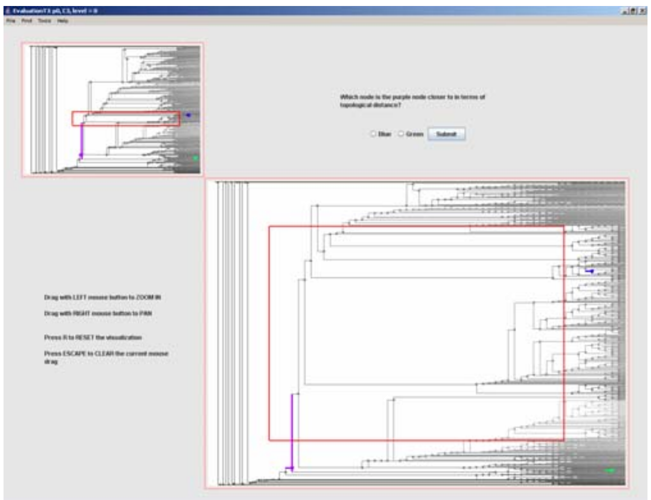

Figure 3: RSN with GV + Overview condition. Results of the "select a long thin horizontal area" strategy are shown.

PZ with GV + Overview – Dragging out a long thin selection area along a horizontal path between the two nodes in question has the effect of zooming the contents of the focus box to fill the entire view. Using this strategy, subjects were shown how to count nodes which were now visually salient (see Figure 4). Subjects were also instructed to use both the overview and detail views for navigation and counting nodes, but were not explicitly told that they had to use either the overview or detail views.

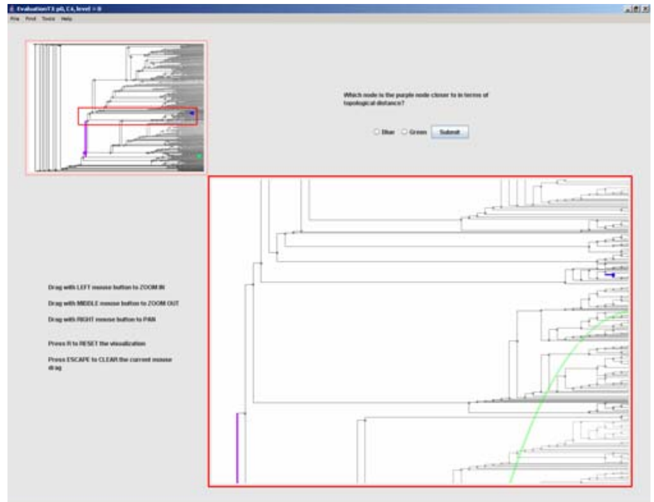

Figure 4: PZ with GV + Overview condition. Results of the "select a long thin horizontal area" strategy are shown.

CONTROLLED STUDY

The goal of this study was to compare the efficiency of the four visualization conditions described earlier in order to determine whether the different types of navigation (PZ and RSN) and the presence or absence of an overview in these visualizations had an impact on their performance and on user preferences.

Design

The design of this study was a between-subjects 2x2 (navigation x presence of overview) design. Participants were randomly assigned to each of the 4 conditions using a round-robin method. A between-subjects design was chosen due to the need for extensive training in order for participants to effectively use each visualization, as elicited in the pilot experiment. For each of the 4 conditions, the level of context that resulted in the best performance for that condition in piloting was used throughout the experiment.

Subjects

Subjects were recruited through advertisements posted throughout campus and through an online experiment reservation system. All subjects were screened right-handedness, colorblindness, and experience with a computer and 3 button mouse. We collected standard demographic data, such as sex, age, and level of education. In total, we tested 40 subjects, consisting of X males and X females. Age levels ranged from X to X, while level of education ranged from X to X. Subjects were compensated $15 for their time.

A total of 44 subjects were used for the experiment. Of these, two were unable to follow the training instructions, and two others followed the instructions but committed 4 or more errors (an error rate of greater than 10%). Since the goal of the experiment was to measure successful performance on the task, these subjects were treated as outliers for the purpose of data analysis.

Procedure

The experiment was designed to fit in a single 1.5-hour session. Participants were first trained on the use of the visualization condition they had been assigned to and provided with strategies that would help them complete the task efficiently. Participants were then given a training session of 5 questions. For each of the first 2 questions, the experimenter demonstrated solving the question using the previously described strategies and then asked the participant to repeat this solution. For the last 3 questions of the session, the participant solved the questions on their own, with the experimented making suggestions for improving the participant's efficiency as needed. At the end of the training session, participants were given a 1 minute break.

Participants were then presented with 7 blocks of questions, each containing 5 questions. The question set as a whole was identical for each participant, and the individual question blocks were verified to be isomorphic in difficulty in pilot experiments. The sequence of blocks was randomly generated for each participant.

At the end of the experiment, participants filled out a feedback questionnaire asking them to evaluate the visualization they had used on a variety of scales, including the NASA-TLX workload assessment scales [Hart and Staveland, 1989], as well as to describe what they liked and disliked about the visualization and any alternative strategies they may have used during the experiment. Brief informal interviews were conducted with some of the participants based on their questionnaire responses.

Hypotheses

Our hypotheses were as follows:

H1. RSN conditions (both with and without overview) perform better than P&Z conditions (both with and without overview).

H2. Conditions without an overview (both RSN and P&Z) perform no worse than conditions with an overview (both RSN and P&Z).

RESULTS

Performance

To determine whether there were differences in performance between the visualization conditions, two-way repeated-measures ANOVA's (navigation x presence of overview, with block as the repeated-measures variable) and post-hoc pairwise comparisons using the Bonferroni adjustment method were performed. Prior to these comparisons, outlier data lying more than 3 standard deviations from the means of each navigation x presence of overview x block cell were removed from the analysis. 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 [Landauer, 1997]. To interpret this value, .02 is a small effect size, .06 is medium, and .14 is large.

Completion times

The overall results for completion times are shown graphically in Figure 5. In order to determine whether subject performance was reaching a plateau, one-way repeated measures ANOVA's with block as the repeated-measures variable were performed for each of the conditions. The results of this analysis are shown in Table 1. (The data for Condition 1 was non-spherical, hence we used the Greenhouse-Geissler adjustment.) Post-hoc pairwise comparisons showed no significant differences between values within blocks 5, 6, and 7 in any of the conditions, p > .5, indicating that performance reached a plateau by the end of the experiment in all conditions.

Figure 5: Mean completion times by block.

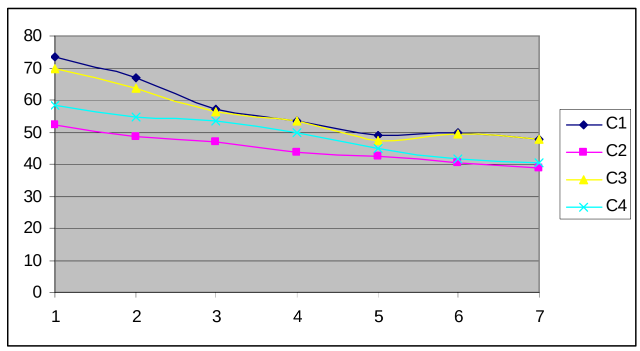

Table 1: All F values have degrees of freedom (6, 54) except Condition 1 (1.687, 15.187)

Condition	F	p	Partial η^2
1	11.032	<.003	.551
2	16.938	<.001	.653
3	15.159	<.001	.627
4	10.196	<.001	.531

In order to compare the impact of navigation and overview at both the novice level and expert level , a 2 navigation by 2 overview by 2 blocks ANOVA was run (blocks 1 and blocks 7 representing novice and expert behaviour, respectively). There was a significant effect of navigation ($F(1,36) = 13.745$, $p < .005$, $\eta^2 = .276$). P&Z conditions were significantly faster than RSN conditions. There was no significant effect of presence of overview (F(1,36) =

.202, $p > .7$, $\eta^2 = .002$) or interaction between navigation and overview.

There was a significant effect of block ($F(3.174,114.26) = 91.71$, $p < .001$, $\eta^2 = .717$) and a borderline significant interaction between block and navigation ($F(3.176,114.35) = 3.647$, $p < 0.07$, $\eta^2 = .064$) showing that the RSN conditions had greater performance improvements between blocks 1 and 7 than the P&Z conditions.

In order to determine whether there were significant differences in performance at the start of the experiment on conditions within each navigation cell, a one-way ANOVA of condition and post-hoc pairwise comparisons were performed on the data from block 1. There was a significant effect of condition ($F(3,36) = 4.068$, $p < .05$, $\eta^2 = .253$) and a significant pairwise difference between conditions 1 and 2, p<.05, but no other pairwise differences, showing that there were no differences in terms of initial performance between either the RSN conditions or the P&Z conditions.

Navigation actions

There was a borderline significant effect of navigation on the number of navigation actions ($F(1,36) = 3.9$, $p < .06$, $\eta^2 = .098$). The P&Z conditions had significantly fewer navigation actions than RSN conditions, consistent with the result reported for completion times above. There was also a significant effect of block ($F(6,216) = 4.395$, $p < .001$, $\eta^2 = .109$). There was no significant main effect of presence of overview for this measure, as well as no interaction effects.

Resets

There was a significant effect of navigation on the number of resets ($F(1,36) = 4.912$, $p < .05$, $\eta^2 = .12$) The P&Z conditions had significantly fewer resets than RSN conditions, consistent with the result reported for completion times above. There were no significant main effects of presence of overview or block or interaction effects for this measure.

Error rate

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 effect of navigation or presence of overview on error rate.

Self-reported measures

To determine whether there were differences in self-reported measures between the visualization conditions, two-way ANOVA's (navigation x presence of overview) were performed for each of the measures reported on the questionnaire.

There was a significant effect of navigation on the TLX mental demand measure ($F(1,36) = 4.214$, $p < .05$, $\eta^2 = .105$). The RSN conditions were reported to be more mentally demanding than the P&Z conditions, confirming the performance results presented above.

There was also a significant effect of presence of overview on the TLX physical demand measure ($F(1,36) = 6.215$, $p < .02$, $\eta^2 = .147$). Conditions without overview were reported to be more physically demanding than those with overview. There were no significant effects on any of the other measures.

Summary of results

We summarize our results according to the experimental hypotheses:

H1. P&Z conditions performed better than RSN conditions in terms of completion times, navigation actions, and resets.

H2. Conditions without an overview did not perform worse than conditions with an overview in terms of completion times, navigation actions, or resets.

FOLLOW-UP EXPERIMENT

As a result of the discrepancy between the performance results, which indicated no significant difference between the conditions with and without an overview, and the effect of presence of overview on the self-reported measure of physical effort, it was decided to run a follow-up experiment to set apart the influence of GV in the detailed view and the presence of overview. Piloting prior to the controlled experiment had determined that RSN with overview and GV in the detailed view was significantly faster than the same visualization without GV in the detailed view. The follow-up experiment therefore compared P&Z with overview and GV in the detailed view to the same visualization without GV in the detailed view. The hypothesis of the follow-up study was as follows:

H3. P&Z with overview and GV in the detailed view is faster than P&Z with overview without GV in the detailed view.

(Discuss results of follow-up study.)

DISCUSSION

Discuss our findings on the impact of overviews for PZ and RSN interfaces

Discuss possible improvements for RSN, such as minimizing global distortion when navigating.

Do we want to make a distinction between integrated and separate visibility now?

Discuss how results can be generalized to other tasks and other kinds of visualizations.

CONCLUSIONS AND FUTURE WORK

FW – Continue our evaluation with different types of tasks

ACKNOWLEDGMENTS

REFERENCES

1. Baudisch, P. and Rosenholtz, R. (2003). Halo: A Technique for Visualizing Off-Screen Locations, *Proc. ACM CHI 2003*, 481-488.
2. 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. *Proc. AVI 2004*, 133-140.
3. Beard, D. B. and Walker, J. Q. (1990). Navigational techniques to improve the display of large two-dimensional spaces. *Behav. Info. Techn.* , 9(6), 451-466.
4. Bederson, B. B. and Hollan, J. (1994). Pad++: A zooming graphical interface for exploring alternate universe physics. *Proc. ACM UIST 1994,* 17-26.
5. Bederson, B. B., Clamage, A., Czerwinski, M. P., & Robertson, G. G. (2004). DateLens: A Fisheye Calendar Interface for PDAs. *ACM ToCHI*, 11(1), 90-119.
6. Card, S. K., Mackinlay, J. D., and Shneiderman, B. (1999). *Readings in Information Visualization: Using Vision to Think,* Morgan Kaufmann Publishers, San Francisco, California, 1999.
7. Card, S. K., Nation, D. (2002). Degree-of-interest trees: a component of an attention-reactive user interface. *Proc. AVI 2002,* 231-246.
8. Furnas, G. W. (1982). The FISHEYE view: A new look at structured files. Bell Laboratories Technical Memorandum, #82-11221-22, Oct 18, 1982.
9. Gutwin, C., and Skopik, A. (2003). Fisheye Views are Good for Large Steering Tasks, *Proc. ACM CHI 2003*, 201-208.
10. Gutwin, C., and Fedak, C. (2004) A comparison of fisheye lenses for interactive layout tasks, *Proc. Graphics Interface 2004*, 213 -220.
11. Gutwin, C., and Fedak, C. (2004) Interacting with Big Interfaces on Small Screens: A Comparison of Fisheye, Zoom, and Panning Techniques, *Proc. Graphics Interface 2004*, 145 -152.
12. Hart, S. G., & Staveland, L. E. (1988). Development of NASA-TLX (Task Load Index): results of empirical and theoretical research. In P. Hancock & N. Meshkati (Eds.), *Advances in psychology: human mental workload* (pp. 139-183). Amsterdam: Elsevier Science.
13. Hornbaek, K., and Frokjaer, E. (2001). Reading of Electronic Documents: The Usability of Linear, Fisheye, and Overview+Detail Interfaces, *Proc. ACM CHI 2001*, 293-300.
14. Hornbæk, K., Bederson, B., & Plaisant, C. (2002). Navigation Patterns and Usability of Zoomable User Interfaces with and without an Overview. *ACM ToCHI,* 9 (4), 362-389.
15. Johnson, J. (1995). A Comparison of User Interfaces for Panning on a Touch-Controlled Display. *Proc. ACM CHI 1995*, 218-225.
16. Kaptelinin, V. (1995). A Comparison of Four Navigation Techniques in a 2D Browsing Task, *Proc. ACM CHI 1995*, 282-283.
17. Landauer, T. (1997). Chapter 9: Behavioral research methods in human-computer interaction. In M.G. Helander, T.K. Landauer, and P.V. Prabhu, (Eds), *Handbook of human computer interaction* (2nd ed.) (pp. 203-227). Amsterdam: Elsevier Science.
18. Kobsa, A. (2004). User Experiments with Tree Visualization Systems. *Proc. IEEE InfoVis 2004*, 9-16.
19. Leung, Y. and Apperley, M. (1994). A review and taxonomy of distortion-oriented presentation techniques. *ACM ToCHI,* 1(2), 126-160.
20. Munzner, T., Guimbretiere, F., Tasiran, S., Zhang, L., Zhou, Y. (2003). TreeJuxtaposer: Scalable Tree Comparison using Focus plus Context with Guaranteed Visibility. *Proc. ACM SIGGRAPH 2003,* 453-462.
21. Plaisant, C., Carr, D., and Shneiderman, B. (1995). Image browsers: Taxonomy, guidelines, and informal specifications. *IEEE Software* 12(2), 21–32.
22. Plaisant, C. Grosjean, J., and Bederson, B. (2002). Spacetree: Supporting exploration in large node link tree, design evolution and empirical evaluation. *Proc. IEEE InfoVis 2002*, 57-64.
23. Sarkar, M., Snibbe, S., Tversky, O., and Reiss, S. (1993). Stretching the Rubber Sheet: A Metaphor for Viewing Large Layouts on Small Screens. *Proc. ACM UIST 1993*, 81-91.
24. Schafer, W. and Bowman, D. (2003). A Comparison of Traditional and Fisheye Radar View Techniques for Spatial Collaboration. *Proc. Graphics Interface 2003*, 39-46.
25. Schaffer, D., Zuo, Z., Greenberg, S., Bartram, L., Dill, J., Dubs, S., and Roseman, M., Navigating Clustered Networks through Fisheye and Full-Zoom Methods, *ACM ToCHI*, 3(2), 1996, 162-188.
26. Shneiderman, B. (1998). *Designing the User Interface*, 3rd ed., Addison-Wesley, Reading, Mass.
27. Slack, J., Hildebrand, K., and Munzner, T. (2005). PRISAD: A Partitioned Rendering Infrastructure for Scalable Accordion Drawing. To appear in *Proc. IEEE InfoVis 2005.*
28. Smith, R., Hixon, R., and Horan, B. (1998). Supporting Flexible Roles in a Shared Space. *Proc. ACM CSCW 1998*, 197-206.
29. Spence, R. and Apperley, M. (1982). Data base navigation: An office environment for the professional. *Behaviour and Information Technology*, 1(l), 43–54.

The columns on the last page should be of approximately equal length.