Effects of Desktop 3D World Design on User Navigation and Search Performance

David Modjeska
Department of Computer Science
University of Toronto
modjeska@dgp.utoronto.ca

John Waterworth
Department of Informatics
Umeå University
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 subjectively preferred naturalistic environments over abstract ones, but that users objectively searched better in environments that had rigorous hierarchical structure, and support for both overview and detail. In comparison with a hypertext interface, the final virtual world design elicited more enjoyment, without worsening search performance significantly, when appropriate training time was allowed. This prototype shows the potential value of 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 techniques, generally known as instances of 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 abstraction. These two 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.

Information visualization addresses the problem of representing various data types for the end-user, so that the data can more easily be understood, managed, and communicated. A variety of techniques has been developed to handle both structured and unstructured data, especially for scientific and industrial tasks. The hallmark of such techniques is a mapping from a source data domain to a destination visual domain, where task-dependent objects and relationships become apparent and available. A secondary goal is to offload part of the burden of conscious information processing to the human perceptual system [1]. While often powerful, existing techniques generally place the user in a third-person perspective with regard to data, or perhaps in a first-person perspective in a limited visual space. Richer possibilities for exploiting human navigational knowledge are open for research and development.

VR uses the techniques of computer graphics to present a model world to the senses through a variety of computer media, ranging from full-body immersion to traditional desktop display. The applications of VR to date have primarily been in engineering and architectural design (CAD), communication (virtual communities), and entertainment (computer games). All of these applications are based more or less directly on real-world phenomena, yet none of them explores in depth the potential of VR as a medium to represent abstract information. Recent research on wayfinding in VR suggests that human navigation skills transfer effectively from real into electronic worlds, in many cases [2]. If this is so, an opportunity exists to use VR for large-scale information visualization tasks. Such usage would have considerable benefits for accessing and communicating large information sets. While proposals and guidelines for such virtual worlds exist [3, 4], few (if any) tested implementations have been described in the research literature.

The best opportunities to use VR for large-scale visualization now lie with hierarchical data, for two 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 [5, 6].

At Umeå University and the University of Toronto, a research project 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. Participants report the research prototypes to be engaging, memorable, and suitable for real work. At this point, the visualization technique 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 discussion of designs and studies, and finally general conclusions.

2. Related Research

Though there exists substantial research on information visualization, only a few of the techniques most relevant to this paper will be discussed. (An overview is provided in [7].) XEROX PARC has developed a suite of 3D tools, the Information Visualizer [1], which includes the *cone tree* [8] and the *hyperbolic browser* [9] 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 [5, 6]. 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 [10]. 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 [2]. 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 [3]. In related work, a layout algorithm was proposed [12], based on hyperbolic visualizations [9] and cone trees [8]. More abstract 3D textual environments, including Apple Computer's "HotSauce" [13], were used to represent hierarchical information.

3. Study 1

3.1. Design

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 [14]. 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 information items and spatial structure. Accordingly, a set of 3 virtual worlds was designed: the worlds had different levels of visual intensity for virtual objects, but consistent textual labels and spatial structure [15]. Each design applied the idea of Information Islands [4] to visualize a filtered subset of a Web index. (Cf. [16].) 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 a 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 (Fig. 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. (A palette of 300 colors was provided by graduate students in industrial design at Umeå University.) Each object had a text label. 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 jump directly to objects of interest.

Partway between object- and text-based representations of information structure, the *Dusk World* resembled the Day World with significant changes in coloring (Fig. 2). In this world, objects were desaturated 90% (as in twilight) and semi-transparent. Labels, however, had bright, saturated colors, which were grouped by sibling relationships in the data hierarchy. The design was intended to support shifts between perceptual modes, here object- and text-based.

At the textual end of the design continuum was the *Night World* (Fig. 3), which is essentially the Dusk World 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 such as HotSauce and others [13]. The Night World lacked directional lighting: it showed 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 ASCII data in a hierarchical format: processing time is about 5 seconds per world, and each VRML file is about 3 megabytes in size. Such 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 user interactivity.

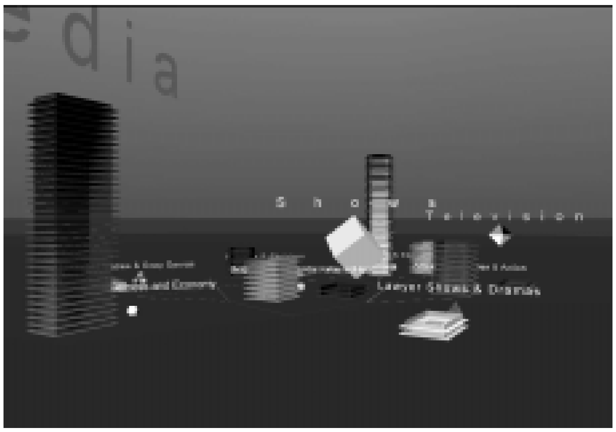

Fig. 1. A view of the Day World (Study 1)

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

Fig. 3. A view of the Night World (Study 1)

Fig. 4. A view of the hypertext design (Study 2)

3.2. Methodology

To compare initial designs on user efficiency and perception, a study was run with 12 participants. The study was a 1-way within-subjects design with 3 levels of the factor, world design. Exposure order was fully counterbalanced. Six dependent variables were used in the study. Five of them were self-reported: world size, exposure duration, sense of presence, ease of use, and enjoyment. The sixth variable was search performance.

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. In each experimental world, participants then explored freely for a couple of minutes. They then performed a 20-minute “scavenger hunt”: a participant received a series of paper cards, each showing the context and name of an information item in the virtual world. Participants were to find as many of the 10 targets as possible in the available time. A participant could skip difficult targets. The hunt tested search performance, as well serving to focus attention on the model worlds. After each world, participants completed a short questionnaire, as well as a summary questionnaire after all scavenger hunts. Participants' comments were also logged.

3.3. Brief Results

In general, self-reported sense of presence ($F = 5.60$, $p < .01$), ease of use ($F = 4.27$, $p < .02$), and enjoyment ($F = 10.2$, $p < .001$) 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 on performance. (The overall mean performance was 4.3 out of 10 targets.) Exposure order had no significant effect on experimental measures.

Ease of use, enjoyment, and sense of presence were all significantly correlated with each other ($p < .05$). Similarly, ease of use, enjoyment, and performance were all significantly correlated ($p < .05$). However, no significant correlation was found between presence and performance.

There was no apparent learning effect during the study sessions. All participants used computers often, but game and Web index usage varied. Analysis showed no significant effect of either factor on experimental measures. (Full results for Studies 1-3 are available in [17].)

3.4. Discussion of Study 1

Most strikingly, this study showed no clear relationship between sense of presence and performance. Performance was not significantly affected by world design or exposure order, but self-reported sense of presence was directly affected by world design. Ease of use and enjoyment correlated significantly with both presence and performance, yet neither of these latter variables correlated with each other. It appears that the proposed VR designs were equivalent in usability. Participants apparently felt more engaged with the Day World, which has potential benefits for concentration, motivation, and long-term learning.

The next study 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.

4. Study 2

Given the results of the first study, it was necessary in a follow-up study: (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.

4.1. Design and Methodology

The second study was carried out using the apparatus of the first one, but with a high-end personal computer. For this study, a cleanly-designed HTML hypertext was implemented using the world-generation software of Study 1 (Fig. 4).

The study was a 2 x 2 fully within-subjects design with 16 participants. The first independent variable was world design: VR (Day World) vs. hypertext. The second independent variable was data set: work vs. leisure subsets of data in Study 1. Participants explored four worlds in 15-minute blocks. After each block, participants were asked to classify (on paper) five new hunt targets. This task assessed learning of domain structure. Also, a questionnaire item about user "preference" was split into "efficiency" and "enjoyment," as these differed in Study 1. Finally, the number of hunt targets was increased to 25 in anticipation of efficient hypertext searching.

4.2. Brief Results

In general, performance ($F = 82.1$, $p < .001$), target classification ($F = 8.644$, $p < .006$), efficiency ($F = 25.8$, $p < .001$), and ease of use ($F = 35.3$, $p < .001$) were all higher for the hypertext interface, and comparably lower for the VR interface. Enjoyment ($F = 4.31$, $p < .010$), however, was generally lower for the hypertext interface and comparably higher for the VR interface.

These measures were significantly correlated ($p < .05$): performance with target classification, ease of use, and efficiency; ease of use with efficiency and enjoyment (inversely); and efficiency with enjoyment (inversely).

4.3. Discussion of Study 2

In considering results of Study 2, the superiority of hypertext for most measures was evident in performance, target classification, self-reported efficiency and ease of use. Hypertext is evidently a more mature technique for information visualization and access, as reflected in better software tools and extensive user experience. At the same time, as in Study 1, participant enjoyment didn't match the most efficient user interface. In fact, participants enjoyed VR more, reflecting greater engagement and motivation. These differences suggest, though, that hypertext's structural benefits should inform further development in visualization of hierarchical information.

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

5. Study 3

5.1. Design and Methodology

After Study 2, many usability and performance problems in the prototype design were fixed, largely based on participants' comments. First, the more structured CityScape layout algorithm was used, to take better advantage of perceptual cues in desktop VR [12]. Such cues included radial axes to trace parent-child relationships; circles in the ground plane to delineate category boundaries; variable circle arcs to reflect category sizes; and circle centers with objects and labels to convey category identities. From a design perspective, these visualization elements correspond to Lynch's recommended urban elements. The CityScape algorithm was thus less naturalistic, but more efficient and perhaps aesthetically satisfying. Second, many usability problems in the prototype were fixed. Third, non-information bearing virtual objects were removed, and the rest geometrically streamlined, to improve graphics performance.

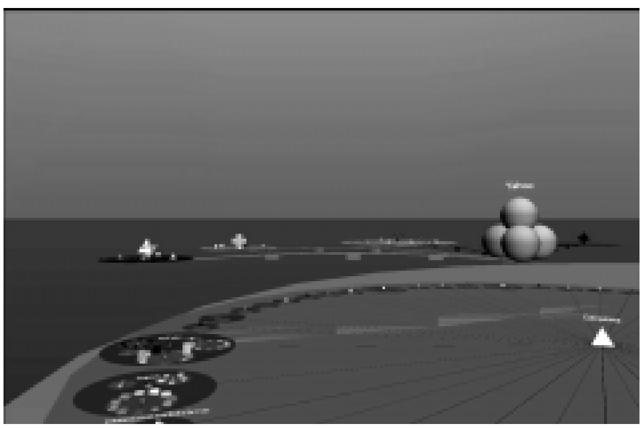

Fig. 5. A view of the "efficient" Day World (Study 3)

The main purpose of Study 3 was to validate a new research prototype (Fig. 5). Accordingly, a variation of Study 2 was run with 8 participants: this variation compared "naturalistic" and "efficient" versions of desktop VR. Since each participant was exposed to only two virtual worlds, hunt blocks grew to 20 minutes each. The study was a 1-way within-subjects design with 2 levels of the factor, world design. Exposure order was fully counterbalanced.

5.2. Results and Discussion of Study 3

In general, ease of use ($F = 12.5$, $p < .010$) and efficiency ($F = 12.5$, $p < .010$) were higher for the "efficient" interface, and comparably lower for the "naturalistic" interface. Ease of use, efficiency, and enjoyment were all significantly correlated ($p < .05$). 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.

Participant responses varied on enjoyment of VR designs. Different participants preferred different designs. Still, efficiency, ease of use, and enjoyment were significantly correlated with each other.

In comparison with results for hypertext in Study 2, the largest difference was in performance. We can speculatively compare 8 efficient VR sessions in Study 3 with 8 hypertext sessions in Study 2, which used the same methodology. (Only the first 15 minutes of sessions from Study 3 are included.) There was no significant difference in performance, though the hypertext mean (40%) was higher than the VR mean (25%).

Observation suggests further benefits of the efficient VR design. Participants often returned to distant virtual locations by recognizing visual structures, when labels were too distant to read. Such user recognition of visual features suggests VR's power for memorability and wayfinding. Moreover, users expressed pleasant surprise at the new design. It struck many as a viable representation of Web structure, which offered some of the experiential engagement of games.

6. Conclusions

To investigate the possible combination of information visualization and large-scale VR, a series of research prototypes was designed, implemented, and subjected to user testing. The final prototype features a structured layout algorithm, geometric abstraction, common landscape and urban metaphors, and controlled variation in color palette and landmark form. It presents a balance of algorithmic structure and naturalistic representation. Experimental results showed the final design to be generally usable and liked by research participants (rated approximately 75 out of 100 for both ease of use and enjoyment).

During these studies, the relationship between subjective reaction and objective performance remained tenuous. An exception to this trend was self-reported efficiency, which was significantly correlated with both attitudinal and behavioral measures. Nevertheless, it is worth recalling that software designs can also benefit from consideration of criteria in addition to efficiency – in particular, a sense of engagement, which increases motivation for task performance.

Further research is necessary to assess the promise and limitations of this visualization approach. First, it would be valuable to investigate efficient use of the third dimension: what VR designs would support such use? Another issue is trade-offs between handcrafted and algorithmic solutions to information representation. For example, to what extent would 2- or 3D icons improve user navigation and learning, and how can such icons be efficiently incorporated into VR visualization? Given the novelty of the medium, another potential line of inquiry concerns the immersion offered by different hardware and software. What effect would such equipment have on the issues in this paper? Finally, it is worth considering possibilities for very large-scale, dynamic structures such as the Web, and for social navigation of shared virtual worlds by physically distant collaborators. It is hoped that the present paper will encourage the pursuit of these possibilities for future research.

7. Acknowledgement

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

8. References

[1] S. K. Card, G. G. Robertson, and J. D. Mackinlay, "The Information Visualizer, an information workspace," in *Proceedings of CHI '91*, ACM, pp. 181-188.

[2] R. P. Darken and J. L. Sibert, "Wayfinding strategies and behaviors in large virtual worlds," in *Proceedings of CHI '96*, ACM, pp. 142-149.

[3] N. G. Vinson, "Design guidelines for landmarks to support navigation in virtual environments," in *Proceedings of CHI '99*, ACM, pp. 278-285.

[4] J. A. Waterworth, "A pattern of islands: exploring public information space in a private vehicle," in *Multimedia, Hypermedia and Virtual Reality: Models, Systems, and Applications*, P. Brusilovsky, P. Kommers, and N. Streitz, Eds. Berlin: Springer, 1996.

[5] M. Chalmers, "Using a landscape metaphor to represent a corpus of documents," in *Proceedings of COSIT '93*, Springer, pp. 377-390.

[6] M. Hemmje, C. Kunkel, and A. Willett, "LyberWorld -- a visualization user interface supporting fulltext retrieval," in *Proceedings of SIGIR '94*, ACM, pp. 249-258.

[7] C. Chen, *Information Visualisation and Virtual Environments*. London: Springer, 1999.

[8] G. G. Robertson, J. D. Mackinlay, and S. K. Card, "Cone trees: animated 3D visualizations of hierarchical information," in *Proceedings of CHI '91*, ACM, pp. 189-194.

[9] J. Lamping, R. Rao, and P. Pirolli, "A focus+context technique based on hyperbolic geometry for visualizing large hierarchies," in *Proceedings of CHI '95*, ACM, pp. 401-408.

[10] K. Lynch, *The Image of the City*. Cambridge, Massachusetts: MIT, 1960.

[11] R. Passini, *Wayfinding in Architecture*. New York: Van Nostrand Reinhold, 1984.

[12] C. Keshkin and V. Vogelmann, "Effective visualization of hierarchical graphs with the Cityscape metaphor," in *Proceedings of the Workshop on "New paradigms in information visualization and manipulation*" (1997), ACM, pp. 52-57.

[13] E. Rennison, "Galaxy of News: an approach to visualizing and understanding expansive news landscapes," in *Proceedings of UIST '94*, ACM, pp. 3-12.

[14] C. McKnight, A. Dillon, and J. Richardson, *Hypertext in Context*. Cambridge, UK: Cambridge University, 1991.

[15] D. Modjeska, "Parallel Worlds: spatial and textual representations of information structure," in *Proceedings of the Workshop on Personalized and Social Navigation in Information Space* (1998), Swedish Institute of Computer Science, pp. 130-137.

[16] A. Dieberger, "Providing object-based navigation for the World Wide Web," in *Proceedings of COSIT '95*, Springer, pp. 93-106.

[17] D. Modjeska, *Desktop virtual-reality modeling of hierarchical data*. Doctoral Dissertation. Toronto: University of Toronto, 2000, forthcoming.