Information Visualization in Desktop Virtual Reality

David Modjeska[1] and John Waterworth[2]

[1] Department of Computer Science, University of Toronto, Toronto, ON M5R 2S2, Canada
modjeska@dgp.utoronto.ca
[2] Department of Informatics, Umeå University, 901 87 Umeå, Sweden
jwworth@informatik.umu.se

Abstract. Desktop virtual reality (VR) offers a powerful environment for visualizing structure in large information sets. In well-designed virtual worlds, users can employ skills from wayfinding in the real world. This paper reports the development and testing of a series of prototype VR worlds, designed to support navigation during information visualization and retrieval. Results indicated that users' sense of presence, ease of use and preference were affected by degree of spatial cueing, but target-hunting performance was not. The virtual world with strongest spatial cues caused poorer user performance than a comparable hypertext interface, but was rated as more enjoyable. The final VR prototype was no worse than hypertext in user performance; the prototype shows the potential for navigable VR that is engaging and useful for everyday information exploration.

1 Introduction

Two recent trends in human-computer interfaces are striking. First, as the volume and complexity of the information sphere grows, more advanced techniques are needed to visualize and manage it. Such approaches, generally known as information visualization, can be seen as moving downward from abstraction to representation. Second, as the power and sophistication of computer technology increases, its ability to simulate the world also increases. Such approaches, generally known as VR, can be seen as moving upward from concreteness to dematerialization. These 2 trends are starting to meet in the domain of information visualization using desktop VR. This meeting enables significant improvements in the everyday accessibility of large information sets.

The best opportunities to use VR for large-scale visualization tasks now lie with hierarchical data, for 2 reasons. First, the world's largest information structure is the Web. Web sites often have a roughly hierarchical structure. While software engines have greatly helped searching, browsing has received less support. Second, several well-known research prototypes of information landscapes for unstructured data exist. These prototypes are often constructed using statistical analysis techniques [2, 5].

A research project at Umeå University in Sweden is investigating the visualization of hierarchical information in VR, with a focus on navigational issues. Several prototypes have been designed, developed, and empirically tested. Experimental tasks mix searching and browsing, with an emphasis on user performance, domain learning, and overall satisfaction. Initial results are satisfactory, and participants report the research prototypes to be engaging, memorable, and suitable for real work. At this point, the VR visualization technique used looks promising for large information environments, e.g., corporate or university intranets, which allow unified design and management.

The paper is organized as follows. Related research is reviewed first, followed by a description of the prototype and then discussions of 3 experiments.

2 Related Research

Though there exists substantial research on information visualization, only techniques most relevant to this paper will be discussed. A suite of 3D tools from XEROX PARC, the Information Visualizer [1], includes the *cone tree* [13] and the *hyperbolic browser* [7] to present hierarchical information. A goal of these tools is to shift part of the burden of conscious cognitive processing to the perceptual system. Animated visual transitions, for example, aid users in tracking visualization changes. A good description of these techniques is "focus+context", because they let the user focus on interesting information, while retaining context in the visual background. It is unclear how these techniques scale up for very large data sets, for which an information landscape may be more navigable and memorable. Such landscapes have been explored in several research projects that lay out unstructured data using statistical analysis [2, 5]. Possibilities for landscapes with structured data, however, are largely unexplored.

Research on navigation in VR began with urban design studies of the physical world. Work on legibility (imageability) showed that resident efficiency and enjoyment is enhanced by a design with landmarks, paths, districts, nodes, and edges in a strong hierarchy [8]. Later work analyzed wayfinding – the conceptual part of navigation – into 3 iterative stages: mental mapping, route planning, and plan execution [11]. Lynch's work mainly concerned strategies to improve mental mapping. Recent research concludes that wayfinding design principles from the real world often apply to large virtual environments [3]. In this work, global structure was recommended to support good wayfinding. What remains is to use these principles for structured information visualization, as this research was conducted with naturalistic landscapes.

Design principles for information landscapes appeared in a summary of wayfinding research with application to virtual worlds [14]. In related work, a layout algorithm was proposed [6], based on hyperbolic visualizations [7] and cone trees [13]. More abstract 3D textual environments, including Apple Computer's "HotSauce" [12], were used to represent hierarchical information.

3 Experimental Prototype

A large design space exists for building a VR landscape to visualize information. In particular, a discrepancy between semantic and spatial structure must be reconciled [9]. Our initial hypothesis was that varying the strength of spatial cueing would significantly affect search performance and environmental perception. More specifically, we expected that more spatial cueing would result in better understanding of the environment's spatial structure and information objects. I.e., spatial cues would aid navigation and be more enjoyable. Accordingly, a series of 3 virtual worlds was designed, taking key points on a continuum between text- and object-based representations of information structure. Each design applied the idea of Information Islands [15] to visualize a filtered subset of a Web index. The 3 worlds were maximally isomorphic in features (e.g., locations, sizes, and labels). The data set was chosen for interest to participants and the research community. The set included about 1500 items over 7 levels of hierarchy, which allowed for rich detail and computational tractability.

The first design was most naturalistic, with colored objects and grayscale labels. This design, the *Day World*, had strong color and lighting cues (Figure 1). Virtual objects were laid out to maximize imageability by Lynch's guidelines: islands, cities, neighborhoods, and buildings ("districts"); mountains and rivers ("edges"); rivers, roads, and bridges ("paths"); and geometric objects ("landmarks" or "nodes"). Objects at each level were clustered around landmarks, according to sibling groups in the data hierarchy. Color assignment grouped buildings in neighborhoods with common palettes, while ground and water objects had naturalistic color. Text labeled each object. To avoid information overload, the distance from which a label was visible varied inversely with the label's depth in the data hierarchy. In general, the best point from which to survey a region was its center. Avoiding indirect navigation, users could fly freely throughout the worlds, or use a UI shortcut to fly directly to desired objects. Navigationally indirect features, e.g., elevators and hallways, were avoided.

Fig.1. A view of the Day World (Experiment 1)

Midway between object- and text-based representations of information structure, the *Dusk World* resembled the Day World with changes in coloring (Figure 2). In this world, objects were desaturated 90% (as in twilight) and semi-transparent. Labels, however, had bright, saturated colors, and were grouped by sibling relationships in the data hierarchy. The design was intended to offer the benefit of optical illusions, where perception varies between modes, here object- and text-based.

At the textual end of the design continuum was the *Night World* (Figure 3), which is similar to Dusk World, but without virtual objects. Here, the user could move in an abstract information space without absolute location or distance. This design was inspired by recent prototypes, including HotSauce and others [12]. The Night World lacked directional lighting. It had only brightly colored text on a black background.

The experimental VR environments were constructed in VRML 2.0, based on a filtered subset of WWW structure. The VRML was generated by C++ code, requiring about 7 months of development effort. The world-generation algorithm takes as input

Fig. 2. A view of the Dusk World (Experiment 1)

Fig. 3. A view of the Night World (Experiment 2)

ASCII data in a hierarchical format: processing time is about 5 seconds per world, and each VRML file is about a megabyte in size. This desktop VR can be viewed with the CosmoPlayer plug-in for Web browsers. Special input/output devices are not needed. A graphics card with a high frame-rate, though, facilitates interactivity.

4 Experiment 1

4.1 Methodology

To compare prototype designs on user efficiency and perception, an experiment was run with 12 participants (8 male, 4 female). The experiment was a two-way (3 x 3) within-subjects design. World design was the first condition: Day, Dusk, or Night World. Ordinal position of exposure was the second condition: first, second, and third. World design and exposure order were fully counterbalanced. Six dependent variables were used in the experiment. Five of these were self-reported on a scale from 1 to 5: world size, exposure duration, sense of presence, ease of use, and overall preference. The sixth variable was search performance: found/total targets.

Experimental sessions were conducted on an SGI Onyx 2 computer in a VR lab over a 2-week period. Participants were first trained in the user interface in a sample world for 15 minutes. They then explored each experimental world for 20 minutes. In a world, participants began by roaming freely for 1-2 minutes. They then performed a "scavenger hunt": they received a series of paper cards, each showing the context and name of an information item (building floor) in the virtual world. Participants were to find as many of the 10 targets as possible in the world, in the available time. They could skip difficult targets. The hunt tested search performance, as well focusing attention on VR. After each world, participants completed a short questionnaire, and a summary one after all worlds. Participants' comments were also logged.

4.2 Results

The effect of world design and exposure order was investigated through a series of one-way ANOVAs. The measures significantly affected by world design were as follows. (Means are scaled between 0.00 and 1.00.)

	Means			
	Day	*Dusk*	*Night*	*Significance*
Presence:	.708	.417	.375	$F(2,33) = 5.60$, $p < .01$
Ease:	.604	.292	.458	$F(2,33) = 4.27$, $p < .02$
Preference:	.729	.292	.438	$F(2,33) = 10.2$, $p < .001$

In general, sense of presence, ease of use, and overall preference were highest for the Day World, and comparably lower for the Dusk World and the Night World. World design had no significant effect on other measures, not even performance. (The mean performance for all participants and sessions was .43 out of 1.0.)

Exposure order had no significant effect on experimental measures.

To investigate direct relationships between measures, we ran a series of Pearson Correlation tests. Ease of use, overall preference, and sense of presence were all significantly related to each other ($p < .05$). Similarly, ease of use, overall preference, and performance were all significantly related ($p < .05$). No significant relationship was found between presence and performance, though.

To determine if there was a learning effect during each session, we analyzed performance timing for the course of a session. No significant changes were observed.

All participants used computers frequently, but game and Web index usage varied. Analysis showed no significant effect of either factor on experimental measures.

4.3 Discussion of Experiment 1

Most strikingly, this experiment showed no clear relationship between sense of presence and performance. Performance was not significantly affected by world design or exposure order, but sense of presence was directly affected by world design. Ease of use and overall preference correlated significantly with both presence and performance, yet neither of these latter variables correlated with each other. It appears that the proposed VR designs are equivalent in usability, counter to experimental hypothesis. By way of explanation, perhaps the interfaces were not usable enough to reveal differences for the given task. In any case, participants felt more engaged with the Day World, which follows the experimental hypothesis. This sense of engagement has potential benefits for concentration, motivation, and long-term learning.

The lack of learning effect on performance was puzzling. Perhaps tasks were too easy or too hard to change a participant's performance during the experiment. Users may already have had the required skills from experience with Web and office software, which they could transfer to information exploration in VR. Alternatively, the UIs may have been so new as to prevent skill improvement during sessions.

The next experiment varied the representation and the input data set independently. It was thus possible to determine the differential impacts of world design and underlying data domain on task performance.

5 Experiment 2

Given the results of the first experiment, it was necessary in a follow-up: (1) to establish a baseline of performance and perception by comparing hypertext with VR representations, since hypertext is the dominant way of presenting information on the Web; and (2) to investigate the effect of different (but comparable) data sets, in order to clarify the relationship between world design and domain data.

5.1 Methodology

The second experiment was carried out using the apparatus of the first one, but in a private office, to reduce distractions. The hardware was a personal computer with a high-end graphics card. For this experiment, a simple HTML hypertext was created using an extension of the world-generation software in Experiment 1.

The experiment was a 2 x 2 within-subjects design. 16 participants were run, all undergraduates (13 male, 3 female) in Informatics. The first independent variable was world design: VR versus hypertext. For VR, the Day World was used, since it was liked best in Experiment 1. The second independent variable was data set: work versus leisure subsets of data in Experiment 1. Earlier procedures were slightly changed. First, participants explored 4 worlds in 15-minute sessions. After each, participants were asked to classify 5 data items like those in the VR. This task assessed learning of domain structure. Answers were scored for accuracy at the top 2 levels of data hierarchy (measures Classification 1 and 2). Also, a question on user "preference" was split into "efficiency" and "enjoyment," as these differed in Experiment 1. Finally, the number of search targets grew to 25, as hypertext efficiency was expected to be high.

5.2 Results

The effect of world design, ordinal position, repeated exposure and elapsed time on dependent variables was examined through several one-way ANOVAs. The measures significantly affected by world design were as follows. (Means are scaled between 0.00 and 1.00):

	Means		
	<u>*Htext.*</u>	<u>*VR*</u>	<u>*Significance*</u>
Performance:	.400	.100	$F(2,61) = 82.1, p < .000$
Classification 1:	.694	.544	$F(2,61) = 4.65, p < .048$
Classification 2:	.450	.331	$F(2,61) = 4.55, p < .050$
Efficiency:	.800	.412	$F(2,61) = 25.8, p < .000$
Ease:	.800	.388	$F(2,61) = 35.3, p < .000$
Enjoyment:	*.562*	.750	$F(2,61) = 4.31, p < .010$

Ordinal position affected mainly Classification 1: this measure increased significantly from .488 (out of 1.00) to .750 over a participant's 4 virtual world sessions ($F = 3.56, p < .044$). Another effect on classification 1 was caused by repeated exploration of a data set. This measure rose significantly from .530 (out of 1.00) on first exploration of a data set to .700 on second exploration ($F = 5.77, p < .03$).

To investigate relationships between measures, we ran a series of Pearson Correlation tests, each of which was significant at the .05 level:

- Performance: classification 1 and 2, ease, and efficiency
- Ease: efficiency, enjoyment (inversely)
- Efficiency: enjoyment (inversely)

Finally, experimental results show signs of improvement in target hunting and classification over the course of an experimental session. For example, in VR conditions only, target hunting improved from .071 (out of 1.00) during the first session to .138 during the second ($F = 17.1, p < .001$).

5.3 Discussion of Experiment 2

In considering the results of Experiment 2, the superiority of hypertext for most measures is evident: target-hunting performance, first- and second-level classification, rated efficiency and ease of use. Hypertext is evidently a more mature technique for information visualization and access, which is reflected in better software tools and prior user experience. At the same time, as in Experiment 1, participants' preferences (enjoyment here) didn't match the most efficient user interface. In fact, participants enjoyed VR more, reflecting greater user engagement and motivation. This difference suggests, at least, that the structural benefits of hypertext should be included in further development of VR visualization for hierarchical information.

Similarly, the usability advantages of hypertext could be considered as potential standards for VR development. Accordingly, users' suggestions for improvements to the research prototype, particularly those inspired by comparison with hypertext, can facilitate the iterative design process. The resulting prototype could be a better tool for task completion, and a more solid and sensitive research platform for future work.

In contrast to Experiment 1, we found evidence of learning in Experiment 2. Participants performed better in target hunting during the second VR exposure, which suggests a learning threshold for the new UI. Hypertext, by contrast, showed no such learning effect, presumably because participants already knew the technique from Web browsing. Top-level classification improved for participants with both interfaces, and this potential was explored in the next experiment, as discussed below.

6 Experiment 3

6.1 Methodology

After Experiment 2, usability and performance problems in the prototype were fixed, largely based on participants' comments. First, the more structured CityScape layout algorithm was used, to take better advantage of perceptual cues in VR [6]. This algorithm was less naturalistic, but more efficient and aesthetically satisfying. Second, many usability problems were fixed, refining the prototype. Third, non-information bearing virtual objects were removed, and the rest geometrically simplified, to improve graphics performance at the cost of some decorative features (e.g., trees).

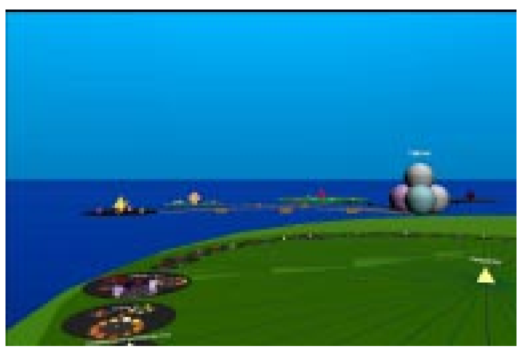

Fig. 4. A view of the "efficient" Day World (Experiment 3)

The main purpose of Experiment 3 was to validate a new research prototype (Figure 5), as well as to compare with hypertext. Accordingly, a variation of Experiment 2 was run with 8 participants, students and instructors (7 male, 1 female); this variation compared "naturalistic" and "efficient" versions of VR. As each participant was exposed to only two virtual worlds, sessions were lengthened to 20 minutes apiece.

6.2 Results and Discussion of Experiment 3

The effect of world design on dependent variables was investigated through a series of one-way ANOVAs. The measures significantly affected were as follows. (Means are scaled between 0.00 and 1.00.)

	Means		
	Naturalistic	*Efficient*	*Significance*
Ease:	.416	.750	$F(2,61) = 12.5, p < .010$
Efficiency:	.416	.750	$F(2,61) = 12.5, p < .010$
Classification 2:	.525	.750	$F(2,61) = 14.5, p < .007$

To examine relationships between measures, we ran several Pearson Correlation tests. Ease of use, efficiency, and enjoyment were all significantly related ($p < .05$).

No other significant effects of world design or exposure order were observed.

Participants reported doing less work in the efficient design than in the naturalistic one, but performance did not reflect this. Participants' second-level classification was better in the efficient design, however; this reflected its effectiveness in conveying domain structure. That top-level classification showed no significant difference is not surprising, because the naturalistic and efficient designs are similar at the top level.

Participant responses varied on enjoyment of VR designs. Some preferred the more organic quality of the naturalistic design. Still, analysis showed rated efficiency, ease of use, and enjoyment were significantly correlated.

In comparison with results for hypertext in Experiment 2, the largest difference was in target hunting. We can compare 8 efficient VR sessions in Experiment 3 with 8 hypertext sessions in Experiment 2, which used the same methodology. (Participant populations were similar, and the new graphics card in Experiment 3 did not affect hypertext.) Only the first 15 minutes of sessions from Experiment 3 should be included. A between-participants analysis (one-way ANOVA) revealed no significant difference in performance, though the hypertext mean of .400 (out of 1.00) was higher than the efficient VR mean of .250.

Observation suggests further benefits of the efficient VR design. Participants often returned to distant virtual locations by recognizing visual structures. This occurred when labels were too distant to read. User recognition of visual features suggests VR's power for memorability and wayfinding. Moreover, users expressed excitement and surprise at the new design, with its structured layout algorithm, careful geometric abstraction, common landscape and urban metaphors, and controlled variation in color palette and landmark form. It struck many as a viable, alternative representation of Web structure, while offering some of the experiential engagement of games.

7 Conclusions

To investigate the combination of information visualization and large-scale desktop VR, a series of research prototypes was developed and empirically tested. The final prototype presents a balance of algorithmic structure and naturalistic representation. Moreover, the prototype seems to engage the human perceptual apparatus to lessen the burden of conscious information processing. Results showed the final design to be usable, efficient, and liked by research participants.

Future research is needed to assess this approach's promise and limitations. First, more browsing experimental tasks might shed light on VR exploration. Another issue is the value of handcrafted versus algorithmic techniques for representation. E.g., how might 2- or 3D icons improve user navigation and learning, and how should they be incorporated into a VR information visualization? Given the novelty of VR, another line of inquiry concerns the degree of user immersion. What effect would different equipment have on the issues in this paper? Finally, it is worth considering possibilities for very large, dynamic structures such as the Web, and for social navigation in shared VR. It is hoped that this paper will encourage research in these directions.

8 Acknowledgements

For work on prototypes, we thank K. Moijson, D. Steeman, W. Frost and R. Kroll. A. Lund gave implementation advice, and S. Zhai commented on a draft of this paper.

9 References

1. Card, S.K., Robertson, G.G., Mackinlay, J.D. The Information Visualizer, an Information Workspace, in *Proceedings of CHI '91*, ACM Press, 181-188.
2. Chalmers, M. Using a landscape metaphor to represent a corpus of documents, in *Proceedings of COSIT '93*, Springer-Verlag, 377-390.
3. Darken, R.P., Sibert, J.L. Wayfinding Strategies and Behaviors in Large Virtual Worlds, in *Proceedings of CHI '96*, ACM Press, 142-149.
4. Dieberger, A. Providing Object-based Navigation for the World Wide Web, in *Proceedings of COSIT '95*, Springer-Verlag, 93-106.
5. Hemmje, M., Kunkel, C., Willett, A. LyberWorld -- A Visualization User Interface Supporting Fulltext Retrieval, in *Proceedings of SIGIR '94*, ACM Press, 249-258.
6. Keshkin, C., Vogelmann, V. Effective Visualization of Hierarchical Graphs with the Cityscape Metaphor, in *Proceedings of the Workshop on "New paradigms in information visualization and manipulation*" (1997), ACM Press, 52-57.
7. Lamping, J., Rao, R., Pirolli, P. A Focus+Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies, in *Proceedings of CHI '95*, ACM Press, 401-408.
8. Lynch, K. *Image of the City*, The. The MIT Press, Cambridge MA, 1960.
9. McKnight, C., Dillon, A., Richardson, J. *Hypertext in Context*. Cambridge University Press, Cambridge England, 1991.
10. Modjeska, D. Parallel Worlds: Spatial and Textual Representations of Information Structure, in *Proceedings of the Workshop on Personalized and Social Navigation in Information Space*, Swedish Institute of Computer Science, 130-137.
11. Passini, R. *Wayfinding in Architecture*. Van Nostrand Reinhold, New York NY, 1984.
12. Rennison, E. Galaxy of News: An Approach to Visualizing and Understanding Expansive News Landscapes, in *Proceedings of UIST '94*, ACM Press, 3-12.
13. Robertson, G.G., Mackinlay, J.D., Card, S.K. Cone Trees: Animated 3D Visualizations of Hierarchical Information, in *Proceedings of CHI '91*, ACM Press, 189-194.
14. Vinson, N.G. Design Guidelines for Landmarks to Support Navigation in Virtual Environments, in *Proceedings of CHI '99*, ACM Press, 278-285.
15. Waterworth, J.A. A pattern of islands: exploring public information space in a private vehicle, in Brusilovsky, P., Kommers, P., and Streitz, N. (eds.), Multimedia, Hypermedia and Virtual Reality: Models, Systems, and Applications. Springer-Verlag, Berlin, 1996.